Consumer Responsiveness to Simple Health Care Prices: Evidence From Tiered Hospital Networks*

Elena Prager†

September 21, 2017

Abstract

This paper shows that consumers price-shop for health care when they can easily assess out-of-pocket prices. Health care cost containment efforts increasingly incentivize price-shopping, despite recent evidence that this does not steer consumers toward lower-priced care. I show that consumers price-shop in the highly simplified price information environment of health insurance plans with tiered hospital networks. These consumers observe a single predictable, well-defined price that applies to a broad range of services within each of at most three tiers of hospitals. The savings from price-shopping are large enough to both compensate for consumer welfare losses and raise insurer profits.

JEL codes: I11, I13, D83

*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. 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.

†Kellogg School of Management, Northwestern University, 2211 Campus Drive, Evanston, IL 60208. Email: elena.prager@kellogg.northwestern.edu

1 Introduction

In an effort to reduce health care spending, health insurers and policy-makers are increasingly incentivizing consumers to shop for health care based on price. However, recent empirical work has largely concluded that consumers do not choose lower-priced health care alternatives in response to such incentives (Brot-Goldberg et al. 2015; Desai et al. 2016; Lieber 2017; Desai et al. 2017). In spite of this early evidence, insurance designs with demand-side incentives continue to proliferate in an effort to control health care spending.

This paper shows empirically that consumers can indeed be incentivized to price-shop for health care. 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 -0.03 to -0.12.[1] While fairly inelastic in an absolute sense, this range shows that consumers are willing to price-shop even for the high-stakes subset of health care received in hospitals. I also find heterogeneity in price responsiveness by income, with higher-income consumers exhibiting smaller price elasticities. These findings suggest that consumers' frequent failure to price-shop for health care may be a result of the complexity of the information environment surrounding health care decisions, rather than an inherent insensitivity to price that is peculiar to health care.

I provide suggestive evidence that consumers become more price-responsive over time, both in response to third-party information and after personally consuming health care with a positive out-of-pocket price. In counterfactual analyses, moving an experienced consumer from a health insurance plan with no out-of-pocket price differentiation to a plan with a \$500 spread between the lowest and highest out-of-pocket prices reduces hospital spending by an average of 8% due to demand steering alone. Increasing the spread in out-of-pocket prices to \$1,250 results in a 15% reduction in hospital spending relative to the baseline of equal out-of-pocket prices. These savings are more than sufficient to compensate consumers for their

[1]This range is less elastic than in the literature on the extensive-margin elasticity of demand for health care (Manning et al. 1987; Chandra et al. 2010; Trivedi et al. 2010; Buntin et al. 2011).

expected welfare losses due to higher out-of-pocket prices and steering away from their preferred hospitals. Taken together, these results provide evidence that well-designed demand-side incentives can reduce health care spending. Moreover, the spending reductions can come from substitution toward lower-priced treatment options rather than outright reductions in the quantity of care consumed.

My finding that consumers can be incentivized to shift toward lower-priced care contrasts with much of the recent work on price-shopping in health care. Brot-Goldberg et al. (2015) 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 price-shopping even when consumers are given a price comparison tool. They find instead an across-the-board reduction in the quantity of care consumed. Desai et al. (2016) and Desai et al. (2017) find both low take-up of a price look-up tool and negligible overall spending reductions from price shopping. A notable recent exception is Lieber (2017), who finds that although price searching substantially reduces out-of-pocket prices, even in this setting only a small minority of consumers ever use the price search tool.

Several crucial features of my empirical setting help to explain these results. I study consumer responses to out-of-pocket prices in insurance plans with tiered hospital networks, which rank providers based on price and place them into mutually exclusive groups, or *tiers*, that determine consumers' out-of-pocket payment for a particular provider. In my setting, a hospital's tier in the insurance plan's network fully determines the out-of-pocket price to the consumer. The structure of tiered networks substantially simplifies the information environment surrounding consumers' health care consumption decisions. My finding that consumers price-shop in this simplified environment is consistent with emerging evidence about health insurance demand, such as an experiment in which showing consumers a highly simplified version of the financial characteristics of health insurance plans nearly eliminates inconsistencies in plan choices (Samek and Sydnor 2016).

In the settings studied in prior literature, consumers must pay a search cost in order to price-shop for each treatment or diagnosis (Brot-Goldberg et al. 2015; Desai et al. 2016; Lieber 2017). Many health care conditions necessitate complicated, multi-part 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 fee, prescription drugs, and anesthesia. Indeed, Lieber (2017) finds that the price reduction from a price look-up tool disappears for complicated episodes of care.[2] 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 defined as absolute dollar values. In other settings, out-of-pocket price is often determined by coinsurance, which is calculated as a percentage of the overall hospital price, a price that is at best imperfectly observable for the reasons discussed above (Brot-Goldberg et al. 2015; Gowrisankaran et al. 2015). These features of tiered networks provide an unusually high degree of ex ante price transparency for hospital care.

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. By contrast, recent survey evidence suggests that three-quarters of consumers would not know where to search for price information even if they tried (Mehrotra et al. 2017).

The final key feature of tiered networks is that they provide stronger marginal incentives than more typical health plans, including the high-deductible plans studied by Brot-Goldberg et al. (2015). The retention of marginal incentives for high-priced care such as inpatient hospitalizations is disproportionately important for overall cost control. In a high-deductible health plan, even a single hospital admission typically causes consumers to exceed the deductible, nullifying the marginal incentive to choose lower-priced care (Desai et al. 2016). In my setting, consumers must pay the out-of-pocket price for the first four hospitalizations in a coverage year, so that

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

marginal incentives are retained for the majority of hospital admissions.[3]

In addition to differences in plan design between tiered networks and the plans studied in the prior literature, my setting follows consumers for a longer time period. Existing papers have had access to at most two years of data after the introduction of the price transparency tools they study. Brot-Goldberg et al. (2015) study two years of enrollment in new high-deductible health plans, and document substantial overall quantity reductions but no effect of consumer substitution toward lower-priced care. Desai et al. (2016) and Lieber (2017) each study consumers for one year after the introduction of a price transparency tool, and find no spending reductions and moderate out-of-pocket price paid reductions, respectively. I observe four years of hospital choices by consumers who have been enrolled in tiered-network plans for as long as six years. Coupled with the relative simplicity of out-of-pocket prices in tiered networks, this longer time period affords consumers the opportunity to learn about the structure of their plan and to begin to price-shop.

Plans with tiered provider networks were introduced in the early 2000s, as insurers sought to bolster their bargaining power with respect to increasingly consolidated providers (Robinson 2003; Sinaiko 2012).[4] Insurers can tier their hospital networks, their physician networks, or both (Sinaiko 2012). The typical tiered hospital network has three tiers, with most or all hospitals in the market included in the network (Fronstin 2003). In my data, out-of-pocket price differentials between the most and least preferred tiers range from $200 up to $1,250.

This paper evaluates the demand-side response to tiered networks and explores whether consumers learn to be more price-responsive over time. I estimate a discrete choice model of demand for hospitals, using 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. Next, I extend the demand model to allow for changes in price-shopping behavior over time, and find suggestive evidence consistent with

[3]In some plan-years, consumers must pay out-of-pocket for at most one hospitalization every three months rather than simply the first four each year. This eliminates the incentive to price-shop for consumers who require hospitalization in quick succession within a three-month window of an initial hospitalization.

[4]A more detailed history of tiered provider networks is presented in the Appendix.

a Bayesian framework of consumers learning to price-shop. Using the hospital demand estimates, I simulate expected hospital shares, spending, and consumer welfare under various counterfactual plan designs and enrollment durations.

My empirical strategy and identification rely on comprehensive data on the private health insurance market in Massachusetts. I combine data on health care 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 et al. 2015).

This paper is related to a large literature on health insurance design and its relationship to health care demand. The paper contributes to the literature on the elasticity of demand for medical care by estimating substitution across providers in response to variation in out-of-pocket prices.[5] The majority of the existing evidence on the elasticity of health care demand, including the landmark estimates from the RAND Health Insurance Experiment, measures elasticities on the extensive margin of whether to purchase any health care. This paper estimates the intensive-margin price elasticity of demand across health care providers in response to price differences borne directly by consumers.[6] This exercise is closely related to the price transparency literature discussed above (Brot-Goldberg et al. 2015; Desai et al. 2016; Lieber 2017; Desai et al. 2017).

The finding that consumers do indeed price-shop in the setting of tiered networks suggests that the demand for health care is not inherently inelastic. Rather,

[5]The landmark estimates of the elasticity of health care 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 mostly fall in the same range (Manning et al. 1987; Chandra et al. 2010; Trivedi et al. 2010; Buntin et al. 2011).

[6]Gowrisankaran et al. (2015) and Ho and Pakes (2013) study 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 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).

health care 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.

The paper proceeds as follows. Section 2 describes the data and empirical setting. Section 3 details the empirical approach and results for the baseline hospital demand estimation, and Section 4 does the same for consumer learning. Counterfactual simulations and welfare implications are discussed in Section 5, and Section 6 concludes.

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

2.1 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) (CHIA 2014). The APCD consists of comprehensive data on interactions with the health care system of all privately insured residents of Massachusetts in the 2009–2012 period.

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 codes of residence. I match patients to zip-level demographic characteristics from the U.S. 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 Appendix Table 8. 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 Appendix.

The APCD reports several key price variables, including allowed amounts. Allowed amounts are actual transaction prices paid health care 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 tiered-network plans. The health care utilization data from the APCD are used to estimate hospital demand in conjunction with the hospital network data described below.

2.2 Hospital Network Data

I have compiled a unique dataset tracking Massachusetts hospitals' placements in several insurers' tiered and narrow networks for the period 2009–2015. 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% and 14% of commercial enrollment, respectively (CHIA 2013).[7] Network data were hand-collected from insurers' current and archived plan documentation, and cover both tiered and narrow networks.[8]

Data on insurers' provider networks are difficult for researchers to obtain, especially retrospective data that can be merged into claims databases, which has limited

[7]The largest insurer in Massachusetts is Blue Cross Blue Shield (BCBS), with 45% of the commercial market (CHIA 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).

[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 2.3. I thank Cindy McGrath at the GIC for sharing these data for the early years in the sample.

the scope of questions the literature has been able to address (Gaynor et al. 2015). To my knowledge, this paper is the first to use longitudinal tiered provider network data from multiple insurers, and indeed among the first to use longitudinal data on any type of provider network. The longitudinal nature of the data provides several sources of identifying variation for estimating demand response to out-of-pocket prices (see Section 3.2).

Massachusetts requires insurers operating tiered-network plans to "clearly and conspicuously indicate" consumers' out-of-pocket prices for each tier (Massachusetts 2012b). Insurers provide this information to enrollees as part of the schedule of benefits documentation for each plan. 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 Appendix Figure 5. This is in contrast to the difficulty of learning out-of-pocket prices for hospital care in advance under traditional plan designs. Price look-up tools require consumers to conduct a new search each time they consume health care, which often involves separate searches for each component of an episode of care.[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 Appendix Figure 4. All Massachusetts hospitals are in-network for these tiered-network plans.[10] Appendix Table 10 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 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.[11] 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 affect tier as-

[9]In the absence of a price search tool, obtaining price information is yet more difficult: even savvy consumers who ask for price quotes typically get poor response rates (Bebinger 2014).

[10]In the latter part of the sample period, the insurers introduce additional plans that use tiering on a narrow network, though take-up of these tiered narrow-network plans remains low.

[11]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.

signments.[12] Appendix Table 11 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.

The tiered network data are used to estimate hospital demand as a function of out-of-pocket price. Clean identification of a price coefficient in health care demand is typically impeded by a lack of data on insurers' provider network arrangements, especially retrospective data that can be merged into data on medical care usage (Gaynor et al. 2015). I overcome this identification challenge using my longitudinal tiered network data, which allow me to infer consumers' out-of-pocket prices at hospitals in which they are not treated.

2.3 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.[13] 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 (GIC 2008, 2009). My sample of GIC enrollees observed in the APCD includes approximately 90,000 employees and 120,000 dependents.

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). Plans on the GIC use copays, which are fixed dollar amounts paid out-of-pocket by consumers when they use health care. 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

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

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

2010, and increase to $275, $500, and $1,500 in 2016.

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 (2014). 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% of employee enrollees throughout the sample period. Additional information about GIC enrollees and insurance plans, including enrollee demographics, is presented in the Appendix.

Of the seven plans offered by other insurers, only one (UniCare) uses a tiered hospital network, and this plan has less than 10% 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 on 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 diagnostic category. By mid-2014, Tufts discontinued tiering by diagnostic category altogether, due to complaints about the complexity from providers and consumers.

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

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.

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 hospital's copay in 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.

3 Consumer Response to Simple Prices

If health care is different from most other goods in that health care 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 health carebut typically stymied by the complexity and unpredictability of prices, then tiered networks will steer consumers to 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.

3.1 Demand 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

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

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, past use of the hospital, and distance. Distance is an important determinant of hospital choice (Kessler and McClellan 2000; Town and Vistnes 2001; Capps et al. 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.[14] A dummy variable for past use of the hospital captures established relationships between patients and health care 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).[15] 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 Appendix Table 8.

[14] Calculated using Bing Maps driving directions.

[15] 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 (CMS 2014). The HCAHPS scores capture patients' perceptions of hospital quality and are highly correlated with other hospital reputation measures such as U.S. News rankings.

I assign each admission to a diagnostic category and severity level using the Clinical Classifications Software (CCS) from the Agency of 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 prevalence are described in Appendix Table 9. The model allows hospital choice to vary according to the hospital's availability of specialized services corresponding to the patient's diagnosis by including relevant interaction terms.[16] 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 2013). 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 hospital choices are subject to unobserved influences not related to price, 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

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

This probability is used to estimate the demand model using maximum likelihood.

