

NON-PARAMETRIC ESTIMATION OF SEARCH COSTS FOR DIFFERENTIATED PRODUCTS: EVIDENCE FROM MEDIGAP*

Haizhen Lin[†]

Matthijs R. Wildenbeest[‡]

December 2017

Abstract

This paper develops a method to estimate search frictions as well as preference parameters in differentiated product markets. Search costs are non-parametrically identified, which means our method can be used to estimate search costs in differentiated product markets that lack a suitable search cost shifter. We apply our model to the Medigap insurance market. We find that search costs are substantial: the estimated median cost of searching for an insurer is \$30. We also find that eliminating search costs could result in price decreases of as much as \$63 (or 3.7 percent), along with increases in average consumer welfare of up to \$236.

Keywords: price dispersion, consumer search, product differentiation, health insurance

JEL Classification: I13, D83, L15

*We thank Mike Baye, Kate Bundorf, Andrew Ching, Leemore Dafny, David Dranove, Hanming Fang, Marty Gaynor, Lorens Helmchen, Claudio Lucarelli, Nicole Maestas, Jeff Prince, Jon Skinner, Kosali Simon, Alan Sorensen, and conference participants at the 10th Annual International Industrial Organization Conference (IIOC) and 2012 American Society of Health Economists Conference (ASHEcon) for helpful comments and suggestions.

[†]Indiana University, Kelley School of Business and NBER, E-mail: hzlin@indiana.edu.

[‡]Indiana University, Kelley School of Business, E-mail: mwildenb@indiana.edu (corresponding author).

1 Introduction

Since Stigler's (1961) seminal article on the economics of information, economists have tried to relate price dispersion to search frictions: if it is costly for individuals to obtain price information, because, for instance, finding out the price means visiting the seller or making phone calls, some individuals will not compare offers, which allows firms to raise prices. As a result firms might set high prices to maximize profits from individuals who do not compare offers, or set a relatively low price and maximize surplus from price comparers, resulting in price dispersion.

An additional explanation for price variation is that even in settings in which a specific product is essentially a homogeneous good, firms selling the product might still differ in terms of quality or service-related characteristics, such as branding and service quality, allowing higher-quality firms to set higher prices. Product differentiation and search can also go together: as shown by Wolinsky (1986) and Anderson and Renault (1999), if consumers are searching not only for low prices, but also for horizontal characteristics such as the whether the product is a good fit, equilibrium prices rise with search costs.

The objective of this paper is to develop a model of search and product differentiation that allows us to non-parameterically identify the contributions of search frictions and product differentiation in explaining price dispersion. In Section 2 of the paper we present a model in which consumers search for products that are both horizontally and vertically differentiated. Consumers determine before they search which of the firms to contact for information about products and prices; they do so by making a trade-off between the expected utility of contacting a subset of firms and corresponding search costs. We show how to estimate the model using aggregate data on prices, market shares, and product characteristics. An attractive feature of the model is that search costs are non-parametrically identified, i.e., we do not have to make any specific assumptions about the distribution of search costs among individuals. Although this means that we have to assume individuals are homogeneous in how they value observed characteristics (i.e., we cannot allow for random coefficients), it does allow us to separately identify search cost and preference heterogeneity without using additional data on search behavior or search cost shifters, which is often not available to researchers.

In Section 3 we use the model to study price dispersion in the Medigap insurance market. Medigap is a form of private insurance designed to supplement Medicare by filling in the coverage gaps in Medicare payments. One important feature of the Medigap market is plan standardization, meaning that within a given plan type, the same set of benefits and coverage needs to be offered

regardless of which insurer sells it. Irrespective of plan standardization, we observe substantial price variation across insurers within a market, which we define at the state level.¹ We first present reduced-form evidence that suggests both product differentiation and search frictions play important roles in explaining observed price variation. We then estimate the model using aggregate data on prices and market shares.

Our estimates indicate that search costs are substantial; the estimated parameters of our main specification indicate that median search costs are \$30 per search. In addition there is a lot of variation in search costs across individuals. We also show how to use our estimates to study the competitive effects of lowering search costs. Assuming insurers set prices to maximize profits, we determine to what extent prices would change if quotes are obtained at no cost. According to our simulation results average prices decrease by \$63, which is 3.7 percent of the average yearly policy premium. Consumer welfare, which also includes savings on search costs as well as the expansion of the market, increases by up to \$236 on average if search costs are zero.

The main contribution of our paper is to develop a framework that allows for the non-parametric identification of search costs in settings in which product differentiation is important. In Section 4 of the paper we use Monte Carlo experiments to show that our estimation method is indeed capable of recovering preference parameters as well as the search cost distribution, without having to make any parametric assumptions on the search cost distribution. In this section we also study the performance of our model when the data is generated using a richer framework in which consumers have different preferences for observed characteristics (as in Moraga-González, Sándor, and Wildenbeest, 2015). We find that our model does reasonably well in terms of the estimation of utility parameters when the variance of the choice-set specific search cost shock, which is included in Moraga-González, Sándor, and Wildenbeest (2015) to keep estimation of their model tractable, is small.

Related literature

Our paper is part of the consumer search literature. Within this literature, several papers have studied search behavior in insurance markets. Brown and Goolsbee (2002) show that increased usage of Internet price comparison sites, which reduces search costs, led to lower prices for term life policies. Cebul et al. (2011) find that search frictions in private health insurance markets lead to high prices and price dispersion, excessive marketing costs, and high insurance turnover. They

¹For instance, a 65-year-old woman living in Indiana could expect to pay anywhere between \$1,223 and \$3,670 for a Medigap Plan F policy in 2009, depending on her choice of insurer. Such a wide range of prices is not limited to this state only: the average coefficient of variation within states is 0.27 for Plan F, which is substantial.

also suggest that government-financed public insurance can reduce distortions created by search frictions.

Several recent papers have developed methods to structurally estimate search models. Hortaçsu and Syverson (2004) develop a model of product differentiation and search and find that both are important determinants of fee dispersion in the retail S&P 500 index funds sector. Wildenbeest (2011) provides a framework for studying price dispersion in markets with product differentiation and search frictions and shows how to estimate search costs using only price data. In both these papers it is not necessary to make any assumptions on the search cost distribution to obtain search cost estimates, as in our model. However, an important difference between these papers and our model is that consumers search for vertically differentiated firms in Hortaçsu and Syverson (2004) and Wildenbeest (2011), whereas firms are horizontally and vertically differentiated in our model. Moreover, in both models consumers search randomly, whereas search is directed in our model. Directed search, or ordered search, is typically thought of as a natural way to search in markets in which certain firms are more prominent than others (see Armstrong, 2017, for a discussion of the theoretical literature on ordered search).

Our paper most closely relates to a study by Moraga-González, Sándor, and Wildenbeest (2015), who add search to the demand estimation framework of Berry, Levinsohn, and Pakes (1995). Moraga-González, Sándor, and Wildenbeest (2015) estimate the model using automobile data from the Netherlands. To be able to separately identify heterogeneity in search costs from heterogeneity in preference parameters, they link search costs to the distances between consumers and dealer locations. Our model builds on their findings; however, by working in a conditional logit rather than a mixed logit framework, search costs can be non-parametrically identified, which is useful if the data lack a suitable search cost shifter, as in our application. We discuss differences between our model and theirs in more detail in Section 4, and use Monte Carlo experiments to test the performance of our model with data that is generated using a framework similar to Moraga-González, Sándor, and Wildenbeest (2015).

Our paper also fits into the literature that studies Medigap prices and price dispersion.² Robst (2006) adopts a hedonic pricing model in order to examine the determinants of Medigap premiums. Our paper is closely related to Maestas, Schroeder, and Goldman (2009), who use the search model of Carlson and McAfee (1983), which is also used in Hortaçsu and Syverson (2004), to model price

²The earlier literature on Medigap has focused mostly on adverse selection (e.g., Wolfe and Goddeeris, 1991; Hurd and McGarry, 1997; Ettner, 1997; Finkelstein, 2004). A recent study by Starc (2014) finds that adverse selection can actually lower markups and thus increase consumer welfare. Fang, Keane, and Silverman (2008) instead find evidence of advantageous selection in Medigap and they identify various sources of advantageous selection such as income and cognitive ability.

dispersion for Medigap plans. They find that the average search cost in this market is \$72. Although we find median search costs to be somewhat lower than this number, unlike Maestas, Schroeder, and Goldman (2009) we model the joint decision of which plan type to obtain and which insurer to choose. To allow for differences in plan-type preferences across individuals, our utility specification includes a stochastic utility shock, which we model as ex-ante unobserved by the decision maker. This means that in our model consumers search for prices as well as a good fit in terms of insurer and plan type. A second difference between our model and that of Maestas, Schroeder, and Goldman (2009) is that search in their model is random, while search in our model is directed, which means that individuals rank insurers according to expected utility, and, depending on the decision-maker's search costs, contact an optimal set of the highest-ranked insurers to obtain information about prices and plan-type matches. In differentiated product markets in which consumers have some prior information on the characteristics of the firms, the optimal way to search is to start with the alternative that provides the highest expected utility (see, for instance, Weitzman, 1979; Chade and Smith, 2006). In Maestas, Schroeder, and Goldman (2009) consumers have no prior information on the parameters of the utility function, which means ex-ante all alternatives are similar and consumers search randomly.³ In a random search model lower ranked firms are as likely to be part of a consumer's choice set as higher ranked firms, whereas in our directed search model a consumer needs to have low search costs to sample a lower ranked firm. To explain demand at lower ranked firms search costs therefore have to be lower in our search model than in the random search model, which means that search cost estimates will likely be higher in a random search model. Finally, whereas Maestas, Schroeder, and Goldman (2009) assume that search costs are uniformly distributed, our model does not require any assumptions on the shape of the search cost distribution, which can be an advantage if one does not know a priori how the search cost distribution looks like. Moreover, search costs are typically interpreted as being related to income—since the income distribution tends to follow a lognormal distribution, assuming search costs follow a uniform distribution may be restrictive.

2 Model

In this section we develop an theoretical and empirical framework that takes search and product differentiation as its main ingredients. As in Anderson and Renault (1999) and Wolinsky (1986),

³Especially at the state level, the Medigap market is dominated by a few insurers with large market shares, which is difficult to explain if product differentiation would not be playing a role. Medigap insurers tend to be active in other markets as well and spend significant amounts on advertising. This in combination with word of mouth effects makes it likely that consumer have some prior knowledge of at least some of the characteristics of the firms.

our model also allows for horizontal product differentiation in order to capture uncertainty about which of the products offered by a firm are best suited for an individual. This means that in our model consumers search to find out price information as well as to figure out which of the products offered by the competing firms is a good match.

2.1 Consumer Optimal Search

Consumers i derive utility from consumption of a product $j \in \{1, 2, \dots, J\}$ sold by firm $f \in \{1, 2, \dots, F\}$ according to:

$$u_{ijf} = x'_{jf}\beta - \alpha p_{jf} + \xi_{jf} + \varepsilon_{ijf} = \delta_{jf} + \varepsilon_{ijf},$$

where x_{jf} are characteristics observed by both the researcher and the consumer, ξ_{jf} is a vertical characteristic observed by the consumer only, and ε_{ijf} is a matching term unobserved by both the researcher and the consumer, which follows a standard Type I Extreme Value distribution. The mean utility of product j sold by firm f is given by δ_{jf} . The utility of not buying any of the products is

$$u_{i0} = \varepsilon_{i0}.$$

We assume that consumers search nonsequentially. Consumers have information about x_{jf} and ξ_{jf} , but not matching parameters ε_{ijf} and prices p_{jf} , and they have to visit different firms to discover these parameters. Search happens at the firm level, i.e., by visiting a firm f the matching parameters and prices of all products sold by this firm are revealed.⁴ Search costs are randomly distributed in $[0, \infty)$ according to the cumulative distribution function (CDF) $G(c)$.

Since consumers do not know the realization of ε_{ijf} and p_{jf} before they search, they rank firms according to their mean utility ϕ_f , where ϕ_f is the expected maximum utility of the set of products

⁴We believe that given the nature of the search process it is more natural to model search as first deciding on which firms to search for, and then deciding which product to purchase. For instance, to get more information consumers typically have to visit a store, which means that search costs tend to occur at the firm level and to a lesser extent at the product level (a store visit would allow consumers to obtain information on all products sold by the store).

