

Estimation of Search Frictions in the British Electricity Market*

Monica Giuliatti
Nottingham University Business School
Nottingham
United Kingdom

Michael Waterson
Department of Economics
University of Warwick
United Kingdom

Matthijs R. Wildenbeest
Kelley School of Business
Indiana University
USA

February 2011

Abstract

This paper studies consumer search and pricing behaviour in the British domestic electricity market following its opening to competition in 1999. We develop a sequential search model in which an incumbent and an entrant group compete for consumers who find it costly to obtain information on prices other than from their current supplier. We use a large data set on prices and input costs to structurally estimate the model. Our estimates indicate that consumer search costs must be relatively high in order to rationalize observed pricing patterns. We confront our estimates with observed switching behaviour and find they match well.

Keywords: electricity, consumer search, price competition

JEL classification: C14, D83, L13

*Amongst others, we are indebted to Mike Baye, Greg Crawford, José Luis Moraga-González, Ali Hortaçsu, Maarten Janssen, Graham Loomes, Régis Renault, Patrick Rey, Michael Rauh and Chris Wilson for useful comments and discussion. Our paper has been improved by comments arising from presentations at Indiana University, the Tinbergen Institute Rotterdam, the University of Warwick, EARIE 2010, Istanbul and the 2010 Safe International Workshop on Quantitative Analysis of Energy Markets, Verona. Our research was facilitated by ESRC grant RES 00-22-1686 to Giuliatti and Waterson that, amongst other things, funded Wildenbeest's visit to Warwick.

1 Introduction

From the consumer's point of view, electricity is one of life's essentials. It is also a very homogeneous product; one company's electricity is the same as any other's.¹ Bertrand-type economic arguments then suggest that if supplier prices differ, all consumers flock to their cheapest. In Britain, this has not happened, despite every consumer having had the opportunity to switch for over 10 years and despite a higher proportion switching than in any other European country or US state (Defeuilley, 2009) and a majority of consumers no longer with their incumbent supplier. In this sense, market competition has not worked well. Faced with six major suppliers offering different prices, significant switching continues to occur² but a majority of consumers still pay over the odds and company prices still diverge considerably, something that has prompted several enquiries. Within the period and bill type we investigate, the price range consistently exceeds 20%, but it continues beyond this. For example, in July 2010 a consumer living in the East Midlands, anticipating use of 3300kWh of electricity a year (a medium quantity), paying by direct debit, would pay between £289 and £494 per year.³ Why do we see such large price ranges?

Our analysis of this significant price divergence focuses on search costs as a candidate explanation. The resulting structural model estimates suggest that even search costs at lower quartile are high, significant as a fraction of the bill, though falling through time. A particular novelty of our paper is that we can independently verify the characteristics of these estimated costs using market information on changes in supplier shares, finding an excellent match between the two. We also investigate several other factors that are of potential importance, for example switching costs, service quality and promotional activity.

¹ This is particularly so if, as in our case, (i) each supplier's electricity would come through the same distributor's wires so that suppliers do not differ amongst themselves in terms of service reliability and (ii) "green" electricity does not command a premium from consumers.

² The proportion of consumers who have "ever switched" by 2005 varies across regions of the country, with the lowest proportion being 31%, the highest 56% and a simple mean of 46% (Ofgem, 2006).

³ Price (i.e. annual bill) quotes retrieved online from uSwitch.com on 22 July 2010, rounded to whole £. For comparison, the default option would result in a bill of £416 per year, so moving to the cheapest gives a 30% saving.

How might search costs explain the remarkably big observed differences in prices across suppliers of electricity? The basic approach is relatively straightforward: if some consumers have more information on prices than others, for example, because these consumers have been shopping around (“shoppers”), firms have an incentive to offer these consumers a lower price than non-shoppers. More efficient firms might then cater to the shoppers, while the firms with higher marginal cost serve the non-shoppers, resulting in price dispersion. Alternatively, with more relevance for our market situation, firms might find it optimal to randomize prices in order to balance the incentive to set a low price to maximize surplus from shoppers against the incentive to set a high price to maximize surplus from the non-shoppers.

We build a model that focuses on optimal price setting behaviour of an incumbent and several entrants under the presence of costly sequential consumer search. Both the incumbent and the set of similar entrants have an endogenous set of loyal customers, the difference being that by definition the incumbent has a bigger set of loyal customers to start with than each of the entrants. Although the model is tailored to the early years of the consumer electricity market in Britain, in principle it applies to many markets across the world that have been recently liberalized (for instance telecommunication and gas). Our theoretical results show that in equilibrium the incumbent supplier sets its price equal to the monopoly price, while the entrants compete by randomly drawing their prices from a price distribution. Simulations indicate that entrant prices are increasing in the number of consumers being served by the entrants.

In the empirical part of our study we use the equilibrium of the theoretical search model to structurally estimate search costs in the British electricity market. For this purpose we use a large data set of prices collected in the period between 2002 and 2005,⁴ as well as information on cost side variables. Our results indicate that even though average margins went up for most of the sampling period, search costs have gone down over time and we are able to reconcile these findings.

⁴ The date range is deliberate, as we explain below. It lies between the final liberalization move and significant structural changes.

The paper builds on both the theoretical and empirical literature on consumer search. Burdett and Judd (1983) show that price dispersion can result as an equilibrium phenomenon without any ex ante heterogeneity in a non-sequential search setting. Stahl (1989) builds a sequential search model with homogeneous firms and a two-type distribution of search cost. Our theory model has its origins in Stahl (1996), who allows for a continuous distribution of search cost in a sequential search setting. As in Baye, Kovenock and De Vries (1992) and Kocas and Kiyak (2006) we assume consumers do not allocate themselves equally across firms: the incumbent in our model has a bigger share of loyal consumers to begin with, which leads to an asymmetry in pricing strategy between the incumbent and entrants.

Related empirical papers include Hong and Shum (2006), Hortaçsu and Syverson (2004), Moraga-González and Wildenbeest (2008), Honka (2010), Hortaçsu et al. (2010) and Waddams and Wilson (2010). Hong and Shum (2006) show how to structurally estimate a non-sequential search and sequential search model using only price data. Hortaçsu and Syverson (2004) estimate search cost as well as the impact of vertical product differentiation in the US mutual fund industry. Moraga-González and Wildenbeest (2008) show how to estimate a non-sequential search model by maximum likelihood. Honka (2010) empirically identifies switching and search costs in the U.S. auto insurance market, based on detailed information about the consumers' search process. Hortaçsu et al. (2010) investigate the determinants of supplier choice for electricity consumers in Texas taking into account non-price differentiation factors, in addition to search and switching costs. Waddams and Wilson (2010) provide sample evidence of misguided choice of supplier by UK electricity consumers, leading to monetary losses rather than gains, which the authors attribute to decision errors and inattention on the part of consumers. Giuliatti et al. (2010a) demonstrate the basic features of the market we investigate here, but focus on description rather than estimation.

There are obviously other potential explanations for price divergence between suppliers. Switching costs can explain price divergence between the incumbent supplier and others, but not the substantial differences between different entrant suppliers. Our identifying assumption on search costs is that the *switching* costs of those who switch *without searching* (see later) have similar distributional

characteristics to the switching costs for those who switch *after* searching. In addition to these, differences in service level or branding might differentiate one provider from another, resulting in suppliers with more favourable characteristics selling at a premium, or gaining greater share. We examine the potential impact of these factors also, but find them of secondary importance.

If, as a result of search costs, consumer prices contain significant divergences, from each other and from costs, why does it matter? Competition was introduced into the electricity supply market largely in order to avoid the burden of regulation that previously existed (and exists still in many parts of the world, given the consumer welfare impact of the market). Our results, echoing concerns expressed more broadly in Stern (2010) and leading to OFGEM's Energy Supply Probe (2008), suggest caution in the assumption that competition will do the job, even if well-engineered to reduce the burden on consumers to a minimum.

In the next section we give an overview of the British electricity market and discuss the several sources of price dispersion in the market. We develop our model in Section 3. In Section 4 we present our estimation method, which we apply to the data on the British electricity market in Section 5. Section 6 contains our test of the implications of our estimated switching costs for switching activity against actual switching behaviour. Section 7 concludes.

2 The British electricity market

The UK was one of the first countries to liberalize its consumer electricity markets, by separating the activities of supply and of distribution and regulating the latter. This led to a significant increase in the number of competitors offering to supply domestic consumers⁵ in each of fourteen geographical regions.⁶ Regional incumbents not only started competing in other regions, becoming entrants in those regions, but US, German, French and Spanish firms have entered or have taken over existing firms as

⁵ We do not consider supply to non-domestic properties in this paper.

⁶ The fourteen regions are: Eastern, East Midlands, London, Midlands, Merseyside & North Wales, Northern, North West, South East, North Scotland, South Scotland, Southern, South Wales, South West, and Yorkshire. Note that this list does not include Northern Ireland, where the whole electricity regime is substantially different.

well. Switching started in 1999 and since March 2002 all supply price regulation has disappeared.⁷ The average number of suppliers in each region has decreased since then, from 11 in the beginning of 2002 to on average 7 near the end of 2005, but since late 2002, the largest six have consistently accounted for over 99% of supplies to domestic consumers (OFGEM, 2003), so we can view the number of effective competitors as essentially fixed.

	2002	2003	2004	2005
<i>Retail prices</i>				
Mean	224.35	229.73	245.50	275.80
(Max – Min)/Min	18.4%	18.7%	20.7%	15.9%
<i>Marginal costs</i>				
Distribution cost	56.97	56.58	54.16	57.25
Wholesale price	63.35	54.77	62.43	90.86
Other cost	52.47	50.79	52.15	66.75
Total	172.79	162.04	168.74	214.87
<i>Margins</i>				
Mean	44.83	60.80	69.39	52.66
Max	66.95	80.71	92.03	75.39
Min	29.77	42.68	47.90	35.80
Max – Min	37.18	38.03	44.14	39.60
Second highest	53.32	69.91	80.02	62.41
Incumbent	66.92	80.39	90.05	69.67
Price-cost margin	29.9%	41.8%	45.6%	29.2%
Market share incumbent	65.7%	61.8%	59.6%	58.2%
Number of firms	11	8	8	7

Notes: Retail prices and costs are denoted in British pounds and calculated for medium users (3,300kWh/year). All figures are averages over all suppliers and regions. Averages over 2002 exclude February.

