Essays on Industrial Organization and Health Care Economics

Zachary Y. Brown

Submitted in partial fulfillment of the
requirements for the degree
of Doctor of Philosophy
in the Graduate School of Arts and Sciences

COLUMBIA UNIVERSITY
2017

ABSTRACT

Essays on Industrial Organization and Health Care Economics

Zachary Y. Brown

A central tenet of industrial organization is that market concentration leads to higher prices. At the same time, there is a growing awareness that the market for health care is unique due to its complexity, and it is often difficult for consumers to make fully informed decisions. Might information frictions exacerbate market power in concentrated markets and lead to higher prices? While a growing literature seeks to address this question in a number of contexts, this dissertation focuses on the lack of price transparency in health care. I argue that the inability of consumers to compare prices is a major factor leading to high prices when health care is provided by the private market.

In Chapter 1, I use a dataset covering all private medical claims in a state to examine the introduction of a state-run website providing detailed information about out-of-pocket prices for a subset of medical procedures. Exploiting plausibly exogenous variation across procedures available on the website as well as the timing of the introduction, I use a difference-in-difference approach and find significant savings for both consumers and insurers. Part of the effect is due to consumers switching to lower cost providers. However, there is a small but significant supply-side effects in the long-run, i.e. there are lower negotiated prices. These lower prices benefit all insured individuals including those that do not use the website. Supply-side effects reduce price dispersion and are especially relevant when medical providers operate in concentrated markets.

A relatively small fraction of consumers actually used the price transparency website when it was available. Therefore, it is important to understand why more consumers aren't using the price transparency tool and what would happen if more consumers were informed about prices. Answering this question requires a structural model, which is the focus of Chapter 2 and Chapter 3.

In Chapter 2, I study demand for health care services when at least some consumers

lack full information about prices. By exploiting the variation from the introduction of the website, I am able to separately identify consumer price sensitivity and the degree of uncertainty about prices. I also explicitly model the decision to use the price transparency website when it is available. This structural approach yields two main advantages over the reduced-form approach. First, the model can be used to examine what would happen if more consumers were incentivized to use the price transparency website. Second, the model provides insight into the welfare effects of price information.

Finally, the reduced-form evidence that there is a supply-side effect of the website when even a small fraction of consumers are informed motivates a more in depth analysis of the supply-side. Chapter 3 combines the demand model of Chapter 2 with a model of bargaining between medical providers and insurers to examine how price transparency affects equilibrium prices. Model estimates and difference-in-differences estimates both imply that the website reduces health care spending by 3 to 4 percent. I then use the model to examine the effects of price transparency more generally. In counterfactual simulations, I find that price transparency would generate a substantial reduction in equilibrium prices if a larger fraction of consumers in the market were informed. Combining the price transparency website with high cost sharing would give individuals more incentive to use the price transparency tool, reducing health care spending by 18 percent.

My research is intended to inform the policy debate surrounding the value of health care price transparency tools. In sum, I argue that while the value of price transparency tools is modest when only a small fraction of consumers are incentivized to use the tools, the savings become quite substantial when enough consumers are informed about prices.

Contents

List of Tables

List of Figures

Acknowledgements

There are large number of people who contributed to this dissertation and I owe a huge debt of gratitude to all of them. First and foremost, I am very grateful to my advisor Kate Ho for motivating me to work on the industrial organization of healthcare, as well as providing invaluable guidance and advice. I benefited greatly from being her student, research assistant, teaching assistant, and advisee. I also benefited greatly from in depth discussions and feedback from Mike Riordan, Chris Conlon, Tobias Salz, Adam Kapor, Doug Almond, Bernard Salanié, Bentley Macleod, Fanyin Zheng, Matt Backus, and Tal Gross. I am positive that this dissertation would not exist without these individuals. I also greatly benefited from the broader Columbia community. In particular, I relish my discussions with Ilton Sores, Paul Piveteau, Evan Riehl, Ildikó Magyari, Joe Hogan, Jessica Van Parys, and Nicolás de Roux, among others. I especially appreciate Chris Hansman being there when I wanted to bounce ideas around, and for shooting down bad ones. I also am thankful for the support of Shane Bordeau and Amy Devine. Outside of Columbia, I thank Mary Fields, Tyler Brannen, Maureen Mustard, and the New Hampshire Department of Health and Human Services for providing data and insight into the New Hampshire HealthCost website. I also benefited from over a decade of economics discussions with Melanie Wasserman. This dissertation was also supported by the National Science Foundation.

I am also thankful for the experience I had at the Council of Economic Advisers in 2014.

Conversations with Jordan Matsudaira, Matt Fiedler, Betsey Stevenson, Gabe Scheffler, Kevin Rinz, John Coglianese, and Martha Gimbel motivated my research and furthered my understanding of economic policy.

I am grateful to numerous seminar participants that provided feedback on parts of this dissertation. At Columbia, I benefited from participants in the Industrial Organization Seminar, Industrial Organization Colloquium, Applied Microeconomics Colloquium, and Econometrics Colloquium. I also greatly benefited from seminar participants at University of Washington, University of Michigan, University of Wisconsin, Pomona College, Wharton, University of Massachusetts Amherst, INSEAD, University of Illinois at Urbana-Champaign, and NBER.

I want to thank my family for constant support regardless of the circumstances. I could not have done this without your advice, guidance, and patience. And finally, there is one person that made me eternally grateful that I was able to come to Columbia. Mica, I am continually in awe of your kindness and intellectual creativity. I value our conversations about economics almost as much as I value our conversations on other topics. Given that we work as a team, perhaps all errors and omissions should not be our own.

Chapter 1

Effects of Health Care Price Transparency: A Reduced-Form Approach

1.1 Introduction

While the price of health care procedures varies widely across medical providers, these prices are often difficult for patients to observe. Consequently, individuals often choose providers without full information about prices.[1] A large theoretical literature, beginning with Stigler (1961) and Diamond (1971), argues that information frictions can impede competition and lead to higher prices. In recent years, technological innovations have made it easier for consumers to compare prices in a number of markets, with important implications for the equilibria in these markets (Clay, Krishnan and Wolff 2001; Morton, Zettelmeyer and Silva-Risso 2001; Brown and Goolsbee 2002; Goldmanis et al. 2010).

In this paper I ask how information about health care prices affects the market for health care services in equilibrium, focusing on both the demand- and supply-side response. While we expect consumers to benefit by choosing low-cost providers, price transparency may also allow insurers to negotiate lower prices with health care providers. This is because increased demand elasticity with respect to prices may generate lower demand for high cost providers, giving them an incentive to lower price. These supply-side effects could benefit all consumers, including those that do not use the information. Effects on negotiated health care prices are especially important given that high health care expenditure in the United States is often attributed to high prices.[2]

I exploit the introduction of a publicly-provided website that allows individuals to access information about their out-of-pocket price for certain medical procedures. While previously studied price transparency efforts have primarily been conducted by specific employers, the website provided information that could be used by all privately-insured individuals in the state. This potentially generated significant supply-side effects in addition to demand-side effects. I employ a difference-in-difference methodology that takes

[1]Surveys show that consumers rarely compare prices for medical services. See "How Much Will it Cost? How Americans Use Prices in Health Care," Public Agenda (2015).

[2]For instance, see Anderson et al. (2003), Koechlin, Lorenzoni and Schreyer (2010), and Cooper et al. (2015).

advantage of two sources of variation, namely the timing of the website introduction and variation among procedures available on the website. I focus on the universe of outpatient radiology visits, which account for over 14 million claims, since information about prices was only available for a subset of radiology procedures. Thus, radiology procedures not on the website can be used as a control group. I argue that whether a radiology procedure was on the website is unlikely to be correlated with demand and cost factors that change over time. This allows me to provide the first evidence on the long-run equilibrium effect of information about out-of-pocket prices. By observing detailed information on copay, coinsurance, deductible, and insurer payments, I also provide new evidence about how insurers and patients split the savings that result from price transparency.

First, I examine transaction prices, which include both demand- and supply-side effects. Over the five year period after the website started, there is a 6.9 percent reduction in out-of-pocket prices for radiology visits on the website relative to radiology procedures not on the website. This effect increases over time and by the fifth year, out-of-pocket prices are 14.6 percent lower relative to the control group. Individuals with the most to gain from using the website—those under their deductible—see more than double the savings over the period. These results are highly significant and robust to the inclusion of procedure category specific time trends, individual fixed effects, and detailed insurer and individual controls.

Next, I examine the mechanisms driving the reduction in transaction prices. On the demand-side, individuals with access to the website are more likely to choose a different provider than their last radiology visit relative to individuals who cannot use the website. In particular, individuals were more likely to choose a provider in the lowest decile of the price distribution in their county relative to when the website is not available. These lower-cost options tend to be non-hospital providers such as imaging centers or clinics.

Although demand-side effects are present throughout the period, I find evidence that supply-side effects are important in the long-run. I analyze the supply-side effects using a difference-in-difference specification that controls for demand-side effects. Specifically, I include provider-procedure-insurer fixed effects that control for transaction price differences due to switching across providers. Since providers and insurers only negotiate a new price schedule at most once a year, the supply-side effect may take time to materialize. The estimates imply that providers reduce their prices in the long-run, defined as more than two years after the introduction of the website. Visit prices decline by 2.2 percent and principal procedure prices decline by 4.7 percent.[3] The effect is greater for providers operating in concentrated markets that are likely to have the highest margins in the absence of the website.

Access to price information may also affect whether individuals decide to receive care, affecting quantity of visits. If individuals know they can choose a low-cost provider, they may be more likely to get a radiology procedure. Conversely, if prices are higher than they expected, price information may decrease quantity. I examine these utilization effects using the same sources of variation. For the majority of procedures, the results imply no effect on the quantity of care. These results are quite precise, and robust to various specifications. The exception is mammograms, which generally are elective screening procedures at the discretion of patients. The quantity of mammograms increases after the introduction of the website relative to radiology procedures not on the website, however it is important to note that I can not rule out confounding factors such as changes in physician recommendations over time. The website may have also affected insurance choice, although there is little evidence that changes in the composition of insurance enrollment explain the effect of the website. In addition, an event study implies little overall effect on insurance choice.

A theoretical literature has found that price dispersion can result from information frictions in the market. Even markets with homogenous products can exhibit price disper-

[3] As explaining in Section 1.3, visit price includes all supplemental line items that are part of a radiology visit while principal procedure price refers to the cost of the individual radiology procedure.

sion in equilibrium. This is true with heterogenous consumers (Salop and Stiglitz 1977) as well as homogenous consumers (Burdett and Judd 1983). In order to test the theoretical prediction that price dispersion in the market for health care services is due in part to information frictions, I use a similar difference-in-difference methodology to directly examine a measure of price dispersion. The estimates imply that access to the website reduces price dispersion, as measured by the interquartile range of negotiated principal procedure prices, by $47 on average relative to the price dispersion of the control procedures.

Prior research has examined the demand-side response to health care price transparency efforts by individual employers and has found relatively small effects in the short run (Lieber 2015; Whaley 2015a; Desai et al. 2016). In particular, Desai et al. (2016) finds no reduction in spending in the year after an employer offers a price transparency tool. Consistent with the this literature, I find modest effects in the initial years. However, I examine the long-run effects and find larger effects on price. Using website traffic data, I show evidence it may take time for individuals to learn about the website. In addition, as I argue below, the previous literature has focused on price transparency efforts that are available to a small subset of consumers. In contrast, the New Hampshire website was publicly available to all consumers in the state, therefore supply-side effects may be important in the long-run.

The evidence on the supply-side effects of market-wide price transparency is quite limited. While the previous literature has focused on price transparency tools for individual employers, I examine the introduction of a first-of-its-kind website that was available to all insured individuals in a state.[4] Sinaiko and Rosenthal (2011) discuss the potential implications of New Hampshire's price transparency website and note that supply-side

[4]Whaley (2015b) focuses on a website providing information to specific employers and finds a reduction in the price of laboratory tests using an event-study methodology. In addition, Christensen, Floyd and Maffett (2015) examine the effect of information about list prices (rather than out-of-pocket prices) and find little evidence of effects on negotiated prices.

effects may be important, but the effects may take time to materialize.[5] Using individual-level data on outpatient radiology visits by all insured individuals in New Hampshire, this paper provides the first evidence quantifying the overall equilibrium effects of price transparency for both individuals and insurers in a state. Understanding the equilibrium effects are particularly relevant given that many states are currently considering price transparency legislation.[6]

This paper is also related to work examining how prices are determined through insurer-provider bargaining. Previous work has focused on network formation, mergers, and insurer competition (Ho 2009; Gowrisankaran, Nevo and Town 2015; Ho and Lee 2017). Gowrisankaran, Nevo and Town (2015) develop a model in which hospital demand is a function of out-of-pocket prices, which then determines negotiated prices in equilibrium. I contribute to this literature by providing evidence that information about prices can also affect health care prices in equilibrium.

The remainder of this paper is as follows. Section 1.2 provides additional background on the website and health care pricing and Section 1.3 describes the data. Section 1.4 describes the main empirical strategy and discusses the demand-side and supply-side effects on prices. Section 1.5 examines the quantity margin while Section 1.6 examines the insurance margin. Section 1.7 concludes.

1.2 Institutional Details and Background

About two-thirds of working-age individuals in the United States have privately provided health insurance, either through their employer or directly purchased.[7] Recent research has documented a large degree of price dispersion in health care, especially in

[5]The authors hypothesize that price transparency could either lower or raise prices, but note that "it is too early to tell what the outcome of experiments with increased transparency will be."

[6]At least 27 states proposed price transparency laws in 2015, although many of these proposals would not mandate disclosure of out-of-pocket prices. Price transparency legislation has also been proposed at the Federal level. See Nicholson (2015).

[7]See Smith, Medalia et al. (2014)

the private sector (Philipson et al. 2010; Newhouse et al. 2013).[8] Differences in observable and unobservable characteristics can only explain a modest portion of the variation in prices (Finkelstein, Gentzkow and Williams 2014). Even relatively homogenous medical services vary in price. For instance, Cooper et al. (2015) find that MRI prices vary by a factor of 12 across the country.

There is a large degree of price dispersion even within a geographically constrained area. For example, the total price of a back MRI in New Hampshire for individuals covered by Anthem, the largest insurer in the state, varies widely, with an upper and lower quartile of $1,085 to $2,472 respectively. Consequently, the out-of-pocket price paid by individuals ranges as well, especially for those under their deductible (the out-of-pocket interquartile range is $143). More generally, Table 1.1 shows the potential savings if all consumers switched to a low cost provider, defined as a provider in the first quartile of the price distribution in the state.[9] Across a range of procedure categories, savings would be between 44 and 73 percent. Even if individuals switched to the provider with the median price they would save 16 to 58 percent on average.

One explanation for why these price differences persist even for relatively homogenous products is that patients lack information about health care prices. Health care prices are determined through bargaining between insurers and providers, and insurers often agree not to publicly disclose the negotiated contracts. This is consistent with research on other markets in which sellers have an incentive for price obfuscation (Ellison and Ellison 2005, 2009). Perhaps for this reason, surveys show that the majority of individuals do not compare prices before receiving medical care.[10]

In order to allow health care consumers to find low cost options, the state of New

[8]Note that a large literature also focuses on variation in Medicare spending (e.g. Fisher et al. 2003; Fisher, Bynum and Skinner 2009)

[9]I calculate the first quartile of the price distribution conditional on individuals' insurance and procedure. I consider the case in which all individuals paying about this price switch to the provider charging the first quartile price.

[10]According to a nationally representative survey, 79 percent of individuals stated that they could not compare prices (or did not even try) before receiving medical care (Public Agenda 2015).

Table 1.1: The Potential Cost Savings if Consumers Switched to Low Price Providers

Procedure Class	Mean Total Visit Price	Consumers Switch to 1st Quartile Provider Mean	% Savings	Consumers Switch to Median Provider Mean	% Savings
Computed Tomography (CT)	1,367	562	58.9%	848	38.0%
Mammogram	260	146	44.0%	183	29.8%
Magnetic Resonance Imaging (MRI)	1,504	843	44.0%	1,094	27.3%
Nuclear Imaging	1,525	815	46.6%	1,108	27.4%
Positron Emission Tomography (PET)	3,304	1,976	40.2%	2,765	16.3%
Ultrasound	464	205	56.0%	275	40.7%
X-Ray	497	135	72.9%	209	58.0%

Notes: The above figures show the average price in 2006 if every consumer paid at most the 25th/50th percentile of visit price in New Hampshire for each procedure (as defined by CPT/HCPCS) given the consumer's insurance company and insurance type. All prices in 2010 dollars.

Hampshire began requiring health insurers operating in the state to submit medical claims to a centralized database in 2005. These data were then used to calculate the median bundled out-of-pocket prices for various medical procedures. In March 2007, New Hampshire launched their HealthCost website.[11] Individuals enter the procedure, their insurance information (including remaining deductible), their zip code, and search radius and obtain information about each provider's expected out-of-pocket price, insurer price, and total price. Results are sorted by out-of-pocket price making it easy to select the least expensive provider from the point of view of the patient. In addition to information for insured individuals, the website also has a separate feature providing information for uninsured individuals.[12] More recently, the website has also added information about provider quality and a guide to health insurance.[13] Although other states have since started price transparency websites of their own, including California, Maryland, Florida, Oregon, and New Jersey, New Hampshire's price transparency efforts are the most comprehensive.[14]

[11] Originally nhhealthcost.org, however the website can now be found at nhhealthcost.nh.gov.
[12] I do not observe uninsured individuals, and therefore I do not examine the effect of this information.
[13] This occurred in early 2016, after my period of analysis.
[14] New Hampshire was the only state to receive an "A" grade from Catalyst for Payment Reform's 2015 Report Card on State Price Transparency Laws.

Figure 1.1: Website Searches for Health Care Prices
By Month

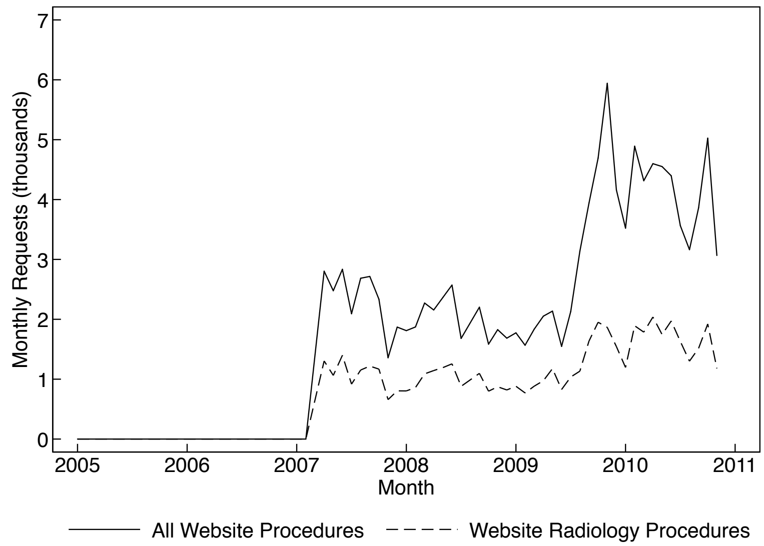

Notes: Includes all searches using "Health Costs for Insured Patients" wizard on either nhhealthcost.nh.gov or nhhealthcost.org. Website traffic data is not available for the period after 2010 due to a change in the website host.

At the time it was introduced, the website had price information for about 35 procedures. In interviews, the administrators of the website noted that a broad subset of common procedures were chosen with the intention of adding addition procedures over time. The website focuses on outpatient procedures since patients often schedule these appointments ahead of time and may have more scope for choosing among providers.

The HealthCost website has received significant attention in the state, with over 40 articles in the local public press. In addition, the New Hampshire Insurance Department promoted the website by encouraging primary care doctors to tell patient about the website. Insurers were also encouraged to inform their enrollees of the website.

A sizable population made use of the website. I construct a measure of website usage using monthly website traffic logs provided by the New Hampshire Insurance Department. Figure 1.1 shows the number of price searches on the website since 2005. When the website began, there were roughly 1,000 searches per month for the price of radiology procedures, which grew over time. Searches for the price of radiology procedures make up about half of all searches using the website.

In addition to immediate demand-side effects from the website, there is anecdotal evidence of supply-side effects. Analysts have noted that "the balance of plan-provider negotiating power began shifting significantly in New Hampshire, a result in large part of public transparency efforts."[15] For instance, Exeter Hospital and Anthem, the largest insurer in New Hampshire, had a public dispute over contract terms in 2010. Anthem argued that prices at Exeter Hospital were too high, pointing to the website as evidence, and was eventually able to negotiate rate cuts.[16]

Why might provider prices respond to price information? First, consider the case in which providers have market power and are able to unilaterally set price. If consumers become more price sensitive due to better information about prices, the profit maximizing price will decline. In the market for private health care, prices are usually determined through bilateral negotiations between providers and insurers rather than set unilaterally. In this case, the same mechanism also applies. However, equilibrium negotiated prices may also depend on insurer incentives (Ho 2009; Gowrisankaran, Nevo and Town 2015; Ho and Lee 2017). To the extent that the website affects either provider or insurer gains from trade, negotiated prices may be affected in equilibrium.[17]

1.3 Data

The main dataset covers the universe of private insurance enrollment and medical claims in the state of New Hampshire from 2005 to 2011. These data were collected as part of the New Hampshire Comprehensive Health Care Information System, which assembled data from all commercial insurers with enrollees that were state residents or who receive services under a policy issued in the state. These are the same data used to construct prices for the website.

[15]See Tu and Gourevitch (2014).

[16]For more information, see "Exeter Hospital Says Costs Being Used as Negotiating Tactic," Seacoastonline.com, Nov. 14, 2010.

[17]I explore these mechanisms in more detail in a related paper.

Table 1.2: The Availability of Outpatient Radiology Procedure Price Information on Website

Procedure Category	Num. Unique Procedures	
	On Website	Not on Website
Computed Tomography (CT)	15	47
Mammogram	12	0
Magnetic Resonance Imaging (MRI)	16	72
Nuclear Imaging	13	121
Positron Emission Tomography (PET)	0	6
Ultrasound	28	46
X-Ray	41	111
Total	125	403

Notes: Unique outpatient procedures are indentified using CPT/HCPCS codes. Procedure codes with updated descriptions are considered separate procedures.

Each outpatient claim has a CPT/HCPCS code which can be used to identify procedures.[18] These codes are very specific (e.g. code 72120 is "x-ray examination, spine, lumbosacral; bending views only, 2 or 3 views"). I limit the sample to the universe of outpatient radiology claims, which includes 528 procedures related to X-rays, computerized tomography (CT) scans, magnetic resonance imaging (MRI) scans, mammograms, ultrasound, positron emission tomography (PET) scans, and other nuclear imaging procedures. These procedures all use imaging to diagnose, and sometimes treat, internal conditions. The number of procedures in each category are listed in Table 1.2.

Note that inpatient radiology procedures, such as those that are part of major surgeries, are excluded from the analysis. Since individuals have little ability to choose a provider when radiology procedures are part of an inpatient episode, the website only includes information about outpatient radiology procedures.[19]

In addition to the principal radiology procedure, there are often supplemental procedures such as contrast agents that are billed along with the main procedure. The quantity

[18] Current Procedural Terminology (CPT) codes are a set of codes developed and maintained by the American Medical Association. Healthcare Common Procedure Coding System (HCPCS) codes are an extension of CPT codes that include additional procedures such as non-physician services.

[19] The website does include prices for a few inpatient procedures that are not related to radiology (e.g. newborn delivery).

and price of these supplemental procedures may also vary across providers. When comparing the cost across medical providers, the relevant price is determined by the entire bundle of procedures. For this reason, the website has information about the cost of an entire visit. For the same reason, my analysis focuses on the price of the visit. However, if a provider wishes to change the price of an individual procedure they will likely change the individual procedure price rather than the price of supplemental procedures since the supplemental procedures are often the same across procedures. I call this price the principal procedure price. The construction of the visit price and principal procedure price are described in more detail in Appendix Section 3.6.

There are 1.1 million individuals under age 65 with at least one radiology claim between 2005 and 2011. Using individuals' zip code, I merge on additional demographic information, including income and education, using the 2007-2011 American Community Survey. I also construct each individual's Charlson Comorbidity Index using International Classification of Diseases (ICD) codes in the claims dataset (Charlson et al. 1987; Stagg 2006). The Charlson Comorbidity Index is an integer score that summarizes comorbid conditions that predict mortality. In addition, I construct the number of medical claims each individual had in the previous year, which I use as a proxy for individuals' experience with the health care system. Individual demographics are summarized in the first panel of Table 1.3.

Each medical claim is also associated with a anonymized provider identifier that can be linked to additional information such as provider zip code and whether the provider is a hospital or non-hospital facility. This information is used to construct provider concentration in each county.

One limitation of the data is that there is no information about patient referrals. Primary care physicians may refer patients to a specific radiology provider, however individuals can still choose a different provider as long as it is in the insurer network. The website may have affected individuals' choice of provider or it may have affected physi-

Table 1.3: Summary of Privately Insured Individuals with Radiology Claims

	Mean	SD	Min	Max
Male	0.39	0.49	0	1
Age	38.4	16.9	0.0	64.0
Charlson Comorbidity Index	0.4	0.7	0	2
Zip income (1000s)	68.7	21.3	4.9	242.2
Zip more than BA Degree	34.0	13.9	0.0	100.0
Claims in previous year	35.0	50.2	0.0	5118.0
Insurance Type:				
PPO	0.32	0.47	0	1
POS	0.13	0.34	0	1
HMO	0.38	0.49	0	1
EPO	0.07	0.25	0	1
Indemnity	0.03	0.17	0	1
Insurance Company:				
Anthem	0.46	0.50	0	1
Cigna	0.24	0.42	0	1
Harvard Pilgrim	0.12	0.33	0	1
Other	0.12	0.32	0	1
Plan Characteristics:				
Plan has Deductible	0.43	0.49	0	1
Plan has Copay	0.80	0.40	0	1
Plan has Coinsurance	0.22	0.41	0	1
Individuals	1,078,892			

Notes: Includes all unique privately insured individuals in the state of New Hampshire over the period 2005 to 2011 with at least one outpatient radiology visit.

cians' referrals. Since I do not observe referrals, I do not differentiate between these mechanisms.

