RAND Journal of Economics
Vol. 0, No. 0, XXXX 2021
pp. 1–28

Mergers and marginal costs: New evidence on hospital buyer power

Stuart V. Craig*

Matthew Grennan*

and

Ashley Swanson**

We estimate the effects of hospital mergers, using detailed data containing medical supply transactions (representing 23% of operating costs) from a sample of US hospitals, 2009–2015. Premerger price variation across hospitals (Gini coefficient 7%) suggests significant opportunities for cost decreases. However, we observe limited evidence of actual savings. In this retrospective study, targets realized 1.9% savings; acquirers realized no significant savings. Examining treatment effect heterogeneity to shed light on theories of "buyer power," we find that savings, when they occur, tend to be local, and potential benefits of savings may be offset by managerial costs of merging.

1. Introduction

■ In the last several decades, hospital systems have consolidated substantially through horizontal mergers (Cutler and Scott Morton, 2013; Gaynor and Town, 2012). Researchers and regulators have raised concerns about these mergers' potential negative welfare effects due to increased concentration, and hospital mergers are heavily represented in Federal Trade Commission investigations (Coate, 2018; Dafny, 2014). A typical justification for these (and many other) horizontal mergers is their potential to generate various "efficiencies," leading to lower prices, improved quality, enhanced service, and/or new product introductions, with a particular emphasis on price (U.S. Department of Justice and the Federal Trade Commission, 2010;

* University of Pennsylvania; stucraig@wharton.upenn.edu, grennan@wharton.upenn.edu.
** Columbia University; ats2180@gsb.columbia.edu.
The data used in this article were generously provided, in part, by the ECRI Institute (www.ecri.org). We gratefully acknowledge financial support from the Leonard Davis Institute, Wharton Dean's Research Fund and Public Policy Initiative, the Wolpow Family, and NSF Grant SES-1559485. We thank Zack Cooper, Martin Gaynor, Charlie Gray, Matthew Schmitt, and John Van Reenen for assistance in constructing the merger panel; Matthew Schmitt, Fiona Scott Morton, Robert Town, Ali Yurukoglu, and multiple conference and seminar audiences for their insights. Harrison Byers, Caleb Diaz-Spatharos, and Gi Heung Kim provided excellent research assistance. Any errors are our own.

Vogt and Town, 2006). A necessary, though not sufficient, condition for mergers to lower prices is that they first lower marginal costs. Whether any true marginal cost reductions should then be counted as "cognizable efficiencies" by antitrust authorities is a contentious issue (Hemphill and Rose (2018)).

In this merger retrospective, we provide new estimates of the effects of hospital mergers on marginal costs, using unique data containing hospital supply purchase orders issued by a large sample of US hospitals from 2009–2015. Hospital supplies and devices accounted for a quarter of the growth in inpatient hospital spending between 2001 and 2006 (Maeda et al., 2012). These estimates are interesting not only as a window into potential downstream price effects, but also in that they allow us to investigate "buyer power." For each of the products in our data, prices are determined in negotiation. It is conventional wisdom that "bigger is better" in bilateral negotiations, and this issue has broad policy relevance, not just for antitrust (Carlton and Israel, 2011) but also as a key issue for evaluating the efficiency of decentralized procurement markets. For example, many policymakers advocate for more centralized procurement of healthcare products and services by federal and state governments (LaVito, 2018), rather than procurement via decentralized bilateral bargaining as is the norm in the United States.

The hospital supply product markets in our full dataset account for 23% of hospital operating costs (34%, excluding labor). Thus, savings on supply input costs represent perhaps the largest potential merger-related savings that are unambiguously *marginal*.[1] In the current context, negotiation can take place directly between a hospital or health system administrator and a representative of the product's manufacturer, or hospitals may rely on group purchasing organizations (GPOs) to negotiate their contracts for some products.

In calculating the potential savings they could achieve as an integrated entity, merging parties typically cite the wide variation in prices paid across hospitals and argue that the merged entity will be able to obtain discounts based on taking the best price among the merging parties, plus leveraging any "buyer power" the larger merged entity might possess. This variation is indeed large, with a Gini coefficient of 0.073 (or a coefficient of variation of 0.219) for the average category, across hospitals for the same exact brand-month. The Gini coefficient (equivalent to one half of the mean absolute difference between any randomly selected pair of hospitals) is a useful benchmark for potential savings, as it translates the pre-merger price variation in our data into the expected savings that might be achieved if the worse off party in a random pair of hospitals were able to improve their price by "splitting the difference."[2,3]

Together, the magnitude and variation in supply spending are also substantial relative to both hospital profit margins and downstream costs of hospital care.[4] However, whereas a hospital's exercise of market power as a supplier of health care services might entail renegotiation over a menu of prices with a handful of commercial insurers, that same hospital's exercise of market power as a buyer of medical and surgical supplies might entail renegotiations with hundreds of

[1] When labor costs are cited as merger "efficiencies," they are often either administrative in nature or due to the shifting of services across facilities. The former are arguably less "marginal," whereas the latter may involve a quality trade-off (Noether and May, 2017).

[2] We note that the Gini coefficient would only precisely align with theoretical expected savings in a very restrictive model of possible sources of contracting heterogeneity. However, it represents the type of expected savings calculation merging hospitals might perform, as input prices are not typically shared during pre-merger due diligence.

[3] In a more extreme example, when the two large insurers Anthem and Cigna proposed to merge, they argued that the integrated entity would obtain the best of their pre-merger prices (United States District Court for the District of Columbia, 2017).

[4] According to the American Hospital Association 2018 Trendbook, the average hospital operating margin in 1995–2016 was 4.4%. The coefficient of variation (CV) for knee replacement procedure prices across hospital markets is 0.32 (Cooper et al., 2019), whereas we estimate the CV across hospitals for knee prosthesis prices is 0.24.

vendors and also might require substantial managerial effort obtaining buy-in from disparate end-users within the hospital.[5] That is, it may be relatively costly to reduce costs.[6]

This article builds on a large body of literature on the effects of hospital mergers, and particularly on recent work that estimates the effects of mergers on overall hospital costs (e.g., Dranove and Lindrooth, 2003; Harrison, 2010; Schmitt, 2017) and on labor costs (Prager and Schmitt forthcoming). The general literature to date on the effects of hospital concentration has not suggested that consolidation improves efficiency. Although an exhaustive review of the evidence is outside the scope of this article, the dominant narrative appears to be one of mergers decreasing quality (Beckert et al., 2012; Capps, 2005; Ho and Hamilton, 2000; Romano and Balan, 2011; Town et al., 2006), increasing prices (Capps and Dranove, 2004; Dafny, 2009; Dafny et al., 2019; Haas-Wilson and Garmon, 2011; Krishnan, 2001; Sacher and Vita, 2001; Tenn, 2011; Thompson, 2011), and weakly decreasing costs (Dranove and Lindrooth, 2003; Harrison, 2010; Schmitt, 2017). We follow this literature in estimating difference-in-differences models that compare cost trends at target and acquirer hospitals to control hospitals. The unique contribution of this article relative to this prior literature is the fine-grained nature of the cost data, with precise prices and quantities paid across nearly all hospital supplies at the sample hospitals.

We find that, for a fixed basket of 37 of the most important hospital supply categories, the average merger target in our sample can expect to save 1.9% or \$214 thousand dollars per year (95% confidence interval [\$79,568, \$349,236]), and the average acquirer can expect no savings (point estimate $-$\$90 thousand "savings"; confidence interval [$-$\$158,518, $-$\$21,968]). To put this in context with a simple example, for a merger with the same size target and acquirer, this would suggest total savings across merging parties of $\frac{1.9-0.9}{2} = 0.5\%$.[7] Based upon our understanding of the information typically available to merging parties regarding input costs, a simple approach via which merging parties might claim expected savings would be based on a measure of price dispersion such as the Gini coefficient (across hospitals, controlling for brand-time), which gives an estimate of what expected savings would be if two randomly chosen hospitals in the data were to merge and obtain the better of their contracts for every item they purchase. Our results would translate into average realized savings that are about 7% of that claimed under this approach. If merging parties claimed higher expected savings based on supposed greater "buyer power" of the integrated entity, this ratio of realized savings to expected savings would be even lower.

In addition to being interesting in their own right, mergers also provide useful variation to examine economic mechanisms underlying "buyer power" at a scale beyond individual case studies, as they represent a shock to hospital system size that is plausibly uncorrelated with trends in any particular supply category market. As discussed in Section 3, the literature on buyer power points to multiple theoretical mechanisms via which increased buyer size might impact input prices, and we examine several of these through triple-difference specifications allowing for heterogeneity in the merger treatment effects.

Much like in markets for hospital services, prices in hospital input markets are typically determined via bilateral negotiations. The literature over time has developed an increasingly sophisticated approach to modeling downstream hospital price negotiations (Town and Vistnes, 2001; Capps et al., 2003; Ho, 2006, 2009), and many of the same principles apply to considering input price negotiations. In such an environment, the effects of mergers can be complex, depending on how they impact market structure and bargaining abilities (Dafny et al., 2019; Gowrisankaran et al., 2015; Grennan, 2013; Lewis and Pflum, 2015). Market power in upstream supply markets

[5] This price variation has been found to be driven by heterogeneity in hospital preferences and bargaining ability (Grennan, 2013, 2014), and by variation in information and contracting frictions (Grennan and Swanson, 2018, 2020).

[6] Following this logic, if savings on "important" product categories in our analysis here are easier to obtain than savings on the more than 3000 medical supply categories in our data, our results are likely an upper bound on total supply savings achieved post-merger.

[7] Most mergers involve acquirers that are larger than targets, so the average merger would tend to involve even smaller savings, and of course these are average treatment effect estimates, so the outcomes for any particular merger could differ.

may decrease input prices directly, but an important countervailing indirect effect may also occur: market power in downstream markets for hospital services may lead to higher downstream prices, and that greater overall pie may be "shared" with suppliers (Ho and Lee, 2017). Finally, as managerial attention, skill, and incentives play an important role in supply contracting, mergers may have disruptive effects in the short run and may be mediated by geographic proximity, and returns to scale may be positive or negative in the long run (Agrawal et al., 1992; Beckmann, 1960; Fulop et al., 2002; Minemyer, 2017; Uysal et al., 2008).

To shed light on mechanisms underlying pricing, we explore heterogeneity in our reduced form merger treatment effects by: merging parties' size and market overlap, supply market concentration, and downstream market power. Our analyses consider whether cost reductions (if any) are achieved through lower negotiated prices, cost-reducing shifts in utilization, or both. We also examine the role of "standardization"—industry terminology used to refer to hospitals' use of restrictive supply sets (e.g., use of one implant vendor for most joint replacement procedures). Recent research has found that restrictive networks of health care providers (Gruber and McKnight, 2016; Ho and Lee, 2019), restrictive drug formularies (Duggan and Scott Morton, 2010), and restrictive pharmacy networks (Starc and Swanson forthcoming) can lead to lower costs for insurers; similarly, hospitals argue that standardization of medical supply purchasing results in large savings (Noether and May, 2017).

We find that target hospitals' post-merger savings are driven by a 3.4% decrease in prices negotiated within physician preference item (PPI) brands.[8] PPIs are expensive implantable devices over which physicians typically have strong brand preferences, and are frequent targets in policy discussions around excessive spending on medical technologies and tensions between physicians and hospitals as coproducers of health care.[9] The within-brand effect is slightly smaller, though not statistically different, from the within-category effect of 3.8% (which accounts for shifts in quantity utilized to different brands within the category). Thus, targets' PPI savings can be almost entirely accounted for by targets negotiating lower prices, rather than cost-saving changes in usage patterns. These savings are significantly larger for local mergers (5.9%). The point estimates are also larger when the acquiring system is large (4.3%). Targets do not show economically or statistically significant savings on relatively inexpensive non-PPI supplies. Finally, there is no effect of merging on targets' standardization rates.

By contrast, acquirers show little evidence of savings post-merger. We document a quite small, though statistically significant, increase (1.7% within brand; 1.2% within category) in acquirers' PPI costs post-merger. These are only slightly counterbalanced by a small, marginally significant 1.2% reduction in acquirers' non-PPI prices. Interestingly, the cost increase point estimate for PPIs is larger for large acquirers than for small acquirers. Positive price effects for large acquirers are consistent with several managerial and economic theories, which we discuss in Section 3. However, we interpret this result with some caution, as it (unlike the result for targets' purchase of PPIs) disappears in analyses of longer time horizons.

In sum, the net effect of merging on any given party's costs is small. Where significant effects do exist, they depend on multiple countervailing forces, and these forces bear out unevenly across targets and acquirers. Our findings are consistent with mergers inducing an increase in buyer power that is (1) driven by local returns to scale, and (2) more influential for merger targets than for (even small) acquirers. The finding of a positive price effect for large acquirers is consistent with the costs of a merger disrupting management outweighing any benefit from improved buyer power, for merging parties experiencing small relative size increases. We find little evidence that savings, where they exist, are mediated by supplier concentration or by a change

[8] The acronym PPI is commonly used in the healthcare industry (e.g., Lagasse, 2018) and we accordingly use it here as well, but note that it refers everywhere in this manuscript to "physician preference item" and not to any price index as is common in economics (e.g., "producer price index.")

