

MAXIMUM LIKELIHOOD ESTIMATION OF SEARCH COSTS*

José Luis Moraga-González[†]
Matthijs R. Wildenbeest[‡]

Working paper: January 2006
Revised: June 2007

Abstract

In a recent paper Hong and Shum (2006) present a structural method to estimate search cost distributions. We extend their approach to the case of oligopoly and present a new maximum likelihood method to estimate search costs. We apply our method to a data set of online prices for different computer memory chips. The estimates suggest that the consumer population can be roughly split into two groups which either have quite high or quite low search costs. Search frictions confer a significant amount of market power to the firms: despite more than 20 firms operating in each of the markets, we estimate price-cost margins to be around 25%. The paper also illustrates how the structural method can be employed to simulate the effects of the introduction of a sales tax.

Keywords: consumer search, oligopoly, price dispersion, structural estimation, maximum likelihood

JEL Classification: C14, D43, D83, L13

*We thank the Editor, Zvi Eckstein, two anonymous referees, Zsolt Sándor and Chris Wilson for their numerous comments and suggestions. Pim Heijnen, Mariëlle Non, Aico van Vuuren and the seminar participants at the Copenhagen Business School, Institute of Economic Analysis (Barcelona), University of Bologna, Universidad Carlos III de Madrid, University of Essex, and Tinbergen Institute Amsterdam also provided us with useful remarks. The paper has benefited from presentations at the EARIE 2005 Meetings (Porto), World Congress of the Econometric Society 2005 (London), EEA 2005 Meetings (Amsterdam), ESRC Centre for Competition Policy PhD workshop (UEA, 2005), and the ENCORE Workshop on Consumers and Competition (Rotterdam, 2006). The second author gratefully acknowledges financial support from the Vereniging Trustfonds Erasmus Universiteit Rotterdam and the Netherlands Organization for Scientific Research (NWO).

[†]University of Groningen, E-mail: j.l.moraga.gonzalez@rug.nl

[‡]Corresponding author. Erasmus University Rotterdam, Department of Economics, P.O. Box 1738, 3000 DR Rotterdam, The Netherlands. Tel. +31 10 408 1479. Fax. +31 10 408 9161. E-mail: wildenbeest@few.eur.nl.

1 Introduction

There is substantial evidence that the prices of seemingly homogeneous consumer goods are quite dispersed (see e.g. Stigler, 1961; Dahlby and West, 1986; Pratt et al., 1979; Sorensen 2000; Brown and Goolsbee, 2002; Lach, 2002; Baye et al., 2004). During the last twenty-five years, economists have dedicated a significant theoretical effort to explain this empirical regularity as an equilibrium phenomenon. One of the findings is that price dispersion can be sustained in equilibrium when some consumers observe several prices while other consumers observe only one price. Such unequal distribution of price information across consumers often arises in the market as a result of costly search (see e.g. Varian, 1980; Burdett and Judd, 1983; Rob, 1985; Stahl, 1989).

In spite of the considerable theoretical effort, somewhat surprisingly, very little empirical work has focused on identifying and measuring search costs in real-world markets. From an applied point of view, this is certainly an omission because predictions and policy recommendations from the various theoretical models are often sensitive to the magnitude of search costs.¹

In a recent paper, Hong and Shum (2006) present structural methods to retrieve information on search costs in markets for homogeneous goods. They show that firm and consumer equilibrium behavior imposes enough structure on the data to allow for the estimation of search costs using only observed prices. Hortaçsu and Syverson (2004) show that when price and quantity data are available, these methods can be extended to richer settings where price variation is not only caused by search frictions but also by quality differences across products.²

The non-sequential search model studied by Hong and Shum (2006) generalizes Burdett and Judd's (1983) seminal paper by introducing search cost heterogeneity. They consider a market operated by a continuum of firms which compete by setting prices. Consumers,

¹See e.g. Janssen and Moraga-González (2004) for the influence of the magnitude of search costs on equilibrium search intensity and market competitiveness.

²There is a well-established literature in labor economics that structurally estimates models of job search. Key contributions in this literature are Eckstein and Wolpin (1990) and Van den Berg and Ridder (1998). This literature, recently surveyed in Eckstein and Van den Berg (2007), has studied, among other issues, wage dispersion, duration of unemployment, minimum wage policies, returns to schooling and earnings inequality. The empirical work using models where search efforts are endogenous is however relatively small. For a first attempt to estimate search cost distributions in labor markets see Gautier et al. (2007).

who have heterogeneous search costs, engage in search to discover prices. Once a consumer has observed the desired number of prices, he/she buys from the cheapest firm in his/her sample. In equilibrium, only a fraction of consumers compare the prices of various firms which leads to price dispersion. Hong and Shum formulate the estimation of the unknown search cost distribution as a two-step procedure. They first estimate the parameters of the equilibrium price distribution by *maximum empirical likelihood* (MEL). To do this, they derive a (potentially infinitely large) number of moment conditions from the equations that describe the equilibrium. The estimates of the parameters of the cumulative distribution function (cdf) of prices give the height of the search cost distribution evaluated at a series of cutoff points. In the second step, these cutoff points are estimated directly from the empirical cdf of prices. While innovative, this method is limited by the ability to solve a computationally demanding high-dimensional optimization problem. Indeed, in practice, only a few parameters of the price distribution can be estimated which can result in the introduction of biases into the estimates.³

In this paper we present an alternative strategy to estimate an oligopoly version of the non-sequential search model of Burdett and Judd (1983) by using *maximum likelihood* (ML). We first estimate the parameters of the price distribution by ML. To do this, we compute the likelihood of a price as a function of the distribution of prices and exploit the equilibrium constancy-of-profits condition to numerically calculate the value of the price cdf. Once we obtain a ML estimate of the price distribution, we introduce a method to calculate the cutoff points of the search cost distribution as a function of the ML estimate of the price cdf. In this way, by the invariance property of ML estimation, the estimates of the cutoff points of the search cost distribution are also ML. The procedure is relatively easy to implement and has the advantage that the asymptotic theory for computing standard errors and for conducting hypothesis tests remains standard.

