Price Discrimination and Bargaining: Empirical Evidence from Medical Devices

Matthew Grennan*
University of Toronto, Rotman School of Management
matthew.grennan@rotman.utoronto.ca

October 26, 2010

Abstract

Many important issues in business-to-business markets involve price discrimination and negotiated prices, situations where theoretical predictions are ambiguous. This paper uses new panel data on buyer-supplier transfers and a structural model to empirically analyze bargaining and price discrimination in a medical device market. While many phenomena that restrict price discrimination are suggested as ways to decrease hospital costs (e.g., mergers, group purchasing organizations, and transparency), I find that: (1) non-discriminatory pricing actually works against hospitals because competition is more intense under price discrimination; and (2) results depend ultimately on a previously unexplored "bargaining effect".

*Data for this research was provided by Millennium Research Group (www.mrg.net). Special thanks for the guidance and encouragement throughout the many phases of this research project go to my dissertation committee: John Asker, Adam Brandenburger, Allan Collard-Wexler, Michael Katz, and Bernie Yeung. I also benefited from discussions with Heski Bar-Isaac, Luís Cabral, Ken Corts, Catherine de Fontenay, JP Eggers, Robin Lee, PT Leger, Hong Luo, Matt Mitchell, Gabriel Natividad, Ariel Pakes, Rob Salomon, Gus Stuart, Ali Yurukoglu, and seminar audiences at Analysis Group, Arizona, Columbia, Georgia, LSE, NYU, Toronto, and Yale. Andy Cron, John Brumleve, Mark Bly, Rob Tumarkin, and many others have been helpful in teaching me about medical device manufacturing and sales. Physicians and administrators at NYU Medical Center and CGH, Sterling, IL were especially generous with their time and knowledge regarding hospital purchasing and the inner workings of a cardiac catheterization lab. Some industry background research was conducted while attending the Transcatheter Cardiovascular Therapeutics 2008 conference, which was made possible by the Cardiovascular Research Foundation.

1 Introduction

In business-to-business markets, different buyers often pay substantially different prices for the exact same product from the same supplier. As a result, price discrimination plays a role in many prominent phenomena. For instance, mergers between buyers force the seller to set a single uniform price to the merged entity where it had previously set different prices to each buyer. Group purchasing organizations, where a group of buyers delegate purchasing authority to a third party, also enforce a single price across buyers. The degree of transparency in a market can also indirectly influence the ability of sellers to price discriminate (Armstrong 2006). Motivated by these phenomena, this paper estimates the welfare effects of price discrimination in the market for coronary stents, a "blockbuster" medical device on which hospitals spend over $5 billion each year.

The price of medical technologies, such as coronary stents, is often cited as one driver of the increasing costs of healthcare. Some of the most commonly discussed policy remedies for controlling healthcare costs have involved precisely the three phenomena noted above: One common justification for hospital mergers is to lower costs. Group purchasing organizations permeate most hospital purchasing, but they play a small role in the market for coronary stents and other "physician preference items". The *Transparency in Medical Device Pricing Act of 2007*, introduced in the U.S. Senate Committee on Finance on October 23, 2007, and other healthcare transparency initiatives have been the topic of a spirited debate in industry trade journals, with many anticipating that such reform would lead to more uniform prices across hospitals (Kyle and Ridley 2007), but no consensus regarding whether this would benefit hospitals.

The lack of consensus regarding the impact of hospital mergers, group purchasing, and transparency on the cost of healthcare inputs is driven at least in part by the fact that economic theory alone offers ambiguous predictions. This paper uses new panel data on the prices and quantities transferred between hospitals and coronary stent manufacturers to estimate a structural model of supply and demand. I then use the estimated model to compare welfare outcomes under the observed price discrimination regime and several different non-discriminatory counterfactual regimes. I find that: (1) in this market, non-discriminatory pricing actually works against hospitals because competition is more intense under price discrimination; and (2) results depend ultimately on a previously unexplored "bargaining effect".

The first result relates directly to the theoretical literature on price discrimination.[1] With oligopoly, whether suppliers will capture more surplus with or without price discrimination (the main focus of the policy questions in medical devices) depends

[1] Stole (2007) offers an excellent review of this large literature. I discuss some specific results in the last section of this paper.

on whether there is *symmetry* or *asymmetry* regarding which buyers are "strong" and "weak" markets (Corts 1998). For example, a buyer with demand for lots of quantity makes it stronger for all suppliers, but with transportation costs, a buyer who is close to one supplier and far from another will be a strong market for the close supplier but a weak one for the far supplier. Both the final results and reduced-form evidence in this paper suggest that there is more asymmetry than symmetry in demand for stents across hospitals, leading to competition being more intense under price discrimination than under uniform pricing. However, this is only part of the story because this literature does not speak to situations where prices are negotiated—as they are in the coronary stent market—and it is only through allowing for bargaining that the second result can be seen.

Incorporating negotiated prices into the analysis is the central challenge in this paper, but it is important because in business-to-business markets, price discrimination tends to go hand-in-hand with bargaining over prices. The forces that determine negotiated prices are related to, yet different from, the forces that determine posted prices. When buyers and suppliers negotiate prices, supplier costs, buyer willingness-to-pay, and competition determine only a range of potential prices (versus a single price) for each buyer and supplier.[2] The final negotiated price will depend on what I refer to as each firm's *bargaining ability*—the firm's ability to reach a more favorable point within the range determined by costs, demand, and competition. Allowing for heterogeneity in firm bargaining abilities turns out to be important in explaining the variation in prices for the same product sold to different buyers—as one hospital purchasing manager put it, "There is a lot of wiggle room [in prices]." Heterogeneity in bargaining ability also underlies the "bargaining effect" in thinking about price discrimination versus uniform pricing: In the case of mergers and group purchasing organizations, the bargaining ability of the "merged" buyers relative to those of the individual buyers will play a key role. And in the case of non-discrimination through transparency reforms or direct regulation, the market could even change to one with posted prices.

Despite the theoretical ambiguity, the empirical literature on price discrimination is still relatively small. This is largely because empirical studies of price discrimination (and bargaining) in business-to-business markets have been limited by the difficulty of accessing data on transfers between buyers and suppliers. Of the recent empirical studies involving price discrimination (Hastings 2008; Villas-Boas 2009), bargaining (Dranove et al. 2008; Dafny 2009; Ho 2009; Crawford & Yurukoglu 2010), and vertical contract-

[2]The simplest example is bilateral monopoly, where the buyer won't pay a price above its willingness-to-pay, and the seller won't sell for a price less than its marginal cost. With a competing supplier, the buyer has an outside option that lowers the top of this range, but as long as the competing product is not a perfect substitute, there will still be a range of prices at which the buyer and supplier could trade.

ing relationships more generally (Asker 2004; Ho, Ho, & Mortimer 2010), only Hastings (2008) and Dafny (2009) have had access to data on the actual buyer-supplier transfers. Hastings (2008) studies gasoline stations and wholesalers and does not consider bargaining. Dafny (2009) studies employers purchasing health insurance and is interested in diagnosing the presence of market power among insurance providers, not analyzing bargaining or price discrimination per se. Lack of access to data on buyer-supplier transfers requires the researcher to either assume some fixed level of bargaining abilities across market participants (Dranove et al. 2008; Ho 2009) or to use downstream data and a structural model of downstream demand and competition to infer what these transfers must have been (Villas-Boas 2009; Crawford & Yurukoglu 2010).

This paper uses an unusually detailed panel data set, providing the quantities purchased and prices paid for all coronary stents sold to 96 U.S. hospitals from January 2004 through June 2007, at the stent-hospital-month level of observation. The stent market is a business-to-business market in which hospitals generate revenue by implanting stents during angioplasty procedures, and the stent is a necessary input that the hospital must purchase from a device manufacturer. Contracts are negotiated, stipulating the price at which the hospital can purchase a given stent over a specified period of time, and different hospitals negotiate different prices for the exact same stent. The price variation across buyers has significant implications for firm profits: moving from the 25th to 75th percentile in price would result in a change of about $300,000 annually (about four nurses' salaries) at the average-sized hospital.

Even with these detailed data, several important variables—cost, willingness-to-pay, and bargaining ability—are unobserved. Further, separating the impact of competition on the range of potential prices from the impact of bargaining abilities within that range requires an explicit model of how competition and bargaining determine prices. I address these challenges using a structural empirical approach, similar to Berry, Levinsohn, & Pakes (1995) and subsequent literature, but using a pricing model that generalizes the standard Bertrand-Nash price-setting model to allow for bargaining over prices. The model has two parts: (1) a model of doctor demand for coronary stents uses the price and quantity data to estimate demand for each stent in each hospital in each month; and (2) a model of how prices emerge from competition and bargaining uses the demand estimates and the price and quantity data to estimate costs and relative bargaining abilities for each stent in each hospital in each month.

On the demand side, a random coefficients discrete choice model incorporates heterogeneity in preferences for stents across hospitals, physicians, and patients. The fact that prices are fixed in long-term contracts provides a new source of identification for demand. Doctor preferences evolve over time while prices remain fixed; thus when prices

are renegotiated, the movement is along the demand curve, analogous to a supply shift. The demand estimates agree with anecdotal evidence that doctors at different hospitals can differ widely in their preferences over the different stents available on the market.

Given the demand estimates, I estimate cost and bargaining ability parameters using a pricing model in which each stent manufacturer and hospital engage in a bilateral Nash Bargaining problem, and all of these bilateral outcomes must be in a Nash Equilibrium with each other at a given hospital. This model relates to the literature on bargaining with externalities (Horn & Wolinsky 1988), and Crawford & Yurukoglu (2010) use a close variant in an empirical setting. One contribution of this paper is to solve the model for the equilibrium "pricing equation" which is useful in two ways: First, it clarifies how this bargaining model is a generalization of the standard Bertrand differentiated products pricing model, connecting the model to all the previous theoretical literature on differentiated goods pricing and price discrimination. Second, the pricing equation shows that price is equal to marginal cost plus a margin that depends on bargaining abilities, elasticities, and the marginal contribution of each product relative to its competitors, making it clear how covariation in price and demand identify bargaining ability parameters separately from marginal costs. The estimates show that allowing for bargaining—and for heterogeneity in bargaining abilities—is critical for explaining the price variation observed in the data.

Finally, I use the estimated model to simulate the effects of a change to uniform pricing under different specifications related to hospital mergers, group purchasing organizations, or transparency initiatives. As previewed above, there are two main forces at work: (1) a "competitive effect", which ties in closely with the existing theory of price discrimination; and (2) a "bargaining effect", which is new.[3] The competitive effect can be seen by holding the effect of bargaining abilities fixed, in which case a change to uniform pricing would lead to a 9% increase in average prices, making hospitals worse off than they were under price discrimination.

However, predictions about the uniform pricing outcome depend critically on whether hospitals are able to bargain collectively and, if so, at what bargaining ability. If hospitals are unable to bargaining collectively, prices nearly triple. However, if the "merged" hospitals under uniform pricing have a bargaining ability above the 85th percentile of the distribution of individual hospital bargaining abilities, then this is enough to overcome the weakened competition and decrease prices relative to price discrimination. Taken together, these results cast some doubt on whether hospital mergers, group purchasing, or transparency initiatives would lower hospital costs for devices such as coronary stents.

[3]The closest analog to this result is a working paper by Myatt & Rasmusen (2009). In their model, a monopolist decides whether to set a uniform price or bargain over buyer-specific prices. They show that a monopolist may not wish to price discriminate if buyers' bargaining power is high enough.

It is worth mentioning that this structural approach is not the only way to address the question of the welfare effects of price discrimination. For example, one could instead look at exogenously occurring mergers and their impact on device prices. However, even if both the detailed price data and plausibly exogenous mergers could be combined into a single study (unfortunately, hospitals are anonymous in the data used here), this would not be able to differentiate empirically or conceptually between the roles of competition and bargaining ability under the different market structures. This tight link between the theory and empirics is an advantage of the approach used here.

While the model and estimation approach used here are flexible enough to be applied to other settings, the credibility of any structural study depends on capturing important industry-specific details. The next section presents the data and industry facts that motivate the model details in this paper. Section 3 lays out a simple version of the model, solving the pricing model and illustrating how it is related to more established models. Section 4 extends the model to capture the important sources of heterogeneity in the data and discusses identification and estimation of the model parameters, including how identification with bargaining is different from when suppliers set prices. Section 5 presents the resulting model parameter estimates and how they explain the observed variation in prices and market shares for the same stent across hospitals. Section 6 uses the estimated model to simulate counterfactuals with non-discriminatory pricing that provide insight into the potential effects of mergers, group purchasing, and transparency.