[9] Past efforts by hospitals to manage utilization of surgical inputs have encountered significant resistance from surgeons (Nugent et al., 1999). Navathe et al. (2017) estimate that one health system's participation in a bundled payment program led to substantial savings on implant costs, perhaps aided by gainsharing arrangements with physicians.

in downstream market power. Notably, in contrast to previous empirical findings on restrictive contracting by insurers, we find no evidence that merger savings are amplified when hospitals are standardized.

2. Data and setting

□ **Hospital purchasing data.** The primary data used in this study come from a unique database of all supply purchases made by over 1000 US hospitals during the period 2009–2015. The data are from the PriceGuide™ benchmarking service (hereafter, "PriceGuide data") offered by the ECRI Institute, a non-profit health care research organization. For each transaction, we observe price, quantity, transaction month, and supplier. The reported data are of high quality because they are typically transmitted as a direct extract from a hospital's materials management database. Hospitals have strong incentives to report accurately because the analytics the benchmarking service's web portal provides are based on comparing the hospital's submitted data to that of others in the database.

Our analyses consider price negotiations between hospitals and suppliers for a large number of important product categories. Throughout this draft, we use the term "product category" to refer to the "Universal Medical Device Nomenclature System (UMDNS)" grouping code included in the transaction files. The UMDNS system generally classifies products by intended purpose and mechanism of action (e.g., drug-eluting coronary stents have UMDNS code 20383). We use the term "brand" to refer to the "product" level at which prices are negotiated; for example, Medtronic Resolute Integrity drug-eluting coronary stent. In practice, we identify brands using an algorithmic approach that groups together stock-keeping-units within manufacturers that vary in terms of factors such as size and color, but not price. The algorithm endeavors to balance our goal of identifying meaningful, pricing-relevant differences in product attributes (e.g., quality) against the potential for sparsity problems. As discussed in online web Appendix A.1, the patterns we document are insensitive to how we classify brands. Finally, we use "product class" to refer to the distinction between FDA risk classes I–II, which tend to be commodities (e.g., dressings) and other medical/surgical products (e.g., catheters), versus FDA risk class III, which are placed in this class because they are deemed "necessary for the sustainment of life" and thus tend to include high-tech physician preference items (e.g., coronary stents).

Our empirical analyses examine products that are among the top 50 product categories by *either* total spending *or* transactions. There are 71 such "top" categories total, but once we omit product categories that are too broad or with missing or inconsistent data, 37 remain. Some categories in the UMDNS grouping are excessively broad and would not necessarily be used in the same procedures or by the same providers. Codes such as "food item," "office supplies," and various "kits" are flagged as too broad. For example, "IVD Kits" include microbial detection kits costing $2.14 on average, as well as tests for antibiotic-resistant bacteria colonization costing $4400 on average. We also excluded codes for which we could not confidently calculate price per unit due to missing conversion factors (e.g., 10 units per box) or inconsistent unit of measure (e.g., "box" vs. "case"). Other categories were omitted based on "reasonableness" of the observed price variation – categories for which the coefficient of variation in price exceeded 200% were excluded. See online web Appendix A and Grennan and Swanson (2020) for further details and examples.

□ **Hospital and merger data.** To perform the analysis in the current study, we obtained permission to contract a trusted third-party to match facilities in the PriceGuide data to outside data from the Centers for Medicare and Medicaid Services (CMS), the American Hospital Association (AHA), and a merger roster. The third-party then provided us with access to the merged data for analysis, with hospital-identifiable information removed.

We obtained merger data from Cooper et al. (2019), which contain nearly all hospital mergers from 2000 to 2014. The data were generated by correcting known problems in the AHA:

errors in timing of mergers due to lagged survey response and erroneous combination of multiple facilities into single observations post-merger. These data were cross-checked against data from Schmitt (2017) and several business intelligence databases: Irving Levin Associates, Factset, and SDC Platinum. For more details on the merger data, see Appendix D of Cooper et al. (2019).

□ **Representativeness.** Each analytic sample includes facilities in the PriceGuide data that merged uniquely by name and location to general acute care hospitals in the AHA data. The top panel of Table 1 describes how the full sample of AHA-surveyed general acute care hospitals compares to our sample of PriceGuide subscribers. The facilities in the purchase order data voluntarily joined a subscription service that allows them to benchmark their own prices and quantities to those of other members in the database and thus may not be a random sample of US hospitals. In particular, subscription is costly, so we expect hospitals with greater concerns about, or attention to, supply costs to be over-represented in the database. In a survey of database members, "cost reduction on PPIs" and "cost reduction on commodities" were the first and second (and nearly tied) most commonly cited reasons for joining. This accords with our own conversations with purchasing managers who cite a broad array of reasons and product areas as motivations for benchmarking.

The PriceGuide data contain a large number (855) of hospitals, covering 17% of the hospitals in the AHA sample.[10] These hospitals tend to be relatively large (in terms of employment, bed count, and admissions), they use more technologies, and they are more often teaching and nonprofit hospitals, relative to the overall AHA sample. The PriceGuide sample also weights the Northeast and West Census regions of the United States, as opposed to the Midwest and South regions, relatively heavily.

The second and third panels of Table 1 summarize the M&A transactions in the AHA and PriceGuide samples from 2009–2015.[11] The full AHA merger panel contains 445 transactions impacting 661 targets and 1753 acquirers. The full PriceGuide merger panel covers a large sample of these: 211 transactions impacting 121 targets and 301 acquirers. However, our analytic samples are limited to the *first* transaction observed for each target and acquirer in the PriceGuide data, and our main specifications include only those transactions for which we observe at least one full year of pre- and post-merger data.[12] Because PriceGuide members join the database in a staggered fashion over time, such that our data include many more hospitals in 2015 than in 2009, this requirement reduces our sample to 23 transactions with 33 target hospitals, 31 transactions with 86 acquirer hospitals, and 433 non-merging controls. The data contain 49 unique merger transactions. Thus, for most of our mergers, we have detailed purchasing data for either the target or acquirer, not both. This provides little ability to explore issues such as price convergence for the different parties involved in a given merger. Although this restriction is costly, our sample contains many merger case studies – whereas many analyses have considered single mergers in isolation (Kwoka, 2015)—and the rich transaction-level cost information across 37 different product categories compensates in detail for these limitations.

Focusing on the third panel, a comparison of columns (2) and (3) to column (1) illustrates that, along most dimensions, the PriceGuide *merger* samples are slightly more representative of the AHA *merger* sample on observables, than the full PriceGuide sample is of the full AHA sample. This is primarily because larger hospitals are more likely to be involved in M&A transactions. However, it remains the case that our analyses are identified from a sample of relatively

[10] The full PriceGuide database includes 1155 facilities coded as "Hospitals" or "Health Systems" for which we observe the date of database join. 891 of these facilities match uniquely to a general acute care hospital in the AHA data, and 855 also purchase at least one of our focal product categories.

[11] All analytic samples used in this study impose that merging facilities have at least one calendar year of transaction data pre- and post-merger. Accordingly, all *mergers* in this study take place during 2010–2014.

[12] Post-merger here refers to years following the calendar year of the merger (i.e., $y > \tau_h$).

TABLE 1 Merger Sample Restrictions

	(1) Full Sample, 2009–2015	(2) PriceGuide Hospitals, 2009–2015	(3) PriceGuide Hospitals, Full Support, 2009–2015
Hospital characteristics			
Beds	166.6	269.7	
FTEs	960.3	1,784.0	
Technologies	45.2	62.9	
Teaching	0.245	0.407	
Admissions	7,367.1	13,077.2	
Non-Profit	0.610	0.784	
Number of HMO contracts	1.1	1.6	
Percent Medicaid	0.172	0.193	
Percent Medicare	0.505	0.461	
Output price	9256	10,731	
Urban	0.617	0.704	
Midwest	0.302	0.238	
Northeast	0.129	0.222	
South	0.386	0.291	
West	0.183	0.248	
Transaction characteristics			
Number of transactions	445	211	49
Number of targets	661	121	33
Number of acquirers	1753	301	86
Number of controls	2560	433	433
Median target size	1	1	2
Median acquirer size	45	8	13
Merging hospital characteristics			
Beds	191.0	260.5	279.8
FTEs	1065.4	1710.6	1790.3
Technologies	48.1	62.2	65.7
Teaching	0.272	0.385	0.369
Admissions	8755.6	12,707.4	14,131.0
Non-Profit	0.636	0.849	0.811
Number of HMO contracts	1.3	1.6	1.9
Percent Medicaid	0.171	0.183	0.183
Percent Medicare	0.499	0.472	0.480
Output price	9415	10,720	10,569
Urban	0.679	0.671	0.721
Midwest	0.297	0.262	0.245
Northeast	0.131	0.233	0.262
South	0.431	0.263	0.234
West	0.141	0.242	0.259

Notes: Each column reports the counts and characteristics of hospitals in the data at varying levels of sample restrictions. Column (1) reports data on AHA hospitals and mergers for 2009–2015. Column (2) presents the overlap between the AHA hospitals and the PriceGuide member hospitals satisfying our inclusion criteria as described in the text. Column (3) presents data on merging and control hospitals in the PriceGuide data for our main analytic sample. Data on beds, FTEs, technologies, admissions, teaching status, non-profit status, number of HMO contracts, and Medicare and Medicaid share come from the AHA Annual Survey. Following Acemoglu and Finkelstein (2008) and Cooper et al. (2019), we measure technologies using the complete list of binary facility indicators available in the AHA. Output price is calculated using data from the CMS HCRIS and Medicare Impact Files as in Dafny et al. (2019).

large hospitals enrolled in a benchmarking service; if, for example, this implies that sample hospitals are especially sophisticated, they may benefit more or less from merging than the average treated hospital in the AHA data.

For one product category—coronary stents—we have been able to compare the PriceGuide sample to an external source that is explicitly intended to provide an accurate picture of market shares and prices by US region: the Millennium Research Group's (MRG) Markettrack™ survey of catheter labs. As discussed in Grennan and Swanson (2020), the prices paid in the MRG and PriceGuide samples during 2010–2013 are statistically close to one another, with the average prices paid (controlling for brand-time trends) in the MRG sample being slightly higher (mean \$1666, s.d. \$149) than those paid by hospitals in the PriceGuide sample (mean \$1631, s.d. \$120) during the period before they joined the benchmarking service. The matching specifications described in Section 4 are intended to ameliorate concerns regarding internal validity of our estimates given sample selection, but cannot speak directly to questions of external validity.

Online web Appendix E addresses external validity by reweighting our regression sample to approximate the national distribution of merging hospitals on various observable characteristics. These results are qualitatively similar, with point estimates that are slightly smaller in magnitude. However, we offer all findings in this study with the caveat that we cannot rule out bias driven by dimensions of sample selection not observable in publicly available sources.

□ **Price variation, by product class.** Each product category in our analytic sample is summarized in Table 2.

The top panel of Table 2 contains non-PPIs. Non-PPI products can be used in a hospital setting by staff members with a variety of roles and scopes of practice. Some of these are essentially commodities (e.g., surgical drapes): Conditional on a few characteristics, such as material, we do not expect particular manufacturers to be strongly preferred. Some are used by physicians in moderately invasive procedures and brands may vary in perceived quality (e.g., surgical staplers), but they tend to be less critically linked to patient outcomes than Class III PPIs. The average non-PPI category is purchased by 483 sample hospitals. A non-PPI product costs \$714 per unit on average, and the average sample hospital spent \$13,459 per month on the average non-PPI. These averages obscure substantial heterogeneity. For example, nylon sutures cost \$8 per unit, whereas bone grafts cost \$2562 per unit.

The bottom panel of Table 2 contains physician preference items. For PPIs, usage is driven by brand preferences of physicians, often surgeons, choosing which brand to use to treat a given patient. PPIs tend to be expensive cardiac and orthopedic surgical implants used in advanced procedures and are not purchased by all hospitals: only 378 sample hospitals purchased the average PPI, and only 254 purchased "Cardiac Valve Prostheses." PPIs are also used less frequently by hospitals that purchase them: the average PPI category sees 15 products used per month vs. 172 for non-PPIs. Nevertheless, purchasing hospitals spend twice as much per month on the average PPI category (\$30,900 vs. \$13,459), due to PPIs' higher average prices. PPIs are more likely to be sold and distributed by highly specialized sales representatives whose relationships and expertise are valued by physicians. In some cases, representatives are even present in the operating room during procedures.

The competitive landscape varies dramatically across these classes. There are more brands to choose from in non-PPIs (253) versus PPIs (129). For PPIs, each brand is typically purchased directly from its manufacturer (there are 17 in the average category), and hospitals/systems tend to negotiate their own prices. By contrast, the average non-PPI is available from 39 vendors, brands produced by a particular manufacturer may be sold by multiple vendors, and hospitals are more likely to rely on GPO pricing (Schneller, 2009).[13] Despite these differences, both classes

[13] We do not directly observe which products are purchased via a GPO in our data, but there is significant price dispersion in all product categories in our data in spite of them. As discussed in Section 3, we consider GPOs to be an unobserved feature of the setting that may mediate the effects of mergers for some products.

