An Empirical Model of Price Transparency and Markups in Health Care

(Job Market Paper)

Zach Y. Brown*
Columbia University

December 20, 2016

Latest version: www.columbia.edu/~zyb2001/research.html

Abstract

In the market for health care services, consumers often do not know exact prices when choosing where to receive care. This makes it difficult to shop around for low price options, potentially reducing effective consumer price elasticity and leading to higher prices. How does price transparency affect equilibrium prices and welfare? To answer this question, this paper develops a demand model that separates underlying consumer preferences from consumer uncertainty about prices. Identification comes from quasi-experimental variation in price information resulting from the introduction of a website aimed at informing consumers. I then combine the model of demand with a model of bargaining between medical providers and insurers to examine how price transparency affects equilibrium prices. Using administrative data on medical imaging claims and website usage, model estimates and difference-in-differences estimates both imply that the website reduces health care spending by 3 to 4 percent. I then use the model to examine the effects of price transparency more generally. In counterfactual simulations, I find that price transparency would generate a substantial reduction in equilibrium prices if a larger fraction of consumers in the market were informed. Combining the price transparency website with high cost sharing would give individuals more incentive to use the price transparency tool, reducing health care spending by 18 percent.

Keywords: health care, price uncertainty, price transparency, information frictions
JEL Classification: I13, L11, L86

---

*I am very grateful to Kate Ho, Mike Riordan, Doug Almond, and Chris Conlon for guidance and support. In addition, Bernard Salanié, Tobias Salz, and Adam Kapor provided valuable comments. I benefited greatly from participants in the Columbia Industrial Organization Seminar, Industrial Organization Colloquium, Applied Microeconomics Colloquium, and Econometrics Colloquium as well as discussions with Adam Sacarny, Fanyin Zheng, Tal Gross, Jean-François Houde, Mike Dickstein, Ashley Swanson, Victor Aguirregabiria, Leemore Dafny, Ben Handel, John Vickers, Fiona Scott Morton, and Robin Lee. The project received support from the National Science Foundation. I also thank Mary Fields, Tyler Brannen, Maureen Mustard, and the New Hampshire Department of Health and Human Services for providing data and insight into the New Hampshire HealthCost website. I obtained IRB approval for the use of the data through Columbia University (#AAAP8656). This paper does not reflect the views of the New Hampshire Department of Health and Human Services. All errors are my own.

1 Introduction

In certain markets, consumers do not know exact prices until they have committed to a purchase. This is often the case for automotive repair, building construction, and financial services, as well as other products with complicated bundling, discounts, or add-ons.[1] Ex-ante uncertainty about prices is particularly common in the U.S. private health care market. Health care prices are determined in private negotiations between insurers and medical providers, and firms are often contractually forbidden from disclosing these negotiated rates. As a result, the vast majority of consumers say they do not compare prices before receiving medical care.[2] In addition to making it difficult to shop around for medical services, the lack of price transparency may increase hospital prices. In response, some policy makers have called for more "price transparency" in health care.[3]

While an influential literature, starting with Stigler (1961) and Diamond (1971), has examined search frictions, there has been little emphasis on markets in which it is not possible to aquire price information. Like search costs, the lack of price transparency may increase prices and lead to price dispersion. Understanding how price transparency affects prices is particularly important for privately-provided health care in the U.S. since the market comprises about 6 percent of GDP and the relatively high level of spending is often attributed to high prices.[4] In addition, a recent literature has documented the large degree of price dispersion in health care, even for relatively standardized procedures (Cooper et al. 2015).

This paper empirically evaluates how price transparency affects markups and welfare in the U.S. health care market. I combine a model of demand that incorporates price uncertainty with a model of bargaining between providers and insurers. Although a relatively small fraction of consumers currently try to obtain price information when price transparency tools are made available, the model allows for an analysis of out-of-sample counterfactual scenarios. I find that there would be a considerable reduction in equilibrium health care prices if all consumers were informed about prices. One way to increase the number of informed consumers is to combine price transparency tools with high cost sharing health plans, incentivizing consumers to become informed about prices. This combination of policies would also lead to a large reduction in equilibrium prices.

First, this paper introduces a discrete-choice model in which consumers choose where to receive medical care with potentially limited information about prices. In the model, consumers

[1] See, for example, Ellison (2005).

[2] See, for instance, "How Much Will it Cost? How Americans Use Prices in Health Care," Public Agenda, March 2015.

[3] More than half of U.S. states have proposed health care price transparency laws in recent years. Price transparency legislation has also been proposed at the federal level. See "2015 Price Transparency Initiative State Survey," The Source on Healthcare Price & Competition, UC Hastings College of the Law, October 21, 2015.

[4] See Martin et al. (2016) for information on private health care spending. For a discussion of high prices in the market for health care services, see, for example, Anderson et al. (2003), Koechlin et al. (2010), and Cooper et al. (2015).

with rational expectations receive noisy signals about prices, and, consequently, are less able to discern which are the low price options. Consumers may choose options they believe to be the best value but are often surprised by the bill. Accounting for the difference between expected prices and actual prices is important for recovering underlying consumer preferences, including price sensitivity, and evaluating the welfare effects of price information.

The estimation strategy makes use of plausibly exogenous variation in consumers' information set stemming from a price transparency website introduced by the New Hampshire state government. In contrast to other price transparency efforts, the website allowed any privately-insured consumer in the state to enter insurance information and easily compare out-of-pocket prices across hospitals and other providers. I exploit difference-in-differences variation based on the fact that the website was introduced in March 2007, and could only be used to obtain price information for a subset of medical imaging procedures. If consumers use the price transparency website when it is available, I assume that they have perfect information about prices.

In the demand model, individuals' beliefs about prices are treated as unknown parameters to be estimated.[5] These high dimensional latent variables complicate the estimation strategy. To address this issue, I take advantage of recent Bayesian techniques and employ a Markov chain Monte Carlo (MCMC) estimator. This approach allows for a feasible estimation strategy that recovers parameter estimates summarizing individuals' beliefs about prices in addition to underlying taste parameters.

Next, I turn to the supply side and present a bargaining model to recover information about marginal cost and examine how price transparency affects negotiated prices in equilibrium. Recent empirical work has used models of bilateral bargaining between insurers and medical providers to gain insight into the effects of hospital and insurer competition (Gowrisankaran et al. 2015; Ho and Lee 2016). While others have suggested that price transparency can affect health care prices, I develop the first model of equilibrium behavior that incorporates consumer price uncertainty.[6] I use the first order condition of the bargaining equation to derive an expression for equilibrium prices and find that price transparency leads to a trade-off. First, price transparency can make residual demand more elastic, decreasing the incentive for providers to negotiate high prices. Second, price transparency ensures that consumers do not choose high cost providers, implying that insurers may be more willing to have high cost providers in their network. This can actually reduce the incentive of insurers to negotiate low prices. Therefore, the effect of price transparency on negotiated prices is theoretically ambiguous. Using the bargaining model, I derive a moment condition that allows me to recover marginal cost and investigate the effect of price transparency empirically.

The model is estimated using detailed administrative data on private health care claims

[5]Specifically, individuals are assumed to receive a price signal that is the true price plus a mean-zero error. The error is unobserved by the researcher.

[6]For a discussion about how price transparency could affect markups see "Health Care Price Transparency: Can It Promote High-Value Care?", Commonwealth Fund, April/May 2012. Also see Section 2.

and price transparency website usage in New Hampshire. The claims data contain information on the actual out-of-pocket price that consumers pay as well as the price paid by insurers.[7] I focus on relatively simple outpatient medical imaging procedures—X-Rays, CT scans, and MRI scans.[8] The negotiated price of these procedures ranges from a few hundred dollars for X-rays to a few thousand for MRI scans. Despite the fact that specific medical imaging procedures are relatively standardized, I find that the price of each procedure varies widely across providers in the state.[9] In addition to individual-level information on the choice of medical provider, I also utilize disaggregated information on usage of the price transparency tool obtained from website traffic logs.

In my first empirical exercise, I use the estimates from both the demand and supply model to evaluate the effect of New Hampshire's price transparency website. I find that the website resulted in overall savings of 3 percent.[10] This intent-to-treat effect is consistent with my reduced-form results in Brown (2016). In particular, I found that the website reduced overall spending by about 4 percent using a difference-in-differences methodology. These savings are primary due to increased price-shopping on the part of consumers, however part of the decline is also due to a small reduction in the equilibrium prices.

Even though New Hampshire's price transparency website was publicly available to all individuals in the state, a relatively small fraction of consumers in the market actually used it. Overall, I find that consumers used the website for about 8 percent of medical imaging visits when the website was available.[11] Unlike the reduced-form analysis, the empirical model allows me to examine the effect conditional on using the price transparency website.[12] Estimates imply that the website primarily benefited individuals most exposed to the full price, i.e. those subject to a deductible. Individuals with a deductible that used the website saved $178 per visit on average, while individuals without a deductible saved $16. It is important to note, however, that price information may cause individuals to switch, for example, from nearby hospitals perceived as high quality to distant imaging centers with lower perceived quality. Taking the change in non-price attributes into account, I show that the gain in consumer surplus for these individuals is $132 and $12, respectively.

Given modest website usage, the effect on equilibrium prices may be larger if more consumers

[7] These are the same data used to calculate prices for New Hampshire's price transparency website.

[8] The claims data cover all privately-insured individuals in the state, over 1 million covered lives. There are 177,995 individuals with medical imaging procedures over the period.

[9] This is consistent with previous research documenting the large degree of price dispersion for these procedures nationally (Cooper et al. 2015). Also note that medical imaging procedures in the U.S. are roughly double the price of the same procedures in other OECD countries with available data. See "The US health system in perspective: a comparison of twelve industrialized nations," Commonwealth Fund Issue Brief, 2011.

[10] Overall savings refers to change in spending for both insurers and consumers.

[11] This is obtained by dividing website traffic for medical imaging procedures by the number of privately-insurers individuals in New Hampshire receiving medical imaging procedures available on the website.

[12] The effect conditional on program take-up is sometimes referred to as the treatment-on-the-treated effect. The presence of spillover effects hinders estimation of the treatment-on-the-treated using reduced-form methods.

are informed about health care prices. There are two factors that make it difficult to extrapolate from reduced-form estimates. First, even though the availability of the website is exogenous, use of the website when it is available is potentially endogenous. If the individuals who find out about the website and choose to use it are those that receive a larger benefit, there may be decreasing savings as more individuals become informed about prices. Second, equilibrium prices are a function of the number of consumers that have price information. By affecting negotiated prices, price transparency generates spillover effects that benefit all consumers, including those that do not have price information.[13] By estimating the individual-specific probability of using the website and deriving a bargaining equation, the empirical model presented in this paper allows me to address both issues when examining out-of-sample counterfactual scenarios.

Counterfactual simulations imply that, while selection is present, the effect on equilibrium prices dominates. As a result, there would be a considerable reduction in equilibrium prices if a larger fraction of consumers had information. If all consumers were informed, equilibrium prices would be 19 percent lower. Prices decline because demand effectively becomes more elastic, allowing insurers to negotiate lower prices with most providers in their network. In addition, consumers would choose lower cost providers in their choice set, resulting in per visit savings of $44 for consumers and $166 for insurers relative to no price transparency. Overall, spending would decline by 28 percent. Savings would come largely at the expense of provider profits, although some of the savings would also be due to individuals switching to providers with lower marginal cost (e.g. imaging centers and clinics rather than hospitals).

Finally, I shed light on policy by simulating the effect of combining price transparency with high cost sharing insurance plans. One potential reason that current price transparency tools are not widely used even when they are available is that many consumers, especially those that pay a small coinsurance rate, have modest private gains from becoming informed and price shopping. High cost sharing plans expose consumers to a larger fraction of negotiated prices, increasing their incentive to use the price transparency website. In counterfactual simulations, I find that high cost sharing, which I define as a 50 percent coinsurance rate, would lead to a 38 percent increase in the number of consumers using the website. In addition, consumers would have more incentive to choose a low cost provider once they had price information. For these reasons, equilibrium prices would be almost as low as the full information case without high cost sharing. Although this would result in higher out-of-pocket spending for consumers, overall health care spending on medical imaging procedures would decline by 18 percent.

1.1 Related Literature

This paper is related to the large literature on search costs and competition, starting with Stigler (1961). Even with homogenous goods and many sellers, search costs can lead to higher prices (e.g. Diamond 1971; Stahl 1989). Search costs have been shown to be empirically important in

[13] This is similar to a search externality. See Salop and Stiglitz (1977).

a large variety of markets.[14] A common assumption in this literature is that individuals make a purchase decision after learning the price of at least some of the options (i.e. the consideration set).[15] In contrast, this paper studies a context in which individuals make decisions under uncertainty. Although there are similarities to search frictions, the welfare consequences of price uncertainty are distinct since individuals may be surprised by their bill. The model presented in this paper has implications for other situations in which it is not possible to observe actual prices when making a purchase decision, such as markets where consumers receive price quotes.

This paper is also related to the literature examining markets with shrouded add-on pricing. The price of add-ons may be shrouded in equilibrium due to consumer lack of self-control (DellaVigna and Malmendier 2004), selection issues (Ellison 2005), bounded rationality (Spiegler 2006), or myopia (Gabaix and Laibson 2006).[16] Empirical work has found that obfuscation of shipping charges affects consumer behavior (Ellison and Ellison 2009; Brown et al. 2010). Related work on bill-shock has examined situations in which consumers are inattentive about the price of the next unit of consumption, such as for cellular phone contracts (Grubb 2014; Grubb and Osborne 2015). Pricing in the market for medical services can be seen as the limit-case of add-on pricing—in the absence of price transparency tools the full price is partially shrouded. Therefore, the model developed in this paper can be seen as a new approach to add-on pricing in which consumers have noisy beliefs about shrouded attributes and maximize expected utility.

Previous work has also examined the case in which consumers lack information about product attributes other than price.[17] For experience goods, consumers may initially lack information about product quality or other non-price attributes (e.g. Erdem and Keane 1996; Ackerberg 2003; Erdem et al. 2008; Allcott 2013). Building on Allcott (2013), as well as theoretical work by Schmeiser (2014), Train (2015) formalizes the calculation of consumer surplus in discrete-choice models when anticipated attributes are different from experienced attributes. I use this approach to calculate welfare when price is different than expected.

While this paper argues that information frictions are important for understanding consumers' choice of medical providers, a broader literature has emphasized frictions in other parts of the health care system. For instance, Handel and Kolstad (2015) find evidence that a variety of frictions affect health insurance choice. Prior literature has also found that consumer inattention or inertia has implications for Medicare Part D (Ericson 2014; Decarolis 2015; Ho et al. 2016). In addition, there is evidence that uncertainty about the effectiveness of different drugs is relevant for pharmaceutical demand (Crawford and Shum 2005; Ching 2010; Dickstein

[14] Empirical work has studied search frictions in a markets such as mutual funds, textbooks, online bookstores, grocery stores, auto insurance, electricity, online hotel booking, cars, and trade-waste (Hortaçsu and Syverson 2004; Hong and Shum 2006; De Los Santos et al. 2012; Seiler 2013; Honka 2014; Giulietti et al. 2014; Koulayev 2014; Moraga-González et al. 2015; Salz 2015). There has also been related work on technology that reduces search costs (e.g. Brown and Goolsbee 2002; Jensen 2007; Jang 2015; Luco 2015)

[15] This is true for models of sequential as well as non-sequential search.

[16] Also see Grubb (2015) for related review.

[17] This is also related to the literature on quality disclosure. For an overview, see Dranove and Jin (2010).

2014). In a similar vein, a literature has examined uncertainty about quality of medical services and medical devices (e.g. Cutler et al. 2004; Kolstad 2013; Grennan and Town 2015). Finally, Grennan and Swanson (2016) find that information affects hospital-supplier bargaining. Despite this growing literature, to my knowledge, there is no evidence on the welfare effects of frictions that affect consumers' choice of hospital.

After estimating a demand model that incorporates price uncertainty, I use the demand parameters to estimate a model of bilateral bargaining between insurers and providers. Using a theoretical framework based on Horn and Wolinsky (1988), empirical models of bilateral bargaining have been applied to a number of vertical markets (e.g. Crawford and Yurukoglu 2012; Grennan 2013; Allen et al. 2014; Soares 2015). A recent literature has also used this approach to examine bargaining between providers and hospitals in order to examine hospital mergers (Gowrisankaran et al. 2015), hospital system bargaining power (Lewis and Pflum 2015), and insurer competition (Ho and Lee 2016). With the exception of Allen et al. (2014) work examining consumer lending, empirical models of bargaining in oligopolistic markets have assumed perfect information.[18] To examine the effect on prices, I use an approach closely related to Gowrisankaran et al. (2015), however I incorporate price uncertainty. I show that price transparency affects equilibrium prices through three distinct channels: it affects provider incentives by changing effective demand elasticity, it affects insurer cost, and it affects the consumers surplus of the insurers' enrollees. Overall, the effect of price transparency is ambiguous, which has implications for other vertical markets where consumers have uncertainty about product characteristics.

I emphasize the fact that providing information to a subset of consumers can affect equilibrium prices, generating an externality for uninformed consumers. This is related to the theoretical literature examining similar search externalities that arise when only a subset of consumers are savvy (Salop and Stiglitz 1977; Armstrong 2015). In an empirical context, Salz (2015) finds evidence that intermediaries can also generate search externalities. The presence of these externalities implies a role for the public provision of information in these markets.

Prior reduced-form work has examined the effect of health care price transparency efforts by individual employers or insurers (Lieber 2015; Whaley 2015a,b; Desai et al. 2016). In particular, Lieber (2015) and Whaley (2015a) find evidence that this information allowed some individuals to shop around for lower cost options, while Desai et al. (2016) finds little effect. In contrast, the state-run price transparency website in New Hampshire was available to all individuals in the state. In Brown (2016), I use reduced-form methods and find that it not only increased price shopping by consumers, it had a small but statistically significant effect on equilibrium prices. In this paper, I argue that there is potential for substantial reductions in equilibrium prices if a larger fraction of consumers are informed about prices. This would be the case if, for instance,

[18] Allen et al. (2014) incorporate search frictions into a model of the mortgage market. Note that while this paper models business-to-business bargaining, Allen et al. (2014) examines negotiations between consumers and lenders.

price transparency tools were combined with high cost sharing health plans.

1.2 Roadmap

The remainder of the paper is organized as follows. Section 2 describes the data and provides background on the price transparency website. Section 3 presents the model of website usage and choice of medical provider. I also discuss the estimation and robustness. Section 4 presents the bargaining model, focusing on the role of consumer information. Section 5 presents the results from the demand model and the supply model. Section 6 uses the estimates to examine the effect of the website while Section 7 presents out-of-sample counterfactual simulations. Section 8 concludes.

2 Data and Background

I utilize an all-payer claims database from New Hampshire that provides detailed information on negotiated prices along with information about how much is paid by the individual versus the insurer. I use these data to construct the individual-specific out-of-pocket price for each option in individuals' choice sets. In Section 2.2 I provide background on the price transparency website in New Hampshire and describe the variation that is used to estimate the model. I also describe the website traffic data which is used to construct information about the fraction of consumers with price information when the website is available.

2.1 New Hampshire Medical Claims

The main dataset contains enrollment and claims for the universe of individuals with private health insurance in New Hampshire for the period January 2005 to November 2010.[19] These data were collected as part of the New Hampshire Comprehensive Health Care Information System (NHCHIS), which assembled data from all commercial insurers in the state. The data were collected by the state in order to analyze health spending and construct prices for the price transparency website.

This paper analyzes the market for outpatient medical imaging services. This includes X-rays, computerized tomography (CT) scans, and magnetic resonance imaging (MRI) scans, all of which are diagnostic procedures that provide internal images of the body. Note that in Brown (2016), I use a broader definition of radiology procedures that includes procedures such as bone density scans and PET scans. In this paper I focus on X-rays, CT scans, and MRI scans because these procedures are relatively common, allowing me to construct accurate prices within individuals' choice sets.[20]

[19] Although the data include information about claims in later years, I focus on the period prior to December 2010 since this is when website traffic data is available.

[20] In Brown (2016), I find evidence that the price transparency website may have affected the quantity of mammograms (but not other procedures). I also exclude mammograms from the analysis, allowing me to assume

I limit the sample to individuals covered by managed care plans under the three main insurers in the state, Anthem, Cigna, and Harvard-Pilgrim. These insurance companies offer a variety of managed care plans, including Health Maintenance Organization (HMO) plans, Preferred Provider Organization (PPO) plans, Point-of-Service (POS) plans, and Exclusive Provider Organization (EPO) plans.[21] In all of these plans, the insurers negotiate lower prices with a selected network of providers, however the plans differ according to the level of cost sharing and the rules for seeing specialists or going to an out-of-network provider. Although all individuals in the NHCHIS dataset are insured by plans in New Hampshire, some live outside the state. I remove these individuals as well as individuals that go to providers in states other than New Hampshire and surrounding states (Massachusetts, New York, Maine, and Vermont).

Each medical claim is associated with an individual procedure, however a medical imaging visit may contain multiple procedures.[22] Since the price of the bundle of procedures is the relevant amount for consumers, the price transparency website displays price aggregated to the visit level. I follow a similar procedure as the website (using the same dataset) in order to calculate visit prices at each provider. In particular, I aggregate to the visit level by summing all procedures on the day of the visit. I exclude visits in which there was a more expensive primary procedure performed on the same day. This ensures that the sample contains only medical imaging visits that are self-contained. The method used to define visits and associated prices is described in greater detail in Brown (2016).

Each visit is categorized by the imaging procedure, defined by a CPT/HCPCS code.[23] These codes are quite specific and refer to relatively standardized procedures. The full list of medical imaging procedures is given in Table A-1. For each visit, I am able to calculate the out-of-pocket price paid by consumers, the price paid by insurers, as well as the list price.[24] The list price is not relevant for individuals in the sample since insurers negotiate prices that are lower than the list prices. This negotiated price is obtained by simply summing the amount paid by consumers and insurers. The ratio of the out-of-pocket price to the negotiated price determines the individual-specific level of cost sharing (e.g. if the individual is under the deductible, then the cost sharing is equal to 1). Prices are inflation-adjusted to 2010 dollars using the Medical Care Services CPI from the U.S. Bureau of Labor Statistics.

For each visit, an identifier allows me to link information about the medical provider that performed the procedure, which includes both hospital and non-hospital facilities. While hospitals offer outpatient medical imaging services, freestanding outpatient facilities (e.g. imaging

that all individuals choose an inside-option.

[21] Less than 2 percent of enrollees are in indemnity (fee-for-service) plans. I remove these individuals.

[22] For instance, a CT scan may contain a charge for the scan itself as well as supplemental charges for oral contrast agent which help highlight specific parts of the body.

[23] The American Medical Association developed and maintains Current Procedural Terminology (CPT) codes. Healthcare Common Procedure Coding System (HCPCS) codes are an extension of CPT codes that include additional procedures and services.

[24] The data also contain information on capitation payments to providers. Over the relevant period in New Hampshire, these payments were very small.

