Targeting Efficiencies? Reducing Costs Through Hospital Consolidation*

Jonathan Arnold, University of Pennsylvania
Atul Gupta, University of Pennsylvania and NBER
Tong Liu, Massachusetts Institute of Technology
Alexander Olssen, University of Pennsylvania

July 2025

PRELIMINARY, DO NOT CIRCULATE

Abstract

Policymakers are concerned about the effects of consolidation in the hospital industry, much of it driven by the expansion of multi-unit hospital "systems." Several studies have shown that system acquisitions lead to net price increases at the target. Price changes occur due to a combination of at least three factors: an increase in market power, the target benefiting from the acquirer's bargaining ability, and the pass-through of cost synergies, if any. Despite the attention that consolidation has received, surprisingly little is known about the relative magnitudes of these channels. We fill this gap using novel data on transaction prices for hospital outpatient care and detailed administrative data on hospital costs from California. We first show that system ownership generates significant net price increases and cost savings for independent hospitals but not for targets that were already organized as systems. We then develop and structurally estimate a model of patient demand and hospital-insurer negotiations to quantify the different mechanisms. We find that systems target for acquisition independent hospitals with weak bargaining ability, suggesting a strategy of pricing arbitrage, and this explains most of the price increase observed at the target. We also detect significant aggregate cost savings from system ownership, of which approximately 50% is passed on to insurers in the form of price reductions. However, since the independent hospitals acquired are small in relation to the system acquirers, the synergies do not meaningfully affect negotiated prices.

*Emails: jearnold@sas.upenn.edu, atulgup@wharton.upenn.edu, tongl@mit.edu, and olssen@wharton.upenn.edu.
We thank Andrew Olenski for thoughtful comments and participants at the Harvard-MIT-BU Health Economics seminar, Whistler Junior Health Economics Summit, and New York Health Economics Day for their suggestions. All remaining errors are our own.

1 Introduction

The hospital industry has experienced significant consolidation over the past few decades, with multihospital firms, commonly known as "systems", steadily acquiring independent hospitals or other systems. This trend has sparked concerns about the effects of consolidation on prices and the need for more antitrust enforcement (Brot et al. 2024). Proponents of consolidation argue that it generates cost savings which are passed through to consumers while critics counter that consolidation leads to higher prices with little evidence of improved quality. Although a large literature has shown that hospital consolidation leads to higher prices (Gaynor, Ho, and Town 2015), there is little evidence quantifying the underlying mechanisms. For example, prices should increase due to greater market power, but could also increase because the acquirer is superior at bargaining with insurers than the target. Similarly, although previous work has shown that consolidation leads to cost savings (Schmitt 2017; Andreyeva et al. 2024), we are not aware of evidence quantifying the pass-through of these savings to consumers. This paper presents novel empirical evidence to fill these important gaps in the literature and inform this policy debate.

System acquisitions can affect prices at the target facilities through at least three channels. The canonical antitrust concern is that the combined entity gains market power, or bargaining leverage, due to the consolidation, which enables them to negotiate higher prices with insurers. However, there are two additional channels. Depending on the nature of the deal, the acquirer may also have superior bargaining *ability* relative to the target. For example, larger systems typically have specialized and better resourced price negotiation teams, potentially armed with deeper market intelligence. This may allow them to obtain higher rates than the target in negotiations, holding all else equal (Grennan 2014). Although cross-sectional variation in hospital bargaining ability has been documented previously (Lewis and Pflum 2015), its role in price increases following acquisitions is as yet unknown. This channel may also partly explain why recent studies have found evidence of price increases following cross-market hospital acquisitions where conventional market power concerns should not apply (Lewis and Pflum 2017; Schmitt 2018).

The third channel relates to cost synergies. Hospital systems often justify consolidation by pointing to potential cost synergies, such as economies of scale, shared administrative resources, and streamlined operations. By pooling resources and standardizing care delivery, hospital systems argue that they can operate more efficiently, potentially reducing the cost of providing care. Although this is often the key argument for a benefit to consumers from hospital consolidation, there is little or no evidence on whether the resulting cost synergies are passed through to insurers and help lower prices. Since this channel will lower prices while the other two increase prices, "reduced form" retrospective analyses of price changes following acquisitions recover the net effect across all three channels. An economic model is therefore needed to disentangle the contribution of each mechanism.

We inform this active debate by empirically examining the effects of hospital consolidation using novel data on hospital prices and operating costs from California. We have obtained novel proprietary hospital outpatient billing claims data, which allows us to observe detailed patient-level claims, including the negotiated prices, during 2013–2020. The data is sourced from a claims processing intermediary and, therefore, allows us to observe data from a wide sample of hospitals and insurers in the state. Comparisons with the universe of hospitals in California indicate that hospitals in the claims sample are representative on key dimensions such as size, ownership type, etc. Similarly, the leading commercial insurers in California, such as Anthem, Blue Shield, and United Healthcare, are well represented in the sample. We obtain detailed data on hospital operating costs during 2013–2020 from annual reports submitted to California's Department of Health Care Access and Information (HCAI). These administrative reports are more detailed and reliable than the corresponding information in the Healthcare Cost Reporting Information System (HCRIS) reports commonly used to study hospital costs. This unique data source on costs is available only for California and motivates us to focus on this state for the empirical analysis.

We study acquisitions by hospital systems that were implemented during 2014–2018, allowing us to observe both prices and costs before and after all deals for all affected hospitals. During this period, 12 independent hospitals and 11 system-owned hospitals were acquired by systems. We proceed in two steps. We first perform a reduced form analysis using a staggered difference-in-differences research design that is standard in the previous literature on consolidation (e.g., Cooper et al. 2019; Eliason et al. 2020). This enables us to transparently present evidence on the causal effects of system acquisitions on the price and operating cost per outpatient care visit for a privately insured patient. The main finding from this step is that when systems acquire independent hospitals, it leads to both a net increase in prices and a reduction in operating costs, which is consistent with recent work on the effects of system acquisitions (Andreyeva et al. 2024). In the second step, we develop and structurally estimate a model of patient demand for outpatient care and hospital-insurer price negotiations that formally links price setting to costs (Gowrisankaran, Nevo, and Town 2015; Ho and Lee 2017). As is conventional in these models, a hospital's estimated bargaining weight represents an empirical measure of its bargaining ability relative to insurers. We use both the model estimates and the reduced-form effect on costs to simulate counterfactual prices under alternate assumptions about market structure and hospital costs, both for the observed deals and for a larger set of hypothetical deals.

There are three main takeaways from the analyses. First, there is substantial heterogeneity in the effects of system ownership, depending on the nature of the target facility. When systems acquire independent hospitals, we observe a large net increase in prices along with cost savings; however, the increase in prices is an order of magnitude higher than the cost savings. In contrast, acquisitions of system-owned targets show no detectable effects on net prices or on costs. These results imply that acquisitions of system-owned hospitals do not even have the potential to improve consumer welfare through cost synergies. Previous studies have typically not examined

heterogeneity in the effects of system ownership and have tended to estimate the average effect across a variety of deals (e.g., Cooper et al. 2019; Schmitt 2017). However, these results highlight the policy relevance of considering heterogeneity in the organization of the target when studying the effects of consolidation.

Second, acquirers and target facilities appear to be strategically matched to obtain a large price increase for the target hospital rather than to prioritize cost savings. This interpretation flows intuitively from both the reduced form and structural analyses. The independent hospitals acquired had low prices before acquisition relative to both their acquirers and to other independent hospitals. Within three years of acquisition, the prices of the target hospitals nearly converge to those of the acquirers. This pattern is consistent with a pricing strategy of systems persuading insurers to reimburse all their member facilities at similar rates. In contrast, the average prices of independent hospitals that did not participate in acquisitions were comparable to those of acquirers at baseline, suggesting that they would not experience much of an increase. Model estimates from the structural analysis confirm that the acquired independent hospitals had low bargaining weights at baseline, implying that they negotiated poorly with insurers as independent entities and benefited from the acquirer's bargaining ability, reflected by their superior bargaining weight. In contrast, the independent hospitals acquired had similar operating costs to their acquirers prior to the deal, suggesting little potential for further cost savings.

Third, we use our model estimates to simulate counterfactual scenarios to decompose the net price effect due to three channels and quantify their contributions: the increase in bargaining leverage which represents the change in market power, the superior bargaining weight of the acquirer, which represents better bargaining ability at baseline, and cost synergies. We first perform this exercise for the observed acquisitions in our sample. Our main finding is that over 90% of the price increase at the target is due to the targeting of independent hospitals with much lower bargaining ability at baseline relative to the acquirer. Changes in bargaining leverage account for the remainder of the price increase. Finally, our model predicts very small changes in price due to the cost savings following system ownership. Although systems pass-through approximately 50% of the aggregate cost savings at the target in the form of reduced insurer reimbursements, the negotiated prices barely change on average. Although this may seem counterintuitive, comparative statics from the model help explain that the pass-through of cost savings to prices increases in the target's share of the system's total patient volume. Intuitively, since the cost savings are only generated at the target hospital and it is typically small relative to the acquirer, average prices at the acquirer are not affected much.

Since the realized acquisitions appear to be selected to maximize the mismatch in baseline bargaining ability between the acquirer and target, we expand our counterfactual analysis to probe this question more generally. We evaluate around 136 hypothetical within-market acquisitions of independent hospitals. Reassuringly, we find similar patterns in this analysis. The predicted

net change in price increases nearly linearly with the difference in the bargaining weights of the acquirer and the target. In contrast, changes in market concentration due to the deals are not very predictive of the price changes. This occurs in large part because acquisitions of independent hospitals typically lead to a small increase in concentration. Cost synergies also fail to produce meaningful price reductions in this larger set of acquisitions, reinforcing the conclusion that system acquisitions of independent hospitals are unlikely to generate meaningful price reductions, even though they generate savings for the acquirers.

Our paper occupies a unique position at the intersection of three strands of the literature. The first strand examines the effects of hospital consolidation on cost synergies. Recent studies have documented the presence of moderate cost savings following hospital acquisitions with substantial heterogeneity depending on the context. Estimated savings at the acquired hospital are larger when a system acquires an independent hospital and are small or not detected when a system acquires or combines with another (Schmitt 2017; Gaynor et al. 2023; Andreyeva et al. 2024).[1] However, these studies have faced data limitations in deciphering the source of the savings, and specifically whether they are primarily due to fixed or variable costs. We leverage the granular nature of the California cost data and confirm that the savings are mainly driven by a reduction in variable costs at the acquired hospital, and within that, mainly by a reduction in employment of non-nursing staff.

The second related strand is the literature examining the effect of horizontal consolidation in healthcare on prices. These studies have consistently found that horizontal consolidation leads to an increase in prices for inpatient care services (Cooper et al. 2019; Andreyeva et al. 2024). A few studies have estimated or predicted much greater price increases (Dafny 2009; Gowrisankaran, Nevo, and Town 2015), including in cases where the merging hospitals are located in different markets (Lewis and Pflum 2017; Dafny, Ho, and Lee 2019). Our study examines the effect on prices in a new setting, outpatient care services, which now contributes more than half of all hospital revenue and yet remains understudied. We show that systems target hospitals with low prices and poor bargaining ability at baseline, suggestive of a strategy of price arbitrage. We subsequently show that the majority of the price increase at the target is explained by the large difference in bargaining weights between the acquirer and target, and not by the increase in leverage, the conventional antitrust concern. This pattern is replicated in a large set of hypothetical acquisitions of independent hospitals.

Finally, we add to a nascent literature that has examined the pass through of merger efficiencies to customers (Miller and Weinberg 2017). In the hospital setting, Andreyeva et al. (2024) showed that system acquisition of independents leads to higher net prices and cost savings at the target facility.

1. Outside the hospital setting, a substantial body of literature similarly finds that firms often realize productivity improvements and cost synergies following mergers and acquisitions (e.g., Andrade, Mitchell, and Stafford 2001, Maksimovic and Phillips 2002, Schoar 2002, Maksimovic, Phillips, and Prabhala 2011, Miller and Weinberg 2017, and Demirer and Karaduman 2024).

However, they did not quantify the pass-through of the savings to insurers. To our knowledge, ours is the first to perform this exercise in the hospital setting. Our analysis confirms that the acquirer passes through a substantial share of aggregate savings to the insurer in the form of price reductions. However, our analysis also provides a novel insight that such cost savings are unlikely to meaningfully affect hospital prices because independent targets are typically very small relative to the acquiring system. The effect on prices would be greater in cases where the acquirer and target are both relatively closer in size. However, we do not detect any cost savings when we study the acquisitions of system-owned targets, where this is more likely to occur.

The rest of the paper is organized as follows. Section 2 provides a brief background on the potential effects of system ownership, recent trends in hospital outpatient care, and the related literature. Section 3 describes the key data sources and sample construction. Section 4 describes the reduced form analysis on the effects of system acquisitions on prices and costs. Section 5 presents a model of patient demand for outpatient care and a Nash-in-Nash bargaining model of hospital-insurer price negotiations. Section 6 presents the model estimates and Section 7 presents results from the counterfactual simulations. Finally, Section 8 concludes.

2 Background

2.1 Hospital system ownership

Multihospital firms, colloquially known as hospital "systems", have steadily expanded their footprint in the US hospital industry over the past few decades. In 2000, systems represented nearly 60% of bed capacity among short-term general acute care hospitals. This share increased to over 80% by 2020 (Andreyeva et al. 2024). Since relatively few new hospital facilities have opened in the US during this period, most of the growth in system capacity has been inorganic, driven by the acquisition of formerly independent or standalone hospitals.

System ownership offers numerous benefits to independent hospitals. One of the most significant benefits is economies of scale, which affects nearly every aspect of operations. In theory, systems can leverage the bulk purchasing power for medical supplies, equipment, and services, significantly reducing the procurement costs for materials and equipment. Scale facilitates greater financial stability. Systems enjoy access to cheaper credit and potentially more sources of capital than standalone facilities. A lower cost of capital gives systems a competitive edge in making the necessary investments for growth. Furthermore, systems can subsidize unprofitable member units for short periods of time, allowing them to make necessary long-term investments or changes in their operations. Systems can also streamline operations by centralizing administrative functions such as billing and accounting, IT, and human resources, minimizing overhead expenses. Finally, greater scale also helps to attract and retain top clinical and managerial talent. This has important implications for reputation, patient demand, and operational efficiency.

In addition to having access to potentially higher quality staff, hospital systems may also provide better patient care due to standardized protocols that ensure consistency in quality and safety across all facilities. Shared electronic health records (EHR) further improve care coordination, especially for patients receiving treatment at multiple locations, leading to better health outcomes.

The preceding advantages of system ownership may be considered beneficial to society as a whole, as these changes improve efficiency in the hospital industry and patient care simultaneously. Transitioning an independent hospital to system ownership may therefore generate a surplus for both the hospital and payers. The latter group is diverse and consists of the federal and state governments, employers, private insurers, and self-insured consumers. However, system ownership may also lead to changes that benefit the hospital at the expense of payers. Specifically, systems may use acquisitions to increase their share of a local market and, consequently, increase their bargaining leverage in negotiations with private insurers leading to higher prices without a commensurate change in quality of care. Alternatively, systems may use their existing leverage and bargaining tactics to negotiate a higher price for the newly acquired hospital without meaningful changes to the target facility. By construction, this will result in a welfare loss for consumers and taxpayers that fund hospital care.

Although the discussion above was with reference to the effects of system ownership for an independent hospital, the same arguments also apply to the case where a system acquires another system, so long as the acquirer is larger and offers scale economy benefits to the acquired units. That is, the effects of system ownership may apply both on the extensive and intensive margins.

2.2 Growth of outpatient care

In general, hospitals provide two types of care services, inpatient and outpatient. Inpatient services refer to care for which patients spend at least one night in the hospital. These services are geared for more intensive procedures that require sophisticated technology or that require supervised postoperative recovery and continuous monitoring for patients. Examples of inpatient care include complex procedures such as cardiac surgery and treatment for highly acute conditions such as pneumonia. In contrast, outpatient services refer to care in which the patient returns home after the procedure without staying overnight in the hospital. Common examples include diagnostic imaging and screening tests and minor procedures such as colonoscopies. Intuitively, hospitals receive much lower reimbursement for outpatient procedures. Traditionally, the vast majority of hospital revenue was sourced from inpatient care. As a result, economists have typically focused on outcomes related to inpatient care, including patient volume, prices, and quality.

With advances in technology, more procedures that were previously considered safe only in the inpatient setting are now feasible in the outpatient setting as well. Insurers have a financial incentive to encourage hospitals to perform more procedures in the outpatient setting once the technology has been proven. A recent example is knee replacement surgery. Starting in 2018, the Centers for

Medicare and Medicaid Services (CMS) allowed hospitals to perform knee replacement surgery, known as total knee arthroplasty (TKO), as an outpatient procedure for Medicare patients. Medicare payments for TKO are approximately 30% lower when it is outpatient, generating significant savings for the program.

As a result of technology and cost considerations, the volume of outpatient care has increased over time, while that of inpatient care has declined. Therefore, the share of total revenue contributed by outpatient care has also increased. Figure A1 presents the trend in the revenue share of outpatient care specifically for privately insured patients at general acute care hospitals. Panel (a) presents the trends during 1995–2019 using data from CMS' HCRIS data for a national sample of hospitals. The revenue share of outpatient care increased from 35% in 1995 to more than 50% at the end of the study. Panel (b) shows a very similar pattern for hospitals in California using more granular data, albeit over a shorter period (2003–2019). Hence, outpatient care now accounts for the majority of revenue at hospitals among privately insured patients. This pattern, coupled with the paucity of evidence on outpatient care, motivates our focus on outpatient care services in this study.

2.3 Related Literature

2.3.1 Hospital Consolidation and Cost Efficiencies

Economists have long debated whether hospital consolidation yields economies of scale and lower operating costs. Early evidence was mixed, but recent studies suggest modest cost savings. In a difference-in-differences analysis of 2000–2010 mergers, Schmitt (2017) finds that acquired hospitals experienced 4—7% reductions in operating cost per discharge following a merger or acquisition. These cost declines were statistically significant and robust to alternative controls, implying that some promised efficiencies do materialize. Andreyeva et al. (2024) examine the effects of system acquisitions over 2013–2017 and find a 5% decrease in costs when the targets are standalone hospitals, but only a small and statistically insignificant effect when the target is already part of a system. Similarly, Gaynor et al. (2023) fail to detect cost savings following the merger of two large hospital systems. Overall, recent literature suggests that hospital mergers can produce some operating cost efficiencies, although far from uniformly or dramatically. These studies have used operating cost data from Medicare cost reports, which limited their ability to shed light on the origins of the cost reductions. Hence, there is little consensus on the mechanisms that produce these efficiencies.

2.3.2 Hospital Consolidation and Prices

A separate body of work analyzes how hospital consolidation affects prices negotiated with private health insurers. Here the evidence is clearer: consolidation leads to higher prices on average. Gowrisankaran, Nevo, and Town (2015) develop a structural bargaining model for hospital-insurer negotiations and use it to simulate a proposed merger in northern Virginia. Their simula-

tions predict that a within-market hospital merger would significantly raise negotiated prices absent large efficiencies, and even with assumed cost synergies, the net price effect remains positive. Using transaction price data from large insurers, Cooper et al. (2019) examine the effects of 366 hospital mergers from 2007—2011. They estimate that transactions between nearby hospitals led to price increases over 6% for the merging facilities, whereas mergers between hospitals in distant markets (> 25 miles) had no significant price effect. Andreyeva et al. (2024) use transaction price data from a different insurer and reach a similar conclusion. Other studies have detected price increases following hospital acquisitions that span markets. Lewis and Pflum (2017) and Dafny, Ho, and Lee (2019) document price increases even when the merging hospitals do not share a market. The latter study theorizes that the price effects arise due to the presence of a common customer, the health insurer. However, their model assumes away the possibility of cost efficiencies following a merger. In addition to market power, Lewis and Pflum (2015) hypothesize that systems obtain higher prices because they are more skilled at negotiating prices with insurers than independent hospitals. Using data from California, they show that hospital systems have much higher bargaining weights than independent hospitals at a point in time. Overall, the net increase in prices following hospital consolidation implies that cost synergies are not passed through sufficiently to offset the incentive to increase prices. The same pattern has also been detected in other sectors of healthcare. For example, Dafny, Duggan, and Ramanarayanan (2012) find that insurer consolidation led to higher premium growth.

2.3.3 The Pass-Through of Cost Synergies

To our knowledge, Andreyeva et al. (2024) is the only study that has examined the effects of hospital acquisitions on prices and costs together. They detect a negative correlation between price and cost effects when systems acquire independent hospitals. This is consistent with a pass-through of savings, as deals that delivered larger cost savings also had smaller price increases. However, they cannot accurately quantify the passthrough of savings, since they cannot simulate prices without cost savings holding market structure constant. Therefore, the pass-through of cost synergies in hospital acquisition remains unexplored. Beyond the hospital setting, two recent studies have empirically examined this question. Blonigen and Pierce (2016) analyze manufacturing plant data across many sectors and find that after mergers and acquisitions, average markups increase without detectable improvements in plant productivity. In other words, mergers increased the pricing power of firms more than they improved efficiency. Miller and Weinberg (2017) estimate a structural model of the beer market and show that a merger of two large manufacturers achieved cost reductions (shipping distances fell by an average of 360 miles). The cost savings were large enough to offset unilateral price increases by the merging firms. However, the study finds that prices increased on net due to price coordination among competing firms after the merger.

3 Data sources and key variables

This section describes the main data sources that underpin our analyses and the construction of key variables.

3.1 Hospital outpatient care use and prices

We use a novel data set that allows us to observe a sample of national patient-level billing claims for outpatient hospital care during the period 2013–20. This was first used by Liu (2022) to study price negotiations between private equity-backed hospitals and private insurers. This data has several important features that suit our application. First, it contains granular details about the characteristics of the patient and the procedures performed, including the ability to follow patients over time across providers. Therefore, we can carefully model the demand for outpatient care from the patient. Second, for a subset of claims, we observe transaction prices for privately insured patients negotiated between hospitals and commercial insurers. Researchers typically have not had access to transaction prices since insurers and providers have protected this information as a trade secret. As a result, much of the prior literature on hospital consolidation has relied on prices imputed from hospital revenue data reported to regulators (Dafny 2009; Lewis and Pflum 2017; Dafny, Ho, and Lee 2019). Third, since the data is sourced from a claims processing intermediary between insurers and providers, it is not limited to a single insurer, provider, or employer, and is more representative of the average prices for hospital care faced by privately insured consumers.

We limit our empirical analysis to California, since it is a large state that experienced several hospital acquisitions during this period, is covered relatively well in the claims sample, and publicly releases detailed hospital accounting information that allows us to examine changes in operating costs for commercially insured patients. With regard to the claims data, we observe prices for 273 general acute care hospitals in the sample, out of a maximum of 299 possible facilities.[2] Panel A of Table A1 shows that the hospitals in our sample are representative of the universe of hospitals in California on important aspects like bed size, owner type, and payer mix. Similarly, Panel B of Table A1 shows that we observe prices for the large commercial insurers in the state. We limit the sample to patients aged 18 to 64 years and further exclude patients in privately administered Medicaid and Medicare plans. The final analysis sample has 6.6 million privately insured outpatient visits over 8 years and we observe the allowed amounts (prices) for approximately 20% of them. Section B.1 describes the claims data and sample construction in more detail.

2. Kaiser Permanente (KP) is the largest hospital system in California with 31 hospitals in 2019. This is not observed in the claims data and we do not include it in the count of hospitals in the state since, for the most part, it is a closed-network vertically integrated firm that offers both insurance and healthcare to its members. Hence, these hospitals do not negotiate contracts with other insurers, nor does the insurance arm purchase care from outside facilities.

Hospital price index

The patient-level claims allow us to observe prices by hospital, procedure, and payer. However, retaining such rich heterogeneity in prices introduces additional complexity into the analysis. Descriptive studies of hospital prices have found that variation in prices across hospitals is much larger in magnitude than that across payers (Panhans, Rosenbaum, and Wilson 2021). Hospitals are also much more likely to be consistently high- or low- priced across insurers than vice versa (Gorman et al. 2016). Therefore, to simplify the analysis, we construct a scalar price index for each hospital year, eliminating the variation across insurers. We deflate prices by the CPI-U series so that they are expressed in 2020 dollars. We also scale observed prices by the procedure APC weight, a measure of the intensity of resource use developed by CMS to adjust Medicare payments for hospital outpatient care. Hence, the variable used in our analysis reflects a standardized price charged by a hospital for a procedure of APC weight one. Finally, we estimate fixed effects models at the patient level that account for differences in patient demographics and health risks. Hospital-year fixed effects recovered from this estimation serve as the corresponding price indices. This approach follows previous work on the effects of hospital consolidation (Gowrisankaran, Nevo, and Town 2015; Cooper et al. 2019). Appendix B.1 describes the steps in more detail.

