

An Empirical Model of Price Transparency and Markups in Health Care*

Zach Y. Brown[†]

August 2019

Abstract

It is usually difficult for patients to compare out-of-pocket prices for medical services. What are the implications for prices and welfare? In order to understand the effects of price transparency, this paper develops an empirical model of demand and supply for medical imaging services that incorporates patients' limited information about prices. Estimation exploits the introduction of a price transparency website that informed a subset of patients. Counterfactual simulations imply a 22 percent reduction in prices if all patients had full information. However, the results also shed light on the barriers to widespread adoption of price transparency tools.

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

*I am very grateful to Kate Ho, Mike Riordan, and Chris Conlon for their support and guidance. I also thank Tobias Salz, Adam Kapor, Bernard Salanié, Robert Town, Ying Fan as well as seminar participants at Columbia, U Washington, Michigan, Wisconsin, Wharton, U Mass-Amherst, Imperial, INSEAD, UIUC, NBER, Princeton, and FTC. This research was supported by the National Institute on Aging, Grant Number T32-AG000186, as well as 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 Health Cost website. This paper does not reflect the views of the New Hampshire Department of Health and Human Services.

[†]Department of Economics, University of Michigan, 266 Lorch Hall, 611 Tappan Ave., Ann Arbor, MI 48109, zachb@umich.edu.

1 Introduction

In many markets, consumers do not know exact prices until they have committed to a purchase. For instance, this is 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 these firms are often contractually forbidden from disclosing negotiated rates. As a result, the vast majority of consumers say they do not compare prices before receiving medical care.² Due to the fact that prices are opaque, hospitals and other providers potentially face more inelastic demand, leading to higher prices. Although there have been some efforts to make price information more available to patients, these efforts have been quite limited in their reach. In response, some policy makers have called for more “price transparency” in health care.³

Despite the fact that the lack of price transparency is a key feature of U.S. healthcare markets, models have generally not accounted for this issue. The issue is particularly important given that privately-provided health care in the U.S. 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. 2018). Like search costs, the lack of price transparency may increase prices and lead to price dispersion.

This paper empirically evaluates how price transparency affects markups and welfare in the U.S. health care market. I estimate a demand model that explicitly accounts for consumer uncertainty about prices. During the sample period, an innovative price transparency website became available, providing information about out-of-pocket costs. Given that only some patients sought out the tool and used it, the model allows for usage of the tool to be endogenous. I then combine the demand system with a model of bargaining between providers and insurers in order to examine how consumer price transparency affects negotiated prices.

The model allows me to answer three questions. First, I evaluate the welfare effects of the price transparency website. The model also allows me to examine the mechanisms and distributional consequences of the tool. Second, I use the model to provide insight into the potential equilibrium effects if more individuals were to use price transparency tools. Doing so allows me to quantify the welfare effects of increased price transparency. Finally, I examine how cost sharing interacts with price transparency. This provides insight into why more individuals are not using price transparency tools despite the large dispersion in negotiated prices.

In order to examine these issues, this paper develops a discrete-choice model in which consumers choose where to receive medical care in the presence of information frictions. The model

¹See, for example, Ellison (2005).

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

³More than half of U.S. states have proposed health care price transparency laws in recent years. Various price transparency initiatives have also been proposed at the federal level.

⁴See 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. (2018).

allows patients to potentially have some limited information about prices even when prices are not posted publicly, an important feature of health care markets. In the model, consumers with rational expectations receive noisy signals about prices, form beliefs about prices based on the signals, then make decisions according to their beliefs. 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 key estimation challenge stems from the fact that it is difficult to determine whether individuals do not care about prices, i.e. have low price sensitivity, or do not know prices. 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 privately-insured consumers in the state to enter insurance information and easily compare accurate out-of-pocket prices across all providers in the state. I exploit 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. Individuals with the most to gain may be more likely to seek out the tool and use it. By leveraging website traffic information, I also estimate a model of website usage. If consumers use the price transparency website when it is available, I assume that they have perfect information about out-of-pocket prices. Otherwise, I allow for a discrepancy between expected prices and actual prices. Estimation of the model makes use of MCMC methods in order to recover individuals' beliefs about the prices of all options, a high dimensional latent variable.

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. 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 2017). 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.⁵ I then derive an expression for equilibrium prices and highlight two countervailing effects of price transparency on prices. Price transparency can make residual demand more elastic, decreasing the incentive for providers to negotiate high prices. This would decrease negotiated rates. On the other hand, 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 could potentially reduce the incentive of insurers to negotiate low prices. Therefore, the effect of price transparency on negotiated prices is theoretically ambiguous and it is necessary to examine these issues empirically.

The model is estimated using detailed administrative data on private health care claims and price transparency website usage in New Hampshire. The claims data contain information on negotiated prices and cost sharing for all privately-insured individuals in the state. These are the same data used to construct plan-specific out-of-pocket prices for the website. I focus on relatively simple outpatient medical imaging procedures—X-rays, CT scans, and MR scans. Despite the fact that these procedures are relatively standardized, I find that the price of each procedure varies

⁵For 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.

widely across providers in the state.⁶ 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 the first empirical exercise, the estimates are used to evaluate the effect of New Hampshire’s price transparency website. Estimates imply that the website resulted in overall savings of 4 percent for medical imaging procedures.⁷ These results are largely consistent with reduced-form results, helping to validate the model. In contrast to the reduced-form approach, the structural model allows me to disentangle the mechanisms and shed light on the welfare effects and distributional consequences. I find that welfare effects are primary due to increased price-shopping on the part of consumers, however part of the welfare gains are also due to a modest reduction in the equilibrium prices. I also use the model to examine the effect for individuals who actually used the website. Perhaps unsurprisingly, estimates show the website primarily benefited individuals subject to a deductible. These individuals saw substantial savings, about \$200 per visit. However, 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, the estimates reveal that welfare gains are substantially smaller than savings.

Next, I use the model to examine what would happen if more individuals used the website. Website traffic data imply that consumers used the website for only about 8 percent of medical imaging visits when the website was available. Given modest take-up, policy makers are interested in the potential effects if more individuals used these tools. In order to answer this question, it is important to take into account two issues that make it difficult to simply extrapolate from reduced-form estimates. First, 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 returns in terms of 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. This implies that there may be increasing returns as more individuals are informed.

Counterfactual simulations imply that, while both mechanisms are present, the effect on equilibrium prices dominates. If all consumers were informed, the model implies that equilibrium prices would be 22 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 \$39 for consumers and \$281 for insurers relative to no price transparency. 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). The results imply that even a quarter of individuals being informed is enough to generate a considerable reduction in equilibrium prices, generating a large externality

⁶This is consistent with previous research documenting the large degree of price dispersion for these procedures nationally (Cooper et al. 2018). 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.

⁷Overall savings refers to change in spending for both insurers and consumers.

for uninformed patients.

The results highlight that there would be large spending reductions if more individuals used price transparency tools. So why is take-up so low? One explanation is that low cost sharing reduces consumers' private benefit of using these tools. To examine this, I analyze how price transparency interacts with cost sharing in counterfactual simulations. Results indicate that if cost sharing is high enough, enough consumers are incentivized to use the price transparency tool to substantially reduce health care prices. In particular, a 50% coinsurance rate applied to medical imaging procedures results in prices that are 15% lower. This suggests that for the price transparency tool to generate large savings for patients and insurers, cost sharing would have to be quite high.

1.1 Related Literature

This paper is related to the large literature on search costs and competition, starting with Stigler (1961) and Diamond (1971). Search costs have been shown to be empirically important in a large variety of markets.⁸ 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. In contrast, this paper studies a context in which individuals make decisions with uncertainty about prices, and true prices are only revealed afterwards. I argue that this has important welfare implications. The model presented in this paper also 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 also contributes 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).⁹ 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 imprecise beliefs about shrouded attributes and maximize expected utility.

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. A number of studies have examined information frictions related to insurance choice (e.g. Ericson 2014; Decarolis 2015; Handel and Kolstad 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 2014). In a similar vein, a literature has examined uncertainty about quality of medical services and medical devices (e.g. Kolstad 2013; Grennan and Town 2015). Finally, Grennan and Swanson (2016) find that informa-

⁸Empirical work has studied search frictions in markets for mutual funds, textbooks, online bookstores, grocery stores, auto insurance, electricity, online hotel booking, and trade-waste (Hortaçsu and Syverson 2004; Hong and Shum 2006; De Los Santos et al. 2012; Seiler 2013; Honka 2014; Giuliotti et al. 2014; Koulayev 2014; Salz 2015).

⁹Also see Grubb (2015) for related review.

tion 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, which I argue is particularly important for understanding high health care spending.

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. 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. 2019). 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), tiered hospital networks (Prager 2016), and insurer competition (Ho and Lee 2017). With the exception of Allen et al. (2019) work examining consumer lending, empirical models of bargaining in oligopolistic markets have assumed perfect information.¹⁰ To examine the effect on prices, I use an approach closely related to Gowrisankaran et al. (2015), however I incorporate consumer uncertainty about prices and examine the various channels through which price transparency affects equilibrium prices. Overall, the effect of price transparency is ambiguous in a bargaining context, which has implications for other vertical markets where consumers have uncertainty about product characteristics.

Prior reduced-form work has examined the effect of health care price transparency efforts by individual employers or insurers. In particular, Lieber (2017) and Whaley (2015a,b) 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 (2019), I use reduced-form methods to examine the effect of the price transparency website on spending. The reduced-form approach provides evidence on the intent-to-treat effect of the price transparency website but remains silent on a number of important issues. First, it provides little insight into the mechanisms and the effect on welfare. Given that price transparency has implications for distance traveled and the quality of medical providers chosen, the effect on welfare may be quite different than the effect on spending. Perhaps most importantly, health care price transparency tools are not yet widely used, and therefore it is difficult to draw general conclusions about the role of information frictions using a reduced-form approach. By developing a model based on theory, counterfactual analysis can be used to examine what would happen if more individuals were informed about health care prices as well as how price transparency interacts with other potential policies such as cost sharing.

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. Section 4 presents the bargaining model, focusing on the role of consumer information. Section 5 presents the results. Section 6 uses the estimates to examine the effect of the website while Section 7 presents out-of-sample counterfactual simulations. Section 8

¹⁰Allen et al. (2019) incorporate search frictions into a model of the mortgage market. Note that while this paper models business-to-business bargaining, Allen et al. (2019) examines negotiations between consumers and lenders.

provides a discussion and Section 9 concludes.

2 Data and Background

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 2010.¹¹ 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. I restrict the sample to the three major insurers in the state and eliminate uncommon procedures. I describe the sample restrictions in more detail in Appendix A. The full list of medical imaging procedures is given in Table A-1.¹²

I focus on this set of procedures for a few reasons. I argue these procedures are particularly important given that medical technology, especially related to medical imaging, is often cited as a key driver of health care cost growth.¹³ Second, these procedures are relatively standardized, mitigating concerns about unobserved quality. Finally, patients have significant discretion over where to receive outpatient medical imaging procedures.¹⁴

The data contain information on out-of-pocket prices and insurer reimbursement amounts, allowing me to construct each patient's cost sharing. Importantly, prices are aggregated to the visit level, which may include multiple procedures. I use a similar methodology to aggregate prices as the price transparency website, which displays the aggregated visit prices. The construction of visit prices is described in more detail in Appendix A.

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 centers) are significantly less expensive. In New Hampshire, the average total cost of an imaging visit is \$1,004 at hospitals but only \$797 at non-hospital providers. The location of these providers, derived from their zip code, is shown in Figure A-1.

For each individuals, I observe age, gender, zip code, insurance enrollment, and whether they are subject to a deductible. 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.

¹¹Although the data include information about claims in later years, I focus on the period prior to December 2010 since this is when detailed website traffic data is available.

¹²In addition, 2011 is excluded since website traffic data is not available.

¹³See Newhouse (1992) and Cutler (1995).

¹⁴Other procedures featured on the website, such as kidney stone removal, physician office visits, and newborn delivery, tend to be less standardized and involve a different set of providers.

Table 1
Summary of Privately Insured Individuals with
Medical Imaging Claims

	Mean	SD	Min	Max
Male	0.47	0.50	0	1
Age 0-18	0.20	0.40	0	1
Age 19-35	0.18	0.39	0	1
Age 36-50	0.31	0.46	0	1
Age 51-64	0.31	0.46	0	1
Charlson Comorbidity Index	0.6	0.8	0.0	2.0
Zip income (\$1,000s)	82.8	24.2	22.0	309.3
Zip BA Degree (%)	33.8	14.0	0.0	100.0
Unique Individuals		174,672		

Notes: Sample includes all privately insured individuals in the state of New Hampshire over the period 2005 to 2010 with at least one outpatient medical imaging visit. The unit of observation is a unique individual.

In addition, patient zip code is used to calculate the distance to each provider. I also construct each patient’s Charlson Comorbidity Index, a measure of chronic diseases or conditions. Given the potential importance of primary care recommendations, I also construct an indicator for likely referrals.¹⁵ Finally, I construct an indicator for whether each individual has the medical imaging procedure in the week following an emergency.¹⁶

Table 1 provides a summary of the 174,672 individuals with outpatient imaging visits over the period. Table A-2 provides additional summary statistics. Half of the individuals are in HMO plans and most of the remainder are in PPO or POS plans. About 43 percent of individuals have a plan with a deductible.

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 each insurer network by examining the providers chosen by individuals in each insurance company-product pair (e.g. Anthem HMO).¹⁷ 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.

The full dataset is summarized for each of the three insurers in Table 2. Anthem is by far

¹⁵The referral indicator is described in more detail in Appendix A.

¹⁶Although 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).

¹⁷In 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. 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.