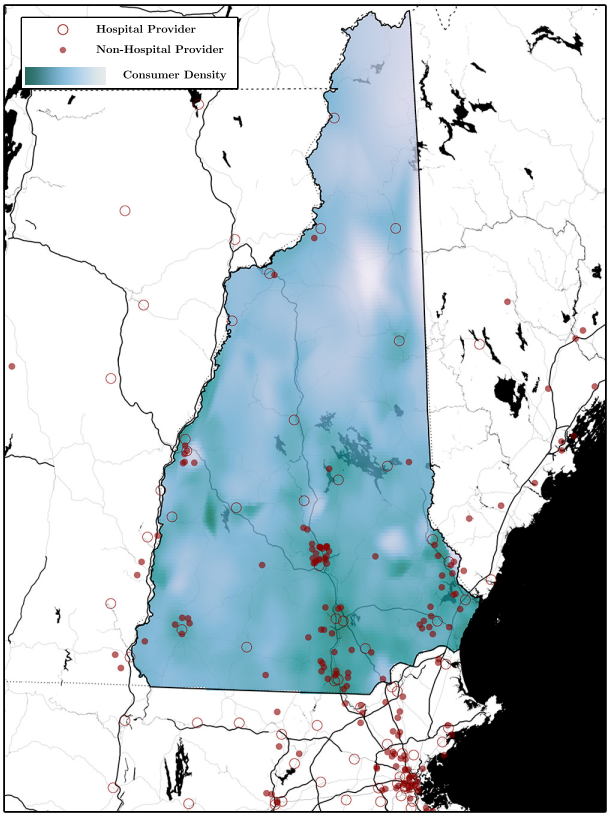
centers) are significantly less expensive. In New Hampshire, the average total cost of imaging procedures is $1,004 at hospitals but only $797 at non-hospital providers. In addition to observing provider type, I also observe the provider zip code.[25] The location of these providers is shown in Figure 1.

For individuals, I observe age, sex, zip code, insurance enrollment, and whether they are subject to a deductible. I also observe a patient identifier. I define 5 different age groups (0-18, 19-35, 36-50, 51-64) and omit individuals over age 65 since they are likely eligible for Medicare. Average income and education using the 2007-2010 American Community Survey is linked to each individual using the zip code. In addition, patient zip code is used to calculate the distance to each provider. Using observed International Classification of Diseases (ICD) codes, I also construct a measure of chronic diseases or conditions that may affect how difficult it is to treat patients. This measure is referred to as the Charlson Comorbidity Index.[26] Finally, I construct

[25]Note it is not possible to obtain the identity of each provider and link additional information.

[26]In particular, the Charlson Comorbidity Index is an integer score that is often used to predict mortality. See Charlson et al. (1987) and Stagg (2006).

Figure 1
Density of Consumers and Location of Medical Imaging Providers



Notes: Map shows the location of providers providing medical imaging services that service privately-insured individuals in New Hampshire.

an indicator for whether each individual has the medical imaging procedure in the week following an emergency. Although these are relatively minor emergency visits since I exclude inpatient admissions, this may affect demand since it may be more time sensitive (e.g. demand for medical imaging procedures after a bone fracture may be different than for routine preventative care).

Table 1
Summary of Privately Insured Individuals with Medical Imaging Claims

	Mean	SD	Min	Max
Male	0.46	0.50	0	1
Age 0-18	0.20	0.40	0	1
Age 19-35	0.19	0.39	0	1
Age 36-50	0.31	0.46	0	1
Age 51-64	0.30	0.46	0	1
Charlson Comorbidity Index	0.6	0.8	0.0	2.0
Zip income ($1,000s)	83.5	24.6	22.0	309.7
Zip BA Degree (%)	33.9	13.8	0.0	100.0
Insurance Type:				
PPO	0.29	0.45	0	1
POS	0.15	0.36	0	1
HMO	0.50	0.50	0	1
EPO	0.06	0.24	0	1
Insurance Company:				
Anthem	0.60	0.49	0	1
Cigna	0.23	0.42	0	1
Harvard Pilgrim	0.17	0.37	0	1
Plan Characteristics:				
Plan has Deductible	0.44	0.50	0	1
Unique Individuals		177,995		

Notes: Includes all unique privately insured individuals in the state of New Hampshire over the period 2005 to 2010 with at least one outpatient medical imaging visit. If an individual switches insurance plans, summary statistics are calculated using only their first plan.

Table 1 provides a summary of individuals in the sample. There are 177,995 unique individuals with outpatient imaging visits over the period. Half of the individuals are in HMO plans, and most of the remainder are in PPO or POS plans. About 44 percent of individuals have a plan with a deductible. New Hampshire is a relatively high-income state, and privately insured individuals have even higher income than the general population.

When an individual needs a specific procedure, the choice set is defined as the providers that are available through the individual's insurance plan that can perform the procedure in the given year. Although I do not observe each insurer's network directly, I construct a proxy

by examining the providers chosen by individuals in each insurance company-product pair (e.g. Anthem HMO). In some cases, individuals may have plans, such as PPO plans, that allow them to choose providers out-of-network. To the extent that individuals actually choose these providers, they are included in the choice set (but have higher prices).[27] For each option in the choice set, I construct procedure prices that vary by insurance company-product pair and year. In addition, out-of-pocket prices vary across individuals with the same insurance product since some individuals are under the deductible and some are not. Within each individual's choice set, I remove providers that cannot perform the procedure as well as those that are more than 75 miles from the individual.

In general, patients are told they need a diagnostic test by their primary care physician or other specialist. They may receive a referral, however consumers are generally free to schedule an appointment for a medical imaging procedure at any provider within their insurer's network.[28] Although the NHCHIS dataset does not have information on referrals, I construct a measure of likely referrals. To do this, I find each individual's primary care physician in each year, defined as the most frequently visited primary care physician. I then find the most common medical imaging provider chosen by the primary care physician's patients. Using this, I construct an indicator for likely referrals.

Table 2
Summary of Medical Imaging Visits by Insurer

	Anthem		Cigna		Harvard Pilgrim	
	Mean	SD	Mean	SD	Mean	SD
Observations	2,142,583		442,836		457,294	
Number of choice situations	200,231		48,938		52,760	
Number of unique patients	115,370		32,259		30,366	
Number of unique non-hospital providers	177		110		88	
Number of unique hospital providers	38		14		0	
Providers in choice set	13.7	5.3	13.1	6.2	11.2	4.5
Total Negotiated Price	924.2	1060.5	680.9	815.5	677.0	761.2
Insurance price	827.3	1015.0	639.8	785.2	601.7	719.5
Out-of-pocket price	96.9	216.2	41.0	91.4	75.3	185.5
Distance to provider	38.2	18.0	35.0	19.0	33.2	17.9
Choose hospital	0.34	0.47	0.21	0.40	0.00	0.02
Choose referral	0.27	0.45	0.36	0.48	0.31	0.46

Notes: Includes all outpatient medical imaging visits for privately insured individuals in the state of New Hampshire over the period 2005 to 2011. All prices in 2010 inflation-adjusted dollars.

[27]For the purposes of the model, I refer to the set of providers that individuals can access given their insurance as the "network" even though this could potentially include providers that are technically out-of-network.

[28]Note I do not include inpatient medical imaging procedures since patients are unlikely to choose their provider when they are already admitted to a hospital.

The full dataset is summarized for each of the three insurers in Table 2. Anthem is by far the largest insurer, with over 200,000 medical imaging visits over the period. On average, the out-of-pocket price is 12 percent of the total negotiated price. However, there is large variation—individuals under the deductible pay the full price. In particular, there is greater cost sharing in the beginning of the year, when individuals have not hit their deductible, then at the end of the year (see Figure A-1). Individuals choose between 13 different providers on average, although, again, there is significant variation. This is partially due to the fact that there are more providers that are capable of performing X-rays than MRI scans. Given large number of observations, I use a 2 percent sample of visits for the main analysis.

Within individual's choice sets there is a large degree of price dispersion, and consequently, significant potential savings if individuals switch to low cost options. Figure 2a shows the distribution of demeaned negotiated prices within individuals' choice sets. The distribution is approximately normal, with standard deviation of $639 (and coefficient of variation of 44.5 percent). If a consumer is under the deductible for the year, the individual is fully exposed to the variation in prices. However, since most patients share cost with an insurer, out-of-pocket price dispersion is smaller, with a standard deviation of $127 (see Figure 2b).[29] Finally, Figure 2c shows the distribution of prices paid by the insurer.

Given the variation in prices, there are large potential savings if consumers switched to cheaper providers in their network. The potential savings for consumers and insurers are summarized in Table A-2. Overall, I find that there would be savings of over 40 percent if consumers switched to providers in the first quartile of the price distribution.[30] The savings are even greater for X-rays and CT scans. Consumers subject to a deductible have large private gains from switching, but much of the potential savings for consumers without a deductible go to insurers. This suggests that, although there are large potential savings for the health care system, these consumers may have little incentive to switch to less expensive providers even if they have price information.

2.2 HealthCost Website

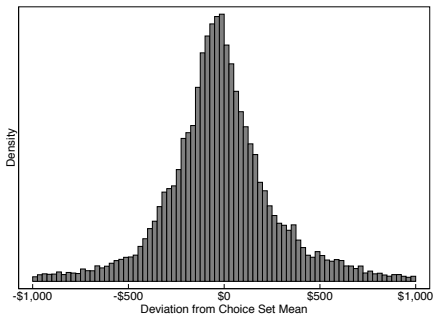
In an effort to increase health care price transparency, the New Hampshire Insurance Department launched the HealthCost website in March 2007.[31] Although other states have implemented health care price transparency initiatives, many only provide information on the hospital list price of each procedure (i.e. charge amount), which has little bearing on the out-of-pocket prices

[29] The coefficient of variation is 49.6 percent.

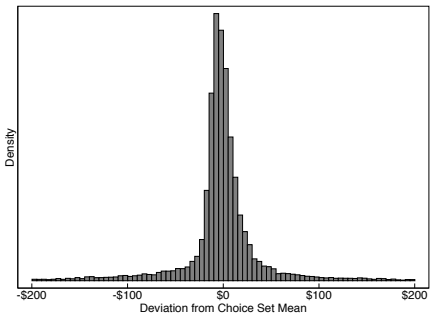
[30] These are the potential consumers and insurer savings if all consumers choosing a provider ranked above the first quartile in their choice set were to switch to the provider in the first quartile of their choice set.

[31] The website can be found at nhhealthcost.nh.gov. Originally the website was nhhealthcost.org. Note that in 2016, after the period of analysis, the website added additional information for consumers, including information provider quality.

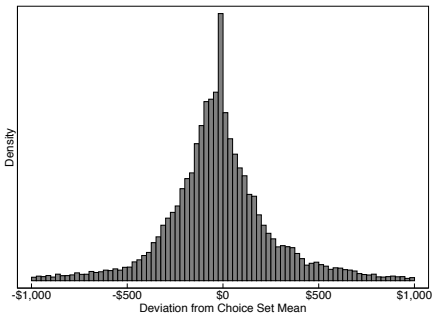
Figure 2
Price Variation within Individuals' Choice Sets



(a) Total Negotiated Price



(b) Patient Out-of-Pocket Price



(c) Insurer Price

Notes: Histograms show distribution of de-meaned prices in individuals' choice sets.

that insured individuals actually pay.[32] New Hampshire's HealthCost website was unique because it provided information about insurer-specific out-of-pocket prices. Although other states, such as Maine and Colorado, have since created tools with similar information, New Hampshire's price transparency efforts remain the most comprehensive.[33] Individuals with private insurance in the state can select one of about 35, mostly outpatient, procedures (see Figure A-2a). In addition to providing information for insured individuals, the website also has a separate tool for uninsured individuals in the state. Since the claims data cover the population of insured individuals, I focus only on the former. In recent years, the website added information about provider quality and a guide to health insurance. This occurred after my period of analysis. It is also important to note that there have been other price transparency efforts by individual insurers, notably Aetna which started its Member Payment Estimator tool in 2010. However, Aetna had a very small presence in New Hampshire and is excluded from the analysis.

[32]Information about list prices may affect uninsured individuals. See Christensen et al. (2015), who examine the effect of information about list prices.

[33]New Hampshire was the only state to receive an "A" grade from Catalyst for Payment Reform's 2015 Report Card on State Price Transparency Laws.

To use the website, consumers enter their insurance information, deductible, zip code, and search radius and the website returns a list of median bundled out-of-pocket prices at each provider calculated using the NHCHIS dataset. Figure A-2b shows an example of prices returned by the website. The table of prices is automatically sorted by out-of-pocket price, making it easy for consumers to schedule an appointment with the lowest cost provider. In addition to the out-of-pocket price, the website also returns the amount paid by insurers and the total negotiated price. For the purposes of analysis, I assume that individuals who use the website are fully informed about prices. I discuss this assumption in greater detail in Section 3.2.[34]

According to discussions with state employees, the website was promoted by encouraging insurers and primary care doctors to inform patients about the website. In addition, there were at least 40 news articles mentioning the website over the period. On average, there were 41,506 searches for price information per year according to website traffic logs, about a third of which were for medical imaging procedures. Furthermore, anecdotal evidence suggests that the website not only let consumers shop around, but may have allowed insurers to negotiate lower rates. One report noted that after the introduction of the website "the balance of plan-provider negotiating power began shifting significantly in New Hampshire."[35] In particular, Anthem, the largest insurer in New Hampshire, had a public battle with an expensive hospital in the state. Local news sources suggest that the price transparency website allowed the insurer to negotiate lower prices.[36]

In order to examine the effect of price transparency, this paper exploits two sources of variation generated by the HealthCost website. First, there is variation due to the timing of the website introduction. In this way, I can examine procedures on the website and compare observed choices from 2005 to February 2007, prior to the introduction of the website, to observed choices in the period starting March 2007. Second, there is variation due to the fact that only a subset of medical imaging procedures were available on the website.[37] The X-ray, CT scan, and MRI scan procedures with and without information available on the website are listed in Table A-1. I argue that imaging procedures on the website tend to be quite similar to procedures not on the website. For example, the price of a knee X-ray is available on the website while the price of a knee/leg CT scan is not. Note that the website also had price information for a few simple surgical procedures (e.g. kidney stone removal), physician office visits, as well as newborn delivery. I do not consider these procedures in the analysis because they tend to be

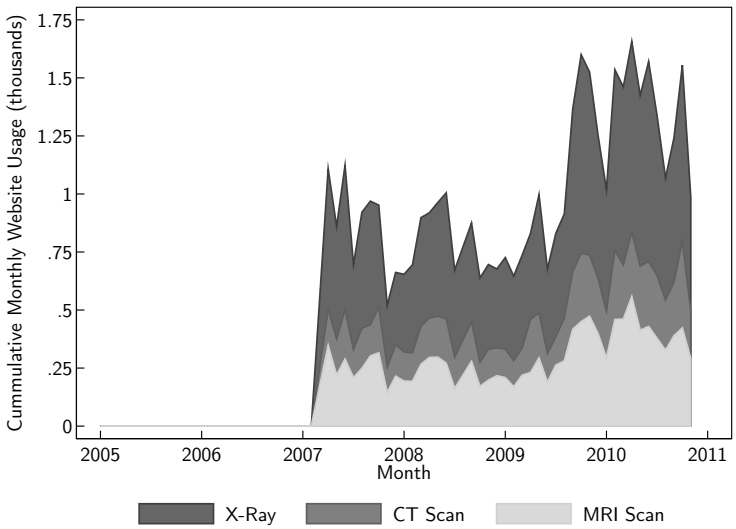
[34]The website also provides information on precision of the cost estimate and typical patient complexity. I argue these are less relevant for medical imaging procedures since the procedures are relatively common (making estimates fairly precise) and relatively standardized (meaning price depends little on patient complexity).

[35]See "Moving Markets: Moving Markets: Lessons from New Hampshire's Health Care Price Transparency Experiment," April 2014. California HealthCare Foundation.

[36]See "Higher costs of services snags Exeter Hospital's new deal with Anthem," *Portsmouth Herald*, November 7, 2010 and "Exeter Hospital says costs being used as negotiating tactic," *Portsmouth Herald*, November 14, 2010.

[37]According to discussions with state employees, only a subset of procedures were chosen because cleaning the data and constructing prices was time consuming and the department had limited resources. Note that after the period of analysis, the website added additional information.

Figure 3
Price Transparency Website Usage for Medical Imaging Procedures
By Month



Notes: Chart shows cumulative searches by procedure group. Includes all searches using "Health Costs for Insured Patients" wizard on either nhhealthcost.nh.gov or nhhealthcost.org. Note the website began in March 2007.

less standardized and involves a different set of providers.

In Brown (2016), I use these two sources of variation in a difference-in-differences framework. In that paper, the key identifying assumption is that the price of procedures on the website would follow a common trend relative to procedures not on the website if the website were never available. I argued that this assumption was plausible given the price trends that exist prior to the introduction of the website. In this paper, I develop an empirical model that relies on an alternative, but related, set of assumptions. One of the key assumptions is that individuals' utility parameters are orthogonal to whether procedures are available on the website. I compare the approach taken in this paper with the reduced-form identification strategy and discuss the structural assumptions in more detail in Section 6.1.

I use website traffic logs obtained from the New Hampshire Insurance Department to calculate the number of website price searches in each month for each procedure listed on the website. Website traffic data is available from March 2007 through November 2010, at which point the website switched hosting companies. Figure 3 shows cumulative monthly price searches for X-rays, CT scans, and MRI scans. When the website was first introduced in 2007 there were about 750 to 1,000 searches per month for the price of medical imaging procedures, however this grew to over 1,500 searches per month by late 2009.

In order to estimate the fraction of informed consumers I divide the number of price searches

Table 3
Monthly Percent of Consumers with Price Information
By Procedure Listed on Price Transparency Website

	Mean %	SD %	Min %	Max %
X-Ray (Ankle)	6.2	3.5	1.5	17.5
X-Ray (Chest)	1.5	0.8	0.6	4.2
X-Ray (Foot)	2.9	1.3	1.4	7.9
X-Ray (Knee)	3.3	1.5	1.7	9.6
X-Ray (Shoulder)	5.2	2.7	2.9	17.3
X-Ray (Spine)	2.4	1.3	0.9	7.9
X-Ray (Wrist)	2.3	1.1	1.0	7.2
CT (Abdomen)	5.3	2.9	2.6	15.2
CT (Chest)	13.4	6.5	6.3	33.1
CT (Pelvis)	15.9	8.6	5.7	50.3
MRI (Back)	9.3	5.0	3.9	29.3
MRI (Brain)	12.0	6.6	5.6	38.0
MRI (Knee)	11.8	5.9	6.1	34.2
MRI (Pelvis)	19.7	11.5	6.2	67.7

Notes: Percent of consumers with price information in each month for each procedure is calculated as website usage (from website traffic logs) divided by visits aggregated across all related CPT codes (from claims data). Period of analysis is March 2007 to November 2010, the period in which website traffic data is available.

per procedure by the total number of visits in New Hampshire from the claims data.[38] In other words, I assume that each use of the website is a unique individual.[39]

Table 3 shows the estimated percent of consumers with price information for each medical imaging procedure listed on the website. The percent of informed consumers is between 2 and 6 percent on average for X-ray procedures. There is a larger fraction of consumers that use the website for CT scans and MRI scans—between 5 and 19 percent on average. CT scans and MRI scans also tend to be more expensive, making the website potentially more valuable for consumers receiving these procedures. There is also temporal variation, potentially due to the fact that there is random variation in the type of individuals that need a procedure in a given month. In addition, more individuals may be learning about the website over time, as seen in Figure 3. This variation is used to help estimate the demand model and recover information about the choice to use the website if it is available.

[38]Note that the website procedures (e.g. knee X-ray) are more broad than the procedures as defined by CPT codes (e.g. knee X-ray with 1 or 2 views). Therefore, I aggregate across all CPT procedure codes related to the website procedure to obtain the total number of visits related to the website procedure in each month.

[39]If the same individual uses the website multiple times prior to a medical visit, the fraction of informed consumers would be lower. This would imply that the estimated savings conditional on using the website are actually larger. For this reason, the assumption that the number of website hits is equivalent to the number of informed consumers results in a conservative estimate of website savings.

3 Demand for Providers and Website Usage

This section presents a model of demand in which individuals have uncertainty about prices unless they use the price transparency website. The model has two parts. First, consumers may choose to use the price transparency website if it is available, in which case they learn actual out-of-pocket prices. I derive an expression for the expected benefit of price information given individuals' beliefs about prices and assume that individuals use the website if this benefit is greater than the cost. Second, consumers choose a medical provider. If individuals do not use the website or it is not available, they choose a provider with uncertainty about prices. However, if individuals use the price transparency website, they choose a provider with knowledge of all prices.

I start backwards and begin by discussing the choice of provider with and without price information in Section 3.2. In Section 3.3 I discuss the model of website usage using results derived from Section 3.2. The two parts of the model are estimated jointly. I present an estimation strategy that relies on revealed preferences in terms of both website usage and provider choices. The high dimensionality of the unobservables makes maximum likelihood estimation computationally infeasible. In Section 3.4, I present a Bayesian estimation strategy that addresses the estimation challenges by utilizing recent advances in MCMC methods. Finally, I discuss identification.

3.1 Model Setup and Timing

There are a set of providers that sell medical imaging services $\mathcal{J}$ indexed by j. The set of providers includes hospitals as well as non-hospital providers (i.e. freestanding outpatient facilities such as imaging centers and clinics). Each year, insurer $k \in \mathcal{K}$ contracts with a subset of providers, $\mathcal{N}_{kmt} \subseteq \mathcal{J}$, that can perform procedure $m \in \mathcal{M}$, where $\mathcal{M}$ is the set of medical imaging procedures.[40] Finally, let $i \in \mathcal{I}$ denote an individual enrolled in an insurance plan who needs a medical imaging procedure.

Each provider has a schedule of negotiated prices that is insurer-specific. In particular, the total price of procedure m at provider j for enrollees in insurer k at time t is given by $p_{jkmt} \in \mathbf{p}_{kmt}$, where $\mathbf{p}_{kmt}$ denotes the vector of prices across all providers. In Section 4, I model the bargaining process that determines these prices in each year. In contrast to the previous literature, it is important to note that I define prices at the visit level (i.e. prices include the cost of supplemental procedures as on the price transparency website).[41]

Individual i pays fraction c_{ikmt} of the negotiated price, which is observed in the claims data. The degree of cost sharing is determined by both the coinsurance rate applied to procedure

[40] Given that insurers contract with a network of providers, their role extends beyond providing insurance. For this reason, they are often referred to as managed care organizations.

[41] Focusing only on the main procedure would likely understate price differences across providers since consumers are in fact purchasing a bundle of procedures. Note that much of the literature focuses on inpatient hospital spending where prices are often defined by diagnosis.

m when enrolled in insurance plan k as well whether the individual is past the deductible for the year. In particular, if the individual is subject to a deductible then $c_{ikmt} = 1$. Therefore, for a given individual, cost sharing can vary over time t. The out-of-pocket price paid by the individual is

$$p_{ijkmt}^{OOP} = c_{ikmt} p_{jkmt}$$

I assume that this is the price internalized by the individual.[42] The remainder is paid by the insurer

$$p_{ijkmt}^{Insur} = (1 - c_{ikmt}) p_{jkmt}$$

After prices are determined via bargaining in each year, individuals that need a medical imaging procedure must choose a provider. I assume that each time an individual needs a medical imaging procedure there is the following timing:

1. The individual forms a prior about prices (i.e. they know the distribution from which prices are drawn)
2. The individual receives a vector of price signals and updates beliefs in a Bayesian fashion
3. The individual evaluates the expected gain from price information and chooses whether to use the website if it is available
4. The individual learns taste shocks and chooses the provider that maximizes expected utility

Previous to potentially using the price transparency website, the individual's taste shocks are unknown. This assumption is required to calculate the expected gain in consumer surplus from price information and tractably model the decision to use the website in the subsequent section. Learning the taste shock after choosing to use the website is consistent with the fact that consumers may evaluate providers based on observable characteristics, choose to use the website if it is available, and only then learn when providers have open appointment times. Under this interpretation, the taste shocks can be interpreted as individuals' idiosyncratic scheduling preferences.

