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

December 31, 2015

Abstract

We empirically examine the role of information in bargaining between hospitals and suppliers. Using a new data set including all purchase orders issued by over ten percent of US hospitals 2009-14, and differences-in-differences identification strategies 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. These reductions are concentrated among hospitals previously paying relatively high prices relative to other hospitals and for products purchased in relatively large volumes, and appear to result from solving asymmetric information problems between hospitals and their suppliers. The results have implications for the emerging role of "information intermediaries" in business-to-business bargaining and calls for transparency in medical device pricing specifically.

*This paper benefited from comments from Robin Lee, Ali Yurukoglu, and audiences at ASSA, 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. Biruk Bekele, Chetna Johry, and Jake Lahr provided excellent research assistance. Any errors are our own.

1 Introduction

Business-to-business markets often lack transparency in the sense that suppliers negotiate different contracts with different buyers, and a buyer typically has limited information regarding other buyers' contracts. As technology has made data easier to collect, distribute, and analyze, many such markets have seen the entry of information intermediaries who facilitate buyers' ability to benchmark the prices they negotiate.[1] Prior research in consumer goods markets (Sorenson 2000; Jin and Leslie 2003; Zettelmeyer et al 2006; Scott Morton et al 2011; Bronnenberg et al 2014) has largely confirmed the economic intuition that information facilitates search and decision making for buyers with imperfect information regarding product quality or costs. However, the implications of this type of increased transparency are not obvious – theoretically or empirically – in a market where both buyers and suppliers have market power and prices are negotiated. In these business-to-business markets, price variation across buyers for the same product need not be due to information frictions (Crawford and Yurukoglu 2012; Grennan 2013, 2014; Gowrisankaran et al 2015; Ho and Lee 2015; Lewis and Pflum 2015), there is often no search across sellers in that a product is only available directly from its manufacturer, and negotiators on both sides are professionals employed by firms and thus with different expertise and incentives than the typical consumer. In this paper, we use a new data set on all purchase orders issued by more than ten percent of US hospitals between 2009-14 to estimate the impact of access to benchmarking information on the prices hospitals negotiate with their suppliers.

We have two primary goals: (1) to estimate the treatment effect of transparency in negotiated prices, where transparency takes the form of benchmarking information on hospital supply prices; and (2) to inform theory development on the role of this type of transparency in negotiated price markets. Hospital supplies and devices are a particularly important case for this analysis as 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 hospital-supplier contracting may hold great potential for reducing health care system cost growth.[2] Indeed, across a broad set of product categories, there is substantial variation in prices across hospitals – for the top fifty hospital supplies by expenditure in our data, the average standard deviation of prices across hospitals for the same exact product and month is ten percent of the mean price. Recent legislation has proposed that the variation in prices across hospitals is at least in part due to a lack of transparency in these input markets, and further that increasing transparency would lower average prices.[3] The policy attention given to these prices reflects both a concern for the financial viability of hospitals and also a

[1]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.

[2]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.

[3]For example, Senator Angus King of Maine recently added an amendment to a tax bill that would increase price transparency for medical devices, stating that "To the extent that prices of implantable medical devices . . . are not disclosed, the ability of hospitals to bring price information to bear in negotiations and decisions is clearly limited." ("King Calls," 2014)

concern that rising supply costs over time filter downstream into higher costs for the health system and consumers.

For many of the most important product categories in medical technology, individual hospitals typically negotiate directly with the product's manufacturer. Hence, any impact of information on the prices other buyers are paying for a product must enter through this negotiation (in contrast with the more well-studied case of price-taking consumers shopping among multiple retailers offering different prices for the same item). Based on the policy and economics literature on this setting (see, e.g., Pauly and Burns, 2008), as well as on conversations with market participants, the most promising candidate mechanisms through which benchmarking information might have an impact in this context are: (1) a model in which hospitals face uncertainty about suppliers' costs or bargaining parameters, so that price transparency reduces the degree of 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. In order to investigate the mechanisms that underly any price effects attributable to benchmarking, we relate the negotiation procedure in this setting to Rubinstein's (1985) model of bargaining with incomplete information and Holmstrom's (1982) model of moral hazard in teams. We empirically test the predictions of each model and provide evidence on the underlying mechanisms. We leave for future research the task of developing a theory that simultaneously incorporates both mechanisms.

Our analysis is based on data from a large hospital supply benchmarking service, covering all purchase orders issued by ten percent of US hospitals between 2009 and 2014. In order to control for a host of differences across product categories, we focus our analysis on price negotiations for coronary stents and thus limit our sample to 508 facilities with cardiac catheterization services. Stents are a desirable category because they are important (one of the largest categories, comprising two percent of hospital supply spend in our sample and about $2 billion annually in the US overall) and typically have simple linear contracts (so the price observed on the purchase order is the price paid). Stents are also physician preference items where doctor usage decisions are insensitive to price, making negotiating lower prices the main mechanism via which a hospital can obtain savings. Given the observed price dispersion for stents alone, potential savings are substantial – if all hospitals paid the minimum price paid by any hospital (for a given stent in a given month), they would save 18 percent on average. The database is generated by monthly submissions from the member hospitals on prices and quantities of each item purchased, at the manufacturer stock-keeping-unit (SKU) level. Importantly, new member hospitals joining the database are asked to submit 12 months of retrospective data, so for any hospital joining during our sample period (about one third of the hospitals in the data) we observe data in pre- and post-information states.

Because different hospitals join at different times, we can construct differences-in-differences estimators based on the prices negotiated by hospitals with and without access to the benchmarking information, controlling for time-invariant differences at the hospital-product level and product-specific 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. The exogeneity of join timing is supported by the qualitative fact that stents are just one of many inputs a hospital purchases (and also one frequently purchased via the catheter lab business unit as opposed to central purchasing) and by quantitative evidence from event studies that show no statistically significant divergence of pre-trends. We provide further support for this assumption by developing a further set of tests focusing on *new* products entering the market during our sample period.

New product introductions provide useful variation for identification along several dimensions. First, new product introduction timing provides even more plausibly exogenous timing, removing any sources of bias at the timing of join that are transient and not persistent over time. Second, and perhaps more importantly, because no information on others' prices is available when a new product first enters the market, comparing prices between hospitals pre- and post-join immediately upon a product's introduction offers a difference between these hospitals that sweeps out any persistent sources of bias of join timing. Third and finally, new product introductions offer a strategy to separate our two theoretical mechanisms of interest: As we argue in our theoretical predictions in Section 3, the asymmetric information mechanism where hospitals use benchmarking information to learn about suppliers relies upon concurrent availability of data on others' prices, but the agency mechanism where hospitals use benchmarking information to create better contracts for their purchasing negotiators relies 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 / contracting effect (plus any persistent bias associated with initial timing of join) and (2) an asymmetric information / learning about supplier type effect.

The estimated average treatment effect across product-hospital-months for coronary stents suggests that simply having access to the information in the database results in small price reductions. This average estimate, however, conceals substantial heterogeneity. Hospital-products whose prices are above the 80^{th} percentile experience price declines of -\$30 per stent upon accessing database information (to give this context, the average size sample hospital uses 700 stents annually and the average stent price is just over \$1,300). The price declines are larger for product-hospital combinations with larger purchase volumes at stake – 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 only -\$20 for hospital-products with lower purchase volumes.

The heterogeneity in results is consistent with the predictions of both models of bargaining under asymmetric information and models in which there is an agency problem in incentivizing effort toward negotiation. Treatment effects are concentrated among hospitals who are least successful in negotiating absent transparency and who therefore learn most when benchmarking data are made available. Further, effects are larger when higher quantities are at stake, consistent with a model in which searching for and utilizing benchmarking data is costly. Importantly, we estimate specifications that estimate separate parameters (by allowing a different parameter for entering products) for price effects attributable to asymmetric information and agency, and demonstrate that asymmetric information consistently explains a substantial portion of the effect (while agency effects are noisier and not statistically different from zero in all specifications).

Finally, in order for benchmarking to have an effect in this setting, the hospital must engage the supplier to negotiate a new contract (the term of the existing contract may not expire for a year or more). Thus, our estimates of the treatment effect of benchmarking on prices will be an underestimate of the treatment effect of benchmarking on prices negotiated in a given contract. For this reason, we also estimate treatment effects of benchmarking on the likelihood of renegotiation and on prices conditioning on months in which we observe renegotiation taking place. These analyses demonstrate that 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 in the form of benchmarking are limited somewhat by stickiness of contracts.

1.1 Related Literature, Public Policy, and Roadmap

This paper relates to literatures on bargaining and on the role of informed buyers on market outcomes. For the latter, much of the prior literature has measured how information affects search and outcomes such as price (Sorenson 2000) and quality (Jin and Leslie 2003; Bronnenberg et al. 2014) in markets where buyers are price-taking consumers, generally finding that effects of information are on average null or beneficial to buyers. Unique among the consumer information literature, Zettelmeyer et al (2006) and Scott-Morton et al (2011) examine how information from information from website research affects the prices consumers negotiate for car purchases. Their studies are quite comprehensive in that they contain data on a variety of consumer characteristics, search, bargaining preferences, and information. Our paper extends this literature to business-to-business bargaining, shutting down the mechanism of search across retailers for the same product, and focusing on the mechanisms via which information affects the price the buying firm is able to negotiate with the same supplying firm.

An emerging empirical bargaining literature (Crawford and Yurukoglu 2012; Grennan 2013, 2014; Gowrisankaran, Nevo, and Town 2014; Ho and Lee 2015; Lewis and Pflum 2015), has thus far modeled business-to-business negotiations of perfect information with exogenously given bargaining parameters.[4] Our tests of the effect of information and the mechanisms of asymmetric information and negotiator agency provide tests of both of these assumptions in our

[4]Larsen (2015) is distinct in estimating a bargaining game of two-sided incomplete information about valuations in the used car wholesale market. Our theoretical motivation differs in that we model uncertainty over bargaining parameters which excludes negotiation breakdown, a desirable feature in ours and other business-to-business contexts.

context. Our finding that information matters suggest that information may be another source of heterogeneity in the bargaining parameters being estimated in those studies, and at the least suggest the information of buyers and sellers – and potential changes to that information – should be thought about carefully when performing empirical estimation and policy analysis.