The vast majority of individuals in the sample are covered by a managed care organization, either a Health Maintenance Organization (HMO) plan, Preferred Provider Organization (PPO) plan, Point-of-Service (POS) plan, or an Exclusive Provider Organization (EPO) plan. The defining feature of managed care plans is that insurers negotiate lower prices with a selected network of providers. The plan types differ according to the standards used when individuals select providers within the network. Only 3 percent of individuals have an indemnity (fee-for-service) plan. Plan type is summarized in the second panel of Table 1.3. There are three main insurers operating in New Hampshire: Anthem, Cigna, and Harvard Pilgrim. Less than a quarter of individuals are enrolled in another plan (see third panel of Table 1.3).[20]

The plans offered in New Hampshire over the period differ in their cost-sharing characteristics. In particular, 43 percent of individuals pay a deductible at some point over the period (see last panel of Table 1.3). In general, individuals are responsible for all health care costs under the deductible amount in a given year. Although I do not observe the deductible amount associated with each plan, I do observe the deductible paid on each visit. Using observed deductible payments, I construct an indicator for whether each individual is under or over her deductible in a given year in order to test whether individuals benefit more from the website when they are subject to a deductible.[21]

Over the period, there are 14.6 million claims that constitute 3.6 million radiology visits (i.e. there are about 3 supplemental procedures on average per radiology visit). For each health claim, I observe the copayment, coinsurance, and deductible paid by the individual, which together makes up the out-of-pocket price. In addition, I observe the insurer

[20]For more detail on the construction of demographic covariates, see Appendix Section 3.6.

[21]Individuals who know they will fulfill their deductible over the course of the year should not be price sensitive. However, to the extent that individuals have uncertainty about their future health care use or their remaining deductible, individuals will be price sensitive even if they are close to hitting their deductible. For this reason, I consider all individuals who have not passed their deductible.

Table 1.4: Summary of Outpatient Radiology Visit Price

	Visits on Website				Visits not on Website			
	Pre-Website		Post-Website		Pre-Website		Post-Website	
	Mean	SD	Mean	SD	Mean	SD	Mean	SD
Patient Cost:								
Copayment	10.5	34.6	12.9	32.5	12.5	45.7	14.6	34.7
Coinsurance	9.2	67.1	15.3	93.3	13.2	81.7	23.1	119.8
Deductible	35.7	155.8	62.0	248.4	49.1	189.0	85.9	293.9
Total out-of-pocket cost	57.2	192.2	90.5	278.1	77.3	232.4	123.9	330.4
Insurance Cost:								
Paid amount	524.5	1,132.5	643.8	1,406.4	691.8	1,532.2	901.1	1,871.1
Total:								
Allowed amount	581.7	1,184.5	734.3	1,469.6	769.1	1,583.1	1,025.0	1,931.7
Charge amount	992.7	2,302.8	1,252.2	2,724.8	1,339.5	2,897.5	1,742.6	3,366.6
Observations (Visits)	865,399		2,019,627		191,576		515,490	
Total Claims	3,235,485		8,626,074		673,696		2,076,856	

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

paid amount. Together, the out-of-pocket price and insurer paid amount constitute the total price received by the provider, often called the allowed amount.[22] This price is often lower than the list price (i.e. charge amount, which is also reported), since insurers negotiate lower rates for their enrollees.

The average price paid by individuals and insurers is presented in Table 1.4. Insurers pay the majority of the cost for radiology procedures. Although out-of-pocket prices are low on average, there is high variance and some individuals are fully exposed to the total price.

The summary statistics presented in Table 1.4 preview the results. Although the price of all procedures is increasing over time, the simple difference-in-difference estimate using the average total price implies that the price of procedures on the website declined by $104 relative to the price of procedures not on the website. However, it is important to control for changes in the composition of procedures and changes across time affecting

[22]Capitation payments for radiology procedures were negligible during the period.

the control group.

1.4 Effect on Prices

I begin by examining the overall effect on transaction prices, both out-of-pocket prices and insurer price. I examine the heterogenous effects and show that results are robust to a number of specifications. Using a similar identification strategy, I show that this effect is due to both demand-side and supply-side factors.

1.4.1 Empirical Strategy

In order to estimate the causal effect of price transparency on prices, I exploit two sources of plausible exogenous variation: the timing of the website introduction and the availability of radiology procedures on the website. In particular, I construct OnWeb_m, which indicates whether procedure m is ever available on the website. I also construct Post_t, which indicates if the website is available at month t. This takes the value of 1 if the date of admission is March 2007 or later. The baseline difference-in-difference specification is given by:

$$\log(1 + p_{imjkt}) = \beta(\text{OnWeb}_m \times \text{Post}_t) + \alpha X_{it} + \lambda_m + \lambda_k + \lambda_t + \varepsilon_{imjkt} \tag{1.1}$$

The outcome of interest is p_{imjkt}, the price of a visit for individual i with insurance k obtaining procedure m from provider j at time t. I consider both the patient's out-of-pocket cost as well as the cost to the insurer. The baseline specification controls for individual covariates X_{it}, which includes age, gender, charlson comorbidity index, income, education, rural classification, and member plan characteristics (deductible, coinsurance, and copay). I also include procedure fixed effects, λ_m, and insurer fixed effects, λ_k, that control for time-invariant factors that may be correlated with prices and the availability of the website.[23] I also include month fixed effects, λ_t, which control for time varying factors that may be correlated with prices and website availability. Finally, ε_{imjkt} is a vector of idiosyncratic random errors. Prices are highly correlated within each month since individuals tend to be subject to a deductible in the beginning of the year, but not at the end of the year. To account for correlation within a month, standard errors are clustered at the month level. The unit of analysis is an individual radiology visit.

The coefficient of interest, β, is interpreted as the change in prices due to the website in log-points. The main identifying assumption is that, in the absence of the website, the procedures on the website and the procedures not on the website would follow common trends. I use a number of methods to examine the validity of this assumption, including a falsification test and directly controlling for procedure category linear trends.

It is important to note that this assumption assumes that website only provides useful information for procedure that are actually featured on the website. In other words, individuals do not use website prices to learn about the prices of procedures not on the website. I argue this would be difficult given that the rank ordering of providers is different even for related procedures. To the extent that individuals do use the website even when they are obtaining procedures not listed, the results will be biased towards zero.

In order to isolate the supply-side effect of price transparency I use a similar identification strategy, but control for the demand-side effects. In particular, I include fixed effects that control for the variation in price of each procedure across providers and insurers. This approach is similar to that of Christensen, Floyd and Maffett (2015). The specification is now

$$\log(1 + p_{imjkt}) = \beta(\text{OnWeb}_m \times \text{Post}_t) + \alpha X_{it} + \lambda_{jmk} + \lambda_t + \varepsilon_{imjkt} \tag{1.2}$$

where the vector λ_{jmk} includes an indicator for each combination of provider, proce-

[23]For the insurer fixed effects, I define an insurance plan as a unique combination of insurance firm and insurance type (e.g. Anthem HMO).

dure, and insurer. If individuals switch to lower-cost providers after the introduction of the website, it is captured by these fixed effects. Note that these fixed effects also absorb changes in prices due to the entry of providers over the period. The remaining variation in price (with insurer-provider-procedure indicators) identifies β, which can now be interpreted as the reduction in transaction prices due to lower negotiated prices.[24]

I examine how the supply-side effect varies by the degree of local competition between providers. In particular, I use the Herfindahl index in each county for each procedure category, which is defined as $\text{HHI}_{cl} = \sum_j s_{jcl}^2$ where s_{jcl}^2 is the market share of provider j in county c among all procedures in procedure category l. The period prior to the introduction of the website is used to calculate HHI_{cl} in order to address concerns that the market structure may have been endogenously affected by the website.

Finally, I examine price dispersion directly as measured by the interquartile rage of prices. I exploit the same sources of variation and estimate

$$IQR_{mt} = \beta(\text{OnWeb}_m \times \text{Post}_t) + \alpha X_t + \lambda_m + \lambda_t + \varepsilon_{mt} \tag{1.3}$$

For this specification, individual visits are aggregated to the month level. In particular, to examine the dispersion in transaction prices, IQR_{mt} is defined as the difference between the third and first quartile of transaction prices for each procedure in each month. To examine the dispersion in provider prices (or negotiated prices), IQR_{mt} is defined as the difference between the third and first quartile of prices for each procedure in each month after aggregating to the provider level. Under the same assumptions as previous specifications, β can be interpreted as the dollar change in the interquartile range of prices due to the website.

[24]It is also possible that providers reduce the list price. This may affect some privately-insured individuals, particularly those going out of network. I do not distinguish between these explanations.

1.4.2 Effect on Transaction Prices

Figure 1.2 presents the main results for total prices by period with the full set of controls and fixed effects.[25] In the periods before the website, there is no significant price difference between procedures that were eventually on the website and those that were not. This provides evidence that the procedures on the website had similar trends in the pre-period as the procedures that were not on the website, supporting the common trends assumption. Once the website launched, the price paid decreased for procedures on the website relative to procedures not on the website. The effect becomes significant in the third year. This effect grows over time, eventually leading to an almost 5 percent reduction in prices five years after the introduction of the website. I argue that this is partially due to the fact that supply-side effects take time to materialize, as prices are renegotiated infrequently. In addition, the website was being used more often in the later period (see Figure 2.3). An alternative interpretation of Figure 1.2 is that the website led to a trend break, decreasing the growth in health care expenditure rather than the level. I explore this interpretation in more detail in Appendix Section 3.6.

The reduction in prices could be due to either lower out-of-pocket prices for consumers or lower insurer reimbursement amount. The first panel of Table 1.5 presents the average effect on out-of-pocket prices over the five-year period. The estimates imply a 6.9 percent reduction in out-of-pocket prices due to the website (from a mean of $86.45) for all individuals in the sample. This result is highly significant. Next I examine how the effect varies by deductible status. Individuals that are not past their deductible are potentially exposed to the full cost of the procedure, and thus have the most to gain from choosing a low-cost provider.[26] Consistent with this fact, I find that individuals subject to

[25]The specification used for Figure 1.2 is $\log(1 + y_{imjkt}) = \beta(\text{OnWeb}_m \times \text{HalfYear}_t) + \alpha X_{it} + \lambda_m + \lambda_k + \varepsilon_{imjkt}$. The interaction with the period before the introduction of the website is omitted.

[26]Depending on the specifics of the plan design, some procedures may not count towards an individual's deductible. In addition, some individuals classified as under their deductible hit their deductible on the current visit. For these reasons, the out-of-pocket cost is often less than the full price when an individual is not past the deductible.

Figure 1.2: The Effect of Price Transparency Website on Visit Price
By Time from Website Introduction

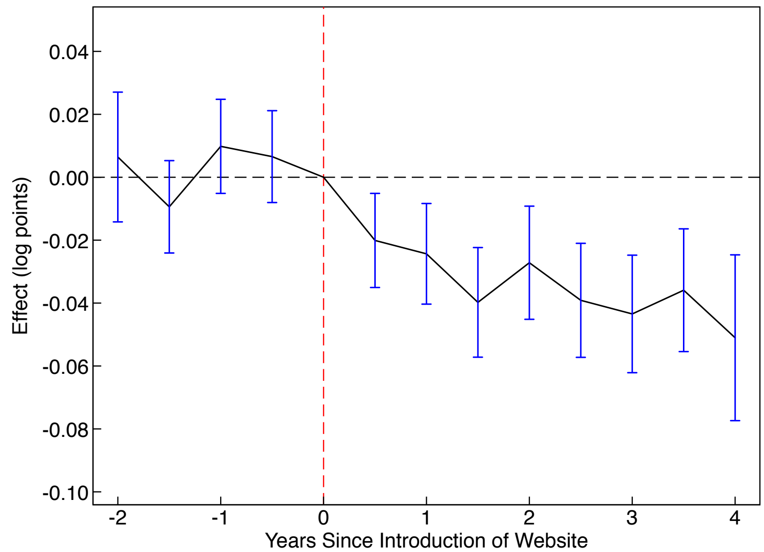

Notes: Chart shows point estimates for each year using the difference-in-difference baseline specification. Error bars indicate 95 percent confidence interval using standard errors clustered at the month-year level.

a deductible see a much larger reduction in price, 14.4 percent (from a mean of $239.70).[27]

Individuals who are not subject to a deductible may still have an incentive to find a low-cost provider. In particular, there are potential gains from switching to a low-cost provider if they are subject to a coinsurance payment.[28] In addition, individuals who do not use the website and find a low-cost provider may still benefit if providers reduce their prices. This is the supply-side effect that I explore in detail in following sections. The results imply that individuals that are not subject to a deductible, both those in a plan without a deductible and those past their deductible, see a significant reduction in prices of about 4 percent (see column 2 and 4 in Table 1.5).

Although individuals are not likely to internalize the cost to the insurer when choosing a provider using the website, there may be an indirect benefit to the insurer. In particular, there is a mechanical correlation between the individual's out-of-pocket price and the insurer price when the individual is subject to a coinsurance payment. To test whether

[27] I conduct a Wald test and determine the difference between the effect for individual subject to a deductible and those not subject to a deductible is significant.

[28] Coinsurance payments are a set percentage of the total price, often between 5 and 25 percent, that are paid by the individual.

Table 1.5: The Effect of Price Transparency Website on Visit Price
Baseline Difference-in-Difference Estimates

	All	No Deductible	Deductible	
			Not Past	Past
	Dep Var: Log(1+Patient Out-of-Pocket Cost)			
$\text{OnWeb}_\text{m} \times \text{Post}_\text{t}$	-0.071^{***}	-0.042^{***}	-0.156^{***}	-0.039^{**}
	(0.010)	(0.009)	(0.019)	(0.018)
Age	-0.007^{***}	-0.007^{***}	-0.006^{***}	-0.004^{***}
	(0.000)	(0.000)	(0.000)	(0.000)
Male	0.024^{***}	0.017^{***}	0.034^{***}	-0.030^{***}
	(0.003)	(0.004)	(0.005)	(0.008)
Mean level	86.45	15.96	239.70	40.89
Adjusted R2	0.373	0.269	0.293	0.149
Observations	3,403,645	1,811,585	1,007,188	584,872
	Dep Var: Log(1+Insurer Paid Amount)			
$\text{OnWeb}_\text{m} \times \text{Post}_\text{t}$	-0.039^{***}	-0.040^{***}	0.009	0.008
	(0.005)	(0.004)	(0.014)	(0.008)
Age	-0.001^{***}	-0.001^{***}	-0.000^{**}	-0.002^{***}
	(0.000)	(0.000)	(0.000)	(0.000)
Male	-0.009^{***}	-0.006^{***}	-0.001	-0.041^{***}
	(0.002)	(0.002)	(0.006)	(0.004)
Mean level	653.06	628.26	588.71	840.68
Adjusted R2	0.290	0.369	0.205	0.377
Observations	3,403,645	1,811,585	1,007,188	584,872
Individual Controls	Yes	Yes	Yes	Yes
Charlson Comorbidity FE	Yes	Yes	Yes	Yes
Insurance FE	Yes	Yes	Yes	Yes
Month-Year FE	Yes	Yes	Yes	Yes
Procedure FE	Yes	Yes	Yes	Yes

Notes: The unit of observation is a patient visit, which may contain multiple medical claims. The sample consists of all commercial claims related to outpatient radiology procedures in the state of New Hampshire over the period 2005 to 2011. OLS regression standard errors clustered at the month-year level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

insurers also benefit from the website, I use the same specification but with the insurer paid amount as the dependent variable. The second panel of Table 1.5 presents the results. The insurers save 3.8 percent on average over the period (from a mean of $653.06). This is statistically significant. Much of this savings is due to individuals without a deductible, consistent with the fact that most of the savings when there is a deductible accrue to the individual.

Although the percentage change in prices is larger for out-of-pocket cost than insurer cost, the insurers actually benefit more from the website in absolute terms. This is because the insurer covers the majority of the cost—88 percent of the total price on average.

1.4.3 Robustness Checks

Appendix Table A2 shows that out-of-pocket price and insurer paid amount results are robust to a number of specifications. The results are not driven by changes in observable characteristics of individuals or changes in insurance plans over the period.[29] In column 4, I include a vector of procedure category specific linear time trends which control for time varying characteristics of procedures that may be correlated with website availability and prices and have a linear trend over the period. The results remain significant with the inclusion of procedure category month trends.

Another concern is unobservable individual characteristics. In particular, individuals that obtain private health insurance after the introduction of the website could be different on unobservable dimensions. In Appendix Table A2 Column 5, I control for individual fixed effects. Identification now comes from the same individuals that received radiology procedures before and after the introduction of the website. The results are robust to this specification.

A related concern is that the website changed the complexity of procedures due to an increase (or decrease) in the probability that an individual has a procedure when informa-

[29] I discuss insurance effects in more detail in Section 1.6.

tion is available. This concern is mitigated by the fact that radiology procedure codes are quite specific, and are standardized across providers. In addition, in Section 1.5, I show that the website did not change the quantity of procedures in general.

The out-of-pocket price may be zero if the insurer pays the full cost due to full insurance. Similarly, the insurer price may be zero if the individual pays the full cost because the individual is under the deductible. Due to this issue, the dependent variable is transformed using $\log(1+y)$. Alternative transformation have been proposed that are arguably less arbitrary (MacKinnon and Magee 1990). In the final column of Appendix Table A2, I show the results are robust to this alternative transformation.

In Appendix Table A3 I conduct a falsification exercise in which I test whether there was an effect on prices in the one year period before the website actually existed. Consistent with the assumption that results are not driven by differential price trends, none of the eight estimates are statistically significant.

1.4.4 Heterogenous Effects

The website may benefit some individuals more than others. In this section I explore how the effect on out-of-pocket prices varies across insurance organization types and individual characteristics by estimating separate models for subpopulations.

Table 1.6 presents results for each of the five insurance organization types. Across all plan types, individual patients have some scope to choose a provider, but the extent to which individuals can choose a provider varies by plan. HMO, POS, and EPO plans have specific rules that seek to guide patients to in-network providers.[30] PPO plans and Indemnity plans (also known as fee-for-service) are more flexible.[31]

[30] In general, individuals with an HMO plan must receive care from a provider within the HMO network, usually with the approval of a primary care physician. EPO plans also require care within a pre-specified network, although without the need for primary care approval. A POS plan is similar, but individuals may be able to go to a provider outside the network with the referral of a primary care physician.

[31] PPO plans allow individuals to go to a provider out of network, but the price may be higher. Indemnity plans do not have a pre-specified network.

Table 1.6: The Effect of Price Transparency Website on Visit Out-of-Pocket Price By Insurer Organization Type

	Restrictive Plans			Non-Restrictive Plans	
	HMO	POS	EPO	PPO	Indemnity
$\text{OnWeb}_m \times \text{Post}_t$	-0.071^{***}	-0.062^{***}	-0.059^{***}	-0.095^{***}	-0.089^{**}
	(0.012)	(0.015)	(0.017)	(0.019)	(0.040)
Full Controls	Yes	Yes	Yes	Yes	Yes
Month-Year FE	Yes	Yes	Yes	Yes	Yes
Procedure FE	Yes	Yes	Yes	Yes	Yes
Adjusted R2	0.380	0.377	0.288	0.364	0.337
Observations	1,380,702	532,512	259,318	1,125,478	105,635

Notes: The dependent variable is Log(1+Patient Out-of-Pocket Cost). OLS regression standard errors clustered at the month-year level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

There is a significant reduction in out-of-pocket prices across all of the insurance organization types, but the effects are larger for non-restrictive plans. Comparing the effect for all restrictive plans (HMO, POS, and EPO) with all non-restrictive plans (PPO and indemnity), the difference is significant.[32] This is evidence that the website is most beneficial for those with the most discretion regarding their choice of provider. One criticism of unrestrictive plans, such as indemnity plans, is that uninformed individuals may make poor choices regarding where to receive care, potentially driving up health care costs. The results presented here suggest this can be mitigated by reducing information frictions in the market.

I examine results by demographic characteristics to examine which groups benefits most from the price information. These results are presented in Table 1.7. First, I examine individuals who received a radiology procedure immediately after having an emergency. Note these are relatively minor emergency visits since I exclude inpatient admissions. I find no statistically significant effect suggesting that this subpopulation was less likely to use the website compare to individuals having non-emergency procedures that could potentially be scheduled further in advance.

[32] A Wald test yields an F-statistics of 4.57.

Table 1.7: The Effect of Price Transparency Website on Visit Out-of-Pocket Price By Patient Characteristics

	Emergency Visit		Previous Healthcare Experience		Charlson Comorbidity Index	
	Yes	No	$\leq$ Median	$>$ Median	$= 0$	≥ 1
$\mathrm{OnWeb_m} \times \mathrm{Post_t}$	-0.021	-0.072^{***}	-0.108^{***}	-0.035^{**}	-0.093^{***}	-0.047^{***}
	(0.014)	(0.010)	(0.018)	(0.015)	(0.013)	(0.011)
F statistic of diff.	8.608^{***}		9.746^{***}		7.268^{***}	
Adjusted R2	0.396	0.362	0.461	0.283	0.419	0.319
Observations	212,310	3,191,335	908,425	888,974	1,856,019	1,547,626
	Age		Education		Income	
	≤ 40	> 40	$\leq 1^{st}$ Quartile	$> 4^{th}$ Quartile	$\leq 1^{st}$ Quartile	$> 4^{th}$ Quartile
$\mathrm{OnWeb_m} \times \mathrm{Post_t}$	-0.090^{***}	-0.056^{***}	-0.063^{***}	-0.104^{***}	-0.071^{***}	-0.091^{***}
	(0.014)	(0.012)	(0.013)	(0.016)	(0.013)	(0.014)
F statistic of diff.	3.432^{*}		4.223^{**}		1.160	
Adjusted R2	0.360	0.367	0.361	0.382	0.360	0.379
Observations	1,237,595	2,166,050	854,013	799,104	858,653	843,406
Full Controls	Yes	Yes	Yes	Yes	Yes	Yes
Month-Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Procedure FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The dependent variable is Log(1+Patient Out-of-Pocket Price). OLS regression standard errors clustered at the month-year level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

There is a literature that suggests that individuals with knowledge of the health care system are able to make more informed decisions about care.[33] Individuals with experience in the health care system may be able to choose low cost providers even without the website. I find evidence that individuals with less previous healthcare experience, defined as those with less than the median number of claims in the previous year, benefit significantly more from the website. While both groups gain, individuals with less experience save 6.8 percentage points more, a statistically significant difference. Individuals without comorbidities also benefit significantly more.[34] This may be because they tend to have less experience with the health care system or because, due to better health, they are better able to shop-around.

In the year the website started, individuals age 18 to 29 were 23 percentage points more

[33] See review by Berkman et al. (2011).

[34] The results of a Wald test are presented in Table 1.7. The null hypothesis is that the two coefficients are equal.

likely to have broadband internet then those age 50 to 64.[35] Consistent with this fact, younger individuals, defined as those age 40 and younger, benefit more from the website, saving 3.2 percentage points more than older individuals (see second panel in Table 1.7). This difference is significant at the 10 percent level. Broadband adoption is also positively correlated with education and income.[36] Estimates imply larger effects for those in the highest education quartile and highest income quartile, although the difference is only significant for education.

Taken together, these results imply that price transparency provides benefits across a range of demographic groups. However, the benefits accrue most to those groups that had access to the internet and had the ability to shop-around, potentially exacerbating health inequality.

Finally, I explore the effect of the website by price quantile. The website has an affect across a large range of prices but is particularly effective for procedures in the medium-high range. I discuss the empirical strategy and results in Appendix Section 3.6.

1.4.5 Demand-Side Effects

The primary motivation for the website was to allow individuals to shop-around for medical care. In Table 1.8 I directly examine provider choice outcomes using the same difference-in-difference specification to provide evidence that price effects are mediated by changes in demand.

When individuals have a radiology procedure, they often go to the same provider as their last radiology visit. Over the period, 34 percent of individuals with repeat visits went to the same provider as their prior visit. Using a difference-in-difference linear probability model where the dependent variable is an indicator for whether the chosen

[35] Pew Internet and American Life Project, Home Broadband Adoption, July 2007.

[36] In 2007, college graduates were 49 percentage points more likely to have broadband than those with less than a high school education. Those with income over $75,000 were 46 percentage points more likely than those with an income under $30,000. See Pew Internet and American Life Project, Home Broadband Adoption, July 2007.

Table 1.8: The Effect of Price Transparency Website on Search Behavior and Provider Type

	Dependent Variable:			
	Same Provider as Last Radiology Visit	Low Cost Provider	Provider Type is Hospital	Provider in New Hampshire
$\text{OnWeb}_{\text{m}} \times \text{Post}_{\text{t}}$	-0.0325^{***}	0.0454^{***}	-0.0031^{**}	0.0045^{***}
	(0.0112)	(0.0059)	(0.0014)	(0.0011)
Full Controls	Yes	Yes	Yes	Yes
Month-Year FE	Yes	Yes	Yes	Yes
Procedure FE	Yes	Yes	Yes	Yes
Mean dep. var.	0.34	0.40	0.12	0.29
Adjusted R2	0.053	0.149	0.134	0.451
Observations	1,571,869	2,882,093	3,404,442	3,404,442

Notes: Results from linear probability model. The unit of observation is a patient visit, which may contain multiple medical claims. The sample consists of all commercial claims related to outpatient radiology procedures in the state of New Hampshire over the period 2005 to 2011. Low cost provider is defined as a provider with an average out-of-pocket cost in the lowest decile in each county conditional on procedure, insurer, and year. OLS regression standard errors clustered at the month-year level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

provider is the same as the prior visit, I examine whether the website made individuals more likely to switch providers. Column 1 in Table 1.8 presents the results. The website significantly reduced the probability of going to the same provider as the previous visit. In particular, individuals with access to the website were 9.6 percent less likely to go to the same provider (3.3 percentage point reduction from the mean of 34 percent).

Over the period, 12 percent of individuals visit hospitals while the rest go to non-hospital providers such as imaging centers or clinics. On average, the out-of-pocket price at hospitals is 31 percent higher for the same procedures.[37] The third column of Table 1.8 shows results for a specification in which an indicator for whether the provider is a hospital is the dependent variable. The website led to a 2.5 percent reduction in the probability of having a radiology procedure at a hospital rather than a non-hospital provider

[37] I regress log out-of-pocket price on an indicator for whether the provider is a hospital and procedure fixed effects. Note that the facility and equipment for outpatient radiology procedures at hospitals is often the same as those used for inpatient procedures.

(0.3 percentage point reduction from the mean). This effect is statistically significant.

New Hampshire is a relatively small state and many resident work in surrounding states, particularly in Massachusetts.[38] Almost a third of the individuals in the sample go to radiology providers outside the state, however the website only provides information about the price of providers within the state. The fourth column of Table 1.8 shows that the website significantly increased the probability that individuals went to a provider within New Hampshire.

1.4.6 Supply-Side Effects

Given that individuals switch to lower-cost providers, did this put downward pressure on prices? I examine supply-side effects by estimating a model with controls for demand-side factors, namely provider-procedure-insurer fixed effects. Rather than examining out-of-pocket prices and insurer price separately, I now consider the total price (i.e. allowed amount).