The model we study is an oligopolistic version of Burdett and Judd's (1983) non-sequential search model. Vis-à-vis the competitive case studied by Hong and Shum (2006), the oligopoly model has the advantage that it captures variation in prices due to variation in the number

³For example, in the empirical examples presented in Hong and Shum (2006) low search cost consumers are ignored because the number of searches a consumer can make is (artificially, by the econometrician) limited.

of competitors in addition to variation in prices due to search frictions; this makes our model useful for the study of competition policy issues. Another advantage is that, if the econometrician knows there are N firms operating in a market, then he/she knows consumers will search up to a maximum of N prices. As a result, no matter the number of prices the econometrician actually observes, he/she can estimate the relevant number of parameters of the price distribution. In this way, we are able to learn about the distribution of search costs at all relevant quantiles.⁴

To estimate the parameters of the price distribution, we need to observe the prices of the firms over some period of market interaction. We perform Monte Carlo simulations and show that, with relatively few data, the estimate of the price distribution is very accurate while the estimate of the search cost distribution is biased towards high search costs. In addition, the simulations reveal that ignoring low search cost consumers, as Hong and Shum (2006) do, leads to significant biases in the estimates: search costs are substantially overestimated and price-cost margins turn out to be much larger than they really are. These biases result in a poor fit of the model to the data and goodness-of-fit tests reject the null hypothesis that the empirical and the estimated distribution of prices are equal.

If the fraction of low search cost consumers were negligible in real-world markets, this would not be a problem. However, it turns out that the fraction of consumers searching intensively in real-world markets is sizable. We apply our method to a data set of prices for four personal computer memory chips. For all the products, we observe significant price dispersion as measured by the coefficient of variation. On average, relative to buying from one of the firms at random, the gains from being fully informed in these markets are sizable, ranging from 21.56 to 32.89 US dollars. Our estimates of the parameters of the price distribution yield an interesting finding: consumers either search very intensively in the market (between 4% and 13% of the consumers) or search very little, namely for at most three prices. Very few consumers search for an intermediate number of prices.⁵ The search cost distribution consistent with these estimates implies that most consumers have quite

⁴Brown and Goolsbee (2002) argue that prices of life insurance policies did not fall with rising Internet usage (which probably meant an upward shift of the search cost distribution) but with the emergence of price comparison sites (which most likely meant a more radical change of the shape of the distribution). Picking up such an effect requires information on how the Internet has affected search costs for all quantiles.

⁵In a study of the consumer click-through behavior online, Johnson et al. (2004) also point out that many consumers search quite little.

high search costs and a few consumers have quite low search costs. Our estimates suggest that the search cost of consumers who search thoroughly in the market is at most 17 US dollar cents.⁶

Consumers' search behavior confers substantial market power to the firms. In spite of the fact that in each of the markets studied we observe more than 20 retailers, we estimate that the average price-cost margin ranges between 23% and 28%. This suggests that demand side characteristics like search frictions might even be more important than market structure to assess market competitiveness (Waterson, 2003).

The validity of the theoretical model is tested, first, by checking whether the data support each of the assumptions of the model and, second, by conducting Kolmogorov-Smirnov tests of the goodness of fit. According to the test results, we cannot reject the null hypothesis that the observed prices are generated by the model.

The paper also illustrates how the structural method can be employed to simulate the effects of policy interventions. In particular, we study how the introduction of a sales tax would affect the equilibrium outcome in the market for one of the memory chips in our data set. We find that sales taxes may affect the equilibrium in non-trivial ways. For example, the introduction of a 15% sales tax may reduce search intensity in such a way that the tax ends up being passed on to the consumers more than proportionately and firms' profits actually rise.

The structure of the paper is as follows. In the next section, we review and modify the non-sequential consumer search model studied in the paper of Hong and Shum (2006). In Section 3 we discuss our maximum likelihood estimation method. Section 4 presents a Monte Carlo study that, among other issues, compares our estimation method with that of Hong and Shum. In Section 5 we estimate the search cost distribution underlying the price data obtained from some online markets for memory chips and show how the market would be affected if a sales tax was introduced. Finally, Section 6 concludes.

⁶Using a different method, Sorensen (2001) finds that between 5% to 10% of the consumers conduct an exhaustive search for prices in the market for prescription drugs.

2 The consumer search model

We study an oligopolistic version of the model proposed in Hong and Shum (2006); their model generalizes the non-sequential consumer search model of Burdett and Judd (1983) by adding search cost heterogeneity.⁷ The details of the model are as follows. There are N retailers selling a homogeneous good. Let r be the common unit selling cost of each retailer. There is a unit mass of identical buyers. Each consumer demands at most one unit of the good. Let \bar{p} be the consumer valuation for the good. Each buyer costlessly learns the price of one of the retailers at random. Beyond the first price, a consumer incurs a search cost c to obtain further price information. Consumers differ in their search costs. Assume that the cost of a consumer is randomly drawn from a distribution of search costs F_c . A consumer with search cost c sampling i firms incurs a total search cost equal to ic .

As in Burdett and Judd (1983), denote the symmetric mixed strategy equilibrium by the distribution of prices F_p , with density $f_p(p)$. Let \underline{p} and \bar{p} be the lower and upper bound of the support of F_p .⁸ Given firm behavior, the number of prices $i(c)$ a consumer with search cost c observes must be optimal, i.e.,

$$i(c) = \arg \min_{i>1} c(i-1) + \int_{\underline{p}}^{\bar{p}} ip(1 - F_p(p))^{i-1} f_p(p) dp. \quad (1)$$

Since $i(c)$ must be an integer, the problem in equation (1) induces a partition of the set of consumers into N subsets of size q_i , $i = 1, 2, \dots, N$, with $\sum_{i=1}^N q_i = 1$; thus, the number q_i is the fraction of buyers sampling i firms and is strictly positive for all i . This partition is calculated as follows. Let $Ep_{1:i}$ be the expected minimum price in a sample of i prices drawn from the price distribution F_p . Then

$$\Delta_i = Ep_{1:i} - Ep_{1:i+1}, i = 1, 2, \dots, N - 1 \quad (2)$$

denotes the search cost of the consumer indifferent between sampling i prices and sampling $i + 1$ prices. Note that Δ_i is a decreasing function of i . Using this property, the fractions of

⁷The oligopoly case is also studied in Janssen and Moraga-González (2004) but with a two-point search cost distribution.

