American Economic Journal: Applied Economics 2020, 12(4): 196–223
https://doi.org/10.1257/app.20180422

Healthcare Demand under Simple Prices: Evidence from Tiered Hospital Networks†

By Elena Prager*

This paper shows that consumers respond to prices for complex healthcare when they can easily assess out-of-pocket prices. Healthcare cost containment efforts increasingly incentivize price shopping despite a dearth of evidence that this steers consumers toward lower-priced care for major medical services. I show that consumers shift toward lower-priced hospitals in the highly simplified price information environment of insurance plans with tiered hospital networks. Consumers observe a single predictable, well-defined price that applies to a broad range of services within each of at most three hospital tiers. Within three years, expected partial-equilibrium savings reach 8–17 percent of baseline spending. (*JEL* G22, H75, I11, I13)

Many healthcare services are not amenable to price shopping. Even with today's price transparency tools, information search costs for consumers remain high, and insurance rarely exposes consumers to marginal price differences for expensive care. Despite evidence of price shopping for relatively simple healthcare services such as imaging and lab tests, researchers remain skeptical that it can work for more complex care.[1]

This paper shows empirically that consumers can be incentivized to choose among alternatives based on price even for complex and expensive care. I focus on inpatient care, which consists of major medical interventions requiring an overnight stay in the hospital. In a setting where out-of-pocket prices are clearly stated, predictable, and simple to understand, I find that consumers substitute toward hospitals for which they face lower out-of-pocket prices. The estimated average elasticity of demand is in the range of -0.04 to -0.16. While fairly inelastic in

* Kellogg School of Management, Northwestern University, 2211 Campus Drive, Evanston, IL 60208 (email: elena.prager@kellogg.northwestern.edu). Ilyana Kuziemko was coeditor for this article. I am indebted to my advisor, Bob Town, and to my dissertation committee, Mark Pauly, Katja Seim, and Ashley Swanson, for their invaluable guidance and support. I received helpful suggestions from numerous seminar and conference participants. I also thank Nora Becker, Allan Collard-Wexler, Guy David, Sunita Desai, David Dranove, Craig Garthwaite, Kate Ho, Shulamite Huang, Tom Hubbard, Brian Finkelman, Cinthia Konichi, Adam Leive, Dan Miller, Chris Ody, Dan Polsky, Preethi Rao, and Amanda Starc, and three anonymous referees for valuable feedback. Financial support from the Agency for Healthcare Research and Quality (dissertation grant R36-HS024164), the Social Sciences and Humanities Research Council of Canada, the Leonard Davis Institute Pilot Grant program, and the Ackoff Doctoral Student Fellowship is gratefully acknowledged. This research does not represent the official views of any of these funders.

† Go to https://doi.org/10.1257/app.20180422 to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

[1] See, e.g., Whaley et al. (2014); Whaley, Guo, and Brown (2017); Desai et al. (2016); Desai et al. (2017).

an absolute sense,[2] this range shows that consumers are willing to make decisions based on price even for inpatient care, contrary to claims in the literature.

This price responsiveness is contingent on the simplicity of the information environment faced by consumers. My main estimation sample consists of consumers enrolled in plans with tiered provider networks, which dramatically lower the search cost for information about out-of-pocket prices relative to typical plan designs. Consumers in these tiered-network plans can perfectly forecast their out-of-pocket price for a hospital inpatient admission, irrespective of the quantity of care they end up receiving or the negotiated prices between the hospital and their insurer. In typical settings studied in the literature, the out-of-pocket price depends on the precise bundle of services the consumer ends up receiving after being admitted to the hospital, the insurer's negotiated prices for each of those services, and the consumer's prior out-of-pocket spending. These features contribute to the widely espoused view is that inpatient care "is not generally amenable to price shopping" (Desai et al. 2017), a view that is supported by a preponderance of evidence (Brot-Goldberg et al. 2017; Desai et al. 2016; Desai et al. 2017; Lieber 2017; Whaley et al. 2014; Whaley, Guo, and Brown 2017).

To check whether the price responsiveness I find can be replicated in a more complex information environment, I also estimate a demand model on a sample of patients in otherwise similar insurance plans that do not use tiered networks to set out-of-pocket prices. These plans instead use coinsurance, which calculates the out-of-pocket price as a fixed percentage of the total price of a hospital admission; the percentage is observable to the consumer, but the total price is difficult to forecast accurately. In these plans, I find no evidence of price responsiveness. These findings suggest that consumers' frequent failure to price shop for healthcare may be a result of the complexity of the information environment surrounding healthcare decisions rather than an inherent price insensitivity that is peculiar to healthcare.

To answer the main question of how consumers respond to simple out-of-pocket prices for inpatient care, I estimate a discrete choice model of demand for hospitals. I use a plausibly exogenous transition of a large plan from a traditional to a tiered network and consumer inertia in insurance plan choices to address the potential endogeneity between plan choice and out-of-pocket hospital price. My empirical strategy and identification rely on comprehensive data on the private health insurance market in Massachusetts. I combine data on healthcare utilization and health insurance enrollment from the 2009–2012 Massachusetts All-Payer Claims Database (APCD); data on insurance plan characteristics and enrollment from the Massachusetts Group Insurance Commission (GIC); and novel, hand-collected longitudinal data on Massachusetts insurers' hospital tiers. I use the longitudinal tiered-network data to cleanly identify a price coefficient in hospital demand, which is typically impeded by a lack of data on provider networks and out-of-pocket prices (Gaynor, Ho, and Town 2015).

This paper is related to a large literature on health insurance design and its relationship to healthcare demand. The paper contributes to the literature on the

[2] This range is less elastic than in the literature on the extensive-margin elasticity of demand for healthcare (Manning et al. 1987; Chandra, Gruber, and McKnight 2010; Trivedi, Moloo, and Mor 2010; Buntin et al. 2011).

elasticity of demand for medical care by estimating substitution across hospitals in response to variation in out-of-pocket prices. There is a large literature measuring demand elasticity on the extensive margin of whether to purchase any healthcare, starting with the landmark estimates from the RAND Health Insurance Experiment.[3] Recent papers in the price transparency literature estimate the price elasticity of demand across healthcare providers in response to transparency in price differences borne directly by consumers (Whaley et al. 2014, Brot-Goldberg et al. 2017, Desai et al. 2016, Lieber 2017, Desai et al. 2017). This paper extends the findings from the price transparency literature by showing that patients can be incentivized to substitute toward lower-priced options when making decisions about complex, expensive care.[4]

The finding that consumers do indeed price shop in the setting of tiered networks suggests that the demand for healthcare, even high-stakes and complex inpatient care, is not inherently inelastic. Rather, healthcare may be a good like any other, one that consumers are willing to trade off against other spending if only they can make sense of its complex pricing.

From a policy perspective, this paper contributes to the debate about mechanisms for containing rapidly rising healthcare costs (Frakt 2016). As health insurance designs that expose consumers to out-of-pocket price variations become more widespread, understanding consumer response to price differences across healthcare providers is increasingly important. The study of tiered-network plans sheds light on policies aimed at containing healthcare costs through price transparency and other demand-side mechanisms that sensitize consumers to healthcare prices. The projected savings from price transparency have been estimated to be as high as $36 billion annually (Coluni 2012).

The paper proceeds as follows. Section I describes the data and empirical setting. Section II details the empirical approach. Section III presents the results and conducts spending simulations. Finally, Section IV concludes.

I. Data and Empirical Setting

My empirical application is the private health insurance market in Massachusetts. The state's largest insurers have substantial enrollment in plans using tiered networks, which provides identifying variation in tier prices and an ample sample size. The data are compiled from multiple sources. Data on healthcare utilization and health insurance enrollment come from the 2009–2012 Massachusetts All-Payer Claims Database (APCD); longitudinal data on hospitals' placement in insurers' tiered and narrow networks were hand-collected from insurers' current and archived

[3] The landmark estimates of the elasticity of healthcare demand provided by the RAND Health Insurance Experiment are in the range of -0.1 to -0.2; more recent estimates for various classes of medical care generally fall in the same range (Manning et al. 1987; Chandra, Gruber, and McKnight 2010; Trivedi, Moloo, and Mor 2010; Buntin et al. 2011).

[4] Gowrisankaran, Nevo, and Town (2015) and Ho and Pakes (2014) study inpatient provider choice under differential pricing, but in their settings, consumers are responding to price via coinsurance or because their choices are mediated by physician referrals. There are also estimates from case studies of consumer response to price transparency initiatives, but these are difficult to generalize because they usually involve a concerted patient information campaign that is not typical in other contexts (Desai et al. 2016).

network lists; and data on insurance plans and choice sets are drawn from the employee benefit guides of the Massachusetts Group Insurance Commission (GIC).

A. *Tiered-Network Plan Designs*

My finding that consumers substitute toward lower-priced hospitals for inpatient care contrasts with the existing literature that documents limited consumer responses to price transparency (Brot-Goldberg et al. 2017; Desai et al. 2016; Desai et al. 2017; Lieber 2017; Whaley et al. 2014; Whaley, Guo, and Brown 2017). Several crucial features of my empirical setting help to explain these results. In the tiered hospital networks I study, out-of-pocket prices are remarkably clearly stated, predictable, and simple to understand. Providers are ranked based on price and placed into mutually exclusive groups, or *tiers*, that fully determine consumers' out-of-pocket payment. This structure substantially simplifies the information environment surrounding consumers' healthcare consumption decisions.

In the settings studied in the prior literature, consumers must pay a search cost in order to obtain price information for each treatment or diagnosis (Whaley et al. 2014, Brot-Goldberg et al. 2017, Desai et al. 2016, Lieber 2017). Online price search tools typically require consumers to navigate a series of menus and enter their plan information; enrollee identity; geographic location; and health condition, service, or physician type (Whaley et al. 2014, Sinaiko and Rosenthal 2016). In Lieber (2017), out-of-pocket price estimates are obtained by placing a phone call to a customer service line and providing similar information. Many healthcare conditions necessitate complicated, multipart episodes of care for which consumers must add up a vector of prices to determine a total for the treatment, such as separate fees for the surgeon, the operating room, prescription drugs, and anesthesia. Indeed, Lieber (2017) finds that the price reductions obtained from a price search tool disappear for complicated episodes of care.[5] Unforeseen complications that occur during treatment can make it impossible for consumers to determine the total price ex ante.

