

Insurer-Provider Networks in the Medical Care Market

By KATHERINE HO*

I use data on the hospital networks offered by managed care health insurers to estimate the expected division of profits between insurers and providers. I include a simple profit-maximization framework and an additional effect: hospitals that can secure demand without contracting with all insurers (e.g., those most attractive to consumers and those that are capacity constrained) may demand high prices that some insurers refuse to pay. Hospital mergers may also affect price bargaining. I estimate that all three types of hospitals capture higher markups than other providers. These results provide information on the hospital investment incentives generated by bargaining. (JEL G22, G34, I11, L25)

The effects of managed care health insurers on the price and quality of medical care have been widely researched.¹ One aspect of their impact, however, has not been addressed in detail: the restriction imposed by each insurance plan on the network of hospitals from which its enrollees can choose. In a previous paper (Ho 2006), I estimate a model of consumer demand for hospitals and health insurers taking these constraints into account. In this paper, I use the demand estimates, together with data on the hospital networks offered by plans in 43 US markets, to analyze the supply side of the market. I model the negotiation process by which plans and hospitals choose their equilibrium networks and which determines the division of the profits generated by each contract. The results provide evidence that hospitals may have long-run incentives to invest in order to increase their attractiveness to patients or to reduce their costs and to merge with other providers, since these strategies may increase their bargaining power with plans. One additional incentive may be detrimental to consumers: hospitals may also benefit from underinvesting in capacity compared to the welfare-maximizing investment level.

I use the demand estimates from my previous paper to calculate the producer surplus (defined as the profit to be divided between the plan and the hospitals in its network) generated by each potential hospital network for each plan in the data. Next I note that consumers' ability to move across plans if necessary to access their preferred providers may prompt certain types of hospitals to demand high prices and to be excluded from some plans' networks. Some positive-surplus contracts may therefore not be agreed upon in equilibrium. My model of the contracting process takes this issue into account by defining the health insurer's expected profits from each potential hospital network as producer surplus less the profits of the relevant hospitals.

* Department of Economics, Columbia University, 1037 IAB, 420 West 118th Street, New York, NY 10027 (e-mail: katherine.ho@columbia.edu). Thanks to Ariel Pakes and David Cutler for their advice. I also thank two anonymous referees, John Asker, Nancy Dean Beaulieu, Oliver Hart, Joy Ishii, and Julie Mortimer for helpful comments and Columbia's Program for Economic Research for financial support.

¹ For example, Robert H. Miller and Harold S. Luft (1997) review 15 studies of the effects of managed care on quality. They find no compelling evidence of a reduction in quality of care, although patients show less satisfaction with managed care than with traditional plans. David M. Cutler, Mark B. McClellan, and Joseph P. Newhouse (2000) consider the causes of the expenditure reductions achieved by managed care plans in the treatment of heart disease. They show that virtually all the difference in spending between indemnity plans and HMOs comes from lower unit prices rather than the quantity of services or a difference in health outcomes.

My methodology is particularly attractive since I am able to estimate the profits secured by different types of hospitals using only data on insurers' choices of hospital networks and insurer and hospital characteristics. I have no information on the prices paid to hospitals by particular plans. The analysis is complicated by the existence of multiple potential equilibria and by problems with endogenous regressors. Several recent papers develop methodologies that address these issues. However, their approaches often make restrictive assumptions regarding the nature of the unobservables, and most are feasible only for problems involving small numbers of firms.² This paper adopts a different approach developed in Ariel Pakes et al. (2006) and discussed further in Pakes (2008), in which plans choose their networks in a two-stage process conditional on their expectations regarding other plan choices and the prices demanded by hospitals. The equilibrium implicitly establishes markups for a hospital's services that are functions of the characteristics of the hospital itself and the distribution of consumer, hospital, and plan characteristics in the particular market.

I estimate the markups of three specific hospital types: "star" hospitals (providers that are very attractive to consumers), those that expect to be capacity constrained, and those that are members of hospital systems. I find that star hospitals capture \$6,700 per patient more than other providers compared to costs of approximately \$11,000 per patient. This, together with my other estimates, implies markups of approximately 25 percent of revenues, in contrast to nonsystem, nonstar providers whose markups over average costs are estimated to be negative. The coefficient on capacity constraints is positive and significant, with a magnitude implying markups \$6,900 per patient higher than those of other hospitals, when this variable is included without the star hospital variable. The star and capacity constraint variables are necessarily quite highly correlated: they capture similar packages of characteristics that are attractive to consumers, rather than measuring single and distinct attributes. Both are calculated using my demand model. However, I find evidence that both variables are important. Star status alone is correlated with high hospital profits. Capacity constraints seem to give the hospital additional leverage in the bargaining process, perhaps by acting as a commitment device to persuade plans that it will choose to contract selectively. I also estimate that the profits of hospitals in systems are approximately \$180,000 per month higher than those of other providers. They also charge high penalties from plans that contract with some but not all of the hospitals in their system. The results are therefore consistent with several recent papers that suggest that hospital mergers may prevent plans from using the threat of exclusions to control prices, and indeed that hospital systems are often formed for exactly this purpose.³ In addition, I find that hospitals with low costs have higher markups than their competitors, consistent with many bargaining models. All of these results point to the importance of bargaining and market power in determining contractual outcomes in this market.

This paper can be considered in the context of the broader literature on vertical relationships between upstream producers and downstream retailers or distributors. A large theoretical literature considers the implications of these relationships for efficiency and welfare.⁴ However, the

² For example, Katja Seim (2001), Donald Andrews, Steven T. Berry, and Panle Jia (2004), and Federico Ciliberto and Elie T. Tamer (2004) all propose methods to analyze market entry problems. Morten Sorensen (2007) and Jeremy T. Fox (2007) set out estimation methods for two-sided matching problems similar to mine. However, Sorensen's nested solution method is infeasible in the large markets studied here, while Fox's simpler methodology does not permit an analysis of the division of profits between upstream and downstream firms.

³ See, for example, Royce D. Luke, Y. Ozcan, and P. Olden (1995), Jason Barro and David M. Cutler (1997), Cara S. Lesser and Paul B. Ginsburg (2001), Glen P. Mays, Robert E. Hurley, and Joy M. Grossman (2003), and Cory S. Capps and David Dranove (2004).

⁴ Jean Tirole (1995) reviews many of these papers. See also, for example, Joseph J. Spengler (1950) and Patrick Rey and Tirole (1986a, b).

empirical literature is quite sparse, limited largely by a lack of data. This paper is among the first to embed demand estimates into a model of the supply side of a vertical market.^{5,6}

Several strands of the health economics literature are also relevant to this paper. A number of authors demonstrate that the prices paid by plans to hospitals are consistent with simple bargaining models.⁷ Esther Gal-Or (1997, 1998) develops a Nash bargaining model in a two-plan, two-hospital setting. Gregory Vistnes (2000) has a model of two-stage competition between hospitals: providers compete first for preferential access to health plans and then for individual patients. Finally, Karen Eggleston, George Norman, and Lynne Pepall (2004) use a similar theoretical framework in a market containing health plans, hospitals, and physician groups to look at the effects of horizontal and vertical integration on prices. However, no previous empirical papers have addressed the determinants of the observed network choices or the effect of the contractual process on long-term investment incentives.

In the next two sections I describe the contractual process between insurers and hospitals and introduce the dataset. Section III outlines the demand estimates from Ho (2006) and uses them to generate a measure of surplus. Section IV discusses the intuition regarding bargaining, and Section V introduces the full empirical model. Results are given in Section VI, and the final section concludes.

I. Industry Background

Each year, every privately insured consumer in the United States chooses a health plan, generally from a menu offered by his employer.⁸ The insurer contracts with hospitals and physicians to provide any care needed during the year. Once the consumer has reached his deductible he may, in general, visit any of the providers listed by the health plan and receive services at zero charge or after making a small out-of-pocket payment.

There is some variety in the restrictiveness of different types of managed care plan. If an individual is insured by a health maintenance organization (HMO) he may visit only the hospitals in that plan's network. Point-of-service (POS) plan enrollees can visit out-of-network hospitals but only if referred by a primary care physician (PCP). Preferred provider organizations (PPOs) and indemnity plans are the least restrictive insurers: enrollees do not need a PCP referral to visit an out-of-network hospital, although PPOs may impose financial penalties for doing so, for example, in the form of increased copayments or deductibles. The focus of this paper is on HMO and POS plans, since their network choices have the strongest effect on both consumers and hospitals; 53 percent of the privately insured population was enrolled in an HMO/POS plan in 2000.

Every HMO/POS plan contracts separately with every hospital in its network. The exact form of the contracts varies but most specify a price to be paid to the provider per unit of care, for

⁵ Other papers that use this approach include Julie Holland Mortimer (forthcoming), which analyzes the welfare effects of different types of vertical contracts in the video rental industry. John Asker (2006) considers the effects of exclusive dealing between beer manufacturers and their distributors.

⁶ This paper could also be thought of as modeling a product-line choice by plans, where hospital networks are considered to be a product characteristic. However, that problem would be somewhat simpler than the one considered here since it would involve plans making choices and hospitals playing a passive role. My model of the plan-hospital bargaining process, in contrast, accounts for both plan and hospital profit maximizations.

⁷ Most of these regress the prices paid to hospitals on measures of hospital and plan bargaining power. Examples are John M. Brooks, Avi Dor, and Herbert S. Wong (1997), Jack Zwanziger and Cathleen Mooney (2000), and Roger Feldman and Douglas Wholey (2001). In addition, Robert J. Town and Gregory Vistnes (2001) and Cory S. Capps, Dranove, and Mark A. Satterthwaite (2004) investigate the effect of the hospital's value to consumers on its profits. They estimate consumer preferences over hospitals and regress hospital profits or prices on variables that summarize consumer demand for the hospital.

⁸ Fifty-seven percent of the population is insured through an employer, compared to 5 percent who purchase insurance independently and 24 percent in Medicare and Medicaid (see www.statehealthfacts.org).

example, a price per inpatient day or per diagnosis-related group (DRG). Prices vary both across providers for a given insurer and across insurers for a given provider; contracts are usually renegotiated annually.⁹ Both parties in the negotiation need to balance consumer demand for services against the price agreed. A health plan would prefer to contract with the hospitals that are valued by its likely customers, particularly the customers on the margin of joining, but must also take into account the fact that hospitals in demand may seek higher prices than their less differentiated counterparts. Hospitals seek to maximize their returns by contracting with plans that both offer high prices and provide a steady flow of patients.

In order to model the contractual process, I need to specify the timing of the different hospital and plan decisions. The stages of my model are as follows:

- Stage 1: Hospitals make price offers to plans.
- Stage 2: Plans choose their hospital networks.
- Stage 3: Plans set premiums.
- Stage 4: Consumers and employers jointly choose plans.
- Stage 5: Sick consumers visit hospitals; plans pay hospitals per service provided.

My main focus is on Stages 1 and 2. Not much is known about the exact form of the bargaining process used in reality, how much it varies across plan-hospital pairs or across markets, or the extent of asymmetric information between insurers and hospitals. Interviews with plan and hospital representatives who are involved in contractual negotiations suggest that plans often have the final decision rights over whether to agree to contracts. The simplest bargaining model with this property has hospitals making simultaneous take-it-or-leave-it offers to all plans in the market and plans choosing whether to accept these offers. I therefore consider this model as the leading case and use it in my empirical estimation.¹⁰

In Stage 3, plans adjust their premiums in order to maximize their profits after a change in hospital networks. I model this premium adjustment in most of my empirical specifications. See Section IIID for details. I analyze Stages 4 and 5 in Ho (2006): my methodology is outlined in Section IIIA and the results of that study are incorporated where necessary in this paper.

I assume that the plan's choice of quality and products, together with the hospital's choices of capacity, location, services, and quality, are made prior to Stage 1. My analysis conditions on these decisions. I therefore do not explicitly model issues such as product-based price discrimination (the plan's choice between HMO and POS products can be seen as a way of dividing the market into segments with different price elasticities of demand) and the hospital's decision regarding investment in new capacity given that offered by its competitors. Similarly, I assume that hospital merger decisions are made prior to the contractual process.¹¹

⁹ Prices paid to hospitals were regulated at the state level in the 1960s and 1970s. However, since Medicare and Medicaid switched from cost-based to prospective payment systems, and managed care encouraged increased price competition between hospitals, rate regulation has virtually disappeared. It remains only in Maryland: markets in this state are excluded from my supply-side analysis.

¹⁰ However, other models with this property may be possible. For example, Ho (2005) discusses and solves a simple model with no uncertainty in which plans make take-it-or-leave-it offers to hospitals.

¹¹ These assumptions seem reasonable because the relevant variables change more slowly over time than hospital-insurer contracts. For example, over 90 percent of hospitals did not alter their offerings of angioplasty, ultrasound, open heart surgery, or neonatal intensive care units over the four-year period 1997–2001; 70 percent of hospitals changed their capacity levels by fewer than 20 beds over the same four-year period. The correlation between market-level bed capacity (beds per thousand population) in 1980 and that in 2001 is 0.63. Plan product offerings and hospital locations are similarly static. Hospital-insurer contracts, in contrast, are usually renegotiated annually. My goal is to estimate the short-term effects of these hospital and plan characteristics on equilibrium contracts.

My dataset contains no exclusive contracts (either hospitals reaching agreement exclusively with a single insurer or vice versa) and few vertically integrated organizations. Many hospitals and insurers attempted vertical integration in the 1990s but this has become increasingly rare in recent years. Papers such as Lawton R. Burns and Darrell P. Thorpe (2000) and Burns and Mark V. Pauly (2002) suggest that the breadth of skills needed to run both a hospital and a plan is too large for the vertically integrated model to be viable, except in very specific circumstances. The key exception to this pattern is Kaiser Permanente, a dominant HMO in California and elsewhere which owns a large number of hospitals. I do not attempt to explain the vertical integration phenomenon in this paper. I condition on the existence of Kaiser health plans and hospitals in my analysis of both the supply and demand sides of the market (since they are important members of the plan and hospital choice sets, particularly in California) but exclude them from my models of firm behavior.

The health plan must take state and federal legislation into account when choosing its providers. Many states have implemented “any willing provider” laws which prohibit health insurers from excluding qualified health care providers willing to accept the plans’ terms and conditions. However, it is argued that these regulations remove the benefits of managed care, since they prevent plans from trading volume for lower provider prices. Perhaps for this reason they apply to hospitals in only seven states (in other areas they are largely limited to pharmacies). I have data covering two markets within these states; I find that plans are just as likely to exclude hospitals in these markets as elsewhere. I therefore assume that these regulations have no impact on plan decisions in the markets I consider.¹²

II. The Dataset

This paper pulls together information from several datasets. I take data on the characteristics of health insurers from two datasets from Atlantic Information Services (AIS).¹³ The data cover all managed care insurers in 43 major markets across the United States for quarters 3 and 4 of 2002¹⁴ and include information such as premiums earned, number of enrollees, and the tax status of each carrier. If a single carrier offers several plans (such as HMO and POS plans) in the same market, my analysis treats them as separate observations. Multiple HMO (or POS) products offered by a single carrier are grouped into a single observation. I supplement the AIS data with information from the *Weiss Ratings’ Guide to HMOs and Health Insurers* for fall 2002. Data on plan performance come from the *Health Employer Data and Information Set* (HEDIS) and the *Consumer Assessment of Health Plans* (CAHPS) 2000, both of which are published by the

¹² In addition, some states have implemented “essential community provider” laws, which require insurers to contract with providers that offer “essential community services,” such as public hospitals and teaching hospitals, and to contract with enough hospitals to serve the needs of the local population. I assume these regulations do not affect the decision of a plan to exclude any particular hospital, since consumer demand forecasts would prevent it from dropping too many hospitals in any case.

¹³ These are *The HMO Enrollment Report and HMO Directory 2002*. Both are based on plan state insurance filings.

¹⁴ The markets are: Atlanta, GA; Austin, TX; Baltimore, MD; Boston, MA; Buffalo, NY; Charlotte, NC; Chicago, IL; Cincinnati, OH; Cleveland, OH; Columbus, OH; Dallas, TX; Denver, CO; Detroit, MI; Fort Worth, TX; Houston, TX; Indianapolis, IN; Jacksonville, FL; Kansas City, MO; Las Vegas, NV; Los Angeles, CA; Miami, FL; Milwaukee, WI; Minneapolis, MN; New Orleans, LA; Norfolk, VA; Oakland, CA; Orange County, CA; Orlando, FL; Philadelphia, PA; Phoenix, AZ; Pittsburgh, PA; Portland, OR; Sacramento, CA; St. Louis, MO; Salt Lake City, UT; San Antonio, TX; San Diego, CA; San Francisco, CA; San Jose, CA; Seattle, WA; Tampa, FL; Washington, DC; and West Palm Beach, FL.