\mathcal{G}_f on offer by firm f ,⁵ i.e.,

$$\begin{aligned}\phi_f &= E \left[\max_{j \in \mathcal{G}_f} \{u_{ijf}\} \right] = \int_{-\infty}^{\infty} u \frac{d}{du} \left(\prod_{j \in \mathcal{G}_f} F(u - \delta_{jf}) \right) du \\ &= \int u \frac{d}{du} \left(\prod_{j \in \mathcal{G}_f} \exp[-\exp[-(u - \delta_{jf})]] \right) du \\ &= \log \left(\sum_{j \in \mathcal{G}_f} \exp[\delta_{jf}] \right).\end{aligned}$$

Note that we use the standard assumption in the theoretical literature on search for differentiated products that consumers form rational expectations about prices (as in Wolinsky, 1986; Anderson and Renault, 1999). This means that even though consumers do not observe realized prices before searching, they use expected (equilibrium) prices to rank firms.

Since the ranking is ex-ante the same for all consumers, we can index firms by their expected maximum utility ϕ_f , i.e, firm $f = 1$ is the most attractive firm, firm $f = 2$ the second-most attractive, and so on. We assume consumers do not observe ε_{i0} before searching so they will not condition their search behavior on the utility of the outside option. Furthermore, we assume that the first observation is free, so that all consumers will search at least once.⁶

An individual consumer with search cost c chooses the number of firms to sample, denoted $k^*(c)$, to maximize her expected utility. That is,

$$k^*(c) = \arg \max_k \{E[\max\{u_{ij1}, u_{ij2}, \dots, u_{ijk}\}] - (k - 1)c\},$$

where $(k - 1)$ reflects that the first observation is obtained for free. As shown by Moraga-González, Sándor, and Wildenbeest (2015), assuming that ε_{ijf} follows a Type I Extreme Value distribution

⁵Note that we have left out the Euler constant from δ_f , since it is common to all firms and therefore does not affect choices.

⁶Should the utility of the outside option be observed before searching, search decisions would have to be conditional on the utility of the outside option, which would greatly complicate the analysis. To make the model tractable, we therefore assume that the outside option is not observed before searching. Also note that the assumption that the first observation is for free does not mean that consumers make decisions conditional on having observed the first observation. This assumption, which is standard in the literature, is only made to assure that everyone in the market participates. Without this assumption, the market share expression for the outside option would include the share of consumers not searching (i.e., μ_0). To separately identify this parameter from consumers who do search but decide not to buy, requires additional information on the proportion of consumers not searching.

allows us to write the optimal number of quotes to obtain as

$$k^*(c) = \arg \max_k \left\{ \log \left(1 + \sum_{f=1}^k \exp \phi_f \right) - (k-1)c \right\}.$$

We define c_1 as the consumer who is indifferent regarding whether to obtain quotes from one firm or from two firms, i.e.,

$$\ln(1 + e^{\phi_1}) = \ln(1 + e^{\phi_1} + e^{\phi_2}) - c_1.$$

Solving for c_1 gives

$$c_1 = \ln \left(1 + \frac{e^{\phi_2}}{1 + e^{\phi_1}} \right).$$

More generally, we define c_k as the consumer who is indifferent regarding whether to search k or $k+1$ times, i.e.,

$$\ln \left(1 + \sum_{f=1}^k e^{\phi_f} \right) - (k-1)c_k = \ln \left(1 + \sum_{f=1}^{k+1} e^{\phi_f} \right) - kc_k.$$

Solving for c_k gives

$$c_k = \ln \left(1 + \frac{e^{\phi_{k+1}}}{1 + \sum_{f=1}^k e^{\phi_f}} \right). \quad (1)$$

Note that by definition ϕ_k is decreasing in k , which means that c_k is decreasing in k as well. Moreover, the above equation shows that critical search cost values cannot be negative.

Using the critical search cost values c_k and the search cost distribution $G(c)$ we can calculate the fraction of individuals searching k times. The fraction of individuals searching once is

$$\mu_1 = 1 - G(c_1). \quad (2)$$

The fraction of individuals searching $k \geq 2$ times is

$$\mu_k = G(c_{k-1}) - G(c_k), \quad k = 2, \dots, N-1; \quad (3)$$

$$\mu_N = G(c_{N-1}). \quad (4)$$

2.2 Consumer Demand

We now move to the discussion of aggregate demand. First consider product j sold by the highest-ranked firm (i.e., the firm with the highest mean utility ϕ_f). Since all consumers will visit this firm,

the market share of product j is given by

$$s_{j1} = \frac{e^{\delta_{j1}}}{1 + e^{\phi_1}} \cdot \mu_1 + \frac{e^{\delta_{j1}}}{1 + \sum_{\ell=1}^2 e^{\phi_\ell}} \cdot \mu_2 + \dots + \frac{e^{\delta_{j1}}}{1 + \sum_{\ell=1}^N e^{\phi_\ell}} \cdot \mu_N = \sum_{k=1}^N \frac{e^{\delta_{j1}}}{1 + \sum_{\ell=1}^k e^{\phi_\ell}} \cdot \mu_k.$$

The second-highest-ranked firm will only attract the consumers who search at least twice, so the market share of product j sold by this firm is

$$s_{j2} = \sum_{k=2}^N \frac{e^{\delta_{j2}}}{1 + \sum_{\ell=1}^k e^{\phi_\ell}} \cdot \mu_k.$$

More generally, the market share of product j sold by the f th-highest-ranked firm is

$$s_{jf} = \sum_{k=f}^N \frac{e^{\delta_{jf}}}{1 + \sum_{\ell=1}^k e^{\phi_\ell}} \cdot \mu_k,$$

while the overall market share of the f th-highest-ranked firm is

$$s_f = \sum_{k=f}^N \frac{e^{\phi_f}}{1 + \sum_{\ell=1}^k e^{\phi_\ell}} \cdot \mu_k.$$

Using the market share of the outside good, i.e.,

$$s_0 = \sum_{k=1}^N \frac{1}{1 + \sum_{\ell=1}^k e^{\phi_\ell}} \cdot \mu_k,$$

we can rewrite s_{j1} as

$$s_{j1} = s_0 \cdot e^{\delta_{j1}}, \tag{5}$$

or, by taking logs and rearranging,

$$\ln s_{j1} - \ln s_0 = \delta_{j1}.$$

Notice that the relation between the market share of the outside option and the market share for the highest-ranked firm is the same as in a standard demand model with no search frictions. This is because we assume the first search is free so by construction, all consumers search the highest-ranked firm. However, market share expressions for lower-ranked firms will be different since we need to take into account that because of search frictions, these firms will not be searched by all consumers. For example, using the definition of s_0 , we can rewrite the market share for the

second-highest-ranked firm s_{j2} as

$$s_{j2} = \left(s_0 - \frac{\mu_1}{1 + e^{\phi_1}} \right) \cdot e^{\delta_{j2}}. \quad (6)$$

Using s_0 , the overall market share of the highest-ranked firm is

$$s_1 = s_0 \cdot e^{\phi_1},$$

which can be rewritten as

$$1 + e^{\phi_1} = \frac{s_0 + s_1}{s_0},$$

which can be plugged into equation (6) to get

$$s_{j2} = s_0 \left(1 - \frac{\mu_1}{s_0 + s_1} \right) \cdot e^{\delta_{j2}}.$$

Taking logs and rearranging gives

$$\ln s_{j2} - \ln s_0 = \delta_{j2} + \ln \left(1 - \frac{\mu_1}{s_0 + s_1} \right).$$

Similarly, the market share equation for a product sold by the third-highest-ranked firm is

$$\begin{aligned} \ln s_{j3} - \ln s_0 &= \delta_{j3} + \ln \left[\left(1 - \frac{\mu_1}{s_0 + s_1} \right) \left(1 - \frac{\mu_2}{s_0 + s_1 + s_2 - \mu_1} \right) \right] \\ &= \delta_{j3} + \ln \left(1 - \frac{\mu_1}{s_0 + s_1} \right) + \ln \left(1 - \frac{\mu_2}{s_0 + s_1 + s_2 - \mu_1} \right). \end{aligned}$$

More generally, the difference between the log market share of product j sold by the f th-highest-ranked firm and the log market share of the outside good can be written as

$$\ln s_{jf} - \ln s_0 = \delta_{jf} + \sum_{k=1}^{f-1} \ln \left(1 - \frac{\mu_k}{s_0 + \sum_{\ell=1}^k s_\ell - \sum_{\ell=2}^k \mu_{\ell-1}} \right). \quad (7)$$

2.3 Estimation

We use equation (7) to estimate the model in the following way:

$$\ln s_{jf} - \ln s_0 = \beta X_{jf} - \alpha p_j + \gamma_1 R_{f1} + \gamma_2 R_{f2} + \dots + \gamma_{N-1} R_{fN-1} + \xi_j, \quad (8)$$

where R_{fk} is a firm ranking-related dummy that is given by

$$R_{fk} = \begin{cases} 1 & \text{if rank}_f > k; \\ 0 & \text{if rank}_f \leq k. \end{cases}$$

Comparing equation (7) to equation (8) yields

$$\gamma_k = \ln \left(1 - \mu_k / \left(s_0 + \sum_{\ell=1}^k s_\ell - \sum_{\ell=2}^k \mu_{\ell-1} \right) \right),$$

which means the estimated γ_k 's can be used to back out the μ_k 's, which represent the shares of consumers searching k times. Specifically, solving for μ_k gives

$$\mu_k = [1 - \exp(\gamma_k)] \left(s_0 + \sum_{\ell=1}^k s_\ell - \sum_{\ell=2}^k \mu_{\ell-1} \right).$$

The μ_k 's can be obtained iteratively by starting with μ_1 : to calculate μ_1 only γ_1 , s_0 , and s_1 are needed; to calculate μ_2 , we can use the estimate of μ_1 , as well as γ_2 , s_0 , s_1 , and s_2 . To derive a non-parametric estimate of the search cost distribution, we combine the estimates of the μ_k 's with estimates of the search cost cutoffs c_k . For this we can use equation (1) as well as estimates of the firm-specific logit inclusive value ϕ_f . The μ_k 's can then be mapped into the quantiles of the search cost distribution that correspond to the search cost cutoff values by inverting equations (2)–(4).

Note that ϕ_2 cannot be larger than ϕ_1 , so the minimum value of $\ln(1 - \mu_1/(s_0 + s_1))$ is $\ln s_2 - \ln s_1$, while the maximum value is 0. Similarly, the minimum value of $\ln(1 - \mu_2/(s_0 + s_1 + s_2 - \mu_1))$ is $\ln s_3 - \ln s_2$, while the maximum value is again 0. This means that we have to estimate equation (8) using the constraint that $\gamma_k \in [\ln s_{k+1} - \ln s_k, 0]$.

We estimate equation (8) using constrained two-stage least squares. In order to correct for the potential endogeneity of prices, we construct two instrumental variables in the spirit of Hausman (1996) and Nevo (2000). The first instrument is the average price of the same product offered by the firm in other markets. The second is the average price of all products offered by the firm in other markets. The identifying assumption is that market-specific valuations are independent across markets.⁷

Note that the ranking-related dummies are market and firm specific. In carrying out our estimation, we pool firms ranked below a certain threshold for the purpose of restricting the total

⁷Whether this identifying assumption is valid or not is more of an empirical question. In our application presented in the next section, we also use other instruments to check the validity of this assumption.

number of these dummies. For example, in our application, we estimate a model in which consumers can either search any number from one to four, or they search five or more times, which allows them to obtain information on all the firms in the market. Our estimation imposes an implicit assumption that consumers either have relatively high search costs, such that they explicitly decide how many firms to search (up to four), or that they have relatively low search costs, such that they obtain information on all firms in the market.

2.4 Identification

In this section we discuss identification of the parameters of the utility function, which are given by α and β , as well as the parameters that reflect the search part of the model, which are given by the vector of ranking-related parameters γ . Variation in product and firm characteristics together with variation in market share allows us to identify the α and β parameters of the utility function. Variation in market share that cannot be attributed to product and insurer characteristics allows us to identify γ . To better explain how search costs parameters γ are separately identified, let us consider two similar products sold by two different firms, with one firm being slightly more desirable than the other.⁸ In a full information model, these two products could have very similar market shares. However, in a search model, the more desirable firm will be ranked higher and therefore visited by more consumers, which gives the product offered by this firm an additional boost in market share. As a result, differences in market share of similar products offered by firms that share similar attributes would allow us to identify the search costs parameters γ . As search costs increase, the market share of the product offered by the more popular firm will increase relative to the product offered by the less popular firm.