In the case of tiered networks, on the other hand, consumers can easily observe the out-of-pocket price associated with any hospitalization since it does not vary by diagnosis or treatment. Furthermore, that price is observed with certainty because the tiered networks in my setting use copays. A copay is an absolute dollar amount that does not vary with the total price of care paid by the insurer.[6] In plans that do not use tiered networks, out-of-pocket price is often determined by coinsurance. Coinsurance is a fixed fraction of the overall hospital price, often ten or fifteen percent. The percentage is known to the consumer, but the total price to which that percentage is applied is at best imperfectly observable for the reasons discussed above (Brot-Goldberg et al. 2017; Gowrisankaran, Nevo, and Town 2015). Online Appendix Figure A1b illustrates the ease with which a consumer can forecast her out-of-pocket price in a tiered network by simply learning the hospital's tier. Online

[5] In particular, access to the price search tool does not lead to reductions in prices paid for patients receiving more than 15 procedures in a day.

[6] In my empirical setting, consumers are exempt from inpatient copays after the fourth hospital admission in a single plan year, so their out-of-pocket price for the fifth and additional hospitalizations is zero. This is documented on the same page of the plan description that contains the copay amounts associated with each hospital tier.

Appendix Figure A1a shows the additional uncertainty that would be faced by a consumer in a coinsurance plan due to the ex ante uncertainty over negotiated prices and the specifics of the treatment.

Further reducing information search costs for consumers, insurance plans in my setting provide their enrollees with a single document that lists the tiers associated with all the hospitals in the network. Insurers in this setting are required by regulation to "clearly and conspicuously indicate" consumers' out-of-pocket prices for each tier, so consumers need not sequentially search for the out-of-pocket price of each hospital or treatment in order to comparison shop. Insurers also publish lists of hospitals' tier assignments each year, which can be easily accessed online for the current year. These lists include each hospital's tier in a single document, so consumers need not sequentially search for each hospital in order to comparison shop. A sample screenshot from the largest tiered-network plan in my data is provided in online Appendix Figure A4.

In addition, for the majority of consumers in tiered-network plans in my sample who are enrolled through a large employer group, the employer actively informs consumers of the structure of their plans. The employer group sends out newsletters, hosts health benefits fairs at multiple work sites, and circulates information directly through benefits coordinators (Sinaiko and Rosenthal 2014).[7] Prior work within the same study population finds that half of these consumers are aware of tiering within their plan a year before the start of my sample period, and one quarter of consumers learned their physician's tier prior to their first physician office visit (Sinaiko and Rosenthal 2010). By contrast, a recent nationally representative survey finds that three quarters of consumers would not know where to begin searching for price information (Mehrotra et al. 2017).

Given the difficulty of accurately forecasting out-of-pocket prices under a typical price transparency scheme, it is perhaps not surprising that the literature has found only modest consumer responses to price transparency, and no response when care is sufficiently complex. Brot-Goldberg et al. (2017) find that switching consumers to a high-deductible health plan (HDHP), in which consumers face the full marginal price of their care until they meet their deductible, does not lead to a shift toward lower-priced providers even when consumers are given a price comparison tool. They find instead an across-the-board reduction in the quantity of care consumed.

The majority of studies that find any reductions in prices paid due to price search focus on non-inpatient care that is viewed as more conducive to price shopping: imaging services, lab tests, physician office visits, and certain common outpatient procedures. Desai et al. (2016) and Desai et al. (2017) find both low take-up of two different price look-up tools and negligible overall spending reductions from price shopping. The latter study finds these negligible overall savings in spite of 14 percent price reductions for consumers who use the tool, reflecting the very low rates of search. Similarly, Whaley et al. (2014) find sizable reductions of 13–14 percent in prices paid by consumers who actively use the search tool, but low rates of search imply that expected overall savings from the tool are still below 1 percent. These

[7] Discussions with the employer group (the Massachusetts Group Insurance Commission) officials suggest that physicians give patients little additional information about out-of-pocket prices.

studies exclude inpatient care from the analysis. A notable exception is Lieber (2017), who shows that price searching can substantially reduce out-of-pocket prices in a variety of care settings, but only for relatively uncomplicated episodes of care. In this paper, I show that consumers can be incentivized to price shop across hospitals even for the complex and expensive subset of care delivered in an inpatient setting.

B. *Hospital Network Data*

Using lists of hospital tier assignments published by the insurers, I have compiled a unique dataset tracking Massachusetts hospitals' placements in several insurers' tiered and narrow networks for the period 2009–2015. Network data were hand collected from insurers' current and archived plan documentation, and cover both tiered and narrow networks.[8] The key insurers of interest are Harvard Pilgrim Health Care and Tufts Health Plan, although other insurers are also included in the analyses. These two insurers are the second- and third-largest in the state, with 20 percent and 14 percent of commercial enrollment, respectively (Massachusetts Center for Health Information and Analysis 2013).[9]

A map of Harvard Pilgrim's and Tufts' network tiers for 2012 (the most recent year for which claims data are available) is shown in online Appendix Figure A3. All Massachusetts hospitals are in-network for the insurers' flagship tiered-network plans. Online Appendix Table A3 reports the distribution of hospitals across tiers for 2012, where tier 1 denotes the insurer's most preferred tier with the lowest out-of-pocket price, and tier 3 denotes the least preferred tier. The analysis is restricted to the state's 61 general acute care hospitals, which have a total of 72 distinct campuses.[10] Hospitals belonging to the same system are not necessarily in the same tier within an insurer. The merger and acquisition activity throughout the sample period does not affect tier assignments.[11] Online Appendix Table A4 reports the distribution of hospital characteristics across tiers. Hospitals in the least preferred tier—tier 3—are disproportionately large. Academic medical centers (AMCs) are more commonly in tier 1 or tier 3 than in the middle tier. A non-negligible fraction of hospitals is found in each tier in both the Boston area and less urban parts of Massachusetts.

Tier placement in an insurer's network is determined by the price negotiated by the insurer with a given hospital. Each insurer-hospital pair negotiates a price schedule, which is a vector of prices for various treatments, and is collapsed to a base price

[8] For three of the insurers—Health New England, Neighborhood Health Plan, and UniCare—data on narrow networks were supplemented with data collected by the Group Insurance Commission (GIC), described in Section ID. I thank Cindy McGrath at the GIC for sharing these data for the early years in the sample.

[9] The largest insurer in Massachusetts is Blue Cross Blue Shield (BCBS), with 45 percent of the commercial market (Massachusetts Center for Health Information and Analysis 2013). BCBS does not participate in the GIC market and is excluded from the analyses. Its tiered hospital network is studied by Frank et al. (2015).

[10] Satellite campuses of hospitals are excluded from these summary statistics, but enter into the demand estimation as separate choice alternatives to account for the fact that their location and available services can differ from the hospital's primary campus.

[11] Almost all the acquired hospitals are low-priced hospitals that begin in the most preferred tier.

according to a formula set by the state of Massachusetts.[12] The formula converts all hospitals' prices to a comparable scale by accounting for cross-hospital variation in the complexity of diagnoses and treatments for each hospital's patient population. Insurers then rank hospitals by their base prices and determine hospitals' tiers based on those rankings.[13] Online Appendix Figure A2 shows the relationship between prices and tiers for a representative network. In principle, insurers can use quality metrics in addition to price in setting hospital tiers.[14] However, I find that in practice, including hospital quality measures does not change hospitals' tier assignments relative to a baseline of using price alone. Discussions with the provider contracting divisions of several Massachusetts insurers indicate that this is an accurate representation of their network design.

C. *Medical Claims and Hospital Price Data*

Medical claims data are drawn from the Massachusetts Center for Health Information and Analysis' (CHIA) All-Payer Claims Database (APCD) (Massachusetts Center for Health Information and Analysis 2014). The APCD consists of comprehensive data on interactions with the healthcare system of all individuals with private insurance in Massachusetts during 2009–2012.

The APCD includes detailed information on physician visits, outpatient hospital visits, inpatient hospital admissions, and prescription drugs. The data also include patient demographic information such as gender, date of birth, and five-digit zip code of residence. I match patients to zip-code-level demographic characteristics from the US Census Bureau and use the patient address information to calculate driving distance from patients to hospitals. The APCD allows me to track patients across years, and often across insurers, using longitudinal patient identifiers.

The analysis focuses on inpatient hospital admissions. Summary statistics for the sample of admissions are reported in online Appendix Table A1. The APCD is supplemented with hospital characteristics data from the American Hospital Association Annual Survey Database; hospital quality data from the Centers for Medicare and Medicaid Services Hospital Compare database; and hospital financial and casemix data from state public use files published by CHIA. Additional data preparation steps are described in the online Appendix.

The APCD reports several key price variables, including allowed amounts. Allowed amounts are actual transaction prices paid to healthcare providers, and they are critical to studying the spending effects of insurance plan design. In addition to amounts paid by insurers, the APCD separately reports patients' out-of-pocket payments for care, a key identifying variable in estimating hospital demand in

[12] This base price is a casemix-deflated average price paid to the hospital for treating the insurer's patients. The state's price adjustment formula uses 3M's All Patient Refined Diagnosis Related Groups (APR-DRGs) for the casemix adjustment (Commonwealth of Massachusetts Act of 2010, Ch. 288; Massachusetts Center for Health Information and Analysis 2015b).

[13] Some insurers make further adjustments to the assigned tiers based on hospitals' geographic isolation or negotiated prices with the hospital system's affiliated physician groups. However, such adjustments are generally minimal, affecting 0–13 percent of hospitals in an insurer's network and only two of the hospitals in online Appendix Figure A2.

