Price Discrimination and Bargaining: Empirical Evidence from Medical Devices

Matthew Grennan

WEB APPENDIX

Data Set Construction

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).[25] The key variables in the data are the price paid and quantity used for each stent in each hospital in each month. In addition, the hospitals report monthly totals for different procedures performed, such as diagnostic angiographies.

There are two main challenges in constructing a usable data set from the raw survey data. First, the survey was not as concerned with collecting price data as it was with collecting quantity data. Second, the survey was concerned with *usage* data, so whenever a stent is not used in a hospital-month that observation is missing (even if it is on the shelf and available for use). Table A1 illustrates how key sample summary statistics have remained stable as I took steps to "clean" the data set. More details are available in the Stata code used to execute these steps.

Table A1—: Data set modifications

	Raw	Remove no p	Impute some $q = 0$	Remove sole	Final with lags
Diagnostic procedures	272	271	273	303	304
	(13)	(19)	(19)	(23)	(23)
Percent receiving stent	29	28	27	29	29
	(0.9)	(1.0)	(1.0)	(1.0)	(1.0)
BMS Price ($)	1011	1019	1026	1010	1009
	(10)	(15)	(15)	(15)	(15)
DES Price ($)	2522	2520	2540	2524	2513
	(11)	(17)	(18)	(21)	(21)
Stent-Hospital-Months	21,035	10,669	14,245	11,301	10,098
Hospital-Months	5867	2902	3038	2196	1973
Hospitals	269	103	101	100	96

Note: Standard errors clustered at the hospital level.

The table rows record the sample mean (and associated standard error) for: number of diagnostic angiographies per hospital-month, percent of these diagnostic procedures that result in a stenting, BMS Price, and DES Price. It also records the total number of stent-hospital-month observations, number of hospital-month

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

markets, and total number of hospitals in each sample. The table columns correspond to the different samples. The first column shows the results for the raw survey data with 21,035 observations across 269 hospitals. Many of the hospitals do not report price data, and removing these cases makes a substantially smaller sample of 10,669 observations across 103 hospitals in column two. Despite the fact that observations are missing whenever $q_{jht} = 0$, there are cases where it is clear that a stent is little-used but present at a hospital. Whenever there are four or less months of no use surrounded by months of use for a stent, I impute the price for that observation. The data set is large, but small enough at this point to look over manually, and doing so reveals some glaring spots where data appears be misrecorded (for example, a hospital that usually performs 300 diagnostic procedure per month that suddenly performs 27), and I delete or impute these hospital-months as well. The result of these modifications is data set with 14,245 observations and 101 hospitals in column three.

There are two further modifications to the data set that result from a combination of data constraints and modeling choices. The first has to do with how to handle hospitals which still only use a single DES or BMS in a given month. There are three possible ways to deal with these cases: (1) leave them, implicitly assuming that no other stents were available in that hospital-month; (2) impute them, implicitly assuming that other stents were available at the imputed price, but no quantity was used; or (3) drop these hospital-months, assuming that these hospital-months are not systematically different from the rest of the sample. In this version of the paper, I choose option (3), dropping these observations. Leaving them as in (1) would not allow for modeling competition from the left out stents, which is unrealistic. Imputing them as in (2) would be an attractive solution if the price imputations were accurate. A previous version of this paper used the imputing approach and obtained results qualitatively the same and quantitatively similar to those obtained here. However, there is always concern that the imputation procedure could drive results, especially those on the firm-specific determinants on prices in demand and bargaining abilities. Given these limitations, I prefer dropping the sole-source cases, as in (3), which does not rely on unrealistic assumptions or "creating data". One hospital and 2,944 observations are dropped in this step. As before, there are no statistically significant changes in the sample means in Table A1, though there is a 10 percent increase in the mean number of diagnostic procedures, consistent with the fact that $q = 0$ cases are more likely to occur in small hospitals for sampling reasons. Subsection A.A1 explores in greater detail whether these dropped hospitals are indeed similar to the remaining sample, or if there is any evidence of these sole-sourcing instances being due to "exclusive dealing".

The final cut of the data occurs because the first observation for each stent-hospital pair is lost in taking the pseudo-differences for the demand unobservables, which are allowed to follow an AR(1) process. 1,203 observations and four hospitals are lost, with no statistically significant differences in the sample means, leaving the final sample used for estimation: an unbalanced panel of 10,098 stent-hospital-month observations over 96 hospitals for eleven stents from January 2004 through June 2007.

