Quantifying the Sources of Price Dispersion: Evidence from Hospital Supplies

PRELIMINARY — PLEASE DO NOT CITE

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

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

October 14, 2016

Abstract

In a wide range of product markets in which prices are negotiated, price dispersion across buyers for similar (or even identical) products can be driven by heterogeneity in brand preferences, search costs, and bargaining abilities. We develop a model that allows for each of these frictions and estimate it using data on hospital purchases of medical devices/supplies in a variety of product categories. While nearly all categories exhibit substantial price dispersion, the drivers of dispersion vary across categories. Among physician preference items, brand preferences are important drivers of price heterogeneity; among more commodity-like products, low price elasticities are driven by high relative bargaining weights of manufacturers and search frictions increase markups even further. We estimate that benchmarking information that improves hospital bargaining for the same products has the potential to reduce average prices by $1.5 - 3.4$ percent. Information that reduces hospital search costs by half could increase hospital surplus by up to 50 percent.

*The data used in this paper were generously provided, in part, by the ECRI Institute (www.ecri.org). We gratefully acknowledge financial support from the Wharton Public Policy Initiative and National Science Foundation. Stuart Craig and Donato Onorato provided excellent research assistance. Any errors are our own.

1 Introduction

The Law of One Price fails to hold in product areas as diverse as coal (Stigler 1961), automobiles (Goldberg and Verboven 2001), retail and wholesale pharmaceuticals (Sorensen 2000; Starc and Swanson 2016), mutual funds (Hortacsu and Syverson 2004), video games (Dinerstein et al. 2014), and Chinese footwear (Roberts et al. 2016). The potential explanations for this failure are numerous. On the demand side, preferences over brands, bundles of characteristics, safety, and price may vary across consumers. On the supply side, costs of production and distribution may vary across supplier-consumer pairs; in cases where prices are negotiated for each pair, heterogeneity in contract structure and relative bargaining weights may also impact prices. Finally, information frictions may play a role by, for example, generating search costs to consumers or incentivizing supplier obfuscation. Notably, the relative importance of demand-side factors, supply-side factors, and information frictions likely varies across markets, implying that the welfare effects of price dispersion (and any policy solutions focused on addressing price dispersion) will depend crucially on the underlying economic environment.

In this paper, we examine price dispersion for a variety of hospital supply markets. Hospital supplies, including medical devices, are estimated to account for 24 percent of the dramatic growth in inpatient hospital costs between 2001 and 2006 (Maeda et al. 2012), and there is substantial variation in prices of inputs across hospitals. For the top fifty hospital supplies in our data, the average standard deviation of prices across hospitals for the same exact product-month is ten percent of the mean price. This is approximately one-third the coefficient of variation estimated for common procedure prices paid to hospitals in different hospital referral regions in Cooper et al. (2015), which could be driven in part by input price variation; it is also approximately the middle of the range of coefficients of variation found in consumer goods markets.[1]

The product markets we consider vary considerably in the complexity of their underlying technologies (e.g., exam gloves vs. implantable cardiac rhythm management devices), the strength of brand preferences, the heterogeneity (perceived or real) in safety and quality, the supplier concentration, the distribution models typically employed, and the relative importance of intermediaries in contracting. The top fifty supply categories in our data include products that are typically known as "physician preference items" or PPIs (e.g., cardiac and orthopedic implants) as well as other products that are more commodity-like (e.g., exam gloves). For physician preference technologies, usage is driven by brand-loyal

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

and price-insensitive physicians choosing which product to use to treat a given patient; for commodity-like products, demand depends to a greater extent on hospitals' procurement and stocking decisions as a function of product characteristics and price. Hospitals typically rely on the services of group purchasing organizations (GPOs) to negotiate contracts for many products, but GPO prices are used only as a starting point for direct hospital-manufacturer negotiations for certain purchases, including physician preference items and capital equipment (Schneller 2009).

Given the above variation, we present a model that allows for heterogeneity in search costs, preferences, and relative bargaining weights across buyer-supplier pairs. We estimate the model using a database containing all non-capital, non-pharmaceutical hospital supply purchases for a large sample of hospitals between 2009 and 2015. The context that makes the data available is that sample hospitals joined a benchmarking database for the explicit purpose of reducing their supply costs, under the assumption that access to benchmarking data would reduce search costs or improve bargaining for member hospitals. We estimate our model separately within each product category in order to explore the relative importance of supply and demand factors and information frictions in determining the degree of observed price dispersion. We then estimate the welfare impact of interventions that focus on addressing information frictions.

In our welfare simulation, we consider the effects of a somewhat extreme intervention: eliminating search costs to buyers entirely. This exercise is motivated by the existence of the benchmarking database, as well as solutions in other product markets that are intended to reduce search costs (e.g., websites like Amazon.com for retail products and Kayak.com for airline tickets). As discussed in our previous research on coronary stents, providing information on prices paid by other buyers for the *same* product and vendor could be useful if the dispersion is due to asymmetric information about seller bargaining abilities (Grennan and Swanson 2016), but such information would not be useful in addressing search costs. On the other hand, providing information on prices paid for the same product *across vendors* or for *similar products* would be helpful for addressing price dispersion due to search costs, but not for addressing price dispersion due to heterogeneity in preferences or bargaining ability. The counterfactual we consider can thus be thought of as capturing an upper bound on what savings or welfare improvement could be achieved by facilitating search in the medical devices space; this is an important starting point for understanding the promise of "big data" interventions in business-to-business markets.

The primary challenge to estimating an empirical model that simultaneously allows for search costs, preference heterogeneity, and bargaining heterogeneity is to separately identify these different mechanisms. This issue is easily seen by considering each separate pair of factors individually. First, the presence of both search costs and preference heterogeneity introduces an identification problem similar to the familiar selection problem in the labor economics literature (beginning with Heckman 1979) and to the problem of "selection on moral hazard" in the insurance economics literature (e.g., Einav et al. 2013): the unobservable shocks in the demand equation may be correlated with the process that generated the *set of products* under consideration. This may lead researchers to obtain biased estimates of demand parameters based on quantities purchases within the observed (endogenous) consideration set. Second, the presence of both search costs and bargaining in a market may lead to bias in models that only account for one or the other. Allen et al. (2013) show in the Canadian mortgage market that the presence of high search costs for some consumers reduces their bargaining leverage and mutes the effect of supply-side concentration on prices; empirically, this means that, when search frictions are present, the average price effect of a merger can underestimate the true underlying increase in market power. Finally, a number of researchers have shown the importance of modeling the bargaining stage in models with differentiated products demand and negotiated prices; in the context of negotiations between hospitals and insurers, Gowrisankaran et al. (2015) show that inelastic demand from end-users (insurance enrollees facing limited out-of-pocket price variation) would imply negative marginal costs when prices are modeled as the outcome of a Bertrand pricing game, but the estimates from a bargaining model imply more reasonable marginal costs. Thus, the central issue we must contend with in our empirical analysis is that the relevant choice set for the buyer (and competitor set for suppliers) is a function of search and bargaining and preference parameters, and each of these may be buyer-supplier specific.

The estimation proceeds in three steps. First, we use observed consideration sets for each hospital-month to estimate a reduced form search model. We then use flexible functions of the estimated residual as a control function selection correction in our demand model (the control function is intended to purge omitted variable bias that results from correlation between the search and demand models). Importantly, the richness of our panel data across many product categories and over time allows us to construct instrumental variables that are correlated with a product's visibility and thus search costs at a given hospital at a point in time, but uncorrelated with demand conditional on search.

In the second step, we jointly estimate differentiated products demand using observed product shares within each hospital-month consideration set, and marginal costs and bargaining parameters in a standard Nash-in-Nash bargaining framework. Market shares of products purchased provide information on relative preferences within the consideration set. Under our assumption of conditional independence (that is, independence of demand preferences from the consideration set formation process, conditional on demand observables and the control function correction term), preferences relative to the outside option are also identified by market shares of products in the consideration set. Estimation of preferences relies on rich variation in consideration sets over time; estimation of price sensitivity is based on price shocks occurring when hospitals subscribed to a benchmarking database, and on periodic price renegotiation of long-term contracts that induces movement along the demand curve. Relative bargaining abilities for each product-hospital pair are identified by the extent to which price changes as the added value of the product changes.