2 Coronary Stents: Industry Description and Data

The coronary stent industry is not only an interesting example of a business-to-business market, it is also interesting and important in and of itself. The coronary stent is a medical device used in angioplasty, an important treatment for blockages in the arteries surrounding the heart (a condition known as coronary artery disease). These blockages can cause pain, loss of mobility, and eventually heart attack, making coronary artery disease the leading cause of death in the United States.[4] Angioplasty is a minimally invasive technique in which the doctor threads a balloon-tipped catheter from an access point (usually the femoral artery near the groin) to the heart. Using imaging devices, the doctor positions the balloon tip across the blockage, and expands the balloon, compressing the blockage to the artery walls. A stent is a small metal tube that is then placed via catheter where the blockage was cleared and left in the body as structural support

[4] U.S. Department of Health & Human Services, National Institutes of Health, National Heart Lung and Blood Institute Diseases and Conditions Index. www.nhlbi.nih.gov/health/dci/Diseases/Cad/CAD_WhatIs.html (June 2008)

for the damaged artery wall. Though angioplasty is attractive due to its minimally invasive nature, traditional stainless steel "bare-metal stents" (BMS) have the drawback that scar tissue growth around this foreign body can lead to significant renarrowing of the artery in about 33% of cases. "Drug-eluting stents" (DES) attempt to remedy this problem by coating the stent with a drug that discourages scar tissue growth, and they have been successful in reducing the incidence of renarrowing to about 9%.[5]

2.1 The "Economics" of the Stent Market

With the introduction of DES, stents became the first medical device to reach revenue levels similar to those of a "blockbuster" drug. The three million stents implanted worldwide each year generate annual revenues of more than $5 billion to stent manufacturers and $30 billion to hospitals and doctors for the stenting procedures.

Hospitals and doctors generate revenue from each angioplasty procedure, usually via reimbursement from a patient's insurer. The reimbursement rates are negotiated by the hospital with each insurer (usually taking Medicare rates, which are not negotiated, as a starting point), so they vary across hospitals and across insurers for each hospital.[6] The average Medicare reimbursement rates for a basic stenting procedure are $812 for the doctor, regardless of the type of stent used; and for hospitals, $10,422 for a BMS and $11,814 for a DES.[7] Reimbursements do not depend on the manufacturer of the stent.

Out of this revenue comes the hospital's costs, including the cost of any stents used. Thus the hospitals keep in profit any price savings they can achieve on the cost of stents. While in many markets there might be some interaction between the costs negotiated with suppliers and the revenues negotiated from buyers, that is not the case here. For Medicare patients, who receive over 50% of all stenting procedures, the reimbursement levels are fixed; and the reimbursements from private insurers are generally negotiated as a markup on Medicare rates across all procedures performed at the hospital. Thus reimbursement levels at each hospital are fixed with respect to the cost of stents.

[5] Robert J. Applegate, "Drug-Eluting Stents: The Final Answer to Restenosis?" Wake Forest University Medical Center, http://www1.wfubmc.edu/articles/CME+Drug-eluting+Stents (September 15, 2008)

[6] Because the data set used in this paper is sold as market research to the device manufacturers, hospitals are anonymous, which, unfortunately, prevents linking this data set with other data sources on the hospitals.

[7] By "basic" I mean single-vessel operations with no "modifiers" for difficulty of the procedure, location of the hospital, etc. These numbers represent the lower bound in revenue for these procedures (Medicare upper bounds are roughly 1.5 times these payments, and private insurers generally reimburse at even higher levels). Numbers from *Federal Register*, Volume 68, No. 216, November 7, 2003; and *Federal Register*, Volume 68, Number 148, August 1, 2003.

2.2 Data Overview

The data set used in this paper is from Millennium Research Group's *Marketrack* survey of catheter labs, the source that major device manufacturers subscribe to for detailed market research. The goal of the survey is to provide an accurate picture of market shares and prices by U.S. region (Northeast, Midwest, South, West).[8] The U.S. market is dominated by four large multinational firms: the Abbott Vascular (formerly Guidant) division of Abbott Laboratories, Boston Scientific, Johnson & Johnson's Cordis division, and Medtronic, which together make up over 99% of U.S. coronary stent sales.[9] These manufacturers offered a total of nine BMS and two DES during the sample period, January 2004 through June 2007. After removing hospitals whose reporting is incomplete and observations without a time dimension, the data set I use for analysis is an unbalanced panel of 14,245 stent-hospital-month observations at 96 U.S. hospitals over the 42 months from January 2004 through June 2007.[10]

Figure 1: Aggregate trends in the market over the sample period. The quantity graph shows the total number of stents implanted, also broken down into DES and BMS. The price graph shows median prices of BMS and DES (the thin lines are the first and third quartiles).

(a) Quantities **(b)** Prices

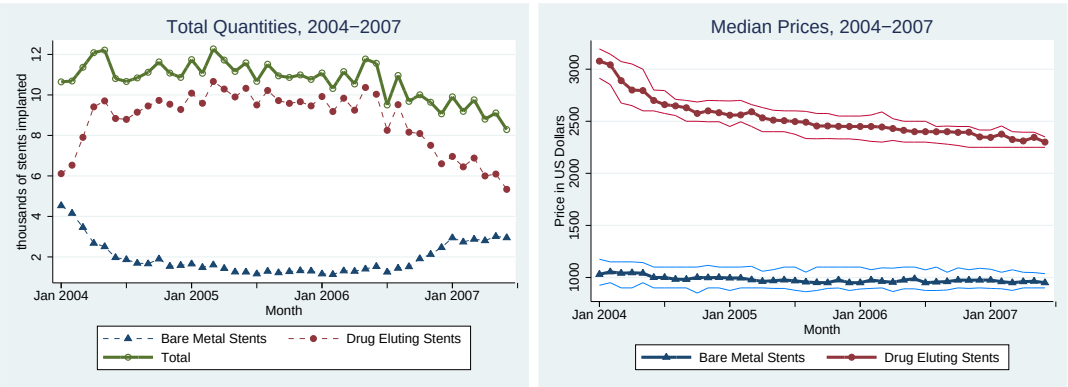

Figure 1 shows the aggregate trends in total quantities and median prices over this sample period. In March 2004, Boston Scientific introduced the second DES to the market, resulting in decreased prices and increased usage of DES. In 2006, some studies questioned the safety of DES, resulting in less DES usage and less stenting overall.

[8] See www.mrg.net for more details on the survey.

[9] iData Research, Inc. "US Markets for Interventional Cardiology" (2006).

[10] Details regarding the data set construction are given in Appendix B.

2.3 Price Variation and Potential Sources

Table 1 provides price summary statistics, documenting the variation in prices across hospitals. The coefficient of variation (standard deviation over mean), a common measure of price dispersion, has a mean of 0.13 in the sample. For example, one of the best-selling stents, DES1, has a mean price of $2508 with a standard deviation of $317.

Table 1: Price variation across hospitals. The table reports summary statistics for the distribution of price ($US) across hospitals for each stent. The sample is restricted to September 2005 (middle of the sample in time) to isolate cross-sectional variation. There are N=54 hospitals sampled, and BMS1-3 have exited the market.

	mean ($)	std. dev. ($)	std.dev./mean	min ($)	max ($)	N
BMS4	1006	175	0.17	775	1500	25
BMS5	926	191	0.21	700	1600	23
BMS6	952	156	0.16	775	1475	26
BMS7	1035	174	0.17	775	1600	39
BMS8	1063	338	0.32	800	1950	11
BMS9	1088	224	0.21	800	1800	47
DES1	2508	317	0.13	2450	3606	54
DES2	2530	206	0.08	2500	2900	54

These per-unit price differences translate into significant dollar amounts. A $300 change in price results in a difference in cost of about $300,000 per year in the mean-volume hospital, or nearly $1 billion per year across the three million stents implanted worldwide. This is about 20% of the annual revenue of the global stent market.

There are many potential explanations for this price variation across hospitals. Revenue for stenting procedures varies across hospitals. The relative strength of the interventional cardiologists versus substitute treatments and the distribution of patient types will vary across hospitals as well. Also, stents are differentiated products, and doctors vary in their preferences over which stent is best to treat a given patient. These variations induce different competitive environments in different hospitals. The variation in the market shares of each stent, the number of diagnostic procedures per hospital, and the frequency with which diagnostic procedures lead to stenting, displayed in Table 2 and Figure 2, all provide a sense of this demand heterogeneity.

Taking a closer look at the market share data also provides some preliminary evidence regarding the amount of "symmetry" in demand across hospitals, which theory suggests will play an important role in determining the effects of competition under price discrimination versus uniform pricing. Regressing the September 2005 (to isolate cross-hospital variation) market shares (percent of diagnostic procedures that are treated with each stent) on stent dummy variables, and then on stent and hospital dummy variables, reveals that hospital effects explain only 12% of the within-stent variation in market

Table 2: Market share variation across hospitals. The table reports summary statistics for the distribution of market share (% of all stents used) across hospitals. (Average shares do not add up to 100% because not all stents are used by all hospitals, as documented in the last column of the table.) The sample is restricted to September 2005 (middle of the sample in time) to isolate cross-sectional variation. There are N=54 hospitals sampled in this month, and BMS1-3 have exited the market.

brand	mean (%)	std. dev. (%)	std.dev./mean	min (%)	max (%)	N
BMS4	5	3	0.7	1	14	25
BMS5	3	2	0.6	1	7	23
BMS6	6	6	1.0	1	25	26
BMS7	4	5	1.1	1	25	39
BMS8	4	4	1.1	1	14	11
BMS9	8	8	1.0	1	32	47
DES1	43	30	0.7	1	88	54
DES2	41	30	0.7	2	93	54

Figure 2: Distribution of procedure volumes across hospitals. All patients must have a diagnostic procedure to locate any blockages and detect their severity. The graph on the left shows the distribution of the average number of these procedures each hospital performs per month. The graph on the right shows the distribution of the average percentage of these procedures that result in a stenting intervention. The table below contains summary statistics for the two distributions.

(a) Diagnostic angiographies **(b)** Percent resulting in stenting

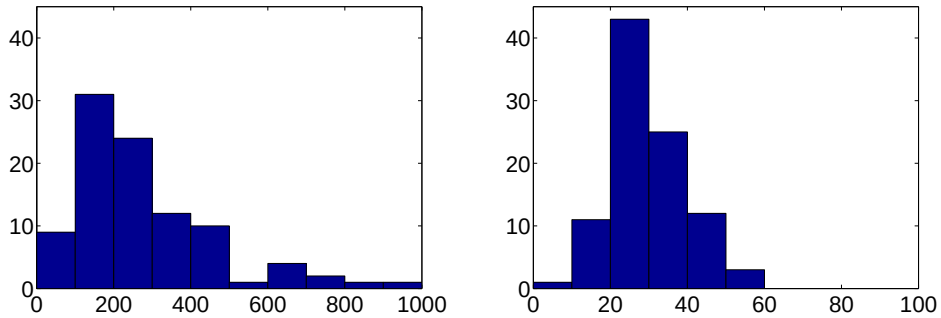

	mean	std. dev.	min	median	max	N
Diagnostic angiographies (procedures/hospital-month)	283	185	58	232	934	96
Share of diagnostic angiographies resulting in stenting (%)	28	9	5	26	52	96

shares. The remaining 88% is stent-hospital specific variation, suggesting more asymmetry than symmetry in demand across hospitals. This is an imperfect test with raw data, but it is a first piece of evidence that competition may be more intense under price discrimination than under uniform pricing.

Prices are usually negotiated directly between each manufacturer and each hospital, and these negotiations are another potential source of the observed price variation. The

typical contract is linear, specifying a price per unit for a given stent over the contract period, often one year.[11] Who is involved in the negotiation and the incentives they face differ across hospitals and manufacturers, and anecdotal evidence suggests that this could also be an important source of variation in the final price.

How much these forces influence price variation, and how they affect welfare with a change to uniform pricing, is ultimately an empirical question. Estimating the unobserved variables and disentangling their effects is the purpose of the rest of this paper.

3 The Theoretical Model

The goal of this section is to lay out the modeling framework in a parsimonious way, showing how the model captures the key features of business-to-business relationships in the coronary stenting industry and relating the model to previous theoretical and empirical work. The model predicts the quantities of each stent used by each hospital and the prices negotiated for each stent by each manufacturer-hospital pair. This section derives those predictions in terms of a set of parameters. Section 4 extends this model to a richer parameterization in order to capture the heterogeneity and dynamic nature of the data and to estimate the parameter values that match the predictions of the model to the values observed in the data.

