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

Matthew Grennan

University of Pennsylvania and National Bureau of Economic Research

Ashley Swanson

Columbia University and National Bureau of Economic Research

Using data on hospitals' purchases across a large number of important product categories, we find that access to information on purchasing by peer hospitals leads to reductions in the prices hospitals negotiate for supplies. These effects are concentrated among hospitals previously paying relatively high prices for brands purchased in large volumes. Evidence from coronary stents suggests that transparency allows hospitals to resolve asymmetric information problems, but savings are limited in part by the stickiness of contracts in business-to-business settings. Savings are largest for physician preference items, where high-price, high-quantity hospital-brand combinations average 3.9% savings, versus 1.6% for commodities.

I. Introduction

Business-to-business markets make up a large part of the economy, but they often lack transparency. Suppliers negotiate different contracts with

The data used in this paper were generously provided, in part, by the ECRI Institute (www.ecri.org). We are grateful to our editors, three helpful reviewers, Allan Collard-Wexler, Michael Dickstein, Glenn Ellison, Robin Lee, Amanda Starc, Robert Town, Ali Yurukoglu, and numerous seminar participants for helpful comments. We gratefully acknowledge financial support from the Leonard Davis Institute, the Wharton Dean's Research Fund and Public

Electronically published March 4, 2020
[*Journal of Political Economy*, 2020, vol. 128, no. 4]

different buyers, potentially with widely varying prices, and a buyer typically has limited information regarding other buyers' contracts. Many business-to-business markets have seen the entry of information intermediaries who facilitate buyers' ability to benchmark the prices they negotiate.[1] In this paper, we empirically examine the effect of transparency in the form of benchmarking information on prices negotiated by hospitals and their suppliers.

There is substantial variation in the prices of hospital supplies, including medical devices, across hospitals. For the top hospital supplies in our data, the average coefficient of variation across hospitals for the same exact brand-month is 0.18. This input price variation is approximately half the coefficient of variation for common procedure prices charged by hospitals in different markets (Cooper et al. 2019). It is also near the top of the range of coefficients of variation found in consumer goods markets (Scholten and Smith 2002).[2] In the short run, these costs come directly from hospitals' profit margins.[3] In the longer run, increasing supply costs tend to increase health care costs (see, e.g., Maeda, Raetzman, and Friedman 2012).

Prior research in consumer goods markets has largely confirmed the intuition that information can facilitate search and better decisions for buyers with imperfect information regarding product quality or cost (Sorensen 2000; Jin and Leslie 2003; Hendricks, Sorensen, and Wiseman 2012; Bronnenberg et al. 2015) or supplier willingness to accept lower prices (Zettelmeyer, Scott Morton, and Silva-Risso 2006; Scott Morton, Silva-Risso, and Zettelmeyer 2011). However, the mechanisms via which information affects consumer goods may not extend to business-to-business markets, where there is often no search across sellers (when products are purchased directly from manufacturers) and negotiators are professionals employed by firms and thus have different expertise

Policy Initiative, the Wolpow Family, and National Science Foundation grant SES-1559485. Biruk Bekele, Stuart Craig, Gi Heung Kim, Donato Onorato, and Mihir Trivedi provided excellent research assistance. Any errors are our own. Replication instructions are provided as supplementary material online.

[1] This trend is due in part to technology improvements that have made data easier to collect, distribute, and analyze. In addition to the hospital purchasing context, we are aware of business-to-business "transparency" benchmarking services emerging in areas as diverse as home appliances and television advertising.

[2] See also Kaplan and Menzio (2015) regarding price variation for consumer goods across stores.

[3] According to the American Hospital Association 2018 "Trendwatch Chartbook," the average hospital operating margin in 1995–2016 was 4.4% (https://www.aha.org/system/files/2018-05/2018-chartbook-table-4-1.pdf). Craig, Grennan, and Swanson (2018) calculate that the supplies in our data represent 26% of hospital operating costs.

and incentives than the typical consumer.[4] Recent empirical research in business-to-business bargaining (Draganksa, Klapper, and Villas-Boas 2009; Crawford and Yurukoglu 2012; Grennan 2013, 2014; Gowrisankaran, Nevo, and Town 2015; Lewis and Pflum 2015; Ho and Lee 2017) explains variation in prices across buyers using full-information models but in doing so also documents substantial heterogeneity in bargaining-ability parameters, which could include variation in information available to negotiators.[5] Our work contributes to these literatures by extending our understanding of transparency to the business-to-business setting and by offering information as one explanation for the large unexplained heterogeneity documented in negotiated prices.

Our empirical analysis utilizes new data containing all purchase orders issued by more than 17% of US hospitals that subscribed to a web-based benchmarking database. Section II details the data and hospital purchasing context. The majority of our analysis focuses on price negotiations for coronary stents in 508 facilities with cardiac catheterization services.[6] We also estimate our main specification with data on 52 product categories that are important in terms of total spending or utilization.

Motivated by discussions with industry experts, section III outlines two candidate mechanisms through which benchmarking information might have an impact in the hospital purchasing context: (1) an asymmetric information model in which hospitals face uncertainty about suppliers' costs or bargaining parameters, so that transparency reduces uncertainty and the

[4] The theoretical and empirical literatures on information disclosure are large and are reviewed in Dranove and Jin (2010). Recent studies from a variety of contexts, including consumer health care spending (Lieber 2017) and choice of college major (Hastings, Neilson, and Zimmerman 2015), have found the effects of information to be increasing in prior uncertainty but attenuated by frictions in consumers' ability and incentives to put information to use. In our context of professionals, one might expect prior uncertainty to be lower but ability and incentives to use information to be higher than in the consumer context.

[5] There are two exceptions of which we are aware. Larsen (2018) estimates a bargaining game of two-sided incomplete information about valuations in the used-car wholesale market; we argue that uncertainty over bargaining parameters better fits our context. Backus, Blake, and Tadelis (2019) study cheap talk in bargaining with asymmetric information for collectibles on eBay, but in our case price seems to dominate all other concerns for negotiators, decreasing the scope for the trade-offs involved in signaling that they document.

[6] Stents are especially useful as a place to focus for several reasons. They are representative of the "physician preference items" (PPIs) central to many policy discussions—PPIs are products where doctors' usage decisions are relatively insensitive to price, making negotiating lower prices the main mechanism via which a hospital can obtain savings. Stents are also important in their own right, constituting 2% of sample-hospital spending. Finally, stents typically have simple linear contracts, so the price we observe is the contracted price. As discussed in apps. C and F (apps. A–F are available online), we observe no evidence of standardization (e.g., exclusive dealing or contracts based on market share) affecting prices or usage in our data, and we find no effect of benchmarking information on quantities. This motivates our focus on the effects of information on linear prices (rather than quantities) in the remainder of the paper.

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.

Section IV clarifies our research design for testing empirical predictions from these models. We use two sources of variation to provide plausibly causal identification. First, the database is generated by monthly submissions of member hospitals' transactions, and new members are asked when they join to submit 12 months of retrospective, prebenchmarking data. We use variation in the timing of hospitals' joining the database to construct difference-in-differences estimators. Exogeneity of join timing with respect to price trends is supported by the institutional details of the setting and by event studies that show no statistically significant divergence of pretrends for coronary stents.

Second, we develop a set of tests focusing on new brands entering the market during our sample period. We compare prices between hospitals before and after join immediately upon a brand's introduction, before either hospital type has access to information, in order to remove any persistent sources of bias around join timing. New-brand introductions also offer a strategy to investigate theoretical mechanisms. The asymmetric information mechanism wherein hospitals use benchmarking information to learn about suppliers relies on concurrent availability of data on others' prices, but the agency mechanism wherein hospitals use benchmarking information to better incentivize their purchasing negotiators relies on the fact that such information will be available in the future. Thus, the delayed release of benchmarking data after new-brand entry allows us to separate the two.

Section V presents our main results. Focusing on stents, we find that access to the database information has heterogeneous effects across hospitals and brands. The average treatment effect of benchmarking information across all hospital-brand-months is small and noisy—we estimate no significant price changes for hospital-brands with low to moderate prices before join—but high-price hospital-brand combinations exhibit unit price declines of $27 upon accessing of database information. The price declines are larger for hospital-brands with larger purchase volumes: $71 for high-quantity hospital-brands that were also high in price before the hospital joined the database, compared to $17 for hospital-brands with lower volumes. The specifications leveraging brand entry suggest that these effects are largely explained by a mechanism wherein benchmarking solves an asymmetric information problem, helping hospitals learn about suppliers. Evidence for the agency mechanism is noisier and less robust across specifications.

Each of the above-described results is an estimate of the treatment effect of benchmarking on prices, which will be an underestimate of the treatment effect of benchmarking on prices negotiated in a given contract. We find that price effects are generated by increasing the likelihood of renegotiation and by generating larger price decreases conditional on renegotiation.

Taken together, our results suggest that transparency can lead to significant savings but that these savings are dampened by the stickiness of contracts and other costs of putting information to use in business-to-business settings. These findings contribute to a growing body of empirical bargaining studies showing evidence that "bargaining costs" play an important role in real-world negotiations (Backus et al. 2018; Jindal and Newberry 2018; Shelegia and Sherman 2018).

In section VI, we extend our difference-in-differences analysis to consider the effects of access to benchmarking information for a wide variety of important product categories, from commodities (e.g., surgical gloves) to other PPIs beyond stents (e.g., prosthetic hips). We find that the effects of information are broadly similar across product categories. Average treatment effects are often negative but tend to be small and statistically insignificant. However, among hospital-brands in the top quartile of quantity and quintile of price at the time of join, treatment effects tend to be larger in magnitude and more often statistically significant. The largest effects we document are for the purchase of expensive PPIs, where our treatment-effect estimates indicate 3%–4% savings due to benchmarking information among hospitals formerly paying the highest prices.[7]

These estimates are of direct interest for considering the impact of information intermediaries on the prices buyers negotiate in previously opaque business-to-business markets. They also provide the first empirical evidence on the potential effects of transparency policies that have been proposed for medical technology markets, though we caution that benchmarking, as implemented in our sample, may have different effects than a broad policy change for many reasons. In section VII, we conclude and also discuss potential directions for future research on information in business-to-business bargaining.