Finally, our estimates provide a first step towards thinking about the transparency policies that have been proposed for medical technology markets. While in our study the transparency provided by the benchmarking service leads to a decrease in the top part of the price distribution for the most used products, it does not eliminate the price variation across buyers that has concerned policy-makers. Further, a full analysis of transparency on a nationwide scale would take into account supply side responses to transparency, which can negate or overturn welfare-positive demand-side effects via greater obfuscation (Ellison and Ellison 2009), facilitating collusion (Albek et al. 1997), or forcing coordination not to price discriminate via secret discounts (Grennan 2013). Our research design and the variation in the data will not allow us to estimate the first two. However, to the extent that suppliers know when buyers join our benchmarking database (and anecdotal evidence suggests that they do), then our estimates will incorporate the net effects of both informed buyers and also the potential reluctance of suppliers to cut any individual buyer a deal when that information will become part of other buyers' future information set.

The paper proceeds by first examining the data, setting, and research design 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 presents our difference-in-differences results on the average treatment effects and also heterogeneous treatment effects at different points in the price and quantity distributions designed to better understand the mechanisms behind the theoretical predictions. Section 5 concludes.

2 Data, Setting, and Research Design

2.1 Hospital Purchase Order Data

In order to investigate this research question, we have transaction data on all supply purchases made by about 10% of US hospitals during the period 2009-2014. 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. There are 1.9m distinct products in almost 3,000 product categories in the data, which are reported monthly. For each transaction, we observe price, quantity (with relevant units), expenditure, transaction date, product (manufacturer SKU and Universal Medical Devices Nomenclature System (UMDNS) code),[5] and supplier. We observe unique (but anonymous) identifiers for each hospital and the data include several coarse hospital characteristics: census region, facility type, and number of beds.

Table 1 displays some summary statistics regarding the transactions data. We observe

[5]UMDNS is a standard international coding system for medical devices developed by the ECRI Institute.

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 of which is dedicated to coronary stents. As expected, hospitals and health systems generate the majority of the spending on stents – 60% of hospitals and health systems purchased stents during 2009-2014, vs. 30% across all members.

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
Total Spend/Month on Stents (\$k)	22.0	55.4	44.9	72.9	80.4	73.8

Figure 1: Distribution of Benchmarking Database vs. Medicare Hospitals

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

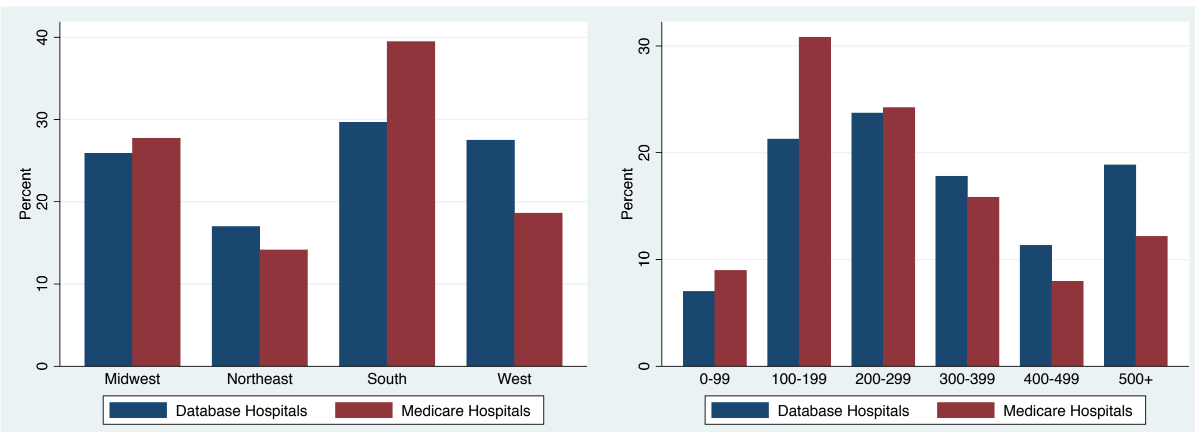

The sample hospitals in the purchase order data voluntarily joined a subscription service that allows them to benchmark purchasing by comparing their own prices and quantities to those of other hospitals in the database. The sample of hospitals joining the database is likely nonrandom. In particular, subscription is costly, so we expect hospitals with greater concerns about supply costs to be overrepresented in the database. The left panel of Figure 1 compares the distribution of sample hospitals across US census regions to that of Medicarecertified hospitals with cardiac catheterization labs.[6] The Figure shows that the west region is overrepresented in the sample data, 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 < 300 beds range and disproportionately more hospitals in the ≥ 300 beds range, relative to Medicare hospitals that would purchase stents. We also compared 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 goal of the MRG survey is to provide an accurate picture of market shares and prices by US region (Northeast, Midwest, South, West). The member facilities in our estimation sample purchase in significantly higher volumes (59 vs. 33 stents per month) and obtain significantly lower prices ($1,530 vs. $1,806 per drug-eluting stent). 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' ability to purchase custom benchmarking services from consulting firms. All of our parameter estimates are internally valid in that they are come only from the sample of hospitals who join the benchmarking database, exploiting the existence of pre/post data and staggered join dates. In the final section of the paper, we return to the issue of representativeness and the external validity of our results, using the MRG and Medicare samples to extrapolate to the population of US hospitals.

2.2 Coronary Stents

As noted above, we focus on coronary stents in our empirical analysis. 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.

Stents are an important product category, both in terms of overall sales and also as a percentage of hospital supply costs. In the US, hospitals spend more than two billion dollars annually on stents used in over 700,000 procedures[7]; in our transactions data, stents comprised two percent of overall supply costs among all members. Table 2 summarizes the stent transactions data for the restricted sample. The average sample hospital submitted stent transactions in 41 months. In a given month, sample hospitals spent $80,000 on 59 stents, 80% of which were

[6] Medicare comparison hospitals obtained by merging data from the Centers for Medicare and Medicaid Services' (CMS) Hospital Compare database from July 2013 with CMS's Healthcare Cost Report Information System (HCRIS) Indirect Medical Education (IME) data for fiscal year 2013. The Hospital Compare files contain data on hospital volume by diagnosis related group (DRG), including DRGs with the description "with cardiac cath" and/or "stent" – this file was used to identify hospitals with cardiac catheterization labs. The IME file contains data on hospital location and total bed capacity.

[7] 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.

drug-eluting (as opposed to bare metal) 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 triple the amount that the smallest hospitals did on stents per month.

Table 2: 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)

Prices for stents have fallen substantially over time as products have proliferated; during 2009-2014, we observe data for twenty branded products sold by four manufacturers – Abbott, Cordis, Medtronic, and Boston Scientific. Between 2009 and 2014, average prices decreased by about 30 percent for both bare metal and drug-eluting stents. Price differences across hospitals are substantial. In Figure 2, we show the distribution of prices across hospitals and products for bare metal and drug eluting stents. The Figure displays the distribution of hospital-product fixed effects, which were obtained from a regression of prices on dummies for hospitals-product combinations, month dummies, and linear product-specific trends; that is, we show here the distribution of prices within product and month, so that price variation is not driven by differences in timing or composition of purchase. As we see in the Figure, drug-eluting stents are far more expensive than bare metal stents ($1,561 vs. $721), and there is substantial dispersion in prices across hospital-products within each category. If all prices were brought down to the minimum price within each product, hospitals would save 15% on average; if, instead, all hospital prices above the mean price within each product were brought down to the mean, hospitals would save 3% on average.

Interestingly, we are unable to associate much of this observed price dispersion with hospital characteristics that would seem a priori to be important for negotiation. For example, in spite of the fact that the largest hospitals spend triple the dollar amount on stent purchases as the smallest hospitals do, we observe no clear relationship between hospital size and stent prices. See Figure 3, in which we display a box plot of bare metal and drug-eluting stent prices for

Figure 2: Distribution of Prices Across Hospitals

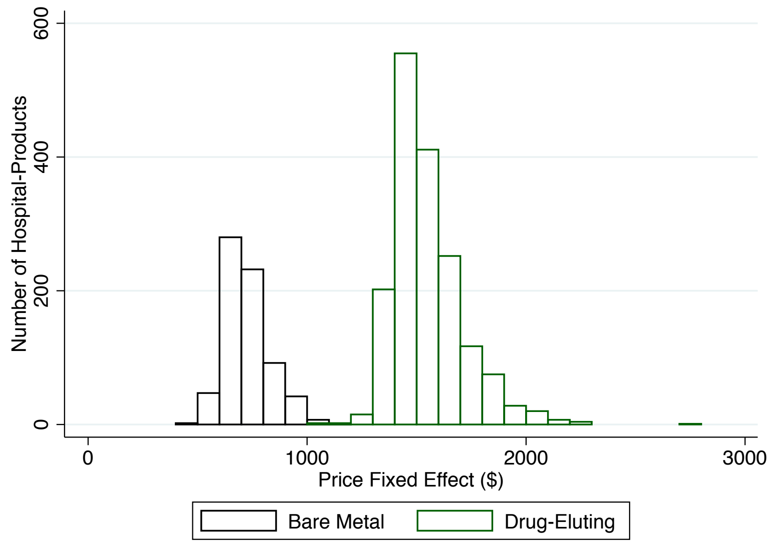

each category of bed count.[8] The price distributions are, if anything, increasing in bed count, though the differences are not statistically significant. Part of this (lack of) relationship is likely due to the heterogeneity in purchasing behavior across hospitals with similar bed counts – small cardiac specialty hospitals purchase stents in greater quantities than similarly-sized acute care hospitals. We cannot directly observe measures of hospital specialization; however, we do observe purchase volume. In Figure 3, we also show box plots of stent prices for each decile of 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 bare metal stents, 10^{th} decile hospitals' prices are 13% lower than those obtained by 1^{st} decile hospitals; the equivalent comparison figure for drug-eluting stents is 6%. 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.