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

3.2 Identification of Hospital Demand

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.

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 a 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 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%) are in different tiers across the two insurers. Of those, one fifth (10% 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.[17] The right panel of Table 2 shows the transition matrix of hospitals' tiers over time within the same insurer. Each cell (i, j) the percentage of hospitals starting in the row i tier in 2010 that have moved to the column j tier by 2014.[18] 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% 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 substantial 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

[17] By law, tier assignments can change at most annually (Massachusetts 2010).

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

Table 2: Variation in hospital tiers

Over time within insurer				Across insurers			
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%

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

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 tiers are a function of hospital quality or prestige in addition to negotiated price. 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.[19] This is unsurprising in light of the documented convergence of hospital quality scores over time, and the fact that Massachusetts legislation proposes a formula for mapping prices directly into networks that serves as a focal point for insurers (Massachusetts 2010). 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.

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

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 consumers' ability to select into health insurance plans. 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 2014). 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, therefore, an inertial consumer's plan characteristics increasingly approximate random assignment. I use the hospital's copay in the first year that a household enrolled in its current plan to deal with the endogeneity in the current copay. The identifying assumption would be violated if, for example, consumers are aware that the insurer intends to raise copays in the future, prior to the publication of those future plan characteristics. An analogous approach is employed by Abaluck et al. (2015) 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.[20] Despite the introduction of two new plans in 2011, 92% 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, at least 90% of enrollees remained. These patterns are consistent with the literature showing that consumers fail to re-optimize their plan choices over time (Handel 2013; Ericson 2014; Shepard 2014). Combined with these findings, the observed inertia motivates the identification strategy.

Table 3: Plan enrollment inertia on GIC, fiscal years 2010–2011

Plan	2010 Enrolt.	2011 Enrolt.	% 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

% 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.

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 et al. 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).[21] This approach requires an exclusion restriction analogous to standard IV methods, namely, that the "instrument" affects hospital choice only through its effect on co-

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

[21]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.

pay. 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 to a fifth-degree polynomial expansion.[22] 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.

3.3 Results of Hospital Demand Estimation

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 (Appendix Table 19). 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. I exclude admissions originating from the emergency department (ED) or via transfers from other hospitals. In such cases, patients have little leeway in choosing a hospital. Furthermore, state legislation prohibits out-of-pocket prices for care originating in the ED from varying by tier (Massachusetts 2010).

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 et al. 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 Appendix Table 12 for these and other additional coefficient estimates). Older patients and patients with chronic conditions are more

[22] Some papers have used two-stage residual inclusion (2SRI), where the residuals are entered into the second stage linearly (see, for example, Terza et al. (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).

willing to travel to their preferred hospital. Hospital fixed effects also display a sensible pattern.[23] 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 on patient-provider relationships, 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 2014).

The primary coefficient of interest is the coefficient 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 to the price of health care, 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.[24]

I do not find evidence of bias from consumers sorting into plans with low out-of-pocket prices for their preferred hospitals. Appendix Table 13 shows the results of the control function estimation described in Section 3.2, using the copays in the first year that a household enrolls in its current plan to instrument for the current copays. Estimates are shown up to a fifth-order polynomial expansion of the control func-

[23] Not shown; detailed results available upon request.

[24] 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% 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 Appendix Table 14).

Table 4: Hospital choice model

	(1) No hospital FEs		(2) Hospital FEs	
Hospital Choice				
Copay ($1,000s)	0.8169***	(0.0549)	-0.1833**	(0.0690)
Copay × std. income	0.1429**	(0.0524)	0.1904***	(0.0540)
Distance (mi)	-0.1817***	(0.0026)	-0.1832***	(0.0028)
Distance^2	0.0005***	(0.0000)	0.0006***	(0.0000)
Past use of hospital	4.1221***	(0.0473)	3.8292***	(0.0498)
Age × distance	0.0001	(0.0000)	0.0000	(0.0000)
Chronic cond × distance	0.0221***	(0.0015)	0.0216***	(0.0015)
Hospital FEs	No		Yes	
Pseudo R^2	0.558		0.605	
Nadmits	29917		29917	

Multinomial logit model of hospital choice. All copay coefficients scaled to $1,000s. Includes additional covariates (more coefficients in Appendix Table 12). Standard errors in parentheses, clustered by patient. Nadmits = number of admissions.

tion 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 even greater price sensitivity than the uninstrumented estimates. Existing findings of limited responsiveness to price transparency warrant caution, so my preferred specification for hospital demand is the conservative and parsimonious model without the control function (Table 4).

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 chosenone. Adding an interaction term between copay and an indicator for the consumer's initial enrollment in the current plan being non-tiered does not change the primary coefficient on copay, nor is the interaction term is statistically significant (column 1 of Appendix Table 14).

The hospital price elasticities implied by the demand model are summarized in Table 5, 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 a fixed $1,000 copay, respectively. Elasticities for specific Boston hospitals are shown in Appendix Table 15. Own-price elasticities of demand range from -0.03 to -0.12. 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 intensive margin of choosing between options, conditional on seeking care in the first place. This result highlights the importance of price transparency for controlling moral hazard on the intensive margin as well as the better-studied extensive margin (Pauly 1968). These estimates are also less elastic than in Gowrisankaran et al. (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 2014).

Table 5: Price elasticities from hospital demand model (at median household income)

Elasticities	Metro Boston	Outside Boston
Own-price (at observed copays)	-0.039 (0.002)	-0.030 (0.002)
Own-price (at $1,000 copays)	-0.117 (0.006)	-0.092 (0.005)
Cross-price (at observed copays)	0.002 (0.000)	0.000 (0.000)
Cross-price (at $1,000 copays)	0.004 (0.001)	0.001 (0.001)

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.

Table 5 also reports hospitals' pairwise cross-price elasticities. They range from essentially zero to approximately 0.05. 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 (Figure 4), allowing consumers to more easily substitute across hospitals in response to copay differences. Appendix Table 15 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). Table 15 also reports elasticities for Cape Cod Hospital, which is geographically isolated in eastern Massachussets 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 I am capturing real patterns in how patients choose hospitals.

4 Consumer Learning

This section explores patterns in price responsiveness over time by mapping a model of Bayesian learning about prices to a reduced-form empirical specification. The finding that consumers respond to differential out-of-pocket prices by substituting toward lower-priced health care providers contrasts with much of the recent literature (Sinaiko and Rosenthal 2014; Brot-Goldberg et al. 2015; Desai et al. 2016, 2017). Relative to these papers, a unique feature of my empirical context is that consumers are observed up to six years after their initial enrollment in a tiered-network plan.[25] Existing papers have had access to at most two years of data

[25] Some consumers in the data may be enrolled for longer still. The earliest enrollments reported in the APCD are in 2006, but some plans began tiering hospitals as early as 2003 and many consumers

after the introduction of a price transparency regime.[26] My finding of a response to out-of-pocket prices may be partially explained by consumers becoming more price-responsive over time.

The potential presence of increasing price sensitivity is suggested by Figure 1. The figure shows the fraction of hospital admissions originating from hospitals in each tier of the largest GIC plan, Harvard Pilgrim's Independence plan. This plan transitioned from a traditional network with flat copays to a tiered network in the 2010 plan year. Since then, admissions to the preferred tier (tier 1) have risen from 22% to 26% of total volume. Admissions to the least preferred tier (tier 3) have fallen from 43% to 33% of volume. The top Boston academic medical centers account for a disproportionately large share of the drop in admissions to tier 3.[27]

The remainder of this section presents suggestive evidence that consumers in tiered-network plans learn to be more price-responsive over time. Price responsiveness increases both as a result of price signals observed from the consumer's own health care consumption history and as a result of third-party information.

4.1 Intuition for Consumer Learning

A consumer who requires inpatient hospital care must choose a hospital at which to be treated. In a tiered-network plan, the consumer has low-cost access to information about the out-of-pocket prices at any hospital in the network. An uninformed consumer ignores out-of-pocket prices when choosing a hospital, but her ex post realized utility nonetheless includes the out-of-pocket price she will have to pay. If the expected gain from a choice under known out-of-pocket prices relative to a choice ignoring prices exceeds the information search cost, the consumer will search.

observed in the data likely began their enrollments prior to the earliest reported 2006 dates.

[26] Brot-Goldberg et al. (2015) study two years of enrollment in new high-deductible health plans, and document substantial overall quantity reductions but no effect of consumer substitution toward lower-priced care. Desai et al. (2016), Desai et al. (2017) and Lieber (2017) each study consumers for one year to fifteen months after the introduction of a price transparency tool, and find no spending reductions, small out-of-pocket prices paid reductions, and moderate out-of-pocket price paid reductions, respectively.

[27] 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 health care system and are widely considered "star" hospitals (Ho 2009; Shepard 2014).

Figure 1: Admissions by tier in the Harvard Pilgrim Independence plan over time

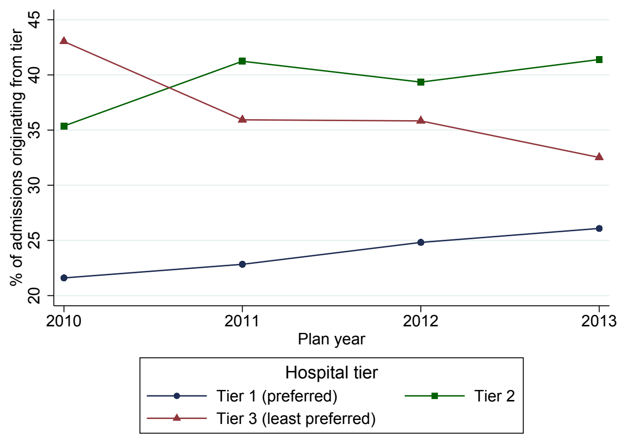

Harvard Pilgrim's Independence plan began as a broad-network plan with no tiering up until the 2009 plan year, and began using a broad tiered hospital network in the 2010 plan year. The fraction of admissions originating from hospitals in the most preferred tier rise over time, while admissions to the least preferred tier decline.

Consumers who are informed but perfectly inelastic are observationally equivalent to consumers who are completely ignorant of tiering. The presence of consumers in an estimation sample who are unaware of differential pricing will therefore bias the price coefficient toward zero.

Consumers who are aware of the differential out-of-pocket prices in their plan may still decide not to search. A consumer with a large gap in utility between her most preferred hospital and the next-best option gains little from learning prices, since price differences are unlikely to be large enough to induce hospital switching. Similarly, a consumer with a low price sensitivity has little to gain from searching. In such a case, low propensity to search is observationally equivalent to a true low price sensitivity.[28] Conversely, the higher the variance of the consumer's priors over prices, the greater the expected gains from searching.

Beliefs about out-of-pocket prices may change due to consumers' own experiences with their plan or due to external factors.[29] For example, consumers may

[28] A consumer may have low price sensitivity if, for example, she is high-income and therefore has a relatively low marginal utility of money.

[29] A third possibility is that referring physicians learn about the structure of their patients' plan and shift their referral patterns toward hospitals with lower out-of-pocket prices.

learn via employer information campaigns or through word of mouth that out-of-pocket prices vary. They may also Bayesian update their beliefs after observing different out-of-pocket prices across their own health care consumption.

A formal model of the learning process by which consumers update their beliefs is presented in the Appendix. The out-of-pocket price for a given hospital can take on one of a discrete set of possible copays in each plan. Consumers start with prior beliefs about the probability of each possible copay. Each time a consumer receives medical care, she observes ex post its out-of-pocket price, and Bayesian updates her priors to arrive at a posterior distribution. Updating from consumers' own health care experiences enters through changes in the updating parameters, while shifts in third-party information enter through changes in the prior parameters.

Consider a consumer whose prior places arbitrarily large weight on a copay value of zero. Zero copays are frequent, including for preventive health care services, 30-day readmissions, or after exceeding the out-of-pocket maximum. With every additional zero-copay health care interaction, the consumer's posterior distribution becomes tighter around a copay value of zero. A consumer with a tight posterior distribution around a zero copay is unlikely to search for out-of-pocket price information. Conversely, the greater the fraction of realized copays that are greater than zero, the higher the posterior probability of non-zero copays and the greater the variance of the posterior distribution. Derivations of the variance and comparative statics are provided in the Appendix.

4.2 Empirical Approach to Learning Estimation

Consumers may learn about out-of-pocket pricing structures from their own past health care consumption, or from third-party information such as employer information campaigns or word of mouth. Distinguishing between potential sources of learning is important for policy and optimal plan design.The effectiveness of public awareness campaigns that educate consumers about differential out-of-pocket pricing hinges on consumers learning to pay attention to out-of-pocket prices in response to third-party information.

In addition, the effect of learning from own consumption experiences can be distinguished from the effect of learning from third-party information. Any employer-wide effects, such as word of mouth, will manifest in a secular trend of increasing price sensitivity over time. They will be observed among consumers with minimal past health care utilization as well as those with frequent interactions with the health care system. Consumer-specific learning as a result of own observations of out-of-pocket prices during past utilization will instead manifest in consumers with a large number of past nonzero copays appearing more price-sensitive.[30] These implications of the learning framework provide an empirical test for whether consumers are learning about out-of-pocket prices.

Due to the sparse nature of the data, I estimate these effects using reduced-form parameters in the discrete choice model rather than explicitly estimating the full Bayesian learning model.[31] To impose the structural learning model would be to ask a lot of the data, particularly in a discrete choice setting with sparse choices among over 70 hospitals. Instead, I rely on the variation in hospital choice behavior across consumers at different durations of enrollment in tiered-network plans. In these models, I drop the control function approach because its primary source of variation, the structure of the plan at the consumer's initial date of enrollment, is determined by the same factors that determine the duration of enrollment.