⁸It will become clear later that the upper bound of the price distribution must be equal to the consumer valuation.

consumers q_i sampling i prices are simply

$$q_1 = 1 - F_c(\Delta_1); \quad (3a)$$

$$q_i = F_c(\Delta_{i-1}) - F_c(\Delta_i), \quad i = 2, 3, \dots, N - 1; \quad (3b)$$

$$q_N = F_c(\Delta_{N-1}). \quad (3c)$$

Given consumers' search behavior it is indeed optimal for firms to mix in prices. The upper bound of the price distribution must be \bar{p} ; this is because a firm that charges the upper bound sells only to the consumers who do not compare prices, i.e. consumers in q_1 , and these consumers would also accept \bar{p} . The equilibrium price distribution follows from the indifference condition that a firm should obtain the same level of profits from charging any price in the support of F_p , i.e.,

$$(p - r) \left[\sum_{i=1}^N \frac{i q_i}{N} (1 - F_p(p))^{i-1} \right] = \frac{q_1(\bar{p} - r)}{N}. \quad (4)$$

From equation (4) it follows that the minimum price charged in the market is

$$\underline{p} = \frac{q_1(\bar{p} - r)}{\sum_{i=1}^N i q_i} + r. \quad (5)$$

As shown in Hong and Shum (2006), equations (2) to (5) provide enough structure to allow for the estimation of the search cost distribution using only price data. Since quantity information is often hard to obtain, the focus of our next section will also be on estimation using the same kind of data.

3 Maximum likelihood estimation

Assume the researcher observes the prices of the N firms operating in the market.⁹ The objective is to estimate the collection of points $\{\Delta_i, q_i\}_{i=1}^N$ of the search cost distribution by maximum likelihood. Once we get these estimates we can construct an estimate of the search cost distribution by spline approximation. A difficulty here is that equation (4) cannot be solved for the equilibrium price distribution F_p and this makes it difficult to calculate the

⁹In practice, sometimes not all the firms are observed by the researcher; our Monte Carlo study in Section 4 examines the implication of this lack of data.

cutoff points

$$\Delta_i = \int_{\underline{p}}^{\bar{p}} p[(i+1)F_p(p) - 1](1 - F_p(p))^{i-1} f_p(p) dp, \quad i = 1, 2, \dots, N-1.$$

Hong and Shum (2006) propose to use the empirical price distribution to calculate the Δ_i 's. Even though this approach is practical, it does not necessarily provide minimal variance estimates. We proceed differently and obtain ML estimates of the cutoff points. To do this, we rewrite Δ_i as a function of the ML estimates of the parameters of the price distribution. This has the advantage that the asymptotic theory for computing the standard errors of Δ_i and for conducting tests of hypotheses remains standard.¹⁰ Integrating by parts, we first rewrite the cutoff points as

$$\Delta_i = \int_{\underline{p}}^{\bar{p}} F_p(p)(1 - F_p(p))^i dp, \quad i = 1, 2, \dots, N-1. \quad (6)$$

Since the distribution function $F(p)$ is monotonically increasing in p , its inverse exists. Using equation (4), we obtain the inverse function:

$$p(z) = \frac{q_1(\bar{p} - r)}{\sum_{i=1}^N i q_i (1 - z)^{i-1}} + r. \quad (7)$$

Using this inverse function, a change of variables in equation (6) yields:

$$\Delta_i = \int_0^1 p(z)[(i+1)z - 1](1 - z)^{i-1} dz, \quad i = 1, 2, \dots, N-1. \quad (8)$$

If we obtain ML estimates of $r, \bar{p}, \underline{p}$ and $q_i, i = 1, 2, \dots, N$, then, by the invariance property of ML estimation (see Greene, 1997), we can use equations (7) and (8) to calculate ML estimates of the cutoff points of the search distribution. This procedure yields a ML estimate of the search cost distribution $F_c(c)$.

We now discuss how to estimate r, \bar{p} and $q_i, i = 1, 2, \dots, N$, by maximum likelihood, assuming that the researcher has only price data. Since the price density cannot be obtained

¹⁰We would like to note now that, for our asymptotic arguments, we shall need that prices are independently and identically distributed in different periods, and since the number of firms is fixed and finite, that the number of periods goes to infinity.

in closed form, we apply the implicit function theorem to equation (4), which yields

$$f_p(p) = \frac{\sum_{i=1}^N i q_i (1 - F_p(p))^{i-1}}{(p - r) \sum_{i=1}^N i(i-1) q_i (1 - F_p(p))^{i-2}}. \quad (9)$$

Let $\{p_1, p_2, \dots, p_M\}$ be the vector of observed prices. Without loss of generality, let $p_1 < p_2 < \dots < p_M$. Following Kiefer and Neumann (1993) we use the minimum price in the sample p_1 and the maximum one p_M as estimates of the lower and upper bounds of the support of the price distribution \underline{p} and \bar{p} , respectively. These estimates of the bounds of the price cdf converge super-consistently to the true bounds.¹¹ Using the estimates of \underline{p} and \bar{p} , equation (5) can be solved to obtain the marginal cost r as a function of the other parameters:

$$r = \frac{p_1 \sum_{i=1}^N i q_i - q_1 p_M}{\sum_{i=2}^N i q_i}. \quad (10)$$

Plugging this formula into equation (9) and using the fact that $q_N = 1 - \sum_{i=1}^{N-1} q_i$ we can solve numerically the following maximum likelihood estimation problem:¹²

$$\max_{\{q_i\}_{i=1}^{N-1}} \sum_{\ell=2}^{M-1} \log f_p(p_\ell; q_1, q_2, \dots, q_N)$$

where

$$F_p(p_\ell) \text{ solves } (p_\ell - r) \left[\sum_{i=1}^N \frac{i q_i}{N} (1 - F_p(p_\ell))^{i-1} \right] = \frac{q_1 (\bar{p} - r)}{N}, \text{ for all } \ell = 2, 3, \dots, M - 1$$

We note that in this formulation the estimate of r is obtained from equation (10) as a function of the estimates of the other parameters. This procedure introduces some dependence between the price observations. Our Monte Carlo study in the next section shows that this approach works reasonably well, as the upper and lower bounds of the price distribution converge to the true values at a super-consistent rate.

The standard errors of the estimates of q_i , $i = 1, 2, \dots, N - 1$ are calculated in the usual way, i.e., by taking the square root of the diagonal entries of the inverse of the negative

¹¹See also Donald and Paarsch (1993) on using order statistics to estimate the lower and upper bound of bid distributions.