TABLE 2 Summary of Product Categories

	% of *spend*	$\overline{spend}_{hmy}$	N_{hjmy}	N_h	N_{tar}	N_{acq}	N_j	HHI_v	$\overline{q}_{hmy}$	$\overline{p}_{hjmy}$	$CV_{h\|jmy}$	$Gini_{h\|jmy}$
Nylon sutures	0.1	1111	45,931	524	29	77	201	0.25	325	8	0.27	0.06
Bone wres	0.1	1658	74,437	511	27	72	123	0.17	42	102	0.38	0.13
Surgical drapes	0.2	2146	94,979	523	31	71	310	0.31	841	11	0.28	0.08
Tracheal tubes	0.1	2558	64,621	530	27	75	176	0.26	443	63	0.54	0.15
Trocars	0.3	4942	65,531	520	26	73	188	0.16	141	76	0.31	0.09
Suture anchors	0.4	6327	52,957	503	24	71	61	0.41	19	381	0.19	0.07
Drill bits	0.4	7104	235,142	509	26	72	335	0.26	51	189	0.22	0.08
Electrosurgical forceps	0.6	9076	42,604	492	22	63	93	0.10	28	905	0.40	0.12
Polymeric mesh	0.5	8867	93,376	528	32	75	385	0.17	16	977	0.20	0.06
Allografts	0.5	10,537	44,545	472	21	56	249	0.84	71	1634	0.19	0.05
Bone nails	0.5	10,938	53,259	480	23	65	123	0.29	8	1558	0.19	0.07
Trauma bone plates	0.6	11,417	187,557	505	27	72	549	0.56	16	787	0.14	0.05
Surgical staplers	0.7	10,897	77,544	511	26	69	238	0.20	51	367	0.24	0.08
Bone implant putty	0.5	12,536	65,885	475	23	60	229	0.22	12	1114	0.16	0.05
Embolization coil	0.5	13,908	54,766	408	20	57	186	0.40	33	955	0.12	0.04
Spinal bone plates	0.5	13,881	48,428	382	16	43	234	0.18	9	1709	0.26	0.08
Guiding cath.	0.8	13,949	251,790	502	26	73	324	0.17	276	226	0.26	0.09
Guide wires	1.0	17,878	352,437	523	30	77	423	0.13	317	122	0.26	0.09
Trauma bone screws	0.9	19,272	369,396	514	29	75	317	0.52	195	154	0.19	0.08
Bone grafts	0.8	20,465	40,058	455	21	53	141	0.98	741	2562	0.16	0.04
Ablation/Mapping cath.	0.8	28,121	64,080	329	13	41	107	0.27	28	1,188	0.18	0.07
Spinal bone screws	2.6	68,515	158,571	420	20	52	568	0.22	130	615	0.31	0.10
Non-PPIs: Total	13.3	240,093	2,537,894	551	33	85	5560		3382			
Non-PPIs: Average	0.6	13,459	115,359	483	24	66	253	0.32	172	714	0.25	0.08
Intraocular lenses	0.2	6786	31,855	327	14	33	39	0.55	34	293	0.13	0.05
Spinal rod implants	0.2	7170	60,302	365	16	43	265	0.18	18	444	0.32	0.09
Mammary prosth.	0.5	14,369	23,212	379	14	46	28	0.45	17	843	0.11	0.04
Acetabular hip prosth.	0.7	16,239	57,128	459	25	56	75	0.21	13	1422	0.30	0.12
Spinal stimulators	0.6	25,193	9722	306	8	27	12	0.34	2	15,693	0.13	0.05
Tibial knee prosth.	1.2	27,858	101,896	467	25	60	206	0.19	23	1371	0.23	0.08
Aortic stents	1.0	27,664	27,442	380	18	49	67	0.33	5	6144	0.09	0.03
Femoral hip prosth.	1.3	31,404	144,116	471	25	60	437	0.21	20	1767	0.30	0.10
Pacemakers	1.3	30,878	41,529	419	16	53	33	0.43	7	4409	0.14	0.07
Cardiac valve prosth.	0.7	31,703	16,842	254	11	42	10	0.40	6	5752	0.15	0.07
Femoral knee prosth.	1.4	34,459	90,243	463	25	60	221	0.21	17	2355	0.21	0.07
Spinal spacers	1.4	43,020	69,580	370	16	39	486	0.14	14	3524	0.22	0.06
Cardioverter defib.	1.6	45,482	16,700	336	12	42	31	0.45	3	15,594	0.13	0.05
Resynchronization defib.	1.7	49,308	11,314	324	12	38	10	0.47	2	20,897	0.12	0.05
Drug eluting stents	2.1	71,965	33,151	348	16	53	15	0.32	49	1543	0.08	0.04
PPIs: Total	15.7	327,278	735,032	506	29	74	1935		172			
PPIs: Average	1.0	30,900	49,002	378	17	47	129	0.33	15	5470	0.18	0.06

Notes: Summary statistics for main analysis sample. Authors' calculations from PriceGuide data. For each product category: "% of *spend*" is percent expenditure in entire PriceGuide database; $\overline{spend}_{hmy}$ is average monthly spending; N_{jhmy}, N_h, N_{tar}, N_{acq} and N_j are total number of observations, hospitals, target hospitals, acquirer hospitals, and brands; HHI_v is vendor Herfindahl–Hirschman Index (HHI); $\overline{q}_{hmy}$ is average monthly quantity; $\overline{p}_{hjmy}$ is average unit price; $CV_{h|jmy}$ is within-brand-month coefficient of variation across hospitals, averaged across all brand-months; $Gini_{h|jmy}$ is within-brand-month Gini coefficient of price, averaged over brand-months. "Total" rows contain aggregate statistics for all categories in each product class; unweighted average statistics across category-level analyses listed in the "Average" rows.

are highly concentrated according to the standards typically applied by the US Department of Justice (DoJ) and Federal Trade Commission (FTC), and there is a great deal of price dispersion: the average coefficient of variation, controlling for brand-month fixed effects, is 0.25 in non-PPIs and 0.18 in PPIs. The analogous Gini coefficient is 0.08 in non-PPIs and 0.06 in PPIs. This variation in prices across hospitals could imply large potential savings to be captured by merging parties, if the merged entity can achieve equivalent or better pricing than the best of the pre-merger contracts. To the extent that a larger merged party will have more "buyer power", savings could be even larger. Whether this will indeed happen depends upon the economic mechanisms at work.

3. Mechanisms of interest

■ The welfare effects of any merger "efficiencies" driven by input cost reductions will depend on the underlying mechanisms (Carlton and Israel, 2011). In evaluating proposed mergers, the FTC and DoJ consider whether cost savings are likely to be large, whether they are likely to pass through to consumers, and whether they are "likely to be accomplished with the proposed merger and unlikely to be accomplished in the absence of either the proposed merger or another means having comparable anticompetitive effects" (i.e., whether they are "merger-specific" (U.S. Department of Justice and the Federal Trade Commission, 2010)). Thus, the agencies' consideration of cost savings focuses for the most part on potential welfare gains in the downstream market.

Input cost savings could also be welfare-neutral—a transfer between upstream and downstream firms—or themselves welfare-reducing. Hemphill and Rose (2018) distinguish cases where mergers increase monopsony power or bargaining leverage from cases where there are real resource savings, such as reduced waste. They conclude that the former cases reduce competition and should not be viewed by regulators as cognizable efficiencies. One potential harm cited that may be particularly relevant for medical technology is dynamic inefficiency, in which upstream firms reduce investment and innovation due to increased downstream monopsony power.

Hospital costs include substantial fixed and variable components. The variable portion of hospital costs scales with the number and severity of patients treated, the quantity of labor and "capital" inputs used per patient, and the prices of those inputs. The prices of inputs are, in turn, determined by brand choice and the price negotiated within each brand. Mergers may in theory impact any part of the hospital's cost function. However, fixed costs are unlikely to pass through to patients in the short run, changes in patient mix raise a battery of questions regarding agency and quality of care, and potential negative effects of monopsony power on labor costs are not rated kindly by antitrust authorities (Gaynor and Town, 2012). Thus, in this study, we focus on variable costs that are truly marginal in the sense that they are incurred along with the provision of additional patient care – those costs most likely to impact downstream prices. Specifically, we examine whether mergers lead to economies of scale in variable supply costs due to changes in negotiated prices and/or input choices.

Analyzing prices requires close attention to the details of hospital procurement. In hospital input markets, prices are determined via bilateral negotiations between suppliers and hospitals. In some cases, a GPO may negotiate on behalf of its hospital members. For products purchased through a GPO, a merger could impact purchasing if it moved the combined entity to a more favorable GPO membership volume tier, or if it induced a change in GPO (which could be favorable for some products and unfavorable for others). In many cases, though, the GPO price acts as a starting point for individual hospital/system negotiations, or there is no relevant GPO contract for a given product. This may explain the presence of significant price dispersion in all product categories in our data in spite of the GPO market being dominated by a few large players (Gooch, 2017). Thus, much of hospital purchasing is the result of direct bilateral bargaining between suppliers and a given hospital/system.

The effect of mergers on bilateral bargaining is ambiguous in the economics literature. A merger could affect a supplier's marginal cost or bargaining position (e.g., via economies of scale in distribution). It could also affect the integrated buyer's bargaining position (e.g., via introducing competition from another supplier if there are fixed or search costs involved in contracting). In general, larger buyer firms may obtain better prices if the bargaining-surplus function is concave, in which case the supplier's surplus in bargaining with two independent firms is smaller at the margin than the average surplus in bargaining jointly with an integrated firm (Chipty and Snyder, 1999; Horn and Wolinsky, 1988; Inderst and Wey, 2007; Stole and Zwiebel, 1996). Further, larger buyer firms may spur competition among multiple suppliers (Dana, 2012; Gans and King, 2002; Marvel and Yang, 2008; Snyder, 1996, 1998), implying an important role for supplier market structure. A merger may also affect the buyer's bargaining power/ability (the share of gains from trade obtained, conditional on bargaining positions). In work on insurer-hospital bargaining, Lewis and Pflum (2015) find that bargaining power is a greater determinant of post-merger markups than bargaining position.

Post-merger changes in bargaining power may be driven by various factors, including firm organizational structure, information, incentives, management, and leadership. These same factors may impact the efficiency of input utilization within firms. It is important to note that these effects may be positive or negative. On the one hand, Bloom et al. (2014) find that larger hospitals have better management practices. Conversely, mergers may have disruptive impacts on management, organizational culture, or earnings (Agrawal et al., 1992; Beckmann, 1960; Fulop et al., 2002; Minemyer, 2017).

Input choice and input pricing may also interact. Dana (2012) posits that buyer groups' primary advantage results from their commitment to purchase from a single supplier in differentiated product markets. We see evidence of this in the hospital–insurer bargaining world: Sorensen (2003) shows that insurers' steering ability impacts pricing more than insurers' size; Gowrisankaran et al. (2015) model how insurers steer patients toward cheaper hospitals; and Ho and Lee (2019) note that restrictive hospital networks could reduce insurers' prices by up to 30%.

In this particular context, it is also important to note that hospital mergers entail changes in market power both upstream (with respect to suppliers) and downstream (with respect to insurers). This creates a linkage between upstream and downstream prices, as both negotiations will depend on total surplus. For example, whereas a merger-induced increase in market power in upstream supply markets may have a direct negative effect on input prices, a contemporaneous increase in market power in downstream markets for hospital services would simultaneously have a direct positive effect on service prices, and that greater overall pie may be "shared" with suppliers. Ho and Lee (2017) document how similar countervailing forces create variation in the effects of insurer competition, when such competition impacts upstream negotiations with hospitals and downstream negotiations with employers.

Many of these mechanisms may depend on details such as the geographic location of the merging hospitals or the extent of competition among suppliers. Prior research has found that geographic proximity is a success factor for mergers, perhaps due to information advantages (Uysal et al., 2008). If transmission of management that might affect bargaining power is similarly impacted by geographic proximity, then mergers involving hospitals with market overlap may be more successful at reducing costs. The geography of medical supply sales and hospital competition may matter as well. For example, economies of scale in distribution might be achieved, or joint post-merger negotiation of supply contracts might be easier, if hospitals are in the same geographic sales territories for suppliers. A countervailing factor would be that mergers of hospitals competing in the same local hospital markets may, as noted above, involve downstream price increases for hospital services that could theoretically be shared with hospitals' suppliers.

Competition among suppliers could likewise interact with the mechanisms via which mergers might affect input prices. In a competitive market where markups are low, mergers are only likely to affect prices via cost mechanisms like potential economies of scale in distribution. In a less competitive market, cost changes might be less likely to be passed through, as other changes to bargaining position or bargaining power/ability would have larger potential effects.