The main results can be seen in Figure 1.3. Panel a shows the baseline total effect on visit prices which includes both demand-side and supply-side effects. Consistent with the results in previous sections, there is no significant effect prior to the introduction of the website, and then a significant effect afterwards. Panel b shows the results after controlling for demand-side factors. The point estimates are smaller in magnitude, especially in the period right after the introduction of the website. However, the estimates are still highly significant, especially in the later period, implying a reduction in provider prices.

Recall that the visit price is determined by the prices of a bundle of procedures. The primary way that the provider can change the price of the bundle is by changing the price of the principal radiology procedure, which makes up more than half of the cost of the bundle on average. In Panel c and d of Figure 1.3 I examine the effect on the principal

[38] About 17 percent of workers living in the New Hampshire work in a different state, one of the highest rates in the nation. See Out-of-State and Long Commutes, Census 2011.

procedure price. Panel c presents the baseline specification while Panel d presents the supply-side effect. Panel d shows a very small effect in the two years after the introduction of the website. However, there is a larger supply-side effect three to five years after the website. I interpret this as evidence that the website reduced negotiated prices in the long-run.

Table 1.9 formalizes these results. There is both a short-run and long-run reduction in transaction prices, where short-run is defined as the two-years after the website started. However, after isolating the supply-side, the short-run effect is quite small. Visit prices declined by 1.7 percent while principal procedure price declined by 1.9 percent. The long-run effects are larger—there is a 2.2 percent reduction in visit prices and 4.7 percent reduction in principle procedure price. These results are all statistically significant.

Table 1.9: The Effect of Price Transparency Website on Supply-Side

	Baseline		Provider-Procedure-Insurer FE	
	Total Visit Price	Principal Procedure Price	Total Visit Price	Principal Procedure Price
$\text{OnWeb}_{\text{m}} \times \text{PostShortRun}_{\text{t}}$	-0.034^{***}	-0.036^{***}	-0.017^{***}	-0.019^{***}
	(0.004)	(0.004)	(0.005)	(0.004)
$\text{OnWeb}_{\text{m}} \times \text{PostLongRun}_{\text{t}}$	-0.038^{***}	-0.055^{***}	-0.022^{***}	-0.048^{***}
	(0.004)	(0.004)	(0.005)	(0.005)
Indiv. Controls	Yes	Yes	Yes	Yes
Insurer FE	Yes	Yes	No	No
Procedure FE	Yes	Yes	No	No
Provider*Procedure *Insurer FE	No	No	Yes	Yes
Month-Year FE	Yes	Yes	Yes	Yes
Mean level	788.64	413.60	791.08	412.24
Adjusted R2	0.357	0.625	0.515	0.790
Observations	3,414,053	3,414,050	3,377,271	3,377,268

Notes: The dependent variable is $\log(1+y)$, where y is either the visit price or principal procedure price. This price includes both patient and insurer payments (i.e. allowed amount). For visit price, the unit of observation is a patient visit, which may contain multiple medical claims. For principal procedure price, the unit of observation is the primary radiology procedure within each visit. The sample consists of all commercial claims related to outpatient radiology procedures in the state of New Hampshire over the period 2005 to 2011. OLS regression standard errors clustered at the month-year level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Overall, this is evidence that there was a significant reduction in negotiated prices. The fact that the principal procedure price is most affected in the long-run is consistent with the fact that prices are renegotiated infrequently.

One caveat is that the changes in demand caused providers to reduce the price of supplemental procedures in addition to principal procedure prices. Since supplemental procedures, such as contrast agents and examinations related to radiology procedures, are common across procedures on the website and procedures not on the website, these reductions in prices would be "differenced-out". This would lead to an underestimate of the supply-side effects of price transparency. Thus, Table 1.9 are conservative estimates.

Providers operating in concentrated markets may be able to negotiate higher prices with insurers (e.g. Dranove, Shanley and White 1993; Town and Vistnes 2001; Gowrisankaran, Nevo and Town 2015). I define the Herfindahl index by county and procedure category in the period prior to the introduction of the website. There is significant variation in competition—some counties have a single provider in the market for certain procedure categories while others are relatively unconcentrated.[39]

Table 1.10 presents the principal procedure effect by market concentration. Column 1 and 2 show the results for transaction prices. In there short-run, there is no statistically significant difference between the effect in low concentration markets (those with Herfindahl index at or below the median) and high concentration markets (those with above median Herfindahl index). However, in the long run, the reduction in prices is 1.5 percentage points larger in high concentration markets. This difference is statistically significant. Results are similar after isolating the supply-side (column 3 and 4 in Table 1.10). While there is no significant difference in the short-run, providers reduce their prices by 1.9 percentage points more in concentrated markets in the long-run. This is evidence that price transparency put the most downward pressure on prices in markets where price cost margins were likely the highest.

One important concern is that the supply-side effects reflect cross-subsidization. In other words, price transparency puts downward pressure on the prices of procedures on

[39]The distribution of Herfindahl index in the sample is shown in Figure A1.

Table 1.10: The Effect of Price Transparency Website on Supply-Side Principal Procedure Price
By Herfindahl Index

	Baseline		Provider-Procedure-Insurer FE	
	$\leq$ Median HHI	$>$ Median HHI	$\leq$ Median HHI	$>$ Median HHI
$\text{OnWeb}_m \times \text{PostShortRun}_t$	-0.038^{***}	-0.035^{***}	-0.019^{***}	-0.028^{***}
	(0.005)	(0.005)	(0.006)	(0.006)
$\text{OnWeb}_m \times \text{PostLongRun}_t$	-0.047^{***}	-0.063^{***}	-0.043^{***}	-0.063^{***}
	(0.004)	(0.006)	(0.007)	(0.007)
F statistic of diff. (SR)	0.130		1.031	
F statistic of diff. (LR)	4.570^{**}		4.253^{**}	
Adjusted R2	0.628	0.623	0.790	0.791
Observations	1,667,383	1,685,266	1,650,459	1,666,555
Indiv. Controls	Yes	Yes	Yes	Yes
Insurer FE	Yes	Yes	No	No
Procedure FE	Yes	Yes	No	No
Provider*Procedure *Insurer FE	No	No	Yes	Yes
Month-Year FE	Yes	Yes	Yes	Yes

Notes: The dependent variable is $\log(1+y)$, where y is the principal procedure price. This price includes both patient and insurer payments (i.e. allowed amount). HHI is calculated for individual's county in each year. OLS regression standard errors clustered at the month-year level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

the website, but providers compensate by increasing prices for procedures that are not on the website or cutting care elsewhere. There is little theoretical justification for why profit-maximizing providers would cross-subsidize (Frakt 2011). However, this may not be the case if hospitals have an objective function other than profit-maximization. For instance, there is evidence that hospitals use profits from procedures with high margins to fund unprofitable services such as psychiatric, substance-abuse, or trauma-care (David et al. 2014). In the sample used for this paper, 80 percent of providers are imaging centers or other non-hospital facilities which do not provide services thought to be unprofitable. I argue that cross-subsidization is unlikely for these providers, although it may be an issue for providers such as Dartmouth-Hitchcock, New Hampshire's only academic hospital.

1.4.7 Price Dispersion

Information frictions can give rise to price dispersion in equilibrium (e.g. Salop and Stiglitz 1977; Burdett and Judd 1983). In this section, I examine whether a reduction in information frictions can reduce price dispersion, as measures by the interquartile range of prices. I find evidence that the website reduced the range of transaction prices as well as the range of negotiated provider prices.

Table 1.11: The Effect of Price Transparency Website on Price Dispersion

	Interquartile Range of Transaction Prices		Interquartile Range of Provider Prices	
	Total Visit Price	Principal Procedure Price	Total Visit Price	Principal Procedure Price
$\text{OnWeb}_{\text{m}} \times \text{Post}_{\text{t}}$	-141.24^{***}	-62.64^{***}	-79.02	-47.40^{**}
	(45.36)	(20.23)	(49.63)	(17.99)
Indiv. Controls	Yes	Yes	Yes	Yes
Procedure FE	Yes	Yes	Yes	Yes
Month-Year FE	Yes	Yes	Yes	Yes
Mean IQR	1034.1	581.6	844.9	471.3
Adjusted R2	0.297	0.466	0.292	0.484
Observations	20,458	20,458	20,458	20,458

Notes: The unit of observation is a visit-month for columns 1 and 3 and procedure-month for columns 2 and 4. The sample consists of all commercial claims related to outpatient radiology procedures in the state of New Hampshire over the period 2005 to 2011. OLS regression standard errors clustered at the month-year level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Column 1 and 2 in Table 1.11 presents the results for transaction prices. There is a \$141 reduction in the interquartile range of visit prices on the website relative to those not on the website (a 13.6 percent reduction from the mean of \$1,043). There is also a \$63 reduction in the interquartile range of principal procedure prices. Both these results are significant.

More importantly, I examine the effect on the interquartile range of negotiated provider prices (see column 3 and 4 in Table 1.11). The interquartile range of visit prices declines by \$79, but is not statistically significant. Focusing on the principal procedure price, which was likely most affected by the website, I find a significant reduction in price dispersion. In particular, the interquartile range of provider prices declined by $47, or 10.1 percent of the mean.

Together with the previous results, these results imply that website decreased both the mean and variance of the distribution of transaction prices. This was due in part to a similar shift in the distribution of negotiated prices.

1.5 Effect on Quantity

In this section I examine whether the website changed the quantity of radiology visits. In most cases, individuals choose whether to get a radiology procedure after a primary care physicians recommends that they receive a specific procedure. To the extent that individuals believe a procedure is optional, they may be more likely to get a procedure if information about price is available and they know they can choose an inexpensive provider. Conversely, the price may be more than the individual expected, leading to lower quantity when price information is available.

1.5.1 Empirical Strategy

I use the similar difference-in-difference specification to examine whether the introduction of the website changed the quantity of procedures on the website relative to procedures not on the website. In particular, the specification is

$$y_{iwt} = \beta(\text{OnWeb}_w \times \text{Post}_t) + \gamma\text{OnWeb}_w + \alpha X_{it} + \lambda_k + \lambda_t + \varepsilon_{iwt} \tag{1.4}$$

where the outcome, y_{iwt}, is either the number of procedures or an indicator for whether the individual ever had the procedure during the year. In the latter case, it becomes a linear probability model. I also include individual covariates, insurer fixed effects, and year fixed effects.[40]

Unlike previous specifications, I use the universe of privately-insured individuals in the state, including those that never had a radiology procedure over the period. The unit of observation is an individual in a year for procedures on the website and not on the website (i.e. procedures are aggregated by whether they are on the website, which is indexed by w).

1.5.2 Results

Columns 1 and 2 of Table 1.12 show the results for the whole sample. Relative to the probability of having a procedure not on the website, the probability of having a procedure on the website increased by 1.2 percentage points. The average number of radiology visits on the website also increased. In particular, individuals had 0.04 more visits after the introduction of the website. Both these results are significant.

Mammogram screening exams are often at the discretion of individuals, whereas other radiology procedures are often needed at a specific point in time to diagnose symptoms or treat an acute illness. Over the period of analysis, the annual probability of having a mammogram in New Hampshire increased from 9.5 percent to 11.8 percent.

In Columns 3 and 4 of Table 1.12 I exclude mammograms from the analysis. Using the same difference-in-difference specification I find no effect on the quantity of procedures due to the website. Both point estimates for the probability of having a procedures and the number of procedures are small and not statistically significant. In particular, I reject the hypothesis that the website increased the probability of having a radiology procedure by more than a percentage point.

I interpret these results as evidence that the quantity of radiology procedures, excluding mammography, was not affected by the introduction of the website. It is important

[40] Individual covariates includes age, gender, charlson comorbidity index, income, education, and rural classification.

Table 1.12: The Effect of Price Transparency Website on Visit Quantity

	Dependent Variable:			
			Excluding Mammograms	
	Any Radiology Visit	Number of Radiology Visits	Any Radiology Visit	Number of Radiology Visits
$\text{OnWeb}_{\text{m}} \times \text{Post}_{\text{t}}$	0.0122**	0.0380**	0.0004	0.0166
	(0.0045)	(0.0111)	(0.0044)	(0.0087)
Full Controls	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
On Website FE	Yes	Yes	Yes	Yes
Mean dep. var.	0.21	0.36	0.21	0.36
Adjusted R2	0.299	0.193	0.217	0.138
Observations	8,948,638	8,948,638	8,948,638	8,948,638

Notes: The unit of observation is an individual-year. The sample consists of all individuals privately insured in New Hampshire or residing in New Hampshire between 2005 and 2011. OLS regression standard errors clustered at the year level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

to note that these results do not necessarily imply that the website caused an increase in mammograms. Many confounding factors, such as changes in unobservable characteristics of the population or changes in physician recommendations regarding mammograms, could explain the increase in mammography.

Past work has stressed the importance of the extensive margin when considering health care costs (e.g. Brot-Goldberg et al. 2015). However, it is important to note that my definition of quantity differs from much of the literature. Often quantity or utilization is quantified in terms of the number of medical claims. However, some providers may be expensive because they bill for superfluous supplementary procedures that do not change the overall quality of care. For instance, two individuals may receive the exact same radiology procedure, but providers submit a different number of medical claims. I argue that if an individual switches to a provider that is lower cost due to fewer claims it is a reduction in price, not a reduction in the quantity of care.

1.6 Insurance Choice and Insurance Plan Characteristics

The website allows users to compare the price of common procedures under different insurance plans. Therefore, the website could be used by individuals or employers when deciding on which insurance plan to purchase. In particular, if individuals can anticipate the health insurance procedures they will need over the coming year, they may be able to use the website to choose a plan that provides low prices for these procedures, magnifying selection effects. However, I do not find evidence that changes in demand for insurance are driving the reduction in radiology prices. [41]

In order to examine whether selection into insurance plans is a factor for radiology prices, I examine whether the difference-in-difference specifications are robust to the inclusion of insurance plan fixed effects. In the baseline specifications, I include fixed effects that control for changes in the composition of insurance plans (see Table 1.5). I also examine a specification without insurance plan fixed effects (see Appendix Table A2, column 2). The difference in the estimated effect between these specifications is not statistically significant implying that the composition of insurance enrollment is not driving the reduction in radiology prices due to the website. I also test whether individuals who switch insurers are different than those that do not switch after the introduction of the website. Table A5 shows that both groups benefit from the website, and there is no significant difference. These results are consistent with the fact that radiology procedures over the coming year are often difficult to predict, and are unlikely to be the primary reason that individuals choose a given insurance plan.

It is still possible that the website affected insurance choice, even though the composition of insurance plans is not responsible for the reduction in radiology prices. The com-

[41] After my period of analysis, the website added separate features for comparing insurance plans. Future research can examine the effect of this information.

position of insurance plans, both insurance firm and plan type, is presented in Appendix Figure A2. There does not appear to be a large shift in the composition of insurance plans after the website launched in March 2007. In order to test this within a regression framework, I use the following specification to examine whether there was a shift in the composition of insurance plans.

$$y_{kt} = \beta \text{Post}_t + \alpha X_t + \gamma t + \varepsilon_{imjkt} \tag{1.5}$$

where y_{kt} is the number of enrollees in each plan type or company in each month. I control for characteristics of individuals in each month, X_t, as well as a linear time trend.[42] This specification required stronger identifying assumptions than the difference-in-difference method, namely constant trends in the absence of the website. I use the period consisting of two years before and after the introduction website in order to avoid confounding factors that influence insurance enrollment.

I examine enrollment for each of the major insurance companies in the state as well as each of the insurance types. Of the ten estimates, one one is statistically significant (see Appendix Table A6). I interpret this as evidence that the website did not have an immediate effect on insurance choice.

A related mechanism of interest is endogenous plan response to price transparency. For instance, plans may change cost-sharing rules in response to the website, further incentivizing individuals to use the website. I modify the baseline regression specification by interacting insurance fixed effects with an indicator for whether the website is available.[43] The estimates from this specification are not statistically different from the baseline specification (see Appendix Table A2, column 3). I interpret this as evidence that changes to plan characteristics after the start of the website that affected all procedures are not

[42]Individual controls include age, gender, income, and education.

[43]In particular, I estimate $\log(1 + p_{imjkt}) = \beta(\text{OnWeb}_m \times \text{Post}_t) + \alpha X_{it} + \lambda_m + \lambda_k \times \text{Post}_t + \lambda_k \times \text{Pre}_t + \lambda_t + \varepsilon_{imjkt}$ where $\text{Pre}_t = (1 - \text{Post}_t)$.

driving the reduction in prices.[44]

1.7 Conclusion

The health care system can be very complicated to navigate, and information frictions are thought to be pervasive (Reinhardt 2012). In this paper, I examine how a publicly available website providing price information to consumers affected the market for radiology procedures. While previous research has focused on the demand-side effect of information supplied by specific employers, I examine the equilibrium effects using the universe of private medical claims in the state.

Overall, I estimate that the HealthCost website reduced the cost for individuals by 6.9 percent and for insurers by 3.8 percent. A simple calculation implies that individuals saved around $20.9 million and insurers saved $86.7 million on radiology visits over the 5 year period.[45] I argue that while demand-side effects are important, there are significant supply-side effects in the long-run when information is available to all consumers in the market. In other words, this is evidence that price opacity softens provider competition. This effect is particularly important given that the average price of radiology procedures in the U.S. is roughly double that of other OECD countries.[46]

Given that website traffic logs reveal that only a small fraction of individuals receiving radiology procedures in New Hampshire use the website, the supply-side effects may be quite large if all consumers were informed about prices. In the following chapters, I examine these counterfactuals using an empirical model of the market for radiology procedures.

[44] Note that after my period of analysis, insurers in the state began offering additional incentives for individuals to choose low-cost providers. These incentives may be due to the presence of the website and could potentially lead to larger long-term effects of the website.

[45] All figures in 2010 dollars.

[46] The average price of an MRI scan is $1,200 in the US but only $569 in other OECD countries with available data. The average price of CT scan is $228 in the US but only $98 in other OECD countries with available data. See Squires (2011).

Figure 1.3: The Effect of Price Transparency Website on Supply-Side
By Time from Website Introduction

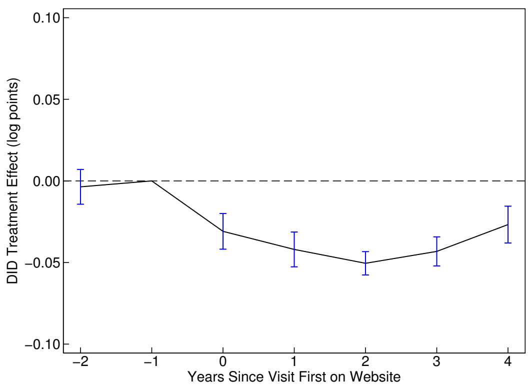

a. Visit price, baseline

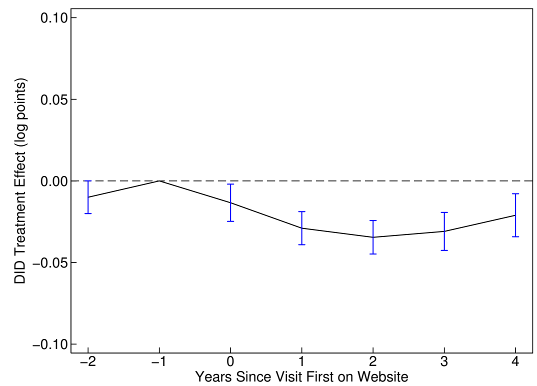

b. Visit price, with provider-procedure-insurer FE

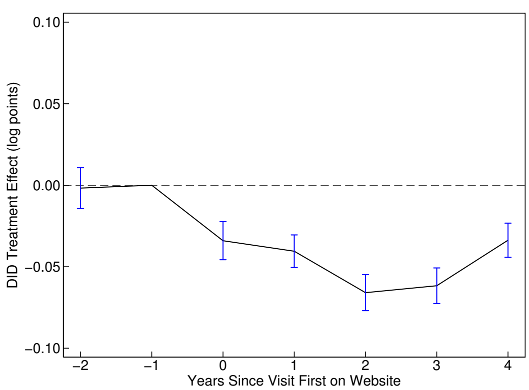

c. Principal procedure price, baseline

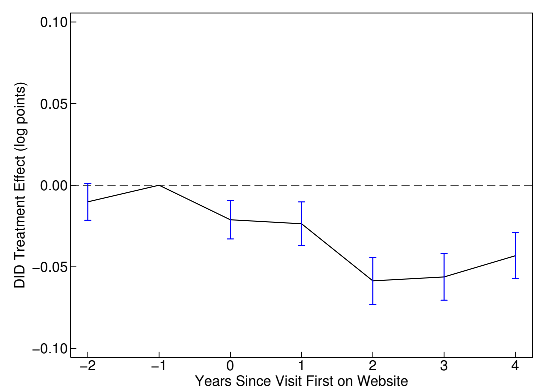

d. Principal procedure price, with provider-procedure-insurer FE

Notes: Chart shows point estimates for each year using the difference-in-difference specifications in Table 1.9. The dependent variable is $\log(1 + p)$, where p is either the principal procedure price or visit price. This price includes both patient and insurer payments (i.e. allowed amount). Error bars indicate 95 percent confidence interval using standard errors clustered at the month-year level.

Chapter 2

An Empirical Model of Price Transparency and Demand for Health Care Services

2.1 Introduction

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

While an influential literature, starting with Stigler (1961) and Diamond (1971), has examined search frictions, there has been little emphasis on markets in which it is not possible to acquire price information. Like search costs, the lack of price transparency may increase prices and lead to price dispersion. Understanding how price transparency affects prices is particularly important for privately-provided health care in the U.S. since the market comprises about 6 percent of GDP.

This paper empirically evaluates how price transparency affects spending in the U.S. health care market, focusing on the demand-side. I combine a model of demand that incorporates price uncertainty with a model of bargaining between providers and insurers. Although a relatively small fraction of consumers currently try to obtain price information when price transparency tools are made available, the model allows for an analysis of

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

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

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

out-of-sample counterfactual scenarios. In Chapter 3, I combine with demand model with model of the supply-side in order to examine how price transparency affects negotiated prices.

I introduce a discrete-choice model in which consumers choose where to receive medical care with potentially limited information about prices. In the model, consumers with rational expectations receive noisy signals about prices, and, consequently, are less able to discern which are the low price options. Consumers may choose options they believe to be the best value but are often surprised by the bill. Accounting for the difference between expected prices and actual prices is important for recovering underlying consumer preferences, including price sensitivity, and evaluating the welfare effects of price information.

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

In the demand model, individuals' beliefs about prices are treated as unknown parameters to be estimated.[4] These high dimensional latent variables complicate the estimation strategy. To address this issue, I take advantage of recent Bayesian techniques and employ a Markov chain Monte Carlo (MCMC) estimator. This approach allows for a feasible estimation strategy that recovers parameter estimates summarizing individuals' beliefs about prices in addition to underlying taste parameters.

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

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 the actual out-of-pocket price that consumers pay as well as the price paid by insurers.[5] I focus on relatively simple outpatient medical imaging procedures—X-Rays, CT scans, and MRI scans.[6] The negotiated price of these procedures ranges from a few hundred dollars for X-rays to a few thousand for MRI scans. Despite the fact that specific medical imaging procedures are relatively standardized, I find that the price of each procedure varies widely across providers in the state.[7] 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.

Estimates from the demand model imply that, in the absence of the price transparency website, consumers have significant uncertainty about prices. In particular, individual beliefs are wrong by 37 percent on average.[8] When individuals use the website and learn true prices, they make different decisions about where to receive care, especially if they are under their deductible and are exposed to the full price. This in turn affects the solution to the bargaining problem between insurers and providers, which is the focus of Chapter 3. The estimates also show that the individuals who use the HealthCost website tend to be those that benefit most from the information, such as those with higher price sensitivity or those in insurance plans with greater cost sharing.

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

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

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

[8] This is the absolute value of the difference as a proportion of the true price. See Section 2.4 for further explanation.

2.1.1 Related Literature

This paper is related to the large literature on search costs and competition, starting with Stigler (1961). Even with homogenous goods and many sellers, search costs can lead to higher prices (e.g. Diamond 1971; Stahl 1989). Search costs have been shown to be empirically important in a large variety of markets.[9] 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).[10] In contrast, this paper studies a context in which individuals make decisions under uncertainty. Although there are similarities to search frictions, the welfare consequences of price uncertainty are distinct since individuals may be surprised by their bill. The model presented in this paper has implications for other situations in which it is not possible to observe actual prices when making a purchase decision, such as markets where consumers receive price quotes.

This paper is also related to the literature examining markets with shrouded add-on pricing. The price of add-ons may be shrouded in equilibrium due to consumer lack of self-control (DellaVigna and Malmendier 2004), selection issues (Ellison 2005), bounded rationality (Spiegler 2006), or myopia (Gabaix and Laibson 2006).[11] Empirical work has found that obfuscation of shipping charges affects consumer behavior (Ellison and Ellison 2009; Brown, Hossain and Morgan 2010). Related work on bill-shock has examined situations in which consumers are inattentive about the price of the next unit of consumption, such as for cellular phone contracts (Grubb 2014; Grubb and Osborne 2015). Pricing in the market for medical services can be seen as the limit-case of add-on pricing—in the

[9]Empirical work has studied search frictions in a variety of markets including prescription drugs, mutual funds, textbooks, online bookstores, grocery stores, auto insurance, electricity, online hotel booking, cars, and trade-waste (Sorensen 2000; Hortaçsu and Syverson 2004; Hong and Shum 2006; De Los Santos, Hortaçsu and Wildenbeest 2012; Seiler 2013; Honka 2014; Giulietti, Waterson and Wildenbeest 2014; Koulayev 2014; Moraga-González, Sándor and Wildenbeest 2015; Salz 2015). There has also been related work on technology that reduces search costs (e.g. Brown and Goolsbee 2002; Jensen 2007; Jang 2015; Luco 2015).

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

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

absence of price transparency tools the full price is partially shrouded. Therefore, the model developed in this paper can be seen as a new approach to add-on pricing in which consumers have noisy beliefs about shrouded attributes and maximize expected utility.

Previous work has also examined the case in which consumers lack information about product attributes other than price.[12] For experience goods, consumers may initially lack information about product quality or other non-price attributes (e.g. Erdem and Keane 1996; Ackerberg 2003; Erdem, Keane and Sun 2008; Allcott 2013). Building on Allcott (2013), as well as theoretical work by Schmeiser (2014), Train (2015) formalizes the calculation of consumer surplus in discrete-choice models when anticipated attributes are different from experienced attributes. I use this approach to calculate welfare when price is different than expected.