Table 2
Summary of Medical Imaging Visits by Insurer

	Anthem		Cigna		Harvard Pilgrim	
	Mean	SD	Mean	SD	Mean	SD
Observations	946,057		216,176		266,747	
Number of choice situations	104,358		26,864		33,275	
Number of unique patients	67,749		19,912		21,114	
Number of unique non-hospital providers	148		82		62	
Number of unique hospital providers	29		13		0	
Providers in choice set	11.7	5.2	11.9	5.9	10.3	4.3
Total negotiated price	777.2	945.9	613.3	764.7	706.3	758.7
Insurance price	697.1	907.5	577.4	739.9	627.6	711.9
Out-of-pocket price	80.1	178.7	35.8	77.9	78.8	193.2
Distance to provider	34.1	17.7	28.9	16.2	28.9	15.1
Choose hospital	0.38	0.49	0.23	0.42	0.00	0.00
Choose referral	0.33	0.47	0.41	0.49	0.33	0.47

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.

the largest insurer, with over 100,000 medical imaging visits over the period. On average, the out-of-pocket price is 13 percent of the total negotiated price. However, there is large variation—individuals under the deductible pay the full price. In general, there is greater cost sharing in the beginning of the year, when individuals have not hit their deductible, than at the end of the year.¹⁸ 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 MR scans. Given large number of observations, I use a 10 percent sample of visits for the main analysis.

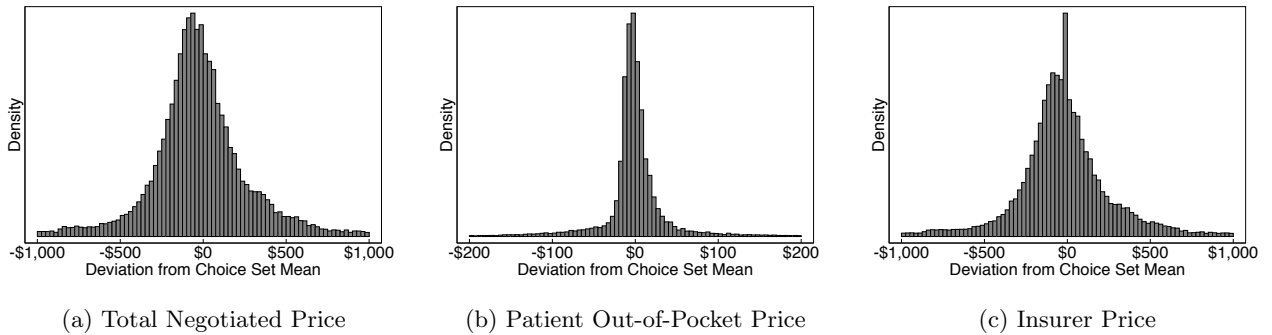
Within individual's choice sets, there is a large degree of price dispersion, and consequently, significant potential savings if individuals switch to low cost options. Figure 1a shows the distribution of demeaned negotiated prices within individuals' choice sets. The distribution is approximately normal, with standard deviation of \$747 (and coefficient of variation of 46 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 \$69 (see Figure 1b). Finally, Figure 1c 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-3. Overall, I find that there would be savings of over 30 percent if consumers switched to providers in the first quartile of the price distribution.¹⁹ This is true for across all three procedures types. Consumers subject to a deductible have large gains from switching, but much

¹⁸In January, out-of-pocket prices are about 20 percent of the total negotiated price on average.

¹⁹These are the potential consumer 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.

Figure 1
Price Variation within Individuals' Choice Sets



Notes: Histograms show distribution of prices across individuals' choice sets. Prices are demeaned since individuals face different choice sets and prices depending on their insurance plan, procedure, year, geographic location, and if they have surpassed their deductible.

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, consumers with low cost sharing 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 2007.²⁰ Other price transparency initiatives 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 that insured individuals actually pay. New Hampshire's HealthCost website was unique because it provided information about insurer-specific out-of-pocket prices for the entire visit. 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.²¹ 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. After my period of analysis, the website added information about provider quality and a guide to health insurance.²²

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 illustrates an example of prices returned by

²⁰The website can be found at nhhealthcost.nh.gov. Originally the website was nhhealthcost.org.

²¹New Hampshire was the only state to receive an "A" grade from Catalyst for Payment Reform's 2015 Report Card on State Price Transparency Laws.

²²It 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.

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 3.2.²³

The website was widely promoted and there were 41,506 searches for price information per year on average according to website traffic logs, about a third of which were for medical imaging procedures.²⁴ Surveys of people in New Hampshire found that 60 percent who used the website reported saving money.²⁵ 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.”²⁶ In particular, Anthem, the largest insurer in New Hampshire, had a public battle with an expensive hospital in the state and local news sources suggest that the price transparency website facilitated lower prices.²⁷

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 January 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.²⁸ The X-ray, CT scan, and MR 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, even though the website includes other CT scans. One concern is that the website indirectly affected prices for procedures not on the website due to cross-subsidization, a concern I address in Brown (2019) by showing that results are robust to exploiting cross-state variation. Using both time variation and cross-sectional variation is important given that there are potential unobserved shocks that affect all imaging procedures.²⁹

I use website traffic logs obtained from the New Hampshire Insurance Department to calculate

²³The 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 (i.e. estimates are fairly precise) and relatively standardized (i.e. price depends little on patient complexity).

²⁴According 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.

²⁵See “How New Hampshire Residents Use Health Care Price Information,” April 2017. Public Agenda Research Brief.


²⁶See “Moving Markets: Lessons from New Hampshire’s Health Care Price Transparency Experiment,” April 2014. California HealthCare Foundation.

²⁷See “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.

²⁸According 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.

²⁹For instance, Medicare reduced payments to imaging procedures starting in 2007. To the extent that this affected private-payer prices, this is unlikely to differentially impact prices for procedures on the website.

Figure 2
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.

the number of website price searches in each month for each procedure listed on the website. Website traffic data, obtained from server logs, is available from March 2007 through November 2010, at which point the website switched hosting companies. Figure 2 shows monthly price searches for X-rays, CT scans, and MR 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, likely due to more individuals learning about the website.

In order to estimate the fraction of informed consumers, I divide the number of price searches per procedure by the total number of visits in New Hampshire from the claims data.³⁰ In other words, I assume that each use of the website is a unique individual.³¹ Overall, this implies about 8 percent of individuals receiving a procedure on the website were informed. This is higher than usage of other price transparency tools, potentially due to the fact that New Hampshire’s price transparency tool is easily accessible and utilizes high quality data on actual negotiated prices, which has not been the case for other price transparency efforts.³²

³⁰Note 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.

³¹If 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.

³²New Hampshire’s tool has consistently received higher marks than other tools. See Catalyst for Payment Reform’s 2015 Report Card on State Price Transparency Laws.

Table A-4 has the estimated percent of consumers with price information for each medical imaging procedure listed on the website. The percent of informed consumers tends to be higher for CT scans and MR scans compare to X-rays. CT scans and MR scans also tend to be more expensive, making the website potentially more valuable. There is also variation across time due in part to the fact that there is random variation in the type of individuals that need a procedure in a given month. This variation is used to help estimate the demand model and recover information about the choice to use the website if it is available.

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, I model the selection of individuals that use the price transparency website if it is available. Second, based on their information set, consumers choose a medical provider.

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. Finally, I discuss estimation and identification.

3.1 Model Setup and Timing

There are a set of providers that perform 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.³³ 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.³⁴

Individual i pays fraction c_{ikmt} of the negotiated price. The cost sharing fraction is observed in the claims data by calculating out-of-pocket costs as a fraction of total cost for each plan. The degree of cost sharing is determined by both the coinsurance rate applied to procedure 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

³³Given 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.

³⁴Focusing 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.

individual, cost sharing can vary over time t . The out-of-pocket price paid by the individual is

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

I assume that individuals do not anticipate whether they will surpass their annual deductible and respond only to this spot price. This is consistent with the fact that, in the reduced-form analysis, much of the price transparency effect is from individuals subject to a deductible, including those that later surpass their deductible. It is also consistent with the previous findings of myopic behavior in health care (Brot-Goldberg et al. 2017).³⁵

The remainder of the cost 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 utility³⁶

After choosing a provider and receiving the procedure, the individual receives a bill and learns the true price.

3.2 Choice of Provider

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. For individual i with insurance k receiving procedure m from medical provider j , indirect utility is assumed to take the additively separable form

$$u_{ijkmt} = -\gamma_i p_{ijkmt}^{OP} + \underbrace{\alpha_1 d_{ij} + \alpha_2 d_{ij}^2 + \alpha_3 r_{ijt} + \xi_{jM} + \beta \mathbf{x}_{ikmt} h_j}_{\delta_{ijkmt}} + \varepsilon_{ijkmt} \quad (1)$$

³⁵In contrast to the evidence from health care, the evidence from prescription drug insurance is mixed. See Aron-Dine et al. (2015), Abaluck et al. (2015), and Sacks et al. (2016).

³⁶This 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 idea that consumers may evaluate providers based on observable characteristics, choose to use the website if it is available, and only then learn about their idiosyncratic shock, such as whether providers have an appointment time that fits their schedule.

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. 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 into more generous plans.³⁷ 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 and may be correlated with cost sharing, i.e. $\gamma_i \sim N(\bar{\gamma} + \rho c_{ik}, (\sigma^\gamma)^2)$. In order to account for correlation in unobserved utility when individuals have multiple medical imaging visits over the period, the random-coefficient on price is assumed to be 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 referred to provider j , r_{ijt} . Controlling for referrals is important as physicians may influence where patients choose to go and there may be benefits to receiving a medical imaging procedures from a provider closely connected to a patient’s specialist.

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, gender, 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 MR 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. 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.³⁸

The observed choice probability of individual i enrolled in insurer k receiving procedure m at

³⁷This 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.

³⁸One concern is that price transparency affects the choice to have a procedure at all. In Brown (2019), I examine the effect of the price transparency website on the probability of having medical imaging procedures for all privately-insured individuals in the state and do not find a statistically significant effect. This is consistent with the fact that X-rays, CT scans, and MR scans tend to be less discretionary than other medical services. This implies that conditioning on individuals that had a medical imaging procedure and assuming they all choose an inside option is unlikely to bias counterfactual estimates. However, price transparency tools may affect quantity for other types of procedures such as preventative care.

time t conditional on price information is then

$$s_{ijkmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt} | \vartheta_{ijkmt} = 1) = \int_{\gamma_i} \frac{\exp(-\gamma_i p_{ijkmt}^{OP} + \delta_{ijkmt})}{\sum_{j' \in \mathcal{N}_{kmt}} \exp(-\gamma_i p_{ij'kmt}^{OP} + \delta_{ij'kmt})} f(\gamma_i) d\gamma_i \quad (2)$$

where ϑ_{ijkmt} 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:³⁹

$$CS_{ijkmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt} | \vartheta_{ijkmt} = 1) = \frac{1}{\gamma_i} \log \left(\sum_{j \in \mathcal{N}_{kmt}} \exp(-\gamma_i p_{ijkmt}^{OP} + \delta_{ijkmt}) \right) \quad (3)$$

Provider Choice with Price Uncertainty

A defining feature of the health care market is that precise price information is, in general, not available to patients. Previous models of hospital demand either assume that individuals do not account for hospital prices at all or have perfect information about prices.⁴⁰ Consistent with the fact that demand for medical services is not perfectly inelastic even when price transparency tools are not available, the model allows individuals to have limited information about prices.⁴¹ Individuals form beliefs using Bayes' rule and then make a decision based on those beliefs. This information structure is related to the empirical literature on consumer learning and investment decisions with imperfect information (e.g. Erdem and Keane 1996; Akerberg 2003; Erdem et al. 2008; Crawford and Shum 2005; Ching 2010; Dickstein 2014; Grennan and Town 2015; Dickstein and Morales 2016).

I assume that individuals know the distribution from which prices are drawn, which is assumed to be normal. In particular, their prior is determined by the true mean and variance of prices in their choice set, \bar{p}_{kmt}^{OP} and \bar{s}_{kmt}^2 respectively:

$$p_{ijkmt}^{OP} \stackrel{iid}{\sim} N(\bar{p}_{kmt}^{OP}, \bar{s}_{kmt}^2) \quad (4)$$

The demeaned distribution of prices across options in individuals' choice sets can be seen in Figure 1b, which is approximately normal. For individuals subject to a deductible, \bar{s}_{kmt}^2 can be quite large (see Figure 1a).

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

³⁹This 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).

⁴⁰For example, Kessler and McClellan (2000), Tay (2003), Ho (2006), and Ho and Lee (2017) 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 utility and assume perfect information.

⁴¹When price transparency tools are not available, demand is still observed to be somewhat elastic, implying individuals have at least some information about prices. This can be seen in Table A-5, as well as other studies examining demand for hospitals (e.g. Gowrisankaran et al. 2015).

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.⁴² I model this by assuming that individuals receive a vector of unbiased signals, where each signal is given by

$$p_{ijkmt}^{OP} + e_{ijkmt} \quad (5)$$

where p_{ijkmt}^{OP} is the true price and e_{ijkmt} is signal noise with density $f(e_{ijkmt})$. I assume the distribution of signal noise is normal:

$$e_{ijkmt} \stackrel{iid}{\sim} N(0, \sigma^2) \quad (6)$$

The key parameter is σ^2 , which can be thought of as a measure of price transparency (or opacity).

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

$$\mathbb{E} \left[\widetilde{p_{ijkmt}^{OP}} \right] = w_{ikmt}(p_{ijkmt}^{OP} + e_{ijkmt}) + (1 - w_{ikmt})\bar{p}_{kmt}^{OP} \quad (7)$$

where the weight given to the signal is defined as $w_{ikmt} = \bar{s}_{kmt}^2 / (\bar{s}_{kmt}^2 + \sigma^2)$. Using the assumption that the prior and signal are normally distributed, the variance of posterior beliefs is $Var[\widetilde{p_{ijkmt}^{OP}}] = w_{ikmt}\sigma^2$.