TABLE 1—DESCRIPTIVE STATISTICS FOR HOSPITALS

	Mean	Standard deviation
Number of beds (set up and staffed)	338.66	217.19
Teaching status	0.195	0.397
For-profit	0.202	0.401
Registered nurses per bed	1.263	0.498
Cardiac services	0.812	0.310
Imaging services	0.539	0.287
Cancer services	0.647	0.402
Birth services	0.857	0.348

Notes: $N = 665$ hospitals. Cardiac, imaging, cancer, and birth services refer to four summary variables defined in Table 4. Each hospital is rated on a scale from 0 to 1, where 0 indicates that no procedures in this category are provided by the hospital, and a higher rating indicates that a less common service is offered.

National Committee for Quality Assurance (NCQA).¹⁵ These data measure clinical performance and patient satisfaction in 1999.

Hospital characteristics are taken from the American Hospital Association (AHA) dataset for 2001. My hospital demand model also uses the MEDSTAT Marketscan Research Database for 1997–1998. This is constituted from privately insured paid medical claims data provided by approximately 50 employer databases across the United States. It provides encounter-level data on all hospital admissions of the relevant enrollees during this two-year period. For each admission, the data include the patient's diagnosis and characteristics, the identity of the hospital, and the type of plan. I focus on inpatient care which, according to the AHA, generated 65 percent of hospital revenues in 2001.

My demand estimation includes all 665 hospitals and all 516 managed care plans in the data. When I consider the supply side, I exclude one of the 43 markets, Baltimore, MD, since the state of Maryland sets hospital prices centrally rather than permitting the plan-hospital bargaining analyzed here. In the remaining 42 markets, I consider only non-Kaiser plans for which premiums are observed; I also exclude a few extremely selective insurers that I regard as outliers.¹⁶ The remaining data contain 441 plans in total. I model these plans' contracts with all non-Kaiser hospitals in each market: there are 633 hospitals in total in the supply-side dataset. I condition on the observed contracts of each excluded plan and hospital in each market.

Descriptive statistics for the hospitals and plans in the data are given in Tables 1 and 2, respectively. The hospitals have 339 beds and 1.26 registered nurses per bed on average; 20 percent are teaching hospitals. The average market share of the HMO/POS plans in the dataset is 3 percent of the nonelderly population in the market. Premiums average \$141 per member per month: 35 percent of insurers are POS plans; 76 percent have been in existence for over 10 years. Plan performance scores vary widely, from an average rating of 0.15 (for the percent of children receiving all required doses of measles, mumps, and rubella (MMR), Hepatitis B, and varicella zoster virus (VZV) vaccines before their thirteenth birthday) to an average of 0.73 (the proportion of women age 52–69 who had received a mammogram within the previous two years). The two most frequently occurring carriers are Aetna and CIGNA, with 15 percent and 10 percent of observations, respectively.

¹⁵ Missing NCQA data represent a significant issue. Dropping plans with missing data could cause selection bias because submission is voluntary. Instead, I include these plans and add dummy variables to my analysis that indicate missing data. See the Appendix for more details.

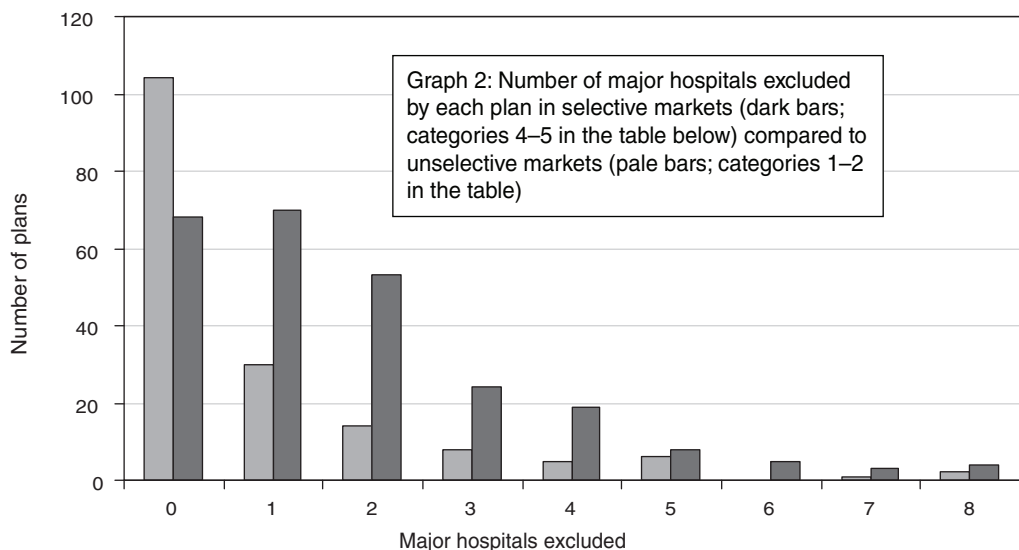
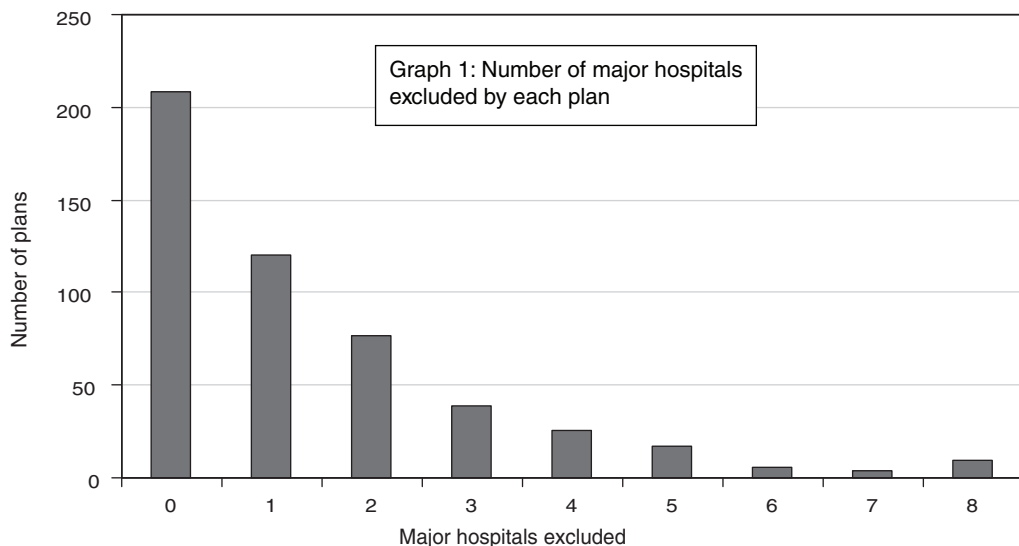
¹⁶ I exclude plans that drop more than four of the top six hospitals because these may have different reasons for their contracting decisions from other plans in the data. I also exclude two specific outliers: Scott and White Health Plan of Austin, TX, and Group Health Cooperative of Puget Sound. These are different from most other plans in the market in that they are locally based, consumer-driven insurers that are heavily focused on primary care.

TABLE 2—DESCRIPTIVE STATISTICS FOR HMO/POS PLANS

Variable	Definition	<i>N</i>	Mean	Standard deviation
Market share	Plan share of nonelderly market	516	0.03	0.04
Premium pmpm (\$)	Premiums earned per member per month	478	140.75	44.27
Physicians per 1,000 population	Number of physician contracts per 1,000 population in markets covered by plan	418	1.56	1.51
Breast cancer screening	Percent of women age 52–69 who received a mammogram within last 2 years	352	0.73	0.05
Cervical cancer screening	Percent of adult women who received pap smear within last 3 years	352	0.72	0.07
Check-ups after delivery	Percent of new mothers receiving a check-up within 8 weeks of delivery	351	0.72	0.11
Eye exams for diabetics	Percent of adult diabetics receiving eye exam within last year	350	0.45	0.11
Adolescent immunization 1	Percent of children receiving all required doses of MMR and Hep B vaccines before 13th birthday	346	0.31	0.16
Adolescent immunization 2	Percent of children receiving all required doses of MMR, Hep B, and VZV vaccines before 13th birthday	313	0.15	0.11
Advice on smoking	Percent of adult smokers advised by physician to quit	213	0.63	0.07
Mental illness checkup	Percent of members seen as outpatient within 30 days of discharge after hospitalization for mental illness	307	0.68	0.15
Care quickly	Composite measure of member satisfaction re: getting care as soon as wanted	304	0.75	0.05
Care needed	Composite measure of member satisfaction re: getting authorizations for needed/desired care	304	0.72	0.06
Age 0–2	Dummy for plans age 0–2 years	516	0.01	0.08
Age 3–5	Dummy for plans age 3–5 years	516	0.06	0.23
Age 6–9	Dummy for plans age 6–9 years	516	0.17	0.37
Aetna	Plan fixed effect	516	0.15	0.36
CIGNA	Plan fixed effect	516	0.10	0.31
Kaiser	Plan fixed effect	516	0.03	0.16
Blue Cross/Blue Shield	Dummy for ownership by BCBS	516	0.16	0.36
POS plan	Dummy for POS plan	516	0.35	0.49

The final dataset analyzed in this paper defines the network of hospitals offered to enrollees by every HMO/POS plan in every market considered in March/April 2003. The information was collected from individual plan Web sites; missing data were filled in by phone. On average, 17 percent of insurer-hospital pairs in my data do not arrange contracts to provide care. The proportion varies from zero in some markets to as many as 40 percent in others. Figure 1 documents the observed variation across both markets and plans in the extent to which plans exclude major hospitals from their networks.¹⁷ Markets are categorized on a scale from 1 to 5, where 1 is the least selective, indicating that each of the five largest plans (by enrollment) contracts with all eight largest hospitals (by number of admissions). In markets ranked 5, at least four of the largest plans

¹⁷ Figure 1 and Table 3 both exclude Baltimore, MD.



Category	Definition	Number of markets	Examples
1	The 5 largest plans (by enrollment) contract with all 8 largest hospitals (by number of admissions)	5	San Antonio, TX Atlanta, GA
2	One plan excludes at least one hospital	10	Boston, MA Columbus, OH
3	Two plans exclude at least one hospital or three plans exclude exactly one hospital each	6	Detroit, MI San Francisco, CA
4	Three plans exclude at least one hospital; one of them excludes more than one	13	Houston, TX Miami, FL
5	Four or more plans exclude at least one hospital each	8	Portland, OR New Orleans, LA

FIGURE 1. VARIATION IN PLAN NETWORKS ACROSS AND WITHIN MARKETS

Note: This figure summarizes the variation in selectivity of plans' hospital networks both across and within markets. Markets are categorized on a scale from 1 to 5, where 1 is the least selective.

TABLE 3—SUMMARY DATA FOR SELECTIVE AND UNSELECTIVE MARKETS

	Unselective markets (category 1 and 2) mean (std dev)	Selective markets (category 4 and 5) mean (std dev)	<i>p</i> -value for difference in means
Market population (million)	2.36 (1.11)	2.36 (1.96)	1.00
Number of HMO/POS plans with over 1 percent market share	6.80 (1.70)	6.57 (1.89)	0.71
Number of hospitals	19.80 (11.40)	21.24 (20.53)	0.78
Beds per 1,000 population	2.78 (1.00)	2.90 (0.99)	0.74
Managed care penetration	0.33 (0.17)	0.35 (0.15)	0.66
Average age of population	34.76 (2.19)	34.31 (1.39)	0.49
Percent of under 65 population age 55–64	0.09 (0.01)	0.09 (0.01)	0.75
Median total family income of population	\$48,890 (\$8,460)	\$46,130 (\$8,642)	0.35
Standard deviation of total family income of population	\$53,687 (\$9,805)	\$52,797 (\$6,511)	0.76
Mean distance between hospitals (miles)	11.71 (5.60)	13.41 (5.12)	0.36
Standard deviation of distances between hospitals (miles)	7.67** (3.37)	10.30** (4.06)	0.04
Number hospitals with open heart surgery	8.07 (3.67)	10.19 (8.59)	0.31
<i>N</i>	15	21	

**Significant difference in means at the 5 percent level.

exclude at least one major hospital; the other categories lie between these extremes. Markets are fairly evenly spread across the five categories: 15 markets are ranked 1 or 2 (not selective) and 21 are ranked 4 or 5 (very selective). The figure also shows the distribution of plans by the number of major hospitals excluded and the variation in this distribution across types of market. Plans' selective behavior varies widely: 208 plans exclude no major hospitals, but 62 plans exclude at least 4 of the 8 major hospitals in their markets. Further details on all these datasets are set out in Appendix A.

Table 3 compares the means of a number of market characteristics in selective and unselective markets. There are few significant differences. Selective markets do not have significantly smaller populations, higher managed care penetration, more hospitals, or more beds per capita than unselective markets, and are not clustered geographically. There are no significant demographic differences. The only difference that is significant at $p = 0.05$ (or, in fact, at $p = 0.2$) is the standard deviation of the distances between hospitals in the market. Plans seem to be more willing to exclude hospitals in areas where hospitals are clustered into several groups, perhaps because each provider in a given group is a reasonable substitute for the others. The raw data therefore do not offer an obvious explanation for the observed variation; however, they do provide a hint that demand effects may be important. These are taken into account in my analysis.

The hospital-level data offer further clues to help explain the observed contracting choices. Table 4 defines four variables that summarize the services offered by each hospital. The summary

TABLE 4—DEFINITION OF HOSPITAL SERVICES

Cardiac	Imaging	Cancer	Births
1. Cardiac catheterization lab	1. Ultrasound	1. Oncology services	1. Obstetric care
2. Cardiac intensive care	2. CT scans	2. Radiation therapy	2. Birthing room
3. Angioplasty	3. MRI		
4. Open heart surgery	4. SPECT		
	5. PET		

Notes: This table sets out the definition of the hospital service variables summarized in Table 1. Hospitals were rated on a scale from 0 to 1 within four service categories, where 0 indicates that no services within this category are provided by the hospital, and a higher rating indicates that less common (assumed to be higher-tech) service in the category is offered. The categories are cardiac, imaging, cancer, and births. The services included in each category are listed in Table 5. The exact methodology for rating hospitals is as follows. If the hospital provides none of the services, its rating = 0. If it provides the least common service, its rating = 1. If it offers some service X but not the least common service, its rating = $(1 - x)/(1 - y)$, where x = the percent of hospitals offering service X and y = the percent of hospitals offering the least common service.

variables cover cardiac services, imaging, cancer, and birth services. Each hospital is rated on a scale from 0 to 1, where 1 implies that the hospital offers the least common of a list of relevant services and 0 implies that it offers none of the services. I interact these variables with consumer characteristics in the model of demand for hospitals. They can also be used to investigate which hospital characteristics are correlated with market share. Table 5 sets out the results of a regression of hospital market shares on hospital characteristics. All four service variables, and the indicator for teaching hospitals, are positively and significantly related to market share. Together with hospital location, they will be key determinants of hospitals' attractiveness to consumers (which, as we shall see, generates market power and the ability to negotiate positive profit margins) later in the analysis.

III. Demand Estimation and Producer Surplus Calculation

In order to understand the equilibrium network outcomes, I need to analyze Stages 4 and 5 of the model, in which consumers choose their health plans taking into account the hospitals they expect to visit in the coming year. The parameter estimates generated in Ho (2006) are used as an input to this paper's supply-side analysis. The demand estimation process has three stages outlined below.

A. Hospital Demand

The first step is to estimate demand for hospitals using a discrete choice model that allows for observed differences across individuals. With some probability, consumer i (whose type is defined by age, gender, and zip code tabulation area (ZCTA)) becomes ill. His utility from visiting hospital h given diagnosis l is given by

$$(1) \quad u_{i,h,l} = \eta_h + \mathbf{x}_h \alpha + x_h v_{i,l} \beta + \varepsilon_{i,h,l}$$

where x_h, η_h are vectors of observed and unobserved hospital characteristics, respectively, $v_{i,l}$ are observed characteristics of the consumer such as diagnosis and location, and $\varepsilon_{i,h,l}$ is an idiosyncratic error term assumed to be i.i.d. Type 1 extreme value.¹⁸ Hospital characteristics include location, the number of beds, the numbers of nurses and doctors per bed, and details

¹⁸ This model was first proposed in Daniel L. McFadden (1973).