Table 1: Retail prices and margins medium user (direct debit)

Table 1 gives an overview of how retail prices and margins have evolved throughout the sample period.⁸ Retail prices of electricity differ by payment method and geographical location. Consumers can generally pay their electricity bill in three different ways: direct debit, standard credit and prepayment, with direct debit being over 40% of the market. Table 1 as well as Figure 1(a) show how average electricity retail prices have evolved over time for consumers who pay by direct debit. Since all retail prices are at least two-part we focus on the charge for medium users, consuming

⁷ Regulation of high voltage transmission and of local distribution prices remains in place--for more detail see Giulietti et al (2010a).

⁸ The data is described in more detail in Appendix A.

3,300kWh/year (hereafter called price). There is quite a bit of variation in prices, even within regions. As shown in Table 1 the percentage difference in retail prices between the most and least expensive suppliers is on average about 18%. Given that a medium user then spent on average about £250 per year on electricity this means that for some consumers substantial savings can be made by changing supplier. For instance, in June 2002 a moderate user of electricity living in Birmingham (Midlands) could save about £41 a year by switching from the incumbent Npower to Scottish Hydro Electric (now SSE), which at that time was the least expensive supplier in the Midlands.

We stop the analysis of price differences at the end of 2005 because after that period, the market became somewhat different in character. By March 2006, only 4% of consumers had arranged their electricity supplier directly through the Internet. Since then, this number has increased rapidly, with reports that in 2009, 26 percent of consumers purchase electricity in this way. The proportion of customers on internet only tariffs reached 12% in 2009 (OFGEM, 2010). Accompanying this change in consumer behavior, the suppliers have vastly increased the range and complexity of tariffs on offer. Whereas in the period we study here, each supplier effectively had a single tariff on offer, now each has a considerable range.⁹

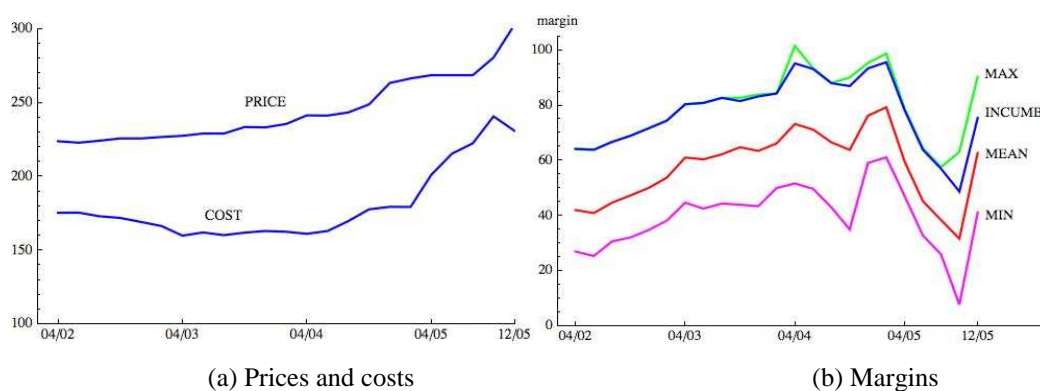


Figure 1: Retail prices and margins over time

During the sample period there were large fluctuations in input fuel prices—wholesale gas prices decreased in 2002, were relatively flat in 2003, but were rising in 2004 and 2005—gas is the most significant fuel, followed by coal. This is also illustrated by the

⁹ In the case noted in the introduction, the website gave a listing of 74 tariffs for the consumer to choose amongst, including the six major companies with many tariffs each and a few additional “virtual suppliers”, who have a tiny market share.

cost curve plotted in Figure 1(a). Here our measure of cost also includes the cost of distributing the electricity as well as costs related to transmission, network losses, balancing, renewables obligation, carbon emission, metering, and energy efficiency commitment. The large fluctuations in marginal cost make the retail price rather uninformative about how the introduction of competition has affected pricing in the retail electricity market. Figure 1(b), however, shows how the average margin has evolved over time for direct debit consumers. Apart from the sharp drop in 2005, which is likely caused by a sudden increase in fuel input prices, margins have been increasing during most of the sample period.¹⁰

There are several explanations for the observed differences in retail prices across electricity providers. One explanation is that although electricity is a homogenous product, consumers might value provider characteristics in different ways. For example, some suppliers might be offering more service than others, or might have a more favourable brand image than other suppliers, resulting in higher prices. To test for this we regress prices on a constant and a set of provider dummies. To control for regional differences we include regional dummies. Moreover, we add time dummies to control for differences over time. The resulting R^2 value indicates that around 39% of variation in prices that is not explained by time and regional fixed effects can be explained by firm fixed effects.¹¹ This means that although part of the variation in price is probably due to firm heterogeneity, still a substantial 61% of the variation cannot be explained by supplier differences. Moreover, as we will explain below, our model does take heterogeneity between the entrants and incumbent into account. If we allow for an additional incumbent dummy in our regression the percentage of variation left that cannot be explained by supplier differences goes up to 71%.

¹⁰ A fuller analysis of the development of average retail margins over time is contained in Giuliotti et al. (2010b).

¹¹ Detailed results are available on request.




Figure 2: Pricing patterns Eastern region

An additional explanation for the price dispersion we observe in the data is that it is costly for consumers to switch. If consumers face a switching cost when changing suppliers a rational consumer will switch to another provider only if the gains from switching more than offset the cost of switching. Providers can increase the benefits of switching by setting lower prices. Although this can provide a reason for the incumbent firm (particularly in early stages) charging higher prices than entrants, it provides no reason for prices to differ across entrant firms. As Figure 2 illustrates, there are very substantial, but changing, differences across entrant suppliers.¹²

Part of this unexplained variation might be due to search frictions. As shown by a large number of theoretical papers (e.g., Varian, 1980; Burdett and Judd, 1983; Stahl, 1989; Baye and Morgan, 2001; and Janssen and Moraga-González, 2004), price dispersion can result as an equilibrium outcome, even in markets with homogenous firms. If some consumers search more than others, firms have incentives to either set a high price, and maximize surplus from consumers who do not search that much, or to set a low price in order to maximize surplus from price comparing consumers. As a result it might be optimal for firms to start randomizing prices, with price dispersion as a result. Consumers in Britain who plan to change electricity provider have on average six major options. Since prices are in general not similar across the available providers, consumers first have to (decide how much to) gather information on prices before making a decision to which supplier to switch. This information gathering can

¹² TXU, succeeded by Powergen (E.On), was the incumbent. Basic Power, one of those depicted, went into administration in 2005.

go through several channels. For example, a consumer could talk with a sales person, phone a company or visit the website of an electricity provider or, more recently, could go to one of the many comprehensive price comparison sites. Collecting information on prices will be time consuming, and as such it can be considered to be costly for consumers. Consumers differ in their in their opportunity cost of time and search method, which will create the heterogeneity on the consumer side needed to make firms willing to charge different prices.¹³ As mentioned above, in a typical search model with homogenous firms the equilibrium is in mixed strategies. In fact, in our sample, it is not possible to reject the hypothesis that price movements are consistent with a mixed strategy equilibrium; see Appendix C.

An important difference between search and switching cost is that search costs are generally made upfront, while switching cost only arises when a consumer actually switches. For instance, by adding switching costs into a search model for differentiated products, Wilson (2009) shows that search costs lead to consistently larger anticompetitive effects than an equivalent level of switching costs. Although it is difficult to distinguish switching costs from search costs empirically, because we know that a significant proportion (30%) of consumers report as switching without searching we can attempt to distinguish between the two.

In the next section we build a model that takes search frictions as an important determinant in price setting behaviour of electricity providers. In addition the model makes a distinction between optimal behaviour of the incumbent and of the entrants to the market. Before liberalization of the market in 1999, each regional market was a local monopoly. Each of the incumbents by definition had a large number of loyal consumers when the market opened up, which made it costly for the incumbent to lower prices enough to prevent consumers from moving to one of the entrants.¹⁴ Entrants, on the other hand, did not have any loyal consumers to start with, so faced zero cost in setting margins at least as low to cover search costs. Therefore, in each market one would expect the prices initially set by entrants to be much lower than the

¹³ In some markets, price advertising would significantly reduce consumer search costs. In this case, as is shown in Appendix D, there is rather little advertising in this market. Moreover, because different consumers have different consumption levels, the price message is not (in principle or in practice) easy to target and firms rely on generalities such as “low prices” and “excellent service”.

¹⁴ In the period under discussion, price discrimination between old and new customers was ruled out by the regulator.

price set by the incumbent. Over time, lower prices at the entrants will induce consumers to switch from the incumbent to the entrants. As a result it becomes more costly for the entrants to set very competitive prices, while at the same time it will be more costly for an incumbent to set a very high price, so one would expect the prices of the incumbent to be higher initially, but to converge over time to those of the entrants. In our data this indeed seems to be the case. Figure 1(b) illustrates that while in the first half of the sampling period the incumbent always has the highest price, in the second half of the sample it is no longer the case that the incumbent is always the highest priced supplier of electricity. Figure 2 illustrates regionally, using the Eastern region. Our model developed in the next section will have this asymmetry as an important element in explaining price patterns across electricity providers.

