

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

PRELIMINARY – PLEASE DO NOT CITE OR CIRCULATE

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

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

February 18, 2015

Abstract

This paper empirically analyzes the role of information in bargaining between hospitals and their suppliers. Hospital supplies account for a large percentage of both the level and growth of health care expenditures, and prices for the same input can vary dramatically across hospitals. This variation has prompted calls for increased transparency as a mechanism to lower prices, but whether such an intervention would be successful depends on the details of what information is provided and how market participants respond. We analyze a new data set including all purchase orders issued by over ten percent of US hospitals over 2009-13. The empirical setting contains an intervention in which sample hospitals gained access to benchmarking data on other hospitals' negotiated prices. Using differences-in-differences identification strategies based on timing of hospitals' access to price information and on new product entry, we find that access to information on purchasing by peer hospitals led to large reductions in prices, concentrated among hospitals learning that they were performing relatively poorly in contracting and for products purchased in relatively large volumes, and that the usefulness of price information is constrained by delays in renegotiation between hospitals and suppliers.

*This paper benefited from comments from audiences at the 2015 ASSA meetings, University of California-Berkeley, Dartmouth Winter IO Conference, iHEA 2014, Indian School of Business, University of Louisville, Queens University School of Business, 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

Suppliers in business-to-business markets often 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. Intuition suggests this information would be valuable to buyers in markets with imperfect or asymmetric 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 this paper, we use a new data set on all purchase orders issued by ten percent of US hospitals from 2009-13 to estimate the impact of access to benchmarking information on the prices hospitals negotiate with their suppliers.

We have two primary goals in this study: (1) to estimate the average treatment effect (ATE) 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 area for this analysis as they have accounted 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.¹ 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.² The policy attention given to these prices reflects both a concern for the financial viability of hospitals and also a concern that rising supply costs over time filter downstream into higher costs for the health system and consumers.

For the most important product categories in this setting, individual buyers 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 in this context are: (1) a model in which hospitals

¹For 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.

²For 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)

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 test the predictions of each model using our empirical analysis 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 2013. 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 the 386 hospitals with cardiac catheterization services. Stents are a desirable category because they are important (one of the largest categories, comprising two percent of supply spend and about \$2 billion annually in the US) 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. And those 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 15.4 percent. 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 in hospital-product prices and product-specific price trends. The assumption underlying this approach is that timing of join is uncorrelated with latent hospital-specific price trends. In particular, this strategy would fail if hospitals join when they are experiencing increases in input prices, or when they are enacting other cost-cutting measures beyond benchmarking. To provide reassurance on this point and to investigate model predictions that differ based on whether the product is in the market when the hospital joins the database, we also perform additional analyses focusing only on *new* products entering the market during 2009-2013: because no information on others' prices is available when a new product first enters the market, this provides exogenous variation in the timing of information availability.

The estimated ATE across product-hospital-months for coronary stents suggests that simply having access to the information in the database results in zero savings. This average estimate,

however, conceals substantial heterogeneity. Hospital-products whose prices are above the 75th percentile experience price declines of -\$21 to -\$34 per stent upon accessing database information; those hospital products above the 90th percentile of the price distribution show even larger declines of -\$34 to -\$80 per stent (to give this context, the average size hospital uses almost 1,000 stents annually and the average stent price is just over \$1,200). The price declines are larger for product-hospital combinations with larger purchase volumes at stake – for hospital-products above the 75th percentile in monthly purchase volume prior to joining the database, price effects increase to -\$38 to -\$46 at the 75th price percentile, and to -\$50 to -\$119 at the 90th price percentile. In each comparison, the treatment effects are larger among entering products than among products already in the market when the hospital joins the database. This suggests that, if anything, treatment effects are biased toward zero due to factors contemporaneous with joining.

The heterogeneity in results is consistent with the predictions of the model of bargaining under asymmetric information. 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. Finally, treatment effects are also larger when we focus attention on entering products, for which the moral hazard mechanism is not likely to lead to price declines when benchmarking data are accessed. This provides strong evidence in favor of the asymmetric information mechanism.

Measurement of the average treatment effect of access to benchmarking information is complicated by a number of frictions that push the short-run effect for a product-hospital pair toward zero. Once a hospital joins the database, someone at the hospital must anticipate sufficient potential gains for a product to search the database for information regarding what other hospitals are paying. The information found must then be useful in the sense that it suggests the hospital has an opportunity to negotiate a better price. Finally, 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 each of these potential frictions have substantive costs in terms of time or resources, a treatment effect measured based on the first step of having access to the benchmarking data will be biased toward zero. For this reason, we also estimate treatment effects conditioning on months in which we observe renegotiation taking place – this limits our sample (and statistical power) substantially, but the 90th percentile treatment effects increase significantly in magnitude. 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

Theory of how benchmarking information matters in business-to-business markets where prices are negotiated may differ from that where price-taking consumers search among retail outlets, but both relate to the broader question of the role of informed buyers on market outcomes. 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), generally finding that demand side effects of information are on average null or beneficial to buyers. Our findings extend to negotiated prices and to models of bargaining under uncertainty and negotiator effort as explanations for how information can matter for buyers.

Our study also relates to the literature on supply side responses to transparency. In this literature, supply side responses can negate or overturn welfare-positive demand-side effects of transparency due to greater obfuscation (Ellison and Ellison 2009), facilitating collusion (Albæk 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. This makes our estimates potentially useful for considering the policy proposals that have put forth imposing greater transparency (in the most recent proposal through requiring quarterly reporting of average, minimum, and maximum prices) as a way to decrease prices for medical inputs.

Finally, the finding of a meaningful role of information on negotiated prices suggests an important future direction for the empirical bargaining literature (Crawford and Yurukoglu 2012; Grennan 2013; Gowrisankaran, Nevo, and Town 2014; Ho and Lee 2014), which has thus far modeled negotiations of perfect information.