The reduced-form implementation of the learning model distinguishes between learning from own past utilization and third-party information. However, it does not allow me to further disentangle the underlying mechanisms . The estimates do not distinguish between information diffusion among coworkers and learning from centralized information campaigns by the employer. Since my data do not include any information on information campaigns or employee position, department, or geographic location, there would be little insight gained from estimating the full

[30]In the data, a Bayesian-updating consumer who learns about the pricing structure in her plan according to the framework in the Appendix will appear less price-sensitive the more zero copays she has encountered to date, and more price-sensitive the more nonzero copays she has encountered.

[31]The estimation of learning models typically relies on repeated observations of the same agent's decision. These repeated observations form a sequence of choices that identifies the shift in the agent's preferences or information set over time. In the context of hospital admissions, the within-consumer data are sparse. The majority of consumers in the data are never admitted to a hospital over the course of the sample period; among those with any inpatient hospital utilization, only 37% have two or more admissions, and 6% have five or more. Consequently, there is insufficient within-person variation over time to identify the learning parameters.

learning model. In addition, unlike the baseline hospital demand model, the learning model implementation proceeds without a control function. I therefore interpret the results in this section as suggestive evidence about consumer learning.

4.3 Results of Learning Estimation

Table 6 reports the results of the hospital demand estimation, now allowing for changes in price-responsiveness. Column 1 estimates the relationship suggested by Figure 1: consumers become more responsive to out-of-pocket price the longer they are enrolled in a tiered-network plan. Columns 2 and 3 further decompose the sources of this apparent learning. All models in Table 6 include the covariates from the baseline demand model (Table 4). These coefficients are similar in magnitude and are omitted for brevity.

In column 1 of Table 6, the key coefficient is on the interaction of copay with the number of months since a consumer first enrolled in a tiered-network plan.[32] The negative estimate indicates that the longer a consumer is enrolled, the more likely she is to choose hospitals with low out-of-pocket prices. Regressions with polynomial terms for time or enrollment duration show no evidence of learning slowing down.[33] Although the point estimates suggest a positive price coefficient for new enrollees, this is an artifact of the linear specification of the time trend. The point estimate for copays is negative in a sample restricted to first-year enrollees in a tiered-network plan (Appendix Table 14, column 3).

Column 2 asks whether the learning suggested by column 1 is explained by third-party information available to all enrollees. The interaction of copay with a secular time trend is negative and significant even after controlling for consumers' enrollment duration in a tiered-network plan. This result is indicative of some learning from third-party information. Unfortunately, in the absence of additional information that would allow me to study information campaigns or word of mouth between coworkers, I cannot disentangle the mechanisms driving this effect.

Column 2 also suggests a role for consumers' own experiences in tiered-network

[32] Separate interaction terms are included for consumers whose enrollment in a tiered-network plan began prior to the start of the data or whose enrollment history is unobserved.

[33] Results available from the author upon request.

Table 6: Hospital choice model with learning

	(1) Enrolt. duration	(2) Enrolt. + time trend	(3) Past use + trend
Hospital Choice			
Copay × months enrolled	-0.0236***	-0.0168***	
	(0.0040)	(0.0044)	
Copay × left-censored enrolt.	0.0735	0.0520	
	(0.1596)	(0.1603)	
Copay × unknown enrolt.	-0.3897**	-0.1155	
	(0.1232)	(0.1459)	
Copay × calendar months		-0.0172***	-0.0255***
		(0.0049)	(0.0043)
Copay × count of $0 claims			0.0049**
			(0.0018)
Copay × count of non-$0 claims			-0.0065*
			(0.0031)
Copay ($1,000s)	0.2133*	0.5363***	0.4894***
	(0.1085)	(0.1431)	(0.1341)
Copay × std. income	0.1937***	0.2031***	0.2164***
	(0.0542)	(0.0545)	(0.0546)
Hospital FEs	Yes	Yes	Yes
Pseudo R^2	0.606	0.606	0.606
Nadmits	29917	29917	29917

Multinomial logit model of hospital choice. All copay coefficients scaled to $1,000s for ease of interpretation. Enrollment variables measure time since first enrolled in a tiered-network plan. Standard errors in parentheses, clustered by patient. Nadmits = number of choice sets (admissions).

plans: the interaction of copay and tiered enrollment duration remains significant when the calendar time trend is added. If consumers learn from their own experiences, the estimated price sensitivity should be larger for consumers who have had many nonzero-price health care encounters, conditional on enrollment duration. Column 3 therefore replaces the interaction of copay and duration with two interaction terms measuring consumers' own experiences with tiered networks. The first interacts current copay with the total number of past health care encounters the consumer has experienced in the tiered-network plan with a $0 out-of-pocket price, and the second is analogous for health care encounters with positive out-of-pocket prices.[34] All major categories of medical care including physician services, outpatient hospital care, and inpatient hospital admissions are included in the counts.[35]

The results in column 3 suggest a nuanced consumer response to experience. Greater past utilization does not unambiguously drive up price sensitivity. Consumers with a large number of $0 past health care encounters appear to become less price-sensitive, but this is more than offset by encounters with positive out-of-pocket prices. A single encounter with a nonzero price has an effect equivalent to more than two additional months of passive enrollment. These patterns are consistent with the Bayesian learning framework outlined in Section 4.1.

This section highlights the importance of accounting for changes in the response to demand-side incentives over time. I find suggestive evidence of consumers learning to price-shop. A reevaluation with a longer sample would shed light on whether the rate of learning eventually slows down and price responsiveness reaches a steady state. In this sample, I find no evidence of plateauing. To understand the long-term effects of recent plan design innovations, such as price look-up tools and high-deductible health plans, existing evidence on their short-term effects will need to be supplemented by studying them several years after implementation.

5 Implications for Spending and Welfare

Demand-side incentives can be an effective cost control tool if the demand response is sufficient for meaningful spending reductions, and if the harm to consumers is not too great. Tiered networks may reduce consumers' welfare via higher out-of-pocket spending and steering away from their preferred hospitals. In this section, I quantify the potential savings on the table, and perform a back-of-the-envelope calculation of the potential for compensating consumers' welfare losses.

[34] Consumers most commonly face $0 out-of-pocket prices for certain services such as preventive care, after exceeding their deductible, or for 30-day hospital readmissions.

[35] The underlying assumption is that consumers extrapolate from the out-of-pocket prices they observe for any class of medical care, and use that information to form beliefs about the structure of out-of-pocket pricing for inpatient hospital care. Due to the sparseness of hospital admissions, this or a similar assumption is required in order to estimate the learning effects. Pharmaceutical drug purchases are not included; many consumers are accustomed to tiered cost-sharing for drugs since well before the start of hospital tiering.

5.1 Average Spending Effects of Tiered Networks

To evaluate the average effect of tiered networks on spending, I simulate inpatient hospital spending under a non-tiered network and various tiered-network and coinsurance 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 same tiered 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 hospitals. In addition, I simulate coinsurance scenarios by setting out-of-pocket prices to 5%, 10%, and 20% of the total diagnosis-adjusted price.

I simulate hospital shares for each patient-diagnosis pair using the hospital demand estimates from Section 3.3, assuming all 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.

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% of their baseline volume moving from the flat network to the tiered network with a \$1,250 spread; tier 1 gains 6.0% of baseline volume. Total spending per hospital admission falls by 1.3%. 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.

Appendix Table 18 shows the the simulated scenarios with coinsurance. Coinsurance allows insurers to fine-tune incentives by passing through all price differences across hospitals, not just price differences across tiers. When consumers pay a low coinsurance rate of 5% of the total price, predicted spending actually rises. The marginal differences in out-of-pocket price are in many cases markedly less than tiered copay differences, which blunts the incentives. A higher coinsurance rate of 10% or 20% is required to generate spending reductions. Importantly, these coinsurance simulations assume that consumers will respond equivalently to copays and coinsurance. In practice, copays are substantially simpler for consumers to understand. Moreover, with coinsurance of 10% or higher, a consumer will likely exceed her out-of-pocket maximum on or soon after her first hospitalization, nullifying the marginal incentive to price-shop. These simulations should therefore be interpreted as an upper bound on the spending reductions achievable with coinsurance.

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 health care among low-income consumers. While assessing these distributional effects is beyond the scope of this paper, this issue remains important for policy. I now turn to a back-of-the-envelope calculation of average welfare effects.

Table 7: Hospital sorting counterfactuals (at median household income)

	Flat copay $250	$250/500/750	$250/500/1,500
Tier 1 hospitals % of volume	27.15	27.96	28.77
Tier 2 hospitals % of volume	36.65	36.76	37.87
Tier 3 hospitals % of volume	36.21	35.28	33.36
Patient spending per admission ($)	250	518	762
Change in patient $ over flat copay	–	107.32%	204.67%
Insurer spending per admission ($)	18,564	18,205	17,814
Change in insurer $ over flat copay	–	-1.93%	-4.04%
Total spending per admission ($)	18,810	18,715	18,564
Change in total $ over flat copay	–	-0.5%	-1.31%

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.

5.2 Spending and Welfare Under Learning

If consumers learn to price-shop, then spending reductionsfrom tiered networks may grow over time. On the other hand, consumers in tiered-network plans may face reduced welfare due to higher out-of-pocket spending and lower utilization of their preferred hospitals. This section simulates hospital utilization, spending, and consumer welfare over various durations of enrollment in tiered networks.

Figure 2a shows the simulated share of volume in tiers 1 and 3, assigning all consumers to enrollment durations between zero and four years. The estimates from column 1 of Table 6 are used, which collapse all learning into a single reduced-form enrollment duration coefficient. The exercise parallels the previous section, simulating tier copay structures of \$250, \$500, and \$750 and copays of \$250, \$500, and \$1,500. After three years in the tiered network with the smaller spread in copays, the preferred tier, tier 1, gains 53.1% relative to its baseline volume at the initial month, and tier 3 loses 36.9%. In the plan with the \$1,250 total spread in copays, tier 1 hospitals gain 89.4% of initial volume and tier 3 hospitals lose 77.6%.

The shift away from high-priced tier 3 hospitals reduces the average price of a hospital admission. Figure 2b shows the simulated spending per admission as a function of tiered-network enrollment duration. 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 \$2,500, or 7.8% of the initial spending. In the plan with the larger \$1,500 tier 3 copay, average spending falls by nearly \$3,000, or 14.9% of the initial spending.

The simulations of long enrollments should be interpreted solely in a partial-equilibrium sense. After four years in a tiered network, the simulated volume of patients at tier 1 hospitals doubles. Hospital capacity constraints would bind if the entire market switched to tiered networks. In my empirical context, the GIC population that enters the simulations makes up approximately 6% of the population of Massachusetts, so a doubling of GIC volume at tier 1 hospitals would amount to a more manageable 3% increase. Nevertheless, it is likely that consumers in tiered networks climb an initial learning curve that makes them appear increasingly price-sensitive over time, but that their price sensitivity ultimately plateaus. As discussed above, no slowdown in learning is apparent in my four years of data. To draw conclusions on longer-term effects, longer samples are needed. In addition, supply-side responses may occur if insurers can successfully leverage the threat of non-preferred tier placement to negotiate lower hospital prices.[36]

Figure 2: Hospital sorting with learning (at median household income)

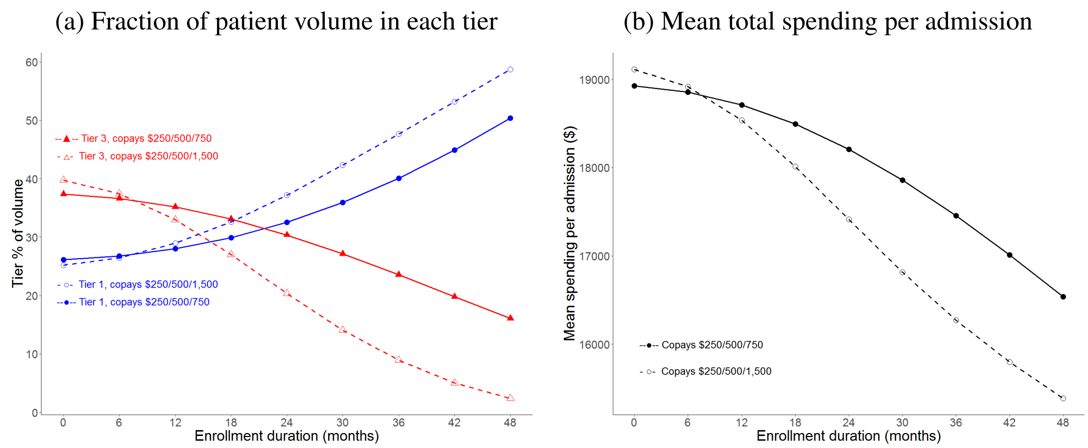

Projected hospital volumes and mean spending per admission over time using the learning estimates, using the Harvard Pilgrim Independence hospital network. Solid lines represent the plan's observed copay regime of $250, $500, and $750 for hospitals in tiers 1, 2, and 3, respectively. Dashed lines represent a regime where the copay for tier 3 is doubled from $750 to $1,500.

The overall expected spending reductions from tiered networks come at the cost of steering consumers away from preferred hospitals and increased out-of-pocket spending. In addition, tiered-network plans offer less financial risk-smoothing due to the spread in potential out-of-pocket prices that consumers may face. However, if the resulting spending reductions exceed consumers' welfare losses, then consumers can be compensated via a transfer in the form of lower insurance premiums.