3 The model

Our model is a unit demand version of Stahl (1996) with asymmetric firms selling a homogenous good. We introduce asymmetry on the supply side of the market by modeling the strategic interaction of one incumbent firm and a set of N symmetric entrants. More specifically, we assume a mass of consumers normalized to one inelastically demands at most one unit of electricity.¹⁵ At the beginning of a period each consumer is local to one electricity provider. Assume the incumbent provider has a market share of λ consumers, whereas all the entrants together have $1-\lambda$ consumers. Though we vary λ over time, throughout the analysis we assume the incumbent has more local consumers than each of the entrants, i.e., $\lambda > (1-\lambda)/N$. Consumers' maximum willingness to pay is equal to v . The common marginal cost of the providers is denoted by r .

At the beginning of a period, consumers only observe the price of the provider that provided them with electricity in the previous period. To obtain additional price quotations consumers have to engage in costly search. We assume consumer search sequentially for additional price information, that is, consumers determine after each search whether to obtain an additional price quotation or not. Furthermore, let the

¹⁵ We are assuming consumption does not vary in the short run as a result of the price differences experienced.

consumers be characterized by a search cost c , which is drawn at the beginning of a period from some search cost distribution $G(c)$, with density $g(c) > 0$.

At the beginning of the period each provider sets a price taking consumer search and switching behaviour as well as pricing strategies of the other firms as given. Let the distribution of tariffs set by the entrants be denoted by $F(p)$, with support between \underline{p} and \bar{p} . As shown by Stahl (1996) for the case of symmetric firms, if there is an (arbitrary small) atom of shoppers there are no pure-strategy symmetric Nash Equilibria. Moreover, for the equilibrium price distribution to be a symmetric Nash equilibrium it needs to be atomless.

First consider optimal consumer search behaviour for consumers who are served by one of the entrants. Let $H(\hat{p})$ denote the gains from searching after a consumer has observed a tariff \hat{p} , i.e.,

$$H(\hat{p}, F) = \int_{\underline{p}}^{\hat{p}} (\hat{p} - p) f(p) dp = \int_{\underline{p}}^{\hat{p}} F(p) dp.$$

The reservation price $\rho(c; F)$ of a consumer is defined as the price at which the gains from searching one more time are equal to the cost of searching one more time, that is, $\rho(c; F)$ is the solution to

$$H(\rho, F) - c = 0. \tag{1}$$

A consumer will continue searching as long as observed prices are higher than its reservation price $\rho(c; F)$. If the consumer finds a firm that has set a lower tariff than her reservation value, she will stop searching and switch to that firm.

We now move on to the price setting behaviour of the firms. Assume for the moment the incumbent firm sets a price equal to the consumers' valuation v .¹⁶ The entrants draw their price each period from some price distribution $F(p)$, with support between \underline{p} and \bar{p} . Firms cannot or are not allowed to price discriminate between consumers.

¹⁶ In Appendix B we show that this is indeed optimal for the incumbent.

An entrant setting a price p serves only those of its own local consumers $(1 - \lambda)/N$ that have a reservation value higher than p , so the group of own locals accepting the current tariff is

$$\frac{1 - \lambda}{N} (1 - G(H(p))).$$

To attract consumers from the other entrants, the entrant should have a tariff lower than the reservation price ρ of a consumer that visits the entrant's supplier. Such a consumer will only start searching if her reservation price is lower than the tariff being set by her local supplier, which happens with probability $1 - F(\rho)$. Conditionally on searching, this consumer might visit the entrant's supplier at her first search, second search, third search, and so on, so the probability of selling to such a consumer is given by $\sum_{k=1}^{N-1} (1 - F(\rho(c, F)))^k$, where k is the number of visits. Summing up over all consumers that have a reservation value higher than p , but lower than the maximum tariff charged by the entrants, \bar{p} , and multiplying by the sampling probability $(1 - \lambda)/N$ gives

$$\frac{1 - \lambda}{N} \int_{H(p)}^{H(\bar{p})} \sum_{k=1}^{N-1} (1 - F(\rho(c, F)))^k g(c) dc.$$

Since consumers served by the incumbent will only switch if the expected gains from switching are bigger than zero, i.e., $v - E[p] - c > 0$, where $E[p]$ is the expected price at the entrants, the expected number of switchers from the incumbent is given by

$$\frac{\lambda}{N} \int_{H(p)}^{v - E[p]} \sum_{k=1}^N (1 - F(\rho(c, F)))^{k-1} g(c) dc.$$

Notice that compared to the previous expression, the sampling probability is now λ/N , because if a consumer local to the incumbent firm faces N entrants there might be consumers that search N times, and if a consumer searches only once the probability that the entrant serves this consumer, conditional on being sampled, is 1.

Finally, the entrant serves the consumers with a reservation price lower than p only if it has the lowest price among all entrants, which happens with a probability of $(1 - F(p))^{N-1}$, so this group of consumers is expected to be

$$G(H(p))(1 - F(p))^{N-1}.$$

This information can be summarized in the following profit equation, i.e., the profit of each entrant is given by

$$\begin{aligned} \pi_E(p) = (p - r) & \left[\underbrace{\frac{1 - \lambda}{N} (1 - G(H(p)))}_{\text{locals accepting current price}} + \underbrace{\frac{1 - \lambda}{N} \int_{H(p)}^{H(\bar{p})} \sum_{k=1}^{N-1} (1 - F(\rho(c, F)))^k g(c) dc}_{\text{switchers from other entrants}} \right. \\ & \left. + \underbrace{\frac{\lambda}{N} \int_{H(p)}^{v - E[p]} \sum_{k=1}^N (1 - F(\rho(c, F)))^{k-1} g(c) dc}_{\text{switchers from incumbent}} + \underbrace{G(H(p))(1 - F(p))^{N-1}}_{\text{consumers with lower } \rho \text{ if lowest price}} \right]. \end{aligned}$$

This equation can be simplified to

$$\pi_E(p) = (p - r) \left[\frac{1}{N} \int_{H(p)}^{\infty} (1 - F(\rho(c)))^{k-1} g(c) dc - \frac{\lambda}{N} [1 - G(H(v))] + G(H(p))(1 - F(p))^{N-1} \right]. \quad (2)$$

If there is no incumbent firm, i.e., if $\lambda = 0$, the profit equation simplifies to Stahl (1996) with unit demand. With an incumbent firm, some of the λ consumers are locked in at the incumbent firm and will never switch to one of the entrants. Each period the entrants make random draws from the price distribution $F(p)$, so even though consumers can be locked in at one of the entrants as well, in expected terms business from these locked in consumers is offset by the decreased probability of attracting consumers from the competing entrants. Therefore, the share of locked-in consumers at a given entrant will not have an impact on the expected profits of this entrant.¹⁷

The upper bound of the price distribution of the entrants $F(p)$ is denoted \bar{p} , where \bar{p} is the minimum of the monopoly price and the maximum a switcher is ever willing

¹⁷ In fact, the common assumption in the search literature that the first observation is for free leads to exactly the same random assignment of consumers across firms.

to pay, i.e., $\bar{p} = \min\{p^m, v\}$. The monopoly price p^m can be found by taking the first order condition of the profits at the upper bound

$$\pi_E(p^m) = (p^m - r) \left[\frac{1}{N} [1 - G(H(p^m))] - \frac{\lambda}{N} [1 - G(H(v))] \right]$$

with respect to p^m , which gives

$$p^m = \frac{[1 - G(H(p^m))] - \lambda[1 - G(H(v))]}{g(H(p^m))} + r.$$

An entrant will charge at most a price equal to \bar{p} . Setting such a price gives a profit

$$\pi_E(\bar{p} - r) \left[\frac{1}{N} [1 - G(H(\bar{p}))] - \frac{\lambda}{N} [1 - G(H(v))] \right]. \quad (3)$$

In a mixed strategy equilibrium, the entrants should be indifferent between charging any price in the support of $F(p)$, which implies $\pi_E(p) = \pi_E(\bar{p})$. This indifference condition implicitly defines the entrants' equilibrium price distribution, $F(p)$, i.e., for each p in the support of $F(p)$, $F(p)$ should solve

$$(p - r) \left[\int_{H(p)}^{\infty} \sum_{k=1}^N (1 - F(\rho(c)))^{k-1} g(c) dc - \lambda[1 - G(H(v))] + NG(H(p))(1 - F(p))^{N-1} \right] = (\bar{p} - r) \left[[1 - G(H(\bar{p}))] - \lambda[1 - G(H(v))] \right]. \quad (4)$$

Now consider optimal behaviour of the incumbent. Since the incumbent only sells to her own local consumers that have a reservation price higher than the gains from searching at a price v , the profits of the incumbent are

$$\pi_I = (v - r) [\lambda[1 - G(H(v))]]. \quad (5)$$

Since no consumers are willing to buy for a price that is higher than their valuation, it can never be optimal for the incumbent to set a price $p > v$. To see if, given the prices

set by the entrants, it is also not optimal to set a price $p < v$ we compare equation (5) to the profits that can be made at any other price $p < v$, i.e.,

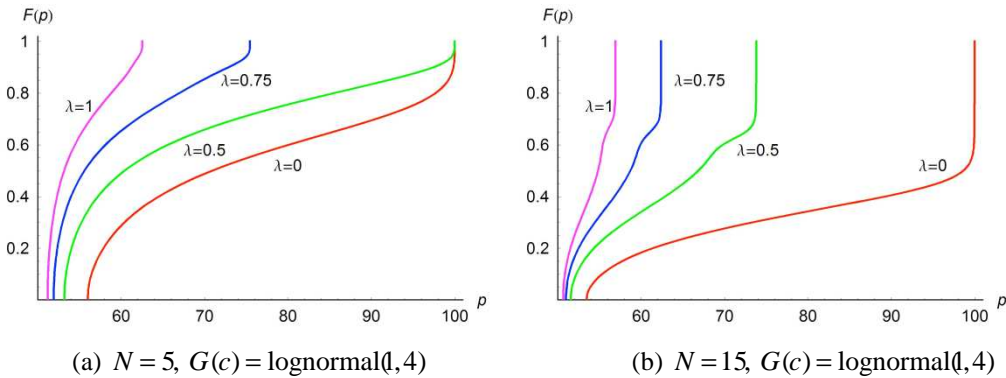
$$\pi_i(p) = (p-r) \left[\lambda(1-G(H(p))) + \frac{1-\lambda}{N+1} \int_{H(p)}^{\bar{p}} \sum_{k=1}^N (1-F(\rho(c,F)))^k g(c) dc + G(H(p))(1-F(p))^N \right]. \quad (6)$$

In Appendix B we show that it is indeed optimal for the incumbent to set its tariff equal to v . Equation (6) shows the tradeoff the incumbent faces: set a high price to maximize surplus from the relatively large number of loyal customers it has, but face the risk that consumers with low enough search cost will switch to one of the entrants, or set a more modest price to prevent consumers from leaving the firm at the cost of making less money per consumer.