One potential explanation for this residual heterogeneity may be that stents are "physician preference items," products whose demand is determined in large part by physician preferences 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. In the following Section, we discuss the potential role of cooperation be-

[8] As before, "prices" are hospital fixed effects obtained from a regression of price on hospital, month, and product fixed effects.

Figure 3: Distribution of Prices Across Hospitals

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

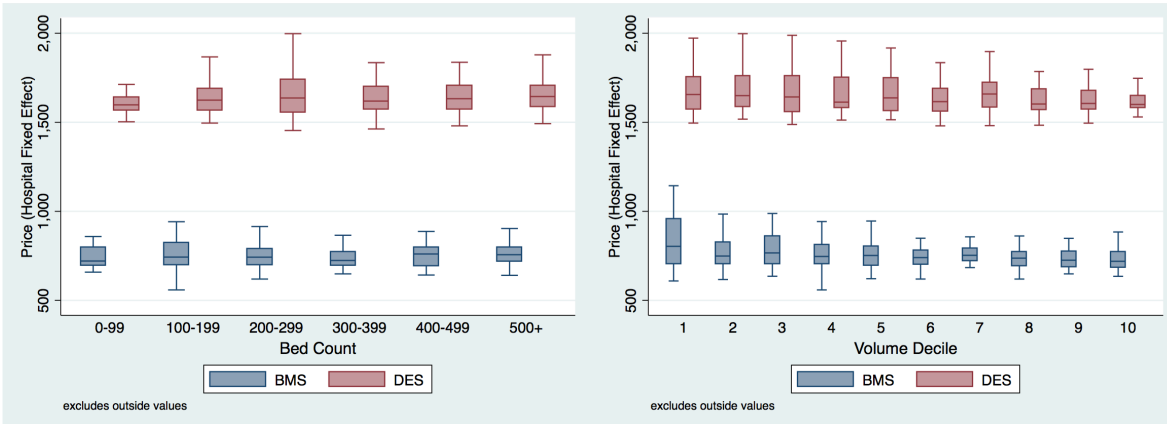

tween physicians and hospitals in determining prices. It is worth noting here, however, that we observe no strong evidence of standardization in our purchasing data. See Appendix A for detail.

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 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. By repeating this final step for each SKU purchased, member hospitals could in principle construct the full purchase order database used in this study, though this process would require a great deal of patience due to the large number of SKUs each hospital purchases and to the daily download restrictions imposed by the web site that hosts the benchmarking service.

In order to analyze the effects of price transparency on negotiations, 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 in addition to the month in which each transaction occurred, 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.

2.4 Research Design: 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 discussion of identification in this Section and of treatment effects in Section 4 focus on the issue of internal validity – consistently estimating 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; and (2) that new members join over time in a staggered (and seemingly random) fashion.

Thus, 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 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.

Our estimation relies on hospitals joining the database at different points in time. While the pattern of hospitals joining over time is relatively flat, we note that the composition of joining hospitals varies somewhat over the period 2009-2014. Figure 5 shows the trend in joining hospitals' potential savings (relative to the mean price for the same product-month) compared to an outside sample chosen to be representative of the set of hospitals with catheterization labs in the US.[9] The Figure shows that on average, those joining the database have slightly higher potential savings in 2010 and 2011, but then lower potential savings in 2012 and 2013. In both samples, however, this is large variation across hospitals in potential savingsso much so that the 95% CI for the means overlap substantially over the entire period 2010-13 (the time horizon for which both samples are available).

[9]Each datapoint is based on a calculation of the potential savings (relative to the mean price) a given hospital would have observed for a given product upon logging into the database. Hence, the magnitude of potential savings is increasing in the joining/comparison hospitals' "pre-join" prices *and* decreasing in the mean price of "post-join" hospitals already in the database. Data are aggregated up to the hospital using hospital-product-specific purchase volumes. Smooth local polynomial approximation of the trends shown for each sample.

Figure 4: Count of Hospitals Joining in Each Quarter

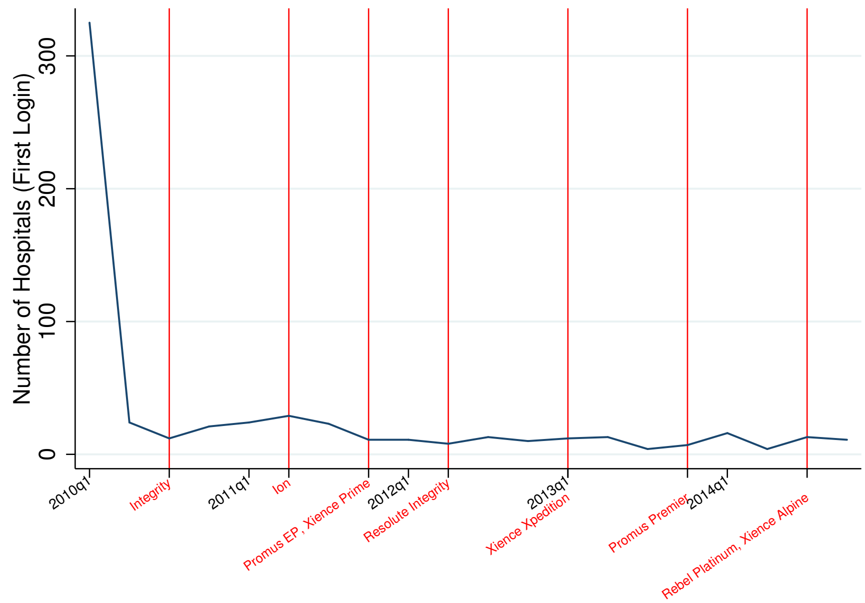

Figure 5: Potential Savings among Joining Hospitals and Comparison Group

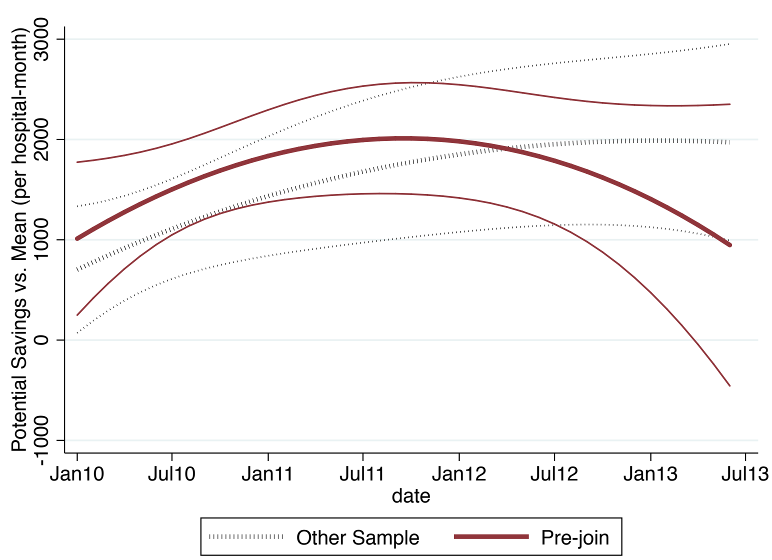

The availability of both pre- and post-join data for hospitals joining the database at different points in time allows us to use a differences-in-differences strategy to estimate the treatment effect of having access to benchmarking information. The logic behind this identification strategy is illustrated in Figure 6. In our sample, there are no pure "control" hospitals – all hospitals by definition access the benchmarking data at some point. However, different hospitals join the database at different points in time. Suppose there are two hospitals, hospital A and hospital B, where A joined the database one period before hospital B. Under the standard differences-in-differences 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 (in practice, we analyze transparency effects across many products, so we capture these using product-hospital fixed effects). In the Figure, these time-invariant differences are identified by the term $\Delta_{t=2}$.[10] Overlapping periods where hospitals are in different information states identify the difference between the two hospitals *plus the difference of access to benchmarking information* – in Figure 6, this difference is $\Delta_{t=1}$, 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, which we capture using time fixed effects and a product-specific linear trend. The time trends are important because prices decreased steadily over time during our period of interest – if we omitted controls for time trends, we would estimate larger effects of access to information based solely on the negative price trend coinciding with hospitals' pre- and post-information periods.

Figure 6: Graphical Illustration of Identification Based on Timing of Join

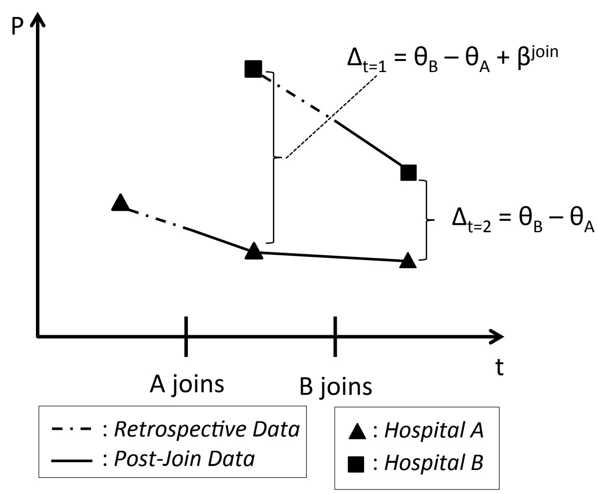

The primary concern with this identification strategy is that timing of a hospital joining the database may be correlated with other contemporaneous factors that may impact price trends at that hospital. For example, a hospital may be inspired to join the database due to particular concerns about price trends, which would bias our results upward by underestimating

[10]In Figure 6, the difference in fixed effects is identified where both hospitals are post-join, but in many cases hospitals that join in different months will have overlapping pre-join data as well.

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, which would bias our results downward by conflating the effects of these other initiatives with the effect of access to the benchmarking information. In Section 4, we bring to bear multiple pieces of qualitative (that stents are only one of many products a hospital purchases) and quantitative (event studies of trends around join timing; comparison to price trends in a different data set of hospital stent purchases; information variation from the introduction of new products) evidence regarding this issue. Our ultimate conclusion is that there is little evidence for timing of join being endogenous with respect to stent price trends for much of our sample, and that even the strongest such bias not ruled out by our tests would leave our main qualitative results unchanged.