[14] Commonwealth of Massachusetts Act of 2010, Ch. 288.

tiered-network plans. The healthcare utilization data from the APCD are used to estimate hospital demand in conjunction with the hospital network data described above.

D. *Insurance Plan Data*

Data on health insurance plans are drawn from the Massachusetts Group Insurance Commission (GIC) for the subset of consumers in the APCD who are insured through the GIC.[15] The GIC is the benefits administrator for the state, some municipalities, and additional public employers. It insures 300,000–350,000 people per year during my sample period, consisting of GIC-covered employees, retirees, and their dependents. The GIC was an early adopter of tiered provider networks, introducing its first tiered hospital network plan in 2003 and rolling out tiered physician networks in 2006 (Group Insurance Commission 2008, 2009). My sample of GIC enrollees observed in the APCD includes approximately 90,000 employees and 120,000 dependents. In addition to the GIC plans, some analyses include other employer plans that are observed in the APCD to use coinsurance or tiered networks with copays for setting out-of-pocket prices for inpatient hospital admissions; these are further described in Section III.

Six insurers offer a total of eleven plans through the GIC, some of which use tiered networks and some of which use narrow networks (Table 1). GIC plans use copays, which are fixed dollar amounts paid out of pocket by consumers when they use healthcare. For example, inpatient copays in the Harvard Pilgrim Independence plan start at a flat $300 per admission in fiscal year 2009, move to a tiered structure of $250, $500, and $750 copays across the three hospital tiers in 2010, and increase to $275, $500, and $1,500 in 2016. Prior to the tiering introduced in the 2010 plan year, the information structure of the plan was identical: consumers could easily observe the fixed copay that would apply to any inpatient admission at any hospital.

The key insurers of interest in this paper, Harvard Pilgrim and Tufts, each offer two plans through the GIC, one using a broad tiered hospital network and the other using a narrow version of their tiered network. The narrow-network plans were introduced in July 2010, and are studied extensively in Gruber and McKnight (2016). The broad tiered-network plans by Harvard Pilgrim and Tufts have the two highest market shares among employees insured through the GIC, with a combined share ranging from 49 to 59 percent of employee enrollees throughout the sample period. Additional information about GIC enrollees and insurance plans, including enrollee demographics, is presented in the online Appendix.

Of the seven plans offered by other insurers, only one (UniCare) uses a tiered hospital network, and this plan has less than 10 percent market share. UniCare does not contribute data to the APCD, so its enrollees are excluded from the analyses. During the sample period, Tufts' tiered plans offered through the GIC use separate hospital tiers for pediatric, obstetric, and general care. Its contemporaneous non-GIC tiered plans use standard tiering at the hospital level irrespective of

[15] I am grateful to GIC Budget Director Catherine Moore for detailed information on the institutional setting and goals of the GIC.

Table 1—Plans Available on the GIC

Plan name	Tiered	Narrow	Copays ($)
Fallon Direct		Yes	200
Fallon Select			250
Harvard Pilgrim Independence	Yes		250/500/750
Harvard Pilgrim Primary Choice	Yes	Yes	250/500/—
Health New England		Yes	250
Neighborhood Health Plan		Yes	250
Tufts Navigator	Yes		300/700/700
Tufts Spirit	Yes	Yes	300/700/—
UniCare Basic			200
UniCare Community Choice	Yes		250/500/750
UniCare PLUS			250

Notes: Hospital network structures of GIC plans for fiscal year 2011 (July 2010–June 2011). Copays are for hospital inpatient services across tiers 1/2/3, respectively.

diagnostic category. By mid-2014, Tufts discontinued tiering by diagnostic category altogether due to complaints about the complexity from providers and consumers.

The GIC plan data are used for identification of the copay coefficient in the hospital demand model. To address the potential endogeneity from selection into plans with low copays for the consumer's preferred hospital, I leverage consumers' high level of inertia in plan choices. When consumers first enroll in insurance, they are in an active-choice setting and may consider copays for their preferred hospitals when choosing a plan. However, due to inertia in plan enrollment, over time a consumer's plan characteristics increasingly approximate random assignment. I leverage this inertia by using the evolution of the hospital's copay from the first year that a household enrolled in its current plan to deal with the endogeneity of the current copay. Information about past characteristics of GIC plans, in some cases prior to the start of the APCD claims data, allows me to operationalize this empirical strategy.

II. Hospital Demand Estimation

If healthcare is different from most other goods in that healthcare demand is inherently inelastic, then consumers will not respond to tiered networks by substituting toward hospitals with lower copays. If, on the other hand, consumers are willing to price shop for healthcare but are typically stymied by the complexity and unpredictability of prices, then tiered networks will steer consumers toward lower-copay hospitals. To distinguish between these possibilities, I estimate a discrete choice model of hospital demand using approximately 30,000 inpatient hospital admissions of nonelderly, privately insured patients in Massachusetts between 2009 and 2012.

A. *Estimation*

Consumers who become sufficiently sick to require hospitalization choose a hospital at which to receive medical care. For consumer i enrolled in health insurance plan m, the set of available hospitals h and their associated out-of-pocket prices c_{mh} are determined by the plan's hospital network. Among these hospitals, the consumer chooses a hospital to maximize her utility, which depends on the consumer's characteristics, the hospital's characteristics, and the out-of-pocket price in her health plan. For consumer i enrolled in plan m who is sick with diagnosis d, utility from seeking treatment at hospital h is given by

$$u_{mhid} = -\alpha_i c_{mh} + \beta x_{hid} + \varepsilon_{mhid}, \quad (1)$$

where c_{mh} is the copay for treatment at hospital h under plan m; α_i is the consumer's out-of-pocket price sensitivity; x_{hid} is a vector of patient, illness, and hospital characteristics and their interactions, including hospital fixed effects; β is the associated coefficient vector; and ε_{mhid} is an idiosyncratic error term that is i.i.d. type 1 extreme value. The key parameter of interest is demand sensitivity to out-of-pocket price α_i. The empirical specification includes an interaction term between copay c_{mh} and the median household income in the consumer's zip code. Thus, the baseline out-of-pocket price coefficient α_i measures price sensitivity for a consumer living in a median-income zip code, while the coefficient on the interaction between copay and income allows price sensitivity to vary by income.

Patient and hospital characteristics in x_{hid} include patient demographics, diagnosis category, hospital characteristics, hospital fixed effects, past use of the hospital, and distance. Distance is an important determinant of hospital choice (Kessler and McClellan 2000; Town and Vistnes 2001; Capps, Dranove, and Satterthwaite 2003). The demand model uses driving distance from the centroid of the patient's zip code to the hospital's street address and the square of the distance.[16] A dummy variable for past use of the hospital captures established relationships between patients and healthcare providers, following Shepard (2014). Patient demographics such as age and gender are also included.

Hospital characteristics include teaching status, number of beds, an indicator for satellite campuses, and hospital quality. Quality is measured as perceived by patients using the Hospital Consumer Assessment of Healthcare Providers and Systems (HCAHPS).[17] Compared to previous work on hospital choice, these measures allow less of the preference heterogeneity to be loaded onto hospital fixed effects. Summary statistics for the sample of admissions are shown in online Appendix Table A1.

I assign each admission to a diagnostic category and severity level using the Clinical Classifications Software (CCS) from the Agency for Healthcare Research and Quality. The CCS classifies diagnoses into approximately 300 mutually exclusive groups, which are further aggregated into eighteen broader categories. The CCS diagnostic categories and their sample prevalence are described in online Appendix Table A2. The model allows hospital choice to vary according to the hospital's availability of specialized services corresponding to the patient's diagnosis by including

[16] Calculated using Bing Maps driving directions.

[17] The HCAHPS is a third-party national survey of patients that asks about their hospital experience, including responsiveness of medical staff, cleanliness, pain control, and overall rating (Centers for Medicare and Medicaid Services 2014). The HCAHPS scores capture patients' perceptions of hospital quality and are highly correlated with other hospital reputation measures such as *US News & World Report* rankings.

relevant interaction terms.[18] In particular, I include: cardiac CCS interacted with catheterization lab; obstetric CCS interacted with neonatal intensive care unit; nervous, circulatory, and musculoskeletal CCS interacted with MRI; and nervous system CCS interacted with neurological services.

This parameterization of hospital choice has several implications. The multinomial logit structure implies the independence of irrelevant alternatives (IIA) property of demand, which I mitigate by including detailed data at the consumer-hospital level, such as driving distance and interactions between diagnosis and hospital facilities. The model also treats choice of hospital as a composite measure of the patient's preferences and other factors. Hospital choice may be mediated by unobserved factors, notably referrals by the patient's physician (Kolstad and Chernew 2009, Ho and Pakes 2014). In this paper, the goal is to estimate the ultimate effect of tiered networks on market outcomes, so I treat the observed choice of hospital as the quantity of interest irrespective of the physician's influence on the decision. If consumers are inferring higher quality from higher out-of-pocket prices, this will bias my estimate of out-of-pocket price sensitivity toward the null.

Conditional on a diagnosis and a set of out-of-pocket prices, consumers choose a hospital to maximize utility as a function of all the choice variables just described. Because the error ε_{mhid} is assumed i.i.d. type 1 extreme value, the consumer's probability σ_{mhid} of choosing hospital h under plan m and diagnosis d is

$$\sigma_{mhid} = \frac{\exp\left(-\alpha_i c_{mh} + \beta x_{hid}\right)}{\sum_{h' \in H_m} \exp\left(-\alpha_i c_{mh'} + \beta x_{h'id}\right)}, \tag{2}$$

where H_m eneumerates the set of hospitals in plan m's network. This probability is used to estimate the demand model using maximum likelihood.

B. *Identification*

Identification of the hospital choice model relies on cross-sectional and longitudinal variation in hospital networks in addition to differences in hospital and patient characteristics. The model includes hospital fixed effects, so identification comes from within-hospital variation across plans, patients, and years. For example, a patient's distance to the hospital and match quality between diagnosis and hospital characteristics vary across admissions. Hospital choice sets vary across plans, with some networks including all hospitals in the state and others using narrow networks. Additional regressions using more detailed fixed effects specifications further decompose the sources of variation.

