Transparency and Negotiated Prices: The Value of Information in Hospital-Supplier Bargaining

Matthew Grennan*
University of Pennsylvania, The Wharton School & NBER
grennan@wharton.upenn.edu

Ashley Swanson
University of Pennsylvania, The Wharton School & NBER
aswans@wharton.upenn.edu

June 6, 2016

Abstract

We empirically examine the role of information in bargaining between hospitals and suppliers of medical technologies. Using a large new dataset of hospitals purchase orders, and identification based on both timing of hospitals joining a benchmarking database and on new products entering the market, we find that access to information on purchasing by peer hospitals leads to reductions in prices. Reductions are concentrated among hospitals previously paying relatively high prices and for products purchased in large volumes, and are consistent with hospitals resolving asymmetric information problems. Achieved savings due to information provision amount to 26 percent of potential savings.

*The data used in this paper were generously provided, in part, by the ECRI Institute (www.ecri.org). The paper benefited from comments from Mike Dickstein, Robin Lee, Amanda Starc, Robert Town, Ali Yurukoglu, and audiences at ASSA Annual Meetings, Bates White, University of California-Berkeley, Cornell, Dartmouth Winter IO Conference, Harvard, iHEA, Indian School of Business, University of Louisville, NBER Spring Health Care, NBER Summer IO, Queens University, Triangle Health Economics Workshop, and Wharton IO Brownbag. We gratefully acknowledge financial support from the Wharton Dean's Research Fund and Public Policy Initiative, the Wolpow Family, and NSF Grant 30067535. Biruk Bekele, Stuart Craig, and Mihir Trivedi provided excellent research assistance. Any errors are our own.

1 Introduction

Business-to-business markets make up a large part of the economy,[1] but they often lack transparency: suppliers negotiate different contracts with different buyers, potentially with widely varying prices, and a buyer typically has limited information regarding other buyers' contracts. As technology has made data easier to collect, distribute, and analyze, many business-to-business markets have seen the entry of information intermediaries who facilitate buyers' ability to benchmark the prices they negotiate.[2] In this paper, we empirically examine the effect of transparency in the form of benchmarking information on negotiated prices between hospitals and their suppliers.

Hospital supplies, including medical devices, are an example of a business-to-business market, but also important in their own right. They are estimated to account for 24 percent of the dramatic growth in inpatient hospital costs between 2001 and 2006 (Maeda et al. 2012), and policymakers have argued that improvements in contracting hold great potential for reducing health care system cost growth.[3] There is substantial variation in prices of inputs across hospitals. For the top fifty hospital supplies in our data, the average standard deviation of prices across hospitals for the same exact product-month is ten percent of the mean price. This is approximately one-third the coefficient of variation estimated for common procedure prices paid to hospitals in different hospital referral regions in Cooper et al. (2015), which could be driven in part by input price variation. It is also approximately the middle of the range of coefficients of variation found in consumer goods markets.[4] Recent legislation has proposed that variation in prices across hospitals is due in part to a lack of transparency in input markets, and that greater transparency would lower average prices.[5]

Prior research in consumer goods markets has largely confirmed the economic intuition that information can facilitate search and decision making for buyers with imperfect information regarding product quality or costs (Sorenson 2000; Jin and Leslie 2003; Hendricks et al. 2012; Bronnenberg et al. 2015) or supplier willingness to accept lower prices (Zettelmeyer

[1]For example, Agicha et al. (2010) report that business-to-business transactions account for 72-86 percent of all payments reported by financial institutions in North America, Europe, and Asia-Pacific markets.

[2]In addition to the hospital purchasing context we study, with product categories ranging from cottons swabs to pacemakers, we are aware of business-to-business "price transparency" benchmarking services emerging in areas as diverse as home appliances and television advertising.

[3]For example, the recent Acute Care Episode demonstration, a bundled payment pilot orchestrated by the Centers for Medicare and Medicaid Services, found that lower costs at demonstration sites were achieved largely due to improved contracting with suppliers. *See* Calsyn and Emanuel (2014) for a discussion.

[4]E.g., Scholten and Smith (2002) report dispersion measures of 1.6% to 20.7% for a variety of retail consumer goods such as cameras, batteries, contact lens solution, and lettuce.

[5]In 2014, Senator Angus King of Maine recently added an amendment to a tax bill that would increase price transparency for medical devices, stating that otherwise "the ability of hospitals to bring price information to bear in negotiations and decisions is clearly limited." (Sarvestani 2014)

et al. 2006; Scott Morton et al. 2011). However, the mechanisms via which information impacts consumer goods may not extend to business-to-business markets where there is often no search across sellers (products are purchased directly from their manufacturers) and negotiators on both sides are professionals employed by firms and thus with different expertise and incentives than the typical consumer.[6] Recent empirical research in business-to-business bargaining (Crawford and Yurukoglu 2012; Grennan 2013, 2014; Gowrisankaran et al. 2015; Ho and Lee 2014; Lewis and Pflum 2015) explains variation in prices across buyers using full information models, but in doing so also documents substantial heterogeneity in *bargaining ability* parameters that intuitively could include variation in information available to negotiators.[7] Our work contributes to both of these literatures by extending our understanding of transparency to the business-to-business setting and by offering information as one explanation for the large unexplained heterogeneity documented in negotiated price contracts.

In order to motivate our empirical analyses, we propose two candidate mechanisms through which benchmarking information might have an impact in this context: (1) a model in which hospitals face uncertainty about suppliers' costs or bargaining parameters, so that transparency reduces uncertainty and the equilibrium dispersion in negotiated prices; and (2) an agency model in which price transparency allows hospital managers to better observe purchasing agents' effort and, in turn, provide improved incentives to purchasing agents to reduce prices. We then empirically test the predictions of each model, using different sources of identifying variation with different informational content.

Our empirical analysis is based on new data containing all purchase orders issued by more than 16 percent of US hospitals. The data set exists because sample hospitals subscribed to a price benchmarking service between 2009-14. Subscription to the service is voluntary, so we control selection in our research design and obtain outside data to address external validity and policy analysis. In order to control for a variety of institutional factors that vary across product categories, we focus our analysis on price negotiations for coronary stents in 508 facilities with cardiac catheterization services. Stents are especially useful as a place

[6]The theoretical and empirical literatures on information disclosure are large and reviewed in Dranove and Jin (2010). Recent studies from a variety of contexts including consumer healthcare spending (Lieber 2015) and choice of college major (Hastings et al. 2015) have found the effects of information to be increasing in prior uncertainty but attenuated by frictions in consumers' ability and incentives to put information to use. In our context of professionals, one might expect prior uncertainty to be lower but ability and incentives to use information to be higher than in the consumer context.

[7]There are two exceptions of which we are aware. Larsen (2015) estimates a bargaining game of two-sided incomplete information about valuations in the used car wholesale market, while uncertainty over bargaining parameters better fits our context. Backus et al. (2015) study cheap talk in consumer-facing bargaining with asymmetric information, but in our case price seems to be the dominant factor of concern for negotiators, decreasing the scope for the trade-offs required for signaling.

to focus because they are representative of the medical technologies central to many policy discussions (so-called "physician preference items"), and they are important in their own right, comprising two percent of sample hospital spend. They also typically have simple linear contracts, so the price we observe is the contracted price.

We rely on two sources of variation to provide plausibly causal identification of the effects of benchmarking information on prices. First, the database is generated by monthly submissions from the member hospitals on prices and quantities of each item purchased, and new members are asked when they join to submit 12 months of retrospective, pre-information data. This data in pre- and post-information states is rare, and it is critical to our research design. We use variation in timing of hospitals' joining the database to construct differences-in-differences estimators, controlling for time-invariant differences at the hospital-product level and product-specific time trends. The assumption underlying this approach is that timing of a hospital joining the benchmarking service is uncorrelated with latent hospital- or hospital-product specific price trends in stents. This strategy would fail and result in an upward bias of information effects if hospitals join when they are experiencing increases in stent prices, or a downward bias if hospitals join when they are enacting other cost-cutting measures (for stents) beyond benchmarking. However, the assumption seems plausible in this institutional setting, as stents are just one of many inputs a hospital purchases that motivate benchmarking. The exogeneity of join timing is also supported by event studies that show no statistically significant divergence of pre-trends.

Second, we develop a set of tests focusing on *new products* entering the market during our sample period. New product introductions provide useful variation for identification along multiple dimensions. Regarding any potential for bias around timing of join, new product introduction provides even more plausibly exogenous timing: assuming new product entry and database join timing are not correlated (which is supported in the data), comparing prices between hospitals pre- and post-join immediately upon a product's introduction, before either hospital type has access to information, offers a difference between these hospitals that sweeps out any persistent sources of bias of join timing.[8] Moreover, new product introductions offer a strategy to separate our two theoretical mechanisms of interest. As we argue in our theoretical discussion in Section 3, the asymmetric information mechanism wherein hospitals use benchmarking information to learn about suppliers relies upon concurrent availability of data on others' prices, but the agency mechanism wherein hospitals use benchmarking information to create better contracts for their purchasing negotiators relies

[8] To be clear, as discussed in Section 4, new products offer opportunities to examine the data for potential bias and to distinguish mechanisms by comparing *trends* in prices within each new product as the information set changes. That is, prices are not compared between existing and newly-entered products for these analyses.

only on the fact that such information will be available in the future. Thus new product entry events allow us to separate the information treatment effects into (1) an agency effect (plus any persistent bias associated with initial join) and (2) an asymmetric information effect.

The estimated average treatment effect across hospital-product-months for coronary stents is small and noisy, suggesting that, consistent with our theoretical predictions, access to the information in the database has heterogeneous effects across hospitals and products. We find that improvements are concentrated among hospital-products with larger potential for savings and with greater quantities of spending at stake. Hospital-products whose prices are above the 80^{th} percentile experience price declines of -\$30 per stent upon accessing database information. The price declines are larger for hospital-product combinations with larger purchase volumes – for hospital-products above the 75^{th} percentile in monthly purchase volume prior to joining the database, price effects increase to -\$70 at the 80^{th} price percentile, compared to -\$20 for hospital-products with lower volumes. Tests suggest that these effects are largely explained by a mechanism where benchmarking solves an asymmetric information problem, helping hospitals learn about their suppliers. Evidence of agency effects is noisier and less robust across specifications.

Each of the above-described results is an estimate of the treatment effect of benchmarking on *prices*, which will be an underestimate of the treatment effect of benchmarking on prices negotiated *in a given contract*. This is because prices are "sticky" in this and other business-to-business markets. In order for benchmarking to have an effect, the buyer must engage the supplier to negotiate a new contract, as the term of the existing contract may not expire for a year or more. Separately estimating the treatment effects on the likelihood of renegotiation and prices conditional upon renegotiation, we find price effects are generated by increasing the likelihood of renegotiation and by generating larger price decreases conditional on renegotiation. This suggests that the benefits of transparency are dampened somewhat by stickiness of contracts and other costs of putting information to use in business-to-business settings.

Ultimately, the estimated effects of benchmarking are modest relative to baseline stent prices, but they are large relative to reported hospital profit margins and the variation in baseline stent prices.[9] Among hospitals achieving low prices before accessing benchmarking data, there is little opportunity for savings and indeed no significant savings are achieved. However, among hospitals learning that there are large opportunities for savings, 12-51 percent of potential savings are achieved after benchmarking data are accessed. Across all

[9] According to the American Hospital Association 2015 trendbook the average hospital operating margin 1993-2013 was 3.8 percent (http://www.aha.org/research/reports/tw/chartbook/ch4.shtml).

hospitals, savings on drug-eluting stents are estimated to be 26 percent of potential savings.[10]

These estimates are of direct interest for considering the impact the entry of information intermediaries can have on the prices buyers negotiate in previously opaque business-to-business markets. They also provide a first step towards thinking about the transparency policies that have been proposed for medical technology markets. Extrapolating our results to all US hospitals with catheterization labs and to the top 50 product categories by expenditure (which make up one third of spending in our data) suggests savings from a broader transparency policy could exceed $480 million annually.[11]

The paper proceeds by first examining the data and providing background on hospital purchasing in Section 2. Section 3 discusses potential theoretical mechanisms and predictions for how benchmarking data might affect negotiated prices, based on existing theory and claims of industry participants. Section 4 clarifies how we use differences-in-differences research designs that leverage plausibly exogenous variation in the data to measure the treatment effect of information. Section 5 presents our results on the treatment effects, underlying mechanisms, and policy implications. Section 6 concludes.

2 Data and Background on Hospital Purchasing

Health care in the hospital setting has high fixed capital costs in facilities and equipment, but it also has high variable costs in the form of skilled labor, pharmaceuticals, and disposable/implantable medical devices. Device costs are of particular interest because, in the short run, hospitals are reimbursed a fixed amount by private or public insurers based on the services they provide, and so device prices come directly from the hospital's bottom line. In this Section, we provide some background on how such devices are used and purchased, and we describe the unique data set and research setting that allow us to analyze the effects of increasing transparency.

The product category we analyze, coronary stents, is an important category in its own right and also an example of the high-tech, high-dollar "physician preference items" at the

[10] As a matter of convenience, we follow the benchmarking database companys example in defining potential savings as being the savings that would be achieved if prices above the mean for a given product were reduced to the same products mean price. We recognize, however, that a complete definition of potential savings would reckon with further economic fundamentals such as preferences and product substitutability.

[11] An important caveat to these policy implications is that our setting does not measure the extent of potential supply side responses to a transparency intervention in the full marketplace. These may include greater obfuscation (Ellison and Ellison 2009), facilitating collusion (Albek et al. 1996), or forcing coordination not to price discriminate via secret discounts (Grennan 2013). The potential for the first two factors will surely depend on purchasing context and implementation of transparency policy. However, to the extent that suppliers know when buyers join our benchmarking database, then our estimates will incorporate the net effects of buyers becoming informed and any reduction in price discrimination possibilities.

center of the policy discussions regarding health care costs and transparency of device pricing. Coronary stents are small metal tubes placed into narrowed coronary arteries to widen them and allow blood flow to the heart. The original technology, the bare metal stent (BMS), was approved in the early 1990s; in the early 2000s, the drug-eluting stent (DES) was introduced as an improvement over the older technology with lower risk of restenosis, a condition that may arise when scar tissue builds up around the stent and restricts blood flow yet again. In the US, hospitals spend more than two billion dollars annually on stents used in over 700,000 procedures.[12]

For stents, as for other physician preference technologies, usage is driven by physicians choosing which product to use to treat a given patient, while prices are determined in negotiation between a hospital administrator and a representative of the product's manufacturer.[13] There is no "search" in the conventional sense, as a given product can only be purchased directly from its manufacturer. The manufacturer holds inventory on-site at the hospital, and the purchase is made when the physician pulls the product off the shelf and implants it into the patient.

Stent contracts typically specify a linear price for the contract duration, often a year. In our conversations with industry participants, the purchasing practices via which these contracts are negotiated vary widely across organizations. Some hospitals have large materials management or purchasing departments with agents who specialize in negotiations, but these departments vary in practices regarding the scope of agent responsibility and incentive contracting. Sometimes a large business unit, such as a catheter lab in the case of stents, will coordinate its own purchasing separately from the rest of the hospital. Finally, even absent accessing benchmarking information, hospitals likely vary in access to information on the prices other hospitals pay via GPOs, hospital system membership, or informal networks of peers.

2.1 Hospital Purchase Order Data

The primary data set used in this study comes from a unique database of all supply purchases made by about 16 percent of US hospitals during the period 2009-2014. The data are from the PriceGuide™ benchmarking service (hereafter, "PriceGuide data") offered by the ECRI Institute, a non-profit healthcare research organization. We observe unique (but

[12] 700,000 estimate from Waldman et al. (2013), referencing stent procedures in Medicare enrollee population. Two billion dollar figure based on authors' calculations using Boston Scientific's reported US revenue in 2012 (BSX 10-K 2012) and Boston Scientific's 2012 market share in purchase order data.

[13] Hospitals typically rely on the services of group purchasing organizations (GPOs) to negotiate contracts for many products, but GPO prices are used as a starting point for direct hospital-manufacturer negotiations for physician preference items and capital equipment (Schneller 2009).

anonymous) identifiers for each hospital and several coarse hospital characteristics: census region, facility type, and number of beds. For each transaction, we observe price, quantity, transaction month, product, and supplier. This includes a wide range of products, encompassing commodities such as cotton swabs and gloves as well as physician preference items such as stents and orthopedic implants.

The reported price and quantity data are of high quality because they are typically transmitted as a direct extract from a hospital's materials management database. Hospitals have strong incentives to report accurately because the analytics the benchmarking service's web portal provides are based on comparing the hospital's submitted data to that of others in the database.[14] Related to its materials management origins, the data is at the stock-keeping-unit (SKU) level, requiring us to use a combination of data collected from manufacturer catalogs and text analysis algorithms to group SKUs that belong to the same manufacturer-product (see Appendix A for this and other sample construction details).

Table 1 displays some summary statistics regarding the transactions data. We observe transactions for 2,111 members, 1,013 of which are hospitals or health systems, and 508 of which are sample facilities that purchase stents. On average, we observe 31 months of transactions for all members, 41 for sample members. We observe purchases in more product categories for sample hospitals than for all members on average (1,143 vs. 462). The average facility in our sample spends $3.4 million per month on all supplies, $80 thousand (2.4 percent) of which is dedicated to coronary stents.

During our sample period 2009-2014, there are twenty branded products sold by four manufacturers – Abbott, Cordis, Medtronic, and Boston Scientific – and in the average month 8.3 branded stents are available from 3.3 manufacturers (with Cordis exiting the market in 2011). The average hospital purchases 59 stents per month (80 percent of which are drug-eluting), distributed among 3.8 brands from 2.1 different manufacturers. Average stent prices decreased by about 30 percent over the sample period. These dynamics mean that controlling for time trends will be important. However, after controlling for trends, price differences across hospitals remain substantial.

2.1.1 Representativeness of the benchmarking database sample

The hospitals in the purchase order data voluntarily joined a subscription service that allows them to benchmark their own prices and quantities to those of other hospitals in the database and thus may not be a random sample of US hospitals. In particular, subscription is costly,

[14]There is no clear incentive for a hospital to misreport data, but of course any misreporting is difficult to verify empirically. We find it reassuring that the distribution of pre-information prices in the benchmarking data is similar to that observed in an external, representative market research dataset; see discussion below. In our empirical analysis, hospital and hospital-product fixed effects will absorb any persistent misreporting.

Table 1: Summary Statistics from Purchase Order Database

	All Members [N=2,111]		Hospitals/Health Systems [N=1,013]		Sample Members [N=508]	
	Mean	SD	Mean	SD	Mean	SD
Months of Data	31.2	21.2	36.8	21.9	41.4	21.4
Product Categories	462.1	502.7	854.7	442.0	1,143.1	313.4
Total Spend/Month ($m)	1.1	2.7	2.2	3.6	3.4	3.2
Purchases Stents?	0.30	0.46	0.59	0.49	1.00	0.00
Stent Spend/Month ($k)					80.4	73.8
Stent Qty/Month					58.7	52.7
Stent Brands/Month					3.8	1.4
Stent Mfrs/Month					2.1	0.7

Authors' calculations from PriceGuide data, 2009-2014.

so we expect hospitals with greater concerns about supply costs to be overrepresented in the database. In a survey of database members, "cost reduction on PPIs" and "cost reduction on commodities" were the first and second (and nearly tied) most commonly cited reasons for joining. This is in accord with our own conversations with purchasing managers who cite a broad array of reasons and product areas as motivations for benchmarking. Along with our analysis of the summary statistics here and pre-trends in the results, this gives us some confidence that entry into the benchmarking database membership is not systematically correlated with price levels or trends in the particular segment of coronary stents that we analyze.

The left panel of Figure 1 compares the distribution of sample hospitals across US census regions to that of American Hospital Association (AHA) member hospitals with cardiac catheterization labs. The Figure also compares our sample to another outside dataset based on Millennium Research Group's (MRG) survey of catheter labs (the source that major device manufacturers subscribe to for detailed market research). The MRG survey is explicitly intended to provide an accurate picture of market shares and prices by US region. The Figure shows that, relative to both comparison samples, the west region is overrepresented in the benchmarking database sample, while the south is underrepresented. We also note that the average sample hospital is larger than the average US hospital with cardiac catheterization capabilities – the right panel of Figure 1 shows that the sample contains disproportionately fewer hospitals in the < 200 beds range and disproportionately more hospitals in the ≥ 500 beds range, relative to AHA hospitals that would purchase stents. We do not have access to bed size for the MRG sample, but we do find that the member facilities in our estimation sample purchased in significantly higher volumes (60 vs. 33 stents per month) and, prior to joining the benchmarking database, obtained slightly lower prices in overlapping periods ($1,615 vs. $1,655 per drug-eluting stent) than the hospitals in the MRG sample. The representation of larger facilities with better negotiation outcomes ex ante in our sample may be due to small hospitals' limited ability to afford access to the database, though we would expect a countervailing effect to come from large hospitals' greater ability to purchase custom benchmarking services from consulting firms.

Figure 1: Distribution of Benchmarking Database vs. Comparison Hospitals

(a) Across Census Regions **(b)** By Bed Size

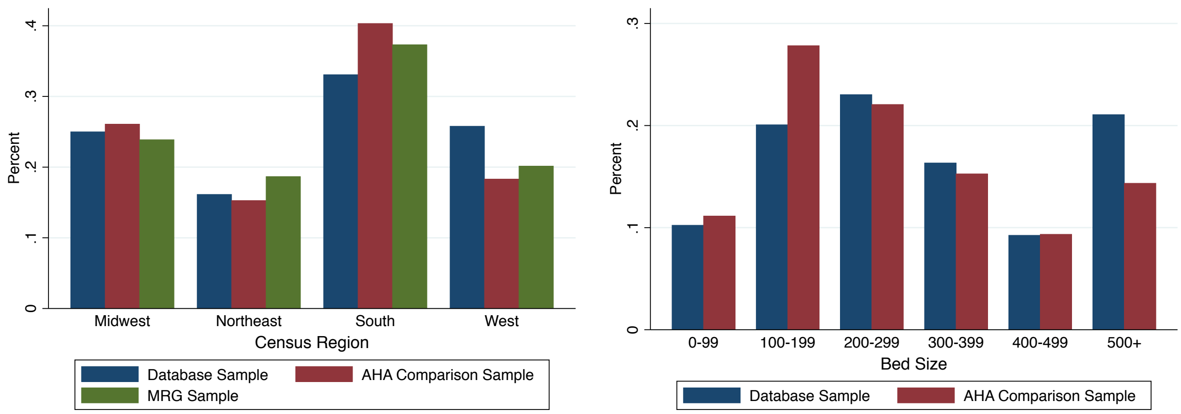

Database sample computations from PriceGuide data, 2009-2014. AHA sample computations from American Hospital Association Annual Survey of Hospitals, 2012; hospitals with catheterization labs defined as those listed as having in-hospital adult or pediatric interventional or diagnostic catheterization services. MRG sample computations from Millennium Research Group survey of catheter labs, Jan. 2010 - June 2013.

Despite potential selection into our sample, the estimation strategy we develop will be internally valid in that it exploits the existence of pre/post data and staggered join dates within the sample of joiners (and uses no non-joiners) to estimate the effect of access to benchmarking information. These estimates are of direct interest, capturing the benefit of benchmarking for facilities that seek out such services. In our later discussion of policy implications, we return to the issue of representativeness and the external validity of our results, using the MRG sample to extrapolate our estimates to the population of US hospitals.

2.2 Price Variation Across Hospitals and Products

Figure 2 displays the distributions of hospital-product and hospital fixed effects, obtained from a regression of prices on dummies for hospital-product combinations (or, respectively, hospitals) and product-month fixed effects. Here and for the remainder of the body of the paper, we focus our analysis on drug-eluting stents only (we examine bare metal stents and find similar results in Appendix F.4). The left panel of the Figure and summary statistics below illustrate the wide cross-hospital variation in prices for the same product at the same point in time in our database hospitals in their pre-join period, with a standard deviation of \$164 and mean of \$1,615, for a coefficient of variation of 0.10. Hospital-product effects explain much of this variation with an $R^2 = 0.89$ for the residual price variation (after product-month detrending). Hospital effects in turn explain almost half of the hospital-product variation, with an $R^2 = 0.44$. Thus our price variation is driven in part by some hospitals consistently paying more than others for drug-eluting stents, and in equal part by some hospitals paying more than others for particular stents.[15]

These patterns are shared by the representative MRG sample shown in the right panel of the Figure. The main difference to note is that the MRG sample has slightly higher prices on average, driven by a slight shift out of the upper part of the distribution relative to the hospitals joining the database. Thus hospitals joining the database have slightly less to gain than the representative sample in terms of raw price differentials.

Our limited observable characteristics do not explain much of this variation in prices. Hospital bed size has no explanatory power. Total volume of stents purchased has some at the extremes: 10^{th} decile hospitals by purchase volume (188 stents per month) achieve prices that are 7% lower than those obtained by 1^{st} decile hospitals (7 stents per month). However, we observe substantial dispersion in prices even conditional on facility size and purchase volume (see Appendix A for details).