The discussion above assumes we use firm characteristics to capture firm heterogeneity. If we allow for firm fixed effects instead, separately identifying the effects of firm differentiation from the effects of search frictions is more challenging. This is because in our search model, the vector of search cost parameters γ varies across markets. Therefore, separately identifying firm fixed effects from search cost parameters requires observing multiple firms that share the “same” ranking in a given market. In principle, this can be achieved in two ways. First, by construction of the model, firm fixed effects can be identified using variation in identities of the highest-ranked firms across markets (and time). Identification works here by taking advantage of the fact that the highest-ranked firm is part of all choice sets, so by construction no ranking-related dummies are used

⁸Note that in our model consumers might value one firm more than the other even if the attributes of the firms are similar. For example, one firm could offer more product choices, which leads to a higher logit inclusive value for that firm and therefore a higher expected mean utility.

for the highest-ranked firms. A second way in which identification can be achieved is due to our assumption that consumers obtain information from all firms if they search more than a certain cutoff number of firms: by pooling together firms below a certain cutoff ranking, these firms share the same “ranking,” which allows us to separately identify firm fixed effects from the effects of ranking-related dummies. This identification strategy relies on whether it is reasonable to lump all firms together that are ranked below a certain cutoff ranking. In our empirical analysis part, we assume that consumers search up to four firms or search all firms, but we find our results to be largely similar to other thresholds such as five and six. In this particular case we therefore feel comfortable to pool firms ranked five and below together, since this allows us to separately identify firm fixed effects and search costs. More generally, we acknowledge that separately identifying firm fixed effects from search costs imposes additional restrictions on the data, which makes it largely an empirical question whether this is feasible. Therefore, our search model fits best in settings where unobserved firm heterogeneity is found not to play a major role or in settings in which there exists a good amount of observed information regarding firm heterogeneity.

3 Application

In this section we apply the model to the market for Medigap insurance plans, which is a market in which there is substantial price dispersion despite products being relatively homogenous. We first describe some institutional details as well as the data we use for estimation. We next discuss how we relate price variation to product differentiation and consumer search frictions, and provide reduced form evidence that suggests search frictions are important in this market. We then estimate the model developed in Section 2 and discuss the results. We also make comparisons between our model and several alternative models in terms of estimates and model fit. We finally use the model to calculate price elasticities and simulate prices when search costs go down to zero.

3.1 Industry Background

Medigap, or Medicare supplemental health insurance, is a form of private insurance designed to supplement Medicare by filling in the coverage gaps in Medicare payment. Basic Medicare has substantial cost-sharing and gaps in coverage, and does not include an upper limit on beneficiaries’ out-of-pocket spending.⁹ As a result, in 2009 over 90 percent of beneficiaries had some kind of

⁹For example, in 2009, Medicare Part A (inpatient insurance) required an enrollee to pay a total deductible of \$1,068 for the first 60 days of a hospital stay, \$267 per day for days 61-90, \$534 per day for days 91-150, and all costs for each day beyond 150 days. Under Medicare Part B (outpatient insurance), enrollees were responsible for a total

supplemental insurance to fill in the coverage gaps in basic Medicare. Of the different sources of supplemental insurance, Medigap provided coverage for about 20 percent of the Medicare beneficiaries in 2009 (Sheingold, Shartzter, and Ly, 2011).

One important feature of the Medigap market is plan standardization that came with the passage of the Omnibus Budget Reconciliation Act (OBRA) in 1990. All new policies sold after July 1992 are required to conform to a set of standardized plans (Massachusetts, Minnesota, and Wisconsin are not subject to the standardization). Each plan must offer the standard set of benefits regardless of which company sells it. Plan standardization was intended to promote competition among firms and to avoid misunderstanding and confusion among seniors regarding coverage. As of 2009, a total of twelve standardized plans, labeled by the letters A through L, were offered. Of these twelve plans, Plan A is considered the basic plan and all other plans include additional coverage options. Not all plans are available in each state. Insurers that operate in a given state are required to offer at least the basic plan (Plan A). If they choose to offer other plans, Plan C or F has to be provided too. So far Plan F, which offers the second-most comprehensive coverage, has been the most popular plan type, enrolling more than 40 percent of Medigap policy holders.

Another important feature of the Medigap market is that during open enrollment insurers cannot deny coverage to an eligible applicant because of health status (guaranteed issue). The OBRA 1990 mandates a six-month open enrollment period for Medigap, starting from the day an individual turns 65 and enrolls in Medicare Part B. Medical underwriting is allowed after the open enrollment period, making it difficult to switch to a new plan. As a result, the majority of seniors subscribes to a Medigap plan during the open enrollment period in which medical underwriting is prohibited.

In addition, Medigap premiums are subject to regulation.¹⁰ Premiums are allowed to vary on the basis of age, gender, and smoking status. Regarding age, three different rating methods can be used: attained-age, issue-age, and community-rated. Attained-age means that premiums are allowed to increase with the age of the policy holder. Issue-age premiums are charged based on the age when the policy holder enrolls. Under issue-age rating, premiums are only allowed to increase to compensate for rising healthcare costs over time, and premium increases cannot be based on the age of the policy holder. Community-rating premiums are uniform across all subscribed individuals in

deductible of \$135 in 2009 as well as a 20 percent coinsurance payment of the Medicare-approved amount for covered services after the deductible had been reached.

¹⁰State departments of insurance or other state agencies have the authority to approve insurance policy forms and premiums (Sheingold, Shartzter, and Ly, 2011). A state's regulation is required to meet the National Association of Insurance Commissioners Model Standards. In addition, all insurers are subject to regulations on a minimum loss ratio of 65 percent for individual policies, which means that plans must spend at least 65 dollars on medical care for every 100 dollars in premium.

the community.¹¹ At the entry age of 65, premiums are higher for issue-age plans in comparison to attained-age plans in order to incorporate increases in utilization as individuals age (Robst, 2006). Premium differences based on gender are minimal, and premiums for smokers are on average higher than for non-smokers.

3.2 Data and Basic Analysis

Data

Our data come from two sources. Data on premiums come from Weiss Ratings, which provide detailed pricing information for Medigap plans offered in 2009. The data include information on insurers, location, plan type, gender, age group, and some other characteristics such as rating method (attained-age, issue-age, and community-rated) and smoking status.

Data on market shares are derived from the National Association of Insurance Commissioners' (NAIC) Medigap Experience Files. The NAIC dataset provides information on the number of active policies as of December 31, 2009, as well as other variables such as total premiums and claim volumes for 2009. This dataset allows us to differentiate between individual plans and group plans that are sold through employers. Since we are interested in how individual consumers make purchase decisions, our focus is on individual plans. The number of policies issued before 2007 and from 2007 to 2009 combined are reported separately.¹² Policies issued before 2007 represent policy holders who subscribed before 2007 and chose to renew their policies up to 2009. For the calculations of the market shares used in our empirical analysis, we only use policies newly issued between 2007 and 2009.¹³ One obvious advantage of this measure is that we can focus our analysis on new policies sold to individuals who turned 65 and became eligible to purchase a Medigap plan. Market shares are constructed in a similar way in Starc (2014) and Fang, Keane, and Silverman (2008).

The merged dataset contains information on prices and sales for each Medigap policy offered by each insurer at the state level. It corresponds to a total of 4,704 observations (at the state-insurer-

¹¹Several states, including Connecticut, New York, and Vermont, mandate community rating.

¹²According to the data reporting instructions, the NAIC dataset provides a snapshot of the number of individuals covered under each policy on December 31 of each year reported. This means that we can only infer from the data the *net* gain in the number of covered lives. For example, suppose an insurer shows 10 covered lives for the most current three years (2007-2009). If a member bought a policy in April 2009 and passed away in November 2009, the person would not be reflected in the year-end covered lives number. The actual number of policies sold for the year is 11, but because of the way in which the data is captured, we do not observe the policy sold to the deceased policy holder.

¹³Ideally, we would like to observe new policies sold only in 2009 to match our pricing data. However, only three years of combined data are available through NAIC, which, to the best of our knowledge, is the most comprehensive source of Medigap market share information.

plan level), accounting for more than 80 percent of the total policies sold between 2007 and 2009. A total of 109 companies are observed in the data, of which 14 operate in more than 30 states. The median number of states in which an insurer operates is 6. Note that the NAIC data only reports covered lives for all ages combined. Since most new policies were sold to individuals at the open enrollment period, and since most premiums do not differ across gender, the prices we use for estimating the model are the premiums for 65-year-old females (following Maestas, Schroeder, and Goldman, 2009).¹⁴

Price dispersion

There exists large price variation for Medigap plans across insurers within a state. Take premiums for Medigap Plan F (the most popular plan) in Indiana as an example: a total of 40 companies offered Plan F in Indiana in 2009, and premiums for a 65-year-old woman using the attained-age rating method (the most popular rating method) range from \$1,223 to \$3,670, with a median of \$1,818.

To get a better picture of the extent of premium variation, Table 1 summarizes the coefficient of variation by rating method, averaged over plan types and states.¹⁵ For all three rating methods the coefficient of variation is substantial. Attained-age and issue-age show slightly higher variation than a community-based rating. Moreover, there is substantial variation in the coefficient of variation across plan types and states; whereas in a few states some plan types show hardly any price variation, the coefficient of variation can be as high as 0.65 for the issue-age rating method. Note that Table 1 is calculated using prices for 65 year old females only; as reported above, prices for different age groups as well as for males follow similar patterns.

Table 1: Coefficient of Variation

Rating	Mean	S.D.	Min	Percentiles					Max
				5%	25%	50%	75%	95%	
Attained-Age	0.217	0.082	0.000	0.071	0.173	0.216	0.267	0.344	0.486
Issue-Age	0.215	0.122	0.002	0.043	0.118	0.201	0.290	0.450	0.651
Community	0.208	0.114	0.009	0.018	0.154	0.197	0.263	0.381	0.559

Notes: All reported values are averaged over plan types and states.

¹⁴For those observations that reveal a difference between male and female premiums, the mean difference is only around 7 percent of the average price (with a 99th percentile of approximately 19 percent).

¹⁵Since insurers use different rating methods for a given plan type within a state, we combine plan type and rating method to identify a product. This means that Plan F attained-age and Plan F issue-age are treated as different products.

Table 2 reports average prices and coefficient of variation by plan type for the attained-age rating plans, the most popular rating method in the data. The average price varies greatly across states. For example, the average premium for plan A is \$1,128 across states, with a standard deviation of 152. Also the average difference between the minimum and maximum premium charged within a state is substantial. For instance, for Plan A the average difference is \$869, which is close to 80 percent of Plan A’s mean price. Table 2 also shows that most plan types have an average coefficient of variation that exceeds 0.20. Plan F, the most popular and the second-most comprehensive plan, has the highest average price variation as measured by the coefficient of variation.

Table 2: Coefficient of Variation by plan type (attained-age)

Plan Type	Price		Max-Min		Coefficient of Variation			
	Mean	S.D.	Mean	S.D.	Mean	S.D.	Min	Max
A	1,128	152	869	488	0.241	0.087	0.000	0.415
B	1,567	261	959	512	0.247	0.078	0.081	0.486
C	1,804	190	1,490	594	0.255	0.052	0.160	0.370
D	1,437	199	927	482	0.203	0.061	0.009	0.278
E	1,509	284	535	446	0.225	0.111	0.014	0.478
F	1,813	194	1,889	760	0.261	0.040	0.194	0.341
G	1,422	197	868	437	0.205	0.046	0.085	0.323
H	1,468	217	278	226	0.111	0.047	0.031	0.226
I	1,515	210	319	359	0.187	0.104	0.044	0.483
J	1,746	181	721	353	0.170	0.055	0.055	0.341
K	1,030	161	215	204	0.223	0.078	0.041	0.377
L	1,455	265	322	334	0.229	0.118	0.004	0.430

Part of the variation in premiums across insurers might be attributed to plans that attract very few individuals due to high prices. To correct for differences in market share, we also calculated the coefficient of variation weighted by market share for each plan. The weighted coefficient of variation for Plan F under attained age rating averages 0.20 for 65-year-old females, indicating that considerable price dispersion among Medigap plans still remains even after controlling for differences in market share.

Cost heterogeneity

One plausible explanation for price variation in Medigap plans involves what is known as “endogenous sorting.” That is, insurers might have a different pool of risk types, with high-priced insurers attracting consumers that are more costly to insure. Though a majority of beneficiaries purchase Medigap during the open enrollment period when insurers cannot decline an application, firms might still target different types of consumers through marketing or advertising. In this case, consumers with different risk types then sort themselves into different insurers, and the equilibrium

might be such that high-risk consumers select an insurer that charges a high premium while low-risk ones select an insurer that offers a low premium.