4. Empirical specification and identification

■ We estimate two difference-in-differences price specifications. First, using a dataset containing unit prices for each product category (UMDNS code) u, hospital h, brand j, month m, year y, we estimate:

$$lnP_{uhjmy} = \alpha_u * \mathbb{1}[y = \tau_h] + \beta_u * \mathbb{1}[y > \tau_h] + \theta_{hj} + \theta_{jmy} + X_{hmy}\theta^X + \varepsilon_{uhjmy}, \tag{1}$$

where τ_h is the year of hospital h's merger (if any), θ_{hj} is a hospital-brand fixed effect, and θ_{jmy} denotes brand-month-year fixed effects (with j implicitly uj as brands do not span categories by construction). Brand-specific time trends are necessary to control for the presence of brands both early and late in their life cycles in these data. Measuring these trends at the monthly, rather than yearly, level becomes important for our event study specifications in Section 5 below; we do so here as well for the sake of consistency. $X_{hmy}\theta^X$ can in principle control for any further time-varying hospital characteristics, but in our baseline analyses it contains a single dummy variable to indicate month-years after the hospital joins the benchmarking database, so that join effects are not conflated with merger effects.[14] We estimate separate regressions for acquirers and targets; the acquirers regression excludes targets, and vice versa.

Because the month of merger is unknown, we estimate separate treatment effects for the merger year (α_u) and the post-merger period (β_u), focusing on β_u as our main coefficient of interest. Prior work has shown that these hospital supply contracts are typically renegotiated roughly annually (Grennan and Swanson, 2020), and we find that the same is true across our focal product categories. The effects of hospital mergers on hospital procedure prices have been shown to manifest immediately post-merger (Cooper et al., 2019). Therefore, focusing on β_u allows us to estimate merger treatment effects that are unlikely to be biased downward by delayed price adjustments due to structured renegotiations.[15] In our baseline results, we report specifications focusing on one year pre-merger, the year in which the merger occurs, and one year post-merger. Online web Appendix I also contains analyses for alternative time horizons. In each specification, we always limit our estimation sample to the set of hospitals with complete pre- and post-merger data over the specified timing support, and use only the specified range of years. This decision enables us to interpret the resulting treatment effect as the effect of merging on the average treated hospital over that time horizon.

Finally, within each hospital-UMDNS code, we weight each brand using quantity share within hospital. This approach allows us to interpret the resulting coefficient of interest β_u as the treatment effect of merging on the average unit price of product u while also including brand-level controls on the right-hand side.[16] For estimation of specification (1), we hold quantity share weights across brands fixed at those observed for the hospital's first year in the analytic sample. Intuitively, this regression examines the weighted average *within-brand* effect of mergers on negotiated prices, for brands purchased both before and after the merger.

Next, using the same dataset, we estimate:

$$lnP_{uhjmy} = \alpha_u * \mathbb{1}[y = \tau_h] + \beta_u * \mathbb{1}[y > \tau_h] + \theta_{uh} + \theta_{jmy} + X_{hmy}\theta^X + \varepsilon_{uhjmy}, \tag{2}$$

where θ_{uh} denotes a set of hospital fixed effects (that vary by category in regressions where we pool categories). To avoid overweighting products purchased in small quantities in this specification, we weight each hospital-brand-year using the brand's quantity share within the hospital-year. Intuitively, this regression examines the effect of mergers on negotiated prices *per unit* across all brands *within-category*. That is, specification (1) uses fixed pre-merger quantity share weights, whereas specification (2) uses contemporaneous quantity share weights; thus, the estimates from (1) measure the extent to which renegotiation leads to lower prices for the same brand at the same

[14] 67% of hospitals are post-join for at least one year prior to the merger event in which they are used. Online web Appendix Table A14 includes estimates using only such hospitals. The results are slightly larger in absolute magnitude, though broadly consistent with our main findings.

[15] In unreported analyses, we estimate the effect of mergers on renegotiation timing and find no statistically significant effects at conventional levels.

[16] Note that we would obtain similar estimates if we just estimated a regression where the unit of observation was the hospital-use case-month-year (for example, a separate observation for each coronary stent used by a hospital in a given year) and ran (unweighted) OLS.

hospital, and (2) will further include the extent to which the hospital switches usage to different brands. We find this specification of interest because switching to cheaper brands could be one mechanism via which savings could be achieved. However, we interpret these results cautiously, as changes in usage patterns could affect welfare via mechanisms other than price changes, if there are average or patient-specific match quality differences across brands within a category. In all regressions where the dependent variable is price, standard errors are clustered by hospital-brand.

In addition to the product category-specific regressions, we also estimate pooled regressions across all categories within each class (PPI vs. non-PPI). We stack all category-specific data within each class and estimate specifications (1) and (2) with a single α and β for the class, weighting by the total expenditure share for each category across all years of the data. This approach allows us to frame our findings in terms of total potential savings associated with horizontal mergers. Implicitly, however, this approach downweights product categories with low spending shares and hospitals that tend to purchase less expensive product categories.

□ **Identification.** Our empirical approach compares input price trends at merging hospitals to those at non-merging hospitals, around the time of the merger. In Table 3 below, we compare merging and non-merging hospitals in our final analytic sample. Columns in each panel of Table 3 compare the full set of controls (1) to merging target (3) and acquirer (5) hospitals.

Relative to all non-merging controls, target hospitals tend to be smaller (lower employment, fewer beds); they are also less likely to be teaching hospitals and more likely to have non-profit ownership. Although they are smaller than controls, they tend to use more technologies, and have higher monthly purchase quantities for the product categories they purchase.[17] Relative to the average control hospital, target relationships with payers are nuanced: They have above-average contracting with managed care organizations (proxied by count of contracts with health maintenance organizations (HMOs)); they rely relatively less on Medicaid and more on Medicare for admissions; and their average case-mix-adjusted price per inpatient admission is significantly lower.

Acquirers show a different pattern. Relative to controls, they: are larger; are more often teaching hospitals and more often non-profit; use more technologies; have more HMO contracts; and have a similar price per admission.

Given these differences in composition, we might be concerned that merging and non-merging hospitals exhibit very different purchasing patterns even prior to the merger, and more importantly, that they might have different latent trends in input purchasing (which would invalidate the core assumption behind the differences-in-differences research design). To address this issue, first we note that the input price indices for merging and non-merging hospitals are not very different. In Table 3, we see that, relative to non-merging control hospitals, targets have about 2% to 3% lower prices pre-merger, whereas acquirers have about 3% to 5% lower prices pre-merger.[18]

We also address observed differences directly in our preferred specifications. We match both target and acquirer hospitals to a subset of non-merging hospitals in order to ensure that "treated" merging hospitals are similar to the "control" non-merging hospitals, at least along observable dimensions. Within each product category, we match each merging hospital to its 10 nearest non-merging neighbors using Mahalanobis distance.[19] Distances are calculated based

[17] Following Acemoglu and Finkelstein (2008) and Cooper et al. (2019), we measure technologies using the complete list of 153 binary facility indicators available in the AHA. These vary widely, encompassing burn care, chemotherapy, Meals on Wheels, psychiatric child/adolescent services, and proton beam therapy.

[18] Input price indices are hospital fixed effects recovered from a stacked regression of log price on brand-month-year fixed effects and hospital fixed effects. Intuitively, they represent hospital-level residual price variation holding the basket of product categories and brands fixed.

[19] Online web Appendix D discusses the performance of alternative matching algorithms, and includes pooled regression results for a subset of matching approaches.

TABLE 3 Comparison of Merging and Non-Merging Hospitals

	(1) Controls	(2) Target Controls (Matched)	(3) Targets	(4) Acquirer Controls (Matched)	(5) Acquirers
Panel A: Non-PPI Purchasers					
FTEs	2540	2533	2246	2893	2713
Technologies	74.26	77.69	76.13	79.28	79.53
Beds	355.08	333.62	306.37	406.51	404.91
Number of unique products	49.09	32.24	34.95	37.77	42.01
Average monthly quantity	244.80	326.99	356.36	229.31	220.11
Admissions	17,431	16,936	15,384	20,458	19,995
Teaching	0.55	0.49	0.39	0.63	0.60
Non-Profit	0.75	0.94	0.91	0.81	0.77
Number of HMO contracts	1.68	1.95	2.05	1.89	1.85
Percent Medicaid	0.20	0.16	0.15	0.19	0.20
Percent Medicare	0.44	0.49	0.51	0.45	0.45
Output price	12,440	12,397	9,518	12,624	12,593
Input price index (θ_h)	4.95	4.96	4.93	4.95	4.92
Number of hospitals	433	286	33	369	85
Panel B: PPI Purchasers					
FTEs	2674	2484	2366	2945	2897
Technologies	75.74	78.78	79.62	78.86	81.46
Beds	372.06	346.59	324.79	417.31	427.50
Number of unique products	23.68	15.08	16.50	15.12	18.73
Average monthly quantity	24.34	27.48	31.90	27.14	30.35
Admissions	18,336	17,258	16,314	20,865	20,964
Teaching	0.58	0.50	0.43	0.63	0.64
Non-profit	0.75	0.95	0.93	0.80	0.75
Number of HMO contracts	1.71	2.01	2.03	1.91	1.77
Percent Medicaid	0.20	0.16	0.15	0.19	0.20
Percent Medicare	0.44	0.48	0.50	0.45	0.44
Output price	12,693	12,250	9538	12,816	12,711
Input price index (θ_h)	7.22	7.21	7.19	7.21	7.17
Number of hospitals	403	242	29	330	74

Notes: Each column reports the counts and characteristics of merging and non-merging hospitals in the data. Column (1) shows characteristics of all non-merging hospitals. Column (2) shows the subset of these controls that serve as the matched sample of controls for target hospitals. Column (3) shows characteristics of target hospitals. Column (4) shows the characteristics of matched controls for acquirer hospitals. Column (5) shows the characteristics of acquirer hospitals. Panel A shows the samples used for estimation for non-PPI products and Panel B shows the samples used for estimation for PPIs. Matching is at the hospital-UMDNS level, so N of matched samples is the superset of controls used in each class-merger type, and variable means are weighted the same as each hospital's weight in the pooled regressions. Data on beds, full time equivalent employees (FTEs), technologies, admissions, teaching status, non-profit status, number of HMO contracts, and Medicare and Medicaid share come from the AHA Annual Survey. Following Acemoglu and Finkelstein (2008) and Cooper et al. (2019), we measure technologies using the complete list of binary facility indicators available in the AHA. Output price is calculated using data from the CMS HCRIS and Medicare Impact Files as in Dafny et al. (2019).

on the hospital's following characteristics as in Dranove and Lindrooth (2003): inputs and outputs (log admissions, log full-time equivalent (FTE) employment, log technologies, number of unique products purchased, and average monthly purchase quantity); number of beds; payer mix (Medicare and Medicaid share of discharges, number of HMO contracts); teaching hospital status; and non-profit ownership. The weighted average characteristics of the matched samples are included in columns (2) and (4) of Table 3. The matched samples for both the target and acquirer samples are closer on most observable dimensions within both PPIs and non-PPIs.

In implementing the preferred specification, we generate a dataset containing a copy of each transaction for each of the 10 neighbors along with the full set of data from each treated hospital. Each of the 10 neighbors is therefore weighted equally in specifications (1) and (2), though some control hospitals are used as a comparison for multiple treated hospitals. For the stacked class-level regressions, matching is performed within each product category.

As discussed in detail in Dafny (2009), we note that this reduced form identification approach cannot address endogenous selection of hospitals into the merger "treatment" on unobserved dimensions. In order to provide greater confidence that our results are not driven by differential trends across merging and control hospitals, we augment our results with detailed monthly event studies with different pre- and post-merger time horizons. The results are reassuring as to our main conclusions, and to the extent that endogeneity bias remains, it must be due to time-varying factors that are precisely contemporaneous with the mergers in our sample.

Lastly, we only observe mergers which were proposed and consummated. Implicitly, this subset of all potential mergers that might take place was deemed to have lower anti-competitive effects by antitrust enforcement agencies. In the event that cost savings were used as a justification for these mergers, the cost savings we estimate are likely an upper bound on what one might expect from the average proposed horizontal merger in this setting.

5. Estimates of merger treatment effects

■ We discuss results in three subsections, beginning with the effect of mergers on prices for each product category. We then consider the pooled effects obtained by stacking the categories into a single regression for each product class. Finally, we use triple-differences versions of the pooled regressions to evaluate treatment effect heterogeneity corresponding with various potential buyer power mechanisms.

□ **Product category-specific price effects.** In Figure 1, product categories are grouped by class (non-PPIs vs. PPIs), then ordered from top to bottom in order of increasing total expenditure in the database. We show the estimated coefficients β_u and corresponding 95% confidence intervals for specifications ($1 - hj$ fixed effects; solid markers) and ($2 - h$ fixed effects; hollow markers).

The first pattern of interest is that, for both targets (left) and acquirers (right), the within-hospital-category estimates (hollow markers) closely mimic the within-hospital-brand estimates (solid markers), and the point estimates of the two are rarely statistically significantly different. This suggests that there is not a prevalence of large changes in composition of products purchased to higher or lower dollar products post-merger. Given this, and given our concerns that any such switching might have ambiguous welfare effects, we focus most of our discussion on within-hospital-brand differences moving forward.