Identifying variation for the coefficient of interest on out-of-pocket price comes from three sources. First, due to differences in negotiated prices, hospitals' tiers vary across insurers. The left panel of Table 2 shows the contemporaneous variation in hospital tiers across Harvard Pilgrim's and Tufts' tiered networks. Each cell (i,j) denotes the percentage of hospitals, among those in Harvard Pilgrim's

[18] Facilities data are drawn from the American Hospital Association Annual Survey of Hospitals.

TABLE 2—VARIATION IN HOSPITAL TIERS

Across insurers				Over time within insurer			
HPHC\Tufts	Tier 1	Tier 2	Tier 3	From\To	Tier 1	Tier 2	Tier 3
Tier 1	81.0%	5.0%	14.0%	Tier 1	68.2%	25.8%	6.1%
Tier 2	67.0%	9.6%	23.4%	Tier 2	31.8%	63.6%	4.5%
Tier 3	23.1%	7.7%	69.2%	Tier 3	3.0%	24.2%	72.7%

Notes: Left panel: Fraction of hospitals in HPHC's tier (rows) that are in Tufts' tier (columns) in the same year. Right panel: Fraction of hospitals transitioning from row tier in 2010 to column tier in 2014. Satellite campuses are excluded.

row i tier, that are in Tufts' column j tier in the same year. Although some hospitals consistently occupy high or low tiers, half (49 percent) are in different tiers across the two insurers. Of those, one-fifth (10 percent of the total) are in the most preferred tier for one insurer and the least preferred tier for the other.

Hospitals also change tiers within an insurer's network over time as price contracts are renegotiated.[19] The right panel of Table 2 shows the transition matrix of hospitals' tiers over time within the same insurer. Each cell (i,j) denotes the percentage of hospitals starting in the row i tier in 2010 that have moved to the column j tier by 2014.[20] Hospitals move across tiers in both directions; this movement is typically not consistent across insurers. Depending on the tier in the baseline year, 27–36 percent of hospitals in an insurer's tiered network switch tiers by the end of the sample period. The majority of tier shifts are movements to an adjacent tier; there is little movement between tiers 1 and 3.

Finally, within a year, there is variation in out-of-pocket price arrangements across plans in the sample. For example, Harvard Pilgrim offers plans with copays for tiers 1, 2, and 3 of \$250, \$500, and \$750, respectively; it also offers plans with copays of \$300, \$300, and \$700. In both cases, the identity of hospitals in each tier is unchanged within an insurer-year, but the associated copay structure varies across plans. Among high-enrollment plans, the largest out-of-pocket price differences across tiers are in Tufts plans with copays of \$250, \$750, and \$1,500 across hospitals in tiers 1, 2, and 3. The inclusion of non-GIC tiered-network plans provides additional identifying variation in copays that helps to identify the price coefficient. The combination of cross-sectional variation in hospital tiers across insurers, variation over time within an insurer, and variation in copays across plans within an insurer-year is used to estimate hospital demand.

Hospital copays may be endogenous to hospital choice if, in addition to being a function of negotiated price as described in Section IB, tiers are a function of hospital quality or prestige. In supplementary analyses, I find no evidence that hospital quality plays a role in determining tier assignments beyond its effect on price negotiations between insurers and hospitals. After accounting for negotiated prices, neither hospital quality metrics nor consumer preferences have any remaining explanatory power for tier assignments.[21] This is unsurprising in light of the

[19] By law, tier assignments can change at most annually (Commonwealth of Massachusetts Act of 2010, Ch. 288).

[20] A handful of hospitals move out of and then back into their initial tier during the sample period.

[21] Detailed results available from the author upon request.

documented convergence of hospital quality scores over time and the fact that Massachusetts legislates a formula for mapping prices directly into networks, which serves as a focal point for insurers.[22] Even if tier assignments are not determined by hospital quality, there may be an endogeneity problem if consumers nonetheless perceive tier assignment as a signal of quality. In a national survey, Mehrotra et al. (2017) find that a large majority of consumers do not believe that prices reflect quality differences across providers. Nonetheless, the demand model includes HCAHPS quality measures and hospital fixed effects to mitigate endogeneity concerns. If consumers' inferences from copays vary systematically over time, this may still bias the demand estimates.

Another potential source of endogeneity arises from the relationship between plan enrollment decisions and hospital choices. If insurers set tiers as a function of unobservable enrollee characteristics that are also correlated with the enrollees' hospital choices, this would bias the coefficient on copays. Alternatively, if consumers are taking their preferences over hospitals into account when choosing a plan, then the copays in their chosen plans will not be exogenous (Shepard 2016). For example, a consumer who places high value on treatment at Massachusetts General Hospital (MGH) for unobservable reasons such as a strong preference for academic hospitals may enroll in plans that cover MGH at a low out-of-pocket price. The copays faced by consumers in the hospital demand stage, c_{mh}, may therefore be correlated with the error term ε_{mhid}. Such sorting would bias the estimate of price sensitivity away from the null in the direction of a more negative coefficient than the true price sensitivity.

To address the potential endogeneity from correlated plan and hospital choices, I leverage inertia in plan choices. Intuitively, the identification strategy uses consumers' past plan choices to deal with endogeneity in current plan characteristics. The identifying assumption is that, conditional on current plan copays and preferences over hospitals, consumers do not anticipate future network or copay changes. When consumers first enroll in insurance through the GIC, they are in an active-choice setting and may consider copays for their preferred hospitals when choosing a plan. In subsequent enrollment periods, although premiums and plan characteristics change, most consumers remain in the same plan without reevaluating their choice sets.

Over time, an inertial consumer's plan characteristics increasingly approximate random assignment. I use the hospital's contemporaneous copay in the plan in which a household first enrolled to deal with the endogeneity in the current copay. To the extent that non-inertial consumers reoptimize their plan enrollments over time in response to changes in copays, instrumenting with copays in their initial chosen plan helps to deal with these enrollment changes. The identifying assumption would be violated if, for example, consumers were aware that the insurer intended to raise copays in the future, prior to the publication of those future plan characteristics, and made their initial enrollment decision based on the anticipated future copays. An

[22] Commonwealth of Massachusetts Act of 2010, Ch. 288.

Table 3—Plan Enrollment Inertia on GIC, Fiscal Years 2010–2011

Plan	2010 enrollment	2011 enrollment	% Inertial
Fallon Direct	3,034	3,913	88.40
Fallon Select	8,109	10,019	91.92
Harvard Pilgrim Independence	70,131	73,486	92.61
Health New England	20,779	21,482	87.43
Neighborhood Health Plan	2,759	3,616	93.33
Tufts Navigator	82,747	85,292	93.39
Mean across plans (weighted)			92.29

Notes: Percent of GIC enrollees remaining in their plans. Two new plans were introduced in 2011 (not shown). Plan enrollments are highly inertial even following a shock to the choice set. This inertia helps to identify the hospital demand model.

analogous approach is employed by Abaluck, Gruber, and Swanson (2018) in the context of pharmaceutical coverage choice.

The use of previous plan choices to identify the effect of current copays is only justified if there is, indeed, a high degree of inertia in plan choice. Table 3 reports the fraction of consumers enrolled in each GIC plan in enrollment year 2010 who remained in the same plan in 2011.[23] Despite the introduction of two new plans in 2011, 92 percent of 2010 enrollees remain in the same plan. In 2010, Harvard Pilgrim's Independence plan switched from a standard network with flat $300 copays to a tiered hospital network for the first time, with copays of $250, $500, and $750 (Table 1). In spite of this substantial change, more than 90 percent of enrollees remained. These patterns are consistent with the literature showing that consumers fail to reoptimize their plan choices over time (Handel 2013; Ericson 2014; Shepard 2016). Combined with these findings, the observed inertia motivates the identification strategy. Given this high degree of inertia, the variation in a hospital's tiers within an insurer's network over time is, among the three key sources of variation described above, least susceptible to issues of hospital tiers being endogenous to hospital choices.

Since hospital choice is not linear in the endogenous variable (copay), the standard IV approach of substituting predicted values of the endogenous regressor into the second-stage equation would induce bias (Terza, Basu, and Rathouz 2008). Instead, I employ a control function approach, which corrects for the correlation between copays c_{mh} and the error term ε_{mhid} by approximating the component of the error that is correlated with copays and including it as a separate regressor (Petrin and Train 2010). In practice, the endogenous variable is regressed on the exogenous variables and the "instrument," and the residuals from this first-stage regression enter into the nonlinear second-stage model. This approach requires an exclusion restriction analogous to standard IV methods; namely, that the "instrument" affects hospital choice only through its effect on copay. Under this assumption, there exists some function of the first-stage residuals that produces consistent coefficient estimates (Wooldridge 2010). Because the true functional form is unknown, I allow the first-stage residuals to enter flexibly into the hospital choice model using up

[23] The GIC's enrollment periods coincide with its fiscal years, which begin on July 1 of the preceding calendar and end on June 30.

to a fifth-degree polynomial expansion.[24] The control function leveraging the high degree of plan choice inertia allows me to obtain a consistent estimate of price sensitivity in a nonlinear setting.

III. Results

Estimates from the multinomial logit hospital choice model are shown in Table 4. The sample consists of approximately 30,000 inpatient hospital admissions of nonelderly, privately insured patients in Massachusetts between 2009 and 2012. The sample includes all observed admissions of GIC enrollees in four tiered and five non-tiered GIC plans (online Appendix Table C1). I also include 4,000 admissions from Harvard Pilgrim's and Tufts' tiered plans offered outside the GIC. The non-GIC enrollees contribute additional variation in hospital tier copays. In analyses comparing consumer response to tiered networks with response to more complex coinsurance pricing, I include approximately 1,600 additional admissions from non-GIC plans that use coinsurance for setting out-of-pocket prices for inpatient care, meaning that the out-of-pocket price is a percentage of the total price for the admission. Admissions originating from the emergency department (ED) or via transfers from other hospitals are excluded from the main analysis.