TABLE 5—RELATION OF HOSPITAL CHARACTERISTICS TO MARKET SHARES

	Coefficient estimate	
Cardiac services	0.732** (0.104)	0.676** (0.072)
Imaging services	0.233** (0.107)	0.224** (0.074)
Cancer services	0.158** (0.079)	0.299** (0.054)
Birth services	0.507** (0.082)	0.394** (0.056)
Teaching hospital	0.243** (0.074)	0.461** (0.051)
Constant	-4.484** (0.097)	-0.005 (0.007)
Market FEs?	No	Yes
Adjusted R^2	0.27	0.69

Notes: Regression of the log of hospital market shares on hospital characteristics. $N = 633$ hospitals (the 665 providers in the full dataset less 14 Kaiser hospitals and 18 hospitals in Baltimore, MD, which were excluded from the supply-side analysis). Standard errors are reported in parentheses. Cardiac, imaging, cancer, and birth services refer to the four hospital service variables defined in Table 4.

**Significant at the 5 percent level.

*Significant at the 10 percent level.

of services offered, ownership, and accreditation. This equation is estimated using standard maximum likelihood techniques and micro (encounter-level) data from the MEDSTAT dataset described above.

I would ideally estimate consumers' hospital choices using data for managed care enrollees as well as for indemnity and PPO enrollees, since this would avoid analyzing the behavior of a self-selected sample. This is not feasible, however, because the MEDSTAT data do not identify the hospital networks offered by each managed care plan, so the choice sets of managed care enrollees are unobserved. Instead, I consider only the choices made by indemnity and PPO enrollees, whose choice set is unrestricted. I assume that indemnity/PPO enrollees have the same preferences over hospitals as managed care enrollees, conditional on their diagnosis, income, and location. This assumption has been made several times in the existing literature¹⁹ and is supported by secondary analysis of a market in which HMO/POS enrollees have nearly complete access across hospitals.²⁰ The average fee-for-service plan enrollee probably has different preferences over hospitals from the average managed care enrollee before he knows his diagnosis: for example, he may have a stronger desire for choice. However, when informed that he has a particular disease, I assume that he would choose the same hospital as the average managed care enrollee of the same age and living in the same zip code.²¹

¹⁹ For example, Town and Vistnes (2001) use data on the hospital selection decisions of Medicare enrollees, assuming that the Medicare population's valuation of hospitals is a reasonable proxy for that of HMO enrollees. Capps, Dranove, and Satterthwaite (2003) make a similar assumption to justify considering patients insured by Medicare, Medicaid, Blue Cross/Blue Shield, and indemnity plans.

²⁰ This analysis estimates the hospital choice model using MEDSTAT data for HMO/POS enrollees in Boston, MA, a market in which I observe that most plans contract with all hospitals. The estimated coefficients are not identical, but are broadly similar, to those estimated using PPO/indemnity enrollee data for Boston only. Only 3 out of 36 hospital dummy coefficients and 2 out of 32 interaction terms are different in sign across the two models, and both are significant at $p = 0.1$. I take this to be sufficient evidence to support the assumption. See Ho (2006) for more details.

²¹ I make a second simplifying assumption with regard to prices. PPO enrollees may be required to pay additional copayments or deductibles if they choose to go out of network. These financial penalties, and the hospitals in the PPO

B. Expected Utility Calculation

I use the estimated coefficients from the hospital demand equation to predict the utility provided by each plan's hospital network. Individual i 's expected utility from the hospital network offered by plan j in market m is calculated as

$$(2) \quad EU_{i,j,m} = \sum_l p_{i,l} \log \left(\sum_{h \in H_j} \exp(\eta_h + \mathbf{x}_h \hat{\alpha} + \mathbf{x}_h v_{i,l} \hat{\beta}) \right),$$

where $p_{i,l}$ is the probability that individual i will be hospitalized with diagnosis l .^{22, 23}

C. Health Plan Demand

Finally, I use aggregate data from AIS, the NCQA, and the AHA to estimate the health plan demand model. I use a methodology similar to that set out in Berry, Levinsohn, and Pakes (1995). The utility of individual i from plan j in market m is given by

$$(3) \quad \tilde{u}_{i,j,m} = \xi_{j,m} + z_{j,m} \lambda + \gamma_1 EU_{i,j,m} + \gamma_2 \frac{prem_{j,m}}{y_i} + \omega_{i,j,m}$$

where $z_{j,m}$ and $\xi_{j,m}$ are observed and unobserved plan characteristics, respectively, $prem_{j,m}$ are plan premiums, y_i is the income of individual i , and $\omega_{i,j,m}$ represents idiosyncratic shocks to consumer tastes, again assumed to be i.i.d. Type 1 extreme value. The characteristics included in z are premium, the size of the physician network, plan age, a list of eight clinical quality variables (taken from the NCQA's HEDIS dataset), and two variables summarizing consumer assessment of plans on dimensions such as availability of needed care and speed with which care is received (from their CAHPS dataset).

The model is completed by defining the outside good. The simplest definition would be a composite of nonmanaged care private coverage and no insurance. However, indemnity coverage and no coverage are at opposite ends of the spectrum in terms of price and many aspects of quality, so this outside good would be nonhomogenous. Instead, I define the outside good as "choosing to be uninsured" and create a separate choice in each market defined as "choosing indemnity or PPO insurance" and assumed to be homogenous in each market. (See Appendix A for details.)

The premium variable is endogenous to the plan demand equation. The instruments used, in addition to the usual set of plan characteristics, are the average hourly hospital wage and the average weekly nurse wage across the markets in which the health plan is observed to be active. The main assumption required for these to be valid instruments is that health plan costs are

network, are not identified in the dataset; that is, the "price" of the hospital at the point of service is unobserved. I therefore assume that out-of-pocket prices charged to patients on the margin are zero for both PPO and indemnity patients. This may be reasonable, particularly where prices take the form of increased deductibles, since many of these patients are likely to have spent beyond their deductible before making their decision. The average copayment for PPO enrollees in my data was \$289 for an average stay of 4.8 days. This is only around 3 percent of the average cost per admission. I test the assumption further by reestimating the hospital choice model using data for indemnity enrollees only (roughly half the total sample). The results were similar to those for the main specification. See Ho (2006) for details.

²² Diagnosis probabilities conditional on age, gender, and admission to hospital were taken from the MEDSTAT data; probabilities of admission to hospital given age and gender come from the National Hospital Discharge Survey 2000.

²³ The expectation over values of ε_{ihl} implies an assumption that each individual's ε is unknown when he chooses his plan. I estimated the model implied by the alternative assumption, that he knows his ε when making the choice, as a robustness test. Using the new expected utility variable in the health plan choice model had little effect on the final results. Ho (2006) contains further details on this robustness test.

TABLE 6—RESULTS OF PLAN DEMAND ESTIMATION

	Coefficient estimate
Premium (\$00 pmpm)	-0.94 (1.13)
Expected utility from hospital network ($EU_{rep_{jm}}$ or EU_{ijm})	0.59** (0.21)
Premium (\$00 pmpm) / Income (\$000 per year)	0.002 (43.9)
Physicians per 1,000 population	0.21** (0.09)
Breast cancer screening	-0.38 (2.66)
Cervical cancer screening	4.40** (2.09)
Check-ups after delivery	0.18 (1.38)
Eye exams for diabetics	-1.19 (1.60)
Adolescent immunization 1	-4.11** (1.17)
Adolescent immunization 2	3.08 (3.76)
Advice on smoking	6.17** (2.08)
Mental illness check-ups	2.70** (1.30)
Care quickly	0.78 (5.63)
Care needed	0.85 (3.99)
Plan age: 0–2 years	1.36 (0.97)
Plan age: 3–5 years	-0.64 (1.97)
Plan age: 6–9 years	-0.25 (0.58)
POS plan	-1.11** (0.13)
Constant	-10.50* (5.65)
Large plan fixed effects	Yes
Market fixed effects	Yes

Notes: $N = 559$ plans (the 516 HMO/POS plans in the full dataset plus one indemnity/PPO option in each market). Standard errors (adjusted for the three-stage estimation process) are reported in parentheses.

**Significant at the 5 percent level.

*Significant at the 10 percent level.

correlated with premiums but not with unobserved plan quality. See Ho (2006) for a discussion of the choice of instruments.

The results of this third stage of the analysis are reproduced in Table 6. Standard errors are adjusted to take account of the variance introduced by the previous stages of estimation. I find that consumers place a positive and significant weight on their expected utility from the hospital network when choosing a plan. The coefficient magnitudes imply that a one-standard-deviation

increase in expected utility is equivalent to a reduction in premium of \$39 per member per month (a little less than one standard deviation).

D. Producer Surplus Generated by the Network

The next step is to use the demand estimates to predict the producer surplus generated by each insurance plan when it contracts with each potential hospital network, that is, the total profit to be divided between the plan and all the hospitals with which it contracts. The producer surplus generated by plan j in market m when it contracts with hospital network H_j is

$$(4) \quad S_{j,m}(H_j, H_{-j}) = \sum_i \left(n_i s_{i,j,m}(H_j, H_{-j}) \left[\text{prem}_{j,m} - p_i \sum_{h \in H_j} s_{i,h}(H_j) \text{cost}_h \right] \right),$$

where n_i is the population in consumer-type cell i (defined by ZCTA, age, and gender), p_i is the probability that a type- i person will be admitted to hospital, cost_h is the average cost of treatment at hospital h , and $\text{prem}_{j,m}$ is plan j 's premium in market m . The quantities $s_{i,j,m}(H_j, H_{-j})$ and $s_{i,h}(H_j)$ are plan j 's and hospital h 's predicted shares of type- i people when networks H_j and H_{-j} are offered by plan j and other plans, respectively. These are predicted using the demand estimates and take account of the flow of consumers across plans, and across hospitals given their choice of plans, in response to network changes.

The surplus calculation incorporates plans' premium adjustments in response to changes in hospital networks. I use a two-step process to allow plan j to predict how much its own premium and those of other plans in the market will change if it deviates from its observed network.²⁴ First I estimate the supply model set out in Section V assuming fixed premiums. Then I allow all plans to simultaneously adjust their premiums to maximize their profits (revenues less prices paid) where prices are determined by the first-stage estimates.²⁵ This premium adjustment is conducted as part of the producer surplus calculation for all networks considered.

The calculation also takes account of hospital capacity constraints. If any network combination implies that any hospital is over 85 percent of its maximum capacity level, I reallocate patients randomly to noncapacity-constrained hospitals in the market. I assume that patients are treated in the order in which they arrive and that the timing of sickness is random: each plan therefore has the same percentage of enrollees reallocated for any given capacity-constrained

²⁴ Unfortunately, I do not have access to panel data and so cannot observe the true reaction of plan premiums to network changes over time. I investigate the magnitudes of premium responses in my cross-sectional data by regressing observed premiums on market fixed effects, dummy variables for the ten largest (nationally active) carriers in the data, and network characteristics. I find evidence that plans contracting with the most popular hospitals (those defined as stars in the main analysis below) have higher premiums than other insurers: the coefficient is significant at $p = 0.051$. Other network characteristics in the regression did not have significant coefficients.

²⁵ The new premiums are the result of an equilibrium price-setting game computed by finding the values that simultaneously minimize all plans' first-order conditions. I impose these predicted premiums for both observed and unobserved contracts. The median difference between the predicted and observed values for observed contracts is \$32.57, or 23 percent of the observed value. The prediction error may be caused by such simplifying assumptions as ignoring plans' nonhospital costs and assuming fixed hospital prices across patients. It is also worth noting that observed premiums are measured with error: only average plan revenue per enrollee is observed (including both employer and employee contributions). Some of the difference between predicted and observed premiums may therefore be due to noise in the observed rather than the predicted values. In any case, the error is most likely to bias the results of the full model if it affects premium changes in response to network changes rather than premium levels. I regress the difference between my predictions and observed premiums on market fixed effects, dummy variables for the ten largest carriers in my data, and network characteristics, and find no evidence that the prediction error is correlated with plans' network choices.

hospital. The adjustment affects patients' hospital choices and therefore their predicted costs of care, but does not affect consumers' choices of plan or premium levels.²⁶

The surplus definition does not include plans' nonhospital variable costs. Each plan faces a number of costs of enrolling consumers: these include payments to PCPs and prescription drug costs, for example, in addition to the costs of treatment at hospitals. Unfortunately, I do not have access to data on plan variable costs and therefore cannot include them in the surplus term.²⁷ I account for this issue later in the analysis by estimating the cost of enrolling consumers directly. The details of this specification are discussed in Section V.

IV. Intuition Regarding the Price Negotiation

The producer surplus variable alone is unlikely to be enough to explain the contracts we observe in the data. The reason relates to the bargaining process that determines hospital prices. My definition of producer surplus measures the effect of hospital-plan contracts on the profits to be divided between the plan and all the hospitals in its network. It takes account of the effect of a particular contract on the plan's profits from other hospitals with which it already has contracts. However, it does not account for the effect of a hospital's contract with one plan on its revenues and profits from other plans in its market. This section provides some intuition on the importance of this issue for different types of hospitals. The goal is to justify my focus on particular hospital characteristics in the empirical analysis.

Consider a simple example of the negotiation in Stages 1 and 2 of the five-stage game set out in Section I. Insured consumers receive two types of services from their plan: acute care from the hospitals in the network and preventive services from the plan's PCPs. In Stage 1 of the game, hospitals make simultaneous take-it-or-leave-it offers to all plans in the market. In Stage 2, plans simultaneously choose which offers to accept. When a hospital-plan pair has agreed on a contract, the hospital is required to treat every enrollee from the plan who requests care, provided it has spare capacity. Each plan then sets a single premium level and consumers choose their plans for the year.

Hospitals that are undifferentiated and that do not expect to be full receive zero profits in this model. Plans capture 100 percent of the surplus created and will include hospital h in their networks provided it generates positive producer surplus. Evidence from the interviews I conducted indicates that this may be a reasonable representation of many markets where managed care is strong and hospitals compete for contracts. The executive director of one hospital system described a potential outcome in such markets: "There are examples where there were too many hospitals in an area and the plans played them off against each other to the point where the price paid was no more than marginal cost."

The more interesting situation arises when hospitals tailor their characteristics in order to capture positive profits. Interviewees noted that the negotiations could be very different in these markets. A hospital director said the following: "In market X [where hospitals are very strong], the prices [the best hospitals] charge are based on their very high patient satisfaction results and

²⁶ A hospital is predicted to be over 85 percent of maximum capacity if predicted admissions \times average length of stay at the hospital is greater than 85 percent of the number of beds \times 365 days. By using the surplus variable without adjusting consumers' choices of plan, I am assuming that the plan does not expect consumers to predict their probability of treatment at each hospital in its network when choosing their insurer. Instead, consumers are expected to assume they will have access to every hospital on the list. Consumers may update this belief if a hospital is consistently capacity constrained (although many of the non-Medicare, non-Medicaid enrollees considered in this paper will have little experience in seeking hospital treatment on which to base their updates). Unfortunately, without a panel dataset, there is no variation in the data to identify the extent of any such updating.

²⁷ The analysis does allow for the existence of additional fixed costs, since these would cancel out when we consider the surplus change from a change in networks.

their strong reputation. They can get high prices from any plan in the market and they don't need them all." The CEO of a small hospital in a different market had a similar story: "Large [hospitals] in this market can dictate whatever prices they want. The bigger names can demand the higher prices."

Consumers may be willing to switch plans if necessary to access one of these popular hospitals. In that case the provider might be able to increase its profits by contracting selectively. By contracting only with the plans willing to pay the highest prices, it can effectively price discriminate, concentrating high-valuation consumers in the high-priced plans. The hospital's revenues per patient, and its total profits, may be higher under this arrangement than under unselective contracting where some high-valuation consumers might choose lower-priced insurers, reducing the hospital's average revenues per patient.²⁸ The effect follows most clearly for three types of hospitals:

- (i) The hospital may be sufficiently differentiated from its competitors that all or most consumers are willing to pay more for its services than for other hospitals in the market and are willing to switch plans to access it. This is most likely for hospitals that offer very high-tech services, teaching hospitals, and hospitals with a high reputation for quality.²⁹ I describe these as star hospitals. I discuss in Section V the characteristics used to define these providers in the empirical estimation.
- (ii) The hospital may be somewhat differentiated from its competitors and may expect to be capacity constrained: that is, it may expect to fill its beds without treating all consumers who wish to access it. In many models this implies two effects. First, capacity-constrained hospitals are particularly likely to choose to contract selectively because this strategy acts as a profitable form of rationing, helping the hospital to avoid lower-valuation consumers who would otherwise displace those with a higher willingness to pay. The capacity constraint may also act as a commitment device, helping persuade the plan that the hospital will contract selectively. Second, the capacity constraint may alter the nature of the bargaining game, essentially forcing plans to compete for contracts with the hospital and pushing its prices up even further.³⁰
- (iii) Finally, if a sufficiently large proportion of the hospitals in the market merge to form a single system, the combined organization may be very attractive to consumers and the number of competitors that remain may be small. This, too, may imply that many consumers are willing to switch plans to access the system and that selective contracting increases its profits.