3.2 Hospital operating costs

We obtain information on operating costs from annual hospital cost reports submitted by hospitals to California during the period 2013–2020. The reports are processed by HCAI and released for public use. We prefer using these reports over the HCRIS or Medicare cost report data more commonly used in health economics research for three reasons. First, states are known to invest more in maintaining the quality of such reports than CMS, such as by performing regular audits. California reports have long been recognized as high quality and have frequently been used by researchers to study hospital operations (Duggan 2000; Ho and Lee 2017). Second, California data allow us to calculate, with reasonable assumptions, the average cost per outpatient visit for privately insured patients, which would require more assumptions to be made if using HCRIS. Third, we can disaggregate costs into more granular categories in California reports than in HCRIS. This enables examining the sources of cost synergies following system ownership.

Outpatient cost per visit

We briefly describe the steps to calculate the operating costs for privately insured patients, the details of which are left for Section B.3. For each hospital, we observe operating cost by granular service line (e.g., coronary care, emergency department, etc.). We also observe the gross revenue for inpatient and outpatient care by payer type (e.g., Medicare, Medicaid, and private) for each service line. We then make two assumptions that have been used in prior research examining hospital costs. First, outpatient care contributes the same share of total cost as its share of gross

revenue for a service line. In effect, we are assuming that the ratio of list prices ("charges") to costs does not vary by service line. Researchers frequently invoke this assumption when they use a hospital-wide cost-to-charge ratio to calculate patient-level costs based on charges. Second, we assume that payers contribute the same share of the total outpatient cost for a service as their share of gross outpatient revenue for that service. This effectively imposes that the cost-to-charge ratio within a service is constant across payers. Using these two assumptions, we calculate the outpatient cost by service line and payer and aggregate to the payer level. Since the report also records the number of outpatient visits by payer, we calculate the mean operating cost per visit for privately insured patients. We scale the mean cost per visit by the mean APC weight corresponding to the hospital in the claims sample and obtain a standardized mean cost per visit for a weight of 1. Finally, as with prices, we deflate the operating cost per visit to be expressed in 2020 dollars.

3.3 Changes in hospital ownership

A key task is to identify which hospitals and systems are involved in acquisitions. We limit the analysis to deals that were executed between 2014 and 2018. This allows us to observe prices and costs for at least one year before and two years after each deal. To identify the facilities involved in the deals, we closely follow the approach used by previous studies on hospital consolidation (e.g., Schmitt 2017; Andreyeva et al. 2024). We identify potential acquisitions by examining changes in system ownership reported in the American Hospital Association (AHA) annual hospital surveys. The AHA assigns a unique ID to each system, allowing us to track changes in hospital-system links. We detect two types of deals involving systems as acquirers. First, systems may acquire independent or standalone facilities. In this case, a previously unaffiliated hospital begins to report being a member of a specific system. Second, one system may acquire all or some facilities owned by another system. This will show in the data as a change in the system ID associated with the acquired hospitals. This exercise results in a starting list of the facilities involved in acquisitions.

We enhance the list from the AHA using hospital acquisitions listed in PitchBook, Preqin, and SDC platinum, which have also been used in prior studies of hospital ownership (Liu 2022). Finally, we manually validate each deal to confirm: 1) it involves a change in ownership, not merely a partnership or clinical affiliation; 2) the acquirer and target(s) involved; and 3) the year the deal closed as opposed to when it was announced.[3] This process yields 12 deals in which systems acquired independent hospitals and 2 deals in which one acquirer, Providence, acquired a total of 11 hospitals from other systems. Hence, a total of 23 hospitals were acquired during this period in California. Table A3 lists the acquisition systems and targets for all the deals identified during this period.

3. Reassuringly, our list of deals overlaps almost perfectly with the independently validated list of hospital mergers and acquisitions released by Oh et al. (2025).

3.4 Supplementary data

In addition to the key data sources described above, we use several publicly available datasets. This includes data from the American Community Survey (ACS) on population and demographic characteristics by zipcode, definitions from California HCAI, and hospital characteristics from CMS Provider of Service files. The travel time to the hospital is a key explanatory variable in the demand analysis. This is calculated using information from the proprietary API of a data provider, HERE. Specifically, we look up the average travel time between the population-weighted centroid of the patient's 3-digit zipcode and the hospital's exact location.

4 Reduced-form analysis

This section describes a reduced-form analysis to quantify the causal effects of acquisitions by systems. The goal is to transparently test whether system ownership leads to changes in prices and operating costs for outpatient care services.

4.1 Research design

We follow previous studies on hospital consolidation and use a staggered difference-in-differences (D-D) research design for this exercise (e.g., Craig, Grennan, and Swanson 2021; Andreyeva et al. 2024). We begin with an approach that mirrors most prior work and estimates an average effect across all acquisitions in the sample. The design compares the trends for target hospitals with those of a set of comparison hospitals that are not involved in any acquisition during this period. The comparison group includes hospitals that remain independent until December 2020 (when studying the effects on independent targets), or system-owned hospitals that are neither acquirers nor targets in a deal themselves (in the case of system-owned targets).

We estimate the model presented in Equation 1 below to quantify the impacts on the acquired hospitals.

$$Y_{ht} = \alpha_h + \alpha_t + \beta\, 1(\text{acquired})_h \cdot 1(t \geq T)_{ht}\, [+X'_{ht}\delta]\, + \epsilon_{ht}. \tag{1}$$

We include hospital fixed effects to eliminate stable unobserved differences between hospitals and ensure that we rely only on within-hospital changes in outcomes to estimate the treatment effect. Since the most important concerns about potential omitted variables bias relate to stable factors (e.g., teaching status, weak management, etc.), this approach mitigates concerns about the endogeneity of acquisitions to a first-order. We also include year fixed effects to account for secular trends in the outcomes, such as changes in operating costs that reflect technological advances. The coefficient of interest, β, is on the interaction of an indicator for acquisition and a hospital-specific time-varying indicator that takes the value of one starting in the year of acquisition. In secondary models, we also include a vector of time-varying hospital characteristics, X_{ht}, such as the number

of beds and the share of patients contributed by Medicare and Medicaid. This helps control for differences in attributes that may change significantly over time and affect outcomes. Reassuringly, including or excluding these covariates does not affect the coefficients of interest, mitigating concerns about omitted variables bias. ϵ_{ht} denotes all remaining unobserved idiosyncratic factors.

Under the standard parallel trends assumption, β is unbiased and recovers the average treatment effect on treated hospitals, or the ATT. In our setting, this means assuming that the acquired and comparison hospitals would continue along their preexisting trends in the absence of acquisitions. This assumption is untestable. However, we present evidence on the dynamic treatment effects to assess the plausibility of the assumption. Equation 2 below represents these models. The coefficients β_s from years before treatment indicate whether the acquired hospitals trended differently from the comparison hospitals even before they were acquired. A lack of differential pre-trends supports the parallel trends assumption.

$$Y_{ht} = \alpha_h + \alpha_t + \sum_{s=-5,\neq -1}^{s=4} \beta_s \, 1(\text{acquired})_h \cdot 1(t = s)_{hs} \, [+X'_{ht}\delta] \, + \epsilon_{ht}. \tag{2}$$

4.2 Estimation

Recent econometric studies have shown that D-D estimates obtained using conventional two-way fixed effects (TWFE) models suffer from bias in the presence of staggered treatment and heterogeneity in the treatment effect (Goodman-Bacon 2021). This threat also extends to the dynamic effects obtained by using Equation 2 (Sun and Abraham 2021). We use the estimator developed by Borusyak, Jaravel, and Spiess (2024) to circumvent this concern. This estimator is consistent for the ATT in our setting even if there is treatment effect heterogeneity, and it produces unbiased dynamic treatment effects in addition to an unbiased ATT. We use the R package "didimputation" to implement this estimator.

4.3 Descriptive evidence

Figure 1 presents California maps to help the reader navigate the geography of hospital markets in the state and the locations of hospitals involved in the acquisitions studied here. Panel (a) depicts the hospitals involved in the acquisitions of independent targets, while Panel (b) does the same for the acquisitions of systems-owned facilities. We use the geographic markets delineated by HCAI based on healthcare use patterns. These are called Health Service Areas (HSA), and they consist of contiguous counties selected so that the residents of these counties are typically observed choosing hospitals within the set. These markets have previously been used by Ho and Lee (2017) to study hospital-insurer price negotiations for inpatient care services. Although the agency defines 14 markets, we combine two small markets into their larger neighbors based on

the hospital choice patterns observed in the claims data sample. Our analysis therefore recognizes 12 hospital markets.[4]

For those familiar with California's geography and population distribution, the market definitions are intuitive. For example, 14 sparsely populated counties in the northern part of the state are combined to form one HSA with Redding as the main city (HSA 1). At the other extreme, Los Angeles (LA) County is coterminous with a hospital market (HSA 11).

The map also shows the locations of the hospitals involved in the acquisitions. Hospitals that were already part of the acquiring system prior to 2014 and targets are represented in triangles and circles, respectively. Intuitively, most target hospitals are located in populous counties in or near the San Francisco Bay area and LA. In our sample, acquirers have more geographic overlap with their independent targets. Eight out of the 12 independent targets share a market with at least one acquirer hospital, while only one of the 11 system-owned targets shares a market (LA) with the acquirer, Providence. Previous studies on hospital consolidation have found that the geographic proximity between merging parties is a strong predictor of price increases (Dafny 2009; Cooper et al. 2019). Furthermore, although Providence is a large chain nationwide, it owned only five hospitals in California prior to 2014 and was therefore a small player relative to other systems in the state (e.g., Sutter). Together, these patterns suggest that the price gains due to acquisition will be larger for independent targets.

Table 1 presents descriptive statistics on the hospitals observed in the sample. The values represent means for the first year that we observe the hospital in the sample, which is 2013 in the case of most hospitals. Column 1 describes all hospitals in the sample. Columns 2–7 describe different subgroups of hospitals based on their system status and role in the analysis. Column 2 describes the 50 hospitals belonging to acquirers that bought independent hospitals. Column 3 describes the 12 independent target hospitals, while Column 4 describes the corresponding comparison group of hospitals that remain independent through the end of the sample period. Column 5 describes the hospitals of the acquirer that bought system-owned hospitals, which are described in Column 6. Finally, Column 7 describes the comparison group of system-owned hospitals that were not involved in any acquisitions during this period.

Panel A presents hospital characteristics such as bed capacity, teaching status, share of patients on Medicaid and Medicare (including private plans such as Medicare Advantage), and the fraction of hospitals owned by for-profits. Panel B presents the mean values of the key outcome variables. This includes the standardized values of the price (allowed amount) and variable cost per outpatient visit for privately insured patients. We refer to these as standardized values since they

4. Specifically, we combine HSA 3 (Napa, Solano, and Sonoma counties) into HSA 4 (San Francisco, Marin, and San Mateo counties) and HSA 8 (Monterey, San Benito, Santa Cruz, and San Louis Obispo) into HSA 10 (Ventura and Santa Barbara). This decision was prompted by the fact that the majority of patients in the claims data residing in the former markets choose hospitals located in the latter markets.

correspond to a procedure of APC weight one. We also present the mean APC weight for privately insured patients observed in the claims sample. Finally, we also present the proportion of total operating cost that we consider to be variable in nature. Table **??** lists all the cost categories observed in the cost reports and their classification into fixed and variable costs.

Columns 2 and 3 compare the acquirer and target independent hospitals at the beginning of the sample period. The targeted hospitals had about 15% smaller bed capacity on average, but had similar shares of patients on public insurance plans and a similar rate of for-profit ownership. The most striking difference is in the price level. The hospitals in the acquiring system enjoyed a standardized price nearly 75% higher than the independent hospitals targeted for acquisition. In contrast, variable operating costs were similar between the two groups of hospitals, implying that the acquirer enjoyed significantly higher mark-ups. Indeed, mean variable (and total) cost per visit was slightly lower at the targeted facilities. The hospitals in Column 4 constitute the comparison group for the independent targets (Column 3). The hospitals in both groups are smaller on average than the acquirer's hospitals. Comparison hospitals enjoy much higher prices than the target hospitals. In fact, their price level is even higher on average than that of system hospitals. These patterns argue against the realization of cost efficiencies as the key motivation for the acquisition of independent hospitals. Instead, a pricing arbitrage opportunity is readily apparent if the acquirer could renegotiate prices for the newly acquired hospital at levels close to those enjoyed by its existing hospitals.

Similarly, comparing the acquirers in Column 5 and the system-owned target hospitals in Column 6 shows that in this case the target hospitals had only slightly lower price levels than the acquirers at baseline. The target hospitals appear to have a lower standardized operating cost at baseline, arguing against the potential for further efficiency improvements by merging into the acquirer. Considering prices and costs together, the acquirer and target hospitals had similar markups. Perhaps the most distinctive feature of the system target facilities is that there is no for-profit hospital among them.

4.4 Results

This section presents the results of the reduced form analysis. We focus on the effects on the prices and costs of outpatient care services. We also investigate the effects on specific cost categories to help explain the changes in operating costs.

4.4.1 Prices

Table 2 presents the effect of system acquisition on prices for outpatient care services. We present difference-in-differences estimates for independent target hospitals (columns 1-2), system-owned target hospitals (columns 3-4), and all targets (columns 5-6). For each group of targets, we present estimates from unweighted models without and with time-varying covariates discussed in Section

4.1. The models presented in even-numbered columns represent our preferred specification. Figure 2 presents the corresponding dynamic effects on outpatient prices estimated using Equation 2 from a model corresponding to Column 2. Panels (a), (b), and (c) present the effects on prices for independent targets, system-owned targets, and all targets, respectively.

Table 2 Column 2 indicates that independent targets experience a sharp increase in price of approximately $221. This represents more than a 100% increase in the baseline price level at the targeted hospitals before acquisition. The choice of specification does not qualitatively affect this conclusion. Figure 2 Panel (a) confirms this interpretation. There is no evidence of differential trends before acquisition and a continuous increase in price in each year of system ownership, which appears to stabilize by the fourth year after the deal. A large increase in price at independent targets due to system ownership is also consistent with unadjusted trends in the price index. The mean price at the acquired hospitals rose approximately $240 from the year before the acquisition to 3 years after the acquisition.

Table 2 Column 4 shows that system-owned targets experience a decrease in price after acquisition, although the coefficient is imprecisely estimated. Figure 2 Panel (b) shows a downward trend in price after acquisition, however, it is not consistently trending down. Overall, we interpret these patterns as representing an ambiguous effect on the target's price index.

Table 2 Column 6 implies that system acquisitions result in an average increase of $45 per outpatient visit on a base of approximately $400. This represents a meaningful 10% increase in prices; however, the coefficient is imprecisely estimated at conventional levels of significance. Figure 2 Panel (c) shows dynamic effects consistent with a modest and imprecise increase in price. For completeness, we also present DD estimates obtained using conventional TWFE models. These results are presented in Table A4 and imply quantitatively similar effects as our main estimates.

The results highlight the heterogeneity in the treatment effects of system acquisition depending on the state of the target hospital. Much of the previous literature has overlooked this heterogeneity and estimated the average effect of hospital consolidation (e.g., Cooper et al. 2019). However, our estimates imply that the consolidation of independent hospitals has a much larger impact on price levels than changes in ownership of hospitals that are already owned by a system.

A potential concern with the effect on prices could be whether the result described above reflects a change in the set of insurers contracting with the target, a change in negotiated prices for the same set of insurers, or a mix of both factors. To shed more light on the role of different channels we directly test for a change in insurer mix at the acquired hospitals. Table A5 presents the corresponding DD coefficients from this analysis. We organize the data at the hospital-insurer-year level for these regressions. The outcome of interest is an indicator for whether the hospital contracts with the insurer in a given year, which we determine based on whether we observe a claim with nonmissing price information for the hospital-insurer pair in that year. Therefore

this test also accounts for potential nonrandom missingness in price information in the outpatient claims sample. Columns 1, 2, and 3 present the effects for independent targets, system-owned targets, and all targets, respectively. For brevity we focus on our preferred models that include time-varying covariates. As the results indicate, we do not detect a change in the insurer mix at any of the target categories. The coefficients are not only imprecisely estimated but also close to zero, implying little change. Hence, we conclude that the change in price at independent targets mainly reflects a change in negotiated prices for the same set of insurers.

4.4.2 Operating costs

Table 3 presents the corresponding difference-in-differences estimates on variable operating cost at the acquired hospital following system acquisition. It is identical in format to Table 2 that examined the effects on prices. Similarly, Figure 3 presents the corresponding dynamic effects of system acquisition on operating costs at the targeted hospitals.

We detect a precisely estimated and economically meaningful reduction in operating cost at independent targets. The coefficient in Table 3 Column 2 indicates a reduction of approximately $32 per visit, approximately 30% of the baseline variable cost. Figure 3 Panel (a) presents the corresponding dynamic effects, which corroborate the DD estimate. The figure shows no evidence of differential trends at the targeted facilities before acquisition and a sharp trend break toward a reduction in costs after system ownership.

The point estimate in Table 3 Column 4 implies a modest decrease in variable costs at system-owned targets after acquisition. However, the coefficient in our preferred specification is not statistically significant. As in the case of prices, these results highlight the importance of heterogeneity in the effects of system acquisitions. Repeated acquisitions of system-owned hospitals do not appear to generate economically meaningful cost savings.

Reflecting this heterogeneity in effects, Table 3 Column 6 reports an average decrease in operating cost per visit of $17. This represents a decrease of approximately 16% relative to the baseline cost at the target hospitals. However, the coefficient is statistically insignificant. Panel (c) of Figure 3 similarly suggests that there is no significant change in costs after acquisition.

We also present the DD estimates of effects on operating costs obtained using conventional two-way fixed effect (TWFE) models in Table A6. These models imply an even larger cost saving of $38 at formerly independent targets.

Sources of cost savings

The previous analysis focused only on variable operating costs. We further investigate the reduction in costs at independent targets by examining the effect on different categories of total operating costs. Figure 4 presents dynamic effects on three mutually exclusive and collectively

exhaustive groupings of costs at independent targets. We focus on this subset since we detect a change in operating costs only for this group. The top panel presents the effect on labor expenses, assigned entirely to variable costs following the approach described in MedPAC (2020). The two bottom panels present the effect on nonlabor expenses bifurcated into fixed and variable, respectively. The panels clearly show that the main reduction in operating cost is driven by labor. The dynamic effects imply a small increase in fixed costs and a small decrease in nonlabor variable costs (e.g., medical supplies). Table A7 presents the corresponding D-D estimates for the effects on each category. The DD coefficients are consistent with the event study plots and show that the decrease in operating costs is mainly due to labor expenses. Note that the point estimates in this table do not line up perfectly with those in Table 3 since the samples are slightly different.

Since the evidence points to a reduction in labor costs, we quantify the changes in employment (Full-time equivalent staff) and mean hourly wages at independent targets following the acquisition. We further examine these two variables separately for administrative and non-administrative staff, since these are reported distinctly in the California reports. Note that our approach attributes all labor costs as variable in nature.

Table 4 Panels (a) and (b) present the corresponding effects on employment and hourly wages, respectively, using our baseline approach. Our measure of employment is the log of full-time equivalent (FTE) staff. We examine employment in logs since there is tremendous heterogeneity in total employment across hospitals of different sizes. Panel (b) examines the effects on hourly wages in dollars. As in the main results tables, Columns 1-2, 3-4, and 5-6 present the results across independent targets, system-owned targets, and all acquisitions, respectively. We focus on the results for independent targets for brevity. Considering the results on employment and wages together, the reduction in labor costs appears to be driven entirely by a reduction in employment, with small and statistically insignificant effects on mean wages. There are relatively large decreases in both administrative and non-administrative staff. Although the percent reduction is greater among administrative staff, a greater absolute number of FTE are reduced from among the non-administrative staff since that group is nearly ten times larger: a reduction of approximately 11 versus 48 FTE.

Table A8 further investigates the reduction in non-administrative employment. We divide non-administrative staff into nurse and non-nurse categories. We pool all subtypes of nurses into the former group while the latter group includes technicians, specialists, social workers, and case managers. The table clarifies that the reduction is entirely driven by cuts to the non-nursing staff. In results not presented here, we find that the biggest reductions in this subgroup are among technicians and specialists. This result is also important because it suggests that system acquisitions of independent hospitals are unlikely to directly affect quality of care, which is disproportionately sensitive to changes in nurse availability.

4.4.3 Patient volume and composition

A reduction in employment at the target hospital naturally raises the question of whether there is also a reduction in patients served by the hospital after the change in ownership. We directly assess this by testing for changes in patient volume at independent targets. We use data on both hospital inpatient stays and outpatient visits reported to California for this exercise. Figure A2 presents the corresponding dynamic estimated effects. We convert the volume variable into logs for this regression analysis to address heterogeneity in hospital size. Panel (a) suggests there may be a slight increase in inpatient stays after the acquisition, although the coefficients are imprecisely estimated. Panel (b) indicates no change in outpatient visits after the acquisition. Table A9 presents the corresponding DD coefficients which are consistent with the event study plots and imply imprecisely estimated increases in volume.

A potential explanation for the decrease in hospital staff while serving the same patient volume could be a decrease in the complexity of the average patient. This could occur if the acquiring system steers the target to specialize in some low-risk services and directs more complex patients to other hospitals in the system. This explanation is unlikely to hold, since our variable cost measure is scaled by the hospital's APC weight, an index reflecting the intensity of the procedure. A decrease in procedure intensity would tend to *increase* the standardized cost per visit rather than decrease it. However, procedure intensity may not accurately reflect patient complexity due to upcoding or prescribing unnecessarily intense medical treatment. We directly test this hypothesis by examining changes in observed patient demographics and risk indicators.

Table A10 presents the corresponding DD coefficients using our baseline model including covariates. We examine two measures of demographics: the fraction of women and the mean age and find small and statistically insignificant effects on both dimensions. We test for changes in two measures of risk: the mean APC weight (a measure of procedure intensity) and a flag that indicates whether the patient had an outpatient hospital visit in the past year. Patients with a recent history of hospital use are likely to be more complex. The nice feature of this flag is that it is not sensitive to upcoding. Once again, we estimate small and statistically insignificant coefficients. Figure A3 presents the corresponding event study plots, which are consistent with the DD coefficients. Overall, the results imply that there is no meaningful change in patient complexity at the target hospital in the period after system ownership.

4.4.4 Quality of care

The previous results imply that systems reduce the number of staff at formerly independent targets without any change in the composition or number of patients. This raises the concern that quality of care may suffer after the change in ownership. We test this hypothesis by examining standard measures of quality used by CMS in its hospital performance pay programs such as the Hospital Readmissions Reduction Program (HRRP) and Hospital Value-Based Purchasing

(HVBP). Specifically, we examine the effect on 30-day risk-adjusted readmission and mortality rates for Medicare fee for service patients admitted with heart attack, heart failure, or pneumonia. Table 5 presents the corresponding DD coefficients using our baseline estimating equation, including time-varying hospital covariates. Each coefficient is obtained from a separate regression. The sample size varies across conditions, since CMS calculates and releases data on risk-adjusted quality metrics for different sets of hospitals by condition. The table shows that the effects on readmission and mortality rates are small and statistically insignificant. They are also not consistently of the same sign across conditions. We infer that there is no detectable change in hospital quality, at least on these measures. Figure A4 presents the corresponding event study plots, which are consistent with the DD point estimates.

Considering the results on operating costs, patient volume, composition, and quality of care together, we conclude that when systems acquire independent hospitals, they reduce labor costs by reducing employment, primarily among administrative and non-nurse patient care employee groups. There is no change in wages. There is no evidence of changes in patient mix, volume, or quality of care that may help explain the decrease in employment. Hence, we interpret these results collectively as a sign of cost savings due to improved operating efficiency.