After choosing a provider and receiving the procedure, the individual receives a bill and learns the true price. Welfare calculations must take into account the fact that realized price may differ from ex-ante beliefs about prices.

3.2 Choice of Provider

In this section I present a discrete choice model of provider demand in which consumers receive noisy signals about prices. If consumers do not become informed about prices, they choose a

[42] It has been suggested that individuals respond to dynamic incentives that arise due to annual deductibles, however the evidence is mixed (e.g Aron-Dine et al. 2015; Sacks et al. 2016; Brot-Goldberg et al. 2015). I assume that individuals do not anticipate whether they will surpass their annual deductible and respond only to the spot price.

provider with uncertainty about prices. However, if consumers use the price transparency website, they know true prices. Previous models of hospital demand either assume that individuals do not account for hospital prices at all or have perfect information about prices.[43] In contrast, I assume that individuals may have some information about prices even if they do not use the website. In this way, the model nests both the full information case as well as the case in which individuals completely ignore prices. In addition to price, the choice of provider is also assumed to depend on the distance from each individual to each provider, referrals, provider quality or amenities, as well as factors that vary with observed differences across individuals.

Individuals may only visit a provider in their network, $j \in \mathcal{N}_{kmt}$. There is no outside option since individuals are assumed to receive a medical imaging procedure if their doctor recommends it. One concern is that price transparency affects the choice to have a procedure at all. In Brown (2016), I use the entire sample of privately-insured individuals and examine the effect of the price transparency website on the probability of having medical imaging procedures and do not find a statistically significant effect.[44] This finding suggests that conditioning on individuals that had a medical imaging procedure and assuming they all choose an inside option is unlikely to bias counterfactual estimates.

Provider Choice When Prices are Known

I start by defining utility for the standard case in which prices are known. This expression is also the ex-post realized utility for the case in which individuals have ex-ante uncertainty.[45] For individual i with insurance k receiving procedure m from medical provider j, indirect utility is assumed to take the additively separable form

$$u_{ijkmt} = -\gamma_i p^{OOP}_{ijkmt} + \underbrace{\alpha_1 d_{ij} + \alpha_2 d_{ij}^2 + \alpha_3 r_{ijt} + \xi_{jM} + \beta \mathbf{x}_{ikmt} h_j}_{\delta_{ijkmt}} + \varepsilon_{ijkmt} \tag{1}$$

I allow for individual-specific heterogeneity in out-of-pocket price sensitivity, γ_i, which is distributed with density $f(\gamma_i)$. This approach has the benefit of not exhibiting the independence from irrelevant alternatives property and allowing for more flexible substitution patterns. It is also important since individuals that are more price sensitive may be more likely to use the price transparency website, which I explicitly account for in Section 3.3. I estimate the mean and variance of the distribution and allow the price coefficient to be correlated with the individual's average cost sharing, c_{ik}, since individuals with greater price sensitivity may differentially select

[43] For example, Kessler and McClellan (2000), Tay (2003), Ho (2006), and Ho and Lee (2016) assume that price does not influence patient choice while Capps et al. (2003), Gaynor and Vogt (2003), Ho and Pakes (2014) and Gowrisankaran et al. (2015) include price in individual utility and assume individuals have perfect information.

[44] In particular, I do not find an effect on quantity when mammograms are excluded. This is one reason why this paper focuses only on X-rays, CT scans, and MRI scans.

[45] Train (2015) refers to this as "experienced" utility and refers to ex-ante utility as "anticipated" utility.

into more generous plans.[46] Accounting for the adverse selection into insurance is important for understanding which individuals benefit from the price transparency website. In particular, I assume that the random coefficient is distributed normally:

$$\gamma_i \sim N(\bar{\gamma} + \rho c_{ik}, (\sigma^{\gamma})^2) \tag{2}$$

Since the same individual may have multiple medical imaging visits over the period, it is important to account for correlation in unobserved utility. Therefore, I assume that the random-coefficient is individual-specific (Revelt and Train 1998).

In addition to price, utility depends on observable non-price attributes, δ_{ijkmt}. This term includes distance from each individual to each provider, d_{ij}, distance-squared, d_{ij}^2, as well as an indicator for whether individual i was likely referred to provider j, r_{ijt}. Demand for hospitals may also differ depending on individual characteristics. Utility includes $\mathbf{x}_{ikmt}h_j$, the interaction between observable individual characteristics and an indicator for whether the provider is a hospital. The vector of individual characteristics, $\mathbf{x}$, includes age categories, sex, income, education, outpatient emergency indicator, and the Charlson Comorbidity Index. The last two are important for accounting for the fact that sicker patients or those in more urgent need of care may have distinct preferences. Utility is also a function of unobserved perceived quality or amenities at each provider, ξ_{jM}. This is allowed to vary according to the three procedure groups, X-rays, CT scans, or MRI scans, which are indexed by M. This accounts for the fact that providers may specialize in certain types of procedures.

Finally, ε_{ijkmt} is an idiosyncratic error distributed i.i.d. type 1 extreme value that is known by the individuals at the time the choice of provider is made. The observed choice probability of individual i enrolled in insurer k receiving procedure m at time t conditional on price information is

$$s_{ijkmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt}|\vartheta_{ikmt} = 1) = \int_{\gamma_i} \frac{\exp(-\gamma_i p_{ijkmt}^{OOP} + \delta_{ijkmt})}{\sum_{j' \in \mathcal{N}_{kmt}} \exp(-\gamma_i p_{ij'kmt}^{OOP} + \delta_{ij'kmt})} f(\gamma_i) d\gamma_i \tag{3}$$

where ϑ_{ikmt} is an indicator for whether the individual used the website and was informed about prices.

The expected consumer surplus, conditional on having price information, for a patient needing a medical imaging procedure is then:[47]

$$CS_{ikmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt}|\vartheta_{ikmt} = 1) = \frac{1}{\gamma_i} \log \left(\sum_{j \in \mathcal{N}_{kmt}} \exp(-\gamma_i p_{ijkmt}^{OOP} + \delta_{ijkmt}) \right) \tag{4}$$

[46] This approach is related to Limbrock (2011), who models selection into HMO plans and pharmaceutical demand. Note that c_{ik} is defined as the individual's average cost sharing for medical imaging procedures over the period of analysis.

[47] This is the consumer surplus before the idiosyncratic error is known. All expressions for expected consumer surplus are up to a constant. See Small and Rosen (1981).

Provider Choice with Price Uncertainty

Next, I model the case in which individuals have uncertainty about prices. Individuals form noisy beliefs using Bayes' rule and then make a decision based on those beliefs. This information structure is related to the empirical work on consumer learning (e.g. Erdem and Keane 1996; Ackerberg 2003; Erdem et al. 2008; Crawford and Shum 2005; Ching 2010; Dickstein 2014; Grennan and Town 2015). I assume that individuals know the distribution from which prices are drawn, which is assumed to be normal.[48] In particular, their prior is determined by the true mean and variance of prices their choice set, $\bar{p}^{OOP}_{kmt}$ and $\bar{s}^2_{kmt}$ respectively:

$$p^{OOP}_{ijkmt} \overset{iid}{\sim} N(\bar{p}^{OOP}_{kmt}, \bar{s}^2_{kmt}) \tag{5}$$

The prior provides no information about relative prices in the choice set, and therefore is not useful for choosing a provider on its own. However, individuals may be able to obtain additional information about individual prices. For instance, they may be able to look up list prices or receive potentially noisy price information from other individuals that had similar procedures. When asked, providers and insurers sometimes provide a price range if they provide any price information at all.[49] I model this by assuming that individuals receive a vector of unbiased signals, where each signal is given by

$$p^{OOP}_{ijkmt} + e_{ijkmt} \tag{6}$$

where p^{OOP}_{ijkmt} is the true price and e_{ijkmt} is signal noise with density $f(e_{ijkmt})$. In particular, I assume the distribution of signal noise is normal:

$$e_{ijkmt} \overset{iid}{\sim} N(0, \sigma^2_h) \tag{7}$$

The key parameter is σ^2_h, which can be thought of as a measure of price transparency (or opacity). The precision of price signals may be different for hospital versus non-hospital providers, therefore σ^2_h is indexed by h, an indicator for whether the provider j is a hospital.

Using Bayes' rule, individuals' posterior beliefs about price, $\widetilde{p^{OOP}_{ijkmt}}$, are also normally distributed. The mean of the posterior (i.e. expected price) given the individual's signal is given by

$$\mathbb{E}\left[\widetilde{p^{OOP}_{ijkmt}}\right] = w_{ikmt}(p^{OOP}_{ijkmt} + e_{ijkmt}) + (1 - w_{ikmt})\bar{p}^{OOP}_{kmt} \tag{8}$$

where the weight given to the signal is defined as

$$w_{ikmt} = \frac{\bar{s}^2_{kmt}}{\bar{s}^2_{kmt} + \sigma^2_h} \tag{9}$$

[48] The true distribution of prices is approximately normal. See Figure 2b.

[49] See, for instance, "How to Research Health Care Prices," *Wall Street Journal*, December 4, 2009.

If $\sigma_h^2 = 0$ then $w_{ikmt} = 1$ and individuals know true prices. Conversely, if $\sigma_h^2 \to \infty$ then $w_{ikmt} \to 0$, implying that individuals place no weight on the price signals. In this way, the prior is important because it disciplines individual's beliefs about price—if individuals receive very noisy signals than they effectively ignore prices.[50] In Section 5.1, I present an alternative model in which individuals have an uninformative prior and take price signals as given.

Using the assumption that the prior and signal are normally distributed, the variance of posterior beliefs is

$$Var\left[\widetilde{p_{ijkmt}^{OOP}}\right] = w_{ikmt}\sigma_h^2 \tag{10}$$

When individuals do not use the price transparency website, I assume they form beliefs about utility, $\widetilde{u_{ijkmt}}$, and choose the provider that maximizes expected utility. In particular, the expected utility of risk neutral individuals is

$$\begin{aligned}\mathbb{E}\left[\widetilde{u_{ijkmt}}\right] &= -\gamma_i \mathbb{E}\left[\widetilde{p_{ijkmt}^{OOP}}\right] + \underbrace{\alpha_1 d_{ij} + \alpha_2 d_{ij}^2 + \alpha_3 r_{ijt} + \xi_{jM} + \beta \mathbf{x}_{ikmt} h_j}_{\delta_{ijkmt}} + \varepsilon_{ijkmt} \\ &= -\gamma_i w_{ikmt}(p_{ijkmt}^{OOP} + e_{ijkmt}) + \delta_{ijkmt} + \varepsilon_{ijkmt}\end{aligned} \tag{11}$$

The second line follows from the fact that $(1 - w_{ikmt})\bar{p}_{kmt}^{OOP}$ is a constant that is the same across choices, and thus can be differenced out.

Focusing on the component of utility that is due to price, it is useful to clarify what is known by the individual and what is known by the researcher. The individual knows her price sensitivity, γ_i, and signal, $p_{ijkmt}^{OOP} + e_{ijkmt}$, but not the true price. However, the researcher observes the true price, p_{ijkmt}^{OOP}, but not the signal noise, e_{ijkmt}, or the individual's price sensitivity. The prior distribution is known by both the researcher and the individual.

Therefore, the observed choice probabilities from the researcher's perspective is given by

$$\begin{aligned}&s_{ijkmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt} | \vartheta_{ikmt} = 0) = \\ &\int_{\gamma_i}\int_{\mathbf{e}_{ikmt}} \frac{\exp(-\gamma_i w_{ikmt}(p_{ijkmt}^{OOP} + e_{ijkmt}) + \delta_{ijkmt})}{\sum_{j' \in \mathcal{N}_{kmt}} \exp(-\gamma_i w_{ikmt}(p_{ij'kmt}^{OOP} + e_{ij'kmt}) + \delta_{ij'kmt})} f(\mathbf{e}_{ikmt}) f(\gamma_i) d^J \mathbf{e}_{ikmt} d\gamma_i\end{aligned} \tag{12}$$

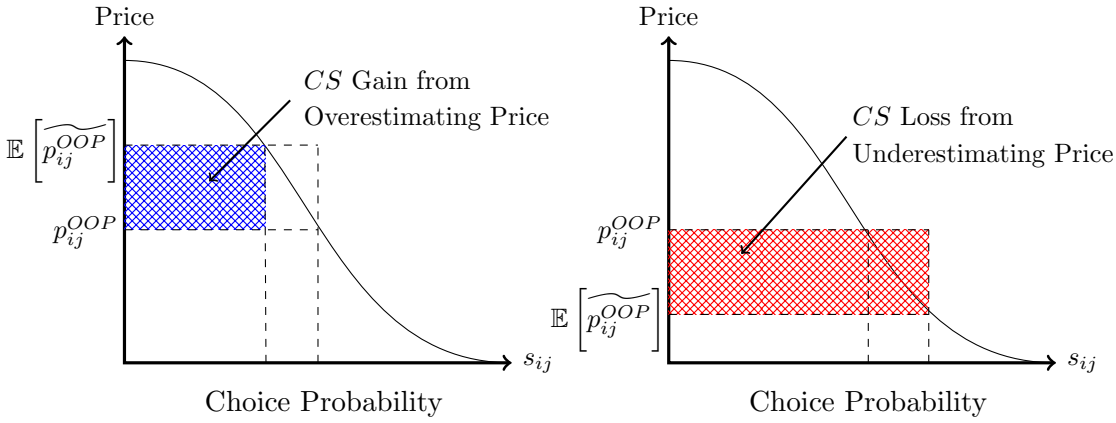
where $\vartheta_{ikmt} = 0$ indicates that the individual did not use the website and is uninformed about prices. It is worth noting that the vector of signal noise, $\mathbf{e}_{ikmt}$, has the same number of elements as $\mathcal{N}_{kmt}$. Therefore, computing the expectation over individual beliefs requires evaluating a potentially high dimensional integral, complicating the estimation strategy. I address this issue in Section 3.4.

The calculation of expected consumer surplus must take into account that, from the per-

[50] This can be seen formally by noting that the choice probability depends only on non-price attributes in the limit:

$$\lim_{\sigma_h^2 \to \infty} \mathbb{E}_e\left[\frac{\exp(-\gamma_i w_{ikmt}(p_{ijkmt}^{OOP} + e_{ijkmt}) + \delta_{ijkmt})}{\sum_{j' \in \mathcal{N}_{kmt}} \exp(-\gamma_i w_{ikmt}(p_{ij'kmt}^{OOP} + e_{ij'kmt}) + \delta_{ij'kmt})}\right] = \frac{\exp(\delta_{ijkmt})}{\sum_{j' \in \mathcal{N}_{kmt}} \exp(\delta_{ij'kmt})}$$

Figure 4
Consumer Surplus when Expected Price Differs from Actual Price



Notes: Blue shaded region shows the gain in consumer surplus relative to expected consumer surplus due to price being less than expected. Red region shows the loss in consumer surplus from price being more than expected. Note that there is a "winner's curse" and the expected loss is larger.

spective of the individual, the expected price, $\mathbb{E}[\widetilde{p_{ijkmt}^{OOP}}]$, may differ from true price, p_{ijkmt}^{OOP}. Train (2015) formalizes the calculation of consumer surplus when individuals misperceive product attributes. In particular, individual's expected ex-post consumer surplus includes a standard term (Small and Rosen 1981) as well as a term that captures the loss from incorrect beliefs:

$$\begin{aligned} CS_{ikmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt} | \vartheta_{ikmt} = 0) = & \underbrace{\frac{1}{\gamma_i} \log \left(\sum_{j \in \mathcal{N}_{kmt}} \exp(-\gamma_i w_{ikmt}(p_{ijkmt}^{OOP} + e_{ijkmt}) + \delta_{ijkmt}) \right)}_{\text{CS evaluated at beliefs of prices}} \\ & + \underbrace{\sum_{j \in \mathcal{N}_{kmt}} \left[\mathbb{E}\left[\widetilde{p_{ijkmt}^{OOP}} \right] - p_{ijkmt}^{OOP} \right] s_{ijkmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt} | \vartheta_{ikmt} = 0)}_{\text{CS gain/loss from incorrect beliefs}} \end{aligned} \tag{13}$$

The second term is the average difference between expected price and true price, weighted by choice probabilities. The "bill shock" for each option is given by

$$\mathbb{E}\left[\widetilde{p_{ijkmt}^{OOP}} \right] - p_{ijkmt}^{OOP} = w_{ikmt} e_{ijkmt} + (1 - w_{ikmt})(\bar{p}_{kmt}^{OOP} - p_{ijkmt}^{OOP}) \tag{14}$$

In general, individuals are more likely to choose a provider they falsely believe to be inexpensive, creating a situation similar to a "winner's curse". This can be seen in Figure 4 which presents two situations, one in which expected price is greater than actual price and one in which expected price is less than actual price. Believing an option to be inexpensive (i.e. receiving a low e_{ijkmt}) results in a higher choice probability, increasing the expected loss from incorrect beliefs.

3.3 Website Usage

In this section I develop a model in which individuals choose to use the price transparency website if it is available. The model seeks to recover information about which individuals use the price transparency website. In Section 7, the estimates from this selection model are used to simulate website usage under counterfactual scenarios such as increased cost sharing.

Although I argue that that the availability of the website is plausibly exogenous, it is not random which consumer use the website conditional on it being available. There may be search moral hazard—consumers with the least to gain from using the website choose not to use it. I assume that individuals evaluate the expected gain in consumer surplus from using the website and compare this to the cost. They then use the website if the net benefit is positive.

Before using the website, individuals use the available information and believe prices are distributed

$$p_{ijkmt}^{OOP} \overset{iid}{\sim} N\left(w_{ikmt}(p_{ijkmt}^{OOP} + e_{ijkmt}) + (1 - w_{ikmt})\bar{p}_{kmt}^{OOP}\,,\, \sigma_h^2 w_{ikmt}\right) \tag{15}$$

With price uncertainty, the ex-ante consumer surplus from the individual's perspective is determined by evaluating Equation 13 at all possible prices.[51] Therefore, given individuals' beliefs, the expected consumer surplus is

$$\frac{1}{\gamma_i}\log\left(\sum_{j\in\mathcal{N}_{kmt}}\exp(-\gamma_i\mathbb{E}\left[\widetilde{p_{ijkmt}^{OOP}}\right] + \delta_{ijkmt})\right) + \int_{\widetilde{\mathbf{p}_{ikmt}^{OOP}}'}\sum_{j\in\mathcal{N}_{kmt}}\left(\mathbb{E}\left[\widetilde{p_{ijkmt}^{OOP}}\right] - \widetilde{p_{ijkmt}^{OOP}}'\right)s_{ijkmt}g(\widetilde{\mathbf{p}_{ikmt}^{OOP}}')d^J\widetilde{\mathbf{p}_{ikmt}^{OOP}}' \tag{16}$$

where $g(\widetilde{\mathbf{p}_{ikmt}^{OOP}}')$ is the joint distribution of beliefs determined by the individual's prior and signals following Equation 15.

In order to evaluate the expected gain from using the website, the individual must compare Equation 16 with the expected consumer surplus after using the website. If individuals use the website, they can re-optimize. In addition, they will no longer be surprised by the bill. Therefore, given individual's beliefs, expected consumer surplus with price information is given by:

$$\frac{1}{\gamma_i}\int_{\widetilde{\mathbf{p}_{ikmt}^{OOP}}}\log\left(\sum_{j\in\mathcal{N}_{kmt}}\exp(-\gamma_i\widetilde{p_{ijkmt}^{OOP}} + \delta_{ijkmt})\right)g(\widetilde{\mathbf{p}_{ikmt}^{OOP}})d^J\widetilde{\mathbf{p}_{ikmt}^{OOP}} \tag{17}$$

The difference between Equation 17 and Equation 16 is the benefit from using the website. Since there does not exist a closed form expression, I derive an approximation using a second-

[51] I assume individuals evaluate the benefit of using the website prior to knowing idiosyncratic shocks.

order multivariate Taylor series around the expectation. This approach is necessary since it is computationally infeasible to use simulation-based methods.[52] See Appendix A for derivation and discussion about the accuracy of this approach. Using this approximation, the expected benefit of using the website, b_{ikmt}, is

$$b_{ikmt} \approx \frac{\gamma_i w_{ikmt} \sum_{j \in \mathcal{N}_{kmt}} \sigma_h^2 \left[\exp(-\gamma_i \mathbb{E}\left[\widetilde{p_{ijkmt}^{OOP}}\right] + \delta_{ijkmt}) \Phi_{ijkmt} \right]}{2 \left[\sum_{j \in \mathcal{N}_{kmt}} \exp(-\gamma_i \mathbb{E}\left[\widetilde{p_{ijkmt}^{OOP}}\right] + \delta_{ijkmt}) \right]^2} \tag{18}$$

where $\Phi_{ijkmt} \equiv \sum_{j' \in \mathcal{N}_{kmt} \setminus j} \exp(-\gamma_i \mathbb{E}[\widetilde{p_{ij'kmt}^{OOP}}] + \delta_{ij'kmt})$.

Unlike Equation 17, the interpretation of the closed-form expression above is relatively straightforward. Holding beliefs about prices fixed, an increase in price uncertainty, as measured by σ_h^2, increases the value of using the website. Similarly, an increase in price dispersion affects w_{ikmt}, also increasing the value of using the website. Note that the benefit of using the website also depends on γ_i, the individual-specific price sensitivity.

Now I turn to the cost of using the website. In practice, the website is free to use and only takes a few minutes. However, there may be large non-pecuniary costs. In 2007, when the website started, only 58 percent of New Hampshire households had high speed internet.[53] In addition, many individuals were likely unaware of the website and had to be motivated enough to discover the website on their own.

I assume cost has both an observable component, which is a function of individual characteristics $\mathbf{x}_{ikmt}$, as well as an unobservable component, ν_{ikmt}. Observable characteristics include age categories, sex, income, eduction, Charlson Comorbidity Index, emergency indicator, and year indicators in order to account for the fact that more individuals may hear about the website over time, reducing the implicit cost.[54] I also include a constant.

Individuals use the website if the net benefit is positive

$$\underbrace{\theta b_{ikmt}}_{\text{Website Benefit}} - \underbrace{\phi \mathbf{x}_{ikmt} + \nu_{ikmt}}_{\text{Website Cost}} > 0 \tag{19}$$

I assume that the distribution of ν_{ikmt} is distributed i.i.d. type 1 extreme value (with normalized variance). Therefore, the observed probability that individual i uses the website for

[52] Numerically integrating the expression by simulating draws for each price and then averaging over the draws is computationally expensive given the high dimensionality of $\widetilde{\mathbf{p}_{ikmt}^{OOP}}$. In addition, $\widetilde{\mathbf{p}_{ikmt}^{OOP}}$ is itself a function of latent variables (i.e. $\mathbf{e}_{ikmt}$). For these reasons, a closed form expression for b_{ikmt} is necessary in practice.

[53] See "State of New Hampshire Broadband Action Plan," New Hampshire Department of Resources and Economic Development & Telecommunications Advisory Board, June 30, 2008.

[54] This also accounts for the fact that more consumers have broadband internet over time.

the price of procedure m at time t takes the logistic form:

$$\vartheta_{ikmt} = \frac{\exp(\theta b_{ikmt} - \phi \mathbf{x}_{ikmt})}{1 + \exp(\theta b_{ikmt} - \phi \mathbf{x}_{ikmt})} \tag{20}$$

where θ and ϕ are parameters to be estimated. Note that θ can be interpreted as the marginal utility of income.