If $\sigma^2 = 0$, then $w_{ikmt} = 1$ and individuals know true prices. Conversely, if $\sigma^2 \rightarrow \infty$ then $w_{ikmt} \rightarrow 0$, 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.⁴³ In Section D, I present an alternative model in which individuals have an uninformative prior and take price signals as given.

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

$$\begin{aligned} \mathbb{E} [\widetilde{u_{ijkmt}}] &= -\gamma_i \mathbb{E} \left[\widetilde{p_{ijkmt}^{OP}} \right] + \underbrace{\alpha_1 d_{ij} + \alpha_2 d_{ij}^2 + \alpha_3 r_{ijt} + \xi_{jM} + \beta \mathbf{x}_{ikmt} h_j}_{\delta_{ijkmt}} + \varepsilon_{ijkmt} \\ &= -\gamma_i w_{ikmt}(p_{ijkmt}^{OP} + e_{ijkmt}) + \delta_{ijkmt} + \varepsilon_{ijkmt} \end{aligned} \quad (8)$$

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

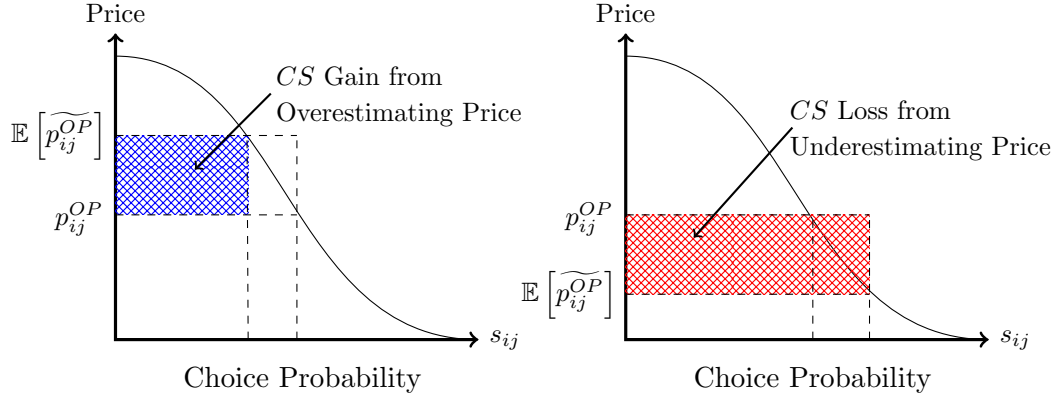
In this model, all individuals receive signals with the same variance, parameterized by σ . This is motivated by the fact that individuals have access to similar public information, and therefore

⁴²See, for instance, “How to Research Health Care Prices,” *Wall Street Journal*, December 4, 2009.

⁴³This can be seen formally by noting that the choice probability depends only on non-price attributes in the limit, i.e.

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

Figure 3
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.

have a similar degree of uncertainty about prices. It is important to note that the model rules out learning over time. Only 7 percent of patients get repeat outpatient procedures in a year, and even then may still have uncertainty due to changes in deductible status.⁴⁴

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, $p_{ijkmt}^{OP} + e_{ijkmt}$, but not the true price. However, the researcher observes the true price, p_{ijkmt}^{OP} , 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

$$s_{ijkmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt} | \vartheta_{ijkmt} = 0) = \int_{\gamma_i} \int_{\mathbf{e}_{ikmt}} \frac{\exp(-\gamma_i w_{ikmt}(p_{ijkmt}^{OP} + e_{ijkmt}) + \delta_{ijkmt})}{\sum_{j' \in \mathcal{N}_{kmt}} \exp(-\gamma_i w_{ikmt}(p_{ij'kmt}^{OP} + e_{ij'kmt}) + \delta_{ij'kmt})} f(\mathbf{e}_{ikmt}) f(\gamma_i) d^J \mathbf{e}_{ikmt} d\gamma_i \quad (9)$$

where $\vartheta_{ijkmt} = 0$ indicates that the individual did not use the website and is uninformed about prices. Since the vector of signal noise, \mathbf{e}_{ikmt} , has the same number of elements as \mathcal{N}_{kmt} , computing the expectation over individual beliefs requires evaluating a potentially high dimensional integral. This is the key estimation challenge, an issue that is discussed in greater detail in Section 3.4.

The calculation of expected consumer surplus must take into account that, from the perspective of the individual, the expected price, $\mathbb{E}[p_{ijkmt}^{OP}]$, may differ from true price, p_{ijkmt}^{OP} . 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

⁴⁴I find that the effect of the website is similar for patients with and without a history receiving medical imaging procedures, indicating that learning about prices over time is not a first order concern for these procedures. See Brown (2019) Appendix.

(Small and Rosen 1981) as well as a term that captures the loss from incorrect beliefs:

$$\begin{aligned}
CS_{ikmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt} | \vartheta_{ikmt} = 0) &= \underbrace{\frac{1}{\gamma_i} \log \left(\sum_{j \in \mathcal{N}_{kmt}} \exp(-\gamma_i w_{ikmt} (p_{ijkmt}^{OP} + e_{ijkmt}) + \delta_{ijkmt}) \right)}_{\text{CS evaluated at beliefs of prices}} \\
&+ \underbrace{\sum_{j \in \mathcal{N}_{kmt}} \left[\mathbb{E} \left[\widetilde{p_{ijkmt}^{OP}} \right] - p_{ijkmt}^{OP} \right] s_{ijkmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt} | \vartheta_{ikmt} = 0)}_{\text{CS gain/loss from incorrect beliefs}}
\end{aligned} \tag{10}$$

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

$$\mathbb{E} \left[\widetilde{p_{ijkmt}^{OP}} \right] - p_{ijkmt}^{OP} = w_{ikmt} e_{ijkmt} + (1 - w_{ikmt}) (\bar{p}_{kmt}^{OP} - p_{ijkmt}^{OP}) \tag{11}$$

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 3 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 seek out the price transparency website and use it if it is available. The model allows for selection by assuming that individuals with the most to gain may be more likely to use the website. Specifically, 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. In Section 7, the estimates from this selection model are used to simulate website usage under counterfactual scenarios.

The expected consumer surplus gain of using the website is the difference between expected consumer surplus when consumers know they will receive price information, but have not yet received it, and expected consumer surplus without price information. The expected benefit of using the website, b_{ikmt} , is given by

$$b_{ikmt} = \frac{\gamma_i w_{ikmt} \sum_{j \in \mathcal{N}_{kmt}} \sigma_h^2 \left[\exp(-\gamma_i \mathbb{E} \left[\widetilde{p_{ijkmt}^{OP}} \right] + \delta_{ijkmt}) \Phi_{ijkmt} \right]}{2 \left[\sum_{j \in \mathcal{N}_{kmt}} \exp(-\gamma_i \mathbb{E} \left[\widetilde{p_{ijkmt}^{OP}} \right] + \delta_{ijkmt}) \right]^2} \tag{12}$$

where $\Phi_{ijkmt} \equiv \sum_{j' \in \mathcal{N}_{kmt} \setminus j} \exp(-\gamma_i \mathbb{E} \left[\widetilde{p_{ij'kmt}^{OP}} \right] + \delta_{ij'kmt})$. Given the lack of a closed-form expression, which is needed for computational tractability, the expression above is an approximation using a second-order Taylor series expansion. In Appendix B, I derive this expression and argue that the approximation is quite accurate.

Holding beliefs about prices fixed, an increase in price uncertainty, as measured by σ^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 is increasing in the magnitude of the individual-specific price sensitivity parameter, γ_i .

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.⁴⁵ 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, gender, income, education, Charlson Comorbidity Index, emergency indicator, and year fixed effects in order to account for the fact that more individuals may hear about the website over time, reducing the implicit cost.⁴⁶ I also include a constant.

Individuals use the website if the net benefit is positive

$$\underbrace{\theta b_{ikmt}}_{\text{Website Benefit}} - \underbrace{\phi \mathbf{x}_{ikmt} + \nu_{ikmt}}_{\text{Website Cost}} > 0 \quad (13)$$

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

$$\vartheta_{ikmt} = \frac{\exp(\theta b_{ikmt} - \phi \mathbf{x}_{ikmt})}{1 + \exp(\theta b_{ikmt} - \phi \mathbf{x}_{ikmt})} \quad (14)$$

where θ and ϕ are parameters to be estimated.

Website traffic logs provide an estimate of the number of individuals that 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 $\vartheta_{mt} = \frac{1}{n_{mt}} \sum_{i \in \mathcal{I}_{mt}} \vartheta_{ikmt}$ where n_{mt} is the number of individuals receiving procedure m in month t .

3.4 Joint Estimation of Demand

Likelihood Function and Estimation

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. If the website is not available, individuals have uncertainty about prices and choice probabilities are given by Equation 9. If the website is available, the choice probabilities are given by Equation 2 if the website is actually used, where the probability of using the website is given by Equation 14. The second component of the

⁴⁵See "State of New Hampshire Broadband Action Plan," New Hampshire Department of Resources and Economic Development and Telecommunications Advisory Board, June 30, 2008.

⁴⁶This also accounts for the fact that more consumers have broadband internet over time.

likelihood function is the probability of actual website traffic given predicted website usage. The likelihood function is described in greater detail in Appendix C.

To estimate the model, I use a Markov chain Monte Carlo (MCMC) estimator to simulate the posterior distribution of $\Theta = (\sigma^\gamma, \rho, \alpha, \xi, \beta, \sigma, \theta, \phi)$. The key estimation challenge is computing high dimensional integrals in order to find the expectation over \mathbf{e}_{ikmt} , individuals' unobserved beliefs. The standard maximum likelihood estimation strategy is to use simulation methods and draw from $f(\mathbf{e}_{ikmt})$, 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. Given the high dimensionality of \mathbf{e}_{ikmt} , this approach is not computationally feasible. The MCMC estimator overcomes the curse of dimensionality by sampling the parameter space conditional on the data.

I employ Hamiltonian Monte Carlo (HMC), a variant of MCMC.⁴⁷ 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 posterior estimation, I use 4 chains, each with 2,000 warmup draws, which are discarded, and 2,000 sample draws. See Appendix C for additional estimation details.

It is important to note that using Bayesian methods for estimation does not impose additional assumptions since I use uninformative priors for all of the structural parameters.⁴⁸ The use of MCMC is primarily motivated by the fact that it is computationally attractive.

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. I overcome this issue by exploiting quasi-random variation due to the introduction of the price transparency website.

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 by the fact that the price of a given provider varies depending on an individual's insurer, if 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.

⁴⁷This approach, developed by Hoffman and Gelman (2014), uses the gradient of the log posterior density to more efficiently sample the posterior distribution and ensure that sampling does not double back on the parameter space. This algorithm is implemented in the Stan programming language, which I use to automatically compute gradients and estimate the model.

⁴⁸In supplemental material, I examine a simplified version of the model with a small choice set and show that point estimates and standard errors obtained via simulated maximum likelihood are very similar to those obtained via MCMC estimation.

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 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 signal variance.⁴⁹

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 observed website traffic helps identify θ , while correlation between observed characteristics of consumers and observed website traffic helps identify ϕ , the cost.

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 extend the bargaining model from Gowrisankaran et al. (2015) by incorporating information frictions. Using the estimates from the demand model given in the previous section, I use the 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.

⁴⁹The 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.

4.1 Bargaining Model

In each year, insurer k negotiates the price of procedure m with each provider in the insurer's network, $j \in \mathcal{N}_{kmt}$.⁵⁰ For the analysis, I assume that each provider negotiates independently.⁵¹ I also take the set of providers \mathcal{J} and networks \mathcal{N}_{kmt} as given.⁵²

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

$$\Pi_{ijkmt}^J(\mathcal{N}_{kmt}, \mathbf{p}_{kmt} | \vartheta_{ikmt}) = s_{ijkmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt} | \vartheta_{ikmt}) [p_{jkmt} - mc_{jkmt}] \quad (15)$$

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

$$TC_{ikmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt} | \vartheta_{ikmt}) = \sum_{j \in \mathcal{N}_{kmt}} p_{jkmt} (1 - c_{ikmt}) s_{ijkmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt} | \vartheta_{ikmt}) \quad (16)$$

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 3). 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 10.

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

$$\Pi_{ikmt}^K(\mathcal{N}_{kmt}, \mathbf{p}_{kmt} | \vartheta_{ikmt}) = \zeta CS_{ikmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt} | \vartheta_{ikmt}) - TC_{ikmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt} | \vartheta_{ikmt}) \quad (17)$$

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

$$\Delta_j \Pi_{ikmt}^K(\mathcal{N}_{kmt}, \mathbf{p}_{kmt} | \vartheta_{ikmt}) = \Pi_{ikmt}^K(\mathcal{N}_{kmt}, \mathbf{p}_{kmt} | \vartheta_{ikmt}) - \Pi_{ikmt}^K(\mathcal{N}_{kmt} \setminus j, \mathbf{p}_{kmt} | \vartheta_{ikmt}) \quad (18)$$

Equation 17 and Equation 18 can be thought of as a stylized approach to modeling the in-

⁵⁰While 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.

⁵¹In contrast, Gowrisankaran et al. (2015) and Ho and Lee (2017) 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.