To measure changes in consumer welfare, I calculate consumer surplus using Appendix Equation 3, which corresponds to the willingness-to-pay (WTP) measure widely used in the health economics literature to measure valuation for provider networks (Capps et al. 2003). In this case, the availability of a direct estimate of price responsiveness from the hospital demand model allows me to dollarize WTP, rather than measuring it in utils as in settings that lack out-of-pocket price

[36] In a companion paper, I model the potential hospital-insurer bargaining responses explicitly.

variation.[37] The calculation of WTP requires each consumer's ex ante probability of each possible diagnosis for the year. I calculate these probabilities for each sex–10-year age band cell and each CCS diagnostic category using data on all non-transfer hospital admissions of Massachusetts residents from the 2010 HCUP State Inpatient Database.[38] Additional details on constructing the probabilities and WTP measure from the data are given in the Appendix. Since patient covariates such as distance to hospitals vary across zip codes, WTP for a given hospital network takes on a separate value for each gender-age group-zip code triplet. Allowing for this granular variation in consumers' preferences and admission probabilities at the diagnostic category level allows the WTP measure to capture rich variation.

Figure 3 shows the distribution of lost WTP due to a tiered network with copays of \$250, \$500, and \$750, subtracted from mean spending reductions per enrollee. The mass of the distribution to the right of \$0 corresponds to the fraction of consumers who can be compensated for their welfare losses at various lengths of enrollment. Initially, consumers are not yet responding to out-of-pocket prices, and most consumers' loss in WTP exceeds the mean spending reduction. After two years, more than 90% of consumers have predicted WTP losses small enough to be compensated via the spending reduction. In addition, Figure 6 in the Appendix shows a more conservative welfare calculation that double-counts consumers' welfare loss from higher expected out-of-pocket spending under tiered networks.[39] An insurer can therefore compensate enrollees for their utility loss from tiered networks in the

[37] Dollarized WTP is equivalent to standard consumer surplus in multinomial logit models.

[38] This is equivalent to the assumption that that a consumer's expectation of her health status for the upcoming year is a consistent predictor of her health status, given only her sex, her 10-year age group, and the fact of residing in Massachusetts. This assumption is more likely to hold for relatively healthy consumers who do not have highly informative personal experience to inform their ex ante expectations of diagnosis (Shepard 2014). Since my data consist of non-elderly, commercially insured, mostly employed individuals, they are healthier than the general population and good candidates for the assumption that their expected health status is approximately equal to the average health status for their age group. To the extent that there are deviations from the average health status, they will load onto the error term in the plan choice model.

[39] WTP already accounts for expected spending, but in Figure 6, I explicitly add each consumer's expected out-of-pocket spending increase to the WTP loss. With this double-counting of projected out-of-pocket, nearly 80% of consumers can be compensated for their "behavioral welfare loss" after two years of enrollment, and by three years, this fraction exceeds 90%. This allows for behavioral biases such as valuing low out-of-pocket spending over and above the expected utility from a hospital network given its out-of-pocket spending.

form of lower premiums, while still reaping higher net profits.

Figure 3: Compensating consumers for welfare losses from tiered networks

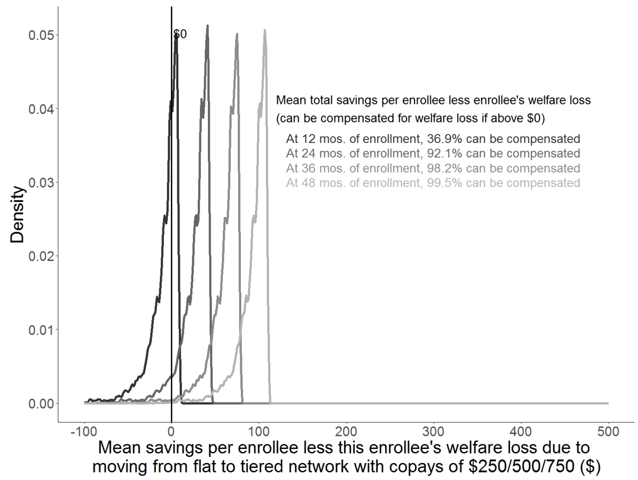

Mean savings per enrollee less each enrollee's welfare loss due to moving from a flat $250 copay to the Harvard Pilgrim Independence hospital network with its observed copay regime of $250, $500, and $750 for hospitals in tiers 1, 2, and 3, respectively. Hospital shares are calculated using copays from the learning model at each listed enrollment duration; utilities and WTP are calculated based on those shares and the copay coefficient at 36 months of enrollment.

6 Conclusion

Reliance on market forces plays a larger role in health care policy in the United States than in most advanced economies. In the last decade, market-based approaches to health care delivery have increasingly focused on demand-side financial incentives as a mechanism for reducing health care spending. This paper shows that, contrary to recent evidence, consumers can be successfully incentivized to price-shop for their health care under certain conditions. I find that in tiered-network health insurance plans where out-of-pocket prices for health care are clearly stated, predictable, and simple to understand, consumers price-shop across hospitals. These findings suggest that consumers' frequent failure to price-shop for health care may be attributable to the complexity of health care decision-making, rather than an inherent insensitivity to health care 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 health care treatments, where conventional wisdom suggests the least elastic consumption (Manning et al. 1987). 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 health care 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 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 health care.

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 (Brot-Goldberg et al. 2015; Desai et al. 2016; Lieber 2017; Desai et al. 2017).

My findings have several implications for health care 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 health care 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. Moreover, at least in the empirical context of this paper, consumers' welfare losses from demand-side incentives can be more than compensated by concomitant spending reductions through lower premiums. 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.

Policy-makers 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 our uncertain and complex medical care environment, 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 health care 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 their intelligibility to those consumers.

References

Abaluck, Jason, Jonathan Gruber, and Ashley Swanson (2015) "Prescription Drug Use under Medicare Part D: A Linear Model of Nonlinear Budget Sets," Working Paper 20976, National Bureau of Economic Research.

Bebinger, Martha (2014) "Mass. Patients Can 'Shop' For Health Care – At Least In Theory," *Kaiser Health News*.

Brot-Goldberg, Zarek C., Amitabh Chandra, Benjamin R. Handel, and Jonathan T. Kolstad (2015) "What does a deductible do? the impact of cost-sharing on health care prices, quantities, and spending dynamics,"Technical report, National Bureau of Economic Research.

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," *The American Journal of Managed Care*, Vol. 17, No. 3, pp. 222–230.

Capps, Cory, David Dranove, and Mark Satterthwaite (2003) "Competition and Market Power in Option Demand Markets," *The RAND Journal of Economics*, Vol. 34, No. 4, p. 737.

Chandra, Amitabh, Jonathan Gruber, and Robin McKnight (2010) "Patient Cost-Sharing and Hospitalization Offsets in the Elderly," *The American Economic Review*, Vol. 100, No. 1, pp. 193–213.

CHIA, (Massachusetts Center for Health Information and Analysis) (2013) "Annual Report on the Massachusetts Health Care Market,"Technical report, Commonwealth of Massachusetts.

——— (2014) "All-Payer Claims Database."

——— (2015) "2015 Annual Report on the Performance of the Massachusetts Health Care System," Annual Report 15-245-CHIA-01, Commonwealth of Massachusetts.

CMS, (Centers for Medicare & Medicaid Services) (2014) "HCAHPS: Patients' Perspectives of Care Survey," September.

Cutler, David M., Mark McClellan, and Joseph P. Newhouse (2000) "How Does Managed Care Do It?" *The RAND Journal of Economics*, Vol. 31, No. 3, p. 526.

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," *JAMA*, Vol. 315, No. 17, p. 1874.

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*, Vol. 36, No. 8, pp. 1401–1407.

DHCFP, (Division of Health Care Finance and Policy) (2010) "Massachusetts Health Care Cost Trends 2010 Final Report,"Technical report, Commonwealth of Massachusetts.

Ericson, Keith M. Marzilli (2014) "Consumer Inertia and Firm Pricing in the Medicare Part D Prescription Drug Insurance Exchanges," *American Economic Journal: Economic Policy*, Vol. 6, No. 1, pp. 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*.

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*.

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

Gaynor, Martin, Kate Ho, and Robert J. Town (2015) "The Industrial Organization of Health-Care Markets," *Journal of Economic Literature*, Vol. 53, No. 2.

GIC, (Group Insurance Commission) (2008) "Fiscal Year 2007 Annual Report," annual Report, Commonwealth of Massachusetts Group Insurance Commission.

——— (2009) "Fiscal Year 2008 Annual Report," annual Report, Commonwealth of Massachusetts Group Insurance Commission.

Gowrisankaran, Gautam, Aviv Nevo, and Robert Town (2015) "Mergers When Prices Are Negotiated: Evidence from the Hospital Industry," *American Economic Review*, Vol. 105, No. 1, pp. 172–203.

Gruber, Jonathan and Robin McKnight (2014) "Controlling Health Care Costs Through Limited Network Insurance Plans: Evidence from Massachusetts State Employees,"Technical report, National Bureau of Economic Research.

Handel, Benjamin R (2013) "Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts," *American Economic Review*, Vol. 103, No. 7, pp. 2643–2682.

HHS (2016) "Federal Register, Part II. Vol. 81, No. 45,"Technical report, Department of Health and Human Services.

Ho, Kate and Ariel Pakes (2013) "Hospital choices, hospital prices and financial incentives to physicians,"Technical report, National Bureau of Economic Research.

Ho, Katherine (2006) "The welfare effects of restricted hospital choice in the US medical care market," *Journal of Applied Econometrics*, Vol. 21, No. 7, pp. 1039–1079.

——— (2009) "Insurer-Provider Networks in the Medical Care Market," *The American Economic Review*, Vol. 99, No. 1, pp. 393–430.

Kessler, Daniel P. and Mark B. McClellan (2000) "Is Hospital Competition Socially Wasteful?" *The Quarterly Journal of Economics*, Vol. 115, No. 2, pp. 577–615.

KFF, (Kaiser Family Foundation) (2014) "Employer Health Benefits: 2014 Annual Survey,"Technical report, Kaiser Family Foundation.

——— (2015) "2015 Employer Health Benefits Survey,"Technical report, Henry J. Kaiser Family Foundation.

——— (2016) "Employer Health Benefits: 2016 Annual Survey,"Technical report, Kaiser Family Foundation.

Kolstad, J. T. and M. E. Chernew (2009) "Quality and Consumer Decision Making in the Market for Health Insurance and Health Care Services," *Medical Care Research and Review*, Vol. 66, No. 1 suppl, pp. 28S–52S.

Lieber, Ethan M.J. (2017) "Does It Pay to Know Prices in Health Care?" *American Economic Journal: Economic Policy*, Vol. 9, No. 1, pp. 154–179.

Manning, Willard J., 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," *The American Economic Review*, Vol. 77, No. 3, pp. 251–277.

Martin, Shawn (2014) "We're Not Gonna Take It: Network Optimization Disrupts Continuity of Care," *American Academy of Family Physicians*.

Massachusetts, Commonwealth of (2010) "An Act To Promote Cost Containment, Transparency And Efficiency In The Provision Of Quality Health Insurance For Individuals And Small Businesses."

——— (2012a) "An Act Improving The Quality Of Health Care And Reducing Costs Through Increased Transparency, Efficiency And Innovation."

——— (2012b) "An Act Relative to Tiered and Selective Network Health Plans."

McKinsey, and Company (2015) "Hospital networks: Evolution of the configurations on the 2015 exchanges,"Technical report, McKinsey Center for U.S. Health System Reform.

——— (2016) "Hospital networks: Perspective from three years ofexchanges,"Technical report, 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*, Vol. 36, No. 8, pp. 1392–1400.

Pauly, Mark V. (1968) "The economics of moral hazard: comment," *The American Economic Review*, pp. 531–537.

Petrin, Amil and Kenneth Train (2010) "A control function approach to endogeneity in consumer choice models," *Journal of Marketing Research*, Vol. 47, No. 1, pp. 3–13.

Robinson, J. C. (2003) "Hospital Tiers In Health Insurance: Balancing Consumer Choice With Financial Motives," *Health Affairs*.

Samek, Anya and Justin Sydnor (2016) "Simplifying Health Insurance Choice with Consequence Graphs," December.

Shepard, Mark (2014) "Hospital Network Competition and Adverse Selection: Evidence from the Massachusetts Health Insurance Exchange," working paper.

Sinaiko, Anna D. (2012) "Tiered networks as strategy to improve health care quality and efficiency,"Technical report, National Institute for Health Care Management.

Sinaiko, Anna D. and Meredith B. Rosenthal (2014) "The Impact of Tiered Physician Networks on Patient Choices," *Health Services Research*, Vol. 49, No. 4, pp. 1348–1363.

Song, Z., D. G. Safran, B. E. Landon, M. B. Landrum, Y. He, R. E. Mechanic, M. P. Day, and M. E. Chernew (2012) "The 'Alternative Quality Contract,' Based On A Global Budget, Lowered Medical Spending And Improved Quality," *Health Affairs*, Vol. 31, No. 8, pp. 1885–1894.

Starc, Amanda (2014) "Insurer pricing and consumer welfare: Evidence from medigap," *The RAND Journal of Economics*, Vol. 45, No. 1, pp. 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*, Vol. 27, No. 3, pp. 531–543.

Town, Robert and Gregory Vistnes (2001) "Hospital competition in HMO networks," *Journal of Health Economics*, Vol. 20, No. 5, pp. 733–753.

Trivedi, Amal N., Husein Moloo, and Vincent Mor (2010) "Increased ambulatory care copayments and hospitalizations among the elderly," *New England Journal of Medicine*, Vol. 362, No. 4, pp. 320–328.