I have access to website traffic logs that provide an estimate of the number of individuals that decide to use the website for each procedure in each month. Since it is not possible to link website usage to individual claims, it is necessary to connect the model's predicted individual website usage to overall website usage in each month for each procedure. Conditional on the parameters, the average predicted website usage for a procedure-month is given by

$$\vartheta_{mt} = \frac{1}{n_{mt}} \sum_{i \in \mathcal{I}_{mt}} \vartheta_{ikmt} \tag{21}$$

where n_{mt} is the number of individuals receiving procedure m in month t.

3.4 Joint Estimation of Demand

Next, I describe the procedure used to estimate the parameters of the demand model. There are two pieces of the demand model, provider choice and website usage, which are estimated jointly. I begin by discussing the likelihood function and the Bayesian estimation procedure. I make use of Markov chain Monte Carlo (MCMC) methods to address the estimation challenges that arise due to the fact that beliefs about prices are unobserved. I then sketch the identification argument and discuss how the model takes advantage of quasi-experimental variation to separately identify uncertainty about prices and underlying consumer preferences.

Likelihood Function

The likelihood function is directly based on the structural equations describing individual provider choices and website usage. The first component of the likelihood function is the probability of choosing the provider that was actually chosen. Equation 3 and Equation 12 are the conditional choice probabilities with and without price information. Therefore, the unconditional choice probability is:

$$s_{ijkmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt} | \vartheta_{ikmt}) = \begin{cases} s_{ijkmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt} | \vartheta_{ikmt} = 0) & \text{if website is not available} \\ \vartheta_{ikmt} \cdot s_{ijkmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt} | \vartheta_{ikmt} = 1) & \text{if website is available} \\ \quad + (1 - \vartheta_{ikmt}) \cdot s_{ijkmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt} | \vartheta_{ikmt} = 0) & \end{cases} \tag{22}$$

If the website is not available for procedure m at time t, either because it is prior to March

2007 or because the procedure is never on the website, then the consumer has uncertainty about prices and choice probabilities are given by Equation 12. If the website is available for procedure m at time t, the consumer is informed about prices if the website is actually used. Therefore, choice probabilities are given by a mixture between Equation 3 and Equation 12, where the mixture weights are determined by the predicted probability of using the website, given by Equation 20.

The second component of the likelihood function is the probability of actual website traffic for each procedure-month given predicted website usage. The likelihood of website usage for procedure m in month t takes the following binomial form[55]

$$\frac{n_{mt}!}{\mathcal{V}_{mt}!(n_{mt}-\mathcal{V}_{mt})!}\left(\vartheta_{mt}\right)^{\mathcal{V}_{mt}}\left(1-\vartheta_{mt}\right)^{n_{mt}-\mathcal{V}_{mt}} \tag{23}$$

where n_{mt} is the number of individuals receiving procedures and $\mathcal{V}_{mt}$ is the observed search traffic for a given procedure-month.

Therefore, the likelihood function is

$$\begin{aligned}\mathcal{L}(\Theta) = &\underbrace{\sum_{t}\sum_{k\in\mathcal{K}}\sum_{m\in\mathcal{M}}\sum_{i\in\mathcal{I}_{kmt}}\sum_{j\in\mathcal{N}_{kmt}}\left[s_{ijkmt}(\mathcal{N}_{kmt},\mathbf{p}_{kmt}|\vartheta_{ikmt})\right]^{y_{ijkmt}}}_{\text{Likelihood of Observed Provider Choices}} \\ &\times\underbrace{\sum_{t}\sum_{m\in\mathcal{M}}\frac{n_{mt}!}{\mathcal{V}_{mt}!(n_{mt}-\mathcal{V}_{mt})!}\left(\vartheta_{mt}\right)^{\mathcal{V}_{mt}}\left(1-\vartheta_{mt}\right)^{n_{mt}-\mathcal{V}_{mt}}}_{\text{Likelihood of Observed Search Traffic}}\end{aligned} \tag{24}$$

where y_{ijkmt} is an indicator for the observed choice.

MCMC Estimation

To estimate the model, I use a Markov chain Monte Carlo (MCMC) estimator to simulate the posterior distribution of Θ. This approach helps circumvent the computational curse of dimensionality caused by the fact that beliefs, specifically the signal noise draws, are a high-dimensional nuisance parameter.[56] Rather than compute high dimensional integrals in order to find the expectation over $\mathbf{e}_{ikmt}$ and calculate the likelihood, the MCMC estimator samples the

[55] This is an approximation. The predicted probability of website usage within a procedure-month is not identically distributed across individuals, therefore the sum of these Bernoulli distributed variables takes a poisson binomial distribution. Since calculating the density of the poisson binomial distribution is computationally expensive, I approximate this distribution with a binomial distribution. For an analysis of the accuracy of this approximation see Ehm (1991).

[56] The standard estimation strategy is to use simulation methods and draw from $f(\mathbf{e}_{ikmt})$ and $f(\gamma_i)$, calculate the log-likelihood for each draw, and average over the results to obtain the simulated log-likelihood for a given value of the parameters. This simulated maximum likelihood approach is computationally infeasible due to the high dimensionality of $\mathbf{e}_{ikmt}$. A very large number of draws from a multivariate distribution would be required in order to accurately approximate the log-likelihood at each iteration.

parameter space conditional on the data.

I take advantage of recent advances in Bayesian estimation and use a variant of MCMC known as Hamiltonian Monte Carlo (HMC) No-U-Turn Sampler (NUTS).[57] This approach, developed by Hoffman and Gelman (2014), uses the gradient of the log posterior density to more efficiently sample the posterior distribution.[58] Relative to standard MCMC algorithms such as Metropolis-Hastings and Gibbs sampling, this approach is known to converge significantly faster for high-dimensional problems, making it well suited for a situation with alternative-specific unobservables. In addition, it does not necessitate the use of conjugate priors, allowing for more flexible modeling assumptions.

For the purposes of estimation, I reformulate the model in terms of a simplified likelihood that is augmented with a set of priors. The likelihood uses the choice probability conditional on unobservables while the priors describe the distribution of website traffic and the distribution of the unobservables. This version of the model is described in more detail in Appendix B.

It is important to note that using this Bayesian hierarchical model for estimation does not impose additional assumptions since I use uninformative priors for all of the structural parameters. In supplemental material, I examine a simplified version of the model with a small choice set and show that the results obtained via simulated maximum likelihood are very similar to those obtained via MCMC estimation.[59] The use of MCMC is primarily motivated by the fact that it is computationally attractive.

In order to estimate the posterior distribution of Θ, the algorithm uses the following approach. At iteration n, the MCMC algorithm returns parameter estimates $\Theta^{(n)}$. As starting values, I use parameter estimates from a standard multinomial logit (see Section 5.1).[60] However, to ensure that initial values do not influence the resulting posterior distribution, samples drawn during a warm-up period are discarded. The remaining collection of samples, $(\Theta^{(1)}, ..., \Theta^{(N)})$, approximately converge to the distribution of the posterior. I report the mean and standard deviation of these samples in the results.

Identification Intuition

Without variation in consumers' information set, it is difficult or impossible to separately identify price sensitivity and the degree of price uncertainty, i.e. the observed choices from a population with low price sensitivity are potentially observationally equivalent to the observed choices from a population with high price sensitivity but limited information about prices. An alternative

[57] This algorithm is implemented in the Stan programming language, which I use to automatically compute gradients and estimate the model. See Carpenter et al. (2016).

[58] In particular, HMC uses gradient information to avoid random walk behavior and sensitivity to correlated sampling. However, this approach suffers from the fact that it must be manual tuned for a given problem and can double back on the parameter space, decreasing efficiency. Hoffman and Gelman (2014) introduce NUTS to address these issues.

[59] Code and Monte Carlo results are available on my website.

[60] For parameters that are not included in the multinomial logit, I use random starting values.

way to see this is to note that price uncertainty is closely related to classical measurement error, except rather than the researcher having noisy information it is the individual decision-maker. It is widely known that classical measurement error can cause biased and inconsistent parameter estimates. This is also true in non-linear models, and, in general, the variance of measurement error and the underlying parameters are not separately identified (e.g. Chen et al. 2011). I overcome this issue due to the fact that some consumers—those that use the price transparency website—do not observe prices with error.

To describe the source of identification, I begin by focusing on individuals with price information. Assuming the researcher can identify a subset of consumers that have price information, identification of demand parameters $(\sigma^\gamma, \bar{\gamma}, \rho, \alpha, \xi, \beta)$ follows the same argument as for the standard mixed logit model. Identification relies on variation in observed provider choices when the characteristics of the providers or the choice set differ. In particular, price sensitivity is identified from the fact that the price of a given provider varies depending on an individual's insurer, whether the individual is under the deductible, and year. In addition, the choice set of consumers varies over insurers, locations, and years. Substitution patterns help identify the variance of the random coefficient on price.

In order to illustrate how underlying tastes and the degree of price uncertainty are separately identified, it is useful to start by describing the ideal experiment. Consider a population that is randomly divided into a treatment group and control group. Although both groups have the same distribution of consumer preferences, the treatment group is given information about prices. If the treatment group appears more price sensitive than the control group, it must be due to the fact that the control group had noisy beliefs about prices. The extent to which individuals in the control group are less price sensitive provides information about the variance of signal noise. However, the mean bias of beliefs is not identified in the case in which all individuals choose an inside option. This is because if individuals underestimate or overestimate the price of *all* options in the choice set, observed choices do not change.

In this paper, I take advantage of a natural experiment in which a price transparency website was available for a subset of consumers. In contrast to the ideal experiment described above, individuals often did not use the website even when it was available. However, conditional on θ and ϕ, the parameters that predict website usage, the observed choices of individuals who used the website when it is available can be compared to the observed choices of similar individuals who would have used the website if it were available. For this population, the identification argument is the same as in the ideal experiment.

Finally, I turn to identification of the website usage parameters (i.e. θ and ϕ). In principle, these parameters can be identified by observing which individuals appear to be more price sensitive when the website is available relative to when the website is not available. In practice, identification is facilitated by using the website traffic data and exploiting variation in website traffic across months and across procedures. In particular, correlation between consumers' benefit of using the website and website usage helps identify θ, while correlation between observed

characteristics of consumers and website usage helps identify ϕ.

4 Bargaining between Providers and Insurers

In a variety of markets, prices are determined through bilateral bargaining. For instance, wholesalers negotiate prices with retailers and unions negotiate wages with employers. Although there is a growing empirical literature that seeks to shed light on how outcomes are determined in these markets, there is little evidence about how information frictions, in particular price transparency, affects equilibrium outcomes when prices are negotiated.

In this section, I examine how price transparency affects bargaining between providers and insurers. I use an approach that is similar to Gowrisankaran et al. (2015), who assume consumers are perfectly informed about prices. In contrast, I show how consumer information frictions affect "gains-from-trade" for both the providers and insurers. Prices are then determined as a Nash equilibrium of bilateral Nash bargaining problems (Horn and Wolinsky 1988).

Using the estimates from the demand model given in the previous section, I use the supply-side model to estimate the marginal cost of each procedure at each provider. These estimates are then used in Section 6 and Section 7 to simulate negotiated prices under various counterfactual scenarios.

4.1 Bargaining Model

I now present the model of bilateral bargaining between medical providers and insurers incorporating consumer price uncertainty. In each year, insurer k negotiates the price of procedure m with each provider in the insurer's network, $j \in \mathcal{N}_{kmt}$.[61] For the analysis, I assume that each provider negotiates independently.[62] I also take the set of providers $\mathcal{J}$ and networks $\mathcal{N}_{kmt}$ as given.[63]

I start by describing the gains from trade for provider j when contracting with insurer k. The provider's profit from individual i enrolled in insurer k receiving procedure m at time t is given by

$$\Pi^{J}_{ijkmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt}|\vartheta_{ikmt}) = s_{ijkmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt}|\vartheta_{ikmt})[p_{jkmt} - mc_{jkmt}] \tag{25}$$

where mc_{jkmt} is the marginal cost of the procedure and $s_{ijkmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt}|\vartheta_{ikmt})$ is the choice

[61] While the previous literature has assumed that insurers negotiate over a price index, I allow insurers to negotiate over the visit price of each procedure $m \in \mathcal{M}$. For outpatient procedures, I believe this to be a more realistic assumption. Note that negotiated prices for a visit may change due to lower individual procedure prices or different supplemental procedures. I do not distinguish between these mechanism.

[62] In contrast, Gowrisankaran et al. (2015) and Ho and Lee (2016) allow hospitals that are part of a system to jointly negotiate with insurers. I am unable to link anonymous provider identifiers to ownership data, and therefore cannot examine hospital systems. To my knowledge, the medical imaging providers in the sample tend to be independently owned.

[63] It is possible that a large increase in price transparency increases entry of low cost outpatient facilities, leading to larger cost savings for consumers. I assume entry and exit are exogenously determined.

probability which depends on whether the individual is informed about prices, ϑ_{ikmt}. Without a contract with the insurer, the provider's profit from a given individual is zero. Therefore, the gains from trade are simply the provider profit summed over individuals and procedures.

Next, I turn to the insurer's gains from trade. For a given individual, the reimbursement amount paid by the insurer across all providers is

$$TC_{ikmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt}|\vartheta_{kmt}) = \sum_{j \in \mathcal{N}_{kmt}} p_{jkmt}(1 - c_{ikmt}) s_{ijkmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt}|\vartheta_{ikmt}) \tag{26}$$

Following Gowrisankaran et al. (2015), I also assume that insurers internalize the consumer surplus of their enrollees. When consumers are informed about prices, consumer surplus takes the standard form (see Equation 4). However, insurers are aware when consumers have uncertainty about prices, and consumer surplus includes a term that accounts for incorrect beliefs. In particular, consumer surplus is given by Equation 13.

The insurer's surplus generated by an individual visit is then the weighted sum of consumer surplus and total cost

$$\Pi^K_{ikmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt}|\vartheta_{kmt}) = \zeta CS_{ikmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt}|\vartheta_{kmt}) - TC_{ikmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt}|\vartheta_{kmt})) \tag{27}$$

where ζ is a parameter reflecting the relative weight on consumer surplus. The insurer gains from trade for an enrollee visit are the difference between the surplus generated with and without provider j in the network:

$$\Delta_j \Pi^K_{ikmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt}|\vartheta_{kmt}) = \Pi^K_{ikmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt}|\vartheta_{kmt}) - \Pi^K_{ikmt}(\mathcal{N}_{kmt} \backslash j, \mathbf{p}_{kmt}|\vartheta_{kmt}) \tag{28}$$

Equation 27 and Equation 28 can be thought of as a stylized approach to modeling the insurer's profit function. The consumer surplus of the insurer's enrollees enters the insurer's surplus function since a larger consumer surplus implies that the insurer can charge higher premiums to consumers, generating profit for the insurer. In contrast, Ho and Lee (2016) explicitly model demand for insurance and insurer competition in order to derive an expression for provider and insurer profits that accounts for the fact that consumers may switch insurers to access their preferred providers. I lack data on insurance premiums, and therefore, I cannot explicitly model insurer competition. Consistent with reduced-form results in my previous work, I assume that price transparency does not affect insurance choice.[64]

I now define the Nash bargaining problem that determines equilibrium prices. Importantly, the equilibrium price at a given provider, p_{jkmt}, also depends on the price of the procedure at other providers. Following Horn and Wolinsky (1988) and the previous empirical bargaining literature, I assume that equilibrium prices are those that solve the Nash bargaining solution

[64] In Brown (2016), I examine whether insurance enrollment changed after the introduction of the price transparency website and do not find a statistically significant effect.

given the equilibrium prices at other providers, $\mathbf{p}^*_{kmt} \backslash p_{jkmt}$. In other words, a hypothetical disagreement is assumed to not affect other prices.[65] Extending Rubinstein (1982), Collard-Wexler et al. (2014) rationalize this model by showing conditions under which the Nash-in-Nash solution is equivalent to a non-cooperative extensive form game with alternating offers.

Therefore, the Nash bargaining solution is the negotiated prices for each provider-insurer-procedure triple in a given year, p^*_{jkmt}, that satisfy

$$p^*_{jkmt} = \arg\max_{p_{jkmt}} \left(\underbrace{\sum_{i \in I_{kmt}} \mathbb{E}_e \left[\Pi^J_{ijkmt}(\mathcal{N}_{kmt}, p_{jkmt}, \mathbf{p}^*_{kmt} \backslash p_{jkmt} | \vartheta_{ikmt}) \right]}_{\text{Provider gains from trade}} \right)^{\tau} \times \left(\underbrace{\sum_{i \in I_{kmt}} \mathbb{E}_e \left[\Delta_j \Pi^K_{ikmt}(\mathcal{N}_{kmt}, p_{jkmt}, \mathbf{p}^*_{kmt} \backslash p_{jkmt} | \vartheta_{kmt}) \right]}_{\text{Insurer gains from trade}} \right)^{1-\tau} \tag{29}$$

where the gains from trade are summed over all individuals enrolled in insurer k receiving procedure m in year t, I_{kmt}. The Nash bargaining weight is $\tau \in [0, 1]$. Since insurers and providers do not know the price signals that consumers will receive, both take the expectation over consumer beliefs.[66]

Empirical models of bilateral bargaining in vertical markets generally assume that the negotiating parties do not have asymmetrical information about the relevant gains from trade.[67] I do not deviate from this assumption.[68] In the model presented in this section, price transparency indirectly affects equilibrium prices since changes in consumer behavior affect the gains from trade. I assume that the price transparency website, which was targeted towards consumers, did not directly affect the information set of the providers or insurers. Further research is needed to understand whether price transparency affects provider-insurer bargaining directly.[69]

[65] This is similar to a contract equilibrium (Cremer and Riordan 1987).

[66] The providers and insurers know the variance of the price signals, σ^2_h. In practice, I simulate beliefs by drawing from the distribution of $\mathbf{e}_{ikmt}$, computing each term, and then averaging over the draws.

[67] To my knowledge, all empirical models of business-to-business bargaining assume perfect information. Note that a sizable theoretical literature, starting with Samuelson (1984), examines bargaining with asymmetric information.

[68] Insurers and providers have uncertainty about the draws that determine consumer beliefs about prices, but have full information about the expected gains-from-trade for all participants.

[69] In the context of hospital-supplier bargaining, Grennan and Swanson (2016) find reduced-form evidence that price transparency affects negotiated prices in a way that is consistent with a theoretical model of bargaining under asymmetric information. It is possible that a similar mechanism is important for provider-insurer bargaining.

4.2 First Order Condition of the Bargaining Problem

I now turn to the equilibrium of the bargaining model. The first order condition of the bargaining problem given by Equation 29 implies that equilibrium prices are determined by marginal cost plus a margin:[70]

$$p_{jkmt} = mc_{jkmt} + \left(-\frac{1-\tau}{\tau} \frac{\frac{\partial}{\partial p_{jkmt}} [\sum_{i \in I_{km}} \mathbb{E}_e \Pi^K_{ikmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt} | \vartheta_{kmt})]}{\sum_{i \in I_{km}} \mathbb{E}_e \Delta_j \Pi^K_{ikmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt} | \vartheta_{kmt})} - \frac{\frac{\partial}{\partial p_{jkmt}} [\sum_{i \in I_{km}} \mathbb{E}_e s_{ijkmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt} | \vartheta_{ikmt})]}{\sum_{i \in I_{km}} \mathbb{E}_e s_{ijkmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt} | \vartheta_{ikmt})} \right)^{-1} \tag{30}$$

I present further detail, including the derivation of $\frac{\partial}{\partial p_{jkmt}} [\sum_{i \in I_{km}} \mathbb{E}_e \Pi^K_{ikmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt} | \vartheta_{kmt})]$, in Appendix C. Given that there are many providers in each network, a single price change has a minimal effect on individuals' prior about the distribution of prices. For tractability, I assume that providers and insurers do not take changes in the prior into account, and therefore hold the prior fixed when solving for the first order condition.

The Nash-in-Nash bargaining model nests the standard Bertrand-Nash pricing assumption when $\tau = 1$. In this case, providers unilaterally set prices and an increase in price transparency that makes demand more elastic leads to lower prices in equilibrium.

In the market for privately-provided health care, insurers negotiate their own rates with each provider that are thought to be lower than what a Bertrand-Nash pricing assumption would imply.[71] This corresponds to the case in which $\tau < 1$. Therefore, it is important to also understand how price transparency affects insurers' incentive to negotiate lower prices.

There are multiple channels through which consumer price transparency can affect equilibrium outcomes in the bargaining model. First, price transparency affects the incentives of the provider. This can be seen by noting that the provider gains-from-trade are a function of the choice probabilities, $s_{ijkmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt} | \vartheta_{ikmt})$, which depend on website usage (ϑ_{ikmt}). In general, demand is more elastic when more consumers are informed about prices. Under a Bertrand-Nash pricing assumption, this implies that providers will choose lower prices when more consumers are informed. Similarly, in the bargaining framework, providers have less incentive to negotiate high prices.

The effect of price transparency on insurers' incentives are more complicated. Price transparency affects insurer cost since consumers tend to switch to lower cost providers. This can be seen by noting that $TC_{ikmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt} | \vartheta_{kmt})$ depends on the choice probabilities. Price transparency also affects the consumer surplus of the insurer's enrollees, $CS_{ikmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt} | \vartheta_{kmt})$, since individuals can switch to lower cost providers and are not surprised by the bill (see Equation 4 and Equation 13).

[70] For simplicity, I omit the $*$ used to indicate equilibrium outcomes.

[71] Under Bertrand-Nash pricing, providers would be able to set prices unilaterally. In the absence of price information, the the effective demand elasticity is about -0.06 on average, implying extremely large markups.

It is important to note that price transparency does not always increase the incentive for insurers to negotiate low prices with all providers in their network. This is because insurers may be willing to have high priced providers in their network if they know that consumers will not choose these options. In other words, when more consumers are informed about prices, insurers find it easier to steer consumers to low-priced providers, and they take this into account when negotiating prices. Therefore, when demand becomes more elastic due to increased price transparency, it is not always the case that all prices decline.

4.3 Estimation and Identification of Bargaining Model

In this section, I describe the estimation strategy for the bargaining model. Following the previous empirical bargaining literature, I parameterize marginal cost and use the bargaining first-order condition to derive a moment condition which is than estimated using GMM.

The marginal cost of a visit is assumed to vary by procedure, provider, and year and is additively separable taking the form

$$mc_{jkmt} = \eta_j + \eta_m + \eta_t + \varepsilon^{MC}_{jkmt} \tag{31}$$

where η_j are provider fixed effects and η_m are procedure fixed effects. Health care prices increased significantly over the six year period, therefore it is important to include year fixed effects, η_t. The unobservable component of marginal cost is ε^{MC}_{jkmt}. I assume providers have constant returns to scale.

Using the parameterized marginal cost above along with the first-order condition given by Equation 30, the marginal cost error is given by

$$\varepsilon^{MC}_{jkmt} = \eta_j + \eta_m + \eta_t - p_{jkmt} + \left(-\frac{1-\tau}{\tau} \frac{\frac{\partial}{\partial p_{jkmt}} \sum_{i \in I_{km}} [\zeta CS_{ikmt} - TC_{ikmt}]}{\sum_{i \in I_{km}} [\zeta \Delta_j CS_{ikmt} - \Delta_j TC_{ikmt}]} - \frac{\frac{\partial}{\partial p_{jkmt}} [\sum_{i \in I_{km}} s_{ijkmt}]}{\sum_{i \in I_{km}} s_{ijkmt}} \right)^{-1} \tag{32}$$

This is used to form the following moment condition:

$$\mathbb{E}[\varepsilon^{MC}_{jkmt} | Z_{jkmt}] = 0 \tag{33}$$

where Z_{jkmt} is a vector of variables assumed to be exogenous. The model assumes that the bargaining participants know mc_{jkmt}, including ε^{MC}_{jkmt}, implying that prices are potentially endogenous. Following the previous literature, I address this issue by including two instruments: predicted willingness-to-pay for each provider at mean price and predicted total provider quantity at mean price.[72] Although these instruments are correlated with price, it is assumed that they are uncorrelated with ε^{MC}_{jkmt}. The instrument set, Z_{jkmt}, also includes all marginal cost fixed effects.