Third and finally, we use observed consideration sets and demand and supply parameter estimates to infer search costs. Search costs rationalize which products are included vs. excluded from hospitals' consideration sets, given supply and demand parameters. Note that, since Goeree (2008) first formalized the modeling and estimation of consideration sets, most work on search has been done in two settings: (1) online, where detailed search data and information on the consideration set are available; and (2) in wage bargaining, where only equilibrium matches are available. Our setting is somewhere in between. We do not observe buyer search in the way that online studies observe it, but we do leverage use of multiple products and the panel nature of our data over time and across buyers and suppliers and related product categories, which provide similar types of identifying information on consideration sets.

Many of our ultimate findings map closely into our ex ante predictions regarding the determinants of price variation across products. There is substantial heterogeneity in preferences for particular products across hospitals for stents, but not for staplers or gloves, reflecting the known importance of physicians' brand preferences in stent purchasing. Price elasticities are quite low for all products, including the commodity-like category of exam gloves, but this result is due to high relative bargaining weights of some glove manufacturers and obtains in spite of a quite high price coefficient in the demand model. Finally, search costs are much larger as a percentage of price for staplers and gloves than for stents, consistent with the fact that stent manufacturers devote extensive resources to marketing directly to cardiac units in hospitals and would be unlikely to be unaware of all available options.

In the last section of the paper, we consider counterfactual information interventions affecting bargaining and search that are motivated by the entry of benchmarking intermediaries and policy discussions regarding medical device transparency. Using our estimated model, we compute new equilibria with improved hospital bargaining abilities and lower search costs. We find that bargaining changes have a modest effect on prices, that is strongest in stents and staplers where markups and bargaining heterogeneity are largest. By contrast, search changes have a large impact on the total surplus and prices, especially in gloves and staplers where the product-vendor space is large and search costs are high.

2 Data and Background on Hospital Purchasing

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

2.1 Hospital Purchase Order Data

The primary data set used in this study comes from a unique database of all supply purchases made by about 16 percent of US hospitals during the period 2009-2015. The data are from the PriceGuide™ benchmarking service (hereafter, "PriceGuide data") offered by the ECRI Institute, a non-profit healthcare research organization. We observe unique (but anonymous) identifiers for each hospital and several coarse hospital characteristics: census region, facility type, and number of beds. For each transaction, we observe price, quantity, transaction month, product, and supplier. This includes a wide range of products, encompassing commodities such as cotton swabs and gloves as well as physician preference items such as stents and orthopedic implants.

The reported price and quantity data are of high quality because they are typically transmitted as a direct extract from a hospital's materials management database. Hospitals have strong incentives to report accurately because the analytics the benchmarking service's web portal provides are based on comparing the hospital's submitted data to that of others in the database.[2] Related to its materials management origins, the data is at the stock-keeping-unit (SKU) level, requiring us to use machine learning and text analysis algorithms to group SKUs that belong to the same manufacturer-product and to identify common product characteristics.[3] We also validate our algorithms against data collected from manufacturer

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

[3]The goal of the machine learning procedure is to identify the level of product at which hospital-supplier contracts are negotiated. E.g., for stents, prices are negotiated separately for each brand, and each brand subsumes a large number of SKUs. The algorithm steps through the characters in each manufacturer's SKU, at each step allowing for each product to be a unique combination of characters. It stops when it finds an inflection point where the predictive power of additional SKU characters (i.e., R^2 in a regression of price on hospital-"product" dummies) falls below a pre-specified threshold.

catalogs.

We link the purchase order data to clickstream data on each member hospital's login activity. We define each member's join date as the date of first observed login. We limit our sample to hospitals joining the database after March 2010 – ECRI reissued invitations to all existing members in January 2010, after they rolled out a new interface, and our join timing data is censored at that point.

2.1.1 Representativeness of the benchmarking database sample

The hospitals in the purchase order data voluntarily joined a subscription service that allows them to benchmark their own prices and quantities to those of other hospitals in the database and thus may not be a random sample of US hospitals. In particular, subscription is costly, so we expect hospitals with greater concerns about supply costs to be overrepresented in the database. In a survey of database members, "cost reduction on PPIs" and "cost reduction on commodities" were nearly tied as the first and second most commonly cited reasons for joining. This is in accord with our own conversations with purchasing managers who cite a broad array of reasons and product areas as motivations for benchmarking.

2.2 Focal Products

For the moment, we focus on three widely-purchased medical supplies: coronary stents, surgical staplers, and exam gloves. For each of these products, prices are determined in negotiation. Negotiation can take place directly between a hospital administrator and a representative of the product's manufacturer, or hospitals may rely on group purchasing organizations to negotiate their contracts for many products.[4]

The first product category we analyze, coronary stents, is described in detail in Grennan and Swanson (2016). Coronary stents are small metal tubes placed into narrowed coronary arteries to widen them and allow blood flow to the heart. The original technology, the bare metal stent (BMS), was approved in the early 1990s; in the early 2000s, the drug-eluting stent (DES) was introduced as an improvement over the older technology with lower risk of restenosis, a condition that may arise when scar tissue builds up around the stent and restricts blood flow yet again. Our analysis focuses on drug-eluting stents, which comprise the vast majority of the market during our sample period. Stents are one example of physician preference items, for which usage is driven by strong brand preferences of physicians choosing which product to use to treat a given patient.

[4]As noted previously, GPO prices are typically used as a starting point for direct hospital-manufacturer negotiations for physician preference items and capital equipment (Schneller 2009).

The sample data for drug-eluting stents is described in the first two columns of Table 1. We analyze data on 332 hospitals, each having an average of 35 months of transactions data. Stents can generally only be purchased directly from one of four manufacturers during our sample period. There are many more SKUs than branded products (denoted by $\mathcal{J}$), as each SKU indicates variation in fine product characteristics (e.g., stent length) that are not pertinent to the price negotiation. On average, each hospital purchased two of seven available branded stents in a given month; as discussed in greater detail in Grennan and Swanson (2016), we see no evidence of hospitals standardizing on brand or manufacturer in these data. The average sample stent cost $1,473, and 23 units of each brand were purchased in the average hospital-month with nonzero purchase.

The second product category considered in this study is surgical staplers. Surgical staplers are used as an alternative to traditional sutures, with advantages to surgeons such as speed, accuracy, and more symmetrical wound closures. Stapler models come in different materials, with different degrees of curvature, and vary in appropriateness for different surgical applications. Some surgeons prefer disposable staplers for use on only one patient; others prefer stainless steel models that are sterilized after each surgery. Surgeons may also have preferences over automatic vs. manually controlled models. Our data contain 31 thousand hospital-month-products for 490 hospitals, each having 31 months of data on average. As we see in the third and fourth columns of Table 1, the market for staplers is highly concentrated, having a Herfindahl-Hirschman Index (HHI) at the manufacturer level of 0.73. Staplers are, however, distributed via a less concentrated supply chain – two manufacturers' products are distributed by nine vendors, and the HHI at the vendor level is only 0.21. For staplers, we observe that there are fewer SKUs per product than we observed for stents – we observe a total market of 17 products covering 31 SKUs in the average month. The average sample box of three staplers cost $498 and the average hospital purchased six boxes of each product chosen in the average month.

Finally, we also analyze exam gloves. Exam gloves can be used in a hospital setting by caregivers, health-care professionals, and surgeons. Gloves are differentiated by their material (e.g., nitrile, latex, vinyl) and by their coatings (e.g., powdered vs. non-powdered, sterile vs. non-sterile, textured). These variants can affect comfort, durability, flexibility, sensitivity to touch, barrier protection, chemical resistance, and risk of allergic reaction. Our data cover 46 thousand hospital-month-product observations from 547 hospitals – as expected, exam gloves are one of the most commonly-purchased products in our data overall. As we see in the last two columns of Table 1, our data include five manufacturers and seven vendors, and the market is fairly concentrated, having a manufacturer HHI of 0.35. There are 82 unique products in the average month of our sample, covering 98 SKUs – the finer categorization of SKUs includes size and color, which don't impact negotiated price. Exam gloves are sold by the case and each case of 2,000 gloves cost \$111 on average; hospitals purchased 32 cases of each product chosen in the average month.