Next, we develop an economic model of insurer and hospital price negotiations and estimate the parameters of the model so that we can simulate counterfactual scenarios and decompose the net price effect into key components of interest, including the pass-through of cost savings in the form of price reductions.

5 Model and estimation

5.1 Model

This section describes our model of the market for private hospital outpatient care. The model features the negotiations between hospital systems and insurers as well as patients' choices over providers. We use this model to quantify how provider reimbursements are determined in equilibrium. Furthermore, we use the model to decompose how reimbursements change after a hospital system makes an acquisition, where we allow for changes in bargaining weights, and bargaining positions (as in Lewis and Pflum 2015) and, importantly for our research question, quantify the pass-through of the cost synergies that we document in our reduced form evidence in the previous sections.

The model consists of two stages. In the first stage, insurers negotiate with the hospitals in their networks over the reimbursement rates for outpatient care. In the second stage, patients requiring outpatient care choose a hospital within their insurer's network. We describe the components of the model by starting with the second stage and then working backward to the first.

5.1.1 Patient demand

In each market year, there exists a set of patients $n = 1, \cdots, N$ (we omit market year subscripts for notational convenience). Patients are covered by a set of insurers $i = 1, \cdots, I$; where applicable, we denote patient n's insurer as $i(n)$. Each insurer offers a network of covered outpatient hospitals $\mathcal{G}_i$, which is a subset of the total set of hospitals operating in the market $h = 1, \cdots, H$.

Patient n's indirect utility from receiving care at hospital h is specified as

$$u_{nhi} = \alpha_0 \cdot d_{nh} + \alpha_1 \cdot d_{nh}^2 + \alpha_2 \cdot \mathbf{X}_{ni}^1 \cdot \mathbf{Y}_h^1 + \alpha_3 \cdot \mathbf{X}_n^2 \cdot \mathbf{Y}_h^2 + \eta_h + \varepsilon_{nhi}. \tag{3}$$

This specification is similar to that used in Liu (2022), with the notable difference that we do not include indicators for hospital ownership by private equity firms, which is a central focus of his paper. Our specification is also similar to Prager and Tilipman (2020). Here, d_{nh} measures the travel time (in hours) between patient n's residence and hospital facility h's location. The term $\mathbf{X}_{ni}^1 \cdot \mathbf{Y}_h^1$ captures interactions between patient observables and hospital characteristics. $\mathbf{X_i^1}$ includes the average income in patients' ZIP code, travel time, relative service-mix weights, patient gender, an indicator for whether the patient visited any hospital in the past year, and indicators for insurance coverage by either UnitedHealth, Anthem, or CA Blue Shield. $\mathbf{Y}_h^1$ includes the number of hospital beds, for-profit status, and teaching status. Similarly, $\mathbf{X}_n^2 \cdot \mathbf{Y}_h^2$ captures additional interaction terms. $\mathbf{X}_n^2$ includes all variables in $\mathbf{X}_{ni}^1$ except those relating to insurance carriers. $\mathbf{Y}_h^2$ includes indicators for the presence of catheterization laboratories, magnetic resonance imaging services, neurological services, or neonatal intensive care units in a hospital. The term η_h represents hospital fixed effects, while ε_{nhi} is assumed to be independently and identically distributed, following the Type 1 Extreme Value distribution.

We model patient choices among hospitals in their local market that are within their insurance's network. The outside option corresponds to seeking care at a hospital outside of the market.[5] We assign the outside option a mean utility of zero.

Let $\delta_{nhi} = u_{nhi} - \varepsilon_{nhi}$, where δ_{nhi} represents patient n's mean utility. Conditional on requiring treatment, the probability that patient n goes to hospital h (among in-network hospitals) is given by the standard logit choice probability:

$$S_{nhi}(\mathcal{G}_i) = \frac{\exp(\delta_{nhi})}{1 + \sum_{\mathcal{G}_i} \exp(\delta_{nhi'})}. \tag{4}$$

Aggregating across enrollees, the total demand from insurer i for hospital h is

$$D_{hi} = \sum_{\{n:i(n)=i\}} \gamma_{vn}\gamma_{dn} S_{nhi}, \tag{5}$$

5. Our sample only consists of patients that live in California and receive outpatient hospital care in California.

where γ_{vn} is the probability that patient n will need outpatient care within the year, and γ_{dn} is the diagnosis probability, conditional on having a visit. Both probabilities are allowed to vary across patients.

The expected utility that patient n derives from their insurer's network $\mathcal{G}_i$ plays an important role in the bargaining model (e.g., Gowrisankaran, Nevo, and Town 2015 and Lewis and Pflum 2015). It is given by

$$W_{ni}(\mathcal{G}_i) = \gamma_{vn}\gamma_{dn} \log\left(\sum_{h\in\mathcal{G}_i} \exp(\delta_{nhi})\right). \tag{6}$$

5.1.2 Bargaining model

In each market-year pair, every hospital h belongs to a system denoted by the set $\mathcal{H}_h$ (for independent hospitals, this set is a singleton).[6] Each insurer i operating in the market negotiates reimbursement rates p_{hi} for every hospitals h in its network $\mathcal{G}_i$.

The payoff to hospital h for being part of insurer i's network is given by

$$\Pi_{hi}(\mathcal{G}_i) = (p_{hi} - c_h)D_{hi}(\mathcal{G}_i), \tag{7}$$

where $D_{hi}(\mathcal{G}_i)$ represents the demand for hospital h from enrollees of insurer i, which depends on insurer i's entire network ($\mathcal{G}_i$). Demand tends to be increasing in the size of the network.

Following Gowrisankaran, Nevo, and Town (2015) and Lewis and Pflum (2015), the payoff to insurer i from offering network $\mathcal{G}_i$ is characterized as

$$V_i(\mathcal{G}_i) = \beta \sum_{\{n:i(n)=i\}} W_{ni}(\mathcal{G}_i) - \sum_{h\in\mathcal{G}_i} p_{hi}D_{hi}(\mathcal{G}_i). \tag{8}$$

The first term on the right-hand side represents the dollar value of the aggregate expected utility of insurer i's enrollees under network $\mathcal{G}_i$, where β reflects the weight insurers place on their enrollees' expected utility from their network. The second term measures the outlays that insurer i must pay to hospitals in network $\mathcal{G}_i$.

We assume that the negotiation over reimbursement rate p_{hi} between hospital h and insurer i follows the Nash-in-Nash bargaining framework.[7] Under this framework, reimbursements are

6. Under this notation, if h and h' belong to the same hospital system, then the sets $\mathcal{H}_h$ and $\mathcal{H}_{h'}$ are identical.

7. If h belongs to a hospital system $\mathcal{H}$, the entire system simultaneously negotiates with insurer i. This case is discussed in Appendix C.1.

determined by solving

$$\max_{p_{hi}} [\Pi_{hi}(\mathcal{G}_i)]^{\alpha_h} [V_i(\mathcal{G}_i) - V_i(\mathcal{G}_i \setminus h)]^{1-\alpha_h} \tag{9}$$

where α_h is hospital h's bargaining weight.

We define the following differences to measure the effects of losing a hospital from the network:

$$\Delta W_{ni}(\mathcal{G}_i, h) := W_{ni}(\mathcal{G}_i) - W_{ni}(\mathcal{G}_i \setminus h),$$

which captures the reduction in expected utility for patient n when hospital h is removed from the network, and

$$\Delta D_{h'i}(\mathcal{G}_i, h) := D_{h'i}(\mathcal{G}_i) - D_{h'i}(\mathcal{G}_i \setminus h),$$

which reflects the change in demand for hospital h' when hospital h is excluded from insurer i's network. The solution to the Nash-in-Nash bargaining problem is characterized by a set of first-order conditions (FOCs). The FOC for hospital h and insurer i is

$$\Pi_{hi}(\mathcal{G}_i) = \alpha_h \left(\beta \sum_{\{n:i(n)=i\}} \Delta W_{ni}(\mathcal{G}_i, h) - \sum_{h' \neq h} p_{h'i} \Delta D_{h'i}(\mathcal{G}_i, h) - c_h D_{hi}(\mathcal{G}_i) \right). \tag{10}$$

This condition is standard in the literature.[8] It states that under Nash-in-Nash bargaining with transferrable utility, hospital h's gains from trade are equal to its bargaining-weighted share of the joint surplus gains (relative to disagreement payoffs). This occurs because negotiated reimbursements are used to transfer any gains from trade between the two parties.

5.2 Estimation

5.2.1 Demand Estimation

We use Maximum Likelihood Estimation (MLE) to estimate the patient demand model specified in Equation 3. The model is estimated separately for each market-year pair, allowing all coefficients to vary across market-year pairs. This approach ensures a high degree of flexibility in our demand estimation. For example, patients in Los Angeles may exhibit a different aversion to travel time than those in San Francisco or the Central Valley. Similarly, patient demographics can interact differently with hospital teaching status across market-year pairs.

Patients' aversion to travel time and their preferences for hospital characteristics are important determinants of the value they place on broader insurer networks. Following the literature, we assume that hospital-specific taste shocks are uncorrelated with hospital characteristics. This assumption eliminates the need for instrumental variables to estimate the coefficient on distance

8. See, for example, Ho and Lee (2017) or Dafny, Ho, and Lee (2019).

(e.g., Ho and Lee 2017 and Gowrisankaran, Nevo, and Town 2015). Under this assumption, patients' aversion to travel time is identified using within-market variation in patients' ZIP codes of residence, which generates variation in travel times and hospital choice probabilities.

To calculate how hospital demand depends on insurer networks, we estimate two additional probabilities showing up in Equation 5. These probabilities are estimated "offline." Specifically, γ_{vn} is derived from Census data, while γ_{dn} is estimated using a combination of Census data and our claims data. Further details are in Appendix C.2.

5.2.2 Supply: A Sufficient Condition for Aggregating Over Insurers

Our hospital price data are measured at the hospital level, which prevents us from directly estimating Equation 10, which uses hospital insurer level price data. To overcome this data limitation, we aggregate the FOCs in Equation 10 across insurers, following the approach in Lewis and Pflum (2015).

Suppose that we had hospital-by-insurer prices. We could still sum Equation 10 over insurers to derive an aggregate FOC that serves as the basis for estimation:

(11)

$$\sum_{\{i:h\in\mathcal{G}_i\}} \Pi_{hi}(\mathcal{G}_i) = \alpha_h \left(\beta \sum_{\{i:h\in\mathcal{G}_i\}} \sum_{\{n:i(n)=i\}} \Delta W_{ni}(\mathcal{G}_i, h) - \sum_{\{i:h\in\mathcal{G}_i\}} \sum_{h'\neq h} p_{h'i} \Delta D_{h'i}(\mathcal{G}_i, h) - c_h \sum_{\{i:h\in\mathcal{G}_i\}} D_{hi}(\mathcal{G}_i) \right).$$

The only difference between Equation 11 and Equation 10 is that here, all terms are summed across insurers whose networks include hospital h.

Now consider the implications of having only hospital prices (rather than hospital-by-insurer prices). For estimation, we use the following modified FOC:

(12)

$$\sum_{\{i:h\in\mathcal{G}_i\}} \Pi_{hi}(\mathcal{G}_i) = \alpha_h \left(\beta \sum_{\{i:h\in\mathcal{G}_i\}} \sum_{\{n:i(n)=i\}} \Delta W_{ni}(\mathcal{G}_i, h) - \sum_{\{i:h\in\mathcal{G}_i\}} \sum_{h'\neq h} p_{h'} \Delta D_{h'i}(\mathcal{G}_i, h) - c_h \sum_{\{i:h\in\mathcal{G}_i\}} D_{hi}(\mathcal{G}_i) \right).$$

We clarify, for the first time to our knowledge, a sufficient condition under which estimation using Equation 12 recovers the same parameters as estimation using Equation 11 or Equation 10. Comparing Equations 12 and 11, it is clear that the middle term in parentheses on the right-hand

side of Equation 11 is critical. The other terms do not include hospital prices and aggregate in a straight forward manner.

Proposition. In the Nash-in-Nash bargaining model specified in Section 5.1.2, estimation using the aggregate FOC (equation 12) and average hospital prices recovers the same parameters as estimation based on the standard Nash-in-Nash FOCs (equation 10) if

$$\frac{D_{h'j}(\mathcal{G}_j)}{\sum_{\{i:h'\in\mathcal{G}_i\}} D_{h'i}(\mathcal{G}_i)} = \frac{D_{h'j}(\mathcal{G}_j \setminus h)}{\sum_{\{i:h'\in\mathcal{G}_i\}} D_{h'i}(\mathcal{G}_i \setminus h)} \quad \text{for all } h' \neq h \text{ and for all insurers } j. \tag{13}$$

Proof. See Appendix C.3.

The most important aspect of this proposition is that it is empirically testable, as all terms in Equation 13 can be computed using estimates from our patient demand model. In Appendix C.4, we quantify the empirical analogs of the left- and right-hand sides of Equation 13. In our application, the correlation between the two exceeds 0.99. Additionally, we show in Appendix C.4 that with different coefficients, our demand model can generate substantial deviations from Equation 13, ensuring that the high correlation we obtain is not merely mechanical.

5.2.3 Supply Estimation

Given that the condition in Equation 13 appears to hold in our context, we proceed to use Equation 12 to estimate the bargaining weights and the weight that insurers place on enrollee utility.

Following Lewis and Pflum (2015) and Liu (2022), we parametrize hospital bargaining weights as a funtion of hospital characteristics:

$$\alpha_h = \omega \mathbf{X}_{ht} + \nu_{ht} \tag{14}$$

where $\mathbf{X}^3_{ht}$ are time-varying covariates (including time subscripts) that leverage the panel nature of our data.

Let GFT^h_t denote hospital h's gains from trade in its negotiations, corresponding to the left-hand side of Equation 12. Let $GFT^T_t(\beta)$ denote the total aggregate gains from trade for negotiations involving hospital h, corresponding to the parenthetical term on the right-hand side of Equation 12.[9] Combining Equations 12 and 14 yields

$$\nu_{ht} = \frac{GFT^h_t}{GFT^T_t(\beta)} - \omega \mathbf{X}_{ht}. \tag{15}$$

We estimate ω and β using the Genralized Method of Moments, which makes the sample analogs

9. Our notation emphasizes the fact that this term depends on β, a parameter to be estimated.

of our moment conditions align closely with the data. Our first set of moments takes the form:

$$\mathbb{E}[\nu_{ht}\mathbf{X}_{ht}] = 0 \tag{16}$$

where we assume that hospital characteristics in the model are exogenous to the error terms in the bargaining weight specification. These moment conditions provide as many equations as there are terms in ω. Importantly, while the model is described using independent hospitals for clarity, we appropriately account for hospital systems; the relevant equations are detailed in Appendix C.1.

To estimate the remaining parameter, β, which captures the weight insurers place on their enrollees' expected utility, we introduce an additional moment condition. This condition equates the medical loss ratios (MLRs) implied by our model with observed MLR data from California.

$$\mathbb{E}\left[MLR_t - \sum_i \gamma_{it} \frac{\sum_{\{h:h\in\mathcal{G}_i\}} p_{ht} D_{hit}}{\beta \sum_{\{n:i(n)=i\}} W_{nit(\mathcal{G}_{it})}}\right] = 0, \tag{17}$$

which is similar to the MLR condition in Ho and Lee (2017). Notably, this last moment condition does not depend on the parameter ω. This independence allows us to estimate the model sequentially: first, we use Equation 17 to pin down β, and then, with the estimate of β in hand, we estimate ω using Equation 16.

6 Results

This section reports our results. We first describe the estimates from our demand model. Then we discuss our bargaining model estimates. Last, we describe counterfactuals where we quantify pass-through of cost synergies from horizontal mergers.

6.1 Demand Estimates

We estimate our demand model separately for 96 HSA-year pairs. This gives our demand model substantial flexibility as the parameters are allowed to vary across both HSAs and years. To illustrate our demand estimates, Figure A5 plots the distribution of mean demand elasticities with respect to travel time at the HSA-year level. Our results indicate that demand elasticities with respect to travel time are negative in over 80% of HSA-year markets. The average elasticity across all HSA-years is -0.17, implying that a hospital located twice as far from a patient's home has a 17% lower probability of being chosen. To our knowledge, other studies have not reported equivalent estimates of elasticity in the case of outpatient care. Our estimates imply that patient demand for outpatient hospital services is less elastic with respect to travel time than for acute inpatient care (Gowrisankaran, Nevo, and Town 2015). This is intuitive since most patient visits in our sample are for elective procedures or consultations for nonacute conditions. In contrast, a large fraction

of inpatient acute hospital stays originate in the emergency department. Furthermore, since outpatient visits do not require an overnight stay, patients may put less weight on convenience.

6.2 Bargaining Model Estimates

This section discusses the estimated bargaining weights and their correlations with hospital characteristics. We estimate the Nash-in-Nash bargaining model for each hospital system in every year. Independent hospitals are treated as singleton systems. The vast majority of our estimated raw bargaining weights lie in the unit interval, although we do not impose this restriction during estimation. We estimate a negative raw bargaining weight for a hospital system if its profits are negative. Conversely, we estimate a raw bargaining weight greater than 1 when a hospital system's profits exceed the total joint surplus from its negotiations with insurers.

We first examine the patterns in estimated bargaining weights for the hospitals in our sample prior to any of the observed acquisitions. This helps discern whether acquirers and targets are well matched in bargaining strength at baseline. The average pre-acquisition bargaining weight in our sample is 0.64, which is similar in magnitude to the mean hospital bargaining weights reported in Lewis and Pflum (2015), which ranged from 0.59 to 0.70.[10]

Figure 5 highlights the difference in the bargaining weight between acquirers and targets involved in the acquisitions of independent hospitals. Recall that we estimate price increases only in these deals. It plots the mean estimated bargaining weights for the acquiring hospitals, targets, and the independent hospitals that remain independent. While the mean weight for targets is 0.25, the average for acquirers is 0.82, a difference of 0.57. In contrast, the mean weight for independent hospitals not acquired is 0.73, very close to that of the acquirers. These are unadjusted values and may partly reflect differences in market features and hospital characteristics. However, this pattern suggests that these deals matched acquirers with targets that would stand to benefit substantially from the acquirer's bargaining skill.

Next, we formally examine the difference in bargaining power between acquirers and targets in these deals. We project the raw bargaining weights on hospital characteristics using data from the pre-acquisition period (2013 and 2014).[11] Table 6 reports the corresponding coefficients. Standard errors are obtained by bootstrapping the estimation procedure 100 times. The excluded category is being a member of an acquiring system, hence all coefficients represent differences relative to this group.

All else equal, teaching hospitals have substantially larger bargaining weights relative to non-teaching hospitals. Having a higher share of Medicaid patients is associated with a lower bargain-

10. Lewis and Pflum (2015) also examine hospitals in California but focus on inpatient hospitals in 2007 and 2008, so the results need not line up closely.

11. Figure A6 plots a histogram showing the distribution of the fitted values of bargaining weights obtained from this projection.

ing weight. The coefficient is imprecisely estimated but implies a large effect: a 10% increase in Medicaid share is associated with a 0.07 lower value of the bargaining weight, about 10% of the overall mean. This is not surprising since higher Medicaid patient share will lead to a lower profit margin, all else equal, and reduce the hospital's capacity to maintain a large and specialized price negotiation team. Acquiring hospital systems have greater bargaining weights than non-acquiring systems, implying selection in which systems successfully execute acquisitions. As suggested by the raw means, the acquirers had much larger bargaining weights than the hospitals they targeted for acquisition. After adjusting for other factors, the estimated difference is 0.62, slightly greater than the unadjusted difference. It is worth emphasizing that the acquisitions in our data should be viewed as matches between two parties, and it is unclear whether the selection is driven primarily by the acquirer's or the target's motives.

Although not reported formally in the table, we obtain a statistically significant estimate of \$755 for β. Recall that this represents the marginal value insurers put on consumer utility from their network relative to their costs.

7 Counterfactuals

7.1 Price Effects Following Observed Acquisitions

We use our model estimates to quantify the magnitude of three mechanisms that underlie the changes in prices at the acquired target due to system ownership. The first mechanism is cost synergies generated by system ownership. Some portion of these savings may be passed through to insurers in the form of price reductions. We quantify these by reestimating equilibrium prices after resetting marginal costs to their predicted pre-acquisition levels using the DD point estimate.

The second mechanism is the increase in the target's bargaining leverage due to system ownership. This operates through a direct effect as well as an indirect effect. The system negotiates with the insurer as a single entity and is either all in or all out of its network. Therefore, there is a change in the disagreement point payoff for the insurer when bargaining with the acquiring system. This is the direct effect. Other hospitals in the market competing for the insurer's business will also change their prices in response to the change in the target's bargaining leverage. This is the indirect effect. The collective magnitude of both these mechanisms is estimated by "unraveling" the acquisition and allowing the target to negotiate with the insurer on its own (as in the pre-acquisition period) while retaining the system's bargaining weight value. We then reestimate the model and obtain equilibrium prices under this scenario.

Finally, the third mechanism is a change in the target's bargaining skill since it now benefits from the system's higher bargaining skill, reflected by its higher bargaining weight. To quantify the magnitude of this mechanism, we change the bargaining weight of the target from that of the acquirer back to its pre-acquisition level and reestimate equilibrium prices.

Figure 6 presents the results from this counterfactual exercise. Panel (a) shows model predicted prices under different scenarios for the acquired independent hospitals, while Panel (b) shows prices for the acquirer's existing hospitals. In both panels, separate columns correspond to distinct counterfactual scenarios. Counterfactual prices are calculated using Equation 12, the same equation we used to estimate the bargaining model. Here, we treat our estimated bargaining weights as data and use them to calculate prices.

Focusing on Panel (a), the first bar plots the mean predicted post-acquisition price at target hospitals.[12] This incorporates the net effect of all three mechanisms. We normalize all values against this net price, which is set to 100, so all changes can directly be interpreted as percent changes. The second bar plots counterfactual prices after eliminating cost savings and represents the gross price increase before offsets due to synergies. The counterfactual predicted price would be higher by a little over 1%. This implies that although system ownership leads to relatively large savings for the target hospital, these do not materially reduce prices set by the system for the target. We return to the reasons behind this seemingly counterintuitive finding in the next section. The third bar plots the mean predicted price at the target after unraveling all-or-nothing bargaining while retaining the system's bargaining weight. The mean price drops by about 2.5%. Although about twice the magnitude of the synergies channel, this is a small change relative to the net price effect discussed in Section 4. Finally, the fourth bar plots the mean price at the target when we also reset the target's bargaining weight to its pre-acquisition level. The predicted price drops dramatically by 36%. Overall, the change in bargaining weight accounts for over 90% of the change in price at the target due to the acquisition and is therefore the dominant mechanism.

Panel (b) plots the corresponding prices for hospitals already owned by the acquirer at the time of the deal. The figure shows that the acquirer's existing hospitals also experience an increase in price relative to the counterfactual of no acquisition, though it is much smaller at around 5%. Again, synergies offset only about 1%. In this case, the bargaining leverage and bargaining power channels contribute nearly equally to the change in price, each accounting for a change of about 2.5%. Intuitively, this occurs because we estimate a very small change in the acquirer's bargaining power following the acquisition.

7.2 Predicting The Pass-Through of Synergies

We now explain why substantial reductions in operating costs translate into modest price effects in our counterfactuals. By implicitly differentiating the hospital system version of the Nash-in-Nash FOC (in Appendix C.1), we derive $\partial p_s / \partial c_h$, the effect of a cost change at hospital h on the system price. Ignoring equilibrium effects of cost changes on other hospital systems' prices provides the

12. Our model predicts lower prices than we observe in the data because we winsorize bargaining weights to ensure they remain within the unit interval. This approach helps prevent our counterfactuals from being unduly influenced by hospitals systems with outlier bargaining weights. Hence, we focus on percent changes here.

following approximation:[13]

$$\frac{\partial p_s}{\partial c_h} \approx \frac{\sum_{\{i:h\in\mathcal{G}_i\}} D_{hi}(\mathcal{G}_i)}{\sum_{h'\in\mathcal{H}_h}\sum_{\{i:h\in\mathcal{G}_i\}} D_{h'i}(\mathcal{G}_i)}(1-\alpha_s) \tag{18}$$

This equation says that system prices are more sensitive to hospital cost changes when: (1) the hospital experiencing the cost change accounts for a larger proportion of system demand (as captured by the fraction on the right-hand side), and (2) the hospital system has lower bargaining power (reflected by $1-\alpha_s$).