II. Data

The primary data used in this study come from a unique database of all supply purchases made by approximately 17% of US hospitals that joined a price benchmarking service during the period 2009–14. The data are from the PriceGuide benchmarking service (hereafter "PriceGuide data")

[7] For some product categories, we also find evidence that prices increase in the bottom of the distribution, consistent with a supplier response to an information externality across buyers in the asymmetric information framework. However, this evidence is sensitive to empirical specification, whereas the evidence of decreases in the top of the price distribution is quite robust.

offered by the ECRI Institute, a nonprofit health care research organization. We observe unique, anonymous hospital identifiers and several hospital characteristics: census region, facility type, and number of beds. For each transaction, we observe price, quantity, transaction month, and supplier for a wide range of product categories, including commodities (e.g., cotton swabs and gloves) as well as PPIs (e.g., stents and orthopedic implants).[8]

Our analyses consider price negotiations between hospitals and suppliers for a large number of important product categories. The contracting environment is described in detail in appendix C. Included products were the top 50 product categories by either total spending or transactions, for a total of 52.[9] As detailed in appendix A, the data are at the stock-keeping-unit (SKU) level, requiring us to use manufacturer catalogs and classification algorithms to group SKUs that belong to the same manufacturer-brand.[10]

Table A5 (tables A1–A14 are available online) summarizes the data for the 52 product categories of interest. Spending per month varies dramatically across product class: hospitals typically spend only $11,000 per year on bandages, versus over $1.1 million on drug-eluting coronary stents. Some product categories are used fairly universally—618 sample hospitals purchased hypodermic injection needles—and some are used only in highly specialized facilities—only 249 sample hospitals purchased biological cardiac valve prostheses. As discussed in greater detail in appendix A.3.1, we document price dispersion for commodities similar to that for other medical/surgical products and PPIs—coefficients of variation within brand-time across hospitals are 0.195, 0.166, and 0.188, respectively. This suggests that opportunities for savings are similar, relative to preinformation prices, across product classes.

Our sample facilities are discussed in detail in appendix A.3. Overall, our regression samples include 775 facilities spending an average of $1.8 million per month across 774 product categories. We return to the full set of important products in section VI, but our main analyses focus

[8] The reported data are of high quality because they are typically transmitted as a direct extract from a hospital's materials-management database. Hospitals have strong incentives to report accurately because the analytics the benchmarking service's web portal provides are based on comparing the hospital's submitted data to those of others in the database.

[9] There are 80 "top" categories total, but we omit product categories that are too broad or with missing or inconsistent data. See app. A for details.

[10] Note that we use the term "brand" to refer to the "product" level at which prices are negotiated—e.g., Medtronic Resolute Integrity drug-eluting coronary stent. The use of "brand" is not meant to connote any particular marketing strategy. We use "product category" to refer to the Universal Medical Device Nomenclature System (UMDNS) code grouping included in the transaction files. The UMDNS system classifies devices by intended purpose and mechanism of action (e.g., drug-eluting coronary stents have UMDNS code 20383). Finally, we use "product class" to refer to broad groupings of product categories: commodities, PPIs, and other medical/surgical products.

on 508 sample facilities that purchase drug-eluting coronary stents.[11] Stent purchasing patterns are presented in table A1. Briefly, the average sample facility spends \$3.4 million per month on 1,143 categories of supplies, \$73,000 of which is dedicated to drug-eluting coronary stents. During our sample period 2009–14, there were 13 branded products sold by four manufacturers—Abbott, Cordis, Medtronic, and Boston Scientific, with Cordis exiting the market in 2011. The average hospital purchases 48 drug-eluting stents per month.

The hospitals in the purchase-order data voluntarily joined a subscription service that allows them to benchmark their own prices and quantities against those of other hospitals in the database. We consider the effect of potential nonrepresentativeness of our sample in our discussion of identification in section IV. In appendix A.1, we also compare our sample hospitals along observable dimensions to two outside samples of US hospitals—all American Hospital Association member hospitals with cardiac catheterization and Millennium Research Group's (MRG) Markettrack survey of a geographically representative sample of catheter labs. The West region is overrepresented in our benchmarking database sample, while the South is underrepresented. The average PriceGuide hospital is larger than the average US hospital with cardiac catheterization. Similarly, member facilities in our PriceGuide estimation sample purchased in higher volumes and obtained slightly lower prices than MRG hospitals in overlapping periods. Finally, the analyses in appendix D.2 show that our main findings are qualitatively and quantitatively similar when we perform our analyses within hospitals in each census region or within hospitals in the same bed size range.

A. *Price Variation across Hospitals and Brands*

Figure 1*A* displays the distributions of drug-eluting stent prices across hospitals and hospital-brands. It illustrates the wide variation in prices across PriceGuide sample hospitals in their prebenchmarking transactions, with a standard deviation of \$164 and a mean of \$1,615, for a coefficient of variation of 0.10. Hospital-product effects explain much of this variation, with $R^2 = 0.89$ for the residual price variation (after brand-month detrending). Hospital effects, in turn, explain almost half of the hospital-brand variation, with $R^2 = 0.44$. Thus, our price variation is driven in part by some hospitals consistently paying more than others for all drug-eluting stents and in equal part by some hospitals paying more for particular stents.

[11] The database includes a few other health care facilities, such as clinics and surgical centers, but hospitals constitute 97% of the stent sample and 73% of the extended sample by count and more by spending. For the remainder of the paper, we use the terms "hospital" and "facility" interchangeably.

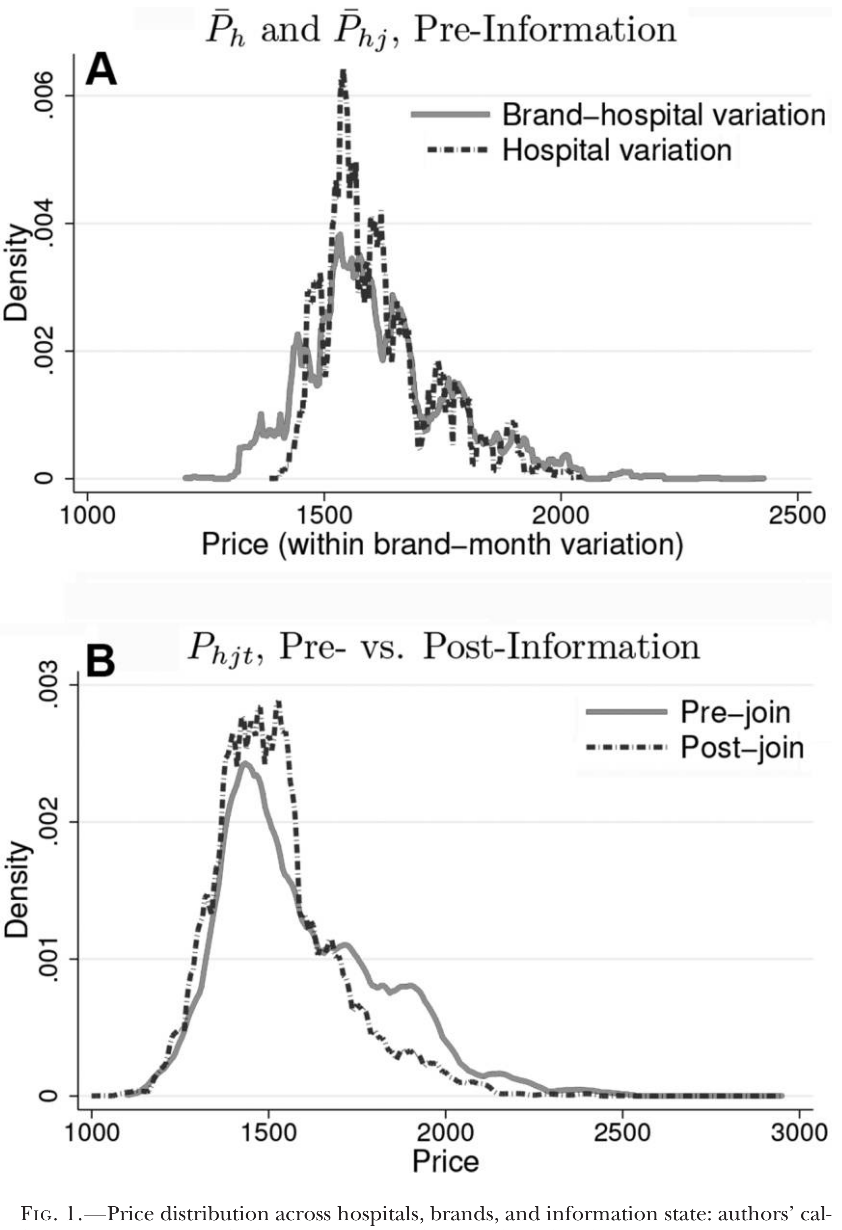

FIG. 1.—Price distribution across hospitals, brands, and information state: authors' calculations from PriceGuide data. *A*, Estimated hospital-brand and hospital fixed effects, obtained from regressions of price per stent, controlling for brand-month fixed effects. *B*, Raw price distributions before and after access to benchmarking information. A color version of this figure is available as an online enhancement.

As shown in appendix A.1, these patterns are shared by the representative MRG sample, though the MRG sample has slightly higher prices on average. This implies that hospitals joining the database have slightly less to gain than the representative sample in terms of raw price differentials.

Observable hospital characteristics do not explain much of the variation in prices. Hospital bed size has no explanatory power. Total volume of stents purchased has more explanatory power: tenth-decile hospitals by purchase volume (188 stents per month) achieve prices that are 7% lower than those obtained by first-decile hospitals (7 stents per month). However, we observe substantial dispersion in prices, even conditional on facility size and purchase volume (see app. A.1 for details).

In a different data set, Grennan (2013, 2014) found evidence that heterogeneity in stent prices across hospitals could be explained in part by heterogeneity in physician brand loyalty, but this left a large residual heterogeneity in hospital-product bargaining ability.[12] Our analysis explores the possibility that part of this heterogeneity in bargaining abilities may be due to heterogeneity in information among hospitals and that transparency in the form of benchmarking information on other hospitals' prices might affect this.

B. The Benchmarking Information Treatment

The information treatment considered in this study is one in which hospitals log in to a database and receive information about their relative performance in purchasing. The basic interface members access upon logging in presents graphical analytics for "potential savings" opportunities at the supplier level, defined as the total dollars that might have been saved in the previous year based on the hospital's volume of purchase and the mean/minimum prices paid by other hospitals at the manufacturer-SKU level. By clicking through, the hospital could observe these potential savings broken down by SKU, filter by geography and bed size, or even access the full (deidentified) purchase-order microdata. We obtained clickstream data on the timing of all members' website log-ins, allowing us to reconstruct each member's benchmarking information set at each point in time.