A. *Baseline Results*

The first column of Table 4 presents estimates without hospital fixed effects; the second column adds fixed effects to control for time-invariant hospital characteristics not already captured by the hospital quality measures. Consistent with the hospital choice literature, the coefficient on distance is negative and significant, implying that consumers dislike travel (Kessler and McClellan 2000; Town and Vistnes 2001; Capps, Dranove, and Satterthwaite 2003; Ho 2006). Patients with cardiac or obstetric diagnoses are more likely to choose a hospital with a catheterization lab or a NICU, respectively (see online Appendix Table A5 for these and other additional coefficient estimates). Older patients and patients with chronic conditions are more willing to travel to their preferred hospital. Hospital fixed effects also display a sensible pattern. The most prestigious hospitals in the state, such as Massachusetts General Hospital and Brigham and Women's Hospital, have among the largest estimated fixed effects, driven by their large share of patients from across the state despite high out-of-pocket prices. Consistent with the literature, patients have a strong preference for hospitals with which they have established relationships, measured by past use of a given hospital (Sinaiko and Rosenthal 2014, Shepard 2016).

The primary coefficient of interest is that on out-of-pocket price, specifically copays. The negative and significant price coefficient indicates that consumers do, indeed, respond to differences in out-of-pocket price when choosing hospitals. This result lends credence to the hypothesis that, rather than being inherently insensitive

[24] Some papers have used two-stage residual inclusion (2SRI), where the residuals are entered into the second stage linearly (see, for example, Terza, Basu, and Rathouz 2008). However, the consistency result for control functions does not generally hold without a flexible specification for the residuals in the second stage (Wooldridge 2010).

Table 4—Hospital Choice Model

	No FEs (1)	Main specif. (2)	+ coinsurance (3)
Hospital choice			
Copay ($1,000s)	0.8305	−0.1953	−0.1874
	(0.0542)	(0.0700)	(0.0689)
Copay × SD income	0.1416	0.1819	0.1811
	(0.0526)	(0.0535)	(0.0539)
Distance (mi)	−0.1821	−0.1828	−0.1829
	(0.0026)	(0.0028)	(0.0027)
Distance squared	0.0005	0.0006	0.0006
	(0.0000)	(0.0000)	(0.0000)
Past use of hospital	4.9481	4.6785	4.6568
	(0.0600)	(0.0625)	(0.0612)
Age × distance	0.0001	0.0000	0.0000
	(0.0000)	(0.0000)	(0.0000)
Chronic cond. × distance	0.0222	0.0221	0.0223
	(0.0015)	(0.0016)	(0.0015)
Coinsurance ($1,000s)			0.1442
			(0.1065)
Coinsurance × SD income			−0.0043
			(0.0578)
Hospital FEs	No	Yes	Yes
Pseudo R^2	0.579	0.625	0.621
Nadmits	29,658	29,658	31,243
Nadmits_coins			1,585

Notes: Multinomial logit model of hospital choice. All price coefficients scaled to $1,000s for ease of interpretation. Consumers dislike distance and high simple out-of-pocket prices (copays). Hospital quality and income variables are standardized. Standard errors in parentheses, clustered by patient. "Nadmits" is the total number of choice sets (admissions). "Nadmits_coins" is the number of admissions using coinsurance.

to the price of healthcare, consumers are willing to price shop when prices are sufficiently clear, predictable, and simple to understand. The implied elasticities from the discrete choice model are reported and discussed below.

The magnitude of price responsiveness is decreasing in income, as indicated by the positive coefficient on the interaction of copay and median household income in the consumer's zip code of residence (measured in standard deviations). A one standard deviation increase in income from Massachusetts' 2010 average household income of $69,750 to $94,676 eliminates the copay responsiveness. These estimates suggest that the negative effect of price is moderated by high income, which is consistent with decreasing marginal returns to wealth or with liquidity constraints.[25]

[25] The point estimates are suggestive of a positive elasticity of demand for high-income consumers. However, this is merely an artifact of the linear specification of the income interaction: only 6 percent of the sample has an income 1.5 or more standard deviations above the mean, and restricting the estimation to these high-income observations yields a statistically insignificant coefficient on copays (column 2 of online Appendix Table A7).

TABLE 5—HOSPITAL CHOICE WITH ADDITIONAL FIXED EFFECTS (BOSTON PATIENTS ONLY)

	Baseline (hospital FEs) (1)	Hospital-insurer FEs (2)	Hospital-insurer-year FEs (3)	Hospital-year FEs, tier FEs (4)
Hospital choice				
Copay ($1,000s)	−0.2916	−0.2612	−0.0077	
	(0.1264)	(0.1466)	(0.1729)	
Copay × SD income	0.0638	0.0912	0.0981	
	(0.1353)	(0.1417)	(0.1455)	
Tier 1 dummy				0.5981
				(0.1390)
Tier 2 dummy				0.3286
				(0.1272)
Tier 3 dummy				0.2784
				(0.1356)
Non-Boston hospital	−38.0839	−40.4281	−43.2159	−35.0346
	(11.8783)	(11.4459)	(11.5192)	(11.5964)
Hospital FEs	Yes	No	No	No
Hospital-insurer FEs	No	Yes	No	No
Hospital-insurer-year FEs	No	No	Yes	No
Hospital-year FEs	No	No	No	Yes
Pseudo R^2	0.389	0.408	0.423	0.397
Nadmits	6,210	6,210	6,210	6,244

Notes: "Nadmits" is the number of choice sets (admissions). All specifications estimated using multinomial logit. All controls from the baseline specification (see online Appendix Table A5) are also included. Standard errors in parentheses, clustered by patient.

B. *Decomposing Sources of Variation*

As discussed in Section IIB, the baseline specification combines three sources of variation: changes in a hospital's tier within an insurer over time, variation in a hospital's tiers across insurers, and differences in copays conditional on tier across insurers' plans. Table 5 isolates each source of variation by conditioning on the appropriate sets of additional fixed effects. Due to the large number of dummies required to represent hospital-year or hospital-insurer fixed effects, these more detailed specifications are estimated on the subset of consumers who live in Boston. Hospitals outside Boston are pooled into the outside option, and the distance and copay at the closest hospital outside Boston are included as additional controls. Column 1 of Table 5 simply repeats the baseline specification from Table 4, column 2. The price coefficient among the Boston patients is larger than for the baseline sample, consistent with the starker price-distance tradeoff for these patients (see online Appendix Table A5).

The source of variation contributing most to the identification is hospitals changing tiers within an insurer over time. Column 2 of Table 5 isolates this source of variation by including hospital-insurer fixed effects so that the coefficient on copay is identified from changes within a hospital-insurer pair. The precision of the estimate deteriorates with the additional fixed effects, but the magnitude is similar to the baseline and the estimate remains significant at the 10 percent level.[26] Column 3 isolates the variation from different copays for the same hospital in the same year across different plans within the same insurer. Since a hospital's tier assignment is fixed across an insurer's plans within a given year, this is equivalent to the variation from different copays across plans within an insurer-tier-year triplet. The coefficient is substantially smaller and not significantly different from zero at any conventional confidence level. This is consistent with the fact that the majority of variation in copays within an insurer comes from differences across tiers rather than differences within a tier across plans, leaving relatively little variation in column 3. Finally, column 4 estimates the contribution of differences in a hospital's tier across insurers within a year by including hospital-year fixed effects and explicitly estimating the coefficients on tier dummies.[27] Within a year, consumers are more likely to choose a hospital if it is in a more preferred tier in their insurer's network than consumers for whom that hospital is in a less preferred tier.

C. *Consumer Sorting*

I do not find evidence of bias from consumers sorting into plans with low out-of-pocket prices for their preferred hospitals. Online Appendix Table A10 shows the results of the control function estimation described in Section IIB, using the contemporaneous copays in the household's first plan at initial enrollment to instrument for the current copays in its current plan. Estimates are shown up to a fifth-order polynomial expansion of the control function residual. If consumers were selecting into plans based on low out-of-pocket prices for their preferred hospitals in the plan's network, then failing to account for this endogeneity would bias the price sensitivity coefficient away from the null. Instead, the control function estimates suggest price sensitivity similar to the uninstrumented estimates.[28] My preferred specification for hospital demand is therefore the baseline model without the control function (Table 4, column 2).

Consumers also do not appear to differentially sort into tiered-network plans as a function of their underlying price sensitivity. I leverage the inertia of consumers who enrolled in the largest Harvard Pilgrim plan before it was tiered and subsequently, upon the plan's conversion to a tiered plan, found themselves in a tiered plan without having actively chosen one. Adding an interaction term between copay and an indicator for the plan being non-tiered at the time of the consumer's initial enrollment does not change the primary coefficient on copay, nor is the interaction term statistically significant (column 1 of online Appendix Table A7).

[26] Its *p*-value is 0.075.
[27] The omitted category is "no tier," defined as hospitals in a plan that does not use tiered networks.
[28] The control function estimates with even-degree polynomial expansions, while not significant at the 5 percent level, are similar in magnitude to the odd-degree polynomial expansions and significant at the 10 percent level.

D. *Importance of Simple Pricing*

The third column of Table 4 presents evidence that consumers' price responsiveness is contingent upon the simplicity of out-of-pocket prices in the tiered-network plans. This column includes approximately 1,600 hospital admissions by consumers in plans that use coinsurance rather than tiers and copays.[29] Like the baseline estimation sample, these consumers are insured through employer-sponsored coverage. Online Appendix Table A6 shows summary statistics comparing the coinsurance sample to the baseline tiered-network estimation sample. Consumers in the coinsurance sample are 1.5 years younger than consumers in the tiered-network sample, with no statistically significant difference in the proportion who are female. The coinsurance sample has a larger fraction of admissions for obstetric and perinatal diagnoses; the mix of remaining diagnostic categories is similar. The largest relative difference between the two samples is in out-of-pocket prices; while consumers in the tiered-network sample have a mean out-of-pocket price of $330 with a within-choice set standard deviation of $104, the coinsurance sample has nearly double the mean out-of-pocket price at $599 and nearly double the within-choice set variation at a standard deviation of $197.