These conditions align with intuition. When a hospital with a substantial share of system demand experiences a cost change, it has a greater impact on the joint surplus of the negotiating parties, thereby influencing prices more significantly. Additionally, in the limit, if hospitals have no bargaining power (i.e., $\alpha_s = 0$), insurers place hospitals at their disagreement payoff, meaning that all cost changes are fully passed through to prices.

Equation 18 provide a useful framework for understanding the modest price effects observed in counterfactuals that undo cost synergies. Figure 7 plots the empirical distribution of the fraction on the right-hand side of Equation 18. In our data, acquired hospitals are typically small relative to their acquiring system, which empirically explains the limited price effects we observe.

The final set of calculations in this section quantifies the pass-through of cost synergies using an *alternative* approach. Instead of examining the effect of cost synergies on negotiated prices, we calculate how the implied reductions in aggregate costs translate into changes in hospital system profits, and ultimately, into changes in insurer reimbursements. The results are presented in Figure 8. Our focus is on acquiring systems post-acquisition, including the acquired hospitals in the calculations.

Panel (a) of Figure 8 shows the distribution of changes in hospital system profits under the counterfactual with cost synergies relative to the counterfactual without cost synergies. While a substantial number of hospital systems experience no change in profits at the system level post-acquisition, some systems see profit increases of 2.5% or even 5%.

Panel (b) plots the distribution of the value given by 100% minus the ratio of change in profits (from Panel (a)) to the change in costs. If hospitals systems retained all cost savings as profits, this measure would equal 0%. Higher values indicate that a greater share of cost savings is passed through to insurers via lower reimbursements. The histogram reveals a somewhat bimodal distribution. In about 30% of system-years, less than 25% of cost savings are passed through to insurers. At the same time, more than 50% of aggregate savings are passed through in about half the system-years. On average, the pass-through rate is 46%, with a median of 51%.

13. This is an approximation as it neglects equilibrium price effects.

7.3 A Systematic Analysis of Potential Deals

Counterfactual simulations of the observed deals imply that the major share of the price increase following the acquisition of independent hospitals is due to a change in the target's bargaining power and not due to change in bargaining leverage. However, we also noted that there was a large gap in the bargaining power of the acquirers and targets in the realized deals. Hence, this pattern may reflect the selected nature of the realized deals rather than being a generalizable result for acquisitions of independent hospitals by systems. In addition, a practical wrinkle with evaluating the price effects of realized acquisitions is that sometimes a market experiences multiple deals in the same year. Hence, the change in bargaining leverage in some cases may reflect the effects of more than just one focal deal. Ideally, we would assess the price effects of acquisitions one case at a time. We address both of these concerns in the counterfactual analysis in this section.

We perform counterfactual simulations of potential system acquisitions of hospitals that were independent in 2014. We evaluate one acquisition at a time to ensure that other aspects of the market are not changing. We make two restrictions when constructing the set of potential deals. First, the acquirer's bargaining weight, as estimated over 2013-14, must be equal to or greater than that of the target independent facility during the same period. This restriction ensures that the target benefits from a higher bargaining weight due to system ownership. Although this restriction rules out several possible deals, it is consistent with what we observe in practice. For brevity, we refer to deals that meet this criterion as being "feasible." Second, we restrict our attention to deals in which the target is located in the same market as at least one of the acquirer's existing hospitals. This ensures that the acquisition also lessens competition among hospitals in negotiations with insurers and we can compare the relative magnitudes of changes in bargaining leverage and bargaining power. To summarize, we now evaluate the model-predicted price effects of all feasible within-market system acquisitions of independent hospitals (as of 2014) in California. There are 136 such feasible deals, including the realized within-market deals. We report average price increases at the target hospitals over 2015–2020 in the figures below.

Figure 9 below presents the CDF of predicted price effects across the 136 deals evaluated in this exercise. The median and mean predicted price increase in this distribution are approximately \$130 and \$150 per visit. The vertical red line indicates the weighted average mean price increase for the realized within-market deals, \$252. The figure shows that the average price effect for the realized deals falls approximately at the 75^{th} percentile of the distribution of price effects across all feasible deals, consistent with the conjecture that the deals were chosen to obtain relatively large price increases. Note that this result understates the true percentile rank of the realized price effects within the universe of all possible deals since we do not evaluate deals expected to have low or negative price effects. If those were to be included, the average price effect for realized deals would likely occupy an even higher percentile value.

Next, we examine the correlation between predicted price effects and the difference in bargaining

weights between the acquirer and target prior to the acquisition. Figure 10 Panel (a) presents a binned scatter plot showing the mean predicted price effect on the Y-axis versus the mean difference in bargaining weight on the X-axis within each (patient-weighted) decile of the acquisitions studied here. The figure shows that the predicted price effect increases nearly linearly with the bargaining power advantage of the acquirer. An OLS regression of the predicted price effect on the difference in bargaining weights in the underlying deal-level sample recovers a precisely estimated slope coefficient of $440. This implies if a target is bought by a system with a bargaining weight higher by 0.55 (the average for the realized deals), the model predicts an average price increase of around $240 just based on this change alone. This is well over 90% of the net price effect of $252 predicted by the model. Hence, even in this larger sample of deals involving independent targets, selection by an acquirer with greater bargaining ability explains most of the price effect.

Figure 10 Panel (b) presents the corresponding relationship between price effects and the change in market concentration, as measured by the predicted change in the Herfindahl-Hirschman Index (HHI). In contrast to the top panel, we find essentially no correlation between the predicted change in HHI and price effect across these deals. The slope coefficient is small and statistically insignificant. This finding also reinforces the result from Section 7.1 that changes in bargaining leverage can explain only a small share of the total price effect. The other noteworthy takeaway from the plot is that most acquisitions of independent hospitals predict small changes in HHI. The mean predicted change in HHI in the top decile of deals is approximately an increase of 200, which was the prevailing statutory threshold for a presumptive anti-competitive acquisition during this period. This implies that the vast majority of potential acquisitions of independent hospitals in California during this period would not have been considered presumptively anti-competitive but are still predicted to generate substantial price increases, mostly due to the acquirer's bargaining power advantage.

Finally, we present a sanity check of our model's theoretical prediction on the pass-through of cost savings to prices presented in Equation 18. Figure 11 below presents a binned scatter plot showing the relationship between the model inferred change in price for each dollar of cost saving (dp/dc) on the Y-axis versus the target's share of total system outpatient visit volume on the X-axis. As usual, we plot the mean value in each patient-weighted decile, ordered by the target's volume share. The figure shows that as the target's share of total system patient volume increases, the pass-through of cost savings also increases. The slope coefficient is precisely estimated at 0.3. This implies that, on average, if an independent target accounts for 10% of the acquirer's total volume, the negotiated price will decrease by 3 cents for every dollar of cost saving at the target. To be clear, this is *not* an empirical confirmation of our model's prediction since we do not directly observe the pass-through of cost savings. However, it is reassuring that the intuitive comparative statics result, which ignores equilibrium effects, largely captures most of the variation in the pass-through of savings.

8 Discussion and Conclusion

This paper uses novel data and structural analysis to shed light on price setting in the hospital market and the pass-through of cost synergies generated from hospital consolidation to insurers. In the absence of quality improvements, potential price reductions due to cost synergies are the main benefit of consolidation for consumers. Hence, the division of these synergies between hospitals and consumers is a necessary input for policymakers concerned about the negative impacts of consolidation.

We examine the effects of acquisitions of independent and system-owned hospitals by systems in California during 2014–18. We find that when systems acquire independent hospitals, it generates cost savings as well as a net price increase. Hence, we estimate a structural model of hospital-insurer bargaining to help decompose the magnitude of different mechanisms that produce the net price effect. Subsequent simulations of counterfactual prices using model estimates show that, although systems pass through approximately 50% of the aggregate cost savings to insurers in the form of reduced prices, system prices hardly budge. Intuitively, this occurs because the acquired independent hospitals tend to be small relative to the acquiring systems and savings from the acquired facility are diluted to be negligibly small when spread across all hospitals of the acquirer. Overall, our results imply that the trend of increasing consolidation of independent hospitals is unlikely to benefit consumers even though it generates substantial cost savings for the acquirers.

Although this was a first step in understanding the importance of synergies in hospital system acquisitions, there are several directions for future work. Due to data limitations, we could not incorporate additional mechanisms, such as price recapture effects, in this analysis. Furthermore, researchers should consider examining these effects on a larger sample of deals, ideally covering other states, to obtain a nationally representative set of estimates.

References

Andrade, Gregor, Mark Mitchell, and Erik Stafford. 2001. "New evidence and perspectives on mergers." *Journal of Economic Perspectives* 15 (2): 103–120.

Andreyeva, Elena, Atul Gupta, Catherine Ishitani, Malgorzata Sylwestrzak, and Benjamin Ukert. 2024. "The corporatization of independent hospitals." *Journal of Political Economy Microeconomics* 2 (3): 602–663.

Blonigen, Bruce A, and Justin R Pierce. 2016. *Evidence for the effects of mergers on market power and efficiency.* Technical report. National Bureau of Economic Research.

Borusyak, Kirill, Xavier Jaravel, and Jann Spiess. 2024. "Revisiting event-study designs: Robust and efficient estimation." *Review of Economic Studies,* rdae007.

Brot, Zarek, Zack Cooper, Stuart V Craig, and Lev Klarnet. 2024. "Is there too little antitrust enforcement in the US Hospital sector?" *American Economic Review: Insights* 6 (4): 526–542.

Cooper, Zack, Stuart V Craig, Martin Gaynor, and John Van Reenen. 2019. "The price ain't right? Hospital prices and health spending on the privately insured." *Quarterly Journal of Economics* 134 (1): 51–107.

Craig, Stuart V, Matthew Grennan, and Ashley Swanson. 2021. "Mergers and marginal costs: New evidence on hospital buyer power." *RAND Journal of Economics* 52 (1): 151–178.

Dafny, Leemore. 2009. "Estimation and identification of merger effects: An application to hospital mergers." *Journal of Law and Economics* 52 (3): 523–550.

Dafny, Leemore, Mark Duggan, and Subramaniam Ramanarayanan. 2012. "Paying a premium on your premium? Consolidation in the US health insurance industry." *American Economic Review* 102 (2): 1161–1185.

Dafny, Leemore, Kate Ho, and Robin S Lee. 2019. "The price effects of cross-market mergers: Theory and evidence from the hospital industry." *RAND Journal of Economics* 50 (2): 286–325.

Demirer, Mert, and Ömer Karaduman. 2024. *Do Mergers and Acquisitions Improve Efficiency? Evidence from Power Plants.* Technical report. National Bureau of Economic Research.

Duggan, Mark G. 2000. "Hospital ownership and public medical spending." *Quarterly Journal of Economics* 115 (4): 1343–1373.

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

Gaynor, Martin, Kate Ho, and Robert Town. 2015. "The industrial organization of health-care markets." *Journal of Economic Literature* 53 (2): 235–84.

Gaynor, Martin, Adam Sacarny, Raffaella Sadun, Chad Syverson, and Shruthi Venkatesh. 2023. "The anatomy of a hospital system merger: the patient did not respond well to treatment." *Review of Economics and Statistics,* 1–28.

Goodman-Bacon, Andrew. 2021. "Difference-in-differences with variation in treatment timing." *Journal of Econometrics.*

Gorman, Bela, Dan Gorman, Jennifer Smagula, John D Freedman, Gabriella Lockhart, Rik Ganguly, Alyssa Ursillo, Paul Crespi, and David Kadish. 2016. "Why Are Hospital Prices Different? An Examination of New York Hospital Reimbursement." *Working Paper.*

Gowrisankaran, Gautam, Aviv Nevo, and Robert Town. 2015. "Mergers when prices are negotiated: Evidence from the hospital industry." *American Economic Review* 105 (1): 172–203.

Grennan, Matthew. 2014. "Bargaining ability and competitive advantage: Empirical evidence from medical devices." *Management Science* 60 (12): 3011–3025.

Ho, Kate, and Robin S Lee. 2017. "Insurer competition in health care markets." *Econometrica* 85 (2): 379–417.

Lewis, Matthew S, and Kevin E Pflum. 2015. "Diagnosing hospital system bargaining power in managed care networks." *American Economic Journal: Economic Policy* 7 (1): 243–74.

———. 2017. "Hospital systems and bargaining power: Evidence from out-of-market acquisitions." *RAND Journal of Economics* 48 (3): 579–610.

Liu, Tong. 2022. "Bargaining with private equity: Implications for hospital prices and patient welfare." *Working Paper.*

Maksimovic, Vojislav, and Gordon Phillips. 2002. "Do conglomerate firms allocate resources inefficiently across industries? Theory and evidence." *Journal of Finance* 57 (2): 721–767.

Maksimovic, Vojislav, Gordon Phillips, and Nagpurnanand R Prabhala. 2011. "Post-merger restructuring and the boundaries of the firm." *Journal of Financial Economics* 102 (2): 317–343.

MedPAC. 2020. *Report to the Congress: Medicare Payment Policy: Chapter 3.* Technical report. U.S. Department of Health and Human Services, March.

Miller, Nathan H, and Matthew C Weinberg. 2017. "Understanding the price effects of the MillerCoors joint venture." *Econometrica* 85 (6): 1763–1791.

Oh, Hyesung, Vincent Mor, Daeho Kim, Andrew Foster, and Momotazur Rahman. 2025. "Hospital Mergers and Acquisitions From 2010 to 2019: Creating a Valid Public Use Database." *Health Services Research,* e14642.

Panhans, Matthew, Ted Rosenbaum, and Nathan E Wilson. 2021. "Prices for medical services vary within hospitals, but vary more across them." *Medical Care Research and Review* 78 (2): 157–172.

Prager, Elena, and Nicholas Tilipman. 2020. "Regulating out-of-network hospital payments: Disagreement payoffs, negotiated prices, and access."

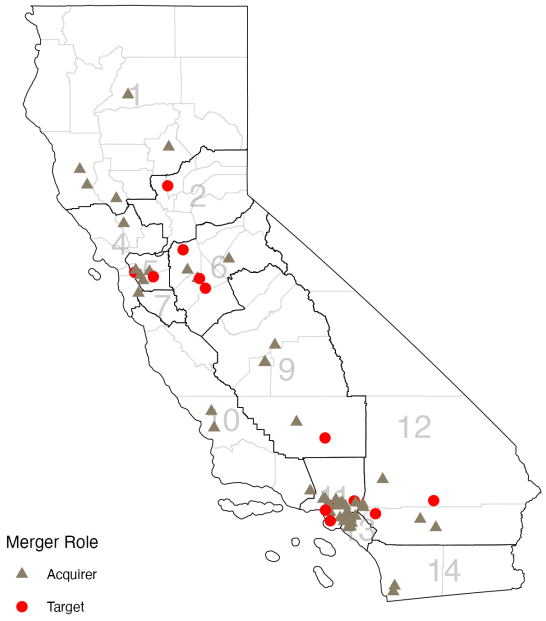
Schmitt, Matt. 2017. "Do hospital mergers reduce costs?" *Journal of Health Economics* 52:74–94.

———. 2018. "Multimarket contact in the hospital industry." *American Economic Journal: Economic Policy* 10 (3): 361–387.

Schoar, Antoinette. 2002. "Effects of corporate diversification on productivity." *Journal of Finance* 57 (6): 2379–2403.

Sun, Liyang, and Sarah Abraham. 2021. "Estimating dynamic treatment effects in event studies with heterogeneous treatment effects." *Journal of Econometrics* 225 (2): 175–199.

(a) Independent



(b) System

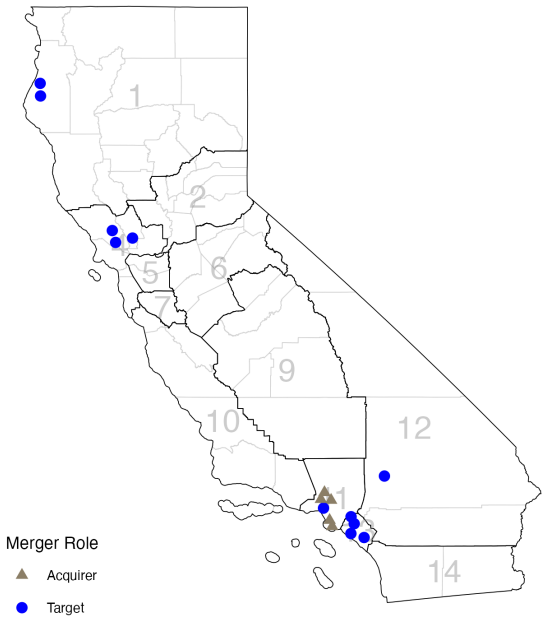


Figure 1: Hospital markets and acquisitions

Notes: This figure depicts the boundaries of the hospital markets in California used in the analysis. These markets are based on Health Service Areas (HSAs) defined and numbered by California Health Care Access and Information (HCAI) to study healthcare markets. The maps also depict the locations of acquiring and target hospitals involved in the deals studied in the article. Panels (a) and (b) present the locations of facilities involved in acquisitions of independent and system-owned hospitals, respectively.

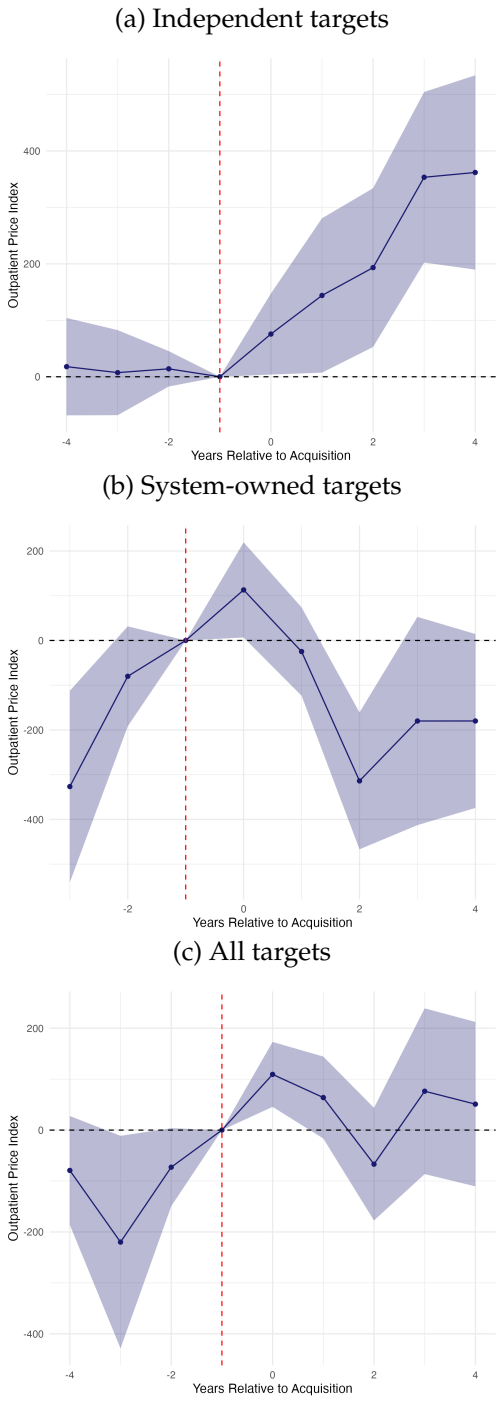
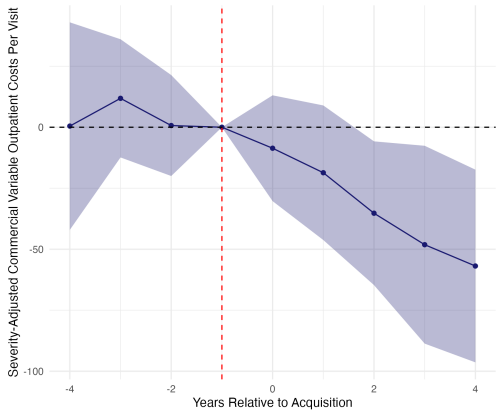


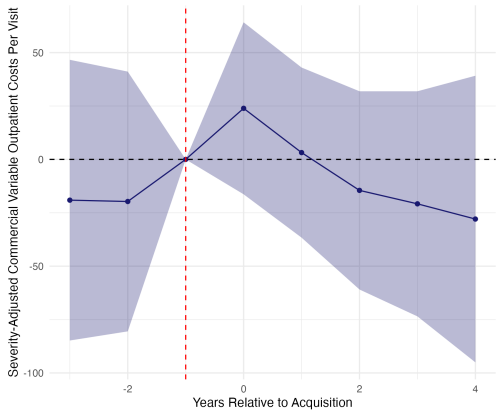
Figure 2: Dynamic Effects on Outpatient Prices

Notes: This figure presents event study estimates of the effect of a system acquisition on hospital outpatient prices. Panels (a), (b), and (c) present the estimates for independent targets, system-owned targets, and all targets, respectively. The outpatient price index is adjusted for severity, demographics, and payer composition, as described in the appendix. The models include time-varying controls such as the Medicare share of visits, the number of beds, and the percentage elderly population in the hospital's HRR. The shaded region indicates 95% confidence intervals. Standard errors are clustered by hospital.

(a) Independent targets



(b) System-owned targets



(c) All targets

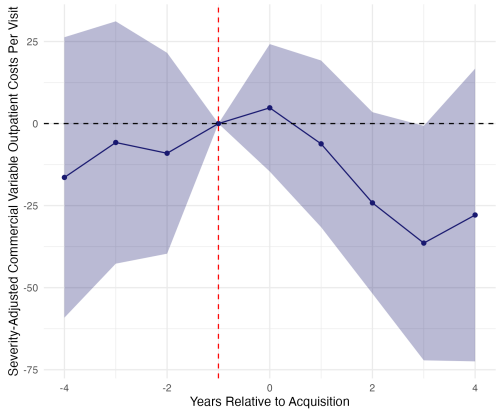
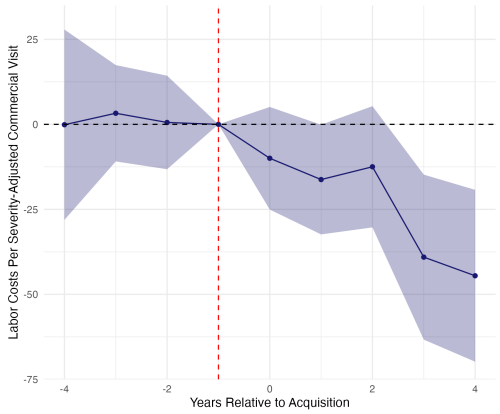


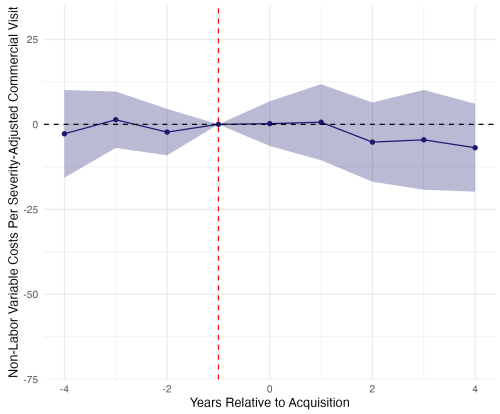
Figure 3: Dynamic Effects on Variable Costs

Notes: This figure presents event study estimates of the effect of a system acquisition on hospital outpatient cost per visit. Panels (a), (b), and (c) present the estimates for independent targets, system-owned targets, and all targets, respectively. The hospital outpatient cost per visit is restricted to variable costs and adjusted for severity, as described in the appendix. The models include time-varying controls such as the Medicaid share of visits, the number of beds, and the percentage elderly population in the hospital's HRR. The shaded region indicates 95% confidence intervals. Standard errors are clustered by hospital.

(a) Labor Expenses



(b) Other Expenses, Variable



(c) Other Expenses, Fixed

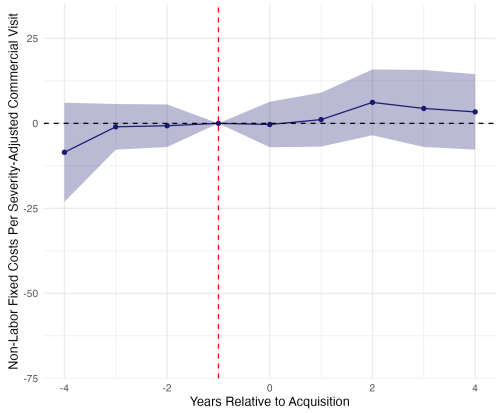


Figure 4: Dynamic Effects on Cost Per Visit, Independent Targets

Notes: This figure presents event study estimates of the effect of a system acquisition on three mutually exclusive and collectively exhaustive components of operating cost at independent target hospitals. Panels (a), (b), and (c) present the results for labor costs, non-labor variable costs, and non-labor fixed costs, respectively. The models include time-varying covariates as described in Section 4.1. The shaded area represents 95% confidence intervals. Standard errors are clustered by hospital.