In order to preview our approach and results in a simple graphical manner, figure 1*B* displays the histograms of prices paid for drug-eluting stents across the entire sample, splitting the sample into pre- and postjoin observations. The raw data clearly suggest the primary impact of access to

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

the benchmarking information: hospitals with information are much less likely to pay the highest prices. In the sections that follow, we consider what theoretical mechanisms might drive this result in business-to-business negotiations as well as the research designs and regression specifications that will allow us to establish causal treatment effects and the mechanisms behind them.[13]

III. Theory: Bargaining and Benchmarking Information

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

Our models build on the Rubinstein (1982) model of alternating-offers bargaining.[14] The model has a single buyer negotiating with a single supplier over a per-unit surplus $V = \mathrm{WTP} - c$ 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.[15] Beginning

[13] We focus on the potential effect of information on negotiated prices. In app. F, we also estimate the effects of information on quantities and find no effect, consistent with stents being PPIs, where physician demand is based on strong preferences and is relatively insensitive to price.

[14] This model underpins a large theoretical literature on bargaining (Rubinstein 1985; Binmore, Rubinstein, and Wolinsky 1986; Horn and Wolinsky 1988; Collard-Wexler, Gowrisankaran, and Lee 2019) as well as a more recent empirical literature on bargaining (Draganksa, Klapper, and Villas-Boas 2009; Crawford and Yurukoglu 2012; Grennan 2013, 2014; Gowrisankaran, Nevo, and Town 2015; Lewis and Pflum 2015; Ho and Lee 2017). The predictions of the model extend well to empirical settings because the "discount factors" that parameterize bargaining strength in the Rubinstein model can be thought of more generally as proxies for a host of factors that might affect a real-world negotiation, such as impatience, opportunity costs of time, laziness, or fear of negotiation breakdown.

[15] Here, V_{hjt} (subscripts suppressed in text) 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 nonstent treatments, given physician preferences over all stents available at time t. For the sake of parsimony, we abstract here from price externalities across negotiations. Appendix D.3 provides a prediction regarding such externalities and an empirical test of that prediction. Fully modeling a multilateral contract equilibrium (e.g., as in Collard-Wexler, Gowrisankaran, and Lee 2019) with the addition of asymmetric information is beyond the scope of this paper.

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 with discount factors $\delta^{B} \in (0, 1)$ for the buyer and $\delta^{S} \in (0, 1)$ for the supplier. In the institutional setting of bargaining over medical devices such as stents, the typical negotiation occurs between a purchasing agent of the hospital and a sales representative of the device manufacturer. Each discount factor should be thought of as coming from a combination of negotiator skill and the incentives negotiators face as agents of their respective employers.

The unique subgame-perfect equilibrium of this game is for it to end immediately, with the buyer making an offer and seller accepting. The resulting complete-information (CI) price is $p^{CI} := c + \delta^{S}[(1 - \delta^{B})/(1-\delta^{B}\delta^{S})]V$. Thus, the observed variation in prices in our data could be generated in a full-information model by wide heterogeneity in discount factors and valuations across buyer-supplier pairs, in which case there may be no effect of benchmarking information.

A. *Asymmetric Information about Bargaining Parameters*

In order to introduce asymmetric information into the baseline bargaining framework, we follow Rubinstein (1985), in which hospital buyers have uncertainty about the bargaining parameter of a given supplier. The model departs from the CI model outlined above in that the supplier is either the weak type, with discount factor δ^{S}_{w}, or the strong type, with discount factor $\delta^{S}_{s}(1 > \delta^{S}_{s} > \delta^{S}_{w} > 0)$. The supplier knows his own type, but the buyer has only a subjective prior ω of the probability that the supplier is the weak type.[16]

Rubinstein (1985) shows that, in this asymmetric information (AI) game, if the buyer is sufficiently pessimistic about the seller being the weak type, then there exists a pooling equilibrium wherein the buyer simply

[16] This model focuses on the case where uncertainty is embodied only in the discount factors and not the value over which negotiations occur, which is not directly testable without data on breakdown or beliefs because the surplus and bargaining parameters enter the price multiplicatively. However, this seems to fit the primary potential source of uncertainty in coronary-stent negotiations, where doctor preferences are typically quite well known by those involved in the negotiation and marginal costs are small relative to the surplus created. It is also consistent with anecdotal evidence of little if any equilibrium breakdown in negotiations or destruction of surplus, which are central predictions of models of incomplete information about values (thanks to Brad Larsen for this observation). See Ausubel, Cramton, and Deneckere (2002) for a review of the literature focused on informational asymmetries in values. Of particular note in that literature is Cramton (1992), which extends a model similar to the one here to a continuum of types and two-sided asymmetric information and where information revelation arises endogenously through the timing of initial offers.

offers what she would offer the strong type in a CI game, p_s^{CI}, and both seller types accept this offer. If the buyer is more optimistic, then there exists a separating equilibrium wherein the buyer offers a low price p_w^{AI}, which the weak seller type accepts. But the strong seller type will reject this offer and counteroffer with p_s^{AI} (where $p_s^{CI} > p_s^{AI} > p_w^{AI}$), which the buyer accepts.

For simplicity, we begin by assuming that access to benchmarking information fully reveals a seller's type. Several empirical predictions for the effects of information on negotiated prices follow directly.

Prediction 1 (Direct effect of complete information on high prices). If information is costless, pessimistic buyers will always become informed. This information will cause the highest prices p_s^{CI} to fall to the CI weak-supplier price p_w^{CI}, for those cases where the supplier was in fact the weak type.

Prediction 2 (Direct effect of complete information on high prices with high quantity). If information is costly to obtain (e.g., searching and analyzing the data take time that could be used on other productive activity), a pessimistic buyer will become informed whenever the expected benefit of information $\omega(p_s^{CI} - p_w^{CI})q$ exceeds the cost.

The above illustrative model and predictions assume that benchmarking information fully reveals the seller's type. However, the transparency introduced by price benchmarking could have a countervailing effect: sellers may be less willing to offer low prices if they know that those prices will then be included in the database and potentially hurt them in subsequent negotiations with other buyers (Duggan and Scott Morton 2006; Grennan 2013). To show how this effect might arise, appendix B provides a more detailed analysis of the baseline AI model, plus a simple extension of the model in which the existence of benchmarking may create externalities across buyers in a multilateral environment.

The extended model allows for benchmarking data to provide a noisy signal of a seller's type to future buyers by revealing the current buyer-seller pair's price. Via this potential information externality, the presence of a future negotiation effectively puts a lower bound on the price the seller would be willing to accept in today's negotiation, and thus some of a seller's lowest prices may increase. Appendix B.1.2 presents one parameterization of the model to illustrate the mechanism of interest. The extended analysis yields an additional prediction under a model of asymmetric information.

Prediction 3 (Effect of information externality across buyers on low prices). Via the benchmarking process, a negotiated price with one buyer becomes information available to other buyers, potentially introducing an externality in future negotiations. If this externality is large enough, it can provide the seller a credible threat not to accept low prices. Thus, some prices may rise with the introduction of benchmarking information, particularly among buyers with relatively low prices and a relatively low quantity under negotiation.

Put plainly, prediction 1 is that exposure to benchmarking information should lead to some of the highest prices falling (cases where the supplier was the weak type). Prediction 2 is that this effect will be more likely among those brands with the highest quantity used. Prediction 3 is that benchmarking may generate an externality such that some of the lowest prices increase, particularly among low-quantity buyers.[17]

B. *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. We expect this mechanism to be relevant in the cardiac unit context. McConnell et al. (2013) present survey data documenting that hospitals' cardiac units vary substantially in their focus on performance measurement, and the Centers for Medicare and Medicaid Services recently found that cardiac and orthopedic units in hospitals responded to bundled payments (which entail higher-powered financial incentives) by improving contracting with suppliers.[18]

Extending the model presented thus far, suppose that instead of the hospital negotiator's bargaining parameter being exogenous, the price will be a function of the hospital agent's choice of discount factor δ^{B}. Further, suppose that in addition to uncertainty as to whether the supplier is a strong type or a weak type, there is an additional independently and identically distributed shock to the supplier's bargaining parameter that is buyer specific (see app. B for details in the case where hospital h faces a supplier bargaining parameter equal to $\delta^{S}_{h} \in \{\delta^{S}_{w}\epsilon_h, \delta^{S}_{s}\epsilon_h\}$ for $\epsilon_h \sim U[0, 1]$). Supplier bargaining strength is then observable to the hospital negotiating agents but not to the principals who manage them.

A moral hazard problem arises in this setting because bargaining effort is costly and provides the agent disutility. Under the usual assumption that the agent is risk averse, the optimal employment contract involves risk sharing between the principal and the agent. Holmstrom (1982) shows that if agents face some common parameter that is uncertain from the principals' perspectives (here, the portion of the bargaining parameter δ^{S} that reflects whether the supplier is a strong or weak type), then relative

[17] Alternative models that could generate similar empirical predictions might include models wherein one party has preferences over relative as well as absolute performance. See, e.g., Card et al. (2012) regarding pay transparency, in which workers learning that they have relatively low salaries have reduced satisfaction and are more likely to leave their jobs.

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

performance evaluation compared to some aggregate sufficient statistic can be used to write a stronger incentive contract with each agent.

In our interviews with industry participants, we did not encounter a single case where purchasing-agent contracts were formal functions of a quantitative performance metric based on benchmarking information or otherwise.[19] However, we did learn that some hospitals use measured price decreases as part of a broader performance review or employee-recognition program. When benchmarking information was available, we also heard cases of such data being used to quantify relative performance and opportunities for savings. This is in keeping with the spirit of the model above, motivating 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) allows the principal to estimate whether the seller is the weak or strong type and thus reduces the risk to which the agent is exposed in a higher-powered contract. The higher-powered contract induces more bargaining effort and a lower price than the case where only the buyer's own price is observed.

Prediction 5 (Monitoring effect on prices with high quantity). Further, information will be used in this way when the expected benefit $\mathbb{E}[p^{\text{NoInfo}} - p^{\text{Info}}]q$ exceeds the cost of information use.

In sum, prediction 4 suggests that prices will decrease, on average, upon the introduction of benchmarking; prediction 5 suggests that this effect will be more likely when greater purchase quantities are at stake.

C. *New-Brand Entry and Timing of Information Effects*

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

This difference in the timing of information required between the two mechanisms is especially relevant when new brands enter the market. There will be no data in the benchmarking database on a brand for the first month or two it is on the market and few data for the first few quarters. Thus, those who engage in their first negotiation for a new brand soon after its release

[19] See app. C.2 for details.

do so without current benchmarking information, even if they have access to the database. This motivates our final theoretical prediction.