In all three cases, a given insurer may choose to exclude the hospital, focusing instead on those consumers whose low valuation for h and higher valuation for its other services prevent them

²⁸ This intuition is similar to a monopolist that restricts volume in order to maximize profits. Perfect price discrimination across consumers is impossible because enrollees are aggregated into plans, each of which charges a single premium, and may choose to move between plans. Price discrimination through selective contracting may be the hospital's most profitable option.

²⁹ Hospital location also plays an important role in differentiating the provider from its competitors. However, location alone may not be sufficient to generate star hospital status.

³⁰ This benefit may prompt particular hospitals to choose to be capacity constrained. The intuition in the case where hospitals make offers to plans is similar to that in David M. Kreps and Jose A. Scheinkman (1983). Allowing firms to choose capacity levels before bargaining begins permits them to move from Bertrand to Cournot competition, implying positive hospital profits. The case where plans make offers to hospitals is demonstrated in Ho (2005).

from switching, if other plans have a higher maximum willingness to pay for the contract.³¹ Thus we may observe a selective outcome even in cases where the producer surplus generated by the contract is positive.

The existence of systems may also help to explain the contracts that are agreed upon despite a negative incremental surplus. I observe in the data that some plans contract with some but not all members of a hospital system, but this practice is infrequent. I rationalize this observation with the idea that, if a hospital system has significant market power (as in point iii), it may impose penalties on plans that contract with some but not all of its members. Even systems with little market power may choose a bundling strategy, charging relatively more for contracts with individual hospitals than for those with the entire organization, to maximize the surplus captured from each plan. In both cases plans may be deterred from cherry-picking from the members of a system.³² The selective equilibrium may have implications for welfare. The simple model in Ho (2005) demonstrates that it may be inefficient for a plan to exclude a hospital even if that provider is full in equilibrium and even if the consumers with the highest value for hospital h are the ones treated. The inefficiency is generated because consumers are forced to make suboptimal choices across health plans in order to gain access to the hospital. The resulting loss of consumer welfare, which may outweigh the gain derived when the highest-valuation patients are given preferential access to h , would be avoided if both plans contracted with it. The intuition is similar for the examples concerning systems and attractive hospitals.³³

V. A Model for Estimation

Having constructed the producer surplus variable, and with an understanding of the bargaining effects that need to be included in the contracting model, the next step is to set out a full model for estimation. I adopt the methodology for analyzing buyer-seller networks that was developed in Pakes et al. (2006) and discussed further in Pakes (2008). Those papers set out conditions under which the inequality constraints generated by a Bayes-Nash equilibrium assumption in multiple-agent games can be used as a basis for estimation and inference. I begin by formally defining the game being played between hospitals and plans.

A. The Bargaining Game

Define the set of hospitals in market m to be $h = \{1, 2, \dots, H_m\}$ and the set of plans to be $j = \{1, 2, \dots, J_m\}$. As discussed above, the game has two stages. In Stage 1 all hospitals make simultaneous take-it-or-leave-it offers to all plans in the market. In Stage 2 all plans simultaneously respond. The information set of agent i before decisions are made is denoted by I_{im} . The set of actions available to hospital h is denoted D_h . Each element is a J_m -tuple of contracts to be offered, one to each plan. One potential offer is the null contract, $d_{h,j} = \phi$, an offer so high that it is never accepted. In the analysis that follows I restrict attention to the hospital's choice of

³¹ A particular plan may be willing to pay less for the contract than other insurers for two reasons. First, it may have a better outside option than other plans due to variation in consumers' preferences for other plan characteristics. Second, its nonswitching enrollees may have a lower valuation for h than those in other plans.

³² The efficient outcome would result in the plan contracting only with hospitals with which it generated positive surplus. A system with high market power could demand a share of the profits generated from nonsystem hospitals. Some friction is required to prevent this outcome: this could be a cost of contracting that is paid once per nonsystem hospital and only once per system, or an inability of one hospital to transfer funds to another member of the same system to compensate it for lost revenues.

³³ Ho (2006) estimates that the restriction of consumer choice of hospitals, and the resulting distortion of their choice of plans, imply a welfare loss of \$1.04 billion per year assuming fixed prices. The model in Ho (2005) shows that a welfare loss is still possible in the case where prices are permitted to adjust.

whether to make a nonnull offer to a particular plan. I do not attempt to model the terms of each offer. The set of actions available to plan j is D_j . Each element is a H_m -tuple of zeros and ones, one for each hospital, where a value of 1 indicates that a contract was accepted.

Offers are private information, known only to the relevant hospital and plan and not to other plans. Firms are assumed to have “passive beliefs”: for example, if a plan gets an alternative offer from a particular hospital, this will not change its beliefs about the offers made to its competitors.³⁴ I use a Bayes-Nash equilibrium concept, implying that the strategy chosen by each agent maximizes its expected profits.

B. The Estimation Strategy

The profit of plan j is the surplus generated, given its chosen network H_j minus its costs:

$$(5) \quad \pi_{j,m}^P(H_j, \mathcal{H}_{-j}, \mathcal{X}, \boldsymbol{\theta}) = S_{j,m}(H_j, \mathcal{H}_{-j}) - c_{j,m}^{HOSP}(H_j, \mathcal{H}_{-j}, \mathcal{X}, \boldsymbol{\theta}) - c_{j,m}^{NONHOSP}(H_j, \mathcal{H}_{-j}, \mathcal{X}, \boldsymbol{\theta}),$$

where \mathcal{X} are plan, hospital, and market characteristics to be defined in Section VF, and $\boldsymbol{\theta}$ is the parameter vector to be estimated. A script typeface is used for variables that the plan views as random. For example, \mathcal{H}_{-j} and \mathcal{X} may not be known to the plan when it makes its choice, but H_j , the realized choice of plan j , is not a random variable. The costs of the plan’s contracts with hospitals (generated by hospital profits) are $c_{j,m}^{HOSP}(H_j, \mathcal{H}_{-j}, \mathcal{X}, \boldsymbol{\theta})$. Its nonhospital costs are denoted $c_{j,m}^{NONHOSP}(H_j, \mathcal{H}_{-j}, \mathcal{X}, \boldsymbol{\theta})$. Details on these hospital and nonhospital cost terms are set out in Section VF.

I allow for two sources of randomness. The first is measurement error on the part of the econometrician: I denote this u_{j,H_j} . I expect this to be important. In particular, the hospital cost data used as an input to the producer surplus calculation provide information only on the average hospital expense per admission, including costs such as depreciation and interest expense and not distinguishing between types of procedure or diagnosis. This is a noisy measure of the true cost of a particular inpatient stay. We can write the plan profits observed by the econometrician as

$$(6) \quad \pi_{j,m}^{P,o}(H_j, \mathcal{H}_{-j}, \mathcal{X}^o, \boldsymbol{\theta}) = \pi_{j,m}^P(H_j, \mathcal{H}_{-j}, \mathcal{X}, \boldsymbol{\theta}) + u_{j,H_j},$$

where \mathcal{X}^o is a random variable whose realization, x^o , contains the observable determinants of profits. I assume that the noise is mean-zero conditional on the plan’s information set: $E(u_{j,H_j} | I_{j,m}) = 0$. Second, the plan may predict its profits from contracting with any particular network with error, perhaps because of uncertainty regarding other plans’ network choices. I denote this error φ_{j,H_j} . The plan’s prediction of its profits from choosing network H_j can therefore be written as

$$(7) \quad E(\pi_{j,m}^P(H_j, \mathcal{H}_{-j}, \mathcal{X}, \boldsymbol{\theta}) | I_{j,m}) = \pi_{j,m}^P(H_j, \mathcal{H}_{-j}, \mathcal{X}, \boldsymbol{\theta}) - \varphi_{j,H_j},$$

where $E(\varphi_{j,H_j} | I_{j,m}) = 0$ by construction.

The primary identifying assumption in Pakes et al. (2006) is a necessary condition for a Bayes-Nash equilibrium: that plan j ’s expected profits from the observed network H_j must be higher than its expected profits from the alternative network formed by reversing its contract with

³⁴ This assumption is needed to construct inequalities using the hospital profit equation. See below for details. Rey and Tirole (2007) note that a passive beliefs equilibrium may not exist in a price-setting game if there is sufficient substitutability between products. I assume existence of an equilibrium in passive beliefs. Pakes (2008) notes that existence of a pure strategy equilibrium is not guaranteed, even in a full information game, and that, given the complexity of the empirical problem, the hospital profit equation would be very difficult to estimate assuming anything other than passive beliefs.

any hospital h in the market and holding all other plan and hospital contracts fixed.³⁵ I denote this alternative network H_j^h . That is, I assume that

$$(8) \quad E(\pi_{j,m}^P(H_j, \mathcal{H}_{-j}, \mathcal{X}, \boldsymbol{\theta}) | I_{j,m}) \geq E(\pi_{j,m}^P(H_j^h, \mathcal{H}_{-j}, \mathcal{X}, \boldsymbol{\theta}) | I_{j,m})$$

for every hospital h in the market.³⁶ Together with equations (6) and (7), this assumption implies that

$$\pi_{j,m}^{P,o}(H_j, \mathcal{H}_{-j}, \mathcal{X}^o, \boldsymbol{\theta}) - \pi_{j,m}^{P,o}(H_j^h, \mathcal{H}_{-j}, \mathcal{X}^o, \boldsymbol{\theta}) - (u_{j,H_j} - u_{j,H_j^h}) - (\varphi_{j,H_j} - \varphi_{j,H_j^h}) \geq 0.$$

Taking expectations and using the fact that $E(u_{j,H_j} | I_{j,m}) = 0$ and $E(\varphi_{j,H_j} | I_{j,m}) = 0$, we find that

$$(9) \quad E(\pi_{j,m}^{P,o}(H_j, \mathcal{H}_{-j}, \mathcal{X}^o, \boldsymbol{\theta}) - \pi_{j,m}^{P,o}(H_j^h, \mathcal{H}_{-j}, \mathcal{X}^o, \boldsymbol{\theta}) | z_{j,m}) \geq 0,$$

for any $z_{j,m} \in I_{j,m}$. Translating expectations into sample means, the equation for estimation is therefore³⁷

$$(10) \quad \frac{1}{M} \sum_m \frac{\sqrt{n_m}}{n_m} \sum_{j=1}^{n_m} [(\pi_{j,m}^{P,o}(H_j, \mathcal{H}_{-j}, \mathcal{X}^o, \boldsymbol{\theta}) - (\pi_{j,m}^{P,o}(H_j^h, \mathcal{H}_{-j}, \mathcal{X}^o, \boldsymbol{\theta})) \otimes g(z_{j,m})] \geq 0,$$

where M is the number of markets in the sample, n_m is the number of plans in market m , \otimes is the Kronecker product operator, and $g(z_{j,m})$ is any positive-valued function of $z_{j,m}$. Each market is weighted by the square root of the number of plans in the market, since we expect less noise in the market average for markets containing many plans. All $\boldsymbol{\theta}$ that satisfy this system of inequalities are included in the set of feasible parameters. If no such $\boldsymbol{\theta}$ exists we find values that minimize the Euclidean norm of the amounts by which the inequalities are violated.

Any feasible alternative networks could be used to generate these moments. I focus on small deviations, affecting just a single contract, because, as discussed below, the assumptions implicit in the bargaining model are most reasonable in this context.³⁸ I consider seven inequality conditions. The first six are defined by reversing the plan's contracts with each of the six largest

³⁵ Hospital offers cannot change in response to the deviation since they were made in Stage 1. See Section VID for a discussion of renegotiation-proofness. Other plans will not respond because their choices are made simultaneously with that of the deviating plan. Reversing a contract means removing the contract if it exists or introducing it if it is not observed in the data. That is, I assume that every observed contract must increase the plan's expected profits. Any contract that does not exist in the data must decrease the expected profits of the plan that turned it down.

³⁶ Several combinations of networks may satisfy this necessary condition; that is, there may be multiple potential equilibria. This does not prevent consistent estimation of the parameter vector $\boldsymbol{\theta}$: I simply search for parameters consistent with the assumption that the observed set of networks constitute a Bayes-Nash equilibrium, without attempting to model how that equilibrium was chosen from the set of potential equilibria.

³⁷ The obvious alternative would be to use a multinomial logit model. That approach would involve using the same identifying assumption, that plans choose their hospital networks to maximize their expected profits. However, we would have to impose the additional assumption that the errors are i.i.d. extreme value. We would then estimate using maximum likelihood, identifying the region of unobservable space that generates the observed choices and choosing parameters to maximize the probability that the unobservables fall within this region. Section VIE discusses the problems with this methodology.

³⁸ The main assumptions are that plans make decisions simultaneously (implying that other plans do not respond to plan j 's deviation) and that hospitals do not make follow-up offers if their initial offers are rejected. The latter assumption is discussed in Section VD. Using alternative networks where the plan reversed two, rather than one, hospital contracts did not significantly change the results but did lead to broader confidence intervals.

hospitals, in turn holding all other contracts fixed.³⁹ The seventh is an average over the analogous inequalities for all remaining hospitals in the market.

The estimation also includes inequalities generated from the hospital's profit equation. Pakes et al. (2006) and Pakes (2008) show how to derive these inequalities. We write hospital profits as $\pi_{h,m}^H(M_h, \mathcal{M}_{-h}, \mathcal{X}, \theta)$, where M_h is the network of plans contracting with hospital h , and \mathcal{M}_{-h} is the set of networks of other hospitals in the market. The assumptions regarding unobservables are analogous to those for the plan profit equation.

Consider the hospital's alternatives to its observed action. As noted above, we restrict attention to the hospital's choice between offering a null contract and offering a nonnull contract to each insurer. In order to construct inequalities we need to predict the network that the hospital expects to be generated if it changes its offer to a particular plan. The assumption of passive beliefs is useful here: the hospital believes that, if it changes its offer, the plan will not change its expectations about the offers made to other plans and therefore about other plans' network choices. However, there is still an issue related to the fact that this is a two-stage game. If a hospital (h) deviates by offering the null contract to a plan (j), then j 's optimal response may include a change in response to the offers made by other hospitals. Pakes et al. (2006) consider two methodologies to account for this issue. First, we assume that the hospital can make an alternative offer to plan j that will prompt j to drop h from its network without changing its contracts with other hospitals.⁴⁰ This leads to a straightforward inequality: that the hospital expects a positive profit difference between a network that includes a plan that accepted its offer and the network that excludes this plan and holds everything else fixed.

The second methodology relaxes this assumption. Now we must allow plan j to change its response to other hospitals if h 's offer deviates to the null contract.⁴¹ One way to model this scenario is to find the minimum of hospital h 's profits over all possible choices plan j could make in response to the deviation, given its other contracts and the networks of all other plans in the market. The difference between the realized profit and this minimized counterfactual profit can then be used as the basis of the inequality, since if the true plan response gives the hospital lower expected profits than the observed network, the inequality must also hold for the minimizing response. This methodology substantially increases the computational burden of the estimator. I therefore begin with the simpler methodology and relax the assumption as a robustness test.^{42, 43}

The instruments are required to be independent of the error terms u_{j,H_i} and φ_{j,H_j} ; that is, they must be perfectly observed by the econometrician and known to the firms when they make their

³⁹ The largest hospitals are defined by numbers of beds. The six largest hospitals cover an average of 57 percent of the admissions to hospitals. There are at least seven hospitals in each market in the database.