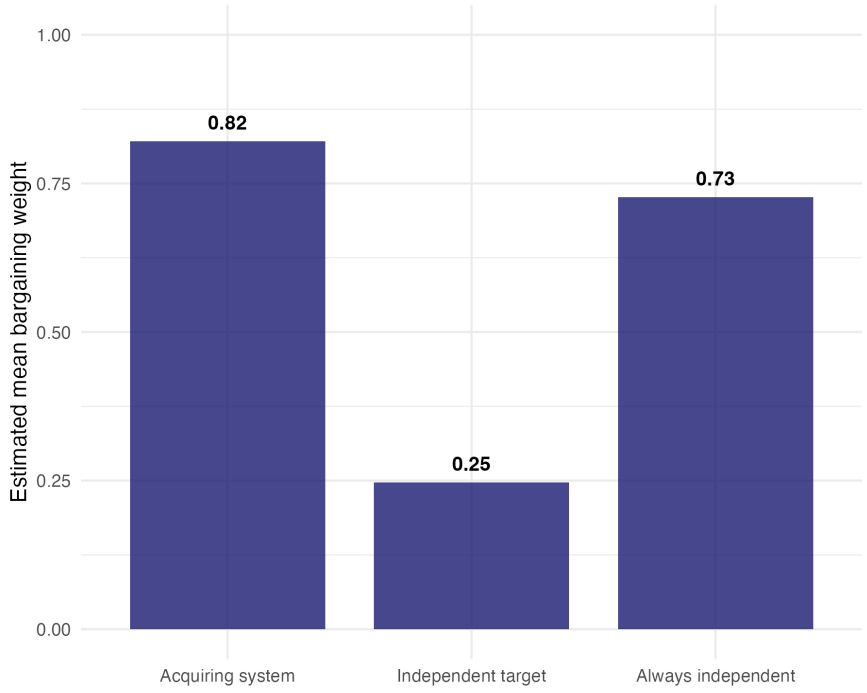
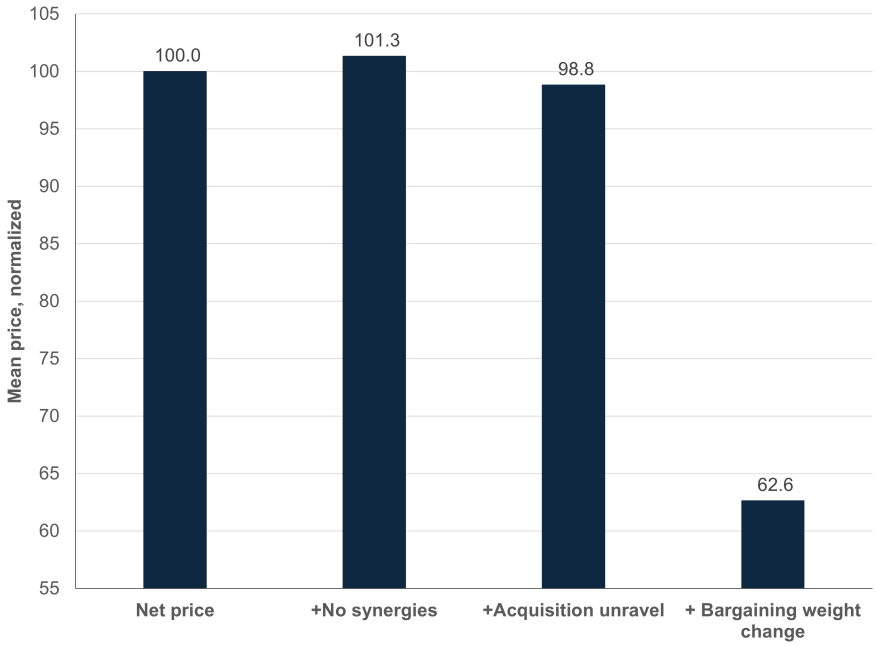


Figure 5: Mean Bargaining Weights at Baseline

Notes: This figure presents mean bargaining weights by system type in the baseline years (2013 and 2014). The figure uses raw bargaining weights, which correspond to to the ratio of hospital gains-from-trade to total gains-from-trade. Bargaining weights are winsorized to lie in the unit interval.

(a) Acquired Independent Hospitals



(b) Acquirer's Existing Hospitals

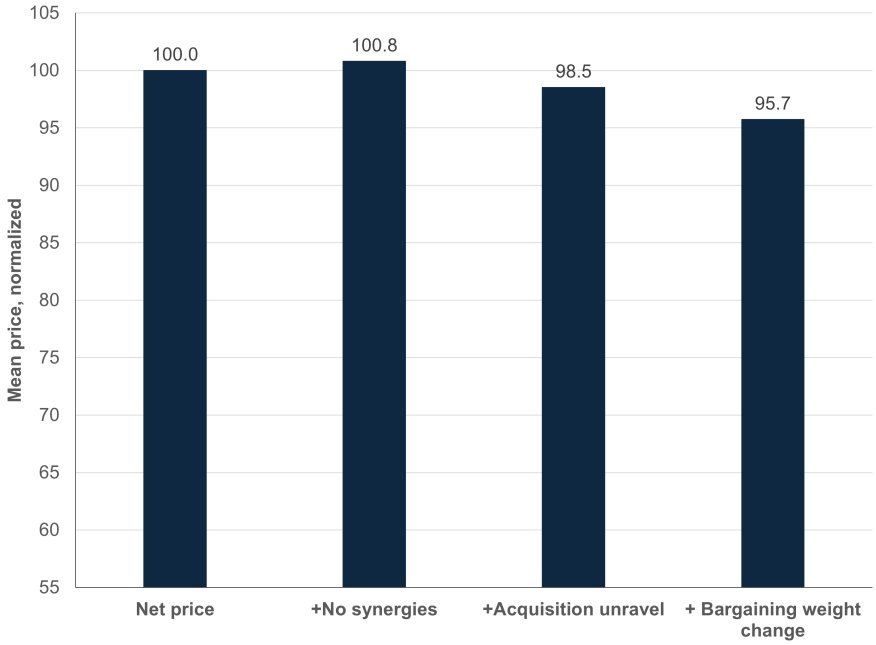


Figure 6: Counterfactual prices for the observed acquisitions

Notes: This figure presents the mean price in each counterfactual scenario, scaled by the predicted price under status quo on the far left as a benchmark. Panel (a) presents the results for the acquired hospital group, and Panel (b) presents the results for hospitals which are members of acquiring systems. The "Net price" column presents the simulated prices predicted by the bargaining model, holding bargaining weights, ownership, and costs fixed. Bargaining weights are winsorized at 0 and 1. The "No synergies" column reverses the cost efficiencies according to the reduced-form average treatment effect on costs. The penultimate column unravels all acquisitions, but keeps bargaining weights the same for formerly acquired hospitals. The final column changes bargaining weights for the acquired hospitals to their pre-acquisition levels.

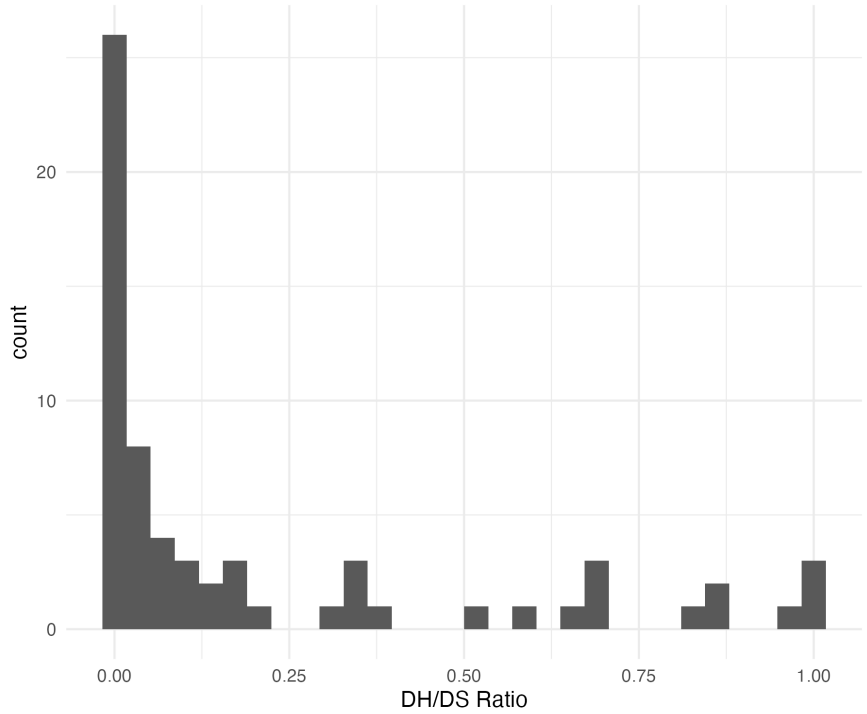
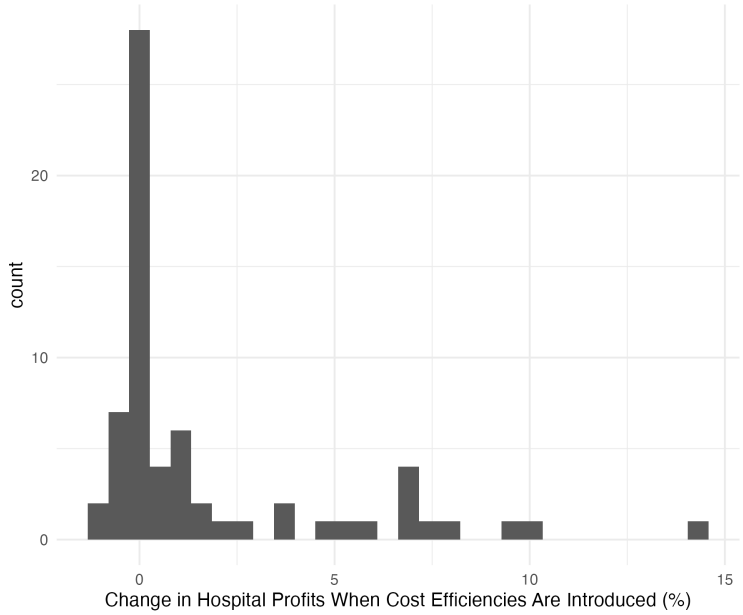


Figure 7: DH/DS Distribution of Acquired Hospitals

Notes: This figure presents the distribution of the fraction on the right hand side of Equation 18. An observation is a hospital year pair and we use data for acquired independent hospitals post acquisition. A ratio of one corresponds to an out–of–market acquisition, where the acquiring system had no presence in the pre–period.

(a) Profit Change After Cost Efficiencies



(b) Pass-Through of Cost Efficiencies to Hospital Profits

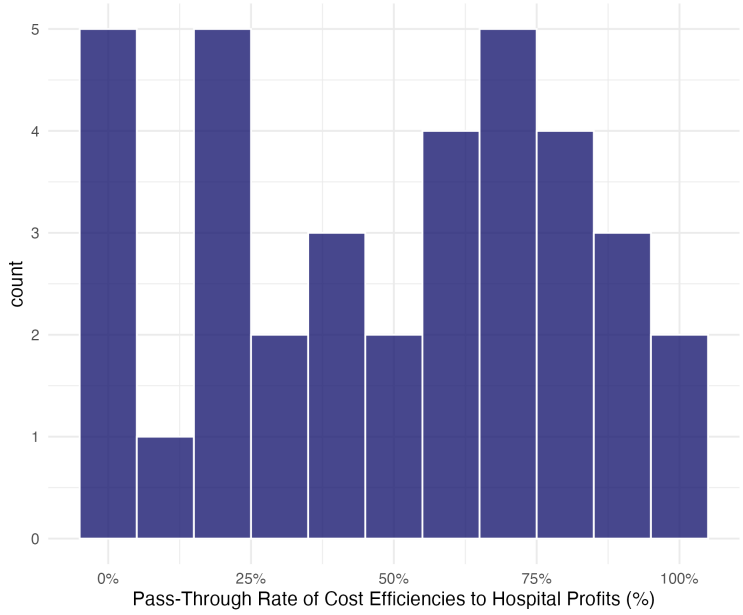


Figure 8: Pass–Through of Cost Savings to Insurers

Notes: This figure presents three measures of the pass-through of cost efficiencies to consumers by hospitals. All three compare hospital profits and costs between the "Model" and "No Cost Efficiencies" counterfactual scenarios. Bargaining weights are winsorized at 0 and 1. All observations are at the system-year level, and all system-years are for acquiring systems post-acquisition. Panel (a) shows the percent change in system profits when cost efficiencies are introduced relative to the no cost efficiencies scenario. Panel (b) plots one minus the ratio, in percentage terms, of the change in hospital profits to the change in hospital costs.

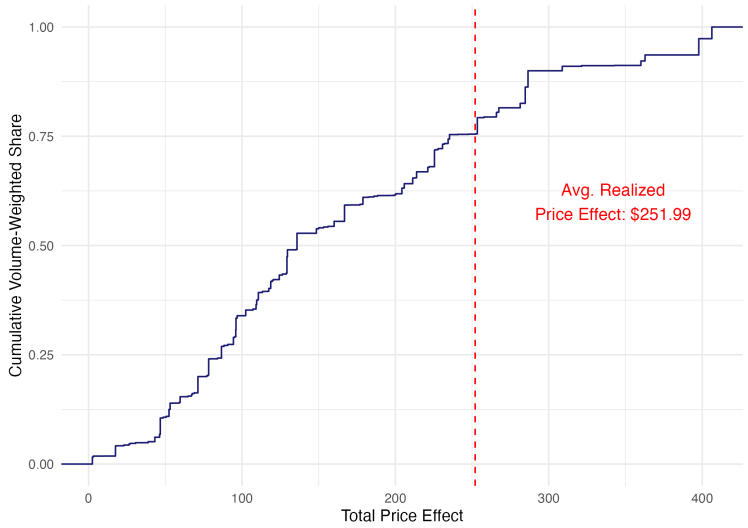
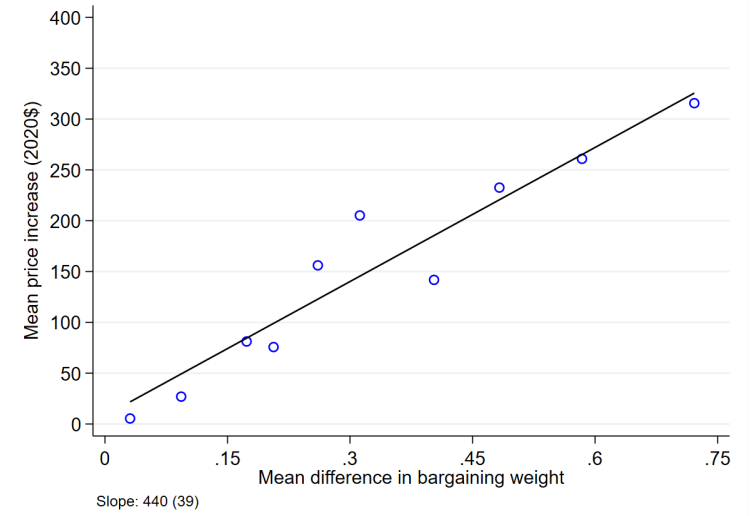


Figure 9: Price Effects of Counterfactual Acquisitions

Notes: This figure presents the volume-weighted distribution of total price effects of counterfactual independent hospital acquisitions in California that could have occurred in 2015, with results shown averaged from 2015–2020. Acquisitions are restricted to be within-market and the acquirer must have a bargaining weight at least as large as the target's. The price effect is the difference between the target's price after acquisition (including cost synergies) and the target's original price. Observations are weighted by the hospital's outpatient visit volume from the California cost reports. The dashed red line represents the average (simulated) price effect for the realized acquisitions, weighted by the volume of the acquired hospitals.

(a) Total Price Effect by Bargaining Weight Difference



(b) Total Price Effect by HHI Change

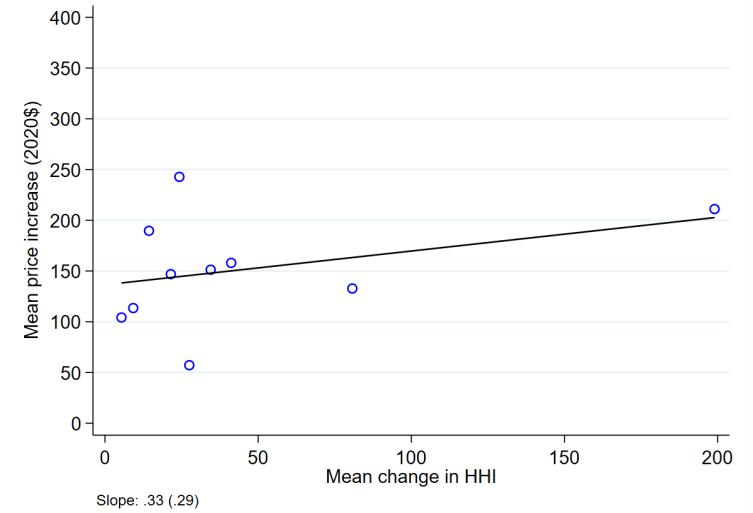


Figure 10: Price Effects from Hypothetical Acquisitions, 2015–2020

Notes: This figure shows bin scatter plots to depict the relationship between the net price effect of system acquisitions of independent hospitals and the acquirer's bargaining power advantage relative to the target (Panel a) or the predicted change in market concentration due to the acquisitions (Panel b). We analyze 136 within-market deals where the system's bargaining weight was the same or higher than that of the target that could have occurred in California in 2015, including the realized deals. The price effect is the difference between the target's price after acquisition (including cost synergies) and the target's original price. Acquisitions are assigned to patient-weighted deciles based on the target's share of total system volume of outpatient visits. The slope coefficient is obtained from an OLS regression of the X-axis variable on the net price effect using the underlying deal-level data.

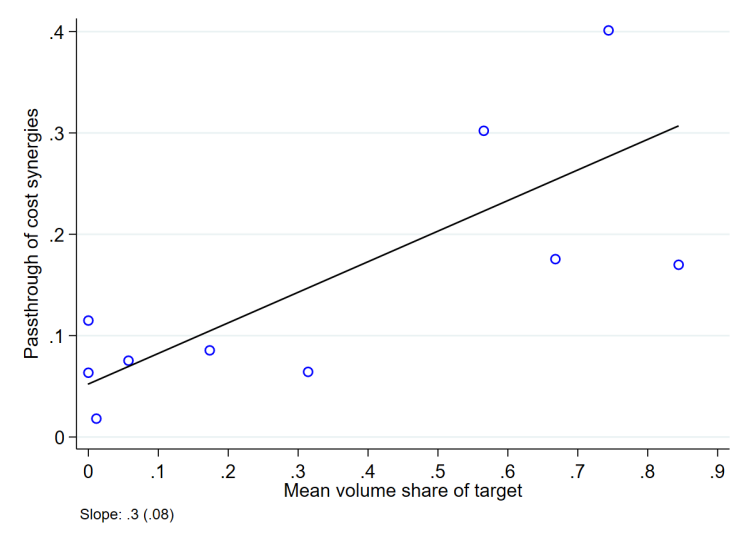


Figure 11: Pass-Through of Cost Savings by Volume Share

Notes: This figure plots mean values of the change in target's price per dollar of cost savings (dp/dc) on the Y-axis against the corresponding mean of the target's share of total system volume on the X-axis in each decile of 136 within-market deals where the system's bargaining weight was the same or higher than that of the target that could have occurred in California in 2015, including the realized deals. Deals are assigned to patient-weighted deciles based on the target's share of total system volume of outpatient visits. The slope coefficient is obtained from an OLS regression of target's volume share on cost pass-through using the underlying deal-level data.

Table 1: Descriptive Statistics

	(1) All Hospitals	(2) Acquiring System (Ind.)	(3) Acquiring System (Sys.)	(4) Independent	(5) Independent Target	(6) System-Owned	(7) System-Owned Target
No. of Hospitals	297	50	5	91	12	128	11
No. of Systems	144	8	1	91	12	31	2
Hospital Characteristics							
Beds	201.33	189.72	328.20	144.76	188.42	239.09	239.18
Fraction Teaching	0.06	0.04	0.00	0.02	0.00	0.10	0.00
Medicare Share of Visits	0.31	0.28	0.36	0.32	0.30	0.32	0.34
Medi-Cal Share of Visits	0.25	0.29	0.21	0.27	0.28	0.24	0.12
Fraction For-Profit	0.22	0.52	0.00	0.11	0.17	0.21	0.00
Fraction Rural	0.05	0.00	0.00	0.12	0.00	0.02	0.00
Prices, Costs, and Service Mix							
Standardized Price Per Visit	397.53	349.10	482.43	509.77	196.85	341.46	449.92
Standardized Cost Per Visit	105.48	73.16	72.24	117.37	100.14	113.92	87.48
Average APC Weight	8.09	8.84	10.84	6.94	9.34	8.36	8.04
Percentage Variable Costs	76.21	71.52	71.52	79.51	78.64	76.18	70.63

Notes: This table presents descriptive statistics on the hospitals observed in the analysis sample. The values are taken from the first year a hospital is observed in the sample, which is 2013 for most hospitals. Column 1 describes all hospitals in the sample. The remaining columns describe the following subgroups: Column 2, acquirers of independent targets; Column 3, acquirer of system-owned targets; Column 4, independent hospitals used as a comparison group for independent targets; Column 5, independent target hospitals; Column 6, system-owned hospitals used as a comparison group for system-owned targets; and Column 7, system-owned target hospitals. The price and variable cost values only reflect outpatient care and are standardized for a procedure of APC weight 1. The price and cost values are expressed in 2020 dollars. The hospital's average APC weight is taken from the claims sample.