One potential explanation may be that stents are "physician preference items": products whose demand is determined in large part by preferences of brand-loyal physicians and which are particularly prominent targets for cost savings by hospital administrators. Policymakers have long argued that the primacy of physician preference in determining demand for such products has limited hospitals' ability to constrain costs using negotiating tools such as standardization (exclusive dealing or contracts based on market share). Consistent with this, we observe no evidence of standardization affecting prices or usage in our purchasing data, and we find no effect of benchmarking information on quantities (see Appendices B and D for details). This motivates our focus on the effects of information on prices (rather than quantities) in the remainder of the paper.

In a different data set and time frame, Grennan (2013, 2014) found evidence that heterogeneity in stent prices across hospitals could be explained in part by heterogeneity in

[15]Though in a different context with likely different mechanisms, these amounts of price variation, distributional shapes, and variance decomposition patterns are remarkably similar to those found in Kaplan and Menzio (2015) in the context of price variation for consumer goods across stores.

Figure 2: Distribution of Prices Across Hospitals and Products

(a) Database Sample (Pre-Join) **(b)** MRG Sample

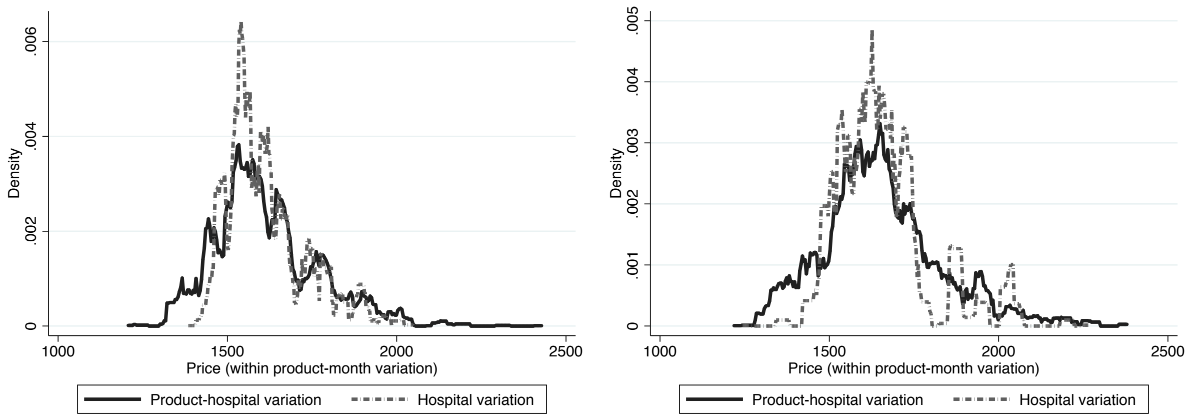

	mean	sd	cv= $\frac{\mu}{\sigma}$	$20^{th}_{\%ile}$	$50^{th}_{\%ile}$	$80^{th}_{\%ile}$	N (unique obs.)
Database sample (pre-join):							
p (within product-month)	1615	164	0.10	1488	1581	1746	3349
$\bar{p}_{hj}$ (within product-month)	1615	154	0.10	1499	1589	1742	561
$\bar{p}_h$ (within product-month)	1613	117	0.07	1526	1586	1704	172
MRG sample:							
p (within product-month)	1655	197	0.12	1497	1636	1794	3444
$\bar{p}_{hj}$ (within product-month)	1655	176	0.11	1525	1637	1776	515
$\bar{p}_h$ (within product-month)	1657	139	0.08	1542	1635	1731	108

Authors' calculations from PriceGuide and MRG data.

physician brand loyalty, but this left a large residual heterogeneity in hospital-product *bargaining ability*.[16] Motivated by policymaker interest in (lack of) transparency in device prices and the existence of the benchmarking intermediary whose data we study, our analysis explores the possibility that part of this heterogeneity in bargaining abilities may be due to heterogeneity in information among hospitals, and that transparency in the form of benchmarking information on other hospitals' prices might affect this.

2.3 The Benchmarking Information Treatment

The information treatment considered in this study is one in which hospitals observe the distribution of other hospitals' prices and quantities and, in so doing, receive information about their relative performance in purchasing. In our empirical setting, sample hospitals were able to access information of this type in several ways. The basic interface members access upon

[16]In these and other studies of empirical bargaining, bargaining ability is parameterized by Nash weights in a structural model of full-information bargaining. These terms capture heterogeneity in prices after controlling for variation in competitive environment, captured by factors such as the outside option.

logging in presents graphical analytics for "potential savings" opportunities at the supplier level. Savings potential is determined by the total dollars that might have been saved in the previous year based on the hospital's volume of purchase and the mean/min prices paid by other hospitals at the manufacturer-SKU level. By clicking through to each manufacturer, the hospital could observe these potential savings broken down by SKU. Further, an interested hospital could filter this comparison to look at only similar hospitals to itself in terms of geography and bed size, and could even click through to access the other hospitals' (de-identified) purchase order data points that were used to construct the analytics.

We obtained clickstream data on the precise timing (to the minute) of all members' website logins. Combined with the purchase order database, which includes the date on which each purchase order was loaded into the database, we are able to reconstruct the analytics a given member would have been presented with upon logging into the database, as well as the more granular data it would have been able to click through to access at each point in time.

In order to preview our approach and results in a simple graphical manner, we begin by simply splitting the sample into pre- and post-join observations and plotting the prices. Figure 3 displays the histograms of prices paid for drug-eluting stents across the entire sample, split between these two groups.

Figure 3: Histograms of Price Distributions: Pre- and Post-Information

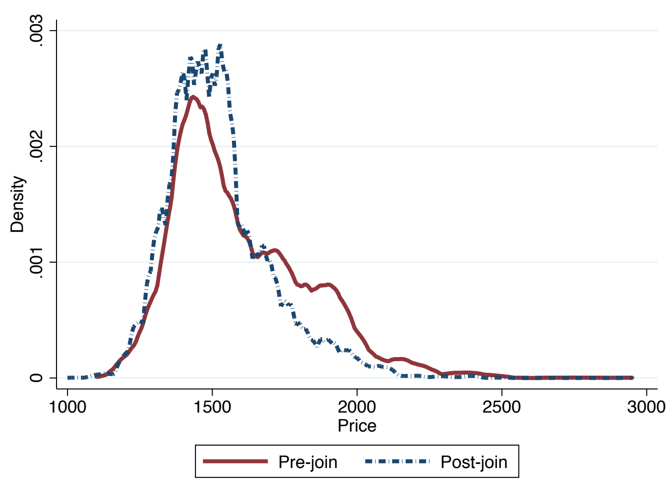

Authors' calculations from PriceGuide data. $N = 32,453$ hospital-product-month observations, classified as pre- or post-join.

The raw data clearly suggests the primary impact of access to the benchmarking information: Hospitals with information are much less likely to pay the highest prices. In the Sections that follow we consider what theoretical mechanisms might drive this result in business-to-business negotiations as well as the research designs and regression specifications that will allow us to credibly establish causal treatment effects and the mechanisms behind them.[17]

3 Theory: Bargaining and Benchmarking Information

While knowledge of others' prices could potentially affect negotiations in many direct and indirect ways, the policy and economics literature on this setting (see, e.g., Pauly and Burns 2008), as well as our conversations with market participants, suggest that there are two primary mechanisms for how benchmarking information could be useful to hospital buyers: (1) in reducing asymmetric information about how low a price the supplier is willing to concede to; and (2) in helping to better solve the agency problem between the hospital and its procurement negotiators by providing a tool for the hospital to monitor negotiator performance relative to the market aggregate. In this Section, we briefly outline two relatively simple theoretical models that capture each of these effects, and use these candidate models to motivate our empirical predictions.

Our models build on the Rubinstein (1982) model of alternating offers bargaining. This model forms the underpinning for a large subsequent literature in theoretical bargaining (Rubinstein 1985; Binmore et al. 1986; Horn and Wolinsky 1988; Collard-Wexler et al. 2014) as well as a recent industrial organization literature in empirical bargaining (Crawford and Yurukoglu 2012; Grennan 2013, 2014; Gowrisankaran et al. 2015; Ho and Lee 2014; Lewis and Pflum 2015). The predictions of the model extend well to empirical settings because the "discount factors" that parameterize bargaining strength in the Rubinstein model can be thought of more generally as proxies for a host of factors that might affect a real-world negotiation such as impatience, opportunity costs of time, laziness, or fear of negotiation breakdown. The model also allows for extension in clear and tractable ways to our mechanisms of asymmetric information about supplier parameters and negotiator agency.

Before we consider incomplete information, we briefly outline the logic of the complete information game as a starting point. The model has a single buyer negotiating with a single supplier over a per-unit surplus $V = wtp - c$ equal to the buyer's willingness-to-pay for a unit of the supplier's product, minus the supplier's marginal cost of manufacturing and distributing a unit of the product.[18] Beginning with the buyer, each player in turn makes a

[17] We focus on the potential effect of information on negotiated *prices*. In Appendix D, we also estimate the effects of information on *quantities* and find no effect, consistent with stents being "physician preference items" where physician demand is based on strong preferences and is relatively insensitive to price.

[18] As noted later in our predictions (and discussed and analyzed in detail in Grennan (2013), V_{hjt} (subscripts suppressed in text) should be thought of as the incremental value created by stent j for the set of

proposal for the division of the surplus. After one player has made an offer, the other must decide to accept or reject it and make a counteroffer in the next round. Players discount continued rounds of bargaining. The buyer has discount factor δ^B and the supplier has a discount factor δ^S, both in $(0, 1)$. In the institutional setting of bargaining over coronary stents, the typical negotiation occurs between a purchasing agent of the hospital and a sales representative of the device manufacturer. For both supply- and demand-side negotiators, discount factors (δ^B, δ^S) should be thought of as coming from some combination of negotiator skill and the incentives they face as agents of their respective employers.

The unique subgame perfect equilibrium of this game is for it to end in the first round with the buyer making an offer that the seller accepts. The resulting price in this complete information (CI) equilibrium is $p^{CI} := c + \delta^S \frac{1-\delta^B}{1-\delta^B\delta^S} V$. Allowing for sufficient heterogeneity in discount factors and valuations across buyer-supplier pairs, this full information model is capable of explaining the observed variation in prices in our data. Thus the null hypothesis of no effect of information remains viable and interesting.

In the Sections that follow, we build off of this baseline model to derive predictions on how benchmarking information might affect prices in alternative information environments. While we do not ultimately estimate the structural parameters of these models, it is important for our economic understanding and for considering policy interventions to examine the mechanisms behind responses to transparency. To that end, we outline two candidate theories, an asymmetric information model (where hospital negotiators are uncertain about the manufacturer negotiator's skill or incentives as embodied in δ^S) and a negotiator agency model (where hospital negotiator δ^B has an effort component which hospital managers cannot observe or infer). We focus on the case where uncertainty is embodied only in the discount factors and not the value over which negotiations occur because this seems to be the primary potential source of uncertainty in coronary stent negotiations, where doctor preferences are typically quite well known by those involved in the negotiation and marginal costs are small relative to the surplus created.[19]

patients for which the doctors at hospital h choose to use j over alternative stents or non-stent treatments, given physician preferences over all stents available at time t.

[19]It is also consistent with anecdotal evidence of little if any equilibrium breakdown in negotiations or destruction of surplus, which are central predictions of models of incomplete information about values. We thank Brad Larsen for this observation. See Ausubel et al. (2002) for a review of the literature focused on informational asymmetries in values. Note this assumption is not directly testable without data on breakdown or beliefs because the surplus and bargaining parameters enter the price multiplicatively.

3.1 Asymmetric Information about Bargaining Parameters

In order to introduce asymmetric information into the baseline bargaining framework, we follow Rubinstein (1985), in which hospital buyers have uncertainty about the bargaining parameter of a given supplier. The model departs from the complete information model outlined above in that the supplier is either of weak type with discount factor δ_w^S or strong type with discount factor δ_s^S ($1 > \delta_s^S > \delta_w^S > 0$). The supplier knows his own type, but the buyer has only a subjective prior ω of the probability that the supplier is the weak type.

Rubinstein (1985) shows that, in this asymmetric information (AI) game, if the buyer is sufficiently pessimistic about the seller being the weak type, then there is a pooling equilibrium where the buyer simply offers what she would offer the strong type in a complete information (CI) game: p_s^{CI}, and both seller types accept this offer. If the buyer is more optimistic, then there is a separating equilibrium where the buyer offers a low price p_w^{AI}, which the weak seller type accepts. But the strong seller type will reject this offer and counteroffer with p_s^{AI} (where $p_s^{CI} > p_s^{AI} > p_w^{AI}$), which the buyer accepts.

For simplicity, assume that access to benchmarking information fully reveals a seller's type. Several empirical predictions for the effects of information on negotiated prices follow directly (Appendix C provides more detailed analysis and further predictions of the theory):

Prediction 1 (Direct Information Effect on High Prices) If information is costless, pessimistic buyers will always become informed. This information will cause the highest prices p_s^{CI} to fall to the complete information weak-supplier price p_w^{CI}, for those cases where the supplier was in fact the weak type. Thus exposure to benchmarking information should lead to some of the highest prices falling.

Prediction 2 (Direct Information Effect on High Prices with High Quantity) If information is costly to obtain (in the sense that searching and analyzing the data takes time that could be used on other productive activity), a pessimistic buyer will become informed whenever the expected benefit of information $\omega(p_s^{CI} - p_w^{CI})q$ exceeds the cost of information. Thus exposure to benchmarking information should lead a proportion of the highest prices p_s^{CI} to fall to p_w^{CI}, for those cases where the supplier was in fact the weak type and among those products with the highest quantity used.

The key feature of this model that we bring to the data is that an information shock for a given negotiating pair may resolve asymmetric information for that same pair going forward. Alternative models that could generate similar empirical predictions might include models wherein one party has preferences over relative as well as absolute performance (e.g., Card et al. (2012) regarding pay transparency, in which workers learning they have relatively low salaries have reduced satisfaction and are more likely to leave their jobs).

3.2 Negotiator Agency

Another mechanism via which benchmarking information could be valuable to buyers would be through providing aggregate information to help the buying firm solve a moral hazard problem with its purchasing agent (who negotiates with the supplier). We expect this mechanism to be relevant in the cardiac unit context. McConnell et al. (2013) present survey data documenting that hospitals' cardiac units vary substantially in their focus on performance measurement,[20] and the Centers for Medicare and Medicaid Services recently found that cardiac and orthopedic units in hospitals responded to bundled payments (which entail higher-powered financial incentives) by improving contracting with suppliers.[21]

Extending the model presented thus far, suppose that instead of the hospital negotiator's bargaining parameter being exogenous, the price will be a function of the hospital agent's *choice* of discount factor δ^B, as well as the discount factor of the supplier. Further, suppose that in addition to uncertainty as to whether the supplier is a strong type or a weak type, there is an additional i.i.d. shock to the supplier's bargaining parameter that is buyer-specific (see Appendix C for details in the case where hospital h faces a supplier bargaining parameter equal to $\delta_h^S \in \{\delta_w^S \epsilon_h, \delta_s^S \epsilon_h\}$ for $\epsilon_h \sim U[0,1]$). Supplier bargaining strength is now observable to the negotiating agents, but not to the principals who manage them at their hospitals.

A moral hazard problem arises in this setting because bargaining effort is costly and provides the agent disutility. Under the usual assumption that the agent is risk averse in money, the optimal employment contract involves risk sharing between the principal and the agent. Holmstrom (1982) shows that if agents face some common parameter which is uncertain from the principals' perspectives (here, the portion of the bargaining parameter δ^S that reflects whether the supplier is a strong or weak type), then relative performance evaluation compared to some aggregate sufficient statistic can be used to write a stronger incentive contract with each agent. This motivates the following Predictions:

Prediction 3 (Monitoring Effect on Prices) If buyer negotiators are imperfect agents of the buying firm, then benchmarking information (observing the distribution of price realizations across hospitals) allows the principal to estimate whether the seller is the weak or strong type, and thus reduces the risk to which the agent is exposed in a

[20] Survey questions that cover performance measurement include "Are new technologies and drugs adopted based on evidence or does no formal process exist for the adoption of new technologies?" and "Is performance reviewed infrequently and only on a success/failure scale, or is performance reviewed continually with an expectation of continuous improvement?"

[21] See Calsyn and Emanuel (2014). The role of incentives in purchasing has also been examined in the broader government contracting context – e.g., in Bandiera et al. (2009), Italian public bodies' prices for generic goods vary with institutional characteristics and poor performance is attributed to passive wastefulness rather than corruption.

higher-powered contract. The higher-powered contract induces more bargaining effort and a lower price than the case where only the buyer's own price is observed.[22]

Prediction 4 (Monitoring Effect on Prices with High Quantity) If buyer negotiators are imperfect agents of the buying firm, but it is costly for hospital managers to search and analyze the data in a way that allows them to write better contracts, then managers will use benchmarking information (observing the distribution of price realizations across hospitals) to write a contract which induces more bargaining effort by the agent and a lower price than in the case where only the buyer's own price is observed if the expected benefit of information $\mathbb{E}[(p^{Info} - p^{NoInfo})]q$ exceeds the cost of information use.

3.3 New Product Entry and Timing of Information Effects

Although the asymmetric information about supplier bargaining type mechanism and the negotiator agency mechanism can generate similar empirical predictions, depending on parameters, an interesting feature that differs between the two mechanisms is the timing during which benchmarking information is valuable to the buyer. In the asymmetric information case, benchmarking is only useful to the extent that data on other buyers' prices for the same product are *currently* available in the database at the time of negotiation. By contrast, even if there is no current data on others' prices for a given product, the agency mechanism allows for managers to incentivize agents today based on performance assessments taking place in the *future* using benchmarking data yet to be collected.

This difference between the timing of information required for the two mechanisms is especially relevant when new products enter the market. By the nature of how the benchmarking database is constructed, there will be no data available on a product for the first month or two it is on the market, and little data for the first few quarters. Thus those who engage in their first negotiation for a product early after its release do so without *current* benchmarking information, even if they have access to the database. This motivates our final theoretical prediction:

Prediction 5 (New Product Entry Separates Asymmetric Information and Agency) For newly introduced products, when they are first released to the market, differences between prices negotiated in the first, uninformed round of negotiation and the second,

[22] In general, the prediction of how the price distribution would move with information depends on where in the model the current heterogeneity is coming from. For example, if the heterogeneity were due to different levels of risk aversion among negotiators, then benchmarking information would tend to decrease the highest prices more than the lowest.

informed round of negotiation must be due to informing negotiators about the seller's bargaining parameter, rather than altering moral hazard. That is, hospital managers can write effort-contingent contracts with purchasing agents in the first round and the second round, but cannot learn about the seller's bargaining parameter until the second round.

3.4 Dynamic Considerations: "Sticky" Contracts, Persistence of Learning, and Supply Responses

In the interest of clearly illustrating the fundamental ideas behind the two theoretical mechanisms of interest, we have abstracted from some realities of hospital purchasing, where contracts are negotiated for a set period of time but sometimes renegotiated before that time, where the same negotiators on the buyer and supplier side may interact repeatedly over time, and where suppliers might change their behavior in response to buyers using benchmarking information. Here we consider how these effects might affect our empirical analysis, if at all.

While a hospital joining the benchmarking database has immediate access to the same data we have on the prices other hospitals are paying for any product, translating that access into differences at the negotiating table still involves a series of steps. In the Predictions above, it was noted that information may be costly to use in the sense that someone at the hospital must anticipate sufficient potential gains for a product to search and analyze the data. Another important friction to consider is that the hospital must engage the supplier to negotiate a new contract (the term of the existing contract may not expire for up to a year or more). To the extent that renegotiation is not frictionless, it will take time and effort to get to the negotiating table and come to a new deal: prices will be "sticky". This will tend to bias the effect of information toward zero. We consider these dynamics in our empirical analyses using event studies and direct examination of recontracting.

The same supplier salesperson may be in charge of negotiating contracts for a bare-metal and a drug-eluting stent. She may also negotiate for the next generation drug-eluting stent when it is released. To the extent that learning about types in the models above captures something that is specific and unchanging over time about that salesperson and the incentives she faces, there will be less asymmetric information and scope for learning, biasing the effect of benchmarking information toward zero.[23]

[23] In specifications where we define information at the manufacturer instead of product level, we obtain nearly identical results. For this reason, we will proceed under the assumption that benchmarking information is potentially relevant for each hospital-product pair, but note that the importance of manufacturer-specific information may vary by context.

While demand side effects of information are generally null or beneficial to buyers, to the extent that suppliers know when buyers join the benchmarking database (or transparency is imposed via public policy), then supply side responses can negate or overturn these effects through greater obfuscation (Ellison and Ellison 2009), facilitating collusion (Albek et al. 1996), or forcing suppliers not to price discriminate via secret discounts (Duggan and Morton 2006; Grennan 2013).[24] Because suppliers will typically know when a hospital is using the benchmarking service, our treatment effects will capture this last effect of reluctance to give discounts when they are no longer secret, but not other supplier obfuscation efforts that might take effect if all buyers had access to benchmarking information, and not collusion that might be facilitated by a public information mechanism. Thus our estimates will be a useful, but potentially incomplete, piece of information in considering large-scale transparency policies.

4 Identification of Information Treatment Effects

The ideal experiment to empirically examine the effect of transparency on prices would be one in which some hospitals were randomly assigned to receive benchmarking data, while others were not. As noted above, the context that allows us to have access to this rare data on business-to-business purchase orders is that the sample hospitals voluntarily joined a subscription database. Our discussions of identification in this Section and of treatment effects in Section 5 focus on the issue of internal validity – consistently estimating causal information effects for the hospitals in our sample. In the final Section, we return to the issue of potential selection into our sample and the external validity of our estimated effects for policies that advocate the rollout of transparency in the form of benchmarking information for all US hospitals. The key features of the data that allow us to estimate causal treatment effects of price transparency for the hospitals in our sample are: (1) that new members submit one year of retrospective data when they first join the benchmarking database, and continue to submit monthly data thereafter; (2) that new members join over time in a staggered (and seemingly random) fashion; and (3) that new products enter the market at points in time that are seemingly uncorrelated with members' information states.

For hospitals that joined during the 2009-14 period, we observe data before and after they were first able to access the benchmarking information available in the database. Figure 4 shows the time series of hospitals joining the database between 2010 and 2014. One technical quirk of the data is that the database vendor rolled out a new version of its database web interface in early 2010 and re-invited all current members to "join" at that point. Thus, for

[24]E.g., Cutler and Dafny (2011) use these potential outcomes to urge caution regarding transparency policies for negotiated prices for medical care.

members "joining" in early 2010, we cannot cleanly identify their pre-period and we exclude those members' "pre-join" data from our analyses. After March 2010, 14 hospitals join the database in each quarter, on average.

Figure 4: Count of Hospitals Joining in Each Quarter

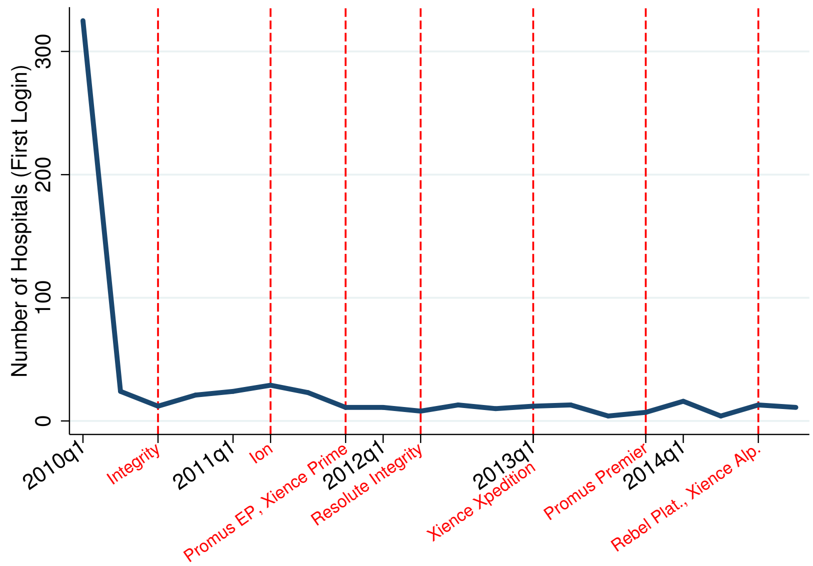

Authors' calculations from PriceGuide data, 2010-2014. "Join" defined by member's first associated login.

4.1 Using Join Date to Identify Price Effects

The availability of both pre- and post-join data for hospitals joining the database at different points in time suggests a differences-in-differences strategy to estimate the treatment effect of having access to benchmarking information. The logic behind this identification strategy is illustrated in Figure 5. In our sample, there are no pure "control" hospitals – all hospitals by definition access the benchmarking data at some point. Suppose there are two hospitals, hospital A and hospital B, where A joined the database one period before hospital B. Under the standard assumption of parallel trends, we can isolate the treatment effect of joining the database on prices by comparing the price trends between the two hospitals for their overlapping time periods. Overlapping periods where both are in the same information state identify any fixed difference between the hospitals unrelated to information access: $\Delta_{t=2}$. This term can be netted out from the difference in period 1 ($\Delta_{t=1}$), taken at the point where A has joined and B has not yet. The difference between these two differences identifies the treatment effect of access to information, $\beta^{join} = \Delta_{t=1} - \Delta_{t=2}$. In our empirical setting, for any given product-month, we observe many hospitals in pre- and post-information states, allowing us to estimate not only time-invariant differences across hospital-products but also product-specific time trends.