The paper proceeds by first examining the data, setting, and research design in Section 2. Section 3 discusses potential theoretical mechanisms for how benchmarking data might affect negotiated prices, based on existing theory and claims of industry participants. Section 4 presents our empirical results on the overall ATEs and also cuts of the data designed to better understand mechanisms. 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 2009-2013. 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 (SKU and Universal Medical Devices Nomenclature System (UMDNS) code),³ and supplier. We do not observe decrypted hospital information, but we observe unique identifiers for each member and the data include several coarse member characteristics: census region, facility type, and

³UMDNS is a standard international coding system for medical devices developed by the ECRI Institute.

number of beds. Table 1 displays some summary statistics regarding the transactions data. We observe transactions for 1,384 members, 652 of which are hospitals. On average, we observe 25 months of transactions for members, 30 for hospitals. We observe purchases in more product categories for hospitals than for all members on average (754 vs. 441). The average hospital in our database spends \$2.2 million per month on all supplies, \$52 thousand of which is dedicated to coronary stents. As expected, hospitals generate the majority of the spending on stents – 60% of hospitals purchased stents during 2009-2013, vs. 33% across all members.

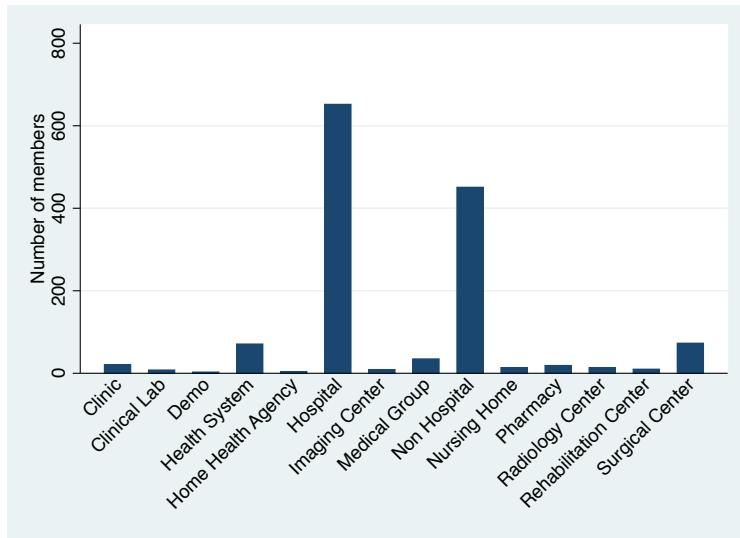
Table 1: Summary Statistics from Purchase Order Database

	All Members [N=1,384]		Hospitals [N=652]	
	Mean	SD	Mean	SD
Months of Data	24.74	14.83	29.30	14.16
Product Categories	441.38	434.61	754.35	339.77
Total Spend/Month (\$m)	1.36	2.57	2.20	2.73
Purchases Stents?	0.33	0.47	0.59	0.49
Total Spend/Month on Stents (\$k)	30.31	69.15	52.33	81.35

Notes: Summary statistics for all members in purchase order database, with hospital members broken out separately. Count of "Product Categories" based on UMDNS code in transaction record.

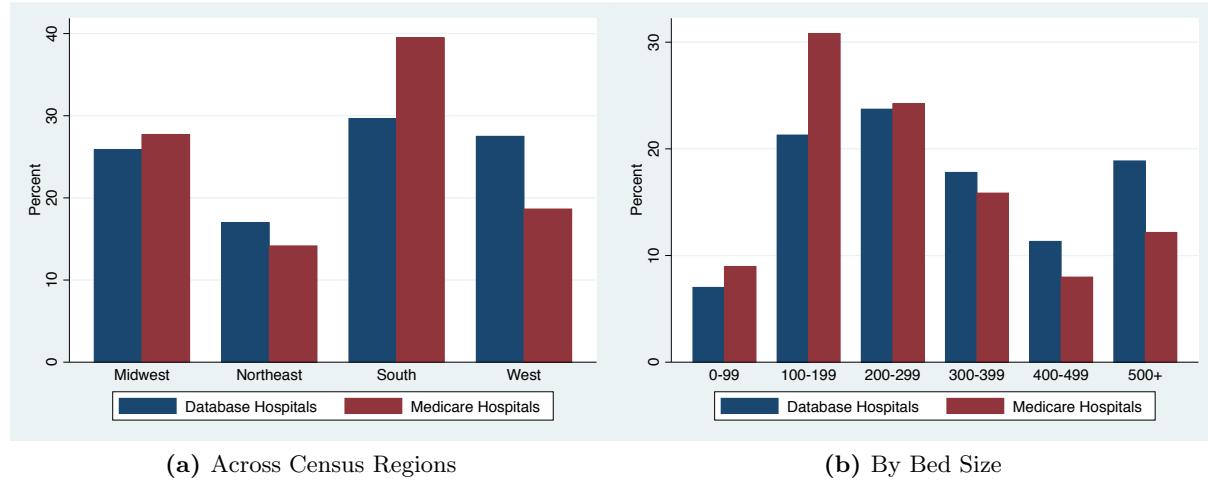
The data include multiple different types of facilities, such as hospitals, nursing homes, and surgical centers. See Figure 1 for the distribution of member types. The vast majority of members are hospitals, but there is also a large mass of members coded as "Non Hospital" or "Surgical Center." In this analysis, we restrict the sample to focus on the 386 hospitals we observe to purchase stents.

Figure 1: Facility Types in Purchase Order Database



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. Indeed, the sample is not perfectly representative of US hospitals with catheterization labs. The left panel of Figure 2 compares the distribution of sample hospitals across US census regions to that of Medicare-certified hospitals with cardiac catheterization labs.⁴ 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 2 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. This 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. In our Results and Conclusions, we will discuss this issue of nonrepresentativeness.

Figure 2: Distribution of Benchmarking Database vs. Medicare 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

⁴Medicare 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.