Table 1: Focal Product Markets

	Stents		Staplers		Gloves	
	Mean	SD	Mean	SD	Mean	SD
N_{hjt}	23,627		30,767		45,681	
N_h	332		490		547	
$Months_h$	35	21	31	21	29	20
$Mfgs$	4		2		5	
HHI_m	0.33		0.73		0.35	
$Vendors$	4		9		7	
HHI_v	0.33		0.21		0.29	
$\|J_{ht}\|$	2	1	2	2	3	3
$\|J_t\|$	7	1	17	3	82	19
$\|SKU_{ht}\|$	24	20	2	2	3	3
$\|SKU_t\|$	301	111	31	5	98	23
P_{hjt}	1,472.7	190.3	498.2	273.8	111.1	75.9
Q_{hjt}	22.8	32.0	6.0	7.5	32.2	84.2

Notes: Sample includes data on hospitals and health systems for each indicated UMDNS code. Hospitals joining the database after Q1 2010 only. Observations with erroneous or extreme price data (prices in exact multiples of other per-unit prices, prices outside the $1^{st} - 99^{th}$ percentile range) or missing item, manufacturer, or vendor information dropped. Niche manufacturers and vendors (comprising $< 1\%$ of total quantity) excluded. Stent prices and quantities are for single units. Stapler prices and quantities are for three units (the modal quantity per box). Glove prices and quantities are for $2,000$ units (the modal quantity per case). Glove data for boxes with fewer than 500 units per box omitted, as manufacturer catalogs indicate those data may correspond to inconsistent numbers of individual gloves.

The products we have included in the analysis thus far were chosen because (a) they are commonly-purchased in our sample, each being in the top fifty product categories according to either total expenditure or total quantity; and (b) they differ from one another dramatically in the (hypothesized) relative importance of brand preferences and search in determining price dispersion.

Stents are physician preference items, for which brand preferences are notoriously strong and frequently noted in policy discussions, while exam gloves are much more commodity-like: conditional on a few characteristics, such as material and coating, we do not expect particular manufacturers to be strongly preferred. Staplers are an intermediate category in this dimension – they are not mentioned as a target physician preference item in policy discussions, but different physicians may have different preferences over models and brands. Based on this broad characterization, we expect brand preferences to be most important in determining stent prices, less so in determining stapler prices, and least important in determining glove prices.

The other dimension of interest is search costs. For stents, there is limited scope for "search" in the conventional sense, as a given product can generally only be purchased directly from one of a few manufacturers, and manufacturers exert substantial effort in marketing their products to hospitals with catheterization labs. In contrast, for gloves and staplers, there are more different products to choose from and hospitals may have limited awareness of the total set of products (and corresponding prices) available from each potential vendor. Accordingly, we expect search costs to be least important in determining stent prices, more so in determining stapler and glove prices.

2.2.1 Focal products: observed price variation

Figure 1 displays the distributions of prices across hospital-item combinations for each of our focal product categories. Each plot is a histogram of fixed effects obtained from a regression of prices on dummies for hospital-item combinations and product-month fixed effects. For each category, we plot three levels of "item" definition, which are decreasing in granularity: SKU, unique branded product, and groupings of unique product characteristics. Here, "product" and SKU are as defined in Table 1 above; e.g., for stents, there are 301 unique SKUs in the average month, but only seven unique branded products. For each product category, we include in the legend the R^2 from a regression of price on product-month and hospital-item fixed effects.

In Figure 1, as in Table 1, we see that stent prices are centered around \$1,500 per unit; the price distribution for staplers is centered around \$500 per box of three; and prices of gloves are centered around just over \$100 per case. However, the price distributions are substantially less spread out than the price standard deviations in Table 1 would indicate. The standard deviation of hospital-item *fixed effects* is around \$100 (rather than \$190) for stents; around \$120 (rather than \$274) for staplers; and around \$16 (rather than \$76) for gloves. This reflects that the full distribution of prices in the transactions data reflects variation across time as well as across hospital-item combinations. That said, the coefficients of variation for these persistent price differences across hospital-items are still quite large, in the range 0.07 to 0.24, indicating substantial potential for savings.

A key objective of this project is to characterize the extent to which lower prices could be achieved by sample hospitals adding lower-priced products with similar characteristics to their consideration sets via facilitated search. Such an improvement is less likely if observed price variation is primarily driven by strong brand preferences or variation in bargaining power. For stents, there are no product characteristics that are commonly shared by brands, so the plot of hospital by characteristics bundle fixed effects in the leftmost panel is simply a plot of hospital fixed effects. These are much less dispersed than the hospital-item fixed

Figure 1: Distribution of Prices Across Hospital-Items

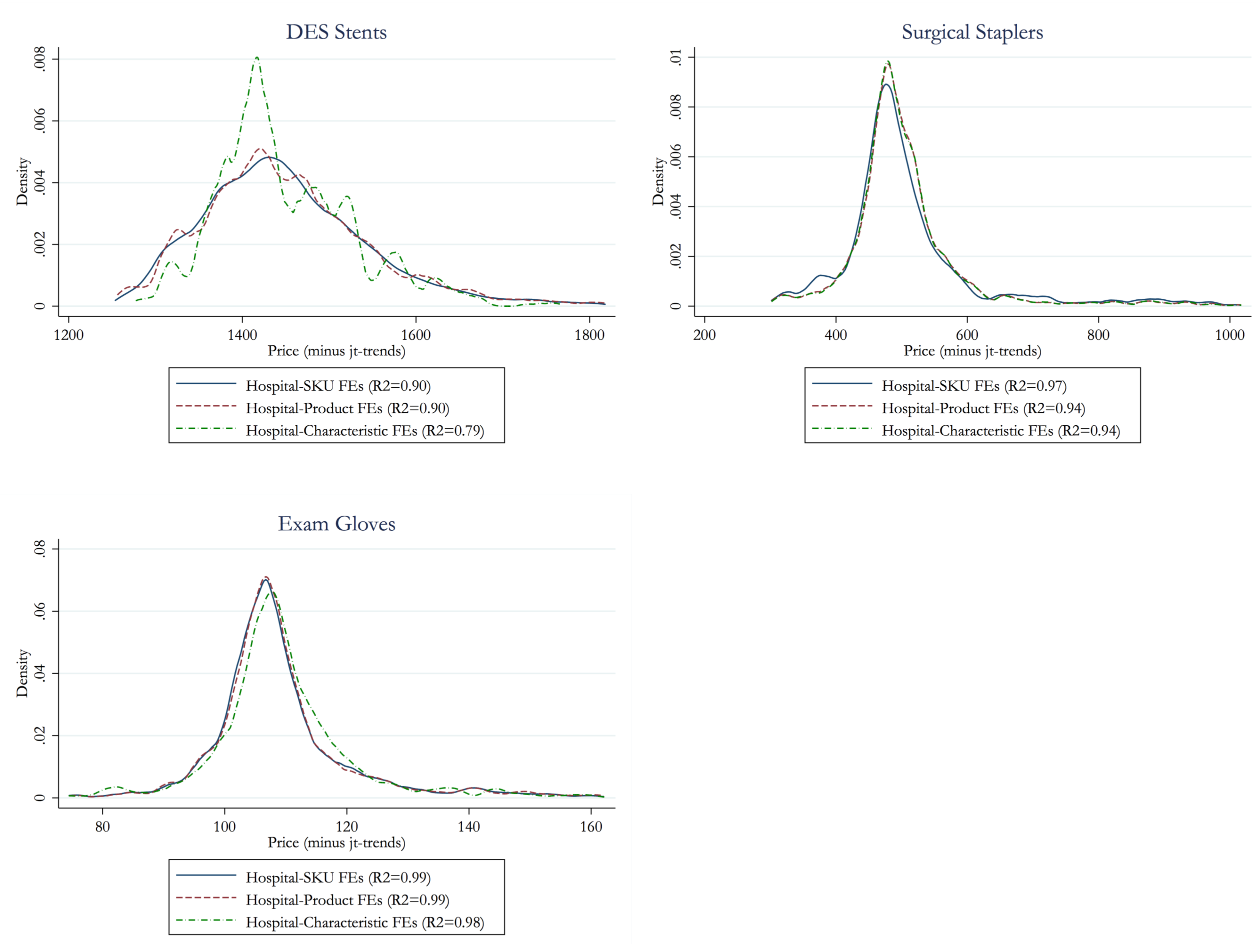

effects (defined by SKU or brand) shown in the same panel, implying that observed price variation is driven not only by variation in hospitals' average bargaining power but by other economic factors such as perhaps hospital-specific brand preferences and hospital-product-specific bargaining power. In contrast, for both surgical staplers (in the middle panel) and exam gloves (in the rightmost panel), hospital-characteristic fixed effects move much more closely with hospital-SKU and hospital-brand fixed effects, indicating that price variation for those products is determined to a greater extent by hospitals' preferences over characteristics (e.g., nitrile vs. latex gloves, circular vs. linear staplers) than by hospitals' preferences for particular brands within a characteristics bundle.