To investigate whether and to what extent price variation across insurers can be explained by cost differentials, we use data on the total amount of claims filed for each observation (at the state-insurer-plan level).¹⁶ We use the average claim, defined as the ratio of total claims to the total number of covered lives, as a measure of the risk type for an average policyholder of a given plan, assuming a larger average claim indicates a higher risk type. We then run a regression of premiums on the average claim, after controlling for state-fixed effects, dummies for plan types, and dummies for rating methods, and we compare the results to a regression without controlling for the average claim.

Table 3: Price Dispersion, Cost Differentials, and Firm Heterogeneity

	Premiums			
	(1)	(2)	(3)	(4)
Average claims		0.052 (0.003)***		0.026 (0.003)***
Number of observations	4,704	4,686	4,704	4,686
Insurer fixed effects	No	No	Yes	Yes
R-squared	0.394	0.425	0.677	0.684

The results are presented in columns 1 and 2 of Table 3. Column 1 is based on a regression without controlling for the average claim. Column 2 adds the average claim. We find that a one-dollar increase in average claim is associated with a 5-cent increase in premium. The *R*-squared increases by about 3 percentage points, indicating that adding controls for the average claim does not contribute much to explaining the observed variation in premiums. Although not reported, we also run regressions separately for each plan type and find similar results. These results indicate that cost differentials are related to Medigap plan pricing but they do not play a major role in explaining the observed price variation across different insurers.

¹⁶Note that according to federal regulation, the loss ratio (defined as total claims divided by total premiums) cannot be lower than 65 percent for individual Medigap plans. If this minimum ratio is violated, an insurer is required to provide transfers to its policyholders to be compliant with the regulation. We observe large differences in loss ratio across plans. The average loss ratio is 0.80, but almost 40 percent of the insurer-plan combinations have a loss ratio below 65 percent.

3.3 Product Differentiation and Search

As documented above, there is substantial price variation in Medigap plans and cost heterogeneity does not seem to help much to explain the observed price variation. In this subsection we provide reduced-form evidence for product differentiation and search friction, two ingredients that motivate our model.

Product differentiation

Although Medigap plans are standardized, premiums might still differ across insurers due to differences in other dimensions than plan characteristics. For example, a consumer might favor an older and more established firm or a firm with a better financial rating. Firms do differ in these attributes: age of the insurers ranges from 3 to 159 years and according to financial safety ratings provided by Weiss Ratings, 22 out of the 109 insurers in our dataset had a financial rating of A, 47 had a rating of B, and 12 were either unrated or rated below C.¹⁷ Firms also differ in billing services, the number of plans that are offered, as well as how many states they are active.

Still, pricing patterns may provide indirect evidence of whether firm-level product differentiation plays an important role in explaining price variation for Medigap plans. If price variation is solely caused by firm-level product differentiation, one would expect insurers to charge consistently high or low prices for all offered plans in a given state: if a better financial rating or reputation allows the insurer to charge a higher premium this is likely to happen for all plan types the firm is selling. To see if this is indeed the case in our data, for each plan type we compare each insurer's price to the overall price distribution in a state and determine in which quartile each price falls. We then calculate the fraction of plans that belong to each quartile for each insurer in a given state. If price variation primarily reflects product differentiation at the firm level, insurers are likely to consistently set premiums in the same quartile for all the plans that they offer in a state. We find that this happens to on average 30 percent of insurers in a state. The remaining insurers on average have an equal number of plans in each of the four quartiles.¹⁸ If we compare each plan to the median of the price distribution, we find that on average 60 percent of insurers in a state charge prices for all offered plans consistently below or above the median. The remaining insurers tend

¹⁷Weiss Ratings rates each insurer as follows. "A" means the company offers excellent financial security; "B" means the company offers good financial security and has the resources to deal with a variety of adverse economic conditions; "C" means the company is currently stable, although during an economic downturn or other financial pressures it may encounter difficulties in maintaining its financial stability. A company is unrated for reasons such as total assets being less than \$1 million or a lack of the information necessary to reliably issue a rating.

¹⁸To be more specific, we find the proportion of plans belonging to each quartile averages 22 percent for the first quartile, 28 percent for the second, 29 percent for the third, and 21 percent for the fourth.

to switch prices back and forth between below and above the median of the corresponding price distribution. These findings suggest that product differentiation is important, but it alone cannot explain all the variation in premiums observed in the data.

To further examine the role of firm heterogeneity in explaining observed price variation, we follow an approach suggested by Sorensen (2000). We compare the R -squared of a regression of premiums on plan dummies, rating methods, and state dummies (column 1 of Table 3) to the R -squared of a similar specification that also includes insurer-fixed effects (column 3 of Table 3). Adding the insurer-fixed effects increases the R -squared from 0.39 to 0.68, accounting for about 48 percent of the variation unexplained by the regression if we leave out the insurer-fixed effects.¹⁹ We find similar results if we add the average claim as an explanatory variable (compare column 4 of Table 3 to column 2). We also conducted a regression with additional controls for state-insurer fixed effects; this increases the R -squared from 0.68 to 0.83, accounting for about 45 percent of the variation unexplained by the regression using only insurer-fixed effects. These results further confirm that although a significant amount of variation in premiums can be absorbed by firm heterogeneity, product differentiation alone does not fully explain price variation across insurers within a state.

Search

An additional explanation for the observed price differences relates to search frictions.²⁰ Given the large number of insurers in each market, it will be costly for consumers to obtain information on the plans and corresponding rates offered by each insurer. Since Medigap is supplemental insurance sold and administered by private insurers, no free and universal pricing information is available to consumers. The most comprehensive guide to choosing a Medigap policy is published annually by the Center of Medicare and Medicaid Services (CMS). The CMS suggests Medicare beneficiaries take the following three steps to acquire information about Medigap plans before subscribing to a plan: (1) decide which plan type to purchase; (2) find out which insurers are selling this plan type; and (3) call these insurers and compare costs. For step 2 the CMS advises beneficiaries to

¹⁹Note that this comparison with and without controlling for firm fixed effects might lead us to overstate the importance of firm heterogeneity, because some price dispersion unrelated to firm differences will nevertheless be absorbed by the inclusion of fixed effects (Sorensen, 2000).

²⁰We study how search friction and product differentiation translate into price dispersion in identical Medigap plans. We do not specifically consider whether firms raise prices on existing enrollees after they are locked in (see, e.g., Ericson, 2014). Note that Ericson (2014) finds that firms introduce new and cheaper plans to attract new enrollees in Medicare Part D, which is prohibited in Medigap due to plan standardization. However, consumer inertia might exist in the Medigap market and it is possible that firms set prices low to attract consumer when they make their initial purchase decisions (typically when they turn to age 65 and enroll in Medicare Part B). This is interesting to explore but beyond the scope of this paper.

call the State Health Insurance Assistance Program or the State Insurance Department, or to visit medicare.gov. Note that these sources only provide a list of insurers that sell Medigap plans by state, which means that no pricing information is readily available. The CMS’s suggestion to call insurers to compare costs (step 3) can therefore be quite time consuming, especially given the large number of insurers in each state.

The presence of several websites that facilitate search further supports the notion that search frictions play an important role in this market. For instance, Weiss Ratings has a comprehensive nationwide Medigap pricing information database. Weiss Ratings sells this information to consumers; in 2015 the cost of acquiring information about all Medigap plans sold in a specific zip code was \$99.




Figure 1: Log price versus log market share (Plan G, Indiana)

As shown by a large theoretical literature on consumer search, search frictions may lead to price dispersion in homogenous product markets (Burdett and Judd, 1983; Stahl, 1989) as well as markets for differentiated products (Anderson and Renault, 1999; Wolinsky, 1986). The results presented in the previous subsection suggest that firm differentiation is important in the Medigap market, so search frictions are unlikely to be the only explanation for the price variation we observe in the data, ruling out a search model in which consumers search for homogenous products.²¹

²¹In a typical homogenous good search model, the pricing equilibrium is in mixed strategies (see, for instance, Burdett and Judd, 1983; Stahl, 1989). We do not find evidence that supports mixed strategies in prices in the data. For instance, using three years of pricing data from 2007, 2008, and 2009, we find that more than 60 percent of insurers consistently set Medigap Plan F prices in the same quartile of the price distribution, and that more than 85 percent stuck to prices that were either below or above the median price. Moreover, around 10 percent of insurers charged the same prices between 2007 and 2009, and 30 percent had only one change in prices during the three-year period.

Additional evidence that rules out search as the sole determinant of price dispersion derives from the relationship between market shares and prices. Were search to be the only explanation for price differences across insurers, one would expect to see a negative and monotonic relationship between price and market share (Hortaçsu and Syverson, 2004). In general, this is not the case in our data. For example, Figure 1 shows how market shares relate to prices for Medigap Plan G in Indiana—although there seems to be a weak negative relationship between the log price and log market share, this relationship is not monotonic. The insurer with the largest market share has the twelfth highest price out of 26 firms that offer plan G, and the insurer with the third largest market share has the fifth highest price. This lack of a monotonic relationship between price and market share also exists if we look at the market share of insurers, regardless of plan type.²² For example, in Indiana the insurer with the largest market share charges the fourth highest price; and the top five insurers in terms of market share all charge prices above the median of the overall price distribution.

3.4 Results

Search model estimates

Table 4 reports the estimation results when estimating our search model using data from the Medigap market. In our main specification, presented in column 1 of the table, we estimate equation (8) by constrained two-stage least-squares. We use Hausman-type instruments: the average price of the same plan type offered by the firm in other markets and the average price of all plans offered by the firm in other markets.²³ The price coefficient is negative and highly significant. Since search cost parameters γ vary across markets, we report these parameters by calculating their average across states. To reduce the number of search cost parameters, we allow consumers to search any number from one to four or five and more. If they search five or more times, we assume they have information on all the firms in a market.²⁴ These coefficients indicate that the lower the ranking of a given firm, the smaller the market share of its products. Our main specification also

²²An insurer’s market share is calculated by summing up market share for all plans offered by an insurer in a given state, whereas price is the average price across the plan offered.

²³Hausman-type instruments assume that prices consists of a common cost shock in addition to a market-specific shock, which makes it valid to use average prices in other markets as an instrument, but only if there are no common demand shocks that are reflected in prices (for instance, as a result of a national advertising campaign). To check robustness to alternative instruments, we adopted three instruments from the Medicare Advantage (MA) market, which offers close substitutes to Medigap plans. These instruments are the number of MA contracts, the number of MA plans, and the benchmark payment (weighted by the MA enrollees in each county of a state). The results are largely similar, which provide some support for the validity of the Hausman type of instrumental variables used for the main specification.

²⁴Our results are robust to using a different cutoff, such as six or seven searches.

controls for a set of insurer-related characteristics, such as financial safety rating, age of a firm, the number of different policies an insurer offers, the number of states in which an insurer is active, and whether the state is the insurer’s domicile. The estimates for the financial ratings indicate that firms with the best safety rating (rating “A”, the omitted group) are the most preferred and that firms with no rating are the least preferred. The coefficient on firms with a “D” rating is imprecisely estimated, due to the fact that only 3.5% of the firms fall into this category. Insurer age has a positive marginal utility, which is likely to reflect the reputation or stability of an insurer. The estimates for the indicator variables related to the number of states in which an insurer is active suggest that consumers prefer insurers that are active in less than 24 states, which suggests that consumers prefer local insurers. Consistent with this finding, we find a positive and highly significant parameter on the dummy variable for domicile state. Consumers also prefer insurers that offer multiple Medigap plans, as indicated by the positive and significant dummy variables related to the number of policies offered by the insurer. The estimated rating-method coefficients indicate that plans with an issued-age or community-rated rating method are more desirable than attained-age plans. This result suggests that policy holders value a rating method that offers some sort of protection against increases in premiums as they get older. All plan dummies have a positive coefficient and are highly significant, suggesting that all plans have a higher marginal utility than the omitted plan (Plan A, which is the basic plan and offers the least coverage of all plans). Not surprisingly, Plan J, the most comprehensive plan, has the highest marginal utility; and Plan F, the second-most comprehensive plan, has the second-highest marginal utility. Plan K has the lowest coefficient estimate of all plans included in the table, due to a higher out-of-pocket limit.