Table 2: Difference-in-Differences Estimates on Outpatient Prices

	Independent Targets		System-owned Targets		All Targets	
	No Controls (1)	Controls (2)	No Controls (3)	Controls (4)	No Controls (5)	Controls (6)
Acquired x Post	239.55^{***}	220.94^{***}	-110.48^{*}	-118.85^{*}	32.77	44.58
	(60.32)	(66.30)	(58.73)	(61.68)	(49.79)	(48.10)
No. Obs.	485	485	421	406	906	891
Pre-acquisition Mean	145.33	145.33	610.84	610.84	424.64	424.64

$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$

Notes: This table presents difference-in-differences estimates of the effect of system acquisition on prices for the acquired hospitals. Columns 1–2, 3–4, and 5–6 present the average effect for independent targets, system-owned targets, and all targets, respectively. We present results from three specifications in each panel. Odd numbered columns present results from models without hospital covariates. Even columns presents results from a models that uses covariates. Standard errors are clustered by hospital. Prices are deflated to 2020 dollars. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$

Table 3: Difference-in-Differences Estimates on Variable Costs

	Independent Targets		System-owned Targets		All Targets	
	No Controls (1)	Controls (2)	No Controls (3)	Controls (4)	No Controls (5)	Controls (6)
Acquired x Post	-35.99^{***}	-31.70^{**}	3.76	-6.97	-16.83	-17.05
	(13.25)	(14.44)	(14.03)	(22.00)	(11.78)	(13.56)
No. Obs.	690	690	528	528	1,218	1,218
Pre-acquisition Mean	110.17	110.17	102.52	102.52	106.51	106.51

$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$

Notes: This table presents difference-in-differences estimates of the effect of system acquisition on variable operating cost at the acquired hospitals. Columns 1–2, 3–4, and 5–6 present the average effect for independent targets, system-owned targets, and all targets, respectively. We present results from three specifications in each panel. Odd numbered columns present results from models without hospital covariates. Even columns presents results from a models that uses covariates. Standard errors are clustered by hospital. Costs are deflated to 2020 dollars. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$

Table 4: Difference-in-Difference Estimates on Employment and Wages

(a) Employment

	Independent Targets		System-owned Targets		All Targets	
	Log Admin FTEs	Log Non-admin FTEs	Log Admin FTEs	Log Non-admin FTEs	Log Admin FTEs	Log Non-admin FTEs
	(1)	(2)	(3)	(4)	(5)	(6)
Acquired x Post	-0.17^{**}	-0.08^{*}	-0.14	-0.03	-0.16^{***}	-0.05
	(0.07)	(0.04)	(0.15)	(0.05)	(0.06)	(0.03)
No. Obs.	714	715	530	530	1,244	1,245
Pre-acquisition Mean	66.90	596.76	127.23	1,033.45	95.76	805.61

$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$

(b) Wages

	Independent Targets		System-owned Targets		All Targets	
	Admin Wages	Non-admin Wages	Admin Wages	Non-admin Wages	Admin Wages	Non-admin Wages
	(1)	(2)	(3)	(4)	(5)	(6)
Acquired x Post	-1.67	0.29	-1.61	1.48	-0.78	1.29^{**}
	(1.31)	(0.72)	(2.42)	(0.95)	(1.02)	(0.59)
No. Obs.	714	715	530	530	1,244	1,245
Pre-acquisition Mean	58.27	37.55	66.10	48.44	62.01	42.76

$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$

Notes: This table presents difference-in-differences estimates of the effect of a system acquisition on employment and wages. Panel (a) uses log employment of administrators and all other workers (in FTEs) as the dependent variable, while Panel (b) uses the respective hourly wages as the dependent variable. Columns 1–2, 3–4, and 5–6 present the average effect for independent targets, system-owned targets, and all targets, respectively. All models include covariates. Standard errors are clustered by hospital. Costs and prices are deflated to 2020 dollars. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$

Table 5: Difference-in-Differences Estimates on Quality

	AMI		Heart Failure		Pneumonia	
	Mortality Rate (1)	Readmission Rate (2)	Mortality Rate (3)	Readmission Rate (4)	Mortality Rate (5)	Readmission Rate (6)
Acquired x Post	0.26	0.19	0.15	-0.07	-0.08	0.33
	(0.18)	(0.30)	(0.27)	(0.29)	(0.36)	(0.27)
No. Obs.	376	320	590	602	628	628
Pre-acquisition Mean	12.93	16.75	11.60	22.52	14.09	17.21

$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$

Notes: This table presents difference-in-differences estimates of the effect of system acquisition on quality at independent acquired hospitals. We source data on risk-adjusted readmission and mortality rates from CMS Hospital Compare. Since CMS uses a 3-year rolling period to compute these measures, we need a rule to consistently assign them to a single year in the regression. The rule is that a value calculated using patient cases from the years T-1, T, and T+1 is assigned to year T in our sample. All specifications include time-varying covariates. Standard errors are clustered by hospital. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$

Table 6: Bargaining Weights and Hospital Characteristics

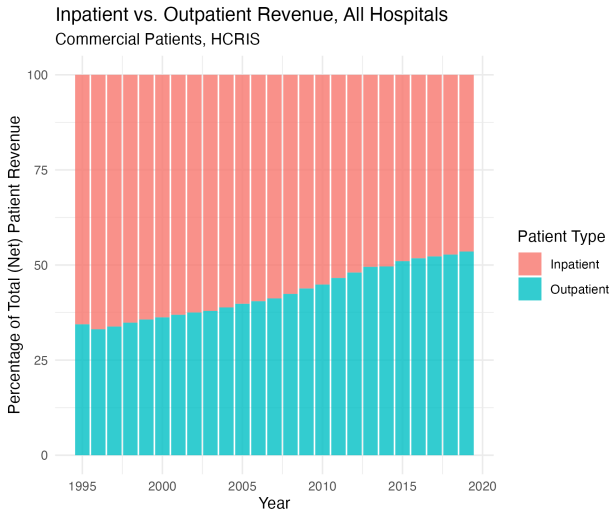
	Bargaining Weight
Teaching	0.86^{***}
	(0.28)
Physician Group	0.02
	(0.15)
For-Profit	-0.08
	(0.32)
Rural	-0.08
	(0.17)
Medicare Share of Visits	-0.18
	(0.59)
Medicaid Share of Visits	-0.73
	(0.56)
Independent Hospital	0.12
	(0.24)
Non-Acquiring System	-0.35^{*}
	(0.20)
Target Hospital	-0.62^{*}
	(0.31)
Num. obs.	134
R^2	0.22
Mean Dep. Var.	0.64
HSA Fixed Effects?	Y
Year Fixed Effects?	Y

$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$

Notes: This table presents the results from a regression of the raw estimated bargaining weight on a set of hospital system characteristics, along with market and year fixed effects. The model is estimated only on data from 2013–2014. The target, independent, and non-acquiring system dummies are estimated relative to an acquiring system baseline. Standards errors are clustered by system.

A Appendix Figures

(a) All states



(b) California

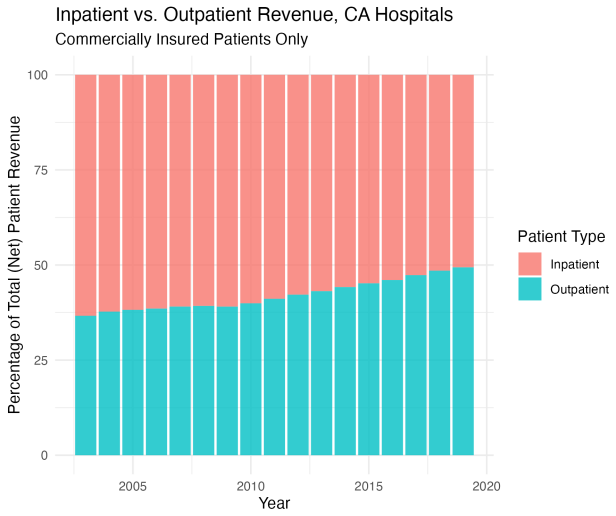
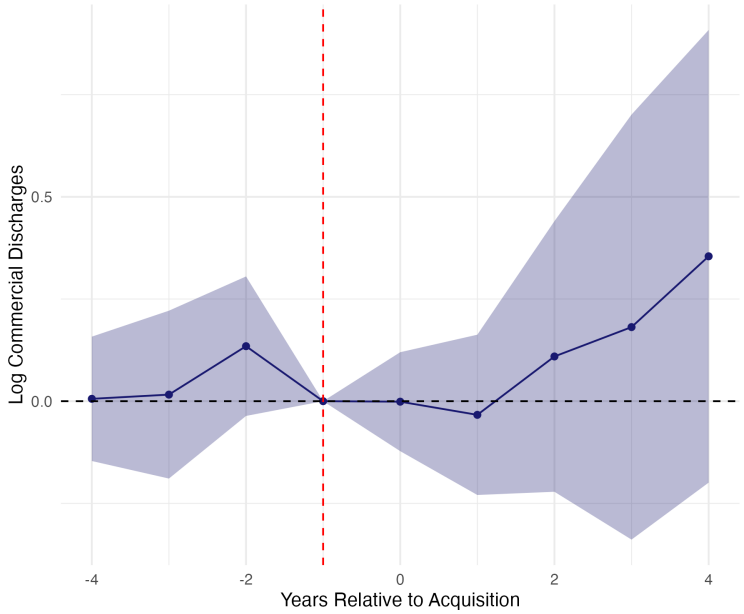


Figure A1: Inpatient Vs. Outpatient Commercial Revenue

Notes: The figure presents the trend in share of hospital revenue from commercial payers attributed to inpatient and outpatient care. Panel (a) presents the trend across all short-term general acute care (GAC) hospitals, as reported in the Medicare cost reports (HCRIS). Panel (b) presents the corresponding trend among GAC California hospitals only. Long-term care revenue is excluded.

(a) Inpatient Discharges



(b) Outpatient Visits

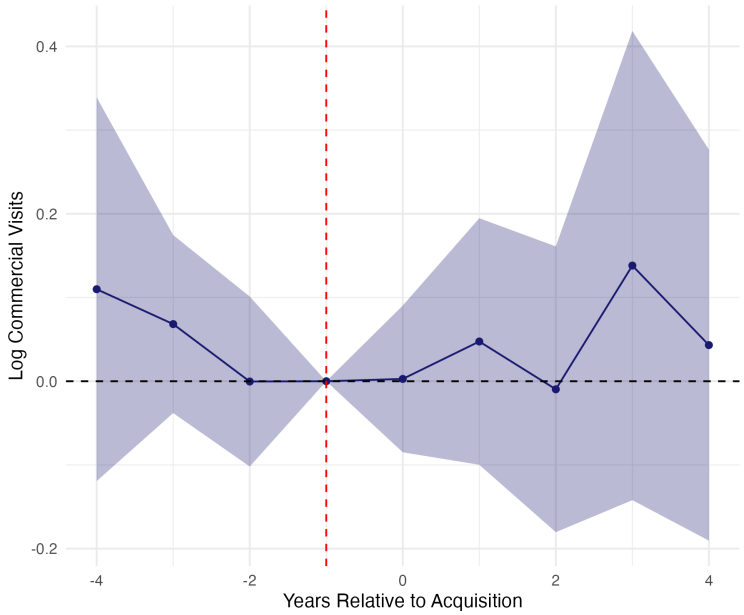


Figure A2: Dynamic Effects on Patient Volume

Notes: This figure presents event study estimates of the effect of a system acquisition on patient volume at independent targets. Panel (a) presents the effects on inpatient stays and Panel (b) presents the effects on outpatient visits. The data on patient volume is obtained from detailed hospital annual reports submitted to California HCAI. The models include time-varying covariates as described in Section 4.1. The gray shaded areas denote 95% confidence intervals. Standard errors are clustered by hospital.

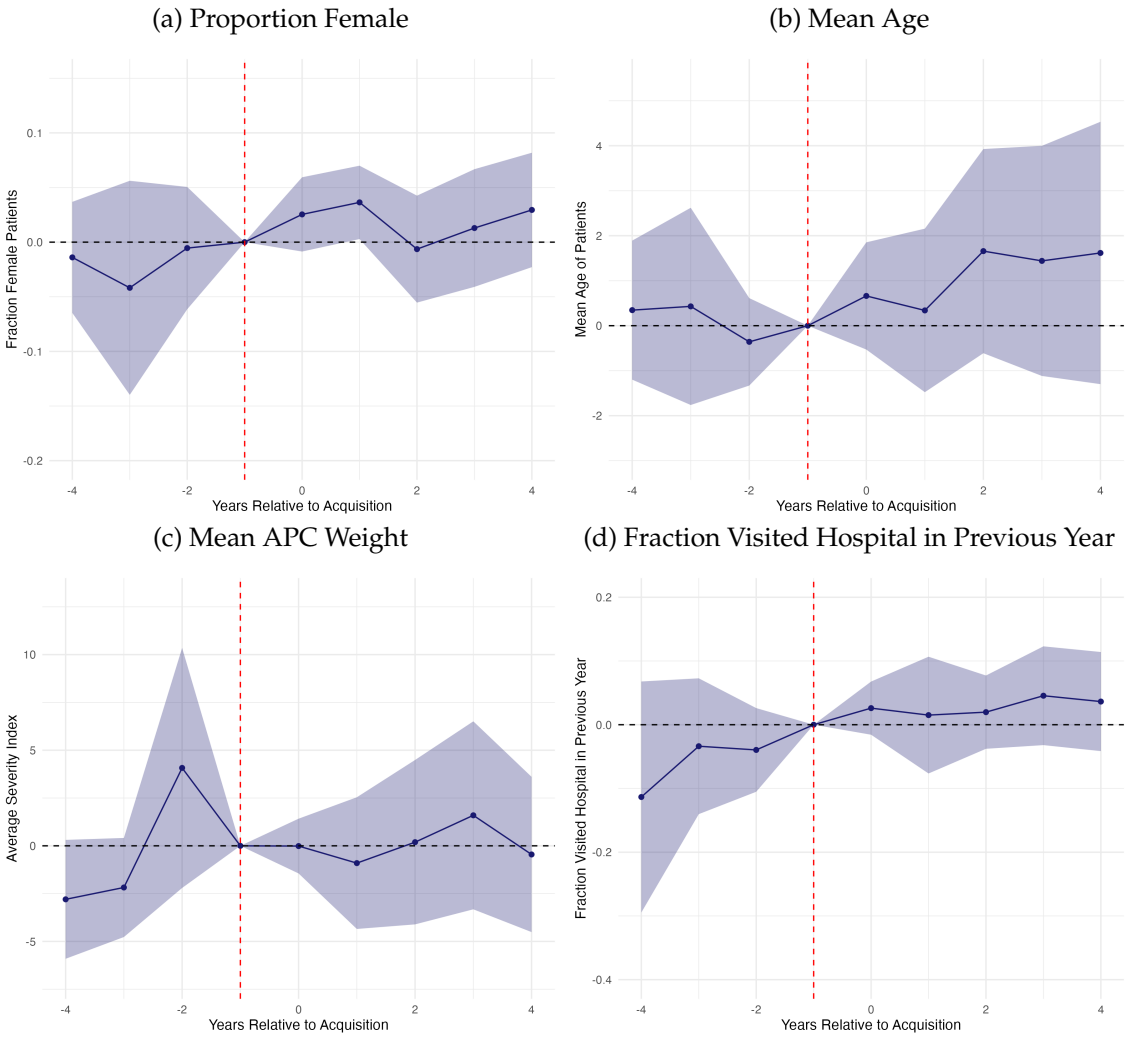
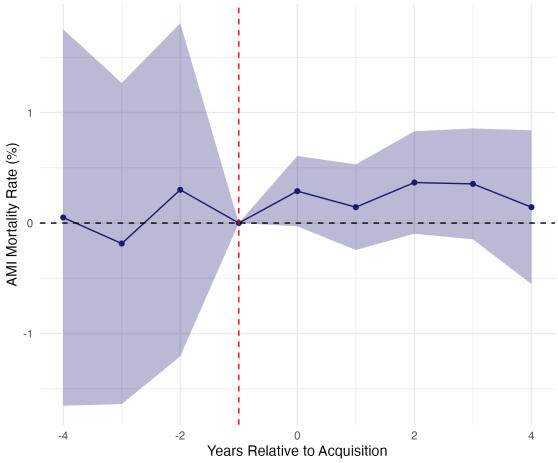


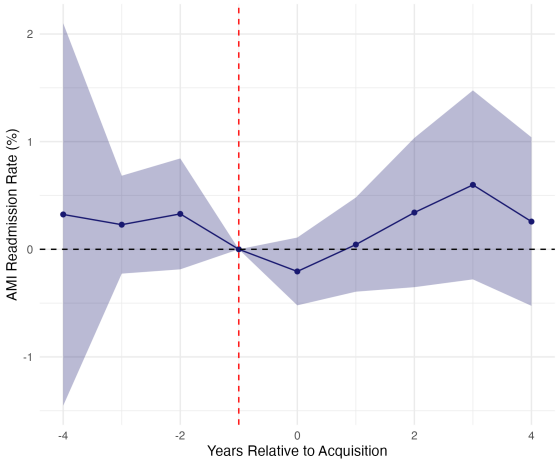
Figure A3: Dynamic Effects on Patient Characteristics

Notes: This figure presents event study estimates of the effect of a system acquisition on patient characteristics of the target hospital. Panel (a) measures the effect on the proportion of the patient population that is female, Panel (b) reports the effect on the mean age of the patient population, Panel (c) reports the effect on the mean APC weight across all outpatient visits at the hospital, and Panel (d) reports the effect on the fraction of patients who visited an outpatient hospital in the previous year. The models include time-varying covariates as described in Section 4.1. The gray shaded areas denote 95% confidence intervals. Standard errors are clustered by hospital.

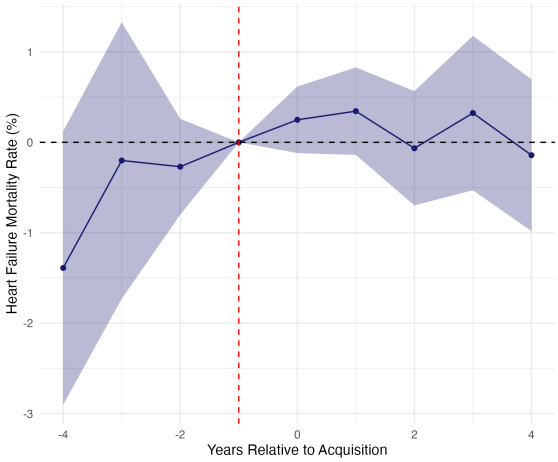
(a) AMI Mortality Rate



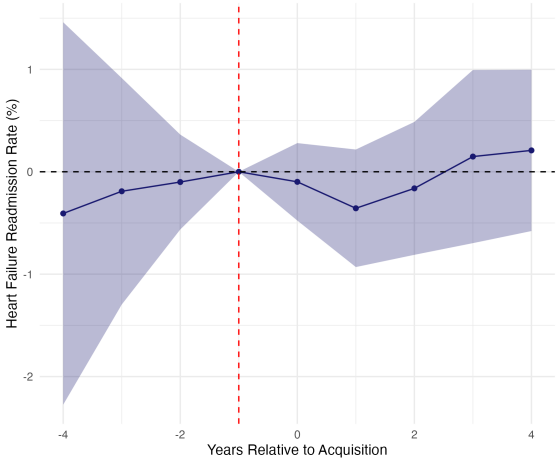
(b) AMI Readmission Rate



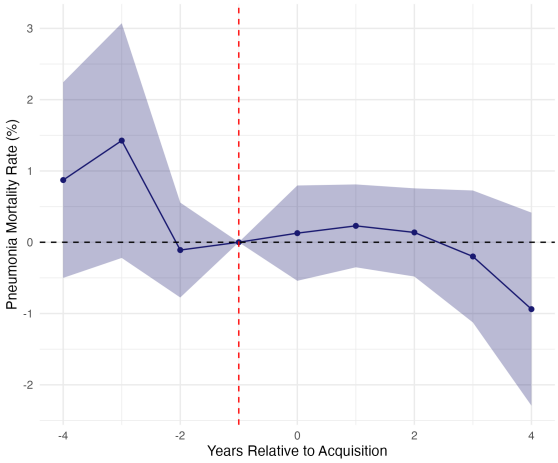
(c) Heart Failure Mortality Rate



(d) Heart Failure Readmission Rate



(e) Pneumonia Mortality Rate



(f) Pneumonia Readmission Rate

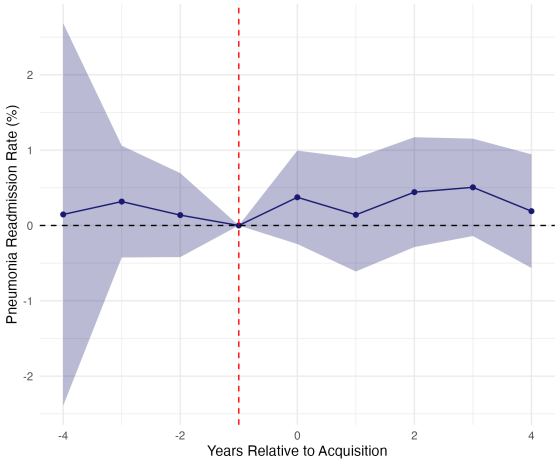


Figure A4: Dynamic Effects on Hospital Quality

Notes: This figure presents event study estimates of the effect of a system acquisition on quality of the target hospital. The models include time-varying covariates as described in Section 4.1. The gray shaded areas denote 95% confidence intervals. Standard errors are clustered by hospital.

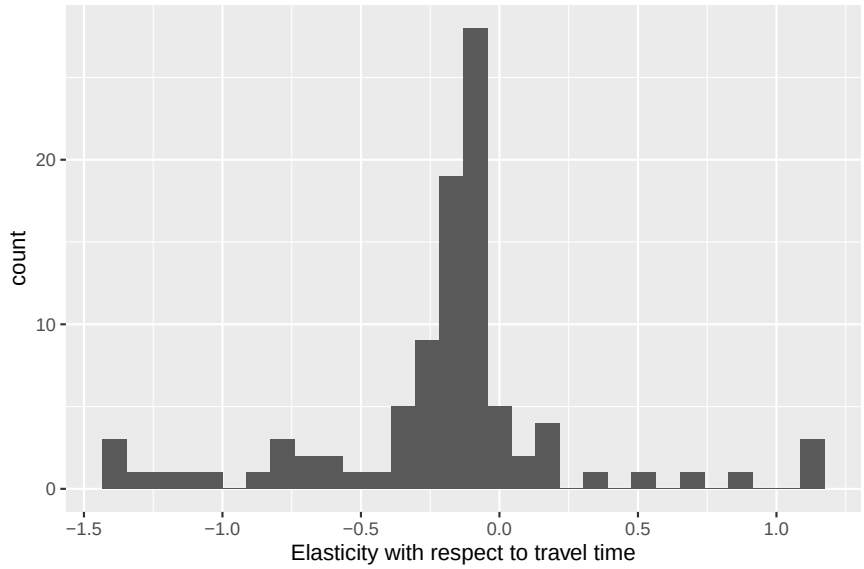


Figure A5: Demand Elasticity with Respect to Travel Time

Notes: This figures presents the distribution of mean demand elasticities with respect to travel time (at the HSA-year level). We calculate elasticities at the individual level using the regular logit expression based on our demand model specified in Equation 3.

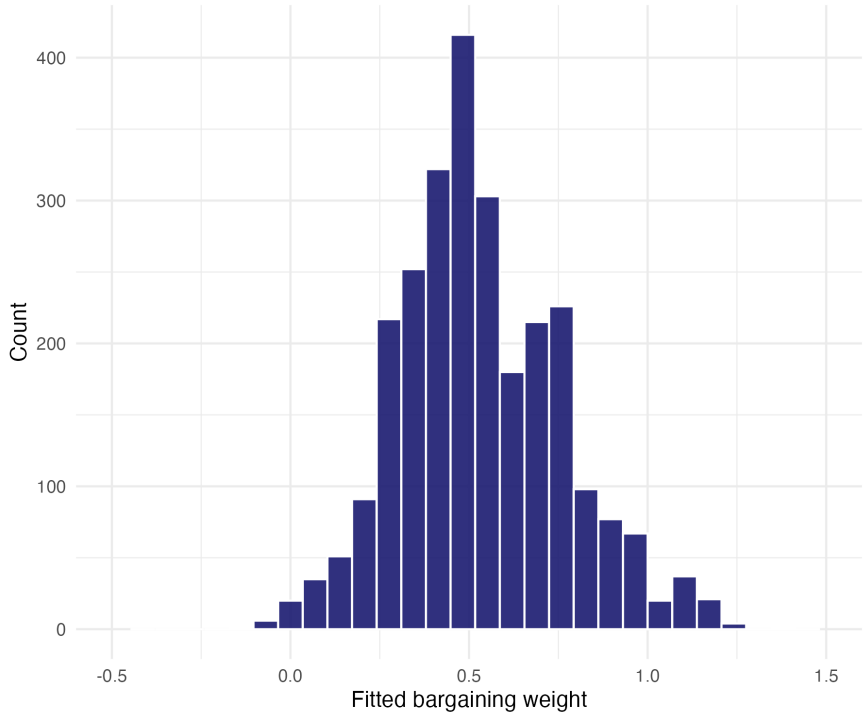


Figure A6: Fitted Bargaining Weight Distribution

Notes: This figure presents the distribution of our estimated bargaining weights. We plot the fitted bargaining weight, which corresponds to $\hat{\alpha}_h$ from Equation 14. Bargaining weights fall outside the unit interval when a hospital systems' profits are negative or when its profits exceed the total joint surplus from its deals with insurers.