Next, we focus on the left panel: merger effects on *target* hospitals. Among non-PPIs (circle markers), there is no discernible pattern of savings post-merger. Point estimates are a near equal mix of positive and negative, with most small in magnitude and not statistically different from zero. Exceptions include significant price decreases within-hospital-brand of 3% and 15%, respectively, for trauma bone plates and polymeric mesh.

In contrast, among PPIs (triangle markers), the majority of within-hospital-brand price effect point estimates are negative. Also, several of these effects are significant at the category level, including decreases on the order of 5% to 9% for cardioverter defibrillators, femoral knee prostheses, mammary prostheses, and pacemakers.

We observe a dramatically different pattern for *acquirers* in the right panel of Figure 1. Price effects are generally more precise, as expected given the larger sample of acquirers. The point estimates are also clustered much closer to zero for both non-PPI products and PPI products. Among non-PPI products, we observe several negative and significant results on the order of 2% to 4% for bone implant putty, drill bits, spinal bone plates, trauma bone plates, and trauma bone screws. Second, in contrast to the target results, the coefficient estimates for some PPIs are positive, ranging from 2% to 5% for acetabular hip prostheses, cardioverter defibrillators, drug eluting stents, and spinal stimulators.

FIGURE 1

MERGER TREATMENT EFFECTS [Color figure can be viewed at wileyonlinelibrary.com]

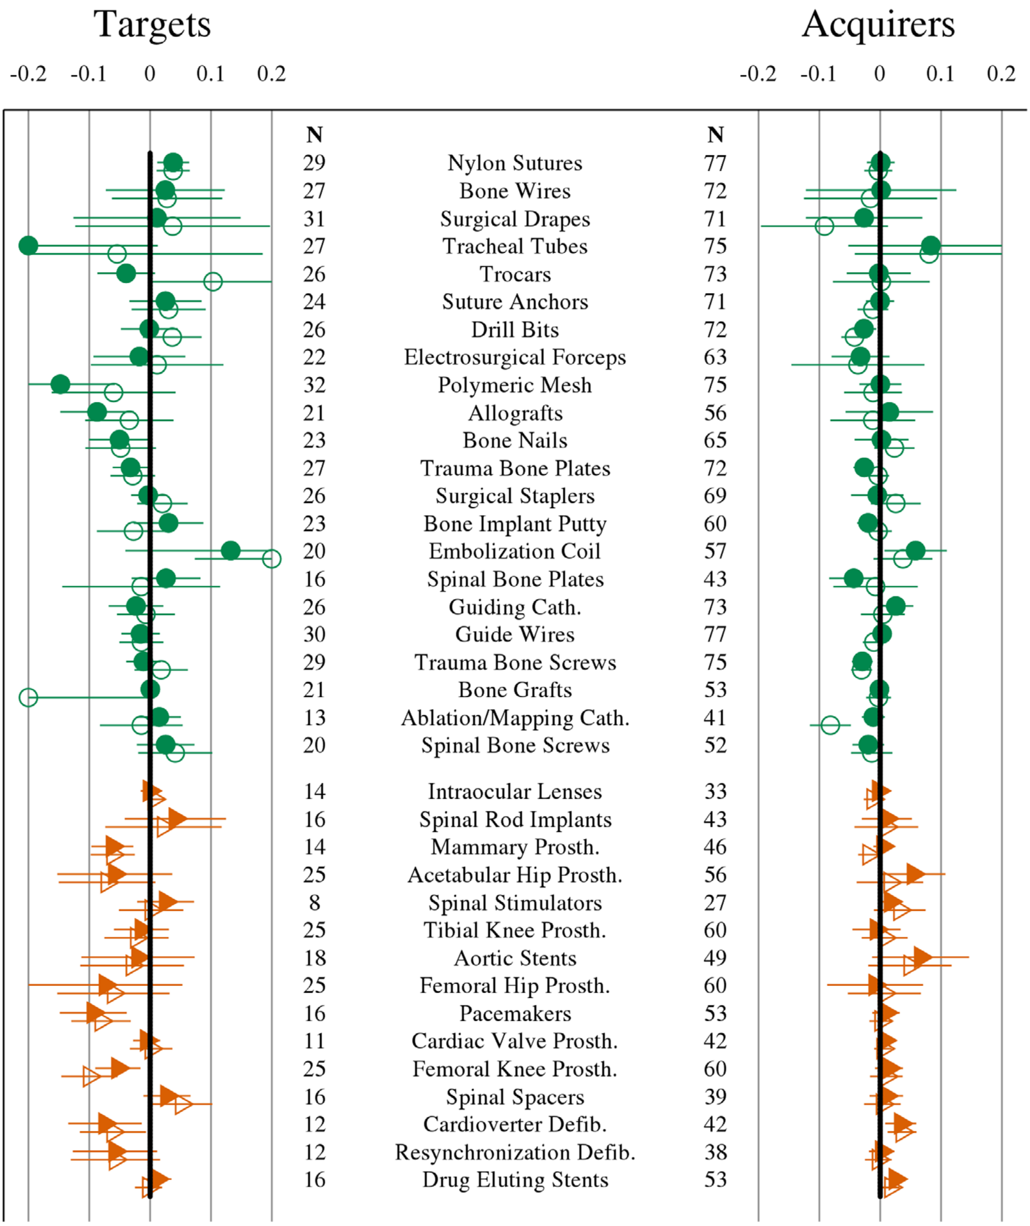

Notes: Regression coefficients from specifications (1) and (2), post-merger year $\tau_h + 1$ only. Authors' calculations from PriceGuide data. Bars indicate 95% confidence interval with standard errors clustered at hospital-brand level. Left panel: Targets. Right panel: Acquirers. Circular/green markers: non-PPIs. Triangular/orange markers: PPIs. Solid markers: specification (1), within-brand price effects. Hollow markers: specification (2), within-category price effects.

Finally, it bears noting that the variation in treatment effects documented in Figure 1 is not driven by the latent relationship between savings *potential* and savings *achieved*, either across or within product class. As shown in online web Appendix Figure A3, the relationship between treatment effects and ex ante potential savings as measured by the Gini coefficient is quite flat. One exception is a negative relationship for targets' purchase of non-PPIs, which is entirely driven by the (large but insignificant) effect for tracheal tubes in Figure 1.

TABLE 4 Merger Treatment Effects—Pooled

Dependent Variable:	$\ln(Price)_{uhjmy}$			
Non-PPIs	−0.006	0.003	−0.004	−0.012**
	(0.008)	(0.011)	(0.004)	(0.005)
PPIs	−0.034†	−0.038†	0.017†	0.012**
	(0.010)	(0.009)	(0.006)	(0.005)
Fixed effects:	$\theta_{uhj} + \theta_{jmy}$	$\theta_{uh} + \theta_{jmy}$	$\theta_{uhj} + \theta_{jmy}$	$\theta_{uh} + \theta_{jmy}$
Treatment:	Targets		Acquirers	

Notes: Authors' calculations from PriceGuide data. $*p < 0.10$, $**p < 0.05$, $\dagger p < 0.01$. Standard errors clustered at the hospital-brand level in parentheses. Coefficients estimated from pooled specifications (1) and (2). The dependent variable ln(Price) is the logged transaction price measured at the hospital-brand-month-year. All price specifications include brand-month-year fixed effects.

□ **Pooled product class price effects.** Given the large number of coefficients estimated across the individual product categories, it is useful to turn to the stacked regressions presented in Table 4 in order to shed light on average patterns at the hospital level. The left columns of Table 4 show pooled coefficient estimates for target hospitals, for each specification and class. These results indicate that targets obtain no significant price decreases on non-PPI product categories post-merger. However, they obtain more meaningful within-hospital-brand savings of 3.4% on PPIs. Finally, the within-hospital-brand coefficient is only slightly smaller than the within-hospital-category coefficient. This indicates that, on average, nearly all savings can be accounted for by renegotiations, rather than brand switching. We have also run specifications examining changes in product usage patterns as well as prices. Unfortunately, what can be done on usage is in part limited by the fact that for most of our mergers, we have detailed purchasing data for either the target or acquirer, not both, so we cannot examine "convergence" between merging parties with any precision.[20]

The pooled acquirer results are summarized in the right columns of Table 4. The non-PPI coefficients are again fairly precise zeros.[21] Prices go up slightly (1.7% within-hospital-brand, 1.2% within-hospital-category) post-merger for acquirers' purchase of PPIs. This result is interesting because, although there are several managerial and economic theories via which mergers might increase input prices (relative to non-merging control hospital trends), we might expect most of these mechanisms to perhaps be less prevalent among acquirers, who are typically the larger (sometimes significantly larger) or more dominant entity involved in the merger. We return to how we interpret this result as we examine event study evidence, robustness to matching and inference decisions, and treatment effect heterogeneity.

Event studies of merger treatment timing. We next examine whether the results above are (1) driven by preexisting differential trends in prices among merging facilities; or (2) biased due to merger effects that develop slowly over time (e.g., due to fixed contracts that take time to renegotiate as in Grennan and Swanson (2020)). In Figure 2 below, the left panels show pooled raw average price trends, controlling only for brand-category fixed effects to account for the fact that hospitals may use very different amounts of various products. The right panels show the pooled event studies for the within-brand version of the above difference-in-differences specification, fully controlling for hospital-brand effects and brand-specific time trends. We focus on PPI prices for targets (top panels) and acquirers (bottom panels); the analogous results for non-PPIs are in online web Appendix Figure A4. In each panel, we show one full calendar year

[20] We explore how standardization mediates price effects in Section 5. In Appendix C, we examine the extent to which there are merger effects on "standardization" of purchasing with a single supplier and find that they are noisy and sensitive to specification.

[21] The results are still small, though statistically significant, if we exclude the "outlier" categories observed in Figure 1: tracheal tubes, polymeric mesh, and embolization coils. The within-hospital-brand effect increases in magnitude to -0.9%, the within-hospital-category effect to -1.5%, and each are significant at the 5% level.

FIGURE 2

MERGER TREATMENT EFFECTS—EVENT STUDIES, PPIS [Color figure can be viewed at wileyonlinelibrary.com]

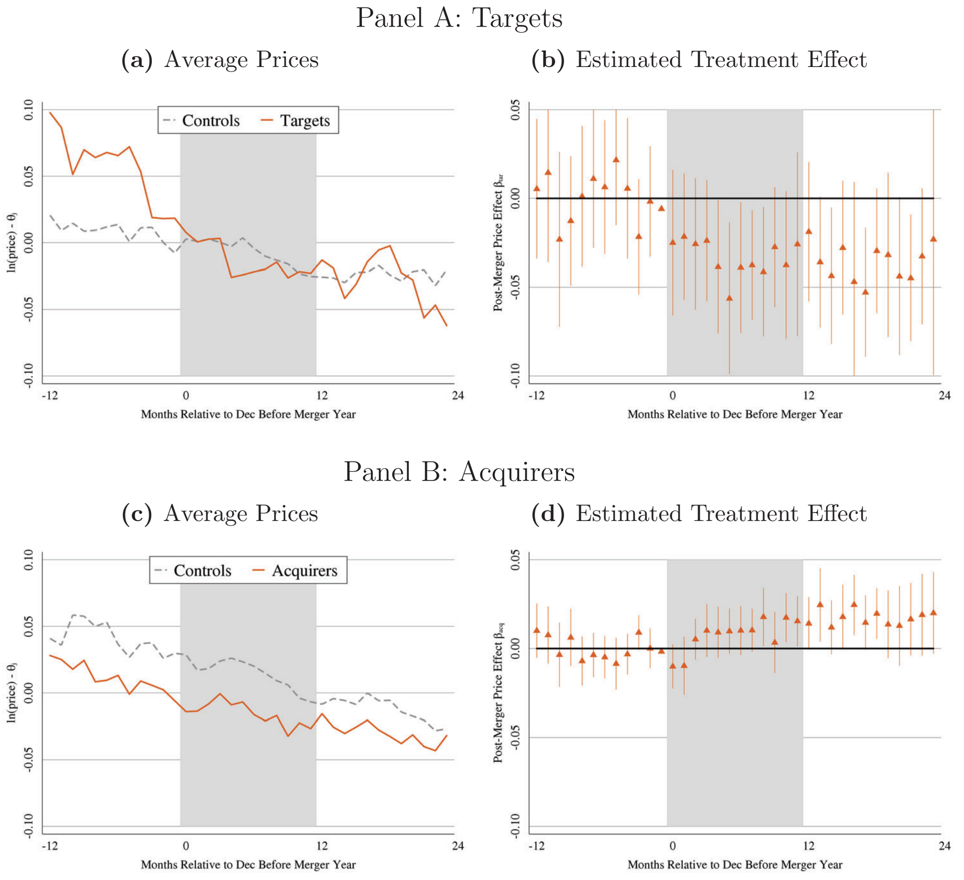

Notes: Authors' calculations from PriceGuide data. The left panels present the raw average price for treated hospitals and matched controls, adjusted for the composition of products using a product-category-brand fixed-effect. The right panels present regression coefficients from pooled event study version of specifications (1), each month within 1 year of merger year τ_h. Hold-out date is December of last pre-merger year; all coefficients represented relative to pre-merger year mean. Bars indicate the 95% confidence interval with standard errors clustered by hospital-brand.

pre- and post-merger; the year of merger is highlighted in gray. The first and second panels of online web Appendix Table A15 compare our baseline difference-in-differences results, in which the treatment effect of merging is identified by comparing the post-merger year $\tau_h + 1$ to the pre-merger year $\tau_h - 1$, to an alternative set of estimates comparing τ_h and $\tau_h + 1$ to $\tau_h - 1$. Intuitively, the latter imposes $\alpha = \beta$ in specification (1). The point estimates are slightly smaller in magnitude in the specification with $\alpha = \beta$, indicating that most, but not all, of the treatment effect of interest is realized in the merger year.