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 procedures⁵; 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 27 months. In a given month, sample hospitals spent \$96,000 on 67 stents, 74% of which were 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 ≥ 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 Count	Members	Months	Monthly Expenditure (\$k)	Monthly Quantity	% DES
0-99	26	21.2	50.0	35.5	76.5
		(16.7)	(48.1)	(33.6)	(21.0)
100-199	79	23.9	49.7	34.7	70.9
		(14.8)	(53.7)	(34.6)	(24.0)
200-299	88	30.1	69.7	48.7	73.5
		(14.4)	(66.3)	(46.4)	(19.5)
300-399	66	27.5	87.9	61.6	74.8
		(14.2)	(56.1)	(39.2)	(19.7)
400-499	42	29.6	140.6	97.8	74.2
		(14.9)	(102.4)	(69.8)	(16.4)
500+	70	29.0	172.8	121.9	78.0
		(13.6)	(203.2)	(152.1)	(12.9)

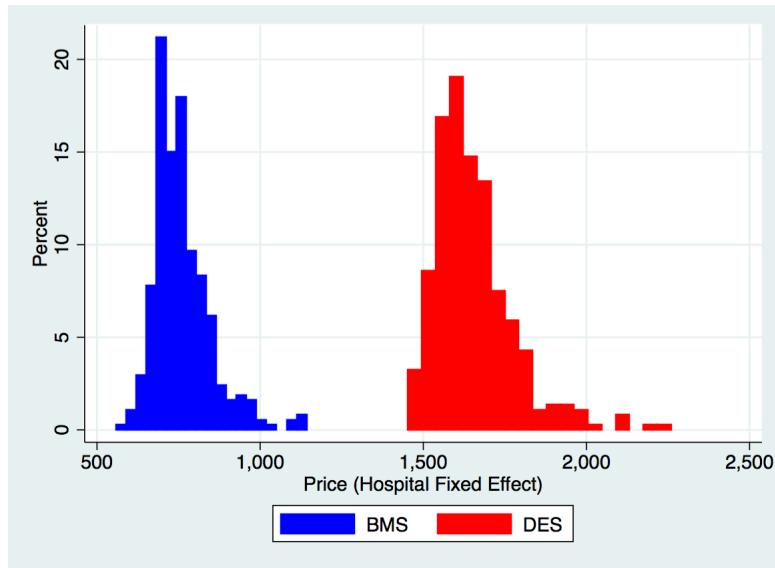
Notes: Summary statistics regarding stent transactions for hospitals in purchase order data. Sample restricted to hospitals with any stent transactions. Standard deviations in parentheses. "% DES" indicates percentage of members' coronary stent purchases devoted to drug-eluting (rather than bare metal) stents.

Prices for stents have fallen substantially over time as products have proliferated; during 2009-Q2 2013, we observe data for eighteen branded products sold by four manufacturers – Abbott, Cordis, Medtronic, and Boston Scientific. Between Q1 2009 and Q2 2013, average prices decreased by 24 percent. Price differences across hospitals are substantial. In Figure 3, we show the distribution of prices across hospitals for bare metal and drug eluting stents. The Figure displays the distribution of hospital fixed effects, which were obtained from a regression of prices on dummies for hospitals, products, and months; that is, we show here the distribution

⁵700,000 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.

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 (\$1,633 vs. \$752), and there is substantial dispersion in prices across hospitals within each category. The interquartile range of prices across hospitals is \$101 for bare metal stents, \$139 for drug-eluting stents. If all prices were brought down to the minimum price within each product category, hospitals would save 12-27% on average; if, instead, all hospital prices above the mean price within each product category were brought down to the mean, hospitals would save 2.4-3.6% on average.

Figure 3: Distribution of Prices Across Hospitals

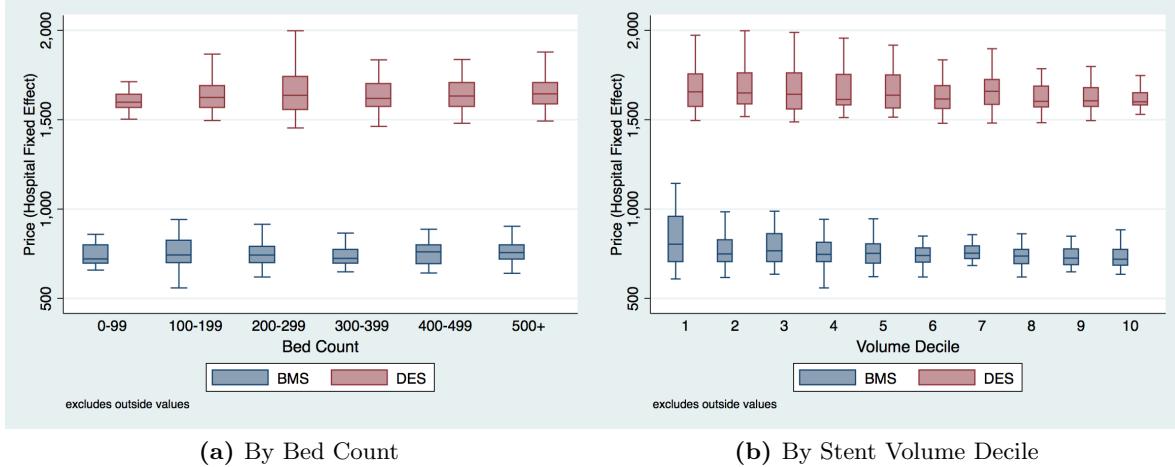


Interestingly, we are unable to associate much of this observed price dispersion with hospital characteristics that would seem a priori to be important for negotiation. For example, in spite of the fact that the largest hospitals spend triple the dollar amount on stent purchases as the smallest hospitals do, we observe no clear relationship between hospital size and stent prices. See Figure 4, in which we display a box plot of bare metal and drug-eluting stent prices for each category of bed count.⁶ 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 4, 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

⁶As before, “prices” are hospital fixed effects obtained from a regression of price on hospital, month, and product fixed effects.

hospitals' prices have larger means and variances than high-volume hospitals. For bare metal stents, 10th decile hospitals' prices are 13% lower than those obtained by 1st 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.

Figure 4: Distribution of Prices Across Hospitals



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 between 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 Information Intervention

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: first, upon logging in to the subscription database, they were presented with a basic interface showing analytics regarding their relative performance. If they so chose, they could see the same analytics for specific products and could click through to access the underlying data used to construct the analytics. By repeating this step for each SKU purchased, member hospitals could 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.