(a) No Multiplier

(b) Payer Effects Multiplied By 10

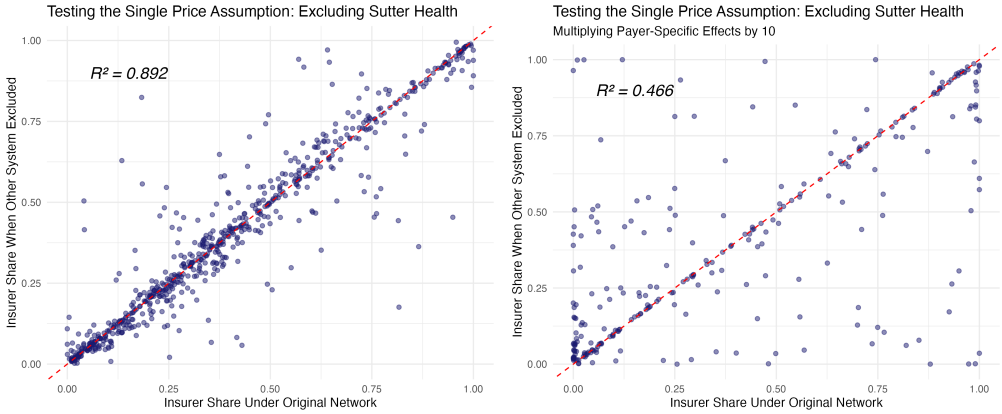


Figure A7: Testing the sufficiency condition

Notes: This figure presents a test of Equation 32, limited to excluding a single system, Sutter Health. All of the demand quantities in Equation 32 can be predicted with our data and demand estimates. In both panels, the x–axis corresponds to the left hand side of Equation 32 and the y–axis corresponds to the right hand side. Panel (a) uses our actual demand estimates. Panel (b) scales our payer-specific demand parameters by 10 before predicting demand.

Table A1: Representativeness of the Claims Sample

Panel A: Hospital Characteristics

	Claims Sample	Full Sample	Diff	t-value
Beds	209.686	206.868	-2.819	(-0.58)
Fraction Teaching	0.059	0.058	-0.001	(-0.12)
Fraction For-Profit	0.207	0.211	0.003	(0.26)
Medicare Share of Visits	0.465	0.465	-0.000	(-0.08)
Medi-Cal Share of Visits	0.278	0.274	-0.004	(-0.92)

Panel B: Payer Share

Payer	Share (%)
Anthem	27.10
Blue Shield	13.53
United	26.66
Others	32.71

Notes: This table compares the characteristics of hospitals in our claims sample with those of the full sample of California short-term general acute care hospitals from 2014 to 2020.

Table A2: Top Clinical Classification Software (CCS) Diagnosis Codes in the Sample

Top-5 CCS Codes			
ED		**Non-ED**	
CCS Codes	Percentage (%)	CCS Codes	Percentage (%)
Abdominal Pain (251)	6.48	Screening for Suspected Conditions (258)	11.96
Nonspecific Chest Pain (102)	5.97	Medical Examination/Evaluation (256)	6.75
Sprains and strains (232)	4.27	Spondylosis/Intervertebral Disorders (205)	3.71
Headache including Migraine (84)	3.51	Non-traumatic Joint Disorders (204)	3.25
Superficial Injury; Contusion (239)	3.35	Connective Tissue Disease (211)	2.91

Notes: This table presents the top five Clinical Classification Software (CCS) diagnosis codes in the claims sample. It reports the corresponding diagnosis names for each CCS code, along with their percentage shares in the subsamples of emergency department (ED) visits and non-ED visits.

Table A3: Hospitals involved in acquisitions

Target Name	Acquirer Name	Year of Acquisition	No. Beds	In-Market
Independent Targets				
Alameda Hospital	Alameda Health System	2014	246	Yes
Olympia Medical Center	Alecto Healthcare	2014	204	No
Glendora Community Hospital	Prime Healthcare Services	2014	118	Yes
Emanuel Medical Center	TENET Healthcare Corporation	2014	209	Yes
Lodi Memorial Hospital	Adventist Health	2015	190	Yes
Stanford Health Care - ValleyCare	Stanford Health Care	2015	207	No
Marina Del Rey Hospital	Cedars-Sinai Health System	2015	145	Yes
Hi-Desert Medical Center	TENET Healthcare Corporation	2015	179	Yes
Adventist Health Medical Center - Tehachapi Valley	Adventist Health	2016	24	Yes
Torrance Memorial Medical Center	Cedars-Sinai Health System	2017	444	Yes
Parkview Community Hospital Medical Center	AHMC & Healthcare, Inc.	2018	193	No
Adventist Health and Rideout	Adventist Health	2018	219	No
System-Owned Targets				
Saint John's Health Center	Providence Health & Services	2014	234	Yes
St. Mary Medical Center	Providence St. Joseph Health	2016	212	No
St. Joseph Hospital	Providence St. Joseph Health	2016	153	No
Santa Rosa Memorial Hospital	Providence St. Joseph Health	2016	278	No
St. Jude Medical Center	Providence St. Joseph Health	2016	320	No
Mission Hospital	Providence St. Joseph Health	2016	249	No
Queen of the Valley Medical Center	Providence St. Joseph Health	2016	120	No
Hoag Memorial Hospital Presbyterian	Providence St. Joseph Health	2016	470	No
St. Joseph Hospital	Providence St. Joseph Health	2016	379	No
Petaluma Valley Hospital	Providence St. Joseph Health	2016	51	No
Redwood Memorial Hospital	Providence St. Joseph Health	2016	25	No

Notes: This table presents details on the hospitals involved in the deals studied in the analysis. For each deal we present the names of the acquiring system and acquired hospitals, the year the deal was executed, whether the target shares a market with at least one member of the acquiring system, and the number of beds of the target.

Table A4: Difference-in-Difference Estimates on Prices Using TWFE

	Independent Targets		System-owned Targets		All Targets	
	No Controls (1)	Controls (2)	No Controls (3)	Controls (4)	No Controls (5)	Controls (6)
Acquired x Post	249.96^{***}	228.62^{***}	-230.56^{***}	-225.33^{***}	-40.30	-7.30
	(55.95)	(57.94)	(64.84)	(65.60)	(60.05)	(59.64)
No. Obs.	485	485	577	556	1,062	1,041
Pre-acquisition Mean	145.33	145.33	610.84	610.84	424.64	424.64

$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$

Notes: This table presents difference-in-differences estimates of the effect of system acquisition on prices at the acquired hospitals using a two–way fixed effects specification. Columns 1–2, 3–4, and 5–6 present the average effect for independent targets, system-owned targets, and all targets, respectively. We present results from two specifications in each panel. Odd numbered columns present results from models without covariates. Even columns presents results from a models that uses covariates. Standard errors are clustered by hospital. Prices are deflated to 2020 dollars. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$

Table A5: Change in Composition of Insurer

	DV: Inclusion of Insurer		
	Acquired Independent	**Acquired System**	**Combined Sample**
	(1)	(2)	(3)
$Acquired \times Post$	-0.16		
	(-1.55)		
$Acquired \times Post$		0.09	
		(1.62)	
$Acquired \times Post$			0.02
			(0.51)
Hospital-Payer FE	Y	Y	Y
Year FE	Y	Y	Y
Controls	Y	Y	Y
N	1,219	1,448	2,667
Pre-acquisition Mean	0.58	0.50	0.57
adj-R^2	0.34	0.34	0.34

Notes: This table presents difference-in-differences estimates of the effect of acquisitions on hospitals' insurer composition. Column 1 presents results for a combined sample including both acquired independent hospitals as well as their controls and acquired hospital systems with their controls, while Column 2 presents results for the sample of acquired independent hospitals plus their controls. Columns 3 presents results for the sample of acquired hospital systems and their controls. The dependent variable, *Inclusion of Insurer*, is an indicator equal to one if there is any claims with nonmissing payment information between a hospital facility and a payer (we classify payers into four categories: Anthem, United, Blue Shield, and others). All models include covariates. Standard errors are clustered by hospital system. Fixed effects are indicated at the bottom of the table. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$

Table A6: Difference-in-Difference Estimates on Costs Using TWFE

	Independent Targets		System-owned Targets		All Targets	
	No Controls (1)	Controls (2)	No Controls (3)	Controls (4)	No Controls (5)	Controls (6)
Acquired x Post	-39.72^{**}	-37.92^{**}	0.92	11.01	-20.36^{*}	-20.76
	(19.39)	(18.64)	(13.27)	(19.55)	(11.70)	(13.82)
No. Obs.	690	690	528	528	1,218	1,218
Pre-acquisition Mean	110.17	110.17	102.52	102.52	106.51	106.51

$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$

Notes: This table presents difference-in-differences estimates of the effect of system acquisition on variable operating costs for the acquired hospitals using a two–way fixed effects specification. Columns 1–2, 3–4, and 5–6 present the average effect for independent targets, system-owned targets, and all targets, respectively. We present results from two specifications in each panel. Odd numbered columns present results from models without hospital covariates. Even columns presents results from a models that uses covariates. Standard errors are clustered by hospital. Costs are deflated to 2020 dollars. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$

Table A7: Difference-in-Differences Estimates on Cost Categories

	Variable Costs		Fixed Costs
	Labor	Non-Labor	Non-Labor
	(1)	(2)	(3)
Acquired x Post	-22.83^{***}	-2.91	2.86
	(8.58)	(5.32)	(4.36)
No. Obs.	690	690	690
Pre-acquisition Mean	83.38	26.55	27.81

$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$

Notes: This table presents difference-in-differences estimates of the effect of a system acquisition on three mutually exclusive collectively exhaustive components of operating cost at independent target hospitals. Column 1 presents results for a model with variable labor costs as the dependent variable, while Column 2 presents results for non-labor variable costs. Columns 3 presents results for fixed costs, which are non-labor by construction. All models include covariates. Standard errors are clustered by hospital. Costs are deflated to 2020 dollars. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$

Table A8: Difference-in-Difference Estimates: Non-Administrative Employees

	Independent Targets		System-owned Targets		All Targets	
	Log Non-nurse FTEs	Log Nurse FTEs	Log Non-nurse FTEs	Log Nurse FTEs	Log Non-nurse FTEs	Log Nurse FTEs
	(1)	(2)	(3)	(4)	(5)	(6)
Acquired x Post	-0.16^{**}	-0.02	-0.04	-0.02	-0.11^{**}	0.00
	(0.08)	(0.04)	(0.06)	(0.05)	(0.05)	(0.03)
No. Obs.	715	715	530	530	1,245	1,245
Pre-acquisition Mean	300.73	296.03	477.18	556.26	385.12	420.49

$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$

Notes: This table presents difference-in-differences estimates of the effect of system acquisitions of independent hospitals on two categories of non-administrative employees. Here nurses includes RNs, LPNs, aides, and non-physician practitioners, and non-nurses includes technician and specialists, food service, and others. All models include covariates. Standard errors are clustered by hospital. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$

Table A9: Difference-in-Difference Estimates on Patient Volume

	Independent Targets		System-owned Targets		All Targets	
	Log Commercial Inpatient Discharges	Log Commercial Outpatient Visits	Log Commercial Inpatient Discharges	Log Commercial Outpatient Visits	Log Commercial Inpatient Discharges	Log Commercial Outpatient Visits
	(1)	(2)	(3)	(4)	(5)	(6)
Acquired x Post	0.11	0.04	0.05	0.01	0.06	-0.03
	(0.15)	(0.08)	(0.08)	(0.09)	(0.09)	(0.05)
No. Obs.	704	712	526	526	1,230	1,238
Pre-acquisition Mean	1,826.04	38,787.47	4,313.82	80,327.55	2,966.27	57,826.67

$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$

Notes: This table presents difference-in-differences estimates of the effect of hospital acquisitions on outpatient (commercially insured) volume at target hospitals. Data on patient stays and visits was sourced from the California annual hospital reports. All models include covariates. Standard errors are clustered by hospital. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$

Table A10: Difference-in-Differences Estimates on Patient Characteristics

	Independent Targets			
	Fraction Female (1)	Mean Age (2)	Service Mix Index (3)	Fraction High-Risk (4)
Acquired x Post	0.02	1.10	0.06	0.03
	(0.02)	(0.96)	(1.72)	(0.03)
No. Obs.	696	696	699	696
Pre-acquisition Mean	0.61	44.53	8.55	0.36

$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$

Notes: This table presents difference-in-differences estimates of the effect of system acquisition on patient demographics and risk at the acquired independent hospitals. Outcome variables are sourced from the patient-level claims data. All specifications include time-varying covariates. Standard errors are clustered by hospital. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$

Table A11: Difference-in-Difference Estimates on Employment and Wages: Acquiring Systems

(a) Employment

	Acquirers of Ind. Targets		Acquirers of Sys.-owned Targets		Acquirers of Any Targets	
	Log Admin FTEs	Log Non-admin FTEs	Log Admin FTEs	Log Non-admin FTEs	Log Admin FTEs	Log Non-admin FTEs
	(1)	(2)	(3)	(4)	(5)	(6)
Acquired x Post	-0.02	-0.01	-0.11	-0.02	-0.04	-0.01
	(0.05)	(0.02)	(0.08)	(0.02)	(0.05)	(0.02)
No. Obs.	1,081	1,081	729	729	1,810	1,810
Pre-acquisition Mean	70.41	766.12	78.67	1,124.17	71.16	798.67

$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$

(b) Wages

	Acquirers of Ind. Targets		Acquirers of Sys.-owned Targets		Acquirers of Any Targets	
	Admin Wages	Non-admin Wages	Admin Wages	Non-admin Wages	Admin Wages	Non-admin Wages
	(1)	(2)	(3)	(4)	(5)	(6)
Acquired x Post	-3.82^{***}	-0.79	0.66	-0.27	-3.28^{***}	-0.73
	(0.89)	(0.51)	(2.36)	(1.03)	(0.85)	(0.49)
No. Obs.	1,081	1,081	729	729	1,810	1,810
Pre-acquisition Mean	61.93	43.01	68.38	45.37	62.52	43.22

$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$

Notes: This table presents difference-in-differences estimates of the effect of involvement in an acquisition as an acquirer on employment and wages. If a hospital (system) is involved in more than one acquisition during the sample period, it is assigned to treatment in the year of the first acquisition. The control group is all system-owned hospitals that are not involved in any acquisitions during the sample period. Panel (a) uses log employment of administrators and all other workers (in FTEs) as the dependent variable, while Panel (b) uses the respective hourly wages as the dependent variable. Columns 1–2, 3–4, and 5–6 present the average effect for acquirers of independent targets, system-owned targets, and any target, respectively. All models include covariates. Standard errors are clustered by hospital. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$

Table A12: Difference-in-Difference Estimates on Costs: Acquiring Systems

	Acquirers of Ind. Targets		Acquirers of Sys.-owned Targets		Acquirers of Any Targets	
	No Controls (1)	Controls (2)	No Controls (3)	Controls (4)	No Controls (5)	Controls (6)
Acquired x Post	14.10	14.26	27.80	27.96*	15.39	15.41
	(10.83)	(13.57)	(17.32)	(16.43)	(10.81)	(12.81)
No. Obs.	1,083	1,083	567	567	1,650	1,650
Pre-acquisition Mean	89.18	89.18	112.52	112.52	90.73	90.73

$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$

Notes: This table presents difference-in-differences estimates of the effect of involvement in an acquisition as an acquirer on variable costs. If a hospital (system) is involved in more than one acquisition during the sample period, it is assigned to treatment in the year of the first acquisition. The control group is all system-owned hospitals that are not involved in any acquisitions during the sample period. Columns 1–2, 3–4, and 5–6 present the average effect for acquirers of independent targets, system-owned targets, and any target, respectively. All models include covariates. Standard errors are clustered by hospital. Costs are deflated to 2020 dollars. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$

B Data appendix

B.1 Outpatient Insurance Claims Data

Our insurance claims data are sourced from the Clarivate Real-World Data (RWD) Product, previously known as Decision Resources Group Real-World Data. This dataset covers over 300 million longitudinal patient records in the United States, drawing from multiple payer sources and encompassing a variety of insurance plans, with records dating back to 2013. The data captures claims submitted through billing software used by healthcare providers, including hospitals, physician offices, long-term care facilities, etc. The open-network structure of this dataset allows for the tracking of patients over time and across different providers, even when patients switch carriers within the observation period. This dataset has previously been used in economic, finance, and health policy research, such as Liu (2022).

The raw claims data are composed of two primary components: submission claims (837s) and remittance claims (835s). Submission claims are records submitted by healthcare organizations and providers to clearinghouses, detailing diagnoses, medical procedures or services rendered, patient demographics, and specifics regarding service dates, locations, charges, and provider and payer information. Once these claims pass through clearinghouses, they are sent to payers, who then return remittance claims back to the providers. Remittance claims provide critical information on remittances, denials, and adjustments, including elements such as allowed amounts, deductibles, and copayments. Only a subset of submission claims can be directly linked to corresponding remittance claims via a linkage table in the dataset.

In the submission data, each claim contains a unique provider identifier (National Provider Identifier or NPI), patient identifier, payer ID, date of service, diagnosis information in the form of International Classification of Disease (ICD) codes, medical procedures in terms of the Healthcare Common Procedure Coding System (HCPCS) codes, patient demographics including the 3-digit ZIP code prefix of residence, gender, and age, payer type (e.g., commercial, Medicare, or Medicaid), and type of insurance plan held by the patient. By matching claims based on the claim payment number, submission claims can link to remittance claims, enabling the identification of hospital charges, negotiated prices (the sum of payer and patient payments), and a detailed breakdown of patient payment components, such as deductibles, coinsurance, and copayments.

Our study focuses on a sample of California hospital outpatient claims from commercial payers between 2013 and 2020. Starting with the submission claims, we follow Liu (2022) in identifying commercial outpatient claims in hospitals. We exclude duplicate claims and claims denied, pended, or suspended by payers. We also exclude any Medicare or Medicaid claims and claims paid by VA Fee Basis Programs in the sample. Additionally, to avoid including Medicare Advantage patient visits, we limit our sample to claims for patients between 18 and 64 years old at the time of medical service. We also exclude some major payers which have large footprints in Med-

icaid Advantage programs, including Centene, Molina, LA Care, Health Plan of San Joaquin and Health Plan of San Mateo.

To incorporate hospital characteristics, we match submission claims data with the American Hospital Association (AHA) annual survey data from 2013 to 2020, using a crosswalk between the National Provider Identifiers (NPIs) and AHA hospital IDs as described in Liu (2022). Our analysis focuses solely on general acute care hospitals within California.

For payer identity, we manually aggregate payer IDs from the submission claims data to align with the National Association of Insurance Commissioners (NAIC) Insurance Holding Company Codes. Each insurance company is assigned a Payer ID, and a single company may have multiple payer IDs. For instance, UnitedHealth is associated with payer IDs such as 87726, 06111, 81400, etc. We categorize payers into four groups based on submission claim volume in California: UnitedHealth, Anthem, CA Blue Shield (the top three payers), and an "other" group that includes all other payer IDs with relatively smaller volumes in the sample.

In our main analysis, we group claim lines to each patient visit. To measure disease severity and resource utilization per visit, following Liu (2022), we assign a relative service-mix weight (or APC weight) for each outpatient visit in the sample. The Ambulatory Payment Classification (APC) system, a fundamental element of the Hospital Outpatient Prospective Payment System (OPPS) developed by CMS, provides the basis for constructing the relative service-mix weights, which are calculated primarily using CPT/HCPCS codes associated with each claim. All APC weights are updated to the 2020Q2 version of the CMS addendum D.

We exclude any patient visits with missing claim numbers, payer ID, patient ID, age, gender, 3-digit ZIP code, service date, or APC weight. Our final sample for patient visits in the submission claims data includes 6,565,412 hospital outpatient visits in California between 2014 and 2020. Of these, 1,344,935 outpatient visits are matched with remittance claims data, containing non-missing and positive payment information.

B.2 Constructing Hospital Price Index

In this section, we describe how we derive the hospital price index—an average negotiated price with private insurers—using patient claims data.

Given the sample of hospital outpatient claims, we aggregate the total allowed amounts (or total paid amounts) Y_{ijmt} in a patient visit (claim encounter) for patient i from insurer m visiting hospital j in year t. We obtain the average price per unit of service-mix weight in a visit by dividing Y_{ijmt} by the relative service-mix weights, w_i, which measures the unit of medical resources used

to treat patient i during his visit. Then we run the following model:

$$\frac{Y_{ijmt}}{w_i} = \gamma_{jt} + \beta X_{it} + \delta_{anthem} + \delta_{bs} + \delta_{united} + \varepsilon_{ijmt}$$

where γ_{jt} is the hospital-year fixed effects, X_{it} is a vector of patient characteristics including patients' gender and the natural logarithm of their age, δ_m for $m \in \{\text{anthem, bs, united}\}$ is an indicator variable equal to one if the insurer is Anthem (Blue Cross), Blue Shield, and UnitedHealth (other insurers are the omitted group), and ε_{ijmt} is the stochastic error term.

To obtain the hospital price, we first recover the vector of hospital-year fixed effects $\hat{\gamma}_{jt}$. We then evaluate the fitted value of patient characteristics and insurer indicators at the sample means, i.e., $\hat{\beta}\bar{X}$ and $\bar{\delta}_m$, with $m \in \{\text{anthem, bs, united}\}$ for each year. Combining all terms give us the average negotiated price per unit of service for hospital j in year t:

$$p_{jt} = \hat{\gamma}_{jt} + \hat{\beta}\bar{X} + \bar{\delta}_{anthem} + \bar{\delta}_{bs} + \bar{\delta}_{united}.$$

In the analysis, we adjust hospital prices p_{jt} to 2020 dollars using the annualized Consumer Price Index (CPI).

B.3 California Cost Report Data

We use annual cost reports from the California Department of Health Care Access and Information (HCAI, formerly OSHPD) for data on hospital outpatient volume, revenue, and costs. These cost reports are referred to as the Hospital Annual Financial Disclosure Reports, and are publicly available going as far back as 1976. Data from these reports has been used in several previous papers on hospital competition (Lewis and Pflum 2015; Ho and Lee 2017). While many states collect and publish similar reports, those available at HCAI are among the most detailed.

There are essentially five types of data included in the reports: 1) basic hospital information, 2) detailed utilization data, 3) the hospital's balance sheet, 4) detailed revenue data, and 5) detailed cost data. We mainly rely on the utilization, revenue, and cost data in this paper. These data are also the most reliable, as the balance sheet data are sometimes reported at the system level with little indication as to when this is being done.

Since the annual reports are required to be submitted by law, coverage is near-complete across general and specialty care hospitals in California. However, until recently, Kaiser hospitals did not submit independent reports and are not comparable with other hospitals in the state. Because of Kaiser's vertically integrated structure and lack of involvement in hospital acquisitions during our sample period, we exclude them from our sample when estimating cost efficiencies and bargaining weights. Thus, the Kaiser reporting issue is immaterial to our results.

The primary benefit of the California data relative to other states is that each type of data (uti-

lization, revenue, and costs) is broken down into detailed categories, which we will refer to as "service lines." Furthermore, these service lines are the same across different sections of the report, enabling the researcher to combine and manipulate all of the reported data at a very detailed level. Examples of service lines include coronary care, emergency department, psychiatric long-term care, and nuclear medicine. Working at the service level provides two primary advantages for this paper: 1) it allows us to exclude portions of hospital operations that are irrelevant to our objectives, such as long-term care, and 2) it enables us to measure heterogeneity in the cost structure across payers, since commercially insured patients may use a different mixture of services compared to Medicare-insured patients.

In this section, we describe the cleaning process for the California cost report data, along with the construction of several key variables used in the paper. These variables include the number of outpatient visits (total and by payer), outpatient fixed and variable costs (total and by payer), and net inpatient and outpatient revenue (total and by payer).

B.4 Cleaning

The main cleaning task required to properly format the California cost report data in a panel structure is to calendarize the data. Hospitals often report their financial data using inconsistent reporting periods. While some hospitals report in January-December intervals, others report using a July-June schedule or any other 12-month schedule. More problematically, some hospitals change their reporting period midway through the sample, which means they have a transitional report that covers less than 12 months. Particularly because most of our variables of interest are flows, these reporting periods need to be standardized.

The basic idea is to aggregate information from the various reports corresponding to a particular calendar year using a weighted sum. Whether those weights need to sum to one or some other number depends upon whether the variable of interest is a stock or a flow. Weights must sum to one for stock variables, but the flow weights may sum to some other number depending upon the length of the reporting periods.

To clarify, consider a hospital in calendar year 2012 with two reporting periods that cover the year: one from January 2012 to June 2012, and another from July 2012 to June 2013. Consider a stock variable such as the number of beds at the hospital during a reporting period, and a flow variable such as the number of discharges. A reasonable procedure for the number of beds would be to calculate a weighted average of the first and second reports, perhaps using weights of 0.5 and 0.5 since each report covers half the year. But this wouldn't make sense for the number of discharges: the first report discharges correspond to only a six-month period, so averaging them with the twelve-month second report would result in a downward-biased estimate. Instead we could sum the total discharges from the first report with half of the second report discharges.