⁴⁰ This assumption seems reasonable for plans that are observed to contract with all or most providers in the area. It is more troubling for plans that contract with just a small subset of hospitals; these may choose to respond to a high offer from one hospital by replacing it with another that was not previously included in the network. I conduct a robustness test for the effect of this assumption by replacing the hospital inequalities with those generated by a different condition: that the profit paid by each plan to each hospital is weakly positive. The estimates are less precise than those from the main specification but the overall results are consistent. It is also reassuring to note that plans do in general contract with the majority of hospitals in their markets. Eighty-three percent of plan-hospital pairs in the data agree on contracts; this number rises to 91 percent when we consider just the 6 largest hospitals in each market. I also exclude from the analysis any plan that drops more than four of the six largest hospitals in its area.

⁴¹ The hospital inequalities consider only contracts that are accepted by plans in the observed data, i.e., that were not null offers. I therefore do not need to consider plans' responses when a hospital deviates from the null contract to a nonnull offer.

⁴² Due to computational constraints, plan premium adjustments following a network change are not modeled in this robustness test.

⁴³ The hospital profit equation models how each hospital's match with a particular plan affects its total profits including those from other plans in the market (as consumers switch plans in response to network changes). That is, I model the externalities faced by hospitals as well as those faced by plans. I also allow system hospitals to account for the effect of each contract on the profits of the system as a whole rather than those of the individual provider.

choices. They must also be positive to ensure that no inequalities are reversed by the interaction with $z_{j,m}$. I use the characteristics included in the fixed cost and markup terms (the x 's) other than the cost per admission, which I omit due to concerns about measurement error. I also include indicator variables and interactions of indicators for several market and plan characteristics. The characteristics included are: a high number of beds per population, a high proportion of the hospitals in the market being in systems, a high proportion of the population age 55–64, whether the plan is local, and whether the plan has good breast cancer screening services.⁴⁴ Each of these instruments is known to the firm when it makes its choice. Each is also correlated with x : for example, plans can more easily exclude hospitals in markets with a younger, less sick population or with more beds per population. System hospitals are more often excluded in markets with a high proportion of hospitals in systems. The logic is similar for the other instruments.

Pakes et al. (2006) provide a proof that the estimator is consistent. It also derives two potential confidence interval definitions for the identified set of parameters. I use both methodologies described there, estimating both inner and outer 95 percent confidence intervals.^{45, 46}

C. A Note on Identification

Identification in this model comes from variation in the probability of agreement across different types of hospitals and plans, both within and across markets. For example, if we observe system hospitals refusing to agree to contracts more frequently than other providers when demand and costs are similar, then the model infers that they must demand higher profits than their competitors. The producer surplus generated by these hospitals when they agree on contracts with particular plans provides an upper bound on the profits they capture. The predicted producer surplus generated when they do not reach agreement offers a lower bound on their profits. We can estimate the average profit of system hospitals by taking averages over these observations. A similar intuition applies to other types of hospital.

D. Robustness to Other Bargaining Environments

The bargaining game set out in Section VA is not renegotiation-proof: it does not allow a hospital that finds many of its offers rejected to make alternative follow-up offers. This issue affects the empirical estimates through the definition of plans' alternative options. In particular, if plan j deviates from its observed network, for example by dropping hospital h , we should ideally allow h to respond by making a "real" (nonnull) offer to a different plan with which it did not previously contract. This second plan may react by changing its own network; the adjustment process could continue for several more steps. Similarly, if plan j adds hospital h , which

⁴⁴ High proportion means more than the mean percentile, except for beds per population where quartiles of the distribution were used.

⁴⁵ The outer confidence interval is asymptotically conservative. The inner confidence interval provides asymptotically correct coverage provided that, in the limit, the number of binding moments at the boundaries of the identified set is equal to the dimension of θ .

⁴⁶ The confidence intervals have not been adjusted to account for variance introduced by the estimated demand parameters. The newness of the literature on inequalities estimators means that some issues, like methodologies to account for this variance, have not yet been worked out. The most feasible option would be to take draws from the joint distribution of the standard errors of the demand coefficients and use each draw to recompute the confidence intervals for the last stage. The ninety-fifth percentile of the distribution would provide an estimate of the full confidence intervals. However, this methodology would probably yield very conservative estimates of the full confidence intervals which may not be very informative. It would also be computationally very burdensome. This issue does not reflect general problems with inequalities estimators: consistency results exist but inference is an ongoing challenge. In my application the adjustment would be unlikely to significantly affect the results since the standard errors in the first stage were relatively low.

was previously excluded, h may respond by deciding not to make a follow-up nonnull offer to a different plan and therefore may drop that plan from its network. Another multistage adjustment process might follow.

There are some reasons to believe that accounting for these adjustments would not significantly affect the results. First, it is easiest to justify my simplifying assumption, that hospitals do not react to plan deviations, when deviations are small, as they are in my model. Second, Pakes (2008) accounts for the adjustments fully in a numerical analysis by solving for the equilibrium in a perfect information game using a hypothetical market with just two plans and two hospitals, and characteristic distributions similar to those in my data. His results are similar to those found using the estimator described above (although the cost coefficient is smaller in magnitude).

However, it is still possible that this issue biases my results. While it is computationally infeasible to model the full adjustment process in response to plan deviations, I can take one step toward a renegotiation-proof model by allowing plan j to predict that, if it deviates by dropping (adding) hospital h , the hospital could respond by adding (dropping) a different plan in the market. I form inequalities using the difference between the plan's predicted profits from the observed network and its predicted profits from the alternative scenario where hospital h chooses the response that minimizes j 's profit. As for the hospital inequalities described above, if the true hospital response generates plan profits that satisfy the inequality condition, then the condition must also hold for the minimizing response. The results, reported in Section VIB, are reassuringly similar to the results of the main specification.⁴⁷

E. An Alternative Estimator: Multinomial Logit

One obvious alternative to the methodology discussed above would be to use a multinomial logit model. This would use the plan profits given by equation (5) and assume that plans chose networks to maximize these profits. We would then assume i.i.d. extreme value errors and estimate using maximum likelihood, identifying the region of unobservable space that generates the observed choices and choosing parameters to maximize the probability that the unobservables fall within this region. As is well known, however, this methodology is problematic in this context. In particular, econometrician measurement error, which is probably important for my application, is not permitted by the logit model since it leads to a correlation between the unobservables and the observed inputs to the plan profit equation. Plan prediction error causes similar problems. In addition, the maximum likelihood framework may lead to biased results if applied to problems with multiple equilibria, where there is no one-to-one mapping between values of the unobservables and firm decisions. In contrast, the estimates from the inequalities methodology are consistent even in the presence of multiple equilibria. The approach is simply to search for parameters that are consistent with the assumption that the observed networks constitute a Bayes-Nash equilibrium without attempting to model the equilibrium selection mechanism.

I show in Section VI (results) that, as expected, the multinomial logit methodology generates inconsistent estimates in this application.

⁴⁷ This robustness test is imperfect. I do not account for the fact that, if hospital h responded to plan j 's deviation by taking plan j^* , that plan might also respond and the adjustment process could continue for several more steps. I also do not allow plans to adjust their premiums in response to network changes due to computational constraints.

F. Specifying the Hospital and Plan Profit Equations

I assume that each hospital has a two-part profit function: a fixed (per-plan) profit and a per-patient markup, $fp_{j,h}(\cdot)$ and $mk_{j,h}(\cdot)$, respectively.⁴⁸ If prices are set by bargaining, both these quantities depend on hospital and plan threat points and are therefore functions of characteristics of the hospital, the plan, and the market as a whole. I would ideally use a model of the plan-hospital bargaining process to estimate these functions directly. However, the fact that each firm's threat point is endogenous (depending on the observed or expected outcome of all other pairs' negotiations), together with the number of insurers and providers bargaining in each market, makes this approach infeasible.⁴⁹ Instead, I adopt a simpler methodology, estimating a reduced-form function for hospital profits as a function of hospital, insurer, and market characteristics. More specifically, the profit of hospital h from plan j is

$$(11) \quad \pi_{j,h,m}^H(\cdot) = fp(x_{j,h,m}) \theta_1 + N_{j,h,m}(H_j, H_{-j}) mk(x_{j,h,m}) \theta_2,$$

where $x_{j,h,m}$ are plan, hospital, and market characteristics, θ_1 and θ_2 are parameters to be estimated, and $N_{j,h,m}(H_j, H_{-j})$ is the number of plan j 's enrollees treated by hospital h .⁵⁰

$$N_{j,h,m}(H_j, H_{-j}) = \sum_i n_i p_i s_{i,j,m}(H_j, H_{-j}) s_{i,h}(H_j).$$

Substituting the sum across hospitals of the costs implied by equation (11) into equation (5), we obtain plan profits as

$$(12) \quad \begin{aligned} \pi_{j,m}^P(H_j, H_{-j}, x, \theta) &= S_{j,m}(H_j, H_{-j}) - \sum_{h \in H_j} \pi_{j,h,m}^H(H_j, H_{-j}, x_{j,h,m}, \theta) - c_{j,m}^{NONHOSP} \\ &= S_{j,m}(H_j, H_{-j}) - \sum_{h \in H_j} fp(x_{j,h,m}) \theta_1 \\ &\quad - \sum_{h \in H_j} N_{j,h,m}(H_j, H_{-j}) mk(x_{j,h,m}) \theta_2 - c_{j,m}^{NONHOSP}, \end{aligned}$$

where $S_{j,m}(H_j, H_{-j})$ accounts for plans' premium adjustments in response to changes in hospital networks as discussed in Section IIID. Hospital profits can be written analogously as

$$(13) \quad \pi_{h,m}^H(M_h, M_{-h}, x, \theta) = \sum_{j \in M_h} fp(x_{j,h,m}) \theta_1 + \sum_{j \in M_h} N_{j,h,m}(H_j, H_{-j}) mk(x_{j,h,m}) \theta_2.$$

⁴⁸ Contracts in reality fall into at least three categories. Many plans pay hospitals on a per diem or case rate basis. The former involves a daily charge plus a separate charge for major procedures such as open heart surgery; the latter implies a single rate, usually for a surgery such as open heart surgery or organ transplants, which includes a specified number of inpatient days. Capitation contracts may also be used: here a hospital receives a fixed payment in return for which it provides or covers the cost of all hospital services needed by a designated population of enrollees.

⁴⁹ The contract between a given plan-hospital pair is affected by the outcome of the negotiations between all other pairs in the market: for example, the willingness to pay of other insurers and the availability and price demands of other hospitals all depend on the outcomes of their own negotiations, and all affect the threat points of both plan and provider. Modeling this set of negotiations explicitly would be very complicated given that there are, on average, 12 plans and 15 insurers in each market.

⁵⁰ Note that both this number of patients and the surplus term also depend on the x 's: the form of the dependence is modeled explicitly using the demand estimates from Ho (2005).

The last term in equation (12) relates to plans' nonhospital costs. I would ideally estimate the cost of enrolling each type of consumer, defined by age and gender, by including the following expression in the profit equation:

$$c_{j,m}^{NONHOSP}(H_j, H_{-j}, c_i) = \sum_i N_{j,i,m}(H_j, H_{-j}) c_i,$$

where $N_{j,i,m}(H_j, H_{-j})$ is the predicted number of enrollees of type i in plan j given the equilibrium hospital networks, and c_i is the cost of insuring that type (to be estimated). Unfortunately, the available data are not rich enough to estimate c_i in addition to the hospital cost parameters. In the main specification, I set $c_i = 0$ for all i , assuming that nonhospital costs have little effect on plans' network choices. I then estimate the average nonhospital cost of enrolling a consumer by including the predicted total number of plan enrollees in the specification. The test is discussed further in Section VIA; it has little effect on the overall results.

The list of variables to include in the expression for hospital profits must be parsimonious: a large number of coefficients is unlikely to be identified given the limited data available and the fairly small variation in plan choice of networks observed. I use the intuition discussed in Section IV to inform the choice of variables:

- (i) Star hospitals: those that are highly differentiated from their competitors in the market. I identify these hospitals using an exogenous variable: an indicator for hospitals whose market share would be above the ninetieth percentile in the data under the thought experiment where all plans contract with all hospitals in the market holding prices fixed. As a robustness test, I also consider indicator variables for hospitals that provide high-tech imaging services.⁵¹
- (ii) A measure of the extent to which particular hospitals are expected to be capacity constrained. I derive an exogenous predictor of this variable by calculating which hospitals would be full under the same thought experiment, that every plan contracts with every hospital in the market.⁵²
- (iii) Hospitals in systems and those for which at least one same-system hospital is excluded.
- (iv) A measure of hospital costs per admission to test whether lower-cost providers, which all else equal generate a higher total surplus, earn higher markups than their competitors.⁵³
- (v) And also a constant term in $mk_{j,k}(\cdot)$: this identifies the average profit per patient received by nonsystem hospitals that are neither stars nor capacity constrained.

Not surprisingly, the star hospital and capacity-constraint variables are quite highly correlated: 61 of the 633 hospitals in the data are stars under this measure; 50 are predicted to be capacity

⁵¹ The imaging service considered is positron emission tomography (PET); 22 percent of hospitals in the dataset provide this service. I also used a market-specific definition of star quality: an indicator for the hospital in each market with the highest predicted market share under the thought experiment described above. The results were very similar to those reported here.

⁵² I define a hospital to be capacity constrained if the predicted number of patients exceeds the number of beds \times 365/average length of stay in the hospital.

⁵³ Joshua S. Gans, Glenn M. MacDonald, and Michael D. Ryall (2005) demonstrate that not all bargaining models generate this result. I also tried using costs per bed per night rather than costs per admission; the results were very similar to those from the main specification.

constrained (all of these are also stars). Star hospitals have significantly (at $p = 0.05$) higher-quality cardiac, imaging, cancer, and birth services than other hospitals where quality is defined, as in Table 4, as offering high-tech services. Stars also have significantly more beds per thousand population than other providers and are significantly more likely to be teaching hospitals. I found few meaningful differences between capacity-constrained and other star hospitals. The only differences that are significant at $p = 0.1$ are a smaller average number of beds per hospital and a lower average bed capacity in the market for capacity-constrained providers, both measured per thousand population in the market.

In reality, the profit received by a particular provider depends not just on its own characteristics, but also on those of the plan and the market. For example, plans that offer higher patient volumes may well receive better prices. I account for this by including the number of patients, $N_{j,h}(H_j, H_{-j})$, in the markup term. More generally, hospitals are likely to demand different prices from different plans depending on the degree to which their services complement those of the hospital (and therefore on the hospital's likely attractiveness to the plans' enrollees). Also, of course, market characteristics like the number of beds per population will probably affect the level of competition between hospitals, and therefore their profits. I include a number of plan and market characteristics in the markup term to identify these effects.⁵⁴ I report the results of several of these specifications in Section VIB below. In most cases, however, these market and plan characteristics do not have significant coefficients. It may be unrealistic to try to estimate more than the most basic effects given my limited data.⁵⁵

There is not enough information in the data to allow for free interactions with both the fixed and the per patient component of the contracts. Most of the results presented below are based on a specification where the fixed component of the contract depends on whether the hospital is in a system and whether another member of that system is excluded by the plan and the variable component includes all other variables. I also include some specifications that have only variable profits.⁵⁶ When I estimated other models where the variables were moved across the marginal and fixed components, the individual coefficients were often insignificant but there was little difference in the implications of the estimates.

VI. Results

A. Overall Results

The results are reported in Table 7. The estimate of θ for every specification was a singleton: that is, there was no parameter vector that satisfied all the inequality constraints.⁵⁷ The first column of the table reports results for the main specification, which include a constant, indicators

⁵⁴ Plan and market characteristics that are not interacted with network attributes do not vary across potential choices for a given plan, and therefore cannot be identified in the fixed-cost term of the hospital profit equation unless interacted with network attributes. It makes more sense to include these variables in the markup term, where they will be interacted with $N_{j,h}$.

⁵⁵ The star and capacity constraints variables are calculated using the predicted allocation of patients across hospitals when all plans offer a free choice: they therefore incorporate information on market characteristics. The other hospital variables, however, do not incorporate plan and market information.

⁵⁶ In particular, I estimated the model where all variables, including the system variables, were part of the markup term. The results were noisier than the main specification. The signs of the coefficients were the same but the system variables were no longer significant. I would ideally investigate this further by interacting the system variables with the number of enrollees, since this would be a more accurate model of a capitation payment than the simple fixed profit term. Unfortunately there is not enough variation in the data to identify these effects.

⁵⁷ As discussed in Pakes et al. (2006), this does not imply that we should reject the specification. The result could easily be caused by the random disturbances in the inequalities. The probability that all inequalities are satisfied can be made arbitrarily small by increasing the number of inequality restrictions.