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

Wrobel, Marian V., David Auerbach, and Sara Sadownik (2014) "2014 Cost Trends Report," annual Report, Commonwealth of Massachusetts Health Policy Commission.

For Online Publication

Appendix A: Additional Tables and Figures

Table 8: Inpatient admissions for hospital demand model

Mean age	41.6	–	–
% female	64.1	–	–
% chronic	34.6	–	–
% in tiered plans	65.6	–	–
	Non-tiered	Tier 1	Tier 2, 3
% of admits	34.4	31.2	68.8
Mean distance	15.1	11.5	15.9
Mean copay ($)	240.2	268	614.8

Summary statistics for admissions used to estimate the hospital demand model. Two-thirds of admissions are from enrollees in tiered plans. First column of second panel reports non-tiered plans' share of admissions and characteristics. Columns 2 and 3 report tiered plan admissions. Patients travel farther to hospitals in higher-copay tiers.

Table 9: Descriptions and prevalence of CCS diagnostic categories

Code	Description	Share
1	Infectious and parasitic diseases	1.9
2	Neoplasms	4.9
3	Endocrine; nutritional; and metabolic diseases and immunity disorders	3.9
4	Diseases of the blood and blood-forming organs	0.9
5	Mental illness	9.8
6	Diseases of the nervous system and sense organs	2.7
7	Diseases of the circulatory system	10.2
8	Diseases of the respiratory system	7.5
9	Diseases of the digestive system	10.0
10	Diseases of the genitourinary system	3.9
11	Complications of pregnancy; childbirth; and the puerperium	13.5
12	Diseases of the skin and subcutaneous tissue	2.1
13	Diseases of the musculoskeletal system and connective tissue	5.4
14	Congenital anomalies	0.5
15	Certain conditions originating in the perinatal period	13.1
16	Injury and poisoning	7.1
17	Symptoms; signs; and ill-defined conditions	2.1
18	Residual codes; unclassified; all E codes	0.3

Clinical Classifications Software (CCS) diagnostic categories. First column is Level 1 code (the broadest level), second column is description, third column is % share of nonelderly hospital discharges in Massachusetts.

Table 10: Distribution of hospitals across tiers, 2012

# of Hospitals in	HPHC	Tufts
Tier 1	28	39
Tier 2	20	2
Tier 3	13	20
Total	61	61

Counts of hospitals in each tier for a sample year. HPHC is Harvard Pilgrim. Satellite campuses are excluded.

Figure 4: Massachusetts insurers' hospital tiers (2012)

(a) Harvard Pilgrim

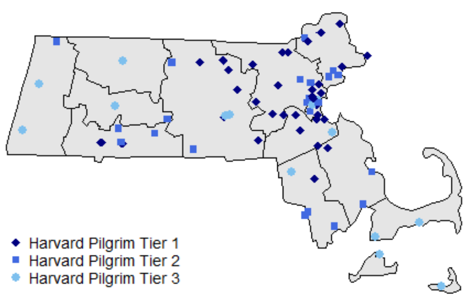

(b) Tufts

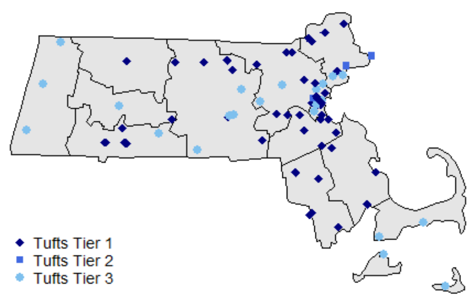

Maps of Harvard Pilgrim's and Tufts' tiered hospital networks in 2012. Each dot represents a general acute care hospital in Massachusetts. Contours represent Massachusetts counties.All hospitals are included in both insurers' tiered networks, but hospitals' tiers are not necessarily consistent across insurers.

Table 11: Hospital characteristics by tier, 2010-2014

	% of All Hospitals	Beds (tier means)	% of System Hospitals	% of AMCs	% of Boston HRR Hospitals	% of Non-Boston HRR Hospitals
Tier 1	51.6	240.9	41.1	32.5	54.7	44.1
Tier 2	23.9	286.7	22.2	30.8	22.5	26.8
Tier 3	24.5	318.2	36.7	36.8	22.8	29.1
Count	61.0	53.0	31.0	14.0	41.0	20.0

Hospital characteristics weighted by tier frequency across insurers and years. Final row reports hospital counts. Hospitals in the least preferred tier (tier 3) are larger and have a higher proportion of academic medical centers (AMCs). Hospitals both in and outside of Boston are present in all three tiers.

Figure 5: Screenshots from Harvard Pilgrim Independence plan documentation

Participating hospitals and their tiers

Massachusetts

Hospital	**Tier**
Addison Gilbert Hospital	2
Anna Jaques Hospital	1
Athol Memorial Hospital	2
Baystate Franklin Medical Center	3
Baystate Mary Lane Hospital	2

(a) Hospital tier assignments

Inpatient Hospital Care: Medical	
Tier 1	$250
Tier 2	$500
Tier 3	$750

(b) Out-of-pocket prices for each tier

Screenshots from the documentation for the highest-enrollment tiered-network plan in the data (Harvard Pilgrim Independence). Figure 5a shows the tier assignments of the first five hospitals, in alphabetical order, taken from Harvard Pilgrim's documentation. Figure 5b shows the copays associated with each tier, taken from the GIC's benefits description (the same information is also available in a slightly different form in Harvard Pilgrim's documentation). Screenshot margins have been modified for figure fit.

Table 12: Hospital choice model: additional coefficients

	(1) No FEs		(2) Hospital FEs	
Hospital Choice				
Copay ($1,000s)	0.8169***	(0.0549)	-0.1833**	(0.0690)
Copay × std. income	0.1429**	(0.0524)	0.1904***	(0.0540)
Distance (mi)	-0.1817***	(0.0026)	-0.1832***	(0.0028)
Distance^2	0.0005***	(0.0000)	0.0006***	(0.0000)
Distance (mi) × Boston	-0.0819***	(0.0042)	-0.0750***	(0.0053)
Distance^2 × Boston	0.0007***	(0.0000)	0.0008***	(0.0000)
Past use of hospital	4.1221***	(0.0473)	3.8292***	(0.0498)
Age × distance	0.0001	(0.0000)	0.0000	(0.0000)
Male × distance	0.0038*	(0.0017)	0.0026	(0.0016)
Chronic cond × distance	0.0221***	(0.0015)	0.0216***	(0.0015)
Teaching × distance	0.0150***	(0.0015)	-0.0032*	(0.0016)
Beds × distance	0.0000***	(0.0000)	0.0000***	(0.0000)
Satellite hosp campus	-0.2835***	(0.0275)	2.0478***	(0.1798)
Cardiac CCS × cath lab	1.0601***	(0.1061)	0.5355***	(0.1092)
Obstetric CCS × NICU	0.7996***	(0.0340)	0.3262***	(0.0393)
Nerv, circ, musc CCS × MRI	0.0437	(0.0601)	-0.0312	(0.0778)
Nerv CCS × neuro	1.6520***	(0.2742)	0.0814	(0.3218)
% good pain control × distance	-0.0022**	(0.0009)	-0.0049***	(0.0008)
% highly recommend × distance	0.0107***	(0.0006)	0.0046***	(0.0007)
Hospital FEs	No		Yes	
Pseudo R^2	0.558		0.605	
Nadmits	29917		29917	

Multinomial logit model of hospital choice. All copay coefficients scaled to $1,000s for ease of interpretation. Consumers dislike distance and high out-of-pocket prices (copays). Hospital quality is standardized and hospital fixed effects are included. Standard errors in parentheses, clustered by patient. Nadmits = number of choice sets (admissions).

Table 13: Hospital choice model (with control function)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Pref. spec	+IV sample	IV deg1	IV deg2	IV deg3	IV deg4	IV deg5
Hospital Choice							
Copay (\$1,000s)	-0.1833**	-0.0730	-0.2994*	-0.3065*	-0.4790***	-0.5161***	-0.6208***
	(0.0690)	(0.0881)	(0.1230)	(0.1255)	(0.1288)	(0.1291)	(0.1302)
Copay $\times$ std. income	0.1904***	0.1965**	0.2060***	0.2079***	0.1927**	0.1910**	0.1966**
	(0.0540)	(0.0641)	(0.0599)	(0.0613)	(0.0625)	(0.0626)	(0.0627)
IVresid_zeroed_1			0.0004*	0.0004*	0.0011***	0.0012***	0.0018***
			(0.0001)	(0.0001)	(0.0002)	(0.0002)	(0.0003)
IVresid_zeroed_2				0.0000	-0.0000*	-0.0000	-0.0000***
				(0.0000)	(0.0000)	(0.0000)	(0.0000)
IVresid_zeroed_3					-0.0000***	-0.0000***	-0.0000***
					(0.0000)	(0.0000)	(0.0000)
IVresid_zeroed_4						-0.0000	0.0000***
						(0.0000)	(0.0000)
IVresid_zeroed_5							0.0000***
							(0.0000)
Distance (mi)	-0.1832***	-0.1813***	-0.1815***	-0.1815***	-0.1821***	-0.1822***	-0.1830***
	(0.0028)	(0.0031)	(0.0020)	(0.0020)	(0.0020)	(0.0020)	(0.0020)
Distance^2	0.0006***	0.0006***	0.0006***	0.0006***	0.0006***	0.0006***	0.0006***
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)
Distance (mi) $\times$ Boston	-0.0750***	-0.0762***	-0.0761***	-0.0761***	-0.0759***	-0.0759***	-0.0758***
	(0.0053)	(0.0057)	(0.0046)	(0.0046)	(0.0046)	(0.0046)	(0.0046)
Distance^2 $\times$ Boston	0.0008***	0.0008***	0.0008***	0.0008***	0.0008***	0.0008***	0.0008***
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)
Past use of hospital	3.8292***	3.7617***	3.7522***	3.7519***	3.7404***	3.7401***	3.7263***
	(0.0498)	(0.0529)	(0.0434)	(0.0435)	(0.0438)	(0.0438)	(0.0441)
Hospital FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Pseudo R^2	0.605	0.611	0.611	0.611	0.611	0.611	0.611
Nadmits	29917	26319	26319	26319	26319	26319	26319

Nadmits = number of choice sets (admissions). All specifications estimated using multinomial logit.
Standard errors in parentheses, clustered by patient. IV columns estimated using a control function with boostrapped standard errors with replications.

Table 14: Specification checks for hospital choice model

	(1) Select. on α	(2) High income	(3) New enrolt.
Hospital Choice			
Copay ($1,000s)	-0.2107**	-0.4086	-0.3123*
	(0.0731)	(0.2393)	(0.1371)
Copay $\times$ std. income	0.1942***		0.1595
	(0.0542)		(0.1181)
Copay $\times$ selected non-tiered	0.1118		
	(0.1100)		
Distance (mi)	-0.1832***	-0.1728***	-0.2044***
	(0.0028)	(0.0168)	(0.0042)
Distance^2	0.0006***	0.0006***	0.0006***
	(0.0000)	(0.0001)	(0.0000)
Hospital FEs	Yes	Yes	Yes
Pseudo R^2	0.605	0.594	0.547
Nadmits	29917	1790	6324

Multinomial logit model of hospital choice. All copay coefficients scaled to $1,000s. Column (1) tests whether consumers select into tiered-network plans by price sensitivity. Column (2) tests whether high-income consumers (1.5 or more standard deviations above the mean) have a positive price coefficient. Column (3) tests whether consumers have a positive price coefficient when they first enroll in a tiered-network plan. Standard errors in parentheses, clustered by patient. Nadmits = number of choice sets (admissions).

Table 15: Own-price elasticities from hospital demand model (at median household income)

Hospitals	At observed copays	At $1,000 copays
Mean across hospitals in metro Boston	-0.039 (0.002)	-0.117 (0.006)
Mean across hospitals outside Boston	-0.030 (0.002)	-0.092 (0.005)
Metro Boston hospitals		
Beth Israel Deaconess Hospital - Milton	-0.035 (0.002)	-0.117 (0.006)
Beth Israel Deaconess Hospital - Needham	-0.037 (0.002)	-0.124 (0.006)
Beth Israel Deaconess Medical Center	-0.040 (0.002)	-0.109 (0.005)
Boston Medical Center	-0.036 (0.002)	-0.120 (0.006)
Brigham and Women's Faulkner Hospital	-0.045 (0.002)	-0.120 (0.006)
Brigham and Women's Hospital	-0.050 (0.003)	-0.106 (0.005)
Cambridge Health Alliance - Cambridge Campus	-0.031 (0.002)	-0.123 (0.006)
Cambridge Health Alliance - Somerville Campus	-0.031 (0.002)	-0.124 (0.006)
Cambridge Health Alliance - Whidden Campus	-0.031 (0.002)	-0.124 (0.006)
Lawrence Memorial Hospital	-0.040 (0.002)	-0.119 (0.006)
Massachusetts General Hospital	-0.053 (0.003)	-0.112 (0.006)
Melrose-Wakefield Hospital	-0.042 (0.002)	-0.114 (0.006)
Mount Auburn Hospital	-0.038 (0.002)	-0.114 (0.006)
Newton-Wellesley Hospital	-0.037 (0.002)	-0.110 (0.005)
Steward Carney Hospital	-0.035 (0.002)	-0.122 (0.006)
Steward St. Elizabeth's Medical Center	-0.045 (0.002)	-0.120 (0.006)
Tufts Medical Center	-0.041 (0.002)	-0.115 (0.006)

Own-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. Standard errors in parentheses, calculated using 100 bootstrap replications.