Note that this is a one-period model. We assume, for analytic tractability, that the static game is unchanged over time so that consumers and firms view it as a repeated finite horizon game, yielding the same equilibrium framework each period. However, the outcome is different across years, since the underlying parameters change.

Illustration

As an illustration we have drawn the equilibrium price distribution for several values of λ , N , and parameters of the search cost distribution in Figure 3. In the examples there is one incumbent, valuation $v = 100$ and marginal cost $r = 50$. A λ equal to one means that all consumers are being served by the incumbent. As a benchmark case we also give the equilibrium distribution for λ equal to zero, which means that there is no incumbent.



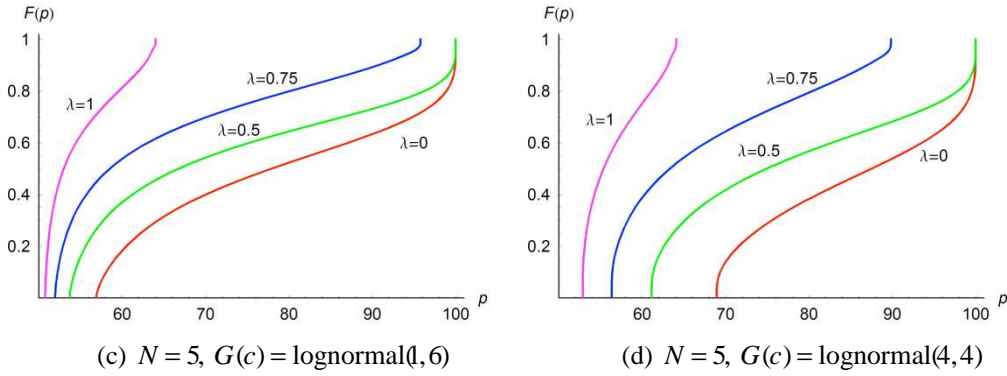


Figure 3: Price CDF of the entrants for several values of λ

All panels of Figure 3 show that a higher share of consumers at the incumbent λ results in lower prices being charged by the entrants. This is in line with intuition since a higher λ means fewer loyal consumers, and therefore more incentives to price aggressively in order to make consumers switch from the incumbent to one of the entrants. A comparison of Figures 3(a) and 3(b) illustrates that for a lognormal distribution with parameters $\mu = 1$ and $\sigma = 4$, an increase in the number of entrants leads to lower prices, although at the same time the entrants put more mass on the upper bound of the price distribution. The intuition for the latter is that it becomes increasingly difficult to attract the consumers with relatively low search costs as the number of firms increases. Firms start putting more mass on high prices to maximize surplus from consumers with relatively high search costs.¹⁸ At the same time competition between the entrants will result in overall lower prices. Only when λ is zero prices go up as a result of more firms, but that is because in the absence of an incumbent it can never be optimal to set the upper bound of the price distribution lower than v . Finally, Figures 3(c) and 3(d) show that both more dispersed search cost as well as higher mean search costs lead to higher prices being charged by the entrants.

4 Estimation strategy

In this section we show how to estimate the model using nonlinear least squares. Equation (4) provides the equilibrium condition for the entrants in the model and as such it provides a starting point for the estimation of the model. Unfortunately this equation only implicitly defines the equilibrium distribution of prices $F(p)$.

¹⁸ See Moraga-González and Janssen (2004) for a similar result for a non-sequential search model.

Moreover, the integral makes it difficult to compute the equilibrium, since this requires the evaluation of the price distribution at reservation prices corresponding to each search cost value over which we integrate. As we show below in more detail, to deal with this we make use of the notion that the model only identifies points on the search cost distribution that correspond to search cost values for which the reservations price is equal to an observed price. This implies that for the calculation of the integral we can simply sum up the integrand over exactly those points of support of the search cost distribution that are identified.

Assume we observe M prices, and that these prices are ordered by increasing price, i.e., $p_1 < p_2 < \dots < p_M$. Let $\tilde{F}(p)$ be the empirical price distribution, i.e., $F(p) = \frac{1}{M} \sum_{i=1}^M \mathbf{1}(p_i < p)$. The upper bound of $\tilde{F}(p)$ is given by $\bar{p} = p_M$. Inspection of equation (1) reveals that there is a one-to-one mapping from prices to search costs via the reservation prices. This can be used in the following way. Let c_i be the search cost of a consumer with reservation value $\rho(c_i, F)$ and let this reservation value be equal to price p_i . This means that the gains from search at price p_i is $H(p_i, F) = c_i$ and that c_i can be calculated directly from the data using the empirical price distribution, i.e.,

$$c_i = \frac{1}{M} \sum_{k=1}^i (p_i < p_k) \quad (7)$$

In addition, it follows from the definition of the empirical price distribution that $\tilde{F}(p) = \tilde{F}(p_j) \forall p \in (p_j, p_{j+1})$, so for each c between $H(p_j)$ and $H(p_{j+1})$ we have $\tilde{F}(p) = \tilde{F}(p_j)$. This means we can rewrite the integral in equation (4) as

$$\sum_{j=1}^M \int_{H(p_j; \tilde{F})}^{H(p_{j+1}; \tilde{F})} \sum_{k=0}^{N-1} [1 - \tilde{F}(p_j)]^k g(c) dc = \sum_{j=i}^M \sum_{k=0}^{N-1} [1 - \tilde{F}(p_j)]^k G(c) \Big|_{H(p_j)}^{H(p_{j+1})} = \sum_{j=i}^M \gamma_j \sum_{k=0}^{N-1} [1 - \tilde{F}(p_j)]^k$$

where $\gamma_j = G(H(p_{j+1})) - G(H(p_j))$, and $\gamma_M = G(v - E[p]) - G(H(p_M))$. Denoting $\gamma_{M+1} = 1 - G(v - E[p])$, we can rewrite the equilibrium condition as

$$(p_i - r) \left[\frac{1}{N} \sum_{j=i}^M \gamma_j \sum_{k=1}^N (1 - \tilde{F}(p_j))^{k-1} + \frac{1-\lambda}{N} \gamma_{M+1} + \left(1 - \sum_{j=i}^{M+1} \gamma_j \right) (1 - \tilde{F}(p_i))^{N-1} \right] = (\bar{p} - r) \left[\frac{1}{N} \gamma_M + \frac{1-\lambda}{N} \gamma_{M+1} \right] \quad (8)$$

According to the equilibrium defined in the previous section, equation (8) should hold for all prices in the support of $\tilde{F}(p)$, so given $\tilde{F}(p)$, N , r , λ , and \bar{p} , our goal is to find $\gamma_1, \gamma_2, \dots, \gamma_{M+1}$ such that this equation holds.

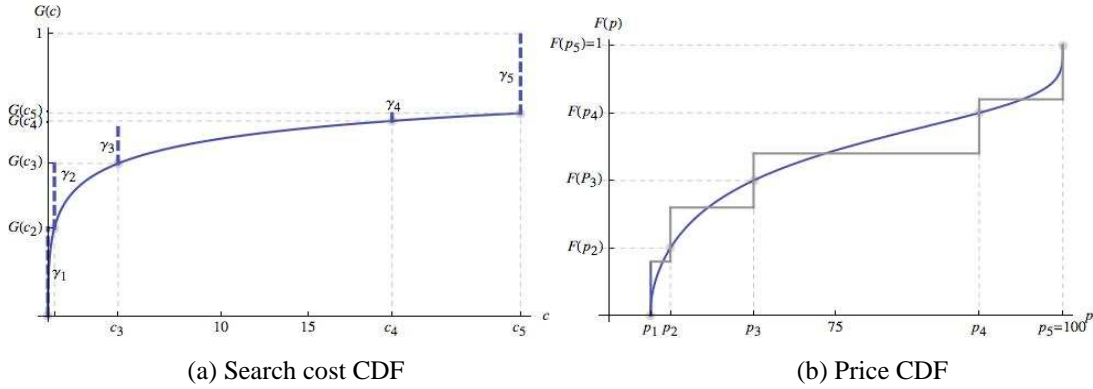


Figure 4: Identification

Figure 4 illustrates the identification for a market with search costs drawn from a log-normal distribution with parameters 1 and 4, the number of firms $N=5$, valuation $v=100$, marginal cost $r=50$, and the share of consumers being served by the incumbent $\lambda=0.25$. Figure 4(b) gives the equilibrium price distribution for these parameters, as well as the empirical price CDF for five draws from the equilibrium price distribution. Figure 4(a) gives the search cost CDF as well as the search cost values corresponding to the five draws from the equilibrium price distribution. These search cost values are determined by equation (7) and each of them represents a search cost such that one of the observed prices in Figure 4(b) is equal to the corresponding reservation price. The γ 's are calculated as the difference between two subsequent evaluations of the search cost CDF at the critical search cost values.¹⁹

¹⁹ Note that the number of points identified on the search cost distribution is equal to the number of price observations. This means the identification of this sequential search model does not suffer from the same identification problems as reported by Moraga-González, Sándor, and Wildenbeest (2010) for the non-sequential search model.