Specification (1) is estimated using constrained two-stage least squares, where we use constraints on the ranking-related coefficients to make sure that the observed ranking of firms in terms of market share are consistent with the underlying search model. In specification (2) of Table 4 we relax those constraints and force the ranking-related coefficients to be the same across states. Notice that the ranking-related coefficients can be interpreted as the marginal effect of moving down the ranking. However, if we restrict these coefficients to be the same across states, there is no guarantee that the estimated ranking-related coefficients are consistent with the observed ranking of the firms in terms of market share in each individual state. As a result, we are unable to use these estimates to derive the search cost distribution. Apart from this, the parameter estimates are very similar to those reported in column 1 of Table 4. Not surprisingly, the fit, as measured by the R -squared, goes up. The release of the restrictions on the ranking-related parameters used for estimating specification (1) has the biggest effect on those firms that are ranked fifth or higher, since that parameter more

Table 4: Estimation Results Search Model

Variable	(1)		(2)	
	Coeff.	Std. Err.	Coeff.	Std. Err.
Price	-1.732	(0.086)***	-1.679	(0.093)***
<i>Ranking-related dummies (average)</i>				
Ranked 2 or higher	-0.732		-0.739	(0.165)***
Ranked 3 or higher	-0.664		-0.489	(0.160)***
Ranked 4 or higher	-0.476		-0.397	(0.157)**
Ranked 5 or higher	-0.321		-1.280	(0.114)***
<i>Insurer characteristics</i>				
Firm Weiss safety rating B	-0.191	(0.081)**	-0.183	(0.080)**
Firm Weiss safety rating C	-0.257	(0.085)***	-0.171	(0.084)**
Firm Weiss safety rating D	0.059	(0.137)	0.171	(0.135)
No Weiss safety rating	-0.491	(0.108)***	-0.463	(0.106)***
Age	0.002	(0.001)**	0.002	(0.001)**
Number of states active 24 – 36	-0.292	(0.072)***	-0.169	(0.072)**
Number of states active > 36	-0.086	(0.081)	0.054	(0.080)
Number of policies offered 4 – 7	0.261	(0.077)***	0.163	(0.076)**
Number of policies offered > 7	0.255	(0.094)***	0.170	(0.093)*
Domicile state	0.531	(0.097)***	0.465	(0.098)***
<i>Rating method</i>				
Issued-age	0.397	(0.092)***	0.444	(0.091)***
Community-rated	1.689	(0.297)***	1.643	(0.293)***
<i>Plan dummies</i>				
Plan B	0.862	(0.112)***	0.861	(0.111)***
Plan C	2.119	(0.114)***	2.133	(0.114)***
Plan D	1.853	(0.102)***	1.867	(0.101)***
Plan E	1.398	(0.146)***	1.387	(0.143)***
Plan F	4.073	(0.105)***	4.087	(0.106)***
Plan G	2.496	(0.103)***	2.496	(0.102)***
Plan H	2.102	(0.181)***	2.093	(0.178)***
Plan I	2.455	(0.175)***	2.445	(0.172)***
Plan J	4.531	(0.135)***	4.519	(0.134)***
Plan K	0.452	(0.195)**	0.502	(0.192)***
Plan L	0.689	(0.178)***	0.753	(0.175)***
Number of Observations	4,704		4,704	
R^2	0.529		0.542	
State Fixed Effects	yes		yes	
Firm Fixed Effects	no		no	
Plan Fixed Effects	yes		yes	
State-specific ranking-related dummies	yes		no	

All specifications include a constant. *** Significant at the 1 percent level. ** Significant at the 5 percent level. * Significant at the 10 percent level.

than quadruples. In addition, the parameters on some of the insurer characteristics change in magnitude, but most other parameters do not change that much.

Figure 2 provides the estimated search cost cumulative distribution function, averaged across states, for specification (1).²⁵ Median search costs are approximately \$30. According to the average search cost distribution, the 25th, 75th, and 90th percentiles are \$13, \$54, and \$151, respectively. On average, 39 percent of consumers searches once, 22 percent searches twice, 14 percent searches three times, 6 percent searches four times, and 19 percent of consumers searches five or more times.




Figure 2: Estimated search costs (averaged across states)

3.5 Comparison to Alternative Models

Insurer fixed effects

To capture heterogeneity across insurers, we use a relatively rich set of controls in our main specification and argue that these controls explain a good amount of heterogeneity across firms. However, it might be possible that part of the insurer-level heterogeneity that affects the decision making of consumers is unobserved by the researcher. To see to what extent our estimates would change if we control for unobserved heterogeneity at the firm level, in column 1 of Table 5 we drop the insurer-specific covariates and include insurer fixed effects instead. As discussed in Section 2.4, to be able to separately identify firm fixed effects from the ranking-related parameters, we exploit variation in firm identities that share the “same” ranking across markets. In conducting this analysis, we add an

²⁵For each state we get non-parametric estimates of the search cost distribution. We then obtain the curves in Figure 2 by fitting a log-normal distribution through the estimated points on the search cost CDF of each state and then taking the average fitted CDF across states.

additional year of data (2010) to our original data. To be able to make a fair comparison between this model and the model that does not include firm fixed effects, we re-estimate the original model using both 2009 and 2010 data.

Table 5: Estimation Results Insurer Fixed Effects

Variable	(1)		(2)	
	Coeff.	Std. Err.	Coeff.	Std. Err.
Price	-1.558	(0.089)***	-1.578	(0.065)***
<i>Ranking-related dummies (average)</i>				
Ranked 2 or higher	-0.539		-0.614	
Ranked 3 or higher	-0.520		-0.668	
Ranked 4 or higher	-0.457		-0.547	
<i>Insurer characteristics</i>				
Firm Weiss safety rating B			-0.257	(0.061)***
Firm Weiss safety rating C			-0.293	(0.065)***
Firm Weiss safety rating D			-0.118	(0.087)
No Weiss safety rating			-0.653	(0.087)***
Age			0.003	(0.001)***
Number of states active 24 – 36			-0.110	(0.056)**
Number of states active > 36			0.061	(0.062)
Number of policies offered 4 – 7			0.374	(0.057)***
Number of policies offered > 7			0.379	(0.067)***
Domicile state			0.697	(0.077)***
<i>Rating method</i>				
Issued-age	0.691	(0.091)***	0.259	(0.077)***
Community-rated	0.702	(0.209)***	2.051	(0.197)***
<i>Plan dummies</i>				
Plan B	0.794	(0.080)***	0.728	(0.086)***
Plan C	1.985	(0.087)***	1.913	(0.085)***
Plan D	2.034	(0.079)***	1.570	(0.081)***
Plan E	1.752	(0.130)***	1.261	(0.140)***
Plan F	4.291	(0.082)***	3.974	(0.076)***
Plan G	2.581	(0.074)***	2.282	(0.075)***
Plan H	2.721	(0.163)***	1.861	(0.176)***
Plan I	2.661	(0.157)***	2.267	(0.170)***
Plan J	4.908	(0.120)***	4.332	(0.125)***
Plan K	-0.263	(0.121)**	0.216	(0.123)*
Plan L	0.331	(0.112)***	0.566	(0.121)***
Plan M	-2.386	(0.219)***	-1.925	(0.238)***
Plan N	1.617	(0.094)***	1.366	(0.104)***
<i>Year dummies</i>				
2010	0.182	(0.038)***	0.321	(0.040)***
Number of Observations	8,601		8,601	
R^2	0.614		0.516	
State Fixed Effects	yes		yes	
Firm Fixed Effects	yes		no	
Plan Fixed Effects	yes		yes	
State-specific ranking-related dummies	yes		yes	

All specifications include a constant. *** Significant at the 1 percent level. ** Significant at the 5 percent level. * Significant at the 10 percent level. We use three ranking-related dummies in both specifications because in 2010 there is one state in which only four firms are active.

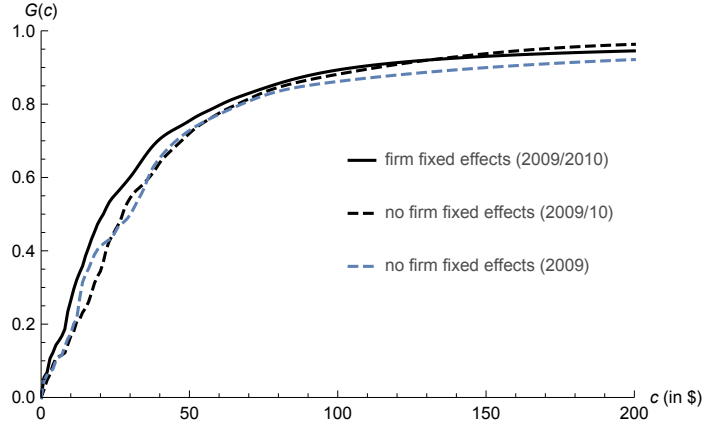


Figure 3: Estimated search costs (firm fixed effects)

As can be seen from comparing columns 1 and 2 of Table 5, most parameter estimates are very similar. The only parameters that are further apart are the estimated rating method parameters—in the model without insurer fixed effects the coefficient on community rating is much larger than the coefficient on issue-age rating. The estimated price coefficients as well as the average ranking-related coefficients, however, are very similar. As a result there is not much difference in estimated search costs, as can also be seen from Figure 3. In this graph the solid black curve represents the search cost CDF when using firm fixed effects, while the search cost CDF for the model without fixed effects is shown by the dashed black curve. Also plotted is the search cost CDF obtained when using just the 2009 data and no firm fixed effects, as in the main specification. Median search costs are smaller in the model with firm fixed effects, which is as expected, since part of the variation in market shares that is attributed to search frictions in our main specification is picked up by the insurer-fixed effects as unobserved heterogeneity. Unfortunately there is no good way of knowing whether the observed variation in market shares is due to unobserved insurer heterogeneity or such variation should be attributed to the ranking-related effects caused by search frictions. However, because of the relatively rich set of firm-level controls we use in our main specification and the fact that we find largely similar results even if we control for firm fixed effects, we are confident that a significant part of the variation in market shares that remains after controlling for observed firm heterogeneity can be attributed to search frictions.