The agents in the model are the device manufacturers who supply the products, the doctors whose decisions determine demand for those products, and the hospitals that negotiate prices with manufacturers. The model is a two-stage game with no information asymmetries:

Stage 1: Pricing Device manufacturers and hospitals contract on prices, taking expected future quantities into account.

Stage 2: Demand Given prices and choice sets, doctors decide on stent purchases as patients arrive at the hospital.

In what follows, I consider the problem of multiple device manufacturers selling to a single hospital. Under the maintained assumptions that hospitals are monopsonists of their own flow of patients and that manufacturer profits are separable across hospitals, this immediately extends to the empirical context in the next section, where each product

[11] Some contracts have discrete non-linearities, offering a lower price if the hospital uses that stent almost exclusively, say 80% of the time. While I do not observe the actual contracts, my empirical approach will allow for this possibility.

is sold to multiple hospitals.[12] Because the first stage pricing equilibrium depends on expected demand, I start from the second stage and work backwards.

3.1 Modeling Demand for Coronary Stents

I model demand using a discrete choice random utility model of how doctors choose which stent to use for each patient. This approach has the benefit of intuitively matching the actual doctors' decision process, and it accommodates the fact that the choice sets of available stents vary across hospitals and over time. It also allows for very flexibly shaped demand curves and the direct computation of consumer surplus measures (Nevo 2000).

The hospital has contracted with a set of stent manufacturers for the set of stent models $j \in \mathcal{J}$. Over the course of a month, patients $i = 1, ..., Q$ arrive at the hospital to undergo a diagnostic procedure. The doctor chooses a treatment for each patient to maximize the following indirect utility function according to the discrete choice problem:

$$\max_{j \in \mathcal{J}} \qquad u_{ij} = \theta_j - \theta^p p_j + \varepsilon_{ij}, \tag{1}$$

where θ_j is the mean quality of product j across all patients, θ^p is the marginal disutility of price p_j (in utils per dollar), and ε_{ij} is a stochastic patient-specific quality component with distribution $f(\varepsilon)$, representing characteristics of the specific patient/doctor combination i that make the patient an especially good candidate for a specific stent. This utility function can be thought of as a reduced form for how a doctor incorporates his own preferences, patient welfare, and hospital profitability into the treatment decision.[13]

The set $\mathcal{J}$ also includes a choice $j = 0$ for a treatment other than stenting, and I normalize $\theta_0 - \theta^p p_0 = 0$ so that the utility for each stent is the utility relative to the next best non-stent treatment.

Given this demand structure, define the set of patients for whom a doctor chooses product j as $\mathcal{A}_j := \{i | j = \arg\max_{k \in \mathcal{J}} u_{ik}\}$. Then expected market shares for each stent are given by the choice probability for each stent:

$$s_j = Pr[j = \arg\max_{k \in \mathcal{J}} u_{ik}] = \int_{\mathcal{A}_j} f(\varepsilon) d\varepsilon. \tag{2}$$

These expected market shares form the basis for the demand estimation procedure in Section 4, which assumes a parameterized distribution for ε and finds the parameter values that match the expected shares generated by the model to the shares observed in

[12] This separability rules out the cases of binding manufacturer capacity constraints and economies of scale in production.

[13] Blomqvist (1991) discusses this unique role of the doctor in healthcare.

the data.

The expected shares are also important because at the time of contracting, the exact set of patients that will show up is uncertain. So expected quantities for any given price vector $p = \{p_j\}_{j \in \mathcal{J}}$ are anticipated via expected market shares by

$$q_j(p) = s_j(p)Q. \tag{3}$$

Given expected demand, expected profits for stent model j at this hospital are

$$\pi_j(p) = (p_j - c_j)q_j(p), \tag{4}$$

and the expected payoff to the hospital is the total expected willingness-to-pay over treated patients minus payments to manufacturers:

$$\pi_h(p) = \sum_{j \in \mathcal{J}} q_j(p) \left(\int_{\mathcal{A}_j} wtp_{ij} f(\varepsilon) d\varepsilon - p_j \right), \tag{5}$$

where willingness-to-pay for patient i is the non-price part of utility scaled into dollars, $wtp_{ij} := \frac{\theta_j + \varepsilon_{ij}}{\theta^p}$. These expected payoffs are taken into account during the first stage of the model: pricing.

3.2 Modeling Pricing with Competition and Bargaining

I consider a model of bargaining and competition similar to that of Horn and Wolinsky (1988). The hospital negotiates with each manufacturer separately and simultaneously, with the outcome of each negotiation satisfying the bilateral Nash Bargaining solution (the product of the manufacturer and hospital payoffs). I impose consistency across the bilateral bargaining problems by requiring that the outcomes form a Nash Equilibrium in the sense that no party wants to renegotiate. Cremer and Riordan (1987) call such an equilibrium in bilateral contracts a *contract equilibrium*. Formally, prices are determined as a Nash Equilibrium of bilateral Nash Bargaining problems. Each bilateral price maximizes the Nash Product of manufacturer profits and hospital surplus, taking the other prices as given, solving

$$\max_{p_j} \quad [\pi_j(p)]^{b_j(h)} \, [\pi_h(p) - d_{jh}]^{b_h(j)} \qquad \forall j \in \mathcal{J}, \tag{6}$$

where the parameters $b_j(h), b_h(j) \geq 0$ represent the bargaining ability of the manufacturer and hospital vis-a-vis each other, respectively, and d_{jh} is the hospital's disagree-

ment payoff when no contract with j is signed.[14] In order to clearly make the links with better-known models, here I write the model with each product negotiated separately.[15]

A variation of this model (with multiple firms on both sides of the market) has been used in prior empirical work by Crawford and Yurukoglu (2010), and many related models have been developed in theoretical work on bilateral negotiations with externalities.[16]

This model "nests" the solutions to many other pricing models of interest. When $b_h = 0$, manufacturers set prices in a Bertrand-Nash price equilibrium. When $b_j = 0$, the manufacturer prices at cost. One point of interest is to see how far bargaining strays from these "take-it-or-leave-it" offer models.

Also, changing the hospital's threat point, d_{jh}, corresponds with different notions of bargaining found in the literature. Setting $d_{jh} = \pi_h(p; \mathcal{J} \setminus \{j\})$, as I do in this paper, is the case where the parties assume that other contracts would not be renegotiated if they did not reach agreement. By contrast, $d_{jh} = \pi_h(\hat{p}; \mathcal{J} \setminus \{j\})$, where $\hat{p}$ is the prices that would be negotiated if j were not in the market, corresponds to the case studied by Stole & Zwiebel (1995). In that case, contracts can be freely renegotiated in the event of a breakdown, and h and j never rejoin negotiations once they have broken down. In the transferable utility (TU) case, where $\frac{\partial q_j}{\partial p_j} = 0$, an appropriate choice of d_{jh} results in the bounds on price specified by the TU Core of this game, where no firm can receive more than its marginal contribution.[17]

Taking the first-order conditions of the model (6) yields the system of equations:

$$0 = b_j(h)\frac{1}{\pi_j}\frac{\partial \pi_j}{\partial p_j} + b_h(j)\frac{1}{\pi_h - d_{jh}}\frac{\partial \pi_h}{\partial p_j} \qquad \forall j \in \mathcal{J}, \tag{7}$$

which state that equilibrium requires that the bargaining-ability-weighted percentage gains and losses to each player are exactly equal. Rearranging with the partial derivatives

[14] The manufacturer's disagreement payoff is zero by the assumptions that the hospital is a monopsonist, the manufacturer is not capacity constrained, and each hospital is small enough that any returns to scale in manufacturing are not affected by inclusion or exclusion from a single hospital.

[15] It is straightforward to allow for multi-product manufacturers, as discussed in Appendix A.

[16] For a recent example, see de Fontenay and Gans (2007).

[17] The specification that achieves this is consumer surplus in the case where j is excluded, plus producer's surplus for all sales that switch from j to $k \neq j$: $d_{jh} = \pi_h(p; \mathcal{J} \setminus \{j\}) + \sum_{k \neq j} \pi_k|_{\{i|j = \arg\max_{k \in \mathcal{J}} u_{ik}\}}$. Intuitively, this could happen if the hospital could negotiate binding contracts with $k \neq j$ specifying that, under the contingency that a contract is not signed with j, the manufacturers would charge only marginal cost for any sale that would have gone to j.

calculated explicitly yields the following "pricing equation":[18]

$$p_j = c_j + \frac{b_j(h)}{b_j(h) + b_h(j)} \underbrace{\left[\left(1 + \frac{\partial q_j}{\partial p_j}\frac{p_j - c_j}{q_j}\right)\right.}_{\text{NTU adjustment}} \underbrace{\left.\frac{\pi_h - d_{jh}}{q_j} + p_j - c_j\right]}_{\text{"Added Value" of } j}, \tag{8}$$

which says that equilibrium prices are equal to cost plus a margin that is the manufacturer's bargaining ability relative to that of the hospital, multiplied by product j's "added value": the additional surplus created when the hospital contracts with product j versus when the hospital doesn't contract with product j. The portion of the added value appropriated by the hospital is adjusted by a term that takes into account that, in this non-transferable utility (NTU) game, a dollar increase in price also results in a decrease in quantity, so it does not transfer linearly into manufacturer profits.

The model requires that the term $\frac{\partial q_j}{\partial p_j}\frac{p_j - c_j}{q_j}$ lies in the interval $[-1, 0]$ (whereas the Bertrand case, where manufacturers set price, requires that it be exactly negative one).[19] This requirement means that, taking the prices in other negotiations as given, equilibrium prices must fall in the range where each manufacturer would prefer to increase price and the hospital would prefer to decrease price. Thus prices are always between marginal cost and the manufacturer's Bertrand best-response price.

Competition between substitutes enters this model in two ways: (1) via the hospital's disagreement point of not contracting with a given product; and (2) via the elasticities. The constraint of the hospital's disagreement point is reminiscent of solutions such as the Core,[20] whereas the elasticities are directly related to standard models of Bertrand price competition with differentiated products.[21] Via these two effects, more "competition," such as lower prices or greater substitutability among products, decreases both the added value and NTU adjustment terms, leaving a smaller piece of the pie for product j to capture. However, conditional on competition, the amount of value captured depends on bargaining via $\frac{b_j(h)}{b_j(h)+b_h(j)}$. In this way, this model achieves the goal of separating the effects of competition and bargaining in determining prices.[22]

[18] In the estimation section, I allow costs to be a function of quantity. This changes the elasticity term in the pricing equation to $\frac{\partial \pi_j/\partial p_j}{\partial \pi_h/\partial p_j} = 1 - \frac{\partial c_j}{\partial p_j} + \frac{\partial q}{\partial p}\frac{p-c}{q}$. However, in the body of the paper I consider only the case where $\partial c/\partial p = 0$ so as not to clutter the notation with extra terms.

[19] Algebraic manipulation of the pricing equation gives $p - c = \frac{b_j}{b_h}\left(1 + \frac{\partial q}{\partial p}\frac{p-c}{q}\right)\frac{\pi - d}{q}$. For price above cost and $d_{jh} = \pi_h(p; \mathcal{J} \setminus \{j\})$, all the components of this equation are positive, requiring that $1 + \frac{\partial q}{\partial p}\frac{p-c}{q} > 0$ as well.

[20] Aumann (1984) refers to the Core as a model of "unbridled competition".

[21] Another way to see the connection between the two models is to look at the "elasticity pricing rule" generated by this model, $\frac{p_j - c_j}{p_j} = \frac{1}{-\frac{\partial q_j}{\partial p_j}\frac{p_j}{q_j} + \frac{b_h(j)}{b_j(h)}\frac{p}{(\pi_h - d_{jh})/q_j}}$, which is the same as the one from the Bertrand model when $b_h = 0$.

[22] The separation can never be quite so clean in this NTU setting because in equilibrium, elasticities, added value, and quantities all depend on prices, which in turn depend on bargaining abilities.

4 The Empirical Model

The main goal of the this section is to estimate the parameters of a structural model that will distinguish among and quantify the various determinants of price variation across hospitals (demand, costs, competition, and bargaining abilities) and provide a "laboratory" in which to conduct the policy experiment of a change to uniform pricing.

I first adjust the theoretical model introduced in Section 3 to allow for heterogeneity across doctors/patients, hospitals, and time. I then discuss the variation in the data used to identify the model parameters, and I develop a generalized method-of-moments (GMM) algorithm that estimates the supply and demand parameters by matching the quantities and prices predicted by the model to the quantities and prices observed in the data. The main challenges in estimation are: (1) find a source of variation to identify the demand curve and solve the simultaneity issue in demand estimation (when prices are negotiated); and (2) separately identify cost and bargaining ability parameters in the pricing model.