[72] These are a similar set of instruments as those used by Gowrisankaran et al. (2015). They also include willingness-to-pay for the hospital system and willingness-to-pay per enrollee for each insurer.

Identification of parameters η, τ, and ζ follows from a similar argument as that presented in Gowrisankaran et al. (2015). The provider choice and website usage parameters from the demand model allow me to construct CS_{ikmt}, TC_{ikmt}, and s_{ijkmt}, as well as their derivatives with respect to price (these are given in Appendix C). In the bargaining model, these are treated like data. Variation in provider incentives (determined by s_{ijkmt} and $\partial s_{ijkmt}/\partial p_{jkmt}$) and insurer incentives (determined by CS_{ikmt}, TC_{ikmt}, $\partial CS_{ikmt}/\partial p_{jkmt}$, and $\partial TC_{ikmt}/\partial p_{jkmt}$) that can explain variation in prices identifies ζ and τ. This variation comes in part from the introduction of the price transparency website. The remaining price variation identifies the marginal cost fixed effects, η. Unlike Gowrisankaran et al. (2015), I take advantage of price variation across individual procedures. This provides an additional source of variation to identify ζ and τ.

5 Results

I now turn to the results. Section 5.1 discusses the results from the demand model, starting with results from a standard multinomial logit model used to motivate the main empirical specification. I then discuss the results from the full demand model (both provider choice and website usage). I also discuss results from an alternative demand model in which consumers do not have an informative prior about the distributions of prices. Section 5.2 presents the results from the supply model (bilateral bargaining).

5.1 Estimates from Demand Model

Estimates from Multinomial Logit Model

In order to examine the effect of the price transparency website, I start by estimating a naive demand model in which I interact the availability of the website and the price coefficient. I assume choice probabilities take the form

$$s_{ijkmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt}) = \frac{\exp(-\gamma_1 q_{mt} p^{OOP}_{ijkmt} - \gamma_2(1-q_{mt})p^{OOP}_{ijkmt} + \delta_{ijkmt})}{\sum_{j' \in \mathcal{N}_{kmt}} \exp(-\gamma_1 q_{mt} p^{OOP}_{ij'kmt} - \gamma_2(1-q_{mt})p^{OOP}_{ijkmt} + \delta_{ij'kmt})} \tag{34}$$

where q_{mt} is an indicator for whether procedure m is available on the website at time t. Therefore, γ_1 is the price coefficient when the website is available and γ_2 is the price coefficient when the website is not available. I also include δ_{ijkmt}, which contains the same non-price characteristics as in Equation 1. Since this simple approach does not model unobserved beliefs, estimation can be performed via maximum likelihood.

Table A-3 presents the coefficient estimates and standard errors from the simple logit model. The magnitude of the price coefficient is larger when consumers have access to the price transparency website, indicating that the website increases the effective demand elasticity of the

population. The difference is statistically significant at the 10 percent level.[73] Complimenting the results in Brown (2016), this provides further evidence that the website had a meaningful impact on consumer behavior.

Consistent with the prior literature on hospital demand, the travel distance is important for understanding consumer demand. I find evidence of non-linearities in travel preferences. The indicator for likely referrals is highly significant, indicating that physician can influence consumer behavior.[74] Finally, there is evidence that income and education affect whether consumers choose a hospital. The results suggest that higher income consumers are more likely to prefer hospitals, which tend to be expensive, rather than medical imaging centers.

It is important to note that the estimates from this model lack a straightforward interpretation and do not allow for the calculation of welfare. When consumers have perfect information, the price coefficient is often interpreted as the marginal utility of income. Since many individuals lack information about price, even when the website is available, γ_1 and γ_2 cannot be interpreted in this way. Given that there are still many uninformed consumers when the website is available, γ_1 is an underestimate of the true price sensitivity (and marginal utility of income). The full model is needed to recover individuals' underlying taste parameters, including price sensitivity, in order to evaluate counterfactuals and conduct welfare analysis.

Provider Choice and Website Usage Estimates from Baseline Model

Table 4 presents estimates for parameters of the full demand model. I focus on specification 1, which reflects the baseline model presented in Section 3. The first column reports the mean of the estimated posterior distribution of each parameter implied by the MCMC estimation procedure. The second column reports the standard deviation of the posterior distribution.[75]

The magnitude of mean price sensitivity, $\bar{\gamma}$, is much larger than the price coefficient in the simple logit model presented in the previous section. This reflects the fact that γ_i can now be interpreted as consumer's underlying price sensitivity when prices are known (i.e. underlying marginal utility of income). There is significant heterogeneity in the price sensitivity parameter, which is negatively correlated with consumer cost sharing. This implies that consumers with high price sensitivity select into generous insurance plans (i.e. those with lower cost sharing). The effect of other explanatory variables, including distance, the likely referral indicator, and hospital interactions, is largely consistent with the results from the logit model presented in

[73] I present results for a small sample for comparison with the results from the full specification. The difference is statistically significant at the 5 percent level for a larger sample. In addition, although the coefficient on price is not statisically significant when the website is not available in Table A-3, it becomes significant with a larger sample.

[74] Since the likely referral indicator is a constructed explanatory variable and is highly significant, there is worry it could bias other coefficients. As a robustness check, I also estimate the model without the likely referral indicator. In both the simple multinomial logit model and the full specification, the other coefficients are qualitatively consistent when the likely referral indicator is not included.

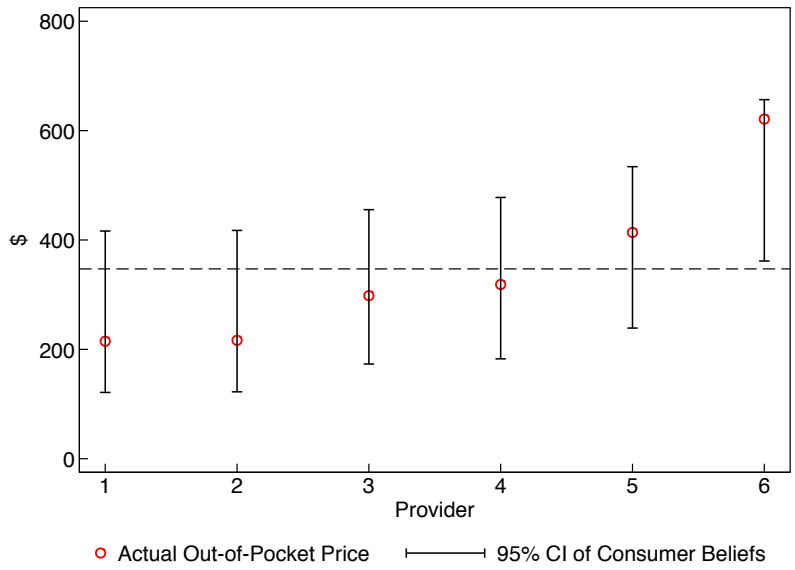
[75] Note that for explanatory variables that overlap with the simple logit model, the standard deviation of the parameter posterior distributions are very similar to the standard errors reported in Table A-3.

Table 4
MCMC Estimates for Demand Model

	Specification 1		Specification 2	
	Estimate	SD	Estimate	SD
Provider Choice Parameters				
OOP Price Mean ($-\bar{\gamma}$)	−0.0099	(0.0003)	−0.0094	(0.0002)
OOP Price SD (σ^{γ})	0.0003	(0.0001)	0.0001	(0.0001)
OOP Price × Cost Sharing (ρ)	−0.0092	(0.0004)	−0.0090	(0.0001)
Distance (α_1)	−0.0320	(0.0028)	−0.0345	(0.0027)
Distance squared (α_2)	0.0024	(0.0003)	0.0025	(0.0003)
Referral Indicator	2.053	(0.032)	2.268	(0.039)
Hospital×Age ≤18	−0.006	(0.161)	−0.021	(0.179)
Hospital×Age 19-35	0.060	(0.161)	0.058	(0.174)
Hospital×Age 36-50	−0.099	(0.155)	−0.113	(0.165)
Hospital×Age 51-64	−0.132	(0.161)	−0.150	(0.171)
Hospital×Male	−0.098	(0.065)	−0.114	(0.072)
Hospital×Income	0.013	(0.002)	0.014	(0.002)
Hospital×BA	−0.036	(0.004)	−0.039	(0.004)
Hospital×Charlson	0.055	(0.043)	0.070	(0.046)
Hospital×Emergency	0.550	(0.085)	0.620	(0.090)
Website Choice Parameters				
Benefit (θ)	0.026	(0.011)	0.003	(0.003)
Cost (ϕ)				
Constant	6.058	(6.441)	6.769	(5.957)
Age 19-35	−5.698	(6.466)	−6.224	(5.956)
Age 36-50	−5.248	(6.432)	−5.765	(5.935)
Age 51-64	−5.261	(6.426)	−5.737	(5.930)
Male	−0.163	(0.187)	−0.175	(0.186)
Income	0.015	(0.007)	0.013	(0.006)
BA	−0.013	(0.010)	−0.014	(0.010)
Charlson Comorbidity	0.001	(0.105)	−0.024	(0.106)
Outpatient Emergency	2.854	(1.158)	2.129	(0.933)
Year: 2007	0.326	(0.126)	0.323	(0.125)
Year: 2008	0.235	(0.108)	0.235	(0.107)
Year: 2009	0.267	(0.119)	0.284	(0.121)
Price Signal: Hospital ($\sigma_{h=1}$)	97.6	(6.3)	83.1	(5.9)
Price Signal: Non-Hospital ($\sigma_{h=0}$)	105.8	(5.6)	82.2	(5.7)
Observations	59,240		59,240	

Notes: Table shows the mean and standard deviation of the posterior distribution estimated via MCMC. Specification 1 refers to model in which consumers know the mean and variance of the price distribution and use this information to form a prior about prices. Specification 2 assumes consumers have an uninformative prior about prices. The provider-choice equation also includes provider-procedure group fixed effects (not shown). For the website choice model, the omitted year is 2010 and the omitted age group is ≤18.

Figure 5
Example of Beliefs for Uninformed Individuals



Notes: Chart shows an example choice situation selected from the data in which an individual is choosing between 6 providers. Confidence interval shows the distribution of beliefs given the estimated price uncertainty from specification 1 in Table 4.

Table A-3.

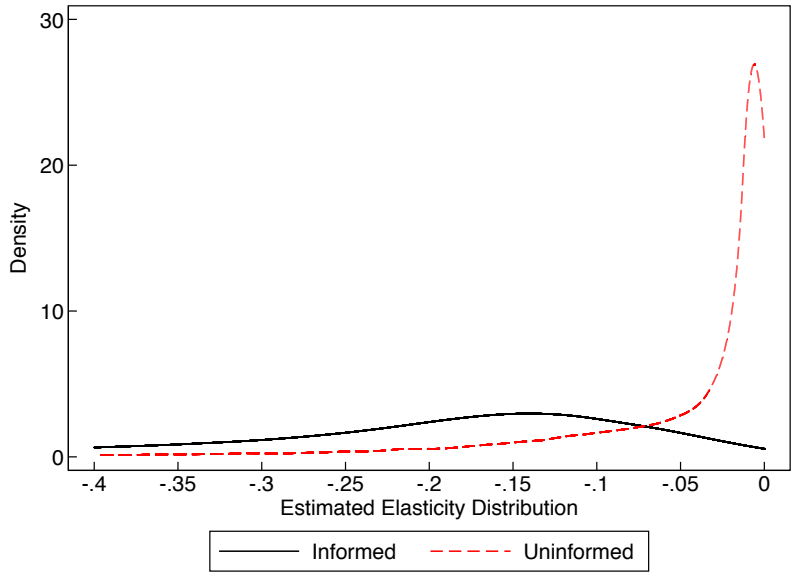
The estimated standard deviation of signal noise, σ_h, is shown at the bottom of Table 4. The estimates imply that, in the absence of price information, individuals have a large degree of uncertainty about prices. The standard deviation of signal noise is larger for non-hospital providers, suggesting that individuals have greater uncertainty about the price of medical imaging centers and other non-hospital providers.[76]

Given that the interpretation of σ_h is complex, it is useful to consider an example from the data. Figure 5 shows a sample individual choosing between six providers that range in price from about $200 to $650. Using the estimate of σ_h, the individual's beliefs about the price of each option can be simulated given different potential draws from the distribution of signal noise. The 95 percent confidence interval for these beliefs is shown for each option in the choice set. Beliefs range by over $200, implying that there is a non-trivial chance that the individual will believe the expensive options (such as option 5), are actually the least expensive. Also note that mean beliefs are not equal to the true price. This is due to the fact that each individual's prior causes shrinkage towards the mean.

I compare uninformed consumers beliefs about prices with the true price. On average, there is a 37 percent absolute difference between beliefs and true prices. The gap is even larger for individuals under the deductible—48 percent. Noisy beliefs about price effectively make demand

[76]However, note that zero lies within the 95 percent credible interval of the difference.

Figure 6
Distribution of Estimated Demand Elasticity
With and Without Price Information



Notes: Solid line shows the estimated distribution of demand elasticity across consumers if they all were to use the price transparency website and were informed about prices. Dashed line shows effective demand elasticity when all consumers do not use the price transparency website and have uncertainty about prices.

more inelastic. The average implied price elasticity of demand when consumers are uniformed is only -0.07. However, average elasticity is -0.64 when consumers are fully informed. This can be seen in Figure 6.

Turning to the website choice parameters, the coefficient on the monetary benefit of using the website, θ, is positive. Furthermore, zero lies outside the 95 percent credible interval. This implies that consumers are more likely to use the website if the potential benefit is large, either because of the potential savings or individual-specific price sensitive. The coefficients on explanatory variables that make up the observable part of the cost of using the website tend to be imprecisely estimated. There is suggestive evidence that higher income consumers have a larger cost, reflecting the higher opportunity cost of time. At the same time, more educated individuals have a lower cost of using the website, perhaps because they are more likely to be proficient internet users. Patients receiving a procedure after an emergency episode have a higher cost. Furthermore, there is lower estimated cost of using the website in 2010, the omitted year. This may reflect the fact that the website became better known over time. Overall, the cost of using the website is estimated to be \$63 on average (see Figure A-3). Note that the magnitude of θ is relatively small, indicating that that ν_{ikmt} is important for understanding website usage. In other words, there are unobserved factors, such as word-of-mouth or internet proficiency, that

determine website usage.

Estimates under Alternative Demand Assumptions

This section considers an alternative assumption regarding consumer beliefs in the absence of price information. Rather than assume individuals know the mean and variance of the distribution from which prices are drawn, I assume that individuals have an uninformative prior about prices. In this case, the distribution of individuals' posterior beliefs is simply:

$$p^{OOP}_{ijkmt} \overset{iid}{\sim} N\left(p^{OOP}_{ijkmt} + e_{ijkmt}\,,\, \sigma^2_h\right) \tag{35}$$

In other words, individuals take the price signal as given and maximize expected utility given by:

$$\mathbb{E}\left[\widetilde{u_{ijkmt}}\right] = -\gamma_i(p^{OOP}_{ijkmt} + e_{ijkmt}) + \delta_{ijkmt} + \varepsilon_{ijkmt} \tag{36}$$

This is equivalent to setting $w_{ikmt} = 1$ in the baseline specification.

Specification 2 in Table 4 presents the results from this alternative model. The estimates that characterize provider choices are broadly consistent with estimates from specification 1. The estimated standard deviation of the price signal is smaller for both hospital and non-hospital providers. This is due to the fact that noisy price signals generate more extreme beliefs than in the baseline model since beliefs are not disciplined by a prior.

The estimates that characterize the cost of using the website are also broadly consistent with the previous estimates. However, the coefficient on the benefit of the website, θ, is very small compared to specification 1. The alternative assumption implies that, in the absence of information, individuals ignore non-price characteristics and choose providers they believe to be inexpensive. Therefore, the expected benefit of using the website is larger than in the baseline model.

I argue that this alternative model is less realistic. Consumers who lack information about prices are likely to ignore prices rather than choose a provider solely because they guess that it is inexpensive. Formally, consumers likely have a prior that disciplines beliefs.

5.2 Estimates from Supply Model

Table 5 provides results from the bilateral bargaining model. The estimated bargaining weight is 0.37, implying that insurer incentives are important for equilibrium prices. This estimate of the bargaining weights is lower than some other estimates in the literature.[77]

The estimated weight on consumer surplus in the insurer's surplus function is 1.76, implying that insurers put more weight on consumer surplus than on cost. This finding is consistent

[77] For comparison, Ho and Lee (2016) estimate provider bargaining weights between 0.50 and 0.88. However, Gowrisankaran et al. (2015) estimate provider bargaining weights that average 0.24.

Table 5
Bargaining Model Estimates

	Estimate	SE
Bargaining Weight (τ)	0.368	(0.024)
Insurer CS Weight (ζ)	1.759	(0.327)
Procedure FE	Yes	
Provider FE	Yes	
Year FE	Yes	
Observations	4,841	

Notes: GMM estimates using results from the baseline demand model. Standard errors in parentheses.

with Gowrisankaran et al. (2015), who also estimate a weight greater than 1. One potential explanation for this result is that enrollees dislike when their usual provider is dropped from the insurer's network. These switching costs may lead to larger estimates of the weight on consumer surplus.

The estimates from the bargaining model can be used to construct the marginal cost of each procedure at each provider in each year. The marginal cost estimates are summarized by procedure group in Table 6. MRI scans have the highest marginal cost, followed by CT scans and X-rays. The marginal cost estimates are consistent with the fact that MRI machines are the most expensive, while CT and X-ray machines are less expensive and require less staff.

In order to examine whether marginal cost estimates are plausible, I compare the estimates from the model to Medicare reimbursement rates. I find that the marginal cost estimates are generally comparable to the Medicare fee schedule. While I estimate the average marginal cost for medical imaging procedures is \$465, the average Medicare non-facility reimbursement for the same procedures is \$526.[78]

I recompute baseline prices using the estimates of marginal cost as well as the other estimates from the demand and supply model. Given the estimates and an initial price schedule, I calculate the solution to each bargaining problem conditional on other prices. I iterate until all solutions of the bargaining problem reflect the best response to all other prices. This procedure is outlined in more detail in Appendix D. The mean and standard deviation of recomputed baseline prices by procedure group are shown in Table 6. Markups are largest for x-ray procedures. The Lerner index, $(p - mc)/p$, is 0.65 for X-rays, but only 0.19 and 0.10 for CT scans and MRI scans respectively. Since the majority of procedures are X-rays, the average Lerner index across all

[78]Note that Medicare reimbursement rates tend to be higher than marginal cost estimates for X-rays but lower for CT and MRI scans. In order to calculate comparable Medicare reimbursement rates, I use the non-facility fee for each procedure in the visit and then average over the population to find the average visit reimbursement rate for each procedure, insurer, year. The Medicare non-facility fee schedule is used as it applies to non-hospital providers (the facility fee schedule is lower since separate hospital fees are not included). Medicare reimbursement rates are also inflation adjusted to 2010 dollars.

procedures is 0.47.

Table 6
Marginal Cost Estimates

	Baseline Price		Estimated Marginal Cost	
	Mean	SD	Mean	SD
CT Scans	1,869.4	1,006.2	1,512.6	1,458.9
MRI Scans	2,123.5	956.2	1,902.3	1,340.4
X-Rays	674.9	548.4	234.2	912.0

Notes: Prices are simulated using the algorithm described in Appendix D. The unit of observation is a provider, procedure, insurer, year. All prices in 2010 dollars.

6 Effect of the Price Transparency Website

In this section I use the estimates from the previous section to perform a policy evaluation and examine the effect of New Hampshire's price transparency website. First, I calculate the overall equilibrium effect of the website for all individuals who potentially used the price transparency website. This estimate can be compared to the intent-to-treat effect estimated using the difference-in-differences approach in my previous paper (see Brown 2016). I show that results from the two methods are qualitatively consistent, providing support for the main counterfactuals presented in Section 7. I also discuss the conceptual differences between the structural approach in this paper and the difference-in-differences approach.

Next, I use the model to examine the effect of the price transparency website in more detail than is possible with the reduced-form approach alone. In particular, I find the effect of the website conditional on using the price transparency website, often referred to as the treatment-on-the-treated effect. Estimating the treatment-on-the-treated using reduced-form methods is hampered by the fact that there are spillovers that effect all consumers.[79] The estimates from the empirical model imply that consumers subject to a deductible benefit the most from the price transparency website. In addition, the empirical model allows me to recover the change in consumer surplus, which I find is less than consumer cost savings.

6.1 Comparison with Difference-in-Differences Estimates

In previous work, I analyzed the effect of New Hampshire's price transparency website using quasi-experimental reduced-form methods.[80] By comparing the price of medical imaging pro-

[79] Without spillover effects, it is possible to use the "Bloom result" to obtain the treatment-on-the-treated effect via an instrumental variable approach (Bloom 1984). However, it is important to account for the fact that consumers who do not use the website still benefit from lower prices.

[80] See Brown (2016).

cedures on the website before and after the website started with the change in price for similar medical imaging procedures that were never on the website, the approach recovers the average effect of access to the website.[81]

For comparison purposes, I construct the average treatment effect using the estimates from the structural model. This is done by using both the demand and supply estimates to simulate prices with and without the price transparency website. The iterative algorithm used to simulate prices is described in Appendix D. These prices are then used to calculate the average percent change in spending due to the website.

Table 7 presents the estimated effect of access to the price transparency website along with the difference-in-differences estimates from Brown (2016). The estimates from the empirical model are similar to the estimates obtained from the difference-in-differences specification. While the estimates from the former imply that the website reduced overall spending by 2.9 percent, the latter implies that there was a 3.8 percent reduction in spending. Focusing on the reduction in cost for consumers and insurer, I also find similar estimates from the two models. In both cases, the estimates from the empirical model are smaller than the difference-in-differences estimates.

While these estimates are qualitatively consistent, there are a number of potential reasons why the estimates from the two approaches are not identical. First, this paper uses a slightly smaller sample of medical imagining procedures than my previous paper (see discussion in Section 2). In addition, 2011 is excluded when estimating the empirical model since website traffic data is not available. Second, the two approaches require different assumptions to identify the causal effect of the website. The main assumption of the difference-in-differences approach is that the price of procedures on the website would follow a common trend relative to procedures that were not on the website in the absence of treatment. In my previous work, I examine trends in the period prior to the introduction of the website to provide evidence consistent with this assumption. In contrast, the empirical model requires functional form assumptions and distributional assumptions, as well as an assumption about the form of imperfect competition. While the reduced form approach assumes that spending is not affected by unobserved covariates that are correlated with the availability of the website, the structural approach implicitly assumes that primitives, such as the distribution of demand parameters, are orthogonal to the availability of the website.

There is a small literature that directly compares quasi-experimental reduced-form estimates and estimates from a structural model.[82] While the reduced-form analysis is useful for policy evaluation and relies on relatively clear assumptions, the remainder of this paper focuses on

[81] In the policy evaluation literature, this is often referred to as the intent-to-treat effect. The baseline difference-in-differences specification is $LogPrice_{imjkt} = \beta(\text{OnWeb}_m \times \text{Post}_t) + \alpha X_{it} + \lambda_m + \lambda_k + \lambda_t + \varepsilon_{imjkt}$ where OnWeb_m is an indicator for whether procedure m is on the website and Post_t is an indicator for whether the website is available at time t.