Finally, we mentioned above that the product definition used in our analyses is determined by an algorithm that is intended to identify the level at which prices are negotiated. Indeed, we see in the R^2 calculations that hospital-product fixed effects explain almost as much variation as hospital-SKU fixed effects, where the latter are the finest product definition available. We find this reassuring that our algorithm is identifying branded products correctly; we also validated the algorithm against a more "brute force" approach for stents, in which we examined manufacturers' product catalogs and verified that the algorithm identified all brands appropriately.

3 A Model of Heterogeneous Preferences, Negotiated Prices, and Search

We model hospital search and demand decisions in a framework with negotiated prices. We assume hospital demand is derived from the preferences of its staff and the needs of its patient population. Those involved in hospital purchasing search for products to add to the consideration set given product observables and the marginal cost of search.[5] Upon termination of search, the hospital contracts on price for a set of products, and purchases from that set according to its demand function.

Below, we describe each step of the search-bargaining-demand model, in reverse order of model timing. Throughout this Section, time subscripts are omitted for brevity.

[5] We allow that search and demand may be jointly determined in that the structural errors in the two processes may be correlated. In our current approach, we use a control function approach to estimate demand parameters and then assume that consideration sets are determined by a simultaneous, directed search process as a function of search costs and observable supply and demand factors. In future work, we extend this analysis to allow for optimal sequential search, and explicitly model search and demand jointly as a function of both observable product-hospital characteristics and unobserved product-hospital preferences. We use a simulated method of moments estimator (see, e.g., Goeree 2008; Gaynor et al. 2016), modified to incorporate the bargaining step.

3.1 Demand model

The utility of product $j \in \mathcal{J} = \{1, ..., J\}$ for use case i (often a doctor/patient/lesion combination for devices) at hospital h is

$$u_{ihj} = \delta_{hj} + \epsilon_{ihj}, \tag{1}$$

where ϵ_{ihj} is a use-specific unobservable i.i.d. type II extreme value error term (with scale normalized to one). The mean utility across use cases is specified as

$$\delta_{hj} = X_{hj}^d \theta^X - \theta^p p_{hj} + \xi_{hj}, \tag{2}$$

where $X_{hj}^d \theta^X$ are observed hospital-product characteristics and their utility weights in demand, $\theta^p p_{hj}$ is the negotiated price for pair hj and its utility weight, and ξ_{hj} is a hospital-product specific unobservable preference heterogeneity term.

Given contracts for a set of products $\mathcal{J}_h$, we assume the hospital chooses the product in the consideration set that maximizes utility for each use case, so that quantities demanded are given by:

$$q_{hj} = Q_h Pr[u_{ihj} > u_{ihk}, \forall k \in \mathcal{J}_h] = Q_h \frac{e^{\delta_{hj}}}{\sum_{k \in \mathcal{J}_h} e^{\delta_{hk}}} \tag{3}$$

and hospital surplus across all contracted products is given by:

$$\pi_h(\mathcal{J}_h) = Q_h \frac{1}{\theta^p} \ln \left(\sum_{k \in \mathcal{J}_h} e^{\delta_{hk}} \right). \tag{4}$$

3.2 Pricing model

In the business-to-business market for a given hospital supply, the price for a given product is buyer-specific. We assume that prices are determined between the hospital and the set of product vendors with which it contracts as a Nash Equilibrium of simultaneous bilateral Nash Bargaining problems. Each price maximizes the bilateral Nash product, taking other prices as given:

$$\begin{aligned} p_{hj} &= \arg\max \left(q_{hj} p_{hj} - C_j(q_{hj})\right)^{b_j(h)} \left(\frac{\pi_h(\mathcal{J}_h) - \pi_h(\mathcal{J}_h \setminus j)}{q_{hj}} \right)^{b_h(j)} \\ &= \frac{C_j(q_{hj})}{q_{hj}} + \frac{b_j(h)}{b_j(h) + b_h(j)} \left[\left(1 + \frac{\partial q_{hj}}{\partial p_{hj}} \frac{p_{hj} - \frac{C_j(q_{hj})}{q_{hj}}}{q_{hj}} \right) \frac{\pi_h(\mathcal{J}_h) - \pi_h(\mathcal{J}_h \setminus j)}{q_{hj}} + p_{hj} - \frac{C_j(q_{hj})}{q_{hj}} \right] \end{aligned} \tag{5}$$

where $C_j(q_{hj}) := X_{hj}^{mc} \gamma(q_{hj})$ is a function capturing the costs of manufacturing and distributing quantity q_{hj} of product j to hospital h. The terms $b_j(h)$ and $b_h(j)$ are relative bargaining ability weights that capture the extent to which the optimal price depends on vendor profits vs. the expected additional hospital surplus in the case that a contract is agreed to for product j: $\pi_h(\mathcal{J}_h) - \pi_h(\mathcal{J}_h \setminus j)$. Denote $\frac{\beta_j(h)}{1-\beta_j(h)} \equiv \frac{b_j}{b_h}\nu_{jh}$. All the terms in this pricing equation are assumed to be known to all market participants at the time of bargaining.

3.3 Search model

The demand and pricing models specified thus far are based upon a consideration set $\mathcal{J}_h$ that has been determined by hospital h's search over the set of all possible products $\mathcal{J}$. As a starting point, we assume a simple and tractable model of *directed, simultaneous* search.[6] The search process is directed in the sense that hospitals rank products based on their ex ante expected contribution to total surplus, less search costs. It is simultaneous in that the set of products to be searched is chosen ahead of time. This is in contrast to sequential search, which would allow hospitals to use a stopping rule as a function of realized parameters for products already searched, expectations of parameters for products not yet searched, and search costs.

With the simultaneous search approach, in general there may be a dimensionality problem to calculating the optimal consideration set in contexts such as ours that have a large number of products. Depending on the modeling assumptions employed, the econometrician may need to calculate the expected surplus under all possible combinations of products. Previous researchers have come up with different solutions to this problem: Chade and Smith (2006) show that assuming first-order stochastic dominance among the supplier-specific distributions over which the buyer is searching makes it optimal to rank suppliers according to expected utility and search only the top N suppliers; Moraga-Gonzalez et al. (2015) use importance sampling to deal with the dimensionality problem. Here, we follow Chade and Smith (2006) and assume that each hospital's belief distribution has product-specific means but the same variance across all products. We can then reduce dimensionality by first ranking all products according to their incremental contribution to hospital surplus, less search costs; this is without loss of generality even in the current setting where each hospital purchases multiple products as, under the assumed demand specification laid out in the previous sections, each product's contribution to surplus is monotonically increasing in its own expected indirect utility.

We operationalize this logic using the following setup. First, the hospital ranks all products by incremental contribution to hospital surplus, less the product's search cost. Because

[6] In future work, we extend this result to allow for search to be directed and sequential.

incremental surplus depends on the total consideration set thus far constructed, this is performed iteratively: the hospital first searches for the product j_1 with the single highest associated surplus less cost $\mathbb{E}[\pi_h(\{j_1\})] - sc_{j_1}$, then searches for the product j_2 with the highest $\mathbb{E}[\pi_h(\{j_1, j_2\}) - \pi_h(\{j_1\})] - sc_{j_2}$, and so on. That is, letting $\mathcal{S}_k^h$ denote the first through k^{th} products according to hospital h's ranking, the product j that is assigned rank $k+1$ will solve

$$j = \arg \max_{j' \in \mathcal{J} \setminus \mathcal{S}_k^h} \left(\mathbb{E}\left[\pi_h(\mathcal{S}_k^h \cup j')\right] - \mathbb{E}[\pi_h(\mathcal{S}_k^h)] \right) - sc_{hj'}. \tag{6}$$

We assume that hospitals observe demand parameters (θ^X, θ^p), supply parameters $(mc_j, \frac{b_j}{b_j + b_h})$, and search costs sc_{hj}; the expectation is taken over the hospital-product demand shocks ξ_{hj} and supply shocks ν_{hj}. The former expectation implies that hospitals have uncertainty regarding their true valuation of each product prior to searching; the latter implies that hospitals formulate expectations regarding the prices that will result from bilateral bargaining once their consideration sets are fixed, but that they have residual uncertainty regarding their relative bargaining position relative to each product/supplier.