4.1 Demand Model: Identification and Estimation

Rewriting the utility function to take full advantage of the available panel data, the indirect utility for patient/doctor i from product j at hospital h in month t is given by

$$u_{ijht} = \theta_{jht} - \theta^p p_{jht} + \varepsilon_{ijht}, \tag{9}$$

where, as before, θ_{jht} is the mean quality of product j across all patients in hospital h in month t; θ^p is the marginal disutility of price p_{jht} (in utils per dollar); and ε_{ijht} is a stochastic patient-specific quality component with distribution $f(\varepsilon)$, representing characteristics of the specific patient/doctor combination i that make the patient an especially good candidate for a specific stent.

4.1.1 Heterogeneity across hospitals and time

The mean quality of product j in hospital h in month t is given by

$$\theta_{jht} = \theta_{jh} + X_{jt}\theta^x + \xi_{jht}, \tag{10}$$

where θ_{jh} is the mean utility of product j in hospital h over the sample period; X_{jt} is a matrix of month-DES interaction dummy variables starting in January 2006 to account for the scare over DES safety during this time; and ξ_{jht} are unobservable time fluctuations in hospital preferences for each stent model.

Including the θ_{jh} fixed effects is important, as doing so controls for persistent unobserved heterogeneity at the product-hospital level (and thus also at the product level and hospital level). This heterogeneity across hospitals comes from different average preferences of doctors due to different opinions regarding the clinical data for each product, different mixes of patients, and different reimbursement levels for stenting procedures.

However, because ξ_{jht} is an average across different doctors with different preferences and different patients with different characteristics, monthly variation occurs when the sample of patients varies, when the month's patients are allocated differently among the hospital's doctors, or when an individual doctor receives information that changes her preferences. Attrition and recruitment of new doctors over time could also lead to changes in these unobserved preferences at the hospital level. To capture this, I model ξ_{jht} as evolving according to a first-order autoregressive (AR(1)) process

$$\xi_{jht} = \rho\xi_{jht-1} + \widetilde{\xi}_{jht}, \tag{11}$$

where $\widetilde{\xi}_{jht}$ is the innovation in hospital preferences for product j at time t, and ρ measures the persistence (of the variation around the mean θ_{jh}) over time.[23]

4.1.2 Doctor/patient heterogeneity

Not all doctor/patient combinations at a given hospital in a given month are the same, and the model captures these differences in the doctor/patient-specific unobservable term, ε_{ijht}. The distribution $f(\varepsilon)$ is an important component of the demand estimation because it directly affects the extent to which different products are substitutes for one another. I model $f(\varepsilon)$ as a mixture of nested logit models to capture three features of doctor/patient heterogeneity: (1) patients will vary in how badly they need a stent versus an alternative treatment; (2) some patients will be especially suited for a DES or BMS; and (3) some doctors have strong preferences for certain stents. Thus I specify

$$\varepsilon_{ijht} = \epsilon^{stent}_{ijht} + (1-\sigma_{stent})\epsilon^{des}_{ijht} + (1-\sigma_{stent})(1-\sigma_{des})\epsilon_{ijht} + \lambda_{ijht}, \tag{12}$$

[23]Note that any drift component of this process is subsumed into θ_{jh}. There are well-known challenges in correctly accounting for and estimating models where dynamic processes are combined with fixed effects (for a nice overview see Blundell and Bond 2000). Not accounting for the appropriate fixed effects can lead to a bias in the persistence estimates, while including fixed effects can introduce an incidental parameters bias because the effects are consistent only in the length of the panel. Regarding the first issue, I tried specifications with product and hospital fixed effects instead of their interaction. The results included a positive coefficient on price, which would be consistent with the bias from leaving out fixed unobserved product-hospital preferences. Regarding the second, the data set here does not suffer from typical short panel issues due to the panel length of 42 months.

where ϵ_{ijht}^{stent} is a random component common to all stents, ϵ_{ijht}^{des} is a random component common to all DES, and ϵ_{ijht} is a random component for stent j. The random mean shifter, λ_{ijht}, takes the value λ with probability ϕ_{jht} and zero otherwise.[24] This allows the distribution of doctor/patient tastes for each stent to be *bimodal*, capturing the fact that a doctor may have a strong preference for a particular stent. Allowing for this possibility is critical because a bimodal distribution allows for a demand curve with multiple groups of consumers, each with similar willingness-to-pay, whereas a unimodal distribution does not; and these two situations have very different implications for pricing.[25]

I assume that ϵ_{ijht} is distributed according to the "logit" extreme-value type I distribution with mean zero and scale parameter one, and that ϵ_{ijht}^{stent} and ϵ_{ijht}^{des} are the distributions derived in Cardell (1998) that converge to a point mass at zero as σ_{stent} and σ_{des} go to zero, so that by the theorem in the same paper, ε takes a two-level nested logit distribution with random mean λ_{ijht} and scale parameter one.

4.1.3 Estimating the utility parameters

I estimate the demand for coronary stents following the procedure suggested in Berry (1994), matching the observed market share data to the expected market shares predicted by the demand model, and inverting this system of equations to obtain an equation that is linear in the parameters, data, and econometric unobservable, $\widetilde{\xi}_{jht}$, allowing the use of linear instrumental variables methods.

Following the customary notation in the literature on random coefficients demand estimation, it is useful to represent the portion of utility that is not patient/doctor-specific using the term δ_{jht}, so that $u_{ijht} = \delta_{jht} + \varepsilon_{ijht}$. Taking the expectation over the distribution of the patient/doctor unobservables, ε, as in (2) yields the market shares predicted by the model for each product, in each hospital, in each month (here each hospital-month is a separate "market"): $s_j(\delta_{ht}; \sigma, \lambda, \phi)$. Where I use the vector notation $\delta_{ht} := (\delta_{1ht}, ..., \delta_{Jht})$ and $s_{ht} := (s_{1ht}, ..., s_{Jht})$.

Setting these predicted shares equal to the observed market shares yields a system of equations, $s_j(\delta_{ht}; \sigma, \lambda, \phi) = s_{jht}$. Berry (1994) proves that there is a unique vector δ_{ht} that solves this system. Therefore, the system can be inverted to obtain the mean utility for a each product in each hospital in each month as a function of market shares and the parameters governing doctor/patient heterogeneity, $\delta_j(s_{ht}; \sigma, \lambda, \phi)$. Under the assumed distribution of doctor/patient heterogeneity, $f(\varepsilon)$, the predicted market shares, $s_j(\delta_{ht}; \sigma, \lambda, \phi)$, have a closed-form solution; thus no simulation is necessary. However, the

[24] The probabilities ϕ_{jht} will vary across hospitals; thus the distribution of patients varies across hospitals both by the means θ_{jht} and by the probabilities ϕ_{jht}.

[25] Hastings (2008) uses a similar setup to characterize "brand loyalty".

inverse, $\delta_j(s_{ht};\sigma,\lambda,\phi)$, must be solved numerically, using the contraction mapping from Berry, Levinsohn, & Pakes (1995).[26]

Setting $\delta_j(s_{ht};\sigma,\lambda,\phi) = \delta_{jht}$ results in a model that is linear in the data and parameters, which can be solved for the econometric unobservables by taking pseudo-differences (i.e., $\tilde{x} := x_t - \rho x_{t-1}$), yielding

$$\widetilde{\xi}_{jht} = \tilde{\delta}_j(s_{ht};\sigma,\lambda,\phi) - \theta_{jh}(1-\rho) + \theta^p \tilde{p}_{jht} - \tilde{X}_{jt}\theta^x. \tag{13}$$

4.1.4 Demand instruments and identification

Long-term pricing contracts offer a new sort of variation to identify the demand curve—variation in demand itself. The reason is that while price is fixed over a period of time, demand shifts, and thus when price is renegotiated the movement is along the demand curve, analogous to a shift in supply (this does not, however, solve the potential problem of a simultaneous shift in demand when prices are renegotiated). This is critical in this study, where the most meaningful source of variation is in monthly demand shocks at each hospital.

More generally, Berry & Haile (2009) show that random coefficients discrete choice demand models are nonparametrically identified under the assumption $E[\widetilde{\xi}_{jht}|Z^d] = 0$ for a set of instrumental variables Z^d with sufficient variation. Thus the critical step is to find valid instruments and consider how their variation identifies the parameters of the model. Because of the product-hospital fixed effects, variation over time is the main source of identification. I set X_{jt} equal to the interaction of a dummy variable for DES with monthly dummy variables beginning in January 2006 in order to capture the market-wide decrease in DES use due to rising concerns about late stent thrombosis during this latter part of the sample. Thus θ^x is identified by the decreasing use, on average, of DES during this time. σ_{stent} parameterizes the degree to which preferences for stents are correlated relative to the outside good, so it is identified by the correlation between changes in price and changes in the market share of the outside good. Similarly, σ_{des} parameterizes the degree to which preferences for each type of stent are correlated relative to stents of the other type or the outside good, so it is identified by the correlation between changes in price and changes in market shares of the same type of stent versus changes in other market shares. The price coefficient, θ^p, is identified by the correlation between price changes and market share changes for each product. Finally, the persistence measure ρ is identified by the correlation over time between ξ_{jht-1} and ξ_{jht}.

The optimal instruments are the first derivatives of the unobservables with respect

[26] The contraction mapping needs to be modified slightly because the i.i.d. logit error term is scaled down by $(1-\sigma_{stent})(1-\sigma_{des})$ with the standard scale normalization for the nested logit.

to the parameters, which in this case are ξ_{jht-1} for ρ and the pseudo-differences for the other parameters. However, the theoretical model states that p_{jht} is a function of the unobservable demand innovation, $\widetilde{\xi}_{jht}$, and thus the respective pseudo-difference does not satisfy the exogeneity condition. Valid instruments must be correlated with these regressors but uncorrelated with the time t unobservable innovation. The lagged value p_{jht-1} satisfies this condition, and this is the instrument I use to identify θ^p.[27] Optimal instruments for the nonlinear parameters (σ, λ, ϕ) are nonlinear functions of market shares and prices. Thus I use nonlinear functions of lagged price and market shares, specifically $s_{jht-1}, s_{jht-1} \cdot p_{jht-1}, s^2_{jht-1}, p^2_{jht-1}$, as additional instruments. I use $\delta_j(s_{ht-1}; \sigma, \lambda, \phi)$ as an instrument to identify ρ, and for the remaining linear parameters I use the pseudo-differences.

This identification strategy will be valid if unobserved demand is imperfectly persistent ($\rho \in (0, 1)$) and if current demand is a sufficient statistic for forecasting future demand (as it is in the AR(1) process here).[28] These instruments would be invalid if, for example, prices were set strategically, taking into account future demand shocks known to the hospitals and/or manufacturers but unobserved by the researcher—e.g., forthcoming clinical data or sales and marketing efforts.

Other instrument strategies, based on changes in costs (Hausman 1996) or competition (Berry, Levinsohn, & Pakes 1995), do not seem promising here. There is little reason to think that significant cost shocks occur during this study, and even if they did, costs are so small compared to price that they would likely be weak instruments. As for competition, this is a relatively stable period in the industry, with two products entering, four (little-used) products exiting, and no changes to characteristics of the products in the market. It is worth noting that another new source of identification is possible here, in that any data on determinants of bargaining ability could serve as instruments for price. Unfortunately, no such data are available here.

4.2 Pricing: Identification and Estimation

This section shows how costs and relative bargaining ability can be estimated at the buyer-supplier transaction (and thus firm) level using the demand estimates and the assumed model of bargaining and competition.

Adding subscripts across hospitals and time, the pricing equation for the Nash Equi-

[27] An added benefit of this IV strategy in business-to-business markets is that it works with nonlinear pricing (where price is a function of quantity) as well.