The basic interface members access upon logging in presents graphical analytics for “potential savings” on each product the given member purchased in the previous year. Savings potential is determined by comparing members’ prices to prices observed for all other members for the same exact product (SKU). These analytics are aggregated up to the supplier level in the main presentation, so that members can see how they are performing relative to peers separately for each manufacturer they have substantial purchasing arrangements with.

In order to analyze the effects of price transparency on negotiations, we obtained clickstream data on all members’ website interactions. The data allow us 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.

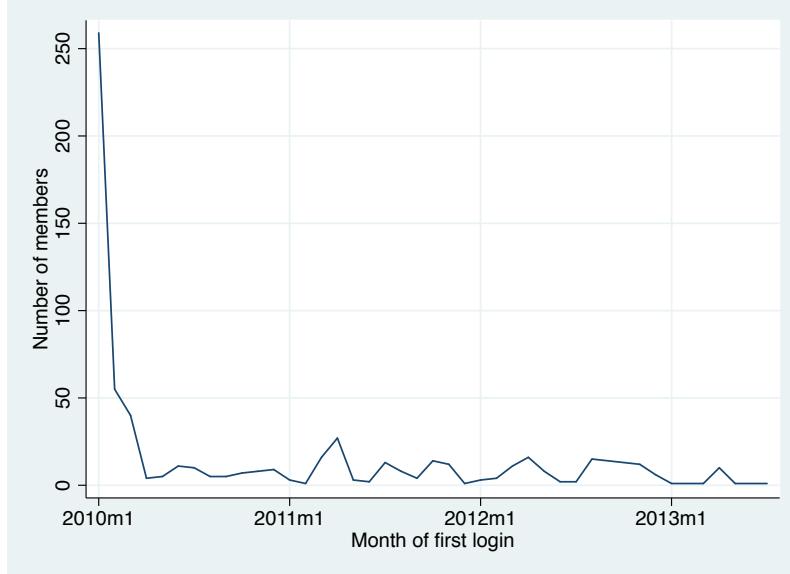
2.4 Identification of Information Effects

The ideal experiment that would enable us 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 this data is that the sample hospitals voluntarily joined a subscription database. The key feature of the data that allows us to estimate causal treatment effects of price transparency is that new members submit one year of retrospective data when they join the database, and continue to submit monthly data thereafter.

Thus, for hospitals that joined during the 2009-13 period, we observe data before and after they were first able to access the benchmarking information available in the database. Figure 5 shows the time series of hospitals joining the database between 2010 and 2013. One technical quirk of the data is that the database vendor rolled out a new version of its database in early 2010 and reinvited 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 “pre-period” data for those members from each of our analyses. After March 2010, seven hospitals join the database in each month, on average.

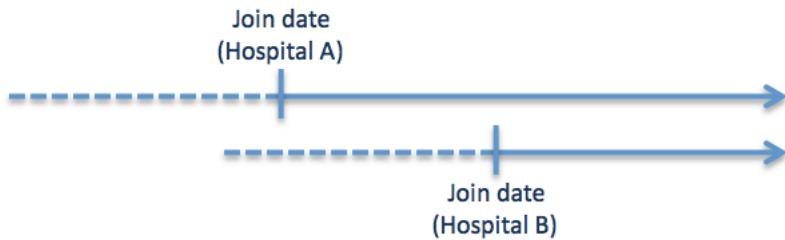
The availability of both pre- and post-join data for hospitals joining the database at different points in time allows us to use a differences-in-differences estimation strategy to estimate the treatment effect of having access to benchmarking information. The logic behind this identification strategy 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. We would be able to isolate the treatment effect of joining the database on prices by comparing the price trends between the two hospitals for their overlapping time periods. In this example, the pre-join price observed for hospital A cannot separately be identified from a common shock to prices

Figure 5: Count of Hospitals Joining in Each Month



in that period. However, if we assume parallel trends between the two hospitals absent any information effects, any difference in price trends in the overlapping periods (one period pre- and post- hospital B joining the database, both of which are post-join for hospital A) can be attributed to the effect of information. To implement this identification approach, we estimate a differences-in-differences model controlling for time-invariant hospital-product price patterns and secular trends in stent prices. The latter control is particularly important, as prices decreased substantially over time during our period of interest; if we omitted controls for time trends, we would estimate large effects of access to information based solely on the large price drops observed between hospitals' pre- and post-information periods, which are likely to be due in large part to time trends.

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



The primary concern with this identification strategy is that timing of hospitals' joining the database may be correlated with other factors going on at the hospital contemporaneously that impact price trends. For example, hospitals may be inspired to join the database due

to particular concerns about price trends, or (perhaps relatedly) they may concurrently be undertaking other initiatives intended to constrain prices, such as hiring new personnel or contracting other outside consulting services. These contemporaneous factors could bias our results upward or downward. To account for these potential issues, we also rely on an additional source of identification: new product entry. When a new product enters the market, initially no member will have data on other members' prices for that product. Benchmarking data will not become available for the new products until members submit their transactions data and they are loaded into the database (generally with a lag). New product entry thus gives us an additional point in time to observe the treatment effect of access to information.