After the ranking has been set, each hospital h increases the number of products k_h to be searched until the expected value of the next search drops below the hospital's search cost. The outside good $j = 0$ and first search are costless. Then the following condition holds:

$$k_h = \arg \max_k \mathbb{E}\left[\pi_h(\mathcal{S}_k^h)\right] - \sum_{j \in \mathcal{S}_k^h, j \notin \mathcal{S}_0^h} sc_{hj} \tag{7}$$

where sc_{hj} is the cost of hospital h searching product j. Note that $k_h = |\mathcal{J}_h|$ and $\mathcal{J}_h = \mathcal{S}_{k_h}^h$.

Finally, we model search costs as a function of hospital-product observables and an idiosyncratic, gamma-distributed shock:

$$sc_{hj} \sim \Gamma(\alpha, e^{X_{hj}^{sc} \phi}). \tag{8}$$

4 Identification and Estimation

As discussed above, we may in general expect that the structural errors in different parts of the model will be correlated, introducing fundamental problems with identification. In the current approach, we allow for correlation between the shocks in the demand and search models. Intuitively, a hospital's demand for a given product may be driven in part by factors that also make that product more likely to appear in the hospital's consideration set, such as the (unobserved) intensity of marketing activity for a particular hospital-product pair. We begin by implementing a control function approach in which we estimate the reduced form error in the search model and allow it to enter flexibly into the demand model.[7]

In the following, we describe our joint estimation of the demand and pricing models, and the separate estimation of the search model. We make the assumption that the set we observe the hospital purchase from is equal to its consideration set. This is consistent with the full-support idiosyncratic preference shocks ϵ_{ihj} in the demand model. It also means that we do not need to estimate the latent searched choice set as in Goeree (2008).

4.1 Demand identification

We follow the procedure in Berry (1994), setting choice probabilities implied by the demand model equal to market shares observed in the data, and inverting the system to yield a linear correspondence between a function of market shares and the mean utility for each product:

$$\ln(s_{hjt}/s_{h0t}) = \delta_{hjt} := X^d_{hjt}\theta^X - \theta^p p_{hjt} + \xi_{hjt}. \tag{9}$$

We specify preferences for each hospital-product combination using a full set of hospital, product, and month fixed effects, plus product-specific linear time trends, for all products in the hospital's consideration set: $X^d_{hjt}\theta^X = \theta_h + \theta_j + \theta_t + \theta_{trend_j}(t - t_{0_j})$. As mentioned above, we also wish to correct for a specific kind of sample selection problem introduced by the search process: we only observe demand realizations for products in each hospital's consideration set, which may introduce bias (e.g., products in the consideration set may specifically be products that are particularly preferred by the hospital). In order to do so, we use a control function selection correction approach as in Petrin and Train (2009).

Specifically, suppose that the search process can be approximated by the following reduced form:

$$\mathbb{1}(s_{hjt} > 0) := X^d_{hjt}\tilde{\phi}^X + Z_{hjt}\tilde{\phi}^Z + \tilde{\varepsilon}_{hjt} \tag{10}$$

where Z_{hjt} is a set of instruments that impact search but not demand, and $\tilde{\varepsilon}$ is a shock to the search process which may in general be correlated with the structural demand shock ξ_{hjt}. Our approach then takes the following steps: first, we estimate the reduced form search model $\mathbb{1}(s_{hjt} > 0)$ as a function of the demand observables $X^d_{hjt}\theta^X$ and instruments Z_{hjt}; second, we predict the search model error term $\hat{\tilde{\varepsilon}}_{hjt}$; finally, we include flexible functions of $\hat{\tilde{\varepsilon}}_{hjt}$ in the demand model.

In each of the included product markets, hospital buyers utilize many products from multiple suppliers, and the relative usage patterns provide information on relative preferences

[7] In ongoing work, we allow for sequential search, which directly links the search and demand models via the unobservable ξ and necessitates joint estimation of the demand, pricing, and search models.

over products purchased. Under our assumption of independence of demand from search, conditional on all demand observables and the control function correction term, preferences relative to the outside option are also identified by market shares of products in the consideration set. Estimation of preferences relies on rich variation in consideration sets over time. Estimation of price sensitivity is based on two additional sources of variation. First, we observe demand and negotiated prices before and after hospitals subscribed to a database containing information on prices paid by other hospitals for the same products. Grennan and Swanson (2016) show that subscription constituted a shock to relative bargaining weights and, in turn, prices. Second, prices are fixed in long-term contracts while hospital and doctor preferences evolve over time, so the periodic contract renegotiation we observe for each hospital-product (which are generally not concurrent with other renegotiations within the same product category) induces movement along the demand curve. Currently, we leverage these shocks using a full set of hospital and product fixed effects; in future iterations, we will more explicitly account for the timing of renegotiation and database subscription using lagged price instruments as in Grennan and Swanson (2016) and Grennan (2013). Finally, a product's likelihood of appearing in a hospital's consideration set may be a function of both anticipated demand *and* of the visibility of that product. We rely on variation in products' visibility (on average and to specific hospitals) over time as instruments in the reduced form search model that we assume are conditionally independent of demand. The set of instruments we currently employ for this purpose are *product* visibility (the percent of hospital-months in the previous three months observed to purchase hospital j in *any* quantity) and *vendor-hospital* visibility (the total spend observed for the given hospital on the current product's vendor *in other product categories* in the last three months).

4.2 Supply and demand estimation

We jointly estimate the above (linearized) demand model with control function correction term and the supply model using a generalized method of moments approach. This enables us to simultaneously recover the demand parameters θ, marginal costs, and mean relative bargaining weights. We model marginal costs as a linear function of product characteristics X_j^{mc}:

$$mc_j = X_j^{mc}\gamma. \tag{11}$$

For drug-eluting stents, there are no observed product characteristics that are expected to significantly influence marginal production or distribution costs. For surgical staplers, we allow for marginal costs to depend on dummies for the following characteristics: *linear, reloadable, disposable, endoscopic,* and *flex.* For exam gloves, we allow for marginal costs to depend on dummies for the following characteristics: *nitrile, powder-free, non-sterile, vinyl, textured,* and *latex-free.*

In order to recover the supply parameters, we rearrange the supply equation and take logs to obtain the following expression:

$$\begin{aligned} ln(\nu_{hjt}) &= ln(b_h) - ln(b_j) + ln(p_{hjt} - mc_j) \\ &- ln\left(\left(1 + \frac{\partial q_{hjt}}{\partial p_{hjt}} \frac{p_{hjt} - mc_j}{q_{hjt}}\right) \frac{\pi_h(\mathcal{J}_{ht}) - \pi_h(\mathcal{J}_{ht} \setminus j)}{q_{hjt}}\right) . \end{aligned}$$

In this expression, prices p_{hjt}, product characteristics X_j^{mc}, and demand observables X_{hjt}^d enter as data, and we condition out the terms $ln(b_h)$ and $ln(b_j)$ using product and hospital fixed effects. That leaves only the marginal cost parameters γ to be recovered from this moment equation.

We use the above expression for $ln(\nu_{hjt})$ and our demand residual

$$\tilde{\xi}_{hjt} = \ln(s_{hjt}/s_{h0t}) - X_{hjt}^d \theta^X + \theta^p p_{hjt} - CF(\hat{\tilde{\varepsilon}}_{hjt}) \tag{12}$$

in a GMM estimator, with hospital, product, and time dummies, as well as product-specific linear time trends, as GMM instruments. In the estimation, we impose the following constraints:

$$mc_j \in [0, p_{hjt}], \tag{13}$$

and

$$\frac{\partial s_{hjt}}{\partial p_{hjt}} \frac{p_{hjt} - mc_j}{s_{hjt}} \in [-1, 0]. \tag{14}$$

Finally, bargaining parameters are a known function of estimated marginal costs and demand parameters:

$$\frac{b_{jt}(h)}{b_{jt}(h) + b_{ht}(j)} = \frac{p_{hjt} - mc_j}{\left(1 + \frac{\partial q_{hj}}{\partial p_{hj}} \frac{p_{hj} - mc_j}{q_{hj}}\right) \frac{\pi_h(\mathcal{J}_h) - \pi_h(\mathcal{J}_h \setminus j)}{q_{hj}} + p_{hj} - mc_j} . \tag{15}$$

Intuitively, relative bargaining abilities for each product-hospital pair are identified by the slope with which price changes as the added value of the product changes, and marginal costs are identified as the intercepts in this relationship as added value approaches zero.