Consumers in the coinsurance sample are exposed to substantially more out-of-pocket price variation than their tiered-network counterparts. Nevertheless, the price coefficient for consumers in coinsurance plans is not significantly different from zero, and the point estimate is positive. In contrast to the consumers enrolled in tiered-network plans with copays, there is no evidence that consumers in coinsurance plans price shop for inpatient care. In these plans, the out-of-pocket price is determined as a percentage, usually five or ten percent, of the total price negotiated between the insurer and the hospital. Consumers in these plans can observe their coinsurance percentages in their plan documentation. However, in order to determine their full out-of-pocket price, consumers must apply the percentage to the total price, which not only depends on the details of treatment but is also unobservable to consumers without a price search tool. The lack of detectable price shopping behavior in coinsurance plans highlights the need for simple out-of-pocket price information in order to induce price shopping.

E. *Heterogeneity in Price Responsiveness*

The baseline specification captures heterogeneity in price responsiveness by income. Other important margins of heterogeneity may shed light on the disparate estimates of price shopping in the literature. As discussed in Section IA, prior work has estimated greater price responsiveness when focusing on simple health-care services. Column 1 of online Appendix Table A8 therefore checks whether

[29] Plans using coinsurance are identified from patient out-of-pocket prices reported in the APCD. Plans are identified as using coinsurance for inpatient care if a consistent coinsurance percentage is observed across admissions and if copay amounts for all admissions are zero. Among GIC insurers, only Harvard Pilgrim has plans that can be reliably identified as using coinsurance in this manner; all of these are included in column 3 of Table 4. Out-of-pocket prices are calculated by multiplying the plan's coinsurance percentage by the insurer's negotiated price with each hospital, scaled by the AP-DRG weight of the admission.

consumers with more complex health conditions, as proxied by chronic condition status, are less responsive to copays. The estimates indicate that price responsiveness is substantially larger among consumers without chronic conditions. In fact, I cannot reject the null hypothesis of no price responsiveness among consumers with chronic conditions.

This pattern is consistent with consumers being less willing to price shop for complex care, mirroring the consensus in the literature. However, it is also consistent with consumers being less willing to switch providers when they have a longstanding relationship with a care team.[30] If this is the case, it may help to explain why services such as imaging and lab tests have been found in the literature to be particularly amenable to price shopping. If an established relationship with a provider is the mechanism driving the lower price responsiveness among consumers with chronic conditions, then a similar pattern should hold for hospital choices for delivery. Women typically have a preexisting relationship with the obstetrician who ultimately delivers their baby. Column 2 of online Appendix Table A8 therefore checks whether consumers going to the hospital for an obstetric diagnosis are less responsive to copays. Although the difference is not significant, the point estimates suggest that women choosing a hospital for delivery are more, rather than less, responsive to price than other consumers. In sum, online Appendix Table A8 is more consistent with the hypothesis that consumers are less likely to choose healthcare providers on the basis of price when the care being sought is more complex.

In addition to the complexity of care, price responsiveness may be a function of whether the hospital admission is planned in advance. Nationwide and within my sample, admissions originating in the emergency department (ED) make up approximately one-third of hospital admissions. Admissions through the ED are unscheduled, perhaps hampering consumers' ability to consider price at the time of deciding among hospital EDs. Moreover, even among admissions through the ED, there is variability in the degree of urgency: acute appendicitis must be treated immediately, whereas an inguinal hernia diagnosed in the ED can be operated on during a scheduled surgery at a later date.

Online Appendix Table A9 estimates the hospital demand model for admissions through the ED. I use the approach of Card, Dobkin, and Maestas (2009) to distinguish between ED-originating hospital admissions that are deferrable and those that are nondeferrable. Card, Dobkin, and Maestas (2009) define nondeferrable emergency admissions as those whose weekend arrival rate does not differ from the arrival rate during the week. Among diagnoses with at least fifty admissions through the ED within my sample, I define nondeferrable diagnoses as those for which the weekend share of arrivals is not statistically distinguishable from $2/7$ at the 10 percent level.[31]

For both categories of diagnoses, the estimated coefficient on copay is larger relative to the coefficient on distance than for the baseline sample, although it is substantially less precisely estimated. The point estimates suggest that consumers in

[30] The indicator for past use of a hospital in the baseline regressions does not distinguish between a consumer who has gone to a hospital once before and a consumer who has had repeated treatments at that hospital.

[31] Admissions for diagnoses occurring fewer than fifty times in the ED sample are excluded from the analysis.

the sample choose lower-priced hospitals even for ED care that leads to a hospital admission. Moreover, I find no evidence that the urgency of the condition affects price responsiveness: the point estimates are nearly identical across deferrable and nondeferrable ED admissions. If they are not merely an artifact of imprecision, these results suggest that at least some consumers may be deliberate in their choice of hospital when seeking emergency care.

F. *Changes over Time*

After the largest GIC plan switched from a standard network to a tiered network in 2010,[32] the share of volume going to hospitals in the preferred tier grew by 6 pp, while the share going to the least preferred tier fell by 10 pp.[33] This pattern in the raw data is consistent with consumers becoming more price responsive over time, perhaps due to learning.

Online Appendix Table A11 more formally checks for evidence of increasing price responsiveness over time. I leverage the long enrollment histories observed in my data to estimate a reduced-form trend of increasing price responsiveness as a function of the number of months since a consumer first enrolled in a tiered-network plan. The negative coefficient on the interaction of copay with the number of months indicates that the longer a consumer is enrolled, the more likely she is to choose hospitals with low out-of-pocket prices.[34] This suggests consumers may be learning about the structure of the tiered network as they gain more experience with it.

G. *Implied Demand Elasticities*

The hospital price elasticities implied by the demand model are summarized in Table 6, calculated as means across all admissions in the demand estimation, separately for hospitals in metropolitan Boston and outside of Boston. The two columns show own-price elasticities with respect to out-of-pocket prices at each hospital's observed mean tiered copay and at a fixed \$1,000 copay, respectively. Elasticities for individual Boston hospitals are shown in online Appendix Table A12. Own-price elasticities of demand range from -0.04 to -0.16. This range is less elastic than the RAND Health Insurance Experiment estimate of approximately -0.2 (Manning et al. 1987). For context, the maximum out-of-pocket price in the RAND experiment was \$1,000 in late 1970s dollars, which is over \$3,000 in 2010 dollars.

The RAND study measures elasticities on the extensive margin of seeking care. My results suggest that consumers also respond to price on the margin of choosing among options, conditional on seeking care in the first place. This result highlights the importance of price transparency for controlling moral hazard on the substitution margin as well as the better-studied extensive margin (Pauly 1968). These

[32] See Section IIB.
[33] A large reduction comes from Brigham and Women's Hospital and Massachusetts General Hospital. These hospitals are the flagship hospitals of the Harvard-affiliated Partners HealthCare system and are widely considered "star" hospitals (Ho 2009; Shepard 2016).
[34] Regressions with polynomial terms for time or enrollment duration show no evidence of learning slowing down. Results available from the author upon request.

Table 6—Price Elasticities from Hospital Demand Model (at Median Household Income)

Elasticities	Metro Boston	Outside Boston
Own-price (at observed copays)	−0.052 (0.002)	−0.037 (0.002)
Own-price (at $1,000 copays)	−0.156 (0.006)	−0.113 (0.006)
Cross-price (at observed copays)	0.002 (0.000)	0.001 (0.000)
Cross-price (at $1,000 copays)	0.006 (0.001)	0.001 (0.001)

Notes: Own-price and cross-price elasticities of demand for hospitals with respect to out-of-pocket price, calculated at the hospitals' observed copays and at a flat $1,000 copay, respectively. Hospital pairs with shorter distance in geographic or characteristics space have larger cross-price elasticities. Standard errors in parentheses, calculated using 100 bootstrap replications.

estimates are also less elastic than in Gowrisankaran, Nevo, and Town (2015), who find own-price elasticities of -0.10 to -0.15. The smaller magnitudes in my context may be driven by the prominent brand effects of Massachusetts hospitals, exemplified by the Harvard-affiliated Partners HealthCare system (Ho 2009; Shepard 2016).

Table 6 also reports hospitals' pairwise cross-price elasticities. They range from essentially zero to approximately 0.01. Hospital pairs that are geographically close have higher cross-price elasticities, indicating that they are good substitutes. The Boston area has a high density of hospitals (online Appendix Figure A3), allowing consumers to more easily substitute across hospitals in response to copay differences. Online Appendix Table 12 also shows cross-price elasticities for select pairs of hospitals. The key academic medical centers in Boston—Brigham and Women's Hospital, Massachusetts General Hospital, Beth Israel Deaconess Medical Center, and Boston Medical Center—are each other's closest substitutes. In addition, many hospitals, including those far from Boston, have a high cross-price elasticity with respect to the top Boston academic medical centers—Brigham and Mass General. That is, the model predicts that patients substituting away from a given hospital are likely to substitute either to its geographic competitors or to the top hospitals, irrespective of geographic proximity. This accords with intuition and with findings that these "star" hospitals are disproportionately attractive to patients (Ho 2009; Shepard 2014). Online Appendix Table A12 also reports elasticities for Cape Cod Hospital, which is geographically isolated in eastern Massachusetts and sends few patients to other hospitals; and for Baystate Medical Center and Cooley Dickinson Hospital, which are in western Massachusetts and compete with each other. These predicted substitution patterns suggest that the demand model is capturing real patterns in how patients choose hospitals.