⁵²It 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.

surer’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. Unlike Ho and Lee (2017), I do not explicitly model insurer competition. Given that I find no reduced-form effect of price transparency on insurance choice, I argue that holding insurer competition fixed is relatively innocuous in this context.⁵³

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 given the equilibrium prices at other providers, $\mathbf{p}_{kmt}^* \setminus p_{jkmt}$. In other words, a hypothetical disagreement is assumed to not affect other prices. 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

$$\arg \max_{p_{jkmt}} \left(\underbrace{\sum_{i \in I_{kmt}} \mathbb{E}_e \left[\Pi_{ijkmt}^J(\mathcal{N}_{kmt}, p_{jkmt}, \mathbf{p}_{kmt}^* \setminus p_{jkmt} | \vartheta_{ikmt}) \right]}_{\text{Provider gains from trade}} \right)^{\tau_h} \left(\underbrace{\sum_{i \in I_{kmt}} \mathbb{E}_e \left[\Delta_j \Pi_{ikmt}^K(\mathcal{N}_{kmt}, p_{jkmt}, \mathbf{p}_{kmt}^* \setminus p_{jkmt} | \vartheta_{ikmt}) \right]}_{\text{Insurer gains from trade}} \right)^{1-\tau_h} \quad (19)$$

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 $\tau_h \in [0, 1]$, which is allowed to vary based on whether the provider is a hospital. Since insurers and providers do not know the price signals that consumers will receive, both take the expectation over consumer beliefs.⁵⁴

Unlike the previous literature, I model bargaining for each procedure separately rather than for a price index. This is important for capturing the fact that there is not a simple scaling of procedure prices for each insurer.⁵⁵ In addition, separate bargaining is needed to explain observed changes in negotiated prices for procedures on the website relative to procedures not on the website. One concern is that bargaining over a given procedure is not independent of other procedures. For instance, hospitals may be able to leverage the fact that they provide services other than medical imaging when negotiating with insurers. One reason for allowing the bargaining parameter, τ_h , to vary by whether the provider is a hospital is that it helps account for this issue without explicitly modeling all inpatient and outpatient procedures.

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. As in the previous literature, I assume that insurers and providers have full information. Price transparency indirectly affects equilibrium prices since changes in consumer behavior affect the gains from trade. In the context of hospital-supplier bargaining, Grennan and Swanson (2016) find

⁵³In Brown (2019), I examine whether insurance enrollment changed after the introduction of the price transparency website and do not find a statistically significant effect.

⁵⁴The providers and insurers know the variance of the price signals, σ^2 . In practice, beliefs are simulated by drawing from the distribution of \mathbf{e}_{ikmt} , computing each term, and then averaging over the draws.

⁵⁵For example, negotiated prices for X-rays are higher for Anthem than for Harvard Pilgrim. However, the opposite is true for MR scans.

evidence that price transparency affects negotiated prices in a way that is consistent with a theoretical model of bargaining under asymmetric information. While hospital-insurer bargaining may also involve asymmetric information, I argue that this concern is mitigated in the context of New Hampshire since the price transparency initiatives allowed insurers and providers access to information about prices for all procedures, not just those on the website.⁵⁶ By exploiting variation in the procedures listed on the website, which was targeted at consumers, I argue that the model isolates the effect of consumer information rather than firm information. Further research is needed to understand whether price transparency affects provider-insurer bargaining directly.

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 19 implies that equilibrium prices are determined by marginal cost plus a margin⁵⁷

$$p_{jkmt} = mc_{jkmt} + \left(\frac{1 - \tau_h}{\tau_h} \frac{\frac{\partial}{\partial p_{jkmt}} [\sum_{i \in I_{km}} \mathbb{E}_e \Pi_{ikmt}^K(\mathcal{N}_{kmt}, \mathbf{p}_{kmt} | \vartheta_{ikmt})]}{\sum_{i \in I_{km}} \mathbb{E}_e \Delta_j \Pi_{ikmt}^K(\mathcal{N}_{kmt}, \mathbf{p}_{kmt} | \vartheta_{ikmt})} - \frac{\frac{\partial}{\partial p_{jkmt}} [\sum_{i \in I_{km}} \mathbb{E}_e s_{ijkmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt} | \vartheta_{ikmt})]}{\sum_{i \in I_{km}} \mathbb{E}_e s_{ijkmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt} | \vartheta_{ikmt})} \right)^{-1} \quad (20)$$

I present further detail, including the derivation of $\frac{\partial}{\partial p_{jkmt}} [\sum_{i \in I_{km}} \mathbb{E}_e \Pi_{ikmt}^K(\mathcal{N}_{kmt}, \mathbf{p}_{kmt} | \vartheta_{ikmt})]$, in Appendix E. 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 into account changes in individuals' priors, and therefore hold the prior fixed when solving the first order condition.

The Nash-in-Nash bargaining model nests the standard Bertrand-Nash pricing assumption when $\tau_h = 1$. 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.⁵⁸ This corresponds to the case in which $\tau_h < 1$. 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

⁵⁶Firms had access to public-use versions of the NHCHIS dataset. Conversations with market participants indicate they were likely informed about prices independently of the website.

⁵⁷For simplicity, I omit the * used to indicate equilibrium outcomes.

⁵⁸Under Bertrand-Nash pricing, providers would be able to set prices unilaterally. In the absence of price information, demand is extremely inelastic, implying extremely large markups. In many cases, Bertrand-Nash pricing would imply negative marginal cost.

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. Insurers always wish to negotiate the lowest prices possible. Holding provider incentives fixed, insurers have the greatest ability to demand lower prices when they can credibly threaten to drop a provider from their network, i.e. when the insurer's gains-from-trade are low. Price transparency has an ambiguous effect on insurer's gains from trade, and therefore has an ambiguous effect on equilibrium prices. This can be seen by noting that insurer gains from trade are, in part, a function of the change in cost, $\Delta_j TC_{ikmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt} | \vartheta_{kmt})$.⁵⁹ Price transparency allows insurers to steer patients to low-price providers, lowering the cost of having expensive providers in the network. This in turn reduces the incentives of insurers to negotiate lower prices with these providers by changing $\Delta_j \Pi_{ikmt}^K$ and Π_{ikmt}^K in Equation 20. 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 then 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

$$mc_{jkmt} = \eta_j + \eta_m + \eta_t + \varepsilon_{jkmt}^{MC} \quad (21)$$

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_{jkmt}^{MC} . I assume providers have constant returns to scale.

Using the parameterized marginal cost above along with the first-order condition given by Equation 20, the marginal cost error is given by

$$\varepsilon_{jkmt}^{MC} = \eta_j + \eta_m + \eta_t - p_{jkmt} + \left(-\frac{1 - \tau_h}{\tau_h} \frac{\frac{\partial}{\partial p_{jkmt}} \sum_{i \in I_{km}} [\zeta CS_{ikmt} - TC_{ikmt}]}{\sum_{i \in I_{km}} [\zeta \Delta_j CS_{ikmt} - \Delta_j TC_{ikmt}]} - \frac{\frac{\partial}{\partial p_{jkmt}} [\sum_{i \in I_{km}} s_{ijkmt}]}{\sum_{i \in I_{km}} s_{ijkmt}} \right)^{-1} \quad (22)$$

This is used to form a moment condition, $\mathbb{E}[\varepsilon_{jkmt}^{MC} | Z_{jkmt}] = 0$, where Z_{jkmt} is a vector of variables assumed to be exogenous. The model assumes that the bargaining participants know mc_{jkmt} , including ε_{jkmt}^{MC} , 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.⁶⁰ Although these instruments are correlated with price, it is assumed that they are uncorrelated with ε_{jkmt}^{MC} . The

⁵⁹Note that 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 3 and Equation 10).

⁶⁰These 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.

instrument set, Z_{jkmt} , also includes all marginal cost fixed effects.

Identification of parameters η , τ_h , 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 E). 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_h . 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_h .

5 Results

5.1 Estimates from Demand Model

Estimates from Multinomial Logit Model

Before presenting the results from the full demand model, I start by examining a naive model in which I interact the availability of the website and the price coefficient rather than explicitly model individuals beliefs.⁶¹

Table A-5 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. This provides further evidence that the website had a meaningful impact on consumer behavior.

It is important to note that the estimates from this model lack a straightforward interpretation and do not allow for welfare calculations. 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

Table 3 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.⁶²

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

⁶¹I assume representative utility takes the form $-\gamma_1 q_{mt} p_{ijkmt}^{OP} - \gamma_2 (1 - q_{mt}) p_{ijkmt}^{OP} + \delta_{ijkmt}$ 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.

⁶²Note that for explanatory variables that overlap with the simple logit model, the standard deviation of the parameter posterior distributions are similar to the standard errors reported in Table A-5.

Table 3
MCMC Estimates for Demand Model

	Specification 1		Specification 2	
	Estimate	SD	Estimate	SD
<i>Provider Choice Parameters</i>				
OOP Price Mean ($-\bar{\gamma}$)	-0.0110	(0.0003)	-0.0172	(0.0002)
OOP Price SD (σ^γ)	0.0003	(0.0002)	0.0004	(0.0001)
OOP Price \times Cost Sharing (ρ)	-0.0081	(0.0012)	-0.0004	(0.0002)
Distance (α_1)	-0.0351	(0.0018)	-0.0351	(0.0017)
Distance Squared (α_2)	0.0027	(0.0002)	0.0027	(0.0002)
Referral Indicator	2.581	(0.025)	2.544	(0.024)
Hospital \times Age 19-35	-0.018	(0.084)	-0.024	(0.084)
Hospital \times Age 36-50	0.166	(0.073)	0.180	(0.074)
Hospital \times Age 51-64	0.113	(0.077)	0.097	(0.076)
Hospital \times Male	-0.018	(0.052)	-0.048	(0.052)
Hospital \times Income	-0.003	(0.001)	-0.002	(0.001)
Hospital \times BA	-0.003	(0.003)	-0.003	(0.003)
Hospital \times Charlson	-0.078	(0.036)	-0.066	(0.035)
Hospital \times Emergency	1.353	(0.076)	1.283	(0.076)
<i>Website Choice Parameters</i>				
Benefit (θ)	0.012	(0.006)	0.004	(0.005)
Cost (ϕ)				
Constant	1.392	(0.482)	1.723	(0.575)
Age 19-35	-1.421	(0.382)	-1.648	(0.489)
Age 36-50	-1.558	(0.342)	-1.954	(0.438)
Age 51-64	-1.739	(0.331)	-2.115	(0.421)
Male	0.085	(0.161)	0.108	(0.179)
Income	0.023	(0.006)	0.024	(0.006)
BA	-0.036	(0.008)	-0.041	(0.009)
Charlson Comorbidity	0.254	(0.088)	0.307	(0.096)
Outpatient Emergency	14.986	(5.811)	15.043	(5.946)
Year: 2007	0.384	(0.081)	0.387	(0.081)
Year: 2008	0.296	(0.073)	0.315	(0.073)
Year: 2009	0.226	(0.075)	0.231	(0.076)
Price Signal (σ)	176.3	(9.3)	96.7	(10.1)
Observations	109,626		109,626	

Notes: Table shows the mean and standard deviation of the posterior distribution estimated via MCMC. Specification 1 refers to the 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 .

interpreted as consumer’s underlying price sensitivity when prices are known, i.e. underlying marginal utility of income, rather than observed price sensitivity. There is 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. Coefficients on other variables are largely consistent with the results from the naive logit model.

The estimate of σ is shown at the bottom of Table 3. The estimates imply that, in the absence of price information, individuals have a large degree of uncertainty about prices. Given that the interpretation of σ is complex, it is useful to consider an example from the data. Figure A-3 shows a sample individual choosing between six providers that range in price from about \$200 to \$600. Using the estimate of σ , 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 are quite heterogenous, implying that there is a non-trivial chance that the individual will believe the expensive options (such as option 6), are actually inexpensive.

I compare uninformed consumers beliefs about prices with the true price. On average, there is a 28 percent absolute difference between beliefs and true prices. The gap is even larger for individuals under the deductible—41 percent. These information frictions effectively make residual demand more inelastic. The implied price elasticity of demand evaluated at mean prices when consumers are uninformed is only -0.19, however it would be -0.38 if they were fully informed.⁶³

Turning to the website choice parameters, the coefficient on the monetary benefit of using the website, θ , is positive. 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. Examining the coefficients on explanatory variables that make up the observable part of the cost of using the website, there is evidence that higher income consumers have a higher cost, perhaps because of higher opportunity cost of time. At the same time, more educated individuals have a lower cost of using the website, consistent with the fact that they are more likely to be proficient internet users. Patients receiving a procedure after an emergency episode have a higher cost. Furthermore, there is a 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 \$48 on average.⁶⁴ Note that the magnitude of θ is relatively small, indicating that, while selection is present, unobserved factors, such as word-of-mouth or internet proficiency, are important for determining website usage.

Specification 2 in Table 3 shows estimates for a model in which individuals do not know the distribution of prices before receiving signals. The parameter estimates are largely consistent with the main specification, however the welfare implications differ. This details of this alternative model, along with the implications, are described in greater detail in Section D. I argue that the alternative model is less realistic given that individuals are likely to ignore prices more as the variance of the signal noise increases.

⁶³For each procedure at each provider, residual demand elasticity for the general case in which consumers are uninformed is calculated as $(p_{jm}^{OP}/s_{jm}) \frac{1}{N} \sum_{i,k,t} \gamma_i w_{ikmt} s_{ijkmt} (1 - s_{ijkmt})$, where all expressions are evaluated at the mean price. The residual demand for each procedure at each provider is then averaged.

⁶⁴Given that this reflects awareness of the website, I do not include the cost of using the website in welfare calculations.