Table 6: Estimation Results Alternative Specifications

Variable	Alternative Rankings								Full Information	
	(1)		(2)		(3)		(4)		(5)	
	Coeff.	Std. Err.	Coeff.	Std. Err.	Coeff.	Std. Err.	Coeff.	Std. Err.	Coeff.	Std. Err.
Price	-1.805	(0.096)***	-1.796	(0.093)***	-2.079	(0.101)***	-2.157	(0.136)***	-2.116	(0.098)***
<i>Ranking-related dummies (average)</i>										
Ranked 2 or higher	-0.774	(0.173)***	-0.927	(0.182)***	0.827	(0.229)***	-0.013	(0.263)		
Ranked 3 or higher	-0.524	(0.170)***	-0.443	(0.175)**	-0.390	(0.239)	-0.020	(0.263)		
Ranked 4 or higher	-0.495	(0.171)***	-0.399	(0.174)**	-0.097	(0.251)	0.037	(0.243)		
Ranked 5 or higher	-1.275	(0.130)***	-1.434	(0.132)***	0.344	(0.190)*	0.184	(0.187)		
<i>Insurer characteristics</i>										
Firm Weiss safety rating B	-0.151	(0.090)*	-0.107	(0.092)	-0.069	(0.091)	-0.053	(0.119)	-0.065	(0.088)
Firm Weiss safety rating C	-0.040	(0.092)	0.049	(0.093)	-0.310	(0.095)***	-0.370	(0.140)***	-0.362	(0.092)***
Firm Weiss safety rating D	0.225	(0.138)	0.288	(0.137)**	-0.044	(0.146)	0.004	(0.236)	-0.007	(0.148)
No Weiss safety rating	-0.422	(0.113)***	-0.237	(0.114)**	-0.854	(0.123)***	-0.599	(0.179)***	-0.629	(0.116)***
Age	0.003	(0.001)***	0.003	(0.001)**	0.008	(0.001)***	0.003	(0.001)**	0.004	(0.001)***
Number of states active 24 – 36	0.008	(0.080)	-0.052	(0.078)	-0.530	(0.081)***	-0.518	(0.105)***	-0.521	(0.078)***
Number of states active > 36	0.263	(0.091)***	0.129	(0.090)	-0.297	(0.092)***	-0.291	(0.136)**	-0.296	(0.088)***
Number of policies offered 4 – 7	0.274	(0.080)***	0.409	(0.081)***	0.453	(0.086)***	0.398	(0.092)***	0.382	(0.083)***
Number of policies offered > 7	0.245	(0.097)**	0.219	(0.100)**	0.450	(0.110)***	0.472	(0.105)***	0.463	(0.101)***
Domicile state	0.175	(0.117)	0.020	(0.120)	0.760	(0.111)***	0.934	(0.294)***	0.977	(0.105)***
<i>Rating method</i>										
Issued-age	0.423	(0.093)***	0.345	(0.094)***	0.465	(0.103)***	0.388	(0.159)**	0.408	(0.100)***
Community-rated	1.714	(0.353)***	1.423	(0.405)***	2.034	(0.314)***	2.004	(0.813)**	2.068	(0.321)***
<i>Plan dummies</i>										
Plan B	0.887	(0.109)***	0.927	(0.109)***	0.953	(0.119)***	0.953	(0.210)***	0.943	(0.122)***
Plan C	2.205	(0.113)***	2.175	(0.113)***	2.289	(0.123)***	2.323	(0.262)***	2.279	(0.125)***
Plan D	1.915	(0.100)***	1.971	(0.100)***	1.882	(0.109)***	1.923	(0.228)***	1.849	(0.111)***
Plan E	1.401	(0.142)***	1.465	(0.142)***	1.517	(0.154)***	1.492	(0.226)***	1.456	(0.158)***
Plan F	4.186	(0.107)***	4.202	(0.106)***	4.290	(0.115)***	4.346	(0.199)***	4.243	(0.117)***
Plan G	2.576	(0.101)***	2.600	(0.101)***	2.638	(0.109)***	2.646	(0.162)***	2.596	(0.112)***
Plan H	2.105	(0.175)***	2.210	(0.175)***	2.183	(0.190)***	2.194	(0.179)***	2.124	(0.196)***
Plan I	2.469	(0.170)***	2.545	(0.170)***	2.534	(0.186)***	2.507	(0.157)***	2.454	(0.190)***
Plan J	4.600	(0.133)***	4.690	(0.132)***	4.581	(0.147)***	4.746	(0.196)***	4.672	(0.147)***
Plan K	0.361	(0.191)*	0.430	(0.190)***	0.356	(0.206)*	0.222	(0.185)	0.324	(0.211)
Plan L	0.662	(0.172)***	0.702	(0.173)***	0.784	(0.187)***	0.656	(0.167)***	0.684	(0.193)***
Number of Observations	4,704		4,704		4,704		4,704		4,704	
R^2	0.565		0.563		0.488		0.490		0.444	
State Fixed Effects	yes		yes		yes		yes		yes	
Firm Fixed Effects	no		no		no		no		no	
Plan Fixed Effects	yes		yes		yes		yes		yes	
State-specific ranking-related dummies	yes		yes		yes		yes		yes	

All specifications include a constant. *** Significant at the 1 percent level. ** Significant at the 5 percent level. * Significant at the 10 percent level. The standard errors in column (4) are bootstrapped.

Alternative rankings

Column 1 of Table 6 gives estimation results for a specification in which we do not impose any constraints on ranking-related parameters in equation (8). These constraints imposed in our search model make the ranking of firms endogenous: consumers form rational expectations of the utility of each firm and search higher valued firms first. A model without imposing these constraints can be interpreted as a model with an exogenous ordering of firms according to decreasing market share. Note that these parameters are still allowed to vary across markets. The estimation results are very close to the main specification, suggesting that those constraints cannot be rejected in the data. The fit of the model, as measured by the R -squared, increases.

Regarding the ordering of firm, it is possible that due to word-of-mouth effect, firms with a larger market share in previous years are more likely to become part of consumers' choice sets. Column 2 gives estimation results for a specification that uses an exogenous ordering based on market shares in previous years. More specifically, we use the ranking of firms according to the number of current policies (active in 2009) that were sold before 2007, and use this as the order of search for our estimation. The estimated price coefficient is similar to that of our main specification.

In column 3 of Table 6 we estimate the model assuming consumers visit firms alphabetically, by firm names. Moreover, the ranking-related parameters (averaged across states) are either positive, or no longer significantly different from zero, which suggests consumers are not very likely to search firms alphabetically. The price coefficient is also very similar to the full information case (see the last column of Table 6), which gives another indication that this particular ranking of firms does not add much in comparison to the model that assumes consumers have full information. As indicated by the lower R -squared, the model also gives a worse fit than the search model.

An important difference between our approach and that of Maestas, Schroeder, and Goldman (2009) is that search is random in Maestas, Schroeder, and Goldman (2009), whereas search is directed in our model. Extending our model to allow for random search is not trivial—if search is random any combination of insurers can be searched, which, given the large number of insurers in a state, leads to a very large number of choice sets. As a result, demand for plans sold by a specific insurer can originate from many different choice sets, which creates a dimensionality problem when deriving overall demand for this insurer.²⁶ To restrict the number of choice sets we assume, as in our

²⁶A similar dimensionality problem does not arise in Maestas, Schroeder, and Goldman (2009). Although consumers in their model know the distribution of realized utility values, they do not know the identify of the firms, which makes consumers search random. Moreover, there is no random utility shock and consumers are searching sequentially, and as a result the firm that offers the lowest mean utility will only sell to consumers that search once. However, in our model such a firm will sell to consumers with relatively low search costs since they would sample all the firms in a market.

model, that all consumers have the same ranking, but now let this ranking be randomly generated. Specifically, the parameter estimates presented in column 4 of Table 6 represent the averages across 100 different estimations, where each estimation is carried out using a different randomly generated ranking of insurers. The estimates indicate that the ranking-related dummies add very little to the model—all ranking-related dummies are small in magnitude and not significantly different from zero, which suggests that a model with randomly generated rankings does a poor job in comparison to our main specification. Note that these results are actually very similar to the case assuming no search frictions, which we discuss next.

Full information

Most existing discrete choice models of demand assume that consumers have full information. To see to what extent our search model gives different estimates in comparison to a model that assumes full information, we estimate a standard conditional logit model of demand. Notice that the full information model is a special case of our search model; by setting all ranking-related coefficients in equation (8) to zero, we get the standard conditional logit model of demand. Column 4 of Table 6 gives estimates for the full information model, which we estimate by two-stage least squares. The estimated price coefficient is higher in absolute sense in comparison to the estimates for the search model, which indicates that assuming full information when individuals have limited information due to search frictions leads to an upward bias in the absolute value of the estimated price coefficient.

Although we allow consumers in our model to differ in terms of search costs, they are assumed to be similar in terms of their utility parameters. Whether this is an important limitation of the model is an empirical question. Unfortunately, our data lacks a suitable search cost shifter, which is a necessary requirement for estimating a richer search model (see Moraga-González, Sándor, and Wildenbeest, 2015). Instead we have estimated a full information model that allows for heterogeneity in some of the preference parameters, to see if in such a setting this factor makes a substantial difference on the estimated parameters. Specifically, we have estimated a random coefficient logit model in which we control for price endogeneity in a similar way as in the search model (following Berry, Levinsohn, and Pakes, 1995; Nevo, 2000). We assume that the price coefficient follows a normal distribution and estimate its mean and standard deviation. The estimated standard deviation parameter is close to zero and very imprecisely estimated. As a result, the parameter estimates are virtually identical to the estimates in column 4 of Table 6 which suggests that allowing for heterogeneity in the price coefficient is not critical.

Advertising and awareness

In our search model the order in which consumers search is driven by how firms are ranked in terms of expected mean utility. Although we include a rich set of insurer characteristics in our main specification, we do not explicitly control for advertising expenditures, nor for factors that may be related to whether the consumer has a prior relationship with an insurer. As shown by Honka, Hortaçsu, and Vitorino (2017) in the context of banking products, these factors may matter not only for purchase decisions, but also for awareness. The first column of Table 7 shows estimates for a model in which we use national advertising spending for each insurer as a utility shifter.²⁷ When adding advertising spending to the utility function, we allow advertising to affect the probability of buying in two ways: directly, since advertising affects utility, but also indirectly because, assuming a positive marginal utility of advertising, firms that advertise more will appear higher in consumers' rankings of firms and are thus more likely to be searched. The estimates of this specification show that even though advertising spending is highly significant and has a positive marginal utility, the estimated price coefficient is very close to our main specification.

In the second column of Table 7 we add a dummy for whether an insurer also offers Medicare Advantage (MA) plans. We focus on the MA market because Medigap and MA products both serve the Medicare population and are therefore closely related. The estimated parameter is positive and significantly different from zero. A firm's offering of MA plans might serve as a proxy for whether consumers have a prior relationship with an insurer, and therefore one would expect it to have a positive effect on utility. Nevertheless, the results are largely similar to the main results shown in the first column of Table 4.

3.6 Price Elasticities

Table 8 reports the estimated demand elasticities for Medigap Plan F for the top-five insurers in Indiana, ranked by decreasing market share, for both the search model and the full information model. The diagonal entries indicate that estimated demand is more inelastic in the search model than in the full information model. In both panels Mutual of Omaha, the insurer with the largest market share in Indiana, has the most inelastic demand for Medigap Plan F. However, while the

²⁷Our advertising data comes from AdSpender, which provides advertising spending in dollars at the firm level for 2009. Our model assumes that national advertising is exogenous to local market unobserved shocks. This might seem a strong assumption since firms likely choose advertising in response to local market differences. We have tried a specification using calculated firm-market level advertising (de-trended using firm-market specific advertising data of 2012, the first year in which firm-market level data is made available to us). This allows us to use Hausman type instrumental variables (i.e., the insurer's average advertising spending in other markets) for advertising. The results are similar, largely due to the lack of variation in advertising across markets within a firm. We therefore choose to use the 2009 firm-level advertising data, but acknowledge our advertising estimates may be affected by this.

Table 7: Estimation Results with Advertising and MA Presence

Variable	(1)		(2)	
	Coeff.	Std. Err.	Coeff.	Std. Err.
Price	-1.719	(0.089)***	-1.820	(0.089)***
<i>Ranking-related dummies (average)</i>				
Ranked 2 or higher	-0.643		-0.651	
Ranked 3 or higher	-0.619		-0.623	
Ranked 4 or higher	-0.452		-0.444	
Ranked 5 or higher	-0.312		-0.309	
<i>Insurer characteristics</i>				
Firm Weiss safety rating B	0.047	(0.086)	0.024	(0.084)
Firm Weiss safety rating C	-0.091	(0.089)	-0.196	(0.090)**
Firm Weiss safety rating D	-0.270	(0.144)*	-0.169	(0.142)
No Weiss safety rating	-0.062	(0.115)	-0.214	(0.116)*
Age	0.001	(0.001)	0.001	(0.001)
Number of states active 24 – 36	-0.171	(0.075)**	-0.135	(0.074)*
Number of states active > 36	-0.302	(0.084)***	-0.276	(0.083)***
Number of policies offered 4 – 7	0.493	(0.081)***	0.469	(0.080)***
Number of policies offered > 7	0.541	(0.099)***	0.554	(0.097)***
Domicile state	0.378	(0.101)***	0.323	(0.100)***
Offers MA plans			0.646	(0.120)***
log(1 + advertising spending)	0.076	(0.005)***	0.060	(0.006)***
<i>Rating method</i>				
Issued-age	0.488	(0.095)***	0.419	(0.095)***
Community-rated	0.815	(0.313)***	0.568	(0.311)*
<i>Plan dummies</i>				
Plan B	0.874	(0.116)***	0.876	(0.114)***
Plan C	2.195	(0.118)***	2.249	(0.116)***
Plan D	2.117	(0.108)***	2.141	(0.106)***
Plan E	1.655	(0.152)***	1.635	(0.150)***
Plan F	4.214	(0.110)***	4.272	(0.108)***
Plan G	2.659	(0.107)***	2.715	(0.106)***
Plan H	2.352	(0.188)***	2.355	(0.185)***
Plan I	2.607	(0.182)***	2.619	(0.179)***
Plan J	4.798	(0.141)***	4.816	(0.139)***
Plan K	0.116	(0.203)	-0.072	(0.203)
Plan L	0.547	(0.185)***	0.363	(0.185)**
Number of Observations	4,704		4,704	
R^2	0.552		0.555	
State Fixed Effects	yes		yes	
Firm Fixed Effects	no		no	
Plan Fixed Effects	yes		yes	
State-specific ranking-related dummies	yes		yes	

All specifications include a constant. *** Significant at the 1 percent level. ** Significant at the 5 percent level. * Significant at the 10 percent level.

percentage change in market share after a competitor’s price change is the same for all insurers in the full information model, in the search model individuals are more likely to switch to insurers that are similarly ranked. This is intuitive: an individual subscribing to Order of United Commercial Travelers must have relatively low search costs, which means that she is likely to have the other insurers in her choice set, making it more likely that she will find a good deal at an insurer that is also relatively low ranked, such as Admiral Life Insurance Company of America. Similarly, Mutual of Omaha gets a more than proportional share of the high-search-cost individuals, which makes it less likely that many individuals will switch to one of the lower-ranked insurers in case of a price increase of Mutual of Omaha plans.