To see how the search cost distribution $G(c)$ can be identified using available data assume $G(c)$ is distributed according to some distribution $G(c; \theta)$ with parameter θ . Using the calculated c_i 's from equation (7) this allows us to calculate all γ_i 's. Next we solve equation (8) for p_i using the empirical price distribution $\tilde{F}(p_i)$, i.e.,

$$\tilde{p}_i = (\bar{p} - r) \frac{\frac{1}{N} \gamma_M + \frac{1-\lambda}{N} \gamma_{M+1}}{\frac{1}{N} \sum_{j=i}^M \gamma_j \sum_{k=1}^N (1 - \tilde{F}(p_j))^{k-1} + \frac{1-\lambda}{N} \gamma_{M+1} + \left(1 - \sum_{j=i}^{M+1} \gamma_j\right) (1 - \tilde{F}(p_i))^{N-1}} + r. \quad (9)$$

We then pick θ in such a way that the difference between the calculated \tilde{p} and the observed p is minimal. This can be done using nonlinear least squares.

The estimation can be done parametrically by starting from a parameterized distribution of search costs. To be able to do so, we need to have prices p , as well as the empirical price distribution $\tilde{F}(p)$, the number of entrants N , unit cost r , the share of consumers at the incumbent λ , and the maximum observed price \bar{p} . Prices are observed, as is the number of entrants, so we can use that to obtain the empirical price distribution. We observe the share of consumers at the incumbent for at least some periods. Furthermore, using data we have collected on input costs, we are able to derive a reasonable approximation of unit cost. Given that according to our model the incumbent will always find it optimal to set its price equal to the valuation of consumers, we can obtain an estimate of the valuation v by simply taking the price of the incumbent. The expected price can be estimated by the mean of the observed prices.

Once we have estimated the critical search cost values c_i from observed prices using equation (7), given the parameter values of the parameterized search cost distribution we can calculate $\gamma_i(\theta)$ and use equation (9) to calculate \tilde{p}_i . By minimizing the distances between the calculated prices \tilde{p}_i and corresponding observed price p_i we obtain a parametric estimate of the search cost distribution. This can be done using nonlinear least squares.

5 Estimation results

In this section we apply our estimation procedure to price collected from British electricity suppliers. We have made several strong assumptions in our theoretical model, which might need additional justification when applying our model to the electricity market. These include the assumptions of sequential search, symmetry across entrants, incumbent price at consumers' valuation, and parametric (lognormal) distribution. In discussing the results, we consider these issues.

We assume consumers search sequentially, which means consumers determine after each search whether to continue searching or not. As shown by Morgan and Manning (1985) sequential search is optimal when price information can be gathered relatively quickly because in that case the decision whether to search further at a given point in time can be conditioned on observed prices. If price quotations are observed with some delay, non-sequential search is typically optimal. It is unlikely that in the specific setting we are studying there is such a delay, but still non-sequential search might be a good approximation of search behavior if consumers use price comparison sites. Over the sampling period, price comparison websites have become more popular. However, the consumer surveys carried out for OFGEM over time show that whilst a significant proportion of consumers "use the internet to find suppliers prices," namely 28% by mid 2005, only 7% report "using comparison website" (Accent, 2006). Thus, for the bulk of consumers, it appears that sequential search is an appropriate model in this market, although that search process will have become more efficient over time.

Although we allow for the incumbent and the entrants to behave differently in our model, we assume entrants in a given region are symmetric. The assumption of symmetry does some violence to reality, but not much. The major outlier is British Gas, which is treated as an entrant in all electricity markets. Otherwise, all the other major players apart from Scottish Power are broadly equally successful in capturing unattached customers, according to rough figures obtained by subtracting their size in the incumbent market from their total market penetration.

We estimate the model using a panel of prices. For our empirical analysis we assume the static equilibrium of the theoretical model described in Section 3 is also the equilibrium of the repeated game with finite horizon. As long as the fundamentals of the model do not change, this allows us to estimate the model by taking several periods together. In line with institutional price-setting habits, we take a period as two months. Because unobserved market characteristics might be different across regions we estimate the model for each different region separately.

For each of the fourteen regions we have estimated the mean and the variance of the natural logarithm of the search costs using the lognormal distribution as our parametric specification. To allow for the mean and variance to vary over time we have included two trend parameters. The estimation results are presented in Table 2. A first observation is that the model explains the data very well: between 94 and 98 percent of the variation in prices can be explained by the model and all estimated parameters are highly significant. Notice that here is substantial variation between the regions in terms of mean log search costs and the variance of the log search costs, with the mean ranging from 5.3 to 8.7 and the variance from 6.8 to 15.6. All estimated trend parameters are smaller than zero, which means that average search costs within the population has been decreasing over time. Moreover, search costs became less dispersed over time.

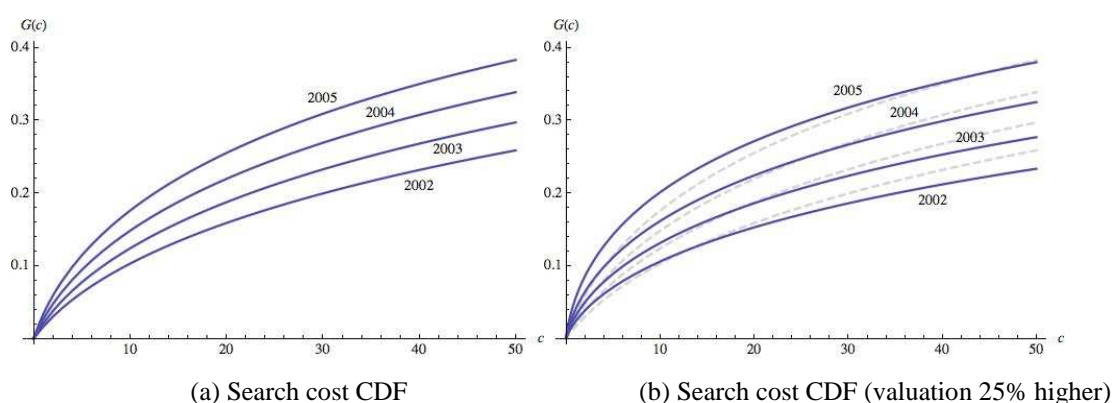


Figure 5: Estimated search costs Midlands region

In Figure 5(a) we plot the estimated search cost distributions for all years between 2002 and 2005 for the Midlands region. According to the model consumers will only

start searching if the expected gains, given by valuation v minus the expected price $E[p]$, are larger than search cost c . This puts an upper bound on the part of the search cost distribution we can identify: search cost cannot be identified beyond $\bar{c} = v - E[p]$. In Table 2 we have put the values of \bar{c} for all regions---for the Midlands this means we can only identify search costs up to £42.16. It is apparent from inspection of the curves that search costs have to be relatively high to rationalize observed prices. In fact, according to our estimates most consumers have search costs that are so high that it is not worthwhile to start searching at all. The relatively high margins we found by comparing observed prices with our measure of marginal cost already gave some indication that firms have a lot of market power, and this can only be rationalized by our model if consumers have relatively high search costs.

Our estimates suggest that there has been a significant leftward shift in the search cost distribution over time. Figure 5(a) shows this for the Midlands, but as indicated by the significant negative trend parameters in Table 2 other regions show the same pattern. As measured by the lower quartile, in the Midlands search costs have gone down from £46.88 in 2002 to £19.17 in 2005.²⁰ Also in other regions search costs have gone down substantially—Table 2 shows that for most regions lower quartile search costs in 2005 are around one-quarter of what they were in 2002. One explanation could be the increased use of price comparison sites over time. Since these sites make searching for the provider with the lowest price a lot easier, this might explain the substantial decrease in search costs. Our estimates also suggest that there are substantial differences in search costs across regions.

In order to estimate search costs we need an estimate of the valuation or willingness to pay of consumers. According to our model, the incumbent will find it optimal to set a price equal to the valuation of consumers, which allows us to use the incumbents price in each region as a (super-consistent) estimate of the valuation. One might be worried that this will underestimate the true willingness to pay of consumers, especially since the demand for electricity is highly inelastic. As a robustness check

²⁰ We report the lower quartile instead of the mean or median of the estimated log-normal search cost distributions, because most of the estimated parameters in Table 2 are such that the median and mean are larger than the upper bound of search costs that we can identify (see \bar{c} in Table 2). Any inference based on the mean or median would therefore rely on parametric extrapolation. Thus our reliance on a particular parametric framework for estimation (the lognormal distribution) is minimized.

we have also estimated the model using higher valuations. In Figure 5(b) we plot the estimated search cost distribution in case consumers' valuation is 25% higher, where the grey dashed curves correspond to the CDFs of Figure 5(a), and the solid curves are the CDFs estimated using the higher valuation. Although search costs appear to be slightly lower, the differences are small.

Note that even though we estimate a static model, our model, and hence our estimation results, do take into account that the share of consumers at the incumbent, denoted by λ , has been going down over time. As can be seen from Figure 3 it is important to take this changing value of λ into account: we can rationalize a given price distribution by a relatively high λ , which in isolation puts downward pressure on prices, in combination with relatively high mean search costs, which has the opposite impact. The fact that we observe λ therefore helps us in identifying search costs.