5.2 Estimates from Supply Model

Table 4 provides results from the bilateral bargaining model. The estimated bargaining weight is 0.29 for non-hospitals and 0.44 for hospitals, implying that insurer incentives are important for explaining equilibrium prices. Consistent with the idea that hospitals have greater leverage in negotiations all else equal, hospitals have a larger bargaining parameter. Overall, these bargaining parameters are reasonable when compared with other estimates of hospital bargaining power in the literature.⁶⁵

Table 4
Bargaining Model Estimates

	Estimate	SE
Bargaining Weight: Base ($\tau_{h=0}$)	0.293	(0.062)
Bargaining Weight: Hospital ($\tau_{h=1} - \tau_{h=0}$)	0.146	(0.091)
Insurer CS Weight (ζ)	0.838	(0.196)
Procedure FE		Yes
Provider FE		Yes
Year FE		Yes
Observations	4,832	

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

Insurers acting as perfect agents for patients would place the same weight on cost and consumer surplus. The estimated weight on consumer surplus in the insurer’s surplus function is 0.84, implying that insurers put more weight on cost than consumer surplus. This suggest that insurers are not fully internalizing the benefits to consumers, perhaps due to market power. However, the estimate is relatively imprecise as seen by the bootstrapped standard error.⁶⁶

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 5. Although MR scans are the most expensive, the estimated marginal cost is actually slightly less than for CT scans. The large markups for MR scans may reflect the fact that there are fewer providers with an MRI machine, implying more concentrated markets. X-rays having lower marginal cost, consistent with the fact that X-rays require less complicated equipment.

6 Effect of the Price Transparency Website

In this section I use the estimates from the previous section to examine the effect of New Hampshire’s price transparency website. First, I calculate the overall equilibrium effect of the website by

⁶⁵For comparison, Ho and Lee (2017) estimate provider bargaining weights between 0.50 and 0.88. However, Gowrisankaran et al. (2015) estimate provider bargaining weights that average 0.24. These papers examine bargaining over an index of hospital services, rather than outpatient medical imaging procedures.

⁶⁶Standard errors are calculated by resampling the set of provider-procedure-insurer-year observations used to estimate the bargaining and marginal cost parameters. In particular, 100 bootstrap samples are used. Accounting for the variance of demand model estimates when estimating the supply model is beyond the scope of this paper due to computational limitations.

Table 5
Marginal Cost Estimates

	Baseline Price		Estimated Marginal Cost	
	Mean	SD	Mean	SD
CT Scans	1,900.0	979.4	1,119.9	672.6
MRI Scans	2,167.6	999.1	994.6	1,297.9
X-Rays	525.7	568.6	442.0	525.4

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

simulating demand and supply with and without the price transparency website. In order to validate the model, the estimated spending change can be compared to reduced-form estimates. The structural model also allows for an analysis of the welfare effects as well as the effect conditional on using the website.

I start by estimating the effect of the website on average spending for all individuals in the sample.⁶⁷ Overall, counterfactual simulations imply that the website reduced overall spending by 3.9 percent. The reduction in out-of-pocket costs is especially important. Out-of-pocket costs declined by 7.6% while insurer cost declined by 3.5%. In Appendix G I compare these estimates to the reduced-form intent-to-treat effect using a difference-in-difference approach. The results are quite similar, helping to validate the model and lending credence to the counterfactual simulations in the remainder of the paper that cannot be analyzed using reduced-form methods.

Next, I examine the effect conditional on using the website, i.e. treatment-on-the-treated effect. The first panel of Table 6 summarizes the effect of using the website holding prices fixed.⁶⁸ 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. These individuals pay a relatively small portion of the total negotiated prices, therefore the savings from using the website are only \$16 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 \$18 on average).

Consumers subject to a deductible benefit most from the price transparency website. Individuals who used the website and have a deductible saved an estimated \$200 per visit, a savings of 36 percent compared 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 3 and Equation 10.⁶⁹ The gain is smaller than the cost savings—\$103 for individuals subject

⁶⁷The iterative algorithm used to simulate prices is described in Appendix F. These prices are then used to calculate the average percent change in spending due to the website.

⁶⁸Demand-side results hold prices fixed using simulated prices for the baseline case in which the website does not exist.

⁶⁹The implied cost of using the price transparency tool, given in Equation 13, is not included in the calculation of consumer surplus since it likely reflects lack of awareness about the website rather than an actual pecuniary cost.

Table 6
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	70.2	53.7	-16.5	9.2	828.5	810.3	-18.1
Under Deductible	561.5	361.9	-199.7	103.2			
(b) Demand- and Supply-Side Effects							
Over/No Deductible	70.2	51.8	-18.4	11.9	828.5	775.4	-53.1
Under Deductible	561.5	349.7	-211.8	135.0			

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.

to a deductible. This is due to the fact that, without price information, individuals place less 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 \$200 when they have price information, the providers they choose are \$97 worse on non-price characteristics.

In the second panel of Table 6, 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 \$18 while individuals subject to a deductible saved \$212. In contrast, the insurers saved \$53 per visit. Overall, the supply-side effects are modest, consistent with the fact that a relatively small fraction of consumers use the price transparency website. However, the supply-side effect also benefited consumers who did not use the website, which is reflected in Table A-10. The next section examines counterfactuals in which a larger fraction of consumers were informed about prices, thus potentially generating larger supply-side effects.

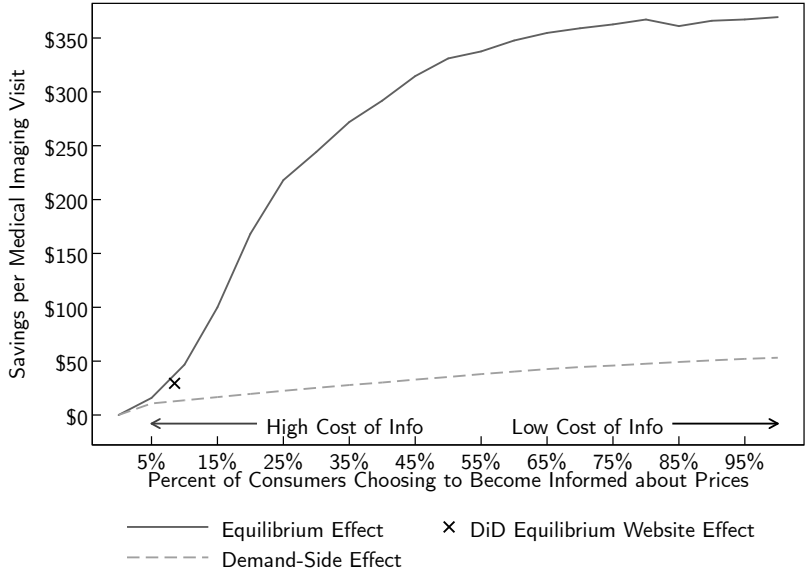
7 Out-of-Sample Counterfactuals

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

7.1 Effect of Increased Price Transparency on Overall Savings

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

Figure 4
Counterfactual Simulation Results
Effect of Price Transparency by Fraction of Individuals with 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. All figures in 2010 dollars.

fraction of informed consumers. A hypothetical policy could even require that all individuals receive a personalized price schedule before choosing where to receive outpatient services, effectively making them fully informed.

In order to analyze increased price transparency, I examine counterfactuals in which I incrementally reduce the cost of becoming informed about prices, increasing the fraction of consumers with price information. Figure 4 summarizes the main effect on total cost for the patient and insurer. 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 4. 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. This is modest given the large dispersion in prices, but is consistent with the fact that many individuals have limited incentive to price-shop even with full information given limited cost sharing.

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 4. I then use these prices to compute consumers choices and overall spending.

Table 7
Counterfactual Simulations for Negotiated Provider Prices

	Mean Price	% Δ Price	Mean Price Dispersion	% Δ Price Dispersion
No Transparency (base)	843		545	
Price Transparency Website	800	-5.1%	535	-1.7%
Full Price Transparency	656	-22.2%	376	-31.0%

Notes: 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. Price dispersion refers to the interquartile range of prices.

The equilibrium effect of increased price transparency is shown by the solid line in Figure 4. This is the primary counterfactual of interest. 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 are increasing returns. These supply-side effects imply a large externality for uninformed patients given the lower prices. Once about half of consumers are informed, savings grow only slightly.

It is important to note that the shape of the curve in Figure 4 could theoretically take almost any form depending on the parameters of the model. The primary reason for the large increase in savings when roughly 10 to 50 percent of individuals is the more vigorous price competition in this range. Not only is residual demand for each provider more elastic, but competitors are also reducing prices, creating a “race to the bottom.” Put another way, insurers are increasingly able to negotiate lower prices given that other hospitals are reducing prices. Once about half of consumers are informed, price-cost margins decline and insurers become limited in their ability to negotiate ever lower prices. Therefore, supply-side effects become less relevant once enough consumers are informed.

7.2 Effect of Increased Price Transparency on Consumers, Insurers, and Providers & Welfare Analysis

I examine the counterfactual simulations in greater detail for consumers, insurers, and providers. I focus on three counterfactuals: no price transparency for all procedures in all years, a price transparency website available for all procedures in all years, and full price transparency.

Table 7 shows the supply-side effect of price transparency, i.e. the effect on unweighted prices. When no individuals have price information, the average price of the medical imaging procedures is \$843. If a price transparency website is available for all medical imaging procedures in all years, the average price declines 5 percent to \$800.⁷⁰ 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 for all options. In this case,

⁷⁰This is broadly consistent with difference-in-differences estimates that isolate the supply side. However, note it is not directly comparable to the estimates from the difference-in-differences model or the results in Section 6 since it is the unweighted effect across procedures, insurers, and years.

Table 8
Counterfactual Simulations for Cost, Welfare, and Expenditure

	Patient					Insurer			Total
	Per Visit		Total			Per Visit		Total	
	OOP Cost	Δ OOP Cost	Δ CS	Δ Spending (millions)	Δ CS (millions)	Insurer Cost	Δ Insurer Cost	Δ Spending (millions)	
No Transparency (baseline)	84.3					756.6			
Website	75.1	-9.2	8.4	-0.3	0.2	713.1	-43.5	-1.2	-1.5
Full Transparency	44.8	-39.4	36.1	-1.1	1.0	475.9	-280.7	-8.0	-9.2

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. Consumer surplus does not include website usage cost. All figures in 2010 dollars.

prices would be 22 percent lower than the baseline case.

In the third and fourth column of Table 7, 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 differs, these results are broadly consistent with the literature stressing that price dispersion can result from search frictions.⁷¹

Panel (a) of Table 8 presents the overall effect on spending taking into account both supply and demand-side effects, i.e. the effect on transaction prices. If the website were available for all procedures in all years, consumers would save \$9 and insurers would save \$44 on average, generating \$1.2 million in total savings on X-ray, CT scans, and MR scans in New Hampshire. Full price transparency generates savings of \$39 for consumers and \$281 for insurers.

The effect for providers is shown in Panel (a) of Table A-6. 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.⁷² The overall net welfare impact for consumers, insurers, and providers is shown in Panel (a) of Table A-8. The net welfare effect of the website is quite small, but becomes economically meaningful when there is full price transparency.

7.3 Effect of Price Transparency Combined with High Cost Sharing

One potential reason that current price transparency tools are not widely used is that many consumers, especially those that pay a small coinsurance rate, have modest private gains from becoming informed and price shopping. Therefore, it is also important to understand how price transparency interacts with 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. 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

⁷¹See, for instance, Stigler (1961), Salop and Stiglitz (1977), and Burdett and Judd (1983).

⁷²It 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 9
Counterfactual Simulation Results
Negotiated Provider Prices with High Cost Sharing

	Mean Price	% Δ Price	Mean Price Dispersion	% Δ Price Dispersion
No Price Transparency	757	-10.2%	585	7.4%
Price Transparency Website	715	-15.2%	468	-14.1%
Full Price Transparency	641	-24.0%	366	-32.7%

Notes: Shows counterfactual unweighted prices across all providers/procedures for the case in which all individuals are enrolled in a high cost sharing plan with a 50% coinsurance rate and no deductible. Percent change is relative to the baseline case in which prices are computed assuming all individuals have uncertainty about prices and cost sharing is fixed at the actual level (see baseline in Table 7). Price dispersion refers to the interquartile range of prices. All prices in 2010 dollars.

not lead consumers to switch to less expensive options. For instance, Brot-Goldberg et al. (2017) do not find evidence that high deductible plans increase price shopping.⁷³

In Table 9, I consider the case in which individuals all have plans with a 50 percent coinsurance rate and no deductible. This hypothetical plan is useful for demonstrating how price transparency interacts with cost sharing. Simulations imply that website usage increases 20 percentage points under the high cost sharing scenario since potential savings from using the website are larger. This puts additional downward pressure on prices, resulting in mean prices that are 11 percent lower than with the price transparency website alone. Comparing full price transparency with baseline cost sharing and full price transparency with high cost sharing, the resulting equilibrium prices are similar. The fact that prices are not lower under full price transparency with high cost sharing is partially 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. Similar to the case in with cost sharing is held fixed, price dispersion decreases as price transparency increases in the high cost sharing case.

When the website is combined with high cost sharing, the annual savings for patients and insurers are over 50% larger than with the website alone. Full price transparency results in \$10 million in savings. Due to the high cost sharing, these savings accrue to the insurer, whereas the consumers have higher out-of-pocket cost.⁷⁴ The impact of high cost sharing on transaction prices, welfare, and total annual spending taking into account both the demand- and supply-side effects can be found in Table A-7. Depending on the incentives of insurers, premiums may adjust or, given that many individuals are in employer-sponsored plans, firms may internalize these costs.

Although much of the reduction in health care spending is due to a transfer from providers to insurers, there are still meaningful net welfare gains according to the model. When full price transparency is combined with high cost sharing, consumers tend to switch to lower marginal cost providers, resulting in a net welfare gain of \$4.3 million. Panel (b) of Table A-8 shows the overall

⁷³Although the sample of consumers examined by Brot-Goldberg et al. (2017) had access to a price transparency tool, they note that only a small fraction of consumers knew about it.

⁷⁴Depending on the level of insurer competition, these savings could potentially be passed on to consumers in the form of lower insurance premiums, however, given data limitations, I do not model insurance premiums.

welfare impact in greater detail.