[28] This same identification strategy was mentioned as an alternative in Berry, Levinsohn, & Pakes (1995) and used effectively in several subsequent studies.

librium of bilateral Nash Bargaining problems introduced in Section 3 becomes

$$p_{jht} = c_{jht} + \frac{b_{jt}(h)}{b_{jt}(h) + b_{ht}(j)} \underbrace{\left[\left(1 + \frac{\partial q_{jht}}{\partial p_{jht}} \frac{p_{jht} - c_{jht}}{q_{jht}}\right)\right.}_{\text{NTU adjustment}} \underbrace{\left.\frac{\pi_{ht} - d_{jht}}{q_{jht}} + p_{jht} - c_{jht}\right]}_{\text{“Added Value” of } j}. \tag{14}$$

The quantities to be estimated are costs, c_{jht}, and the ratio $\frac{b_{ht}(j)}{b_{jt}(h)}$.[29] The model presents two main challenges for estimation. First, the full distributions of cost and relative bargaining ability are not separately identified—there are $J \times H \times T$ equations and $2 \times J \times H \times T$ unknowns. This is in contrast to the case when suppliers set prices in a Bertrand-Nash Equilibrium, in which case $b_h = 0$ by assumption, and one can solve the system of equations for the $J \times H \times T$ c_{jht} terms (the full distribution of costs).[30] Second, the model is nonlinear in that there is no way to have *both* cost and bargaining ability enter the model linearly. However, the terms can be arranged so that *one* of the two does enter linearly. In general, specifying the statistical model from this point depends on the data available and the details of the industry being studied.

A full statistical model requires specifications for costs and bargaining in terms of data, parameters, and unobservables. Because the full distributions of c_{jht} and $\frac{b_{ht}(j)}{b_{jt}(h)}$ are not separately identified, one of these specifications must be entirely in terms of data and parameters—no unobservables.[31]

4.2.1 Cost specification

I specify costs by

$$c_{jht} = \gamma_j + \gamma^q q_{jht}, \tag{15}$$

where γ_j is a fixed effect for product j; and q_{jht} is the quantity of product j delivered to hospital h in month t, intended to capture and potential economies of scale in distribution. I assume that there are no unobservable determinants of costs. This assumption seems reasonable in this context because marginal costs of production and distribution are thought to be quite low and to vary little for a given product across hospitals and time. Also, it allows me to estimate the full distribution of relative bargaining abilities, which I am specifically interested in for this study.

[29] Note that $\frac{b_{jt}(h)}{b_{jt}(h)+b_{ht}(j)} = \frac{1}{1+\frac{b_{ht}(j)}{b_{jt}(h)}}$.

[30] Crawford and Yurukoglu (2010) benefit from the opposite case. In their context, marginal costs are assumed equal to zero, so they can simply solve this system of equations for the full distribution of bargaining parameters.

[31] This is not a desirable feature of this model, but it is a fundamental difficulty that arises from relaxing assumptions about price-setting and price-taking to allow for bargaining.

4.2.2 Relative bargaining ability specification

For relative bargaining ability, I specify

$$\frac{b_{ht}(j)}{b_{jt}(h)} = \frac{\beta_h}{\beta_j}\nu_{jht}, \tag{16}$$

where β_j will measure the bargaining ability associated with product j; β_h will measure the bargaining ability associated with hospital h; and ν_{jht} is the econometric unobservable term that measures the extent to which bargaining outcomes in the data deviate from the outcomes suggested by the firm-specific bargaining abilities. ν_{jht} could represent pair-specific bargaining abilities, the evolution of bargaining abilities over time, or the possibility that bargaining outcomes are simply random. Making the β terms constant across bargaining partners and time measures the extent to which a firm consistently captures a certain portion of the value created across all bargaining partners and time—the extent to which bargaining outcomes are firm-specific. To the extent that bargaining outcomes are simply random, this specification will set $\beta_j = \beta_h = 1$ (or, equivalently, $\ln(\beta_j) = \ln(\beta_h) = 0$ in the log-linear specification below), and all variation will be due to the random unobservable term ν_{jht}.

4.2.3 Estimation of costs and bargaining abilities

Here I describe my approach to isolating the econometric unobservable term, ν_{jht}. Combining the cost and bargaining specifications with the pricing equation and rearranging so that the unobservable enters linearly gives the statistical model

$$\frac{(\pi_{ht} - d_{jht})/q_{jht}}{p_{jht} - \gamma_j - \gamma^q q_{jht}}\left(1 + \frac{\partial q_{jht}}{\partial p_{jht}}\frac{p_{jht} - \gamma_j - \gamma^q q_{jht}}{q_{jht}}\right) = \frac{\beta_h}{\beta_j}\nu_{jht}, \tag{17}$$

and taking logarithms gives

$$\ln\left(g(X^s_{jht};\gamma)\right) \;=\; \ln(\beta_h) - \ln(\beta_j) + \ln(\nu_{jht}) \;=\; \hat{\beta}_h - \hat{\beta}_j + \hat{\nu}_{jht}, \tag{18}$$

where $g(X^s_{jht};\gamma) := \frac{(\pi_{ht}-d_{jht})/q_{jht}}{p_{jht}-\gamma_j-\gamma^q q_{jht}}\left(1 + \frac{\partial q_{jht}}{\partial p_{jht}}\frac{p_{jht}-\gamma_j-\gamma^q q_{jht}}{q_{jht}}\right)$ is the ratio of the amount of per-unit added value that goes to the hospital to the amount that goes to the manufacturer, adjusted by the elasticity term to account for NTU. Bargaining abilities are identified by the extent to which a given firm systematically captures a larger share of the value being created.

4.2.4 Pricing instruments and identification

The natural identifying assumption for costs and bargaining is $E[\hat{\nu}' Z^s] = 0$ for a set of instruments Z^s. The firm-specific bargaining parameters are identified by including manufacturer and hospital dummy variables in Z^s (one hospital dummy variable is dropped to preserve the full rank condition–this normalizes bargaining ability, which has no inherent scale, as relative to this hospital). For the cost parameters, the first derivatives with respect to the parameters (the optimal GMM instruments) are

$$\frac{\partial \hat{\nu}_{jht}}{\partial \gamma} = \frac{\partial \ln(g(X^s_{jht};\gamma))}{\partial c_{jht}} \frac{\partial c_{jht}}{\partial \gamma} = \frac{-1}{1 + \frac{\partial q_{jht}}{\partial p_{jht}} \frac{p_{jht} - c_{jht}}{q_{jht}}} \frac{\partial c_{jht}}{\partial \gamma}. \tag{19}$$

However, the term $\frac{-1}{1+\frac{\partial q_j}{\partial p_j}\frac{p_j - c_j}{q_j}}$ is a function of price, and thus ν, and so these are not valid instruments. Valid instruments are correlated with $\frac{\partial \hat{\nu}_{jht}}{\partial \gamma}$ but not with ν_{jht}. A natural choice would be $\frac{\partial c_{jht}}{\partial \gamma}$; but the manufacturer dummy variables ($\frac{\partial c_{jht}}{\partial \gamma_j}$) are unavailable because they are already being used for the bargaining parameters.

To find instruments for the cost parameters, I again use a "dynamic panel data" style approach, assuming that ν_{jht} evolves according to a dynamic process. I assume

$$\nu_{jht} = \nu_{jht-1} \widetilde{\nu}_{jht}, \tag{20}$$

where $\widetilde{\nu}_{jht}$ is the time t innovation in the relative bargaining abilities of product j and hospital h. The idea that the outcomes of current negotiations depend on the outcomes of past negotiations is consistent with intuition, as well as with case studies which have found that negotiated outcomes depend on relationships and processes that evolve over time. This approach is valid if the innovation, $\widetilde{\nu}_{jht}$, is uncorrelated with equilibrium prices and quantities at time $t-1$. It would be invalid if, for example, firms invest in bargaining ability for a specific partner in response to a perceived unfavorable price at $t-1$.

Treating $\widetilde{\nu}_{jht}$ as the econometric unobservable, taking first-differences gives

$$\hat{\widetilde{\nu}}_{jht} \;:=\; \ln(\widetilde{\nu}_{jht}) \;=\; \ln\left(g(X^s_{jht};\gamma)\right) - \ln\left(g(X^s_{jht-1};\gamma)\right), \tag{21}$$

which removes the time-invariant bargaining parameters (but not the γ_j because $g(\cdot)$ is nonlinear). The identifying assumption for costs becomes $E[\hat{\widetilde{\nu}}' Z^s] = 0$ for instruments Z^s. The optimal instruments are now the *differences* $\frac{\partial \hat{\nu}_{jht}}{\partial \gamma} - \frac{\partial \hat{\nu}_{jht-1}}{\partial \gamma}$, and while the month t terms are all still endogenous, the lagged $t-1$ terms are valid instruments, and so I use these lagged terms to provide the 12 moments needed to identify the cost parameters,

γ. Given the cost parameters, I then recover the full distribution of relative bargaining parameters via $\frac{b_{ht}(j)}{b_{jt}(h)} = \ln\big(g(X^s_{jht};\gamma)\big)$. I can then estimate the firm-specific bargaining abilities via the ordinary least squares regression $\ln\big(g(X^s_{jht};\gamma)\big) = \hat{\beta}_h - \hat{\beta}_j + \hat{\nu}_{jht}$ and the transformations $\beta_h = e^{\hat{\beta}_h}$ and $\beta_j = e^{\hat{\beta}_j}$.

4.3 Estimation Algorithm and Computation

In this section, I describe how I use the identifying assumptions $E[\widetilde{\xi}' Z^d] = 0$ and $E[\widetilde{\hat{\nu}}' Z^s] = 0$ to construct a method-of-moments algorithm to estimate the parameters $(\theta, \lambda, \sigma, \phi, \rho, \gamma, \beta)$.

I estimate the demand and supply parameters jointly, finding the parameters that minimize the GMM criterion $\widetilde{\xi}' Z^d (Z^{d\prime} Z^d)^{-1} Z^{d\prime} \widetilde{\xi}$ and $\widetilde{\hat{\nu}}' Z^s (Z^{s\prime} Z^s)^{-1} Z^{s\prime} \widetilde{\hat{\nu}}$, subject to the parameter constraints implied by the model: $\theta^p \geq 0$, $\lambda, \sigma \leq 1$, $\rho \in [0,1]$, $\beta > 0$, and $c_{jht} \in [0, p_{jht}]$. The pricing model requires a further constraint on the elasticities and cost parameters, $\left(1 + \frac{\partial q_{jht}}{\partial p_{jht}} \frac{p_{jht} - \gamma_j - 2\gamma^q q_{jht}}{q_{jht}}\right) \geq 0$.

I simplify the computational burden of estimation dramatically by fixing the probabilities, ϕ_{jht}, associated with the doctor/patient mean shifter, λ_{ijht}. I fix the probability, ϕ_{jht}, of the shock taking the value λ (as opposed to zero) to be zero for all BMS and equal to the market share among the two DES, $s_{jht|_{j=des}}$, for each of the DES. Although, in principle, the full distribution of ϕ_{jht} could be estimated, this introduces a large number of nonlinear parameters to an already difficult nonlinear minimization problem and asks a lot of the data, which are already being pushed with fixed effects and AR(1) processes.[32] Also, I "concentrate out" the linear parameters (θ_{jh}, θ^x) by rewriting them as a function of the data and $(\theta^p, \lambda, \sigma, \rho)$.[33] Finally, I minimize the objective functions over the five demand parameters $(\theta^p, \lambda, \sigma_{stent}, \sigma_{des}, \rho)$ and twelve cost parameters in γ, using the Price and Storn (2005) *Differential Evolution* global optimization algorithm.

4.3.1 Standard errors

The parameter restrictions imposed in this estimation procedure make it difficult to compute asymptotic standard errors directly. Instead I use a nonparametric "paired" bootstrap, constructing random samples by drawing with replacement from the original

[32] This is not really an assumption when either $\lambda = 0$ or $\lambda >> 0$. Fixing the probabilities has no effect if the best-fit model is unimodal ($\lambda = 0$); and as $\lambda \to \infty$, the probability that a doctor who prefers DES1 (in the sense that $\lambda_{iDES1} = \lambda$) chooses DES2 goes to zero, so the probabilities converge to the market shares of each DES, among the DES.

[33] Conditional on values for the parameters $(\theta^p, \lambda, \sigma, \rho)$, this is a linear regression problem, so the estimator for (θ_{jh}, θ^x) must satisfy the first-order conditions for that linear regression. Instead of searching over (θ_{jh}, θ^x), I "concentrate out" these parameters, replacing them by their estimators as functions of $(\theta^p, \lambda, \sigma, \rho)$.

data set. I sample hospitals instead of individual observations to allow for arbitrary correlation among the unobservables within a hospital (analogous to clustering standard errors at the hospital level).

5 Estimation Results

In this section, I discuss the estimates obtained via the framework developed in Sections 3 and 4. I first present the demand and cost parameters and compare these to external data sources as a way to check that the model captures the industry in a realistic way. The results show that heterogeneity in demand and bargaining ability both play an important role in the observed price variation.