¹²The numerical procedure is as follows. We take arbitrary starting values $\{q_i^0\}_{i=1}^{N-1}$. Then, for every price p_ℓ in the data set, we calculate $F_p(p_\ell)$ using the equilibrium condition (4), which in turn allows us to calculate $f_p(p_\ell)$ using (9). We use a trust region PCG method, which proceeds by changing the q_i 's until the log-likelihood function is maximized.

Hessian matrix evaluated at the optimum. Since $q_N = 1 - \sum_{i=1}^{N-1} q_i$, we can calculate the standard error of the estimate of q_N using the Delta method. The same applies to the standard errors of the estimates of the marginal cost r and the Δ_i 's, since they are obtained as transformations of the estimated q_i 's.

4 A Monte Carlo study

The study in this section has various purposes. First, we investigate how precise the maximum likelihood estimates of the price and search cost distributions are when the number of price observations is limited. In particular, we are interested in the type of bias that the estimation of the upper and lower bound of the price distribution by the maximum and the minimum prices observed in the data may cause. Secondly, we investigate the implications of underestimating the number of firms N that operate in the market, which may be a problem in real-world applications. Finally, we compare our estimation method to that of Hong and Shum (2006).

4.1 Performance of the estimates

The general setup of the Monte Carlo experiment is as follows. We assume that consumers' search costs are drawn independently from a log-normal distribution with parameters $\nu = 0.5$ and $\sigma = 5$. Moreover, the value of the product \bar{p} is assumed to be 100 and the unit cost r to be 50. To solve for equilibrium, we compute numerically the fractions $\{q_1, q_2, \dots, q_N\}$ for which equations (3a)-(3c) and (8) hold simultaneously. Next, we use these parameters to construct the equilibrium price distribution, implicitly defined by equation (4). After this, we draw prices randomly from the cdf of prices, which serve as input for the maximum likelihood estimation procedure described in the previous section. We replicate each of our experiments 1000 times and report the mean and the 90% confidence interval of the estimates we obtain.

In this subsection we set $N = 25$. The first column of Table 1 gives the true parameter (equilibrium) values. We see that the primitives chosen lead to an equilibrium where price dispersion is substantial. In particular, the lowest price of the equilibrium price distribution is 51.68, which is about half the maximum price, 100. Thus, in equilibrium gains from search

are quite significant. We also note that a firm charging the minimum price has a relative price-cost margin (Lerner index) of only 3.36% while for the firm charging the maximum price the same index is 50%.

	TRUE	EXP1	EXP2	EXP3
T	-	4	10	20
N	25	25	25	25
M	-	100	250	500
r	50.00	47.50 (6.53)	49.16 (2.03)	49.47 (1.31)
p	51.68	52.06 (0.42)	51.83 (0.16)	51.75 (0.07)
v	100.00	99.91 (0.09)	99.97 (0.03)	99.98 (0.02)
q_1	0.380	0.451 (0.154)	0.417 (0.098)	0.405 (0.072)
q_2	0.032	0.037 (0.017)	0.035 (0.012)	0.035 (0.009)
q_3	0.026	0.025 (0.032)	0.024 (0.026)	0.025 (0.022)
q_4	0.022	0.025 (0.041)	0.024 (0.033)	0.027 (0.032)
q_5	0.020	0.022 (0.042)	0.027 (0.042)	0.022 (0.034)
q_6	0.018	0.022 (0.046)	0.018 (0.036)	0.016 (0.030)
q_7	0.016	0.018 (0.046)	0.015 (0.035)	0.016 (0.032)
q_8	0.015	0.011 (0.038)	0.015 (0.036)	0.016 (0.034)
q_9	0.014	0.013 (0.046)	0.012 (0.033)	0.013 (0.032)
q_{10}	0.013	0.013 (0.049)	0.014 (0.041)	0.013 (0.033)
q_{11}	0.013	0.011 (0.051)	0.010 (0.035)	0.011 (0.031)
q_{12}	0.012	0.013 (0.054)	0.011 (0.037)	0.011 (0.033)
q_{13}	0.011	0.010 (0.044)	0.010 (0.032)	0.009 (0.027)
q_{14}	0.011	0.009 (0.043)	0.011 (0.036)	0.010 (0.032)
q_{15}	0.010	0.013 (0.052)	0.017 (0.055)	0.010 (0.033)
q_{16}	0.010	0.011 (0.045)	0.015 (0.046)	0.013 (0.042)
q_{17}	0.009	0.012 (0.057)	0.013 (0.043)	0.014 (0.041)
q_{18}	0.009	0.016 (0.072)	0.014 (0.044)	0.017 (0.045)
q_{19}	0.008	0.012 (0.053)	0.014 (0.044)	0.019 (0.050)
q_{20}	0.008	0.011 (0.052)	0.017 (0.053)	0.020 (0.052)
q_{21}	0.008	0.010 (0.049)	0.017 (0.051)	0.022 (0.056)
q_{22}	0.007	0.011 (0.058)	0.016 (0.050)	0.021 (0.054)
q_{23}	0.007	0.012 (0.069)	0.018 (0.056)	0.018 (0.050)
q_{24}	0.007	0.012 (0.060)	0.029 (0.066)	0.019 (0.047)
q_{25}	0.314	0.199 (0.216)	0.186 (0.175)	0.197 (0.182)

Notes: Standard errors in parenthesis.

Table 1: True and estimated parameter values

In equilibrium a great deal of the consumers, about 38%, search for only one price; another important group of buyers searches for all the prices in the market (about 31% of the consumers). The fractions of consumers searching for an intermediate number of prices (from 2 to 24 firms) are pretty small, in all cases less than about 3% and often close to zero.¹³

As discussed above, for the estimation of the model we need to assume that market

¹³This feature of the equilibrium partition of the set of consumers that few consumers search for an intermediate number of firms is somewhat special and has to do with the choice of search cost distribution. For example, in a 10 firm market where the search cost distribution is a twenty-eighty percent mixture of a log-normal with parameters 0.5 and 2 and a gamma distribution with parameters 0.5 and 0.2, the equilibrium has most of the consumers searching intensively (around 75% more than 8 times) and very few consumers not searching at all (around 4%).

interaction evolves over a finite number of $T \geq 2$ periods. We take the equilibrium of the static game described in Section 2 as the equilibrium of the repeated game with finite horizon. Our first set of estimations assumes the market evolves over $T = 4$ periods so we draw a total of 100 price observations each time we run the estimation procedure.