A1. Potential Sole-Sourcing and Exclusivity

Because the data is recorded for stents *used* by a given hospital in a given month, it does not contain data on the set of stents available but not used. Further, the price data does not include any information besides price, such as exclusivity arrangements. Despite the fact that exclusive arrangements which impact prices paid are common in business-to-business markets, including many medical supplies, my understanding from talking with industry participants is that "exclusivity" did not play a major role in coronary stent pricing during the time of this study (2004-07). However, because the model used in this paper does not explicitly allow for strategic choices regarding "who contracts with whom", it is important to verify this omission empirically.

The analysis in this Section looks at the effects of exclusive (100 percent market share among similar type stents) and near-exclusive (over 80 percent) situations on prices paid for two stents: DES2 and BMS8.[26] The results indicate that neither exclusive nor nearly exclusive contracts seem to play a role in driving the observed price variation across hospitals.

Tables A2 and A3 show the results of several regressions of price on dummy variables for exclusivity for DES2 and BMS8. In each case, the first four columns present evidence regarding full exclusivity using the data set before the sole-sourcing cases are cut, and the next four for near exclusivity using the data set used in the paper. In each of these first two specifications are: (1) a regression of price on a dummy variable for exclusivity only (equivalent to a t-test of means between the two samples), and (2) the same regression with the addition of time dummy variables to account for the fact that prices decrease and observations of sole-sourcing increase over time, creating what could be a spurious effect. The next two specifications, (3) and (4), look at the same regressions, but using only within-hospital changes for the subsample of hospitals with both sole-sourcing and non-sole-sourcing months for that stent.

The point estimates for DES2 in Table A2 tell a story of exclusivity potentially being correlated with an average price decrease of $42-94, but these impacts going away once time dummies are included. This is consistent with the facts that prices decrease over time, and doctors may tend to settle on a preferred stent over time. It is also consistent with increased use of exclusive contracts over time, but even if that is the case, the remaining evidence suggests that this is not a systematically important phenomenon.

[26]These are chosen because they are the stents from each category where the most sole-sourcing is observed, suggesting that they would be the first place to look for any evidence of exclusivity. The results reported here are representative of those for other stents and for changing the threshold for near-exclusive to 70 and 90 percent, which are available upon request.

Table A2—: Prices of DES2: Exclusivity and Near-exclusivity

parameter	E1	E2	E3	E4	NE1	NE2	NE3	NE4
Exclusive, $s_{jht}\|_{g_j} = 1$	-42	-11	-94	28				
	(39)	(36)	(66)	(37)				
Nearly-exclusive, $s_{jht}\|_{g_j} > 0.8$					43	4	65	5
					(37)	(36)	(26)	(16)
Month Fixed Effects	-	Y	-	Y	-	Y	-	Y
Hospital Fixed Effects	-	-	Y	Y	-	-	Y	Y
N	2805	2805	742	742	1960	1960	1184	1184
Number "Sole-source"	451	451	451	451	624	624	517	517
$N_{Hospitals}$	101	101	24	24	94	94	52	52
R^2	0.005	0.26	0.32	0.65	0.008	0.26	0.59	0.79

Note: Standard errors in parentheses, clustered at the hospital level.

The point estimates are all very noisy, with none having a t-statistic greater than 1.4, and the R^2 suggest that exclusivity does little to explain the price variation observed in the data. Relatedly, beyond the regression results regarding the two sample means, there is no discernible difference in the sample standard deviations either, at $221 for sole-sourcers and $225 for non. Combined with the further evidence that these sole-sourcing cases comprise only 16 percent of the hospital-month observations for DES2 (and this is the largest percentage observed for any stent), it seems difficult to make a case for an important role of full exclusivity. Results for near exclusivity are similar in every way except for the fact that the sample mean differences for the specifications without time dummy variables suggest that those with high market shares pay about $43-65 *more* on average than others, which is more consistent with the standard problem of a positive correlation between price and market share as a result of unobserved quality than a story of exclusivity.

Table A3—: Prices of BMS8: Exclusivity and Near-exclusivity

parameter	E1	E2	E3	E4	NE1	NE2	NE3	NE4
Exclusive, $s_{jht}\|_{g_j} = 1$	15	52	-23	10				
	(41)	(41)	(16)	(28)				
Nearly-exclusive, $s_{jht}\|_{g_j} > 0.8$					-37	-8	-40	-0.8
					(40)	(43)	(16)	(17)
Month Fixed Effects	-	Y	-	Y	-	Y	-	Y
Hospital Fixed Effects	-	-	Y	Y	-	-	Y	Y
N	2260	2260	516	516	1597	1597	925	925
Number "Sole-source"	168	168	130	130	173	173	173	173
$N_{Hospitals}$	89	89	21	21	82	82	39	39
R^2	0.0003	0.11	0.68	0.75	0.003	0.07	0.65	0.71

Note: Standard errors in parentheses, clustered at the hospital level.