This also helps to explain why our estimates suggest search costs are decreasing over time, while retail margins are broadly increasing (Figure 1(b)). The apparent paradox can be resolved by observing that as the market share of the incumbent (λ) falls, mean price above marginal cost would tend to rise (see Figure 3) because the entrants, having gained a larger market share, compete less aggressively. That is, the tradeoffs represented in equation (2) are different, recalling that we re-evaluate λ each year.

	eastrn	emid	lndn	midlands	manweb	nrthrn	nweb	seebdr	shydro	scpwr	sthrn	swlec	sweb	yorks	mean
Mean log search costs	6.31 (0.04)	6.40 (0.22)	6.91 (0.19)	5.60 (0.10)	6.31 (0.15)	5.28 (0.10)	5.46 (0.05)	6.33 (0.13)	7.88 (0.36)	6.63 (0.22)	6.53 (0.05)	8.69 (0.30)	6.48 (0.16)	5.45 (0.11)	6.45 (0.16)
Variance log search costs	8.18 (0.05)	10.50 (0.12)	10.03 (0.09)	6.77 (0.08)	7.76 (0.09)	7.21 (0.10)	6.88 (0.05)	8.25 (0.08)	9.90 (0.14)	9.35 (0.12)	8.07 (0.05)	15.62 (0.10)	7.98 (0.10)	7.97 (0.10)	8.89 (0.09)
Trend mean search costs	-0.62 (0.02)	-0.60 (0.08)	-0.52 (0.07)	-0.31 (0.04)	-0.53 (0.05)	-0.28 (0.04)	-0.43 (0.02)	-0.41 (0.05)	-0.62 (0.13)	-0.39 (0.08)	-0.45 (0.03)	-0.99 (0.10)	-0.48 (0.06)	-0.28 (0.04)	-0.49 (0.06)
Trend variance search cost	-0.29 (0.04)	-0.39 (0.07)	-0.28 (0.06)	-0.11 (0.04)	-0.26 (0.05)	-0.09 (0.07)	-0.16 (0.03)	-0.24 (0.05)	-0.32 (0.09)	-0.22 (0.08)	-0.21 (0.04)	-0.48 (0.07)	-0.26 (0.06)	-0.22 (0.06)	-0.25 (0.06)
R^2	0.95	0.95	0.98	0.96	0.96	0.94	0.98	0.98	0.94	0.95	0.96	0.96	0.97	0.96	0.96
<i>Lower quartile</i>															
2002	80.14	67.54	118.12	46.88	84.07	32.17	40.18	80.62	315.35	95.85	101.38	411.83	96.80	34.78	114.69
2003	44.49	38.69	72.43	34.79	51.24	24.68	26.66	55.03	175.34	66.23	66.18	159.33	62.13	26.91	64.58
2004	24.72	22.18	44.43	25.82	31.25	18.94	17.70	37.58	97.55	45.77	43.21	61.68	39.90	20.82	37.97
2005	13.74	12.73	27.27	19.17	19.07	14.53	11.75	25.68	54.30	31.64	28.23	23.90	25.63	16.12	23.13
\bar{c}	54.44	52.89	40.16	42.16	45.94	42.42	54.95	39.50	42.39	38.15	41.15	46.42	40.75	42.53	44.59

Notes: Standard errors in parenthesis.

Table 2: Estimation results

6 Independent verification of search costs

One issue that faces modellers attempting to estimate the size of search costs largely through the distribution of observed prices is verification of whether the estimates are reasonable. In our present context, this issue is brought sharply into relief. Our estimates (Table 2) indicate what may appear to be very high lower quartile values for search cost. At the same time, it must be presumed that substantial search has been taking place in the market, because there is a considerable degree of switching between electricity providers. Can these two findings be reconciled? Another potentially questionable finding is the very significant estimated reduction in these search costs over a rather short period of time.

To examine these questions, we make use of a source of data, independent of those previously used in estimation, to confront our estimates. This is information on the degree of switching that has actually taken place. We have available two pieces of statistical information from OFGEM/ DECC, one being the total number of switches per year (unfortunately not broken down by region), the other the market share of the incumbent supplier by region.

Before turning to our calculations, two preliminary points need to be made. First, switching involves costs that are conceptually different from search costs. However, in a report based on a very extensive consumer survey for OFGEM in June 2005, Accent (2005) reported substantial findings on the market, as perceived by consumer respondents to their survey. One finding is that over 80% of electricity customers had, by then, been approached by sales people (most by a “cold-call” visit to the home). The tactics used by these sales people, influenced by their remuneration structure, involve them attempting to sign customers up to switch at the time of their visit. Almost 30% of consumers were first prompted to switch supplier by contact with a sales person, according to the survey respondents. Most remarkably, 30% of those who had switched supplier admitted *never* to have tried to compare suppliers’ prices! We draw two conclusions from this. One is that switching does *not* necessarily involve search, contrary to what might be thought. The second is that the alternative methods of switching supplier are, essentially, acting of one’s own volition

in searching and selecting a new supplier, and alternatively, responding to a sales person's invitation to switch.²¹

These two methods of switching supplier differ in that the first incurs search costs, whilst both incur switching costs. Responding positively to a sales person's invitation can be modelled as switching essentially at random to another (entrant) supplier. Thus the background level of switching is given by those who respond to the sales person's invitation.²² Searching then switching is likely to yield a much superior tariff. The simplest way to measure this difference is as the whole benefit of the difference between the mean entrant tariff and the lowest available tariff, in the relevant area. As an alternative, we also evaluate results based on the difference being between the mean entrant tariff and the second lowest available tariff, because we know that in searching, many people carry out only a limited search. It is these values that we alternately use as a measure of the benefit of search-then-switching.

The second point to make is that this calculation is still an approximation—we work with the most popular tariff type (Direct Debit, Electricity only) but we take its value at the mean consumption level and we ignore what happens with other tariff types (although they are likely to be highly correlated). The exercise is not exact.

Conceptually, our calculations amount to the following. Take the estimated CDFs as illustrated by those in Figure 5. Using the particular values suggested by the methods above for search costs (the X axis), generate the implied proportion of consumers switching, because they can make savings of at least this amount, by reading off the Y axis. Apply this to the number of customers (MPans) in the region. Confront this proportion with actual numbers switching.

These calculations are reported on in Table 3, culminating (through weighting by market sizes in the respective areas) in estimates of the actual number of predicted switchers within each region and across the time period through these means. Panel

²¹ A further implication is that some consumers have switched to a supplier that offers them a poor, or even a worse deal. Using two relatively small samples of consumers, Waddams and Wilson (2010) show that this is indeed true, and a remarkably high proportion of their consumers has chosen poorly.

²² The proportion of people who report having been approached by a salesperson by June 2005 varies rather little by region, between 74% (Northern) and 89% (East Midlands) (Accent, 2006). In addition there is very little difference indeed by demographic- they are almost a random sample of households.

A relates to the total switching pattern over time, panel B to the differential result across the fourteen regions accumulated over the whole time period studied.

Several positive features of these results are worth drawing out. First, taking the largest possible expected gain from search, namely the difference between mean entrant price and lowest entrant price, leads to estimates roughly in line with actual total switches. Specifically, the estimates for the years 2002 to 2005 yield values between 75% and 125% of the true value. The alternative estimates, with the lower calculated gain, lead to estimates between 55% and 90% of the actual switches (see Panel A). Taking into account that around 30% of switches have involved no search, this latter set of figures is broadly more in line with what may be expected,. What has not happened is that our estimates of the size of search costs compared to the available savings led to absurdly low predicted levels of switching compared with actuality.

A second positive feature is that our estimates, of significantly falling search costs over time, do not look outrageous either. The estimates of switcher proportions in 2005 as against earlier years are not too different, on either measure. 2004 if anything is the outlier in our estimates. This arises mainly because the benefits of search fall off in 2005, whereas for various reasons, largely relating to changes in (anticipated) input costs, they are unusually high in 2004.

A third positive feature is that we are able to predict the broad pattern across regions in terms of keenness to search (and switch), whereas simply using the proportion of people approached by a salesperson (to account for sales-initiated switches) would not do this.²³ In panel B we aggregate across all four years, but do so separately by area, and then compare the ranks of our predicted search and switching propensity with total shares lost across areas by the incumbent by December 2005 (taken from DECC figures). This exercise shows a high rank correlation, significant at 1% level. Those areas where we predict more customers will remain with the incumbent do indeed have more retained incumbent customers. We are able to predict the small loss in the Scottish areas, which is resultant from relatively small calculated savings, for

²³ Notice for example that the Northern region, where the lowest proportion of people report salesperson contact, shows the highest predicted, and third highest actual switching across the regions.

example. Of course, as we discussed in Section 2, other factors will be important in determining incumbent losses, for example supplier reputation, and here the clear outlier is SSE, with a consistently lower complaint level (and higher retained share-see Scottish Hydro, Southern and SWALEC areas, which are those where it is the incumbent) than the other five. However, it is remarkable how much can be explained based solely upon our search model.²⁴

In sum, we suggest that this calibration exercise on actual switching behaviour is significant independent verification of our estimates of search costs, as estimated from the price distributions.

Table 3: Predicted and actual shares of direct debit customers based on years 2002-2005

Panel A

Total switchers Millions	Year			
	2002	2003	2004	2005
Prediction 1	2.80	3.76	5.28	4.17
Prediction 2	2.18	2.34	3.78	3.27
Actual	3.70	4.20	4.21	4.31
percent pred 1	75.60	89.60	125.54	96.85
percent pred 2	58.79	55.76	89.78	75.92

Panel B

Region (group)	Predicted		Incumbent % share (DECC) Actual	
	% stay	Rank	Dec-05	Rank
Eastern (E.On)	38	10	40	11
Emid (E.On)	28	13	41	9
Lond (EdF)	45	6	47	5
Mid (npower)	39	9	43	8
Manweb (Scottish Power)	46	5	47	5
Northern (npower)	17	14	39	12
Norweb (E.On)	33	11	28	14
Seeboard (EdF)	49	3	46	7
Schydro (SSE)	68	1	75	1
Scpower (Scottish Power)	50	2	59	3
Southern (SSE)	48	4	59	3

²⁴ Clearly, a multivariate regression approach is infeasible here for degrees of freedom reasons.

Swalec (SSE)	44	7	65	2
Sweb (EdF)	42	8	41	9
Yorks (npower)	29	12	35	13
average	41		46	

Rank correlation 0.835

Note: Prediction 1 refers to the search yielding the largest possible gain, prediction 2 relates to the lower gain.