[82] One example is Hausman and Leonard (2002), who compare the effect of a product introduction estimated with a difference-in-differences approach with estimates from an empirical model based on theory. Einav et al. (2013) also uses difference-in-differences estimates in conjunction with an empirical model of selection on moral hazard. Also see Peters (2006), Fu and Gregory (2016), and Salz and Vespa (2016).

Table 7
Intent-to-Treat Effect of Price Transparency Website
Comparison with Reduced Form Results

	%Δ Patient Out-of-Pocket Cost	%Δ Insurer Cost	%Δ Total Cost
Difference-in-Difference Estimates	−6.9	−3.8	−3.8
Empirical Model Estimates	−5.5	−2.6	−2.9

Notes: Figures show the percent change in transaction prices for all individuals who could have used the website. Difference-in-differences estimates, which are from Table 5 and second row of Table 9 in Brown (2016), are converted from log-points to percent change. The total effect is not necessarily strictly between the out-of-pocket cost effect and insurer cost effect due to heterogenous quantile treatment effects and heterogenous cost sharing.

results that can only be obtained from an empirical model based on economic theory. In particular, the empirical model allows for an analysis of wider-scale adoption of health care price transparency.[83] I argue that the difference-in-differences estimates help validate the policy-invariant parameter estimates from the empirical model, lending credence to the counterfactual simulations in the remainder of the paper.

6.2 Effect Conditional on using the Price Transparency Website

The first panel of Table 8 summarizes the effect of using the website holding prices fixed.[84] First, I examine individuals that are not subject to a deductible, either because they have surpassed their annual deductible amount or because they have a plan that does not have a deductible.[85] These individuals pay a relatively small portion of the total negotiated prices (cost sharing is about 6 percent on average), therefore the savings from using the website are only $15 per visit on average. Although consumers only take into account the out-of-pocket price, there is correlation between the provider out-of-pocket price and the insurer price. Therefore, insurers also benefit from the increased price shopping (insurers save $12 on average).

Consumers subject to a deductible benefit most from the price transparency website. Individuals who used the website and have a deductible saved $127 per visit, a savings of 26 percent compare to prices they would have paid in the absence of the website. Given that these individuals paid the full negotiated price, there are no insurer savings.

I compute the change in consumer surplus for individuals who use the website using Equation 4 and Equation 13. The gain is smaller than the cost savings—$90 for individuals subject to a deductible. This is due to the fact that, without price information, individuals place less

[83] Although there has been significant policy interest in health care price transparency, it is only recently that price transparency tools have become available to consumers. While these tools are available from private companies, insurers, as well as state governments, these efforts generally affect only a small fraction of consumers.

[84] Demand-side results hold prices fixed using simulated prices for the case in which the website does not exist.

[85] Individuals with a deductible that does not apply to a given medical imaging procedures are also included in this group.

Table 8
Effect for Individuals Predicted to have Used the Transparency Website

	Patient				Insurer		
	OOP Price wo/ Website	OOP Price w/ Website	Δ Price	Δ CS	Insurer Price wo/ Website	Insurer Price w/ Website	Δ Price
	(a) Demand-Side Effects Only						
Over/No Deductible	68.1	53.1	−15.0	11.1	829.8	818.0	−11.8
Under Deductible	486.9	360.2	−126.7	89.5			
	(b) Demand- and Supply-Side Effects						
Over/No Deductible	68.1	51.9	−16.1	12.0	829.8	814.6	−15.2
Under Deductible	486.9	309.2	−177.7	131.8			

Notes: Weighted prices calculated using estimated probability of using website. Demand-side effects are calculated holding observed prices fixed. When analyzing supply-side effects, prices are recomputed for the baseline case (with the website) and the counterfactual scenario in which the website did not exist. All prices in 2010 dollars.

weight on price (since w_{ikmt} is low) and choose providers based on non-price characteristics, such as distance and perceived quality, that are known. With price information, individuals tend to choose less expensive providers, however these providers tend to have worse non-price attributes. Although individuals with a deductible save $127 when they have price information, the providers they choose are $37 worse on non-price characteristics.

In the second panel of Table 8, I account for the fact that the website changed negotiated prices in addition to consumer choices. Rather than hold prices fixed, I re-simulate prices for the case in which some individuals used the price transparency website.

Accounting for the equilibrium effects, the savings from the website were slightly larger. Consumers without a deductible saved $16 while individuals subject to a deductible saved $178. In contrast, the insurers saved only $15 per visit. Overall, the supply-side effects are modest, consistent with my previous reduced-form results. This is due to the fact that a relatively small fraction of consumers use the price transparency website. The next section examines counterfactuals in which a larger fraction of consumers were informed about prices.

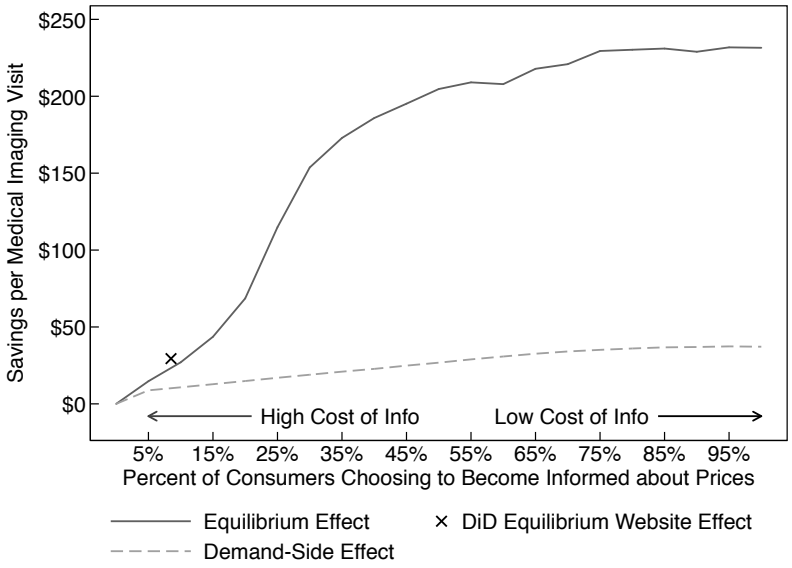
7 Out-of-Sample Counterfactuals

I now use the estimates from the demand and supply model to examine counterfactual policy simulations and explore the broader implications of price uncertainty.

First, I simulate the impact when more consumers are informed about prices holding insurance plan attributes fixed. The results imply a significant supply-side effect when a large fraction of consumers are informed about prices. If all consumers were informed about prices, equilibrium prices would be 20 percent lower. I also find that price transparency leads to less price dispersion.

Next, I examine the effect of a price transparency website that is combined with high cost

Figure 7
Effect of Price Transparency on Medical Imaging Spending
By Fraction on Individuals with Price Information



Notes: Demand-side effect holds prices fixed at distribution simulated with no price transparency. Equilibrium effect re-simulates equilibrium prices for each level of price transparency. Difference-in-differences estimates from Brown (2016). All figures in 2010 dollars.

sharing plans, defined as a 50 percent coinsurance rate. Since out-of-pocket prices are higher on average, individuals have more incentive to use the website and find low cost providers. Although high cost sharing increases patient out-of-pocket spending, overall health care spending declines more than with the price transparency website alone.

7.1 Effect of Increased Price Transparency

Effect on Overall Savings

I begin by examining the effect of reducing the cost of acquiring price information. Only about 8 percent of consumers used the price transparency tool when it was available, implying that there is a large cost of using the website. Much of this cost is likely non-pecuniary, i.e. individuals may not have even known that the website existed. Interventions that reduce this implicit cost, such as advertising the website or even subsidizing usage, would increase the fraction of informed consumers.[86]

Figure 7 shows the effect of reducing the implicit cost of becoming informed about prices, increasing the fraction of consumers with price information. In a similar way as in Section 6,

[86]Other policies that would further increase price information include requiring posted prices or guaranteed price quotes.

I begin by simulating prices for the case in which no individuals have price information, then simulating demand for various cases holding the distribution of prices fixed. This demand-side effect is given by the dashed line in Figure 7. As more individuals choose to use the price transparency effort, average savings increases. However there are decreasing returns due to the fact that website usage is endogenous—the benefit for the marginal consumers is smaller when the cost of using the website is low. There are savings of less than $50 per visit if all consumers are informed.

As more consumers become informed about prices, the demand curve facing providers effectively becomes more elastic. The change in demand affects equilibrium prices, as determined by the bargaining first order condition, potentially generating a positive externality for consumers even if they do not use the price transparency website. I examine the equilibrium effect by simulating prices at each point in Figure 7. I then use these prices to compute consumers choices and overall spending.

The equilibrium effect of increased price transparency is shown by the solid line in Figure 7. As more individuals are informed, the amount saved per visit is highly non-linear. Initially, the supply-side effects are modest—when only a few consumers are informed about prices, equilibrium prices remain relatively constant. When a larger fraction of consumers are informed, there is a meaningful effect on overall demand, changing the incentives of both providers and insurers when negotiating prices. This leads to lower equilibrium prices, generating large savings. As price-cost margins decline, insurers become limited in their ability to negotiate ever lower prices. Once about half of consumers are informed about prices, the supply-side effects become less relevant.

Effect on Consumers, Insurers, and Providers & Welfare Analysis

I examine increased price transparency in greater detail by examining the effect for consumers, insurers, and providers. I focus on two counterfactuals: the equilibrium effect of the price transparency website if it was available for all procedures in all years and the equilibrium effect of full price transparency.

Panel (a) of Table 9 shows the impact on negotiated provider prices relative to no price transparency. When no individuals have price information, the average price of the medical imaging procedures in $1,014. If a price transparency website is available for all medical imagining procedures in all years, the average price declines 2.5 percent to $988.[87] Finally, I examine the counterfactual scenario in which all individuals are fully informed about prices. This would be the case if, for instance, primary care providers were required to provide a price schedule after recommending a medical imaging test. In this case, prices would be 18.9 percent lower than the baseline case. The distribution of prices under various counterfactual scenarios is shown in

[87] This is broadly consistent with difference-in-differences estimates that isolate the supply side. However, note that since this counterfactual examines the effect of the website for all procedures in all years, it is not directly comparable to the estimates from the difference-in-differences model or the results in Section 6.

Table 9
Counterfactual Negotiated Provider Prices

	Mean Price	%Δ Price	Mean Price Dispersion	%Δ Price Dispersion
	(a) Holding Cost Sharing Fixed			
No Transparency (base)	1,014		735	
Transparency Website	988	−2.5%	669	−9.0%
Full Transparency	822	−18.9%	621	−15.5%
	(b) With High Cost Sharing			
No Transparency	908	−10.4%	548	−25.5%
Transparency Website	856	−15.6%	562	−23.5%
Full Transparency	829	−18.2%	619	−15.7%

Notes: Chart shows unweighted prices across all providers/procedures. For the baseline case, prices are computed assuming all individuals have uncertainty about prices. For the price transparency website case, I analyze the case in which the website is available for all procedures in all years. Website usage probabilities are recomputed and then prices are simulated. Full price transparency refers to the case in which all individuals know prices. Price dispersion refers to the interquartile range of prices. High cost sharing refers to 50% coinsurance rate. All prices in 2010 dollars.

Figure A-4.

In the second and third column of Table 9, I examine the effect on price dispersion, as measured by the interquartile range of prices. An increase in price information reduces the degree of price dispersion. Although the mechanism is different, these results are broadly consistent with the literature stressing that price dispersion can result from search frictions.[88]

Panel (a) of Table 10 presents the overall effect on spending taking into account both supply and demand-side effects. If the website were available for all procedures in all years, consumers would save $6 and insurers would save $21 on average, generating $2.1 million in total savings per year on X-ray, CT scans, and MRI scans in New Hampshire. Full price transparency leads to $16.3 million in savings, of which 21 percent accrues to consumers.

The effect for providers is shown in Panel (a) of Table A-4. The savings that accrue to individuals and insurers are, in large part, a result of smaller markups for the provider. However, the change in provider markups is smaller than the savings for consumers and insurers. This is due to the fact that individuals with price information switch to providers that have lower estimated marginal cost, e.g. from hospitals to medical imaging centers.[89] The overall welfare impact for consumers, insurers, and providers is shown in Panel (a) of Table 11. Full price transparency would result in an overall welfare gain of $2.5 million annually.

[88]See, for instance, Stigler (1961), Salop and Stiglitz (1977), and Burdett and Judd (1983).

[89]It is also important to note that, on average, providers still have positive markups even with full price transparency. The fact that there are positive margins helps mitigate concerns about exit from the market.

Table 10
Counterfactual Cost, Welfare, and Expenditure

	Patient					Insurer			Total
	Per Visit			Annual		Per Visit		Annual	Annual
	OOP Cost	Δ OOP Cost	Δ CS	Δ Spending (millions)	Δ CS (millions)	Insurer Cost	Δ Insurer Cost	Δ Spending (millions)	Δ Spending (millions)
	(a) Holding Cost Sharing Fixed								
No Transparency (base)	95.2					871.7			
Website	89.2	−6.1	5.3	−0.5	0.4	850.8	−20.9	−1.6	−2.1
Full Transparency	51.4	−43.9	38.1	−3.4	3.0	705.4	−166.3	−12.9	−16.3
	(b) With High Cost Sharing								
No Transparency	418.5	323.3	−325.9	25.1	−25.3	418.5	−453.1	−35.2	−10.1
Website	396.7	301.4	−310.8	23.4	−24.2	396.7	−475.0	−36.9	−13.5
Full Transparency	249.7	154.4	−208.1	12.0	−16.2	249.7	−622.0	−48.4	−36.3

Notes: Transaction prices are calculated using recomputed prices for each counterfactual. Counterfactual with price transparency website assumes website is available for all imaging procedures in all years. High cost sharing refers to 50% cost sharing. All figures in 2010 dollars.

7.2 Effect of Price Transparency Combined with High Cost Sharing

Health insurance plans with high cost sharing, such as high-deductible plans, potentially give consumers more "skin in the game", increasing the incentive to make cost-effective decisions.[90] Partially for this reason, policies such as tax-advantaged Health Savings Accounts have encouraged high cost sharing plans. However, if consumers cannot observe prices, high cost sharing alone may not lead consumers to switch to less expensive options. For instance, Brot-Goldberg et al. (2015) do not find evidence that high deductible plans increase price shopping.[91] Instead, deductibles may simply lead to larger bill shock. In this section, I examine the effect of combining high cost sharing with the New Hampshire price transparency website. I find that the interaction of these two policies generates larger savings than the price transparency alone.

In Table 9 panel (b), I consider that case in which individuals all have plans with a 50 percent coinsurance rate. In this case, individuals have more incentive to use the website since the potential savings are larger. Simulations imply that website usage increases 38 percent under the high cost sharing scenario. This puts additional downward pressure on prices, resulting in mean prices that are 16 percent lower than with the price transparency website alone. Under high cost sharing, the market is relatively efficient when individuals have access to the price transparency website. Comparing this with full price transparency combined with high cost sharing, the additional supply-side effects are small. In fact, the prices estimates with full price transparency and high cost sharing are slightly higher than those estimates with full transparency alone. The effect on prices is due in part to the fact that insurers have less incentive to negotiate lower prices if they incur a smaller portion of the negotiated price.

[90] See, for instance, discussion in Gowrisankaran et al. (2015).

[91] Although the sample of consumers examined by Brot-Goldberg et al. (2015) had access to a price transparency tool, they note that only a small fraction of consumers knew about it.

Table 11
Counterfactual Net Welfare Impact for Consumers, Providers, and Insurers

	Per Visit Δ Welfare	Annual Δ Welfare (millions)
	(a) Holding Cost Sharing Fixed	
No Transparency (baseline)		
Website	1.9	0.1
Full Transparency	31.7	2.5
	(b) With High Cost Sharing	
No Transparency	32.4	2.5
Website	53.9	4.2
Full Transparency	231.8	18.0

Notes: Figures refer to overall welfare effects for consumers, providers, and insurers due to the fact that consumers switch to providers with lower marginal cost. High cost sharing refers to 50% cost sharing. All figures in 2010 dollars.

High cost sharing leads to a large reduction in price dispersion. However, under high cost sharing, price dispersion actually grows as more individuals are informed about prices. When negotiating prices, the incentives of the insurer are multifaceted, resulting in a complex effect on the distribution of prices (see Figure A-4). As discussed above, high cost sharing means that insurers have less incentive to negotiate low prices in some cases, especially if they know that consumers can shop around. Therefore, it is not always the case that price dispersion declines with increased price transparency.

Table 10 reports the impact on transaction prices, welfare, and total annual spending taking into account both the demand- and supply-side effects. When the website is combined with high cost sharing, the annual savings total $13.5 million, over 6 times larger than with the website alone. Full price transparency results in $36.3 million in savings. Due to the high cost sharing, these savings accrue to the insurer, whereas the consumers have higher out-of-pocket cost. However, it is possible that insurer savings are passed on to consumers in the form of lower insurance premiums.[92]

Panel (b) of Table 11 shows the overall welfare impact of combining price transparency with high cost sharing. Although much of the reduction in health care spending is due to a transfer from providers to insurers, there are still significant net welfare gains. When full price transparency is combined with high cost sharing, many consumers switch to lower marginal cost providers, resulting in a welfare gain of $18.0 million annually.

[92] This would require an explicit model of insurer competition. See Ho and Lee (2016).

8 Conclusion

In this paper, I study the effects of price transparency in the market for medical imaging procedures. I contribute to the literature by developing an empirical model of competition in the market for medical procedures that separately accounts for consumer preferences and consumer uncertainty about prices. To estimate the model, I use detailed administrative medical claims data and exploit a plausibly exogenous policy change—the introduction of a price transparency website— that allowed about 8 percent of consumers to become informed about prices. The degree of uncertainty about prices is identified by comparing observed consumer choices with and without the website. I then use these estimates to examine the broader implications of price transparency on prices, spending, and welfare.

The results imply that price transparency, if used widely, leads to lower provider markups, generating large savings for consumers and insurers. In particular, counterfactual simulations imply that there are considerable spending reductions when roughly half of consumers are informed about prices. The savings are due in large part to the fact that demand effectively becomes more elastic when a large fraction of consumers are informed, allowing insurers to negotiate lower prices with providers. Given that price transparency leads to a positive externality for uninformed consumers by reducing prices, the website is underused. Overall, I find that spending on medical imaging procedures would decline by $16.3 million per year if all consumers were informed about prices.

I also find that price transparency tools would be more widely used if consumers were more exposed to the full price of medical services. In addition, consumers would have more incentive to choose the lowest cost provider once they had price information. This in turn would allow insurers to negotiate prices that would be almost as low as the full information case without high cost sharing. Overall, price transparency tools combined with high cost sharing, defined as 50 percent cost sharing, would result in a $13.5 million annual reduction in medical imaging spending in the state.

Although the empirical analysis in this paper focuses on X-rays, CT scans, and MRI scans, the lack of price transparency likely has implications for other medical procedures that are "shoppable". This includes other outpatient procedures and services such as primary care office visits, simple elective surgeries, and diagnostic testing procedures, as well as some inpatient procedures such as newborn delivery. Previous analysis of health care expenditures has estimated that over 40 percent of health care spending is on services that are potentially shoppable.[93]

Moreover, these results have implications for other markets in which it is not possible to search for prices, such as markets in which consumers receive potentially noisy price quotes (e.g. automotive repair, building contractors, and law services). Even if these quotes are correct on average, the lack of price transparency makes demand less elastic, potentially leading to larger

[93]See "Spending on Shoppable Services in Health Care," Health Care Cost Institute, Issue Brief #11, March, 2016.

markups. A similar mechanism may be present in other markets with shrouded add-on prices, surcharges, or delivery fees.

Firms sometimes justify opaque pricing by noting that there is uncertainty about cost that is only resolved once the service is provided. Although I argue this is not the case for standardized medical imaging procedures, this may be an issue for other medical services. For instance, hospitals may not know how complex a surgery will be until a patient is on the operating room table. While research has examined the incentives to increase search costs (e.g. Ellison and Ellison 2009; Ellison and Wolitzky 2012), there is little evidence on the incentives of firms in this context. One possible direction for future research would be to examine firm incentives for price transparency, especially when marginal cost is individual-specific.

I conclude that information frictions are important for understanding the effect of competition in the market for health care services. Similar information frictions may lead to higher prices in other markets. From a policy perspective, I argue that health care price transparency has substantial promise that has not yet been fully realized.

References

Ackerberg, Daniel A., "Advertising, Learning, and Consumer Choice in Experience Good Markets: An Empirical Examination," *International Economic Review*, 2003, *44* (3), 1007–1040.

Allcott, Hunt, "The welfare effects of misperceived product costs: Data and calibrations from the automobile market," *American Economic Journal: Economic Policy*, 2013, *5* (3), 30–66.

Allen, Jason, Robert Clark, and Jean-François Houde, "Search Frictions and Market Power in Negotiated Price Markets," Technical Report, National Bureau of Economic Research February 2014.

Anderson, Gerard F., Uwe E. Reinhardt, Peter S. Hussey, and Varduhi Petrosyan, "It's the prices, stupid: why the United States is so different from other countries," *Health Affairs*, 2003, *22* (3), 89–105.

Armstrong, Mark, "Search and Ripoff Externalities," *Review of Industrial Organization*, 2015, *47* (3), 273–302.

Aron-Dine, Aviva, Liran Einav, Amy Finkelstein, and Mark Cullen, "Moral Hazard in Health Insurance: Do Dynamic Incentives Matter?," *Review of Economics and Statistics*, 2015, *97* (4), 725–741.

Bloom, Howard S., "Accounting for no-shows in experimental evaluation designs," *Evaluation review*, 1984, *8* (2), 225–246.

Brot-Goldberg, Zarek C., Amitabh Chandra, Benjamin R. Handel, and Jonathan T. Kolstad, "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 2015.

Brown, Jeffrey R. and Austan Goolsbee, "Does the Internet Make Markets More Competitive? Evidence from the Life Insurance Industry," *Journal of Political Economy*, June 2002, *110* (3), 481–507.

Brown, Jennifer, Tanjim Hossain, and John Morgan, "Shrouded Attributes and Information Suppression: Evidence from the Field," *The Quarterly Journal of Economics*, 2010, *125* (2), 859–876.

Brown, Zach Y., "Equilibrium Effects of Health Care Price Information," July 2016. Working Paper.

Burdett, Kenneth and Kenneth L. Judd, "Equilibrium Price Dispersion," *Econometrica*, 1983, *51* (4), 955–969.

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

Carpenter, Bob, Andrew Gelman, Matt Hoffman, Daniel Lee, Ben Goodrich, Michael Betancourt, Michael A. Brubaker, Jiqiang Guo, Peter Li, and Allen Riddell, "Stan: A probabilistic programming language," *Journal of Statistical Software*, 2016, *(forthcoming)*.

Charlson, Mary E., Peter Pompei, Kathy L. Ales, and C. Ronald MacKenzie, "A new method of classifying prognostic comorbidity in longitudinal studies: development and validation," *Journal of chronic diseases*, 1987, *40* (5), 373–383.

Chen, Xiaohong, Han Hong, and Denis Nekipelov, "Nonlinear Models of Measurement Errors," *Journal of Economic Literature*, December 2011, *49* (4), 901–37.

Ching, Andrew T., "A dynamic oligopoly structural model for the prescription drug market after patent expiration," *International Economic Review*, 2010, *51* (4), 1175–1207.

Christensen, Hans Bonde, Eric Floyd, and Mark G. Maffett, "The Effects of Price Transparency Regulation on Prices in the Healthcare Industry," 2015.

Collard-Wexler, Allan, Gautam Gowrisankaran, and Robin S. Lee, ""Nash-in-Nash" Bargaining: A Microfoundation for Applied Work," Technical Report, National Bureau of Economic Research 2014.