Looking to BMS8 and Table A3 shows similar small and noisy point estimates comparing sample means, little in sample standard deviation ($193 for sole and $221 for non), and infrequency of sole-sourcing in general (eight percent of observations).

As discussed above, not modeling exclusivity amounts to an assumption that whenever the data shows little or no use of a particular stent at a particular hospital, then this is because the doctors at that hospital do not prefer that stent, not because the stent was excluded for a strategic pricing reason. Despite the empirical checks here and discussions with industry insiders, there is no way to guarantee that no hospital has an exclusive agreement which affects pricing. To the extent that this occurs, those hospitals will show up as "high bargaining ability" hospitals in the analysis. This would be consistent with the broader interpretation (discussed in Section 3.1.3 and Section 5.1.2) of bargaining ability as potentially capturing administrator power vis-a-vis doctors in addition to pure negotiating skill with manufacturers.

Estimation Details

The estimation approach used in this paper makes some small departures from the well-known GMM algorithms developed in Berry, Levinsohn, and Pakes (1995) and related research. As such, I include a description of the algorithm here to aid in replication of this study or the use of such a model in other contexts with brand-loyalty in demand and/or negotiated prices. 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 separately estimate the demand $(\theta, \lambda, \sigma, \phi, \rho)$ and supply (γ, β) parameters . Although joint estimation would be more efficient, it would also constrain the demand parameters to be consistent with the bargaining model, while estimating the demand system separately allows the demand results to provide a check on the appropriate supply side model.

B1. Demand Estimation Details

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"): $s_j(\boldsymbol{\delta}_{ht}; \sigma, \lambda, \phi)$. Where I use the vector notation $\boldsymbol{\delta}_{ht} := (\delta_{1ht}, ..., \delta_{Jht})$ and $\mathbf{s}_{ht} := (s_{1ht}, ..., s_{Jht})$.

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 $\boldsymbol{\delta}_{ht}$ 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, $\delta_j(\mathbf{s}_{ht}; \sigma, \lambda, \phi)$. Under the assumed distribution of doctor/patient heterogeneity, $f(\varepsilon)$, the predicted market shares, $s_j(\boldsymbol{\delta}_{ht}; \sigma, \lambda, \phi)$, have a closed-form solution where each is a linear combination of the L "brand-loyal" mixture types, $s_j(\boldsymbol{\delta}_{ht}; \sigma, \lambda, \phi) = \sum_{l=1}^{L} \phi_{ht}^l s_j^l(\boldsymbol{\delta}_{ht}; \sigma, \lambda)$ and (note the equation below is written for a DES; for a BMS these two labels would switch places):

$$\begin{aligned}
s_j^l(\boldsymbol{\delta}_{ht}; \sigma, \lambda) &= s_{j|des}^l(\boldsymbol{\delta}_{ht}; \sigma, \lambda) s_{des|stents}^l(\boldsymbol{\delta}_{ht}; \sigma, \lambda) s_{stents}^l(\boldsymbol{\delta}_{ht}; \sigma, \lambda) \\
\text{where} & \\
s_{j|des}^l(\boldsymbol{\delta}_{ht}; \sigma, \lambda) &= \frac{I_{jht}^l}{\sum_{k \in des} I_{kht}^l} \\
s_{des|stents}^l(\boldsymbol{\delta}_{ht}; \sigma, \lambda) &= \frac{\left(\sum_{k \in des} I_{kht}^l\right)^{1-\sigma_{des}}}{\left(\sum_{k \in des} I_{kht}^l\right)^{1-\sigma_{des}} + \sum_{k \in bms} I_{kht}^l} \\
s_{stents}^l(\boldsymbol{\delta}_{ht}; \sigma, \lambda) &= \frac{\left[\left(\sum_{k \in des} I_{kht}^l\right)^{1-\sigma_{des}} + \sum_{k \in bms} I_{kht}^l\right]^{1-\sigma_{stent}}}{1 + \left[\left(\sum_{k \in des} I_{kht}^l\right)^{1-\sigma_{des}} + \sum_{k \in bms} I_{kht}^l\right]^{1-\sigma_{stent}}} \\
\text{and where} & \\
I_{jht}^l &= \exp\left(\frac{\boldsymbol{\delta}_{jht} + \lambda_{des} \mathbf{1}_{\{j=l\}}}{(1-\sigma_{stent})(1-\sigma_{des})}\right)
\end{aligned}$$

Because shares take a closed form, no simulation is necessary. However, the inverse, $\delta_j(\mathbf{s}_{ht}; \sigma, \lambda, \phi)$, must be solved numerically, using the contraction mapping from Berry, Levinsohn, and Pakes (1995) (modified slightly because the i.i.d. logit error term is scaled down by $(1-\sigma_{stent})(1-\sigma_{des})$).