The second column of Table 1 gives the results of our first set of estimations. The numbers reported are the mean of the 1000 estimates of the parameters with corresponding standard errors in parenthesis. We observe that the estimate of the fraction of consumers who search for one price only is about 45% and highly significant. This estimate is about 7% higher than the true value so the fraction of consumers who do not compare prices at all is overestimated. The estimate of the fraction of consumers searching for two prices is also significant and again overestimated (3.7% instead of 3.2%). The estimate of the fraction of consumers searching for all prices in the market is about 20%, somewhat lower than the true parameter (31.4%). The estimates of the rest of the parameters are not significantly different from zero (at the 5% level). Since the true parameters are close to zero anyway, it turns out that this is not a problem for the estimate of the price distribution to exhibit a good fit. In sum, we see that the fractions of consumers searching little are overestimated while the fractions of consumers searching a lot are underestimated. Arguably the implication of these biases is that the estimate of the search cost distribution will be biased towards high search costs.

The first column of Table 2 reports the true cutoff points of the search cost distribution. The second column gives the estimated ones when we set $T = 4$. We see that all the estimates of the cutoff points are highly significant, and quite close to the true ones.

The price and search cost distributions as well as their mean estimates are plotted in Figure 1(a) and 1(b) respectively. In these graphs the solid curves are the true distributions while the thick dashed curves show the mean of the 1000 estimated distributions. The thin dashed curves are respectively the 5% percentile and the 95% percentile of the estimates. We observe that the estimate of the price distribution is remarkably close to the true price cdf. However, the estimate of the search cost distribution lies below the true one. In spite of this, the true distribution falls (for its most part) within the 90% confidence interval.¹⁴

¹⁴The last cutoff point of the search cost distribution we can estimate is c_1 and therefore we do not have information about search costs beyond that point.

	TRUE	EXP1	EXP2	EXP3
T	-	4	10	20
N	25	25	25	25
M	-	100	250	500
Δ_1	7.60	7.40 (0.63)	7.54 (0.40)	7.57 (0.28)
Δ_2	5.01	4.85 (0.34)	4.95 (0.21)	4.98 (0.15)
Δ_3	3.59	3.47 (0.21)	3.55 (0.13)	3.57 (0.09)
Δ_4	2.71	2.62 (0.16)	2.68 (0.09)	2.69 (0.07)
Δ_5	2.12	2.04 (0.13)	2.09 (0.08)	2.10 (0.06)
Δ_6	1.69	1.64 (0.12)	1.67 (0.07)	1.68 (0.05)
Δ_7	1.38	1.33 (0.11)	1.36 (0.07)	1.36 (0.05)
Δ_8	1.14	1.11 (0.11)	1.12 (0.07)	1.13 (0.05)
Δ_9	0.95	0.93 (0.10)	0.94 (0.07)	0.94 (0.05)
Δ_{10}	0.80	0.79 (0.10)	0.80 (0.06)	0.80 (0.05)
Δ_{11}	0.69	0.68 (0.09)	0.68 (0.06)	0.68 (0.04)
Δ_{12}	0.59	0.58 (0.09)	0.59 (0.06)	0.59 (0.04)
Δ_{13}	0.51	0.51 (0.08)	0.51 (0.05)	0.51 (0.04)
Δ_{14}	0.45	0.45 (0.08)	0.44 (0.05)	0.44 (0.04)
Δ_{15}	0.39	0.39 (0.07)	0.39 (0.05)	0.39 (0.04)
Δ_{16}	0.34	0.35 (0.07)	0.34 (0.04)	0.34 (0.03)
Δ_{17}	0.31	0.31 (0.06)	0.31 (0.04)	0.30 (0.03)
Δ_{18}	0.27	0.28 (0.06)	0.27 (0.04)	0.27 (0.03)
Δ_{19}	0.24	0.25 (0.06)	0.24 (0.04)	0.24 (0.03)
Δ_{20}	0.22	0.22 (0.05)	0.22 (0.04)	0.22 (0.03)
Δ_{21}	0.20	0.20 (0.05)	0.20 (0.03)	0.20 (0.02)
Δ_{22}	0.18	0.19 (0.05)	0.18 (0.03)	0.18 (0.02)
Δ_{23}	0.16	0.17 (0.05)	0.16 (0.03)	0.16 (0.02)
Δ_{24}	0.15	0.15 (0.04)	0.15 (0.03)	0.15 (0.02)

Notes: Standard errors in parenthesis.

Table 2: True and estimated critical search cost values

The fact that the estimate of the search cost distribution lies below the true cdf might be a reflection of the fact that the fractions of consumers searching little are overestimated. Since the estimate of the upper bound of the price cdf with a limited amount of data will be lower than the true one, while, at the same time, the estimate of the lower bound of the price cdf will be higher than the true one, the price distribution might very well be less dispersed than the true one. This implies that gains from search might be lower than in equilibrium, which is consistent with estimates of the search cost distribution being biased towards higher search costs.

To see whether the downward bias of the estimated search cost distribution can be attributed to the biased estimation of the upper and lower bounds of the price cdf, we conduct the following experiment. We assume that the econometrician knows the true upper and lower bounds and then re-estimate the model. The thick dashed curve in Figure 2(a) shows the new estimate of the search cost distribution. To compare with the previous estimates, we also plot in gray the search cost distribution when the upper and the lower bounds are estimated by the minimum and the maximum price. The graph reveals that the new esti-

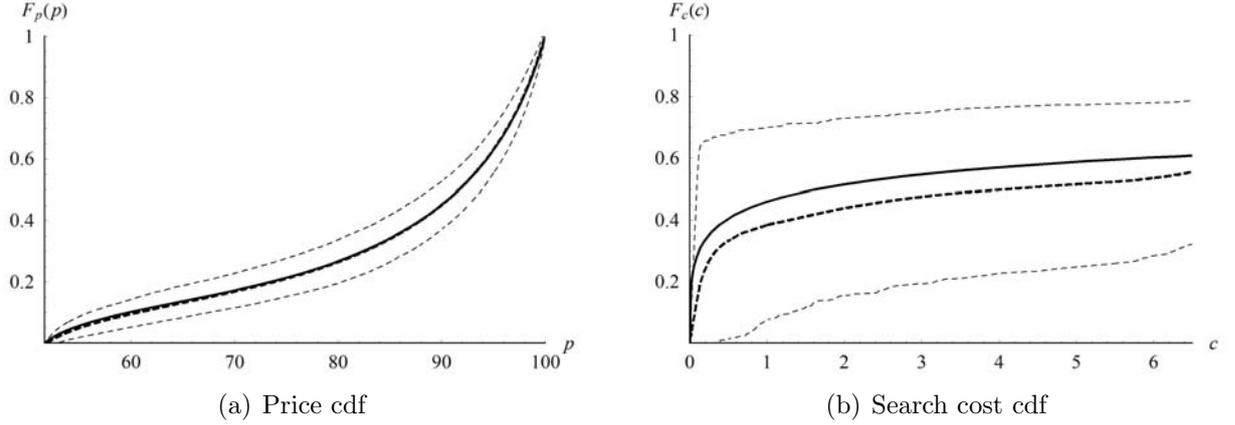


Figure 1: Estimated price and search cost cdf's (# periods $T = 4$; # obs. $M = 100$)