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

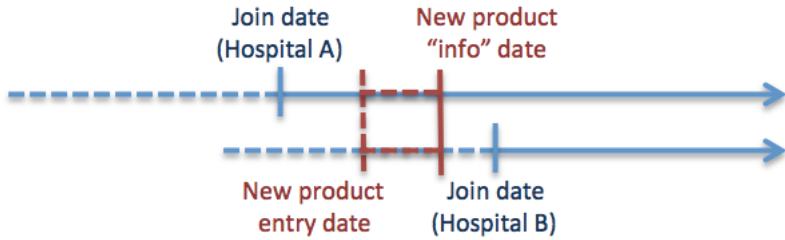


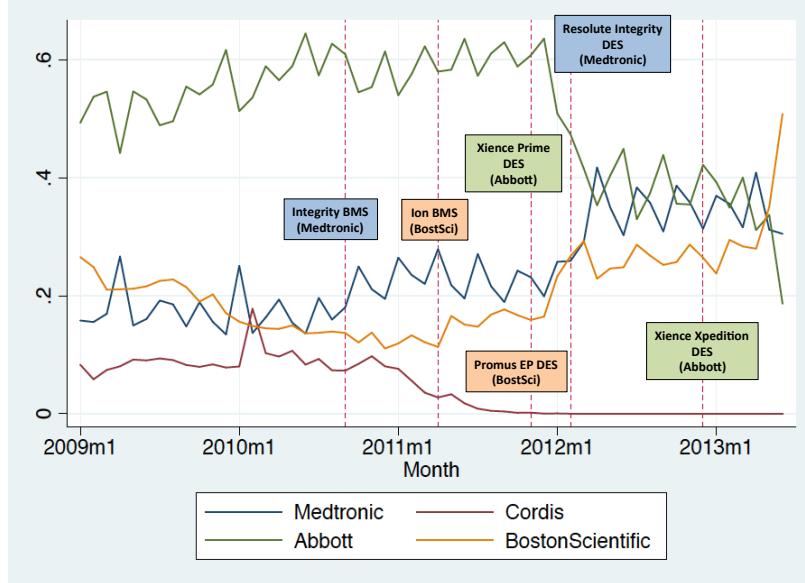
Figure 7 illustrates this identification strategy graphically. If timing of join is nonrandom based on the previously-described factors, such as other contemporaneous management changes, then we may question the differences-in-differences analysis above. However, if timing of join is plausibly random with respect to new product entry (and the effects of contemporaneous factors at the time of join are transitory), then we can identify the treatment effect of information by comparing pre-information and post-information prices to hospitals such as hospital A which are already in the database when a new product enters.

In Figure 8, we display the market share of each stent manufacturer over time, and note the timing of entry of six new products between 2010 and 2013. Assuming timing of join is not endogenous with respect to new product entry (an assumption supported by Figure 5, in which we do not see spikes of join around the timing of entry), we can examine price trends around the timing of product entry to help identify causal treatment effects and to investigate sources of bias.

3 Theory: Negotiated Prices and Using 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

Figure 8: Market Share of Each Manufacturer, Relative to New Product Entry



of benchmarking services offers one hope of doing so. There are two primary mechanisms that the literature and market participants 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). 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. 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 $\delta^B \in (0, 1)$ and the supplier has a discount factor δ^S .

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:

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

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 and negotiator agency.

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 $1 > \delta_s^S > \delta_w^S > 0$. 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: It can be shown 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):

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

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:

$$p^0 := c + \delta_w^S \frac{1 - \delta^B(1 - \omega_w) - \delta^B \omega_w}{1 - \delta^B(1 - \omega_w) - \delta^B \delta_w^S \omega_w} V, \quad (3)$$

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^0 , but that the strong seller strictly prefers:

$$p^1 := c + \frac{1 - \delta^B(1 - \omega_w) - \delta^B \omega_w}{1 - \delta^B(1 - \omega_w) - \delta^B \delta_w^S \omega_w} V, \quad (4)$$

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 $p_s^{CI} > p^1 > p^0 > p_w^{CI}$ (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.

Proposition 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.

Proposition 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 $\omega_w(p_s^0 - p_w^0)q$ exceeds the cost of information. This information will cause a proportion of the highest prices p_s^0 to fall to p_w^0 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.

Proposition 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. 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. In this case, the price will be a function of the 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 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 $v(b_h)$. 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. This motivates the following Proposition:

Proposition 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.⁷

Proposition 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.

⁷The 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.

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 *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:

Proposition 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.

Finally, new product introductions are also of interest because, at the time of new product entry, benchmarking data contain new information only on the specific entering product. Because product entry happens over time and is uncorrelated with the timing of hospitals joining the database, the average member will have already gained any informational advantage about previously available products prior to the new product's entry. This motivates yet another prediction:

Proposition 7 (New Product Entry Isolates Own-product Effects) For newly introduced products, any difference in prices between informed and uninformed buyers must be due to own-product effects, as no new information about previously existing products has been introduced.

3.4 “Stickiness” in Negotiated Prices

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.

4 Estimating How Information Affects Negotiated Prices

Our baseline research design estimates the average treatment effect of access to the database on negotiated prices, across all product-hospital combinations in the data. The estimate is based on the difference observed between pre-join and post-join prices, controlling for time-invariant price differences across hospital-product pairs and product-specific trends in stent prices. After the first quarter of 2010, the join date is relatively uniform over our sample period, allowing us to use variation in the timing of join to identify the price effect attributable to joining.

Letting P_{jht} denote the price observed for hospital h , product j , and month t , and controlling (as in all other subsequent specifications), for hospital-product fixed effects $[\theta_{jh}]$, month fixed effects $[\theta_t]$, and separate linear time trends for each product $[\gamma_{jt}]$, we estimate:

$$P_{jht} = \beta_{post} * \mathbb{1}_{\{post_{jht}\}} + \theta_{jh} + \theta_t + \gamma_{jt} * (t - t_{min}) + \varepsilon_{jht}$$

where the coefficient β_{post} is the treatment effect of access to price information in the benchmarking database. The variable of interest, $\mathbb{1}_{\{post_{jht}\}}$, equals one in all periods after the given member first logs into the database, zero otherwise.⁸ This is the price effect of simply having access to the database. It may underestimate the effect of information on price negotiation if there is a delay in price changes due to sticky contracts; we explore price effects conditional on renegotiation in subsequent analyses.