Cooper, Zack, Stuart V. Craig, Martin Gaynor, and John Van Reenen, "The Price Ain't Right? Hospital Prices and Health Spending on the Privately Insured," Technical Report, National Bureau of Economic Research December 2015.

Crawford, Gregory S. and Ali Yurukoglu, "The Welfare Effects of Bundling in Multichannel Television Markets," *American Economic Review*, April 2012, *102* (2), 643–85.

_ **and Matthew Shum**, "Uncertainty and learning in pharmaceutical demand," *Econometrica*, 2005, *73* (4), 1137–1173.

Cremer, Jacques and Michael H. Riordan, "On governing multilateral transactions with bilateral contracts," *The RAND Journal of Economics*, 1987, pp. 436–451.

Cutler, David M., Robert S. Huckman, and Mary Beth Landrum, "Information, Pricing, and Productivity of Medical Technology-The Role of Information in Medical Markets: An Analysis of Publicly Reported Outcomes in Cardiac Surgery," *American Economic Review*, 2004, *94* (2), 342–346.

Decarolis, Francesco, "Medicare Part D: Are Insurers Gaming the Low Income Subsidy Design?," *American Economic Review*, April 2015, *105* (4), 1547–80.

DellaVigna, Stefano and Ulrike Malmendier, "Contract design and self-control: Theory and evidence," *The Quarterly Journal of Economics*, 2004, pp. 353–402.

Desai, Sunita, Laura Hatfield, Andrew Hicks, Michael Chernew, and Ateev Mehrotra, "The impact of price transparency tool on outpatient spending," *JAMA*, May 2016, *315* (17), 1842–1843.

Diamond, Peter A., "A model of price adjustment," *Journal of Economic Theory*, June 1971, *3* (2), 156–168.

Dickstein, Michael J., "Efficient Provision of Experience Goods: Evidence from Antidepressant Choice," April 2014.

Dranove, David and Ginger Zhe Jin, "Quality Disclosure and Certification: Theory and Practice," *Journal of Economic Literature*, December 2010, *48* (4), 935–63.

Ehm, Werner, "Binomial approximation to the Poisson binomial distribution," *Statistics and Probability Letters*, 1991, *11* (1), 7 – 16.

Einav, Liran, Amy Finkelstein, Stephen P. Ryan, Paul Schrimpf, and Mark R. Cullen, "Selection on moral hazard in health insurance," *American Economic Review*, 2013, *103* (1), 178–219.

Ellison, Glenn, "A Model of Add-On Pricing," *The Quarterly Journal of Economics*, 2005, *120* (2), 585–637.

_ and Alexander Wolitzky, "A search cost model of obfuscation," *The RAND Journal of Economics*, 2012, *43* (3), 417–441.

_ and Sara Fisher Ellison, "Search, Obfuscation, and Price Elasticities on the Internet," *Econometrica*, 03 2009, *77* (2), 427–452.

Erdem, Tülin and Michael P. Keane, "Decision-Making Under Uncertainty: Capturing Dynamic Brand Choice Processes in Turbulent Consumer Goods Markets," *Marketing Science*, 1996, *15* (1), 1–20.

_ , _ , and Baohong Sun, "A Dynamic Model of Brand Choice When Price and Advertising Signal Product Quality," *Marketing Science*, 2008, *27* (6), 1111–1125.

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

Fu, Chao and Jesse Gregory, "Estimation of an Equilibrium Model with Externalities: Combining the Strengths of Structural Models and Quasi-Experiments," March 2016.

Gabaix, Xavier and David Laibson, "Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets," *The Quarterly Journal of Economics*, 2006, *121* (2), 505–540.

Gaynor, Martin and William B. Vogt, "Competition among Hospitals," *The RAND Journal of Economics*, 2003, *34* (4), 764–785.

Giulietti, Monica, Michael Waterson, and Matthijs R. Wildenbeest, "Estimation of Search Frictions in the British Electricity Market," *The Journal of Industrial Economics*, 2014, *62* (4), 555–590.

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

Grennan, Matthew, "Price Discrimination and Bargaining: Empirical Evidence from Medical Devices," *American Economic Review*, February 2013, *103* (1), 145–77.

_ and Ashley Swanson, "Transparency and Negotiated Prices: The Value of Information in Hospital-Supplier Bargaining," Technical Report, National Bureau of Economic Research February 2016.

_ and Robert Town, "Regulating Innovation with Uncertain Quality: Information, Risk, and Access in Medical Devices," Working Paper 20981, National Bureau of Economic Research February 2015.

Grubb, Michael D., "Consumer Inattention and Bill-Shock Regulation," *The Review of Economic Studies*, 2014.

Grubb, Michael D, "Failing to choose the best price: Theory, evidence, and policy," *Review of Industrial Organization*, 2015, *47* (3), 303–340.

Grubb, Michael D. and Matthew Osborne, "Cellular Service Demand: Biased Beliefs, Learning, and Bill Shock," *American Economic Review*, January 2015, *105* (1), 234–71.

Handel, Benjamin R. and Jonathan T. Kolstad, "Health Insurance for "Humans": Information Frictions, Plan Choice, and Consumer Welfare," *American Economic Review*, August 2015, *105* (8), 2449–2500.

Hausman, Jerry A. and Gregory K. Leonard, "The Competitive Effects of a New Product Introduction: A Case Study," *The Journal of Industrial Economics*, 2002, *50* (3), 237–263.

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

_ , Joseph Hogan, and Fiona Scott Morton, "The Impact of Consumer Inattention on Insurer Pricing in the Medicare Part D Program," 2016. Working Paper.

Ho, Katherine, "The welfare effects of restricted hospital choice in the US medical care market," *Journal of Applied Econometrics*, 2006, *21* (7), 1039–1079.

_ and Robin Lee, "Insurer Competition in Health Care Markets," June 2016. Forthcoming, Econometrica.

Hoffman, Matthew D and Andrew Gelman, "The No-U-turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo.," *Journal of Machine Learning Research*, 2014, *15* (1), 1593–1623.

Hong, Han and Matthew Shum, "Using price distributions to estimate search costs," *RAND Journal of Economics*, 06 2006, *37* (2), 257–275.

Honka, Elisabeth, "Quantifying search and switching costs in the US auto insurance industry," *RAND Journal of Economics*, December 2014, *45* (4), 847–884.

Horn, Henrick and Asher Wolinsky, "Bilateral Monopolies and Incentives for Merger," *The RAND Journal of Economics*, 1988, *19* (3), 408–419.

Hortaçsu, Ali and Chad Syverson, "Product Differentiation, Search Costs, and Competition in the Mutual Fund Industry: A Case Study of S&P 500 Index Funds," *The Quarterly Journal of Economics*, 2004, *119* (2), 403–456.

Jang, Youngjun, "The Effects of the Internet and Mobile Search Technologies on Retail Markets: Evidence from the Korean Gasoline Market," 2015. Working Paper.

Jensen, Robert, "The digital provide: Information (technology), market performance, and welfare in the South Indian fisheries sector," *The quarterly journal of economics*, 2007, pp. 879–924.

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

Koechlin, Francette, Luca Lorenzoni, and Paul Schreyer, "Comparing price levels of hospital services across countries," *Organisation for Economic Co-operation and Development Publishing*, 2010.

Kolstad, Jonathan T., "Information and quality when motivation is intrinsic: Evidence from surgeon report cards," *American Economic Review*, 2013, *103* (7), 2875–2910.

Koulayev, Sergei, "Search for differentiated products: identification and estimation," *RAND Journal of Economics*, 09 2014, *45* (3), 553–575.

Lewis, Matthew S. and Kevin E. Pflum, "Diagnosing hospital system bargaining power in managed care networks," *American Economic Journal: Economic Policy*, 2015, *7* (1), 243–274.

Lieber, Ethan M.J., "Does it Pay to Know the Prices in Health Care?," 2015. Working Paper.

Limbrock, Frank, "Pecuniary and non-pecuniary incentives in prescription pharmaceuticals: The case of statins," *The BE Journal of Economic Analysis & Policy*, 2011, *11* (2).

Luco, Fernando, "Who Benefits from Information Disclosure? The Case of Retail Gasoline," 2015. Working Paper.

Martin, Anne B., Micah Hartman, Joseph Benson, Aaron Catlin, and National Health Expenditure Accounts Team, "National health spending in 2014: faster growth driven by coverage expansion and prescription drug spending," *Health Affairs*, 2016, *35* (1), 150–160.

Moraga-González, José-Luis, Zsolt Sándor, and Matthijs Wildenbeest, "Consumer Search and Prices in the Automobile Market," CEPR Discussion Papers 10487, C.E.P.R. Discussion Papers March 2015.

Peters, Craig, "Evaluating the Performance of Merger Simulation: Evidence from the U.S. Airline Industry," *The Journal of Law and Economics*, 2006, *49* (2), 627–649.

Revelt, David and Kenneth Train, "Mixed Logit With Repeated Choices: Households' Choices Of Appliance Efficiency Level," *The Review of Economics and Statistics*, November 1998, *80* (4), 647–657.

Rubinstein, Ariel, "Perfect equilibrium in a bargaining model," *Econometrica*, 1982, pp. 97–109.

Sacks, Naomi C., James F. Burgess, Howard J. Cabral, and Steven D. Pizer, "Myopic and Forward Looking Behavior in Branded Oral Anti-Diabetic Medication Consumption: An Example from Medicare Part D," *Health economics*, January 2016.

Salop, Steven and Joseph Stiglitz, "Bargains and Ripoffs: A Model of Monopolistically Competitive Price Dispersion," *The Review of Economic Studies*, 1977, *44* (3), 493–510.

Salz, Tobias, "Intermediation and Competition in Search Markets: An Empirical Case Study," December 2015.

_ **and Emanuel Vespa**, "Estimating Dynamic Games of Oligopolistic Competition: An Experimental Investigation," August 2016.

Samuelson, William, "Bargaining under asymmetric information," *Econometrica: Journal of the Econometric Society*, 1984, pp. 995–1005.

Santos, Babur De Los, Ali Hortaçsu, and Matthijs R. Wildenbeest, "Testing Models of Consumer Search Using Data on Web Browsing and Purchasing Behavior," *American Economic Review*, October 2012, *102* (6), 2955–80.

Schmeiser, Steven, "Consumer inference and the regulation of consumer information," *International Journal of Industrial Organization*, 2014, *37*, 192 – 200.

Seiler, Stephan, "The impact of search costs on consumer behavior: A dynamic approach," *Quantitative Marketing and Economics*, 2013, *11* (2), 155–203.

Small, Kenneth A. and Harvey S. Rosen, "Applied Welfare Economics with Discrete Choice Models," *Econometrica*, 1981, *49* (1), 105–130.

Soares, Ilton G., "Mergers under Exclusive Dealing: An Empirical Analysis of the Fuel Industry," November 2015.

Spiegler, Ran, "Competition over agents with boundedly rational expectations," *Theoretical Economics*, June 2006, *1* (2), 207–231.

Stagg, Vicki, "Charlson: Stata module to calculate Charlson index of comorbidity," Statistical Software Components, Boston College Department of Economics April 2006.

Stahl, Dale O., "Oligopolistic pricing with sequential consumer search," *American Economic Review*, 1989, pp. 700–712.

Stigler, George J., "The Economics of Information," *Journal of Political Economy*, 1961, *69*, 213.

Tay, Abigail, "Assessing Competition in Hospital Care Markets: The Importance of Accounting for Quality Differentiation," *The RAND Journal of Economics*, 2003, *34* (4), 786–814.

Train, Kenneth, "Welfare calculations in discrete choice models when anticipated and experienced attributes differ: A guide with examples," *Journal of Choice Modelling*, 2015, *16* (C), 15–22.

Whaley, Christopher M., "Searching for Health: The Effects of Online Price Transparency," October 2015a.

_ , "Provider Responses to Online Price Transparency," December 2015b.

Woodford, Michael, "Inflation Stabilization and Welfare," Technical Report, National Bureau of Economic Research January 2001.

APPENDIX

A Derivation of Website Usage Benefit

Here I derive the approximate benefit of using the price transparency website using a second-order multivariate Taylor series around the expectation.[94] The primary reason for using this approximation is that the expression must be computed in every iteration of the estimation procedure. Therefore, simulating the high-dimensional integral is computationally infeasible.

If we wish to approximate the first moment of the function $f(x_1, x_2, \ldots, x_N)$ given mean values $(\mu_1, \mu_2, \ldots, \mu_N)$, the second-order Taylor series is

$$\mathbb{E}[f(x_1, x_2, \ldots, x_N)] \approx$$
$$f(\mu_1, \mu_2, \ldots, \mu_N) + \sum_{n=1}^{N} \frac{\partial f(\mu_1, \mu_2, \ldots, \mu_N)}{\partial x_n} \mathbb{E}(x_n - \mu_n) + \frac{1}{2!} \sum_{n=1}^{N} \sum_{k=1}^{N} \frac{\partial^2 f(\mu_1, \mu_2, \ldots, \mu_N)}{\partial x_n \partial x_k} \mathbb{E}(x_n - \mu_n)(x_k - \mu_k)$$

In this case, I wish to approximate the expected value of consumer surplus if individuals use the price transparency website and know prices

$$\mathbb{E}_{\widetilde{\mathbf{p}_{ikmt}^{OOP}}} \left[\frac{1}{\gamma_i} \log(\sum_{j \in \mathcal{N}_{kmt}} \exp(-\gamma_i \widetilde{p_{ijkmt}^{OOP}} + \delta_{ijkmt})) \right]$$

where $\widetilde{\mathbf{p}_{ikmt}^{OOP}}$ is the vector of beliefs. The individual believes each price to be distributed

$$\widetilde{p_{ijkmt}^{OOP}} \overset{iid}{\sim} N\left(\mathbb{E}\left[\widetilde{p_{ijkmt}^{OOP}} \right], Var\left[\widetilde{p_{ijkmt}^{OOP}} \right] \right)$$

Since price signals are independent

$$\mathbb{E}_{\mathbf{p}_{ikmt}^{OOP}} \left[(\widetilde{p_{ijkmt}^{OOP}} - \mathbb{E}[\widetilde{p_{ijkmt}^{OOP}}])(\widetilde{p_{ij'kmt}^{OOP}} - \mathbb{E}[\widetilde{p_{ij'kmt}^{OOP}}]) \right] = \begin{cases} Var\left[\widetilde{p_{ijkmt}^{OOP}} \right] & \text{if } j = j' \\ 0 & \text{if } j \neq j' \end{cases}$$

Furthermore, note that

$$\frac{\partial^2 \log(\sum_{j' \in \mathcal{N}_{kmt}} \exp(-\gamma_i \widetilde{p_{ij'kmt}^{OOP}} + \delta_{ijkmt}))}{\partial \widetilde{p_{ijkmt}^{OOP}}^2} =$$

[94] In the context of standard errors, a similar approach is often referred to as the delta method.

$$\frac{\gamma_i^2 \exp(-\gamma_i \widetilde{p^{OOP}_{ij'kmt}} + \delta_{ijkmt}) \sum_{j' \in \mathcal{N}_{kmt} \setminus j} \exp(-\gamma_i \widetilde{p^{OOP}_{ij'kmt}} + \delta_{ijkmt})}{\left[\sum_{j' \in \mathcal{N}_{kmt}} \exp(-\gamma_i \widetilde{p^{OOP}_{ij'kmt}} + \delta_{ijkmt})\right]^2}$$

Using this, the second-order taylor series evaluated at the expectation is

$$\frac{1}{\gamma_i} \log\Big(\sum_{j \in \mathcal{N}_{kmt}} \exp(-\gamma_i \mathbb{E}\left[\widetilde{p^{OOP}_{ijkmt}}\right] + \delta_{ijkmt})) + \frac{\gamma_i \sum_{j \in \mathcal{N}_{kmt}} \left[Var\left[\widetilde{p^{OOP}_{ijkmt}}\right] \exp(-\gamma_i \mathbb{E}\left[\widetilde{p^{OOP}_{ijkmt}}\right] + \delta_{ijkmt}) \sum_{j' \in \mathcal{N}_{kmt} \setminus j} \exp(-\gamma_i \mathbb{E}\left[\widetilde{p^{OOP}_{ij'kmt}}\right] + \delta_{ijkmt})\right]}{2\left[\sum_{j' \in \mathcal{N}_{kmt}} \exp(-\gamma_i \mathbb{E}\left[\widetilde{p^{OOP}_{ij'kmt}}\right] + \delta_{ijkmt})\right]^2}$$

Now turn to the consumer surplus without using the website. Note that

$$\frac{1}{\gamma_i} \log\left(\sum_{j \in \mathcal{N}_{kmt}} \exp(-\gamma_i \mathbb{E}\left[\widetilde{p^{OOP}_{ijkmt}}\right] + \delta_{ijkmt})\right) + \mathbb{E}_{\mathbf{p}^{OOP}_{ikmt}}\left[\sum_{j \in \mathcal{N}_{kmt}} \left(\mathbb{E}\left[\widetilde{p^{OOP}_{ijkmt}}\right] - \widetilde{p^{OOP}_{ijkmt}}\right) s_{ijkmt}\right]$$
$$= \frac{1}{\gamma_i} \log\left(\sum_{j \in \mathcal{N}_{kmt}} \exp(-\gamma_i \mathbb{E}\left[\widetilde{p^{OOP}_{ijkmt}}\right] + \delta_{ijkmt})\right)$$

since the expected bill shock from the individual's perspective is zero.

The value of the website is the difference between expected consumer surplus with and without using the website. Therefore, the approximate benefit of the website in dollars is

$$\frac{\gamma_i \sum_{j \in \mathcal{N}_{kmt}} \left[Var\left[\widetilde{p^{OOP}_{ijkmt}}\right] \exp(-\gamma_i \mathbb{E}\left[\widetilde{p^{OOP}_{ijkmt}}\right] + \delta_{ijkmt}) \sum_{j' \in \mathcal{N}_{kmt} \setminus j} \exp(-\gamma_i \mathbb{E}\left[\widetilde{p^{OOP}_{ij'kmt}}\right] + \delta_{ijkmt})\right]}{2\left[\sum_{j' \in \mathcal{N}_{kmt}} \exp(-\gamma_i \mathbb{E}\left[\widetilde{p^{OOP}_{ij'kmt}}\right] + \delta_{ijkmt})\right]^2}$$

I test this approximation by simulating draws from the distribution of beliefs, computing consumers surplus, and then averaging over the draws to compute the expectation. I find that the simulated expectation is within 5 percent of the second-order approximation using reasonable parameter values. For a more detailed discussion of the validity of a Taylor series approximation for calculating welfare in a macroeconomic context, see Woodford (2001).

B Details on Bayesian Interpretation of the Demand Model and Estimation

In this section, I present the model as it is estimated in a Bayesian framework. As I describe in Section 3.4, reformulating the model in this way allows for a computationally feasible estimation strategy that takes advantage of recent advances in Bayesian estimation but does not change the structural interpretation of parameters.

In general, the posterior is defined as

$$P(\Theta|D) \propto \breve{\mathcal{L}}(D|\Theta)P(\Theta)$$

where $\breve{\mathcal{L}}(D|\Theta)$ is the likelihood given data D and $P(\Theta)$ is the distribution of the parameter prior.

Start by defining an individual's choice probabilities conditional on unobservables, γ_i and $\mathbf{e}_{ikmt}$:

$$\begin{aligned} s_{ijkmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt}|\gamma_i, \mathbf{e}_{ikmt}) = \\ \vartheta_{ikmt} \cdot \frac{\exp(-\gamma_i p^{OOP}_{ijkmt} + \delta_{ijkmt})}{\sum_{j' \in \mathcal{N}_{kmt}} \exp(-\gamma_i p^{OOP}_{ij'kmt} + \delta_{ij'kmt})} \\ + (1 - \vartheta_{ikmt}) \cdot \frac{\exp(-\gamma_i w_{ikmt}(p^{OOP}_{ijkmt} + e_{ijkmt}) + \delta_{ijkmt})}{\sum_{j' \in \mathcal{N}_{kmt}} \exp(-\gamma_i w_{ikmt}(p^{OOP}_{ij'kmt} + e_{ij'kmt}) + \delta_{ij'kmt})} \end{aligned}$$

The simplified likelihood, conditional on unobservables, is then

$$\breve{\mathcal{L}}(D|\Theta) = [s_{ijkmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt}|\gamma_i, \mathbf{e}_{ikmt})]^{y_{ijkmt}}$$

Note that, unlike Equation 24, this is a closed-form expression. This likelihood function is than augmented with the following:

$$\mathcal{V}_{mt} \sim Binomial\left(n_{mt}, \frac{\exp(\theta b_{ikmt} - \phi \mathbf{x}_{ikmt})}{1 + \exp(\theta b_{ikmt} - \phi \mathbf{x}_{ikmt})}\right) \quad \text{(Website Usage)}$$

$$\gamma_i \sim N(\bar{\gamma} + \rho c_{ik}, (\sigma^\gamma)^2) \quad \text{(Price Sensitivity)}$$

$$e_{ijkmt} \sim N(0, \sigma_h^2) \quad \text{(Signal Noise)}$$

The remaining parameters are given uninformative priors.

C Details on Bargaining First Order Condition

In this section, I describe the first order condition of the bargaining problem:

$$\frac{\partial}{\partial p_{jkmt}} \left(\sum_{i \in I_{kmt}} [s_{ijkmt}[p_{km} - mc_{mkj}]] \right)^{\tau} \left(\sum_{i \in I_{kmt}} [\Delta_j \Pi^K_{ijkmt}] \right)^{1-\tau} = 0$$

For brevity, I have simplified the notation and omitted the expectation over $\mathbf{e}_{ikmt}$.

Using the fact that $\frac{\partial \Delta_j \Pi^K_{ijkmt}}{\partial p_{jkmt}} = \frac{\partial \Pi^K_{ijkmt}}{\partial p_{jkmt}}$,

$$\sum_{i\in I_{kmt}} \tau\left(s_{ijkmt}+\frac{\partial s_{ijkmt}}{\partial p_{jkmt}}[p_{jkmt}-mc_{jkmt}]\right)\left(s_{ijkmt}[p_{jkmt}-mc_{jkmt}]\right)^{\tau-1}\left(\Delta_j\Pi^K_{ijkmt}\right)^{1-\tau}$$
$$+\sum_{i\in I_{kmt}}(1-\tau)\left(s_{ijkmt}[p_{jkmt}-mc_{jkmt}]\right)^{\tau}\left(\Delta_j\Pi^K_{ijkmt}\right)^{-\lambda}\frac{\partial\Pi^K_{ikmt}}{\partial p_{jkmt}}=0$$

Now solving for the markup:

$$p_{jkmt}-mc_{jkmt}=-\left[\frac{1-\tau}{\tau}\sum_{i\in I_{kmt}}\frac{\partial\Pi^K_{ikmt}}{\partial p_{jkmt}}\left(\Delta_j\Pi^K_{ikmt}\right)^{-1}+\sum_{i\in I_{kmt}}\frac{\partial s_{ijkmt}}{\partial p_{jkmt}}\frac{1}{s_{ijkmt}}\right]^{-1}$$

I now derive $\frac{\partial s_{ijkmt}}{\partial p_{jkmt}}$ and $\frac{\partial\Pi^K_{ikmt}}{\partial p_{jkmt}}$ for the case in which there is price uncertainty ($\vartheta_{kmt}=0$). Given the derivations below, the case in which there is full information ($\vartheta_{kmt}=1$) is easily derived by setting $e_{ijkmt}=0$ and $w_{ikmt}=1$.