TABLE 7—RESULTS OF FULL MODEL FOR ESTIMATION

	Star hospitals	Predicted cap con	Add number of enrollees	Multinomial logit
<i>Hospital characteristics</i>				
Number of enrollees			0.011	
Simulated 95 percent C.I.			(-0.06, 0.09)	
Conservative 95 percent C.I.			(-0.17, 0.11)	
<i>Fixed component (unit = \$ million per month)</i>				
Hospital in system	0.18	0.16	0.17	-0.60
Simulated 95 percent C.I.	(0.05, 0.37)	(0.03, 0.34)	(0.02, 0.34)	(0.09)
Conservative 95 percent C.I.	(0.18, 0.91)	(0.15, 0.75)	(0.15, 0.66)	
Drop same system hosp	0.09	0.08	0.08	0.62
Simulated 95 percent C.I.	(0.01, 0.16)	(0.01, 0.14)	(0.00, 0.15)	(0.06)
Conservative 95 percent C.I.	(0.04, 0.35)	(0.04, 0.28)	(0.06, 0.29)	
<i>Per patient component (unit = \$ thousand per patient)</i>				
Constant	5.51	14.90	6.08	-1.41
Simulated 95 percent C.I.	(-1.93, 16.6)	(1.03, 19.3)	(-2.70, 13.2)	(3.29)
Conservative 95 percent C.I.	(-24.0, 21.5)	(-22.7, 27.9)	(-12.8, 31.4)	
Star hospital	6.77		4.69	7.73
Simulated 95 percent C.I.	(3.27, 15.8)		(0.78, 16.4)	(1.41)
Conservative 95 percent C.I.	(6.40, 27.0)		(4.68, 33.3)	
Cost per admission	-0.80	-1.47	-0.74	-0.39
Simulated 95 percent C.I.	(-1.96, -0.32)	(-2.02, -0.41)	(-1.77, -0.31)	(0.19)
Conservative 95 percent C.I.	(-3.40, 0.16)	(-3.55, -0.32)	(-3.09, 0.04)	
Capacity constrained		6.92		
Simulated 95 percent C.I.		(3.64, 15.7)		
Conservative 95 percent C.I.		(5.36, 24.4)		

Notes: $N = 441$ insurance plans (the 516 in the full dataset less 9 plans in Baltimore, MD; 13 Kaiser plans; 42 with unobserved premiums; and 8 selective plans that I regard as outliers). Figures in parentheses represent 95 percent confidence intervals (C.I.). Plan premium adjustments are incorporated as described in Section IIID. Coefficients represent the predicted profits to the hospital. “Drop same system hosp” is an indicator for hospitals for which a same-system hospital has been excluded. Star hospitals have above ninetieth percentile market share when all plans contract with all hospitals; capacity-constrained hospitals are full under the same thought experiment. “Number of enrollees” is the model’s prediction of the plan’s number of enrollees given its premium and network choice. The logit results report standard errors calculated using the usual maximum likelihood estimation methodology.

for star hospitals and system hospitals, an indicator for hospitals for which a same-system hospital has been excluded by the plan, and the hospital’s cost per admission. I report the point estimate and both inner and outer confidence intervals.⁵⁸

The estimates all have the expected sign. Three of the five coefficients are significant at the 5 percent level; the cost coefficient is significant at $p = 0.08$ using the conservative confidence intervals. The confidence intervals are reasonably large, making statements about precise magnitudes difficult. The overall picture, however, is very clear. Hospitals in systems take a larger fraction of the surplus and also penalize plans that do not contract with all members. Hospitals that are very attractive to consumers (star hospitals) also capture high markups, and hospitals with higher costs per patient receive lower markups per patient than other providers.

The coefficients of the fixed component of hospital profits are measured in millions of dollars per month; those in the per patient component are in thousands of dollars per patient. To help interpret the magnitudes of the results, note that the average cost per admission for hospitals in

⁵⁸ The model is generically set identified rather than being point identified, so the confidence intervals may be more informative than the point estimates. The point estimates do, however, provide information on where the bulk of the distribution, which is not symmetrical, lies within the confidence interval.

the data is around \$11,000.⁵⁹ The estimated coefficient on costs is -0.80 , implying that hospitals on average bear 80 percent of the burden of any increase in their costs per admission.⁶⁰ For hospitals that neither are in a system nor are stars, the point estimates imply negative markups of around $-\$3,000$ per patient.⁶¹ Star hospitals receive an extra \$6,770 per patient which, when their costs are taken into account, translates into an average markup of approximately 25 percent of revenues.⁶² Hospitals that are not stars but are in systems capture \$180,000 in incremental profits per month per plan, which, given their average patient load, translates into a markup of about \$3,200 per patient. When costs per admission are also taken into account, system hospitals are predicted to have average profits of around 1.4 percent of revenues.⁶³ I estimate a penalty of \$90,000 per month, or \$1,600 per patient, for excluding a hospital from a system.

Column 2 of Table 7 replaces the star variable with the exogenous measure of capacity constraints described in Section VF. The results are qualitatively similar to those in column 1, with a positive and significant capacity constraints variable replacing that for star hospitals. Its magnitude implies that capacity-constrained providers receive an additional \$6,900 per patient compared to other hospitals. The cost coefficient is more negative here than in column 1, probably because the capacity constraint and cost variables are negatively correlated.⁶⁴ Column 3 takes account of the nonhospital costs of enrolling consumers by including the number of enrollees in the plan profit equation. The cost per enrollee is imprecisely estimated but has the expected positive sign. The magnitude implies a cost of \$11 per member per month. The other coefficients are similar to the main specification.

Finally, column 4 shows the results of the multinomial logit analysis. As expected, the estimates are quite different from those of the main analysis and make much less sense. The coefficient on system hospitals and the constant term are both negative, implying that all hospitals except stars have large negative revenues. System members do worse than other providers in spite of their higher bargaining power. The biases introduced by the logit model assumptions are clearly important in this application.

The results should be interpreted with the caveat that, since identification in this model comes solely from cross-sectional variation across markets, plans, and hospitals, I assume that there are no unobserved characteristics that are correlated with the variables of interest and that in fact cause the results.⁶⁵ For example, it is possible that stars or capacity-constrained hospitals are

⁵⁹ As noted in Section VB, the cost variable, taken from the AHA Survey 2001, is a noisy measure of true per patient costs. It is defined as total hospital expenses including items such as depreciation and interest expense.

⁶⁰ The producer surplus calculation assumes that plans cover 100 percent of each hospital's costs. The negative coefficient in the hospital markup term indicates the proportion of these costs that are in fact borne by the hospital.

⁶¹ This is \$5,500 per patient less 80 percent of \$11,000 per patient. This negative number is caused largely by the high negative cost coefficient of -0.80 . It may well be realistic. Numerous press articles describe hospitals falling into financial difficulty due to plan-hospital bargaining. For example, "Many see NYPHRM reform as cure for rising care costs" in *The Business Review* (Albany), July 22, 1996, noted that price deregulation in New York the following year would introduce a market-driven health care system and was likely to force hospitals to control their costs. "Hospital Business in New York Braces for a Crisis" in *The New York Times*, April 11, 2005, reported on the aftermath of the deregulation. Twelve New York hospitals had closed in the previous 27 months, largely as a result of the move to price bargaining and plans' refusal to pay high enough prices to cover hospitals' costs. In any case, the broad confidence interval on the constant term, together with the measurement error in the cost per admission data, prevents us from drawing any firm conclusion about the sign of the profits of these hospitals.

⁶² Nonsystem star hospital costs average \$10,786 per patient. These hospitals' profits are therefore predicted to be $\$5,500 + \$6,770 - 0.80 \times \$10,786 = \$3,641$, approximately 25 percent of revenues.

⁶³ On average, system hospitals treat 56 privately insured patients per plan per month. Nonstar system hospital costs average \$10,701 per patient. These figures, together with the estimated coefficients, imply system hospital profits of $\$5,500 + \$180,000/56 - (0.8 \times \$10,701) = \$153$ per patient, or 1.4 percent of revenues.

⁶⁴ On average, the costs of capacity-constrained hospitals in my data are \$110 per patient lower than those of other hospitals in their markets.

⁶⁵ Pakes et al. (2006) develop a methodology to control for these unobservables if valid instruments exist using particular combinations of moment inequalities. The results are qualitatively similar to those set out here, although the

excluded by some plans, not because they demand high prices but because they have unusually high costs. A similar argument could apply to system hospitals. However, the data do not support this idea. Average costs per admission are lower for star and capacity-constrained hospitals than for other hospitals in their markets. System hospitals also have lower costs on average than other providers.⁶⁶ One alternative explanation for the result for systems is that only low-quality providers form systems and that they are less likely to be offered by plans due to quality concerns. However, the services provided by system hospitals are not significantly different from those offered by other providers; system members are slightly more likely to be stars than are nonsystem hospitals.⁶⁷ The bargaining explanation, in addition to being supported by the previous literature, is much more consistent with the data.

Finally we could tell other stories that are consistent with the capacity-constraints result. For example, teaching hospitals might prefer to concentrate on research rather than treating patients; suburban hospitals might contract only with the plans that covered their particular geographical areas. If these hospitals also tended to be capacity constrained this would confound our results. However, neither teaching hospital status nor distance from the city center is significantly correlated with predicted capacity constraints. Alternatively, it is possible that capacity-constrained hospitals turn down contracts with plans, not because they wish to drive up prices or because they have high average costs, but because their costs increase when they reach full capacity. In that case, however, we should observe capacity-constrained hospitals investing in new beds at a faster rate than other providers. In fact, there is no significant difference between the investments made by the two types of hospitals in my data.⁶⁸

Of course, there may still be other valid interpretations of my results. For example, I cannot rule out the possibility that small hospitals can maintain high quality but that increasing capacity may lead to quality reductions, perhaps because there is a limited supply of good doctors in the market. This would create a barrier to capacity-constrained hospitals investing in new beds in addition to the potential disincentive created by bargaining.

B. Results of Alternative Models and Robustness Tests

Table 8 reports the results when we remove the system variables and add market and plan characteristics to the hospital markup term.⁶⁹ Column 1 adds the number of beds per 1,000 population in the market. The coefficient is negative, implying that hospital profits fall with increased competition, and significant at $p = 0.05$ according to the conservative confidence intervals. I return to this result in my investigation of the star and capacity-constraint variables below. Column 2 includes the number of the plan's enrollees admitted to the hospital, and is included to allow plans that offer high patient volumes to negotiate low prices. The negative coefficient is consistent with this idea. Its magnitude implies a price reduction of \$540 per patient for each 1,000 patients sent to the hospital per month. Again, however, the coefficient is not significant. Finally, the last two columns add an indicator for local plans (carriers other than the ten largest

cost coefficient is smaller.

⁶⁶ On average, star hospitals, capacity-constrained hospitals, and system providers have costs per admission \$30, \$110, and \$205 lower than other hospitals in their markets, respectively.

⁶⁷ The four measures of hospital services that are defined in Table 4 have values that are not significantly different in system hospitals from those in nonsystem hospitals. Eleven percent of system hospitals, compared to only 7 percent of nonsystem providers, are stars.

⁶⁸ Capacity-constrained hospitals increased their number of beds between 1997 and 2001 by 5.4 beds (3.7 percent) on average; the equivalent value for other hospitals was 2.5 beds (3.2 percent). The difference is not significant at $p = 0.1$.

⁶⁹ I report the results of the specification which includes the star variable. The results with the capacity-constraints variable are similar.

TABLE 8—ADDING PLAN AND MARKET CHARACTERISTICS

Hospital characteristics	Beds per population	Number of patients	Local plan	Plan: good screening
<i>Per patient component (unit = \$ thousand per patient)</i>				
Constant	8.48	6.26	5.60	7.33
Simulated 95 percent C.I.	(-3.95, 31.9)	(-3.64, 15.9)	(-2.43, 15.5)	(-1.18, 13.6)
Conservative 95 percent C.I.	(-8.05, 58.4)	(-6.86, 20.9)	(-5.78, 21.5)	(-8.16, 21.8)
Star hospital	8.48	6.93	6.36	8.56
Simulated 95 percent C.I.	(1.97, 16.1)	(2.40, 14.9)	(2.12, 12.6)	(2.69, 11.7)
Conservative 95 percent C.I.	(7.36, 35.2)	(5.69, 19.7)	(5.59, 18.9)	(7.67, 23.1)
Cost per admission	-0.76	-0.76	-0.74	-0.93
Simulated 95 percent C.I.	(-2.56, -0.02)	(-1.69, 0.02)	(-1.85, -0.23)	(-1.65, -0.24)
Conservative 95 percent C.I.	(-4.36, 0.52)	(-2.70, -0.16)	(-2.96, -0.24)	(-3.0, -0.27)
Market beds per population	-1.19			
Simulated 95 percent C.I.	(-5.14, 1.70)			
Conservative 95 percent C.I.	(-15.3, -1.17)			
Number of patients		-0.54		
Simulated 95 percent C.I.		(-4.66, 11.0)		
Conservative 95 percent C.I.		(-3.31, 16.2)		
Local plan			1.21	
Simulated 95 percent C.I.			(-0.37, 11.4)	
Conservative 95 percent C.I.			(-7.06, 15.2)	
Breast cancer screen				135.7
Simulated 95 percent C.I.				(-79.9, 251)
Conservative 95 percent C.I.				(-104, 439)

Notes: Results from including plan and market characteristics in per patient profits. $N = 441$ insurance plans. “Beds per population” is the number of hospital beds in the market per 1,000 population. “Number of patients” is the number of the plan’s enrollees admitted to the hospital. “Local plan” is an indicator for all plans other than the ten large national chains in the data. “Breast cancer screen” is the difference between the plan’s breast cancer screening rate and the average in the market.

national chains in the data) and a measure of the difference between the plan’s breast cancer screening rate and the average in the market. Both coefficients are positive, perhaps implying that local plans have less bargaining power than national carriers, and that plans with high preventive care quality (which are often not-for-profit plans) are less focused on profit-maximization than their peers.⁷⁰ Overall, however, these market and plan characteristics rarely have significant coefficients. Including them has little effect on the rest of the results.

Table 9 considers several alternatives to the model set out in Section VB. Column 1 lists the results of the main specification when plan premium adjustments in response to network changes are not permitted. Columns 2 and 3 allow plans and hospitals, respectively, to respond by changing networks after a deviation by the providers or plans with which they have contracts. (See Sections VB and VD for details.) Due to computational constraints I do not allow for premium adjustments in these specifications; the results should therefore be compared to those in column 1. There is very little difference between the results of these three models, implying that the simplifying assumptions made in the model may not significantly affect the results.

Next I investigate the effect of relaxing the model’s assumption that all plans and hospitals (including not-for-profit entities) are profit-maximizers. Evidence from the previous literature is divided on this point but, as Frank A. Sloan (2000) notes, the majority of studies that use data after the mid-1980s do not find a significant difference in pricing behavior between for-profit and private not-for-profit hospitals. This may be explained by increasing levels of competition in

⁷⁰ The large estimated coefficient on the breast cancer screening variable is explained by the fact that the magnitude of the variable is small. Its mean is essentially zero and its standard deviation is 0.05.