4.3 Search estimation

Our strategy for identifying search costs is based upon a *ranking rule* and a *stopping rule* for optimal search – in each period, the hospital chose to stop searching at the observed choice set when it could have, at a cost, continued to search for more suppliers. Assuming optimality of search implies that, given each hospital h's ranking of all products in $\mathcal{J}_t$, search costs and beliefs must be such that the observed choice set (defined by the number of products searched $k_{ht} \leq |\mathcal{J}_t|$) must be better than all other possibilities $k \neq k_{ht}$:

$$k_{ht} = \arg\max_k \mathbb{E}\left[\pi_{ht}(\mathcal{S}_k^{ht})\right] - \sum_{j \in \mathcal{S}_k^{ht}, j \notin \mathcal{S}_0^{ht}} sc_{hjt} \tag{16}$$

where sc_{hjt} is the cost of hospital h searching product j at time t. Note that $k_{ht} = |\mathcal{J}_{ht}|$ and $\mathcal{J}_{ht} = \mathcal{S}_{k_{ht}}^{ht}$.

We estimate search costs using the following simulated method of moments procedure: First, we obtain all demand and bargaining parameters for each hospital-product-month. Next, for each simulation iteration and set of parameters, we draw search cost values sc_{hjt} from the distribution

$$\Gamma(\alpha, e^{X_{hjt}^{sc}\phi}). \tag{17}$$

In the current estimates, we use $\alpha = 7$, which fits the data well in practice. We allow search costs to be a function of the several variables capturing visibility of products to hospitals over time: *product* visibility (the percent of hospital-months in the previous three months observed to purchase hospital j in *any* quantity) and *vendor-hospital* visibility (the total spend observed for the given hospital on the current product's vendor *in other product categories* in the last three months).

Then, for each ht, we order all products in J_{ht} by iteratively finding j such that:

$$j = \arg\max_{j' \in \mathcal{J} \setminus \mathcal{S}_k^{ht}} \left(\mathbb{E}\left[\pi_{ht}(\mathcal{S}_k^{ht} \cup j')\right] - \mathbb{E}[\pi_{ht}(\mathcal{S}_k^{ht})]\right) - sc_{hj't}. \tag{18}$$

In this expression, we simulate the expectation using draws of demand shocks ξ_{hjt} from the normal distribution $\mathcal{N}(0, \sigma_\xi^2)$, and relative bargaining weight shocks ν_{hjt} from the lognormal $ln\mathcal{N}(0, \sigma_\nu^2)$. Hospital h then chooses the consideration set k_{ht} at time t if

$$k_{ht} = \arg\max_{k \leq |\mathcal{J}_{ht}|} \mathbb{E}\left[\pi_{ht}(\mathcal{S}_k^{ht})\right] - \sum_{j \in \mathcal{S}_k^{ht}, j \notin \mathcal{S}_0^{ht}} sc_{hjt}. \tag{19}$$

For every draw s, we simulate a consideration set $\widetilde{CS}_{hts}$ for each hospital-month. This yields a set of dummy variables for each product $j \in \mathcal{J}_t$ appearing in each simulated choice set:

$$\widetilde{Pr}_{hjt}(\phi, \sigma_\xi, \sigma_\nu) = \frac{1}{S} \sum_s \mathbb{1}(j \in \widetilde{CS}_{hts}). \tag{20}$$

Finally, the model is estimated using a simulated method of moments estimator where we set

$$\sum_{hjt} \left[\widetilde{Pr}_{hjt}(\phi, \sigma_\xi, \sigma_\nu) - \mathbb{1}(j \in \mathcal{J}_{ht}) \right] z_{hjt} = 0. \tag{21}$$

The set of instruments is simply a vector of hospital dummies, product dummies, and product time trends.

The search costs are identified by the following intuition: We model the search process as directed, so that hospitals search products with high expected added value first. Search costs must then rationalize which products are included vs. excluded from hospitals' consideration sets. Mean search costs are identified by the distribution of the added value foregone when a hospital chooses to stop searching, as a function of total surplus of products already searched and added value of remaining products not searched. For example, if the expected additional surplus from adding the k^{th} product to a consideration set containing $k-1$ products is high, then we must infer that the average search cost is high.

5 Quantifying Sources of Price Dispersion

Table 2: Model Estimates

	Stents		Staplers		Gloves	
	Mean	SD	Mean	SD	Mean	SD
$\theta^p(\frac{utils}{\$1000})$	-0.49		-0.73		-2.49	
η	-0.65	0.13	-0.34	0.20	-0.27	0.19
AV	3,110	284	1,433	89	425	33
$mc(\$)$	900	0	15	3	12	1
β_{jh}	0.26	0.09	0.33	0.19	0.24	0.19
$sc(\$)$	83	36	264	112	28	11
$sc(\%)$	0.06		0.71		0.31	

5.1 Demand estimates

Table 2 shows the estimated parameters from our demand, pricing, and search models. The top panel shows, for each product category, the estimated utility weights on price (θ^p), the means and standard deviations of the price elasticity (η) and the total surplus under negotiation across all hospital-product-months.[8]

The first object of interest is the degree of preference heterogeneity across product categories. Figure 2 shows the dispersion of the estimated mean utility parameters θ_{hj} for each product. In order to generate a measure that is comparable across product markets, we compare the dispersion in preferences across hospitals within product to the dispersion of preferences across product. Products j are ordered from left to right in each plot by the median of θ_{hj} across h; in order to compare product categories side by side, the medians of each set of θ_{hj}'s are scaled to lie on the range [0,1] (zero representing the minimum, one representing the maximum), while the distributions of θ_{hj}'s across products are centered around those scaled medians. The $10^{th} - 90^{th}$ percentile ranges are shown, as are the interquartile ranges. The main takeaway from Figure 2 is that there is substantially more variation in

Figure 2: Estimated Preference Heterogeneity in θ_{hj}

(a) Stents **(b)** Staplers **(c)** Gloves

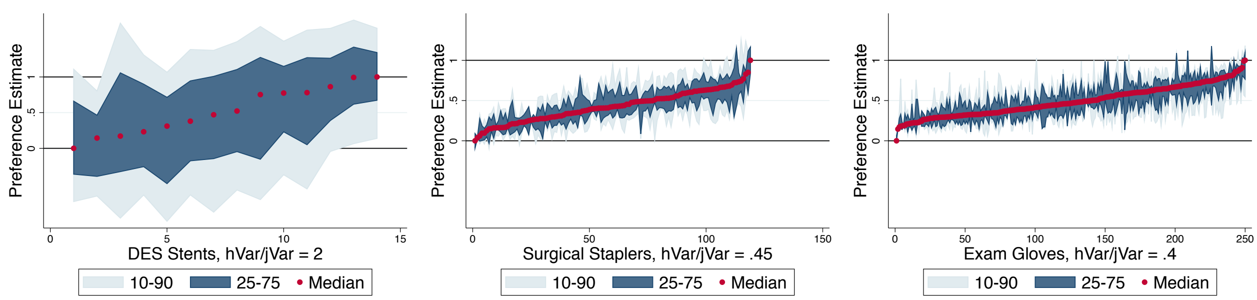

preferences across hospitals for all DES than there is for either staplers or gloves. This is consistent with our prediction that brand preferences are stronger determinants of price variation in stents (physician preference items) than in other products.