While this paper argues that information frictions are important for understanding consumers' choice of medical providers, a broader literature has emphasized frictions in other parts of the health care system. For instance, Handel and Kolstad (2015) find evidence that a variety of frictions affect health insurance choice. Prior literature has also found that consumer inattention or inertia has implications for Medicare Part D (e.g. Ericson 2014; Decarolis 2015; Ho, Hogan and Morton 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. Cutler, Huckman and Landrum 2004; Kolstad 2013; Grennan and Town 2015). Finally, Grennan and Swanson (2016) find that information affects hospital-supplier bargaining. Despite this growing literature, to my knowledge, there is no evidence on the welfare effects of frictions that affect consumers' choice of hospital.

Finally, this paper seeks to integrate quasi-experimental and structural approaches.[13]

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

[13] For a related discussion see Angrist and Pischke (2010) and response by Einav and Levin (2010) and Nevo and Whinston (2010).

There is a small but growing literature that directly compare results from quasi-experimental reduced-form methods with estimates from a structural model (e.g.). I contribute to this literature by demonstrating that an empirical model can exploit a natural experiment to perform out-of-sample counterfactual simulations. I compare these methods in greater detail in Chapter 3.

2.1.2 Roadmap

The remainder of the paper is organized as follows. Section 2.2 describes the data and provides additional background on the price transparency website not provided in Chapter 1. Section 2.3 presents the model of website usage and choice of medical provider. I also discuss the estimation and robustness. Section 2.4 presents the results from the demand model. Section 2.6 concludes.

2.2 Data and Background

I utilize an all-payer claims database from New Hampshire that provides detailed information on negotiated prices along with information about how much is paid by the individual versus the insurer. I use these data to construct the individual-specific out-of-pocket price for each option in individuals' choice sets. In Section 2.2.2 I provide background on the price transparency website in New Hampshire and describe the variation that is used to estimate the model. I also describe the website traffic data which is used to construct information about the fraction of consumers with price information when the website is available.

2.2.1 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.[14] These data were collected as part of the New Hampshire Comprehensive Health Care Information System (NHCHIS), which assembled data from all commercial insurers in the state. The data were collected by the state in order to analyze health spending and construct prices for the price transparency website.

This paper analyzes the market for outpatient medical imaging services. This includes X-rays, computerized tomography (CT) scans, and magnetic resonance imaging (MRI) scans, all of which are diagnostic procedures that provide internal images of the body. Note that in Chapter 1, I use a broader definition of radiology procedures that includes procedures such as bone density scans and PET scans. In this paper I focus on X-rays, CT scans, and MRI scans because these procedures are relatively common, allowing me to construct accurate prices within individuals' choice sets.[15]

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.[16] 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

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

[15] In Chapter 1, I find evidence that the price transparency website may have affected the quantity of mammograms (but not other procedures). I also exclude mammograms from the analysis, allowing me to assume that all individuals choose an inside-option.

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

imaging visit may contain multiple procedures.[17] 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 Chapter 1.

Each visit is categorized by the imaging procedure, defined by a CPT/HCPCS code.[18] These codes are quite specific and refer to relatively standardized procedures. The full list of medical imaging procedures is given in Table A7. For each visit, I am able to calculate the out-of-pocket price paid by consumers, the price paid by insurers, as well as the list price.[19] The list price is not relevant for individuals in the sample since insurers negotiate prices that are lower than the list prices. This negotiated price is obtained by simply summing the amount paid by consumers and insurers. The ratio of the out-of-pocket price to the negotiated price determines the individual-specific level of cost sharing (e.g. if the individual is under the deductible, then the cost sharing is equal to 1). Prices are inflation-adjusted to 2010 dollars using the Medical Care Services CPI from the U.S. Bureau of Labor Statistics.

For each visit, an identifier allows me to link information about the medical provider that performed the procedure, which includes both hospital and non-hospital facilities. While hospitals offer outpatient medical imaging services, freestanding outpatient facilities (e.g. imaging centers) are significantly less expensive. In New Hampshire, the aver-

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

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

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

age total cost of imaging procedures is $1,004 at hospitals but only $797 at non-hospital providers. In addition to observing provider type, I also observe the provider zip code.[20] The location of these providers is shown in Figure 2.1.

For individuals, I observe age, sex, zip code, insurance enrollment, and whether they are subject to a deductible. I also observe a patient identifier. I define 5 different age groups (0-18, 19-35, 36-50, 51-64) and omit individuals over age 65 since they are likely eligible for Medicare. Average income and education using the 2007-2010 American Community Survey is linked to each individual using the zip code. In addition, patient zip code is used to calculate the distance to each provider. Using observed International

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

Figure 2.1: Density of Consumers and Location of Medical Imaging Providers

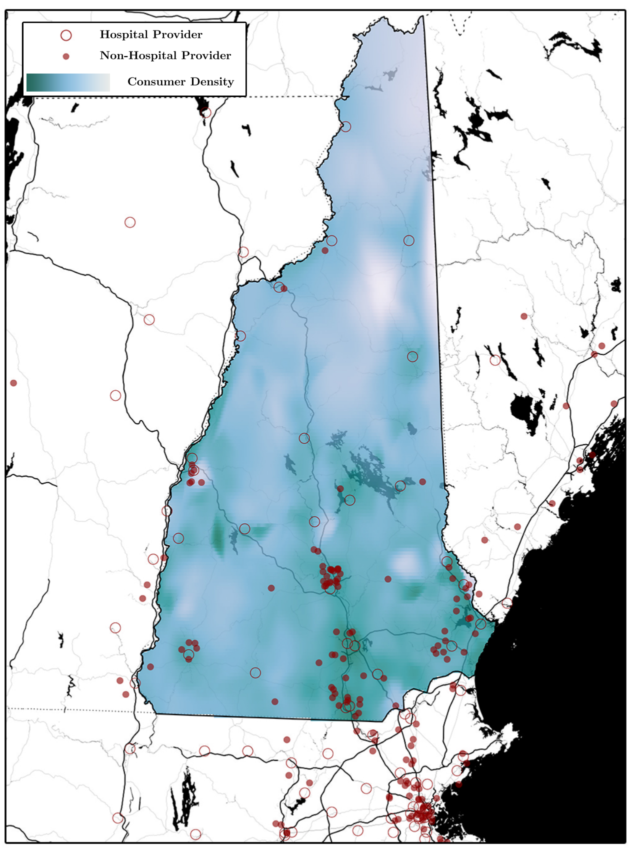

Notes: Map shows the location of providers providing medical imaging services that service privately-insured individuals in New Hampshire.

Classification of Diseases (ICD) codes, I also construct a measure of chronic diseases or conditions that may affect how difficult it is to treat patients. This measure is referred to as the Charlson Comorbidity Index.[21] Finally, I construct an indicator for whether each individual has the medical imaging procedure in the week following an emergency. Although these are relatively minor emergency visits since I exclude inpatient admissions, this may affect demand since it may be more time sensitive (e.g. demand for medical imaging procedures after a bone fracture may be different than for routine preventative care).

Table 2.1: Summary of Privately Insured Individuals with Medical Imaging Claims

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

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

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

Table 2.1 provides a summary of individuals in the sample. There are 177,995 unique individuals with outpatient imaging visits over the period. Half of the individuals are in HMO plans, and most of the remainder are in PPO or POS plans. About 44 percent of individuals have a plan with a deductible. New Hampshire is a relatively high-income state, and privately insured individuals have even higher income than the general population.

When an individual needs a specific procedure, the choice set is defined as the providers that are available through the individual's insurance plan that can perform the procedure in the given year. Although I do not observe each insurer's network directly, I construct a proxy by examining the providers chosen by individuals in each insurance company-product pair (e.g. Anthem HMO). In some cases, individuals may have plans, such as PPO plans, that allow them to choose providers out-of-network. To the extent that individuals actually choose these providers, they are included in the choice set (but have higher prices).[22] For each option in the choice set, I construct procedure prices that vary by insurance company-product pair and year. In addition, out-of-pocket prices vary across individuals with the same insurance product since some individuals are under the deductible and some are not. Within each individual's choice set, I remove providers that cannot perform the procedure as well as those that are more than 75 miles from the individual.

In general, patients are told they need a diagnostic test by their primary care physician or other specialist. They may receive a referral, however consumers are generally free to schedule an appointment for a medical imaging procedure at any provider within their insurer's network.[23] 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

[22]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.

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

care physician in each year, defined as the most frequently visited primary care physician. I then find the most common medical imaging provider chosen by the primary care physician's patients. Using this, I construct an indicator for likely referrals.

Table 2.2: Summary of Medical Imaging Visits by Insurer

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

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

The full dataset is summarized for each of the three insurers in Table 2.2. Anthem is by far the largest insurer, with over 200,000 medical imaging visits over the period. On average, the out-of-pocket price is 12 percent of the total negotiated price. However, there is large variation—individuals under the deductible pay the full price. In particular, there is greater cost sharing in the beginning of the year, when individuals have not hit their deductible, then at the end of the year (see Figure A4). Individuals choose between 13 different providers on average, although, again, there is significant variation. This is partially due to the fact that there are more providers that are capable of performing X-rays than MRI scans. Given large number of observations, I use a 2 percent sample of visits for the main analysis.

Within individual's choice sets there is a large degree of price dispersion, and consequently, significant potential savings if individuals switch to low cost options. Figure 2.2a shows the distribution of demeaned negotiated prices within individuals' choice sets. The distribution is approximately normal, with standard deviation of $639 (and coefficient of variation of 44.5 percent). If a consumer is under the deductible for the year, the individual is fully exposed to the variation in prices. However, since most patients share cost with an insurer, out-of-pocket price dispersion is smaller, with a standard deviation of $127 (see Figure 2.2b).[24] Finally, Figure 2.2c shows the distribution of prices paid by the insurer.

Given the variation in prices, there are large potential savings if consumers switched to cheaper providers in their network. The potential savings for consumers and insurers are summarized in Table A8. Overall, I find that there would be savings of over 40 percent if consumers switched to providers in the first quartile of the price distribution.[25] The savings are even greater for X-rays and CT scans. Consumers subject to a deductible have large private gains from switching, but much of the potential savings for consumers without a deductible go to insurers. This suggests that, although there are large potential savings for the health care system, these consumers may have little incentive to switch to less expensive providers even if they have price information.

2.2.2 HealthCost Website

In an effort to increase health care price transparency, the New Hampshire Insurance Department launched the HealthCost website in March 2007.[26] Although other states have implemented health care price transparency initiatives, many only provide information on the hospital list price of each procedure (i.e. charge amount), which has little

[24] The coefficient of variation is 49.6 percent.

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

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

Figure 2.2: Price Variation within Individuals' Choice Sets

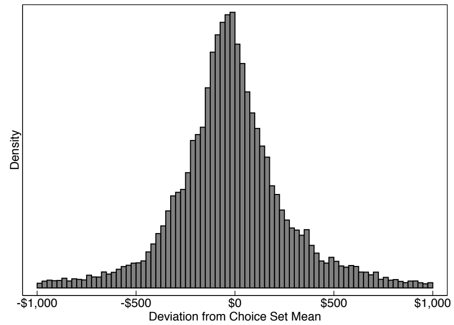

(a) Total Negotiated Price

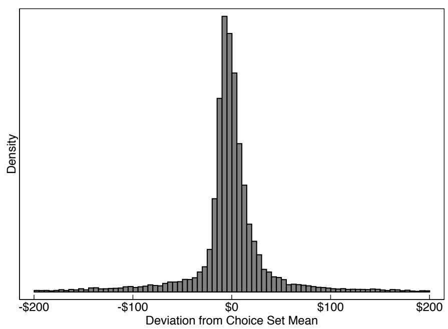

(b) Patient Out-of-Pocket Price

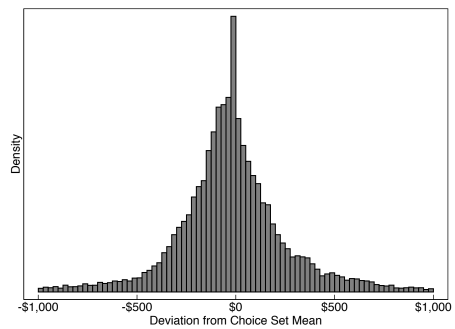

(c) Insurer Price

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

bearing on the out-of-pocket prices that insured individuals actually pay.[27] New Hampshire's HealthCost website was unique because it provided information about insurer-specific out-of-pocket prices. Although other states, such as Maine and Colorado, have since created tools with similar information, New Hampshire's price transparency efforts remain the most comprehensive.[28] Individuals with private insurance in the state can select one of about 35, mostly outpatient, procedures (see Figure A5a). In addition to providing information for insured individuals, the website also has a separate tool for uninsured individuals in the state. Since the claims data cover the population of insured individuals, I focus only on the former. In recent years, the website added information

[27] Information about list prices may affect uninsured individuals. See Christensen, Floyd and Maffett (2015), who examine the effect of information about list prices.

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

about provider quality and a guide to health insurance. This occurred after my period of analysis. It is also important to note that there have been other price transparency efforts by individual insurers, notably Aetna which started its Member Payment Estimator tool in 2010. However, Aetna had a very small presence in New Hampshire and is excluded from the analysis.

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 A5b shows an example of prices returned by the website. The table of prices is automatically sorted by out-of-pocket price, making it easy for consumers to schedule an appointment with the lowest cost provider. In addition to the out-of-pocket price, the website also returns the amount paid by insurers and the total negotiated price. For the purposes of analysis, I assume that individuals who use the website are fully informed about prices. I discuss this assumption in greater detail in Section 2.3.2.[29]

According to discussions with state employees, the website was promoted by encouraging insurers and primary care doctors to inform patients about the website. In addition, there were at least 40 news articles mentioning the website over the period. On average, there were 41,506 searches for price information per year according to website traffic logs, about a third of which were for medical imaging procedures. Furthermore, anecdotal evidence suggests that the website not only let consumers shop around, but may have allowed insurers to negotiate lower rates. One report noted that after the introduction of the website "the balance of plan-provider negotiating power began shifting significantly in New Hampshire."[30] In particular, Anthem, the largest insurer in New Hampshire, had a public battle with an expensive hospital in the state. Local news sources suggest that

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

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

the price transparency website allowed the insurer to negotiate lower prices.[31]

In order to examine the effect of price transparency, this paper exploits two sources of variation generated by the HealthCost website. First, there is variation due to the timing of the website introduction. In this way, I can examine procedures on the website and compare observed choices from 2005 to February 2007, prior to the introduction of the website, to observed choices in the period starting March 2007. Second, there is variation due to the fact that only a subset of medical imaging procedures were available on the website.[32] The X-ray, CT scan, and MRI scan procedures with and without information available on the website are listed in Table A7. I argue that imaging procedures on the website tend to be quite similar to procedures not on the website. For example, the price of a knee X-ray is available on the website while the price of a knee/leg CT scan is not. Note that the website also had price information for a few simple surgical procedures (e.g. kidney stone removal), physician office visits, as well as newborn delivery. I do not consider these procedures in the analysis because they tend to be less standardized and involves a different set of providers.

In Chapter 1, I use these two sources of variation in a difference-in-differences framework. In that paper, the key identifying assumption is that the price of procedures on the website would follow a common trend relative to procedures not on the website if the website were never available. I argued that this assumption was plausible given the price trends that exist prior to the introduction of the website. In this paper, I develop an empirical model that relies on an alternative, but related, set of assumptions. One of the key assumptions is that individuals' utility parameters are orthogonal to whether procedures are available on the website. I compare the approach taken in this paper with the reduced-form identification strategy and discuss the structural assumptions in more

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

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

Figure 2.3: Price Transparency Website Usage for Medical Imaging Procedures
By Month

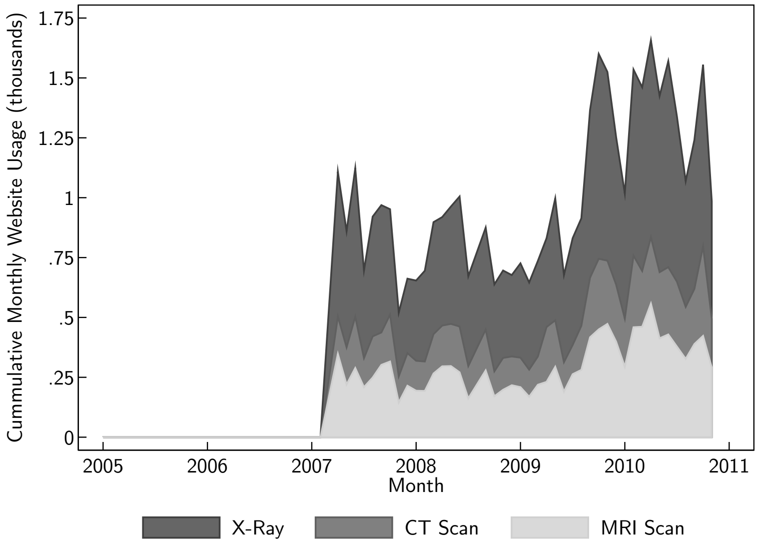

Notes: Chart shows cumulative searches by procedure group. Includes all searches using "Health Costs for Insured Patients" wizard on either nhhealthcost.nh.gov or nhhealthcost.org. Note the website began in March 2007.

detail in Section 3.4.1.

I use website traffic logs obtained from the New Hampshire Insurance Department to calculate the number of website price searches in each month for each procedure listed on the website. Website traffic data is available from March 2007 through November 2010, at which point the website switched hosting companies. Figure 2.3 shows cumulative monthly price searches for X-rays, CT scans, and MRI scans. When the website was first introduced in 2007 there were about 750 to 1,000 searches per month for the price of medical imaging procedures, however this grew to over 1,500 searches per month by late 2009.

In order to estimate the fraction of informed consumers I divide the number of price searches per procedure by the total number of visits in New Hampshire from the claims

Table 2.3: Monthly Percent of Consumers with Price Information
By Procedure Listed on Price Transparency Website

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

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

data.[33] In other words, I assume that each use of the website is a unique individual.[34]

Table 2.3 shows the estimated percent of consumers with price information for each medical imaging procedure listed on the website. The percent of informed consumers is between 2 and 6 percent on average for X-ray procedures. There is a larger fraction of consumers that use the website for CT scans and MRI scans—between 5 and 19 percent on average. CT scans and MRI scans also tend to be more expensive, making the website potentially more valuable for consumers receiving these procedures. There is also temporal variation, potentially due to the fact that there is random variation in the type

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

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

of individuals that need a procedure in a given month. In addition, more individuals may be learning about the website over time, as seen in Figure 2.3. This variation is used to help estimate the demand model and recover information about the choice to use the website if it is available.

2.3 Demand for Providers and Website Usage

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

I start backwards and begin by discussing the choice of provider with and without price information in Section 2.3.2. In Section 2.3.3 I discuss the model of website usage using results derived from Section 2.3.2. The two parts of the model are estimated jointly. I present an estimation strategy that relies on revealed preferences in terms of both website usage and provider choices. The high dimensionality of the unobservables makes maximum likelihood estimation computationally infeasible. In Section 2.3.4, I present a Bayesian estimation strategy that addresses the estimation challenges by utilizing recent advances in MCMC methods. Finally, I discuss identification.

2.3.1 Model Setup and Timing

There are a set of providers that sell medical imaging services $\mathcal{J}$ indexed by j. The set of providers includes hospitals as well as non-hospital providers (i.e. freestanding outpatient facilities such as imaging centers and clinics). Each year, insurer $k \in \mathcal{K}$ contracts with a subset of providers, $\mathcal{N}_{kmt} \subseteq \mathcal{J}$, that can perform procedure $m \in \mathcal{M}$, where $\mathcal{M}$ is the set of medical imaging procedures.[35] 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 3.2, 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).[36]

Individual i pays fraction c_{ikmt} of the negotiated price, which is observed in the claims data. The degree of cost sharing is determined by both the coinsurance rate applied to procedure m when enrolled in insurance plan k as well whether the individual is past the deductible for the year. In particular, if the individual is subject to a deductible then $c_{ikmt} = 1$. Therefore, for a given individual, cost sharing can vary over time t. The out-of-pocket price paid by the individual is

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

I assume that this is the price internalized by the individual.[37] The remainder is paid by

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

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

[37] It has been suggested that individuals respond to dynamic incentives that arise due to annual deductibles, however the evidence is mixed (e.g Aron-Dine et al. 2015; Sacks et al. 2016; Brot-Goldberg et al.

the insurer

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

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

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

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

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

2015). I assume that individuals do not anticipate whether they will surpass their annual deductible and respond only to the spot price.

2.3.2 Choice of Provider

In this section I present a discrete choice model of provider demand in which consumers receive noisy signals about prices. If consumers do not become informed about prices, they choose a provider with uncertainty about prices. However, if consumers use the price transparency website, they know true prices. Previous models of hospital demand either assume that individuals do not account for hospital prices at all or have perfect information about prices.[38] In contrast, I assume that individuals may have some information about prices even if they do not use the website. In this way, the model nests both the full information case as well as the case in which individuals completely ignore prices. In addition to price, the choice of provider is also assumed to depend on the distance from each individual to each provider, referrals, provider quality or amenities, as well as factors that vary with observed differences across individuals.

Individuals may only visit a provider in their network, $j \in \mathcal{N}_{kmt}$. There is no outside option since individuals are assumed to receive a medical imaging procedure if their doctor recommends it. One concern is that price transparency affects the choice to have a procedure at all. In Chapter 1, I use the entire sample of privately-insured individuals and examine the effect of the price transparency website on the probability of having medical imaging procedures and do not find a statistically significant effect.[39] This finding suggests that conditioning on individuals that had a medical imaging procedure and assuming they all choose an inside option is unlikely to bias counterfactual estimates.

[38]For example, Kessler and McClellan (2000), Tay (2003), Ho (2006), and Ho and Lee (2017) assume that price does not influence patient choice while Capps, Dranove and Satterthwaite (2003), Gaynor and Vogt (2003), Ho and Pakes (2014) and Gowrisankaran, Nevo and Town (2015) include price in individual utility and assume individuals have perfect information.

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

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

$$u_{ijkmt} = -\gamma_i p_{ijkmt}^{OOP} + \underbrace{\alpha_1 d_{ij} + \alpha_2 d_{ij}^2 + \alpha_3 r_{ijt} + \xi_{jM} + \beta \mathbf{x}_{ikmt} h_j}_{\delta_{ijkmt}} + \varepsilon_{ijkmt} \tag{2.1}$$

I allow for individual-specific heterogeneity in out-of-pocket price sensitivity, γ_i, which is distributed with density $f(\gamma_i)$. This approach has the benefit of not exhibiting the independence from irrelevant alternatives property and allowing for more flexible substitution patterns. It is also important since individuals that are more price sensitive may be more likely to use the price transparency website, which I explicitly account for in Section 2.3.3. 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.[41] Accounting for the adverse selection into insurance is important for understanding which individuals benefit from the price transparency website. In particular, I assume that the random coefficient is distributed normally:

$$\gamma_i \sim N(\bar{\gamma} + \rho c_{ik}, (\sigma^{\gamma})^2) \tag{2.2}$$

Since the same individual may have multiple medical imaging visits over the period, it is important to account for correlation in unobserved utility. Therefore, I assume that the

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

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

random-coefficient is individual-specific (Revelt and Train 1998).

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

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

$$s_{ijkmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt}|\vartheta_{ikmt} = 1) = \int_{\gamma_i} \frac{\exp(-\gamma_i p_{ijkmt}^{OOP} + \delta_{ijkmt})}{\sum_{j' \in \mathcal{N}_{kmt}} \exp(-\gamma_i p_{ij'kmt}^{OOP} + \delta_{ij'kmt})} f(\gamma_i) d\gamma_i \qquad (2.3)$$

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

The expected consumer surplus, conditional on having price information, for a patient needing a medical imaging procedure is then:[42]

$$CS_{ikmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt} | \vartheta_{ikmt} = 1) = \frac{1}{\gamma_i} \log \left(\sum_{j \in \mathcal{N}_{kmt}} \exp(-\gamma_i p^{OOP}_{ijkmt} + \delta_{ijkmt}) \right) \tag{2.4}$$

Provider Choice with Price Uncertainty

Next, I model the case in which individuals have uncertainty about prices. Individuals form noisy beliefs using Bayes' rule and then make a decision based on those beliefs. This information structure is related to the empirical work on consumer learning (e.g. Erdem and Keane 1996; Ackerberg 2003; Erdem, Keane and Sun 2008; Crawford and Shum 2005; Ching 2010; Dickstein 2014; Grennan and Town 2015). I assume that individuals know the distribution from which prices are drawn, which is assumed to be normal.[43] In particular, their prior is determined by the true mean and variance of prices their choice set, $\bar{p}^{OOP}_{kmt}$ and $\bar{s}^2_{kmt}$ respectively:

$$p^{OOP}_{ijkmt} \overset{iid}{\sim} N(\bar{p}^{OOP}_{kmt}, \bar{s}^2_{kmt}) \tag{2.5}$$

The prior provides no information about relative prices in the choice set, and therefore is not useful for choosing a provider on its own. However, individuals may be able to obtain additional information about individual prices. For instance, they may be able to look up list prices or receive potentially noisy price information from other individuals that had similar procedures. When asked, providers and insurers sometimes provide a price range if they provide any price information at all.[44] I model this by assuming that individuals receive a vector of unbiased signals, where each signal is given by

$$p^{OOP}_{ijkmt} + e_{ijkmt} \tag{2.6}$$

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

[43]The true distribution of prices is approximately normal. See Figure 2.2b.

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

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

$$e_{ijkmt} \overset{iid}{\sim} N(0, \sigma_h^2) \tag{2.7}$$

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

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

$$\mathbb{E}\left[\widetilde{p^{OOP}_{ijkmt}}\right] = w_{ikmt}(p^{OOP}_{ijkmt} + e_{ijkmt}) + (1 - w_{ikmt})\bar{p}^{OOP}_{kmt} \tag{2.8}$$

where the weight given to the signal is defined as

$$w_{ikmt} = \frac{\bar{s}^2_{kmt}}{\bar{s}^2_{kmt} + \sigma_h^2} \tag{2.9}$$

If $\sigma_h^2 = 0$ then $w_{ikmt} = 1$ and individuals know true prices. Conversely, if $\sigma_h^2 \to \infty$ then $w_{ikmt} \to 0$, implying that individuals place no weight on the price signals. In this way, the prior is important because it disciplines individual's beliefs about price—if individuals receive very noisy signals than they effectively ignore prices.[45] In Section **??**, I present an alternative model in which individuals have an uninformative prior and take price signals as given.