Setting $\delta_j(\mathbf{s}_{ht}; \sigma, \lambda, \phi) = \delta_{jht}$ 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

$$\widetilde{\xi}_{jht} = \tilde{\delta}_j(\mathbf{s}_{ht}; \sigma, \lambda, \phi) - \theta_{jh}(1-\rho) + \theta^p \tilde{p}_{jht} - \tilde{\mathbf{X}}_{jt} \boldsymbol{\theta}^x. \tag{B1}$$

I then use the Price and Storn (2005) *Differential Evolution* global optimization algorithm to find the parameters that minimize the GMM criterion $\widetilde{\boldsymbol{\xi}}' \mathbf{Z}^d (\mathbf{Z}^{d\prime} \mathbf{Z}^d)^{-1} \mathbf{Z}^{d\prime} \widetilde{\boldsymbol{\xi}}$, subject to the parameter constraints implied by the model: $\theta^p \geq 0$; $\lambda, \sigma \leq 1$; $\rho \in [0,1]$. The instruments used are

$$\begin{aligned}
Z_{jht}^d = \Big[& \delta_{jht-1} \quad p_{jht-1} \quad \textstyle\sum_{k \neq j} p_{kht-1}/K_{ht-1} \quad \ln(s_{jht-1|stents}) \quad \ln(s_{jht-1|des}) \\
& p_{jht-1}^2 \quad \left(\textstyle\sum_{k \neq j} p_{kht-1}/K_{ht-1}\right)^2 \quad s_{jht-1} p_{jht-1} \quad s_{jht-1} \textstyle\sum_{k \neq j} p_{kht-1}/K_{ht-1} \quad s_{jht-1}^2 \Big].
\end{aligned}$$

I simplify the computational burden of estimation dramatically in two ways.

First, I fix the probability, ϕ_{jht}, of each stent-specific shock λ_{ijht} taking the value λ_{bms} or λ_{des} (as opposed to zero) to be equal to the market share of that stent among the stents actually implanted in each hospital-month, $s_{jht|_{j=stent}}$. 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 to the limit with the stent-hospital fixed effects and AR(1) process. Note also that fixing the probabilities equal to market shares is not really an assumption when either $\lambda = 0$ or $\lambda >> 0$ (with the latter being the case here). 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 stent j (in the sense that $\lambda_{ij} = \lambda$) chooses stent k goes to zero, so the probabilities converge to the market shares of each stent.

Also, conditional on values for the parameters $(\theta^p, \lambda, \sigma, \rho)$, estimation of $(\theta_{jh}, \boldsymbol{\theta}^x)$ is a linear regression problem, and their estimators must satisfy the first-order conditions for that linear regression. Thus instead of searching over $(\theta_{jh}, \boldsymbol{\theta}^x)$, I "concentrate out" these parameters, replacing them by their estimators as functions of $(\theta^p, \lambda, \sigma, \rho)$.

B2. Supply Estimation Details

With demand estimated, I then estimate the supply parameters by finding the parameters that minimize the GMM criterion $\ln(\boldsymbol{\nu})'\mathbf{Z}^s(\mathbf{Z}^{s\prime}\mathbf{Z}^s)^{-1}\mathbf{Z}^{s\prime}\ln(\boldsymbol{\nu})$, subject to the demand parameter estimates from the first stage and the parameter constraints implied by the model: $\beta > 0$; $c_{jht} \in [0, p_{jht}]$; and $-1 \geq \left(1 + \frac{\partial q_{jht}}{\partial p_{jht}} \frac{p_{jht} - \gamma_j}{q_{jht}}\right) \geq 0$.

The supply unobservable is given by

$$\ln(\nu_{jht}) \quad = \quad \ln\left(g(X^s_{jht}; \gamma)\right) - \ln(\beta_{jh}), \tag{B2}$$

where $g(X^s_{jht}; \gamma) := \frac{p_{jht} - \gamma_j}{\left(1 + \frac{\partial q_{jht}}{\partial p_{jht}} \frac{p_{jht} - \gamma_j}{q_{jht}}\right) \frac{\pi_{ht} - d_{jht}}{q_{jht}}}$ 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.