We estimate this specification on the sample of all hospitals in the purchase order data that are observed to purchase stents. As discussed in Section 2.4, we drop “pre-join” data for hospitals that we observe to first log in to the database in Q1 2010, given that the vast majority of those hospitals most likely joined prior to 2010 and were simply re-invited to the new system. This leaves us with a full sample of 349,613 transactions for 385 hospitals and 18 branded stents. We also perform the same regression focusing on the set of stents that

⁸For products that enter the market after a given hospital has joined the database, the member cannot access any benchmarking data on those products until the first round of transactions occurs, the transactions are submitted to the subscription database organization, and the database organization loads the data into the database. Thus, the “post” variable for entering products does not equal one until the given hospital logs into the database after benchmarking data for those products has been loaded in. See further discussion of the “post” variable for new products below for detail.

entered the market between 2010 and 2013. We do this for two reasons. First, as discussed in Section 2.4, we may be concerned that our baseline analysis, leveraging timing of join for identification, may be biased due to other factors occurring at hospitals contemporaneous with them joining the subscription database. As long as any differential price trends occurring at hospitals due to factors concurrent with them joining the subscription database are transitory, the results for entering products will not be biased by endogenous timing of join. Second, as discussed in Section 3, the moral hazard model predicts that prices will not be altered by access to benchmarking data for entering products, while the asymmetric information model predicts that they will. Even if hospitals alter their contracting relationships with purchasing agents when they first join the database, we do not expect stent prices for a given hospital to respond to other hospitals’ prices for newly entering stents in equilibrium. Thus, to the extent that we observe a change in prices for newly entering products after data are loaded in and hospitals access the database, this provides evidence of asymmetric information in bargaining. For our estimation on entering products only, we modify the specification to be the following:

$$P_{jht} = \beta_{post} * \mathbb{1}_{\{post_{jht,entry}\}} + \theta_{jh} + \theta_t + \gamma_{jt} * (t - t_{min}) + \varepsilon_{jht}$$

$$\mathbb{1}_{\{post_{jht,entry}\}} = \mathbb{1}_{\{t \geq \min\{t(\text{login}_h) | t(\text{login}_h) \geq X \text{ months after } j \text{ entry}\}\}}.$$

Limiting the sample to entering stents reduces the number of observations to 130,855. We set $\mathbb{1}_{\{post_{jht,entry}\}} = 0$ for the first X months after the product enters the market; in the below, we present results with $X \in \{2, 6\}$. Directly after product entry, data are entered into the database somewhat slowly. Two months after entry, at least ten hospitals’ prices for the entering product would be available for benchmarking in the database. Six months after entry, there is far greater price information available – hospitals would be able to observe prices for at least 70 other hospitals for the entering product by that point. After X months post-entry, $\mathbb{1}_{\{post_{jht,entry}\}} = 1$ after the given hospital logs in to the database. At this point, we remain agnostic about how many hospitals’ data are needed for a hospital logging in to believe they are observing actionable information, but it is likely that the six month cutoff would be more informative than the two month cutoff. To the extent that hospitals are in fact receiving actionable information about prices before the “post” variable equals one (that is, to the extent that they find a few hospitals’ data informative enough to motivate renegotiation), our results will be biased toward zero.

The results are shown in Table 3. The first row displays the estimated coefficients of interest for all products, entering products with $X = 2$, and entering products for $X = 6$.⁹ The specification is also estimated on a subsample where we only include observations within a one year window of the date of join; these results are displayed in the second row.

The estimates in the first column indicate that, across all hospital-product combinations,

⁹These and all subsequent regressions are estimated on prices in levels (dollars). Results are similar in direction and significance whether we estimate the regressions in logs or levels, so log results are omitted for brevity. The price distribution plots in Section 2 indicate that negotiated prices are slightly skewed, but do not appear log-normal.

Table 3: Differences-in-Differences Results: ATE of Information on Prices

	All Products	Entering Products (2 Mo. Buffer)	Entering Products (6 Mo. Buffer)
<i>Full Sample</i>			
β_{post}	-1.35 (4.66)	2.44 (4.80)	-5.63* (3.02)
N	349,613	130,855	130,855
Num. Hosps.	385	337	337
<i>One Year Pre-Post</i>			
β_{post}	1.32 (4.77)	0.56 (4.52)	-6.54** (3.03)
N	222,708	103,402	103,402
Num. Hosps.	360	332	332

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

simply logging into the database has zero effect on average prices, and standard errors are small enough that, in both samples, we can reject an ATE of -\$11 of joining the database on average prices. Among entering products, we see, if anything, slightly more negative effects; for buffer $X = 2$, the estimated effect of accessing the database is +2 in the full sample and +0.56 in the one-year pre-post analysis; neither estimate is significantly different from 0. For buffer $X = 6$, the effect is -\$5 to -\$7 and statistically significant. The fact that the results for entering products are slightly larger indicates that treatment effects of price information in the first column of Table 3 are, if anything, biased slightly toward zero due to contemporaneous factors at the timing of join.¹⁰ The small effect sizes on average and relatively large standard errors for the average treatment effects are not entirely surprising, given that theory suggests that effects would most likely be found for products where the hospital is paying an especially high price and products the hospital uses often, and that effects would potentially be biased towards zero due to price stickiness. The specifications that follow analyze each of these dimensions in turn.

4.1 Mechanisms: Where Does Information Matter Most?

We next consider differential transparency effects for specific hospital-product combinations that we predict ex ante to experience price changes upon the introduction of benchmarking. We test the predictions from Section 3 by allowing for heterogeneous treatment effects based on the position in the pre-transparency price distribution. We then further allow transparency 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. Finally, we control for any stickiness in prices by estimating solely on the sample of product-hospital-time observations for which there is a price change.

First, we estimate a model in which we allow for different treatment effects for products

¹⁰Results of this and subsequent specifications are not sensitive to the sample restriction to a one year window around the date of join. We omit these results in subsequent specifications for brevity.

and hospitals in different parts of the price distribution. Proposition 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. 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:

$$\begin{aligned}
P_{jht} &= \beta_{post} * \mathbb{1}_{\{post_{jht}\}} + \beta_{post}^{highY} * \mathbb{1}_{\{post_{jht}\}} * \mathbb{1}_{\{high_{jh,pre}^Y\}} \\
&\quad + \theta_{jh} + \theta_t + \gamma_{jt} * (t - t_{min}) + \varepsilon_{jht} \\
\mathbb{1}_{\{high_{jh,pre}^Y\}} &= \mathbb{1}_{P_{jh,pre} \geq \text{prctile}Y\{P_{jh',pre}\}_{h'=1}^H}
\end{aligned}$$

for various values of Y indicating different percentiles of the pre-join price distribution. That is, the coefficient β_{post}^{high25} will capture the differential treatment effect of transparency on prices for product-hospitals which were above the 25th percentile of the price distribution in the year just prior to joining. The results in Table 4 show separate estimates for the 10th, 25th, 50th, 75th, and 90th percentiles. The left panel shows the results for all products; the right panel shows entering products only.