Using the assumption that the prior and signal are normally distributed, the variance

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

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

of posterior beliefs is

$$Var\left[\widetilde{p_{ijkmt}^{OOP}}\right] = w_{ikmt}\sigma_h^2 \tag{2.10}$$

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

$$\begin{aligned}\mathbb{E}\left[\widetilde{u_{ijkmt}}\right] &= -\gamma_i\mathbb{E}\left[\widetilde{p_{ijkmt}^{OOP}}\right] + \underbrace{\alpha_1 d_{ij} + \alpha_2 d_{ij}^2 + \alpha_3 r_{ijt} + \xi_{jM} + \beta\mathbf{x}_{ikmt}h_j}_{\delta_{ijkmt}} + \varepsilon_{ijkmt} \\ &= -\gamma_i w_{ikmt}(p_{ijkmt}^{OOP} + e_{ijkmt}) + \delta_{ijkmt} + \varepsilon_{ijkmt}\end{aligned} \tag{2.11}$$

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

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

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

$$\begin{aligned}&s_{ijkmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt}|\vartheta_{ikmt} = 0) = \\ &\qquad \int_{\gamma_i}\int_{\mathbf{e}_{ikmt}} \frac{\exp(-\gamma_i w_{ikmt}(p_{ijkmt}^{OOP} + e_{ijkmt}) + \delta_{ijkmt})}{\sum_{j'\in\mathcal{N}_{kmt}} \exp(-\gamma_i w_{ikmt}(p_{ij'kmt}^{OOP} + e_{ij'kmt}) + \delta_{ij'kmt})} f(\mathbf{e}_{ikmt}) f(\gamma_i) d^J\mathbf{e}_{ikmt} d\gamma_i\end{aligned} \tag{2.12}$$

where $\vartheta_{ikmt} = 0$ indicates that the individual did not use the website and is uninformed about prices. It is worth noting that the vector of signal noise, $\mathbf{e}_{ikmt}$, has the same number of elements as $\mathcal{N}_{kmt}$. Therefore, computing the expectation over individual beliefs requires evaluating a potentially high dimensional integral, complicating the estimation strategy. I address this issue in Section 2.3.4.

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

$$\begin{aligned} CS_{ikmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt} | \vartheta_{ikmt} = 0) = & \underbrace{\frac{1}{\gamma_i} \log \left(\sum_{j \in \mathcal{N}_{kmt}} \exp(-\gamma_i w_{ikmt}(p_{ijkmt}^{OOP} + e_{ijkmt}) + \delta_{ijkmt}) \right)}_{\text{CS evaluated at beliefs of prices}} \\ & + \underbrace{\sum_{j \in \mathcal{N}_{kmt}} \left[\mathbb{E}\left[\widetilde{p_{ijkmt}^{OOP}} \right] - p_{ijkmt}^{OOP} \right] s_{ijkmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt} | \vartheta_{ikmt} = 0)}_{\text{CS gain/loss from incorrect beliefs}} \end{aligned} \tag{2.13}$$

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

$$\mathbb{E}\left[\widetilde{p_{ijkmt}^{OOP}}\right] - p_{ijkmt}^{OOP} = w_{ikmt} e_{ijkmt} + (1 - w_{ikmt})(\bar{p}_{kmt}^{OOP} - p_{ijkmt}^{OOP}) \tag{2.14}$$

In general, individuals are more likely to choose a provider they falsely believe to be inexpensive, creating a situation similar to a "winner's curse". This can be seen in Figure 2.4 which presents two situations, one in which expected price is greater than actual price and one in which expected price is less than actual price. Believing an option to be inexpensive (i.e. receiving a low e_{ijkmt}) results in a higher choice probability, increasing the expected loss from incorrect beliefs.

2.3.3 Website Usage

In this section I develop a model in which individuals choose to use the price transparency website if it is available. The model seeks to recover information about which

Figure 2.4: Consumer Surplus when Expected Price Differs from Actual Price

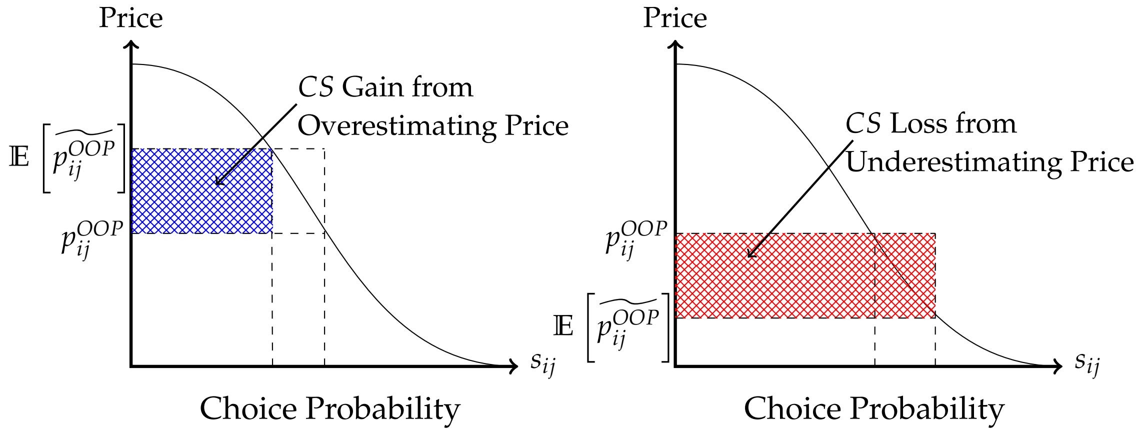

Notes: Blue shaded region shows the gain in consumer surplus relative to expected consumer surplus due to price being less than expected. Red region shows the loss in consumer surplus from price being more than expected. Note that there is a "winner's curse" and the expected loss is larger.

individuals use the price transparency website. In Section 3.5, the estimates from this selection model are used to simulate website usage under counterfactual scenarios such as increased cost sharing.

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

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

$$p_{ijkmt}^{OOP} \overset{iid}{\sim} N\left(w_{ikmt}(p_{ijkmt}^{OOP} + e_{ijkmt}) + (1 - w_{ikmt})\bar{p}_{kmt}^{OOP}\, ,\, \sigma_h^2 w_{ikmt}\right) \tag{2.15}$$

With price uncertainty, the ex-ante consumer surplus from the individual's perspective is determined by evaluating Equation 2.13 at all possible prices.[46] Therefore, given individuals' beliefs, the expected consumer surplus is

$$\frac{1}{\gamma_i} \log \left(\sum_{j \in \mathcal{N}_{kmt}} \exp(-\gamma_i \mathbb{E}\left[\widetilde{p_{ijkmt}^{OOP}}\right] + \delta_{ijkmt}) \right) + \int_{\widetilde{\mathbf{p}_{ikmt}^{OOP}}'} \sum_{j \in \mathcal{N}_{kmt}} \left(\mathbb{E}\left[\widetilde{p_{ijkmt}^{OOP}}\right] - \widetilde{p_{ijkmt}^{OOP}}' \right) s_{ijkmt} g(\widetilde{\mathbf{p}_{ikmt}^{OOP}}') d^J \widetilde{\mathbf{p}_{ikmt}^{OOP}}' \quad (2.16)$$

where $g(\widetilde{\mathbf{p}_{ikmt}^{OOP}}')$ is the joint distribution of beliefs determined by the individual's prior and signals following Equation 2.15.

In order to evaluate the expected gain from using the website, the individual must compare Equation 2.16 with the expected consumer surplus after using the website. If individuals use the website, they can re-optimize. In addition, they will no longer be surprised by the bill. Therefore, given individual's beliefs, expected consumer surplus with price information is given by:

$$\frac{1}{\gamma_i} \int_{\widetilde{\mathbf{p}_{ikmt}^{OOP}}} \log \left(\sum_{j \in \mathcal{N}_{kmt}} \exp(-\gamma_i \widetilde{p_{ijkmt}^{OOP}} + \delta_{ijkmt}) \right) g(\widetilde{\mathbf{p}_{ikmt}^{OOP}}) d^J \widetilde{\mathbf{p}_{ikmt}^{OOP}} \quad (2.17)$$

The difference between Equation 2.17 and Equation 2.16 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.[47] See Appendix 3.6 for derivation and discussion about the accuracy of this approach. Using

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

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

this approximation, the expected benefit of using the website, b_{ikmt}, is

$$b_{ikmt} \approx \frac{\gamma_i w_{ikmt} \sum_{j \in \mathcal{N}_{kmt}} \sigma_h^2 \left[\exp(-\gamma_i \mathbb{E}\left[\widetilde{p_{ijkmt}^{OOP}}\right] + \delta_{ijkmt}) \Phi_{ijkmt} \right]}{2 \left[\sum_{j \in \mathcal{N}_{kmt}} \exp(-\gamma_i \mathbb{E}\left[\widetilde{p_{ijkmt}^{OOP}}\right] + \delta_{ijkmt}) \right]^2} \tag{2.18}$$

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

Unlike Equation 2.17, the interpretation of the closed-form expression above is relatively straightforward. Holding beliefs about prices fixed, an increase in price uncertainty, as measured by σ_h^2, increases the value of using the website. Similarly, an increase in price dispersion affects w_{ikmt}, also increasing the value of using the website. Note that the benefit of using the website is increasing in the absolute value 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.[48] In addition, many individuals were likely unaware of the website and had to be motivated enough to discover the website on their own.

I assume cost has both an observable component, which is a function of individual characteristics $\mathbf{x}_{ikmt}$, as well as an unobservable component, ν_{ikmt}. Observable characteristics include age categories, sex, income, eduction, Charlson Comorbidity Index, emergency indicator, and year indicators in order to account for the fact that more individuals may hear about the website over time, reducing the implicit cost.[49] I also include a constant.

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

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

Individuals use the website if the net benefit is positive

$$\underbrace{\theta b_{ikmt}}_{\text{Website Benefit}} - \underbrace{\phi \mathbf{x}_{ikmt} + \nu_{ikmt}}_{\text{Website Cost}} > 0 \tag{2.19}$$

I assume that the distribution of ν_{ikmt} is distributed i.i.d. type 1 extreme value (with normalized variance). Therefore, the observed probability that individual i uses the website for the price of procedure m at time t takes the logistic form:

$$\vartheta_{ikmt} = \frac{\exp(\theta b_{ikmt} - \phi \mathbf{x}_{ikmt})}{1 + \exp(\theta b_{ikmt} - \phi \mathbf{x}_{ikmt})} \tag{2.20}$$

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

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

$$\vartheta_{mt} = \frac{1}{n_{mt}} \sum_{i \in \mathcal{I}_{mt}} \vartheta_{ikmt} \tag{2.21}$$

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

2.3.4 Joint Estimation of Demand

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

Likelihood Function

The likelihood function is directly based on the structural equations describing individual provider choices and website usage. The first component of the likelihood function is the probability of choosing the provider that was actually chosen. Equation 2.3 and Equation 2.12 are the conditional choice probabilities with and without price information. Therefore, the unconditional choice probability is:

$$s_{ijkmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt}|\vartheta_{ikmt}) = \begin{cases} s_{ijkmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt}|\vartheta_{ikmt} = 0) & \text{if website is not available} \\ \\ \vartheta_{ikmt} \cdot s_{ijkmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt}|\vartheta_{ikmt} = 1) & \text{if website is available} \\ \quad +(1-\vartheta_{ikmt}) \cdot s_{ijkmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt}|\vartheta_{ikmt} = 0) & \end{cases} \tag{2.22}$$

If the website is not available for procedure m at time t, either because it is prior to March 2007 or because the procedure is never on the website, then the consumer has uncertainty about prices and choice probabilities are given by Equation 2.12. If the website is available for procedure m at time t, the consumer is informed about prices if the website is actually used. Therefore, choice probabilities are given by a mixture between Equation 2.3 and Equation 2.12, where the mixture weights are determined by the predicted probability of using the website, given by Equation 2.20.

The second component of the likelihood function is the probability of actual website traffic for each procedure-month given predicted website usage. The likelihood of website usage for procedure m in month t takes the following binomial form[50]

$$\frac{n_{mt}!}{\mathcal{V}_{mt}!(n_{mt}-\mathcal{V}_{mt})!}\left(\vartheta_{mt}\right)^{\mathcal{V}_{mt}}\left(1-\vartheta_{mt}\right)^{n_{mt}-\mathcal{V}_{mt}} \tag{2.23}$$

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

Therefore, the likelihood function is

$$\begin{aligned}\mathcal{L}(\Theta) = & \underbrace{\sum_t \sum_{k\in\mathcal{K}} \sum_{m\in\mathcal{M}} \sum_{i\in\mathcal{I}_{kmt}} \sum_{j\in\mathcal{N}_{kmt}} \left[s_{ijkmt}(\mathcal{N}_{kmt},\mathbf{p}_{kmt}|\vartheta_{ikmt})\right]^{y_{ijkmt}}}_{\text{Likelihood of Observed Provider Choices}} \\ & \times \underbrace{\sum_t \sum_{m\in\mathcal{M}} \frac{n_{mt}!}{\mathcal{V}_{mt}!(n_{mt}-\mathcal{V}_{mt})!}\left(\vartheta_{mt}\right)^{\mathcal{V}_{mt}}\left(1-\vartheta_{mt}\right)^{n_{mt}-\mathcal{V}_{mt}}}_{\text{Likelihood of Observed Search Traffic}}\end{aligned} \tag{2.24}$$

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

MCMC Estimation

To estimate the model, I use a Markov chain Monte Carlo (MCMC) estimator to simulate the posterior distribution of Θ. This approach helps circumvent the computational curse of dimensionality caused by the fact that beliefs, specifically the signal noise draws, are a high-dimensional nuisance parameter.[51] Rather than compute high dimensional integrals in order to find the expectation over $\mathbf{e}_{ikmt}$ and calculate the likelihood, the MCMC estimator samples the parameter space conditional on the data.

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

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

I take advantage of recent advances in Bayesian estimation and use a variant of MCMC known as Hamiltonian Monte Carlo (HMC) No-U-Turn Sampler (NUTS).[52] This approach, developed by Hoffman and Gelman (2014), uses the gradient of the log posterior density to more efficiently sample the posterior distribution.[53] Relative to standard MCMC algorithms such as Metropolis-Hastings and Gibbs sampling, this approach is known to converge significantly faster for high-dimensional problems, making it well suited for a situation with alternative-specific unobservables. In addition, it does not necessitate the use of conjugate priors, allowing for more flexible modeling assumptions.

For the purposes of estimation, I reformulate the model in terms of a simplified likelihood that is augmented with a set of priors. The likelihood uses the choice probability conditional on unobservables while the priors describe the distribution of website traffic and the distribution of the unobservables. This version of the model is described in more detail in Appendix 3.6.

It is important to note that using this Bayesian hierarchical model for estimation does not impose additional assumptions since I use uninformative priors for all of the structural parameters. In supplemental material, I examine a simplified version of the model with a small choice set and show that the results obtained via simulated maximum likelihood are very similar to those obtained via MCMC estimation.[54] The use of MCMC is primarily motivated by the fact that it is computationally attractive.

In order to estimate the posterior distribution of Θ, the algorithm uses the following approach. At iteration n, the MCMC algorithm returns parameter estimates $\Theta^{(n)}$. As starting values, I use parameter estimates from a standard multinomial logit (see Section 2.4.1).[55] However, to ensure that initial values do not influence the resulting posterior

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

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

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

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

distribution, samples drawn during a warm-up period are discarded. The remaining collection of samples, $(\Theta^{(1)}, ..., \Theta^{(N)})$, approximately converge to the distribution of the posterior. I report the mean and standard deviation of these samples in the results.

Identification Intuition

Without variation in consumers' information set, it is difficult or impossible to separately identify price sensitivity and the degree of price uncertainty, i.e. the observed choices from a population with low price sensitivity are potentially observationally equivalent to the observed choices from a population with high price sensitivity but limited information about prices. An alternative way to see this is to note that price uncertainty is closely related to classical measurement error, except rather than the researcher having noisy information it is the individual decision-maker. It is widely known that classical measurement error can cause biased and inconsistent parameter estimates. This is also true in non-linear models, and, in general, the variance of measurement error and the underlying parameters are not separately identified (e.g. Chen, Hong and Nekipelov 2011). I overcome this issue due to the fact that some consumers—those that use the price transparency website—do not observe prices with error.

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

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

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

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

2.4 RESULTS

I now turn to the results. Section 2.4.1 motivates the full demand model by presenting results from a standard multinomial logit model. Section 2.4.2 then discuss the results from the full demand model (both provider choice and website usage). In Section 2.4.3, I also discuss results from an alternative demand model in which consumers do not have an informative prior about the distributions of prices.

2.4.1 Estimates from Multinomial Logit Model

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

$$s_{ijkmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt}) = \frac{\exp(-\gamma_1 q_{mt} p^{OOP}_{ijkmt} - \gamma_2(1 - q_{mt}) p^{OOP}_{ijkmt} + \delta_{ijkmt})}{\sum_{j' \in \mathcal{N}_{kmt}} \exp(-\gamma_1 q_{mt} p^{OOP}_{ij'kmt} - \gamma_2(1 - q_{mt}) p^{OOP}_{ijkmt} + \delta_{ij'kmt})} \tag{2.25}$$

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 2.1. Since this simple approach does not model unobserved beliefs, estimation can be performed via maximum likelihood.

Table A9 presents the coefficient estimates and standard errors from the simple logit model. The magnitude of the price coefficient is larger when consumers have access to the price transparency website, indicating that the website increases the effective demand elasticity of the population. The difference is statistically significant at the 10 percent level.[56] Complimenting the results in Chapter 1, this provides further evidence that the

[56] I present results for a small sample for comparison with the results from the full specification. The difference is statistically significant at the 5 percent level for a larger sample. In addition, although the

website had a meaningful impact on consumer behavior.

Consistent with the prior literature on hospital demand, the travel distance is important for understanding consumer demand. I find evidence of non-linearities in travel preferences. The indicator for likely referrals is highly significant, indicating that physician can influence consumer behavior.[57] Finally, there is evidence that income and education affect whether consumers choose a hospital. The results suggest that higher income consumers are more likely to prefer hospitals, which tend to be expensive, rather than medical imaging centers.

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

2.4.2 Provider Choice and Website Usage Estimates from Baseline Model

Table 2.4 presents estimates for parameters of the full demand model. I focus on specification 1, which reflects the baseline model presented in Section 2.3. The first column reports the mean of the estimated posterior distribution of each parameter implied by the

coefficient on price is not statisically significant when the website is not available in Table A9, it becomes significant with a larger sample.

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

Table 2.4: MCMC Estimates for Demand Model

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

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

Figure 2.5: Example of Beliefs for Uninformed Individuals

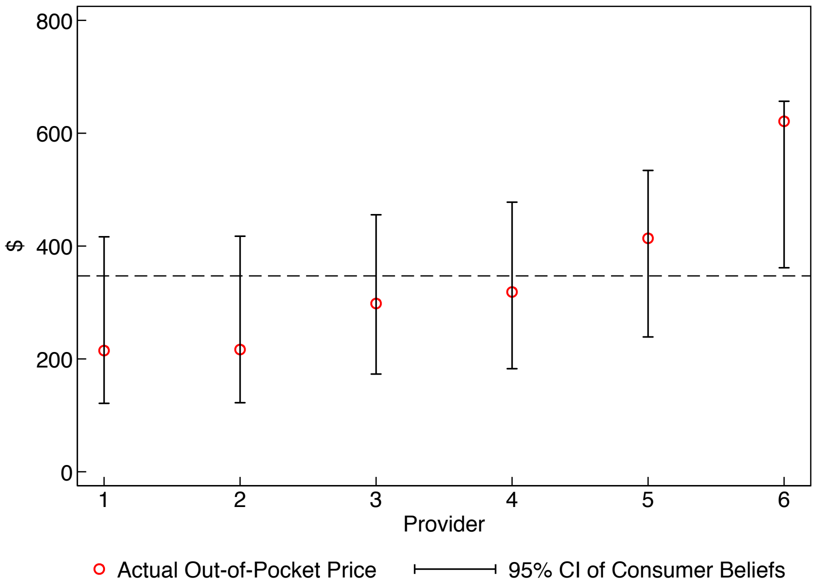

Notes: Chart shows an example choice situation selected from the data in which an individual is choosing between 6 providers. Confidence interval shows the distribution of beliefs given the estimated price uncertainty from specification 1 in Table 2.4.

MCMC estimation procedure. The second column reports the standard deviation of the posterior distribution.[58]

The magnitude of mean price sensitivity, $\bar{\gamma}$, is much larger than the price coefficient in the simple logit model presented in the previous section. This reflects the fact that γ_i can now be interpreted as consumer's underlying price sensitivity when prices are known (i.e. underlying marginal utility of income). There is significant heterogeneity in the price sensitivity parameter, which is negatively correlated with consumer cost sharing. This implies that consumers with high price sensitivity select into generous insurance plans (i.e. those with lower cost sharing). The effect of other explanatory variables, including distance, the likely referral indicator, and hospital interactions, is largely consistent with the results from the logit model presented in Table A9.

The estimated standard deviation of signal noise, σ_h, is shown at the bottom of Table

[58]Note that for explanatory variables that overlap with the simple logit model, the standard deviation of the parameter posterior distributions are very similar to the standard errors reported in Table A9.

2.4. The estimates imply that, in the absence of price information, individuals have a large degree of uncertainty about prices. The standard deviation of signal noise is larger for non-hospital providers, suggesting that individuals have greater uncertainty about the price of medical imaging centers and other non-hospital providers.[59]

Given that the interpretation of σ_h is complex, it is useful to consider an example from the data. Figure 2.5 shows a sample individual choosing between six providers that range in price from about \$200 to \$650. Using the estimate of σ_h, the individual's beliefs about the price of each option can be simulated given different potential draws from the distribution of signal noise. The 95 percent confidence interval for these beliefs is shown for each option in the choice set. Beliefs range by over \$200, implying that there is a non-trivial chance that the individual will believe the expensive options (such as option 5), are actually the least expensive. Also note that mean beliefs are not equal to the true price. This is due to the fact that each individual's prior causes shrinkage towards the mean.

I compare uninformed consumers beliefs about prices with the true price. On average, there is a 37 percent absolute difference between beliefs and true prices. The gap is even larger for individuals under the deductible—48 percent. Noisy beliefs about price effectively make residual demand more inelastic. The implied price elasticity of demand evaluated at mean prices when consumers are uninformed is only -0.08, however it would be -0.47 if they were fully informed.[60]

Turning to the website choice parameters, the coefficient on the monetary benefit of using the website, θ, is positive. Furthermore, zero lies outside the 95 percent credible interval. This implies that consumers are more likely to use the website if the potential benefit is large, either because of the potential savings or individual-specific price sensitive. The coefficients on explanatory variables that make up the observable part of the cost of using the website tend to be imprecisely estimated. There is suggestive evidence

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

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

that higher income consumers have a larger cost, reflecting the higher opportunity cost of time. At the same time, more educated individuals have a lower cost of using the website, perhaps because they are more likely to be proficient internet users. Patients receiving a procedure after an emergency episode have a higher cost. Furthermore, there is lower estimated cost of using the website in 2010, the omitted year. This may reflect the fact that the website became better known over time. Overall, the cost of using the website is estimated to be \$63 on average (see Figure A6). Note that the magnitude of θ is relatively small, indicating that that ν_{ikmt} is important for understanding website usage. In other words, there are unobserved factors, such as word-of-mouth or internet proficiency, that determine website usage.

2.4.3 Estimates under Alternative Demand Assumptions

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

$$p_{ijkmt}^{OOP} \overset{iid}{\sim} N\left(p_{ijkmt}^{OOP} + e_{ijkmt}\,,\, \sigma_h^2\right) \tag{2.26}$$

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

$$\mathbb{E}\left[\widetilde{u_{ijkmt}}\right] = -\gamma_i(p_{ijkmt}^{OOP} + e_{ijkmt}) + \delta_{ijkmt} + \varepsilon_{ijkmt} \tag{2.27}$$

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

Specification 2 in Table 2.4 presents the results from this alternative model. The estimates that characterize provider choices are broadly consistent with estimates from specification 1. The estimated standard deviation of the price signal is smaller for both hospital and non-hospital providers. This is due to the fact that noisy price signals generate more extreme beliefs than in the baseline model since beliefs are not disciplined by a prior.

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

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

2.5 Demand-Side Out-of-Sample Counterfactuals

I now use the estimates from the demand model to examine counterfactual policy simulations. In particular, I simulate the demand-side impact when more consumers are informed about prices holding insurance plan attributes fixed. The results imply decreasing returns as more consumers are informed about prices. However, a full examination of the relevant counterfactuals must take into equilibrium effects. This supply-side effect is the focus of Chapter 3.

Only about 8 percent of consumers used the price transparency tool when it was available, implying that there is a large cost of using the website. Much of this cost is likely non-pecuniary, i.e. individuals may not have even known that the website existed. Interventions that reduce this implicit cost, such as advertising the website or even subsidizing usage, would increase the fraction of informed consumers.[61]

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

Figure 2.6: Demand-Side Effect of Price Transparency on Spending
By Fraction on Individuals with Price Information

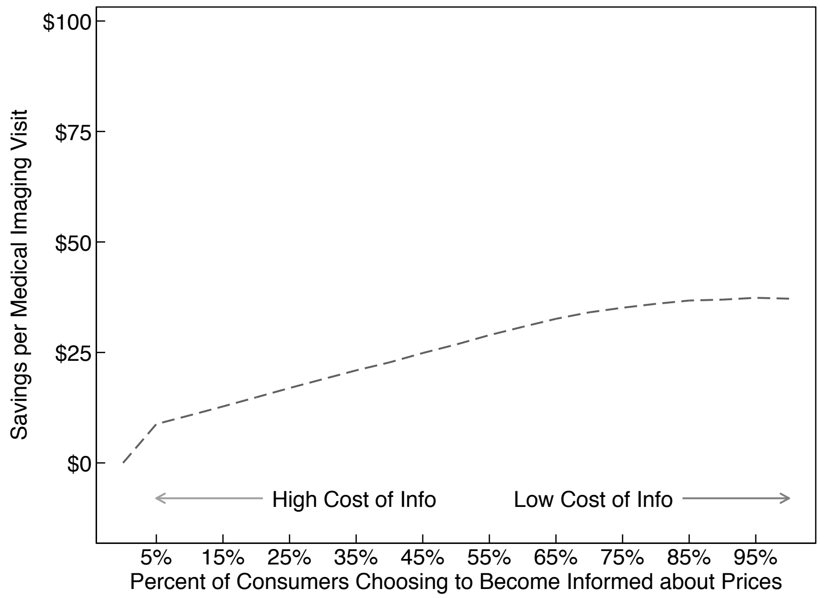

Notes: Demand-side effect holds prices fixed at distribution simulated with no price transparency. All figures in 2010 dollars.