5.1 Demand Parameters

The demand parameters are a critical piece of the model because they give the distribution of preferences (willingness-to-pay) for each stent across hospitals and patients. These preferences directly influence competition among the stents in a given hospital-month. Table 3 presents the estimates of the key demand parameters.[34] The month-to-month persistence in the evolution of demand, ρ, is 0.14, indicating that changes from the product-hospital specific mean are correlated over time. The parameter on price, θ^p, and those on the correlation in demand among all stents, σ_{stent}, and among DES/BMS, σ_{des}, show that price differences affect doctor decisions only when all else is close to equal, and prices are more likely to influence the choice between stents than the choice between using a stent or an alternative treatment. For example, the estimate $\sigma_{stent} = 0.82$ indicates that doctors very rarely make a decision based on price about whether or not a patient receives a stent—an encouraging result that matches what doctors say regarding their decision making. This also means that manufacturer price competition will be more about business-stealing than changing the size of the market.

The significant random coefficient on type of stent, $\sigma_{des} = 0.31$, matches the fact that some patients are better suited for DES or BMS. With regard to the two DES, however, it appears that there is, in fact, little competition. The large and significant estimate on the mean shifting component, $\lambda = 0.82$ (a $\lambda/\theta^p = 0.82/0.091 = \$9,000$ shift), indicates that there is a bimodal distribution of some doctors who have a strong preference for one DES over the other. This is consistent with what both doctors and industry experts say regarding doctors' preferences.

[34] See Appendix C for the results across different demand specifications, illustrating: the importance of the flexible specification used here; the effectiveness of the lagged instrumental variables; and the impact of imposing the supply-side elasticity constraints.

Table 3: Utility function parameter estimates. Estimates of key parameters from the utility function $u_{ijht} = \theta_{jh} - \theta^p p_{jht} + X_{jt}\theta^x + \rho\xi_{jht-1} + \widetilde{\xi}_{jht} + \epsilon^{stent}_{ijht} + (1 - \sigma_{stent})\epsilon^{des}_{ijht} + (1 - \sigma_{stent})(1 - \sigma_{des})\epsilon_{ijht} + \lambda_{ijht}$ for patient/doctor i using product j at hospital h in month t. Doctor/patient i random terms are integrated out to obtain predicted market shares. Estimation via GMM under the assumption $E[\widetilde{\xi}'_{jht} Z^d] = 0$ for instrumental variables Z^d, including lagged functions of price and market shares to correct for simultaneity bias in price and quantity. Standard errors computed using a nonparametric bootstrap at the hospital level.

parameter	point estimate	standard error
ρ (persistence in demand unobservable)	0.14^{***}	(0.04)
θ^p (price sensitivity in $\frac{utils}{\$1000}$)	0.091^{***}	(0.033)
σ_{stent} ("correlation" in demand for stents)	0.82^{***}	(0.06)
σ_{des} ("correlation" in demand for DES)	0.31^{***}	(0.13)
λ (shift in utility for loyal user of DES1 vs. DES2)	0.82^{*}	(0.46)

N=14,245. All s.e. clustered by hospital ($N = 96$). * 90%, ** 95%, *** 99%

5.1.1 Demand elasticities

Table 4 shows the means and standard deviations of the elasticities for each type of stent across stents, hospitals, and months.[35] The mean (standard deviation) own-elasticities are -0.48 (0.10) for BMS and -0.44 (0.20) for DES, meaning that many of the own-elasticities are less than -1. This is consistent with the fact that prices are negotiated, not set by the manufacturers to price-taking hospitals. With negotiated prices, elasticities combine demand, competition, *and* bargaining abilities.

Table 4: Own- and cross-elasticity estimates. $\frac{\partial q_j}{\partial p_k}\frac{p_k}{q_j}$ means and standard deviations across hospitals, months, and stents of that type. Own-elasticities less than -1 are consistent with negotiated prices and inconsistent with suppliers setting prices to price-taking buyers.

price elasticity of q_j:	with respect to p_k:	mean	std. dev.
BMS	own	-0.48	0.10
	other BMS	0.03	0.04
	DES	0.43	0.30
DES	own	-0.44	0.20
	BMS	0.03	0.04
	other DES	0.04	0.04
Outside Alternative	BMS	0.00	0.00
	DES	0.03	0.02

The cross-elasticities show substitution patterns within stent types, between types, and to the outside alternative of no stent. As noted in the utility parameter results,

[35] Standard errors on the elasticities are very tight. This same table with standard errors included is in Appendix B.

substitution to the outside alternative of no stent is very insensitive to stent prices, with mean cross-elasticities of 0.00 with respect to BMS price and 0.03 with respect to DES price. Also, as noted in the utility parameter results, there is very little substitution between the two DES, with an average cross-elasticity of 0.04. A related result that is evident in from the elasticity estimates, but is perhaps less obvious from the utility parameter estimates, is that there is more substitution between BMS and DES than within each type. In particular, the average cross-elasticity for a BMS with respect to DES price is 0.43 versus 0.03 with respect to BMS price (note that there are two DES and on average four BMS in each hospital-month—and DES price and quantity levels are on average much higher than those of BMS).

5.1.2 Willingness-to-pay and added value

The demand parameters also give the distribution of willingness-to-pay across doctor/patient types, products, hospitals, and months via $wtp_{ijht} = u_{ijht}/\theta^p + p_{jht}$. The way that willingness-to-pay enters the bargaining model is through a product's "added value"—the amount of extra value that is created when a hospital contracts with that product. Table 5 provides summary statistics for the distribution of expected added value (expectation over doctor/patient types) per unit, $\frac{\pi_h - d_{jh}}{q_{jh}} + p_{jh}$ (for now without subtracting manufacturer marginal costs), for each product across hospitals in September 2005. The BMS have average added values near \$3000—over three times their average prices—but the standard deviations in BMS added values across hospitals of over \$150 are similar in magnitude to the observed price variation. For the DES, though, the average added values of just over \$6000 are slightly less than three times their average prices, while the standard deviations of over \$700 are two to three times the standard deviations observed in DES prices across hospitals.

Table 5: "Added value" estimates. $\frac{\pi_h - d_{jh}}{q_{jh}} + p_{jh}$ across hospitals for each stent. The sample is restricted to September 2005 (middle of the sample in time) to isolate cross-sectional variation. There are N=54 hospitals sampled in this month; BMS1-3 have exited the market.

	mean (\$)	std. dev. (\$)	std.dev./mean	min (\$)	max (\$)	N
BMS4	3045	169	0.06	2801	3441	25
BMS5	2946	178	0.06	2732	3445	23
BMS6	3006	163	0.05	2807	3451	26
BMS7	3072	175	0.06	2802	3570	39
BMS8	3083	299	0.10	2814	3728	11
BMS9	3161	245	0.08	2819	3884	47
DES1	6272	705	0.11	4879	7808	54
DES2	6319	700	0.11	4914	7738	54

5.1.3 More Evidence of Demand Asymmetry

The finding in the raw data that hospital fixed effects explained only 12% of the cross-sectional variation in market shares for a given stent did not control for differences in prices and unobserved preferences, but the elasticity and added value estimates here do. Again looking at September 2005, regressing the product of own-elasticity and added value (the quantity the theory suggests matters for pricing) on stent dummy variables, and then on stent and hospital dummy variables, reveals that hospital effects still explain only 26% of the within-stent variation. The remaining 74% is stent-hospital specific variation, still suggesting more asymmetry than symmetry in demand across hospitals. This is still only suggestive, as there is no guide for "how much" symmetry or asymmetry leads to a given outcome. Completing the analysis requires estimating costs and bargaining abilities and computing the counterfactual equilibrium prices.

5.2 Cost Estimates

Table 6 presents the cost parameter estimates. The product-specific cost parameters—on average \$405 for BMS and \$1107 for DES—are all close to the range that industry experts report (thought the BMS are on the high end). The cost parameters are all fairly imprecisely estimated. This may be due to the first-differencing approach used for identification, which takes a great deal of the information out of the data. Interestingly, the parameter allowing for economies of scale or returns to more intense buyer-supplier relationships, γ^q, is rather small at about \$0.56/stent/(stent/month). This indicates that marginal costs for a given product vary little across hospitals (or at least not in a way that correlates with quantity sold). Perhaps economies of scale manifest through fixed costs that do not enter pricing, or perhaps the costs of investing in relationships with the largest customers offset any benefits.

Costs are identified in this framework by combining demand estimates with a model of how prices emerge in equilibrium. Table 7 shows how the credibility of these cost estimates depends upon the pricing model used. The first column gives a range of manufacturing and distribution cost estimates from industry experts.[36] From a manufacturing perspective, a DES is essentially a BMS with a polymer-drug coating. The added cost of a DES is a result of the royalty paid to the drug patent owner (typically about \$100 per stent); the added cost of the process of adding the drug coating; and the quality of the process of adding the drug coating. This last point can be particularly important, as some industry engineers quoted yields from the coating process as 15-20%, meaning that only about one in six DES passes quality inspection after the coating process. The

[36] Sources are interviews with current and former industry employees as well as Burns (2003).

Table 6: Cost parameter estimates. Parameters for the cost function, $c_{jht} = \gamma_j + q_{jht}\gamma^q$, estimated from the pricing equation $p_{jht} = c_{jht} + \frac{\beta_j}{\beta_h}\nu_{jht}\left[\left(1 + \frac{\partial q_{jht}}{\partial p_{jht}}\frac{p_{jht}-c_{jht}}{q_{jht}}\right)\frac{\pi_{ht}-d_{jht}}{q_{jht}}\right]$ using GMM, assuming that changes in the bargaining outcome are mean independent of a matrix of instrumental variables, $E[\ln\left(\frac{\nu_{jht}}{\nu_{jht-1}}\right)' Z] = 0$. Lagged bargaining splits interacted with product dummy variables and lagged quantities are used as instruments.

parameter	point estimate	standard error
baseline cost of BMS1, γ_{BMS1} (\$/stent)	188^{***}	(70)
baseline cost of BMS2, γ_{BMS2} (\$/stent)	334^{***}	(94)
baseline cost of BMS3, γ_{BMS3} (\$/stent)	169^{***}	(58)
baseline cost of BMS4, γ_{BMS4} (\$/stent)	480^{***}	(60)
baseline cost of BMS5, γ_{BMS5} (\$/stent)	608^{***}	(87)
baseline cost of BMS6, γ_{BMS6} (\$/stent)	280^{***}	(111)
baseline cost of BMS7, γ_{BMS7} (\$/stent)	161^{***}	(41)
baseline cost of BMS8, γ_{BMS8} (\$/stent)	419^{***}	(147)
baseline cost of BMS9, γ_{BMS9} (\$/stent)	557^{***}	(107)
baseline cost of DES1, γ_{DES1} (\$/stent)	1032^{***}	(381)
baseline cost of DES2, γ_{DES2} (\$/stent)	1221^{***}	(207)
economies of scale, γ^q (\$/stent/(stent/month))	-0.56^{*}	(0.32)

N=14,245. All s.e. clustered by hospital ($N = 96$). * 90%, ** 95%, *** 99%

variation in these ranges reflects different experts' assumptions regarding this and other aspects of what they think should enter marginal costs.

Table 7: Cost estimate comparison. The first column reports industry expert estimates for per-unit costs. The ranges reflect different experts' assumptions about what should enter "cost". Column two reports marginal cost estimates implied by the model if manufacturers were assumed to set prices. Column three reports marginal cost estimates for the bargaining model used in this paper. The numbers are means (std. dev. in parentheses) across hospitals, months, and stents of that type.

	industry expert estimates	assuming Bertrand, $b_h = 0$		allowing hetero $\frac{b_h(j)}{b_j(h)}$	
		mean	std. dev.	mean	std. dev.
cost of BMS in \$	100-400	-1140	364	405	189
cost of DES in \$	400-1600	-4601	3025	1107	225

The second column in Table 7 gives the cost estimates implied by assuming that manufacturers set prices in a Bertrand-Nash Equilibrium, and these results point out two ways in which this model falls short. First, the mean cost estimates are unrealistically small because prices are negotiated, and to assume that manufacturers set price is equivalent to assuming that hospitals have zero bargaining ability, $b_h = 0$, which is not the case on average. Second, the variation in cost estimates across hospitals is unrealistically large because the Bertrand model fails to allow for variation in relative bargaining abilities. Thus all the variation that cannot be explained by willingness-to-pay and com-

petition must be due to costs.[37] Any model with fixed bargaining abilities will produce similarly unreasonable variation in costs.

The third column reports the results from the model estimated in this paper, which allows for bargaining and heterogeneity in bargaining abilities, yielding more reasonable cost estimates.

This comparison also gives a sense of how the assumption of no unobserved variation in costs affects the robustness of the bargaining distribution and counterfactual estimates to come. Any unobserved cost variation would have to be unrealistically large to materially affect the results.