PREDICTION 6 (New-brand entry separates asymmetric information and agency). For newly introduced brands, 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.

Our empirical implementation of this idea identifies the effects of any contract redesign that negotiators are made aware of upon the firm joining the database and that incentivizes effort before benchmarking realizations. This structure, in which today's performance affects tomorrow's information and, accordingly, compensation, is the approach taken in most explicit pay-for-performance schemes in health care markets (see James 2012). In more general compensation schemes, relative performance evaluation can be part of employee compensation contracts with explicit bonuses (e.g., sales force compensation), or rewards can be focused on raises and promotion (see Lazear and Oyer 2013 for a review).

D. *Other Considerations*

In the interest of clearly illustrating the fundamental ideas behind the two theoretical mechanisms of interest, we have abstracted from some realities of hospital purchasing. Here we touch on some key features omitted from the model and how they affect the empirical analysis that is the focus of the paper.

First, 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 short-run effect of information toward zero. We consider these dynamics in our empirical analyses using event studies and direct examination of recontracting.

Second, the same supplier salesperson may be in charge of negotiating contracts for a bare-metal and a drug-eluting stent or for subsequent generations of a branded drug-eluting stent. 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.

Finally, 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), the model may omit supply-side responses that may negate or overturn these effects through greater obfuscation (Ellison and Ellison 2009) or by facilitating collusion (Albæk, Møllgaard, and Overgaard 1996; Cutler and Dafny 2011). We return to this issue, and the extent to which our empirical estimates may capture these various supplier responses, in our discussion of results.

IV. Identification of Information Treatment Effects

The key features of the data that allow us to estimate causal treatment effects of price transparency for the hospitals in our sample are (1) that new members submit one year of retrospective data when they first join the benchmarking database and continue to submit monthly data thereafter, (2) that new members join over time in a staggered (and seemingly random with respect to price trends) fashion, and (3) that new brands enter the market at points in time that are seemingly uncorrelated with members' information states.

For hospitals that joined during the 2009–14 period, we observe data before and after they were first able to access the benchmarking information available in the database. Figure 2*A* shows the time series of stent-purchasing hospitals joining the database between 2010 and 2014. Beginning in Q2 2010, 14 such hospitals joined the database in each quarter, on average.[20]

A. Using Join Date to Identify Price Effects

We leverage this variation by using a series of difference-in-differences strategies. In our sample, all hospitals, by definition, access the benchmarking data at some point. The "control hospitals" for analyzing the price trend of hospital h in a window around h's joining the database are those hospitals $k \in \mathcal{H} \backslash h$ that subscribe either before or after that window. Under the standard assumption of parallel trends, we can isolate the treatment effect of joining the database on prices by comparing the price trends between treatment and control hospitals for their overlapping time periods.

The primary concern with this identification strategy is that timing of a hospital joining the database may be correlated with other contemporaneous factors that affect price trends at that hospital.[21] However, there are several institutional features that one might expect to dampen potential

[20] This figure focuses on hospitals purchasing coronary stents. Over the same period, 36 hospitals purchasing any important product in our data joined in each quarter.

[21] For example, a hospital may be inspired to join the database because of concerns about upward price trends, which could induce a positive bias in our results—we would be underestimating the counterfactual prices joining hospitals would face if they did not join. On the other hand, a joining hospital might concurrently be undertaking other initiatives intended to reduce input prices, such as hiring new personnel or contracting other outside consulting services. Conflating the effects of these other initiatives with the effect of access to the benchmarking information could induce a negative bias in our results.

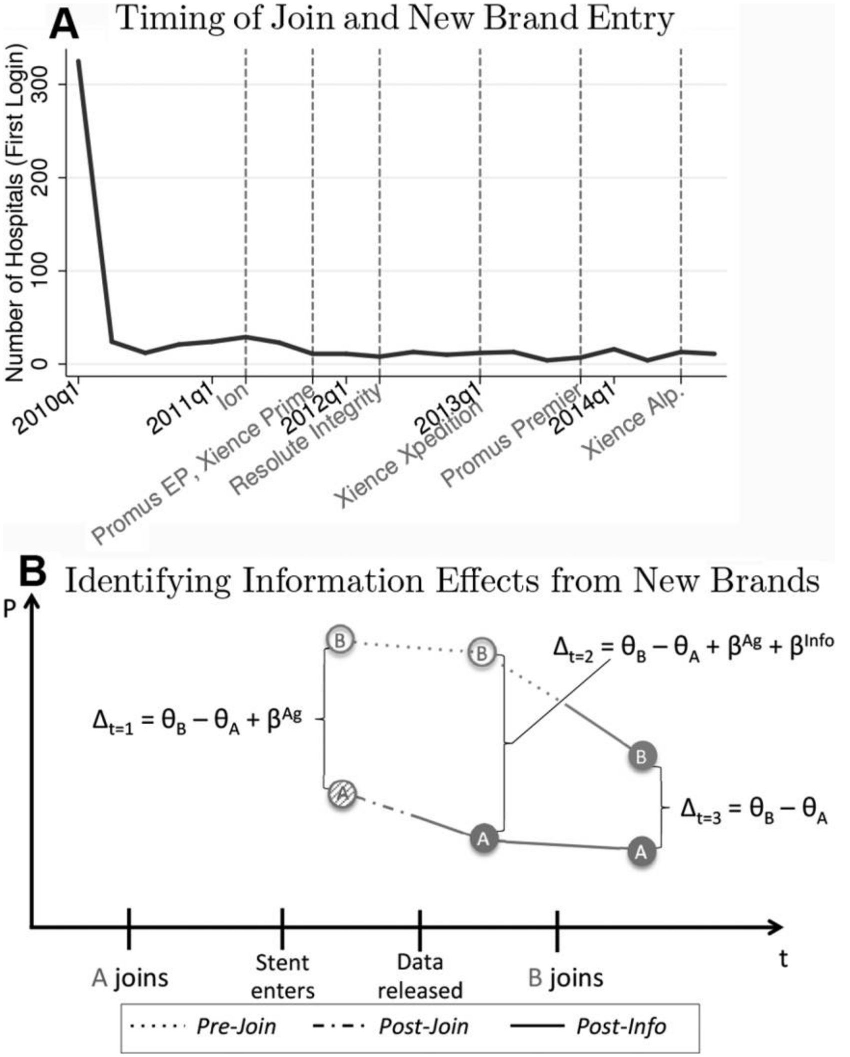

Fig. 2.—Identifying variation and graphical example of identification. *A*, Authors' calculations from PriceGuide data, 2010–14. "Join" is defined by member's first associated log-in. New-brand entry is indicated by vertical lines. PriceGuide rolled out a new version of its web interface in the beginning of 2010 and reinvited all current members to "join"; members' whose first associated log-in is in Q1 2010 may have subscribed in 2009 or earlier. All such members' "prejoin" data are excluded from empirical analyses. Alp. = Alpine. *B*, Graphical illustration of new-brand identification strategy. A color version of this figure is available as an online enhancement.

join-time bias for any particular focal brand. First, our conversations with industry participants indicate that it is unlikely for a hospital to join the database because of any single product category. PPIs such as stents are important purchases for hospitals, but subscription is costly and meant to cover a large number of product categories. Second, many determinants of price trends are specific to a product-category market, limiting the degree of correlation in price trends across, for example, coronary stents and knee prostheses. Finally, hospital purchasing tends to be separated across groups of product categories, implying that, for example, organizational changes regarding purchasing in the catheter lab need not correlate with changes in orthopedic surgery. Thus, any particular product category or brand within a category is likely to be a "bystander" to the join timing. This is consistent with the event studies in section V, which show no evidence of differential pretrends in price in the months before hospitals join the database.

B. New-Brand Entry, Mechanisms, and Bias

New-brand entry provides another opportunity to identify the above information effect and further allows us to identify a treatment effect of having joined the database but not yet having access to concurrent data on other hospitals' purchases. After new-brand entry, there is a lag before members may access benchmarking data on the new brand because of lags in data submission and loading: in the months following new-brand entry, the count of members purchasing that brand exceeds the count of members with transactions loaded in the database by 56, on average. Moreover, we observe transactions for new brands for some members before and after they join the benchmarking database: in the year following brand entry, 9% of members whose transactions are observed in the average month are before join (details in app. A.2). The time period for our study contains many drug-eluting stent brand introductions. In figure 2*A*, we note the timing of entry of seven new brands between 2010 and 2014, of the 13 brands sold during this time period overall.

This variation allows us to identify a treatment effect of access to benchmarking information via a mechanism that does not require concurrent access to data on other hospitals' purchases. In our analysis, we term this the "agency effect" to denote its relation to the mechanism outlined in section III, in which the benchmarking database allows hospitals to resolve a negotiator agency problem. Figure 2*B* illustrates this identification strategy graphically. In this stylized example, we have hospital A joining the database well before the brand enters the market; hospital B joins after the brand enters. Once the brand enters, each hospital negotiates prices; hospital B is untreated, while hospital A is treated (superscript "Ag"), in the sense that it has joined but has no concurrent data on other hospitals. In the next period, after price data are submitted, loaded, and released to database members, nonmember hospital B remains untreated, but hospital A receives another treatment (superscript "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.[22]

Entering brands also allow us to investigate potential bias due to timing of join. Any persistent bias associated with factors beyond information at hospitals that have joined will be included in the difference between pre- and postjoin hospitals in the first few months after new-brand introduction (β^{Ag}). Thus, our estimate of any "AI" effect where hospitals use information concurrently available in the database to negotiate better prices (β^{Info}) would be free of such bias.

V. Results: How Information Affects Negotiated Prices

In this section, we estimate regressions based on the research design just described to more carefully measure and understand the effect of information suggested by figure 1, accounting for time-invariant differences across hospitals (or hospital-brand combinations) and trends in prices over time. All of the regressions we present are extensions of a baseline specification implementing our difference-in-differences design around the timing of hospital access to benchmarking information.[23] Letting P_{hjt} denote the price observed for hospital h, brand j, and month t, our preferred specification controls for hospital-brand fixed effects [θ_{hj}], month fixed effects [θ_t], and separate linear time trends for each brand [$\gamma_j \times (t - t_{\min_j})$]:[24]

$$P_{hjt} = \beta^{\text{Info}} \times \mathbf{1}_{\{\text{post}_{hjt}\}} + \theta_{hj} + \theta_t + \gamma_j \times (t - t_{\min_j}) + \varepsilon_{hjt}. \tag{1}$$

Here, $\mathbf{1}_{\{\text{post}_{hjt}\}}$ is an indicator equal to one after a hospital first accesses benchmarking information for the given brand and zero before that,

[22] Formally, 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 ($\beta^{\text{Ag}} = \Delta_{t=1} - \Delta_{t=3}$) and information ($\beta^{\text{Info}} = \Delta_{t=2} - \Delta_{t=3} - \beta^{\text{Ag}}$) effects.