7 Conclusions

In the late nineties the British electricity market was one of the first to completely open up for competition, leading to a substantial inflow of new suppliers. Moreover, a greater proportion of consumers has switched than perhaps anywhere else in the world. As such it provides an interesting case of how competition has affected pricing strategies, both those of the incumbent supplier of electricity and those of entrants. On the consumer side of the market the inflow of new suppliers has led to changes as well, because since the opening of the market consumers have the opportunity to look for better deals and eventually switch.

Our sequential search model distinguishes between an incumbent and a group of entrants. Using the structure of the model we have estimated a search cost distribution. Estimates indicate that in order for the model to rationalize observed price patterns, search costs have to be relatively high, implying that most consumers do not search or switch. Our estimates do indicate that search costs have decreased over time. But at the same time, because entrants have gained a greater foothold in each market, their incentives to price keenly have been muted.

Taking the implications of the estimates of search costs for the amount of switching that should have been observed, together with the actual amount of switching behaviour that occurred, shows a substantial correlation between the two, which leads us to greater confidence that our search cost estimates, which might be thought to be high, are indeed borne out in consumer switching behaviour.

More recently, the tariff structure has become more complex, rendering consumer decisions more difficult, at the same time as some consumers' search costs have fallen, other things equal, as a result of using Internet search engines. In consequence, substantial differences across tariffs remain, indeed may have become greater still.

The general lesson we draw from this analysis is that, however apparently homogenous the product, however straightforward search and switching appear to be, and even if substantial numbers switch, because consumer search for products such as an electricity supplier is perceived as costly (or, perhaps, simply boring), the results can be poor. Firms, recognizing the limited nature of search, can exploit this in their charging structures. Of course, time is a precious resource, and we would not expect search opportunities to be captured exhaustively. However, for a product on which people spend a substantial proportion of their incomes the result is surprising, and shows the importance of consumer, as well as firm, behaviour, in determining competitive outcomes.

Appendices

Appendix A: the Data

The analysis of search costs in the British electricity market is based on data from a variety of public and commercial sources, in one case proprietary.

1. The **domestic retail annual bills** for electricity consumers and the names of different suppliers serving each region (incumbents and entrants) were collected from the consumer watchdog's 'energywatch' website for the time period between April 2001, when the final price controls on direct debit prices were removed, and December 2005 inclusive. The data have been collected on a bi-monthly basis for consumers using the direct debit payment method and using a medium amount of electricity (3300 kWh) in a year, for all suppliers in each of the fourteen regions across which prices differ. The fourteen electricity regions in Great Britain are: Eastern, East Midlands, London, Midlands, Merseyside and North Wales, Northern, North West, South East, North Scotland, South Scotland, Southern, South Wales, South West and Yorkshire. The final bill figures include VAT at a 5% rate. Our data excludes Internet only tariffs since these represented only a small proportion of the subscribed tariffs even by the end of our period.
2. The cost information used to calculate suppliers' margins comes from a variety of sources to account for the main cost components affecting final consumers' bills, given the lack of access to commercial information about suppliers' costs. To capture the effect of **wholesale electricity prices** on domestic bills we used set of proprietary wholesale data kindly supplied by Platts, one of the three major energy data information companies. Given that all energy suppliers have a portfolio of forward contracts for delivery at various time horizons, we used the year-ahead price as our measure of electricity wholesale costs. These prices are determined at a national level so this variable has no geographical differentiation but only time variation at bimonthly frequency, in line with the frequency of observations for domestic bills.

Transmission costs vary by point of entry (generation plant) and exit (distribution point for delivery to final consumers) so that in order to calculate approximate values of transmission costs at a regional level a series of assumption about flows of power across different parts of the country and the percentage of electricity subject to the transmission charge were required. Average values of nation-wide transmission costs were obtained from Cornwall Energy Associates. To these aggregate values we applied weights for the different electricity regions. These weights have been calculated on the basis of information provided by company managing electricity transmission in the UK, National Grid (based on power flows in 2006).

Distribution costs, i.e. the costs of transmitting low voltage electricity at the regional level, on the other hand, are subject to price regulation for each region set by the energy regulator OFGEM. We have obtained data on distribution costs (both fixed and unit rate) for each of the 14 regions, from the UK Department of Energy (currently designated as DECC and formerly from BERR and DTI) over the entire period 2002-2005.

Estimates relating to **other costs** faced by energy suppliers (balancing costs, network losses, metering costs, supplier costs to serve and environment related levies such as renewable obligations, energy efficiency commitment)²⁵ are based on the Cornwall Energy Associates report on energy costs to consumers, which covers the years 2003 and 2006.

3. On the demand side, our analysis relied on information on the incumbents' **market share** of domestic direct debit consumers for each region over the whole time period considered, which are published on a quarterly basis by the UK Department of energy (DECC). The **number of electricity users** (consumers) for each region was based on the number of metering points (meter point administration numbers or MPANs) published by Ofgem as meters in existence at September 2005. The **total number of switchers** at the national level for each of the years in our analysis was obtained from Ofgem's (approximately) annual Domestic Retail Market Reports from 2003 to 2006.

Data on **advertising expenditure** by the 'big 6' energy suppliers were obtained from the NMR digest for the years between 1999 and 2005.

²⁵ Charges for carbon emissions were introduced in April 2005.

Data on **proportion of consumers using the Internet** was obtained from OFGEM reports (2008, 2010).

Appendix B: Derivation of the optimal upper bound for the incumbent

Start from a situation where the incumbent has as many loyal customers as each of the entrants, i.e., $\lambda = (1 - \lambda)/N$. In this case it is always optimal for the entrant to set $\bar{p} = v$. Incumbent's profit at v is

$$\pi_I(v) = (v - r)[\lambda[1 - G(v - E[p])]] = (v - r) \left[\frac{1 - \lambda}{N} [1 - G(v - E[p])] \right] > \pi_E(\bar{p}).$$

We know that the entrants are indifferent between charging any price in the support of $F(p)$. When $\lambda = (1 - \lambda)/N$ profits of the incumbent for $p < \bar{p}$ are

$$\pi_I(p) = (p - r) \left[\frac{1 - \lambda}{\lambda} (1 - G(H(p))) + \frac{1 - \lambda}{N + 1} \int_{H(p)}^{H(\bar{p})} \sum_{k=1}^N (1 - F(\rho(c, F)))^k g(c) dc + G(H(p))(1 - F(p))^N \right].$$

Comparing this with equation (2) shows that $\pi_I(p) < \pi_E(p) = \pi_E(\bar{p}) < \pi_I(v)$, so it can never be optimal for the incumbent to set a price lower than \bar{p} when $\lambda = (1 - \lambda)/N$. Given this result, if the incumbent's profit at v increases faster in λ than the profits evaluated at any other price p in the support of $F(p)$, we can be sure that the incumbent will find it profitable to set a price equal to v . Therefore, we have to evaluate whether $d\pi_I(v)/d\lambda > d\pi_I(p)/d\lambda$, i.e.,

$$(v - r)[1 - G(H(\bar{p}))] > (p - r) \left[1 - G(H(p)) - \frac{1}{N + 1} \int_{H(p)}^{H(\bar{p})} \sum_{k=1}^N (1 - F(\rho(c, F)))^k g(c) dc \right].$$

Note that both sides of the inequality do not depend on λ . Multiplying both sides by λ gives

$$(v - r)\lambda[1 - G(H(\bar{p}))] > (p - r) \left[\lambda[1 - G(H(p))] - \frac{\lambda}{N + 1} \int_{H(p)}^{H(\bar{p})} \sum_{k=1}^N (1 - F(\rho(c, F)))^k g(c) dc \right],$$

which is equivalent to

$$\pi_1(v) > \pi_1(p) - \int_{H(p)}^{H(\bar{p})} \sum_{k=1}^N (1 - F(\rho(c, F)))^k g(c) dc - G(H(p))(1 - F(p))^N.$$

Since we know that $\pi_1(v) > \pi_1(p)$ at $\lambda = (1 - \lambda)/N$ it has to be that $d\pi_1(v)/d\lambda > d\pi_1(p)/d\lambda$ at $\lambda = (1 - \lambda)/N$. Moreover, since both derivatives are constant in λ , profits at the upper bound are increasing faster in λ than at any other price p , so for $\lambda \geq (1 - \lambda)/N$ it is always optimal for the incumbent to set a price equal to v .

Appendix C: Do the firms play mixed strategies?

A fundamental feature of our model, as with most search models, is that the firms play mixed strategies with respect to pricing. In our case, we are able to assess whether this appears to be true within the data. The means by which we operationalise this is to convert the bi-monthly pricing data for the “big 6” into ranks, going from 1 to 6.²⁶ One aspect of this where theory is unclear is *how often* prices are reassessed. Given the institutional regularities in this industry, we suggest that every two months is an appropriate period. (Although consumers would not be likely to reassess their situation this frequently, so it is probably on the short side.) This gives us 24 observations, across 14 GB regions. The question we ask, for each firm, is whether the strategy it plays is concordant across the time period and across regions. If it is not significantly concordant, then this is indicative that a mixed strategy is being pursued. For the five firms that are incumbents, we do this with and without their incumbent areas (since it is fairly clear that, for the most part, they charged

²⁶ Arguably, the rank of the firm in the set of prices is a convenient way of assessing its strategy- if a consumer is searching amongst a subset of prices, including for some exhaustive search, then the consumer will choose the lowest, all other things equal, whether it is significantly lower or slightly lower than other offers. In the earlier years, the firms had not necessarily merged or consolidated pricing across divisions of the firm. Here, our procedure is to take the lead partner value, except for the local area, in these cases. To illustrate, a clear cut case is Scottish Power, where there was a “Manweb specific” offering in that area, whilst Scottish Power did not offer itself to serve the region. However, in other cases, for example EdF, there were sometimes two separate offerings in the area in some early periods.

prices at the top of the distribution for this group. Our method of testing concordance, or its absence, is using Kendall's coefficient of concordance (W).²⁷

An illustrative tabulation of ranks is given below, using the example of British Gas's electricity prices; the other five are similar. Of course, British Gas is not an electricity incumbent in any area. We can calculate W using either the columns (Is the strategy concordant, or mixed, across the 14 regions?) or the rows (Is the strategy concordant or mixed across bimonthly time periods?). Then, given the degrees of freedom in these cases, we compare the relevant calculated value (a transformation of W) with the χ^2 distribution with $n - 1$ degrees of freedom, n being the number of cases. As can be seen, there is no evidence to reject the null hypothesis of mixed strategy in either dimension. This is obvious by inspection in the case of different regions, but less obvious in the time dimension. One thing this test does not make use of is the sequencing of time, and it is clear that there is some time dependence in at least some regions.²⁸ The other regions give the same answer on the concordance test, once incumbency has been accounted for.