Table 4: Heterogeneity in Results: Treatment Effect of Price Information on Different Parts of the Price Distribution

	All Products					Entering Products (6 Month Buffer)						
	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
	$\beta[Post]$	-1.4 (4.7)	-3.9 (5.2)	1.7 (5.2)	5.3 (5.2)	6.8 (5.1)	3.1 (4.7)	-5.6 (3.0)	-11.3 (10.0)	-1.8 (5.7)	3.0 (4.1)	2.8 (3.4)
$\beta[Post*HighPrice]$		5.6 (6.4)	-1.6 (6.2)	-11.8 (8.6)	-34.9*** (9.5)	-39.6** (19.4)		7.2 (10.1)	-4.4 (6.3)	-16.0*** (5.8)	-36.7*** (9.4)	-82.2*** (16.6)
N	349,613		235,637		130,855			83,882				

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

Here, we observe striking evidence that the average treatment effects described above concealed substantial heterogeneity. The differential treatment effect for higher parts of the price distribution is monotonically increasing as we move to higher percentiles. Hospital-products above the 25th percentile pre-join have a differential treatment effect of -\$1.6 across all products and -\$4.4 across entering products, compared to -\$12 and -\$16 for the 50th percentile, -\$35 and -\$37 for the 75th percentile, and -\$40 and -\$82 for the 90th percentile. The results for the 75th and 90th percentiles are statistically significant at conventional levels. This evidence is consistent with Proposition 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). We also note that the presence of large treatment effects in the upper part of the price distribution for *entering* products is strong evidence that asymmetric information impacts prices in the absence of benchmarking – as described in Proposition 6, the moral hazard model does not predict that access to benchmarking data will lead to price changes for entering products.¹¹

Next, we estimate in which we interact our “post” variable, alone and interacted with flags for position in the pre-join price distribution, with a dummy equalling one for hospital-product combinations with high purchase volumes in the period prior to join. To implement this, we generate a “High Quantity” flag equalling one for hospital-products with monthly purchase volumes above the 75th percentile in the months prior to join. To the extent that search and renegotiation are costly, Proposition 2 predicts that benchmarking data will be sought and used most effectively for hospitals and products purchased in high quantities. The specification we estimate is

$$\begin{aligned}
P_{jht} &= \beta_{post} * \mathbb{1}_{\{post_{jht}\}} + \beta_{post,high^q} * \mathbb{1}_{\{post_{jht}\}} * \mathbb{1}_{\{high_{j_h}^q\}} + \beta_{post,high^Y} * \mathbb{1}_{\{post_{jht}\}} * \mathbb{1}_{\{high_{j_h}^Y\}} \\
&\quad + \beta_{post,high^Y,high^q} * \mathbb{1}_{\{post_{jht}\}} * \mathbb{1}_{\{high_{j_h}^Y\}} * \mathbb{1}_{\{high_{j_h}^q\}} + \theta_{jh} + \theta_t + \gamma_{jt} * (t - t_{min}) + \varepsilon_{jht} \\
\mathbb{1}_{\{high_{j_h}^Y\}} &= \mathbb{1}_{P_{jh,pre} \geq \text{prctile}Y\{P_{jh',pre}\}_{h'=1}^H} \\
\mathbb{1}_{\{high_{j_h}^q\}} &= \mathbb{1}_{Q_{jh,pre} \geq \text{prctile}75\{Q_{jh',pre}\}_{h'=1}^H}
\end{aligned}$$

where $\mathbb{1}_{\{high_{j_h,pre}\}}$ is a dummy for $\{jh\}$ being a high-quantity hospital-product (here, above the 75th percentile) based on pre-join purchasing. The results for the full sample and for entering products only are shown in Table 5.¹²

Table 5: Heterogeneity in Results: Treatment Effect of Price Information on Different Parts of the Price Distribution and Different Purchase Quantities

	All Products					Entering Products (6 Month Buffer)						
	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	9.2 (5.7)	3.5 (6.7)	11.6 (7.4)	12.4* (6.9)	11.4** (5.7)	12.1** (5.8)	-5.8 (3.6)	7.9 (7.3)	0.3 (6.2)	2.6 (4.2)	1.7 (4.0)
Post*High Q	-5.2 (6.2)	-7.7 (9.5)	-6.8 (8.7)	-2.8 (7.1)	2.4 (6.7)	-4.1 (5.8)	1.5 (5.8)	-31.8** (15.5)	-3.5 (9.8)	0.3 (5.7)	1.5 (4.9)	4.0 (4.6)
Post*High P	8.2 (7.0)	-2.8 (7.4)	-6.7 (10.0)	-9.1 (10.9)	-27.4* (14.8)			-15.2* (7.8)	-7.9 (7.1)	-15.4** (6.7)	-28.1*** (10.8)	-54.9*** (20.8)
Post*High Q*High P	1.8 (12.4)	1.8 (11.3)	-6.5 (11.7)	-42.7** (18.9)	-31.0 (30.4)			38.2** (17.1)	6.1 (12.0)	-1.1 (11.7)	-21.1 (19.4)	-68.5** (30.6)
N	135,243		129,813		73,722			68,391				

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

¹¹The results are similar when we look separately at drug-eluting and bare metal stents; see Appendix Table 7. The strongest difference between the two sets of results is that the drug-eluting stents estimates are more precise, as would be expected given their relative popularity in the sample period.

¹²This analysis requires that we use pre-join data to generate the high-quantity flag – accordingly, hospitals “joining” in Q1 2010 are omitted and the regression sample is smaller.

The estimates show that the price treatment effect is largest for high-volume hospital-products in the upper part of the price distribution. The triple interaction term “Post * High Quantity * High Price” varies from -\$20 to -\$70 and is statistically significant for the 75th percentile across all products, and for the 90th percentile across all entering products. 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. See Appendix Table 8 for results of regressions that allow for different treatment effects for each quintile of “potential savings” based on what the member would have seen for each product upon first logging into the database; as expected given the above results, treatment effects are highest where potential savings are highest. For the first quintile of potential savings, the differential treatment effect ranges from -\$2 to \$7 (and is not statistically different from zero); for the fifth quintile, the differential treatment effect is -\$30 to -\$40.