Another way of considering how preference parameters vary (and, ultimately, determine observed price dispersion) is to plot the distribution of the added value for a given product to a given hospital's choice set. Figure 3 plots a histogram of added value estimates across product-hospital-months. Note that added value captures information about demand for the items in the choice set relative to the other items, and thus is particularly large when the choice set contains a single option, in which case the relative comparison is to the outside

[8]The results of our reduced form search model first stage are currently omitted. While our proposed instruments (capturing visibility of products and vendors across hospitals and time) were found to be powerful shifters of consideration sets, the coefficient on the estimated search residual (or on a flexible functional form of the residual) was found to be highly dependent on functional form and weighting/restrictions on the estimation sample using the propensity score. In ongoing work, we are searching for a more robust control function approach.

good. The relative comparison of the added values distributions across product categories are

Figure 3: Added Values Estimates $(1 + \frac{\partial s}{\partial p}\frac{p-mc}{s})(\frac{\pi_h(\mathcal{J}_h)-\pi_h(\mathcal{J}_h \backslash j)}{s}) + p - mc$

(a) Stents **(b)** Staplers **(c)** Gloves

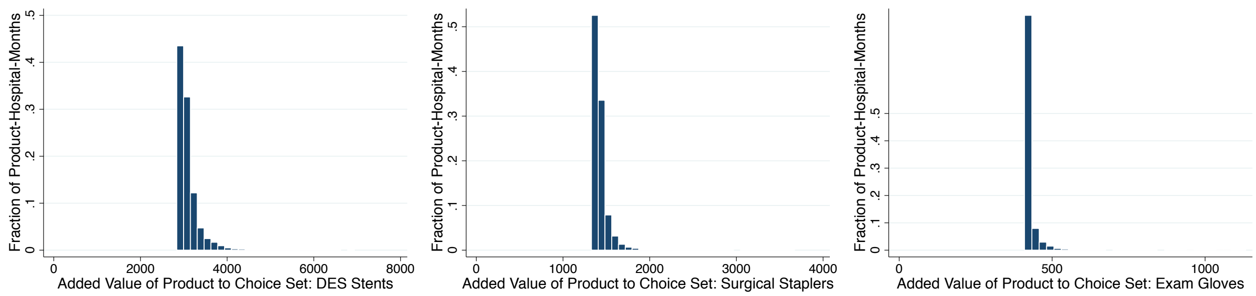

roughly similar to what we saw in comparing the price distributions in Figure 1, but in each case the mass of the distribution is shifted substantially to the right. This is consistent with our expectation that manufacturers do not extract all surplus from hospitals in bargaining. We explore the extent to which surplus is captured by manufacturers below in our analysis of relative bargaining weights.

5.2 Demand+supply estimates

Histograms of the price elasticities implied by the model parameters are shown for each product category in Figure 4. Elasticities are highest for stents (and similar on average to elasticities found in Grennan 2013) and lower for both gloves and staplers. The estimated

Figure 4: Price Elasticities $\frac{\partial q_{hj}}{\partial p_{hj}}\frac{p_{hj}}{q_{hj}}$

(a) Stents **(b)** Staplers **(c)** Gloves

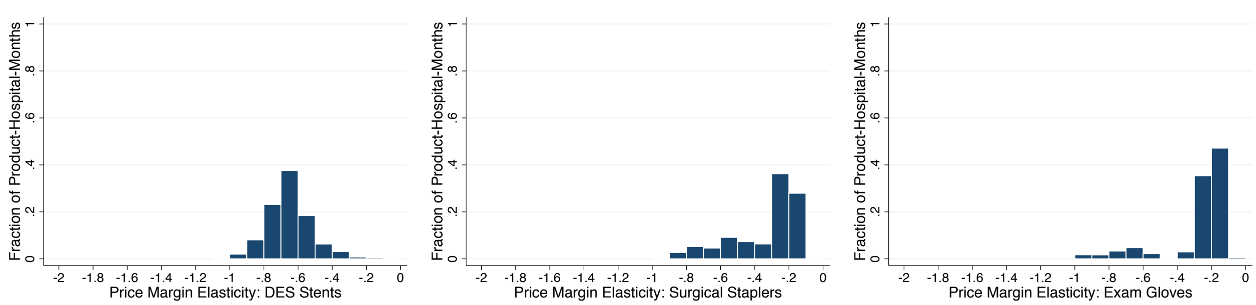

elasticities are quite low for all products. For stents and staplers, this finding is consistent with low price-sensitivity in usage for physician preference items; for gloves, it is perhaps surprising given that gloves are closer to a commodity product. However, it is important to note that, unlike in typical differentiated products models, elasticities in this setting are a function of demand, competition, and bargaining, and accordingly, the small elasticities shown could reflect low buyer price-sensitivity, low relative manufacturer bargaining ability, or both. Indeed, the estimated price coefficient for gloves shown in Table 2 was much *larger* than for stents or staplers. We next turn to estimating relative bargaining weights.

The bargaining split distribution is shown for each product category in Figure 5. These plots show some intriguing patterns. First, the distribution of relative bargaining weights of manufacturers for staplers is shifted higher than for stents or gloves, indicating that stapler manufacturers capture more surplus on average than stent or glove manufacturers. Second, the bargaining split distribution for gloves is bimodal with a large mass near zero and another mass around 0.7 – some glove manufacturers appear to be close to perfectly competitive, but others capture substantial surplus. The latter fact in part explains the quite low price elasticities for gloves in spite of our prediction that gloves are fairly homogenous and relatively price-sensitive within the set of hospital supplies.

Figure 5: Estimated Bargaining Splits $\frac{b_{jt}(h)}{b_{jt}(h)+b_{ht}(j)}$

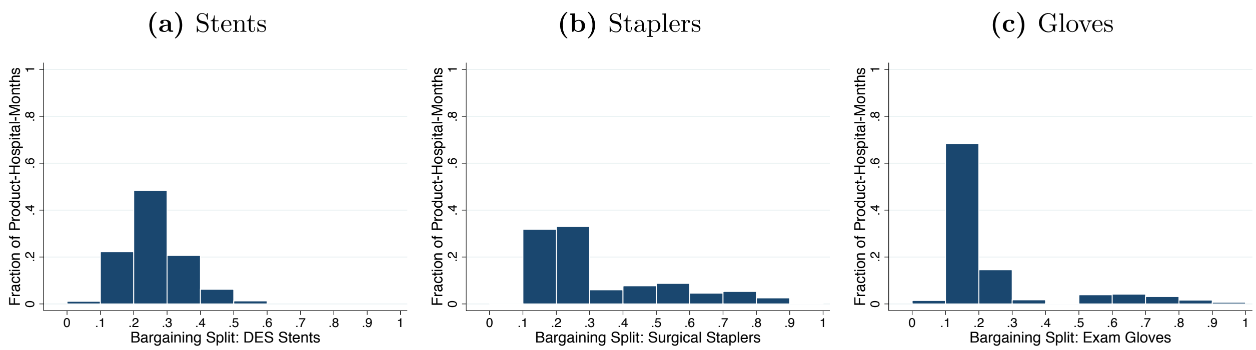

5.3 Search estimates

Figure 6 shows the distribution of choice set sizes across hospitals and months. In each product category, the model currently performs quite well in matching the shape of consideration sets. The estimated search model parameters map into our ex ante intuitions regarding these different product categories nicely as well. As we see in Table 2, as a percentage of product price, search costs are lowest for stents and highest for staplers. Thus, the potential for facilitated search to improve welfare seems greatest for products like gloves and staplers where search frictions are an important determinant of prices, and fairly small for products like stents where price heterogeneity is driven to a greater extent by preferences and bargaining.

Figure 6: Histogram of $|\mathcal{J}_{ht}|$

(a) Stents **(b)** Staplers **(c)** Gloves

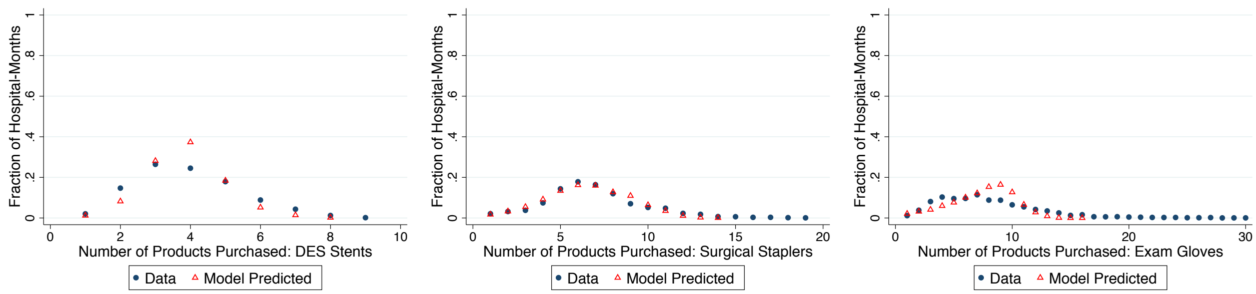

6 Information Interventions

The existence of price dispersion requires market power, and so understanding the sources of price dispersion is important per se in understanding the the economic frictions which give rise to this market power. Understanding these frictions is critical for understanding measures that might mitigate them. One such measure has been the rise of information technology that provides tools to facilitate product search and comparison. In the retail space, tools that indicate "shoppers who liked A also liked B" and "this item available from N other sellers at these prices" have become commonplace. With the increasing amount of data being collected by intermediaries in hospital purchasing and other business-to-business markets, there are a number of analytics that might potentially assist buyers in their purchasing efforts.