[23] This includes information upon joining and when new brands enter. We show results estimated only from the "timing-of-join" variation in app. A.3 and find our discussion unaffected.

[24] The term $t_{\min_j}$ represents the first period in which we observe data for brand j: either the beginning of our sample or the month of entry of brand j into the market. To address concerns that linear trends do not adequately account for price trends at the beginning of a brand's life cycle, refer to app. D.1 for results with brand-month fixed effects, which are qualitatively similar.

making the coefficient β^{Info} an estimator for the treatment effect. All of the regressions and results below extend this specification to allow for varying types of heterogeneity in this treatment effect. Results with alternative fixed effects and time trends are discussed as well.

A. *Effects of Information throughout the Price Distribution*

Our first result, shown in table 1, regards the average treatment effect of information across all hospital-brand-months. Results are shown for a variety of different specifications of control variables. The estimates are significantly smaller when we control for hospital × brand (rather than hospital-plus-brand) fixed effects. This may be due to our effectively controlling for an unknown source of hospital-brand-specific heterogeneity, or hospital-brand fixed effects may introduce attenuation bias toward zero, as there are some hospital-brands for which there are relatively few observations. We generally find that version 3 treatment effects are smaller in magnitude and more precise than those for version 4, so we focus on these results in the main text for the sake of brevity.[25]

The preferred specification finds that prices decrease by only \$3, on average, when benchmarking data are accessed. This average treatment effect (ATE) is also imprecisely estimated, with a standard error of \$3.[26] In keeping with the empirical predictions derived from theory, the remainder of our analyses allow for heterogeneity in treatment effects, depending on each hospital-brand pair's place in the price distribution for that brand at the time the hospital gains access to information. Upon member h's first log-in to the database ($t = \tau_h$), we compare P_{hj,pre_h} for each brand j purchased in the year before log-in ($\text{pre}_h := \{t \in [\tau_h - 13, \dots, \tau_h - 1]\}$) to the full distribution of prices $\{P_{h'j,\text{pre}_h}\}_{h' \in \mathcal{H}}$ across all hospitals $\mathcal{H}$ during (pre_h). We assign each pair hj to a quintile k of the prejoin price distribution. We then estimate a version of equation (1) with separate treatment effects for each quintile, such that coefficient β_k^{Info} represents the estimated treatment effect of accessing information in the benchmarking service for quintile k of the preinformation price distribution $\mathbf{1}_{\{P_{hj,\text{pre}_h} \in \text{quintile}(k, \{P_{h'j,\text{pre}_h}\}_{h' \in \mathcal{H}})\}}$.

Figure 3*A* shows the results. The treatment effects exhibit substantial heterogeneity, depending on the preinformation price the hospital was paying for a brand relative to others. The treatment effects are statistically zero in all but the top quintile of the preinformation price distribution, where the effect is −\$27. This evidence is consistent with prediction 1 that,

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

[26] Detailed tables and figures on the timing of the effect are available in app. D.1.

TABLE 1
Average Treatment Effects of Information across All Hospital-Brand-Months

	Version of Controls			
	1	2	3	4
β^{Info}	−12***	−21***	−3	−7
	(5)	(7)	(3)	(5)
Hospital + brand fixed effects	Y	Y	N	N
Hospital × brand fixed effects	N	N	Y	Y
Linear brand trends	Y	N	Y	N
Brand × month fixed effects	N	Y	N	Y

Note.—Authors' calculations from PriceGuide data, 2009–14. N = 32, 453 member-brand-months; includes 508 members. Standard errors clustered at the hospital (ver. 1, 2) or hospital-brand (ver. 3, 4) level are shown in parentheses.
*** Significant difference from zero at the 1% level.

absent benchmarking, pessimistic hospitals would pay suppliers high prices regardless of those suppliers' true bargaining parameters, leading those hospitals to negotiate lower prices after joining. Under the baseline AI mechanism, there is little reason to expect transparency to affect prices that are relatively good. In fact, the treatment effect for the bottom quintile of the price distribution is small and negative, providing little support for the externality discussed in prediction 3.[27]

We also performed an event study analysis separately for each quintile of the price distribution. The results for the top quintile of the preinformation price distribution are shown in figure 3*B*.[28] The pretrends in the 6 months before information access 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 estimates for the 6–12 months before information access are negative, though not significant. If one were to lend weight to the noisy point estimates, pretrends before joining the database would lead the difference-in-differences estimates to be an understatement of the effects of information on price. This lack of pretrends is strong suggestive evidence that the estimated treatment effects are due to accessing the information in the benchmarking data rather than to any potential sources of bias due to join timing.[29]

[27] A positive shift in the lowest part of the price distribution could be evidence of suppliers' reduced willingness to provide discounts in the presence of benchmarking or of mean reversion. These results suggest that neither is a significant factor for stents.

[28] Results using alternative fixed effects are available in figs. A11 and A12.

[29] Indeed, the evidence of steeper negative price trends after join in the top quintile of the price distribution than in average prices suggests that, if there are factors that cause prices to decrease after join that are unrelated to information access, they must disproportionately affect hospital-brands whose prices are relatively high in the prejoin period, a fact that would be unknown to parties whose behavior affects prices without them accessing the database. For the reader who prefers a more skeptical interpretation, any remaining bias due to timing of join will be absorbed with our measure of the agency effect in our mechanism test, so that we are able to obtain a "clean" AI effect.

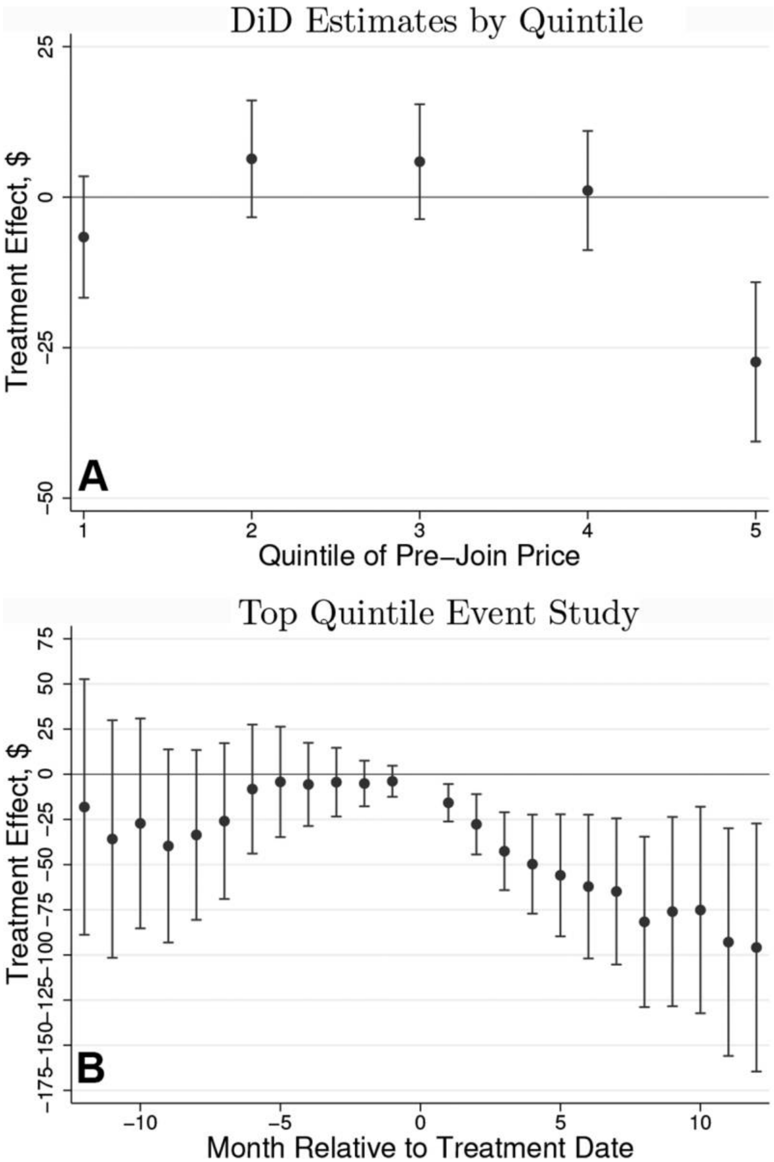

FIG. 3.—Treatment-effect estimates throughout the price distribution: authors' calculations from PriceGuide data, 2009–14. *A*, $N = 32{,}453$ member-brand-months, for 508 members. DiD = difference-in-differences. *B*, $N = 23{,}016$ member-brand-months, for 507 members, 12 months before and after join only. Bars represent 95% confidence intervals; standard errors are clustered at the hospital-brand level. A color version of this figure is available as an online enhancement.

For the sake of statistical power and for expositional simplicity, we return to estimating pre-/posttreatment effects, rather than breaking them down by month relative to information access. However, it is noteworthy that treatment effects become larger over the course of the year after information access. We see this as evidence of price "stickiness," a friction that limits gains from transparency, and we return to this issue in section V.B.3.

B. Mechanisms: Where and Why Does Information Matter Most?

The above results establish that transparency in the form of access to benchmarking information leads to lower prices for hospital-brand cases where the hospital is in the upper quintile of the price distribution for that brand. In this section, we test the further predictions from section III to better understand the mechanisms behind these price reductions. We first allow for treatment effects to vary with purchase volume so that we may investigate whether hospital-brands with high expenditures at stake experience larger changes, in keeping with predictions 2 and 5. Next, we utilize the fact that, for new brands, no benchmarking information is available in the database until several months after brand entry; this allows us to separate the AI mechanism from the agency mechanism (prediction 6). Finally, we decompose the estimated price effects into price effects conditional on renegotiation and price effects due to greater likelihood of renegotiation. The estimates are summarized in table 2 here and discussed in turn below.[30]

1. Costs of Putting Information to Use: Treatment Effects and Quantity

To the extent that using benchmarking information to identify opportunities and then engage in renegotiation (of supply contracts or employment contracts) is costly, predictions 2 and 5 suggest that transparency will have the largest effect for hospitals and brands with high quantities involved. To investigate these predictions, we generate dummy variables $\mathbf{1}_{\{\text{low}^q_{hj,\text{pre}}\}}$ and $\mathbf{1}_{\{\text{high}^q_{hj,\text{pre}}\}}$ that divide the sample into hospital-brands with monthly purchase volume below and above the 75th percentile in the months before join, and we estimate a model that allows for treatment effects to vary by prejoin price and quantity, where $\beta^{\text{Info}}_{k,\text{low}^q}$ now estimates the treatment effect, for quintile k, for lower-volume brands and $\beta^{\text{Info}}_{k,\text{high}^q}$ now estimates the treatment effect, for quintile k, for higher-volume brands.