Figure 5: Graphical Illustration of Identification Based on Timing of Join

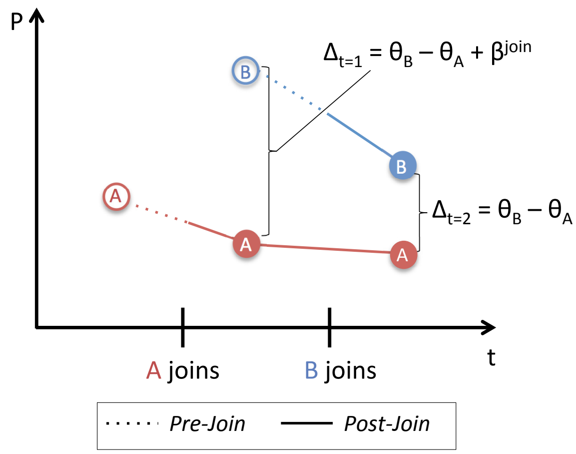

The primary concern with this identification strategy is that timing of a hospital joining the database may be correlated with other contemporaneous factors that impact price trends at that hospital. For example, a hospital may be inspired to join the database due to concerns about upward price trends, which could induce a positive bias in our results – we would be underestimating the counterfactual prices joining hospitals would face if they did not join. On the other hand, a joining hospital might concurrently be undertaking other initiatives intended to constrain prices, such as hiring new personnel or contracting other outside consulting services. Conflating the effects of these other initiatives with the effect of access to the benchmarking information could induce a negative bias in our results. However, the totality of the qualitative and quantitative evidence presented earlier in Section 2 and next in Section 5 suggest that this potential confound has little to no effect in our setting.

Thus we prefer to interpret estimates using this strategy as causal treatment effects. In that case, as we outline next, new product entry can be used as a way to separately identify our theoretical mechanisms. For the reader who remains more skeptical regarding join timing bias, new product entry will also provide a strategy to estimate an asymmetric information effect that is free of such bias.

4.2 Using New Product Entry to Investigate Bias and Identify Mechanisms

New product entry provides another opportunity to identify the above information effect, and further allows us to identify a treatment effect of having joined the database but not yet having access to *concurrent* data on other hospitals' purchases. After new product entry, there is a lag before existing members receive access to benchmarking data on the new product due to the lag between purchase date, data submission, and loading.[25] Moreover, we observe transactions for new products for some members before and after they join the benchmarking database (in the year following new product entry, nine percent of members whose transactions are observed in the average month are pre-join).[26]

This variation allows us to identify a treatment effect of access to benchmarking information *via a mechanism that does not require concurrent access to data on other hospitals' purchases* – Section 3 outlined one such mechanism, in which *joining* the benchmarking database allows hospitals to resolve a negotiator agency problem *even before* benchmarking data are available. We term this the agency ("Ag") effect for the sake of exposition. Once information for the new product becomes available in the database, the same logic as for non-entering products applies: overlapping periods where one hospital is post-join (treated) and the other is pre-join (untreated) identify an overall treatment effect of access to benchmarking information, which is the combination of the agency effect and an information ("Info") effect that requires other hospitals' data. The time period for our study contains many product introductions. In Figure 4, we note the timing of entry of nine new products between 2010 and 2014 (of the twenty products sold during this time period overall).

Figure 6 illustrates this identification strategy graphically. Again, we have hospital A joining the database before hospital B; in this example, hospital A joins well before the product enters the market and hospital B joins after the product enters. Once the product enters, each hospital negotiates prices; hospital B is untreated, while hospital A is treated ("Ag") in the sense that it has joined but has no concurrent data on other hospitals (for example, hospital A may have resolved the agency problem). In the next period, after price data are submitted, loaded, and released to database members, hospital B remains untreated, but hospital A receives another treatment ("Info") in the form of information on other hospitals' prices. In the final period, hospital B has joined the database and received the full treatment effect of access to benchmarking data ("Ag" + "Info"); hospital

[25] In any given month following new product entry, we observe an average of 56 more members having transactions for new products than there are members whose transactions data are actually loaded into the database.

[26] See Appendix A for detail.

A retains both treatments in the final period as well. We thus now have three differences that identify three different objects: in the final period, we identify the fixed hospital differences ($\Delta_{t=3}$); in the penultimate period, we identify the fixed differences plus the "agency" and "information" effects ($\Delta_{t=2}$); and in the first period, we identify the fixed differences plus the "agency" effect only ($\Delta_{t=1}$). These three differences allow us to separately identify the agency ($\beta^{Ag} = \Delta_{t=1} - \Delta_{t=3}$) and information ($\beta^{Info} = \Delta_{t=2} - \Delta_{t=3} - \beta^{Ag}$) effects.

Figure 6: Graphical Illustration of Identification Based on Timing of Join and New Product Entry

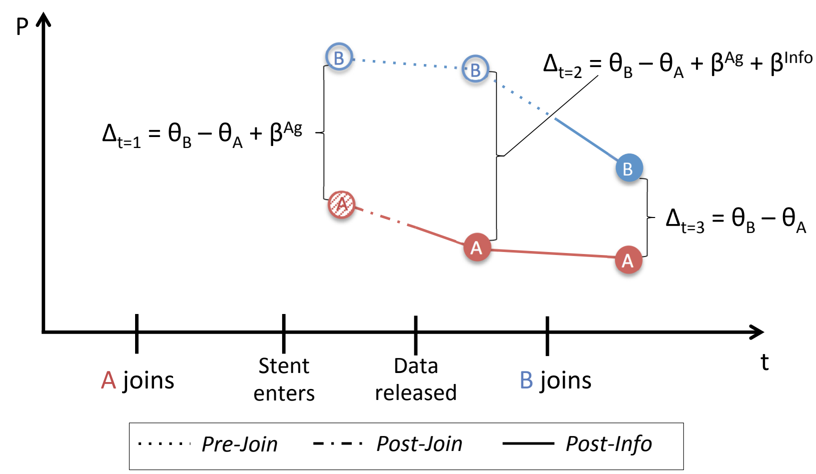

As mentioned previously, entering products also allow us to investigate potential bias due to timing of join. Specifically, any persistent bias associated with factors beyond information at hospitals who have joined or not will be included in the difference between pre- and post-join hospitals in the first few months after new product introduction (β^{Ag}). Thus our estimate of any "asymmetric information" effect where hospitals use information concurrently available in the database to negotiate better prices (β^{Info}) would be free of such bias. It is important to note that this analysis is performed within new products as illustrated in the Figure – a comparison of newly-entered products to existing products is not employed and thus it is not necessary for us to assume that newly-entered products are comparable to existing products.

5 Results: How Information Affects Negotiated Prices

In this Section, we estimate regressions based on the research design just described to more carefully measure and understand the effect of information suggested by Figure 3, accounting for time-invariant differences across hospitals (or hospital-product combinations), imbalance in the sample before and after information shocks, and trends in prices over time. All of the regressions we present are extensions of a baseline specification implementing the difference-in-differences approach based around the timing of hospital access to benchmarking information.[27] Letting P_{hjt} denote the price observed for product j, hospital h, and month t, our preferred specification controls for hospital-product fixed effects $[\theta_{hj}]$, month fixed effects $[\theta_t]$, and separate linear time trends for each product $[\gamma_j * (t - t_{min_j})]$ (where t_{min_j} is the first period in which we observe data for product j: either the beginning of our sample or the month of entry of product j into the market):

$$P_{hjt} \quad = \quad \beta^{Info} * \mathbb{1}_{\{post_{hjt}\}} + \theta_{hj} + \theta_t + \gamma_j * (t - t_{min_j}) + \varepsilon_{hjt}.$$

Here, $\mathbb{1}_{\{post_{hjt}\}}$ is an indicator equal to one after a hospital first accesses information in the benchmarking database (and data for the given product are available in the database) and zero prior, making the coefficient β^{Info} an estimator for the treatment effect.[28] All of the regressions and results below extend this specification to allow for varying types of heterogeneity in this treatment effect.

Recall our empirical predictions are that information will lead to price decreases, perhaps concentrated in the upper part of the pre-information price distribution (indeed, the suggestive evidence in Figure 3 indicated this would be the case). Furthermore, both candidate theoretical models predicted that, if search or re-contracting are costly, price effects would be concentrated in products purchased in higher quantities. Our analyses in this Section will accordingly focus on understanding the heterogeneity in opportunities to benefit from benchmarking information, the frictions in taking advantage of these opportunities, and the theoretical mechanisms at work. We conclude by extrapolating to consider the potential overall impact of transparency on the US hospital expenditures on physician preference items.

[27]For now, the analysis includes all products – entering products as well as products that were present in the market at the beginning of the sample. The timing of information for entering products is here defined as the first date at which the member logs into the database when there are meaningful data on other hospitals' purchases loaded into the database. In the current results, this is the first login after six months post-entry. We consider this to be a pooled "information" effect. We show results estimated only from "timing of join" variation in Appendix F.2 and find our discussion unaffected; hence we defer further consideration of "information" vs. "agency" effects until we arrive at the results that separately identify mechanisms.

[28]See Appendix E for results with alternative sets of controls. We discuss any meaningful differences across specifications in the results below.

5.1 Effects of Information throughout the Price Distribution

Our first result, shown in Table 2, regards the average treatment effect of information across all hospital-product-months. Results are shown for a variety of different specifications of control variables (alternating between hospital and product fixed effects vs. hospital-product fixed effects, and between product-month fixed effects vs. product-specific linear trends). The estimated treatment effect varies somewhat according to which controls are included in the specification – specifically, the estimates are significantly smaller when we control for hospital-by-product (rather than hospital *and* product) fixed effects. This may be due to our effectively controlling for an unknown source of hospital-product specific heterogeneity. The hospital-product fixed effects may also introduce attenuation bias towards zero, as there are some hospital-products for which there are relatively few observations. We consider the preferred specification (Version 3) to be a conservative estimate of the effects of information on negotiated price within hospital-product, and we focus on these results in the main text.[29]

Table 2: Average Treatment Effects of Information across All Hospital-Product-Months

Version of Controls	1	2	3	4
β^{Info}	-12†	-21†	-3	-7
SE	(5)	(7)	(3)	(5)
Hospital+Product FEs	Y	Y	N	N
Hospital-Product FEs	N	N	Y	Y
Linear Product Trends	Y	N	Y	N
Product$x$Month FEs	N	Y	N	Y

Authors' calculations from PriceGuide data, 2009-2014.
$N = 32,453$ member-product-months. Includes 508 members.
Standard errors clustered at hospital (Versions 1 and 2) or hospital-product (Versions 3 and 4) level shown in parentheses.
Superscript (†) indicates significant difference from zero at the 1% level; (**) at the 5% level; (*) at the 10% level.

The preferred specification finds that prices decrease by only -\$3 *on average* when benchmarking data are accessed. The ATE is also imprecisely estimated, with a standard error of \$3.[30] Neither the magnitude nor the variability of this estimate is entirely surprising: the average treatment effect at the hospital-product-month level is not weighted based on potential total savings – which vary substantially across observations – and does not take into account frictions in renegotiation.

[29]See Appendix E for all versions of heterogeneous treatment effect and mechanism results. The results are similar, with the primary difference being that effects in the top of the price distribution roughly double in size with hospital instead of hospital-product fixed effects. This difference is due to a significant negative "agency" effect in the hospital fixed effect specifications, which does not appear in the specifications that control for hospital-product fixed effects.

[30]Detailed tables and figures on the timing of the effect are available in Appendix F.1.

In keeping with the empirical predictions derived from theory, the remainder of our analyses will allow for heterogeneity in treatment effects depending on each hospital-product pair's place in the price distribution for that product at the time the hospital gains access to information. For each member's first login to the database, we compare the member's price for each product purchased in the year prior to login to the full distribution of prices for the same product across all hospitals during the same period. We then flag each hospital-product pair based on its pre-join price relative to percentiles of the price distribution. In regression form, we interact the indicator for a hospital having access to information in the database, $\mathbb{1}_{\{post_{hjt}\}}$, with dummy variables for each pre-join price quintile, $\mathbb{1}_{\left\{quintile_{hj,pre}\right\}}$, allowing for heterogeneous treatment effects depending on whether the hospital was paying a high or low price (relative to other hospitals) for the product at the time of "information":

$$
\begin{aligned}
P_{hjt} &= \beta^{Info}_{quintile} * \mathbb{1}_{\{post_{hjt}\}} * \mathbb{1}_{\left\{quintile_{hj,pre}\right\}} + \theta_{hj} + \theta_t + \gamma_j * (t - t_{min_j}) + \varepsilon_{hjt} \\
\mathbb{1}_{\left\{quintile_{hj,pre}\right\}} &= \mathbb{1}_{P_{hj,pre} \in quintile\left(\{P_{hj',pre}\}_{h'=1}^{H}\right)}
\end{aligned}
$$

where the coefficient $\beta^{Info}_{quintile}$ is the treatment effect of accessing information in the benchmarking service, for each quintile of the pre-information price distribution. Figure 7 shows the results.

The treatment effects exhibit substantial heterogeneity depending on the pre-information price the hospital was paying for a product relative to others. The treatment effects are statistically zero in all but the top quintile of the pre-information price distribution, where the effect is -\$27. This evidence is consistent with Prediction 1 that, absent benchmarking, pessimistic hospitals would pay suppliers high prices regardless of those suppliers' true bargaining parameters, leading those hospitals to negotiate lower prices after joining. Under the asymmetric information mechanism, there is little reason to expect transparency to affect prices that are *relatively good*. It is also worth noting that we do not see evidence that the lower part of the distribution shifts upward significantly, as might happen in the presence of mean reversion (in fact, the point estimate for the bottom quintile is negative, though not significant).

We also performed an event study analysis separately for each quintile of the price distribution. The results for the top quintile of the pre-information price distribution are shown in Figure 8. The pre-trends in the six months pre-information are essentially zero, while there is a steady decline in prices after information access – a year after join, the treatment effect is -\$96 relative to the "info" date. The effects in the 6-12 months prior to information access are negative, though not significant. If one were to lend weight to the noisy point estimates, pre-trends prior to joining the database would lead the differences-in-differences

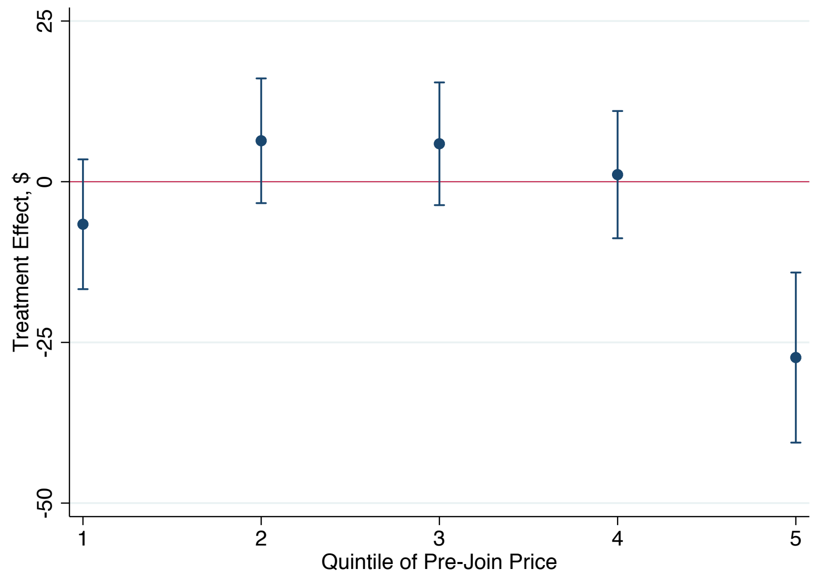

Version	Pre-info price quintiles ($\beta^{Info}_{quintile}$=)				
	1	2	3	4	5
3	-7	6	6	1	$-27^{\dagger}$
	(5)	(5)	(5)	(5)	(7)

Authors' calculations from PriceGuide data, 2009-2014. $N = 32,453$ member-product-months. Includes 508 members. Standard errors clustered at hospital-product level shown in parentheses. Superscript (†) indicates significant difference from zero at the 1% level; (**) at the 5% level; (*) at the 10% level.

Figure 7: Treatment Effect Estimates Throughout the Price Distribution

estimates to be an understatement of the effects of information on price.

We consider these results as strong suggestive evidence that the estimated treatment effects are due to accessing the information in the benchmarking data rather than to any potential sources of bias due to join timing. The evidence of steeper negative price trends after join in the top quintile of the price distribution than there are in average prices suggests that, if there are indeed factors that cause prices to decrease after join that are unrelated to information access, *they must disproportionately impact hospital-products whose prices are relatively high in the pre-period*, a fact which would be unknown to parties whose behavior impacts prices without them accessing the database. In subsequent results, we will proceed under the assumption that any bias due to join timing is small.[31]

For the sake of statistical power and for expositional simplicity, we return to estimating pre-/post-treatment effects, rather than breaking them down by month relative to information access. However, it is noteworthy that treatment effects become larger in magnitude

[31] For the reader who prefers a more skeptical interpretation, any remaining bias due to timing of join will be absorbed with our measure of the agency effect in our mechanism test, so that we are able to obtain a "clean" asymmetric information effect.

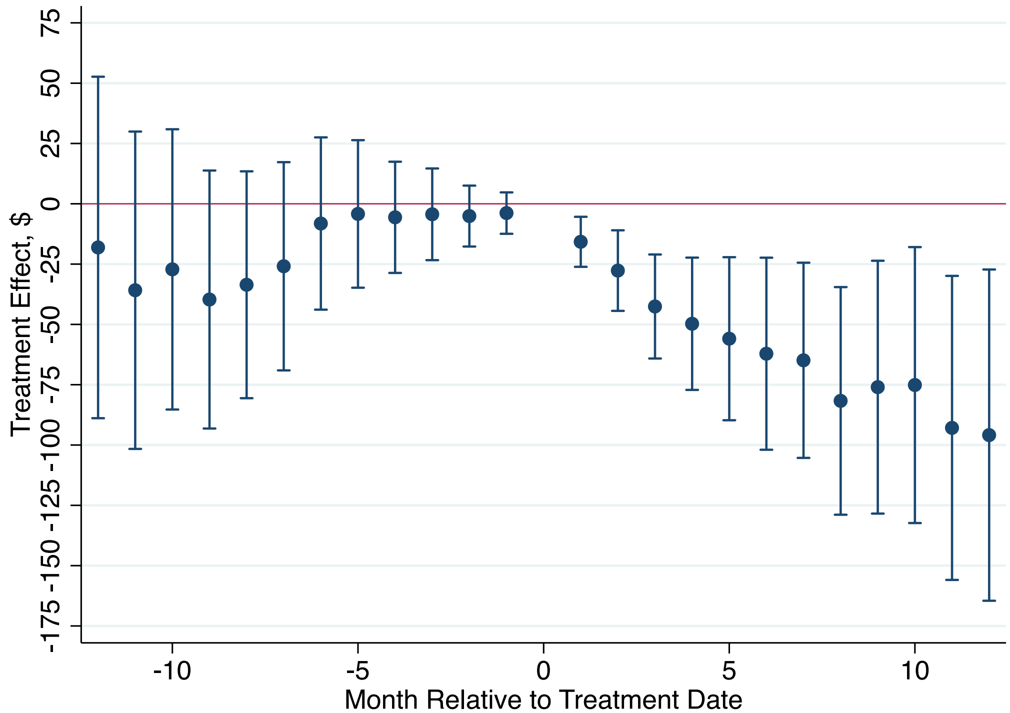

Version	Month relative to join date ($\beta_5^{Join,mo} =$)																							
	-12	-11	-10	-9	-8	-7	-6	-5	-4	-3	-2	-1	1	2	3	4	5	6	7	8	9	10	11	12
3	-18	-36	-27	-40	-34	-26	-8	-4	-6	-4	-5	-4	$-16^{\dagger}$	$-28^{\dagger}$	$-43^{\dagger}$	$-50^{\dagger}$	$-56^{\dagger}$	$-62^{\dagger}$	$-65^{\dagger}$	$-82^{\dagger}$	$-76^{\dagger}$	-75^{**}	$-93^{\dagger}$	$-96^{\dagger}$
	(36)	(34)	(30)	(27)	(24)	(22)	(18)	(16)	(12)	(10)	(6)	(4)	(5)	(9)	(11)	(14)	(17)	(20)	(21)	(24)	(27)	(29)	(32)	(35)

Authors' calculations from PriceGuide data, 2009-2014. $N = 23,016$ member-product-months. Includes 507 members, twelve months pre- and post-join only. Standard errors clustered at hospital-product level shown in parentheses. Superscript (†) indicates significant difference from zero at the 1% level; (**) at the 5% level; (*) at the 10% level.

Figure 8: Event Studies of Treatment Effect of Access to Benchmarking Information, Top Quintile of Price Only

over the course of the year after information access. We see this as evidence of price "stickiness" as a friction that limits gains from transparency, and we return to this issue in Section 5.2.3.

5.2 Mechanisms: Where and Why Does Information Matter Most?

The above results established that transparency in the form of access to benchmarking information leads to lower prices for hospital-product cases where the hospital is in the upper quintile of the price distribution (across hospitals) for that product. In this Section, we test the further predictions from Section 3 to better understand the mechanisms behind these price reductions. We first allow for treatment effects to vary with purchase volume so that we may investigate whether hospital-products with high expenditures at stake experience larger average price changes, in keeping with a model with effort cost of information use and renegotiation (Predictions 2 and 4). Next, we utilize the fact that for new products no benchmarking information is available in the database until several months after product entry to separate the asymmetric information mechanism from the agency mechanism (Prediction 5). Finally, we decompose the estimated price effects into price effects conditional on renegotiation and price effects due to greater likelihood of renegotiation.

5.2.1 Costs of putting information to use: treatment effects vary with quantity

To the extent that using benchmarking information to identify opportunities and then engage in renegotiation (of supply contracts or employment contracts) is costly, Predictions 2 and 4 predict that transparency will have the largest effect for hospitals and products purchased in high quantities. To investigate these predictions, we interact our "post" variable (again separately for each pre-information price quintile), with a dummy equaling one for hospital-product combinations with high purchase volumes in the period prior to information access. We generate a dummy variable $\mathbb{1}_{\{high^q_{hj,pre}\}}$ equal to one for hospital-products with monthly purchase volume above the 75^{th} percentile in the months prior to join, and we estimate the specification:

$$\begin{aligned} P_{hjt} &= \beta^{Info}_{quintile,low^q} * \mathbb{1}_{\{post_{hjt}\}} * \mathbb{1}_{\left\{quintile_{hj,pre}\right\}} \\ &\quad + \beta^{Info}_{quintile,high^q} * \mathbb{1}_{\{post_{hjt}\}} * \mathbb{1}_{\left\{quintile_{hj,pre}\right\}} * \mathbb{1}_{\{high^q_{hj}\}} \\ &\quad + \theta_{hj} + \theta_t + \gamma_j * (t - t_{min_j}) + \varepsilon_{hjt} \\ \mathbb{1}_{\left\{high^q_{hj}\right\}} &= \mathbb{1}_{Q_{hj,pre} \geq prctile75\{Q_{hj',pre}\}^H_{h'=1}} \end{aligned}$$

where $\beta^{Info}_{quintile,low^q}$ now estimates the treatment effect, by price quintile, for lower volume products; and $\beta^{Info}_{quintile,low^q} + \beta^{Info}_{quintile,high^q}$ now estimates the treatment effect, by price quintile, for higher volume products.