Figure 2.6 shows the effect of reducing the implicit cost of becoming informed about prices, increasing the fraction of consumers with price information. 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. 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.

The evidence that consumer decisions are significantly affected by price transparency motivates a model examining the firm-response to increasing consumer price transparency. This is the focus of the last chapter.

2.6 Conclusion

I develop a demand model for health care services that separates underlying consumer price sensitivity and the degree of uncertainty about prices in order to explore how consumer choices would change if more consumers were informed about prices. I also show how to examine the welfare implications of information in this context.

This model of demand and price transparency has implications for other markets in which it is not possible to search for prices, such as markets in which consumers receive potentially noisy price quotes (e.g. automotive repair, building contractors, and law services). Even if these quotes are correct on average, the lack of price transparency makes demand less elastic. A similar mechanism may be present in other markets with shrouded add-on prices, surcharges, or delivery fees.

The fact that price transparency affects the elasticity of demand suggests that bargaining between health insurers and providers may be affected by increased price transparency. Indeed, this is one of the main results from Chapter 1. In Chapter 3, I use the results from the empirical model of demand to estimate a model of the supply-side that can then be used to examine the equilibrium welfare implications of increased price transparency.

Chapter 3

An Empirical Model of Price Transparency and Bargaining

3.1 Introduction

The relatively high level of spending on privately-provided health care in the U.S. is often attributed to high prices.[1] In addition, a recent literature has documented the large degree of price dispersion in health care, even for relatively standardized procedures (Cooper et al. 2015). This suggests that is important to understand the determinants of health care prices, and interventions that can allow insurers to negotiate lower prices with providers.

This paper empirically evaluates how price transparency affects markups and welfare in the U.S. health care market. I combine the model of demand that incorporates price uncertainty developed in Chapter 2 with a model of bargaining between providers and insurers. The model allows for an equilibrium analysis of out-of-sample counterfactual scenarios. I find that there would be a considerable reduction in health care prices if all consumers were informed about prices. One way to increase the number of informed consumers is to combine price transparency tools with high cost sharing health plans, incentivizing consumers to become informed about prices. This combination of policies would also lead to a large reduction in equilibrium prices.

In particular, I present a bargaining model to recover information about marginal cost and examine how price transparency affects negotiated prices in equilibrium. Recent empirical work has used models of bilateral bargaining between insurers and medical providers to gain insight into the effects of hospital and insurer competition (Gowrisankaran, Nevo and Town 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.[2] I use the first order condition of the bar-

[1]See, for example, Anderson et al. (2003), Koechlin, Lorenzoni and Schreyer (2010), and Cooper et al. (2015).

[2]For a discussion about how price transparency could affect markups see “Health Care Price Transparency: Can It Promote High-Value Care?”, Commonwealth Fund, April/May 2012. Also see Section 2.2.

gaining equation to derive an expression for equilibrium prices and find that price transparency leads to a trade-off. First, price transparency can make residual demand more elastic, decreasing the incentive for providers to negotiate high prices. Second, price transparency ensures that consumers do not choose high cost providers, implying that insurers may be more willing to have high cost providers in their network. This can actually reduce the incentive of insurers to negotiate low prices. Therefore, the effect of price transparency on negotiated prices is theoretically ambiguous. Using the bargaining model, I derive a moment condition that allows me to recover marginal cost and investigate the effect of price transparency empirically.

In my first empirical exercise, I use the estimates from both the demand and supply model to evaluate the effect of New Hampshire's price transparency website. I find that the website resulted in overall savings of 3 percent.[3] This intent-to-treat effect is consistent with my reduced-form results in Chapter 1. In particular, I found that the website reduced overall spending by about 4 percent using a difference-in-differences methodology. These savings are primary due to increased price-shopping on the part of consumers, however part of the decline is also due to a small reduction in the equilibrium prices.

Even though New Hampshire's price transparency website was publicly available to all individuals in the state, a relatively small fraction of consumers in the market actually used it. Overall, I find that consumers used the website for about 8 percent of medical imaging visits when the website was available.[4] Unlike the reduced-form analysis, the empirical model allows me to examine the effect conditional on using the price transparency website.[5] Estimates imply that the website primarily benefited individuals most exposed to the full price, i.e. those subject to a deductible. Individuals with a deductible that used the website saved \$178 per visit on average, while individuals without a de-

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

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

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

ductible saved \$16. It is important to note, however, that price information may cause individuals to switch, for example, from nearby hospitals perceived as high quality to distant imaging centers with lower perceived quality. Taking the change in non-price attributes into account, I show that the gain in consumer surplus for these individuals is \$132 and \$12, respectively.

Given modest website usage, the effect on equilibrium prices may be larger if more consumers are informed about health care prices. There are two factors that make it difficult to extrapolate from reduced-form estimates. First, even though the availability of the website is exogenous, use of the website when it is available is potentially endogenous. If the individuals who find out about the website and choose to use it are those that receive a larger benefit, there may be decreasing savings as more individuals become informed about prices. Second, equilibrium prices are a function of the number of consumers that have price information. By affecting negotiated prices, price transparency generates spillover effects that benefit all consumers, including those that do not have price information.[6] By using the demand model from Chapter 2 and estimating the individual-specific probability of using the website and deriving a bargaining equation, the empirical model presented in this paper allows me to address both issues when examining out-of-sample counterfactual scenarios.

Counterfactual simulations imply that, while selection is present, the effect on equilibrium prices dominates. As a result, there would be a considerable reduction in equilibrium prices if a larger fraction of consumers had information. If all consumers were informed, equilibrium prices would be 19 percent lower. Prices decline because demand effectively becomes more elastic, allowing insurers to negotiate lower prices with most providers in their network. In addition, consumers would choose lower cost providers in their choice set, resulting in per visit savings of \$44 for consumers and \$166 for insurers relative to no price transparency. Overall, spending would decline by 28 percent. Sav-

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

ings would come largely at the expense of provider profits, although some of the savings would also be due to individuals switching to providers with lower marginal cost (e.g. imaging centers and clinics rather than hospitals).

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

3.1.1 Roadmap

The remainder of the paper is organized as follows. Section 3.2 presents the bargaining model, focusing on the role of consumer information. Section 3.3 presents the results from the supply model. Section 3.4 uses the estimates to examine the effect of the website while Section 3.5 presents out-of-sample counterfactual simulations. Section 3.6 concludes.

3.2 Bargaining between Providers and Insurers

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

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

Using the estimates from the demand model given in the previous section, I use the supply-side model to estimate the marginal cost of each procedure at each provider. These estimates are then used in Section 3.4 and Section 3.5 to simulate negotiated prices under various counterfactual scenarios.

3.2.1 Bargaining Model

I now present the model of bilateral bargaining between medical providers and insurers incorporating consumer price uncertainty. In each year, insurer k negotiates the price of procedure m with each provider in the insurer's network, $j \in \mathcal{N}_{kmt}$.[7] For the analysis, I assume that each provider negotiates independently.[8] I also take the set of providers $\mathcal{J}$ and networks $\mathcal{N}_{kmt}$ as given.[9]

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

[8]In contrast, Gowrisankaran, Nevo and Town (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.

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

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

$$\Pi^{J}_{ijkmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt}|\#_{ikmt}) = s_{ijkmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt}|\vartheta_{ikmt})[p_{jkmt} - mc_{jkmt}] \tag{3.1}$$

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

$$TC_{ikmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt}|\#_{kmt}) = \sum_{j \in \mathcal{N}_{kmt}} p_{jkmt}(1 - c_{ikmt})s_{ijkmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt}|\vartheta_{ikmt}) \tag{3.2}$$

Following Gowrisankaran, Nevo and Town (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 2.4). However, insurers are aware when consumers have uncertainty about prices, and consumer surplus includes a term that accounts for incorrect beliefs. In particular, consumer surplus is given by Equation 2.13.

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

$$\Pi^{K}_{ikmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt}|\#_{kmt}) = \zeta CS_{ikmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt}|\#_{kmt}) - TC_{ikmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt}|\#_{kmt})) \tag{3.3}$$

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

$$\Delta_j \Pi^K_{ikmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt}|\#_{kmt}) = \Pi^K_{ikmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt}|\#_{kmt}) - \Pi^K_{ikmt}(\mathcal{N}_{kmt} \backslash j, \mathbf{p}_{kmt}|\#_{kmt}) \tag{3.4}$$

Equation 3.3 and Equation 3.4 can be thought of as a stylized approach to modeling the insurer's profit function. The consumer surplus of the insurer's enrollees enters the insurer's surplus function since a larger consumer surplus implies that the insurer can charge higher premiums to consumers, generating profit for the insurer. In contrast, Ho and Lee (2017) explicitly model demand for insurance and insurer competition in order to derive an expression for provider and insurer profits that accounts for the fact that consumers may switch insurers to access their preferred providers. I lack data on insurance premiums, and therefore, I cannot explicitly model insurer competition. Consistent with reduced-form results in my previous work, I assume that price transparency does not affect insurance choice.[10]

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, $\mathbf{p}^*_{kmt} \backslash p_{jkmt}$. In other words, a hypothetical disagreement is assumed to not affect other prices.[11] Extending Rubinstein (1982), Collard-Wexler, Gowrisankaran and Lee (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-

[10] In Chapter 1, I examine whether insurance enrollment changed after the introduction of the price transparency website and do not find a statistically significant effect.

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

insurer-procedure triple in a given year, p^*_{jkmt}, that satisfy

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

$$\times \left(\underbrace{\sum_{i \in I_{kmt}} \mathbb{E}_e \left[\Delta_j \Pi^K_{ikmt}(\mathcal{N}_{kmt}, p_{jkmt}, \mathbf{p}^*_{kmt} \backslash p_{jkmt} | \#_{kmt}) \right]}_{\text{Insurer gains from trade}} \right)^{1-\tau} \quad (3.5)$$

where the gains from trade are summed over all individuals enrolled in insurer k receiving procedure m in year t, I_{kmt}. The Nash bargaining weight is $\tau \in [0, 1]$. Since insurers and providers do not know the price signals that consumers will receive, both take the expectation over consumer beliefs.[12]

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.[13] I do not deviate from this assumption.[14] In the model presented in this section, price transparency indirectly affects equilibrium prices since changes in consumer behavior affect the gains from trade. I assume that the price transparency website, which was targeted towards consumers, did not directly affect the information set of the providers or insurers. Further research is needed to understand whether price transparency affects provider-insurer bargaining directly.[15]

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

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

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

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

3.2.2 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 3.5 implies that equilibrium prices are determined by marginal cost plus a margin:[16]

$$p_{jkmt} = mc_{jkmt} + \left(-\frac{1-\tau}{\tau} \frac{\frac{\partial}{\partial p_{jkmt}} [\sum_{i \in I_{km}} \mathbb{E}_e \Pi^K_{ikmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt} | \#_{kmt})]}{\sum_{i \in I_{km}} \mathbb{E}_e \Delta_j \Pi^K_{ikmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt} | \#_{kmt})} - \frac{\frac{\partial}{\partial p_{jkmt}} [\sum_{i \in I_{km}} \mathbb{E}_e s_{ijkmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt} | \vartheta_{ikmt})]}{\sum_{i \in I_{km}} \mathbb{E}_e s_{ijkmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt} | \vartheta_{ikmt})} \right)^{-1} \tag{3.6}$$

I present further detail, including the derivation of $\frac{\partial}{\partial p_{jkmt}} [\sum_{i \in I_{km}} \mathbb{E}_e \Pi^K_{ikmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt} | \#_{kmt})]$, in Appendix 3.6. Given that there are many providers in each network, a single price change has a minimal effect on individuals' prior about the distribution of prices. For tractability, I assume that providers and insurers do not take changes in the prior into account, and therefore hold the prior fixed when solving for the first order condition.

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

In the market for privately-provided health care, insurers negotiate their own rates with each provider that are thought to be lower than what a Bertrand-Nash pricing assumption would imply.[17] This corresponds to the case in which $\tau < 1$. Therefore, it is important to also understand how price transparency affects insurers' incentive to negotiate lower prices.

There are multiple channels through which consumer price transparency can affect equilibrium outcomes in the bargaining model. First, price transparency affects the in-

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

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

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

The effect of price transparency on insurers' incentives are more complicated. Price transparency affects insurer cost since consumers tend to switch to lower cost providers. This can be seen by noting that $TC_{ikmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt}|\#_{kmt})$ depends on the choice probabilities. Price transparency also affects the consumer surplus of the insurer's enrollees, $CS_{ikmt}(\mathcal{N}_{kmt}, \mathbf{p}_{kmt}|\#_{kmt})$, since individuals can switch to lower cost providers and are not surprised by the bill (see Equation 2.4 and Equation 2.13).

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

3.2.3 Estimation and Identification of Bargaining Model

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

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

$$mc_{jkmt} = \eta_j + \eta_m + \eta_t + \varepsilon^{MC}_{jkmt} \tag{3.7}$$

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

Using the parameterized marginal cost above along with the first-order condition given by Equation 3.6, the marginal cost error is given by

$$\varepsilon^{MC}_{jkmt} = \eta_j + \eta_m + \eta_t - p_{jkmt} + \left(-\frac{1-\tau}{\tau} \frac{\frac{\partial}{\partial p_{jkmt}} \sum_{i \in I_{km}} [\zeta CS_{ikmt} - TC_{ikmt}]}{\sum_{i \in I_{km}} [\zeta \Delta_j CS_{ikmt} - \Delta_j TC_{ikmt}]} - \frac{\frac{\partial}{\partial p_{jkmt}} [\sum_{i \in I_{km}} s_{ijkmt}]}{\sum_{i \in I_{km}} s_{ijkmt}} \right)^{-1} \tag{3.8}$$

This is used to form the following moment condition:

$$\mathbb{E}[\varepsilon^{MC}_{jkmt} | Z_{jkmt}] = 0 \tag{3.9}$$

where Z_{jkmt} is a vector of variables assumed to be exogenous. The model assumes that the bargaining participants know mc_{jkmt}, including ε^{MC}_{jkmt}, implying that prices are potentially endogenous. Following the previous literature, I address this issue by including two instruments: predicted willingness-to-pay for each provider at mean price and predicted total provider quantity at mean price.[18] Although these instruments are correlated with price, it is assumed that they are uncorrelated with ε^{MC}_{jkmt}. The instrument set, Z_{jkmt}, also includes all marginal cost fixed effects.

Identification of parameters η, τ, and ζ follows from a similar argument as that presented in Gowrisankaran, Nevo and Town (2015). The provider choice and website usage

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

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 3.6). In the bargaining model, these are treated like data. Variation in provider incentives (determined by s_{ijkmt} and $\partial s_{ijkmt}/\partial p_{jkmt}$) and insurer incentives (determined by CS_{ikmt}, TC_{ikmt}, $\partial CS_{ikmt}/\partial p_{jkmt}$, and $\partial TC_{ikmt}/\partial p_{jkmt}$) that can explain variation in prices identifies ζ and τ. This variation comes in part from the introduction of the price transparency website. The remaining price variation identifies the marginal cost fixed effects, η. Unlike Gowrisankaran, Nevo and Town (2015), I take advantage of price variation across individual procedures. This provides an additional source of variation to identify ζ and τ.

3.3 Estimates from Supply Model

Table 3.1 provides results from the bilateral bargaining model. The estimated bargaining weight is 0.37, implying that insurer incentives are important for equilibrium prices. This estimate of the bargaining weights is lower than some other estimates in the literature.[19]

Table 3.1: Bargaining Model Estimates

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

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

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

[19] For comparison, Ho and Lee (2017) estimate provider bargaining weights between 0.50 and 0.88. However, Gowrisankaran, Nevo and Town (2015) estimate provider bargaining weights that average 0.24.

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

The estimates from the bargaining model can be used to construct the marginal cost of each procedure at each provider in each year. The marginal cost estimates are summarized by procedure group in Table 3.2. MRI scans have the highest marginal cost, followed by CT scans and X-rays. The marginal cost estimates are consistent with the fact that MRI machines are the most expensive, while CT and X-ray machines are less expensive and require less staff.

In order to examine whether marginal cost estimates are plausible, I compare the estimates from the model to Medicare reimbursement rates. I find that the marginal cost estimates are generally comparable to the Medicare fee schedule. While I estimate the average marginal cost for medical imaging procedures is \$465, the average Medicare non-facility reimbursement for the same procedures is \$526.[20]

I recompute baseline prices using the estimates of marginal cost as well as the other estimates from the demand and supply model. Given the estimates and an initial price schedule, I calculate the solution to each bargaining problem conditional on other prices. I iterate until all solutions of the bargaining problem reflect the best response to all other prices. This procedure is outlined in more detail in Appendix 3.6. The mean and standard deviation of recomputed baseline prices by procedure group are shown in Table 3.2. Markups are largest for x-ray procedures. The Lerner index, $(p - mc)/p$, is 0.65 for X-rays, but only 0.19 and 0.10 for CT scans and MRI scans respectively. Since the majority of procedures are X-rays, the average Lerner index across all procedures is 0.47.

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

Table 3.2: Marginal Cost Estimates

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

Notes: Prices are simulated using the algorithm described in Appendix 3.6. The unit of observation is a provider, procedure, insurer, year. All prices in 2010 dollars.

3.4 Effect of the Price Transparency Website

In this section I use the estimates from both the demand and supply models to perform a policy evaluation and examine the effect of New Hampshire's price transparency website. First, I calculate the overall equilibrium effect of the website for all individuals who potentially used the price transparency website. This estimate can be compared to the intent-to-treat effect estimated using the difference-in-differences approach in Chapter 1. I show that results from the two methods are qualitatively consistent, providing support for the main counterfactuals presented in Section 3.5. I also discuss the conceptual differences between the structural approach in this paper and the difference-in-differences approach.

Next, I use the model to examine the effect of the price transparency website in more detail than is possible with the reduced-form approach alone. In particular, I find the effect of the website conditional on using the price transparency website, often referred to as the treatment-on-the-treated effect. Estimating the treatment-on-the-treated using reduced-form methods is hampered by the fact that there are spillovers that effect all consumers.[21] The estimates from the empirical model imply that consumers subject to a deductible benefit the most from the price transparency website. In addition, the empiri-

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

cal model allows me to recover the change in consumer surplus, which I find is less than consumer cost savings.

3.4.1 Comparison with Difference-in-Differences Estimates

In Chapter 1, I analyzed the effect of New Hampshire's price transparency website using quasi-experimental reduced-form methods. By comparing the price of medical imaging procedures on the website before and after the website started with the change in price for similar medical imaging procedures that were never on the website, the approach recovers the average effect of access to the website.[22]

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 3.6. These prices are then used to calculate the average percent change in spending due to the website.

Table 3.3 presents the estimated effect of access to the price transparency website along with the difference-in-differences estimates from Chapter 1. The estimates from the empirical model are similar to the estimates obtained from the difference-in-differences specification. While the estimates from the former imply that the website reduced overall spending by 2.9 percent, the latter implies that there was a 3.8 percent reduction in spending. Focusing on the reduction in cost for consumers and insurer, I also find similar estimates from the two models. In both cases, the estimates from the empirical model are smaller than the difference-in-differences estimates.

While these estimates are qualitatively consistent, there are a number of potential reasons why the estimates from the two approaches are not identical. First, this paper uses

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

Table 3.3: Intent-to-Treat Effect of Price Transparency Website
Comparison with Reduced Form Results

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

Notes: Figures show the percent change in transaction prices for all individuals who could have used the website. Difference-in-differences estimates, which are from Chapter 1, 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.

a slightly smaller sample of medical imagining procedures than my previous paper (see discussion in Section 2.2). In addition, 2011 is excluded when estimating the empirical model since website traffic data is not available. Second, the two approaches require different assumptions to identify the causal effect of the website. The main assumption of the difference-in-differences approach is that the price of procedures on the website would follow a common trend relative to procedures that were not on the website in the absence of treatment. In my previous work, I examine trends in the period prior to the introduction of the website to provide evidence consistent with this assumption. In contrast, the empirical model requires functional form assumptions and distributional assumptions, as well as an assumption about the form of imperfect competition. While the reduced form approach assumes that spending is not affected by unobserved covariates that are correlated with the availability of the website, the structural approach implicitly assumes that primitives, such as the distribution of demand parameters, are orthogonal to the availability of the website.

There is a small literature that directly compares quasi-experimental reduced-form estimates and estimates from a structural model.[23] While the reduced-form analysis is useful for policy evaluation and relies on relatively clear assumptions, the remainder of this paper focuses on results that can only be obtained from an empirical model based on

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

economic theory. In particular, the empirical model allows for an analysis of wider-scale adoption of health care price transparency.[24] I argue that the difference-in-differences estimates help validate the policy-invariant parameter estimates from the empirical model, lending credence to the counterfactual simulations in the remainder of the paper.

3.4.2 Effect Conditional on using the Price Transparency Website

The first panel of Table 3.4 summarizes the effect of using the website holding prices fixed.[25] 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.[26] These individuals pay a relatively small portion of the total negotiated prices (cost sharing is about 6 percent on average), therefore the savings from using the website are only $15 per visit on average. Although consumers only take into account the out-of-pocket price, there is correlation between the provider out-of-pocket price and the insurer price. Therefore, insurers also benefit from the increased price shopping (insurers save $12 on average).

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

I compute the change in consumer surplus for individuals who use the website using Equation 2.4 and Equation 2.13. The gain is smaller than the cost savings—$90 for individuals subject to a deductible. This is due to the fact that, without price information,

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

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

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

Table 3.4: Effect for Individuals Predicted to have Used the Transparency Website

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

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

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 \$127 when they have price information, the providers they choose are \$37 worse on non-price characteristics.

In the second panel of Table 3.4, I account for the fact that the website changed negotiated prices in addition to consumer choices. Rather than hold prices fixed, I re-simulate prices for the case in which some individuals used the price transparency website.

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

3.5 Out-of-Sample Counterfactuals

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

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

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

3.5.1 Effect of Increased Price Transparency

Effect on Overall Savings

Figure 3.1 shows the effect of reducing the implicit cost of becoming informed about prices, increasing the fraction of consumers with price information. 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 3.1. I then use these prices to compute consumers choices and overall spending.

The equilibrium effect of increased price transparency is shown by the solid line in

Figure 3.1: Effect of Price Transparency on Medical Imaging Spending
By Fraction on Individuals with Price Information

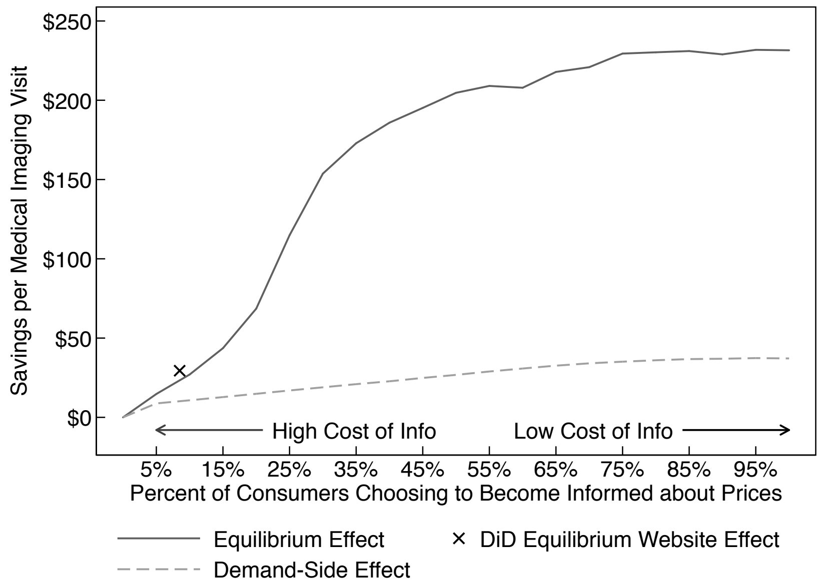

Notes: Demand-side effect holds prices fixed at distribution simulated with no price transparency. Equilibrium effect re-simulates equilibrium prices for each level of price transparency. Difference-in-differences estimates from Chapter 1. All figures in 2010 dollars.

Figure 3.1. As more individuals are informed, the amount saved per visit is highly non-linear. Initially, the supply-side effects are modest—when only a few consumers are informed about prices, equilibrium prices remain relatively constant. When a larger fraction of consumers are informed, there is a meaningful effect on overall demand, changing the incentives of both providers and insurers when negotiating prices. This leads to lower equilibrium prices, generating large savings. As price-cost margins decline, insurers become limited in their ability to negotiate ever lower prices. Once about half of consumers are informed about prices, the supply-side effects become less relevant. For reference, the demand-side effect from Chapter 2 is given by the dashed line in Figure 3.1.

Effect on Consumers, Insurers, and Providers & Welfare Analysis

I examine increased price transparency in greater detail by examining the effect for consumers, insurers, and providers. I focus on two counterfactuals: the equilibrium effect

Table 3.5: Counterfactual Negotiated Provider Prices

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

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

of the price transparency website if it was available for all procedures in all years and the equilibrium effect of full price transparency.

Panel (a) of Table 3.5 shows the impact on negotiated provider prices relative to no price transparency. When no individuals have price information, the average price of the medical imaging procedures in $1,014. If a price transparency website is available for all medical imagining procedures in all years, the average price declines 2.5 percent to $988.[27] Finally, I examine the counterfactual scenario in which all individuals are fully informed about prices. This would be the case if, for instance, primary care providers were required to provide a price schedule after recommending a medical imaging test. In this case, prices would be 18.9 percent lower than the baseline case. The distribution of prices under various counterfactual scenarios is shown in Figure A7.

In the second and third column of Table 3.5, I examine the effect on price dispersion, as measured by the interquartile range of prices. An increase in price information reduces

[27] This is broadly consistent with difference-in-differences estimates that isolate the supply side. However, note that since this counterfactual examines the effect of the website for all procedures in all years, it is not directly comparable to the estimates from the difference-in-differences model or the results in Section 3.4.

the degree of price dispersion. Although the mechanism is different, these results are broadly consistent with the literature stressing that price dispersion can result from search frictions.[28]

Panel (a) of Table 3.6 presents the overall effect on spending taking into account both supply and demand-side effects. If the website were available for all procedures in all years, consumers would save $6 and insurers would save $21 on average, generating $2.1 million in total savings per year on X-ray, CT scans, and MRI scans in New Hampshire. Full price transparency leads to $16.3 million in savings, of which 21 percent accrues to consumers.

The effect for providers is shown in Panel (a) of Table A10. 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.[29] The overall welfare impact for consumers, insurers, and providers is shown in Panel (a) of Table 3.7. Full price transparency would result in an overall welfare gain of $2.5 million annually.