5.3 Bargaining Distribution Estimates

Given demand and cost estimates, the distribution of relative bargaining abilities, $\frac{b_{ht}(j)}{b_{jt}(h)}$, is given by equation (20). This distribution is easiest to interpret when each ratio is normalized to $\frac{b_{jt}(h)}{b_{jt}(h)+b_{ht}(j)} = \frac{1}{1+\frac{b_{ht}(j)}{b_{jt}(h)}}$, which takes the value 0 when the hospital gets all the surplus and the manufacturer prices at cost; and it takes the value 1 when the manufacturer sets its Bertrand best-response price.

Figure 3: Distribution of bargaining ability of manufacturers relative to hospitals, $\frac{b_{jt}(h)}{b_{jt}(h)+b_{ht}(j)}$. Over all product-hospital-time observations. The measure takes the value 0 in the case where the hospital gets all the surplus (conditional on disagreement points) and the manufacturer prices at cost; and it takes the value 1 in the case where the manufacturer gets all the surplus, pricing at the highest price consistent with competition.

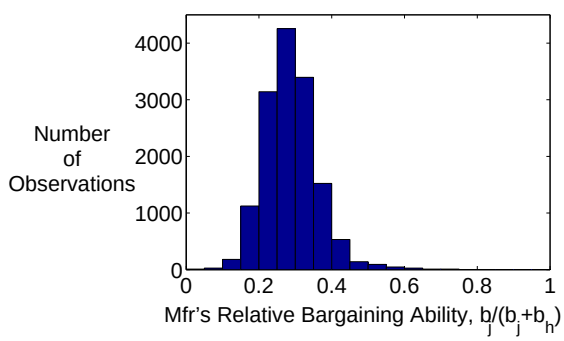

	mean	std. dev.	std. dev. / mean	min	max	N
$\frac{b_{jt}(h)}{b_{jt}(h)+b_{ht}(j)}$	0.31	0.13	0.41	0.02	0.96	14,245
	(0.02)	(0.01)	(0.04)	(0.01)	(0.03)	

Standard errors, clustered by hospital, in parentheses.

[37] One alternative explanation is that there are other sources of heterogeneity that my model is failing to capture. While any model is an imperfect representation of reality, this result has been robust across many specifications.

Figure 3 shows that both of these special cases are nearly always rejected (76 product-hospital-month observations < 0.05 and two observations > 0.95, respectively, out of 14,245 observations). The mean of 0.31 indicates that, on average, the hospital is a more powerful negotiator. This is in addition to the fact that, as a monopsonist, the hospital extracts surplus via competition between the stents. However, with standard deviation of 0.13, there is significant variation around this mean.

Thus the observed price variation across hospitals for a given stent comes from variation in both demand and bargaining abilities. Both of these sources of heterogeneity will also play an important role in thinking about the welfare effects of price discrimination versus non-discriminatory pricing in the next section.

6 The Welfare Effects of Price Discrimination

This section uses the data, model, and parameter estimates to produce welfare estimates of the total social surplus and how it is split among manufacturers and hospitals under the current price discrimination regime and under counterfactual scenarios with non-discriminatory pricing. The goal is to both illustrate and quantify the forces that affect upstream prices when hospitals merge, when group purchasing organizations are formed, and when market transparency initiatives are implemented.

The welfare effects of price discrimination (as opposed to non-discrimination) for buyers, suppliers, and society are theoretically ambiguous. A monopoly supplier is always better off discriminating, but the effect on buyers (those who paid high prices under discrimination are better off, and those who paid low are worse off) and total welfare depends on the curvature of demand (Aguirre, Cowan, and Vickers 2010). With oligopoly, however, suppliers can be worse off (and buyers better off) with discrimination because of business-stealing incentives (Holmes 1989). Sellers can even find themselves in a prisoner's dilemma where they would all be better off not discriminating, but nevertheless each finds it a dominant strategy to discriminate (Thisse and Vives 1988). Whether discriminating is likely to make suppliers better or worse off hinges critically on whether there is *symmetry* or *asymmetry* among suppliers on which buyers are "strong" and "weak" markets (Corts 1998). And whether price discrimination tends to exacerbate or alleviate allocative distortions (whether each consumer is matched with the product that creates the most value for him) also hinges on this same symmetry/asymmetry in demand (Stole 2007). Bargaining introduces another source of ambiguity: the difference between the bargaining abilities of the individual buyers and the "merged" buyers. The aggregation of these effects determines the welfare effects of price discrimination. The purpose of the counterfactual simulations here is to quantify and disentangle these effects.

6.1 Counterfactuals: Competitive and Bargaining Effects

In the counterfactual scenarios, prices are set according to a Nash Equilibrium of Nash Bargaining problems, as before; however, now there is only one price for each stent across all hospitals, so product and hospital profits are aggregated over all hospitals. This has an interpretation of the hospitals bargaining collectively with each manufacturer to solve

$$\max_{p_j} \quad \left[\sum_h \pi_{jh}\right]^{b_j(\mathcal{H})} \left[\sum_h (\pi_h - d_{jh})\right]^{b_{\mathcal{H}}(j)} \qquad \forall j \in \mathcal{J}, \tag{22}$$

where $b_{\mathcal{H}}$ is a bargaining parameter for all the hospitals collectively. Recall that an important special case of this model is $b_{\mathcal{H}} = 0$ where hospitals are price-takers and manufacturers set prices in a Bertrand-Nash Equilibrium.

Table 8 compares the aggregate outcomes from the current price discrimination regime to counterfactual predictions under uniform pricing for three different values of the hospital group bargaining ability—$b_{\mathcal{H}} = 0$, $\overline{b_h}$, and $\max(b_h)$.[38] These three different values of $b_{\mathcal{H}}$ were chosen both to illustrate the economic forces at work and to correspond to outcomes that might result from merger, group purchasing, or transparency.

Table 8: Aggregate changes. Equilibrium outcomes under the current price discrimination regime compared to those under non-discriminatory pricing for September 2005. Columns 2-4 are estimated percent changes under non-discriminatory pricing scenarios with different assumptions regarding hospitals' ability to bargain collectively. Column 2 sets $b_{\mathcal{H}}$ to zero, the case where hospitals do not bargain collectively and manufacturers set prices. Column 3 sets bargaining ability of the group of hospitals, $b_{\mathcal{H}}$, to $\overline{b_h}$ in order to isolate the change to competition. Column 4 sets $b_{\mathcal{H}}$ to the maximum bargaining ability of any individual hospital.

	Discrimination	% change with No Discrimination		
		$b_{\mathcal{H}} = 0$	$b_{\mathcal{H}} = \bar{b}_h$	$b_{\mathcal{H}} = \max(b_h)$
mean BMS price ($/stent)	1023	163	-2.2	-23
		(18)	(1.6)	(3)
mean DES price ($/stent)	2517	152	9.2	-14
		(17)	(3.4)	(4)
total stentings (stents/hospital/year)	969	-20	-1.7	1.5
		(5)	(0.5)	(0.7)
manufacturer profits ($M/hospital/year)	1.17	193	20	-19
		(26)	(4)	(8)
hospital surplus ($M/hospital/year)	14.3	-23	-2.0	1.8
		(5)	(0.6)	(0.9)
total surplus ($M/hospital/year)	15.5	-5.6	-0.25	0.14
		(1.9)	(0.11)	(0.09)

Standard errors, clustered by hospital, in parentheses.

[38] I find the equilibrium price vector by minimizing the sum of squared deviations from the first-order conditions.

The most dramatic change occurs if hospitals are unable to bargain collectively, $b_{\mathcal{H}} = 0$, leaving manufacturers free to set prices. This could result from efforts to increase price transparency. As Armstrong (2006) suggests, it is exactly the lack of transparency that allows sellers to cut the "secret discounts" that lead to different hospitals paying different prices, and increasing transparency could provide manufacturers a mechanism to commit to take-it-or-leave-it uniform pricing. In this case, I estimate that prices and manufacturer profits nearly triple; hospital surplus decreases by 23% (profits 77%);[39] and total surplus decreases by more than 5%. Prices at triple the observed level are well outside the observed range of data, so these exact numbers should be taken with some skepticism. However, I think the robust takeaway is that any policy that removes the hospitals' power to negotiate would be quite bad for hospitals. To the extent that price transparency would lead to this outcome, it would have exactly the opposite effect that policy-makers concerned with hospital costs are looking for.[40]

With regard to the effects of mergers or group purchasing, it seems more likely that the "merged" group of hospitals would have a bargaining ability nearer to the estimated range of individual hospital bargaining abilities. In these cases the results are more nuanced. First, the effects on the total number of stentings and total welfare are very small. Second, the changes that do occur in total welfare are completely explained by the changes in the total number of stentings, so there is no challenge in allocative efficiency of matching patients with the right stent. As a result, the interesting changes are in the way the surplus is split between the device manufacturers and hospitals, and these changes are driven by competitive and bargaining effects.

The results when hospitals bargain collectively at the average bargaining ability of all the hospitals, $b_{\mathcal{H}} = \overline{b_h}$, show how competition is actually less intense without price discrimination. Prices increase by 9% on average; manufacturer profits increase by 20%; and hospital surplus decreases by 2% (profits 14%).[41] This competitive effect is consistent with the theoretical results on best-response asymmetry (Corts 1998), and it confirms the reduced-form evidence regarding the amount of asymmetry across hospitals in the market share data and the elasticity and added value estimates.[42]

[39] Using detailed accounting data for hospitals in New York state, Huckman (2006) finds that marginal profits for angioplasty are on average 30% of revenues.

[40] There has been an active policy debate on transparency in device pricing, with much of issue 27, 2008, of *Health Affairs* devoted to the topic. While there are many theoretical discussions on what might happen, to my knowledge this is the first empirical analysis of the issue.

[41] Setting the group bargaining ability to the average across hospitals is not a perfect way to isolate the change due to competition because there are still two changes. A cleaner measure is to do the change in two steps: First, let all hospitals negotiate their own prices, but with their bargaining abilities fixed at the average; and second, have them negotiate as a group. The difference between the results in steps one and two isolates the true competitive effect. When I computed this, I found that the pure competitive effect accounts for over 90% of the change in prices.

[42] Hastings (2008) also finds that prices increase with a change to non-discrimination, perhaps due to

However, the result under uniform pricing depends critically on the bargaining ability of the group of hospitals relative to the bargaining abilities of the individual hospitals. When the group of hospitals has higher bargaining ability, $b_{\mathcal{H}} = \max(b_h)$, this is enough to overcome the competitive disadvantage and then some: Prices and manufacturers' profits fall by 14% and 19%; and hospital surplus increases by 1.8% (profits 13%).

One instructive way to see these "competitive" and "bargaining" effects is to look at the pricing equations under price discrimination and uniform pricing (the bar above a term denotes the quantity-weighted average over hospitals, e.g. $\overline{p_{jh}} := \sum_h \frac{q_{jh}}{\sum_h q_{jh}} p_{jh}$):

$$\textbf{Discriminatory:} \quad p_{jh} = c_{jh} + \frac{b_j(h)}{b_j(h) + b_h(j)} \left[\left(1 + \frac{\partial q_{jh}}{\partial p_{jh}} \frac{p_{jh} - c_j}{q_{jh}} \right) \frac{\pi_h - d_{jh}}{q_{jh}} + p_{jh} - c_{jh} \right]; \tag{23}$$

$$\textbf{Uniform:} \quad p_j = \overline{c_{jh}} + \underbrace{\frac{b_j(\mathcal{H})}{b_j(\mathcal{H}) + b_{\mathcal{H}}(j)}}_{\text{bargaining effect}} \overline{\left[\underbrace{\left(1 + \overline{\frac{\partial q_{jh}}{\partial p_j} \frac{p_j - c_{jh}}{q_{jh}}} \right)}_{\text{demand/competitive effect}} \frac{\pi_h - d_{jh}}{q_{jh}} + p_j - c_{jh} \right]}. \tag{24}$$

These pricing equations illustrate how the uniform case differs from the price discrimination case in two important ways. First, under price discrimination what matters is the product-hospital elasticities, whereas under uniform pricing the relevant elasticity is a quantity-weighted average of these elasticities. Second, in the discriminatory case the product-hospital bargaining ratio is what matters, whereas under uniform pricing, the bargaining ratio is the same across all hospitals.

The balance of these two effects is especially important for thinking about hospital mergers and group purchasing. The competitive effect of merging demand across hospitals with asymmetric preferences works to raise prices, but the final price will depend on the bargaining ability of the merged hospitals or group purchasing association. Figure 4 shows these two relationships graphically over the range of hospital bargaining abilities observed in the data.