Overall, these results highlight that the degree of cost sharing is an important factor affecting an individual's decision to use the price transparency website. In other words, moral hazard effects impede take-up of the price transparency tool. I find that an increase in cost sharing can incentivize more individuals to become informed, generating large positive spillovers in the form of lower prices. The results suggest that a rather high cost sharing policy is needed to incentive enough patients to become informed and shop. However, greater cost sharing also potentially exposes patients to greater risk.⁷⁵ One potential policy implication is that high cost sharing be applied only to "shoppable" procedures such as medical imaging when price transparency tools are available.

8 Discussion and Robustness

8.1 Marginal Cost Measures and Supply-Side Effects

Negotiations between hospitals and insurers are complex and the Nash-in-Nash bargaining model, while theoretically well-grounded, is ultimately a somewhat stylized modeling device. It is reassuring that the bargaining model fairly accurately recovers the supply-side effect of the website estimated using reduced-form methods.

As an additional check, I compare marginal cost estimates to the medicare reimbursement rates for the same bundle of procedures that make up each visit, which, in theory, are meant to reflect procedure cost. Medicare rates also reflect the opportunity cost for providers. While I estimate that the average marginal cost for medical imaging procedures is \$546, the average Medicare reimbursement for the same set of procedures is \$474 or \$501 depending on the provider type.⁷⁶ Medicare rates for X-rays, CT scans, and MR scans are all somewhat lower than estimated marginal cost, but are largely comparable. See Table A-9.

Given that there are some differences with Medicare rates, I also examine how counterfactual simulations would change if Medicare reimbursement rates were used in place of estimated marginal cost. Using Medicare rates, the nonlinear relationship between spending and the fraction of informed patients is quite similar to the baseline result in Figure 4 even though the model can theoretically emit a very different relationship depending on the underlying parameters. See Figure A-5. Overall, these results suggest that increasing returns to price transparency are robust to alternative marginal cost measures.

8.2 Information Frictions With Respect to Quality and Price as a Signal of Quality

In general, patients may have uncertainty about provider quality, and price may act as a signal of quality (Wolinsky 1983). The demand model includes provider fixed effects and patient char-

⁷⁵The model assumes risk-neutrality on the part of consumers, and therefore welfare calculations do not take this into account.

⁷⁶These are the non-facility and facility reimbursement rates respectively. I use the fee schedule for the bundle of procedures in the visit and then average after aggregating to the procedure-provider-insurer-year level. Medicare reimbursement rates are also inflation adjusted to 2010 dollars.

acteristics interactions, which I interpret as a measure of observed quality or amenities such as whether a provider is in a convenient location. The primary concern is that patients use the price transparency website to choose expensive options perceived to be high quality, especially if they know insurers will pay most of the cost. This would mean that perceived quality is a function of price transparency.

I address this issue in two ways. First, I find no evidence that individuals with low cost sharing choose more expensive providers when the price transparency website becomes available.⁷⁷ Second, I directly examine whether the website changes perceived quality in Appendix H. The distribution of provider fixed effects is statistically identical before and after the introduction of the website, providing additional evidence that price information does not change patients' perception of quality.

These results are consistent with the fact that medical imaging procedures are relatively standardized and individuals do not view higher prices as reflecting higher quality. Indeed, surveys find that the vast majority of New Hampshire residents say higher prices do not reflect higher provider quality.⁷⁸ Nevertheless, unobserved quality may play a more important role for complicated procedures. Analogous to information frictions related to price, information frictions related to quality may generate muted provider incentives and could have important implications for consumer surplus (e.g. Dranove 1995). With variation in information about provider quality, the model developed in this paper could be used to estimate the welfare effects of information frictions with respect to quality.

9 Conclusion

This paper examines the effects of price transparency in the market for medical imaging procedures. I develop an empirical model of competition in the market for medical procedures that separately accounts for consumer preferences and consumer uncertainty about prices. A key feature of the model is the fact that consumers choose options they believe to be optimal, but are often surprised by the bill they receive. The model provides insight into the welfare effects of these information frictions, as well as the mechanisms by which price transparency affects provider competition. In particular, I use the model to examine the broader implications of price transparency on prices, spending, and welfare.

These results highlight the importance of information frictions as a cause of high prices and price dispersion in health care. 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 most providers. In general, this also decreases price dispersion.

The results also shed light on two barriers that can explain low take-up of price transparency tools. First, individuals do not internalize the supply-side effects when choosing whether to

⁷⁷See Brown (2019).

⁷⁸In a 2016 survey, 79 percent of New Hampshire residents say higher prices are not a sign of better quality medical care. See "How New Hampshire Residents Use Health Care Price Information," April 2017. Public Agenda Research Brief.

shop for medical services, implying that price information is underutilized relative to the social optimum. Second, patients do not internalize the savings to health insurers. High cost sharing is predicted to increase take-up of price transparency tools, but with important drawbacks.

Although this paper focuses on medical imaging, consumers likely have uncertainty about prices for a wide array of medical procedures. However, price transparency tools may only be effective for procedures that are “shoppable”, estimated to comprise 30 to 40 percent of health care spending.⁷⁹ It is very difficult to shop for procedures that have prices that differ on a case-by-case basis, such as complicated surgeries. In these cases, price transparency tools are unlikely to provide much useful information unless prices are standardized and providers take on the risk of unexpected costs. In any case, information frictions likely lead to higher prices for these procedures as well.

I conclude that information frictions are important for understanding the effect of competition in the market for health care services. Similar methods as those used for this analysis may be used to examine the consequences of information frictions in other contexts, whether with respect to price, quality, or other product characteristics.

References

- Abaluck, Jason, Jonathan Gruber, and Ashley Swanson**, “Prescription drug use under Medicare Part D: A linear model of nonlinear budget sets,” Technical Report, National Bureau of Economic Research 2015.
- Akerberg, Daniel A.**, “Advertising, Learning, and Consumer Choice in Experience Good Markets: An Empirical Examination,” *International Economic Review*, 2003, 44 (3), 1007–1040.
- Allen, Jason, Robert Clark, and Jean-François Houde**, “Search Frictions and Market Power in Negotiated-Price Markets,” *Journal of Political Economy*, 2019, 127 (4), 000–000.
- 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.
- 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.
- 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,” *The Quarterly Journal of Economics*, 2017, 132 (3), 1261–1318.
- Brown, Zach Y.**, “Equilibrium Effects of Health Care Price Information,” *Forthcoming, Review of Economics and Statistics*, 2019.
- 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.
- Ching, Andrew T.**, “A dynamic oligopoly structural model for the prescription drug market after patent expiration,” *International Economic Review*, 2010, 51 (4), 1175–1207.
- 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,” *The Quarterly Journal of Economics*, 2018, 134 (1), 51–107.
- 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.

⁷⁹Examples of shoppable procedures include 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. See White and Eguchi (2014) and Health Care Cost Institute Issue Brief 11 (2016).

- Cutler, David M.**, “Technology, health costs, and the NIH,” *National Institutes of Health Roundtable on the Economics of Biomedical Research*, 1995.
- 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**, “Association Between Availability of a Price Transparency Tool and 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 and Eduardo Morales**, “What do Exporters Know?,” Working Paper 2016.
- Dickstein, Michael J.**, “Efficient Provision of Experience Goods: Evidence from Antidepressant Choice,” April 2014.
- Dranove, David**, “A problem with consumer surplus measures of the cost of practice variations,” *Journal of Health Economics*, 1995, *14* (2), 243–251.
- Ehm, Werner**, “Binomial approximation to the Poisson binomial distribution,” *Statistics and Probability Letters*, 1991, *11* (1), 7 – 16.
- Ellison, Glenn**, “A Model of Add-On Pricing,” *The Quarterly Journal of Economics*, 2005, *120* (2), 585–637.
- 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.
- 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.
- Health Care Cost Institute Issue Brief 11**, “Spending on Shoppable Services in Health Care,” March 2016.
- 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 **Ariel Pakes**, “Hospital Choices, Hospital Prices, and Financial Incentives to Physicians,” *American Economic Review*, December 2014, *104* (12), 3841–84.
- and **Robin S. Lee**, “Insurer Competition in Health Care Markets,” *Econometrica*, 2017, *85* (2), 379–417.
- , **Joseph Hogan, and Fiona Scott Morton**, “The Impact of Consumer Inattention on Insurer Pricing in the Medicare Part D Program,” 2016. Working Paper.
- 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.
- 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 MJ**, “Does It Pay to Know Prices in Health Care?,” *American Economic Journal: Economic Policy*, 2017, 9 (1), 154–179.
- Limbrock, Frank**, “Pecuniary and non-pecuniary incentives in prescription pharmaceuticals: The case of statins,” *The BE Journal of Economic Analysis & Policy*, 2011, 11 (2).
- 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.
- Newhouse, Joseph P.**, “Medical care costs: how much welfare loss?,” *The Journal of Economic Perspectives*, 1992, 6 (3), 3–21.
- Prager, Elena**, “Tiered Hospital Networks, Health Care Demand, and Prices,” 2016. Working Paper.
- 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.
- 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.
- 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.
- 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.
- Spiegler, Ran**, “Competition over agents with boundedly rational expectations,” *Theoretical Economics*, June 2006, 1 (2), 207–231.
- 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. Working Paper.
- , “Provider Responses to Online Price Transparency,” December 2015b. Working Paper.
- White, Chapin and Megan Eguchi**, “Reference Pricing: A Small Piece of the Health Care Price and Quality Puzzle,” National Institute for Health Care Reform October 2014.
- Wolinsky, Asher**, “Prices as Signals of Product Quality,” *The Review of Economic Studies*, 1983, 50 (4), 647–658.

ONLINE APPENDIX

A Data and Sample Selection

This paper analyzes the market for outpatient medical imaging services. This includes X-rays, computerized tomography (CT) scans, and magnetic resonance imaging (MR) scans, all of which are diagnostic procedures that provide internal images of the body. In [Brown \(2019\)](#), I also use this set of procedures to estimate the effect of the price transparency website using reduced-form methods. However, the empirical model requires that I make additional sample restrictions. In particular, I remove rare medical imaging procedures, defined as those with less than 500 visits per year in the state.⁸⁰ I limit the sample to individuals that live in New Hampshire and choose providers in New Hampshire. Finally, 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.⁸¹ 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.⁸² 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 \(2019\)](#).

Each visit is categorized by the imaging procedure, defined by a CPT/HCPCS code.⁸³ These codes are quite specific and refer to relatively standardized procedures. For each visit, I am able

⁸⁰When a procedure is rare, it becomes difficult to accurately construct a choice set and prices.

⁸¹Less than 2 percent of enrollees are in indemnity (fee-for-service) plans. I remove these individuals.

⁸²For 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.

⁸³The 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.

to calculate the out-of-pocket price paid by consumers, the price paid by insurers, as well as the list price.⁸⁴ 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).

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.⁸⁵ 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.

B Derivation of Website Usage Benefit

Here I derive the benefit of using the price transparency website. In order to get a closed-form expression, I approximate it using a second-order multivariate Taylor series around the expectation.⁸⁶ 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.

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

$$p_{ijkmt}^{OP} \stackrel{iid}{\sim} N(w_{ikmt}(p_{ijkmt}^{OP} + e_{ijkmt}) + (1 - w_{ikmt})\bar{p}_{kmt}^{OP}, \sigma_h^2 w_{ikmt}) \quad (\text{A-1})$$

With price uncertainty, the ex-ante consumer surplus from an individual’s perspective is determined by evaluating Equation 10 at all possible prices.⁸⁷ Therefore, given individuals’ beliefs, the expected consumer surplus is

$$\begin{aligned} \frac{1}{\gamma_i} \log \left(\sum_{j \in \mathcal{N}_{kmt}} \exp(-\gamma_i \mathbb{E} \left[\widetilde{p_{ijkmt}^{OP}} \right] + \delta_{ijkmt}) \right) \\ + \int_{\widetilde{\mathbf{p}_{ikmt}^{OP}}} \sum_{j \in \mathcal{N}_{kmt}} \left(\mathbb{E} \left[\widetilde{p_{ijkmt}^{OP}} \right] - \widetilde{p_{ijkmt}^{OP}} \right) s_{ijkmt} g(\widetilde{\mathbf{p}_{ikmt}^{OP}}) d^J \widetilde{\mathbf{p}_{ikmt}^{OP}} \quad (\text{A-2}) \end{aligned}$$

where $g(\widetilde{\mathbf{p}_{ikmt}^{OP}})$ is the joint distribution of beliefs determined by the individual’s prior and signals

⁸⁴The data also contain information on capitation payments to providers. Over the relevant period in New Hampshire, these payments were very small.

⁸⁵Note I do not include inpatient medical imaging procedures since patients are unlikely to choose their provider when they are already admitted to a hospital.

⁸⁶In the context of standard errors, a similar approach is often referred to as the delta method.

⁸⁷I assume individuals evaluate the benefit of using the website prior to knowing idiosyncratic shocks.

following Equation A-1.

In order to evaluate the expected gain from using the website, the individual must compare Equation A-2 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

$$\frac{1}{\gamma_i} \int_{\widetilde{\mathbf{p}}_{ikmt}^{OP}} \log \left(\sum_{j \in \mathcal{N}_{kmt}} \exp(-\gamma_i \widetilde{p}_{ijkmt}^{OP} + \delta_{ijkmt}) \right) g(\widetilde{\mathbf{p}}_{ikmt}^{OP}) d^J \widetilde{\mathbf{p}}_{ikmt}^{OP} \quad (\text{A-3})$$

The difference between Equation A-3 and Equation A-2 is the benefit from using the website. Since there does not exist a closed form expression, I derive an approximation using a second-order multivariate Taylor series around the expectation. This approach is necessary since it is computationally infeasible to use simulation-based methods.⁸⁸

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

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

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

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

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

$$\widetilde{p}_{ijkmt}^{OP} \stackrel{iid}{\sim} N \left(\mathbb{E} \left[\widetilde{p}_{ijkmt}^{OP} \right], \text{Var} \left[\widetilde{p}_{ijkmt}^{OP} \right] \right)$$

Since price signals are independent

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

Furthermore, note that

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

⁸⁸Numerically integrating the expression by simulating draws for each price and then averaging over the draws is computationally expensive given the high dimensionality of $\widetilde{\mathbf{p}}_{ikmt}^{OP}$. In addition, $\widetilde{\mathbf{p}}_{ikmt}^{OP}$ 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.