Table 8: Demand Elasticities Plan F (Indiana)

	Omaha	Anthem	Royal Neighbors	Admiral Life	United Commercial
<i>Search</i>					
Mutual of Omaha	-1.957	0.067	0.006	0.006	0.003
Anthem Insurance	0.294	-2.685	0.028	0.028	0.015
Royal Neighbors Of America	0.161	0.164	-3.037	0.097	0.053
Admiral Life Insurance Company of America	0.133	0.136	0.081	-2.244	0.063
Order of United Commercial Travelers	0.113	0.115	0.069	0.099	-2.724
<i>Full information</i>					
Mutual of Omaha	-2.429	0.100	0.017	0.020	0.013
Anthem Insurance	0.440	-3.546	0.017	0.020	0.013
Royal Neighbors Of America	0.440	0.100	-3.812	0.020	0.013
Admiral Life Insurance Company of America	0.440	0.100	0.017	-2.864	0.013
Order of United Commercial Travelers	0.440	0.100	0.017	0.020	-3.404

Notes: Percentage change in market share of insurer i with a 1 percent change in the price of insurer j , where i indexes rows and j columns. Obtained using the estimates in Column 1 of Table 4 for the search model and Column 5 of Table 6 for the full information model.

Table 9 looks in more detail at substitution patterns among the plans provided by Anthem Insurance, the second-largest provider in Indiana. Note that although the cross-price elasticities are uniformly higher in the search model, as in the full information model a percentage increase in price leads to the same percentage increase in all plans’ market shares. The reason for this is that we assume a single search cost for getting information on all plans sold by an insurer, making searching less of an issue when substituting within an insurer.

3.7 Counterfactuals

Our model can be used to study what happens to prices and market shares as a result of changes in search costs.²⁸ To proceed, we first model the supply side to obtain marginal cost estimates.

²⁸As in most of the existing literature on consumer demand, we do not allow firms to change the plans they offer.

Table 9: Demand Elasticities Anthem Insurance (Indiana)

	Plan A	Plan B	Plan C	Plan D	Plan E	Plan F	Plan G
<i>Search</i>							
Plan A	-2.070	0.001	0.005	0.001	0.001	0.300	0.071
Plan B	0.032	-2.803	0.005	0.001	0.001	0.300	0.071
Plan C	0.032	0.001	-2.852	0.001	0.001	0.300	0.071
Plan D	0.032	0.001	0.005	-3.086	0.001	0.300	0.071
Plan E	0.032	0.001	0.005	0.001	-2.973	0.300	0.071
Plan F	0.032	0.001	0.005	0.001	0.001	-2.685	0.071
Plan G	0.032	0.001	0.005	0.001	0.001	0.300	-2.533
<i>Full information</i>							
Plan A	-2.557	0.000	0.002	0.000	0.000	0.100	0.024
Plan B	0.011	-3.425	0.002	0.000	0.000	0.100	0.024
Plan C	0.011	0.000	-3.488	0.000	0.000	0.100	0.024
Plan D	0.011	0.000	0.002	-3.771	0.000	0.100	0.024
Plan E	0.011	0.000	0.002	0.000	-3.632	0.100	0.024
Plan F	0.011	0.000	0.002	0.000	0.000	-3.546	0.024
Plan G	0.011	0.000	0.002	0.000	0.000	0.100	-3.157

Notes: Percentage change in market share of plan i with a 1 percent change in the price of plan j , where i indexes rows and j columns. Obtained using the estimates in Column 1 of Table 4 for the search model and Column 5 of Table 6 for the full information model.

The profits of an insurer f supplying a subset \mathcal{G}_f of J plan types are given by

$$\Pi_f = \sum_{r \in \mathcal{G}_f} (p_r - mc_r) M s_r(p),$$

where M is the number of consumers in the market. Assuming insurers set prices to maximize profits, the following first-order conditions should be satisfied:

$$s_j(p) + \sum_{r \in \mathcal{G}_f} (p_r - mc_r) \frac{\partial s_r(p)}{\partial p_j} = 0.$$

To obtain the marginal cost of each plan we solve for mc , i.e.,

$$mc = p - \Delta(p)^{-1} s(p), \tag{9}$$

where Δ is a J by J matrix with the element in row j and column r given by

$$\Delta_{jr} = \begin{cases} -\frac{\partial s_r}{\partial p_j}, & \text{if } r \text{ and } j \text{ are supplied by the same insurer;} \\ 0, & \text{otherwise.} \end{cases}$$

For the derivation of the derivatives of the market shares with respect to prices we assume that consumers choose the insurers from which they will obtain quotes before they observe prices (see the Appendix). This means we use the standard assumption that consumers form rational expectations about prices but do not observe price deviations before searching (as in Wolinsky, 1986; Anderson and Renault, 1999).

We obtain marginal cost estimates using equation (9) as well as the demand estimates from Table 4. Once we have obtained marginal cost estimates, we simulate to what extent prices would change if search costs are eliminated. We find that the average price decreases by the amount of \$63, which corresponds to a change of 3.7 percent.²⁹

Table 10: Change in Coefficient of Variation

Plan Type	(A) Observed prices			(B) Simulated prices (full info)			(C) Simulated prices (full info, no firm heterogeneity)		
	average price	max-min	coeff. of variation	average price	max-min	coeff. of variation	average price	max-min	coeff. of variation
A	1,128	869	0.173	1,106	831	0.168	1,095	857	0.182
B	1,567	959	0.181	1,506	920	0.168	1,504	914	0.173
C	1,804	1,490	0.173	1,730	1,430	0.166	1,725	1,439	0.166
D	1,437	927	0.131	1,372	869	0.132	1,371	894	0.134
E	1,509	535	0.110	1,458	498	0.102	1,443	491	0.099
F	1,813	1,889	0.196	1,728	1,804	0.182	1,730	1,800	0.177
G	1,422	868	0.155	1,369	833	0.142	1,366	829	0.134
H	1,468	278	0.049	1,393	272	0.048	1,378	264	0.045
I	1,515	319	0.071	1,450	315	0.072	1,436	306	0.071
J	1,746	721	0.102	1,661	673	0.099	1,656	676	0.105
K	1,030	215	0.079	997	182	0.069	986	208	0.074
L	1,456	322	0.079	1,397	314	0.080	1,398	335	0.084

Notes: Obtained using estimates in Column 1 of Table 4. Plans only include those rated according to attained-age. The coefficient of variation is weighted by sales. Prices are in US\$.

Table 10 summarizes changes in prices and price dispersion by plan type. A comparison of observed prices, which are used to estimate the model and are summarized in column (A), to the simulated prices in column (B) shows that average prices decrease for all plans when moving to a full information equilibrium. Price dispersion also decreases, as measured by either the difference between the maximum and minimum price or the coefficient variation (weighted by market share of each plan),³⁰ although the differences are small.³¹ Note that price dispersion is not expected

²⁹The case of eliminating search costs can be interpreted as consumers being offered a price menu for all plans sold in a state at no cost. Currently no such information is made publicly available. Some aggregate data on premium ranges at the state-plan type level have become accessible through medicare.gov starting in 2012.

³⁰Using an unweighted coefficient of variation gives similar results, although the differences between observed and simulated prices are larger for unweighted measures.

³¹The results in Table 10 are obtained using the estimates from Column 1 of Table 4.

to disappear in the case of zero search costs: firms with higher marginal cost or firms with more favorable characteristics will find it optimal to set higher prices, which results in price dispersion even in the case of full information. To see how insurer differentiation affects prices, in column (C) of Table 10 we show simulated prices in case of full information but with insurer heterogeneity eliminated. Although the average prices decrease is with \$67 now slightly larger, there is not much difference with the results shown in column (B) in terms of average prices and price variations at the plan-type level. A large part of this reflect marginal cost heterogeneity: firms with higher marginal cost charge higher prices, resulting in substantial price dispersion, even in the absence of both insurer heterogeneity and search costs.

To analyze changes in consumer welfare when moving to full information, we use compensating variation (see also Small and Rosen, 1981; Nevo, 2000), which in our search model corresponds to

$$CV = \frac{\sum_{k=1}^N \mu_k^{\text{post}} \left[\ln \left(1 + \sum_{\ell=1}^k \exp \phi_{\ell}^{\text{post}} \right) - (k-1) \tilde{c}_k^{\text{post}} \right] - \sum_{k=1}^N \mu_k^{\text{pre}} \left[\ln \left(1 + \sum_{\ell=1}^k \exp \phi_{\ell}^{\text{pre}} \right) - (k-1) \tilde{c}_k^{\text{pre}} \right]}{\alpha}, \quad (10)$$

where \tilde{c}_k is the average search cost of consumers searching k times. This amounts to the expected maximum utility when searching k times (measured in dollars by normalizing by the price coefficient), averaged over all k groups of consumers, taking average search costs for individuals in each group into account.³² The average change in consumer surplus when moving to the simulated zero search cost equilibrium is \$236, which is close to four times the magnitude of the average price decrease (\$63). Such changes in consumer welfare are driven by multiple factors such as savings on search costs, reduction in premiums, and expansion of the total market.

It is worthwhile to note that we do not specifically model consumer selection into Medigap plans. The existing literature has found mixed evidence of selection in Medigap plans (Fang, Keane, and Silverman, 2008). Due to regulations in Medigap such as product standardization and guarantee issues, we suspect that firms are limited in their capability to engage in cream-skimming. Nevertheless, we do offer some discussion on how our results would be affected if firms are able to take selection into account when setting prices. One obvious complication is that marginal costs may be related to prices, which leads to an additional channel through which prices affect profits. For example, an increase in price may scare away price-sensitive consumers, who may be more likely to have worse than average health conditions. In such a scenario marginal costs would be negatively related to prices, which means our current model is likely to overestimate marginal costs

³²For the case in which the new search costs are zero, we set $\mu_N = 1$ while all other μ_k 's are zero.

as well as the new equilibrium prices in the counterfactual in which we assume there are no search frictions. Moreover, if this would be the case, our counterfactual results should be considered as a lower bound on the predicted price decreases.

4 Monte Carlo Experiments

In this section we use several Monte Carlo experiments to study the performance of our estimation procedure. Our first goal is to see whether our estimation procedure is able to recover the search cost distribution as well as the parameters of the utility specification. Next, we analyze the performance of our estimation procedure in situations where there is heterogeneity in how consumers value observed characteristics. Our model is not able to capture this, so we proceed by simulating data using a richer framework similar to Moraga-González, Sándor, and Wildenbeest (2015), which allows us to assess to what extent search cost and preference parameter estimates will be biased when using our model for estimation instead of the procedure outlined in Moraga-González, Sándor, and Wildenbeest (2015).

The setup of the Monte Carlo experiments is as follows. We simulate data for 50 different markets, where each market has three different firms, each selling one product. Prices are randomly drawn from a distribution that is uniform distribution between 3 and 4. We allow for a constant as well as two product attributes that are randomly drawn from a uniform distribution between 1 and 2, with parameter values as given by the first column in Table 11. The unobserved characteristic is drawn from a normal distribution with mean zero and standard deviation 0.1.

To see whether our estimation procedure is able to recover the parameters of the model outlined in Section 2, we let consumers' search costs be drawn from a lognormal distribution with mean and standard deviation of the associated normal distribution set to zero and three, respectively. Column (A) of Table 11 gives the mean and standard deviation of the parameter estimates across 100 replications. All preference parameters are close to the true coefficients and have low standard deviation. The reported search cost parameters are obtained by fitting a lognormal distribution through the estimated points on the search cost distribution and then averaged across the 100 replications. Both the mean and the standard deviation parameters are close to the true coefficients. This is also shown in Figure 4(a), which plots the estimated search cost CDF (solid curve), its 90% confidence interval, and the actual CDF (black dashed curve).