Table 16: Cross-price elasticities from hospital demand model for select hospitals (at median household income)

	Brigham	MGH	Beth Israel	BMC	Cape Cod	Baystate	Cooley
Brigham and Women's Hospital	–	0.0052	0.0053	0.0011	0.0002	0.0001	0.0001
		(0.0003)	(0.0003)	(0.0001)	(0.0000)	(0.0000)	(0.0000)
Massachusetts General Hospital	0.0067	–	0.0044	0.0009	0.0003	0.0002	0.0003
	(0.0003)		(0.0002)	(0.0000)	(0.0000)	(0.0000)	(0.0000)
Beth Israel Deaconess Medical Center	0.0084	0.0053	–	0.0011	0.0001	0.0001	0.0001
	(0.0004)	(0.0003)		(0.0001)	(0.0000)	(0.0000)	(0.0000)
Boston Medical Center	0.0083	0.0055	0.0056	–	0.0001	0.0001	0.0001
	(0.0004)	(0.0003)	(0.0003)		(0.0000)	(0.0000)	(0.0000)
Cape Cod Hospital	0.0024	0.0033	0.0009	0.0002	–	0.0002	0.0001
	(0.0001)	(0.0002)	(0.0000)	(0.0000)		(0.0000)	(0.0000)
Baystate Medical Center	0.0003	0.0005	0.0001	0.0000	0.0000	–	0.0057
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)		(0.0003)
Cooley Dickinson Hospital	0.0005	0.0009	0.0002	0.0000	0.0000	0.0098	–
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0005)	

Coss-price elasticities of demand for row hospitals with respect to out-of-pocket price for column hospitals, calculated at the hospitals' observed copays. The first four (Brigham, MGH, Beth Israel, and BMC) are the key academic medical centers in Boston and are each other's closest substitutes. Cape Cod is geographically isolated in the eastern Massachussets and sends few patients to other hospitals. Baystate and Cooley are in western Massachusetts and compete with each other. All hospitals, even those outside Boston, are affected by prices at the flagship hospitals of the "star" Partners HealthCare system, Brigham and MGH. Standard errors in parentheses, calculated using 100 bootstrap replications.

Table 17: Enrollment in GIC plans

Plan	Share (%)	New policies	New enrollees	2009-2012 enrolt.
Fallon Direct	1.52	891	1,543	7,177
Fallon Select	3.78	1,286	2,684	11,167
Harvard Pilgrim Independence	36.42	16,358	36,444	96,103
Harvard Pilgrim Primary Choice	3.04	2,079	4,472	22,208
Health New England	9.54	3,443	6,451	29,312
Neighborhood Health Plan	1.71	924	1,645	7,552
Tufts Navigator	41.97	10,137	20,438	120,519
Tufts Spirit	1.16	1,228	2,577	13,775
UniCare Basic				
UniCare Community Choice				
UniCare PLUS				

GIC plan enrollment for employees and their dependents, excluding UniCare plans.
Share is market share is at the end of fiscal year 2011 (June 2011).
Enrollee and policy holder counts are for first-time GIC enrollees in 2009–June 2011.
Final column is total number of unique enrollees in 2009–2012.

Table 18: Hospital sorting counterfactuals (at median household income)

	Flat copay $250	5% coins.	10% coins.	20% coins.
Tier 1 hospitals % of volume	27.15	27.92	28.79	30.55
Tier 2 hospitals % of volume	36.65	36.89	37.03	37.22
Tier 3 hospitals % of volume	36.21	35.19	34.18	32.23
Patient spending per admission ($)	250	943	1,861	3,621
Change in patient $ over flat copay	–	277.27%	644.26%	1,348.43%
Insurer spending per admission ($)	18,564	17,920	16,746	14,484
Change in insurer $ over flat copay	–	-3.47%	-9.79%	-21.98%
Total spending per admission ($)	18,810	18,864	18,606	18,105
Change in total $ over flat copay	–	0.29%	-1.08%	-3.74%

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. Columns 2-4 set the "copays" equal to percentages of the total negotiated price to mimic coinsurance (assuming fully informed consumers).

Figure 6: Compensating consumers for welfare losses from tiered networks, double-counting out-of-pocket spending

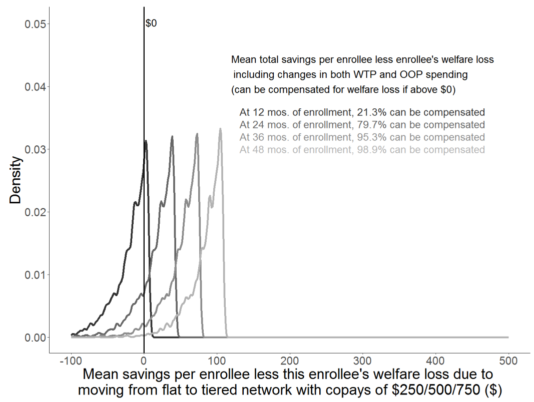

Mean savings per enrollee less each enrollee's welfare loss due to moving from a flat $250 copay to the Harvard Pilgrim Independence hospital network with its observed copay regime of $250, $500, and $750 for hospitals in tiers 1, 2, and 3, respectively. Hospital shares are calculated using copays from the learning model at each listed enrollment duration; utilities and WTP are calculated based on those shares and the copay coefficient at 36 months of enrollment.

Appendix B: Data Preparation Details

Aggregation to the claim level: Like other medical claims databases, the unit of observation in the APCD is the claim line, which is the smallest unit of service for which an insurer or patient is billed separately from other units of service. A single hospital visit, for example, can have many claim lines for drugs, operating room supplies, anesthesia, and physician fees. In the analysis, I aggregate information across claim lines to the level of the hospital admission.

Identifying hospitals: The APCD includes provider identifiers as reported by insurers. An insurer typically uses multiple provider codes for each hospital, corresponding to various departments, facilities, or physician or nurse groups. I build a crosswalk between insurer-reported identifiers and a master list of hospitals using fuzzy matching on hospital names and addresses reported in the APCD. In addition, I conduct a final round of manual checks to correct errors and exclude mistakenly attributed onsite facilities or physician groups that are not involved in inpatient care. Insurers report identifiers for the service provider (where the patient is treated) and the billing provider (the business entity that submits the claim to the insurer); I use the service provider identifiers.

Identifying tiered-network plans: The APCD includes insurer identifiers that can be mapped to names of Massachusetts insurers. In addition, it includes a variable indicating whether an insurance plan is a GIC plan. This indicator is defined correctly for most GIC plans, but fails to label as GIC-affiliated some plans whose characteristics and enrollment reveal them to be GIC plans. I label as GIC plans any plans offered by GIC insurers that match known GIC plan characteristics (networks and copays), share a large number of other GIC plans' enrollees across years, and have enrollment totals that match enrollments from the GIC. This procedure is sufficient to identify all GIC plans in the APCD and matches GIC enrollments within a small margin of error from the GIC's annual reports. For plans offered by HPHC and Tufts outside the GIC, I label as tiered any plan whose observed hospital copays are in round dollar amounts and match hospital tiers in those insurers' contemporaneous tiered networks.

Diagnosis classification: In the data, diagnoses and procedures are reported in the International Classification of Diseases, Clinical Modification (ICD-9) classification system, which consists of approximately 14,000 distinct diagnosis codes and 4,000 procedure codes. For each claim, the principal diagnosis is reported along with up to twelve secondary diagnoses. Similarly, for visits involving procedures, a principal procedure code is reported along with up to six secondary procedures. I assign diagnoses to diagnostic categories and severity levels using the Clinical Classifications Software (CCS) categorizations from the Agency of Healthcare Research and Quality. The CCS classification system assigns diagnosis codes to approximately 300 mutually exclusive diagnosis groups, which are further aggregated into eighteen broad diagnostic categories. The CCS diagnostic categories are described and their prevalence in the Massachusetts nonelderly population given in Appendix Table 9.

Defining copays for unobserved alternatives: The APCD includes information on the out-of-pocket payment from the patient to the hospital. This information is sufficient to identify the copays for hospitals that consumers are observed to choose. To construct the copay vector for the other hospitals in a consumer's choice set, I assign each hospital to its corresponding tier in that insurer-year network. The copay corresponding to the tier is drawn from plan documentation (for GIC plans) or imputed tier copays (for plans outside the GIC).

Calculating distance: The APCD reports five-digit zip codes for patient home address. I geocode the patient zip codes and use them to calculate the driving distance from the centroid of the zip code to the hospital's full address. Driving distances are calculated using Bing Maps driving directions.

Appendix C: Additional Background

History of Tiered Provider Networks

Plans with tiered provider networks were introduced in the early 2000s, as insurers sought new mechanisms for bolstering their bargaining power with respect to increasingly consolidated providers (Robinson 2003; Sinaiko 2012). Tiered networks allowed insurers to maintain some of the bargaining leverage associated with health maintenance organizations (HMOs), which used the threat of contract termination to drive down negotiated prices but which experienced a backlash of public opinion in the 1990s (Cutler et al. 2000; Town and Vistnes 2001; Ho 2009). Detractors argued that HMOs' savings came at the expense of patient choice, access to care, and continuity of care (Martin 2014).

Tiered provider networks combine the cost control mechanisms of narrow networks with patient choice and explicit price information for consumers. In a tiered network, almost all providers in the market remain in the consumer's choice set, but a higher out-of-pocket price is associated with the use of higher-priced providers. Providers are placed into non-overlapping groups, or *tiers*, that determine consumers' out-of-pocket prices for treatment. The out-of-pocket price faced by enrollees is then constant among providers within a tier, but varies across tiers. Throughout the paper, I distinguish between the out-of-pocket price faced by insured consumers and the full price negotiated between providers and insurers, which I call simply "price".

Advocates of tiered networks argue that they reduce health care spending through two mechanisms: the direct effect of steering consumers toward lower-priced providers (Sinaiko 2012), and an indirect effect on prices (Fronstin 2003; Robinson 2003). If consumers indeed respond to the incentives in tiered provider networks, then non-preferred tier placement becomes an additional bargaining lever that insurers can use in price negotiations with providers. The usefulness of tiered networks as a bargaining chip for insurers therefore hinges on consumer responsiveness to out-of-pocket prices in tiered networks.

Since their introduction in the early 2000s, the penetration of tiered-network plan designs has continued to rise. Health care system experts, insurers and employers increasingly see the use of tiered networks and other value-based plan designs as integral to cost control (Robinson 2003; KFF 2014). Among the highest-enrollment health plans offered by very large employers, 38% of the highest-enrollment health plans now include a tiered provider network, with 54% of all employers expecting tiered networks to be a very effective or somewhat effective measure for health care cost reduction (KFF 2014, 2015, 2016). On the Affordable Care Act insurance marketplaces, 5–6% of plans each year have used tiered networks (McKinsey 2016), yet these plans accounted for a disproportionate 14% of enrollment in 2015 (HHS 2016). Multiple states expect growth in tiered-network plans (KFF 2014; McKinsey 2015; KFF 2015); moreover, some states have been directly involved in promoting the adoption of tiered provider networks.

The Massachusetts Health Care Market

In 2006, Massachusetts passed a landmark health care overhaul which aimed to expand health insurance coverage and access to care. The Massachusetts reform subsequently served as the blueprint for the federal Patient Protection and Affordable Care Act (ACA) passed in 2010. Although the 2006 legislation succeeded in broadening insurance coverage in Massachusetts, policymakers remained concerned about the state's high overall health care spending. Not only was the state's per capita health care spending 15% higher than the national average, driven largely by high hospital spending, it had also grown faster than national health care spending since 2002 (DHCFP 2010). Based on recommendations by the Massachusetts Division of Health Care Finance and Policy, the state implemented additional reforms aimed at measuring and reducing health care spending in 2010 and again in 2012 (Massachusetts 2010, 2012a; Wrobel et al. 2014; CHIA 2015). These reforms included, among other provisions,[40] the creation of the All-Payer Claims Database used in this paper and requirements for insurers to offer value-based insurance designs (DHCFP 2010).

Since 2011, Massachusetts legislation has required all large insurers to offer at least one narrow- or tiered-network plan in at least one geographic area (Massachusetts 2010). The regulation does not require insurers to offer tiered-network plans; they may instead offer narrow-network plans. However, all three of the state's largest insurers—Blue Cross Blue Shield of Massachusetts, Harvard Pilgrim Health Care, and Tufts Health Plan—have offered both tiered- and narrow-network plans since before the regulation went into effect in 2011. These insurers now have 10–35% of their commercial enrollees in tiered-network plans. State regulation also outlines a method for insurers to calculate comparable prices across providers by adjusting for disease and patient mix; insurers are required to report these prices to the state's Center for Health Information and Analysis (CHIA) and are expected to use them for determining providers' network status.

Outside of state legislation, the push toward tiered networks in Massachusetts has been led by the Massachusetts Group Insurance Commission (GIC), which administers health insurance and other benefits for state and municipal employees, retirees, and their dependents.[41] The GIC insures some 300,000–350,000 individ-

[40] Other notable pieces of the legislation consisted of health care price transparency requirements and the encouragement of vertical integration between providers in the form accountable care organizations (created under the moniker "Alternative Quality Contract" (Song et al. 2012)).

[41] This is the same employer group studied by Gruber and McKnight (2014) in evaluating the impact of narrow networks and by Sinaiko and Rosenthal (2014) in studying patient response to physician tiering.

uals per year throughout my sample period, corresponding to approximately 8% of the total commercially insured population in Massachusetts. The volume of covered lives on the GIC, along with the substantial fraction of the state budget devoted to it, makes the GIC an important and active player in the Massachusetts health insurance landscape (DHCFP 2010; Wrobel et al. 2014). The GIC was among the earliest adopters of tiered provider networks, introducing its first tiered hospital network plan in July 2003 and rolling out tiered physician networks in July 2006 (GIC 2008, 2009).