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

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

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

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

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

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

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

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.⁸⁹

C Details on Likelihood and Bayesian Estimation

In this section, I discuss the likelihood function and present the model as it is estimated in a Bayesian framework.

The likelihood function has two parts. The first, is the probability of provider choices, which are given by:

⁸⁹For a more detailed discussion of the validity of a Taylor series approximation for calculating welfare in a macroeconomic context, see Woodford, M., 2002. Inflation stabilization and welfare. Contributions in Macroeconomics, 2(1).

$$s_{ijkmt}(\mathcal{N}_{kmt}, \mathbf{P}_{kmt} | \vartheta_{ikmt}) = \begin{cases} s_{ijkmt}(\mathcal{N}_{kmt}, \mathbf{P}_{kmt} | \vartheta_{ikmt} = 0) & \text{if website is not available} \\ \vartheta_{ikmt} \cdot s_{ijkmt}(\mathcal{N}_{kmt}, \mathbf{P}_{kmt} | \vartheta_{ikmt} = 1) & \text{if website is available} \\ + (1 - \vartheta_{ikmt}) \cdot s_{ijkmt}(\mathcal{N}_{kmt}, \mathbf{P}_{kmt} | \vartheta_{ikmt} = 0) & \end{cases} \quad (\text{A-4})$$

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 9. 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 2 and Equation 9, where the mixture weights are determined by the predicted probability of using the website, given by Equation 14.

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 form⁹⁰

$$\frac{n_{mt}!}{\mathcal{V}_{mt}!(n_{mt} - \mathcal{V}_{mt})!} (\vartheta_{mt})^{\mathcal{V}_{mt}} (1 - \vartheta_{mt})^{n_{mt} - \mathcal{V}_{mt}} \quad (\text{A-5})$$

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

$$\begin{aligned} \mathcal{L}(\Theta) = & \underbrace{\sum_t \sum_{k \in \mathcal{K}} \sum_{m \in \mathcal{M}} \sum_{i \in \mathcal{I}_{kmt}} \sum_{j \in \mathcal{N}_{kmt}} [s_{ijkmt}(\mathcal{N}_{kmt}, \mathbf{P}_{kmt} | \vartheta_{ikmt})]^{y_{ijkmt}}}_{\text{Likelihood of Observed Provider Choices}} \\ & \times \underbrace{\sum_t \sum_{m \in \mathcal{M}} \frac{n_{mt}!}{\mathcal{V}_{mt}!(n_{mt} - \mathcal{V}_{mt})!} (\vartheta_{mt})^{\mathcal{V}_{mt}} (1 - \vartheta_{mt})^{n_{mt} - \mathcal{V}_{mt}}}_{\text{Likelihood of Observed Search Traffic}} \end{aligned} \quad (\text{A-6})$$

where y_{ijkmt} is an indicator for the observed choice. The standard approach is find the parameters $\Theta = (\sigma^\gamma, \rho, \alpha, \xi, \beta, \sigma, \theta, \phi)$ that maximize the simulated log-likelihood using numerical integration to compute choice probabilities.

Given the high dimensionality of the integral in Equation 9, it is useful to reformulate the above 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.

In general, the posterior is defined as

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

⁹⁰This 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 apply a commonly used approximation and use a binomial distribution. See Ehm (1991).

where $\check{\mathcal{L}}(D|\Theta)$ 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}

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

The simplified likelihood, conditional on unobservables, is then

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

Note that, unlike Equation A-6, this is a closed-form expression. This likelihood function is then augmented with the following:

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

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

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

The remaining parameters are given uninformative priors.

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).⁹¹ 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, $(\Theta^{(1)}, \dots, \Theta^{(N)})$, approximately converge to the distribution of the posterior. I report the mean and standard deviation of these samples in the results.

D Alternative Assumptions About Individuals' Information Set

The main analysis assumes that individuals know the distribution of prices in their choice set. This prior disciplines their beliefs, generating shrinkage of beliefs towards mean prices, which is motivated by the idea that consumers who lack information about prices are likely to ignore prices rather than choose a provider solely because they guess that it is inexpensive. Nevertheless, it is important to consider an alternative assumption in which individuals have an uninformative prior.

If individuals do not know the mean and variance of prices in their choice set, the distribution

⁹¹For parameters that are not included in the multinomial logit, I use random starting values.

of individuals' posterior beliefs is simply

$$p_{ijkmt}^{OP} \stackrel{iid}{\sim} N(p_{ijkmt}^{OP} + e_{ijkmt}, \sigma^2) \quad (\text{A-7})$$

This is equivalent to setting $w_{ikmt} = 1$ in the baseline specification. Therefore, individuals maximize expected utility given by $-\gamma_i(p_{ijkmt}^{OP} + e_{ijkmt}) + \delta_{ijkmt} + \varepsilon_{ijkmt}$.

Specification 2 in Table 3 presents the results from this alternative model. The estimates that characterize provider choices are broadly consistent with estimates from specification 1. The estimated variance of the price signal is smaller 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.

This has an implications for the welfare effects of price information. In the baseline model, individuals with very noisy price signals know that their signals provide little informational content. Therefore, they essentially ignore price when choosing a provider, instead choosing a provider based on observable characteristics such as distance. Conversely, in the alternative model, individuals who receive very noisy price signals will take the signals as given, implying that information provides larger welfare effects all else equal. To the extent that individuals have less information about the distribution of prices than what is assumed in the main results, the estimated welfare effects of information will be larger. In other words, estimated welfare effects are a lower bound.

E Details on Bargaining First Order Condition

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

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

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

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

$$\begin{aligned} \sum_{i \in I_{kmt}} \tau_h \left(s_{ijkmt} + \frac{\partial s_{ijkmt}}{\partial p_{jkmt}} [p_{jkmt} - mc_{jkmt}] \right) (s_{ijkmt} [p_{jkmt} - mc_{jkmt}])^{\tau_h - 1} (\Delta_j \Pi_{ijkmt}^K)^{1-\tau_h} \\ + \sum_{i \in I_{kmt}} (1 - \tau_h) (s_{ijkmt} [p_{jkmt} - mc_{jkmt}])^{\tau_h} (\Delta_j \Pi_{ijkmt}^K)^{-\lambda} \frac{\partial \Pi_{ijkmt}^K}{\partial p_{jkmt}} = 0 \end{aligned}$$

Now solving for the markup:

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

I now derive $\frac{\partial s_{ijkmt}}{\partial p_{jkmt}}$ and $\frac{\partial \Pi_{ikmt}^K}{\partial p_{jkmt}}$ 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}_{kmt}^{OP} and \bar{s}_{kmt}^2 , 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:

$$\frac{\partial \Pi_{ikmt}^K}{\partial p_{jkmt}} = \frac{\partial \zeta C S_{ijkmt}}{\partial p_{jkmt}} - \frac{T C_{ijkmt}}{\partial p_{jkmt}}$$

Where the partial derivative of consumer surplus is:

$$\begin{aligned} \frac{\partial \zeta C S_{ijkmt}}{\partial p_{jkmt}} &= \zeta \frac{\partial}{\partial p_{jkmt}} \left[\frac{1}{\gamma_i} \log \left[\sum_{j' \in \mathcal{N}_{kmt}} \exp(-\gamma_i w_{ikmt} (p_{ij'kmt}^{OP} + e_{ij'kmt}) + \delta_{ij'kmt}) \right] \right. \\ &\quad \left. + \sum_{j' \in \mathcal{N}_{kmt}} [w_{ikmt} e_{ij'kmt} + (1 - w_{ikmt})(\bar{p}_{kmt}^{OP} - p_{ij'kmt}^{OP})] s_{ij'kmt} \right] \\ &= -c_{ikmt} s_{ijkmt} - \gamma_i w_{ikmt} c_{ikmt} s_{ijkmt} \left[w_{ikmt} \left(1 - \sum_{j' \in \mathcal{N}_{kmt}} e_{ij'kmt} s_{ij'kmt} \right) - c_{ikmt} (1 - w_{ikmt}) \left(1 - \sum_{j' \in \mathcal{N}_{kmt}} p_{j'kmt} s_{ij'kmt} \right) \right] \end{aligned}$$

And the partial derivative of insurer cost is:

$$\begin{aligned} \frac{\partial T C_{ijkmt}}{\partial p_{jkmt}} &= \frac{\partial}{\partial p_{jkmt}} \left[\sum_{j' \in \mathcal{N}_{kmt}} p_{j'kmt} (1 - c_{ikmt}) s_{ij'kmt} \right] \\ &= (1 - c_{ikmt}) s_{ijkmt} (1 - p_{jkmt} \gamma_i w_{ikmt} c_{ikmt}) + \sum_{j' \in \mathcal{N}_{kmt}} [p_{j'kmt} (1 - c_{ikmt}) \gamma_i w_{ikmt} c_{ikmt} s_{ijkmt} s_{ij'kmt}] \end{aligned}$$

F 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, $C S_{ikmt}(\mathbf{p}_{kmt})$, and insurer cost, $T C_{ikmt}(\mathbf{p}_{kmt})$ which are also a function of price transparency.

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

$$p_{jkm t}^* = \hat{m}c_{jkm t} + \left(-\frac{1 - \hat{\tau}_h}{\hat{\tau}_h} \frac{\frac{\partial}{\partial p'_{jkm t}} \sum_{i \in I_{km}} [\hat{\zeta} C S_{ikm t}(\mathbf{p}'_{kmt}) - T C_{ikm t}(\mathbf{p}'_{kmt})]}{\sum_{i \in I_{km}} [\hat{\zeta} \Delta_j C S_{ikm t}(\mathbf{p}'_{kmt}) - \Delta_j T C_{ikm t}(\mathbf{p}'_{kmt})]} - \frac{\frac{\partial}{\partial p'_{jkm t}} [\sum_{i \in I_{km}} s_{ijkm t}(\mathbf{p}'_{kmt})]}{\sum_{i \in I_{km}} s_{ijkm t}(\mathbf{p}'_{kmt})} \right)^{-1}$$

The iterative algorithm proceeds as follows:

1. Find $p_{jkm t}^* \forall j, k, m, t$
2. Update \mathbf{p}'_{kmt} ⁹²
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 then be used to simulate spending and welfare given demand.

In general, the distribution of simulated prices and actual prices is quite similar at baseline. This can be seen in Table A-4, which shows the empirical distribution of actual prices compared with the distribution of simulated prices holding website usage fixed at actual levels. In some limited cases involving rare procedures, the model predicts either negative prices or very high prices. For the purposes of counterfactual simulation, I bound prices to be between \$0 and \$3,000 per visit.

G Comparison with Difference-in-Differences Estimates

In previous work, I analyzed the effect of New Hampshire’s price transparency website using a difference-in-differences approach, focusing on the average effect of access to the website (Brown 2019).⁹³ 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 F. These prices are then used to calculate the average percent change in spending due to the website.

Table A-10 presents the estimated effect of access to the price transparency website along with reduced-form estimates from Brown (2019). 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 3.9 percent, the latter implies that there was a 3.1 percent reduction in spending. Focusing on the reduction in cost for consumers and insurer, I also find similar estimates from the two models. The estimated effect on out-of-pocket cost is slightly larger for the empirical model, but the effect on insurer cost is slightly smaller.

Overall, the results are qualitatively consistent with the reduced-form estimates, despite different underlying assumptions. The main assumption of the difference-in-differences approach

⁹²In rare cases, it is possible for the model to predict negative prices. I bound prices to be positive.

⁹³This is often referred to as the intent-to-treat effect. The baseline difference-in-differences specification is $\text{LogPrice}_{imjkt} = \beta(\text{OnWeb}_m \times \text{Post}_t) + \alpha X_{it} + \lambda_m + \lambda_k + \lambda_t + \varepsilon_{imjkt}$ 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 .

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 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. Also see Appendix A.

H Effect of Price Transparency on Perceived Provider Quality

The baseline demand model includes provider-procedure class fixed effects, ξ_{jM} , which control for observed quality of the providers. In counterfactual simulations, it is assumed that the perceived quality of medical imaging providers remains fixed. This would be violated if, for instance, price information was used as a signal of quality (Wolinsky 1983). In particular, individuals with low cost sharing could potentially use the price transparency website to choose expensive providers perceived to be high quality.

I examine whether perceived provider quality changes after the introduction of the price transparency website. If patients use price as a signal of quality, one would expect that the price transparency website would increase the variance in provider fixed effects. To test this hypothesis, I estimate a similar model as in Section 5.1, but allow for different provider fixed effects after the introduction of the website. In particular, representative utility is assumed to take the form

$$-\gamma_1 q_{mt} p_{ijkmt}^{OP} - \gamma_2 (1 - q_{mt}) p_{ijkmt}^{OP} + \alpha_1 d_{ij} + \alpha_2 d_{ij}^2 + \alpha_3 r_{ijt} + q_{mt} \xi_{jM}^1 + (1 - q_{mt}) \xi_{jM}^2 + \beta \mathbf{x}_{ikmt} h_j$$

where q_{mt} is an indicator for whether procedure m is available on the website at time t . Therefore, ξ_{jM}^1 are provider fixed effects after the introduction of the website and ξ_{jM}^2 are provider fixed effects before the introduction of the website. As in the main specification, provider fixed effects are allowed to differ for X-rays, CT scans, and MRI scans.

Table A-6 shows the distribution of ξ_{jM}^1 and ξ_{jM}^2 . The distributions are quite similar. A two-sample Kolmogorov-Smirnov equality-of-distributions test yields a p-value of 0.87, indicating that one cannot reject the hypothesis that both are drawn from the same distribution. If anything, the variance of provider fixed effects is actually smaller for the period when the price transparency website is available.