3.5.2 Effect of Price Transparency Combined with High Cost Sharing

Health insurance plans with high cost sharing, such as high-deductible plans, potentially give consumers more "skin in the game", increasing the incentive to make cost-effective decisions.[30] Partially for this reason, policies such as tax-advantaged Health Savings Accounts have encouraged high cost sharing plans. However, if consumers cannot observe prices, high cost sharing alone may not lead consumers to switch to less

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

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

[30]See, for instance, discussion in Gowrisankaran, Nevo and Town (2015).

Table 3.6: Counterfactual Cost, Welfare, and Expenditure

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

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

expensive options. For instance, Brot-Goldberg et al. (2015) do not find evidence that high deductible plans increase price shopping.[31] Instead, deductibles may simply lead to larger bill shock. In this section, I examine the effect of combining high cost sharing with the New Hampshire price transparency website. I find that the interaction of these two policies generates larger savings than the price transparency alone.

In Table 3.5 panel (b), I consider that case in which individuals all have plans with a 50 percent coinsurance rate. In this case, individuals have more incentive to use the website since the potential savings are larger. Simulations imply that website usage increases 38 percent under the high cost sharing scenario. This puts additional downward pressure on prices, resulting in mean prices that are 16 percent lower than with the price transparency website alone. Under high cost sharing, the market is relatively efficient when individuals have access to the price transparency website. Comparing this with full price transparency combined with high cost sharing, the additional supply-side effects are small. In fact, the prices estimates with full price transparency and high cost sharing are slightly higher than those estimates with full transparency alone. The effect on prices is due in part to the fact that insurers have less incentive to negotiate lower prices if they incur a

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

Table 3.7: Counterfactual Net Welfare Impact for Consumers, Providers, and Insurers

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

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

smaller portion of the negotiated price.

High cost sharing leads to a large reduction in price dispersion. However, under high cost sharing, price dispersion actually grows as more individuals are informed about prices. When negotiating prices, the incentives of the insurer are multifaceted, resulting in a complex effect on the distribution of prices (see Figure A7). As discussed above, high cost sharing means that insurers have less incentive to negotiate low prices in some cases, especially if they know that consumers can shop around. Therefore, it is not always the case that price dispersion declines with increased price transparency.

Table 3.6 reports the impact on transaction prices, welfare, and total annual spending taking into account both the demand- and supply-side effects. When the website is combined with high cost sharing, the annual savings total $13.5 million, over 6 times larger than with the website alone. Full price transparency results in $36.3 million in savings. Due to the high cost sharing, these savings accrue to the insurer, whereas the consumers have higher out-of-pocket cost. However, it is possible that insurer savings are passed on to consumers in the form of lower insurance premiums.[32]

[32] This would require an explicit model of insurer competition. See Ho and Lee (2017).

Panel (b) of Table 3.7 shows the overall welfare impact of combining price transparency with high cost sharing. Although much of the reduction in health care spending is due to a transfer from providers to insurers, there are still significant net welfare gains. When full price transparency is combined with high cost sharing, many consumers switch to lower marginal cost providers, resulting in a welfare gain of $18.0 million annually.

3.6 Conclusion

In this paper, I study how price transparency in the market for medical imaging procedures affects negotiated health care prices. I contribute to the literature by developing an empirical model of competition in the market for medical procedures that separately accounts for consumer preferences and consumer uncertainty about prices.

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

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

Although the empirical analysis in this paper focuses on X-rays, CT scans, and MRI scans, the lack of price transparency likely has implications for other medical procedures that are "shoppable". This includes other outpatient procedures and services such as primary care office visits, simple elective surgeries, and diagnostic testing procedures, as well as some inpatient procedures such as newborn delivery. Previous analysis of health care expenditures has estimated that over 40 percent of health care spending is on services that are potentially shoppable.[33]

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

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

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

Bibliography

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

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

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

Angrist, Joshua D., and Jörn-Steffen Pischke. 2010. "The Credibility Revolution in Empirical Economics: How Better Research Design Is Taking the Con out of Econometrics." *Journal of Economic Perspectives*, 24(2): 3–30.

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

Athey, Susan, and Guido W. Imbens. 2006. "Identification and Inference in Nonlinear Difference-in-Differences Models." *Econometrica*, 74(2): 431–497.

Berkman, Nancy D., Stacey L. Sheridan, Katrina E. Donahue, David J Halpern, and Karen Crotty. 2011. "Low health literacy and health outcomes: an updated systematic review." *Annals of internal medicine*, 155(2): 97–107.

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

Brot-Goldberg, Zarek C., Amitabh Chandra, Benjamin R. Handel, and Jonathan T. Kolstad. 2015. "What Does a Deductible Do? The Impact of Cost-Sharing on Health Care Prices, Quantities, and Spending Dynamics." National Bureau of Economic Research.

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

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

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

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

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

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

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

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

Christensen, Hans Bonde, Eric Floyd, and Mark G. Maffett. 2015. "The Effects of Price

Transparency Regulation on Prices in the Healthcare Industry."

Clay, Karen, Ramayya Krishnan, and Eric Wolff. 2001. "Prices and Price Dispersion on the Web: Evidence from the Online Book Industry." In *E-commerce*. *NBER Chapters*, 521–539. National Bureau of Economic Research, Inc.

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

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

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

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

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

David, Guy, Richard C. Lindrooth, Lorens A. Helmchen, and Lawton R. Burns. 2014. "Do hospitals cross-subsidize?" *Journal of health economics*, 37: 198–218.

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

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

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

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

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

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

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

Dranove, David, Mark Shanley, and William D. White. 1993. "Price and Concentration in Hospital Markets: The Switch from Patient-Driven to Payer-Driven Competition." *Journal of Law and Economics*, 36(1): 179–204.

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

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

Einav, Liran, and Jonathan Levin. 2010. "Empirical Industrial Organization: A Progress Report." *Journal of Economic Perspectives*, 24(2): 145–62.

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

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

Ellison, Glenn, and Sara Fisher Ellison. 2005. "Lessons About Markets from the Internet." *Journal of Economic Perspectives*, 19(2): 139–158.

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

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

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

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

Finkelstein, Amy, Matthew Gentzkow, and Heidi Williams. 2014. "Sources of geographic variation in health care: Evidence from patient migration." National Bureau of Economic Research.

Fisher, Elliott S., David E. Wennberg, Thrse A. Stukel, Daniel J. Gottlieb, F. Lee Lucas, and Etoile L. Pinder. 2003. "The implications of regional variations in Medicare spending. Part 1: the content, quality, and accessibility of care." *Annals of internal medicine*, 138(4): 273–287.

Fisher, Elliott S., Julie P. Bynum, and Jonathan S. Skinner. 2009. "Slowing the growth of health care costs—lessons from regional variation." *New England Journal of Medicine*, 360(9): 849–852.

Frakt, Austin B. 2011. "How much do hospitals cost shift? A review of the evidence." *Milbank Quarterly*, 89(1): 90–130.

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

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

Gaynor, Martin, and William B. Vogt. 2003. "Competition among Hospitals." *The RAND*

Journal of Economics, 34(4): 764–785.

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

Goldmanis, Maris, Ali Hortaçsu, Chad Syverson, and Onsel Emre. 2010. "E-Commerce and the Market Structure of Retail Industries." *Economic Journal*, 120(545): 651–682.

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

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

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

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

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

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

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

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

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

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

Ho, Kate, and Robin S. Lee. 2017. "Insurer Competition in Health Care Markets." *Econometrica*, 85(2): 379–417.

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

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

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

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

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

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

Hortaçsu, Ali, and Chad Syverson. 2004. "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*, 119(2): 403–456.

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

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

Kessler, Daniel P., and Mark B. McClellan. 2000. "Is Hospital Competition Socially

Wasteful?" *The Quarterly Journal of Economics*, 115(2): 577–615.

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

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

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

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

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

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

MacKinnon, James G., and Lonnie Magee. 1990. "Transforming the Dependent Variable in Regression Models." *International Economic Review*, 31(2): 315–339.

Melly, Blaise, and Giulia Santangelo. 2015. "The changes-in-changes model with covariates." *Working Paper*.

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

Morton, Fiona Scott, Florian Zettelmeyer, and Jorge Silva-Risso. 2001. "Internet Car Retailing." *Journal of Industrial Economics*, 49(4): 501–19.

Nevo, Aviv, and Michael D. Whinston. 2010. "Taking the Dogma out of Econometrics: Structural Modeling and Credible Inference." *Journal of Economic Perspectives*, 24(2): 69–82.

Newhouse, Joseph P., Alan M. Garber, Robin P. Graham, Margaret A. McCoy, Michelle Mancher, Ashna Kibria, et al. 2013. *Variation in Health Care Spending:: Target Decision*

Making, Not Geography. National Academies Press.

Nicholson, Elizabeth. 2015. “2015 Price Transparency Initiative State Survey: Price Disclosures.” The Source on Healthcare Price and Competition.

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

Philipson, Tomas J., Seth A. Seabury, Lee M. Lockwood, Dana P. Goldman, Darius N. Lakdawalla, and David M. Cutler. 2010. “Geographic Variation in Health Care: The Role of Private Markets.” *Brookings Papers on Economic Activity*, 325–361.

Reinhardt, Uwe. 2012. “Divide et impera: protecting the growth of health care incomes (COSTS).” *Health economics*, 21(1): 41–54.

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

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

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

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

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

Salz, Tobias, and Emanuel Vespa. 2016. “Estimating Dynamic Games of Oligopolistic Competition: An Experimental Investigation.”

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

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

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

Sinaiko, Anna D., and Meredith B. Rosenthal. 2011. "Increased price transparency in health care—challenges and potential effects." *New England Journal of Medicine*, 364(10): 891–894.

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

Smith, Jessica C, Carla Medalia, et al. 2014. *Health insurance coverage in the United States: 2013.* US Department of Commerce, Economics and Statistics Administration, Bureau of the Census.

Sorensen, Alan T. 2000. "Equilibrium price dispersion in retail markets for prescription drugs." *Journal of Political Economy*, 108(4): 833–850.

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

Squires, David A. 2011. "The US health system in perspective: a comparison of twelve industrialized nations." *Issue Brief (Commonwealth Fund)*, 16: 1–14.

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

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

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

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

Town, Robert, and Gregory Vistnes. 2001. "Hospital competition in HMO networks." *Journal of Health Economics*, 20(5): 733–753.

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

Tu, Ha, and Rebecca Gourevitch. 2014. "Moving Markets: Moving Markets: Lessons from New Hampshire's Health Care Price Transparency Experiment." California HealthCare Foundation.

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

Whaley, Christopher M. 2015b. "Provider Responses to Online Price Transparency."

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

Appendix

Appendix A: Additional Details on Data Construction

This section describes the construction of the data in further detail.

Sample Selection

In the main specification, I consider the universe of private-health insurance claims for individuals under age 65. Individual age 65 and older may have Medicare in addition to their private insurance, making them a very selected population. I also remove claims that were denied or refunded.

I identify the universe of radiology procedure by linking procedure descriptions from Centers for Medicare & Medicaid Services (CMS) CPT/HCPCS Database to the procedure codes in the claims database. Procedure codes that changed description are assigned separate identifiers.

I also use the CMS CPT/HCPCS Database to identify the subset of radiology procedures with price information available on the HealthCost website. A procedure on the website may refer to multiple procedure codes (e.g. the website does not provide separate price information for a knee X-ray with 2 views versus a knee X-ray with 3 views, even though these have different procedure codes.[34] The list of procedures on the website

[34] The HealthCost website often only used one CPT/HCPCS code to calculate prices, however prices are

and associated CPT/HCPCS codes are given in Appendix Table A7.

Visit Price and Principal Procedure Price

Prices on the website are calculated by aggregating all health claims related to the radiology procedure, which I call the visit price. I use a similar method as the website, with some important differences. In this section, I describe how visit prices are calculated and how my method differs from that of the website.[35]

Like the website, I sum the price of all medical claims on the day of the visit to get the visit price. I exclude visits in which the individual had inpatient claims as well as visits in which the most expensive procedure was not a radiology procedure.

The New Hampshire Insurance Commission removed claims when the total charge was above the 95th percentile. They also remove visits in which the patient out-of-pocket price was in the lowest one percentile or highest fifth percentile. Unlike the website, I do not remove any visits based on price. Although outliers may be an issue for surgical procedures in which the complexity of a given surgery can vary widely, radiology procedures are fairly homogenous.

The website generally calculates the prices using a single procedure code. For instance, in order to calculate the price of a foot X-ray, the website consider the price of procedure code 73620, defined as an X-ray exam of the foot with 2 views. However, individuals receiving an X-ray exam with 1 view may also use the website and are likely to be unaware that the posted prices are for 2 views. Since price of these procedures is highly correlated for a given provider, the individual may still benefit from the information. For this reason, I consider all procedures with a description matching a procedure on the website to be treated, even if it was not included when calculating prices for the website.

---

highly correlated.

[35]The detailed method used by the New Hampshire Insurance Department is described on the website.

Additional Data Sources

Nominal prices of radiology procedures are generally trending upwards over the period. All prices are inflation-adjusted to 2010 dollars using the Medical Care Services CPI from the U.S. Bureau of Labor Statistics. Even so, real prices are still generally increasing over the period.

I use a number of additional sources to construct demographic covariates. In order to construct the Charlson comorbidity index, I use Stagg (2006). The Charlson comorbidity index is calculated for each individual for each year. I obtain information about the urbanization/rurality of each zip code using the 2014 Rural-Urban Chartbook from the Rural Health Research Center. Finally, I use the CMS CPT/HCPCS Database to identify procedures that are likely a results of an emergency or urgent care episode.

Appendix B: Effect on Price Trends

I find evidence that the effect of the website grows consistently over the five year period after the website's introduction. I argue that this is likely due in part to the delayed supply-side response and the increase in website usage over time. An alternative explanation is that the website affected health expenditure growth.

I estimate an alternative model that allows the effect of the website to vary linearly over the period after the website began. This specification is:

$$\log(1 + p_{imjkt}) = \beta(\text{OnWeb}_m \times \text{YearsSinceWebsite}_t) + \alpha X_{it} + \lambda_m + \lambda_k + \lambda_t + \varepsilon_{imjkt} \quad (3.10)$$

Rather than include dummy variable Post_t as in the baseline specification, I now include $\text{YearsSinceWebsite}_t$, the number of years since the introduction of the website which is defined as

$$\text{YearsSinceWebsite}_t = \begin{cases} (t - t_0)/12 & \text{if } t \geq t_0 \\ 0 & \text{if } t < t_0 \end{cases} \tag{3.11}$$

where t_0 is the month of the website introduction, March 2007.

In this specification, β is now interpreted as the annual effect of the website. Or more precisely, it is the log-point change in transaction prices due to the website being available for an additional year.

The results are presented in Appendix Table A4. Each year the website is available decreases out-of-pocket prices by 3.8 log points. This is statistically significant at the 1 percent level. Therefore, if the trend continues, the website is predicted to lower out-of-pocket prices by 32 percent in a decade. For the insurer paid amount, there is no evidence of an increasing annual effect of the website. The coefficient is negative but insignificant.

Although I interpret the effect as a change in price levels in the main specifications, future research is needed to examine whether price transparency can reduce health care spending growth in the long run.

APPENDIX C: TREATMENT HETEROGENEITY BY PRICE QUANTILE

Empirical Strategy

The price of radiology procedures varies widely. For instance, ultrasound examinations can be relatively inexpensive while PET scans can be thousands of dollars. In this section, I examine how the treatment effect of the website varies by price quantile. This provides additional insight into the population affected by the introduction of the website.

In particular, I am interested in the effect on the τth quantile

$$Q_\tau(LogPrice_{imjkt}|\text{OnWeb}_m, \text{Post}_t, X_{it}) = \beta(\tau)(\text{OnWeb}_m \times \text{Post}_t) + \lambda_1(\tau)\text{OnWeb}_m + \lambda_2(\tau)\text{Post}_t + \alpha(\tau)X_{it} \tag{3.12}$$

In general, the quantile treatment effect implied by the difference-in-difference estimator is asymptotically biased. This is because the quantile difference-in-difference estimator assumes that the time effects and group affects are additive, a strong assumption in a non-linear model. In order to recover the full distribution of counterfactual outcomes when there is both panel variation and time variation, Athey and Imbens (2006) develop the "changes-in-changes" model. Their approach has relatively mild assumption, namely that the distribution of unobservables in the treated and untreated groups remains the same across time. Another benefit of their approach is that it is invariant to the scale of the dependent variable.

Note that the model is simplified in order to make it computationally tractable. Rather than individual procedure fixed effects and month fixed effects, I use an indicator variables for whether procedures are on the website and an indicator variable for whether the website has been introduced. I include covariates using the approach developed by Melly and Santangelo (2015).

Results

The results from both the quantile difference-in-difference model and the changes-in-changes model are presented in Figure A3. Both approaches yield similar conclusions, but I focus on the results from the changes-in-changes model.

I estimate the effect of the website by price decile. All but the lowest price procedures are significantly affected by the website. This is consistent with the intuition that individuals having inexpensive procedures are less likely to use the website and shop around since they have less to gain from switching to a low cost provider. Prices at the 6th decile are the most affected, and see a price decline of more than 10 percent.

There is a significant affect for the highest price procedures, but the point estimate is relatively small. This may be due to the fact that expensive procedures can be highly-specialized, and it may be more difficult to switch to a low price provider. Consequently, the benefits from the website are smaller.

Appendix D: Additional Figures for Reduced-Form Analysis

Figure A1: Distribution of Herfindahl index Across Counties

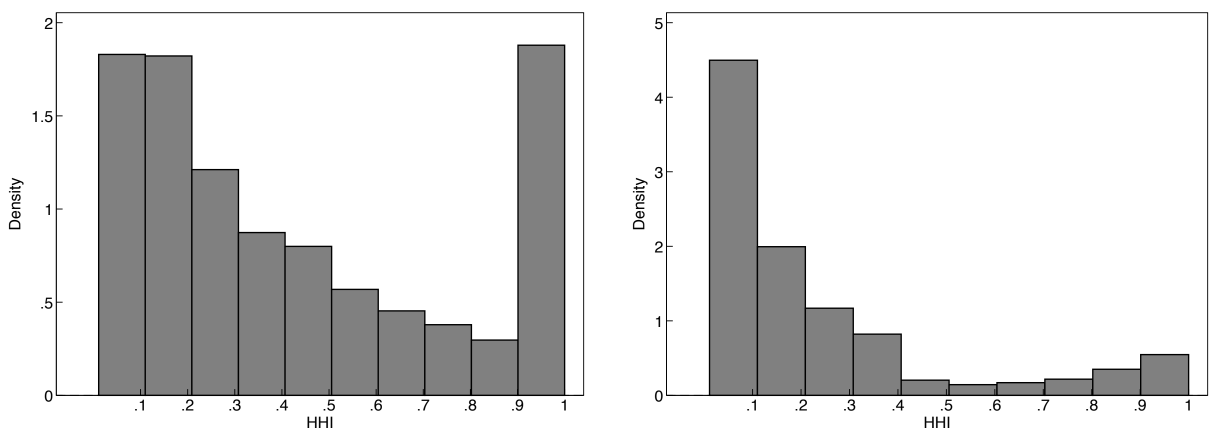

a. Unweighted

b. Weighted by County Population

Figure A2: Composition of Insurer Firm and Insurance Type by Month

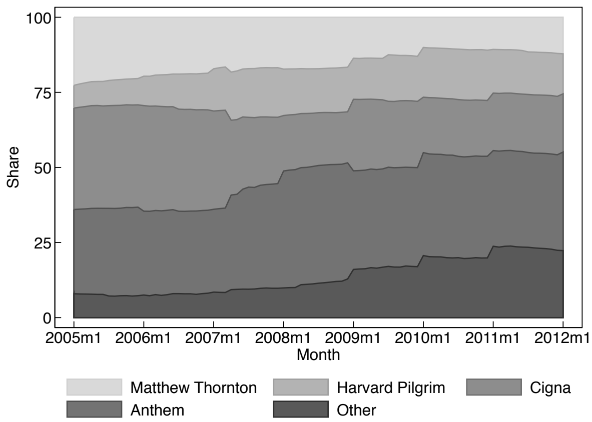

a. Composition of Insurer Firm

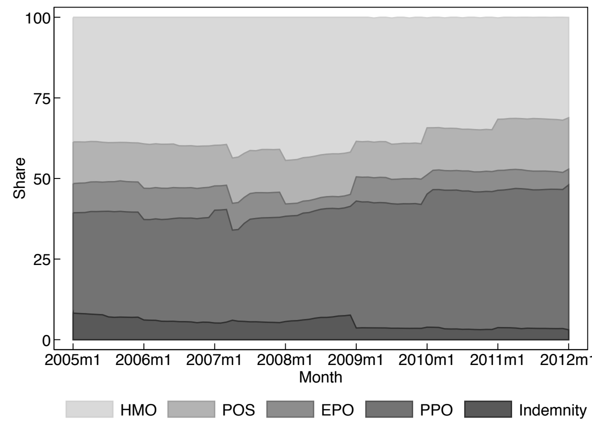

b. Composition of Insurer Type

Notes: Sample is the universe of privately-insured individuals in New Hampshire, including those with no medical claims.

Figure A3: The Effect of Price Transparency Website on Visit Price by Price Quantile

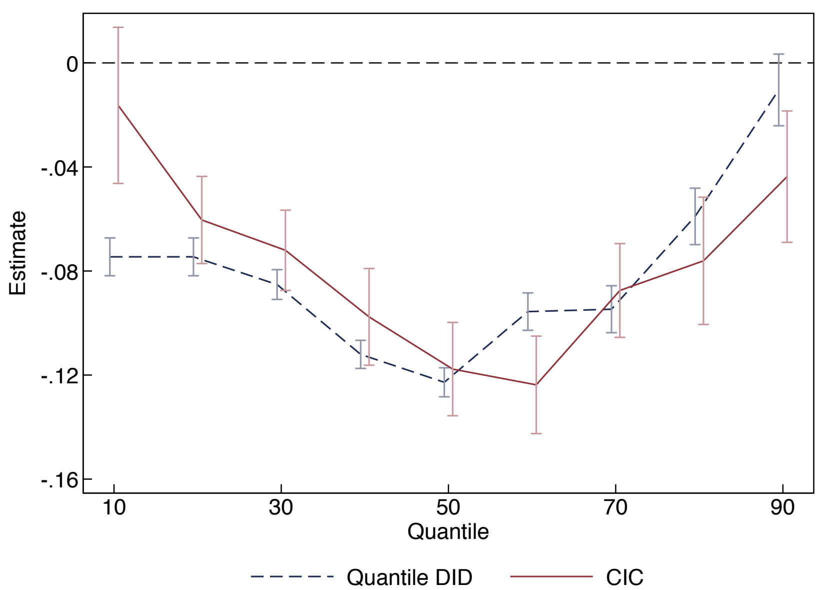

Notes: Error bars indicate 95 percent confidence interval. Quantile difference-in-difference (DID) model estimated using full sample. Changes-in-changes (CIC) model estimated using quarter sample.

Appendix E: Additional Tables for Reduced-Form Analysis

Table A1: Radiology Procedures with Price Information Available on HealthCost Website

Website Procedure Name	CPT/HCPCS Code	Detailed Description
Ankle X-Ray	73610	X-ray ankle
Bone Density Scan	76075	Dual X-ray absorptiometry, axial skeleton
	77080	Dual X-ray absorptiometry, axial skeleton
	77082	Dual X-ray absorptiometry, vertebral fracture
	76076	Dual X-ray absorptiometry, appendicular skeleton
	76071	CT bone density, peripheral
Chest X-Ray	71020	Chest x-ray, two views, frontal and lateral
	71035	Chest x-ray, special views
	71022	Chest x-ray, two views, frontal and lateral (with oblique)
	71010	Chest x-ray, single view frontal
	71101	Chest/ribs X-ray exam, three views
CT - Abdomen	74160	CT abdomen w/dye
	74150	CT abdomen w/o dye
	74170	CT abdomen w/o & w/dye
CT - Chest	71260	CT thorax w/dye
	71250	CT thorax w/o dye
	71270	CT thorax w/o & w/dye
CT - Pelvis	74177	CT Pelvis and Abdomen w/dye
	72193	CT pelvis w/dye
	72192	CT pelvis w/o dye
	72194	CT pelvis w/o & w/dye
	74178	CT Pelvis and Abdomen w/o & w/dye
	74176	CT Pelvis and Abdomen w/o dye
Foot X-Ray	73630	Foot X-ray, $\geq$ three views
	73650	Foot/heel X-ray
	73620	Foot X-ray, two views
Knee X-Ray	73562	X-ray exam of knee, 3
	73564	X-ray exam, knee, $\geq$ 4 views
	73560	X-ray exam of knee, 1 or 2
	73565	X-ray exam of knees
Mammogram	77057	Screening mammogram, bilateral
	G0202	Screening mammogram, digital bilateral
	77063	Screening mamogram, tomosynthesis bilateral
	76090	Mammogram, unilateral
	76091	Mammogram, bilateral
	76092	Screening mammogram, film bilateral
MRI - Back	72148	MRI lumbar spine w/o dye
	72158	MRI lumbar spine w/o & w/dye
	72141	MRI neck spine w/o dye
	72156	MRI neck spine w/o & w/dye
MRI - Brain	70553	MRI brain w/o & w/dye
	70551	MRI brain w/o dye
MRI - Knee	73721	MRI knee w/o dye
	73723	MRI knee w/o & w/dye
MRI - Pelvis	72197	MRI pelvis w/o & w/dye
	72195	MRI pelvis w/o dye
	74181	MRI abdomen w/o dye
	74182	MRI abdomen w/dye
	74183	MRI abdomen w/o & w/dye
	72196	MRI pelvis w/dye
	72198	MRI angiography pelvis w/o & w/dye
Myocardial Imaging	78465	Myocardial Imaging, multiple studies
	78451	Myocardial Imaging, single studies
	78452	Myocardial Imaging, single studies rest
	78453	Myocardial Imaging, single studies planar
	78454	Myocardial Imaging, multiple studies planar
	78460	Myocardial Imaging, single studies planar quantification
	78461	Myocardial Imaging, multiple studies planar quantification
	78464	Myocardial Imaging, single studies
Shoulder X-Ray	73030	X-ray shoulder
	73000	X-ray collar bone
	73010	X-ray shoulder blade
	73050	X-ray shoulder joints
Spine X-Ray	72100	X-ray lower spine, 2 or 3 views
	72040	X-ray neck spine, 2 or 3 views
	72050	X-ray neck spine, $\geq$ 4 views
	72052	X-ray neck spine, complete
	72110	X-ray lower spine, $\geq$ 4 views
	72114	X-ray lower spine, complete
	72120	X-ray lower spine, bending complete
	72090	X-ray exam of trunk spine
	72072	X-ray exam of thoracic spine
Ultrasound - Breast	76645	Ultrasound breast, real time
	76641	Ultrasound breast, complete
	76642	Ultrasound breast, limited
Ultrasound - Pelvic	76856	Ultrasound pelvic, complete
	76700	Ultrasound abdomen, complete
	76705	Ultrasound abdomen
	76857	Ultrasound pelvic, limited
	76830	Ultrasound transvaginal non-obstetrics
Ultrasound - Pregnancy	76805	Obstetrics ultrasound $\geq$ 14 wks, sngl fetus
	76817	Obstetrics ultrasound, transvaginal
	76816	Obstetrics ultrasound, follow-up, per fetus
	76802	Obstetrics ultrasound < 14 wks, addl fetus
	76801	Obstetrics ultrasound < 14 wks, single fetus
	76811	Obstetrics ultrasound, detailed, sngl fetus
	76810	Obstetrics ultrasound, $\geq$ 14 wks, addl fetus
	76812	Obstetrics ultrasound, detailed, addl fetus
Wrist X-Ray	73110	X-ray exam of wrist, $\geq$ 3 views
	73130	X-ray exam of wrist/hand
	73090	X-ray exam of forearm
	73100	X-ray exam of wrist, 2 views

Notes: Outdated procedure codes and codes not used during the sample period are not shown.