TABLE 9—ALTERNATIVE SPECIFICATIONS

Hospital characteristics	No prem adjustments	Allow plans to adjust	Allow hosps to adjust	NFP plan obj. finance	Public hosp obj. finance
<i>Fixed component (unit = \$ million per month)</i>					
Hospital in system	0.20	0.07	0.17	0.17	0.12
Simulated 95 percent C.I.	(0.05, 0.41)	(0.02, 0.32)	(0.04, 0.35)	(0.04, 0.33)	(0.04, 0.30)
Conservative 95 percent C.I.	(0.20, 1.03)	(0.07, 0.62)	(0.17, 0.75)	(0.14, 0.82)	(0.12, 0.64)
Drop same system hospital	0.09	0.04	0.08	0.08	0.07
Simulated 95 percent C.I.	(0.01, 0.18)	(0.01, 0.15)	(0.01, 0.15)	(0.01, 0.15)	(0.02, 0.13)
Conservative 95 percent C.I.	(0.04, 0.41)	(0.01, 0.26)	(0.05, 0.32)	(0.02, 0.31)	(0.04, 0.27)
<i>Per patient component (unit = \$ thousand per patient)</i>					
Constant	9.65	9.26	8.15	5.58	9.23
Simulated 95 percent C.I.	(-1.73, 16.7)	(-0.69, 14.5)	(1.41, 18.0)	(-4.17, 17.1)	(-3.62, 16.4)
Conservative 95 percent C.I.	(-20.6, 15.3)	(-14.2, 17.0)	(-14.3, 19.1)	(-22.0, 22.7)	(-21.4, 23.0)
Star hospital	2.76	2.21	2.28	6.64	4.14
Simulated 95 percent C.I.	(1.37, 13.1)	(1.17, 10.4)	(1.39, 13.3)	(0.95, 12.7)	(2.83, 11.8)
Conservative 95 percent C.I.	(2.68, 23.1)	(2.19, 15.6)	(2.19, 20.6)	(6.34, 24.3)	(4.04, 20.7)
Cost per admission	-0.99	-0.90	-0.85	-0.84	-0.85
Simulated 95 percent C.I.	(-1.94, -0.42)	(-1.65, -0.26)	(-1.87, -0.48)	(-1.7, -0.23)	(-1.56, -0.17)
Conservative 95 percent C.I.	(-2.74, -0.07)	(-2.28, 0.04)	(-2.5, -0.18)	(-3.1, -0.05)	(-2.59, 0.07)
NFP plan × enrollees				0.31	
Simulated 95 percent C.I.				(-2.24, 3.15)	
Conservative 95 percent C.I.				(-3.23, 6.13)	
Public hospital × patients					12.38
Simulated 95 percent C.I.					(-19.4, 542)
Conservative 95 percent C.I.					(-30.7, 523)

Notes: Results with alternative specifications. $N = 441$ insurance plans. Figures in parentheses represent 95 percent confidence intervals (C.I.). Column 1 gives the results when plans do not adjust premiums in response to network changes. Columns 2 and 3 allow plans and hospitals, respectively, to respond by changing networks after a deviation by their negotiating partners. Plans are not permitted to adjust premiums in these scenarios, so the results should be compared to column 1. See Sections VA and VD for details. Columns 4 and 5 allow not-for-profit plans and public hospitals to maximize a weighted sum of profits and the number of enrollees/number of patients treated, respectively.

hospital markets that have forced not-for-profit hospitals, like for-profit providers, to focus largely on revenues. The previous literature does, however, find some differences in behavior between public and private hospitals. Mark Duggan (2000) finds that for-profit and private not-for-profit hospitals responded to changes in financial incentives in California but public facilities did not, probably because “any increase in their revenues were taken by the local governments that own them.” There is less previous literature on not-for-profit plans, but some studies (such as Eric C. Schneider, Alan M. Zaslavsky, and Arnold M. Epstein 2005) find that for-profit plans may deliver lower-quality care than not-for-profit insurers. I investigate these issues by allowing not-for-profit plans and public hospitals to place a positive weight on the number of enrollees and number of patients treated, respectively, as well as on profits. The results are reported in columns 4 and 5 of Table 9. The coefficients are positive, as expected, but are not significant. The other coefficients are again very similar to the main specification.

Finally, I consider the robustness of the results for star and capacity-constrained hospitals. First, inaccuracies in the estimated demand system could cause problems with these variables. For example, if the demand for hospital h is biased up, so that the surplus increase when the hospital is added to plan j 's network is inflated, this would also imply an upward bias on the estimated hospital profit.⁷¹ I test for this by replacing the star hospitals variable with an indicator

⁷¹ This is particularly likely to be a problem if plan j is horizontally differentiated on a dimension not identified by the model. In that case, the plan's estimated average quality would be biased up (to explain its ability to exclude

TABLE 10—ROBUSTNESS TO CHOICE OF HOSPITAL CHARACTERISTICS

Hospital characteristics	Main specification	Star and cap con	Imaging services	Cap con last year
<i>Fixed component (unit = \$ million per month)</i>				
Hospital in system	0.18	0.12	0.33	0.39
Simulated 95 percent C.I.	(0.05, 0.37)	(0.04, 0.35)	(0.04, 0.54)	(0.06, 0.62)
Conservative 95 percent C.I.	(0.18, 0.91)	(0.12, 0.64)	(0.28, 1.13)	(0.31, 1.4)
Drop same system hospital	0.09	0.06	0.15	0.17
Simulated 95 percent C.I.	(0.01, 0.16)	(0.01, 0.15)	(0.03, 0.24)	(0.03, 0.25)
Conservative 95 percent C.I.	(0.04, 0.35)	(0.02, 0.24)	(0.10, 0.48)	(0.12, 0.57)
<i>Per patient component (unit = \$ thousand per patient)</i>				
Constant	5.51	9.09	3.14	1.37
Simulated 95 percent C.I.	(-1.93, 16.6)	(-5.36, 15.8)	(-9.77, 17.0)	(-13.6, 17.1)
Conservative 95 percent C.I.	(-24.0, 21.5)	(-17.4, 17.2)	(-25.6, 19.3)	(-53.4, 18.8)
Star hospital	6.77	3.48		
Simulated 95 percent C.I.	(3.27, 15.8)	(-0.72, 14.0)		
Conservative 95 percent C.I.	(6.40, 27.0)	(0.35, 27.0)		
Cost per admission	-0.80	-1.10	-0.38	-0.35
Simulated 95 percent C.I.	(-1.96, -0.32)	(-1.89, -0.28)	(-1.56, 0.27)	(-1.63, 0.70)
Conservative 95 percent C.I.	(-3.40, 0.16)	(-2.76, -0.16)	(-2.23, 0.97)	(-2.41, 2.10)
Capacity constrained		4.23		
Simulated 95 percent C.I.		(-6.21, 12.6)		
Conservative 95 percent C.I.		(-12.9, 16.7)		
Imaging services			2.85	
Simulated 95 percent C.I.			(-2.51, 7.90)	
Conservative 95 percent C.I.			(-4.30, 14.4)	
Last year cap con				1.72
Simulated 95 percent C.I.				(-3.20, 8.07)
Conservative 95 percent C.I.				(-7.09, 21.4)

Notes: Results of robustness tests. $N = 441$ insurance plans. The first test includes both star and capacity-constrained hospitals. The second replaces the star variable with an indicator for the 22 percent of hospitals that offer positron emission tomography. The third replaces predicted capacity constraints with an indicator for capacity constraints in the previous year.

for providers that offer high-tech imaging services: I use PET, which is offered by 22 percent of hospitals in my data. I replace the predicted capacity constraints measure with an indicator for hospitals that were full in the previous year.⁷² I expect less significant results here than in the main specifications. In particular, the PET coefficient may not be significant because this single characteristic may not be enough to predict the package of attributes that make the hospital attractive to patients. These attributes are not highly correlated with one another: for example, the correlation between PET provision and each of cardiac quality, cancer quality, birth quality, and teaching status is less than 0.3. The results of both specifications are reported in Table 10. They are comparable to the main model: the star and capacityconstraints coefficients are smaller than those in the main specification and not significant, but are still positive.

I then investigate whether the effects of the star and capacity-constraint variables can be distinguished from one another. I start by including both variables in the specification. The results are in Table 10. Both coefficients remain positive, but only that for star hospitals is significant.

hospitals); its increase in surplus when excluded hospitals are added would also be inflated. I exclude the plans that are most clearly horizontally differentiated from the dataset: these are Kaiser Permanente, Group Health Cooperative of Puget Sound, and Scott & White Plan of Austin, TX.

⁷² I choose not to include this variable in the main specification because it is endogenous: any serial correlation in the disturbance from the model would induce a bias in its coefficient. However, the endogeneity implies a negative bias, so a positive coefficient is still meaningful.

TABLE 11—INVESTIGATING FIT OF THE MODEL

	Contract observed	Contract not agreed upon
Number of contracts	5,587	1,160
Percent explained by $\Delta surplus$	70.2 percent	35.8 percent
Percent explained by full model	72.5 percent	54.0 percent

Notes: Summary of ability of the model to explain the data. “Percent explained by $\Delta surplus$ ” is the percent of observed contracts for which $\Delta surplus > 0$ (i.e., for which the producer surplus generated by the contract is positive) and the percent of contracts not agreed upon for which $\Delta surplus < 0$. “Percent explained by full model” has analogous meaning, with $\Delta surplus$ replaced by $\Delta \pi$, the profit change predicted by the full model.

The high correlation between the two variables makes it difficult to interpret this result. The results in column 1 of Table 8, which cover the specification where the star variable is included along with the market bed capacity per 1,000 population, are probably more informative. As noted above, hospitals that are predicted to be capacity constrained are located in lower-capacity markets on average than other star hospitals, and this is one of very few distinguishing features of capacity-constrained providers. The results in Table 8 indicate that, conditional on the star variable (which has a positive and significant effect on profits), hospitals in low-capacity markets have significantly higher markups than those elsewhere.⁷³ The available evidence therefore suggests that expected capacity constraints provide star hospitals with additional leverage in the bargaining process.⁷⁴

Though none of the robustness tests changes the qualitative nature of the results, some of the coefficients do change in magnitude. This is consistent with the results of a number of other tests not reported here, and again implies that conclusions about overall effects can be drawn from the results, but that it is difficult to make statements about precise magnitudes.

One further issue should be mentioned here: I have no data on the physician networks of the plans and therefore cannot account for them in the model. It is possible that a plan might decide not to contract with a particular hospital because this would involve establishing new physician contracts. There is no obvious reason why these physician contracting costs should be particularly high for star hospitals or for providers that expect to be capacity constrained, but this point might go some way toward explaining the result for hospital systems. If the physician networks associated with two hospital systems do not overlap, this provides an additional incentive for a plan to contract with all of one system or all of another, rather than taking some hospitals from each. In the absence of relevant data, it is difficult to say more on this issue; it may mean that the monetary costs of excluding a same-system hospital are overstated.

C. The Fit of the Model

I investigate the ability of the model to explain the data by using equation (12) to predict the change in plan profits when each hospital is added to each plan’s network, holding other plan choices fixed. I find that this expected profit variable is positive (that is, it rationalizes the

⁷³ This specification does not include the system variables. However, the correlation between market bed capacity and the system indicator is low at -0.06 . Excluding these variables, while it improves the precision of the estimates, is therefore unlikely to bias the coefficient on beds per population.

⁷⁴ When I include the number of beds in the hospital per 1,000 population in place of market-level bed capacity, the star coefficient is positive and significant and the capacity coefficient is negative but not significant at $p = 0.1$. Capacity at the market level seems to be more important for hospital profits than the number of beds in the individual hospital. Both of these variables are likely to influence the probability that the particular hospital will be full.

observed choices) for 73 percent of the observed contracts. It is negative (and therefore explains why a contract was not agreed upon) for 54 percent of the contracts that are not observed in the data. One way to interpret the fit of the model is to calculate a pseudo- R^2 measure. If we place equal weight on correctly predicting the set of observed contracts and the set that is not observed, the pseudo- R^2 is 0.60.⁷⁵ Table 11 compares these figures to those generated by the estimated producer surplus term considered alone. This variable explains 70 percent of observed contracts and just 36 percent of those that are not agreed upon. The pseudo- R^2 value is 0.29.⁷⁶

It is worth noting here that one explanation for the inability of the producer surplus term to explain the data is that the model used to estimate the surplus term is unrealistic. For example, the demand model does not fit the data perfectly; I also assume zero nonhospital costs. However, the substantial improvement in fit when we move to the full model indicates that bargaining effects are important to help us understand the determinants of the observed contracts.

The estimates are also consistent with evidence gathered in interviews. A number of interviewees noted that the dominant influence on the division of the surplus belonged to insurers in some markets and to providers in others. As the Director of Operations Analysis in one hospital chain put it: "There are counteracting effects here: the outcome [of any plan decision, like excluding a particular hospital] depends on where the balance of power lies." This makes sense in light of the estimation results. For example, hospitals are likely to dominate both in markets that include centers of excellence, such as those with top-tier medical schools, and in areas where many hospitals have merged to form systems. (In Salt Lake City, for example, two systems own six of the nine largest hospitals; we would expect hospitals to have high leverage here.) Plan power should be high in other areas.

VII. Discussion and Conclusion

The analyses in this paper investigate the determinants of hospital-insurer networks. Five factors are important: consumer demand for a particular hospital; hospital costs of care; "star" status; expected capacity constraints; and the existence of hospital systems. Together these rationalize the majority of the observed contracts and those that are not observed.

If the results of the model identify causal relationships, then they have clear implications for hospital investment incentives. The cost coefficient of -0.8 indicates that hospitals bear most of the burden of any cost increase, implying a strong incentive for them to control their costs. Providers have incentives to merge to form systems. They would also benefit by increasing their attractiveness to patients and by limiting their bed numbers in order to become capacity constrained. The second of these effects could have significant negative implications for consumer welfare.

The estimates also relate to a fairly substantial literature on the challenges faced by HMOs and POS plans in controlling costs. The original rationale for managed care was that the threat of selective contracting could be used as a lever to prevent hospitals demanding high prices. A number of recent papers have set out interview and other evidence suggesting that health

⁷⁵ The pseudo- R^2 is defined as $1 - [\sum (y_i - \hat{y}_i)^2 / \sum (y_i - \bar{y}_i)^2]$, where y_i is the observed outcome, \hat{y}_i is its predicted probability, and \bar{y}_i is the mean value in the data. Placing equal weights on the sets of observed and unobserved contracts implies the following definition for the pseudo- R^2 : $1 - [\sum (1/N_i)(y_i - \hat{y}_i)^2 / \sum (1/N_i)(y_i - \bar{y}_i)^2]$, where N_i is the number of times the i 'th alternative is observed in the data.

⁷⁶ The distribution of the expected profit change variable generated by the full model for observed contracts is quite different from that for contracts that are not agreed upon. The means are \$0.154 million and $-\$0.026$ million per month, respectively; the standard deviations are \$0.615 million and \$0.460 million. The difference in means is significant at $p = 0.01$. If we consider only the producer surplus term, the two distributions are much more similar. The means are \$0.051 million and \$0.052 million, respectively, with standard deviations of \$0.040 million and \$0.041 million. The difference in means is not significant at $p = 0.02$.

plans' leverage has declined in recent years prompting them to move away from selective contracting toward offering more choice.⁷⁷ The major causes of the reduced leverage suggested by these papers are a rising consumer demand for choice and an extensive consolidation of hospitals resulting in increased provider market power. Capacity constraints are also mentioned as a source of hospital leverage. The evidence set out in my first paper (Ho 2006) supports the first hypothesis: consumers do have a significant preference for choice. If this has developed recently, in response to experience of the restrictions imposed by managed care, it explains some of the move away from selective contracts. The results of this paper are among the first to support the other two hypotheses. Without access to data on actual prices paid, it is impossible to know whether the reduced-form function estimated here has changed as a result of plans' selective contracting, that is, whether high-priced hospitals would demand even more if no plans turned them down. However, I do show that hospitals in systems and those that are particularly attractive to patients or expect to be full are the most often excluded when the surplus they generate is positive, consistent with the theory that they have the highest leverage. Further research would be useful, particularly in a setting where price data were available, to investigate these issues in more detail.

APPENDIX: DETAILS ON HOSPITAL AND PLAN DATASETS

A. Hospital Dataset

The MEDSTAT MarketScan Research Database lists the date, the patient's age, sex, zip code, and primary diagnosis (defined using ICD-9-CM codes), the identity of the hospital, and the type of plan (managed care; indemnity; PPO, etc.) for every admission. I include six diagnosis categories in my demand analyses, defined using ICD-9-CM codes. These are: cardiac; cancer; labor; newborn baby; digestive diseases; and neurological diseases. The seventh diagnosis category, "other diagnoses," comprises all other diagnoses included in the data. I also identify emergency admissions using the place of service and the type of service for each admission. Patient income is not included in the MEDSTAT data; I approximate it using the median income of families in the Zip Code Tabulation Area (ZCTA) taken from Census 2000 data.

I supplement the MEDSTAT data with hospital characteristics data from the American Hospital Association (AHA) for 1997 and 1998. I define distance from the patient's home to the hospital using the five-digit zip codes of each; distance of the hospital from the city center is defined as distance from the hospital's five-digit zip code to that of the City Hall.

The set of hospitals operating in each market is defined by the zip codes of consumers considered to reside in the market and the distance they are likely to be willing to travel to a hospital. I consider patients whose home zip code is within the Primary Metropolitan Statistical Area (PMSA) and general medical/surgical nonfederal hospitals within 30 miles of the city center.