H. *Effects on Spending*

Demand-side incentives can be an effective cost control tool if the demand response is sufficient for meaningful spending reductions. In this section, I quantify the potential savings on the table. To evaluate the effect of tiered networks on spending, I simulate inpatient hospital spending under a non-tiered network and various tiered-network designs. I examine a tiered network with copays of $250, $500, and $750 (the observed network of the highest-enrollment tiered-network plan); and the

Table 7—Hospital Sorting Counterfactuals (at Median Household Income)

	Flat copay $250	$250/500/750	$250/500/1,500
Tier 1 hospitals % of volume	27.13	27.94	28.74
Tier 2 hospitals % of volume	36.70	36.82	37.94
Tier 3 hospitals % of volume	36.17	35.24	33.32
Patient $ per admission ($)	250	518	761
Δ patient $ over flat copay	–	107.31%	204.56%
Insurer $ per admission ($)	18,584	18,226	17,835
Δ insurer $ over flat copay	–	−1.93%	−4.03%
Total $ per admission ($)	18,830	18,736	18,584
Δ total $ over flat copay	–	−0.5%	−1.31%

Notes: Demand-side effects of tiered networks, holding prices and enrollments fixed. Column 1 is the baseline scenario: a traditional hospital network with a flat copay across all hospitals. Column 2 is Harvard Pilgrim's largest tiered-network plan in 2011, with tier copays of $250, $500, and $750 across its three tiers, respectively. Column 3 uses the same tier structure but raises the tier 3 copay to $1,500.

same tiered network but with the $750 copay doubled to $1,500. The $1,500 copay is motivated by an actual increase of the tier 3 copay by the largest plan in my data after the sample period in 2015, an effort to further steer demand away from the highest-priced Partners HealthCare hospitals.

I simulate hospital shares for each patient-diagnosis pair using the hospital demand estimates from Section III, assuming all GIC consumers with an inpatient admission are enrolled in the largest tiered-network plan in the data, Harvard Pilgrim Independence. In the flat network simulation, all hospitals are assigned an identical copay of $250. These simulations hold negotiated hospital prices, hospital tiers, and non-inpatient spending fixed.[35]

Table 7 presents the results of the three copay scenarios. From left to right, the spread in out-of-pocket price across tiers rises from $0 to $1,250. Hospitals in the more preferred tiers (1 and 2) gain volume as consumers face higher out-of-pocket price spreads. Tier 3 hospitals collectively lose 7.9 percent of their baseline volume moving from the flat network to the tiered network with a $1,250 spread; tier 1 gains 5.9 percent of baseline volume. Total spending per hospital admission falls by 1.3 percent. The savings from a tiered network are small, under $300 per hospital admission on average. By comparison, the total annual premium for individual coverage in this plan is in the range of $6,000 to $8,000 over the sample period.

These savings projections use the average price sensitivity from the baseline demand estimates. If consumers become more price sensitive as they become more experienced with tiered networks, as suggested in Section IIIF, then the potential savings on the table may grow over time. Online Appendix Figure A5 traces out the potential savings as a function of consumers' experience with tiered networks.

[35] In a companion paper, I model the potential price responses explicitly by estimating a model of insurer-hospital bargaining under tiered networks (Prager 2017). In that paper, I find that the spending effects of the demand response are likely to be amplified via the mechanism of lower negotiated prices in response to larger differences in out-of-pocket prices across tiers.

After three years in the tiered-network plan with tier copays of \$250, \$500, and \$750, the average total price paid by the insurer and consumer combined falls by approximately \$1,500, or 8.4 percent of the initial spending. In the plan with the larger \$1,500 tier 3 copay, average spending falls by approximately \$3,700, or 17.4 percent of the initial spending.

Although demand-side incentives can successfully steer consumers toward lower-priced care, this comes at the expense of higher out-of-pocket spending and muted risk smoothing. The incidence of spending changes is not symmetric across consumers and the insurer. Consumers' mean out-of-pocket spending rises as copay differentials increase, even as total spending falls. Moreover, because low income is correlated with high price sensitivity, demand-side incentives may have distributional consequences by discouraging the use of high-quality but high-priced healthcare among low-income consumers. While assessing these distributional effects is beyond the scope of this paper, this issue remains important for policy.

IV. Conclusion

Reliance on market forces plays a larger role in healthcare policy in the United States than in most advanced economies. In the last decade, market-based approaches to healthcare delivery have increasingly focused on demand-side financial incentives as a mechanism for reducing healthcare spending. This paper shows that, contrary to recent evidence, consumers can be successfully incentivized to price shop for major healthcare services under certain conditions. I find that in tiered-network health insurance plans where out-of-pocket prices for healthcare are clearly stated, predictable, and simple to understand, consumers price shop across hospitals. These findings suggest that consumers' frequent failure to price shop for healthcare may be attributable to the complexity of healthcare decision-making, rather than an inherent insensitivity to healthcare prices.

That consumers of inpatient hospital care in Massachusetts are responsive to price is notable for two reasons. Inpatient care is typically required only for fairly severe conditions or serious healthcare treatments, where conventional wisdom suggests the least elastic consumption (Manning et al. 1987; Desai et al. 2017). If consumers price shop for care in the high-stakes environment of inpatient care, there is room for optimism about price shopping for less consequential healthcare services. Furthermore, the Massachusetts hospital market is characterized by strong brand effects and customer loyalty, exemplified by the Harvard-affiliated Partners HealthCare system (Ho 2009; Shepard 2014). A sizable fraction of the volume shifts in this paper is attributable precisely to lower utilization of flagship Partners HealthCare hospitals and other prestigious hospitals. On the dimensions of brand loyalty and high stakes of care, then, the sample in this paper is relatively unfavorable for finding an effect of price shopping. My estimates can be construed as a lower bound for the degree of price responsiveness that is, at least in principle, achievable in healthcare.

In other ways, my setting represents a best-case scenario for price shopping. Out-of-pocket prices in these tiered-network plans are particularly transparent and simple to understand, and the information search cost is minimal. In addition, the tiered networks in my setting provide stronger marginal incentives than most health plans. This setting represents an unusually high degree of ex ante price transparency for hospital care, even compared to the recent wave of price search tools. Along with the longer time horizon studied in this paper, these features help to explain my finding of substantially larger effects of price shopping than other recent work (Desai et al. 2016, Brot-Goldberg et al. 2017, Lieber 2017, Desai et al. 2017).

My findings have several implications for healthcare policy and optimal health insurance plan design. Consumers can learn to be more responsive to demand-side financial incentives over time and through repeated interactions with the healthcare system. It is therefore possible that some plan designs that have not been found to reduce spending through price shopping will become more effective over time as consumers adjust. This is an argument against rolling back recent insurance innovations, such as high-deductible health plans and some price transparency tools, that have not yet proved effective. However, the success of such plan designs is likely to hinge on the ease and certainty with which consumers can predict out-of-pocket prices across treatment options.

Policymakers and plan designers face a trade-off between out-of-pocket pricing schemes that are simple but blunt versus more sophisticated ones that aim to sensitize consumers to detailed price variation but may be inscrutable to consumers. High-deductible health plans, which have greatly gained in market share, fall on the sophisticated end of this spectrum. Consumers in these plans essentially pay every marginal dollar of price increases out of pocket, which preserves fine variation in prices across treatment options. However, this fine variation impedes consumers' ability to make sense of prices ex ante, especially for complex treatments with many price components. Perversely, these complex treatments are often precisely the ones with the highest overall prices.

On the simple-but-blunt end of the spectrum are plan designs like tiered networks. Consumers in these plans face only two or three distinct out-of-pocket price levels, making any raw price variation within a tier irrelevant to the consumer. In a healthcare environment characterized by uncertainty and complexity, however, this simplicity can make it possible for consumers to act on out-of-pocket price differences. Little is known empirically about the right balance between the comprehensibility and sophistication of demand-side incentives, and identifying the optimal trade-off remains an important question for policy and health insurance design.

As health insurance plan designs that encourage price shopping continue to gain market share, understanding their effects on the overall healthcare landscape will become increasingly important. The success of demand-side incentives in fostering price competition will depend not only on their passing through sufficient marginal incentives to consumers but also on consumers' ability to make sense of them.

REFERENCES

Abaluck, Jason, Jonathan Gruber, and Ashley Swanson. 2018. "Prescription Drug Use under Medicare Part D: A Linear Model of Nonlinear Budget Sets." *Journal of Public Economics* 164: 106–38.

Bebinger, Martha. 2014. "Mass. Patients Can 'Shop' for Health Care—At Least in Theory." *Kaiser Health News*. https://khn.org/news/massachusetts-price-transparency-part-of-law/ (accessed January 21, 2015).

Brot-Goldberg, Zarek C., Amitabh Chandra, Benjamin R. Handel, and Jonathan T. Kolstad. 2017. "What Does a Deductible Do? The Impact of Cost-Sharing on Health Care Prices, Quantities, and Spending Dynamics." *Quarterly Journal of Economics* 132 (3): 1261–1318.

Buntin, Melinda Beeuwkes, Amelia M. Haviland, Roland McDevitt, and Neeraj Sood. 2011. "Healthcare Spending and Preventive Care in High-Deductible and Consumer-Directed Health Plans." *American Journal of Managed Care* 17 (3): 222–30.

Capps, Cory, David Dranove, and Mark Satterthwaite. 2003. "Competition and Market Power in Option Demand Markets." *RAND Journal of Economics* 34 (4): 737–63.

Card, David, Carlos Dobkin, and Nicole Maestas. 2009. "Does Medicare Save Lives?" *Quarterly Journal of Economics* 124 (2): 597–636.

Chandra, Amitabh, Jonathan Gruber, and Robin McKnight. 2010. "Patient Cost-Sharing and Hospitalization Offsets in the Elderly." *American Economic Review* 100 (1): 193–213.

Centers for Medicare and Medicaid Services (CMS). 2014. *HCAHPS: Patients' Perspectives of Care Survey*. https://www.cms.gov/Medicare/Quality-Initiatives-Patient-Assessment-Instruments/HospitalQualityInits/HospitalHCAHPS (accessed April 2, 2015).

Coe, Erica, Jason Bello, and Jessica Lamb. 2016. *Hospital Networks: Perspective from Three Years of Exchanges*. New York: McKinsey Center for U.S. Health System Reform.

Coluni, Bobbi. 2012. *Save $36 Billion in U.S. Healthcare Spending through Price Transparency*. Toronto: Thomson Reuters.