The estimates in Figure 9 show that the price treatment effect is largest for high-volume hospital-products in the upper part of the price distribution. At -\$71, the treatment effect for high-quantity hospital-products is more than triple the effect for low-quantity hospital-products in the preferred specification. These results are consistent with a positive effort cost of using information and renegotiation.[32]

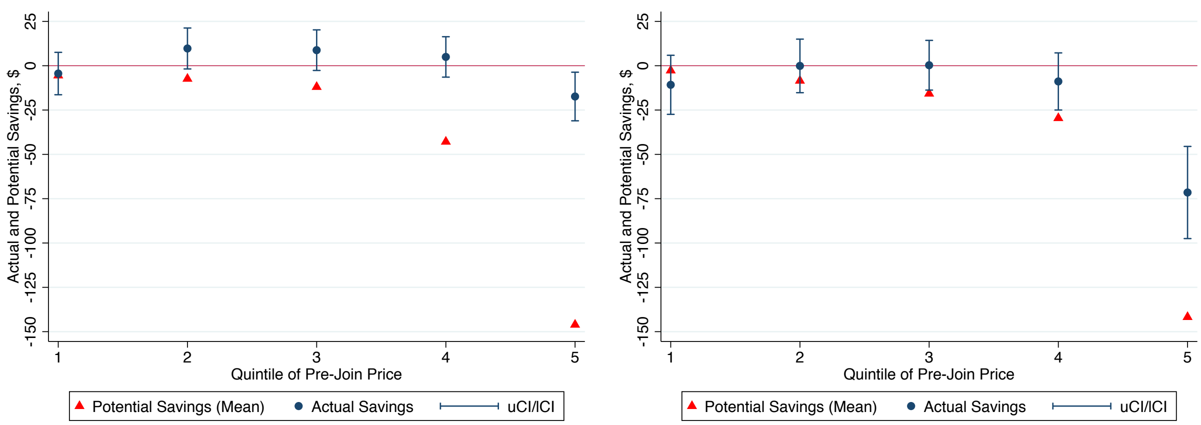

(a) Low Quantity **(b)** High Quantity

Version	Achieved savings: Pre-info price quintiles ($\beta^{Info}_{quintile,low^q}$ =)					Pre-info price quintiles ($\beta^{Info}_{quintile,low^q} + \beta^{Info}_{quintile,high^q}$ =)				
	1	2	3	4	5	1	2	3	4	5
3	-4	9	9	5	-17^{**}	-11	0	0	-9	$-71^{\dagger}$
	(6)	(6)	(6)	(6)	(7)	(9)	(8)	(7)	(8)	(13)
	Potential savings:									
	-6	-7	-12	-43	-146	-3	-8	-16	-30	-142

Authors' calculations from PriceGuide data, 2009-2014. $N = 32,453$ member-product-months. Includes 508 members. Standard errors clustered at hospital-product level shown in parentheses. Superscript (†) indicates significant difference from zero at the 1% level; (**) at the 5% level; (*) at the 10% level.

Figure 9: Treatment Effect Estimates Across the Price and Quantity Distributions

To put this in perspective, recall that high-price, high-volume products are those that would be flagged by the benchmarking database interface as targets for renegotiation according to the "potential savings" analytic. In order to relate achieved savings to potential savings, we compare the estimates above to the heterogeneity in prices in each price and quantity category. Specifically, using the subset of our data containing each hospital's year

[32]In Appendix F.3, we show the results of this specification for a variety of alternative samples and our results are unchanged.

of pre-information data, we extracted a hospital-product specific fixed effect, controlling for product-time fixed effects ($\hat{p}_{hj}$ from the regression $p_{hjt} = \hat{p}_{hj} + \hat{p}_{jt} + \hat{\varepsilon}_{hjt}$). Potential savings is then defined, for above-mean hospital-products, as the difference between this and the mean of the distribution across hospitals for each product: $PS_{hj} := \max\{0, \hat{p}_{hj} - \bar{p}_j\}$.

Figure 9 displays average potential savings per stent for each price/quantity category. High-price (and particularly high-price, high-quantity) hospital-products achieved substantial savings – in the top quintile of the price distribution, hospitals achieved 12-51% of potential savings. Savings are not substantial for lower points in the price distribution, but this is consistent with the fact that potential savings are mechanically smaller for hospital-products already achieving lower prices. Across all hospitals, savings on drug-eluting stents are estimated to be 26% of potential savings.

In sum, the heterogeneity results indicate that the treatment effects of information are largest exactly where we most expect to see them – among hospital-products in the upper part of the price distribution pre-join, among products with the largest budgetary impact on hospitals ex ante, and in hospital-products with the largest potential savings. The fact that quantity matters suggests costs of information analysis or action as frictions that sustain price variation. That all "potential savings" are not realized, even when large quantities are at stake, suggest further frictions independent of information, such as strong physician brand preferences, are important as well. As shown in Appendix F.3, we see similar patterns when we consider different sets of fixed effects and time trends, when we modify the regression sample to focus on only the twelve months pre- and post-information, when we identify treatment effects based only on the information shock of database "join", and when we limit the sample to hospitals only. We also see similar patterns for the bare metal stent product category, though the top quintile treatment effects are smaller as would be expected given the lower volumes at stake. See Appendix F.4 for detail.

5.2.2 Differentiating between agency and asymmetric information mechanisms

The β^{Info} estimates thus far have provided a treatment effect of access to the benchmarking information, subsuming both the agency and asymmetric information mechanisms that market participants put forth, as outlined in our Section 3. In this Section, we will separate these two theories. The key insight that we rely upon is that the different theories require different *timing* of access to information – using the benchmarking data to resolve asymmetric information about the seller's bargaining type requires concurrent access to the data, while using the benchmarking data to better resolve agency problems within the hospital by designing negotiator contracts with higher powered incentives and less risk only requires the knowledge that the data will eventually be available for the negotiator's performance review.

New product introductions offer variation in the timing of access to information, allowing us to separate these theoretical mechanisms. The fact that no information is available in the database on prices hospitals negotiate for a new product during the first several months after its introduction means that, during this time, differences between prices negotiated for that product by hospitals post- and pre-join must be attributable to the agency mechanism, not asymmetric information.

In practice we implement this separation of the two mechanisms by including two separate indicator variables regarding join and information. The first term is simply an indicator for all hospital-months after the hospital joins the benchmarking database. We also add an interaction term with our join variable that is equal to one for hospital-product-months more than six months after the introduction of that product. Almost all hospitals negotiate their first contract with a new product by the first or second month after its introduction, but the resulting purchase order data will not begin to show up in the benchmarking database until month three or four. By month six, there are enough observations in the database for a hospital to develop a useful estimate of its place in the price distribution for that product. The specification we estimate is:

$$\begin{aligned} P_{hjt} &= \beta^{Agency}_{quintile} * \mathbb{1}_{\{post^{join}_{ht}\}} * \mathbb{1}_{\left\{quintile_{hj,pre}\right\}} \\ &\quad + \beta^{Info}_{quintile} * \mathbb{1}_{\{post^{join}_{hjt}\}} * \mathbb{1}_{\left\{quintile_{hj,pre}\right\}} * \mathbb{1}_{\{(t-t_{min_j})>6\}} \\ &\quad + \theta_{hj} + \theta_t + \gamma_j * (t - t_{min_j}) + \varepsilon_{hjt} \end{aligned}$$

where $\mathbb{1}_{\{(t-t_{min_j})>6\}}$ is a dummy equal to one greater than six months after a product's entry date and zero during the first six months when zero to little concurrent benchmarking information is available.[33] The results are shown in Figure 10.

The estimates consistently suggest that the asymmetric information effect explains a substantial portion of the effect of information on prices. For the hospital-product fixed effects model shown in the Figure, the effect of information on price is approximately -$30 in the top 20% of the price distribution, which is nearly identical to our main results. The estimated effect of agency on price is extremely noisy, and not statistically significantly different from zero after controlling for unobserved differences across hospital-product combinations.[34]

[33]The interpretation of the interactions between the post-join, or "agency," effect and the indicators for position in the price distribution is slightly different than that of the interactions in our previous results. For the "info" effect, we estimate the effect of information on prices for observations that were previously high- (or low-) priced. For the "agency" effect, we instead estimate the effect of having *joined* the database on the price level within each price quintile – the regression determines whether joining the database, absent information, shifts the upper, or lower, part of the price distribution.

[34]In Appendix E, we see that in specification Versions 1 and 2 with hospital and product fixed effects only, β^{Ag} and β^{Info} each explain approximately half of the total effect. It is possible that the difference

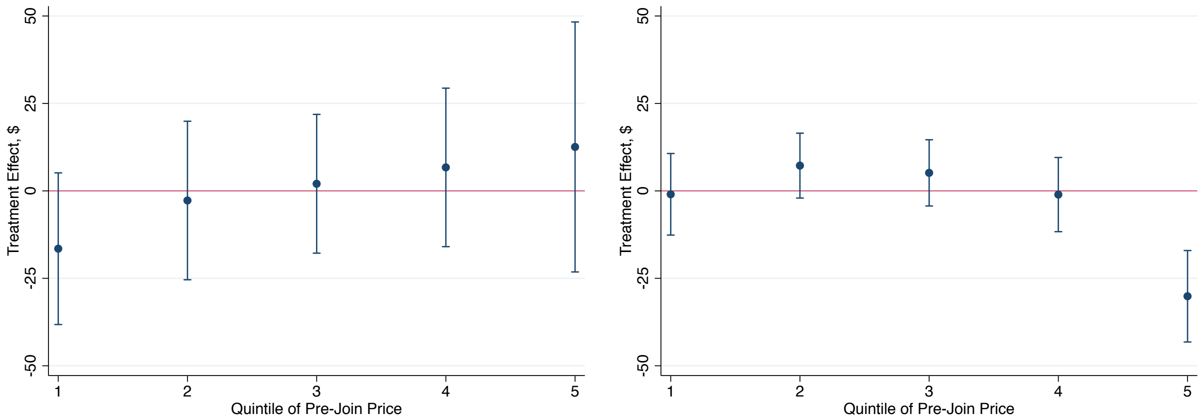

(a) Agency **(b)** Asymmetric Info

Version	Pre-info price quintiles ($\beta^{Agency}_{quintile}=$)					Pre-info price quintiles ($\beta^{Info}_{quintile}=$)				
	1	2	3	4	5	1	2	3	4	5
3	−17	−3	2	7	13	−1	7	5	−1	$-30^{\dagger}$
	(11)	(12)	(10)	(12)	(18)	(6)	(5)	(5)	(5)	(7)

Authors' calculations from PriceGuide data, 2009-2014. $N = 32,453$ member-product-months. Includes 508 members. Standard errors clustered at hospital-product level shown in parentheses. Superscript (†) indicates significant difference from zero at the 1% level; (**) at the 5% level; (*) at the 10% level.

Figure 10: Treatment Effect Estimates Across the Price Distribution, Separating Agency and Asymmetric Information Mechanisms

While our interpretation of the event study evidence is that bias due to endogenous timing of join is unlikely to be large, it is important to note that in the most pessimistic case where the timing of join correlates with other hospital activities unrelated to benchmarking information that decrease prices, this bias will be captured in β^{Ag} but not β^{Info}. This is because in our research design, β^{Ag} is identified by any differences between pre- and post-join hospitals that are *not* due to contemporaneous access to information. Consistent with the discussion of the event studies by quintile, these estimates would suggest that, if anything, the total effect of information on price estimated in the first quintiles specification is biased toward zero.

Our most robust finding is that for the presence of asymmetric information in these negotiations. Our finding of a statistically and economically significant (and free of join timing bias) β^{Info}—concentrated among those paying the highest prices before obtaining information—is consistent with the theory of asymmetric information bargaining based on Rubinstein (1985).

One implication of this result is that asymmetric information may be among the effects

between the results with and without hospital-product fixed effects points to challenges with attenuation bias in the hospital-product fixed effect models, which leave very little identifying variation, especially for β^{Ag}.

driving the heterogeneity found in bargaining parameter estimates in studies using full information Nash Equilibrium of Nash Bargaining models, such as Crawford and Yurukoglu (2012) and Grennan (2013, 2014). It suggests that these information and incentive issues should be kept in mind when thinking about the factors driving bargaining outcomes. A corollary to this is that when considering counterfactuals with negotiated prices, it may be important to consider how information might change in the counterfactual regime.

5.2.3 Price changes with "sticky" contracts

All of the price coefficient estimates reported above can be described as capturing the combined effect of information on the probability that price negotiation occurs *and* on prices arrived at during each price negotiation. We consider this to be the treatment effect of interest for policy, as it estimates the overall value of access to benchmarking information for decreasing the total spend of hospitals on medical inputs over time, taking into account the stickiness of prices in real-world hospital-supplier contracting. That is, the above estimates measure the treatment effect of information on prices paid, whereas another object of interest would be the treatment effect of information on prices negotiated (which requires that renegotiation take place). In this Section, we separately consider the effects of information on price *conditional on renegotiation* and on the *likelihood of renegotiation*.

In order to estimate these two effects, we flag hospital-product-month observations in which renegotiation is observed.[35] We then estimate the usual price quintiles specification on the subset of months in which renegotiation takes place:

$$P_{hjt} \quad = \quad \beta^{Info}_{quintile} * \mathbb{1}_{\{post_{hjt}\}} * \mathbb{1}_{\{quintile_{hj,pre}\}} + \theta_{hj} + \theta_t + \gamma_j * (t - t_{min_j}) + \varepsilon_{hjt}$$

as well as a specification where the dependent variable is a dummy for renegotiation:

$$\mathbb{1}_{\{reneg_{hjt}\}} \quad = \quad \beta^{Info}_{quintile} * \mathbb{1}_{\{post_{ht}\}} * \mathbb{1}_{\{quintile_{hj,pre}\}} + \theta_{hj} + \theta_t + \gamma_j * (t - t_{min_j}) + \varepsilon_{hjt}.$$

In the main estimation sample, renegotiations take place in 9% of observations (member-

[35] We sort transactions for each hospital-product by month and group observations with the same price together within month. We then flag each hospital-product-month as including a renegotiation event if we observe that prices change *and* that the price change "sticks" for two cumulative months after the renegotiation event (or until the final observed purchase for that member-product). This is a relatively conservative method for flagging renegotiations; the results are qualitatively similar (though larger in magnitude) using a less conservative method that flags all months in which average prices change. Of course with transactions data we cannot observe if a renegotiation took place and price remained the same. We take some comfort that our measure results in frequency of renegotiations similar to the annual contract structure that is common in the industry.

product-months with any transactions). Transactions do not occur in every month for every hospital-product, so this corresponds to a little under one renegotiation per year for the average hospital-product (for which we observe a time horizon of at least one year). In the pre-information sample, prices decrease on average by \$25 at each renegotiation. Hence, we would expect small price changes to occur if information led to larger price decreases at each renegotiation *or* if information increased the likelihood of renegotiation. The results are shown in Figure 11.

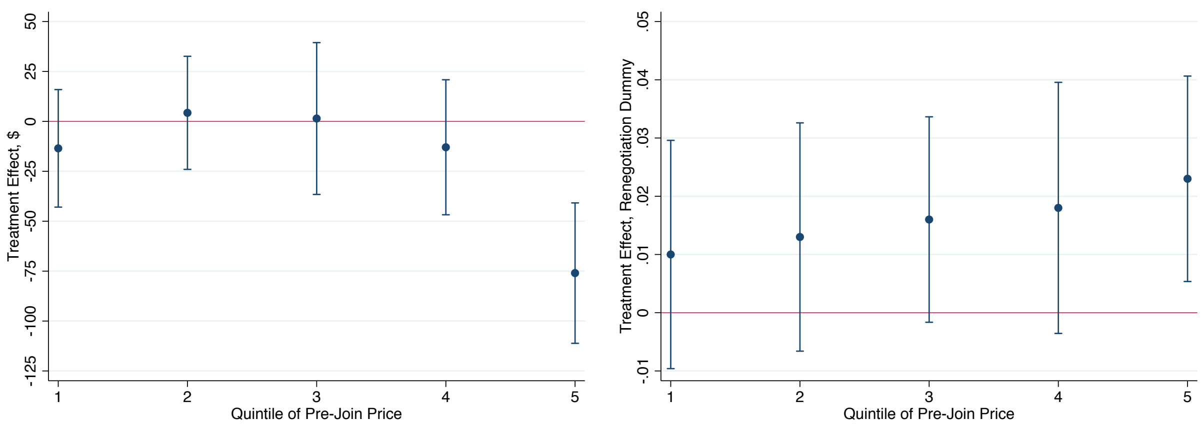

(a) Conditional on Renegotiation

(b) $\mathbb{1}_{\{reneg_{hjt}\}}$

Version	Pre-join price quintiles ($\beta^{Info}_{quintile}$ =)					Pre-join price quintiles ($\beta^{Info}_{quintile}$ =)				
	1	2	3	4	5	1	2	3	4	5
3	-14	4	1	-13	$-76^{\dagger}$	.01	.013	.016*	.018	.023**
	(15)	(14)	(19)	(17)	(18)	(.01)	(.01)	(.009)	(.011)	(.009)

Authors' calculations from PriceGuide data, 2009-2014. $N = 6,510$ member-product-months in regressions conditional on renegotiation. $N = 32,453$ member-product-months in renegotiation dummies regression. Includes 508 members. Standard errors clustered at hospital-product level shown in parentheses. Superscript (†) indicates significant difference from zero at the 1% level; (**) at the 5% level; (*) at the 10% level.

Figure 11: Treatment Effects Conditional on Renegotiation and on Occurrence of Renegotiation

The left panel of Figure 11 shows that the price decrease at renegotiation is about \$75 larger when hospitals have access to benchmarking information and learn that their previous prices were relatively high. This effect of information on price conditional on renegotiation is approximately three times larger than the \$27 effect of information on price paid. Thus the impact of transparency in the form of benchmarking information is substantially affected by the friction of getting back to the negotiating table.

The right panel of the Figure shows the effect of information on the likelihood of renegotiation throughout the price distribution. The only effect that is statistically significant at the 95 percent level is again the top quintile of the price distribution, where information increases the probability of renegotiation by 2.3 percentage points, which is about one quarter the baseline probability of renegotiation. Point estimates in other quintiles are positive but smaller and not significant at conventional levels. To the extent this is not simply a statistical coincidence, it could be due to a slight increase in efforts to get to the negotiating table or change in the frequency with which renegotiation results in zero price change among those with information.

5.3 Aggregate Estimates and Policy Implications

In this Section, we use our treatment effect estimates – properly weighted according to the observed volume and price distributions – to calculate the savings achieved due to access to benchmarking information. Figure 12 displays the distribution of savings achieved by the hospitals in our sample, in terms of total savings per hospital per year. On average hospitals gaining access to benchmarking information save $4,152 per year on stents, but this conceals substantial heterogeneity. 10 percent of sample hospitals are expected to save more than $30,000 on stents alone. For context in the scheme of catheter lab input costs, this is approximately half of a catheter lab nurse's annual salary.[36] At the other end of the distribution, the point estimates indicate that some hospitals may be worse off after joining the database, but these effects are not statistically significant at conventional levels.

These numbers speak directly to the value of the benchmarking service whose data we study. They also provide a first step toward projecting aggregate savings that might be achieved by transparency policies proposed by policymakers. To perform this thought experiment, we extrapolate based upon our treatment effect estimate for coronary stents and its relation to the "potential savings" that could be achieved if all hospitals paying above the mean price were to pay the mean price for a given product at a given time. We first reweight our estimates for stents in our sample based on a more representative set of US hospitals with catheter labs; we then extrapolate the percent of potential savings achieved from benchmarking information on stents across the top 50 hospital supply categories in our database (which together represent one third of total supply spending observed in the database).[37]

To extrapolate to a representative sample, we flag each MRG sample hospital based on the position in the price and quantity distributions it would have held for its most recent observation month. We then apply the treatment effects we estimated, allowing for price and quantity heterogeneity, to all MRG sample hospital observations. As discussed above, the

[36] http://www.payscale.com/research/US/Job=Cardiac_Catheterization_Laboratory_Registered_Nurse/Hourly_Rate. Last accessed June 2016.

[37] Of the top 50 product categories, we omit one – "office supplies" – as being too broad.

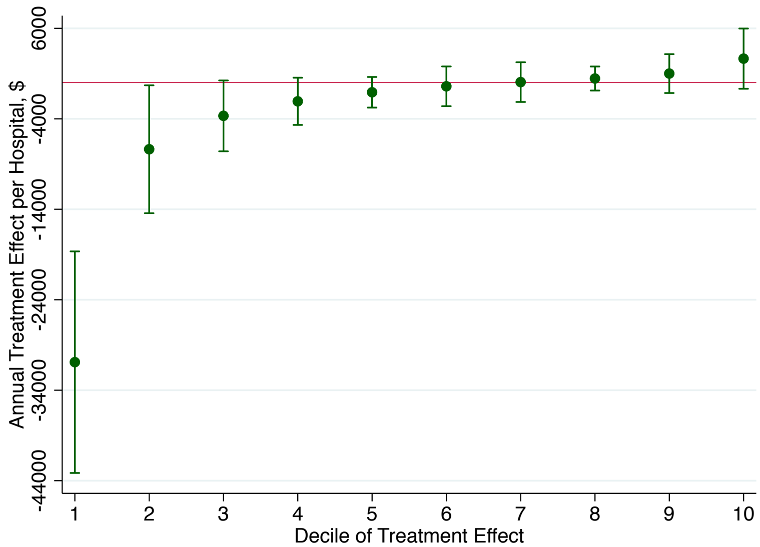

Treatment Effect by Decile ($TE_{decile}/year/hospital$=)									
10	20	30	40	50	60	70	80	90	100
$-30,902$	$-7,364$	$-3,675$	$-2,067$	$-1,066$	-406	58	456	$1,001$	$2,652$
(6,250)	(3,606)	(1,998)	(1,333)	(863)	(1,121)	(1,121)	(678)	(1,096)	(1,699)

Authors' calculations from PriceGuide data, 2009-2014. Bootstrapped standard errors based on $1,000$ draws from full variance-covariance matrix of parameter estimates shown in parentheses. Original standard errors clustered at hospital-product level. Achieved savings weighted up to hospital level using purchase quantities of each product per month within each hospital's reporting horizon.

Figure 12: Achieved Savings Across Hospitals ($/year/hospital)

estimation sample hospitals are larger and have ex ante lower prices than the MRG sample hospitals. We find higher savings for high price and high quantity hospitals, so there are two countervailing effects that cause this counterfactual savings calculation to differ from the above estimates. On balance, we estimate that representative hospitals would have achieved $5,052 in average savings on stents per year (22 percent more than our regression sample) with access to benchmarking data.

Thus, adjusting for prices and quantities suggests that our estimates are conservative with respect to the US population of facilities overall. This cannot, of course, adjust for unobservable ways in which our sample might differ from the average US hospital with a catheterization lab (e.g., joiners may have better management practices and/or be better able to utilize databases). Another possible caveat to considering the rollout of a broader transparency policy at the national level is the extent to which such a policy might change supply side price cuts, collusion, or obfuscation. However, to the extent manufacturers know when a hospital has access to this information (which we understand they do), then our estimates do incorporate effects on hospitals who get "bad news" when joining the database or reluctance of manufacturers to lower prices to database members.

Finally, to obtain an estimate of the aggregate dollar savings that transparency in device pricing might deliver, we estimate the savings that would be achieved across the top 50 product categories, assuming that each would save the same 26 percent of potential savings as on drug-eluting coronary stents. Many of the top 50 product categories are known physician preference items, such as implantable cardiac rhythm management devices and orthopedic implants (and coronary stents). Others are perhaps more commodity-like, such as bone screws. We estimate that the average US hospital with a catheter lab would save $256,476 annually on top product categories after accessing database information. This is over $480 million per year across all such hospitals.[38] While arriving at this precise number requires a number of simplifying assumptions, as an approximation we find it to be compelling evidence that substantial savings are at stake in hospital purchasing more generally.

6 Conclusion

This paper conducts the first empirical study of the impact of transparency on price negotiations in business-to-business markets. Our empirical context is hospital supply purchasing, an area where there has been keen interest in information as a way to decrease hospital supply costs. We use new data on all purchase orders issued by over 16 percent of US hospitals from 2009-14 and differences-in-differences research designs to compare the prices negotiated by hospitals with and without access to benchmarking information. Hospitals who gain access to benchmarking information see subsequent savings in the products for which they were paying relatively high prices. The savings amount to 26 percent of the potential savings available (relative to the mean price paid by other hospitals, holding product and time fixed). These estimates provide evidence on the potential economic impacts of the rise in benchmarking data services marketed towards buyers in business-to-business markets, as well as calls for greater transparency in these markets by policymakers. Extrapolating our estimates to all US hospitals with catheter labs and the top 50 hospital supply categories suggests such policies could save over $480 million annually.

Our tests of the mechanisms behind these information effects imply that the value of information is attenuated by the costs of putting the information to use. The evidence suggests that there are costs consistent with time-constrained negotiators (gains are focused in high quantity items where the most money is at stake) and also stickiness of business-to-business contracts (long term contracts may not be renegotiated for some time). The

[38] This figure based on the approximately 1,875 US hospitals listed as having catheterization capabilities in the 2012 American Hospital Association annual survey. It is an underestimate in the sense that it does not account for the fact that many products in the top 50 are purchased by hospitals beyond those with catheter labs.

latter friction is a fundamental feature of many business-to-business markets. However, the time and effort cost of accessing and/or using information could be reduced as technology improves. As both information and analytics are increasingly important products in the modern economy, this suggests a path for future research.

We examined two potential theories for how benchmarking information might be used in a business-to-business setting – asymmetric information about seller bargaining parameters and buyer-side negotiator agency. We found robust evidence for the asymmetric information theory, but noisy evidence for agency. We see modeling frictions in the use of information and the potential for information to affect within-firm agency frictions in negotiation as two especially interesting areas for future theory development. Fully developed theories might suggest more precise tests to better understand these mechanisms. A more structural model, combined with variation in market structure, could also be useful in determining the extent to which theories of supply side phenomenon such as obfuscation or collusion might be important for considering the policy proposals for greater transparency in medical device markets.