[30] Detailed results with alternative fixed effects specifications are available for each of the table 2 panels in figs. A13–A15 in app. D.1.

TABLE 2
Treatment Effects of Information: Mechanisms

P_{quintile}					P_{quintile}				
1	2	3	4	5	1	2	3	4	5
A. Treatment-Effect Variation with Quantity Purchased									
Low Quantity: $\beta^{\text{Info}}_{\text{quintile,low}^q}$					High Quantity: $\beta^{\text{Info}}_{\text{quintile,high}^q}$				
−4	9	9	5	−17**	−11	0	0	−9	−71***
(6)	(6)	(6)	(6)	(7)	(9)	(8)	(7)	(8)	(13)
B. Agency versus AI Mechanisms									
Future Information: $\beta^{\text{Ag}}_{\text{quintile}}$					Concurrent Information: $\beta^{\text{info}}_{\text{quintile}}$				
−17	−3	2	7	13	−1	7	5	−1	−30***
(11)	(12)	(10)	(12)	(18)	(6)	(5)	(5)	(5)	(7)
C. Renegotiation									
Pr Renegotiation: $\mathbf{1}(\{\text{reneg}_{hjt}\})$					Upon Renegotiation: $\beta^{\text{Info}}_{\text{quintile},\mathbf{1}_{\{\text{reneg}_{hjt}\}}}$				
.01	.013	.016*	.018	.023**	−14	4	1	−13	−76***
(.01)	(.01)	(.009)	(.011)	(.009)	(15)	(14)	(19)	(17)	(18)

Note.—Authors' calculations from PriceGuide data, 2009–14. N = 32, 453 member-brand-months; includes 508 members. Standard errors clustered at hospital-brand level are shown in parentheses.
* Significant difference from zero at the 10% level.
** Significant difference from zero at the 5% level.
*** Significant difference from zero at the 1% level.

The estimates in panel A of table 2 show that the price treatment effect is largest for high-volume hospital-brands in the upper part of the price distribution. At −\$71, the top-quintile treatment effect for high-quantity hospital-brands is more than four times the effect for low-quantity hospital-brands.

The fact that quantity matters suggests that costs of attention, analysis, or action act as frictions that sustain price variation. That more savings are not realized, even when large quantities are at stake, suggests that further frictions independent of information, such as strong physician brand preferences, could be important as well. As shown in appendix D, we see similar patterns when we consider different sets of fixed effects and time trends, when we modify the regression sample to focus on only the 12 months before and after information access, when we limit the sample to hospitals only, and when we estimate effects within similar sets of hospitals. The results are also similar when we identify treatment effects based only on the information shock of database "join" as part of the expanded analysis discussed in section VI. All told, these results imply partial price convergence: removing time trends and applying our treatment-effect estimates to the prejoin price distribution decreases the standard deviation of price by 3.7% among low-quantity hospital-brands and by 6.4% among high-quantity hospital-brands.

2. New Brands: Agency and AI Mechanisms

The β^{Info} estimates thus far have provided a treatment effect of access to the benchmarking information, subsuming both the agency and AI mechanisms that market participants put forth, as outlined in section III. To separate these theories, we leverage the fact that almost all hospitals negotiate their first contract with a new brand 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 3 or 4. By month 6, there are enough observations in the database for a hospital to develop a useful estimate of its place in the price distribution for that brand. We use this to estimate a regression in the spirit of figure 2 that allows for heterogeneous treatment effects by price quintile and by these two information states. We estimate β_k^{Ag} by interacting the price quintile k treatment effect with an indicator for all hospital-months after the hospital joins and logs in to the benchmarking database $\mathbf{1}_{\{\text{post}_{ht}^{\text{join}}\}}$, and we estimate a "clean" β_k^{Info} by including a further interaction with an indicator that is equal to one upon a hospital's first log-in more than 6 months after the introduction of that brand $\mathbf{1}_{\{(t-t_{\min_j})>6\}}$.

The estimates in panel B of table 2 suggest that the AI effect explains a substantial portion of the effect of information on prices. To the extent that one remains concerned about endogenous timing of join, recall from section IV that any associated bias will be captured in β^{Ag} but not in β^{Info}. Thus, our most robust finding is that of a statistically and economically significant β^{Info}, concentrated among those paying the highest prices before obtaining information and consistent with the use of concurrent information in bargaining. These results are most consistent with the theory of asymmetric information in bargaining based on Rubinstein (1985).

One implication of this result is 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. It suggests that these information and incentive issues should be kept in mind when thinking about the factors driving bargaining outcomes, including as potential sources of changes to bargaining parameters in counterfactuals with negotiated prices.

3. Price Changes with "Sticky" Contracts

All of the price coefficient estimates reported thus far can be described as capturing the combined effects 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 spending of hospitals on medical inputs over time, taking into account the stickiness of prices in real-world hospital-supplier contracting. However, in the main estimation sample, renegotiations take place in 9% of observations (member-brand-months with any transactions), and prices decrease, on average, by $25 at each renegotiation, meaning that 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.

We consider these two effects separately by flagging hospital-brand-month observations in which renegotiation is observed $\mathbf{1}_{\{\text{reneg}_{hjt}\}}$.[31] We then estimate the effect of information on the rate of renegotiation, using the usual price-quintile specification but with the indicator for renegotiation as the dependent variable. The effect of information on price, conditional on renegotiation, is obtained from the same regression, run only on the subset of data where the renegotiation indicator equals one.

The results in panel C of table 2 show that the effect of information on the likelihood of renegotiation is statistically significant (at the 5% level) only in the top quintile of price, where information increases the probability of renegotiation by 2.3 percentage points, or about one-quarter the baseline probability of renegotiation. Point estimates in other quintiles are positive but smaller and are not significant at conventional levels. To the extent that this is not simply a statistical coincidence, it could be due to a slight increase in efforts to get to the negotiating table or change in the frequency with which renegotiation results in zero price change among those with information.

By contrast, the effect of information on price conditional on renegotiation is about −$75, nearly three times the −$27 effect on price paid. Thus, the impact of transparency in the form of benchmarking information is substantially affected by renegotiation frictions.

VI. Generalizing the Results: All Important Products

While the above results from the coronary-stent sample are useful for investigating mechanisms via which savings are achieved for an important

[31] We sort transactions for each hospital-brand by month and group observations with the same price together within month. We then flag each hospital-brand-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-brand). The results are qualitatively similar (though larger in magnitude) using a less conservative method that flags all months in which average prices change. Of course, with transactions data, we cannot observe whether a renegotiation took place and price remained the same, but the baseline level of these events is differenced out in our estimation strategy. We also take some comfort that our measure results in frequency of renegotiations similar to the annual contract structure that is common in the industry.

product category, a natural question arises: What happens to the remaining 98% of hospital spending? To investigate this question, we extend our analyses to 52 product categories that are important in terms of high spending or transaction counts. We organize product categories on the basis of the likely importance of physician preferences in determining their usage. Class 1 is "commodities": relatively inexpensive products that are unlikely to be chosen primarily by physicians; for example, surgical gloves. At the other extreme, class 3 is PPIs: high-tech medical devices, mainly coronary and orthopedic products that are the primary implanted device in their corresponding surgical procedures; for example, coronary stents. Class 2 is intermediate: other medical/surgical products used during invasive procedures, but explicitly excluding PPIs.[32] See appendix A.3 for sample details.

In this expanded set of analyses, we then estimate the same regressions as in table 1 and figure 3 within each product category, using join timing to identify treatment effects of information. Appendix E reports estimated ATEs and treatment effects by price and quantity for each product category, including several alternative regression specifications, and for the full sample as well as a restricted sample focusing on products purchased in at least 100 hospitals, on average, in each year (where price quintile estimates have greater statistical power). Here, we summarize the key insights, using the full sample. Figure 4 plots the ATE (*A*) and high-price, high-quantity treatment effect (*B*) for each of the 52 product categories, normalized by mean price to facilitate comparisons across groups and organized by product class. Within each product class, treatment effects are displayed in ascending order, followed by an "X" indicating the spending-weighted ATE for the class. Table 3 shows spending-weighted averages of the estimated treatment effects for the top and bottom quintiles of the price distribution across all product categories within each class, in percentage terms.

We observe several general patterns similar to those from stents. In figure 4*A*, we see that the ATEs are negative for the majority of product categories but relatively small and rarely statistically significant. As shown in table 3, the aggregate ATE is largest for PPIs, at −0.5%. In figure 4*B*, we see that among hospital-brands in the top quartile of quantity and quintile of price at the time of joining, the treatment effects are larger in magnitude, almost always negative, and sometimes statistically significant, particularly among PPIs. In aggregate, we observe the largest price decreases

[32] Our typology overlaps substantially with the Food and Drug Administration's classification system. Class I devices, such as gloves, are deemed to be low risk and are therefore subject to the fewest regulatory controls. Class II devices, such as catheters, are higher-risk devices with greater regulatory controls to provide reasonable assurance of the devices' safety and effectiveness. Class III devices, such as replacement heart valves and coronary stents, are the highest-risk devices and must typically be approved by FDA before they are marketed (FDA 2018). For product categories that did not obviously fit into one of our classes, we relied on the FDA class directly.