The elasticities and added value terms are obtained from the demand estimates. The mixture of nested logits allows for closed form solutions, which dramatically speeds up estimation relative to cases when they must be simulated (e.g. normally distributed random coefficients). The elasticities are given by $\frac{\partial q_{jht}}{\partial p_{kht}} \frac{p_{kht}}{q_{jht}} = \sum_{l=1}^{L} \phi^l_{ht} \frac{\partial q^l_{jht}}{\partial p_{kht}} \frac{p_{kht}}{q_{jht}}$ where (suppressing the hospital and time subscripts):

$$\frac{\partial q^l_j}{\partial p_k} = |\theta^p| q_j \left(s_k + s_k|_{stent} \frac{\sigma_{stent} \mathbf{1}_{\{j,k \in stent\}}}{1 - \sigma_{stent}} + s_k|_{des} \frac{\sigma_{des} \mathbf{1}_{\{j,k \in des\}}}{(1 - \sigma_{des})(1 - \sigma_{stent})} - \frac{\mathbf{1}_{\{j=k\}}}{(1 - \sigma_{des})(1 - \sigma_{stent})} \right)$$

and the hospital surplus is given by $\pi_{ht} = \sum_{l=1}^{L} \phi_{ht}^{l} \pi_{ht}^{l}$ where:
(B3)

$$\pi_{ht}^{l} = \frac{1}{|\theta^p|} \ln \left(1 + \left[\left(\sum_{j=des} e^{\frac{\delta_{jht}+\lambda_{des}\mathbf{1}_{\{j=l\}}}{(1-\sigma_{stent})(1-\sigma_{des})}} \right)^{1-\sigma_{des}} + \sum_{j=bms} e^{\frac{\delta_{jht}+\lambda_{bms}\mathbf{1}_{\{j=l\}}}{(1-\sigma_{stent})(1-\sigma_{des})}} \right]^{1-\sigma_{stent}} \right).$$

The hospital disagreement point d_{jht} for each stent is calculated as the hospital surplus when that stent is removed from the choice set and prices of other stents remain the same (this is the "Nash" or "passive beliefs" assumption on disagreement points used in much of the bargaining with externalities literature, including the original Horn and Wolinsky (1989) and recent empirical work by Crawford and Yurukoglu (2011)).

The instruments used are the first derivatives of the unobservables with respect to the parameters, lagged by one month:

$$Z_{jht}^{s} = \begin{bmatrix} \mathbf{1}_{\{bms\}} \frac{p_{jht-1}}{1+\frac{\partial q_{jht-1}}{\partial p_{jht-1}} \frac{p_{jht-1}}{q_{jht-1}}} & \mathbf{1}_{\{des\}} \frac{p_{jht-1}}{1+\frac{\partial q_{jht-1}}{\partial p_{jht-1}} \frac{p_{jht-1}}{q_{jht-1}}} \end{bmatrix}$$

The search is only over the cost parameters because again, instead of searching over (β_{jh}), I "concentrate out" these parameters by taking the "within" transformation, subtracting stent-hospital means.

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

$$\text{(B4)} \qquad \max_{\{p_j\}_{m_j=m}} \quad [\pi_m(\mathbf{p})]^{b_m} [\pi_h(\mathbf{p}) - d_{mh}]^{b_h} \qquad \forall m \in \mathcal{M},$$

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. This second reason is why I choose not to use the multi-product manufacturer setup in this paper—several hospitals in the data use a subset of a given manufacturers' products. Combined with the low cross-elasticities, which makes externalities between products less of a concern, the stent-specific pricing model seems more appropriate for this application.

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

$$\text{(B5)} \qquad \pi_m = \frac{b_m}{b_m + b_h}\left[\underbrace{\left(-\frac{\partial \pi_m/\partial p_j}{\partial \pi_h/\partial p_j}\right)}_{\text{NTU adjustment}}\underbrace{(\pi_h - d_{mh}) + \pi_m}_{\text{“Added Value” of } m}\right] \qquad \forall j \text{ s.t. } m_j = m.$$

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.

B3. Standard Errors

The parameter restrictions and multiple stages in the estimation procedure make it difficult to compute asymptotic standard errors directly; so I use a delete-one jackknife, constructing 96 sub-samples, each with one hospital deleted from the original 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). For each sample, I compute the demand estimates, supply estimates, and counterfactuals; and I then use the standard deviation in these estimates across the samples as the standard errors.