As expected in markets with evolving technology and new product entry, PPI prices are decreasing for both targets and acquirers. In panel (a), it appears that targets have a steeper negative trend pre-merger than their matched controls. However, this may be driven by any number of features that differ between targets and controls, such as different patterns in when expensive brands are purchased throughout the year. Indeed, in panel (b), which controls for such compositional differences, there is little evidence of remaining pre-trends in our preferred specification, and there is strong evidence that targets' PPI prices decrease more steeply in the merger and post-merger year. For acquirers, we observe in panels (c) and (d) that acquirers' and control hospitals' prices are on a parallel downward trend in the pre-merger year, consistent with our identifying assumption. Interestingly, it appears that the positive price effects for acquirers purchasing PPIs are driven by a slightly flatter trend among acquirers in the merger year than we observe for non-merging controls. Finally, in both panels (b) and (d), although the point estimates suggest that savings begin accruing early in the merger year, the standard errors are such that we cannot rule out zero price effects for the first five to eight months of the merger year. This could be due to price effects being realized slowly over time, or to mergers being consummated during the later part of the grayed-out "merger year." Unfortunately, we do not observe the exact timing of the merger. However, the trend in treatment effect estimates is flat in the post-merger year, indicating that, where merger effects exist, they are not continuing to evolve at the end of the time horizon observed. In event studies using alternative time horizons (online web Appendix Figures A5 and A6, discussed in more detail below), we find some evidence of differences in the pattern of merger-year monthly treatment effects, but a similar qualitative and quantitative pattern of treatment effect point estimates in the post-merger year. This emphasizes that it is more appropriate to focus on the post-merger year than the merger year.

In all of these results, we have focused on the largest sample for which we have complete pre- and post-merger years: a panel of treated and matched control hospitals in the years $\{\tau_h - 1, \tau_h, \tau_h + 1\}$. These results are, for the most part, robust to estimation on (smaller) samples of hospitals with longer pre- and post-merger periods. First, online web Appendix Figure A5 confirms that there are no differential pre-trends in targets' prices in the 2 years prior to the merger year; online web Appendix Figure A6 confirms analogously that there is no evidence of differential pre-trends in acquirers' prices for PPIs. There is a stronger negative pre-trend in acquirers' non-PPI prices, but it appears to be contained within the year $\tau_h - 2$ and would not be a source of bias in our main specifications.

Online web Appendix Figures A5 and A6 also examine whether price effects are continuing to evolve after $\tau_h + 1$. The strongest evidence of this is in panel (b) of online web Appendix Figure A5, in which targets' non-PPI prices exhibit some larger negative point estimates 2 years after the merger.

We summarize the estimated treatment effects with alternative timing supports in online web Appendix Table A15. Although the results are qualitatively similar to our baseline results, we note a few key differences. First, the subsample of targets for which we observe two pre-merger years exhibits larger non-PPI and PPI savings. Second, neither of our subsamples of acquirers with extended timing support shows evidence of positive, significant PPI price effects; thus, we interpret this result with some caution.

□ Robustness.

Quantity effects. In this section, we consider "quantity" effects that may have implications for welfare or for the validity of our research design. The results are briefly summarized here and shown in detail in online web Appendix G.

First, an important issue for thinking about the welfare effects of mergers is whether they impact the composition of products purchased. To investigate this issue, we estimate whether and how merging hospitals changed the set of brands they purchased around the time of acquisition, relative to a set of matched controls. In contrast to the theory that merging would lead to greater switching (due to standardization or to some other alteration of procurement practices), we find that brand switching is significantly *less likely* for merging hospitals than for matched controls, with similar results for targets and acquirers. This would be consistent with mergers disrupting the regular procurement activities of merging firms, but other mechanisms may also be at work here. Taken together with our finding in Figure 1 and Table 4 that the treatment effects of merging on prices within hospital-category are very similar to the treatment effects within hospital-brand, these results suggest that, although we do observe a different pattern of brand switching post-merger among treated hospitals (relative to controls), this pattern does not significantly change the composition of brands purchased in favor of higher or lower dollar products post-merger. Thus, switching patterns have little implication for savings via switching to cheaper brands, but they may be related to documented price effects through the threat of switching or propensity to renegotiate. We also caution that merging hospitals' relative lack of switching could have positive or negative effects on the quality of products purchased.

Second, we examined the effect of mergers on purchase volume, motivated by the logic that targets may have enjoyed price reductions due to increases in their quantity purchased rather than changes in negotiated prices.[22] In the online web Appendix, we report the results of differences-in-differences specifications with $ln(Q_{uhy})$ as the dependent variable. The summary results mostly suggest that there is no consistent effect of mergers on volume of supplies purchased. Importantly, comparing online web Appendix Figure A2 to Figure 1, we do not find evidence that savings are correlated with quantity effects.

Lastly, we investigate whether there is nonrandom attrition in our sample. For example, if merging hospitals are more likely to disappear from our sample after a merger takes place, then the price treatment effects we estimate may be systematically biased upward or downward. Similarly, mergers could impact the attribution of transactions to specific facilities; for example, a target facility's transactions might be attributed to the acquiring system due to centralized reporting post-merger. We find that there is some differential attrition of targets. However, as discussed in online web Appendix G, this is greatly driven by targets merging at the end of our sample, and dropping those targets has no effect on our documented pattern of price effects. We speculate that this differential attrition may be driven by a longer delay in data submission for targets post-merger.

Sensitivity of price treatment effects. We have examined the sensitivity of our results to several decisions regarding modeling, regression sample, and inference. In online web Appendix Table A5, we present estimates from specifications (1) and (2) using different matching approaches. Panel A presents the baseline estimates for reference. Panel B presents the non-matched results, using all non-merging hospitals as controls. Panel C uses a 10 neighbor Probit version of the match as in Dranove and Lindrooth (2003). Panel D uses a 1-to-1 Mahalanobis match as in Schmitt (2017) – these results are the most notable in that all merger effect estimates are significantly noisier, with the target PPI savings no longer statistically significant. The alternative matching approaches generally track our preferred estimates, with largest savings for targets' purchase of PPIs and positive treatment effects for acquirers' purchase of PPIs. However, none of the estimates is statistically significantly different from those in our main results, indicating that observed compositional differences do not generate large differences in input price or trends between treated (merging) and control (non-merging) hospitals.

Next, although the matching exercises described above focus on selecting the best comparison groups for our in-sample mergers to ensure internal validity, they do not address external validity: our data only include hospitals that voluntarily joined a benchmarking database, which may be different in observable and unobservable ways from the average merging hospital. In online web Appendix E, we estimate our main merger specification from equation (1), with sample treated hospitals re-weighted to be representative of the distribution of the full sample of targets and acquirers in the AHA based on (a) bed size, or (b) ownership and teaching status. These results are qualitatively similar, with point estimates that are slightly smaller in magnitude. We also investigated this issue using an alternative source of hospital cost data—total cost per adjusted discharge in the HCRIS data reported by hospitals to CMS (as in Schmitt (2017))—to investigate how cost treatment effects change as we impose the sample limitations that lead us to our final analytic sample. The results are shown in Appendix Table A17. We do not observe

[22] We thank an anonymous referee for this comment.

significant changes in the merger treatment effects as we narrow the sample from all mergers in 2009–2015 to mergers in 2009–2015 among hospitals in the PriceGuide data, to our analytic sample of PriceGuide hospitals with "full support" merging in 2009–2015.[23] However, we note that this is not a particularly high-powered test due to the large standard errors. For example, the point estimate of the treatment effect of mergers on costs for targets in the post-merger year is -0.011 in the full sample and in the PriceGuide, full support sample, but the standard error for the latter treatment effect is 0.046.

We also attempt to directly address any potential confounding of merger effects and database join effects. Our baseline analyses contain a dummy variable to indicate month-years after the hospital joins the benchmarking database, so that join effects are not conflated with merger effects. Online web Appendix Table A14 shows a slightly cleaner specification, estimated only on hospitals whose three focal periods $\{\tau_h - 1, \tau_h, \tau_h + 1\}$ are entirely post-join.[24] The results are slightly larger in absolute magnitude, but confirm our main findings.

Next, we address potential bias introduced by hospitals' involvement in multiple mergers. Our main specification identifies the first merger for each of our treated hospitals over the sample period of 2009–2015. At baseline, we impose that treated hospitals have no merger in $\tau_h - 1$, and that matched control hospitals have no merger in $\{\tau_h - 1, \tau_h, \tau_h + 1\}$. In Panel C of online web Appendix Table A14, we implement a stricter version of this restriction, ensuring that no mergers occur between $\tau_h - 2$ and $\tau_h + 1$ except for the focal merger in τ_h, applying this rule to both treated and matched control hospitals. The results are qualitatively similar to our main estimates in Table 4.

Lastly, online web Appendix Table A16 explores various alternative approaches to standard errors: a wild bootstrap method as well as alternative clustering at the hospital-vendor and system-UMDNS levels. Our main findings are stable across approaches to standard errors.

□ **Price treatment effect heterogeneity and mechanisms.** In this section, we examine heterogeneity in treatment effects along several dimensions in order to explore mechanisms. This is intended not to be an exhaustive exploration, but rather to shed light on those mechanisms highlighted in Section 3 for which we have relevant data.[25] For the sake of brevity, we continue to focus discussion on within-hospital-brand price effects, as our previous results indicated that these were where the strongest evidence of merger-driven savings were concentrated. Within-hospital-category results are available in online web Appendix Table A18.

Size effects. As noted previously, much of the literature regarding mergers and cost savings focuses on advantages associated with firm size. Within our sample, we observe substantial variation in the (absolute and relative) buyer firm size change induced by the merger: with one exception, all of our transactions involve 1–2 target hospitals, but our acquirer systems range from very small (1 or 2 hospitals) to large (over 70 hospitals). The effect of target and acquirer size on purchasing is *ex ante* ambiguous. Theories such as that of Chipty and Snyder (1999) and others would predict that—if the surplus function is concave—we should see larger effects when a merger entails a larger change in the size of a hospital system. This would predict the largest effects on targets' prices when acquirers are large. On the other hand, price decreases may be driven by improved management practices, and there may be economies or diseconomies of scale in sharing management between merging hospitals (Beckmann, 1960).

[23] Online web Appendix Table A17 presents results for 434 targets and 964 acquirers; Table 1 column (1) has 661 targets and 1,753 acquirers. This difference is due to the fact that we follow the approach in our price regressions of limiting the regression sample to the *first* transaction observed for each target and acquirer.

[24] We also remove matched controls when their associated treated hospital is removed from the data based on this restriction.

[25] There are numerous interesting questions for which we do not have relevant data, for example, regarding the supplier's cost function, regarding convergence of target and acquirer prices for a given merging pair, etc.

The top two rows in each panel of Table 5 show separate results for mergers involving small (1–3 hospitals) versus large (4+ hospitals) acquirers.[26] For both targets and acquirers, point estimates of merger price effects for non-PPIs are small and negative (2.1% for targets and 0.6% for acquirers) when the merger involves small acquirers. The positive treatment effect previously documented for acquirers' PPI prices appears to be driven by large acquirers. The savings on PPIs for targets is slightly larger for large acquirer mergers (average acquirer system size of 41.8 hospitals) than for small acquirer mergers (average size 1.8 hospitals). These point estimates are consistent with mergers involving countervailing effects of improved buyer power and managerial disruption. The net effect is small and negative for merging parties with the largest relative system size change (small targets of large acquirers), but small and positive for merging parties with the smallest relative system size change (large acquirers of small targets). However, these differences are not statistically significant.

Geographic proximity. Next, as noted in Schmitt (2017), many of the mergers in the recent "great reconsolidation" involve hospital systems acquiring hospitals in distant geographic markets. We next split the treatment effects according to whether any of the merging hospitals share a hospital referral region (HRR). (See online web Appendix F for alternative market definitions.) Heterogeneity in merger effects by market overlap may be due to local economies of scale in management or distribution of inputs, to local diffusion of management practices, to the countervailing effects of changes in upstream and downstream market power, or to the relative roles of bargaining power versus bargaining position in mediating merger-related cost savings.