mate is much closer to the true distribution than the previous one. The average estimate of q_1 goes down 0.04 points and this results in an upward shift of the estimated search cost cdf.

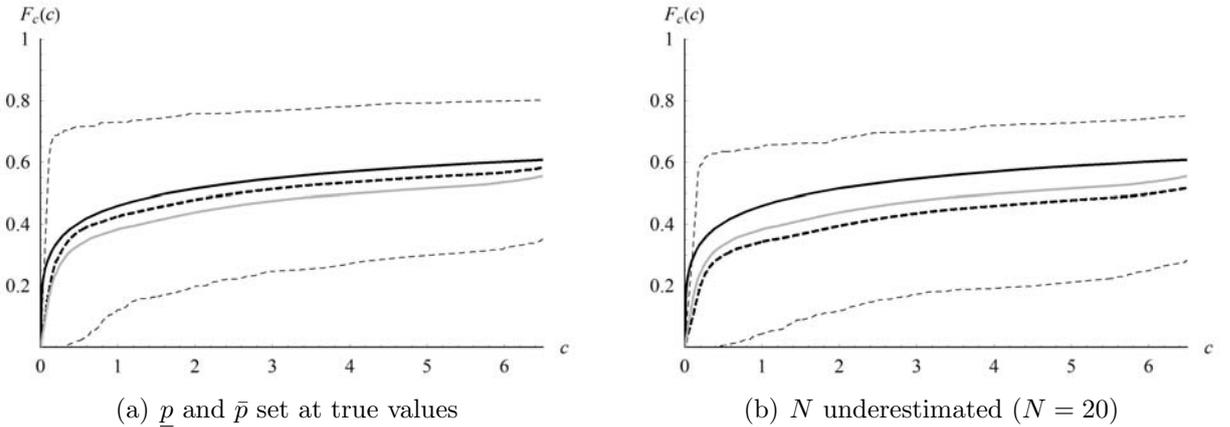


Figure 2: Estimated search cost cdf (# periods $T = 4$; # obs. $M = 100$)

An alternative explanation for the bias we observe is simply based on the fact that the maximum likelihood estimator is biased for finite samples. Indeed, we see that the gap between the true and the estimated search cost distributions becomes smaller as we increase the number of observations. This can be seen in columns 3 and 4 of Tables 1 and 2, where the number of periods over which the market develops is set equal to 10 and 20, respectively (correspondingly, the number of observations in each simulation goes up to 250 and 500). The tables show that the estimates of the parameters of the price distribution (including the upper and lower bounds) become more precise and this leads to more accurate estimates of the search cost cdf. The effect of an increase in the number of observations on the estimates

can be seen in Figures 3(a) and 3(b).

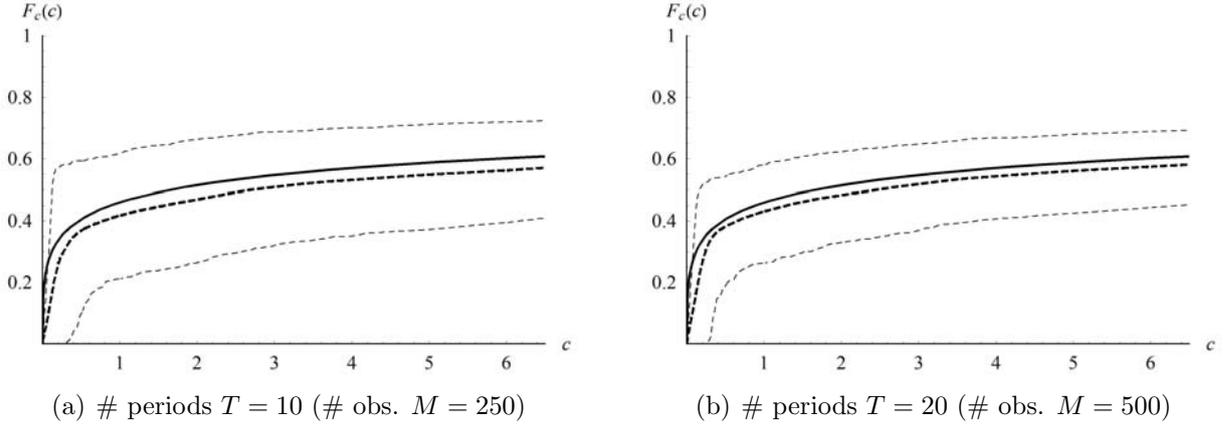


Figure 3: Estimated search cost cdf

4.2 Measurement error in the number of firms

The econometrician might often encounter the problem that he/she does not know the exact number of retailers operating in a market. To investigate the impact of measurement error in the number of firms N , we conducted two experiments. In the first experiment we set N equal to 20 instead of equal to the true value 25. This experiment captures a situation where the econometrician observes N with some, but not very large, error (20% fewer firms). The results can be seen in Figure 2(b), where again the gray curve depicts the estimated search cost distribution using the true number of firms. As the graph shows, the underestimation of the number of firms did not change the shape of the search cost distribution. The average estimate of q_1 went up from 0.451 to 0.487 and this led to a greater downward bias in the estimate of the search cost distribution.