Robustness

This Appendix conducts several specification and robustness checks, focusing especially on the demand estimates, which are critical for the analysis in this paper. C.1.1 estimates a series of specifications using a simple logit demand system in order to verify that the basic identification approach works. C.1.2 demonstrates the importance of allowing for more flexibility in the demand curve with the nested logit random coefficients and the mixture terms which allow for brand loyalty. C.1.3 checks the robustness of the demand estimates to estimating from a subsample of the data and including time dummy variables. C.2 checks robustness of the paper's results to various assumptions on stent marginal costs.

C1. Demand Estimation Specification and Robustness

Identifying the Effect of Price on Demand

Table C1 illustrates how the stent-hospital fixed effects, AR(1) error process, and instrumental variables identify the price sensitivity coefficient in the context of a simple logit model of demand: $\ln(s_{jht}/s_{0ht}) = \theta^p p_{jht} + X_{jht}\theta^x + \xi_{jht}$. Though the logit restricts the shape of the demand curve and thus does a poor job of estimating own and cross-elasticities, it will consistently estimate the average price effect, and it provides a simple context that focuses on this effect in order to see the identification strategy at work.

Table C1—: Identifying the Effect of Price on Demand

parameter	OLS	stent-hospital FE	FE & AR(1)	IV
Persistence in demand unobservable, ρ	-	-	0.26	0.26
			(0.004)	(0.004)
Price sensitivity in $\frac{utils}{\$1000}$, θ^p	0.98	-0.63	-0.67	-0.73
	(0.04)	(0.02)	(0.02)	(0.03)

Note: Logit demand estimates from: $\ln(s_{jht}/s_{0ht}) = \theta^p p_{jht} + X_{jht}\theta^x + \xi_{jht}$ for different specifications to illustrate how the fixed effects, AR(1) term, and instrumental variables identify the effect of price on demand. $N = 10,098$. Standard errors in parentheses, clustered by hospital, $N_{Hospitals} = 96$. First stage F-test for instrument strength: $F = 664$.

OLS results in a positive price coefficient, consistent with the standard problem of unobserved demand heterogeneity that is correlated with price. Both the institutional accounts of demand heterogeneity and the economics of identifying demand with negotiated prices suggest adding stent-hospital fixed effects and relying on within stent-hospital variation over time. The resulting negative price coefficient suggests that this approach is well-founded. Institutional knowledge also suggests that even within a stent-hospital, demand may evolve over time with some amount of persistence, and the result of adding an AR(1) component in addition to the fixed effects suggests that this is indeed the case.

If prices are always set at the beginning of the month (and do not incorporate future changes to demand that are not incorporated into current demand), then there may be no further endogeneity/simultaneity problem. To avoid this potentially strong assumption, the paper's analysis of the economics of negotiated prices suggests that both lagged own price and mean lagged other prices would be valid instrumental variables. Using these instruments increases the magnitude of the price coefficient by approximately nine percent. The results of the first-stage regression of price on these instruments and the other regressors shown below in Table C2 indicate that both are strongly correlated with price; and under the timing assumption discussed in the paper—that price does not incorporate known changes in future demand that are not already captured in current demand—the instruments are also uncorrelated with the unobservable innovation in demand ($\widetilde{\xi}_{jht}$).

Allowing for Nonlinearities in the Demand Curve

Whereas the stent-hospital fixed effects and AR(1) term capture heterogeneity in demand across hospitals and time, institutional knowledge suggests that there is significant heterogeneity across patients and doctors within a hospital. While

Table C2—: First-stage IV Regression

p_{jht-1}	$\sum_{k\neq j} p_{kht-1}/K_{ht-1}$	F(2,95) statistic
0.68	0.033	664
(0.02)	(0.016)	

Note: Price (p_{jht}) regressed on instrumental variables of lagged own price (p_{jht-1}) and lagged average price of other stents at the same hospital ($\sum_{k\neq j} p_{kht-1}/K_{ht-1}$) and the other regressors. $N = 10,098$. Standard errors in parentheses, clustered by hospital, $N_{Hospitals} = 96$.

the logit can identify average price effects, it does so by fitting a demand curve that has relatively little curvature and thus restricts substitution patterns between products. Providing a demand specification that is flexible enough to "allow the data to speak" is especially important for a study such as this one where so much hinges on the nature of demand. Table C3 shows estimates for the logit, for a nested logit with random coefficients on the stent versus no stent and DES versus BMS, and for a mixture of nested logits that allows each stent to have its own mean-shifter for some set of patients/doctors.