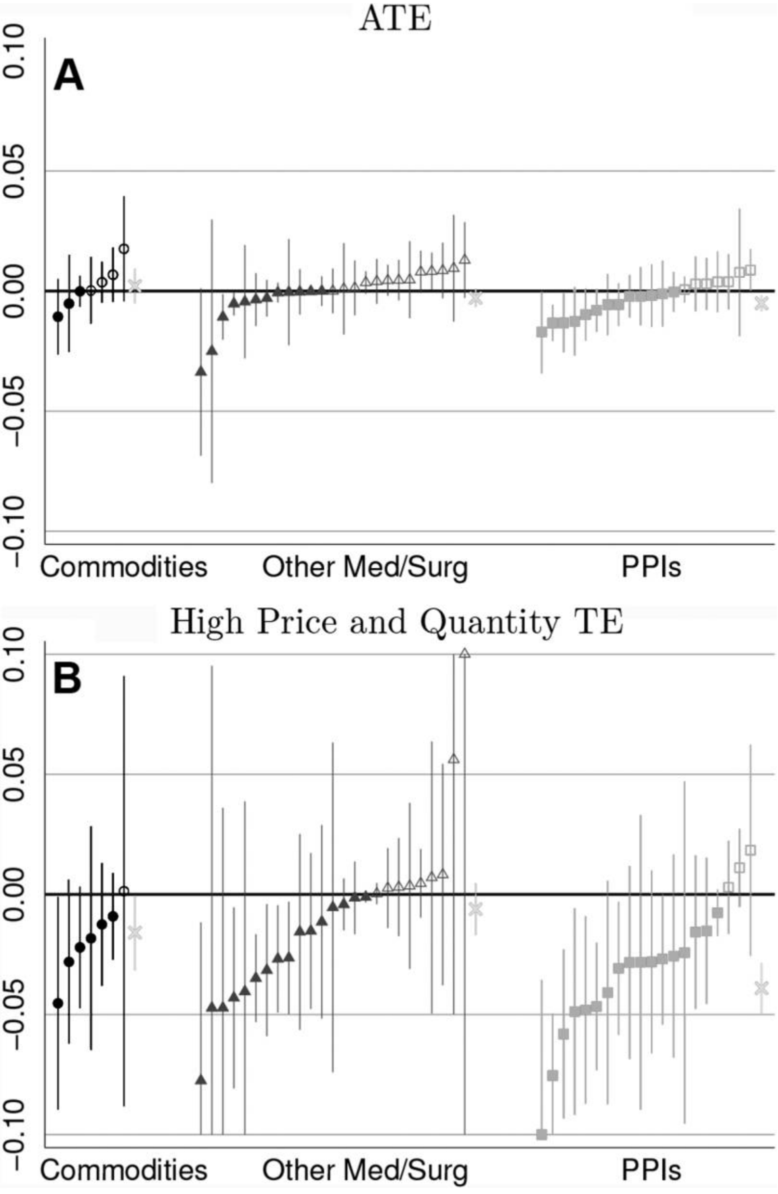

FIG. 4.—Treatment-effect estimates for important product categories: authors' calculations from PriceGuide data, 2009–14. $N_{hjt}^{(1)} = 516{,}582$; $N_{hjt}^{(2)} = 1{,}344{,}515$; $N_{hjt}^{(3)} = 703{,}544$; $N_{h}^{(1)} = 748$; $N_{h}^{(2)} = 701$; $N_{h}^{(3)} = 601$; superscripts (1)–(3) refer to the three product classes, from left to right. Reported specifications include hospital-brand and brand-year fixed effects; alternative fixed effects are shown in the appendix. Bars represent 95% confidence intervals; standard errors in the category-specific regressions are clustered at the hospital-brand level. A color version of this figure is available as an online enhancement.

TABLE 3

Spending-Weighted Averages of Estimated Treatment Effects (TEs)

	ATE	TE by Price Quintile (All Quantities)		TE by Price Quintile (Low Quantity)		TE by Price Quintile (High Quantity)		Expected Savings	
		1	5	1	5	1	5	μ (\$/h-year)	σ (\$/h-year)
Commodities (1)	.002	.013*	−.013***	.021	−.012**	.005	−.016**	63	2,828
	(.004)	(.008)	(.005)	(.016)	(.006)	(.008)	(.008)	(306)	(1,028)
Other medical/surgical (2)	−.003	.003	−.017***	.004	−.021***	.002	−.006	−254	1,974
	(.002)	(.003)	(.004)	(.004)	(.005)	(.004)	(.005)	(428)	(279)
PPIs (3)	−.005***	.014***	−.033***	.014***	−.033***	.014***	−.039***	−1,869	5,492
	(.002)	(.003)	(.004)	(.003)	(.004)	(.005)	(.005)	(1,281)	(893)

Note.—Authors' calculations from PriceGuide data, 2009–14. $N_{hjt}^{(1)} = 516{,}582$; $N_{hjt}^{(2)} = 1{,}344{,}515$; $N_{hjt}^{(3)} = 703{,}544$; $N_h^{(1)} = 748$; $N_h^{(2)} = 701$; $N_h^{(3)} = 601$; superscripts (1)–(3) refer to the three product classes in the first column. Reported specifications include hospital-brand and brand-year fixed effects; alternative fixed effects are shown in the appendix. Standard errors in the category-specific regressions, clustered at the hospital-brand level, are shown in parentheses.

* Significant difference from zero at the 10% level.

** Significant difference from zero at the 5% level.

*** Significant difference from zero at the 1% level.

under transparency for hospitals purchasing PPIs at formerly high prices (3.3% savings) and in relatively high quantities (3.9% savings).

Our finding of negative treatment effects in the top of the price distribution for PPIs is quite robust to decisions regarding sampling and regression specification. As discussed in greater detail in the appendix, the numerous specifications reported in tables A10–A12 consistently document a negative and significant average PPI estimate for high-price, high-quantity hospital-brands, usually document a negative and significant estimate for high-price hospital-brands, and report treatment-effect estimates that are broadly similar to our preferred estimate of 3.3%–3.9% savings on PPIs under transparency.[33]

In contrast to that for stents, the expanded analysis also reveals price increases in the bottom part of the price distribution for some product categories. As summarized in table 3, we document a marginally significant price increase of 1.3% in the bottom quintile of the price distribution for the average commodity; the analogous result for PPIs is 1.4% and is estimated more precisely. The detailed category-level results in table A10 show that this pattern is particularly strong for several prosthesis categories in the PPI class. These results are consistent with an externality effect of transparency, as set forth in prediction 3, and accordingly merit further discussion. Focusing on PPIs, the price increases in the bottom of the distribution are smaller than the price decreases in the top of the distribution (3.3%). The PPI-average point estimates are the same for buyers purchasing in relatively low quantities and for buyers purchasing in relatively high quantities (1.4%); our theoretical framework predicted larger effects for low-quantity buyers, so this is suggestive evidence against the externality model, but standard errors are such that we cannot rule out an economically meaningful difference in effects.

Also in contrast to the robust results we report for the top of the price distribution, the results for the bottom part of the price distribution are somewhat fragile. As detailed in table A12, many of the treatment-effect estimates in the first price quintile are sensitive to whether they are estimated in a pooled regression of all price quintiles together (*left*; indicating an average PPI effect of 1.4% for the bottom price quintile) versus separate regressions of each price quintile (*right*; indicating an average PPI effect of −0.7% for the bottom price quintile). The primary difference between these specifications is that the pooled regression imposes the same

[33] Table A10 compares our preferred full-sample specification with all price quintiles estimated simultaneously and hospital-brand and brand-year fixed effects (*left*), to the same specification estimated on a restricted sample of brands with at least 100 hospitals purchasing in the average year (*right*). Table A11 presents summary versions of the full-sample (*left*) and restricted-sample (*right*) regressions in table A10, aggregated to the product-class level, with alternative control sets. Table A12 compares our preferred full-sample specification (*left*) to an alternative full-sample specification with hospital-brand and brand-year fixed effects but with each price quintile estimated in a separate regression (*right*).

brand-year trends across hospitals in all quintiles of the price distribution, whereas the nonpooled version allows brand-year trends to vary on the basis of hospitals' position in the price distribution. Thus, the difference in the first quintile results may be due to differential brand-specific price trends for low-price hospitals. For example, if low-price hospitals are those that have relatively high bargaining parameters, it may be the case that those hospitals achieve high discounts regardless of changes in the market environment and accordingly exhibit less steep downward trends than other hospitals.

The broad robustness of the treatment effects among high-price (and high-price, high-quantity) hospital-brands to sample and specification choices leads us to focus on them as the primary effect of transparency that we find in this study. Thus, our main takeaway across all of these top categories is that high-price hospital-brands within PPIs experience 3.3% savings (or 3.9% savings among high-quantity hospital-brands) after hospitals obtained access to benchmarking data; savings are limited for other hospitals and products. We offer these results with the caveat that we observe suggestive evidence of a positive externality among previously low-price hospital-brands for some PPIs, though this evidence is sensitive to specification, and the secondary hypothesis of larger increases of low prices in cases with low quantities is not borne out in any specification.

The final two columns in table 3 show the mean and standard deviation of expected annual savings across all hospitals. Savings are calculated from the richest specification, with separate treatment effects for each hospital-brand based on its position in the prejoin price and quantity distributions. Savings are aggregated to the hospital-product category, then averaged over product categories for each class. The results reinforce that, on average, hospitals can expect modest savings of $1,869 on PPIs, but there is a large amount of heterogeneity across hospitals. A one standard deviation improvement takes expected savings to $7,361 per hospital–product category–year. The favorable parts of the savings distributions for commodities and other surgical supplies offer substantial opportunities as well, with annual category savings of over $2,000 one standard deviation from the mean.

VII. Conclusion

This paper conducts the first empirical study of the impact of transparency on price negotiations in business-to-business markets. Our empirical context is hospital-supply purchasing, an area where there has been keen interest in information as a way to decrease hospital-supply costs. Using new data on all purchase orders issued by over 17% of US hospitals from 2009–14 and difference-in-differences research designs, we find that hospitals that gain access to benchmarking information see subsequent savings on the brands for which they were previously paying relatively high prices. These estimates provide evidence on the potential economic impacts of the rise in benchmarking data services marketed toward buyers in business-to-business markets. They also suggest that information is a potentially important driver of heterogeneity in negotiated prices, with implications for the growing structural empirical literature in bargaining.

Our tests of the mechanisms behind these information effects imply that the value of information is attenuated by the costs of putting the information to use. The evidence suggests that there are costs consistent with time-constrained negotiators (gains are focused in high-quantity items, where the most money is at stake) and also with the stickiness of business-to-business contracts (long-term contracts may not be renegotiated for some time). The latter friction is a fundamental feature of many business-to-business markets. However, the time and effort cost of accessing and/or using information could be reduced as technology improves. As both information and analytics are increasingly important in the modern economy, this suggests a path for future research.