Our general conclusion from this analysis is that a model in which firms play mixed strategies does not do empirical violence to the actual experience. Whether this is conscious behaviour on the part of the firms involved or not is unclear and we do not claim that the firms deliberately pursue mixed strategy pricing, simply that the outcome is as if they in fact do so.

²⁷ The literature on testing for mixed strategies in play is not very helpful in suggesting approaches. The best-known strand of it relates to sports activities (e.g. penalty kicks in serves, behaviour in tennis serves), which is not very relevant here. A more promising approach in the present context is adopted by Lach (2002)—this is the approach used in Giulietti et al (2010a). What we report on above is essentially an extension of this idea, with the same results as in that paper.

²⁸ In a market like electricity, where there is continuous purchasing, we would expect to see only gradual movements to being a better or worse buy, rather than sudden changes.

COMPANY:	BG Dates																							
REGION	f02	a02	j02	a02	o02	d02	f03	a03	j03	a03	o03	d03	f04	a04	j04	a04	o04	d04	f05	a05	j05	a05	o05	d05
EASTRN	5	5	4	4	4	4	4	5	5	5	5	5	5	4	4	4	6	5	4	4	4	4	6	6
EMID(Powergen)	4	4	4	4	3	3	2	3	2	2	2	2	5	3	3	3	6	5	4	4	4	4	4	4
LNDN	3	3	3	3	3	3	3	5	5	5	5	4	5	5	3	2	4	2	2	2	2	2	5	4
MEB	4	4	3	3	2	2	2	4	3	3	2	2	4	3	3	2	5	4	3	3	3	3	6	5
MANWEB	3	3	3	2	2	2	2	1	1	1	1	4	4	2	2	1	5	4	3	3	3	3	5	4
NRTHRN	4	4	4	3	3	3	3	4	4	4	4	3	5	5	5	3	5	4	4	4	4	4	6	5
NWEB	5	5	5	5	4	4	4	5	4	4	4	3	5	3	4	4	6	5	5	5	5	5	6	6
SEEBRD	5	5	5	5	5	5	5	5	5	5	5	4	5	5	5	4	4	4	4	4	4	4	5	4
SHYDRO	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	5	4	2	2	2	2	2	4
SCPWR	2	2	2	1	1	1	1	1	1	1	1	1	2	1	1	1	4	3	2	1	1	1	5	4
STHRN	4	4	3	3	3	3	3	5	3	3	3	2	4	2	2	2	6	5	4	3	3	3	6	5
SWLEC	5	5	4	4	4	4	4	5	3	3	3	2	6	5	5	4	6	6	5	4	4	4	6	6
SWEB	4	4	4	4	4	4	4	5	4	4	4	3	4	3	2	2	4	3	3	3	3	3	3	4
YORKS	4	4	4	3	3	3	3	5	5	5	5	4	4	4	4	4	5	4	4	4	4	4	6	5

Region	k=24	n=14	critical
chisq	4.7		20
Time	k=14	n=24	
chisq	11		32

Appendix D: Advertising as a source of price information

In addition to consumers being able to glean information for themselves, from information providers such as energywatch and online switching services, and from sales people, information (or perhaps, invitations to obtain information) is also available via advertisements. To investigate this aspect, we obtained information from NMR Digest for our range of years from liberalization. This is a quarterly publication that checks all significant (non-digital) media activity. Its methodology appears to be to scan a very wide range of print, poster, TV and radio media outlets throughout the UK, recording the presence of advertisements and applying published rate-card information to get a figure for spend²⁹ (subject to a lower cut-off). Since categorization of the advertisement series is sometimes ambiguous (for example, as to whether it applies to electricity or gas, or both), we aggregated all relevant cases to get an annual figure per firm, with the results shown in the table below.

²⁹ This will tend to be an overestimate, due to discounts being applied.

Media spend, energy sales total, £'000

Year/Co	1999	2000	2001	2002	2003	2004	2005	Average per year
EdF	2986	2390	2069	4827	1119	3537	1891	2688
NPOWER	4010	10944	5809	2370	583	5015	8184	5274
EON	2973	4983	10168	8580	5491	4518	8372	6441
SSE	717	162	0	236	88	92	88	198
Scottish P	4549	837	666	600	733	2419	2219	1718
BG	17868	5682	13345	10306	810	4466	5012	8213
Average	5517	4166	5343	4487	1471	3341	4294	

(Source, NMR Digest, various issues, after aggregation. Data are in £'000s)

Outside the “big 6”, there is remarkably little expenditure, right from the outset. Even within this group, SSE seems to promote very little indeed by these means. Indeed, in total the expenditure can be viewed as modest, less than £1 per customer per year (so maybe ½% of turnover), on average. On a more qualitative note, given the pricing structures employed by all of these players, it is not feasible to produce a simple comparative information advertisement. Thus the media tend to be selective at best in quoting pricing reasons for a switch. All also claim excellent service, in some cases a potentially dubious claim! We conclude from this investigation that this form of promotional activity is not particularly important in the electricity industry relative to those we investigate.

References

Accent: *Customer Experience Survey Report (Key Findings)*, Accent, London, May 2005. Appendix 6 to Ofgem (2006).

Michael R. Baye, Dan Kovenock, and Casper G. de Vries: "It Takes Two to Tango: Equilibria in a Model of Sales," *Games and Economic Behavior* 4, 493-510, 1992.

Michael R. Baye and John Morgan: "Information Gatekeepers on the Internet and the Competitiveness of Homogeneous Product Markets," *American Economic Review* 91, 454-74, 2001.

Kenneth Burdett and Kenneth L. Judd: "Equilibrium Price Dispersion," *Econometrica* 51, 955-69, 1983.

Cornwall Energy Associates: "Gas and Electricity Costs to Consumers: a Report for the National Right to Fuel Campaign," 2008.

Christophe Defeuilley: "Retail Competition in Electricity Markets", *Energy Policy* 37, 377-386, 2009.

Monica Giuliatti, Jesús Otero, and Michael Waterson: "Pricing Behaviour under Competition in the UK Electricity Supply Industry," *Oxford Economic Papers* 62, 478-503, 2010a.

Monica Giuliatti, Luigi Grossi and Michael Waterson: "Price Transmission in the UK Electricity Market: Was NETA Beneficial?" *Energy Economics*, 2010b, doi:[10.1016/j.eneco.2010.01.008](https://doi.org/10.1016/j.eneco.2010.01.008).

Han Hong and Matthew Shum: "Using Price Distributions to Estimate Search Costs," *Rand Journal of Economics* 37, 257-75, 2006.

Elisabeth Honka: "Quantifying Search and Switching Costs in the U.S. Auto Insurance Industry," Mimeo, Booth School of Business, University of Chicago, 2010.

Ali Hortaçsu and Chad Syverson: “Product Differentiation, Search Costs, and Competition in the Mutual Fund Industry: a Case Study of S&P 500 Index Funds,” *Quarterly Journal of Economics* 119, 403-56, 2004.

Ali Hortaçsu, Seyed H. Madanizadeh and Steven L. Puller: “Power to Choose: An Analysis of Consumer Behaviour in the Texas Retail Electricity Market,” Mimeo, 2010.

Maarten C. W. Janssen and José Luis Moraga-González: “Strategic Pricing, Consumer Search and the Number of Firms,” *Review of Economic Studies* 71, 1089-118, 2004.

Cenk Kocas and Tunga Kiyak: “Theory and Evidence on Pricing by Asymmetric Oligopolies,” *International Journal of Industrial Organization* 24, 83-105, 2006.

José Luis Moraga-González, Zsolt Sándor and Matthijs R. Wildenbeest: “Semi-Nonparametric Estimation of Consumer Search Costs,” Mimeo, 2010.

José Luis Moraga- González and Matthijs R. Wildenbeest: “Maximum Likelihood Estimation of Search Costs,” *European Economic Review* 52, 820-48, 2008.

Peter Morgan and Richard Manning: “Optimal Search,” *Econometrica* 53, 923-55, 1985.

Ofgem: *Domestic Gas and Electricity Supply Competition*, Ofgem, London, 2003.

Ofgem: *Domestic Retail Market Report- June 2005*, Ofgem, London, 2006.

Ofgem: *Energy Supply Probe: Initial Findings Report*, No. 140/08, October 2008.

Ofgem: *Update on Probe Monitoring: Tariff Differentials and Consumer Switching*, No. 79/10, July 2010.

Dale O. Stahl: “Oligopolistic Pricing with Sequential Consumer Search,” *American Economic Review* 79, 700-12, 1989.

Dale O. Stahl: “Oligopolistic Pricing with Heterogeneous Consumer Search,” *International Journal of Industrial Organization* 13, 243-68, 1996.

Nicholas Stern: “Imperfections in the Economics of Public Policy, Imperfections in Markets, and Climate Change”, *Journal of the European Economic Association* 8, 253-288, 2010.

Hal Varian: “A Model of Sales,” *American Economic Review* 70, 651-59, 1980.

Catherine Waddams and Chris Wilson: “Do Consumers Switch to the Best Supplier?”,
Oxford Economic Papers 62, 647-68, 2010.

Chris Wilson: “Market Frictions: A Unified Model of Search and Switching Costs,”
ESRC Centre for Competition Policy Working Paper # 09-3, 2009.