In the second experiment, we set N equal to 4 instead of equal to 25. In this case, the econometrician measures the number of firms with a pretty large error. Note that this is equivalent to assuming that low search cost consumers are ignored altogether, which is something that should show up in the estimate of the search cost distribution. Figure 4 gives the estimation results. In panel 4(a) the thick dashed curve shows the estimated price cdf, while the true cdf is depicted by the solid curve. Clearly, the fit is much worse than if we had estimated N more or less correctly (see Figures 1(a) and 2(b)). This has a large impact on the estimates of search cost and marginal cost parameters. As can be seen in Figure 4(b),

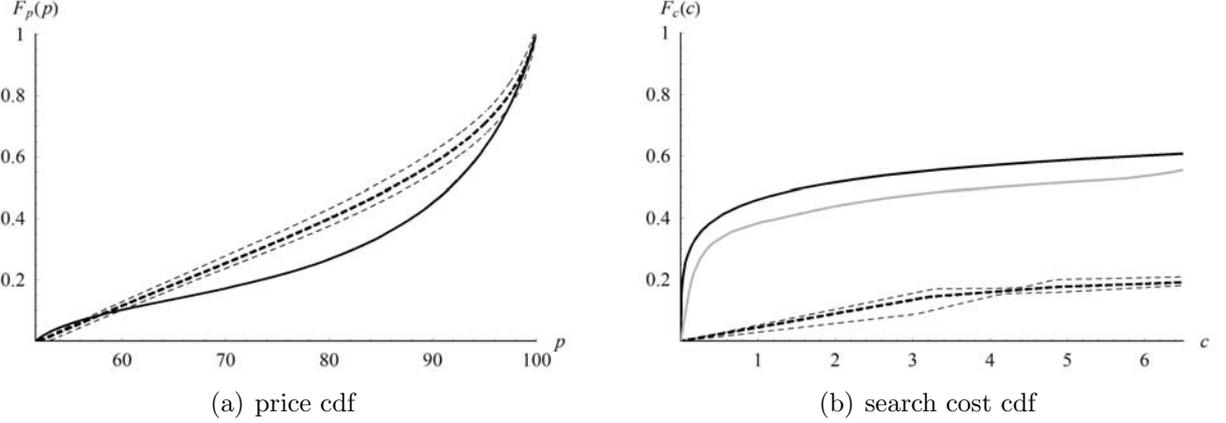


Figure 4: Estimated equilibrium price and search cost cdf ($N = 4$; # periods $T = 4$)

the estimated search cost cdf (thick dashed curve) is far from the true one (solid curve). In particular, the estimates lead to the wrong conclusion that search costs are much higher than what they actually are. Likewise, this translates into an average price-cost margin being largely exaggerated; in particular it is around 100%, while the true price-cost margin of a typical firm is 42%.

In sum, this subsection suggests that the estimates of the search cost cdf are meaningful even when the econometrician does not know the exact number of firms operating in the market but has a fair estimate of it; in addition, it clarifies the nature of the bias introduced by this type of measurement error.

4.3 Comparison of the ML estimation method and Hong and Shum's (2006) method

Hong and Shum (2006) propose to estimate the parameters of the price distribution by maximum empirical likelihood (MEL) and the cutoff points of the search cost distribution by using the empirical cdf of prices. In this subsection we compare the performance of our method relative to the one of Hong and Shum.

Before proceeding with the simulation results, let us review their approach briefly.¹⁵ Suppose we have a data set containing M prices. Consider the discrete price distribution $\hat{F}_p(p) = \sum_{j=1}^M \pi_j \mathbf{1}(p_j \leq p)$ with M mass points, each price p_j being charged with probability

¹⁵For details, we refer to the Appendix of Hong and Shum (2006).

π_j . Using the equilibrium condition (4), each price $i = 1, 2, \dots, M - 1$ in the data set satisfies

$$(p_i - r) \left[\sum_{k=1}^N k q_k \left(1 - \left[\frac{1}{M} \sum_{j=1}^M \pi_j \mathbf{1}(p_j \leq p_i) \right] \right)^{k-1} \right] = (\bar{p} - r) q_1, \quad (11)$$

where, as before, r and q_N can be eliminated from these expressions using the formula for the lower bound of the price distribution and the summing up condition $\sum_{k=1}^N q_k = 1$, respectively. The equations in (11) can be transformed into moment conditions as follows. For $s_\ell \in [0, 1]$, $\ell = 1, 2, \dots, L$ we have

$$F_p^{-1}(s_\ell) = r + \frac{(\bar{p} - r) q_1}{\sum_{k=1}^N k q_k (1 - s_\ell)^{k-1}} \equiv g_{s_\ell}(q_1, q_2, \dots, q_N).$$

Hong and Shum (2006) write these population quantile restrictions as

$$\frac{1}{M} \sum_{j=1}^M \pi_j \left[\mathbf{1} \left(p_j \leq r + \frac{(\bar{p} - r) q_1}{\sum_{k=1}^N k q_k (1 - s_\ell)^{k-1}} \right) - s_\ell \right] = 0. \quad (12)$$

The empirical likelihood problem consists of maximizing $\sum_{j=1}^M \log \pi_j$ subject to the constraints in (12) and the condition $\sum_{j=1}^M \pi_j = 1$ with respect to the probabilities π_j 's and the parameters $\{q_1, q_2, \dots, q_{N-1}\}$. It turns out that the MEL estimates of the parameters can be obtained from solving the saddle-point problem:

$$\max_{\{q_i\}_{i=1}^{N-1}} \min_{\{t_m\}_{m=1}^{M-1}} \sum_{j=1}^M \log \left(1 + \mathbf{t}' \left[\mathbf{1} \left(p_j \leq r + \frac{(\bar{p} - r) q_1}{\sum_{k=1}^N k q_k (1 - s_\ell)^{k-1}} \right) - s_\ell \right] \right),$$

where \mathbf{t} denotes the Lagrange multipliers associated with the constraints in equation (11). The MEL estimates of the parameters of the price distribution form the ordinates of the cutoff points of the search cost distribution. To find the abscissas of the cutoff points, Hong and Shum propose to use the empirical cdf of prices in equation (6) above.

The maximum empirical likelihood method of Hong and Shum requires to solve a constrained optimization problem where the number of constraints equals, potentially, the number of price observations. Essentially, this requires the optimization of a Lagrangian function in $N - 1 + M$ parameters, where M is the number of Lagrange multipliers. When N and/or M is relatively large, the dimension of the problem makes it computationally difficult. In fact, we have often witnessed the algorithm not to converge, unless the starting point was

very close to the true vector of parameters. The same sort of numerical problems have been reported by Hong and Shum themselves (see footnote *a* of their Table 2) and by other authors (see e.g. Owen, 1990 and Qin and Lawless, 1994). To overcome the numerical problems, Hong and Shum (2006) suggest not to use all moment conditions but a small subset of them; we followed this approach in our initial set of simulations and still encountered difficulties. The reason is that in our initial set of simulations we considered a market operated by a large number of firms, in particular $N = 25$.