2.4.1 Using New Product Entry to Identify Mechanisms

As noted above, we also rely on an additional source of identification: new product entry. New products (at the time of their entry, when they are in fact "new") provide another opportunity to identify the above information effect, and further allow us to identify a treatment effect of joining the database but not having concurrent data on other hospitals' purchases. This is because of the timing of information availability. When a new product is first introduced, no information on other hospitals' purchases of that product will be available in the database for several months,[11] so that *during these first months after new product introduction*, we have overlapping periods where one hospital is post-join (treated, but without concurrently available data) and the other is pre-join (untreated). This 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*; in Section 3, we outline 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.[12]

Figure 7 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

[11] Benchmarking data will not become available for the new products until members submit their purchase order data and they are loaded into the database, a process that will take several months, depending on the timing of purchasing, data sharing lag after purchase, and data entry lag after sharing.

[12] For new products, we also note the minor difference that the identification of fixed differences between hospitals (product-hospital fixed effects) is driven entirely by overlapping periods where both hospitals are post-join (treated), as there can be no overlapping pre-treatment periods before the product is introduced.

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 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. For each entering product, in the months close to the timing of join, we observe many hospitals in each treatment/control state, allowing us to estimate product-specific time trends as well.

Figure 7: Graphical Illustration of Identification Based on Timing of Join and New Product Entry

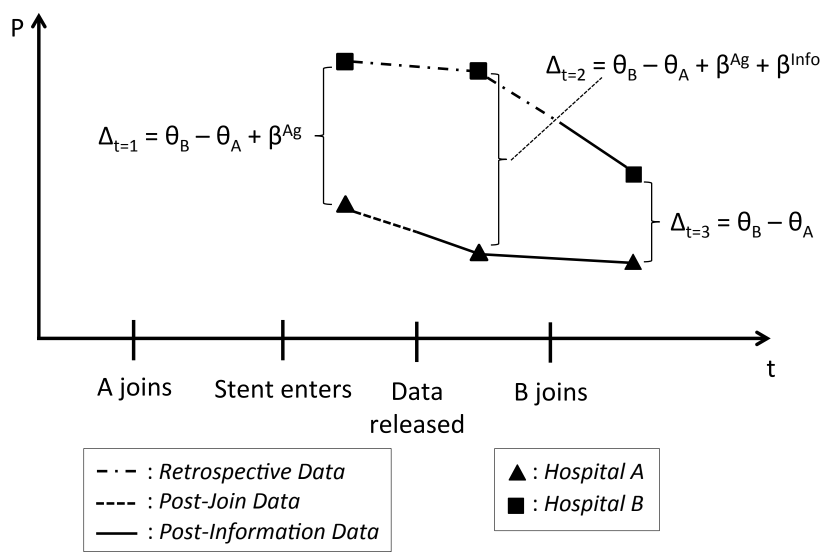

Though it is not the primary way in which we prefer to think about the usefulness of entering products, it is worth noting that entering products also provide two types of robustness checks regarding any potential bias due to timing of join. First, for hospitals that have joined some time ago, new product introductions offer another point in time – a point in time not "near" when they decided to join – at which to compare them to untreated hospitals. The assumption that timing of join is not endogenous with respect to new product entry is supported by Figure 4, in which we do not see spikes of joining around product entry times. Second, any persistent bias associated with something different besides 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 (labeled β^{Ag} in the previous discussion). Thus even in the most extreme case, our estimate of any "asymmetric information" effect where hospitals use information concurrently available in the database to negotiate better prices (labeled β^{Info} above) would be free of such bias.

Critical for this strategy is that the time period for our study contains multiple meaningful product introductions. In Figure 4, we also note the timing of entry of nine new products between 2010 and 2014 (of the twenty products sold during this time period).

3 Theory: Negotiated Prices and Benchmarking Information

Hospitals are typically reimbursed a fixed amount by private or public insurers for services they provide, and the inputs in our purchase order data are required to perform these services. Thus, these prices reflect costs that, at least in the short run, come directly from the hospital's bottom line. For this reason, hospitals are keen to find ways to reduce input costs, and the availability of benchmarking services offers one hope of doing so. As mentioned previously, there is typically no search mechanism in that a given product can only be purchased directly from its manufacturer, especially for the highest spend physician preference items like coronary stents. There are two primary mechanisms that market participants and economic theory suggest 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. Below we outline simple theoretical models that capture each of these effects, and use the models to generate testable predictions we can then take to the empirical analysis.

Our models are built from the baseline of the Rubinstein (1982) model of alternating offers bargaining. This model is useful because it allows for extension in clear and tractable ways to our mechanisms of asymmetric information about supplier parameters and negotiator agency. It is also useful because it forms the underpinning for a large subsequent literature in theoretical bargaining (Rubinstein 1985; Binmore, Rubinstein, and Wolinsky 1983; Horn and Wolinsky 1988; Collard-Wexler, Gowrisankaran, and Lee 2014) as well as a recent industrial organization literature in empirical bargaining studies (Crawford and Yurukoglu 2012; Grennan 2013, 2014; Gowrisankaran, Nevo, and Town 2014; Ho and Lee 2014; Lewis and Pflum 2015). The predictions of the model are useful to map into 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.

Before we consider incomplete information, it is helpful to briefly outline the logic of the Rubinstein (1982) 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.[13] Beginning with the buyer, each player in turn makes a 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 $(0, 1)$.

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 intuition for this equilibrium is that the buyer offers just enough so that the seller is indifferent between accepting the offer and rejecting, incurring a period's discounting, and making a counteroffer (which would in turn be just enough for the buyer to be indifferent between accepting and continuing). The resulting price in this equilibrium is:

$$p^{CI} := c + \delta^S \frac{1-\delta^B}{1-\delta^B\delta^S} V. \tag{1}$$

It will be useful for intuition in both the theory and empirics that follow to map this model into the institutional setting of negotiations over coronary stent prices. The typical negotiation occurs between agents/employees of the hospital and device manufacturer, negotiating on behalf of their employers. On the hospital side, the negotiator will typically be either the catheter lab business unit manager responsible solely for catheter lab operations or a purchasing / materials management professional in the hospital operations department, who may be responsible for a variety of product categories across the hospital. On the device manufacturer side, the negotiator will typically be a regional sales manager. Thus for both negotiators, their respective 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 potential for uncertainty among hospital negotiators (and the managers responsible for the incentives they face) regarding the skill and/or incentives faced by manufacturer negotiators will be the primary focus of our theorizing as to the potential mechanisms via which transparency in the form of price benchmarking information might impact prices. In the Sections that follow, we build off of this baseline model to derive predictions on how benchmarking information might affect prices in cases of asymmetric information (where hospital negotiators are uncertain about the manufacturer negotiator's skill or incentives as embodied in δ^S) and negotiator agency (where hospital negotiator δ^B has an effort component which hospital managers cannot directly observe).[14]

[13] As noted later in our predictions (and discussed and analyzed in detail in Grennan (2013,2014)), V_{jht} should be thought of as the incremental value created by stent j for the set of 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.

[14] We focus on the case where uncertainty is embodied only in the discount factors and not the value over which negotiations occur for simplicity and also the fact that in our discussions with market participants, this seems to be the primary source of potential 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. Because the surplus and bargaining parameters enter the price multiplicatively, similar types of uncertainty regarding either would yield similar predictions.

3.1 Asymmetric Information about Supplier Bargaining Parameters

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 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_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):

$$p_s^{CI} := c + \delta_s^S \frac{1 - \delta^B}{1 - \delta^B \delta_s^S} V, \tag{2}$$

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

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

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:

$$p_s^{AI} := c + \frac{1 - {\delta^B}^2(1 - \omega_w) - \delta^B \omega_w}{1 - {\delta^B}^2(1 - \omega_w) - \delta^B \delta_w^S \omega_w} V, \tag{4}$$

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 $p_s^{CI} > p_s^{AI} > p_w^{AI} > p_w^{CI}$ (where p_w^{CI} is the equilibrium price for the 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 be interested in benchmarking information that would reveal the seller's type, and what would happen to price in such a case. For simplicity, we will assume that this information would fully reveal a seller's type, though the qualitative results should extend 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 attempt 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 3 (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.[15] Thus 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.

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. 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 $1 > \delta_s^S > \delta_w^S > 0$ is greater than that of the weak type. ϵ_h is a random term distributed uniform on $[0, 1]$. It is important to note that 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

[15] 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 primitive of willingness-to-pay and costs.

are 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 $v(\delta_h^B)$. 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 how 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. In our context, the bargaining parameter of the supplier plays the role of an uncertainty (from each principal's view) faced by each purchasing agent. And thus price benchmarking data provides exactly the sort of information that would be useful to a hospital in designing better incentive contracts for its purchasing agents. 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 4 (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.[16]

Prediction 5 (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.

3.3 New Product Entry and the Timing of Benchmarking Information Effects

An interesting feature that differs between the asymmetric information about supplier bargaining type mechanism and the negotiator agency mechanism 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

[16] 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.

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 1-2 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 next theoretical predictions:

Prediction 6 (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, 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 as well as 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 the reality 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 would likely show up (or not show up) in our empirical analysis.

While a hospital joining the benchmarking database has immediate access to the same data we do 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 Propositions 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.

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.

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 2004), facilitating collusion (Albek et al. 1997), or forcing suppliers not to price discriminate via secret discounts (Duggan and Scott Morton 2006; Grennan 2013). 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 (at least in part, though the information will only be available to the benchmarking members, not the entire market), 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, yet not complete, piece of information in considering large-scale transparency policies.