Quantifying the extent to which further savings due to greater information are possible in our setting (that is, rather than other sources of price heterogeneity such as physician brand preferences or negotiator expertise accounting for the remaining price variation) would require more structural assumptions than we make here. However, an important takeaway from our study for structural empirical work in bargaining is that information can play a role in the observed heterogeneity in prices, and thus changes in information may be important to consider in counterfactual analyses in markets with negotiated prices.

While our results suggest that, on net, policies or intermediaries that increase transparency may indeed lower the prices hospitals pay for medical supplies, our detailed analysis focuses on coronary stents. Stents are an example of the "physician preference" medical technologies that receive the greatest policy attention (and are important in their own right), but other categories with different demand and supply structures might see different results. Variation across product markets in terms of supply side competition and complexity of contracts (for example, nonlinearities or multi-product bundling) thus represents rich opportunities for future empirical analysis of information in business-to-business bargaining. Such work would require expanding the empirical toolkit to analyze complex contracts when the contract terms themselves may not be observed.

Appendices

A Data Appendix

The raw transactions data contain 82.5 million observations for 2,111 members across 3,375 product categories and 1.9 million SKUs. In this project, we focus on coronary stents – for this subset of products, we observe 716,752 observations for 624 members across 20 brands.

In order to construct each member's information set upon joining the database and later upon new products' entry, we linked the transactions data with several additional datasets relating each individual login ID with associated members and login activity. The first of these is a "clickstream" dataset containing timestamped observations of unique IDs' login activity, to the minute.[39] The second is a membership "hierarchy" file linking individual database members with parent accounts for those cases where members are part of a set of health care organizations purchasing membership jointly. The third file associates each login ID with the given individual's direct-linked member organization and broader access level – i.e., the individual with ID X may work with member 1 but have access to the data for all members 1, 2, and 3 under the same parent organization. For individuals with higher-level access, data for all associated members is automatically reported to them when they log in to the database. Accordingly, when we observe a click for a particular login ID, we associate that click with all linked member organizations for which that login has access.

We use the above-described datasets to determine each date of login for each member-product. The goal of this exercise is to determine when benchmarking data on a given product would enter each member's information set. For non-entering products, this is the date of the first observed login for each member. For entering products, this is the date of the first observed login for each member after six months post-entry. This is to account for the lag between transactions occurring for new products and transactions being submitted by member facilities, loaded into the database, and viewed by members logging in. The left panel of Figure 13 displays the steady increase over time in the count of members for which transactions for the average entering product are observed, and demonstrates the lag with which members' transactions become available in the database for benchmarking purposes.[40] In any given month in the year following new product entry, there are on average 56 more members we observe having transactions for new products than there are members whose

[39]Each login ID corresponds to a unique individual's account within a member. For example, a given database member may have login accounts for a number of different purchasing managers, administrators, and department clinicians.

[40]There may be an additional lag before joining hospitals become informed if they do not frequently log in to the database.

transactions data are actually loaded into the database. We also observe transactions data for members that have not yet joined the benchmarking database – in the year following product entry, nine percent of members observed in the average month are pre-join. To see this more concretely, the right panel of Figure 13 displays the trend in the number of hospitals purchasing the average new product, overall and for pre-join hospitals in particular. For each new product, we observe 10-15 hospitals in the pre-join state within a short window after product entry, and the time period for our study contains many meaningful product introductions. This is precisely what allows us to separately identify effects of *joining* the database per se, versus actually having access to information *on a particular product*, on prices. In Figure 4 in the main text, we note the timing of entry of nine new products between 2010 and 2014 (of the twenty products sold during this time period overall).

Figure 13: Transactions Observed After New Product Entry

(a) Cumulative Members Purchasing **(b)** Members Purchasing, by Join Status

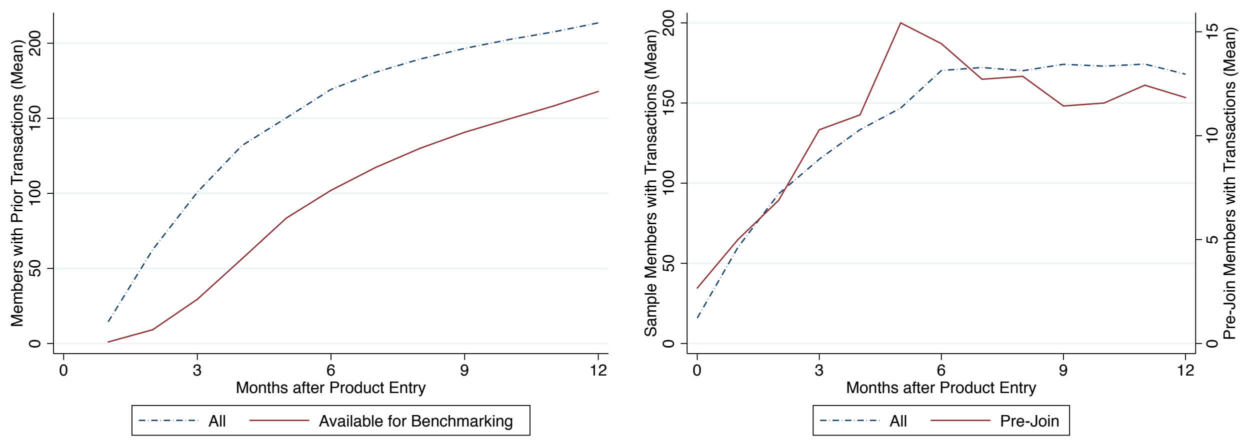

Authors' calculations from PriceGuide data.

We use the linked login and transactions data to calculate each member-product's position in the pre-information price and quantity distributions. All calculations are specific to the first *informed* login for the given member-product, as defined above. Following the approach used in the database to aggregate data across all members' transactions and present summary statistics to members logging in at a point in time, we calculate percentiles of the price distribution using all members' *most recent* transactions for the same product, in the past year, that would have been loaded into the database prior to the login. We calculate percentiles of the quantity distribution using all members' total quantity purchased per month for the same product in the past year. Across all specifications, we consistently include only those observations that can be used to estimate the richest specification with interactions based on point in the price and quantity distributions – this requires that we observe pre-information data for the given member-product. The final analysis sample of drug-eluting stents includes 395,271 transactions for 508 members and thirteen brands in 72 months between January 2009 and December 2014. Seven of the included drug-eluting stent brands entered the market during this time horizon. We collapse the transaction-level data to perform all analyses at the member-product-month level (with mean price as the dependent variable). We do this in order to avoid overweighting member-products that tend to have multiple transactions per month. The analytic sample contains 32,453 observations.

Table 3 summarizes the stent transactions data for the sample on which we perform our estimation. The average sample hospital submitted stent transactions in 41 months. In a given month, sample hospitals spent $80,000 on 59 stents. The Table shows each statistic separately by hospital bed count; larger hospitals generally submitted more months' data and, as logic would indicate, purchased more stents per month for a greater total monthly expenditure. Hospitals with $\geq$ 500 beds spent more than double the amount that the smallest hospitals did on stents per month. The vast majority of transactions in our data are for drug-eluting (as opposed to bare metal) stents; in the remainder of our description of the data and in our results, we focus on drug-eluting stents. See Appendix F.4 for bare metal stent data and results.

Table 3: Summary Statistics – Stent Hospitals Only

Bed Size	Members	Months of Data	Monthly Exp. ($ k)	Monthly Quantity	% DES
0-99	52	31.4	59.0	45.0	82.0
		(19.4)	(56.6)	(44.4)	(10.8)
100-199	102	40.0	45.5	33.5	81.6
		(20.3)	(43.3)	(31.5)	(12.1)
200-299	117	43.4	55.6	40.7	77.4
		(22.0)	(45.9)	(33.5)	(14.3)
300-399	83	41.0	73.5	53.6	79.7
		(20.7)	(46.9)	(33.1)	(11.6)
400-499	47	41.4	128.9	93.5	79.6
		(21.3)	(92.1)	(65.5)	(12.2)
500+	107	45.9	135.2	97.7	81.1
		(22.0)	(94.2)	(65.6)	(9.5)

Authors' calculations from PriceGuide data.

The heterogeneity in prices observed in our sample is not well-explained by hospital characteristics that might seem a priori to be important for negotiation. For example, we observe no clear relationship between hospital size and stent prices. See the left panel of

Figure 14, in which we display a box plot of drug-eluting stent prices for each category of bed count.[41] Mean prices are, if anything, increasing in bed count, though the differences are not statistically significant. Part of this (lack of) relationship is likely due to the heterogeneity in purchasing behavior across hospitals with similar bed counts – e.g., small cardiac specialty hospitals may purchase stents in greater quantities than similarly-sized acute care hospitals. In the right panel of Figure 14, we also show box plots of stent prices for each decile of monthly stent purchasing volume. Here, we do see a relationship between "size" and price – the hospitals with the smallest purchasing volumes have price distributions which are spread slightly upward relative to that of the hospitals with the largest volumes, so that low-volume hospitals' prices have larger means and variances than high-volume hospitals. For drug-eluting stents, 10^{th} decile hospitals' prices are 7% lower than those obtained by 1^{st} decile hospitals. These differences are economically and statistically significant; however, the price distributions for the high-volume and low-volume hospitals overlap substantially, so that there is a great deal of unexplained hospital price heterogeneity conditional on purchasing volume.

A.1 MRG Sample Comparison

For Jan 2010 - Jun 2013 we have access to additional data on hospital stent purchasing from the MRG Markettrack™ survey, which is intended to provide a representative sample of usage and price patterns across the US. This allows us to further check the representativeness of the sample of hospitals joining the price benchmarking database. In the text we compared quantities used and average prices. Figure 15 provides further details on the full distribution of prices across hospitals in the two samples of 143 pre-join vs 107 MRG hospitals. The prices paid in the two samples are statistically close to one another, with the average prices paid (controlling for product-time trends) in the MRG sample slightly higher (mean \$1631, s.d. \$120) than those paid by hospitals in the estimation sample (mean \$1666, s.d. \$149) during the period before they joined the benchmarking service. These pre-join hospitals do have a slightly larger upper tail of high prices, with a 80^{th} percentile of 1743 versus 1730 in the MRG sample, but this difference is not statistically significant.

[41] "Prices" are hospital fixed effects obtained from a regression of price on hospital and product-month fixed effects.

Figure 14: Distribution of Prices Across Hospitals

(a) By Bed Count **(b)** By Stent Volume Decile

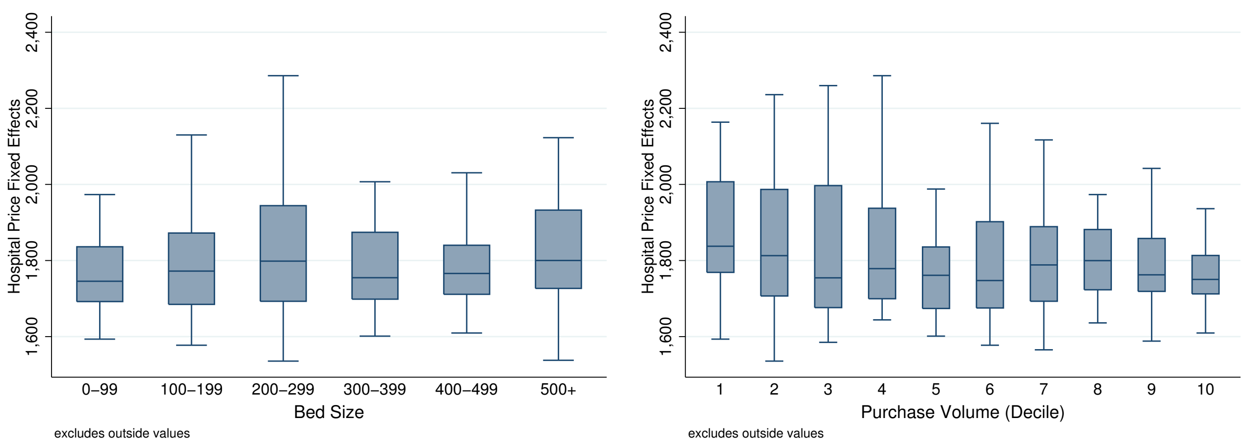

Drug-Eluting Stent Prices by Size Category (Regression Results)

$\beta^{\mathbb{1}\{h \in \text{ bed size } x\}} =$						$\beta^{\mathbb{1}\{h \in \text{ decile } x\}} =$									
$0-99$	$100-199$	$200-299$	$300-399$	$400-499$	500+	1	2	3	4	5	6	7	8	9	10
1,785	1,801	1,852	1,778	1,805	1,829	1,916	1,864	1,830	1,832	1,766	1,784	1,801	1,807	1,787	1,774
(38)	(27)	(30)	(25)	(32)	(22)	(54)	(40)	(44)	(36)	(26)	(31)	(29)	(25)	(25)	(32)

Authors' calculations from PriceGuide data. Estimated mean hospital fixed effects within bed size categories and decile of monthly purchase volume. Hospital fixed effects obtained from regression of price on hospital and product-month fixed effects, pre-join data only. Mean estimates from regression of fixed effects on indicators for size. Standard errors from nonparametric bootstrap of entire procedure, resampling at hospital level.

Figure 15: Pre-Join Distribution of Prices Across Hospitals: Comparison to MRG Sample

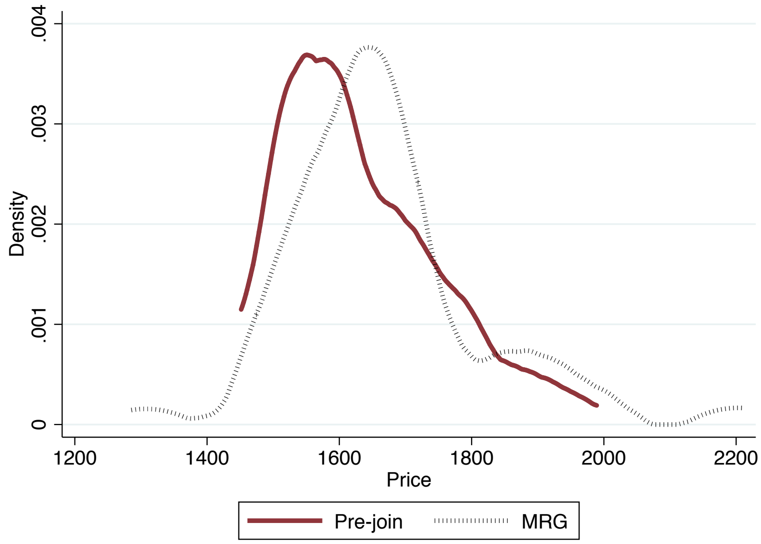

	mean	sd	cv= $\frac{\mu}{\sigma}$	$20^{th}_{\%ile}$	$50^{th}_{\%ile}$	$80^{th}_{\%ile}$	N_H
Pre-join $\bar{p}_h$ (within product-month)	1631	120	0.07	1539	1607	1743	143
MRG $\bar{p}_h$ (within product-month)	1666	149	0.09	1563	1651	1730	107

Authors' calculations from PriceGuide data and MRG survey. Sample restricted to Jan. 2010 - Jun. 2013 when MRG data available. Hospital average prices with product-month trends removed.

B Checks for Standardization and Share-based Contracts

It is important for our analysis that the prices we observe are comparable across observations in the sense that there are not important contract dimensions that we do not observe (e.g., bundling, exclusivity, or market share based contracting). Our conversations with industry participants indicate that stents tend to have simple linear price contracts, so we assume that transactions data capture real prices. Here, we examine this assumption in the data.

In panels (a) and (c) of Figure 16, we show histograms of total unique manufacturers and stent products (brands) purchased over each quarter by each hospital in the sample. The vast majority of hospitals purchase multiple brands from multiple manufacturers, rather than purchasing a single most-preferred product for the whole facility. Panels (b) and (d) show these histograms for only hospitals above the 25^{th} percentile in total stent volume, and show even less evidence of "exclusivity" – fewer than three percent of hospital-quarters involve a single product and fewer than seven percent involve a single manufacturer. The fact that the majority of the already small amount of observed "exclusivity" occurs at hospitals with lower utilization is consistent with the anecdotal evidence that exclusivity does not play a systematic role in stent contracting – with true linear price contracts, "exclusive" purchasing patterns are more likely to be observed among hospitals with low utilization due to random variation and costly contracting.

As a further check, we look at the pricing consequences of the observed sole sourcing in the usage data. For the minority of hospitals that do happen to use only a single product or manufacturer in a given quarter, we create an indicators for $\mathbf{1}_{\{|\mathcal{J}_{ht}|=1\}}$ and $\mathbf{1}_{\{|\mathcal{M}_{ht}|=1\}}$ and regress price on each indicator and product-month fixed effects θ_{jt}. The resulting small, positive point estimates are not statistically different from zero, again suggesting that the small amount of sole sourcing observed is due to other factors besides contracting concerns.

Finally, we check for any evidence of near-exclusivity in the form of market share based contracts (which we are told are commonly used for many medical products, but not stents). Figure 17 plots the cumulative density of observations by product market share at the hospital-quarter level. We do not observe the bunching that we would expect if contracts commonly specified market share thresholds in either the full sample (panel (a)) or restricting to the most used product at each hospital (panel (b)).

Figure 16: Histograms – Number of Unique Products/Manufacturers per Hospital-Quarter

(a) Products

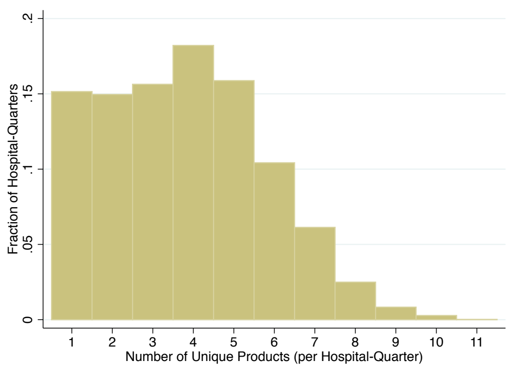

(b) Products (excluding smallest hospitals)

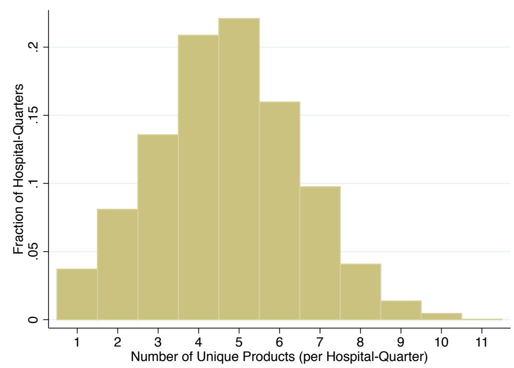

(c) Manufacturers

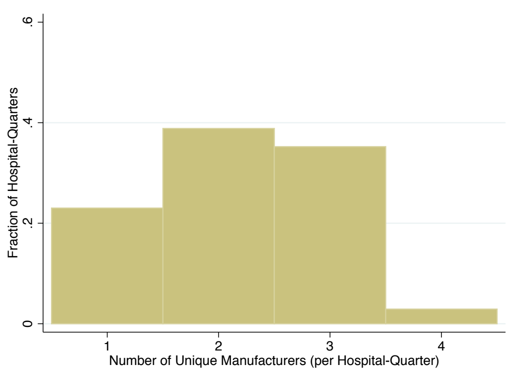

(d) Manufacturers (excluding smallest hospitals)

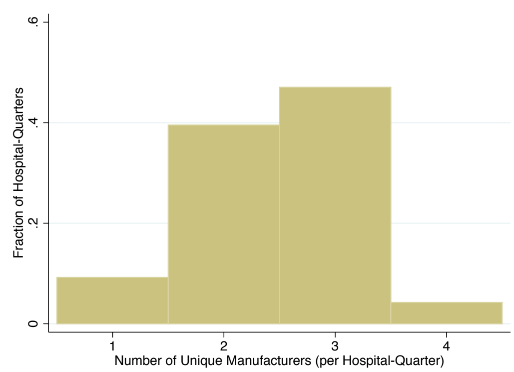

Regressions of price on indicators for exclusivity:		
Specification	β^{Excl}	(s.e.)
$p_{hjt} = \beta^{Excl}\mathbf{1}_{\{\lvert\mathcal{J}_{ht}\rvert=1\}} + \theta_{jt} + \epsilon_{hjt}$	25.3	(9.0)
$p_{hjt} = \beta^{Excl}\mathbf{1}_{\{\lvert\mathcal{J}_{ht}\rvert=1\}} + \theta_{jt} + \theta_h + \epsilon_{hjt}$	13.2	(9.5)
$p_{hjt} = \beta^{Excl}\mathbf{1}_{\{\lvert\mathcal{J}_{ht}\rvert=1\}} + \theta_{jt} + \theta_{hj} + \epsilon_{hjt}$	-2.9	(8.0)
$p_{hjt} = \beta^{Excl}\mathbf{1}_{\{\lvert\mathcal{M}_{ht}\rvert=1\}} + \theta_{jt} + \epsilon_{hjt}$	8.6	(9.4)
$p_{hjt} = \beta^{Excl}\mathbf{1}_{\{\lvert\mathcal{M}_{ht}\rvert=1\}} + \theta_{jt} + \theta_h + \epsilon_{hjt}$	-15.9	(6.5)
$p_{hjt} = \beta^{Excl}\mathbf{1}_{\{\lvert\mathcal{M}_{ht}\rvert=1\}} + \theta_{jt} + \theta_{hj} + \epsilon_{hjt}$	-4.2	(6.8)

Authors' calculations from PriceGuide data, 2009-2014. $N = 32,223$. Standard errors clustered at hospital level, $N_h = 507$ in first two specifications and hospital-product level $N_{hj} = 2227$ in product hospital fixed effects specification.

Figure 17: Cumulative distributions by market share.

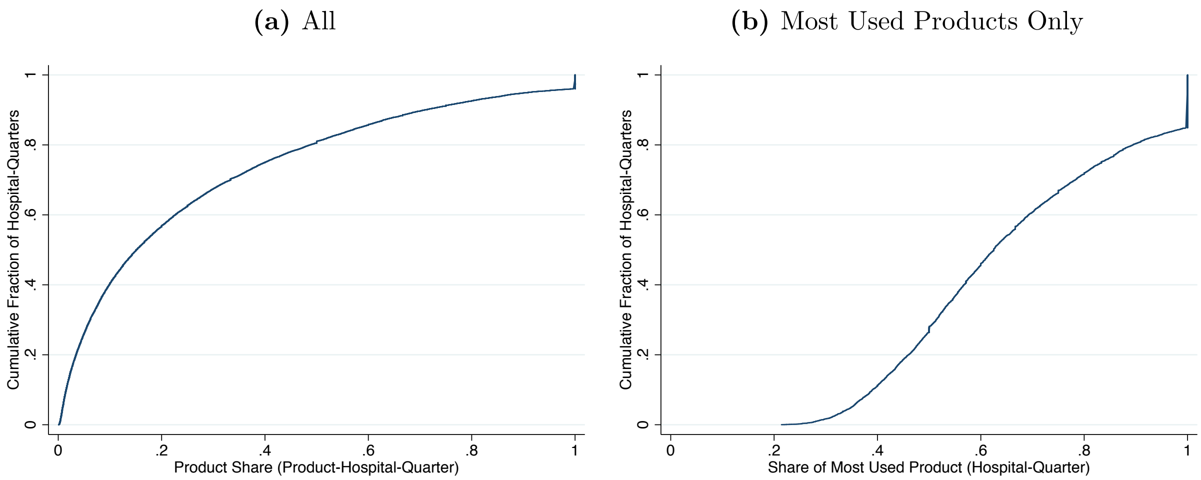

Regressions of price on indicators for exclusivity (share-based contracts):		
Specification	β^{Excl}	(s.e.)
$p_{hjt} = \beta^{Excl}\mathbf{1}_{\{\lvert\mathcal{J}_{ht}\rvert=1\}} + \theta_{jt} + \epsilon_{hjt}$	12.2	(11.5)
$p_{hjt} = \beta^{Excl}\mathbf{1}_{\{\lvert\mathcal{J}_{ht}\rvert=1\}} + \theta_{jt} + \theta_h + \epsilon_{hjt}$	-9.4	(10.7)
$p_{hjt} = \beta^{Excl}\mathbf{1}_{\{\lvert\mathcal{J}_{ht}\rvert=1\}} + \theta_{jt} + \theta_{hj} + \epsilon_{hjt}$	-2.9	(8.0)
$p_{hjt} = \beta^{Excl}\mathbf{1}_{\{\lvert\mathcal{M}_{ht}\rvert=1\}} + \theta_{jt} + \epsilon_{hjt}$	18.3	(13.5)
$p_{hjt} = \beta^{Excl}\mathbf{1}_{\{\lvert\mathcal{M}_{ht}\rvert=1\}} + \theta_{jt} + \theta_h + \epsilon_{hjt}$	-7.5	(9.9)
$p_{hjt} = \beta^{Excl}\mathbf{1}_{\{\lvert\mathcal{M}_{ht}\rvert=1\}} + \theta_{jt} + \theta_{hj} + \epsilon_{hjt}$	-4.2	(6.8)

Authors' calculations from PriceGuide data, 2009-2014. $N = 32,223$. Standard errors clustered at hospital level, $N_h = 507$ in first two specifications and hospital-product level $N_{hj} = 2227$ in product hospital fixed effects specification.