Cutler, David M., Mark McClellan, and Joseph P. Newhouse. 2000. "How Does Managed Care Do It?" *RAND Journal of Economics* 31 (3): 526–48.

Department of Health and Human Services (HHS). 2016. "Department of Health and Human Services." *Federal Register* 81 (45): 12204–352.

Desai, Sunita, Laura A. Hatfield, Andrew L. Hicks, Michael E. Chernew, and Ateev Mehrotra. 2016. "Association between Availability of a Price Transparency Tool and Outpatient Spending." *Journal of the American Medical Association* 315 (17): 1874–81.

Desai, Sunita, Laura A. Hatfield, Andrew L. Hicks, Anna D. Sinaiko, Michael E. Chernew, David Cowling, Santosh Gautam, Sze-jung Wu, and Ateev Mehrotra. 2017. "Offering a Price Transparency Tool Did Not Reduce Overall Spending among California Public Employees and Retirees." *Health Affairs* 36 (8): 1401–07.

Division of Health Care Finance and Policy (DHCFP). 2010. *Massachusetts Health Care Cost Trends 2010 Final Report*. Boston: Commonwealth of Massachusetts. https://archives.lib.state.ma.us/bitstream/handle/2452/50099/ocn649853990-1-MainReport.pdf.

Ericson, Keith M. Marzilli. 2014. "Consumer Inertia and Firm Pricing in the Medicare Part D Prescription Drug Insurance Exchanges." *American Economic Journal: Economic Policy* 6 (1): 38–64.

Ericson, Keith M. Marzilli, and Amanda Starc. 2015. "Pricing Regulation and Imperfect Competition on the Massachusetts Health Insurance Exchange." *Review of Economics and Statistics* 97 (3): 667–82.

Frakt, Austin. 2016. "Price Transparency Is Nice. Just Don't Expect It to Cut Health Costs." *The New York Times*. https://www.nytimes.com/2016/12/19/upshot/price-transparency-is-nice-just-dont-expect-it-to-cut-health-costs.html (accessed December 19, 2016).

Frank, Matthew B., John Hsu, Mary Beth Landrum, and Michael E. Chernew. 2015. "The Impact of a Tiered Network on Hospital Choice." *Health Services Research* 50 (5): 1628–48.

Fronstin, Paul. 2003. "The Impact of a Tiered Network on Hospital Choice." Employee Benefit Research Institute Issue Brief 260.

Gaynor, Martin, Kate Ho, and Robert J. Town. 2015. "The Industrial Organization of Health-Care Markets." *Journal of Economic Literature* 53 (2): 235–84.

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

Group Insurance Commission (GIC). 2008. *Fiscal Year 2007 Annual Report*. Boston: Commonwealth of Massachusetts Group Insurance Commission.

Group Insurance Commission (GIC). 2009. *Fiscal Year 2008 Annual Report*. Boston: Commonwealth of Massachusetts Group Insurance Commission.

Gruber, Jonathan, and Robin McKnight. 2016. "Controlling Health Care Costs through Limited Network Insurance Plans: Evidence from Massachusetts State Employees." *American Economic Journal: Economic Policy* 8 (2): 219–50.

Handel, Benjamin R. 2013. "Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts." *American Economic Review* 103 (7): 2643–82.

Health Policy Commission. 2015. *2014 Cost Trends Full Report*. Boston: Massachusetts Health Policy Commission.

Ho, Katherine. 2006. "The Welfare Effects of Restricted Hospital Choice in the US Medical Care Market." *Journal of Applied Econometrics* 21 (7): 1039–79.

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

Ho, Kate, and Ariel Pakes. 2014. "Hospital Choices, Hospital Prices, and Financial Incentives to Physicians." *American Economic Review* 104 (12): 3841–84.

Kaiser Family Foundation (KFF). 2014. *Employer Health Benefits: 2014 Annual Survey*. San Francisco: Kaiser Family Foundation.

Kaiser Family Foundation (KFF). 2015. *Employer Health Benefits: 2015 Annual Survey*. San Francisco: Kaiser Family Foundation.

Kaiser Family Foundation (KFF). 2016. *Employer Health Benefits: 2016 Annual Survey*. San Francisco: Kaiser Family Foundation.

Kessler, Daniel P., and Mark B. McClellan. 2000. "Is Hospital Competition Socially Wasteful?" *Quarterly Journal of Economics* 115 (2): 577–615.

Kolstad, Jonathan T., and Michael E. Chernew. 2009. "Quality and Consumer Decision Making in the Market for Health Insurance and Health Care Services." *Medical Care Research and Review* 66 (S1): 28S–52S.

Lieber, Ethan M.J. 2017. "Does It Pay to Know Prices in Health Care?" *American Economic Journal: Economic Policy* 9 (1): 154–79.

Manning, Willard G., Joseph P. Newhouse, Naihua Duan, Emmett B. Keeler, and Arleen Leibowitz. 1987. "Health Insurance and the Demand for Medical Care: Evidence from a Randomized Experiment." *American Economic Review* 77 (3): 251–77.

Martin, Shawn. 2014. "We're Not Gonna Take It: Network Optimization Disrupts Continuity of Care." American Academy of Family Physicians. https://www.aafp.org/news/blogs/inthetrenches/entry/we_re_not_gonna_take.html (accessed March 13, 2014).

Massachusetts Center for Health Information and Analysis (CHIA). 2013. *Annual Report on the Massachusetts Health Care Market*. Boston: Massachusetts Center for Health Information and Analysis.

Massachusetts Center for Health Information and Analysis (CHIA). 2014. *All-Payer Claims Database*. Boston: Massachusetts Center for Health Information and Analysis.

Massachusetts Center for Health Information and Analysis (CHIA). 2015a. *Performance of the Massachusetts Health Care System Annual Report September 2015*. Boston: Commonwealth of Massachusetts Center for Health Information and Analysis.

Massachusetts Center for Health Information and Analysis (CHIA). 2015b. *Massachusetts Hospital Profiles: Data through Fiscal Year 2013*. Boston: Commonwealth of Massachusetts Center for Health Information and Analysis.

McKinsey and Company. 2015. *Hospital Networks: Evolution of the Configurations on the 2015 Exchanges*. New York: McKinsey Center for U.S. Health System Reform.

Mehrotra, Ateev, Katie M. Dean, Anna D. Sinaiko, and Neeraj Sood. 2017. "Americans Support Price Shopping for Health Care, but Few Actually Seek Out Price Information." *Health Affairs* 36 (8): 1392–1400.

Pauly, Mark V. 1968. "The Economics of Moral Hazard: Comment." *American Economic Review* 58 (3): 531–37.

Petrin, Amil, and Kenneth Train. 2010. "A Control Function Approach to Endogeneity in Consumer Choice Models." *Journal of Marketing Research* 47 (1): 3–13.

Prager, Elena. 2017. "Restricting the Bargaining Agenda: Evidence from Health Insurance." Working Paper.

Prager, Elena. 2020. "Replication Data for: Health Care Demand Under Simple Prices: Evidence from Tiered Hospital Networks." American Economic Association [publisher], Inter-university Consortium for Political and Social Research [distributor]. https://doi.org/10.3886/E111666V1.

Robinson, James C. 2003. "Hospital Tiers in Health Insurance: Balancing Consumer Choice with Financial Motives." *Health Affairs* 22 (1 Supplement): W3-135–46.

Shepard, Mark. 2016. "Hospital Network Competition and Adverse Selection: Evidence from the Massachusetts Health Insurance Exchange." NBER Working Paper 22600.

Sinaiko, Anna D. 2012. *Tiered Provider Networks as Strategy to Improve Health Care Quality and Efficiency*. Washington, DC: National Institute for Health Care Management (NIHCM).

Sinaiko, Anna D., and Meredith B. Rosenthal. 2010. "Consumer Experience with a Tiered Physician Network: Early Evidence." *American Journal of Managed Care* 16 (2): 123–30.

Sinaiko, Anna D., and Meredith B. Rosenthal. 2014. "The Impact of Tiered Physician Networks on Patient Choices." *Health Services Research* 49 (4): 1348–63

Sinaiko, Anna D., and Meredith B. Rosenthal. 2016. "Examining a Health Care Price Transparency Tool: Who Uses It, and How They Shop for Care." *Health Affairs* 35 (4): 662–70.

Song, Zirui, Dana Gelb Safran, Bruce E. Landon, Mary Beth Landrum, Yulei He, Robert E. Mechanic, Matthew P. Day, and Micahel E. Chernew. 2012. "The 'Alternative Quality Contract,' Based on a Global Budget, Lowered Medical Spending and Improved Quality." *Health Affairs* 31 (8): 1885–94.

Starc, Amanda. 2014. "Insurer Pricing and Consumer Welfare: Evidence from Medigap." *RAND Journal of Economics* 45 (1): 198–220.

Terza, Joseph V., Anirban Basu, and Paul J. Rathouz. 2008. "Two-Stage Residual Inclusion Estimation: Addressing Endogeneity in Health Econometric Modeling." *Journal of Health Economics* 27 (3): 531–43.

Town, Robert, and Gregory Vistnes. 2001. "Hospital Competition in HMO Networks." *Journal of Health Economics* 20 (5): 733–53.

Trivedi, Amal N., Husein Moloo, and Vincent Mor. 2010. "Increased Ambulatory Care Copayments and Hospitalizations among the Elderly." *New England Journal of Medicine* 362 (4): 320–28.

Whaley, Christopher, Jennifer Schneider Chafen, Sophie Pinkard, Gabriella Kellerman, Dena Bravata, Robert Kocher, and Neeraj Sood. 2014. "Association between Availability of Health Service Prices and Payments for These Services." *Journal of the American Medical Association* 312 (16): 1670–76.

Whaley, Christopher M., Chaoran Guo, and Timothy T. Brown. 2017. "The Moral Hazard Effects of Consumer Responses to Targeted Cost-Sharing." *Journal of Health Economics* 56: 201–21.

Wooldridge, Jeffrey M. 2010. *Econometric Analysis of Cross Section and Panel Data*. Cambridge: MIT Press.