We compare treatment effects for in- versus out-of-market mergers in the second pair of rows in each panel of Table 5. The strongest merger savings previously documented—for targets' purchase of PPIs—are concentrated in in-market mergers, where we see a price decrease post-merger of 6%, relative to the control trend. Targets also achieve larger reductions on non-PPIs when there is market overlap (2.5% savings, versus price increases of 1.1% for out of HRR mergers). Acquirers show price increases for both in- and out-of-market mergers; point estimates are smaller, but not significantly so, for in-market mergers. These results stand in contrast to the large out-of-market merger effects documented in Schmitt (2017), in which merger effects were strongest for *targets* in out-of-market acquisitions, perhaps due to the differing nature of the marginal costs of inputs in our purchasing data from hospital costs more broadly construed. Instead, they echo Dranove and Lindrooth (2003), in which cost savings are greatest when previously independent hospitals integrate under a single license and consolidate facilities. They are also consistent with theories of concavity and economies of scale, given the qualitative fact that some non-PPIs and almost all PPIs tend to be sold by highly specialized, regional sales representatives who spend large amounts of time with a few local accounts. Finally, our results are consistent with Farrell and Shapiro (2001)'s argument that the agencies should give consideration to "efficiencies based upon the close integration of specific, hard-to-trade assets owned by the merging parties," while noting that "the same conditions that tend to make synergies more merger-specific and more beneficial to consumers also tend to make the merger itself more problematic." That is, we find evidence of greater savings associated with local mergers; unfortunately, Cooper et al. (2019) and others also find evidence of greater anticompetitive effects of local hospital mergers in the downstream markets for hospital services.

Supplier market structure. We also examine whether merger effects are mediated by supply-side market structure. To this end, we separate UMDNS codes within each product class into those above or below the median HHI for the class. As noted in Table 2, the product categories analyzed in this article are almost all moderate-high concentration according to typical FTC and DoJ standards. That said, the mean "High HHI" non-PPI has an HHI of 0.419, versus 0.179 among the "Low HHI" non-PPIs; the same measures among PPIs are 0.497 and 0.227, respectively. The

[26] Our average sample target is acquired by a 26-hospital system; our average sample acquirer is part of an 11-hospital system prior to the focal merger.

TABLE 5 Merger Treatment Effects—Heterogeneity, Within Brand

	Targets			Acquirers		
	N_{tar}	β	SE	N_{acq}	β	SE
Panel A: Non-PPIs						
Acquirer size						
Small	13	−0.021*	(0.012)	26	−0.006	(0.006)
Large	20	0.007	(0.012)	59	−0.000	(0.006)
Difference		0.028	(0.019)		0.006	(0.009)
Market exposure						
In HRR	14	−0.025**	(0.011)	36	−0.007	(0.005)
Out of HRR	19	0.011	(0.012)	49	0.003	(0.007)
Difference		−0.036**	(0.018)		−0.010	(0.008)
Vendor market structure						
High HHI	33	−0.004	(0.012)	85	−0.004	(0.006)
Low HHI	33	−0.008	(0.010)	85	−0.004	(0.006)
Difference		0.003	(0.016)		0.001	(0.008)
Controlling for output price						
Post-Merger	33	0.000	(0.008)	85	−0.005	(0.004)
ln(Output Price)		0.023	(0.018)		0.003	(0.005)
Standardization interaction						
Post-Merger	30	−0.013	(0.009)	80	0.000	(0.005)
Post X Std.		0.015	(0.016)		−0.009	(0.008)
Panel B: PPIs						
Acquirer size						
Small	12	−0.023	(0.015)	26	0.009	(0.006)
Large	17	−0.043**	(0.018)	48	0.026**	(0.010)
Difference		−0.021	(0.026)		0.016	(0.012)
Market exposure						
In HRR	12	−0.059†	(0.018)	35	0.014**	(0.007)
Out of HRR	17	−0.012	(0.011)	39	0.023**	(0.009)
Difference		−0.047**	(0.021)		−0.008	(0.010)
Vendor market structure						
High HHI	29	−0.052†	(0.012)	74	0.011**	(0.004)
Low HHI	29	−0.027*	(0.014)	74	0.021**	(0.009)
Difference		−0.025	(0.019)		−0.009	(0.010)
Controlling for output price						
Post-Merger	29	−0.037†	(0.011)	74	0.018†	(0.006)
ln(Output Price)		0.034*	(0.018)		0.020†	(0.008)
Standardization interaction						
Post-Merger	28	−0.054**	(0.024)	65	0.007	(0.009)
Post X Std.		0.031	(0.027)		0.016	(0.012)

Notes: Authors' calculations from PriceGuide data. *$p < 0.10$, **$p < 0.05$, †$p < 0.01$. Standard errors clustered at the hospital-brand level in parentheses. Coefficients estimated from pooled specification (1). The dependent variable is the logged transaction price measured at the hospital-brand-month-year. Small acquirers are hospital systems consisting of 1-3 hospitals pre-merger, and large acquirers are hospital systems with more than 3 hospitals. A target is categorized as "In HRR" if there is at least one hospital in the acquiring system in the same HRR, and vice versa. A product category is classified as "High" concentration if its vendor HHI is above the median within its product class. ln(Output Price) is estimated using the HCRIS as in Dafny et al. (2019). Standardization is an indicator for whether the hospital purchased at least 75% of all units in a product category from a single vendor in its first sample year.

third pair of rows in each panel of Table 5 show that there is no economically or statistically significant difference in price effects as a function of supplier competition.[27]

Downstream hospital–insurer market power. The fourth pair of rows in each panel of Table 5 examines whether the cost effects documented above are muted due to mergers causing hospitals' supply side and demand side market power to increase concurrently. For example, if merger-enabled market power allowed hospitals to exercise monopoly power and increase procedure prices, then most bargaining models would predict that some of that pie could be shared with suppliers, mitigating cost decreases due to increased monopsony power (e.g., see the discussion of insurer-hospital bargaining in Ho and Lee (2017)). To that end, we estimate our same input price regression specifications, *controlling for* output prices.[28] We employ the method described in Dafny et al. (2019) to infer hospital prices from HCRIS reports. The results indicate that, although hospitals' downstream price changes tend to be positively correlated with upstream price changes, this does not change the estimated merger treatment effect.

Standardization and renegotiation. The final set of rows in each panel of Table 5 examines the interaction between merger effects and standardization. We estimate a simple modification of the above specifications, in which the year-of and post-merger dummies are interacted with a dummy for pre-merger standardization at the hospital-category-level. That is, this specification indicates whether merger-induced savings are larger for hospitals that were standardized pre-merger.

The results confirm our previous result that targets receive savings on PPIs after merging. This correlation is significant for acquirers' non-PPI prices, and marginally significant for targets' PPI prices. However, the merger price effect is not significantly amplified for hospital-categories that are standardized, for any combination of product class and type of merging entity. For targets, standardization appears to diminish post-merger savings, if anything.

6. Conclusion

■ The US hospital industry has experienced a large amount of contentious consolidation via mergers over the last several decades. Marginal cost savings have been perhaps the most common justification offered for these mergers, often appealing to the large input price variation across hospitals and notions that "buyer power" is increasing in hospital system size. Prior research examining aggregated accounting measures of hospital costs has found mixed results.

In this study, we use data on all purchase orders issued by a large set of US hospitals, 2009–2015, in order to conduct a detailed examination of the effects of mergers on the prices paid for medical/surgical supplies, an important component of hospital marginal costs. The most robust finding is target savings of 3.4% on targets' purchase of physician preference items. Across our 37 product categories, targets save an estimated \$214,402 per year (1.9%) due to within-brand price decreases after horizontal mergers, whereas acquirers experience an (insignificant) average net price increase of \$90,243.[29] Perhaps the simplest way to summarize these findings is that, *given the precision of our estimates, we can rule out average input price savings of greater than 3.1% at the 95% level for both targets and acquirers*. This seems modest relative to the cross-sectional price variation across hospitals and claims of potential savings via increased"buyer power."

[27] Online web Appendix Figure A7 presents this in richer detail, in the form of a scatterplot of the treatment effect estimates from Figure 1 versus vendor HHI for each UMDNS code. Each scatterplot is essentially a cloud, with no clear relationship between merger price effect and vendor HHI.

[28] The goal of this regression is simply to test for the "mitigation" effect described above. In a model like Ho and Lee (2017), input and output prices will be codetermined and thus we would need to use extreme care in interpreting any parameter estimates in this regression.

[29] Calculation details in online web Appendix H. For comparison, a recent AHA-sponsored study documented a decrease in operating expenses of 2.5% for acquired hospitals (Noether and May, 2017); our estimate is lower, though our 95% confidence interval would include 2.5% .

The variety of product categories in the data allows us to look more closely at merger effects and examine mechanisms underlying "buyer power" (which has previously been studied in theory and in case studies of specific product markets). We examine heterogeneity in merger treatment effects across different product categories, and by acquirer size, market overlap, and vendor market concentration. We find that the observed target savings on PPIs is driven by local mergers. These savings may be consistent with local returns to scale in sales and distribution or transfer of managerial practices. Merger treatment effects on targets are also larger when acquirers are larger, consistent with savings driven by concavity in the surplus function as in Chipty and Snyder (1999), though the size comparison is not statistically significant. These findings are echoed in the results for acquirers' purchase of PPIs, in which price increases are smallest for small acquirers (where the relative size increase is larger) and for local mergers. Although there are multiple factors that may drive cost increases after a merger—e.g., managerial attention—the countervailing force of increased buyer power is most powerful for local mergers involving larger relative size changes. Our remaining explorations suggest that our average treatment effects are not obscuring a great deal of heterogeneity: merger effects on marginal costs are small regardless of standardization, supplier concentration, or downstream prices.

Antitrust agencies consider a merger's "efficiencies" to be cognizable if they are likely to occur if the merger proceeds, and unlikely to occur if it does not. The agencies also ask whether efficiencies are large and/or likely to pass through to consumers (Farrell and Shapiro, 2001). We have limited ability in our data to speak to the merger-specificity of the savings we document, or to potential pass-through. However, whether cognizable or not, our estimates of post-merger savings are small, indicating little effect of mergers on buyer power. Moreover, estimated savings are largest for local mergers where hospitals' market power vis à vis insurers is also likely to increase (Cooper et al., 2019). Finally, the largest estimated savings, by targets on PPIs, can entirely be attributed to renegotiation, rather than brand switching, in that savings estimates within hospital-brand are statistically equivalent to estimates within hospital-category. This transfer of surplus from device manufacturers to hospitals is suggestive of increased monopsony power and may not increase efficiency. For example, it may negatively impact dynamic incentives of suppliers to innovate or maintain product quality or manufacturing reliability (see discussion in Hemphill and Rose (2018)). Each proposed merger should certainly be judged on its own merits, given its specific context. However, each of these features of our findings urges caution regarding the use of expected hospital purchasing efficiencies as justification for horizontal hospital mergers.

We offer these and all results with the caveat that our sample size of mergers is smaller than we would like due to the relative newness of detailed purchasing order data availability. Another drawback of our data is that we do not observe which products are purchased through group purchasing organizations, which are an important feature of the setting that may mediate the effects of mergers for some products. However, our data cover a larger sample than that of many merger retrospective case studies from which economists have learned a great deal (e.g., Miller and Weinberg (2017)). We also believe the detail and breadth of the purchasing data brings new light to the study of hospitals and buyer power broadly, and mergers specifically.

For hospital mergers in particular, another important phenomenon to consider is the simultaneity of input market negotiation and output market negotiation. We control for this using a proxy for hospital output prices. However, a more detailed study would require matching hospital purchasing data with private insurer claims, and modeling demand and negotiated prices explicitly in both upstream and downstream markets. We see this as an important area for future research.

References

Acemoglu, D. and Finkelstein, A. "Input and Technology Choices in Regulated Industries: Evidence from the Health Care Sector." *Journal of Political Economy*, Vol. 116(5) (2008), pp. 837–880.

Agrawal, A., Jaffe, J. F. and Mandelker, G. N. "The Post-Merger Performance of Acquiring Firms: A Re-Examination of an Anomaly." *Journal of Finance*, Vol. XLVII(4) (1992), pp. 1605–1621.

Beckert, W., Christensen, M. and Collyer, K. "Choice of NHS-Funded Hospital Services in England." *The Economic Journal*, 122(560) (2012).

Beckmann, M. J. "Some Aspects of Returns to Scale in Business Administration." *Quarterly Journal of Economics*, Vol. 74(3) (1960), pp. 464–471.

Bloom, N., Sadun, R.and Van Reenen, J. "Does Management Matter in Healthcare?" *Stanford Mimeo*, 2014.

Capps, C. "The Quality Effects of Hospital Mergers." *Economic Analysis Group Discussion Paper 5-6*, 2005.

Capps, C. and Dranove, D. "Hospital Consolidation and Negotiated PPO Prices." *Health Affairs*, Vol. 23(2) (2004), pp. 87–109.

Capps, C., Dranove, D. and Satterthwaite, M. "Competition and Market Power in Option Demand Markets." *RAND Journal of Economics*, Vol. 34(4) (2003), pp. 737–763.

Carlton, D. W. and Israel, M. "Proper Treatment of Buyer Power in Merger Review." *Review of Industrial Organization*, Vol. 39(1/2) (2011), pp. 127–136.