C Mapping of Bargaining Setting into Models of Asymmetric Information and Agency

C.1 Asymmetric Information about Supplier Bargaining Parameters

As discussed in the main text, we follow Rubinstein (1985) to model uncertainty of hospital buyers about the bargaining parameter of a given supplier. The model departs from the complete information model in Rubinstein (1982) in that the supplier is either of weak type with discount factor δ_w^S or strong type with discount factor δ_s^S ($1 > \delta_s^S > \delta_w^S > 0$). The supplier knows his own type, but the buyer has only a subjective prior ω_w of the probability that the supplier is the weak type.

The equilibrium split of this surplus depends on both the type of the supplier and the prior of the buyer as follows: Rubinstein (1985) shows that there exists a cutoff prior ω^* such that if the buyer is sufficiently pessimistic about the seller being the weak type $\omega_w < \omega^*$, then the buyer simply offers what she would offer the strong type in a complete information game of Rubinstein (1982):

$$p^{AI}(\omega_w < \omega^*) := c + \delta_s^S \frac{1-\delta^B}{1-\delta^B\delta_s^S} V, \tag{1}$$

and both seller types accept this offer. However, if the buyer is more optimistic about the probability that the seller is the weak type ($\omega_w \geq \omega^*$), then the buyer offers:

$$p_w^{AI}(\omega_w \geq \omega^*) := c + \delta_w^S \frac{1-\delta^{B^2}(1-\omega_w)-\delta^B\omega_w}{1-\delta^{B^2}(1-\omega_w)-\delta^B\delta_w^S\omega_w} V, \tag{2}$$

which the weak seller type accepts. The strong seller type will reject this offer, and counteroffer with a price that would make a weak seller no better off than p_w^{AI}, but that the strong seller strictly prefers:

$$p_s^{AI}(\omega_w \geq \omega^*) := c + \frac{1-\delta^{B^2}(1-\omega_w)-\delta^B\omega_w}{1-\delta^{B^2}(1-\omega_w)-\delta^B\delta_w^S\omega_w} V, \tag{3}$$

which the buyer accepts.

This equilibrium has direct implications for what we would expect to happen to prices in a move from this type of asymmetric information to complete information. First, note that $p_s^{CI} = p^{AI}(\omega_w < \omega^*) > p_s^{AI}(\omega_w \geq \omega^*) > p_w^{AI}(\omega_w \geq \omega^*) > p_w^{CI}$ (where p_s^{CI} (p_w^{CI}) is the equilibrium price for the strong (weak) supplier type with complete information). Thus the weak type seller is strictly better off with asymmetric information. The strong type seller is weakly worse off (strictly whenever the buyer's prior is sufficiently optimistic). A sufficiently pessimistic buyer is also weakly worse off without information. For more optimistic buyers, whether information would make them better off ex-ante depends on parameter values.

In our context we are interested in when a buyer might benefit from benchmarking information that reveals the seller's type, and what would happen to price in such a case. For simplicity, we assume that this information *fully* reveals a seller's type, though the qualitative results can be extended to a signal extraction problem where the information moves the buyer's prior in the direction of the truth. The intuition for how this unfolds in practice is a scenario where a manufacturer sales representative says "This is the best price I can offer. Corporate won't let me go any lower." Benchmarking information allows the hospital negotiator to perform the due diligence of checking the prices at other hospitals in order to verify or refute this statement.

Prediction 1 (Direct Information Effect on High Prices) If information is costless, pessimistic buyers will always become informed. This information will cause a proportion of the highest prices p_s^{CI} to fall to p_w^{CI} for those cases where the supplier was in fact the weak type. Thus exposure to benchmarking information should lead to some of the highest prices falling.

Prediction 2 (Direct Information Effect on High Prices with High Quantity) If information is costly to obtain (in the sense that searching and analyzing the data takes time that could be used on other productive activity), a pessimistic buyer will become informed whenever the expected benefit $\omega_w(p_s^{CI} - p_w^{CI})q$ exceeds the cost of information. This information will cause a proportion of the highest prices p_s^{CI} to fall to p_w^{CI} for those cases where the supplier was in fact the weak type. Thus exposure to benchmarking information should lead to some of the highest prices falling, among those products with the highest quantity used.

Prediction 2b (Indirect Information/Competition Effect on All Prices) With imperfect substitute products, under reasonable assumptions on how the negotiation for one product affects the disagreement payoff of other product negotiations, a fall in price of substitute product j will decrease the surplus up for negotiation for other products $-j$, leading to a decrease in the prices of other products $-j$, all else equal.[42] Thus

[42]This will be the case in any model where disagreement payoffs are a function of the prices agreed to with other manufacturers, which has been the case in the empirical bargaining literature thus far and much of the negotiation with externalities theory. It would not be the case in a model such as the Core, where disagreements are based on the primitives of willingness-to-pay and costs.

exposure to benchmarking information that leads to a fall in a high price for j should also lead to a fall in any price for other products $-j$, and the size of this fall will be increasing to the extent the products are good substitutes for j. We do not have the statistical power to test this prediction in the coronary stent sample. E.g., although we find that high-priced hospital-product hj's price falls after benchmarking information is accessed, we do not find that there is a statistically significant spillover effect on hospital-product hk's price.

C.2 Negotiator Agency

The other candidate mechanism via which we propose benchmarking information could be valuable to buyers would be through providing aggregate information to help the buying firm solve a moral hazard problem with its purchasing agent. Here we provide a specific model of information in the bargaining context that generates predictions in our setting. Modifying Holmstrom (1982) to our context, let price p_h at hospital h be as in the full information Rubinstein bargaining game. However, instead of the hospital negotiator's bargaining parameter being exogenous, the price will be a function of the hospital agent's *choice* of discount factor δ_h^B and the discount factor of the supplier, which takes value $\delta_w^S \epsilon_h$ with probability ω_w and $\delta_s^S \epsilon_h$ with probability $1-\omega_w$. As before, the discount factor of the strong supplier type is greater than that of the weak type ($1 > \delta_s^S > \delta_w^S > 0$). ϵ_h is a random term distributed uniform on $[0, 1]$. Importantly, the realization of ϵ_h is independent across hospital buyers, but whether the seller is weak or strong is common to all buyers. The realizations of both of these random variables are observable to the negotiating agents, but not to the principals who manage them.

A moral hazard problem arises in this setting because bargaining effort is costly and provides the agent disutility $v(\delta_h^B)$. The agent is compensated by some contract based on the price $m(p_h)$. The agent is risk averse in money, so the optimal solution to the agency problem involves risk sharing between the principal and the agent. Holmstrom (1982) shows that if agents face some common parameter which is uncertain from the principals' perspectives, then relative performance evaluation compared to some aggregate sufficient statistic can be used to write a better contract with each agent. The intuition in our real-world setting is one where with the benchmarking data, hospital administrators can make their negotiators' performance reviews contingent on the prices they negotiate relative to other hospitals for the same product. This motivates the following Predictions:

Prediction 3 (Monitoring Effect on Prices) If buyer negotiators are imperfect agents of the buying firm, then benchmarking information (observing the distribution of price realizations across hospitals $\{p_h\}_{h=1}^{H}$) allows the principal to estimate whether the seller is the weak or strong type, and thus reduce the risk to which the agent is exposed and write a contract which induces more bargaining effort and a lower price than in the case where only p_h is observed.[43]

Prediction 4 (Monitoring Effect on Prices with High Quantity) If buyer negotiators are imperfect agents of the buying firm, but it is costly for hospital managers to search and analyze the data in a way that allows them to write better contracts, then managers will use benchmarking information (observing the distribution of price realizations across hospitals $\{p_h\}_{h=1}^{H}$) to write a contract which induces more bargaining effort by the agent and a lower price than in the case where only p_h is observed if $(p_h(m) - p_h(m(\{p_h\}_{h=1}^{H})))q_h$ exceeds the cost of information use.

[43] The model as written has a strong prediction that this effect will be independent of price. However, in general the prediction of how the price distribution would move with information depends on where in the model the current heterogeneity is coming from. For example, if the heterogeneity were due to different levels of risk aversion among negotiators, then benchmarking information would tend to decrease the highest prices more than the lowest.

D Usage Pattern Changes and Demand

Stents and certain other expensive medical technologies are "physician preference items" where physician demand is based on strong preferences and is relatively insensitive to price – Grennan (2013) estimates own-elasticities centered around -0.32 for bare metal stents and -0.52 for drug-eluting stents. In general, however, the price benchmarking information treatment could influence quantities as well. Here, we perform a set of analyses to provide a check of this hypothesis and also provide proof of concept for how this analysis might be incorporated in the case of products where demand is more sensitive to price.

There are two primary ways in which quantities might adjust to benchmarking price information and subsequent renegotiations: (1) In a context where contracts specify quantities or market shares in addition to price, renegotiations to obtain better prices might also involve large quantity or share commitments—this effect was tested and ruled out in our analysis in Appendix B. (2) In a context where quantity is responsive to price, negotiation of better prices would lead to increased usage in the products with the largest relative price decreases. We analyze this second case here.

We run the regression specifications allowing for heterogeneous treatment effects of information depending on pre-join prices and quantities, but here with quantity q_{hjt} as the dependent variable (results are qualitatively similar and so unreported for markets shares and log transformations):

$$\begin{aligned} Q_{hjt} &= \beta^{Info}_{quintile} * \mathbb{1}_{\{post_{hjt}\}} * \mathbb{1}_{\{quintile_{hj,pre}\}} \\ &\quad + \theta_{hj} + \theta_t + \gamma_j * (t - t_{min_j}) + \varepsilon_{hjt} \end{aligned}$$

where $\beta^{Info}_{quintile}$ estimates the treatment effect, by price quintile. The results are shown in Figure 18.

If quantity were responsive to price (with downward sloping demand), then we would expect quantity/share increases in exactly the areas we see relative price decreases. Because information leads to decreases in prices for products in the high price, high quantity part of the pre-information distribution, we should expect potential quantity increases for those products and decreases for other products (whose prices haven't changed, but have become higher relative to the products with price decreases). This is not the case in Table 18, where no specification shows significantly different effects (economically or statistically) across the pre-join price distribution.

Version	Pre-info price quintiles ($\beta^{Info}_{quintile}=$)				
	1	2	3	4	5
$\theta_h + \theta_j + \theta_t + \gamma_j * (t - t_{min_j})$	0.2	1.4	1.3	0.7	2.4
	(1.5)	(1.9)	(1.7)	(1.7)	(1.1)
$\theta_h + \theta_{jt}$	-0.1	2.0	1.6	1.0	2.9
	(1.6)	(2.0)	(1.9)	(1.8)	(1.3)
$\theta_{hj} + \theta_t + \gamma_j * (t - t_{min_j})$	3.2	2.7	3.8	2.9	3.0
	(1.3)	(1.7)	(1.3)	(1.3)	(1.1)
$\theta_{hj} + \theta_{jt}$	4.0	3.9	5.2	4.2	5.0
	(1.5)	(2.0)	(1.5)	(1.5)	(1.3)

Authors' calculations from PriceGuide data, 2009-2014. $N = 32,453$ member-product-months. Includes 508 members. Version 1 includes hospital, product, and month fixed effects, plus linear product-specific time trends. Version 2 includes hospital and product-month fixed effects. Version 3 includes hospital-product and month fixed effects, plus linear product-specific time trends. Version 4 includes hospital-product and product-month fixed effects. Figure shows Version 3 results. Standard errors clustered at hospital (Versions 1 and 2) or hospital-product (Versions 3 and 4) level shown in parentheses. Superscript (†) indicates significant difference from zero at the 1% level; (**) at the 5% level; (*) at the 10% level.

Figure 18: Treatment Effect *on Quantity* Estimates Across the Price Distribution

E Alternative Versions of Controls

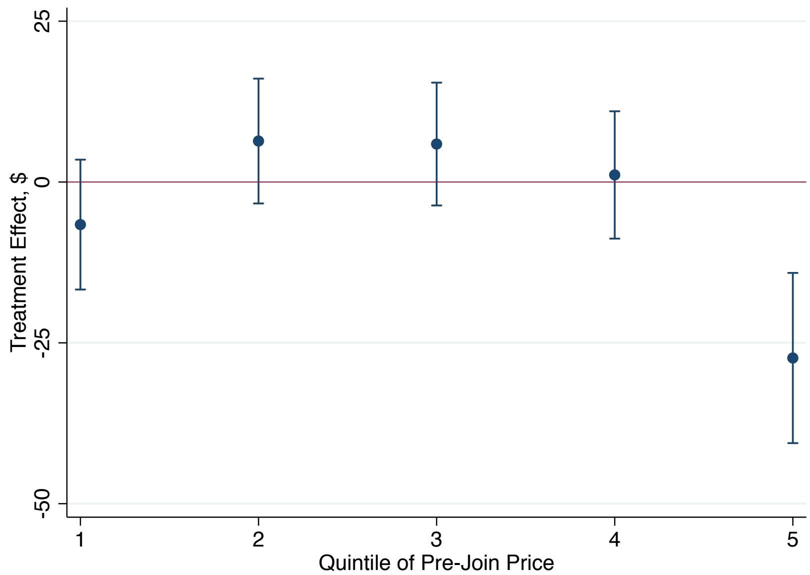

Version	Pre-info price quintiles ($\beta^{Info}_{quintile}$=)				
	1	2	3	4	5
1	4	$17^{\dagger}$	9	-2	$-55^{\dagger}$
	(6)	(6)	(6)	(6)	(9)
2	-1	11	4	-7	$-63^{\dagger}$
	(7)	(7)	(7)	(8)	(10)
3	-7	6	6	1	$-27^{\dagger}$
	(5)	(5)	(5)	(5)	(7)
4	-10	1	2	-3	$-34^{\dagger}$
	(6)	(6)	(6)	(6)	(8)

Authors' calculations from PriceGuide data, 2009-2014. $N = 32,453$ member-product-months. Includes 508 members. Version 1 includes hospital, product, and month fixed effects, plus linear product-specific time trends. Version 2 includes hospital and product-month fixed effects. Version 3 includes hospital-product and month fixed effects, plus linear product-specific time trends. Version 4 includes hospital-product and product-month fixed effects. Figure shows Version 3 results. Standard errors clustered at hospital (Versions 1 and 2) or hospital-product (Versions 3 and 4) level shown in parentheses. Superscript (†) indicates significant difference from zero at the 1% level; (**) at the 5% level; (*) at the 10% level.

Figure 19: Treatment Effect Estimates Throughout the Price Distribution

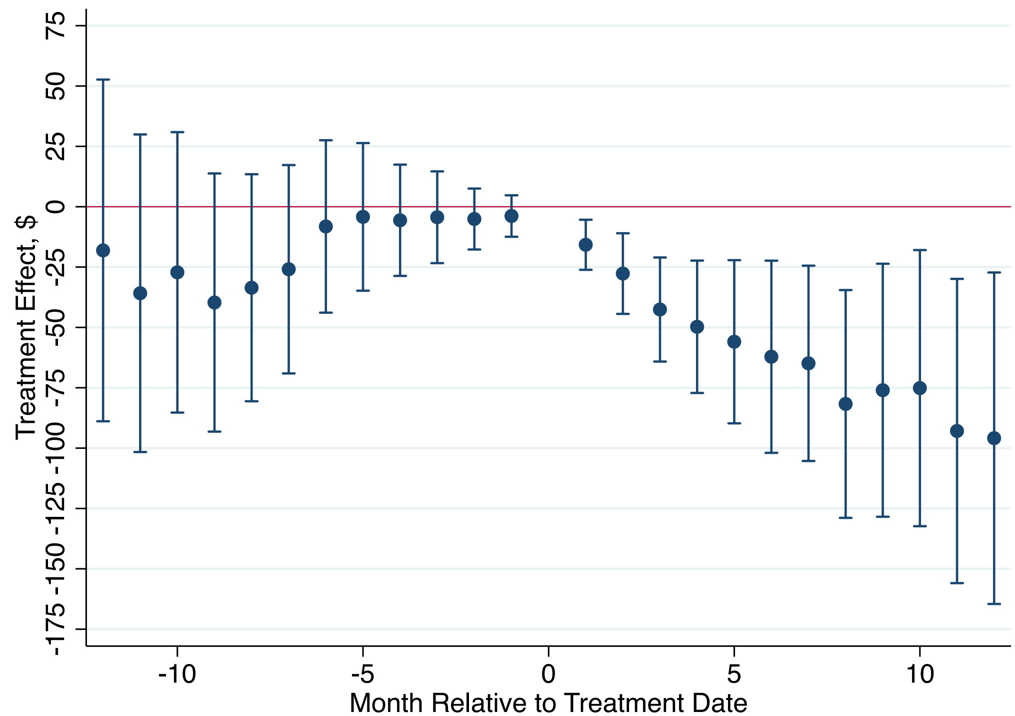

Version	Month relative to join date ($\beta_5^{Join,mo}=$)																							
	-12	-11	-10	-9	-8	-7	-6	-5	-4	-3	-2	-1	1	2	3	4	5	6	7	8	9	10	11	12
1	-24	-33	-20	-24	-14	-14	-21^{*}	$-31^{\dagger}$	$-27^{\dagger}$	$-22^{\dagger}$	-14^{**}	-8	$-23^{\dagger}$	$-40^{\dagger}$	$-52^{\dagger}$	$-61^{\dagger}$	$-64^{\dagger}$	$-74^{\dagger}$	$-77^{\dagger}$	$-88^{\dagger}$	$-83^{\dagger}$	$-78^{\dagger}$	$-104^{\dagger}$	$-104^{\dagger}$
	(20)	(25)	(22)	(19)	(16)	(17)	(12)	(10)	(8)	(7)	(7)	(6)	(7)	(8)	(9)	(9)	(11)	(12)	(10)	(13)	(13)	(11)	(13)	(13)
2	-17	-15	-10	-11	-6	-8	-14	$-29^{\dagger}$	-22^{**}	-15^{*}	-9	-6	$-26^{\dagger}$	$-43^{\dagger}$	$-54^{\dagger}$	$-65^{\dagger}$	$-70^{\dagger}$	$-80^{\dagger}$	$-80^{\dagger}$	$-89^{\dagger}$	$-83^{\dagger}$	$-75^{\dagger}$	$-104^{\dagger}$	$-102^{\dagger}$
	(19)	(24)	(22)	(18)	(15)	(17)	(13)	(10)	(9)	(8)	(7)	(7)	(7)	(9)	(9)	(10)	(11)	(13)	(11)	(13)	(14)	(12)	(13)	(13)
3	-18	-36	-27	-40	-34	-26	-8	-4	-6	-4	-5	-4	$-16^{\dagger}$	$-28^{\dagger}$	$-43^{\dagger}$	$-50^{\dagger}$	$-56^{\dagger}$	$-62^{\dagger}$	$-65^{\dagger}$	$-82^{\dagger}$	$-76^{\dagger}$	-75^{**}	$-93^{\dagger}$	$-96^{\dagger}$
	(36)	(34)	(30)	(27)	(24)	(22)	(18)	(16)	(12)	(10)	(6)	(4)	(5)	(9)	(11)	(14)	(17)	(20)	(21)	(24)	(27)	(29)	(32)	(35)

Authors' calculations from PriceGuide data, 2009-2014. $N = 23,016$ member-product-months. Includes 507 members, twelve months pre- and post-join only. Version 1 includes hospital, product, and month fixed effects, plus linear product-specific time trends. Version 2 includes hospital and product-month fixed effects. Version 3 includes hospital-product and month fixed effects, plus linear product-specific time trends. Figure shows Version 3 results. Standard errors clustered at hospital (Versions 1 and 2) or hospital-product (Version 3) level shown in parentheses. Superscript (†) indicates significant difference from zero at the 1% level; (**) at the 5% level; (*) at the 10% level.

Figure 20: Event Studies of Treatment Effect of Access to Benchmarking Information, Top Quintile of Price Only

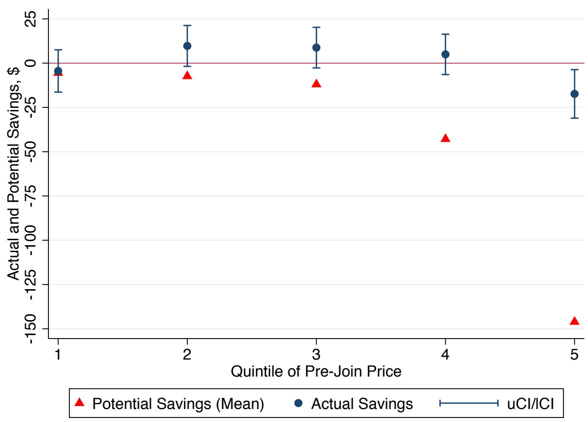

(a) Low Quantity

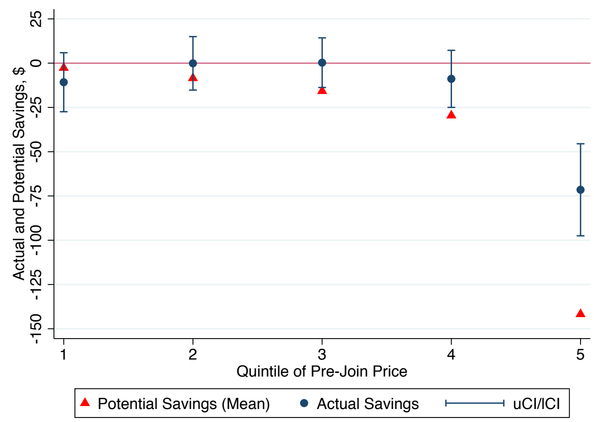

(b) High Quantity

	Achieved savings:									
Version	Pre-info price quintiles ($\beta^{Info}_{quintile,low^q} =$)					Pre-info price quintiles ($\beta^{Info}_{quintile,low^q} + \beta^{Info}_{quintile,high^q} =$)				
	1	2	3	4	5	1	2	3	4	5
1	5	$19^{\dagger}$	16^{**}	1	$-51^{\dagger}$	3	13	−4	−10	$-73^{\dagger}$
	(7)	(6)	(7)	(7)	(9)	(10)	(9)	(9)	(10)	(15)
2	−1	12	11	−4	$-60^{\dagger}$	0	7	−12	−14	$-79^{\dagger}$
	(8)	(8)	(8)	(8)	(11)	(11)	(10)	(10)	(11)	(16)
3	−4	9	9	5	-17^{**}	−11	0	0	−9	$-71^{\dagger}$
	(6)	(6)	(6)	(6)	(7)	(9)	(8)	(7)	(8)	(13)
4	−9	4	6	1	$-23^{\dagger}$	−12	−4	−5	−12	$-78^{\dagger}$
	(7)	(7)	(7)	(7)	(8)	(9)	(8)	(8)	(9)	(13)
	Potential savings:									
	−6	−7	−12	−43	−146	−3	−8	−16	−30	−142

Authors' calculations from PriceGuide data, 2009-2014. $N = 32,453$ member-product-months. Includes 508 members. Version 1 includes hospital, product, and month fixed effects, plus linear product-specific time trends. Version 2 includes hospital and product-month fixed effects. Version 3 includes hospital-product and month fixed effects, plus linear product-specific time trends. Version 4 includes hospital-product and product-month fixed effects. Figure shows Version 3 results. Standard errors clustered at hospital (Versions 1 and 2) or hospital-product (Versions 3 and 4) level shown in parentheses. Superscript (†) indicates significant difference from zero at the 1% level; (**) at the 5% level; (*) at the 10% level. Potential savings calculated using pre-information data only and Version 3 controls.