Table A2: The Effect of Price Transparency Website on Visit Price
Robustness Estimates

	No Controls	No Insurer Controls	Insurer FE $\times$ Post Interaction	With Trends	Individual Fixed Effects	Inverse Hyperbolic Sine
			Dep Var: Patient Out-of-Pocket Price			
$\text{OnWeb}_m \times \text{Post}_t$	-0.069^{***}	-0.062^{***}	-0.072^{***}	-0.031^{***}	-0.058^{***}	-0.077^{***}
	(0.010)	(0.010)	(0.010)	(0.012)	(0.012)	(0.011)
Mean level	86.62	86.64	86.50	86.50	86.50	86.50
Adjusted R2	0.145	0.358	0.376	0.376	0.406	0.367
Observations	3,592,076	3,579,071	3,404,442	3,403,645	3,415,266	3,404,442
			Dep Var: Insurer Paid Amount			
$\text{OnWeb}_m \times \text{Post}_t$	-0.030^{***}	-0.039^{***}	-0.037^{***}	-0.050^{***}	-0.026^{***}	-0.039^{***}
	(0.005)	(0.005)	(0.005)	(0.012)	(0.006)	(0.006)
Mean level	655.30	86.64	86.50	653.90	653.90	653.90
Adjusted R2	0.151	0.260	0.292	0.295	0.415	0.283
Observations	3,592,076	3,579,071	3,404,442	3,403,645	3,415,266	3,404,442
Dep Var Transformation	$\log(1+y)$	$\log(1+y)$	$\log(1+y)$	$\log(1+y)$	$\log(1+y)$	$\sinh^{-1} y$
Indiv. Controls	No	Yes	Yes	Yes	No	Yes
Individual FE	No	No	No	No	Yes	No
Insurer FE	No	No	Yes	Yes	No	Yes
Insurer FE $\times$ Post	No	No	Yes	No	No	No
Month-Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Procedure FE	Yes	Yes	Yes	Yes	Yes	Yes
Procedure Category Trend	No	No	No	Yes	No	No

Notes: The unit of observation is a patient visit, which may contain multiple medical claims. The sample consists of all commercial claims related to outpatient radiology procedures in the state of New Hampshire over the period 2005 to 2011. OLS regression standard errors clustered at the month-year level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A3: The Effect of Price Transparency Website on Visit Price
Falsification Test

	All	No Deductible	Deductible	
			Not Past	Past
	Dep Var: Log(1+Patient Out-of-Pocket Cost)			
$\text{OnWeb}_{\text{m}} \times \text{Post}_{\text{t}}$	0.006	0.006	-0.007	-0.050
	(0.014)	(0.014)	(0.024)	(0.031)
Mean level	60.09	12.92	154.92	33.17
Adjusted R2	0.336	0.252	0.237	0.110
Observations	978,787	545,312	307,175	126,300
	Dep Var: Log(1+Insurer Paid Amount)			
$\text{OnWeb}_{\text{m}} \times \text{Post}_{\text{t}}$	0.007	0.014^{*}	0.021	-0.005
	(0.007)	(0.007)	(0.016)	(0.013)
Mean level	549.57	532.51	517.16	702.04
Adjusted R2	0.299	0.366	0.255	0.415
Observations	978,787	545,312	307,175	126,300
Geographic Controls	Yes	Yes	Yes	Yes
Charlson Comorbidity FE	Yes	Yes	Yes	Yes
Insurer FE	Yes	Yes	Yes	Yes
Insurance Org Type FE	Yes	Yes	Yes	Yes
Month-Year FE	Yes	Yes	Yes	Yes
Procedure (CPT) FE	Yes	Yes	Yes	Yes

Notes: In the above specification, post is defined as the period beginning March 2006, a year before the website actually launched. The unit of observation is a patient visit, which may contain multiple medical claims. The sample consists of all commercial claims related to outpatient radiology procedures in the state of New Hampshire over the period 2005 to 2007. OLS regression standard errors clustered at the month-year level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A4: The Effect of Price Transparency Website on Price Trend

	Patient Out-of-Pocket Price	Insurer Paid Amount
$\text{OnWeb}_m \times \text{YearsSinceWebsite}_t$	-0.038^{***}	-0.001
	(0.003)	(0.003)
Indiv. Controls	Yes	Yes
Procedure FE	Yes	Yes
Month-Year FE	Yes	Yes
Mean level	86.50	5.56
Adjusted R2	0.373	0.291
Observations	3,404,442	3,404,442

Notes: The sample is the same as the baseline specification. OLS regression standard errors clustered at the month-year level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A5: The Effect of Price Transparency Website on Out-of-Pocket Price by Insurer Switching

	Individual Switches Insurers Post-Website	Individual Does Not Switch Insurers Post-Website
$\text{OnWeb}_m \times \text{Post}_t$	-0.064^{***}	-0.086^{***}
	(0.010)	(0.018)
Full Controls	Yes	Yes
Month-Year FE	Yes	Yes
Procedure FE	Yes	Yes
F statistic of diff.	1.102	
Adjusted R2	0.374	0.346
Observations	2,844,724	559,718

Notes: The dependent variable is Log(1+Patient Out-of-Pocket Price). The sample of individuals that switch insurers includes those that gain private insurance post-website. OLS regression standard errors clustered at the month-year level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Appendix F: Derivation of Website Usage Benefit

Here I derive the approximate benefit of using the price transparency website using a second-order multivariate Taylor series around the expectation.[36] The primary reason for

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

Table A6: The Price Transparency Website and Insurance Composition by Insurance Type and Insurance Firm

	Dependent Variable:				
	Anthem	Cigna	Harvard Pilgrim	Matthew Thornton	Other
Post_t	$8,569$	$-4,624$	$3,475$	$-6,430$	-987
	$(14,916)$	$(10,769)$	$(14,996)$	$(6,701)$	$(4,173)$
Mean level	163,215	102,042	69,845	100,457	50,582
Adjusted R2					
Observations	48	48	48	48	48
	PPO	POS	HMO	EPO	Indemnity
Post_t	$10,782$	$-1,572^{**}$	$-4,233$	$-4,551$	-422
	$(2,272)$	(53)	$(1,907)$	$(4,266)$	$(4,602)$
Mean level	116,592	66,921	229,916	41,063	31,650
Adjusted R2	0.987	0.911	0.868	0.938	0.809
Observations	48	48	48	48	48
Indiv. Controls	Yes	Yes	Yes	Yes	Yes
Month Trend	Yes	Yes	Yes	Yes	Yes

Notes: Outcome is number of individuals enrolled in each insurance type or insurance firm in each month. OLS regression with robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

using this approximation is that the expression must be computed in every iteration of the estimation procedure. Therefore, simulating the high-dimensional integral is computationally infeasible.

If we wish to approximate the first moment of the function $f(x_1, x_2, \dots, x_N)$ 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

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

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

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

Since price signals are independent

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

Furthermore, note that

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

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

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

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

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

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

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

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

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

APPENDIX G: DETAILS ON BAYESIAN INTERPRETATION OF THE DEMAND MODEL AND ESTIMATION

In this section, I present the model as it is estimated in a Bayesian framework. As I describe in Section 2.3.4, reformulating the model in this way allows for a computationally feasible estimation strategy that takes advantage of recent advances in Bayesian estimation but does not change the structural interpretation of parameters.

In general, the posterior is defined as

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

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 $\mathbf{e}_{ikmt}$:

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

The simplified likelihood, conditional on unobservables, is then

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

Note that, unlike Equation 2.24, this is a closed-form expression. This likelihood function is than augmented with the following:

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

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

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

The remaining parameters are given uninformative priors.

APPENDIX H: DETAILS ON BARGAINING FIRST ORDER CONDITION

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

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

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

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

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

Now solving for the markup:

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

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

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

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

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

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

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

Where the partial derivative of consumer surplus is:

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

$$+ \sum_{j'\in\mathcal{N}_{kmt}} [w_{ikmt} e_{ij'kmt} + (1-w_{ikmt})(\bar{p}^{OOP}_{kmt} - p^{OOP}_{ij'kmt})] s_{ij'kmt}$$

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

And the partial derivative of insurer cost is:

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

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

APPENDIX I: DETAILS ON PROCEDURE FOR SIMULATING EQUILIBRIUM PRICES

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

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

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

The iterative algorithm proceeds as follows:

1. Find p^*_{jkmt} $\forall j, k, m, t$
2. Update $\mathbf{p}'_{kmt}$[37]
3. Iterate until all prices are optimal given all other prices (i.e. until prices converge within a mean tolerance of \$1).

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

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

APPENDIX J: STRUCTURAL MODEL APPENDIX TABLES

Table A7: List of Medical Imaging Procedures

CPT Code	Short Description	On Website
70100	X-ray exam of jaw	
73610	X-ray exam of ankle	×
73615	X-ray of ankle contrast	
70160	X-ray exam of nasal bones	
70030	X-ray eye for foreign body	
73050	X-ray exam of shoulders	×
75680	X-rays artery, neck	
74415	X-ray, urinary tract contrast	
73520	X-ray exam of hips	
71035	X-ray chest	×
71110	X-ray exam of ribs	
73630	X-ray exam of foot	×
73010	X-ray exam of shoulder blade	×
72110	X-ray exam of lower spine	×
71021	X-ray chest	
75519	X-ray heart/catheterization	
75731	X-rays artery, adrenal gland	
70210	X-ray exam of sinuses	
75724	X-rays artery, kidneys	
72220	X-ray exam of tailbone	
77072	X-rays for bone age	
73525	X-ray of hip with contrast	
70330	X-ray exam of jaw joints	
73080	X-ray exam of elbow	
71130	X-ray exam of breastbone	
73562	X-ray exam of knee, 3	×
73660	X-ray exam of toe(s)	
72010	X-ray exam of spine	
74405	X-ray urinary tract contrast	
74320	X-ray of bile ducts contrast	
73580	X-ray of knee joint contrast	
72170	X-ray exam of pelvis	
75665	X-rays artery, head & neck	
72090	X-ray exam of trunk spine	×
75716	X-ray artery, arms/legs	
73040	X-ray of shoulder w/contrast	
72100	X-ray exam of lower spine	×
75805	X-ray lymph vessel, trunk	
74301	X-rays at surgery add-on	
74305	X-ray bile ducts/pancreas	
73000	X-ray exam of collar bone	×
72120	X-ray exam of lower spine	×
70200	X-ray exam of eye sockets	
71120	X-ray exam of breastbone	
70130	X-ray exam of mastoids	
70260	X-ray exam of skull	
75774	X-ray artery, each vessel	
71010	X-ray chest	×
72080	X-ray exam of trunk spine	
74000	X-ray exam of abdomen	
71020	X-ray chest	×
75671	X-ray artery, head & neck	
73110	X-ray exam of wrist	×
75660	X-ray artery, head & neck	
75722	X-ray artery, kidney	
73030	X-ray exam of shoulder	×
73092	X-ray exam of arm, infant	
75500	X-ray cinema heart vessels	
75676	X-ray artery, neck	
72040	X-ray exam of neck spine	×
74400	X-ray urinary tract contrast	
75741	X-ray artery, lung	
75662	X-ray artery, head & neck	
73542	X-ray exam, sacroiliac joint	
75705	X-ray artery, spine	
73560	X-ray exam of knee, 1 or 2	×
75807	X-ray lymph vessel, trunk	
73565	X-ray exam of knees	×
70220	X-ray exam of sinuses	
74300	X-ray bile ducts/pancreas	
75650	X-ray artery, head & neck	
75523	X-ray heart /catheterization	
73540	X-ray exam of pelvis & hips	
73650	X-ray exam of heel	×
75756	X-ray artery, chest	
75736	X-ray artery, pelvis	
75658	X-ray artery, arm	
73140	X-ray exam of finger(s)	×
73115	X-ray of wrist contrast	×
77071	X-ray stress view	
70360	X-ray exam of neck	
71100	X-ray exam of ribs	
71022	X-ray chest	×
73130	X-ray exam of hand	×
75743	X-ray artery, lungs	
75726	X-ray artery, abdomen	
73600	X-ray exam of ankle	
73085	X-ray of elbow contrast	
75803	X-ray lymph vessel,arms/legs	
73060	X-ray exam of humerus	
73090	X-ray exam of forearm	×
71111	X-ray exam of ribs/chest	
70150	X-ray exam of facial bones	
75685	X-ray artery, spine	
71030	X-ray chest	
73550	X-ray exam of thigh	
73564	X-ray exam, knee, 4 or more	×
70110	X-ray exam of jaw	
75733	X-ray artery, adrenals	
70250	X-ray exam of skull	
72069	X-ray exam of trunk spine	
72050	X-ray exam of neck spine	×
75801	X-raylymph vessel, arm/leg	
73100	X-ray exam of wrist	×
73592	X-ray exam of leg, infant	
71101	X-ray exam of ribs/chest	×
75710	X-ray artery, arm/leg	
72052	X-ray exam of neck spine	×
72072	X-ray exam of thoracic spine	×
73510	X-ray exam of hip	
72114	X-ray exam of lower spine	×
75527	X-ray heart/catheterization	
70332	X-ray exam of jaw joint	
73620	X-ray exam of foot	×
72200	X-ray exam sacroiliac joints	
73590	X-ray exam of lower leg	
74020	X-ray exam of abdomen	
75573	CT heart w/3d image congen	
75572	CT heart w/3d image	
72128	CT chest spine w/o dye	
75574	CT angio hrt w/3d image	
72130	CT chest spine w/o & w/dye	
72127	CT neck spine w/o & w/dye	
74150	CT abdomen w/o dye	×
70498	CT angiography, neck	
73702	CT lwr extremity w/o&w/dye	
72129	CT chest spine w/dye	
72133	CT lumbar spine w/o & w/dye	
73701	CT lower extremity w/dye	
70490	CT soft tissue neck w/o dye	
70480	CT orbit/ear/fossa w/o dye	
72125	CT neck spine w/o dye	
74178	CT abd & pelv 1/> regns	×
70482	CT orbit/ear/fossa w/o&w/dye	
73201	CT upper extremity w/dye	
73202	CT uppr extremity w/o&w/dye	
71260	CT thorax w/dye	×
73206	CT angio upr extrm w/o&w/dye	
71270	CT thorax w/o & w/dye	×
71275	CT angiography, chest	
73700	CT lower extremity w/o dye	
70486	CT maxillofacial w/o dye	
70460	CT head/brain w/dye	
70488	CT maxillofacial w/o & w/dye	
70487	CT maxillofacial w/dye	
75571	CT hrt w/o dye w/ca test	
72194	CT pelvis w/o & w/dye	×
72193	CT pelvis w/dye	×
73200	CT upper extremity w/o dye	
72192	CT pelvis w/o dye	×
74175	CT angio abdom w/o & w/dye	
70450	CT head/brain w/o dye	
76360	CT scan for needle biopsy	
75635	CT angio abdominal arteries	
74177	CT abd & pelv w/contrast	×
74170	CT abdomen w/o & w/dye	×
73706	CT angio lwr extr w/o&w/dye	
76497	CT procedure	
70481	CT orbit/ear/fossa w/dye	
70492	CT sft tsue nck w/o & w/dye	
71250	CT thorax w/o dye	×
74174	CT angio abd&pelv w/o&w/dye	
72126	CT neck spine w/dye	
70491	CT soft tissue neck w/dye	
72191	CT angiograph pelv w/o&w/dye	
74176	CT abd & pelvis w/o contrast	×
74160	CT abdomen w/dye	×
70470	CT head/brain w/o & w/dye	
70496	CT angiography, head	
72132	CT lumbar spine w/dye	
72131	CT lumbar spine w/o dye	
72158	MRI lumbar spine w/o & w/dye	×
70544	MRI angiography head w/o dye	
70557	MRI brain w/o dye	
75559	MRI cardiac w/stress img	
72157	MRI chest spine w/o & w/dye	×
72148	MRI lumbar spine w/o dye	×
73222	MRI joint upr extrem w/dye	
70540	MRI orbit/face/neck w/o dye	
72146	MRI chest spine w/o dye	×
72196	MRI pelvis w/dye	×
72198	MRI angio pelvis w/o & w/dye	×
73225	MRI angio upr extr w/o&w/dye	
75553	MRI heart for morph w/dye	
71550	MRI chest w/o dye	
73721	MRI jnt of lwr extre w/o dye	×
71552	MRI chest w/o & w/dye	
70558	MRI brain w/dye	
74183	MRI abdomen w/o & w/dye	×
70552	MRI brain w/dye	
72149	MRI lumbar spine w/dye	
70545	MRI angiography head w/dye	
73718	MRI lower extremity w/o dye	×
73720	MRI lwr extremity w/o&w/dye	×
71551	MRI chest w/dye	
70551	MRI brain w/o dye	×
70547	MRI angiography neck w/o dye	
70543	MRI orbt/fac/nck w/o & w/dye	
73719	MRI lower extremity w/dye	
73219	MRI upper extremity w/dye	
73723	MRI joint lwr extr w/o&w/dye	×
70553	MRI brain w/o & w/dye	×
75562	MRI card flow/vel w/dye	
72197	MRI pelvis w/o & w/dye	×
75556	MRI cardiac flow mapping	
75558	MRI cardiac flow/velocity	
70559	MRI brain w/o & w/dye	
76390	MRI spectroscopy	
75561	MRI cardiac for morph w/dye	
75552	MRI heart for morph w/o dye	
75555	MRI cardiac/limited study	
74182	MRI abdomen w/dye	×
75557	MRI cardiac for morph	
72147	MRI chest spine w/dye	
70546	MRI angiograph head w/o&w/dye	
73218	MRI upper extremity w/o dye	×
72159	MRI angio spine w/o&w/dye	
72141	MRI neck spine w/o dye	×
73220	MRI uppr extremity w/o&w/dye	×
70549	MRI angiograph neck w/o&w/dye	
73725	MRI ang lwr ext w or w/o dye	
73223	MRI joint upr extr w/o&w/dye	
73722	MRI joint of lwr extr w/dye	
74185	MRI angio, abdom w orw/o dye	
76400	MRI bone marrow	
72195	MRI pelvis w/o dye	×
76498	MRI procedure	
70542	MRI orbit/face/neck w/dye	
72156	MRI neck spine w/o & w/dye	×
75563	MRI cardiac w/stress img & dye	
70336	MRI, jaw joint	
70541	MRI, head (MRA)	
72142	MRI neck spine w/dye	
75565	MRI cardiac veloc flow mapping	
74181	MRI abdomen w/o dye	×
70548	MRI angiography neck w/dye	
75560	MRI cardiac flow/vel/stress	
73221	MRI joint upr extrem w/o dye	
75564	MRI hrt w/flo/vel/strs & dye	
71555	MRI angio chest w or w/o 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 A8: Potential Savings from Switching Providers

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

Notes: Chart shows potential savings from switching to providers with lower negotiated prices within individual's choice set relative to observed choices. See Section 2.2.1 for definition of choice set. If provider in first quartile is more expensive then chosen provider, simulation assumes individuals do not switch.

Table A9: Estimates from Multinomial Logit Demand Model

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

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

Table A10: Counterfactual Provider Markups and Surplus

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

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

APPENDIX K: STRUCTURAL MODEL APPENDIX FIGURES

Figure A4: Cost Sharing and Deductible Status
By Month of Year

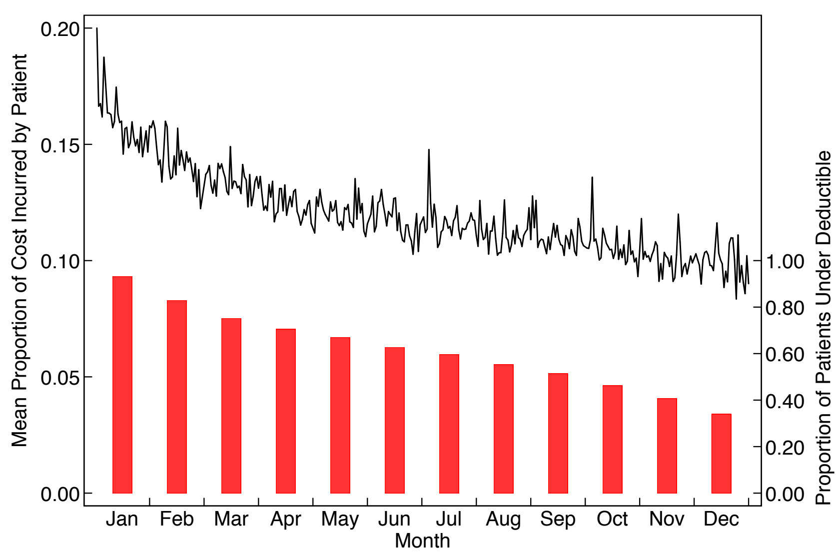

Notes: Bars show the percent of medical imaging consumers that have not hit their deductible conditional on having a plan with a deductible. Line shows the unconditional mean percent of total cost paid by the patient for each day of the year.

Figure A5: Prices Variation within Individuals' Choice Sets

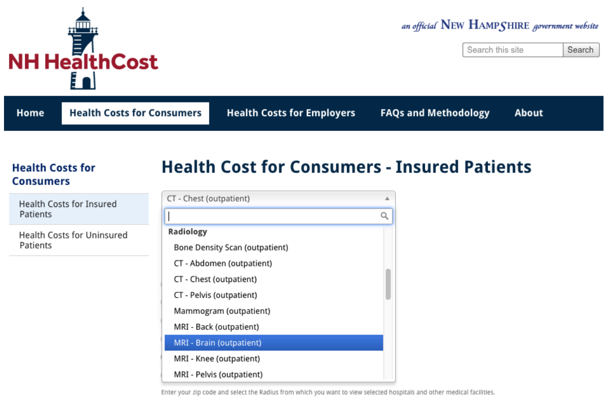

(a) Entering Procedure Information

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

(b) List of Prices

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

Figure A6: Distribution of Estimated Website Usage Cost

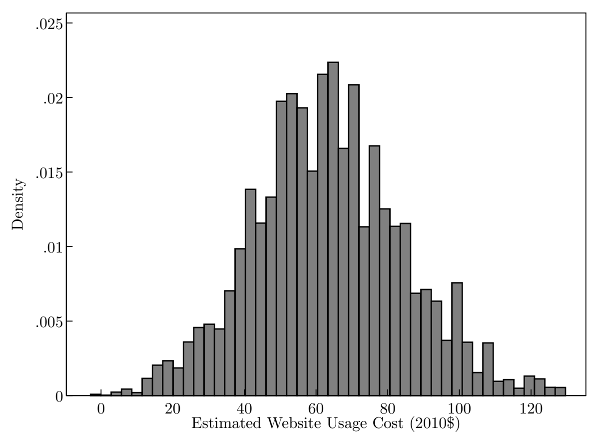

Figure A7: Distribution of Counterfactual Negotiated Prices

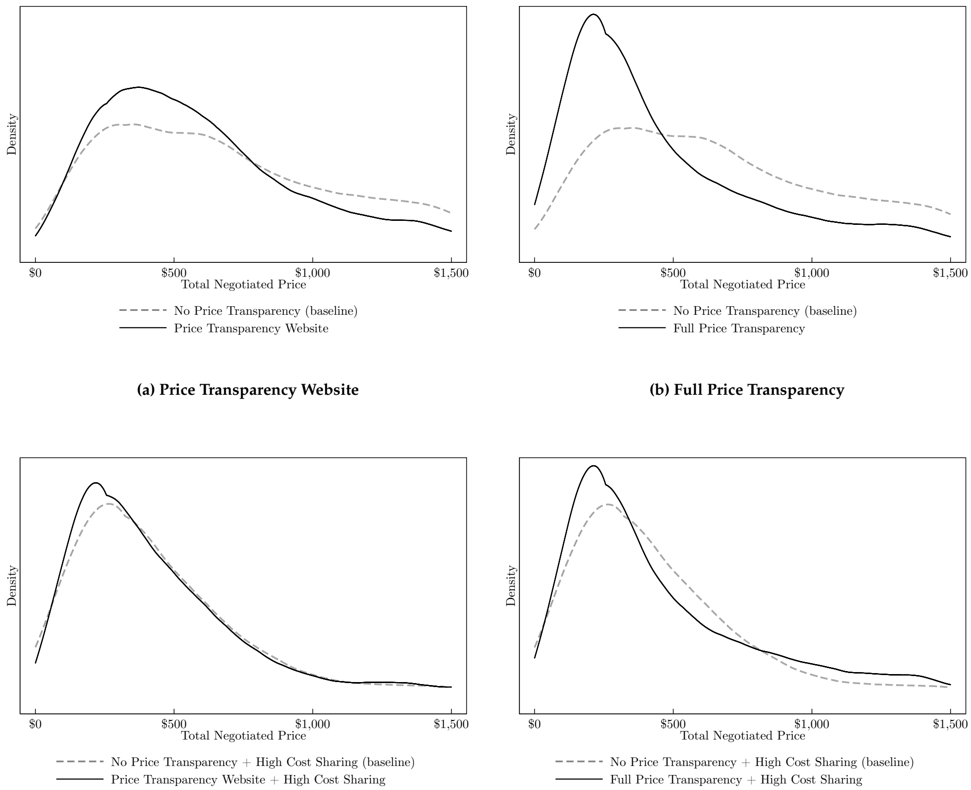

(a) Price Transparency Website

(b) Full Price Transparency

(c) Price Transparency Website + High Cost Sharing

(d) Full Price Transparency + High Cost Sharing

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