Table C3—: Demand Specifications: Nonlinear Demand Parameters

parameter	Logit	Nested Logit	Mixture of NL (Paper)
Persistence in demand unobservable, ρ	0.26	0.10	0.08
	(0.004)	(0.002)	(0.002)
Price sensitivity in $\frac{utils}{\$1000}$, θ^p	-0.73	-0.29	-0.27
	(0.03)	(0.02)	(0.02)
"Correlation" in demand for stents, σ_{stent}	-	0.56	0.38
		(0.04)	(0.05)
"Correlation" in demand for DES, σ_{des}	-	0.31	0.29
		(0.02)	(0.02)
Shift for loyal user of each DES, λ_{des}	-	-	3.3
			(0.3)
Shift for loyal user of each BMS, λ_{bms}	-	-	2.0
			(0.2)
mean BMS own-elasticity	-0.61	-0.56	-0.32
mean DES own-elasticity	-1.38	-2.05	-0.52
mean outside option cross-elasticity	0.08	0.04	0.03
GMM criterion	161.2	16.25	15.19

Note: $N = 10,098$. Standard errors in parentheses, clustered by hospital, $N_{Hospitals} = 96$.

The results show that allowing for a more flexible demand curve is important for explaining the data. The random coefficient on stents versus the outside good captures the fact that some patients need a stent while others don't. The random coefficient on DES captures the fact that some patients (or their blockage type) may not be appropriate for a DES or the fact that some doctors may favor DES more than others at a given hospital. The random mean shifters capture the fact that some stents can be especially appropriate for a specific type of patient and the (now confirmed) institutional belief that doctors can be intensely loyal to their preferred stent(s).

These nonlinearities in demand are especially important in their implications for pricing. The "brand-loyalty" evident here provides an incentive to keep prices high to extract surplus from loyal customers, as shown in Figure C1.

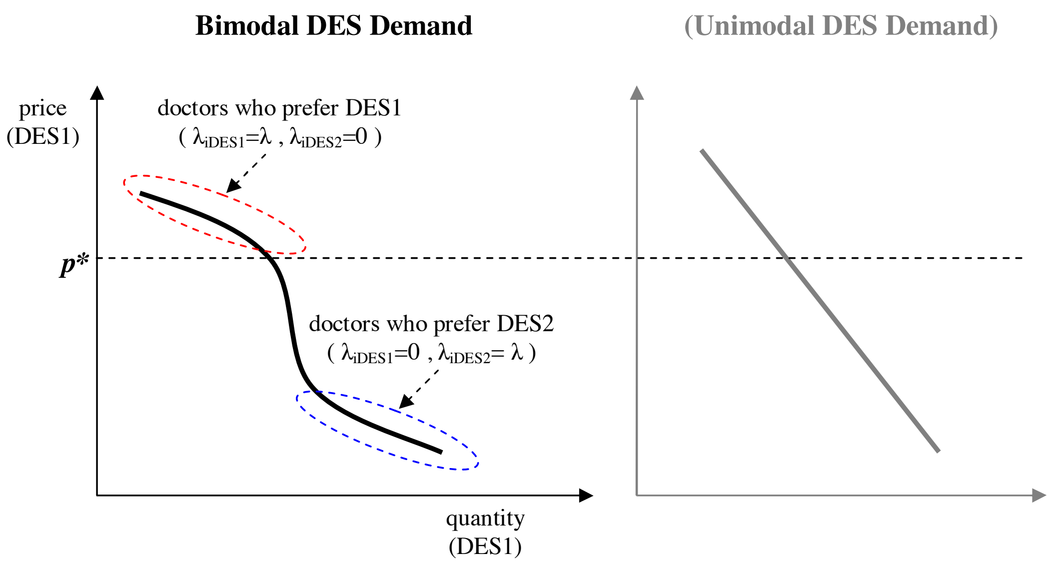

Figure C1. : Bimodal versus unimodal demand for DES

Note: The random mean, λ_{ijht}, allows the distribution of doctor/patient tastes to be *bimodal*. A bimodal distribution implies 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—in particular near a price such as p^* in the figure.

Robustness to Sample Time and Control Variables

The demand model used in the paper represents my preferred specification, balancing parsimony with flexibility in capturing the heterogeneity across hospitals and patients. Table C4 shows the results of robustness checks that (1) estimate the same model on the subset of the data before the DES safety scare, and (2) estimate the same model with month fixed effects added.

The results across the robustness checks are all qualitatively similar. In particular, demand is relatively inelastic, consistent with the institutional facts about doctor price-sensitivity and negotiated prices. Quantitatively, the results of the two robustness checks are close to those of the main specification from the paper, though they differ in some ways that make sense.