We propose a weighting scheme that accomplishes both of these goals in an intuitive way. Start with a year of interest y, which has begin date t_y^b and end date t_y^e. Call the reporting period $p = (t_p, t_p + T_p)$, where t_p is the begin date and T_p is the period length in days. Here we only focus on reporting periods that overlap with y, where the set of such periods is called P_y. We first calculate the following two weights:

$$\nu_p^A = \begin{cases} \frac{t_p + T_p - t_y^b}{T_p} & \text{if } t_p < t_y^b \\ \min\left\{\frac{t_y^e - t_p}{T_p}, 1\right\} & \text{if } t_p > t_y^b, \end{cases}$$

$$\nu_p^B = \begin{cases} \frac{T_p}{365} & \text{if } t_p \geq t_y^b, t_p + T_p \leq t_y^e \\ \frac{t_p + T_p - t_y^b}{365} & \text{if } t_p < t_y^b, t_p + T_p \leq t_y^e \\ \frac{t_y^e - t_p}{365} & \text{if } t_p \geq t_y^b, t_p + T_p > t_y^e \end{cases}$$

Weights A calculate the proportion of the reporting period that falls within year y. Weights B calculate the size of the within-y portion of the reporting period relative to the length of a year. Then, the stock and flow weights are as follows:

$$\omega_p^{stock} = \frac{\nu_p^B}{\sum_{p' \in P_y} \nu_{p'}^B}$$

$$\omega_p^{flow} = \frac{\nu_p^A}{\sum_{p' \in P_y} \nu_{p'}^B}$$

The intuition here is that for the stocks, the main concern is weighting observations according to how much of the year they account for. Thus we focus on the B weights, adjusted by the sum of the weights to account for the possibility that a portion of the year is missing and so the weights must still sum to one. For the flows, we want to keep track of how much of each report corresponds to the year in concern, so we use the A weights. Still, we adjust using the sum of the B weights to account for missing portions of the year. Note that the flow weights need not sum to one.

Then, the calendarized data points are calculated as follows:

$$x_y^{stock} = \sum_{p \in P_y} \omega_p^{stock} x_p^{stock}$$

$$x_y^{flow} = \sum_{p \in P_y} \omega_p^{flow} x_p^{flow}$$

To implement this process, we begin with the raw complete data sets for each year during our sample period. These files are in long format, with each row corresponding to a hospital-reporting period-report entry, and report entries are characterized by a page number, column number, and line number. We classify every report entry as either a stock or a flow[B1]. In the data, examples of

B1. Some entries, such as the hospital's street address, are placed into a third category called "attributes." Since these

stocks are assets, liabilities, and number of beds. Examples of flows are visits, revenues, and costs. We then implement the calendarization process as described above to create the hospital-calendar year-report entry panel.

B.5 Constructing the Revenue Variables

Now that the data is calendarized, we construct the relevant variables for our analysis (insofar as they are not reported directly). On the revenue side, this means obtaining net inpatient and outpatient revenue, both total and by payer. By payer, we do not mean specific insurance companies (UnitedHealth, Aetna, etc.), at which level the data is not reported. Rather, "payer" here refers to categories such as Medicare, Medi-Cal (the California version of Medicaid), and commercial payers.

The primary obstacle to obtaining inpatient and outpatient net revenue by payer is that while gross patient revenue is reported separately by payer and inpatient/outpatient, deductions are not always reported separately for inpatient and outpatient care. The same problem exists for capitation premium revenue. So to convert gross patient revenue to net patient revenue, we have to split the deductions and capitation revenue across inpatient and outpatient care. The method we use for this is described below.

In the following, j refers to a payer, s refers to a service line (e.g. pediatric acute), R refers to gross revenue, D refers to deductions, and G refers to capitation revenue. $\tilde{R}$ refers to net revenue. S is the set of all service lines, while $S' \subset S$ refers to the set of service lines that we include in our analysis (we exclude long-term care service lines). Inpatient/outpatient status will be denoted using superscripts i and o, while payers and services lines will be denoted using subscripts.

The formulae for net inpatient and outpatient revenue for a payer j where inpatient and outpatient deductions are reported separately are as follows:

$$\tilde{R}_j^i = \left(1 - \frac{D_j^i}{R_j^i}\right)\left(\sum_{s \in S'} R_{js}^i\right)$$
$$\tilde{R}_j^o = \left(1 - \frac{D_j^o}{R_j^o}\right)\left(\sum_{s \in S'} R_{js}^o\right)$$

If capitation revenue exists, the formulae are as follows:

$$\tilde{R}_j^i = \left(1 - \frac{D_j^i - G_j^i}{R_j^i}\right)\left(\sum_{s \in S'} R_{js}^i\right)$$

cannot be averaged, we assign the value coming from the reporting period that covers the most time in the calendar year. If this is ambiguous, we take the earliest report that covers part of the calendar year.

$$\tilde{R}_j^o = \left(1 - \frac{D_j^o - G_j^o}{R_j^o}\right)\left(\sum_{s \in S'} R_{js}^o\right)$$

In the case of managed care payers and Medi-Cal, we only observe total deductions. So the formulae become:

$$\tilde{R}_j^i = \left(1 - \frac{D_j}{R_j^i + R_j^o}\right)\left(\sum_{s \in S'} R_{js}^i\right)$$
$$\tilde{R}_j^o = \left(1 - \frac{D_j}{R_j^i + R_j^o}\right)\left(\sum_{s \in S'} R_{js}^o\right)$$

And if capitation revenue exists:

$$\tilde{R}_j^i = \left(1 - \frac{D_j - G_j}{R_j^i + R_j^o}\right)\left(\sum_{s \in S'} R_{js}^i\right)$$
$$\tilde{R}_j^o = \left(1 - \frac{D_j - G_j}{R_j^i + R_j^o}\right)\left(\sum_{s \in S'} R_{js}^o\right)$$

The idea of these formulae is to calculate net inpatient/outpatient revenue overall as a proportion of gross inpatient/outpatient revenue, and apply this proportion to impute net inpatient/outpatient revenue for the specific services in which we are interested.

In the reports, payers are categorized at a slightly more detailed level than we need. For instance, commercial payers are separated into managed care and traditional insurers. We finish the revenue variable construction by aggregating the calculated net patient revenue to the level of payer category we use.

B.6 Constructing the Cost Variables

The next set of variables we derive from the cost reports are inpatient and outpatient costs, both total and by payer. We additionally allocate these costs to fixed and variable categories, of which we focus on the variable component in order to avoid fixed costs contaminating our measure of average costs. We also separate the costs into other broad categories, such as administrative costs and clinical labor costs, mainly for descriptive purposes.

The cost reports neither allocate costs by payer nor by inpatient and outpatient care. However, the cost reports do allocate cost by service line, C_s. We use this breakdown, along with the proportions of gross revenue for each service line by payer, to estimate inpatient and outpatient costs by payer. The underlying assumption is that within each service line, inpatient and outpatient charges by payer differ in proportion to the costs of providing care.

We will demonstrate the formula for estimating inpatient and outpatient costs by payer for the

inpatient case, though the same calculations apply for the outpatient case. First, we calculate the proportion of gross revenue corresponding to inpatient admissions for service line s:

$$\alpha_s^i = \frac{\sum_{j \in J} R_{js}^i}{\sum_{j \in J} R_{js}^i + \sum_{j \in J} R_{js}^o}$$

Next, we compute the proportion of gross inpatient revenues in service line s that can be attributed to payer j:

$$\gamma_{js}^i = \frac{R_{js}^i}{\sum_{k \in J} R_{ks}^i}$$

Finally, we adjust costs per service line by these two proportions and sum across all the service lines of interest to form our estimate of inpatient costs by payer:

$$C_j^i = \sum_{s \in S'} \gamma_{js}^i \alpha_s^i C_s$$

Following this approach gives us our estimates of total inpatient and outpatient costs by payer. However, these total costs include fixed costs and variables costs, because the costs for each service line C_s include both fixed and variable costs. To separate fixed and variable costs, we need to adjust each service line cost. To adjust each service line cost, we use follow the approach outlined in MedPAC (2020). Essentially, we assign building and capital costs (interest, depreciation, insurance, equipment, plant maintenance, and utilities) as fixed. Additionally, research and education are less relevant to patient care expenses, so we treat them as fixed costs. We assign all labor costs as variable.

Summing the fixed and variable cost categories within each service line, we obtain fixed costs by service line C_s^f and variable costs by service line C_s^v. Then, we reapply our inpatient and payer proportions to obtain fixed and variable inpatient costs by payer:

$$C_j^{if} = \sum_{s \in S'} \gamma_{js}^i \alpha_s^i C_s^f$$
$$C_j^{iv} = \sum_{s \in S'} \gamma_{js}^i \alpha_s^i C_s^v$$

These measures can then be divided by the number of outpatient visits (by payer) to obtain our measure of average costs per visit. Overall, as reported in Table 1, variable costs are about 75-80% of total costs, with little variation across different payers.

B.7 Constructing the Visits and Average Costs

To estimate average costs, we need to have a measure of total outpatient visits by payer to put in the denominator. For this, we take the measure provided in the cost report as given. Total

outpatient visits in the cost reports include several types of visits: emergency services, clinic visits, observation care, psychiatric care, home health care, outpatient hospice care, outpatient surgeries, referred patients, and other visits. We sum all these visits to obtain our measure of outpatient visits. These visits are also reported by payer, and we perform the usual aggregation of adding the traditional and managed care categories within each payer type. This produces our measures of total outpatient visits, v, and total outpatient visits by payer, v_j. Average outpatient costs for payer j are therefore:

$$c_j^o = \frac{C_j^o}{v_j}$$

C Model Appendix

C.1 Hospital Systems

In the main text, for expositional clarity, we wrote our Nash–in–Nash bargaining model for an independent hospital (equivalently a singleton hospital system). Here, we write down the analogous equations for a hospital system.

The first place where the model differs for hospital systems relative to independent hospitals is Equation 9. Suppose that hospital system $\mathcal{H}$ and insurer i negotiate over hospital prices.[B2]

The system version of Equation 9 is

$$\max_{\{p_{hi}:h\in\mathcal{H}\}} \left[\sum_{h\in\mathcal{H}} \Pi_{hi}(\mathcal{G}_i)\right]^{\alpha_{\mathcal{H}}} \left[V_i(\mathcal{G}_i) - V_i(\mathcal{G}_i \setminus \mathcal{H})\right]^{1-\alpha_{\mathcal{H}}}, \tag{19}$$

where $\alpha_{\mathcal{H}}$ is the bargaining weight for hospital system $\mathcal{H}$. Here we maintain the standard assumption in the literature that if system $\mathcal{H}$ and insurer i disagree, then i will have to drop all of the hospitals in $\mathcal{H}$ from its network, so its resulting network will be $\mathcal{G}_i \setminus \mathcal{H}$.

Using the Δ notation introduced in the main text, the Nash–in–Nash FOC is

$$\sum_{h\in\mathcal{H}} \Pi_{hi}(\mathcal{G}_i) = \alpha_{\mathcal{H}} \left(\beta \sum_{\{n:i(n)=i\}} \Delta W_{ni}(\mathcal{G}_i,\mathcal{H}) - \sum_{h'\notin\mathcal{H}} p_{h'i}\Delta D_{h'i}(\mathcal{G}_i,\mathcal{H}) - \sum_{h\in\mathcal{H}} c_h D_{hi}(\mathcal{G}_i)\right). \tag{20}$$

As in Ho and Lee (2017), this FOC is the same for every hospital $h \in \mathcal{H}$ (see their Equation A.3) and its preceding text). So, we impose a single system wide price $p_{\mathcal{H}i}$ for every hospital in $\mathcal{H}$.

We aggregate hospital system level Nash–in–Nash FOCs over insurers in the same way as we aggregate independent hospitals (described in the main text).

C.2 Visit and Diagnosis Probabilities

While our estimated demand model alone implies choice probabilities for each hospital, translating these probabilities into market shares and volume requires estimates of market size and patient demographics. We follow Ho and Lee (2017) by using visit and diagnosis probabilities to convert Census statistics on total population by market to the population of individuals who are actively choosing a hospital for outpatient care.

We use the 1-year American Community Survey (ACS) for our Census estimates. First restricting to all respondents in California, we then filter our sample to only those who had commercial health

B2. Here we omit any subscript labeling the system, e.g., if $h, h' \in \mathcal{H}$, then according to the notation in the main text, the system could also be referred to by $\mathcal{H}_h$ or equivalently $\mathcal{H}_{h'}$.

insurance and were adults below the age of 65. This matches our claims data sample. We then calculate the total commercially insured population each year by gender and age group, where age groups are the following four groups: ages 20-34, ages 35-44, ages 45-54, and ages 55-64. Call these population statistics $N_t(\text{age}, \text{gender})$.

We then compute the corresponding statistics by age group, gender, and year using the claims data. Call these statistics $C_t(\text{age}, \text{gender})$. We then estimate the probability of an individual needing hospital outpatient care using the following frequency estimator:

$$\hat{\gamma}_{vnt} = \frac{C_t(\text{age}_n, \text{gender}_n)}{N_t(\text{age}_n, \text{gender}_n)}.$$

The diagnosis probability γ_{dn} is a conditional probability of an APC weight given that a patient is seeking hospital outpatient care. Though the APC weight is a continuous variable, for ease of estimation we discretize its distribution to deciles and use a frequency estimator to calculate the diagnosis probabilities.

Similar to the visit probabilities, we calculate the probability of an APC weight being in a particular decile conditional on gender and age group. This allows different types of patients to require treatment at different severity levels. If we consider diagnosis d to represent the average APC weight within a particular decile of the distribution, our estimator is as follows:

$$\hat{\gamma}_{dn} = \frac{C(\text{age}_n, \text{gender}_n, \text{decile}_d)}{C(\text{age}_n, \text{gender}_n)},$$

where $C(\cdot)$ is the number of claims in the respective group.

Once we have these probabilities, we simulate a set of patients in California that we use to estimate patient demand in all counterfactual scenarios. Within each market-year, our set of patients represents all combinations of the categorical variables and discretized APC weight distribution used in the demand model. We then assign each patient a weight, according to the combination of the shares of each of their attributes in the claims data and the population in the Census data. Then we adjust that weight using the visit and diagnosis probabilities. The total demand for a hospital in all counterfactual scenarios is therefore a weighted sum of its choice probability for each of these simulated patients.

C.3 Proving Sufficiency of Equation 12

Assume

$$\frac{D_{h'j}(\mathcal{G}_j)}{\sum_{\{i:h'\in\mathcal{G}_i\}} D_{h'i}(\mathcal{G}_i)} = \frac{D_{h'j}(\mathcal{G}_j \setminus h)}{\sum_{\{i:h'\in\mathcal{G}_i\}} D_{h'i}(\mathcal{G}_i \setminus h)} \quad \text{for all } h' \neq h \text{ and for all } j. \tag{21}$$

We want to show that this implies

$$\sum_{\{i:h\in\mathcal{G}_i\}}\sum_{h'\neq h} p_{h'i}\Delta D_{h'i}(\mathcal{G}_i,h) = \sum_{\{i:h\in\mathcal{G}_i\}}\sum_{h'\neq h} p_{h'}\Delta D_{h'i}(\mathcal{G}_i,h), \tag{22}$$

which allows us to estimate the Nash–in–Nash bargaining model that we specify in section 5.1.2 using data on mean hospital prices (average over insurers within a hospital).

Swapping the order of summation on both sides of Equation 22, it is sufficient for us to show that

$$\sum_{h'\neq h} a_{h'i} = \sum_{\{i:h'\in\mathcal{G}_i\}} p_{h'i}\Delta D_{h'i}(\mathcal{G}_i,h) = \sum_{h'\neq h}\sum_{\{i:h'\in\mathcal{G}_i\}} p_{h'}\Delta D_{h'i}(\mathcal{G}_i,h) = \sum_{h'\neq h} b_{h'i} \tag{23}$$

We prove that if we assume Equation 21 holds, then $a_{h'i} = b_{h'i}$ for all h' and i, which completes the proof.

For each $h' \neq h$, by definition

$$a_{h'i} = \sum_{\{i:h'\in\mathcal{G}_i\}} p_{h'i}\Delta D_{h'i}(\mathcal{G}_i,h) = \sum_{\{i:h'\in\mathcal{G}_i\}} p_{h'i}(D_{h'i}(\mathcal{G}_i) - D_{h'i}(\mathcal{G}_i\setminus h)) \tag{24}$$

Note that for any insurer i we can write

$$p_{h'}\sum_{\{i:h'\in\mathcal{G}_i\}} D_{h'i}(\mathcal{G}_i) = p_{h'i}D_{h'i}(\mathcal{G}_i) + \sum_{j\neq i} p_{h'j}D_{h'j}(\mathcal{G}_j) \tag{25}$$

because the left hand side is the total revenue for hospital h' (it's mean price multiplied by it's total demand across all insurers that cover it) and the right hand side is the sum of insurer–specific revenues for hospital h' (split out between i and the rest — $j' \neq i$).

The previous equation can be rearranged to give

$$p_{h'i} = \frac{p_{h'}\sum_{\{i:h'\in\mathcal{G}_i\}} D_{h'i}(\mathcal{G}_i) - \sum_{j\neq i} p_{h'j}D_{h'j}(\mathcal{G}_j)}{D_{h'i}(\mathcal{G}_i)} \tag{26}$$

Substituting Equation 26 into Equation 24 gives (for any h' and any i)

$$\begin{aligned} a_{h'i} &= \sum_{\{i:h'\in\mathcal{G}_i\}} p_{h'i}(D_{h'i}(\mathcal{G}_i) - D_{h'i}(\mathcal{G}_i\setminus h)) \\ &= \sum_{\{i:h\in\mathcal{G}_i\}} \left(\frac{p_{h'}\sum_{\{i:h'\in\mathcal{G}_i\}} D_{h'i}(\mathcal{G}_i)}{D_{h'i}(\mathcal{G}_i)} D_{h'i}(\mathcal{G}_i) - \frac{\sum_{j\neq i} p_{h'j}D_{h'j}(\mathcal{G}_j)}{D_{h'i}(\mathcal{G}_i)} D_{h'i}(\mathcal{G}_j)\right) - \\ &\qquad \sum_{\{i:h'\in\mathcal{G}_i\}} \left(\frac{p_{h'}\sum_{\{i:h'\in\mathcal{G}_i\}} D_{h'i}(\mathcal{G}_i)}{D_{h'i}(\mathcal{G}_i)} D_{h'i}(\mathcal{G}_i\setminus h) + \frac{\sum_{j\neq i} p_{h'j}D_{h'j}(\mathcal{G}_j)}{D_{h'i}(\mathcal{G}_i)} D_{h'i}(\mathcal{G}_i\setminus h)\right) \end{aligned} \tag{27}$$

$$= b_{h'i} + \sum_{\{i:h' \in \mathcal{G}_i\}} \frac{\sum_{j \neq i}(p_{h'j} - p_{h'})D_{h'j}(\mathcal{G}_j)}{D_{h'i}(\mathcal{G}_i)} D_{h'i(\mathcal{G}_i)}(\mathcal{G}_i \setminus h)$$
$$= b_{h'i} + B.$$

All that remains is to show B is zero. The mean hospital price for h' is

$$p_{h'} = \frac{\sum_{\{i:h' \in \mathcal{G}_i\}} p_{h'i} D_{h'i}}{\sum_{\{i:h' \in \mathcal{G}_i\}} D_{h'i}}. \tag{28}$$

Substituting for $p_{h'}$ in B and rearranging gives

$$B = \frac{1}{\sum_{\{i:h' \in \mathcal{G}_i\}} D_{h'i}(\mathcal{G}_i)} \sum_{\{i:h' \in \mathcal{G}_i\}} \frac{D_{h'i(\mathcal{G}_i)}(\mathcal{G}_i \setminus h)}{D_{h'i}(\mathcal{G}_i)} \sum_{j \neq i} D_{h'j}(\mathcal{G}_j) \sum_{k \neq j} (p_{h'j} - p_{h'k}) D_{h'k}(\mathcal{G}_k). \tag{29}$$

After a lot of simplification, we reach

$$B = \frac{1}{\sum_{\{i:h' \in \mathcal{G}_i\}} D_{h'i}(\mathcal{G}_i)} \sum_{\{i:h' \in \mathcal{G}_i\}} \sum_{j > i} (p_{h'j} - p_{h'k})(D_{h'j}(\mathcal{G}_j) D_{h'k}(\mathcal{G}_j \setminus h) - D_{h'k}(\mathcal{G}_k) D_{h'j}(\mathcal{G}_j \setminus h)). \tag{30}$$

Clearly, if there is no variation in the prices that h' receives across insurers, then B is zero, the proof is finished and we can aggregate as we wish.

Alternatively, B is zero if

$$\frac{D_{h'j}(\mathcal{G}_j)}{D_{h'k}(\mathcal{G}_k)} = \frac{D_{h'j}(\mathcal{G}_j \setminus h)}{D_{h'k}(\mathcal{G}_k \setminus h)}, \tag{31}$$

but this equation is implied by the assumption in Equation 21.

C.4 Aggregation Assumption

Here, we empirically investigate support for the following equation

$$\frac{D_{h'j}(\mathcal{G}_j)}{\sum_{\{i:h' \in \mathcal{G}_i\}} D_{h'i}(\mathcal{G}_i)} = \frac{D_{h'j}(\mathcal{G}_j \setminus h)}{\sum_{\{i:h' \in \mathcal{G}_i\}} D_{h'i}(\mathcal{G}_i \setminus h)} \quad \text{for all } h' \neq h \text{ and for all } j, \tag{32}$$

which would allow us to aggregate our Nash–in–Nash FOCs and to use a hospital–level price for estimation.

Each quantity in Equation 32 can be calculated with our data and demand estimates. To proceed, we predict demand for every hospital insurer year triple taking using our observed networks. Then, we predict demand again for every hospital insurer year triple, but excluding a hospital system from each insurers' network. Each time we exclude a system, we can calculate the ratios

on both sides of Equation 32.

Figure C1 presents the results. Panel (a) uses our actual demand estimates. The x–axis corresponds to the left hand side of Equation 32 and the y–axis corresponds to the right hand side. The R–squared exceeds 0.99. Panel (b) repeats this exercise, but we multiply our demand estimates by 10 before predicting hospital system demand. The figure shows that with different demand estimates, there is a weaker relationship between the two ratios.

Our takeaway is that Equation 32 presents an empirically testable condition that rationalizes aggregating the FOCs from our Nash–in–Nash model and using a mean hospital prices (as done previously by Lewis and Pflum 2015).

(a) Actual Demand Estimate (b) Multiply Demand Estimates by 10

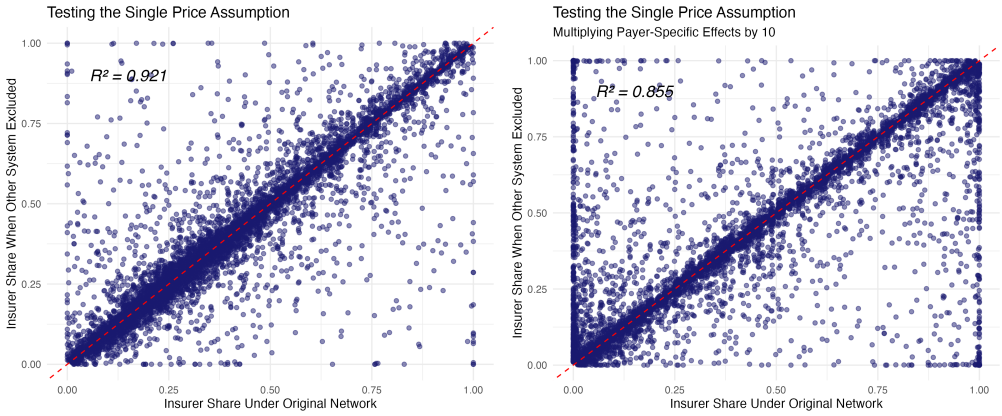


Figure C1: Testing Equation 13

Notes: This figure presents a test of Equation 32. All of the demand quantities in Equation 32 can be predicted with our data and demand estimates. In both panels, the x–axis corresponds to the left hand side of Equation 32 and the y–axis corresponds to the right hand side. Panel (a) uses our actual demand estimates. Panel (b) scales our payer-specific demand parameters by 10 before predicting demand.