Massachusetts requires insurers operating tiered-network plans to "clearly and conspicuously indicate" consumers' out-of-pocket prices for each tier (Massachusetts 2012b). Insurers provide this information to enrollees as part of the schedule of benefits documentation for each plan. At the insurer level, they also publish lists of hospitals and their network tiers each year, which can be easily accessed through their websites for the current year. These lists include each hospital's tier, so consumers do not need to search for multiple providers' network status in order to comparison-shop. This is in contrast to the difficulty of learning out-of-pocket prices for hospital care in advance in traditional plan types: even savvy consumers who ask for price quotes typically get poor response rates (Bebinger 2014).

The Massachusetts Group Insurance Commission

The Massachusetts Group Insurance Commission (GIC) is the benefits administrator for the state of Massachusetts, some municipalities, and a number of other public entities. It insures some 300,000–350,000 people per year during my sample period, consisting of GIC-covered employees, retirees, and their dependents. My sample of GIC enrollees observed in the APCD includes approximately 90,000 state and municipal employees and 120,000 dependents. The remaining individuals insured through the GIC are retired government employees and their surviving spouses. The demographic characteristics for the GIC enrollees in my sample are shown in Table 19. Approximately 60% of primary enrollees insure their dependents as well. The majority of the primary enrollees live in the Boston area or elsewhere in eastern Massachusetts. Approximately half of the enrollees are first observed in the GIC prior to the start of the medical claims data in 2009. The remaining individuals insured through the GIC are retired government employees and their surviving spouses.

The demographic characteristics for the GIC enrollees in my sample are shown in Table 19. Approximately 60% of primary enrollees insure their dependents as well. The majority of the primary enrollees live in the Boston area or elsewhere in eastern Massachusetts. Approximately half of the enrollees are first observed in the GIC prior to the start of the medical claims data in 2009.

Table 19: Characteristics of GIC health insurance enrollees

	Individuals	Families
% of households	39.5	60.5
% of total enrollment	17.8	82.2
Median family size	1	3
Mean family size	1	3.2
% female	59.5	50.3
Mean age	48.1	35.7
Median age	49	39
% entering before 2009	47.3	56.2
% Western Mass.	19.8	18.2
% Central Mass.	12.2	13.1
% Northeast Mass.	28.1	29.4
% Metro Boston	25.4	20
% Southeast Mass.	14.6	19.3

Summary statistics for Massachusetts Group Insurance Commission (GIC) health insurance enrollees. Column 1 is single enrollees; column 2 is enrollees with dependents. 60% of enrolled households include dependents, who are typically younger than primary enrollees. Approximately half of households are enrolled in the GIC prior to the start of the data in 2009.

I use data on the GIC's health plan offerings, premiums, and plan characteristics such as deductibles for GIC fiscal years 2009–2011, which cover the calendar period July 2008–June 2012.[42] The plan offerings and their premiums for a sample enrollment year are described in Table 1. The employee portion of premium contributions is 25% of the total premium.[43] Two levels of premiums are set for each plan: one for individual coverage and another for family coverage (defined as two or more enrollees), with no variation in these two premium amounts across the entire state for each fiscal year. Plan characteristics, such as out-of-pocket prices and hospital networks, change over time. Plans on the GIC use copays, which are fixed dollar amounts paid out-of-pocket by consumers when they use health care. 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/$750 across the three hospital tiers in 2010, and increase to $275/$500/$1,500 in 2016.

Plans on the GIC market are fairly standardized: deductible levels, prescription drug copays, and some other plan characteristics vary little or not at all across plans within a fiscal year. This type of standardization is found in many health insurance markets, including Medigap, state health insurance exchanges, and large employers (Starc 2014; Ericson and Starc 2015; Handel 2013). Such markets can shed light on plan competition on the health insurance exchanges set up under the Affordable Care Act. The primary differences between plans on the GIC come from the insurer brands, provider networks, and copay structures for physician and hospital care.

Appendix D: Willingness-to-Pay for Hospital Networks

Patients will value more highly those hospital network arrangements that set low out-of-pocket prices c_{mh} for nearby, high-quality, or otherwise desirable hospitals. Consumer valuation of a hospital network is measured by willingness-to-pay (WTP). An individual consumer's ex ante dollarized valuation of plan m's tiered hospital network is the expected utility of seeking care at various hospitals at the

[42] Data from July 2012 onward excluded because the GIC implemented a premium discount program that affected employees differently depending on characteristics I do not observe in the APCD (Gruber and McKnight 2014). The plan demand analysis therefore relies on GIC data through June 2012.

[43] Employees hired prior to July 2003 only pay 20% of the total premium cost. In the analyses, I therefore exclude GIC enrollees who were enrolled prior to 2007 (the earliest enrollment data in the APCD) in order to reduce noise in premium measurement.

out-of-pocket prices dictated by the tiers in m:

$$W_{mi} = \frac{1}{\alpha_i} \sum_{d \in D} f_{id} \ln \left(\sum_{h \in H} \exp\left(-\alpha_i c_{mh} + \beta x_{hid}\right) \right). \tag{3}$$

This expression is the familiar log-sum equation for expected consumer surplus for a logit model, modified in that an additional expectation is taken over the probability of consuming any care, expressed in f_{id}. This modification gives rise to the willingness-to-pay for a hospital network as defined in Capps et al. (2003), here with the additional complication that networks can vary in out-of-pocket prices across hospitals. The availability of a direct estimate of the price responsiveness parameter α allows the WTP to be expressed in dollars, rather than in utils as is the case in settings that lack out-of-pocket price variation.

The calculation of WTP requires each consumer's ex ante distribution of diagnosis probabilities f_{id} for the upcoming year. I calculate these probabilities separately for each sex–10-year age band cell and each CCS diagnostic category using data on all non-transfer hospital admissions of Massachusetts residents from the 2010 HCUP State Inpatient Database.[44] Since patient covariates such as distance to hospitals also vary across zip codes, WTP for a given hospital network takes on a separate value for each gender-age group-zip code triplet. The geographic variation in WTP is driven by the fact that some consumers are geographically closer to a larger number of hospitals or more desirable hospitals. Allowing for this granular variation in consumers' preferences and admission probabilities at the diagnostic category level allows the WTP measure to capture rich variation across consumers.

[44] This is equivalent to the assumption that that a consumer's expectation of her health status for the upcoming year is a consistent predictor of her health status, given only her sex, her 10-year age group, and the fact of residing in Massachusetts. This assumption is more likely to hold for relatively healthy consumers who do not have highly informative personal experience to inform their ex ante expectations of diagnosis (Shepard 2014). Since my data consist of non-elderly, commercially insured, mostly employed individuals, they are healthier than the general population and good candidates for the assumption that their expected health status is approximately equal to the average health status for their age group. To the extent that there are deviations from the average health status, they will load onto the error term in the plan choice model.

Appendix E: Learning Model Framework

Framework for Consumer Learning

A consumer who falls sick with a diagnosis that requires inpatient hospital treatment must choose a hospital at which to receive treatment. In a tiered-network plan, the consumer has low-cost access to information about the out-of-pocket prices she would face for receiving medical care at any hospital in the market. In traditional plans without a tiered network, including high-deductible health plans, determining out-of-pocket prices ex ante is costly and often impossible, even in the presence of a price look-up tool. Many health care conditions necessitate complicated, multi-part episodes of care for which consumers must add up a vector of prices to come up with a total for the treatment, such as separate fees for the surgeon, the operating room fee, prescription drugs, and anesthesia. In the context of a tiered-network plan, the consumer faces a single fixed search cost for obtaining information about out-of-pocket prices at all the hospitals in her plan's network. The plan provides a single document that lists the tiers associated with all the hospitals in the network, so consumers need not sequentially search for the out-of-pocket price of each hospital in order to comparison-shop.

Consumer i enrolled in plan m who becomes sick with diagnosis d has utility $u_{mhid} = -\alpha_i c_{mh} + \beta x_{hid} + \varepsilon_{mhid}$ from receiving care at hospital h. An informed consumer chooses a hospital h' such that $h' = \arg\max_h \{u_{mhid}\}$. A consumer who does not learn out-of-pocket prices would instead choose a hospital h'' that solves $h'' = \arg\max_h \{v_{mhid}\}$, where $v_{mhid} = \beta x_{hid} + \varepsilon_{mhid}$. The ex post utility of choosing hospital h'' will be $u_{mh''id} = -\alpha_i c_{mh''} + \beta x_{h''id} + \varepsilon_{mh''id}$, even if the price is ignored at the ex ante point of choice. If the expected gain in realized utility from a choice with known out-of-pocket prices relative to a choice ignoring prices exceeds her search cost of obtaining price information, the consumer will pay the search cost.

That is, the consumer will engage in search for out-of-pocket price information if

$$\begin{aligned}\kappa_i \leq & \mathbb{E}\left[\max_{h'}\left\{-\alpha_i c_{mh'} + \beta x_{h'id} + \varepsilon_{mh'id}\right\}\right. \\ & \left. - \left\{-\alpha_i c_{mh''} + \beta x_{h''id} + \varepsilon_{mh''id} | h'' = \arg\max_h \left\{\beta x_{hid} + \varepsilon_{mhid}\right\}\right\}\right]\end{aligned} \tag{4}$$

where κ_i is the consumer's search cost for obtaining price information.

If expected out-of-pocket prices are constant across all hospitals in a plan, the expected gain in utility from searching is zero because maximizing v_{mhid} is equivalent to maximizing u_{mhid}. Thus, a consumer who is completely unaware of the tiered structure of her plan will not search for prices. This behavior is observationally equivalent to that of a consumer who is aware of out-of-pocket price differences but has perfectly inelastic demand with $\alpha_i = 0$. The presence of consumers in an estimation sample who are unaware of differential pricing will therefore bias the price coefficient toward zero.

Consumers who are aware of the differential out-of-pocket prices in their plan may still decide not to search for prices. To see this, consider a consumer who has identical uniform priors over out-of-pocket prices for all hospitals, $c_h \sim U[\underline{c}, \overline{c}]$. Without searching, the consumer's expected realized utility is $-\alpha_i \frac{\underline{c}+\overline{c}}{2} + \beta x_{h''id} + \varepsilon_{mh''id}$, where $h'' = \arg\max_h \left\{\beta x_{hid} + \varepsilon_{mhid}\right\}$. Therefore, the consumer will search if

$$\begin{aligned}\kappa_i \leq & \mathbb{E}\left[\max_{h'}\left\{-\alpha_i c_{mh'} + \beta x_{h'id} + \varepsilon_{mh'id}\right\}\right] \\ & + \alpha_i \frac{\underline{c}+\overline{c}}{2} - \left\{\beta x_{h''id} + \varepsilon_{mh''id} | h'' = \arg\max_h \left\{\beta x_{hid} + \varepsilon_{mhid}\right\}\right\}\end{aligned}$$

The consumer is less likely to search the smaller is her price sensitivity α_i, the greater is the difference between $\max_h \left\{\beta x_{hid} + \varepsilon_{mhid}\right\}$ and the next-best option, and the tighter is the distribution of the prior over prices.

A consumer with a large gap in utility between $\max_h \left\{\beta x_{hid} + \varepsilon_{mhid}\right\}$ and the next-best option has less to gain from learning prices, since the expected utility gain from paying a lower price must exceed a large loss in utility from substituting away from the most preferred hospital. A strong preference for the most preferred hospital may be the result of several factors, such as other hospitals being substantially farther from the consumer's home or a strong established relationship with a given hospital. The latter is especially likely if the consumer has a chronic condition, since this implies more regular treatment and more complex disease management. Similarly, a consumer who simply has a low price sensitivity α_i has less to gain from searching since she is less likely to change her hospital selection as a result of learning prices. In such a case, low propensity to search is observationally equivalent to a true low price sensitivity. A consumer may have low price sensitivity if, for example, she is high-income and therefore has a relatively low marginal utility of money.

The higher the variance of the distribution over priors, the more the consumer expects to gain from searching. Consumers' beliefs about the structure of their plan's out-of-pocket pricing may change as a result of their own experiences with the plan or as a result external factors.[45] For example, consumers may learn via benefits information campaigns or through workplace word of mouth that out-of-pocket prices vary across hospitals. They may also Bayesian update their beliefs about the distribution of out-of-pocket prices as a result of consuming medical care and observing different out-of-pocket prices across their own care.

I model the learning process by which consumers update their beliefs about the distribution of out-of-pocket prices across hospitals. The out-of-pocket price for a given hospital h in plan m can take on one of a discrete set of values for possible copays in that plan, indexed by $c_k \in \{c_1, \ldots, c_K\}$ with the addition of a zero copay $c_0 = 0$. The copays are modeled as being distributed according to the discrete categorical distribution. The probability mass function for the categorical distribution with parameters $\vec{\phi}$ is simply $f(c_h = c_k|\vec{\phi}) = \phi_k$, where $\vec{\phi} = (\phi_0, \phi_1, \ldots, \phi_K)$ is the vector of probabilities that the copay is equal to the kth value of the possible set of copays. Consumers start with prior beliefs about $\vec{\phi}$ that are distributed according to a Dirichlet distribution, $\vec{\phi}|\vec{a} \sim Dir(K, \vec{a})$, where $\vec{a} = (a_0, a_1, ..., a_K), a_k > 0 \forall k$ is the concentration hyperparameter and $K = (\gamma_0, \gamma_1, ..., \gamma_K)$ is (with slight abuse of notation) a vector with $\gamma_k \in (0, 1)$ whose elements sum to one. The probability mass

[45] A third possibility is that referring physicians learn about the structure of their patients' plan and shift their referral patterns toward hospitals with lower out-of-pocket prices.

function of the Dirichlet is given by

$$f(K,\vec{a}) = \frac{1}{B(\vec{a})}\prod_{k=0}^{K}\gamma_k^{a_k-1}$$

where $B(\vec{a})^{-1}$ is the Beta function and acts as the normalizing constant.