4 Estimation and Results: How Information Affects Negotiated Prices

Figure 8: Histograms of Price Distributions: Pre- and Post-Information

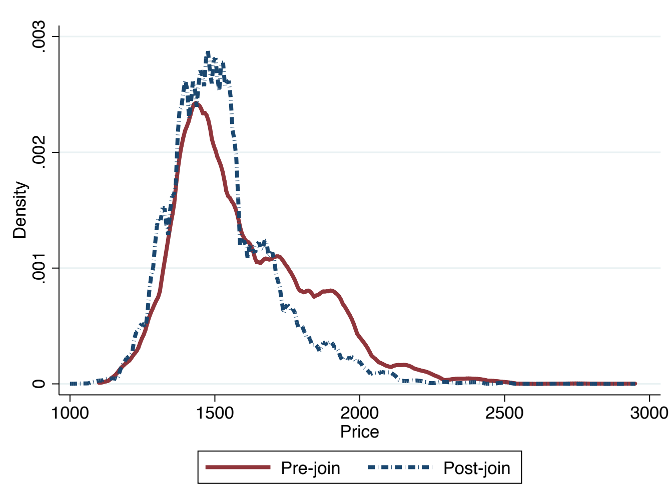

Recall that our baseline research design estimates the treatment effect of access to the benchmarking information on negotiated prices, leveraging the fact that hospitals join the benchmarking in a staggered fashion over time, and each hospital submits a year of retrospective purchase order data at the time of join in addition to ongoing data thereafter. This allows us to construct difference-in-differences style estimators based on the difference in prices paid between pre-join and post-join hospitals. Figure 8 displays the histograms of prices paid of drug eluting stents across the entire sample, split between these two groups. The raw data clearly shows the primary impact of access to the benchmarking information—hospitals with information are much less likely to pay the highest prices. In this Section, we explicitly specify regressions to measure this effect of information, controlling in multiple ways for time-invariant hospital-product differences and product-specific trends in stent prices. We then go on to unpack the mechanisms suggested by the theory. In all cases, we focus our analysis on drug-eluting stents only, as they make up the majority of the usage and spending in the market. Bare metal stents are available in the Appendix.

We begin with an event study of the differences between treated and untreated groups around the time of information access. The event study allows us to be as transparent as possible in establishing the effects we find and in discussing any potential biases around join timing. We then conduct a series of analyses aimed directly at testing the theoretical predictions of Section 3: examining effects conditional on pre-join price and quantity distributions, and using new product entry to disentangle asymmetric information and agency mechanisms (again noting that any remaining worries about bias due to endogenous join timing will be captured in our measure of the agency mechanism). We also examine the underlying drivers of the average price effects by separately considering the effects of information on the likelihood of renegotiation and on price changes conditional on renegotiation. Finally, we use our estimates to extrapolate to the overall effect of access to benchmarking information on the hospitals in our sample and to consider the potential effect of the types of transparency being called for by policymakers.

All of the regressions we present are extensions of a baseline specification implementing the difference-in-differences around the timing of join. Letting P_{jht} denote the price observed for product j, hospital h, and month t; and controlling for hospital-product fixed effects $[\theta_{jh}]$, 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), our preferred specification[17] estimates regressions of the form:

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

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. All of the regressions

[17] In each set of results, we also display estimates with alternative fixed effect controls. In particular, we show estimates with hospital *and* product fixed effects, rather than hospital-by-product fixed effects, and estimates with product-time fixed effects, rather than month fixed effects and product-specific linear trends.

and results below extend this specification to allow for varying types of heterogeneity in this treatment effect.

4.0.1 Event studies around timing of joining database

In our first analysis, we estimate a flexible version of the treatment effect – rather than regressing price on a dummy variable for having access to information, we use an event study specification that includes indicators for each month relative to the hospital's "info" date:

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

The value in the event study versus the baseline regression is that it allows us to examine differences in trends between our treatment and control hospitals that provide evidence regarding the presence of potential biases around timing of join as well as any lags in the treatment effect due to sticky prices.[18]

Figure 9 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.[19] After the hospital accesses the benchmarking information, there is a steady downward trend in the price coefficients. The downward trend in the post-period may be due to price stickiness in that 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. The pooled "post" effect from a differences-in-differences specification is close to zero (-$3) and statistically insignificant. Moreover, estimated patterns are similar across the different specifications of controls, though standard errors are larger in the richest specification shown (Version 3).[20]

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

[18]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 "info" 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 – on average, ten hospitals' data would be available two months after entry, vs. seventy hospitals' data six months after entry. Results are similar for non-entering products only and for different definitions regarding how much data needs to be available to provide meaningful information – accordingly, we consider this to be a pooled "info" effect across products and defer further discussion until the results that separately identify mechanisms.

[19]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.

[20]It was not possible to estimate the monthly event studies with hospital-product and product-month fixed effects. However, the quarterly event study with control Versions 3 and 4 are essentially identical.

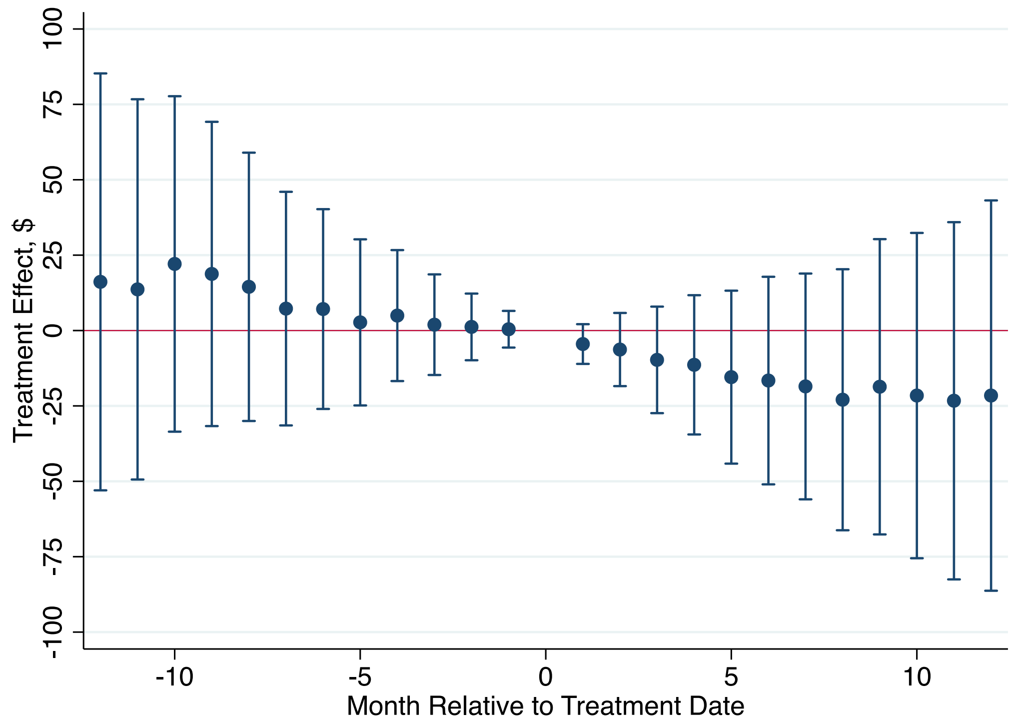

Version	-12	-11	-10	-9	-8	-7	-6	-5	-4	-3	-2	-1	1	2	3	4	5	6	7	8	9	10	11	12
	Month relative to info date ($\beta^{Info,mo}=$)																							
1	17	15	20^*	20^*	17^*	3	-1	-5	3	0	-1	3	$-8^\dagger$	$-10^\dagger$	$-14^\dagger$	$-16^\dagger$	$-20^\dagger$	$-21^\dagger$	$-22^\dagger$	$-24^\dagger$	$-22^\dagger$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

$N = 23,016$ 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 9: Event Studies of Treatment Effect of Access to Benchmarking Information

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).[21] Finally, the average treatment effect of 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.

4.0.2 Treatment heterogeneity across the price distribution

Prediction 1 of Section 3 predicted that, in a model with asymmetric information regarding the supplier's bargaining parameter, benchmarking would lead to price decreases in the upper part of the price distribution (for agency, whether the effect would be at the top or throughout the price distribution depended on the specifics of the model). Accordingly, 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 product-hospital 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 joined the database, $\mathbb{1}_{\{post_{hjt}\}}$, with dummy variables for each pre-join price quintile, $\mathbb{1}_{\{quintile_{jh,pre}\}}$, 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":

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

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 10 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 in our preferred specification. This evidence is consistent with Prediction 1, which predicted that, absent benchmarking, pessimistic hospitals would pay suppliers high prices regardless of those hospitals' true bargaining parameter, so that benchmarking would lead those hospitals to negotiate lower prices after joining. It is also worth noting that we do not see evidence that the lower part of the distribution shifts upward significantly, which would be suggestive evidence of mean reversion (we define the interaction term based on previous periods' prices).

[21] We investigate this issue in some of the following results, 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.

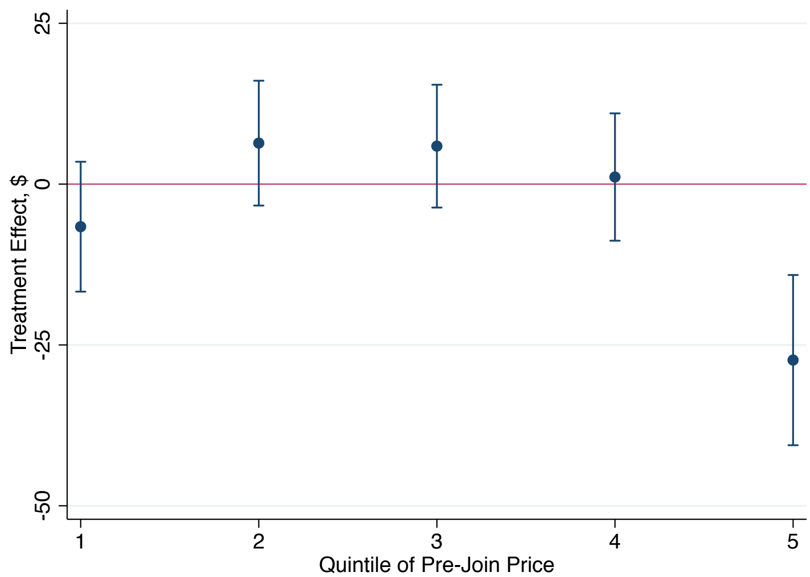

Version	Pre-info price quintiles ($\beta^{Info}_{quintile}$=)				
	1	2	3	4	5
1	4	$17^{\dagger}$	9	-2	$-55^{\dagger}$
	(6)	(6)	(6)	(6)	(9)
2	-1	11	4	-7	$-63^{\dagger}$
	(7)	(7)	(7)	(8)	(10)
3	-7	6	6	1	$-27^{\dagger}$
	(5)	(5)	(5)	(5)	(7)
4	-10	1	2	-3	$-34^{\dagger}$
	(6)	(6)	(6)	(6)	(8)

$N = 32,453$ 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 10: Treatment Effect Estimates Throughout the Price Distribution

The estimated treatment effects do not vary substantially across specification *except* in the top quintile of prices. There, the estimates are significantly smaller when we control for hospital-by-product (rather than hospital and product) fixed effects. This may be evidence of hospitals' shifting volume to less expensive products after accessing the database; we consider the preferred specification to be a conservative estimate of the effects of information on negotiated price within hospital-product.