Here we consider two types of information interventions: one that could affect search across products/vendors, and one that could affect bargaining outcomes conditional on search. The latter is motivated by the price benchmarking information that is currently provided by the data source for this project. Grennan and Swanson (2016) used staggered join dates of hospitals to this service to estimate that access to the information provided by the service leads to price reductions among those previously paying relatively high prices for these products, and further find evidence consistent with an asymmetric information bargaining mechanism where hospitals with benchmarking information changed their estimates of supplier bargaining parameters $b_j(h)$. We implement this idea in our structural model here by holding all parameters fixed except for hospital bargaining parameter estimates, where we adjust all estimates above the average for a given product down to the average for that product and recompute equilibrium outcomes.

Motivated by the fact that the hospital benchmarking data also could provide information on other vendors selling the same or similar products, we also consider an intervention that would affect search costs. We compute equilibrium outcomes where, in addition to the bargaining information effect, search costs are reduced to one half their estimated. Besides their potential to impact prices paid, the search information interventions are interesting in their ability to change equilibrium choice sets and total surplus generated—in principle, if lower search costs induce better demand choices, they could lead to increased prices, but also increased surplus through use of better products. Our estimated results of these interventions are summarized in Table3.

Table 3: Counterfactual Estimates: Information Interventions

	DES Stents		Surg. Staplers		Exam Gloves	
	Δp	$\Delta \pi$	Δp	$\Delta \pi$	Δp	$\Delta \pi$
Bargaining Info $b_{hj} \searrow \overline{b_j}, \mathcal{J}_h(sc_{hj})$	-3.4	2.2	-5.6	1.3	-1.5	0.3
Barg. + Search Info $b_{hj} \searrow \overline{b_j}, \mathcal{J}_h(sc_{hj} \searrow \frac{sc_{hj}}{2})$	-3.1	38.1	10.2	49.0	0.1	50.0

The results show how the returns to different types of information depend a great deal on the market structure and frictions present in each product market. The price effects of bargaining information are all negative, but their magnitude depends on the size or price-cost margins as well as the estimated amount of variation in bargaining outcomes. The impact of this type of information is non-trivial but also modest, representing an increase of 0.3-2.2 percent in surplus to the hospital.

By contrast, information that cuts search costs in half is estimated to increase hospital surplus by 38.1-50.0 percent, and the pattern of this impact is greatest in the areas of staplers and gloves, where the product-vendor space is largest and search was estimated to be a larger friction in the market. It is important to note that these (preliminary) estimated effects on mean surplus are largely driven by a tail of hospitals with extremely large gains, with more modest effects on the median hospital.

7 Conclusion

The pervasive failure of the Law of One Price across a multitude of heterogeneous product markets is difficult to ascribe to any single factor due to the massive heterogeneity across markets in preferences, supply factors, information, and the practical realities of contracting. We explore this issue using data from a large and policy-relevant market: hospital supply contracting. In spite of the fact that the demand side in this application solely includes hospitals, there is large variation across products in the preferences of end users, the concentration and bargaining power of suppliers, and the potential importance of information frictions. In this first set of results, we present estimates from three product markets that vary in their supply, demand, and information characteristics: coronary stents, surgical staplers, and exam gloves.

As has been demonstrated in a rich empirical industrial organization literature (and related literatures in labor and insurance), the presence even of any two of the three above mentioned factors generates difficulties for modeling and identification. Our approach models the market in three stages: a search stage in which hospitals rank products and choose a consideration set that includes all products whose incremental contribution to surplus exceeds their search costs; a bargaining stage in which prices are negotiated in a set of simultaneous Nash bargaining games for each product in a hospital's consideration set; and, finally, a demand stage in which products are chosen for each end use case as a function of negotiated prices and preferences over each product in the consideration set.

The results of our empirical approach are enlightening. First, they demonstrate the importance of accounting for each of the three model features (search, bargaining, and demand) in order to understand equilibrium price dispersion. For example, we anticipated ex ante that exam gloves would be our most commodity-like product category and thus that markups would be limited. Indeed, the price coefficient in the demand equation is quite large for exam gloves. Instead, we find evidence of substantial price dispersion in this market, which is somewhat perplexing, but ultimately can be attributed to high relative bargaining power of some glove manufacturers and large search costs relative to glove prices.

Second, our approach allows us to perform counterfactual analyses to demonstrate how the effects of interventions that affect these market frictions will manifest unevenly across product markets. Motivated by the presence of benchmarking databases and policy proposals focused on greater transparency in medical device markets, we consider the effects of information that increase hospitals' relative bargaining abilities and also hospital search costs in the product-vendor space. Consistent with previous research, information that affects bargaining results in worthwhile but perhaps modest reductions in prices, with a larger impact in the stent market and smaller impact in the exam glove market. Interestingly, we find the impact of information that reduces search costs to be larger in magnitude and the opposite direction across markets, with a larger impact on staplers and gloves. These results are driven by the relative importance of economic factors across product markets; in particular, price heterogeneity in stents is driven by preferences and search costs are of less importance, while for gloves and staplers search costs are relatively large.

Appendices

A Data Appendix

References

Allen, J., Clark, R., and Houde, J.-F. (2013). The effect of mergers in search markets: Evidence from the canadian mortgage industry. *American Economic Review*, 104(10):3365–3396.

Berry, S. T. (1994). Estimating discrete-choice models of product differentiation. *RAND Journal of Economics*, 25(2).

Chade, H. and Smith, L. (2006). You have full text access to this content simultaneous search. *Econometrica*, 74(5):1293–1307.

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

Dinerstein, M., Einav, L., Levin, J., and Sundaresan, N. (2014). Consumer price search and platform design in internet commerce. Revise and resubmit, Review of Economic Studies.

Einav, L., Finkelstein, A., Ryan, S., Schrimpf, P., and Cullen, M. R. (2013). Selection on moral hazard in health insurance. *American Economic Review*, 103(1):178–219.

Gaynor, M., Propper, C., and Seiler, S. (2016). Free to choose? reform, choice, and consideration sets in the english national health service. forthcoming in the American Economic Review.

Goeree, M. (2008). Limited information and advertising in the u.s. personal computer industry. *Econometrica*, 76(5):1017–1074.

Goldberg, P. and Verboven, F. (2001). The evolution of price dispersion in the european car market. *Review of Economic Studies*, 68(4):811–848.

Gowrisankaran, G., Nevo, A., and Town, R. (2015). Mergers when prices are negotiated: Evidence from the hospital industry. *American Economic Review*, 175(1):172–203.

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

Grennan, M. and Swanson, A. (2016). Transparency and negotiated prices: The value of information in hospital-supplier bargaining. Under review.

Heckman, J. (1979). Sample selection bias as a specification error. *Econometrica*, 47(1):153–162.

Hortacsu, A. and Syverson, C. (2004). Product differentiation, search costs, and competition in the mutual fund industry: A case study of s&p 500 index funds. *Quarterly Journal of Economics*, 119(2):403–456.

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

Moraga-Gonzalez, J. L., Sandor, Z., and Wildenbeest, M. (2015). Consumer search and prices in the automobile market. Tinbergen Institute Discussion Paper.

Petrin, A. and Train, K. (2009). A control function approach to endogeneity in consumer choice models. *Journal of Marketing Research*, xlvi.

Roberts, M., Xu, D. Y., Fan, X., and Zhang, S. (2016). The role of firm factors in demand, cost, and export market selection for chinese footwear producers. Revise and resubmit, Review of Economic Studies.

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

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

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

Starc, A. and Swanson, A. (2016). What do private firms bring to the (bargaining) table? vertical structure and negotiated prices of generic pharmaceuticals. Working paper.

Stigler, G. (1961). The economics of information. *Journal of Political Economy*, 69(3):213–225.