Chipty, T. and Snyder, C. M. "The Role of Firm Size in Bilateral Bargaining: A Study of the Cable Television Industry." *The Review of Economics and Statistics*, Vol. 81(2) (1999), pp. 326–340.

Coate, M. B. "The Merger Review Process at the Federal Trade Commission from 1989 to 2016." Available at SSRN: https://ssrn.com/abstract=2955987 or http://doi.org/10.2139/ssrn.2955987, 2018.

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

Cutler, D. M. and Scott Morton, F. "Hospitals, Market Share, and Consolidation." *JAMA*, Vol. 310(18) (2013), pp. 1964–70.

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

Dafny, L. "Hospital Industry Consolidation - Still More to Come?" *New England Journal of Medicine*, Vol. 370 (2014), pp. 198–199.

Dafny, L., Ho, K. and Lee, R. S. "The Price Effects of Cross-Market Mergers: Theory and Evidence from the Hospital Industry." *Rand Journal of Economics*, Vol. 50(2) (2019), pp. 286–326.

Dana, J. "Buyer Groups as Strategic Commitments." *Games and Economic Behavior*, Vol. 74(2) (2012), pp. 470–485.

Dranove, D. and Lindrooth, R. "Hospital Consolidation and Costs: Another Look at the Evidence." *Journal of Health Economics*, Vol. 22(6) (2003), pp. 983–997.

Duggan, M. and Scott Morton, F. "The Impact of Medicare Part D on Pharmaceutical Prices and Utilization." *American Economic Review*, Vol. 100(1) (2010), pp. 590–607.

Farrell, J. and Shapiro, C. "Scale Economies and Synergies in Horizontal Merger Analysis." *Antitrust Law Journal*, Vol. 68 (2001), pp. 685–710.

Fulop, N., Protopsaltis, G., Hutchings, A., King, A., Allen, P., Normand, C. and Walters, R. "Process and Impact of Mergers of NHS Trusts: Multicentre Case Study and Management Cost Analysis." *BMJ*, Vol. 325 (2002).

Gans, J. S. and King, S. P. "Exclusionary Contracts and Competition for Large Buyers." *International Journal of Industrial Organization*, Vol. 20(9) (2002), pp. 1363–1381.

Gaynor, M. and Town, R. J. "Competition in Health Care Markets." In McGuire, T. G., Pauly, M. V., and Barros, P., (Eds.), *Handbook of Health Economics*, volume 2, pages 499–637.Elsevier B. V., 2012.

Gooch, K. "4 of the Largest Gpos." Becker's Hospital CFO Report, 2017.

Gowrisankaran, G., Nevo, A. and Town, R. "Mergers when Prices are Negotiated: Evidence from the Hospital Industry." *The American Economic Review*, Vol. 105(1) (2015), pp. 172–203.

Grennan, M.. "Price Discrimination and Bargaining: Empirical Evidence from Medical Devices." *American Economic Review*, Vol. 103(1) (2013), pp. 145–177.

Grennan, M. "Bargaining Ability and Competitive Advantage: Empirical Evidence from Medical Devices." *Management Science*, 60(12) (2014), pp. 3011–3025.

Grennan, M. and Swanson, A. "Diagnosing Price Dispersion." Working paper, 2018.

Grennan, M. and Swanson, A. "Transparency and Negotiated Prices: The Value of Information in Hospital-Supplier Bargaining." *Journal of Political Economy*, Vol. 128(4) (2020), pp. 1234–1268.

Gruber, J. and McKnight, R. "Controlling Health Care Costs Through Limited Network Insurance Plans: Evidence from Massachusetts State Employees." *American Economic Journal: Economic Policy*, Vol. 8(2) (2016), pp. 219–50.

Haas-Wilson, D. and Garmon, C. "Hospital Mergers and Competitive Effects: Two Retrospective Analyses." *International Journal of the Economics of Business*, Vol. 18(1) (2011), pp. 17–32.

Harrison, T. D. "Do Mergers Really Reduce Costs? Evidence from Hospitals." *Economic Inquiry*, Vol. 49(4) (2010), pp. 1054–1069.

Hemphill, C. S. and Rose, N. L. "Mergers that Harm Sellers." *The Yale Law Journal*, 127(1) (2018).

Ho, K. "The Welfare Effects of Restricted Hospital Choice in the us Medical Care Market." *Journal of Applied Econometrics*, Vol. 21(7) (2006), pp. 1039–1079.

Ho, K. "Insurer-Provider Networks in the Medical Care Market." *American Economic Review*, Vol. 99(1) (2009), pp. 393–430.

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

Ho, K. and Lee, R. S. "Equilibrium Provider Networks: Bargaining and Exclusion in Health Care Markets." *American Economic Review*, vol. 109(2) (2019), pp. 473–522.

Ho, V. and Hamilton, B. H. "Hospital Mergers and Acquisitions: Does Market Consolidation Harm Patients?" *Journal of Health Economics*, Vol. 19(3) (2000), pp. 408–419.

Horn, H. and Wolinsky, A. "Bilateral Monopolies and Incentives for Merger." *RAND Journal of Economics*, Vol. 19 (1988), pp. 408–419.

Inderst, R. and Wey, C. "Buyer Power and Supplier Incentives." *European Economic Review*, Vol. 51 (2007), pp. 647–667.

Krishnan, R. "Market Restructuring and Pricing in the Hospital Industry." *Journal of Health Economics*, Vol. 20 (2001), pp. 213–237.

Kwoka, J. *Mergers, Merger Controls, and Remedies: A Retrospective Analysis of U.S. Policy*. Cambridge, MA, The MIT Press, 2015.

Lagasse, J. "Here's a Big Supply Chain Cost-Savings Opportunity: Physician Preference Items." *Healthcare Finance* (2018).

LaVito, A. "Trump Administration Opens Up More Drugs Covered by Medicare to Price Negotiations." *CNBC.com* (2018).

Lewis, M. and Pflum, K. "Diagnosing Hospital System Bargaining Power in Managed Care Networks." *American Economic Journal: Economic Policy*, Vol. 7(1) (2015), pp. 243–74.

Maeda, J. L., Raetzman, S. O. and Friedman, B. S. "What Hospital Inpatient Services Contributed the Most to The 2001-2006 Growth in the Cost Per Case?" *Health Services Research*, Vol. 47(5) (2012), pp. 1814–1835.

Marvel, H. P. and Yang, H. "Group Purchasing, Nonlinear Tariffs, and Oligopoly." *International Journal of Industrial Organization*, Vol. 26(5) (2008), pp. 1090–1105.

Miller, N. and Weinberg, M. "Understanding the Price Effects of the Millercoors Joint Venture." *Econometrica*, Vol. 85(6) (2017), pp. 1763–91.

Minemyer, P. "Geisinger's Turnaround Offers Lessons on Integrated Delivery." *FierceHealthcare* (2017).

Navathe, A. S., Troxel, A. B., Liao, J. M., Nan, N., Zhu, J., Zhong, W. and Emanuel, E. J. "Cost of Joint Replacement using Bundled Payment Models." *JAMA Internal Medicine*, Vol. 177(2) (2017), pp. 214–222.

Noether, M. and May, S. "Hospital Merger Benefits: Views from Hospital Leaders and Econometric Analysis." Technical report, Charles River Associates, 2017.

Nugent, W. C., Kilo, C. M., Ross, C. S., Marrin, A. S., Berwick, D. M. and Roessner, J. *Improving Outcomes and Reducing Costs in Adult Cardiac Surgery*. Boston, MA, Institute for Healthcare Improvement, 1999.

Prager, E. and Schmitt, M. "Employer Consolidation and Wages: Evidence from Hospitals." *American Economic Review* (forthcoming).

Romano, P. S. and Balan, D. J. "A Retrospective Analysis of the Clinical Quality Effects of Acquisition of Highland Park Hospital by Evanston Northwestern Healthcare." *International Journal of the Economics of Business*, Vol. 18(1) (2011), pp. 45–64.

Sacher, S. and Vita, M. "The Competitive Effects of a Not-for-Profit Hospital Merger: A Case Study." *Journal of Industrial Economics*, Vol. 49(1) (2001), pp. 63–84.

Schmitt, M. "Do Hospital Mergers Reduce Costs?" *Journal of Health Economics*, Vol. 52 (2017), pp. 74–94.

Schneller, E. S. "The Value of Group Purchasing - 2009: Meeting the Need for Strategic Savings." *Health Care Sector Advances, Inc.*, (2009).

Sela, R. J. and Simonoff, J. S. "Reemtree: Regression Trees with Random Effects." R package version 0.90.3, 2011.

Snyder, C. M. "A Dynamic Theory of Countervailing Power." *Rand Journal of Economics*, Vol. 27 (1996), pp. 747–769.

Snyder, C. M. "Why do Larger Buyers Pay Lower Prices? Intense Supplier Competition." *Economics Letters*, Vol. 58 (1998), pp. 205–209.

Sorensen, A. "Insurer-Hospital Bargaining: Negotiated Discounts in Post-Deregulation Connecticut." *Journal of Industrial Economics*, Vol. 51(4) (2003), pp. 469–490.

Starc, A. and Swanson, A. "Preferred Pharmacy Networks and Drug Costs." *American Economic Journal: Economic Policy*, (forthcoming).

Stole, L. and Zwiebel, J. "Intra-firm Bargaining Under Non-Binding Contracts." *Review of Economic Studies*, Vol. 63 (1996), pp. 375–410.

Tenn, S. "The Price Effects of Hospital Mergers: A Case Study of the Sutter-Summit Transaction." *International Journal of the Economics of Business*, Vol. 18(1) (2011), pp. 65–82.

Thompson, E. "The Effect of Hospital Mergers on Inpatient Prices: A Case Study of the New Hanover-Cape Fear Transaction." *International Journal of the Economics of Business*, Vol. 18(1) (2011), pp. 91–101.

Tousi, K. and Kawalek, C. "Designing for Profitability: A New Approach to Territory Alignment." *ZS Associates: Sales and Marketing Insights*, (2009).

Town, R. J. and Vistnes, G. "Hospital Competition in hmo Networks." *Journal of Health Economics*, Vol. 20(5) (2001), pp. 733–752.

Town, R. J., Wholey, D. R., Feldman, R. and Burns, L. R. "The Welfare Consequences of Hospital Mergers." National Bureau of Economic Research Working Paper 12244, 2006.

United States District Court for the District of Columbia. "Memorandum Opinion, United States of America vs. Anthem, Inc., et al." Civil Action No. 16-1493 (ABJ), 2017.

U.S. Department of Justice and the Federal Trade Commission. "Horizontal Merger Guidelines." https://www.ftc.gov/sites/default/files/attachments/merger-review/100819hmg.pdf, 2010.

Uysal, V. B., Kedia, S. and Panchapagesan, V. "Geography and Acquirer Returns." *Journal of Financial Intermediation*, Vol. 17 (2008), pp. 256–275.

Vogt, W. B. and Town, R. J. "How has Hospital Consolidation Affected the Price and Quality of Care?" Synthesis Report 9, Robert Wood Johnson Foundation, 2006.

Supporting information

Additional supporting information may be found online in the Supporting Information section at the end of the article.

Table A1: Sensitivity of Main Results to Product Granularity
Table A2: Merger Treatment E.ects . Standardization
Table A3: Merger Treatment E.ects on Standardization . Pooled, Alternate Standard Error Calculations
Table A4: Comparison of Matching Results
Table A5: Merger Treatment E.ects on Price . Pooled, Alternative Matching Approaches
Table A6: Merger Treatment E.ects . Pooled, Re-weighting for Generalizability
Table A7: Merger Treatment E.ects For Targets . Heterogeneity By Market Exposure
Table A8: $\tau + 1$ Market Share of Products Purchased in τ - 1
Table A9: Merger Treatment E.ects on Purchase Quantity
Table A10: Evidence on Di.erential Attrition of Merging Hospitals
Table A11: Main Price Regressions Excluding 2014 Mergers
Figure A2: Merger Treatment E.ects on Purchase Quantity
Table A12: Estimated Savings Using Within-Brand Merger E.ects
Table A13: Estimated Savings Using Pooled Coefficients
Table A14: Merger Treatment E.ects . Pooled, Alternative Sample Restrictions
Figure A1: Distribution of Beds by Treatment Type
Figure A3: Price Dispersion and Treatment E.ects, Across UMDNS Categories
Figure A4: Merger Treatment E.ects . Event Studies, Non-PPIs
Figure A5: Merger Treatment E.ects for Targets . Event Studies using Alternative Timing Supports
Figure A6: Merger Treatment E.ects For Acquirers . Event Studies using Alternative Timing Supports
Figure A7: Relationship Between Vendor HHI and Treatment E.ect
Table A15: Merger Treatment E.ects on Price . Pooled, Alternative Timing
Table A16: Merger Treatment E.ects on Price . Pooled, Alternate Standard Error Calculations
Table A17: Cost Regressions Using the HCRIS Measure
Table A18: Merger Treatment E.ects . Heterogeneity, Within Category