Taken together, this provides evidence that price information does not affect perceptions of provider quality, at least for medical imaging procedures. This is consistent with the fact that these procedures are quite standardized and there is little evidence to suggest that different medical imaging providers generate meaningful differences in health outcomes.

I 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 orb/fac/neck 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/i 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 (MIRA)	
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
 Additional Summary of Privately Insured
 Individuals with Medical Imaging Claims

	Mean	SD	Min	Max
<i>Insurance Type:</i>				
PPO	0.28	0.45	0	1
POS	0.14	0.35	0	1
HMO	0.52	0.50	0	1
EPO	0.06	0.23	0	1
<i>Insurance Company:</i>				
Anthem	0.58	0.49	0	1
Cigna	0.22	0.41	0	1
Harvard Pilgrim	0.20	0.40	0	1
<i>Plan Characteristics:</i>				
Plan has Deductible	0.43	0.50	0	1
Unique Individuals		174,672		

Notes: Sample includes all privately insured individuals in the state of New Hampshire over the period 2005 to 2010 with at least one outpatient medical imaging visit. The unit of observation is a unique individual.

Table A-3
 Potential Savings from Switching Providers

	Patient				Insurer			
	Switch to Lowest Price Provider		Switch to Provider in 1 st Quartile		Switch to Lowest Price Provider		Switch to Provider in 1 st Quartile	
	Δ Price	% Δ Price	Δ Price	% Δ Price	Δ Price	% Δ Price	Δ Price	% Δ Price
<i>CT Scans</i>								
Over/No Deductible	32	63.6%	29	55.8%	1,039	46.5%	804	32.0%
Under Deductible	1,052	41.2%	904	32.7%				
<i>MR Scans</i>								
Over/No Deductible	38	71.1%	34	62.3%	924	39.2%	658	27.7%
Under Deductible	886	39.8%	743	32.8%				
<i>X-Rays</i>								
Over/No Deductible	17	50.2%	16	44.6%	355	50.8%	309	39.5%
Under Deductible	471	52.0%	420	40.8%				

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 than chosen provider, simulation assumes individuals do not switch.

Table A-4
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.

Table A-5
Estimates from Multinomial Logit Demand Model

	Estimate	SE
Patient Out-of-Pocket price		
× Website Available ($-\gamma_1$)	-0.0014***	(0.0002)
× Website Not Available ($-\gamma_2$)	-0.0006***	(0.0002)
Distance (α_1)	-0.0370***	(0.0022)
Distance squared (α_2)	0.0003***	(0.0000)
Referral Indicator (α_3)	2.5238***	(0.0236)
Hospital × Age 19-35	-0.0243	(0.0832)
Hospital × Age 36-50	0.1777**	(0.0722)
Hospital × Age 51-64	0.0946	(0.0749)
Hospital × Male	-0.0500	(0.0518)
Hospital × Income	-0.0024*	(0.0014)
Hospital × BA	-0.0029	(0.0028)
Hospital × Charlson	-0.0647*	(0.0348)
Hospital × Emergency	1.2666***	(0.0742)
Procedure group FE × provider		
		Yes
Log Likelihood	-20,742	
Pseudo-R2	0.309	
Observations	109,626	

Notes: MLE estimates from a 10% sample of visits. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A-6
Counterfactual Simulation Results
Provider Markups and Surplus

	Provider		
	Markup	Δ Markup	Δ Surplus (millions)
(a) Holding Cost Sharing Fixed			
No Transparency (base)	384.8		
Website	332.9	-51.9	-8.6
Full Transparency	69.4	-315.4	-52.4
(b) With High Cost Sharing			
No Transparency	262.1	-122.7	-20.4
Website	222.9	-161.9	-26.9
Full Transparency	86.8	-298.0	-49.5

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.

Table A-7
Counterfactual Simulation Results
Cost, Welfare, and Expenditure with High Cost Sharing

	Patient					Insurer			Total
	Per Visit			Total		Per Visit		Total	
	OOP Cost	Δ OOP Cost	Δ CS	Δ Spending (millions)	Δ CS (millions)	Insurer Cost	Δ Insurer Cost	Δ Spending (millions)	
No Transparency	357.2	273.0	-274.4	7.8	-7.9	357.2	-399.4	-11.4	-3.6
Website	335.4	251.1	-255.1	7.2	-7.3	335.4	-421.3	-12.1	-4.9
Full Transparency	245.1	160.8	-187.5	4.6	-5.4	245.1	-511.5	-14.7	-10.0

Notes: Shows counterfactual transaction prices for the case in which all individuals are enrolled in a high cost sharing plan with a 50% coinsurance rate and no deductible. Change is relative to the baseline case in which prices are computed assuming all individuals have uncertainty about prices and cost sharing is fixed at the actual level (see baseline in Table 8) 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. All figures in 2010 dollars.

Table A-8
Counterfactual Simulation Results
Net Welfare Impact for Consumers, Providers, and Insurers

	Per Visit Δ Welfare	Δ Welfare (millions)
(a) Holding Cost Sharing Fixed		
No Transparency (baseline)		
Website	0.0	0.0
Full Transparency	1.4	0.2
(b) With High Cost Sharing		
No Transparency	2.3	0.4
Website	4.3	0.7
Full Transparency	26.0	4.3

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.

Table A-9
Marginal Cost Estimates

	Estimated Marginal Cost	Medicare Rate	
		Facility	Non-Facility
CT Scans	1,119.9	778.5	753.5
MRI Scans	994.6	893.7	886.2
X-Rays	442.0	281.3	218.1

Notes: Average is over provider-procedure-insurer-year. All figures reflect the cost of the bundle of procedures that make up a medical imaging visit. All prices in 2010 dollars.


Table A-10
 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	-5.4	-3.7	-3.1
Empirical Model Estimates	-7.6	-3.4	-3.9

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 Brown (2019), 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.

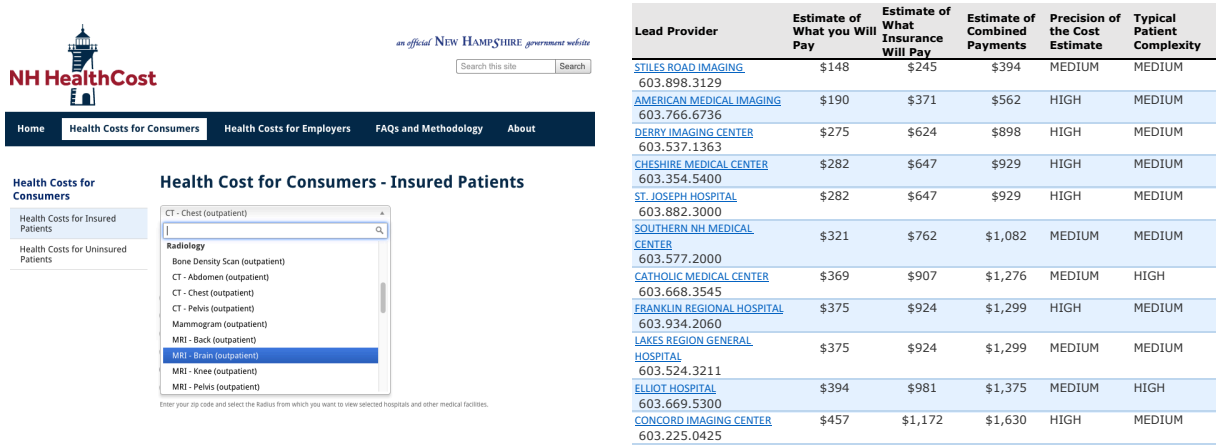
J Appendix Figures

Figure A-1
 Density of Consumers and Location of Medical Imaging Providers



Notes: Map shows the location of providers providing medical imaging services in the New Hampshire area along with the density of consumers.

Figure A-2
Prices Variation within Individuals' Choice Sets

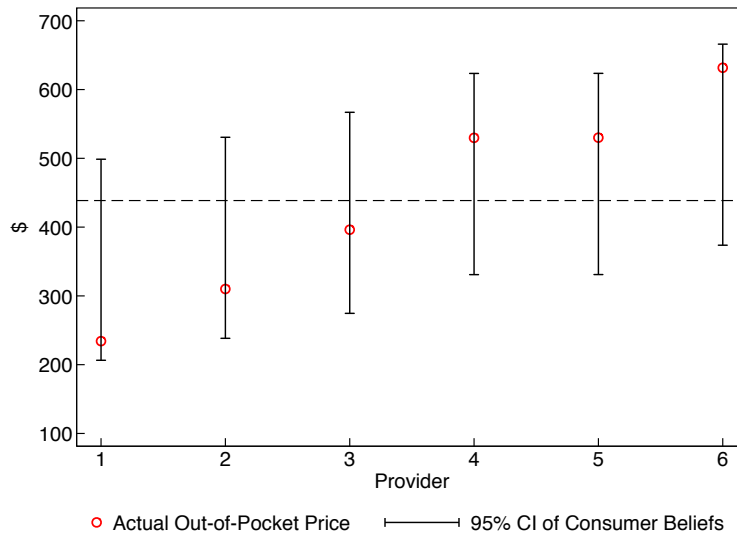


(a) Entering Procedure Information

(b) List of Prices

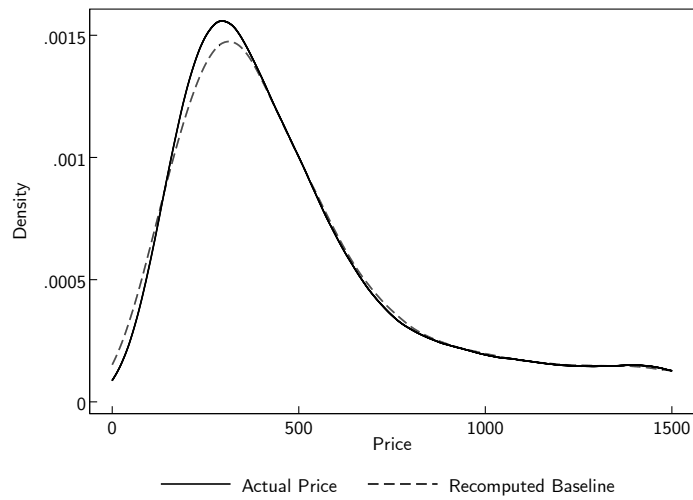
Notes: Screenshots show the New Hampshire HealthCost website as it looked in 2015. Over the period of analysis (2007 to 2010) the website had the same information.

Figure A-3
Example of Beliefs for Uninformed Individuals



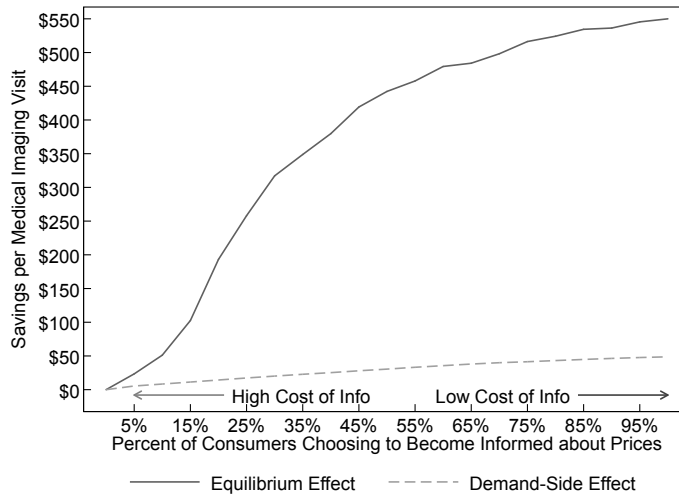
Notes: Chart shows an example choice situation selected from the data in which an individual is choosing between 6 providers. Confidence interval shows the distribution of beliefs given the estimated price uncertainty from specification 1 in Table 3. Mean beliefs are not equal to the true price since individuals' prior causes shrinkage towards the mean.

Figure A-4
Distribution of Counterfactual Negotiated Prices and Observed Prices



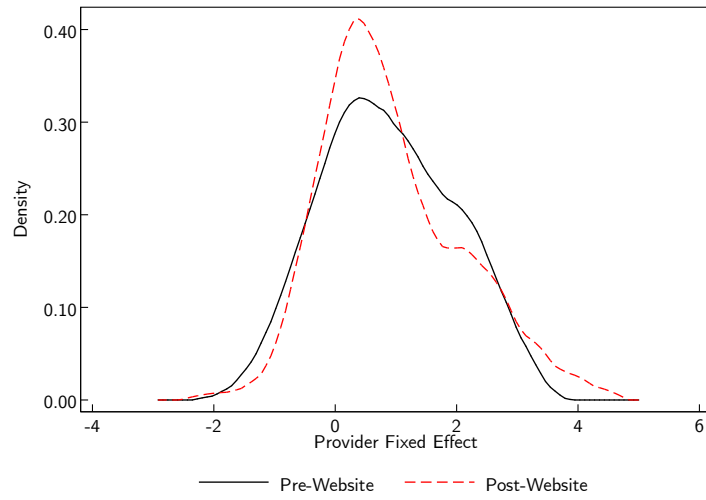
Notes: Histogram shows distribution of unweighted counterfactual prices across providers, procedures, insurers, and years.

Figure A-5
Robustness of Counterfactual Simulations to Marginal Cost Estimates
Effect of price Transparency on Medical Imaging Spending by Fraction
of Individuals with Price Information



Notes: Simulation uses Medicare reimbursement rates as a measure of marginal cost. 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.

Figure A-6
Distribution of Perceived Provider Quality Before and After
Introduction of Price Transparency Website



Notes: Shows distribution of provider-procedure class fixed effects before and after introduction of the price transparency website using multinomial logit model.