First, the own-price and cross-price derivative of the choice probabilities is given by:

$$\frac{\partial s_{ij'kmt}}{\partial p_{jkmt}}=\begin{cases}-\gamma_i w_{ikmt}c_{ikmt}s_{ijkmt}(1-s_{ijkmt}) & \text{if } j'=j\\ \gamma_i w_{ikmt}c_{ikmt}s_{ijkmt}s_{ij'kmt} & \text{if } j'\neq j\end{cases}$$

Note that the above assumes that each individual's prior, which is determined by $\bar{p}^{OOP}_{kmt}$ and $\bar{s}^2_{kmt}$, is a constant that is not affected by a price change (i.e. $\frac{\partial w_{ikmt}}{\partial p_{jkmt}}=1$). This is justified by the fact that each provider likely has a negligible effect on consumer priors, especially when the choice set is large. Therefore, I argue providers are unlikely to internalize this effect. For the same reason, I also use this assumption when deriving the insurer's problem.

I now turn to the partial derivative of the insurer surplus, which has two parts:

$$\frac{\partial\Pi^K_{ikmt}}{\partial p_{jkmt}}=\frac{\partial\zeta CS_{ijkmt}}{\partial p_{jkmt}}-\frac{TC_{ijkmt}}{\partial p_{jkmt}}$$

Where the partial derivative of consumer surplus is:

$$\frac{\partial\zeta CS_{ijkmt}}{\partial p_{jkmt}}=\zeta\frac{\partial}{\partial p_{jkmt}}[\frac{1}{\gamma_i}\log[\sum_{j'\in\mathcal{N}_{kmt}}\exp(-\gamma_i w_{ikmt}(p^{OOP}_{ij'kmt}+e_{ij'kmt})+\delta_{ij'kmt})]$$
$$+\sum_{j'\in\mathcal{N}_{kmt}}[w_{ikmt}e_{ij'kmt}+(1-w_{ikmt})(\bar{p}^{OOP}_{kmt}-p^{OOP}_{ij'kmt})]s_{ij'kmt}$$

$$=-c_{ikmt}s_{ijkmt}-\gamma_i w_{ikmt}c_{ikmt}s_{ijkmt}[w_{ikmt}(1-\sum_{j'\in\mathcal{N}_{kmt}}e_{ij'kmt}s_{ij'kmt})-c_{ikmt}(1-w_{ikmt})(1-\sum_{j'\in\mathcal{N}_{kmt}}p_{j'kmt}s_{ij'kmt})]$$

And the partial derivative of insurer cost is:

$$\frac{\partial TC_{ijkmt}}{\partial p_{jkmt}}=\frac{\partial}{\partial p_{jkmt}}[\sum_{j'\in\mathcal{N}_{kmt}}p_{j'kmt}(1-c_{ikmt})s_{ij'kmt}]$$

$$= (1 - c_{ikmt}) s_{ijkmt} (1 - p_{jkmt} \gamma_i w_{ikmt} c_{ikmt}) + \sum_{j' \in \mathcal{N}_{kmt}} [p_{j'kmt} (1 - c_{ikmt}) \gamma_i w_{ikmt} c_{ikmt} s_{ijkmt} s_{ij'kmt}]$$

D Details on Procedure for Simulating Equilibrium Prices

In order to find counterfactual equilibrium prices, I start by finding choice probabilities as a function of the vector of prices, $s_{ijkmt}(\mathbf{p}_{kmt})$, which depend on whether consumers know prices and the degree of cost sharing. This can then be used to find the consumer surplus, $CS_{ikmt}(\mathbf{p}_{kmt})$, and insurer cost, $TC_{ikmt}(\mathbf{p}_{kmt})$ which are also depend on policy-relevant parameters.

With estimates of the bargaining parameter, $\hat{\tau}$, weight on consumer surplus, $\hat{\zeta}$, and the marginal cost for each option, $\hat{mc}_{jkmt}$, in hand, the negotiated best response, p^*_{jkmt}, given price vector $\mathbf{p}'_{kmt}$ is

$$p^*_{jkmt} = \hat{mc}_{jkmt} + \left(-\frac{1-\hat{\tau}}{\hat{\tau}} \frac{\frac{\partial}{\partial p'_{jkmt}} \sum_{i \in I_{km}} [\hat{\zeta} CS_{ikmt}(\mathbf{p}'_{kmt}) - TC_{ikmt}(\mathbf{p}'_{kmt})]}{\sum_{i \in I_{km}} [\hat{\zeta} \Delta_j CS_{ikmt}(\mathbf{p}'_{kmt}) - \Delta_j TC_{ikmt}(\mathbf{p}'_{kmt})]} - \frac{\frac{\partial}{\partial p'_{jkmt}} [\sum_{i \in I_{km}} s_{ijkmt}(\mathbf{p}'_{kmt})]}{\sum_{i \in I_{km}} s_{ijkmt}(\mathbf{p}'_{kmt})} \right)^{-1}$$

The iterative algorithm proceeds as follows:

1. Find p^*_{jkmt} $\forall j, k, m, t$
2. Update $\mathbf{p}'_{kmt}$[95]
3. Iterate until all prices are optimal given all other prices (i.e. until prices converge within a mean tolerance of \$1).

This equilibrium vector of prices can them be used to simulate spending and welfare given demand.

[95] In rare cases, it is possible for the model to predict negative prices. I bound prices to be positive.

E Appendix Tables

Table A-1
List of Medical Imaging Procedures

CPT Code	Short Description	On Website	CPT Code	Short Description	On Website	CPT Code	Short Description	On Website
70100	X-ray exam of jaw		71100	X-ray exam of ribs		72126	CT neck spine w/dye	
73610	X-ray exam of ankle	×	71022	X-ray chest	×	70491	CT soft tissue neck w/dye	
73615	X-ray of ankle contrast		73130	X-ray exam of hand	×	72191	CT angiograph pelv w/o&w/dye	
70160	X-ray exam of nasal bones		75743	X-ray artery, lungs		74176	CT abd & pelvis w/o contrast	×
70030	X-ray eye for foreign body		75726	X-ray artery, abdomen		74160	CT abdomen w/dye	×
73050	X-ray exam of shoulders	×	73600	X-ray exam of ankle		70470	CT head/brain w/o & w/dye	
75680	X-rays artery, neck		73085	X-ray of elbow contrast		70496	CT angiography, head	
74415	X-ray, urinary tract contrast		75803	X-ray lymph vessel,arms/legs		72132	CT lumbar spine w/dye	
73520	X-ray exam of hips		73060	X-ray exam of humerus		72131	CT lumbar spine w/o dye	
71035	X-ray chest	×	73090	X-ray exam of forearm	×	72158	MRI lumbar spine w/o & w/dye	×
71110	X-ray exam of ribs		71111	X-ray exam of ribs/chest		70544	MRI angiography head w/o dye	
73630	X-ray exam of foot	×	70150	X-ray exam of facial bones		70557	MRI brain w/o dye	
73010	X-ray exam of shoulder blade	×	75685	X-ray artery, spine		75559	MRI cardiac w/stress img	
72110	X-ray exam of lower spine	×	71030	X-ray chest		72157	MRI chest spine w/o & w/dye	×
71021	X-ray chest		73550	X-ray exam of thigh		72148	MRI lumbar spine w/o dye	×
75519	X-ray heart/catheterization		73564	X-ray exam, knee, 4 or more	×	73222	MRI joint upr extrem w/dye	
75731	X-rays artery, adrenal gland		70110	X-ray exam of jaw		70540	MRI orbit/face/neck w/o dye	
70210	X-ray exam of sinuses		75733	X-ray artery, adrenals		72146	MRI chest spine w/o dye	×
75724	X-rays artery, kidneys		70250	X-ray exam of skull		72196	MRI pelvis w/dye	×
72220	X-ray exam of tailbone		72069	X-ray exam of trunk spine		72198	MRI angio pelvis w/o & w/dye	×
77072	X-rays for bone age		72050	X-ray exam of neck spine	×	73225	MRI angio upr extr w/o&w/dye	
73525	X-ray of hip with contrast		75801	X-raylymph vessel, arm/leg		75553	MRI heart for morph w/dye	
70330	X-ray exam of jaw joints		73100	X-ray exam of wrist	×	71550	MRI chest w/o dye	
73080	X-ray exam of elbow		73592	X-ray exam of leg, infant		73721	MRI jnt of lwr extre w/o dye	×
71130	X-ray exam of breastbone		71101	X-ray exam of ribs/chest	×	71552	MRI chest w/o & w/dye	
73562	X-ray exam of knee, 3	×	75710	X-ray artery, arm/leg		70558	MRI brain w/dye	
73660	X-ray exam of toe(s)		72052	X-ray exam of neck spine	×	74183	MRI abdomen w/o & w/dye	×
72010	X-ray exam of spine		72072	X-ray exam of thoracic spine	×	70552	MRI brain w/dye	
74405	X-ray urinary tract contrast		73510	X-ray exam of hip		72149	MRI lumbar spine w/dye	
74320	X-ray of bile ducts contrast		72114	X-ray exam of lower spine	×	70545	MRI angiography head w/dye	
73580	X-ray of knee joint contrast		75527	X-ray heart/catheterization		73718	MRI lower extremity w/o dye	×
72170	X-ray exam of pelvis		70332	X-ray exam of jaw joint		73720	MRI lwr extremity w/o&w/dye	×
75665	X-rays artery, head & neck		73620	X-ray exam of foot	×	71551	MRI chest w/dye	
72090	X-ray exam of trunk spine	×	72200	X-ray exam sacroiliac joints		70551	MRI brain w/o dye	×
75716	X-ray artery, arms/legs		73590	X-ray exam of lower leg		70547	MRI angiography neck w/o dye	
73040	X-ray of shoulder w/contrast		74020	X-ray exam of abdomen		70543	MRI orbt/fac/nck w/o & w/dye	
72100	X-ray exam of lower spine	×	75573	CT heart w/3d image congen		73719	MRI lower extremity w/dye	
75805	X-ray lymph vessel, trunk		75572	CT heart w/3d image		73219	MRI upper extremity w/dye	
74301	X-rays at surgery add-on		72128	CT chest spine w/o dye		73723	MRI joint lwr extr w/o&w/dye	×
74305	X-ray bile ducts/pancreas		75574	CT angio hrt w/3d image		70553	MRI brain w/o & w/dye	×
73000	X-ray exam of collar bone	×	72130	CT chest spine w/o & w/dye		75562	MRI card flow/vel w/dye	
72120	X-ray exam of lower spine	×	72127	CT neck spine w/o & w/dye		72197	MRI pelvis w/o & w/dye	×
70200	X-ray exam of eye sockets		74150	CT abdomen w/o dye	×	75556	MRI cardiac flow mapping	
71120	X-ray exam of breastbone		70498	CT angiography, neck		75558	MRI cardiac flow/velocity	
70130	X-ray exam of mastoids		73702	CT lwr extremity w/o&w/dye		70559	MRI brain w/o & w/dye	
70260	X-ray exam of skull		72129	CT chest spine w/dye		76390	MRI spectroscopy	
75774	X-ray artery, each vessel		72133	CT lumbar spine w/o & w/dye		75561	MRI cardiac for morph w/dye	
71010	X-ray chest	×	73701	CT lower extremity w/dye		75552	MRI heart for morph w/o dye	
72080	X-ray exam of trunk spine		70490	CT soft tissue neck w/o dye		75555	MRI cardiac/limited study	
74000	X-ray exam of abdomen		70480	CT orbit/ear/fossa w/o dye		74182	MRI abdomen w/dye	×
71020	X-ray chest	×	72125	CT neck spine w/o dye		75557	MRI cardiac for morph	
75671	X-ray artery, head & neck		74178	CT abd & pelv 1/¿ regns	×	72147	MRI chest spine w/dye	
73110	X-ray exam of wrist	×	70482	CT orbit/ear/fossa w/o&w/dye		70546	MRI angiograph head w/o&w/dye	
75660	X-ray artery, head & neck		73201	CT upper extremity w/dye		73218	MRI upper extremity w/o dye	×
75722	X-ray artery, kidney		73202	CT uppr extremity w/o&w/dye		72159	MRI angio spine w/o&w/dye	
73030	X-ray exam of shoulder	×	71260	CT thorax w/dye	×	72141	MRI neck spine w/o dye	×
73092	X-ray exam of arm, infant		73206	CT angio upr extrm w/o&w/dye		73220	MRI uppr extremity w/o&w/dye	×
75500	X-ray cinema heart vessels		71270	CT thorax w/o & w/dye	×	70549	MRI angiograph neck w/o&w/dye	
75676	X-ray artery, neck		71275	CT angiography, chest		73725	MRI ang lwr ext w or w/o dye	
72040	X-ray exam of neck spine	×	73700	CT lower extremity w/o dye		73223	MRI joint upr extr w/o&w/dye	
74400	X-ray urinary tract contrast		70486	CT maxillofacial w/o dye		73722	MRI joint of lwr extr w/dye	
75741	X-ray artery, lung		70460	CT head/brain w/dye		74185	MRI angio, abdom w orw/o dye	
75662	X-ray artery, head & neck		70488	CT maxillofacial w/o & w/dye		76400	MRI bone marrow	
73542	X-ray exam, sacroiliac joint		70487	CT maxillofacial w/dye		72195	MRI pelvis w/o dye	×
75705	X-ray artery, spine		75571	CT hrt w/o dye w/ca test		76498	MRI procedure	
73560	X-ray exam of knee, 1 or 2	×	72194	CT pelvis w/o & w/dye	×	70542	MRI orbit/face/neck w/dye	
75807	X-ray lymph vessel, trunk		72193	CT pelvis w/dye	×	72156	MRI neck spine w/o & w/dye	×
73565	X-ray exam of knees	×	73200	CT upper extremity w/o dye		75563	MRI cardiac w/stress img & dye	
70220	X-ray exam of sinuses		72192	CT pelvis w/o dye	×	70336	MRI, jaw joint	
74300	X-ray bile ducts/pancreas		74175	CT angio abdom w/o & w/dye		70541	MRI, head (MRA)	
75650	X-ray artery, head & neck		70450	CT head/brain w/o dye		72142	MRI neck spine w/dye	
75523	X-ray heart /catheterization		76360	CT scan for needle biopsy		75565	MRI cardiac veloc flow mapping	
73540	X-ray exam of pelvis & hips		75635	CT angio abdominal arteries		74181	MRI abdomen w/o dye	×
73650	X-ray exam of heel	×	74177	CT abd & pelv w/contrast	×	70548	MRI angiography neck w/dye	
75756	X-ray artery, chest		74170	CT abdomen w/o & w/dye	×	75560	MRI cardiac flow/vel/stress	
75736	X-ray artery, pelvis		73706	CT angio lwr extr w/o&w/dye		73221	MRI joint upr extrem w/o dye	
75658	X-ray artery, arm		76497	CT procedure		75564	MRI hrt w/flo/vel/strs & dye	
73140	X-ray exam of finger(s)	×	70481	CT orbit/ear/fossa w/dye		71555	MRI angio chest w or w/o dye	
73115	X-ray of wrist contrast	×	70492	CT sft tsue nck w/o & w/dye				
77071	X-ray stress view		71250	CT thorax w/o dye	×			
70360	X-ray exam of neck		74174	CT angio abd&pelv w/o&w/dye				

Notes: Procedures listed on website are more broad than procedures as defined by CPT codes (i.e. a selected website procedure may refer to multiple CPT codes). CPT codes that had modified descriptions over the period are assigned separate identifiers in the analysis.

Table A-2
Potential Savings from Switching Providers

	Patient				Insurer			
	Switch to Lowest Price Provider		Switch to Provider in 1st Quartile		Switch to Lowest Price Provider		Switch to Provider in 1st Quartile	
	Δ Cost	%Δ Cost	Δ Cost	%Δ Cost	Δ Cost	%Δ Cost	Δ Cost	%Δ Cost
CT Scans								
Over/No Deductible	33	66.2%	30	57.8%	1,281	50.7%	1,015	37.0%
Under Deductible	1,114	42.2%	914	33.1%				
MR Scans								
Over/No Deductible	33	68.1%	35	69.8%	927	37.0%	676	25.9%
Under Deductible	796	36.0%	590	26.6%				
X-Rays								
Over/No Deductible	16	46.1%	15	41.3%	375	46.8%	331	36.5%
Under Deductible	457	50.8%	396	38.6%				

Notes: Chart shows potential savings from switching to providers with lower negotiated prices within individual's choice set relative to observed choices. See Section 2.1 for definition of choice set. If provider in first quartile is more expensive then chosen provider, simulation assumes individuals do not switch.

Table A-3
Estimates from Multinomial Logit Demand Model

	Estimate	SE
Patient Out-of-Pocket price		
× Website Available ($-\gamma_1$)	-0.0014^{***}	(0.0004)
× Website Not Available ($-\gamma_2$)	-0.0004	(0.0004)
Distance (α_1)	-0.0287^{***}	(0.0030)
Distance squared (α_2)	0.0002^{***}	(0.0000)
Referral Indicator (α_3)	1.9939^{***}	(0.0313)
Hospital × Age ≤18	0.0769	(0.1667)
Hospital × Age 19-35	-0.0270	(0.1638)
Hospital × Age 36-50	-0.0141	(0.1568)
Hospital × Age 51-64	-0.1080	(0.1618)
Hospital × Male	-0.0863	(0.0671)
Hospital × Income	0.0112^{***}	(0.0022)
Hospital × BA	-0.0313^{***}	(0.0040)
Hospital × Charlson	0.0361	(0.0444)
Hospital × Emergency	0.5407^{***}	(0.0861)
Procedure group FE × provider	Yes	
Log Likelihood	-11,820	
Pseudo-R2	0.224	
Observations	59,240	

Notes: MLE estimates from a 2% sample of visits. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

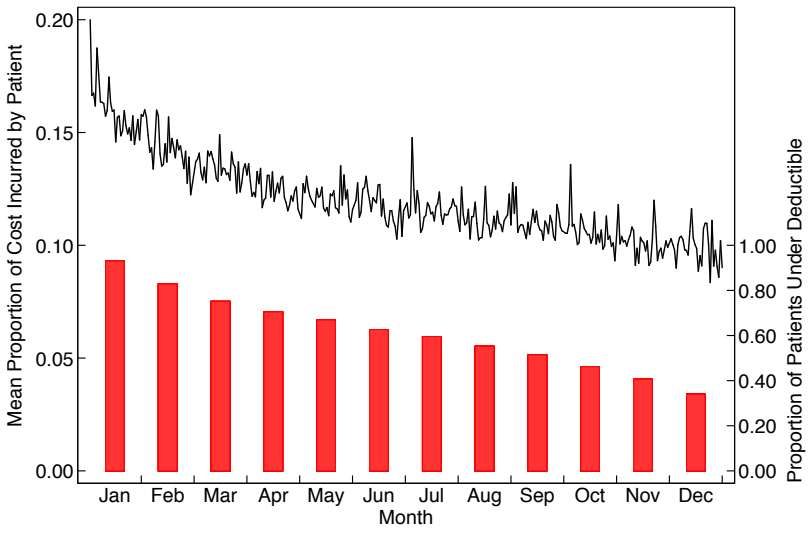
Table A-4
Counterfactual Provider Markups and Surplus

	Provider		
	Markup	Δ Markup	Δ Annual Surplus (millions)
	(a) Holding Cost Sharing Fixed		
No Transparency (base)	309.3		
Website	285.0	−24.3	−1.9
Full Transparency	136.7	−172.6	−13.4
	(b) With High Cost Sharing		
No Transparency	214.5	−94.8	−7.4
Website	199.0	−110.3	−8.6
Full Transparency	127.1	−182.1	−14.2

Notes: Transaction prices are calculated using recomputed prices for each counterfactual. Counterfactual with price transparency website assumes website is available for all imaging procedures in all years. High cost sharing refers to 50% cost sharing.

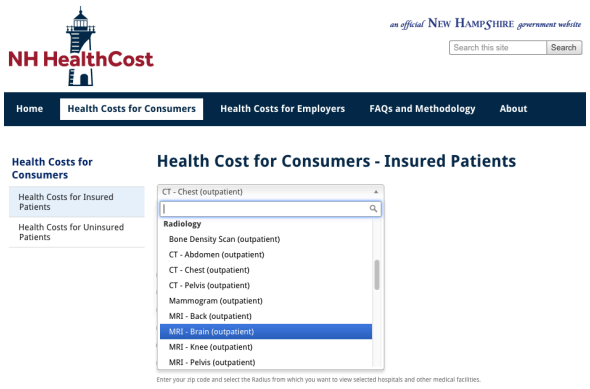
F Appendix Figures

Figure A-1
Cost Sharing and Deductible Status
By Month of Year



Notes: Bars show the percent of medical imaging consumers that have not hit their deductible conditional on having a plan with a deductible. Line shows the unconditional mean percent of total cost paid by the patient for each day of the year.

Figure A-2
Prices Variation within Individuals' Choice Sets



(a) Entering Procedure Information

Lead Provider	Estimate of What you Will Pay	Estimate of What Insurance Will Pay	Estimate of Combined Payments	Precision of the Cost Estimate	Typical Patient Complexity
STILES ROAD IMAGING 603.898.3129	$148	$245	$394	MEDIUM	MEDIUM
AMERICAN MEDICAL IMAGING 603.766.6736	$190	$371	$562	HIGH	MEDIUM
DERRY IMAGING CENTER 603.537.1363	$275	$624	$898	HIGH	MEDIUM
CHESHIRE MEDICAL CENTER 603.354.5400	$282	$647	$929	HIGH	MEDIUM
ST. JOSEPH HOSPITAL 603.882.3000	$282	$647	$929	HIGH	MEDIUM
SOUTHERN NH MEDICAL CENTER 603.577.2000	$321	$762	$1,082	MEDIUM	MEDIUM
CATHOLIC MEDICAL CENTER 603.668.3545	$369	$907	$1,276	MEDIUM	HIGH
FRANKLIN REGIONAL HOSPITAL 603.934.2060	$375	$924	$1,299	HIGH	MEDIUM
LAKES REGION GENERAL HOSPITAL 603.524.3211	$375	$924	$1,299	MEDIUM	MEDIUM
ELLIOT HOSPITAL 603.669.5300	$394	$981	$1,375	MEDIUM	HIGH
CONCORD IMAGING CENTER 603.225.0425	$457	$1,172	$1,630	HIGH	MEDIUM

(b) List of Prices

Notes: Screenshots show the New Hampshire HealthCost website as it looked in 2015. Over the period of analysis (2007 to 2010) the website had the same information.

Figure A-3
Distribution of Estimated Website Usage Cost

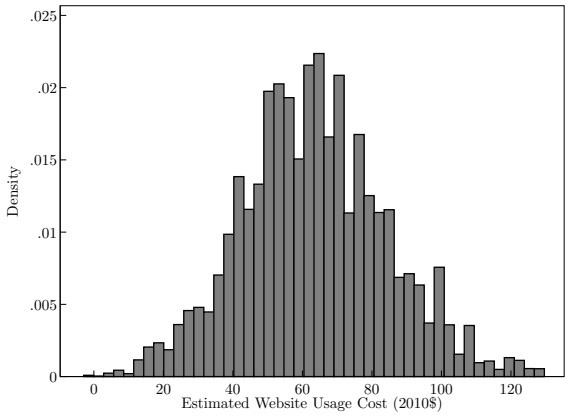


Figure A-4
Distribution of Counterfactual Negotiated Prices



(a) Price Transparency Website

(b) Full Price Transparency

(c) Price Transparency Website + High Cost Sharing

(d) Full Price Transparency + High Cost Sharing

Notes: Histograms show distribution of unweighted counterfactual prices across providers, procedures, insurers, and years.