We examined two potential theories for how benchmarking information might be used in a business-to-business setting—asymmetric information about seller bargaining parameters and buyer-side negotiator agency. We found robust evidence for the AI theory but noisy evidence for agency. Within the AI framework, we find strong evidence that transparency leads to partial price convergence through decreases in the top of the price distribution. We also find some evidence that prices increase in the bottom of the distribution, though these are sensitive to empirical specification. We see modeling frictions in the use of information, the potential for information to interact with within-firm agency frictions in negotiation, and externalities in multilateral bargaining with asymmetric information as three especially interesting areas for future theory development.

While our results suggest that intermediaries who increase transparency may indeed lower the prices hospitals pay for a wide variety of medical supplies, our detailed analysis of mechanisms focuses on coronary stents. Variation across product markets in terms of supply-side competition, complexity of contracts (e.g., nonlinearities or multiproduct bundling), and the particular mechanisms through which information affects prices thus represents a rich opportunity for future empirical analysis of information in business-to-business bargaining. Such work would require expanding the empirical toolkit to analyze complex contracts when the contract terms themselves may not be observed.

Finally, while this paper takes a small step toward understanding the implications of information in business-to-business markets, more work is necessary to evaluate what we would expect as benchmarking information diffuses into wider use or in policy proposals for greater transparency in medical device markets. For example, a more structural model, combined with variation in market structure and information penetration, could explore the potential roles of supply-side phenomena such as obfuscation or collusion.

References

Albæk, Svend, H. Peter Møllgaard, and Per Overgaard. 1996. "Law-Assisted Collusion? The Transparency Principle in the Danish Competition Act." *European Competition Law Rev.* 17:339–43.

Ausubel, Lawrence M., Peter Cramton, and Raymond J. Deneckere. 2002. "Bargaining with Incomplete Information." In *Handbook of Game Theory*, vol. 3, edited by Robert Aumann and Sergiu Hart, 1897–945. Amsterdam: North-Holland.

Backus, Matthew, Thomas Blake, Bradley Larsen, and Steven Tadelis. 2018. "Sequential Bargaining in the Field: Evidence from Millions of Online Bargaining Threads." Working Paper no. 24306 (February), NBER, Cambridge, MA.

Backus, Matthew, Thomas Blake, and Steven Tadelis. 2019. "On the Empirical Content of Cheap-Talk Signaling: An Application to Bargaining." *J.P.E.* 127 (4): 1599–628.

Bandiera, Oriana, Andrea Prat, and Tommaso Valletti. 2009. "Active and Passive Waste in Government Spending: Evidence from a Policy Experiment." *A.E.R.* 99 (4): 1278–308.

Binmore, Ken, Ariel Rubinstein, and Asher Wolinsky. 1986. "The Nash Bargaining Solution in Economic Modelling." *RAND J. Econ.* 17 (2): 176–88.

Bronnenberg, Bart J., Jean-Pierre Dubé, Matthew Gentzkow, and Jesse M. Shapiro. 2015. "Do Pharmacists Buy Bayer? Informed Shoppers and the Brand Premium." *Q.J.E.* 130 (4): 1669–726.

Calsyn, Maura, and Ezekiel J. Emanuel. 2014. "Controlling Costs by Expanding the Medicare Acute Care Episode Demonstration." *JAMA Internal Medicine* 174 (9): 1438–39.

Card, David, Alexandre Mas, Enrico Moretti, and Emmanuel Saez. 2012. "Inequality at Work: The Effect of Peer Salaries on Job Satisfaction." *A.E.R.* 102 (6): 2981–3003.

Collard-Wexler, Allan, Gautam Gowrisankaran, and Robin S. Lee. 2019. "'Nash-In-Nash' Bargaining: A Microfoundation for Applied Work." *J.P.E.* 127 (1): 163–95.

Cooper, Zack, Stuart V. Craig, Martin Gaynor, and John Van Reenen. 2019. "The Price Ain't Right? Hospital Prices and Health Spending on the Privately Insured." *Q.J.E.* 134 (1): 51–107.

Craig, Stuart, Matthew Grennan, and Ashley Swanson. 2018. "Mergers and Marginal Costs: New Evidence on Hospital Buyer Power." Working Paper no. 24926 (August), NBER, Cambridge, MA.

Cramton, Peter C. 1992. "Strategic Delay in Bargaining with Two-Sided Uncertainty." *Rev. Econ. Studies* 59:205–25.

Crawford, Gregory S., and Ali Yurukoglu. 2012. "The Welfare Effects of Bundling in Multichannel Television." *A.E.R.* 102 (2): 643–85.

Cutler, David, and Leemore Dafny. 2011. "Designing Transparency Systems for Medical Care Prices." *New England J. Medicine* 364:894–95.

Draganksa, Michaela, Daniel Klapper, and Sofia B. Villas-Boas. 2009. "A Larger Slice or a Larger Pie? An Empirical Investigation of Bargaining Power in the Distribution Channel." *Marketing Sci.* 29 (1): 57–74.

Dranove, David, and Ginger Zhe Jin. 2010. "Quality Disclosure and Certification: Theory and Practice." *J. Econ. Literature* 48 (4): 935–63.

Duggan, M., and Fiona M. Scott Morton. 2006. "The Distortionary Effects of Government Procurement: Evidence from Medicaid Prescription Drug Purchasing." *Q.J.E.* 121:1–30.

Ellison, Glenn, and Sara Fisher Ellison. 2009. "Search, Obfuscation, and Price Elasticities on the Internet." *Econometrica* 77 (2): 427–52.

FDA (Food and Drug Administration). 2018. "Classify Your Medical Device." https://www.fda.gov/MedicalDevices/DeviceRegulationandGuidance/Overview/ClassifyYourDevice/default.htm.

Gowrisankaran, Gautam, Aviv Nevo, and Robert Town. 2015. "Mergers When Prices Are Negotiated: Evidence from the Hospital Industry." *A.E.R.* 105 (1): 172–203.

Grennan, Matthew. 2013. "Price Discrimination and Bargaining: Empirical Evidence from Medical Devices." *A.E.R.* 103 (1): 145–77.

———. 2014. "Bargaining Ability and Competitive Advantage: Empirical Evidence from Medical Devices." *Management Sci.* 60 (12): 3011–25.

Hastings, Justine, Christopher A. Neilson, and Seth D. Zimmerman. 2015. "The Effects of Earnings Disclosure on College Enrollment Decisions." Working Paper 21300 (June), NBER, Cambridge, MA.

Hendricks, Kenneth, Alan Sorensen, and Thomas Wiseman. 2012. "Observational Learning and Demand for Search Goods." *American Econ. J.: Microeconomics* 4 (1): 1–31.

Ho, Kate, and Robin S. Lee. 2017. "Insurer Competition in Health Care Markets." *Econometrica* 85 (2): 379–417.

Holmstrom, Bengt. 1982. "Moral Hazard in Teams." *Bell J. Econ.* 13 (2): 324–40.

Horn, Henrick, and Asher Wolinsky. 1988. "Bilateral Monopolies and Incentives for Merger." *RAND J. Econ.* 19:408–19.

James, Julia. 2012. "Health Policy Brief: Pay-for-Performance." *Health Affairs*, October 11. https://doi.org/10.1377/hpb20121011.90233.

Jin, Ginger Zhe, and Phillip Leslie. 2003. "The Effect of Information of Product Quality: Evidence from Restaurant Hygiene Grade Cards." *Q.J.E.* 118 (2): 409–51.

Jindal, Pranav, and Peter Newberry. 2018. "To Bargain or Not to Bargain: The Role of Fixed Costs in Price Negotiations." *J. Marketing Res.* 55 (6): 832–51.

Kaplan, Greg, and Guido Menzio. 2015. "The Morphology of Price Dispersion." *Internat. Econ. Rev.* 56 (4): 1165–206.

Larsen, Bradley. 2018. "The Efficiency of Real-World Bargaining: Evidence from Wholesale Used-Auto Auctions." Working Paper no. 20431, NBER, Cambridge, MA.

Lazear, Edward P., and Paul Oyer. 2013. "Personnel Economics." In *The Handbook of Organizational Economics*, edited by Robert Gibbons and John Roberts, 479–519. Princeton, NJ: Princeton Univ. Press.

Lewis, Matthew S., and Kevin E. Pflum. 2015. "Diagnosing Hospital System Bargaining Power in Managed Care Networks." *American Econ. J.: Econ. Policy* 7 (1):243–74.

Lieber, Ethan M. J. 2017. "Does It Pay to Know the Prices in Healthcare?" *American Econ. J.: Econ. Policy* 9 (1): 154–79.

Maeda, Jared Lane K., Susan O. Raetzman, and Bernard S. Friedman. 2012. "What Hospital Inpatient Services Contributed the Most to the 2001–2006 Growth in the Cost per Case?" *Health Services Res.* 47 (5): 1814–35.

McConnell, K. John, Richard C. Lindrooth, Douglas R. Wholey, Thomas M. Maddox, and Nick Bloom. 2013. "Management Practices and the Quality of Care in Cardiac Units." *JAMA Internal Medicine* 173 (8): 684–92.

Pauly, Mark V., and Lawton R. Burns. 2008. "Price Transparency for Medical Devices." *Health Affairs* 27 (6): 1544–53.

Rubinstein, Ariel. 1982. "Perfect Equilibrium in a Bargaining Model." *Econometrica* 50 (1): 97–110.

———. 1985. "A Bargaining Model with Incomplete Information about Time Preferences." *Econometrica* 53 (5): 1151–72.

Scholten, Patrick A., and S. Adam Smith. 2002. "Price Dispersion Then and Now: Evidence from Retail and E-tail Markets." In *The Economics of the Internet and E-Commerce*, edited by Michael R. Baye. Advances in Applied Microeconomics, no. 11, 63–88. Amsterdam: Elsevier Science.

Scott Morton, Fiona, Jorge Silva-Risso, and Florian Zettelmeyer. 2011. "What Matters in a Price Negotiation: Evidence from the U.S. Auto Retailing Industry." *Quantitative Marketing and Econ.* 9 (4): 365–402.

Shelegia, Sandro, and Joshua Sherman. 2018. "Bargaining at Retail Stores: Evidence from Vienna." Working paper 1606, Dept. Econ. and Bus., Univ. Pompeu Fabra.

Sorensen, Alan T. 2000. "Equilibrium Price Dispersion in Retail Markets for Prescription Drugs." *J.P.E.* 108 (4): 833–50.

Zettelmeyer, Florian, Fiona Scott Morton, and Jorge Silva-Risso. 2006. "How the Internet Lowers Prices: Evidence from Matched Survey and Automobile Transaction Data." *J. Marketing Res.* 43 (2): 168–81.