We also performed the 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 11. The Figure repeats the patterns from the ATE: Again, 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; it appears that, if anything, pre-trends prior to joining the database would lead the differences-in-differences 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 benchmarking data rather than to any potential sources of bias. 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 (though again we note for the most skeptical interpretation that 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). Also, 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.

4.1 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 product-hospital 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 product-hospitals with high expenditures at stake experience larger average price changes, in keeping with a model with effort cost of search and renegotiation (Predictions 2 and 5). Next, we use 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 6). Finally, we decompose the

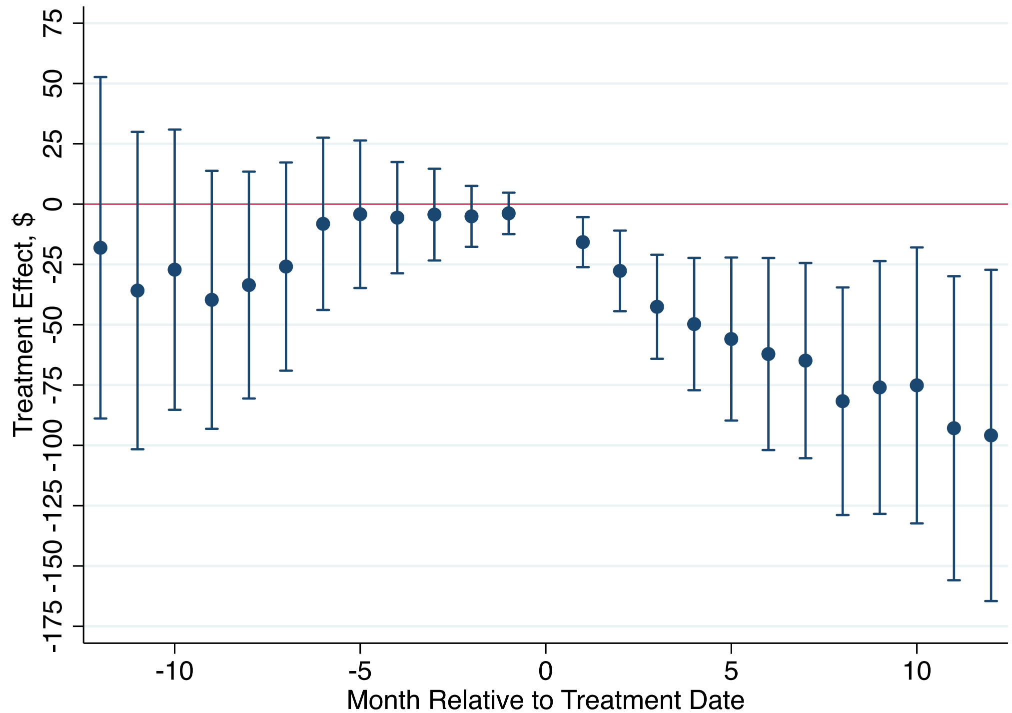

Version	Month relative to join date ($\beta_5^{Join,mo} =$)																							
	-12	-11	-10	-9	-8	-7	-6	-5	-4	-3	-2	-1	1	2	3	4	5	6	7	8	9	10	11	12
1	-24	-33	-20	-24	-14	-14	-21^{*}	$-31^{\dagger}$	$-27^{\dagger}$	$-22^{\dagger}$	-14^{**}	-8	$-23^{\dagger}$	$-40^{\dagger}$	$-52^{\dagger}$	$-61^{\dagger}$	$-64^{\dagger}$	$-74^{\dagger}$	$-77^{\dagger}$	$-88^{\dagger}$	$-83^{\dagger}$	$-78^{\dagger}$	$-104^{\dagger}$	$-104^{\dagger}$
	(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)

$N = 23,016$ 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 11: Event Studies of Treatment Effect of Access to Benchmarking Information, Top Quintile of Price Only

estimated price effects into price effects conditional on renegotiation and price effects due to greater likelihood of renegotiation.

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

To the extent that search and renegotiation are costly, Predictions 2 and 5 predict that benchmarking data will be sought and used most effectively 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. To implement this, we generate a dummy variable based on the quantity used of each product at each hospital: $\mathbb{1}_{\{high^q_{jh,pre}\}}$ is equal to one for hospital-products with monthly purchase volume above the 75^{th} percentile in the months prior to join. The specification we estimate is:

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

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 results are shown in Figure 12.

The estimates 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 double the effect for low-quantity hospital-products in the preferred specification. These results are consistent with a positive effort cost of search and renegotiation leading to decreases in high prices for high-volume purchase combinations in particular. It is worth noting 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 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.

4.1.2 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 several potential theoretical mechanisms, in particular 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

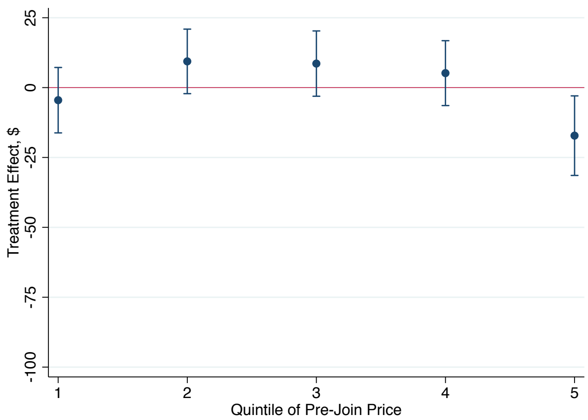

(a) Low Quantity

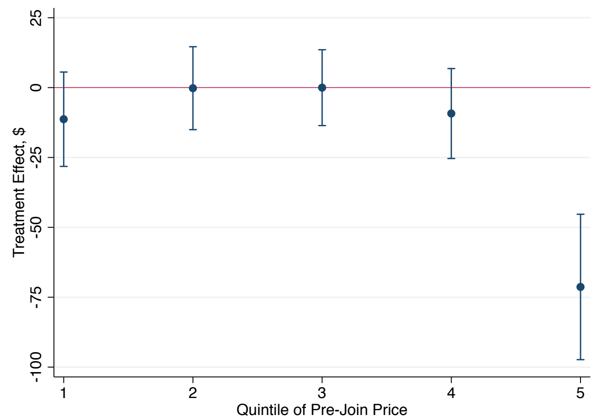

(b) High Quantity

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	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)

$N = 32,453$ 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 12: Treatment Effect Estimates Across the Price and Quantity Distributions

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 product-hospital-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:

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

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. The results are shown in Figure 13.

While the separate results are estimated more imprecisely, the point estimates consistently suggest that the asymmetric information effect explains a substantial portion of information on prices. For the product-hospital 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, but not statistically significantly different from zero after controlling for unobserved differences across hospital-product combinations. However, in specifications versions 1 and 2 with hospital fixed effects only, β^{Ag} and β^{Info} each explain approximately half of the total effect. It is possible that they difference between these two studies points to challenges with attenuation bias in the product-hospital fixed effect models, which leave very little identifying variation, especially for β^{Ag}.

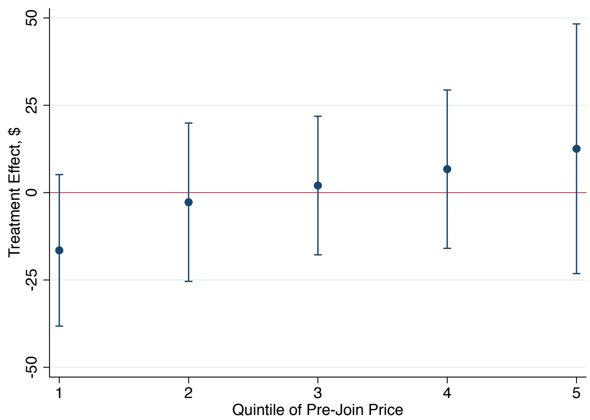

(a) Agency

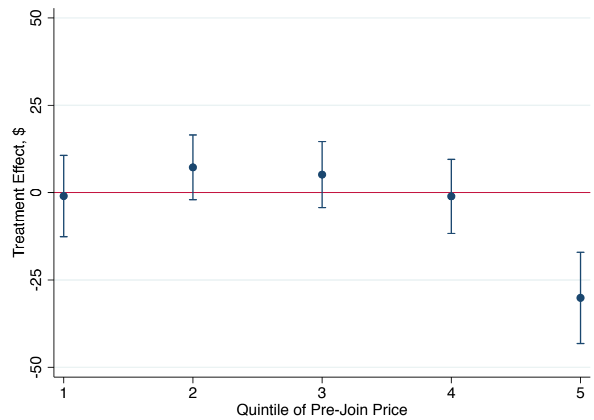

(b) Asymmetric Info

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)

$N = 32,453$ 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 13: 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 that the timing of join correlates with other hospital activities unrelated to benchmarking information that decrease prices, this bias will be captured in β^{Agency} but not β^{Info}. This is because in our research design, β^{Agency} 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).

For empirical work, this suggests that asymmetric information may be among the effects 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). In the most innocuous case, it suggests 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, and in what way any information changes might induce changes in the relevant bargaining parameters used in estimating the negotiated outcomes.