Figure 21: Treatment Effect Estimates Across the Price and Quantity Distributions

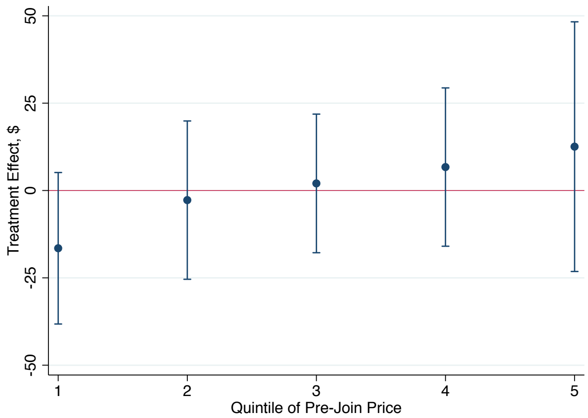

(a) Agency

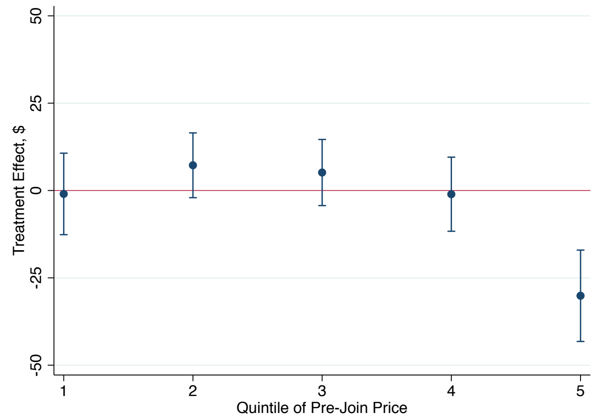

(b) Asymmetric Info

Version	Pre-info price quintiles ($\beta^{Agency}_{quintile} =$)					Pre-info price quintiles ($\beta^{Info}_{quintile} =$)				
	1	2	3	4	5	1	2	3	4	5
1	8	14	-9	-21^{*}	$-60^{\dagger}$	3	12^{**}	$15^{\dagger}$	8	$-33^{\dagger}$
	(12)	(11)	(10)	(12)	(21)	(6)	(5)	(5)	(7)	(7)
2	19	31^{**}	1	-11	-53^{**}	-8	-2	4	-2	$-46^{\dagger}$
	(13)	(12)	(11)	(14)	(22)	(9)	(8)	(8)	(10)	(10)
3	-17	-3	2	7	13	-1	7	5	-1	$-30^{\dagger}$
	(11)	(12)	(10)	(12)	(18)	(6)	(5)	(5)	(5)	(7)
4	-3	16	16	18	25	-15^{*}	-11	-11	-15^{**}	$-47^{\dagger}$
	(12)	(12)	(11)	(12)	(19)	(8)	(7)	(7)	(7)	(9)

Authors' calculations from PriceGuide data, 2009-2014. $N = 32,453$ member-product-months. Includes 508 members. Version 1 includes hospital, product, and month fixed effects, plus linear product-specific time trends. Version 2 includes hospital and product-month fixed effects. Version 3 includes hospital-product and month fixed effects, plus linear product-specific time trends. Version 4 includes hospital-product and product-month fixed effects. Figure shows Version 3 results. Standard errors clustered at hospital (Versions 1 and 2) or hospital-product (Versions 3 and 4) level shown in parentheses. Superscript (†) indicates significant difference from zero at the 1% level; (**) at the 5% level; (*) at the 10% level.

Figure 22: Treatment Effect Estimates Across the Price Distribution, Separating Agency and Asymmetric Information Mechanisms

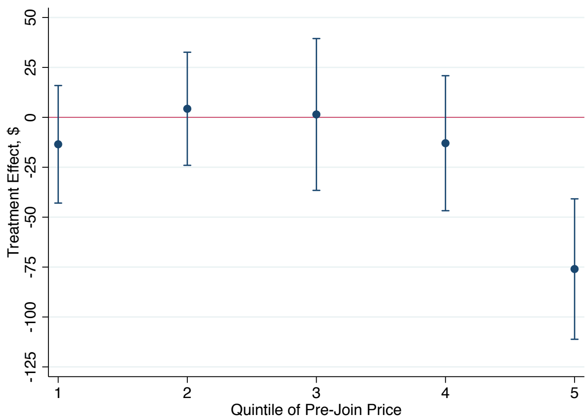

(a) Conditional on Renegotiation

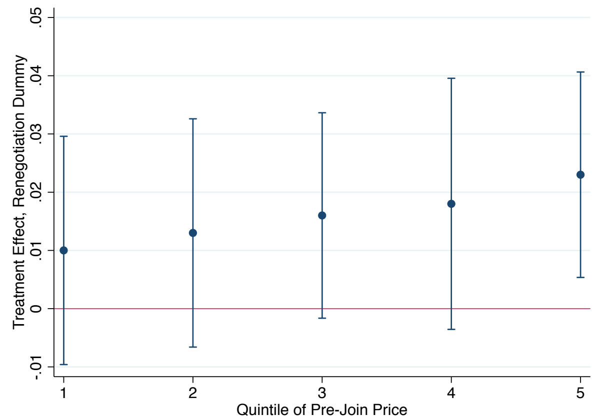

(b) $\mathbb{1}_{\{reneg_{hjt}\}}$

Version	Pre-join price quintiles ($\beta^{Info}_{quintile}$ =)					Pre-join price quintiles ($\beta^{Info}_{quintile}$ =)				
	1	2	3	4	5	1	2	3	4	5
1	−7	13	11	−22	$-91^{\dagger}$	.008	.002	$.018^{**}$	$.018^{*}$	$.034^{\dagger}$
	(12)	(12)	(14)	(14)	(19)	(.008)	(.009)	(.008)	(.01)	(.008)
2	−11	5	1	−25	$-111^{\dagger}$	$.017^{*}$	.01	$.023^{**}$	$.025^{**}$	$.04^{\dagger}$
	(16)	(15)	(17)	(17)	(23)	(.009)	(.01)	(.009)	(.011)	(.009)
3	−14	4	1	−13	$-76^{\dagger}$	.01	.013	$.016^{*}$	.018	$.023^{**}$
	(15)	(14)	(19)	(17)	(18)	(.01)	(.01)	(.009)	(.011)	(.009)
4	−8	4	−3	−11	$-80^{\dagger}$	$.022^{**}$	$.022^{*}$	$.024^{**}$	$.029^{**}$	$.032^{\dagger}$
	(19)	(19)	(22)	(20)	(22)	(.011)	(.012)	(.01)	(.012)	(.011)

Authors' calculations from PriceGuide data, 2009-2014. $N = 6,510$ member-product-months in regressions conditional on renegotiation. $N = 32,453$ member-product-months in renegotiation dummies regression. Includes 508 members. Version 1 includes hospital, product, and month fixed effects, plus linear product-specific time trends. Version 2 includes hospital and product-month fixed effects. Version 3 includes hospital-product and month fixed effects, plus linear product-specific time trends. Version 4 includes hospital-product and product-month fixed effects. Figure shows Version 3 results. Standard errors clustered at hospital (Versions 1 and 2) or hospital-product (Versions 3 and 4) level shown in parentheses. Superscript (†) indicates significant difference from zero at the 1% level; (**) at the 5% level; (*) at the 10% level.

Figure 23: Treatment Effects Conditional on Renegotiation and on Occurrence of Renegotiation

F Additional Specifications

F.1 Average Treatment Effect

The differences in differences estimate of the average treatment effect of access to information is small, statistically and economically, at -$3 (s.e. 3). Here, we also use an event study specification that includes indicators for each month relative to the hospital's "info" date:

$$P_{hjt} = \sum_{mo=-12}^{+12} \beta^{Info,mo} * \mathbb{1}_{\{mo=t-t_{info_{hj}}\}} + \theta_{hj} + \theta_t + \gamma_j * (t - t_{min_j}) + \varepsilon_{hjt}$$

Figure 24 shows results for these estimated differences between treated and untreated prices. The plot shows evidence of a slight decline in prices prior to accessing information, though the pre-trends in price in the six months leading up to the timing of information are essentially zero.[44] After the hospital accesses the benchmarking information, there is a steady downward trend in the coefficients. The downward trend in the post-period may be due to price stickiness – it may take newly-informed hospitals some time to arrive at the bargaining table.

In general, estimates for each relative month effect are insignificant and there is not strong evidence of a trend break. Moreover, estimated patterns are similar across the different specifications of controls, though standard errors are large in the richest specification (Version 3).[45]

When interpreting these results, it is important to note that this is the price effect of simply having *access* to information in the database. It may understate the effect of access to information on stents if, for example, the hospital joins the database because of an interest in benchmarking its orthopedic implant prices and never considers the stent information. It could also underestimate the effect of information on price *negotiation* if there is a delay in price changes due to sticky contracts (which both institutional knowledge and the post-period trend noted above suggest is the case).[46] Finally, the average treatment effect of

[44]It should be noted that there are fewer "pre-info" observations available 6-12 months prior to accessing information because of the presence of entering products and because some hospitals do not submit retrospective data until a few months after joining the database. Accordingly, the earlier relative month effects are less precisely estimated.

[45]It was not possible to estimate the monthly event studies with hospital-product and product-month fixed effects. However, the quarterly event study with hospital-product and product-month fixed effects is essentially identical to the quarterly event study with hospital-product and product-specific linear trends.

[46]We investigate this issue in the main text, but we argue that the treatment effect we estimate – the combined effect of information on a particular price negotiation and the probability that price negotiation occurs – is the more important treatment effect of interest for policy as it estimates an overall value of access to benchmarking information for decreasing the total spend of hospitals on medical inputs over time.

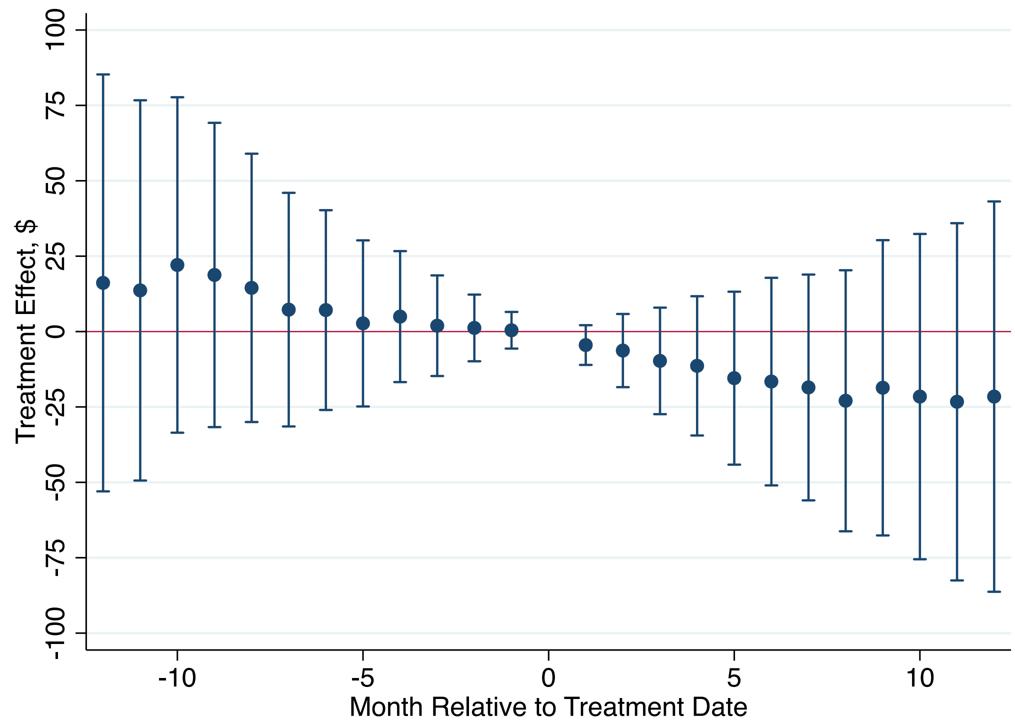

Version	Month relative to info date ($\beta^{Info,mo}=$)																							
	-12	-11	-10	-9	-8	-7	-6	-5	-4	-3	-2	-1	1	2	3	4	5	6	7	8	9	10	11	12
1	17	15	20*	20*	17*	3	−1	−5	3	0	−1	3	−8†	−10†	−14†	−16†	−20†	−21†	−22†	−24†	−22†	−20†	−22†	−21†
	(14)	(14)	(11)	(11)	(9)	(7)	(6)	(5)	(5)	(4)	(3)	(3)	(3)	(3)	(4)	(4)	(5)	(5)	(6)	(6)	(7)	(7)	(7)	(8)
2	22	19	19*	22**	16*	7	5	1	9	8	4	6	−13†	−17†	−18†	−21†	−27†	−27†	−25†	−24†	−24†	−21**	−23†	−18*
	(14)	(14)	(11)	(10)	(9)	(8)	(8)	(7)	(6)	(6)	(5)	(5)	(4)	(4)	(5)	(6)	(7)	(7)	(7)	(8)	(9)	(8)	(9)	(9)
3	16	14	22	19	14	7	7	3	5	2	1	0	−4	−6	−10	−11	−15	−17	−19	−23	−19	−22	−23	−22
	(35)	(32)	(28)	(26)	(23)	(20)	(17)	(14)	(11)	(9)	(6)	(3)	(3)	(6)	(9)	(12)	(15)	(18)	(19)	(22)	(25)	(28)	(30)	(33)

Authors' calculations from PriceGuide data, 2009-2014. $N = 23,016$ member-product-months. Includes 507 members, twelve months pre- and post-join only. Version 1 includes hospital, product, and month fixed effects, plus linear product-specific time trends. Version 2 includes hospital and product-month fixed effects. Version 3 includes hospital-product and month fixed effects, plus linear product-specific time trends. Figure shows Version 3 results. Standard errors clustered at hospital (Versions 1 and 2) or hospital-product (Version 3) level shown in parentheses. Superscript (†) indicates significant difference from zero at the 1% level; (**) at the 5% level; (*) at the 10% level.

Figure 24: Event Studies of Treatment Effect of Access to Benchmarking Information

information on price is pooled across all hospital-products in the database, some of which have substantial opportunities for savings and some of which do not.

F.2 Isolating Identification Based on Timing of Join

We also estimated our regression specifications focusing only on timing of join. To implement this, we limited the regression sample to non-entering products and to hospital-products where the hospital joined the database at least six months after product entry. The event study results for the average treatment effect across all hospitals and products are shown below in Figure 25. The event study results for the top quintile of the price distribution are shown in Figure 26.

As noted above, the date of "information" in many of the results in the main text may be when a hospital joins the database or when it, as an existing member, receives information on a new product. Our discussion of potential bias has primarily focused on the potential existence of contemporaneous factors affecting prices around the timing of join. Thus, it is useful to note that, when we estimate our average treatment effects using only timing of join for identification, we see a nearly identical pattern as in our pooled information treatment results, albeit with larger standard errors due to the smaller sample. Specifically, the results indicate no evidence of a pre-trend in the months just before hospitals join the database, and point estimates drift downward after join but are not statistically significantly different from zero in any month reported.

Figure 26 shows a similar pattern as our information treatment event study, focusing on the top quintile of prices. There is a lack of trend leading up to the join date and a steep and steady decline in prices after join – the point estimate 12 months after join is larger than in Figure 8 (-$154) but not statistically significantly so (recall that the estimation sample is much smaller in the "join-only" regressions).

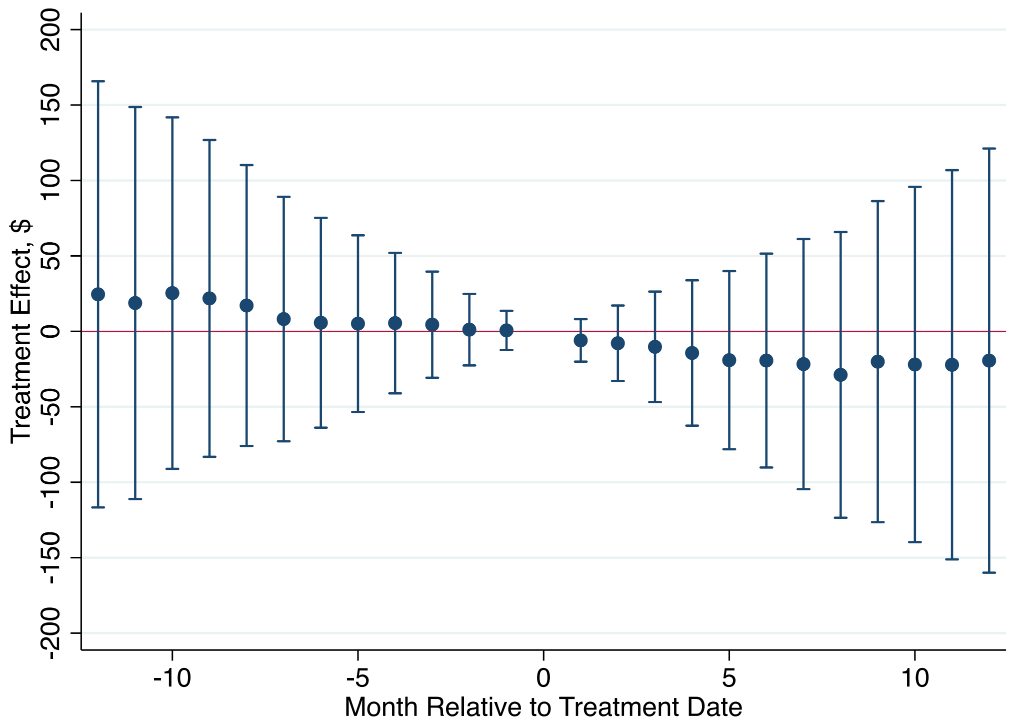

Version	-12	-11	-10	-9	-8	-7	-6	-5	-4	-3	-2	-1	1	2	3	4	5	6	7	8	9	10	11	12
	Month relative to info date ($\beta^{Info,mo} =$)																							
1	44	31	29	27	23	10	2	1	7	1	−2	2	−12	−14	−19	−24	−31	−29	−33	−35	−26	−30	−28	−20
	(78)	(72)	(64)	(58)	(51)	(45)	(38)	(32)	(26)	(20)	(14)	(8)	(8)	(15)	(21)	(27)	(33)	(40)	(47)	(54)	(60)	(67)	(73)	(79)
2	−30	−38	−36	−27	−26	−31	−33	−32	−17	−16	−14	−5	−7	−2	−1	0	−2	6	12	15	26	30	34	50
	(160)	(147)	(132)	(119)	(105)	(92)	(79)	(66)	(53)	(40)	(27)	(14)	(15)	(28)	(41)	(55)	(68)	(81)	(94)	(108)	(121)	(134)	(147)	(160)
3	25	19	25	22	17	8	6	5	5	4	1	1	−6	−8	−10	−14	−19	−19	−22	−29	−20	−22	−22	−19
	(72)	(66)	(59)	(54)	(47)	(41)	(35)	(30)	(24)	(18)	(12)	(7)	(7)	(13)	(19)	(25)	(30)	(36)	(42)	(48)	(54)	(60)	(66)	(72)

Authors' calculations from PriceGuide data, 2009-2014. $N = 9,786$ member-product-months. Includes 327 members, twelve months pre- and post-join only. Version 1 includes hospital, product, and month fixed effects, plus linear product-specific time trends. Version 2 includes hospital and product-month fixed effects. Version 3 includes hospital-product and month fixed effects, plus linear product-specific time trends. Figure shows Version 3 results. Standard errors clustered at hospital (Versions 1 and 2) or hospital-product (Version 3) level shown in parentheses. Superscript (†) indicates significant difference from zero at the 1% level; (**) at the 5% level; (*) at the 10% level.

Figure 25: Event Study of Treatment Effect of Joining Benchmarking Database

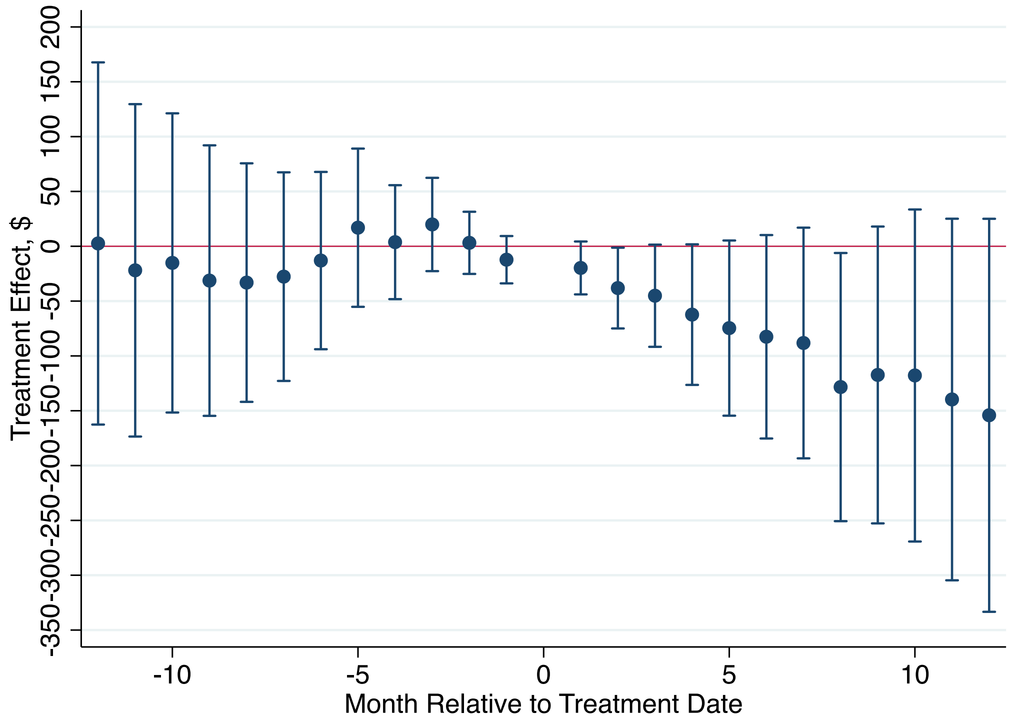

Version	Month relative to join date ($\beta_5^{Join,mo} =$)																						
	-12	-11	-10	-9	-8	-7	-6	-5	-4	-3	-2	-1	1	2	3	4	5	6	7	8	9	10	11	12
1	3	-13	-3	-19	-13	-10	-6	17	17	21	7	-1	-43^{**}	$-61^{\dagger}$	-65^{**}	-88^{**}	-102^{**}	-116^{**}	-123^{**}	-154^{**}	-141^{*}	-153^{*}	-172^{*}	-176^{*}
	(93)	(86)	(78)	(71)	(63)	(55)	(46)	(40)	(30)	(24)	(19)	(17)	(17)	(23)	(28)	(38)	(47)	(54)	(60)	(71)	(77)	(86)	(92)	(100)
2	-94	-94	-79	-81	-72	-65	-52	-33	-11	-1	-10	-13	-35^{*}	-42	-36	-51	-58	-68	-66	-86	-66	-64	-80	-76
	(167)	(154)	(140)	(126)	(112)	(98)	(83)	(69)	(55)	(43)	(30)	(21)	(21)	(32)	(45)	(60)	(74)	(88)	(101)	(116)	(130)	(144)	(157)	(171)
3	3	-22	-15	-31	-33	-28	-13	17	4	20	3	-12	-20	-38^{**}	-45^{*}	-62^{*}	-75^{*}	-83^{*}	-88	-128^{**}	-117^{*}	-118	-140^{*}	-154^{*}
	(84)	(77)	(70)	(63)	(56)	(49)	(41)	(37)	(27)	(22)	(14)	(11)	(12)	(19)	(24)	(33)	(41)	(47)	(54)	(62)	(69)	(77)	(84)	(91)

Authors' calculations from PriceGuide data, 2009-2014. $N = 23,016$ member-product-months. Includes 507 members, twelve months pre- and post-join only. Version 1 includes hospital, product, and month fixed effects, plus linear product-specific time trends. Version 2 includes hospital and product-month fixed effects. Version 3 includes hospital-product and month fixed effects, plus linear product-specific time trends. Figure shows Version 3 results. Standard errors clustered at hospital (Versions 1 and 2) or hospital-product (Version 3) level shown in parentheses. Superscript (†) indicates significant difference from zero at the 1% level; (**) at the 5% level; (*) at the 10% level.

Figure 26: Event Study of Treatment Effect of Joining Benchmarking Database, Top Quintile of Price Only

F.3 Alternative Samples

The following Figures show the results of our richest regression specification, allowing for different treatment effects for different parts of the price and quantity distributions, for specifications that (1) focus only on the twelve months before and after information (Figure 27); (2) focus on identification based only on timing of database join (Figure 28); and (3) limit the sample to those facilities registered with the database as “hospitals” (Figure 29). Results are qualitatively and quantitatively similar across all specifications.