I consider the 11 largest markets in the MEDSTAT data (those with over 1,000 observations per market-year). The markets are located in Massachusetts, Illinois, Arizona, Washington, Florida, and Michigan; five of the eleven markets are in Michigan.⁷⁸ There are a total of 237 hospitals and 29,657 encounters⁷⁹ in these market-years. A number of observations are lost from the analysis because of missing hospital or individual data: if a variable is missing for a given individual, that

⁷⁷ See, for example, Cara S. Lesser and Paul B. Ginsburg (2001) and Lesser, Ginsburg, and Kelly Devers (2003).

⁷⁸ The markets are Boston, MA; Chicago, IL; Dayton-Springfield, OH; Orlando, FL; Phoenix, AZ; Seattle, WA; and Detroit, Grand Rapids, Kalamazoo-Battle Creek, Lansing, and Saginaw-Bay, all in MI.

⁷⁹ Each encounter is an individual admission; 51 percent of these are for PPO enrollees. If a single patient is admitted more than once in the two-year time period, I assume that the admissions represent independent choices.

individual is excluded from the analysis. If a variable (such as services provided or location) is missing for a given hospital, the missing data are filled in using surrounding years of AHA data where possible; otherwise, that hospital, and all individuals who chose it, are excluded. The final dataset comprises 217 hospitals, 434 hospital-years, and 28,666 encounters in total.

B. Plan Networks

The 516 plans for which network data were collected comprise all HMO/POS plans in 43 markets, as defined by the Atlantic Information Services data discussed in Section II.⁸⁰ The list of potential hospitals comprises all general medical/surgical hospitals listed by AHA 2001 that have more than 150 beds, are not owned by the federal government, and are located in the relevant PMSA. In smaller PMSAs, where there were fewer than ten such hospitals, facilities with over 100 beds were included.

I categorize consumers by gender, age group (0–17; 18–34; 35–44; 45–54; 55–64), and zip code tabulation area (ZCTA) of residence. There are 10 cells per ZCTA (5 age groups and 2 sexes within each), and a total of 6,363 ZCTAs across the 43 markets (an average of 148 in each market). The number of people in each ZCTA-age-sex cell is found using Census 2000 data from GeoLytics. Diagnosis probabilities given age, sex, and admission to hospital are estimated from the MEDSTAT data using probit analysis; probabilities of admission to hospital given age and sex are taken from the National Hospital Discharge Survey 2000.

AHA data for 2001 (the most recent year for which data were available) was used in the expected utility calculation. A number of hospitals have missing data for AHA 2001. To avoid dropping these from the choice set, the missing data were filled in using previous years of AHA data where possible, and if necessary (for 16 hospitals) using information provided in individual hospital Web sites.⁸¹

C. Plan Characteristics

The two datasets from Atlantic Information Services are *The HMO Enrollment Report* and *HMO Directory 2003*. Both are based on plan state insurance filings. The enrollment data give detailed enrollment for every HMO and POS plan in 40 major markets in the United States.⁸² The characteristic data cover all commercial health plans in the country.

The National Committee for Quality Assurance (NCQA) 2000 data measure clinical performance and patient satisfaction in 1999. Individuals choosing health plans in 2002 would in fact have been informed by NCQA 2001 data, to which I do not have access. I therefore chose the variables that were most highly correlated with their 1999 counterparts, assuming that their

⁸⁰ I consider all the markets covered by the AIS data except New York. The markets are: Atlanta, GA; Austin, TX; Baltimore, MD; Boston, MA; Buffalo, NY; Charlotte, NC; Chicago, IL; Cincinnati, OH; Cleveland, OH; Columbus, OH; Dallas, TX; Denver, CO; Detroit, MI; Fort Worth, TX; Houston, TX; Indianapolis, IN; Jacksonville, FL; Kansas City, MO; Las Vegas, NV; Los Angeles, CA; Miami, FL; Milwaukee, WI; Minneapolis, MN; New Orleans, LA; Norfolk, VA; Oakland, CA; Orange County, CA; Orlando, FL; Philadelphia, PA; Phoenix, AZ; Pittsburgh, PA; Portland, OR; Sacramento, CA; St. Louis, MO; Salt Lake City, UT; San Antonio, TX; San Diego, CA; San Francisco, CA; San Jose, CA; Seattle, WA; Tampa, FL; Washington, DC; and West Palm Beach, FL.

⁸¹ Data were taken from hospital Web sites in 2003. The same data were not used to fill in 1997/98 characteristics since hospitals are likely to have changed the services offered over the intervening five-year period.

⁸² AIS works with individual plans to disaggregate their base data. The data include some Medicare-only and Medicaid-only insurers, as well as commercial plans; my analysis excludes the former and examines only plans that accept commercial business. I also exclude plans with fewer than 100 enrollees and/or no hospitals in the relevant market; I assume that these plans primarily serve neighboring areas. AIS publishes the data for 40 markets. I disaggregate this to 43 markets; see later in this section for methodology.

correlation with the 2001 data would also be high. All variables used have correlations with the 1999 data of over 0.65.

The unit of observation for the NCQA data is the NCQA plan identifier, which does not correspond exactly to the identifier for the AIS enrollment data. I matched the NCQA and AIS datasets at the carrier-plan-market level; in cases where multiple NCQA plans correspond to one AIS plan, I used the mean rating over NCQA plans. Similarly, the Weiss data and AIS characteristic data do not correspond perfectly to AIS enrollment data plan identifiers: both contain more aggregated data (for example, characteristics are provided for Aetna Florida rather than Aetna Jacksonville, Aetna Miami, etc.) and often cover only HMOs. I matched the two datasets to the AIS enrollment data at the carrier-plan-market level where possible, and at the carrier-market level otherwise. I matched aggregate data to all plans within the geographic area, and if no POS data were given separately, I matched the plan's HMO characteristics to both HMO and POS plan types.

Missing data represent a significant issue. Of the 516 observations considered, 162 (31.4 percent) do not have HEDIS data and 212 (41.1 percent) do not have CAHPS data. Most of the plans without data did not respond to NCQA data requests; many did not provide information for any of the HEDIS and CAHPS categories used in this paper. There are similar problems with the price measure used: premium earned per enrolled member per month. In most cases (354 observations), this measure is calculated from AIS data on both premium and enrollment. Both inputs come from Weiss for 120 observations (where one or both pieces of information were missing in the AIS data). Price data are missing for 42 plans (8.2 percent of observations). Dropping plans with missing data (particularly NCQA data) could cause selection bias because the plans that failed to respond to NCQA requests are likely to be smaller or have lower quality than those that provided data. Instead I include these plans and add dummy variables that indicate missing premium and characteristic data.⁸³

I convert the AIS enrollment data to market shares for my demand analysis using the total nonelderly population of the MSA as defined by the Census Bureau in 2000 as the denominator.⁸⁴ The share uninsured and the share in PPO/indemnity plans in each market are also needed for the analysis. Census Bureau data are used to find the number of nonelderly uninsured; the difference between the total nonelderly population and the sum of uninsured and insured by HMO/POS plans is assumed to be indemnity/PPO coverage. One assumption is implicit in this methodology. The publicly insured, nonelderly (Medicaid) population should ideally be excluded explicitly from these groups. When this was done using Census data, some markets had very low or negative implied indemnity/PPO market shares. I therefore ignore the existence of the nonelderly publicly insured. The problem is caused by errors in the AIS enrollment data; I assume that the errors are randomly distributed across plans and markets and therefore will not bias the results.

I assume that the indemnity/PPO option is homogenous in each market. None of the data sources provides information on nonmanaged care coverage, so assumptions must be made to complete the dataset. I assume that indemnity plans are over ten years old, have premiums equal to the highest managed care premium in the relevant market, and offer a physician network size equal to the largest offered by a managed care plan in the market. NCQA performance ratings are assumed to equal the average of managed care plans in the market. (Assuming zero NCQA ratings had little effect on the results.)

⁸³ The other option would be to fill in the missing observations with previous years' data. This is not attempted because plans with missing data in 2000 often had missing data in previous years, and because cross-year correlations in reported data for a given plan are low (much lower than for the hospital data where this approach was used).

⁸⁴ I exclude people over the age of 64 from the plan demand equation in order to exclude Medicare enrollees.

Several assumptions are needed to link the datasets. In order to apply the estimated parameters from the eleven markets in the hospital dataset to the 43 markets used for the plan choice model, I assume that hospital preferences given demographics and diagnosis are fixed across markets and over time. In addition, because patient diagnosis and severity of illness influence hospital choice but are not observed in the health plan data, I assume invariance of both factors across markets and over time.

REFERENCES

- Andrews, Donald, Steven T. Berry, and Panle Jia.** 2004. "Confidence Regions for Parameters in Discrete Games with Multiple Equilibria, with an Application to Discount Chain Store Location." Unpublished.
- Asker, John.** 2006. "Diagnosing Foreclosure due to Exclusive Dealing." http://pages.stern.nyu.edu/~jasker/Paper1_ExclusiveForeclosure.pdf.
- Barro, Jason R., and David Cutler.** 1997. "Consolidation in the Medical Care Marketplace: A Case Study from Massachusetts." National Bureau of Economic Research Working Paper 5957.
- Berry, Steven, James Levinsohn, and Ariel Pakes.** 1995. "Automobile Prices in Market Equilibrium." *Econometrica*, 63(4): 841–90.
- Brooks, John M., Avi Dor, and Herbert S. Wong.** 1997. "Hospital-Insurer Bargaining: An Empirical Investigation of Appendectomy Pricing." *Journal of Health Economics*, 16(4): 417–34.
- Burns, Lawton R., and Mark V. Pauly.** 2002. "Integrated Delivery Networks: A Detour on the Road to Integrated Health Care?" *Health Affairs*, 21(4): 128–43.
- Burns, Lawton R., and Darrell P. Thorpe.** 2001. "Why Provider-sponsored Health Plans Don't Work." *Healthcare Financial Management 2001 Resource Guide*, 12–16.
- Capps, Cory S., and David Dranove.** 2004. "Hospital Consolidation and Negotiated PPO Prices." *Health Affairs*, 23(2): 175–81.
- Capps, Cory S., David Dranove, and Mark Satterthwaite.** 2003. "Competition and Market Power in Option Demand Markets." *RAND Journal of Economics*, 34(4): 737–63.
- Ciliberto, Federico, and Elie T. Tamer.** 2004. "Market Structure and Multiple Equilibria in the Airline Industry." Unpublished.
- Cutler, David M., Mark McClellan, and Joseph P. Newhouse.** 2000. "How Does Managed Care Do It?" *RAND Journal of Economics*, 31(3): 526–48.
- Duggan, Mark G.** 2000. "Hospital Ownership and Public Medical Spending." *Quarterly Journal of Economics*, 115(4): 1343–73.
- Eggleston, Karen, George Norman, and Lynne Pepall.** 2004. "Pricing Coordination Failure and Health Care Provider Integration." <http://www.bepress.com/bejeap/contributions/vol3/iss1/art20>.
- Feldman, Roger, and Douglas Wholey.** 2001. "Do HMO's Have Monopsony Power?" *International Journal of Health Care Finance and Economics*, 1(1): 7–22.
- Fox, Jeremy T.** 2007. "Estimating Matching Games with Transfers." http://home.uchicago.edu/~fox/_Media/foxmatching-2.pdf.
- Gal-Or, Esther.** 1997. "Exclusionary Equilibria in Health Care Markets." *Journal of Economics and Management Strategy*, 6(1): 5–43.
- Gal-Or, Esther.** 1999. "Mergers and Exclusionary Practices in Health Care Markets." *Journal of Economics and Management Strategy*, 8(3): 315–50.
- Gans, Joshua S., Glenn M. MacDonald, and Michael D. Ryall.** 2005. "Operationalizing the Value-Based Business Strategy." Melbourne Business School Working Paper Number 2005–04.
- Ho, Katherine.** 2005. "Essays on the Industrial Organization of Medical Care." Phd diss. Harvard University.
- Ho, Katherine.** 2006. "The Welfare Effects of Restricted Hospital Choice in the US Medical Care Market." *Journal of Applied Econometrics*, 21(7): 1039–79.
- Kaiser Family Foundation Report.** 2004. *Trends and Indicators in the Changing Health Care Marketplace, 2004 Update*. <http://www.kff.org/insurance/7031/index.cfm>.
- Kreps, David M., and Jose A. Scheinkman.** 1983. "Quantity Precommitment and Bertrand Competition Yield Cournot Outcomes." *Bell Journal*, 14(2): 326–37.
- Lesser, Cara S., and Paul B. Ginsberg.** 2001. "Back to the Future? New Cost and Access Challenges Emerge." Center for Studying Health System Change Issue Brief No. 35.

- Lesser, Cara S., Paul B. Ginsberg, and Kelly J. Devers.** 2003. "The End of an Era: What Became of the 'Managed Care Revolution' in 2001?" *Health Services Research*, 38(1, Pt. 2): 337–55.
- Luke, Roice D., Yasar Ozcan, and Peter C. Olden.** 1995. "Local Markets and System Hospital Consolidations in Metropolitan Areas." *Health Services Research*, 30(4): 555–75.
- Manning, Willard G., Joseph P. Newhouse, Naihua Duan, Emmet B. Keeler, Arleen Leibowitz, and Susan M. Marquis.** 1987. "Health Insurance and the Demand for Medical Care: Evidence from a Randomized Experiment." *American Economic Review*, 77(3): 251–77.
- Mays, Glen P., Robert E. Hurley, and Joy M. Grossman.** 2003. "An Empty Toolbox? Changes in Health Plans' Approaches for Managing Costs and Care." *Health Services Research*, 38(1, Pt. 2): 375–93.
- McFadden, Daniel L.** 1973. "Conditional Logit Analysis of Qualitative Choice Behavior." In *Frontiers in Econometrics*, ed. Paul Zarembka, 105–42. New York: Academic Press.
- Miller, Robert H., and Harold S. Luft.** 1997. "Does Managed Care Lead to Better or Worse Quality of Care?" *Health Affairs*, 16(5): 7–25.
- Mortimer, Julie H.** Forthcoming. "Vertical Contracts in the Video Rental Industry." *Review of Economic Studies*.
- Pakes, Ariel.** 2008. "Theory and Empirical Work on Imperfectly Competitive Markets." Unpublished.
- Pakes, Ariel, Jack R. Porter, Katherine Ho, and Joy Ishii.** 2006. "The Method of Moments with Inequality Constraints." http://www.economics.harvard.edu/faculty/pakes/files/Moments_November-06.pdf.
- Rey, Patrick, and Jean Tirole.** 1986a. "The Logic of Vertical Restraints." *American Economic Review*, 76(5): 921–39.
- Rey, Patrick, and Jean Tirole.** 1986b. "Vertical Restraints from a Principal-Agent Viewpoint." In *Marketing Channels: Relationships and Performance*, ed. L. Pellegrini and S. K. Reddy, 1–30. Lexington, MA: Lexington Books.
- Rey, Patrick, and Jean Tirole.** 2007. "A Primer on Foreclosure." In *Handbook of Industrial Organization Volume 3*, ed. Mark Armstrong and Robert Porter, 2145–2220. Amsterdam: Elsevier.
- Schneider, Eric C., Alan M. Zaslavsky, and Arnold M. Epstein.** 2005. "Quality of Care in For-Profit and Not-for-Profit Health Plans Enrolling Medicare Beneficiaries." *The American Journal of Medicine*, 18(12): 1392–1400.
- Seim, Katja.** 2001. "Spatial Differentiation and Market Structure: The Video Retail Industry." PhD diss. Yale University.
- Sloan, Frank A.** 2000. "Not-for-Profit Ownership and Hospital Behavior." In *Handbook of Health Economics Volume 1b*, ed. Anthony J. Culyer and Joseph P. Newhouse, 1141–74. Amsterdam: Elsevier Science, North-Holland.
- Sorensen, Morten.** 2007. "How Smart Is Smart Money? A Two-Sided Matching Model of Venture Capital." *Journal of Finance*, 62(6): 2725–62.
- Spengler, Joseph J.** 1950. "Vertical Integration and Anti-Trust Policy." *Journal of Political Economy*, 58(4): 347–52.
- Tirole, Jean.** 1995. *The Theory of Industrial Organization*. Cambridge, MA: The MIT Press.
- Town, Robert, and Gregory Vistnes.** 2001. "Hospital Competition in HMO Networks." *Journal of Health Economics*, 20(5): 733–53.
- Vistnes, Gregory.** 2000. "Hospitals, Mergers, and Two-Stage Competition." *Antitrust Law Journal*, 67(3): 671–92.
- Zwanziger, Jack, and Cathleen Mooney.** 2000. "What Factors Influence the Prices Negotiated by HMOs and Hospitals?" Unpublished.