4.1.3 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 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. We then estimate the usual price quintiles specification on the subset of months in which renegotiation takes place:

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

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

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

In the main estimation sample, renegotiations take place in 9% of observations (member-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 14.

The results in the left panel of Figure 14 show that the effect of information on price conditional on renegotiation is substantially larger than the effect of information on price paid. In the top quintile of the price distribution, we see 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. The effects are estimated imprecisely, but the 95% confidence intervals indicate that price decreases at renegotiation are two to three times larger with information than without. The right panel of the Figure shows that information increases the likelihood of renegotiation throughout the price distribution – in the top quintile of the price distribution, information increases the probability of renegotiation by 2.3 percentage points, which is nearly one-third the baseline probability of renegotiation.

Taken together, these results show that information affects price paid through two channels, by making renegotiation more likely to occur and by leading to larger price decreases when renegotiation takes place.

4.2 Achieved Savings from Information

Thus far we have considered the magnitude of the information effects in dollar terms at the product-hospital-transaction level, allowing for heterogeneity in treatment effects. In this Section, we use those treatment effect estimates – properly weighted according to the observed volume and price distributions – to calculate the savings achieved due to access to benchmarking information. We compare this percent "achieved savings" to the "potential savings"

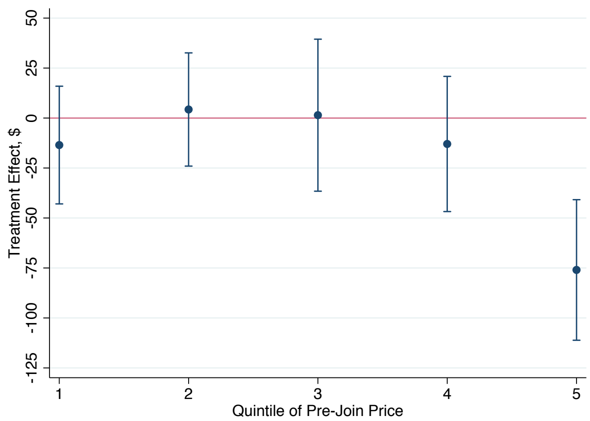

(a) Conditional on Renegotiation

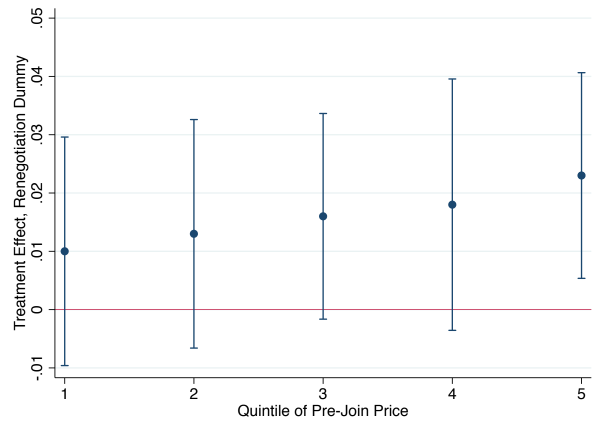

(b) $\mathbb{1}_{\{post_{ht}\}}$

Version	Pre-join price quintiles ($\beta^{Join}_{quintile} =$)					Pre-join price quintiles ($\beta^{Join}_{quintile} =$)				
	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)

$N = 6,510$ 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 14: Treatment Effects Conditional on Renegotiation and on Occurrence of Renegotiation

numbers that are based on the pre-information heterogeneity in prices across hospitals.[22] We then display the distribution of savings across hospitals in dollars per year. These numbers are informative for examining the value of the benchmarking service whose data we study, and also as a step towards projecting the potential aggregate savings in the case that a transparency policy such as the ones proposed by policymakers were to achieve the same treatment effect as the benchmarking service we study. We take care to interpret these projections with caution due to the potentially selected nature of our sample and the potential supply side responses to a nationwide policy that may not be captured in our treatment effect.

Recall that, in Section 2, we constructed "potential savings" based on the heterogeneity in the prices paid by different hospitals for the same product over the same period of time. We constructed this metric using the subset of our data containing each hospital's year of pre-information data, and extracting a product-hospital specific fixed effect, controlling for product-time fixed effects ($\hat{p}_{jh}$ from the regression $p_{jht} = \hat{p}_{jh} + \hat{p}_{jt} + \hat{\varepsilon}_{jht}$). Potential savings is defined as the difference between this and the mean of the distribution across hospitals for each product: $PS_{jh} := \max\{0, \hat{p}_{jh} - \bar{p}_j\}$.

Similarly, we can use our estimated treatment effects across the price and quantity distributions to construct "achieved savings" due to access to the benchmarking information service, in dollars per hospital-product: $AS^{Info}_{jh} := \hat{\beta}^{Info}(quintile_{jh}, high^q_{jh})$.

Figure 15 displays potential and achieved savings per stent across product-hospitals based on their position in the price/quantity distributions. Each calculation is based only on the data used to estimate the regression specifications (e.g., excluding pre-join data for facilities joining the database before Q1 2010). 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 (defined as savings that would accrue if all prices $\hat{p}_{jh}$ were altered to $\tilde{p}_{jh} = min\{\hat{p}_{jh}, mean_h\{\hat{p}_{jh}\}\}$. Savings are not substantial for lower points in the price distribution, but it should be noted that potential savings are mechanically not substantial for hospital-products already achieving lower prices.

Figure 16 displays the distributions of savings achieved by the hospitals in our sample, in terms of total savings per hospital per month. The average hospital achieves \$346 in savings on stents per month, but this average effect conceals substantial heterogeneity. 10% of sample hospitals save \$2,500 per month on stents alone, and savings are statistically significant below the 20^{th} percentile. At the top of the distribution, the estimates indicate that 10% of hospitals *lose* \$221 after joining the database, but none of the positive effects are statistically significant at conventional levels.

[22]Depending on the specific model of the world, achieved savings need not be bounded above by potential savings, and further the potential savings may be due largely to hospital- or product-hospital-specific factors that have nothing to do with information. However, among many reasonable models of the world, potential savings is exactly an upper bound for what might be achieved by information, and given the policy interest based upon the observed variation in prices across hospitals, this seems like a natural and useful benchmark.

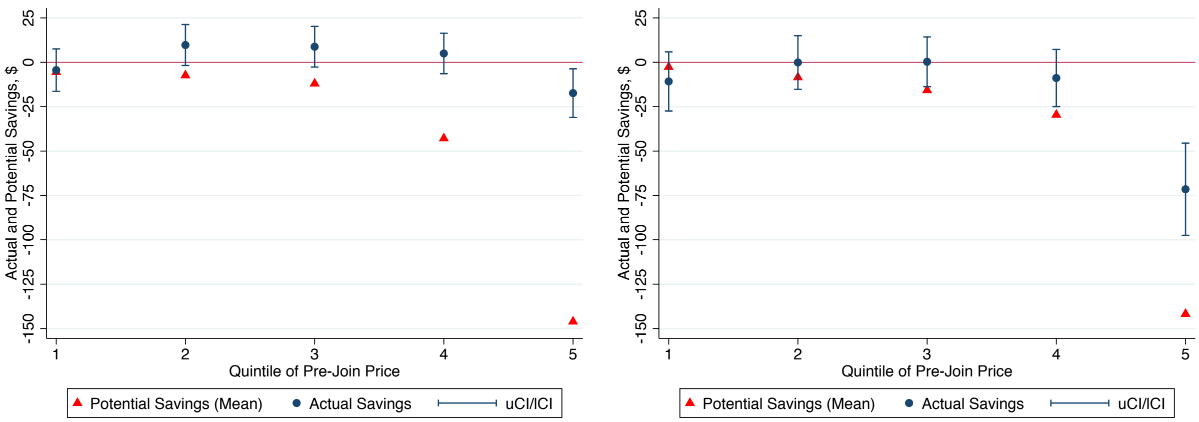

(a) Low Quantity **(b)** High Quantity

Version	Pre-join price quintiles $(PS/AS_{quintile,low^q} =)$					Pre-join price quintiles $(PS/AS_{quintile,high^q} =)$				
	1	2	3	4	5	1	2	3	4	5
Potential savings	-6	-7	-12	-43	-146	-3	-8	-16	-30	-142
Achieved savings	-4	10^{*}	9	5	-17^{**}	-11	0	0	-9	$-72^{\dagger}$
	(6)	(6)	(6)	(6)	(7)	(8)	(8)	(7)	(8)	(13)

$N = 32,453$ member-product-months. Includes 508 members. Both actual and potential savings based on regressions using "Version 3" controls. Potential savings calculated using pre-information data only. 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. Superscript (†) indicates significant difference from zero at the 1% level; (**) at the 5% level; (*) at the 10% level.

Figure 15: Achieved Savings vs. Potential Savings per Stent

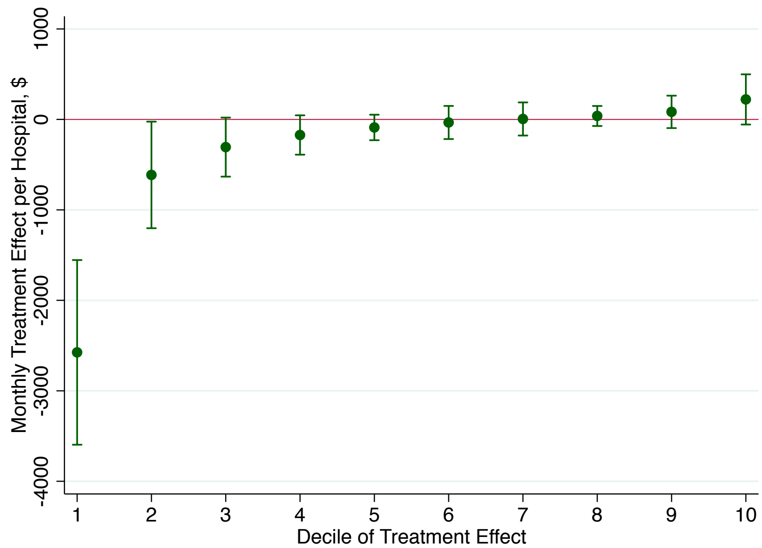

Percentile of Treatment Effect ($TE_{decile}/month/hospital$=)									
10	20	30	40	50	60	70	80	90	100
−2575	−614	−306	−172	−89	−34	5	38	83	221
(521)	(300)	(166)	(111)	(72)	(93)	(93)	(56)	(91)	(142)

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 16: Achieved Savings Across Hospitals ($/month/hospital)

4.2.1 Policy Implications

The "achieved savings" figures above are limited in their applicability for policy analysis by two factors. First, they are based on a selected sample of hospitals that voluntarily joined a purchasing database. Joining hospitals may differ from other hospitals in the US in both observable and unobservable ways; thus, our results may not generalize to US hospitals overall. Second, our results may be "partial equilibrium" results in that they may not capture the effects of a policy of full transparency applied to all US hospitals.