Figure 4 shows a different perspective on the competitive and bargaining effects—the group of hospitals would need a bargaining ability more than 20% larger than the average hospital in order to overcome the disadvantage due to softer competition. Below this, hospitals would be worse off (in the cost of their stents) merging or group purchasing; above this, they would be better off.

The fact that only 14% of hospitals have such a high bargaining ability speaks to

the same mechanism via the brand loyalty she finds in the gasoline market.

Figure 4: Competitive and bargaining effects. The vertical axis is the average DES price, and the horizontal axis is the bargaining ability of the "merged" hospitals as a ratio of the mean hospital bargaining ability. The downward sloping curve shows the relationship between the predicted non-discriminatory price and hospital bargaining ability. The dotted horizontal line is the mean DES price observed under price discrimination. Comparing this to the predicted price at the average hospital bargaining ability shows how competition is more intense under price discrimination.

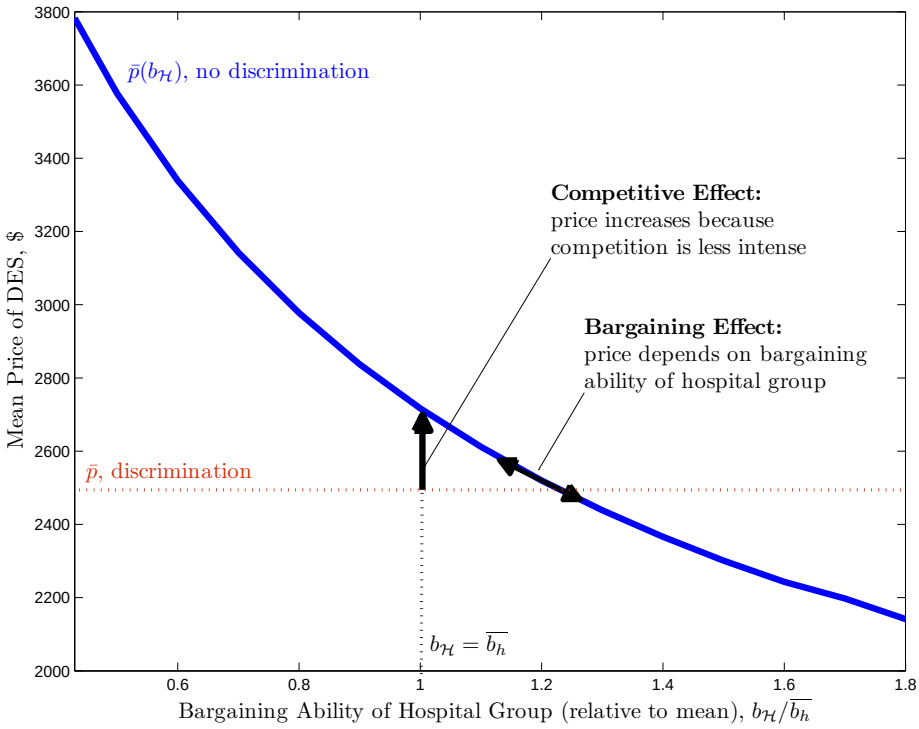

how difficult it might be to obtain. This could help explain why group purchasing plays a small role in contracting for physician preference items such as coronary stents. It also casts doubt on whether hospital mergers will be effective at decreasing costs for stents and other physician preference items.

7 Conclusion

This paper combines new panel data on the prices and quantities transferred between medical device manufacturers and hospitals with a structural model of supply and demand to estimate the welfare effects of price discrimination in the coronary stent market. The context highlights the fact that buyer mergers, group purchasing, and transparency—in addition to direct regulation—all offer reasons to study price discrimi-

nation. The major empirical challenge is that prices in the coronary stent market are negotiated (as they are in many business-to-business markets). I capture this using a model that generalizes the standard price-setting model to allow for bargaining, and I show how bargaining affects identification of both supply and demand parameters. The raw data and counterfactual estimates provide evidence that asymmetry in demand across hospitals leads to a softening of competition under non-discriminatory pricing, consistent with theory. This competitive effect increases stent prices by 9% relative to the average prices under price discrimination. However, final prices under non-discrimination also depend on the collective bargaining ability of the merged hospitals, which must be above the 85th percentile of the distribution of hospital bargaining abilities to overcome the disadvantage of softened competition. Taken together, these results cast doubt on whether hospitals mergers, group purchasing, or transparency would be effective in reducing the prices hospitals pay for medical devices such as coronary stents.

Though the details of the model in this paper are tailored to the specific context, the general modeling and estimation framework should be applicable to other contexts in which outcomes are negotiated. The results here suggest that thinking about bargaining ability—where firms end up within the range determined by costs, demand, and competition—matters. As more detailed data on vertical contracting relationships become available, it would be interesting to see how much of a role bargaining ability plays in other contexts.

Relatedly, while I document significant heterogeneity in bargaining abilities across firms, due to data limitations, I am unable to find any firm characteristics that seem to determine bargaining ability. Anecdotal evidence from talking to industry professionals suggests that better understanding the determinants of bargaining ability could lead to interesting links among individual performance, organizational structure, and market outcomes. Pursuing this research topic would require detailed data related to the price negotiation process in addition to the type of transfer data used here.

In the long run, price discrimination, like anything that affects firm profitability, could impact market entry and exit on both sides of the market. In the medical device market, this is particularly important because the buyer side represents the availability of medical care and the supplier side represents the availability of new medical technologies. Future research that takes a step back to endogenize the choices of who contracts with whom and market entry and exit would extend our understanding of the economics of business-to-business markets.

A Multi-product manufacturers

The model in the paper treats pricing for each product independently, but optimal behavior for a multi-product device manufacturer would be to take into account the externalities between its products. Let $m \in \mathcal{M}$ denote the manufacturers contracting with hospital h, with m_j denoting the manufacturer of product j. The new pricing equilibrium must then solve

$$\max_{\{p_j\}_{m_j=m}} \quad [\pi_m(p)]^{b_m} [\pi_h(p) - d_{mh}]^{b_h} \qquad \forall m \in \mathcal{M}, \tag{25}$$

where $\pi_m = \sum_{j \text{ s.t. } m_j=m} \pi_j$ is the total profits to manufacturer m and now negotiation occurs at the manufacturer level, so the relevant bargaining ability parameter is b_m, and the relevant outside option is d_{mh}. Note this has two effects: (1) the profit function of the manufacturer now takes into account externalities between its product's prices and (2) the hospital's outside option now reflects failure of bargaining with *all* of the manufacturer's products.[43]

The first order conditions of this optimization problem now yield a vector of equations that relate the profits of a manufacturer to its "added value" via

$$\pi_m = \frac{b_m}{b_m + b_h} \left[\underbrace{\left(-\frac{\partial \pi_m / \partial p_j}{\partial \pi_h / \partial p_j} \right)}_{\text{NTU adjustment}} \underbrace{(\pi_h - d_{mh}) + \pi_m}_{\text{"Added Value" of } m} \right] \qquad \forall j \text{ s.t. } m_j = m. \tag{26}$$

Note that the NTU adjustment here now changes the requirement that $\frac{\partial q_j}{\partial p_j} \frac{p_j - c_j}{q_j} \in [-1, 0]$ by taking the cross partials into account, making the requirement $\frac{\partial q_j}{\partial p_j} \frac{p_j - c_j}{q_j} + \sum_{k \neq j, m_k = m_j} \frac{\partial q_k}{\partial p_j} \frac{p_k - c_k}{q_j} \in [-1, 0]$ where the cross partial terms will be positive because the products are (imperfect) substitutes.

References

Aguirre, Inaki, Simon Cowan, and John Vickers, "Monopoly Price Discrimination and Demand Curvature," *American Economic Review*, 2010, *100* (4), 1601–1615.

Armstrong, Mark, "Recent Developments in the Economics of Price Discrimination," in Blundell, Newey, and Persson, eds., *Advances in Economics and Econometrics: Theory and Applications: Ninth World Congress*, Cambridge University Press, 2006.

[43] Note that the hospital may in fact have outside options that consist of excluding any subset of the manufacturer's products.

Asker, John, "Measuring Cost Advantages from Exclusive Dealing: An Empirical Study of Beer Distribution," 2004. Unpublished manuscript.

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

___ **and Philip A. Haile**, "Identification of Discrete Choice Demand From Market Level Data," May 2009. Unpublished manuscript.

___ **, James Levinsohn, and Ariel Pakes**, "Automobile Prices in Market Equilibrium," *Econometrica*, July 1995, *63* (4), 841–890.

Blomqvist, Ake, "The doctor as double agent: Information asymmetry, health insurance, and medical care," *Journal of Health Economics*, 1991, *10*, 411–432.

Blundell, Richard and Stephen Bond, "GMM Estimation with Persistent Panel Data: an application to production functions," *Econometric Reviews*, 2000, *19*, 321–340.

Burns, Lawton R., ed., *The Business of Healthcare Innovation*, Cambridge, 2005.

Cardell, N. Scott, "Variance Components Structures for the Extreme-Value and Logistic Distributions with Applications to Models of Heterogeneity," *Econometric Theory*, 1997, *13*, 185–213.

Corts, Kenneth S., "Third-degree Price Discrimination in Oligopoly: All-Out Competition and Strategic Commitment," *RAND Journal of Economics*, 1998, *29*, 306–323.

Crawford, Greg and Ali Yurukoglu, "The Welfare Effects of Bundling in Multichannel Television," 2010. Unpublished manuscript.

Cremer, Jacques and Michael Riordan, "On Governing Multilateral Transactions with Bilateral Contracts," *RAND Journal of Economics*, 1987, *18*, 436–451.

Dafny, Leemore, "Competition in Health Insurance Markets," *forthcoming in the American Economic Review*, 2009.

de Fontenay, Catherine and Joshua Gans, "Bilateral Bargaining with Externalities," 2007. Unpublished manuscript.

Dranove, David, Mark Satterthwaite, and Andrew Sfekas, "Boundedly Rational Bargaining in Option Demand Markets: An Empirical Application." Unpublished manuscript.

Hastings, Justine, "Wholesale Price Discrimination and Regulation: Implications for Retail Gasoline Prices," 2008. Unpublished manuscript.

Hausman, J., "Valuation of New Goods Under Perfect and Imperfect Competition," in T. Bresnahan and R. Gordon, eds., *The Economics of New Goods*, Vol. 58, National Bureau of Economic Research, 1996.

Ho, Justin, Katherine Ho, and Julie Mortimer, "The Use of Full-line Forcing Contracts in the Video Rental Industry," September 2010. Unpublished manuscript.

Ho, Katherine, "Insurer-Provider Networks in the Medical Care Market," *American Economic Review*, 2009, *99(1)*, 393–430.

Holmes, Thomas, "The Effects of Third-Degree Price Discrimination in Oligopoly," *American Economic Review*, March 1989, *79*, 244–250.

Horn, Henrick and Asher Wolinsky, "Bilateral Monopolies and Incentives for Merger," *RAND Journal of Economics*, 1988, *19*, 408–419.

Huckman, Robert, "Hospital integration and vertical consolidation: An analysis of acquisitions in New York State," *Journal of Health Economics*, 2006, *25*, 58–80.

Kyle, Margaret and David Ridley, "Would Greater Transparency and Uniformity of Health Care Prices Benefit Poor Patients?," *Health Affairs*, 2007, (5), 1384–1391.

Myatt, David and Eric Rasmusen, "Posted Prices vs. Haggling: The Economics of Isoperfect Price Discrimination," April 2009. Unpublished Manuscript.

Nash, J.F., "The Bargaining Problem," *Econometrica*, 1950, *18*, 155–162.

Nevo, Aviv, "A Practitioner's Guide to Estimating Random Coefficients Logit Models of Demand," *Journal of Economics & Management Strategy*, 2000, *9*, 513–548.

Price, K., R. Storn, and J. Lampinen, *Differential Evolution - A Practical Approach to Global Optimization*, Springer, Berlin, 2005.

Stole, Lars, *Handbook of Industrial Organization, Vol. 3* 2007.

—— **and Jeffrey Zwiebel**, "Intra-firm Bargaining under Non-binding Contracts," *Review of Economic Studies*, 1996, *63*, 375–410.

Thisse, J.F. and X. Vives, "On the strategic choice of spatial price policy," *American Economic Review*, 1988, *78*, 122–137.

Villas-Boas, Sofia B., "An Empirical Investigation of the Welfare Effects of Banning Wholesale Price Discrimination," *RAND Journal of Economics*, 2009, *40*, 20–46.