The results from running the model on the period before the DES safety scare (Jan. 2004 - Feb. 2006) show slightly more elastic demand estimates, and in particular less brand loyalty among BMS. This makes sense because the DES safety scare provided exactly the type of variation that was useful in pinning

Table C4—: Demand Robustness

parameter	Paper	2004-06	Month FE
Persistence in demand unobservable, ρ	0.08	0.09	0.08
	(0.002)	(0.006)	(0.003)
Price sensitivity in $\frac{utils}{\$1000}$, θ^p	-0.27	-0.31	-0.15
	(0.02)	(0.03)	(0.03)
"Correlation" in demand for stents, σ_{stent}	0.38	0.26	0.46
	(0.05)	(0.03)	(0.14)
"Correlation" in demand for DES, σ_{des}	0.29	0.23	0.41
	(0.02)	(0.02)	(0.09)
Shift for loyal user of each DES, λ_{des}	3.3	3.95	3.25
	(0.3)	(0.3)	(1.0)
Shift for loyal user of each BMS, λ_{bms}	2.0	0.0	2.0
	(0.2)	(0.1)	(0.8)
mean BMS own-elasticity	-0.32	-0.41	-0.17
mean DES own-elasticity	-0.52	-0.62	-0.28
mean outside option cross-elasticity	0.03	0.07	0.03

Note: $N = 10,098$. Standard errors in parentheses, clustered by hospital, $N_{Hospitals} = 96$.

down how inelastic demand really was, especially the substitution patterns to and between BMS.

The results from adding month fixed effects to the model show elasticities almost half of those in the main specification, driven entirely by a decrease in the price sensitivity parameter, θ^p. This move is in the opposite direction of what would be expected if there were residual correlation between the demand unobservable and price variable in the main specification (month fixed effects will soak up any month-specific unobserved variation in the value of stenting versus alternative options that affects all stents and all hospitals). A perhaps more plausible explanation for the decrease in the price coefficient with month fixed effects is attenuation bias—because of the stent-hospital fixed effects and AR(1) process, identification comes from within stent-hospital variation over time, and including month fixed effects absorbs some of this variation, biasing the price coefficient towards zero. The fact that standard errors increase dramatically in this specification is also consistent with attenuation from the time fixed effects absorbing useful variation over time in the data.

C2. Robustness to Cost Estimates

Cost parameters are not tightly identified in this application because the large amount of product differentiation leads to added values that are always much larger than marginal costs. However, even large changes to the cost numbers induce relatively small changes in bargaining ability and counterfactual estimates. Table C5 shows the results of these estimates for costs fixed at zero, the estimated costs in the paper ($c_{bms} = 34, c_{des} = 1103$), and costs fixed at the minimum observed prices in the data ($c_{bms} = 240, c_{des} = 1540$).

The results of varying the cost parameters show that, as expected, bargaining ability estimates change, but less dramatically than the cost changes. The

Table C5—: Robustness to Various Cost Assumptions

	$c_{bms} = 0$ $c_{des} = 0$	Paper $c_{bms} = 34$ $c_{des} = 1103$	$c_{bms} = 240$ $c_{des} = 1540$
Mean bargaining split, $\frac{b_j(h)}{b_j(h)+b_h(j)}$, $(0,1)$	0.43	0.33	0.25
Std. dev. bargaining split, $\frac{b_j(h)}{b_j(h)+b_h(j)}$, $(0,1)$	0.15	0.07	0.07
Mfr profits, ($M/hospital/year)	2.18	1.24	0.84
Hospital surplus, ($M/hospital/year)	4.32	4.32	4.32
Mean DES price, ($/unit)	2509	2509	2509
Mfr profit change for $b_{\mathcal{H}} = \bar{\beta}_h$, (percent)	5.5	8.0	10.7
Hospital surplus change for $b_{\mathcal{H}} = \bar{\beta}_h$, (percent)	-3.1	-1.4	-1.2
Mean DES price change for $b_{\mathcal{H}} = \bar{\beta}_h$, (percent)	5.2	1.7	0.7

level of manufacturer profits are directly related to costs and thus sensitive to price changes, but manufacturer profit *changes* under the counterfactuals are less sensitive to the cost changes. The different manufacturer bargaining abilities implied by the different costs does lead to different price increases under the more uniform pricing counterfactual, which leads to different hospital surplus changes. Overall, these robustness checks confirm that, even under these two extreme cost possibilities, the results are quantitatively similar and qualitatively identical.