To the first issue, we can extend our results to a "more" representative sample of US hospitals using the previously mention MRG survey of catheter labs (the source that major device manufacturers subscribe to for detailed market research). To perform this analysis, we flag each outside 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 outside sample hospital observations. As discussed above, the regression sample hospitals are larger and have ex ante lower prices; 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 an average outside sample hospital would have achieved $421 in savings on stents per month (vs. $346 in our regression sample) if it accessed benchmarking data. This would suggest that our estimates are conservative with respect to the US population of facilities overall. However, we cannot control for unobservable ways in which our sample differs from the average US hospital with a catheterization lab (e.g., joiners may have better management practices and be better able to utilize databases). We hope to explore this issue in ongoing work with better data on database hospitals.

To the second issue, our analysis is one of the effect of introducing benchmarking information to a hospitals in a way that we expect captures supply responses in a limited fashion. To the extent manufacturers know when a hospital has access to this information (which we understand they do), then our estimates incorporate effects on hospitals who get "bad news" when joining the database or reluctance of manufacturers to lower prices to database members. We also do not find any evidence of increasing "obfuscation" in the form of proliferation of SKUs as more hospitals adopt the benchmarking database over time. However, our data does not allow us to be certain the extent to which the incentives for and activities regarding supply side price cuts, collusion, or obfuscation might change with a rollout of a broader transparency policy.

5 Conclusion

This paper conducts one of the first studies of the impact of information in negotiated price markets, motivated by the rise in benchmarking data services marketed towards buyers in business-to-business markets and calls for greater transparency in these markets by policymakers. Our empirical study is done in the context of 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 ten percent of US hospitals from 2009-14 and a differences-in-differences research design to compare the prices negotiated by hospitals with and without benchmarking information on what other hospitals pay. The estimated average treatment effect of this type of information across all product-hospital-months for coronary stents ranges is small, but masks dramatic heterogeneity. We estimate that the conditional average treatment effects are large for hospitals paying especially high prices for a given product, and even larger when these products are also used in large volume. There is also evidence that stickiness in price renegotiation mutes the value hospitals can reap from benchmarking information.

While our results suggest that on net policies or intermediaries that increase transparency may indeed lower the prices hospitals pay for medical supplies, our hope is that this study opens more doors than it closes. Coronary stents are just one product category (albeit an important one), and the results are likely to be different for different medical products, let alone for different industries. While our data contains purchase orders for nearly 3,000 categories and 2 million product SKUs, analysis of other product categories using this price and quantity data alone may be complicated by the impact of unobserved nonlinearities or bundling in contracts. We believe this reinforces the need for more data collection and theory development.

In the large existing theory on bargaining and incomplete information, we were surprised that no model quite captured the main phenomena of interest here. 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 suggested by our analysis for future theory development.

Appendices

A Appendix

A.1 Checks for Standardization and Share-based Contracts

In Figure 17, we show histograms of total manufacturers and total stent brands purchased 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.

Figure 17: Histograms – Number of Brands/Manufacturers per Hospital

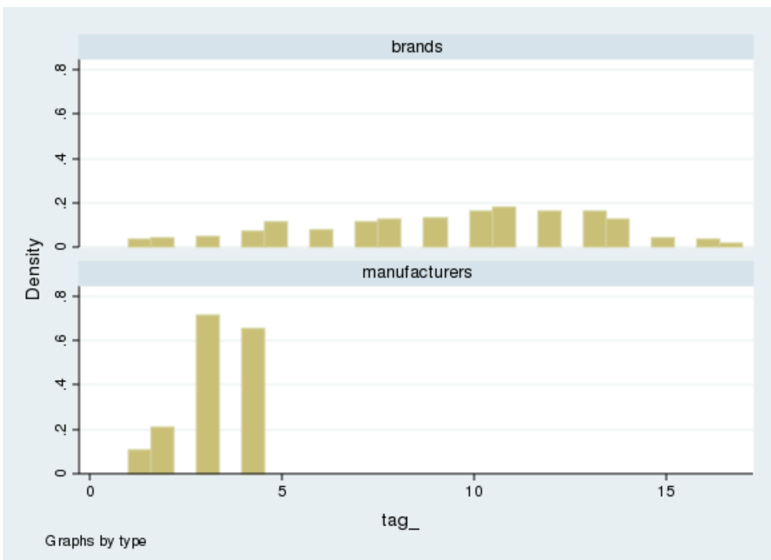

As a practical matter, stents tend to have simple contracts, so that we can be confident our transactions data captures real prices – we cannot observe volume or market share rebates in the data, so to the extent such rebates were common we would have substantial measurement error. Figure 18 gives some reassurance on this point, showing a histogram of the market share of each hospital's most-preferred product, grouping hospitals by quintile of overall stent volume purchased. We observe very little evidence of hospitals bunching at market share thresholds where we would expect to see market share rebates be awarded (for example, 80 percent, 75 percent, etc.), as would be expected if nonlinear contracts were popular in the sample.

A.2 Additional Tables

Figure 18: Market Share of Most-Preferred Brand, by Quintile of Purchase Volume

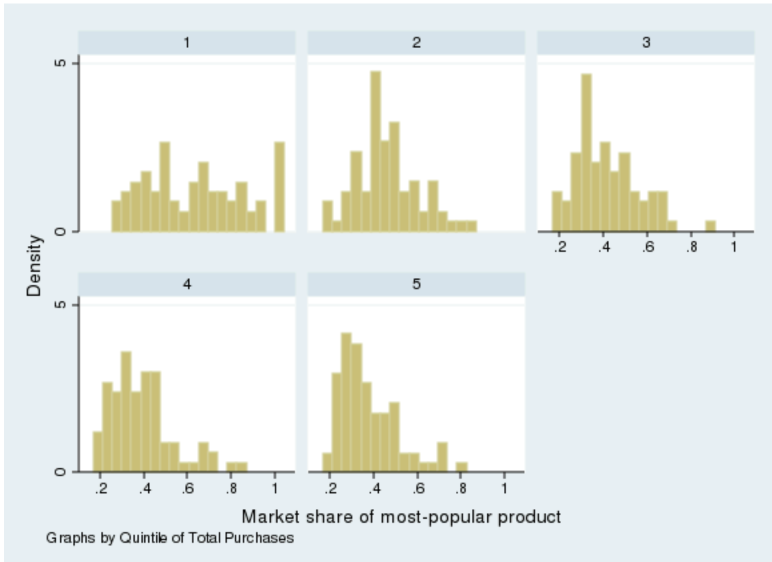

Table 3: Treatment Effect of Price Information on Different Parts of the Price Distribution: Bare Metal vs. Drug-Eluting Stents

	Bare Metal Stents						Drug-Eluting Stents					
		Differential Effect for Hospital-Products above Yth Percentile						Differential Effect for Hospital-Products above Yth Percentile				
	Mean	10th	25th	50th	75th	90th	Mean	10th	25th	50th	75th	90th
Post	-0.3	-10.1	-4.8	2.0	5.6	4.2	-6.4	-10.0	-3.6	2.3	1.9	0.1
	(4.0)	(8.9)	(7.9)	(5.9)	(5.5)	(5.0)	(4.9)	(7.6)	(5.8)	(5.6)	(5.2)	(4.9)
Post*High Price		12.2	6.7	-4.9	-39.2**	-64.9**		7.0	-0.7	-14.4**	-34.3†	-75.1†
		(8.4)	(7.8)	(6.7)	(17.6)	(28.5)		(7.8)	(6.1)	(7.0)	(8.9)	(18.2)
N	95,635	65,656	65,656	65,656	65,656	65,656	253,976	169,981	169,981	169,981	169,981	169,981
N	95,635			65,656			253,976			169,981		

Standard errors clustered at hospital-product level shown in parentheses. Superscript (†) indicates significance at the 1% level; (**) indicates significance at the 5% level; (*) indicates significance at the 10% level.

Table 4: Heterogeneity in Results: Treatment Effect of Price Information for Different Quintiles of "Potential Savings"

	All Products						Entering Products (6 Month Buffer)					
		Quintile of Potential Savings						Quintile of Potential Savings				
	Mean	1st	2nd	3rd	4th	5th	Mean	1st	2nd	3rd	4th	5th
Post	-1.4	1.1	-2.8	-1.6	-0.7	6.9	-1.4	-6.7**	-9.1†	-7.0**	-6.1*	3.4
	(4.7)	(5.2)	(5.7)	(5.5)	(5.0)	(4.9)	(4.7)	(3.4)	(3.5)	(3.4)	(3.3)	(3.4)
Post*Xth Quintile		-2.3	13.4*	10.9	11.4*	-29.6†		7.2	18.0†	11.0*	6.9	-39.7†
		(9.7)	(7.5)	(7.0)	(6.6)	(10.3)		(6.6)	(5.7)	(5.7)	(5.8)	(9.7)
N	349,613			235,634			130,855			83,882		

Standard errors clustered at hospital-product level shown in parentheses. Superscript (†) indicates significance at the 1% level; (**) indicates significance at the 5% level; (*) indicates significance at the 10% level.

References

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

Bronnenberg, B., Dube, J.-P., Gentzkow, M., and Shapiro, J. (2014). Do pharmacists buy bayer? informed shoppers and the brand premium.

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

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.

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.

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.

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.

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.

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).

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.