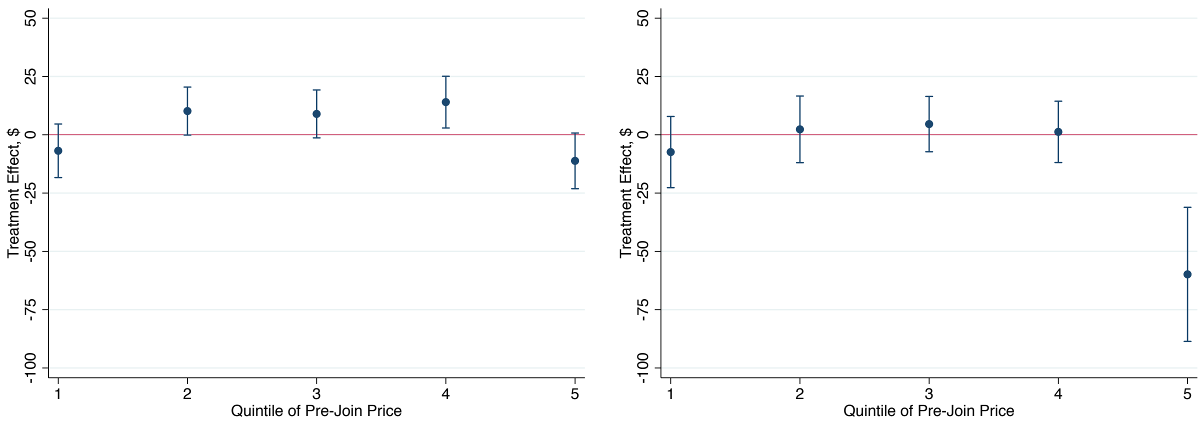

(a) Low Quantity **(b)** High Quantity

Version	Pre-info price quintiles ($\beta^{Info}_{quintile,low^q} =$)					Pre-info price quintiles ($\beta^{Info}_{quintile,low^q} + \beta^{Info}_{quintile,high^q} =$)				
	1	2	3	4	5	1	2	3	4	5
1	-7	15^{**}	14^{**}	8	$-30^{\dagger}$	-11	7	7	1	$-54^{\dagger}$
	(7)	(6)	(6)	(7)	(9)	(9)	(8)	(7)	(8)	(16)
2	-13	9	6	3	$-38^{\dagger}$	-16	0	-2	-3	$-61^{\dagger}$
	(8)	(8)	(7)	(8)	(9)	(10)	(8)	(8)	(9)	(16)
3	-7	10^{*}	9^{*}	14^{**}	-11^{*}	-7	2	5	1	$-60^{\dagger}$
	(6)	(5)	(5)	(6)	(6)	(8)	(7)	(6)	(7)	(15)
4	-11^{*}	7	5	11^{*}	-16^{**}	-10	-1	1	-2	$-66^{\dagger}$
	(7)	(6)	(6)	(6)	(7)	(8)	(8)	(6)	(7)	(15)

Authors' calculations from PriceGuide data, 2009-2014. $N = 23,016$ member-product-months. Includes 507 members. Version 1 includes hospital, product, and month fixed effects, plus linear product-specific time trends. Version 2 includes hospital and product-month fixed effects. Version 3 includes hospital-product and month fixed effects, plus linear product-specific time trends. Version 4 includes hospital-product and product-month fixed effects. Figure shows Version 3 results. Standard errors clustered at hospital (Versions 1 and 2) or hospital-product (Versions 3 and 4) level shown in parentheses. Superscript (†) indicates significant difference from zero at the 1% level; (**) at the 5% level; (*) at the 10% level.

Figure 27: Treatment Effect Estimates Across the Price and Quantity Distributions – Twelve Months Pre/Post

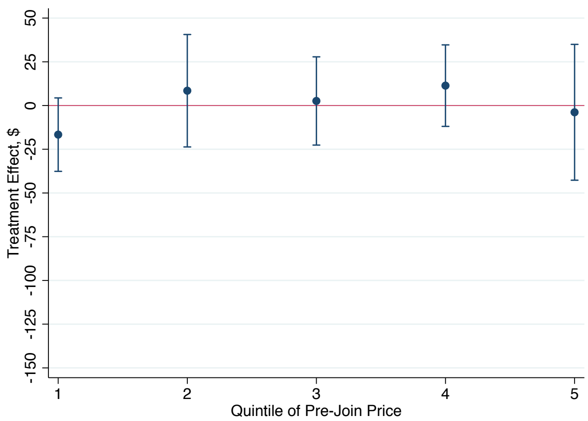

(a) Low Quantity

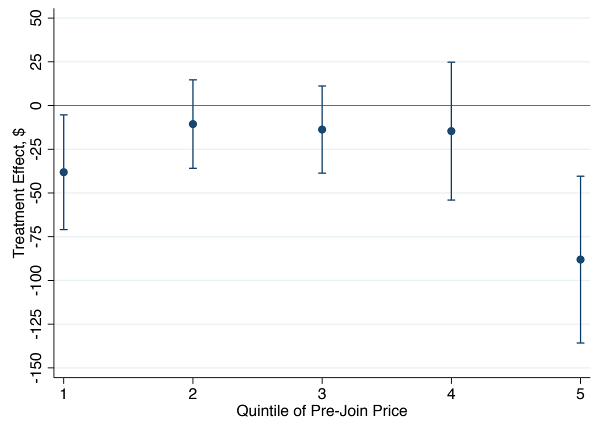

(b) High Quantity

Version	Pre-info price quintiles ($\beta^{Info}_{quintile,low^q} =$)					Pre-info price quintiles ($\beta^{Info}_{quintile,low^q} + \beta^{Info}_{quintile,high^q} =$)				
	1	2	3	4	5	1	2	3	4	5
1	-15	-8	13	-3	$-91^{\dagger}$	-12	-7	$-46^{\dagger}$	-23	$-131^{\dagger}$
	(14)	(15)	(15)	(13)	(23)	(18)	(15)	(18)	(21)	(30)
2	-16	-6	16	-3	$-98^{\dagger}$	-6	-7	$-56^{\dagger}$	-18	$-128^{\dagger}$
	(13)	(17)	(15)	(14)	(22)	(19)	(15)	(18)	(22)	(31)
3	-17	8	3	11	-4	-38^{**}	-11	-14	-15	$-88^{\dagger}$
	(11)	(16)	(13)	(12)	(20)	(17)	(13)	(13)	(20)	(24)
4	-21^{**}	10	7	7	-4	-30^{*}	-12	-20^{*}	-6	$-84^{\dagger}$
	(11)	(17)	(14)	(13)	(19)	(16)	(13)	(12)	(21)	(21)

Authors' calculations from PriceGuide data, 2009-2014. $N = 14,701$ member-product-months. Includes 331 members. Version 1 includes hospital, product, and month fixed effects, plus linear product-specific time trends. Version 2 includes hospital and product-month fixed effects. Version 3 includes hospital-product and month fixed effects, plus linear product-specific time trends. Version 4 includes hospital-product and product-month fixed effects. Figure shows Version 3 results. Standard errors clustered at hospital (Versions 1 and 2) or hospital-product (Versions 3 and 4) level shown in parentheses. Superscript (†) indicates significant difference from zero at the 1% level; (**) at the 5% level; (*) at the 10% level.

Figure 28: Treatment Effect Estimates Across the Price and Quantity Distributions – "Join" Variation Only

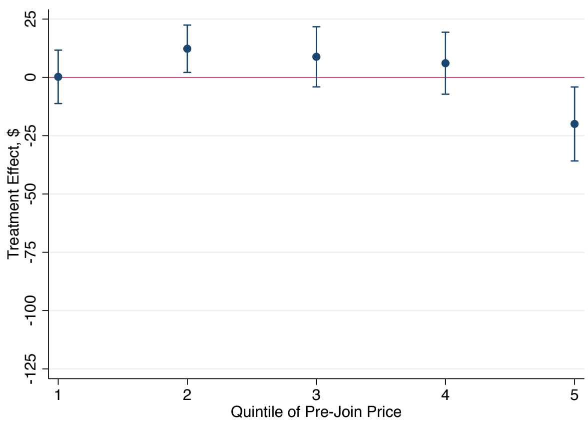

(a) Low Quantity

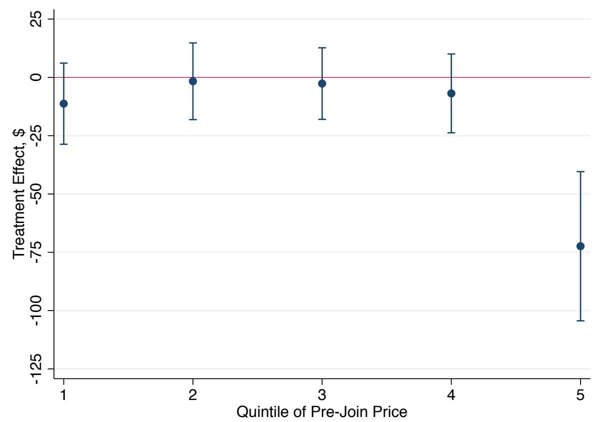

(b) High Quantity

Version	Pre-info price quintiles ($\beta^{Info}_{quintile,low^q} =$)					Pre-info price quintiles ($\beta^{Info}_{quintile,low^q} + \beta^{Info}_{quintile,high^q} =$)				
	1	2	3	4	5	1	2	3	4	5
1	8	$21^{\dagger}$	15^{**}	-2	$-53^{\dagger}$	-1	8	-10	-7	$-80^{\dagger}$
	(7)	(6)	(7)	(8)	(11)	(10)	(10)	(10)	(12)	(18)
2	2	13	11	-5	$-64^{\dagger}$	-5	2	-18	-13	$-84^{\dagger}$
	(8)	(8)	(8)	(9)	(12)	(11)	(10)	(11)	(12)	(18)
3	0	12^{**}	9	6	-20^{**}	-11	-2	-3	-7	$-72^{\dagger}$
	(6)	(5)	(7)	(7)	(8)	(9)	(8)	(8)	(9)	(16)
4	-3	7	7	2	$-26^{\dagger}$	-12	-5	-7	-10	$-77^{\dagger}$
	(7)	(6)	(8)	(7)	(9)	(9)	(9)	(9)	(9)	(16)

Authors' calculations from PriceGuide data, 2009-2014. $N = 27,698$ member-product-months. Includes 436 members. Version 1 includes hospital, product, and month fixed effects, plus linear product-specific time trends. Version 2 includes hospital and product-month fixed effects. Version 3 includes hospital-product and month fixed effects, plus linear product-specific time trends. Version 4 includes hospital-product and product-month fixed effects. Figure shows Version 3 results. Standard errors clustered at hospital (Versions 1 and 2) or hospital-product (Versions 3 and 4) level shown in parentheses. Superscript (†) indicates significant difference from zero at the 1% level; (**) at the 5% level; (*) at the 10% level.

Figure 29: Treatment Effect Estimates Across the Price and Quantity Distributions – Hospitals Only

F.4 Bare Metal Stents

In the following, we display select summary statistics and regression results for bare metal stent transactions. As with the drug-eluting stents that are the focus of this paper, there is substantial dispersion in price outcomes after conditioning on hospital size or volume. The left panel of Figure 30 shows a box plot of bare metal stent prices for each category of bed count. The right panel of Figure 30 shows box plots of stent prices for each decile of monthly stent purchasing volume. In the latter, we see a slight relationship between "size" and price – 10^{th} decile hospitals' prices are 11% lower than those obtained by 1^{st} decile hospitals. However, there is a great deal of unexplained hospital price heterogeneity conditional on purchasing volume. Further evidence of this heterogeneity is shown in the left panel of Figure 31 – there is substantial dispersion in both hospital and hospital-product fixed effects after controlling for product-specific time trends.

Figure 30: Distribution of Prices Across Hospitals

(a) By Bed Count

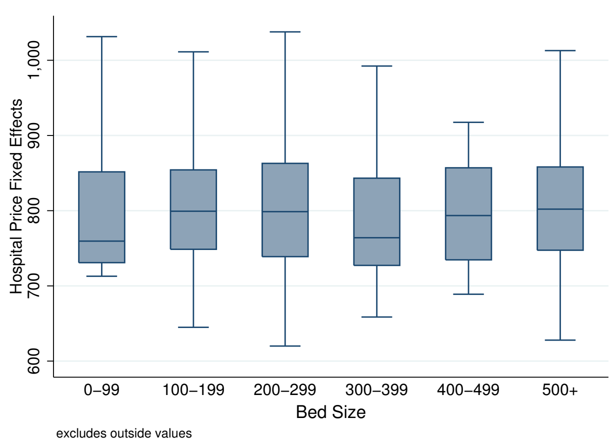

(b) By Stent Volume Decile

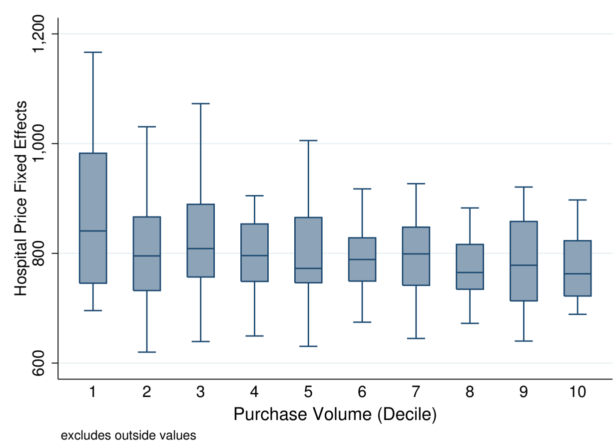

Bare Metal Stent Prices by Size Category (Regression Results)

$\beta^{\mathbb{1}\{h\in \text{ bed size } x\}} =$						$\beta^{\mathbb{1}\{h\in \text{ decile } x\}} =$									
0 – 99	100 – 199	200 – 299	300 – 399	400 – 499	500+	1	2	3	4	5	6	7	8	9	10
835	814	813	784	793	805	868	813	846	800	802	790	802	794	786	776
(36)	(18)	(14)	(17)	(17)	(12)	(29)	(23)	(35)	(19)	(19)	(17)	(17)	(16)	(16)	(15)

Authors' calculations from PriceGuide data. Estimated mean hospital fixed effects within bed size categories and decile of monthly purchase volume. Hospital fixed effects obtained from regression of price on hospital and product-month fixed effects, pre-join data only. Mean estimates from regression of fixed effects on indicators for size. Standard errors from nonparametric bootstrap of entire procedure, resampling at hospital level.

The right panel of Figure 31 shows reduced form evidence of our overall results: after accessing benchmarking information, hospitals are somewhat less likely to pay the highest prices. Figure 32 shows regression results demonstrating the same phenomenon: after accessing benchmarking data, hospitals previously paying the highest prices experience price decreases, and these decreases are larger among hospitals with greater quantities at stake. In both the reduced form and regression evidence, the results are less pronounced than they were for drug-eluting stents. Given that bare metal stents are substantially less popular than drug-eluting stents, this may be yet another illustration that information effects are mediated by the total volume of purchase at stake. Furthermore, bare metal stents are an older technology and are declining in popularity over our period of study, with sample market share of 22% in 2009 and only 14% in 2014. This implies that the *future* volume at stake for bare metal stents is relatively small, in addition to the *current* volume at stake being relatively small.

Figure 31: Histograms of Price Distributions: Hospital and Hospital-Product Variation and Pre- and Post-Information

(a) h and hj Price Variation

(b) pre and post Price Variation

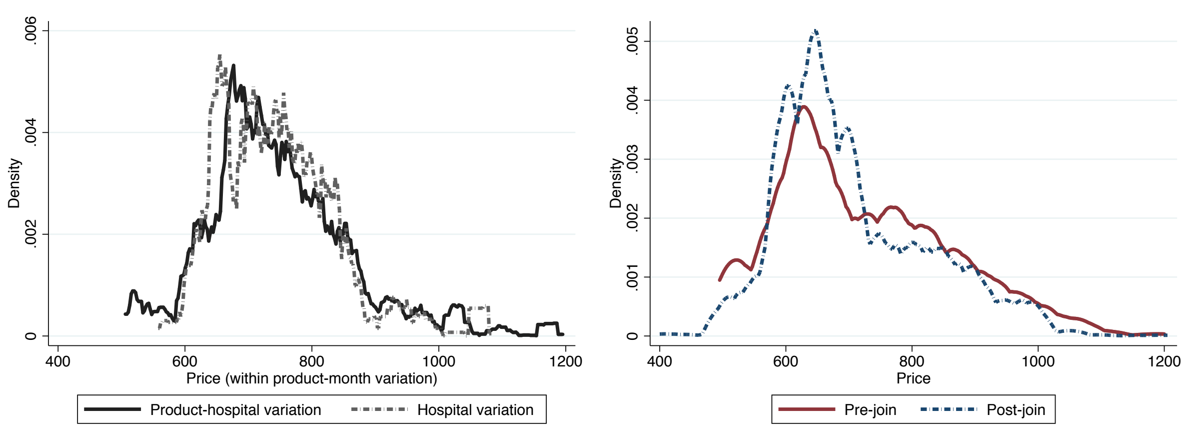

Authors' calculations from PriceGuide data, 2009-2014.

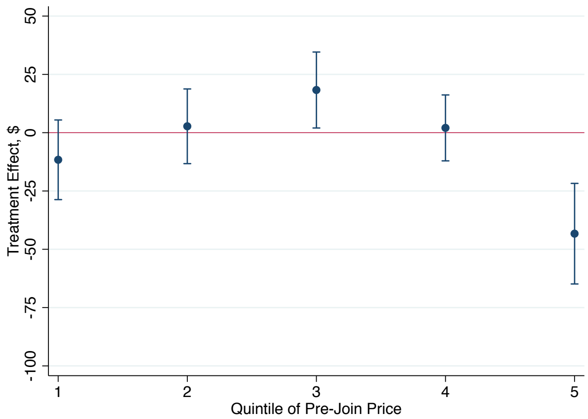

(a) Low Quantity

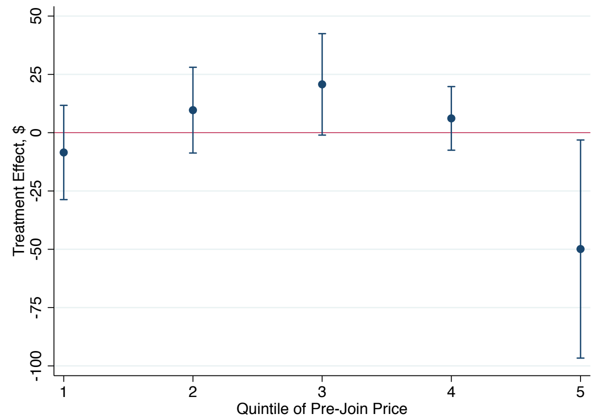

(b) High Quantity

Version	Pre-info price quintiles ($\beta^{Info}_{quintile,low^q} =$)					Pre-info price quintiles ($\beta^{Info}_{quintile,low^q} + \beta^{Info}_{quintile,high^q} =$)				
	1	2	3	4	5	1	2	3	4	5
1	11	10	13	−3	−59†	15	19*	13	−14	−81†
	(8)	(8)	(8)	(9)	(13)	(12)	(10)	(12)	(10)	(17)
2	12	15*	13	−3	−57†	14	23**	15	−14	−83†
	(8)	(8)	(9)	(10)	(13)	(12)	(10)	(12)	(10)	(16)
3	−12	3	18**	2	−43†	−8	10	21*	6	−50**
	(9)	(8)	(8)	(7)	(11)	(10)	(9)	(11)	(7)	(24)
4	−10	8	19**	3	−41†	−7	13	21*	6	−50**
	(9)	(8)	(9)	(8)	(11)	(11)	(10)	(11)	(7)	(23)

Authors' calculations from PriceGuide data, 2009-2014. $N = 19,106$ member-product-months. Includes 410 members. Version 1 includes hospital, product, and month fixed effects, plus linear product-specific time trends. Version 2 includes hospital and product-month fixed effects. Version 3 includes hospital-product and month fixed effects, plus linear product-specific time trends. Version 4 includes hospital-product and product-month fixed effects. Figure shows Version 3 results. Standard errors clustered at hospital (Versions 1 and 2) or hospital-product (Versions 3 and 4) level shown in parentheses. Superscript (†) indicates significant difference from zero at the 1% level; (**) at the 5% level; (*) at the 10% level.

Figure 32: Treatment Effect Estimates Across the Price and Quantity Distributions – Bare Metal Stents

References

Agicha, V., Denecker, O., Dias, Y., and Sarvady, G. (2010). The mckinsey global payments map: A global eye on local opportunities. Technical report, McKinsey and Company.

Albek, S., Mollgaard, P., and Overgaard, P. (1996). Law-assisted collusion? the transparency principle in the danish competition act. *European Competition Law Review*.

Ausubel, L., Cramton, P., and Deneckere, R. (2002). Bargaining with incomplete information. In Aumann, R. and Hart, S., editors, *Handbook of Game Theory*, volume 3, chapter 50, pages 1897–1945. Amsterdam: Elsevier Science.

Backus, M., Blake, T., and Tadelis, S. (2015). Cheap talk, round numbers, and the economics of negotiation. NBER Working Paper No. 21285.

Bandiera, O., Prat, A., and Valletti, T. (2009). Active and passive waste in government spending: Evidence from a policy experiment. *American Economic Review*, 99(4):1278–1308.

Binmore, K., Rubinstein, A., and Wolinsky, A. (1986). The nash bargaining solution in economic modelling. *The RAND Journal of Economics*, 17(2):176–188.

Bronnenberg, B., Dube, J.-P., Gentzkow, M., and Shapiro, J. (2015). Do pharmacists buy bayer? informed shoppers and the brand premium. *Quarterly Journal of Economics*, 130(4):1669–1726.

Calsyn, M. and Emanuel, E. J. (2014). Controlling costs by expanding the medicare acute care episode demonstration. *JAMA Intern Med*, 174(9):1438–1439.

Card, D., Mas, A., Moretti, E., and Saez, E. (2012). Inequality at work: the effect of peer salaries on job satisfaction. *American Economic Review*, 102(6):2981–3003.

Collard-Wexler, A., Gowrisankaran, G., and Lee, R. S. (2014). An alternating offers representation of the nash-in-nash bargaining solution. *working paper*.

Cooper, Z., Craig, S., Gaynor, M., and Van Reenen, J. (2015). The price ain't right? hospital prices and health spending on the privately insured. NBER Working Paper 21815.

Crawford, G. and Yurukoglu, A. (2012). The welfare effects of bundling in multichannel television. *American Economic Review*, 102(2).

Cutler, D. and Dafny, L. (2011). Designing transparency systems for medical care prices. *The New England Journal of Medicine*, 364:894–895.

Dranove, D. and Jin, G. (2010). Quality disclosure and certification: Theory and practice. *Journal of Economic Literature*, 48(4):935–63.

Duggan, M. and Morton, F. S. (2006). The distortionary effects of government procurement: Evidence from medicaid prescription drug purchasing. *Quarterly Journal of Economics*, 121:1–31.

Ellison, G. and Ellison, S. (2009). Search, obfuscation, and price elasticities on the internet. *Econometrica*, 77(2):427–452.

Gowrisankaran, G., Nevo, A., and Town, R. J. (2015). Mergers when prices are negotiated: Evidence from the hospital industry. *American Economic Review*, 105(1):172–203.

Grennan, M. (2013). Price discrimination and bargaining: Empirical evidence from medical devices. *American Economic Review*, 103(1):145–177.

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

Hastings, J., Neilson, C. A., and Zimmerman, S. D. (2015). The effects of earnings disclosure on college enrollment decisions. NBER working paper 21300.

Hendricks, K., Sorensen, A., and Wiseman, T. (2012). Observational learning and demand for search goods. *American Economic Journal: Microeconomics*, 4(1):1–31.

Ho, K. and Lee, R. S. (2014). Insurer competition and negotiated hospital prices. NBER working paper 19401.

Holmstrom, B. (1982). Moral hazard in teams. *The Bell Journal of Economics*, 13(2):324–340.

Horn, H. and Wolinsky, A. (1988). Bilateral monopolies and incentives for merger. *RAND Journal of Economics*, 19:408–419.

Jin, G. and Leslie, P. (2003). The effect of information of product quality: Evidence from restaurant hygiene grade cards. *Quarterly Journal of Economics*, 118(2):409–451.

Kaplan, G. and Menzio, G. (2015). The morphology of price dispersion. *International Economic Review*, 56(4).

Larsen, B. (2015). The efficiency of real-world bargaining: Evidence from wholesale used-auto auctions. *working paper*.

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

Lieber, E. M. J. (2015). Does it pay to know the prices in healthcare? working paper.

Maeda, J. L. K., Raetzman, S. O., and Friedman, B. S. (2012). What hospital inpatient services contributed the most to the 2001-2006 growth in the cos per case? *Health Services Research*, 478:1814–1835.

McConnell, K. J., Lindrooth, R. C., Wholey, D. R., Maddox, T. M., and Bloom, N. (2013). Management practices and the quality of care in cardiac units. *JAMA Internal Medicine*, 173(8):684–692.

Pauly, M. V. and Burns, L. R. (2008). Price transparency for medical devices. *Health Affairs*, 27:1544–1553.

Rubinstein, A. (1982). Perfect equilibrium in a bargaining model. *Econometrica*, 50(1):97–110.

Rubinstein, A. (1985). A bargaining model with incomplete information about time preferences. *Econometrica*, 53(5):1151–1172.

Sarvestani, A. (2014). No more secrecy in medtech pricing, says sen. king. http://www.massdevice.com/no-more-secrecy-medtech-pricing-says-sen-king/.

Schneller, E. S. (2009). The value of group purchasing - 2009: Meeting the need for strategic savings. *Health Care Sector Advances, Inc.*

Scholten, P. A. and Smith, S. (2002). Price dispersion then and now: Evidence from retail and e-tail markets. *The Economics of the Internet and E-commerce*, 11:63–88.

Scott Morton, F., Silva-Rosso, J., and Zettelmeyer, F. (2011). What matters in a price negotiation: Evidence from the u.s. auto retailing industry. *Quantitative Marketing and Economics*, 9(4):365–402.

Sorenson, A. (2000). Equilibrium price dispersion in retail markets for prescription drugs. *Journal of Political Economy*, 108(4).

Waldman, P., Armstrong, D., and Freedberg, S. P. (2013). Deaths linked to cardiac stents rise as overuse seen. *Bloomberg Business*.

Zettelmeyer, F., Scott Morton, F., and Silva-Risso, J. (2006). How the internet lowers prices: Evidence from matched survey and automobile transaction data. *Journal of Marketing Research*, XLIII:168–181.