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.

4.2 Contract “Stickiness” and Renegotiation

As noted above, we are concerned that contracts may be “sticky” and hospitals may accordingly be constrained in their ability to use benchmarking data effectively due to delays in renegotiation. If so, our above estimates will be biased toward zero due to the “post” period containing several months of data for which the prices are effectively fixed. In order to investigate this possibility, we next estimate specifications wherein we only include months in which a renegotiation takes place – we thus identify our treatment effects based on whether or not renegotiations in the post-join period result in lower prices than renegotiations in the pre-join period. This reduces our sample substantially, as we only have up to twelve months of pre-join data.

Table 6: Treatment Effect of Price Information on Different Parts of the Price Distribution, Conditional on Renegotiation

	All Products					All Products, Renegotiation Months Only						
	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	-1.4 (4.7)	-3.9 (5.2)	1.7 (5.2)	5.3 (5.2)	6.8 (4.7)	3.1 (13.8)	-6.7 (32.8)	53.5 (28.5)	22.1 (18.4)	-8.4 (15.0)	-0.6 (13.9)	-4.7
Post*High Price	5.6 (6.4)	-1.6 (6.2)	-11.8 (8.6)	-34.9*** (9.5)	-39.6** (19.4)		-61.8* (31.7)	-36.4 (30.9)	4.9 (22.4)	-39.5 (24.0)	-100.9*** (29.0)	
N	349,613		235,637			30,702			21,749			

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

The results are shown in Table 6. The left panel shows our original results from Table 4; the right panel shows the same results conditional on renegotiation taking place. Each panel shows results for all products (entering and other). Due to the sample reduction, we have limited ability to compare the results with and without renegotiation; however, we do note that the result for the 90th percentile product-hospitals is significantly larger once we condition on renegotiation. This result indicates that stickiness in contracts limits the benefits of benchmarking, at least for some products and hospitals.

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 2009-13 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.

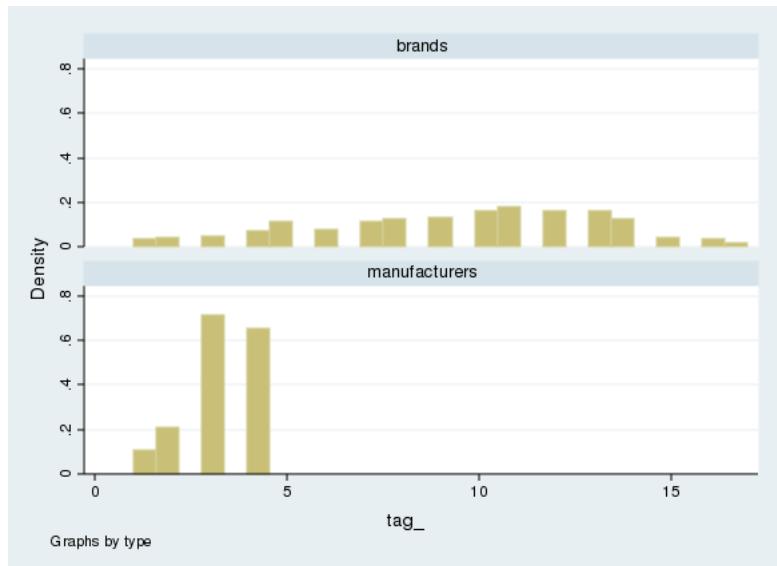
Appendices

A Appendix

A.1 Checks for Standardization and Share-based Contracts

In Figure 9, 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 9: 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 10 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 10: Market Share of Most-Preferred Brand, by Quintile of Purchase Volume

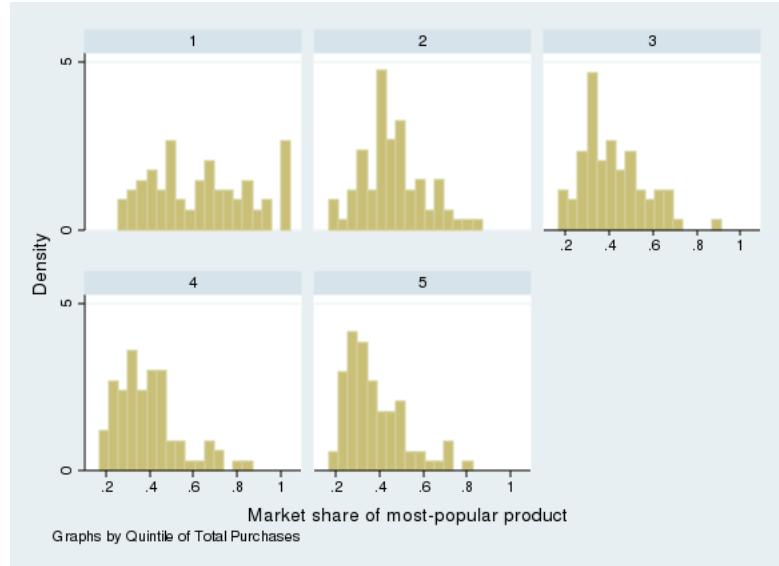


Table 7: Treatment Effect of Price Information on Different Parts of the Price Distribution: Bare Metal vs. Drug-Eluting Stents

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

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

Table 8: Heterogeneity in Results: Treatment Effect of Price Information for Different Quintiles of “Potential Savings”

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

Standard errors clustered at hospital 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 et al., 1996] Albek, S., Mollgaard, P., and Overgaard, P. (1996). Law-assisted collusion? the transparency principle in the danish competition act. *European Competition Law Review*.

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

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

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

[Gowrisankaran et al., 2014] Gowrisankaran, G., Nevo, A., and Town, R. J. (2014). Mergers when prices are negotiated: Evidence from the hospital industry. forthcoming in the American Economic Review.

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

[Grennan, 2014] Grennan, M. (2014). Bargaining ability and competitive advantage: Empirical evidence from medical devices. forthcoming in Management Science.

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

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

[Jin and Leslie, 2003] 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.

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