Each time a consumer receives medical care, she observes ex post the out-of-pocket price associated with that care. Let y_k denote the number of copays of value c_k that are observed by the consumer. These observations are drawn from the categorical distribution $\vec{y}|\vec{\phi} \sim Cat(K,\vec{\phi})$. The consumer incorporates observed data on copays using Bayes' rule to arrive at an updated posterior distribution. Since the Dirichlet is the conjugate prior to the categorical distribution, the posterior distribution of beliefs after updating are also distributed according to the Dirichlet, now with the parameters

$$\vec{\phi}|\vec{y} \sim Dir(K,\vec{a}+\vec{y}) = Dir(K,a_0+y_0,a_1+y_1,\ldots,a_K+y_K) \tag{5}$$

The larger is $a_k + y_k$ relative to other $k' \neq k$, the more probable the consumer considers copay c_k to be. Equation 5 describes the process by which consumers update their beliefs about the distribution of copays. Updating that occurs in response to consumers' own health care experiences enters through changes in $\vec{y}$, while changes due to third-party information enter through changes in $\vec{a}$.

To see the effect of updating based on own health care experiences, consider a consumer who starts with a prior that places arbitrarily large weight on a copay value of zero. Consumers face zero copays in many cases, such as for preventive health care services, 30-day readmissions, or after exceeding their out-of-pocket maximum spending for the year. Such a consumer's prior beliefs are parameterized as $\vec{\phi}|\vec{a} \sim Dir(K,A-\sum_{k=1}^{K}\varepsilon_k,\varepsilon_1,\ldots,\varepsilon_K)$, where $\varepsilon_1,\ldots,\varepsilon_K > 0$ are arbitrarily close to zero and $A \gg 0$ is the sum of the concentration parameters. The consumer's prior expected probability of a copay being equal to zero is arbitrarily close to one, $\mathbb{E}[\phi_0] = \frac{A-\sum_{k=1}^{K}\varepsilon_k}{A} \to 1$. With every count of y_0 realized copays equal to zero and every $y_1,\ldots,y_K$ realized copays that are greater than zero, the updated posterior becomes

$$\vec{\phi}|\vec{y} \sim Dir(K, A - \sum_{k=1}^{K} \varepsilon_k + y_0, \varepsilon_1 + y_1, ..., \varepsilon_K + y_K)$$

which has a posterior expected probability of a copay being equal to zero of

$$\mathbb{E}[\phi_0] = \frac{A - \sum_{k=1}^{K} \varepsilon_k + y_0}{A + \sum_{k=1}^{K} y_k}$$

The posterior expected probability of a zero copay is increasing in y_0, as the greater the number of realized copays that are zero, the tighter the posterior becomes around a copay value of zero. Conversely, it is decreasing in $Y := \sum_{k=1}^{K} y_k$, as the greater the fraction of realized copays that are greater than zero, the higher the posterior probability of non-zero copays. A consumer with a tight posterior distribution around a zero copay is unlikely to engage in search for out-of-pocket price information, since the posterior expected probability that price information will change her hospital selection is small.

More generally, it is useful to characterize the comparative static on the variance of the posterior probability distribution over copays. As discussed earlier, higher variance results in more price search. The variance of the posterior expected copay is given by

$$\mathbb{V}[c_h] = \sum_{k=1}^{K} \frac{\varepsilon_k + y_k}{A + y_0 + Y} c_k^2 - \left[\sum_{k=1}^{K} \frac{\varepsilon_k + y_k}{A + y_0 + Y} c_k \right]^2$$

This variance is decreasing in the observed number of zero copays y_0. It is increasing in the number of observed nonzero copays $Y = \sum_{k=1}^{K} y_k$. Derivations of the variance and comparative statics are provided in the next section. In the data, a Bayesian-updating consumers who learns about the pricing structure in her plan according to Equation 5 will appear less price-sensitive the more zero copays she has encountered to date, and more price-sensitive the more nonzero copays she has encountered. These implications of the learning framework provide an empirical test for whether consumers are learning about out-of-pocket prices as a result of their own health care utilization.

Derivations of Comparative Statics

This section derives the variance and comparative statics for the learning framework. Consider the effects of consumers observing zero versus non-zero copays for their past episodes of care.

The posterior expected probability of a zero copay is

$$\mathbb{E}\left[\phi_0\right] = \frac{A - \sum_{k=1}^{K} \varepsilon_k + y_0}{A + \sum_{k=1}^{K} y_k}$$

which is increasing in y_0 (the greater the number of realized copays that are zero, the tighter the posterior becomes around a copay value of zero); and decreasing in $Y := \sum_{k=1}^{K} y_k$ (the greater the fraction of realized copays that are greater than zero, the higher the posterior probability of non-zero copays that make it make sense to price-shop).

The posterior mean expected copay is

$$\begin{aligned}
\mathbb{E}\left[c_h\right] = \sum_{k=0}^{K} \mathbb{E}\left[\phi_k\right] c_k =& \frac{A - \sum_{k=1}^{K} \varepsilon_k + y_0}{A + y_0 + Y} c_0 + \frac{\varepsilon_1 + y_1}{A + y_0 + Y} c_1 + \ldots + \frac{\varepsilon_K + y_K}{A + y_0 + Y} c_K \\
=& \frac{\varepsilon_1 + y_1}{A + y_0 + Y} c_1 + \ldots + \frac{\varepsilon_K + y_K}{A + y_0 + Y} c_K \\
=& \sum_{k=1}^{K} \frac{\varepsilon_k + y_k}{A + y_0 + Y} c_k
\end{aligned}$$

where the second equality obtains from the fact that $c_0 = 0$.

To find the variance of the posterior expected copay, which determines gains from search:

$$\mathbb{V}\left[c_h\right] = \mathbb{E}\left[c_h^2\right] - \mathbb{E}\left[c_h\right]^2$$

where

$$
\begin{aligned}
\mathbb{E}\left[c_h^2\right] &= \sum_{k=0}^{K} c_k^2 \phi_k \\
&= \frac{A - \sum_{k=1}^{K} \varepsilon_k + y_0}{A + y_0 + Y} c_0^2 + \sum_{k=1}^{K} \frac{\varepsilon_k + y_k}{A + y_0 + Y} c_k^2 \\
&= \sum_{k=1}^{K} \frac{\varepsilon_k + y_k}{A + y_0 + Y} c_k^2
\end{aligned}
$$

so that the total variance is given by

$$
\begin{aligned}
\mathbb{V}\left[c_h\right] &= \sum_{k=1}^{K} \frac{\varepsilon_k + y_k}{A + y_0 + Y} c_k^2 - \left[\sum_{k=1}^{K} \frac{\varepsilon_k + y_k}{A + y_0 + Y} c_k\right]^2 \\
&= \frac{1}{K\left(A + y_0 + Y\right)} \left(\sum_{k=1}^{K} c_k^2 \left(\varepsilon_k + y_k\right)\right) - \frac{1}{\left[K\left(A + y_0 + Y\right)\right]^2} \left(\sum_{k=1}^{K} c_k \left(\varepsilon_k + y_k\right)\right)^2 \\
&= \frac{1}{K\left(A + y_0 + Y\right)} \left[\left(\sum_{k=1}^{K} c_k^2 \left(\varepsilon_k + y_k\right)\right) - \frac{\left(\sum_{k=1}^{K} c_k \left(\varepsilon_k + y_k\right)\right)^2}{K\left(A + y_0 + Y\right)}\right]
\end{aligned}
$$

which is decreasing in y_0 (note that the difference inside the square brackets must be positive, since variance is always positive). Consumers with many draws of zero copays should be less inclined to search, since the variance of copays falls with zero-copay realizations. To determine whether the variance is increasing or decreasing in $y_{k>1}$, we need to know whether $\sum_{k=1}^{K} c_k^2 \left(\varepsilon_k + y_k\right)$ is increasing faster than $-\frac{\left(\sum_{k=1}^{K} c_k(\varepsilon_k+y_k)\right)^2}{K(A+y_0+Y)}$ is decreasing. Without loss of generality, check this for y_1:

$$
\begin{aligned}
\partial\left(-\frac{\left(\sum_{k=1}^{K} c_k\left(\varepsilon_k+y_k\right)\right)^2}{K\left(A+y_0+Y\right)}\right)/\partial y_1 =&\partial\left(-\frac{\left(c_1\left(\varepsilon_1+y_1\right)\right)^2+\left(\sum_{k=2}^{K} c_k\left(\varepsilon_k+y_k\right)\right)^2}{K\left(A+y_0+y_1+\sum_{k=2}^{K} y_k\right)}\right)/\partial y_1 \\
=&\frac{-c_1\left(2y_1+2\varepsilon_1\right)K\left(A+y_0+y_1+\sum_{k=2}^{K} y_k\right)-K\left(-\left(c_1\left(\varepsilon_1+y_1\right)\right)^2\right)}{K^2\left(A+y_0+y_1+\sum_{k=2}^{K} y_k\right)^2} \\
=&\frac{-c_1\left(2y_1+2\varepsilon_1\right)\left(A+y_0+y_1+\sum_{k=2}^{K} y_k\right)+\left(c_1\left(\varepsilon_1+y_1\right)\right)^2}{K\left(A+y_0+y_1+\sum_{k=2}^{K} y_k\right)^2} \\
=&\frac{-c_1\left(2y_1+2\varepsilon_1\right)\left(A+y_0+y_1+\sum_{k=2}^{K} y_k\right)+c_1^2\left(y_1^2+2\varepsilon_1 y_1+\varepsilon_1^2\right)}{K\left(A+y_0+y_1+\sum_{k=2}^{K} y_k\right)^2}
\end{aligned}
$$

which is increasing if and only if

$$
\begin{aligned}
c_1^2\left(y_1^2+2\varepsilon_1 y_1+\varepsilon_1^2\right) >&c_1\left(2y_1+2\varepsilon_1\right)\left(A+y_0+y_1+\sum_{k=2}^{K} y_k\right) \\
c_1^2\left(\varepsilon_1+y_1\right)^2 >&2c_1\left(\varepsilon_1+y_1\right)\left(A+y_0+y_1+\sum_{k=2}^{K} y_k\right) \\
c_1\left(\varepsilon_1+y_1\right) >&2\left(A+y_0+y_1+\sum_{k=2}^{K} y_k\right) \\
c_1\left(\varepsilon_1+y_1-\frac{2}{c_1}y_1\right) >&2\left(A+y_0+\sum_{k=2}^{K} y_k\right) \\
\left(1-\frac{2}{c_1}\right)y_1 >&\frac{2}{c_1}\left(A+y_0+\sum_{k=2}^{K} y_k\right)-\varepsilon_1
\end{aligned}
$$

which will hardly ever happen, because in practice $\left(1-\frac{2}{c_1}\right)$ will almost always be large and negative. So in fact, the term $-\frac{\left(\sum_{k=1}^{K} c_k(\varepsilon_k+y_k)\right)^2}{K(A+y_0+Y)}$ is almost always decreasing in in any given $y_{j\in\{1,\ldots,K\}}$. It remains to check whether this quantity decreasing faster than $\sum_{k=1}^{K} c_k^2 (\varepsilon_k + y_k)$ is increasing. To check this, note that:

$$\partial\left(\sum_{k=1}^{K} c_k^2 (\varepsilon_k + y_k)\right) / \partial y_1 = c_1^2 (2y_1 + 2\varepsilon_1)$$

so the variance is increasing in y_1 iff

$$\begin{aligned}
c_1^2 (2y_1 + 2\varepsilon_1) > & \frac{-c_1 (2y_1 + 2\varepsilon_1) \left(A + y_0 + y_1 + \sum_{k=2}^{K} y_k\right) + c_1^2 \left(y_1^2 + 2\varepsilon_1 y_1 + \varepsilon_1^2\right)}{K \left(A + y_0 + y_1 + \sum_{k=2}^{K} y_k\right)^2} \\
c_1^2 2 (\varepsilon_1 + y_1) > & \frac{-2c_1 (\varepsilon_1 + y_1) \left(A + y_0 + y_1 + \sum_{k=2}^{K} y_k\right) + c_1^2 (\varepsilon_1 + y_1)^2}{K \left(A + y_0 + y_1 + \sum_{k=2}^{K} y_k\right)^2} \\
2c_1 > & \frac{-2 \left(A + y_0 + y_1 + \sum_{k=2}^{K} y_k\right) + c_1 (\varepsilon_1 + y_1)}{K \left(A + y_0 + y_1 + \sum_{k=2}^{K} y_k\right)^2}
\end{aligned}$$

for which a sufficient (but not necessary) condition is that

$$\begin{aligned}
2c_1 > & \frac{c_1 (\varepsilon_1 + y_1)}{K \left(A + y_0 + y_1 + \sum_{k=2}^{K} y_k\right)^2} \\
2 > & \frac{(\varepsilon_1 + y_1)}{K \left(A + y_0 + y_1 + \sum_{k=2}^{K} y_k\right)^2}
\end{aligned}$$

which holds everywhere because $\varepsilon_1 \to 0$ and $A + y_0 + \sum_{k=2}^{K} y_k > 1$ so that $\varepsilon_1 + y_1 \ll K \left(A + y_0 + y_1 + \sum_{k=2}^{K} y_k\right)^2$. So we have that the variance of the posterior expected copay is increasing in the number of observed non-zero copays, which should motivate more search among those who have observed more non-zero copays