To get a better sense to what extent the estimates that are obtained using our procedure are sensitive to preference heterogeneity in observed characteristics, Columns (B) and (C) of Table 11

Table 11: Results Monte Carlo Experiments

Variable	True Coeff.	Main model		Preference heterogeneity				
		(A)		(B)	$\sigma_\lambda = 1.00$		(C)	$\sigma_\lambda = 0.01$
		Coeff.	Std.dev.	Coeff.	Std.dev.	Coeff.	Std.dev.	
<i>Preference parameters</i>								
Constant	-1.000	-1.086	(0.176)	-2.109	(0.388)	-2.063	(0.391)	
Attribute 1	4.000	4.011	(0.079)	3.368	(0.115)	4.127	(0.237)	
Attribute 2	2.000	2.012	(0.050)	1.651	(0.084)	1.992	(0.152)	
Price	-3.000	-2.985	(0.050)	-2.583	(0.104)	-3.098	(0.182)	
<i>Search cost parameters</i>								
Mean	0.000	-0.026	(0.100)	-6.449	(0.627)	-4.209	(0.289)	
St.dev.	3.000	2.982	(0.110)	3.392	(0.499)	1.932	(0.256)	

give results of Monte Carlo experiments in which we use a more general model to generate the data, and use that data to estimate our model. The more general model we use to generate the data is similar to the model in Moraga-González, Sándor, and Wildenbeest (2015). An important difference between this model and our model is the way in which search costs are modeled. Specifically, Moraga-González, Sándor, and Wildenbeest (2015) let search costs be consumer-choice-set specific, i.e., consumer i 's search cost for choice set S is

$$c_{iS} = t'_{iS}\gamma + \sigma_\lambda \lambda_{iS},$$

where t'_{iS} contains a vector of search cost shifters, λ_{iS} is a consumer-choice-set specific search cost shifter with scale parameter σ_λ . According to Moraga-González, Sándor, and Wildenbeest (2015), by assuming $-\lambda_{iS}$ is IID Type I Extreme Value distributed, the probability that product j is purchased is

$$s_j = \int s_{ij} dF_\tau(\tau_i), \quad (11)$$

where F_τ is the CDF of the vector of consumer-specific random variables τ_i and s_{ij} is the probability that consumer i buys product j , i.e.,

$$s_{ij} = \sum_{S \in S_j} \frac{\exp[m_{iS}/\sigma_\lambda]}{\sum_{S' \in S} \exp[m_{iS'}/\sigma_\lambda]} \frac{\exp[\delta_{ij}]}{1 + \sum_{r \in S} \exp[\delta_{ir}]},$$

where $m_{iS} = \log\left(1 + \sum_{j \in S} \exp[\delta_{ij}]\right) - t'_{iS}\gamma$ and S_j is the set of choice sets that contains product j .

We use the market share expression in equation (11) to generate data. We use the same




Figure 4: Search cost distribution Monte Carlo experiments

preference parameters as we used to simulate the data for the main model, but in addition we also allow for a random price coefficient. Specifically, we assume the distribution for the price coefficient is normal with mean -3 and standard deviation 0.2 . We assume one search cost shifter, which, as in the main specification, is drawn from a log-normal distribution with parameters zero and three. An important difference with our model is that consumers no longer have a similar ranking of firms: because of the random price coefficient, a consumer's price sensitivity is going to affect her ordering of firms. Moreover, the choice-set specific search cost shock creates additional randomness in consumers' rankings of firms.

Moraga-González, Sándor, and Wildenbeest (2015) assume $\sigma_\lambda = 1$ when estimating their model, which is also what we use to generate the data that is used in column (B) of Table 11. As shown in the table, all the preference parameters as well as the mean search cost parameters are more than two standard deviations away from the true coefficients. One reason for why the estimates are very different is because of the relatively large variance of the choice-set specific search cost error: the higher the scale parameter, the more likely it is that lower ranked options are searched and the higher their market shares. To rationalize this using our model, search cost estimates need to be lower than they really are. Column (C) shows results when the data is generated using $\sigma_\lambda = 0.01$, which implies that the choice-set specific search cost shock has a much smaller impact on the generated market shares, making the model that is used to generate the data more similar to our model. Indeed, all preference parameter estimates except for the constant are close to the true parameter values. However, estimated search costs are still lower than the search costs that are used to generate the model. This is also illustrated in Figure 4(b). The reason for this is that even with a low variance of the search cost shock, lower ranked alternatives are still more likely to be bought than in our model, so if our model is used for estimation then relatively low search costs

are needed to explain the higher market shares for these products.

5 Conclusion and Discussion

In this paper we develop a model of consumer search to explain price and price variation for differentiated products. Our model allows for both vertical and horizontal product differentiation. The model can be estimated using constrained two-stage least squares. When applying our model to the Medigap market, we find that our search model provides a better fit than a model that assumes consumers have full information. Assuming full information when consumers in fact have limited choice sets leads to a higher estimated price coefficient in absolute sense. The estimated search costs are substantial: median search costs are estimated to be \$30 in this market. When moving to a zero-search-cost equilibrium, we find average prices fall by \$63 and average consumer welfare increases by \$236.

One of the main contributions of our paper is to provide a methodology that allows us to non-parametrically estimate search costs for differentiated goods when using aggregate data on consumer demand. Our model is general enough to be applied to markets in which both search frictions and product differentiation coexist, and that lack suitable search cost shifters. Monte Carlo simulations suggest estimation of utility parameters works reasonably well even in situations where there is some heterogeneity in how consumer value observed characteristics, although the results with respect to the estimated search cost distribution should be interpreted cautiously since search costs may be underestimated.

References

- ANDERSON, S. AND RENAULT, R. “Pricing, Product Diversity, and Search Costs: A Bertrand-Chamberlin-Diamond Model.” *RAND Journal of Economics*, Vol. 30 (1999), pp. 719–735.
- ARMSTRONG, M. “Ordered Consumer Search,” *Journal of the European Economic Association*, Vol. 15 (2017), pp. 989–1024.
- BERRY, S., LEVINSOHN, J., AND PAKES, A. “Automobile Prices in Market Equilibrium.” *Econometrica*, Vol. 63 (1995), pp. 841–890.
- BROWN, J. AND GOOLSBEE, A. “Does the Internet Make Markets More Competitive? Evidence from the Life Insurance Industry.” *Journal of Political Economy*, Vol. 110 (2002), pp. 481–507.
- BURDETT, K., JUDD, K. “Equilibrium Price Dispersion.” *Econometrica*, Vol. 51 (1983), pp. 955–969.
- CARLSON, J. AND MCAFEE, R. “Discrete Equilibrium Price Dispersion.” *Journal of Political Economy*, Vol. 91 (1983), pp. 480–493.
- CEBUL, R., REBITZER, J., TAYLOR, L., VOTRUBA, M. “Unhealthy Insurance Markets: Search Frictions and the Cost and Quality of Health Insurance.” *American Economic Review*, Vol. 101 (2011), pp. 1842–1871.
- CHADE, H. AND SMITH, L. “Simultaneous Search.” *Econometrica*, Vol. 74 (2006), pp. 1294–1307.
- De los Santos, Babur, Ali Hortaçsu, and Matthijs R. Wildenbeest: “Testing Models of Consumer Search Using Data on Web Browsing and Purchasing Behavior.” *American Economic Review* 102, 2955–2980, 2012.
- ERICSON, K. “Consumer Inertia and Firm Pricing in the Medicare Part D Prescription Drug Insurance Exchange.” *American Economic Journal: Economic Policy*, Vol. 6 (2014), pp. 38–64.
- ETTNER, S. “Adverse Selection and the Purchase of Medigap Insurance by the Elderly.” *Journal of Health Economics*, Vol. 16 (1997), pp. 543–562.
- FANG, H., KEANE, M., AND SILVERMAN, D. “Sources of Advantageous Selection: Evidence from the Medigap Insurance Market.” *Journal of Political Economy*, Vol. 116 (2008), pp. 303–350.

- FINKELSTEIN, A. “Minimum Standards, Insurance Regulation and Adverse Selection: Evidence from the Medigap Market.” *Journal of Public Economics*, Vol. 88 (2004), pp. 2515–2547.
- FOX, P., SNYDER, R., AND RICE, T. “Medigap Reform Legislation of 1990: A 10-Year Review.” *Health Care Financing Review*, Vol. 24 (2003), pp. 1–17.
- HONKA, E., HORTAÇSU, A., AND VITORINO, M.A. “Advertising, Consumer Awareness, and Choice: Evidence from the U.S. Banking Industry.” *RAND Journal of Economics*, Vol. 48 (2017), pp. 611–646.
- HORTAÇSU, A. AND SYVERSON, C. “Product Differentiation, Search Costs, and Competition in the Mutual Fund Industry: a Case Study of S&P 500 Index Funds.” *Quarterly Journal of Economics*, Vol. 119 (2004), pp. 403–456.
- HURD, M. AND MCGARRY, K. “Medical Insurance and the Use of Health Care Services by the Elderly.” *Journal of Health Economics*, Vol. 16 (1997), pp. 129–154.
- MAESTAS, N., SCHROEDER, M., AND GOLDMAN, D. “Price Variation in Markets with Homogeneous Goods: The Case of Medigap.” Working Paper No. 14679, NBER, Cambridge, 2009.
- ROBST, J. “Estimation of a Hedonic Pricing Model for Medigap Insurance.” *Health Services Research*, Vol. 41 (2006), pp. 2097–2113.
- MORAGA-GONZÁLEZ, J.L., SÁNDOR, Z., AND WILDENBEEST, M. “Consumer Search and Prices in the Automobile Market.” *Mimeo*, 2012.
- NEVO, A. “Mergers with Differentiated Products: the Case of the Ready-to-Eat Cereal Industry.” *RAND Journal of Economics*, Vol. 31 (2000), pp. 395–421.
- SHEINGOLD, S., SHARTZER, A., AND LY, D. “Variation and Trends in Medigap Premiums.” ASPE Report, U.S. Department of Health and Human Services, 2011.
- SMALL, K. AND ROSEN, H. “Applied Welfare Economics with Discrete Choice Models.” *Econometrica*, Vol. 49 (1981) 105–130.
- SORENSEN, A. “Equilibrium Price Dispersion in Retail Markets for Prescription Drugs.” *Journal of Political Economy*, Vol. 108 (2000), pp. 833–850.
- STAHL, D. “Oligopolistic Pricing with Sequential Consumer Search.” *American Economic Review*, Vol. 79 (1989), pp. 700–712.

- STARC, A. “Insurer Pricing and Consumer Welfare: Evidence from Medigap.” *RAND Journal of Economics*, Vol. 45 (2014), pp. 198–220.
- STIGLER, G. “The Economics of Information.” *Journal of Political Economy*, Vol. 69 (1961), pp. 213–225.
- WEITZMAN, M. “Optimal Search for the Best Alternative.” *Econometrica*, Vol. 47 (1979), pp. 641–654.
- WILDENBEEST, M. “An Empirical Model of Search with Vertically Differentiated products.” *RAND Journal of Economics*, Vol. 42 (2011), pp. 729–757.
- WOLINSKY, A. “True Monopolistic Competition as a Result of Imperfect Information,” *Quarterly Journal of Economics*, Vol. 101 (1986), pp. 493–511.
- WOLFE, J. AND GODDEERIS, J. “Adverse Selection, Moral Hazard and Wealth Effects in the Medigap Insurance Market.” *Journal of Health Economics*, Vol. 10 (1991), pp. 433–459.

APPENDIX

The derivative of the market share of firm f 's product j with respect to its own price is given by

$$\frac{\partial s_{jf}}{\partial p_{jf}} = -\alpha \cdot \sum_{k=f}^N \frac{e^{\delta_{jf}} \left(1 + \sum_{\ell=1}^k e^{\phi_{\ell}} - e^{\delta_{jf}}\right)}{\left(1 + \sum_{\ell=1}^k e^{\phi_{\ell}}\right)^2} \cdot \mu_k.$$

The derivative of the market share of firm f 's plan j with respect to the price of firm f 's product g is

$$\frac{\partial s_{jf}}{\partial p_{gf}} = \alpha \cdot \sum_{k=f}^N \frac{e^{\delta_{jf}} e^{\delta_{gf}}}{\left(1 + \sum_{\ell=1}^k e^{\phi_{\ell}}\right)^2} \cdot \mu_k.$$

The derivative of the market share of firm f 's product j with respect to the price of firm h 's product g is

$$\frac{\partial s_{jf}}{\partial p_{gh}} = \alpha \cdot \sum_{k=\max[f,h]}^N \frac{e^{\delta_{jf}} e^{\delta_{gh}}}{\left(1 + \sum_{\ell=1}^k e^{\phi_{\ell}}\right)^2} \cdot \mu_k.$$