To be able to perform a comparison of the two methods, we studied a market with fewer firms, in particular we set $N = 10$; the rest of the parameters were kept the same as in the main set of simulations. Even in this case of 10 firms, we experienced some numerical difficulties but, fortunately, they became salvageable in limited time by trying several starting values. The results of the simulations are reported in Tables 2(a) and 2(b).

(a) True and estimated parameter values

	TRUE	MEL	ML
T	-	10	10
N	10	10	10
M	-	100	100
r	50.00	48.37 (3.68)	47.97 (7.27)
\underline{p}	53.29	53.56 (0.27)	53.56 (0.27)
v	100.00	99.89 (0.11)	99.89 (0.11)
q_1	0.370	0.376 (0.142)	0.421 (0.111)
q_2	0.038	0.040 (0.029)	0.043 (0.018)
q_3	0.032	0.060 (0.042)	0.033 (0.039)
q_4	0.029	0.066 (0.046)	0.025 (0.045)
q_5	0.026	0.070 (0.048)	0.027 (0.060)
q_6	0.023	0.071 (0.048)	0.036 (0.088)
q_7	0.021	0.079 (0.050)	0.050 (0.107)
q_8	0.020	0.079 (0.050)	0.058 (0.114)
q_9	0.018	0.080 (0.050)	0.051 (0.102)
q_{10}	0.422	0.080 (0.051)	0.257 (0.221)

Notes: Standard errors in parenthesis.

(b) True and estimated Δ_i 's

	TRUE	Empirical cdf	ML
T	-	10	10
N	10	10	10
M	-	100	100
Δ_1	8.640	8.556 (0.464)	8.470 (0.463)
Δ_2	5.264	5.193 (0.197)	5.135 (0.196)
Δ_3	3.484	3.432 (0.178)	3.392 (0.154)
Δ_4	2.428	2.393 (0.184)	2.366 (0.153)
Δ_5	1.756	1.734 (0.177)	1.716 (0.147)
Δ_6	1.309	1.295 (0.162)	1.284 (0.136)
Δ_7	0.999	0.992 (0.146)	0.985 (0.123)
Δ_8	0.779	0.776 (0.129)	0.773 (0.110)
Δ_9	0.619	0.619 (0.114)	0.617 (0.098)

Notes: Standard errors in parenthesis.

Table 3: Estimation results (ML vs Hong and Shum's (2006) method)

Table 2(a) reports the true parameters of the price distribution along with the MEL and ML estimates. As in the above experiments, the equilibrium has the features that the group of consumers searching only once is fairly large (about 37%), and that the fraction of consumers searching thoroughly is also sizable (about 42%). The second column of the Table shows that the MEL method is unable to capture the effect on prices of the consumers who search intensively; in fact, this parameter is quite poorly estimated, which would wrongly suggest that search costs are higher than what they really are. By contrast, as it can be

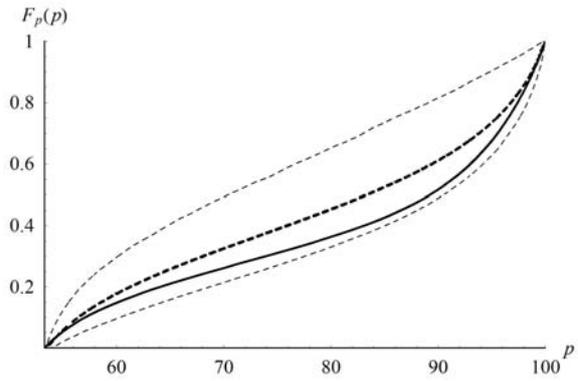
seen in the third column of the Table, our ML procedure yields a pretty good estimate of this parameter, as well as of the others; in fact, it can be seen that except for the estimates of q_1 and q_2 , all ML estimates are closer to the true parameters than the MEL estimates.

Table 2(b) shows the true cutoff points of the search cost distribution, along with the corresponding estimates obtained using the empirical cdf of prices and by maximum likelihood. It can be seen that the estimates using the empirical price cdf are closer to the true parameters but they generate larger standard errors than the ML estimates.

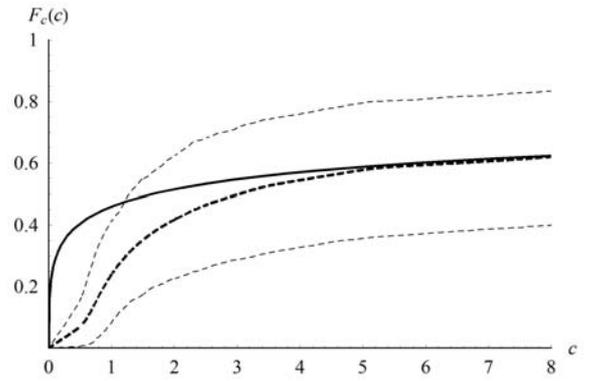
These differences between the two methods are reflected both in the estimate of the price cdf as well as in the estimate of the search cost distribution, which can be seen in Figure 5. Figure 5(a) shows the MEL estimate of the price cdf (thick dashed curve) along with the corresponding 90% confidence interval (thin dashed curves) and the true equilibrium price distribution (solid curve). The graph reveals that the estimated price cdf has prices larger than what they really are; this is consistent with the fact that the MEL method underestimated the extent of price comparison in the market. The fit of our ML estimate of the price cdf can be seen in Figure 5(c); the graph reveals that the fit is remarkably good and clearly outperforms the alternative MEL estimate.

Figure 5(b) shows the estimate of the search cost distribution using Hong and Shum's procedure (thick dashed curve), along with the true lognormal distribution with parameters 0.5 and 5 (solid curve). While the fit of the estimate is pretty good at relatively high quantiles, the estimate suggests search costs are higher than what they really are at low quantiles. Again, this is consistent with the fact that the MEL method underestimated the amount of search in the market. The alternative ML method generates an estimate of the search cost cdf that can be seen in Figure 5(d). This estimate resembles closely the estimate we obtained for the earlier simulations with 25 firms. Again, search costs are underestimated but to a lower extent than in the case of the estimate based on Hong and Shum's method.

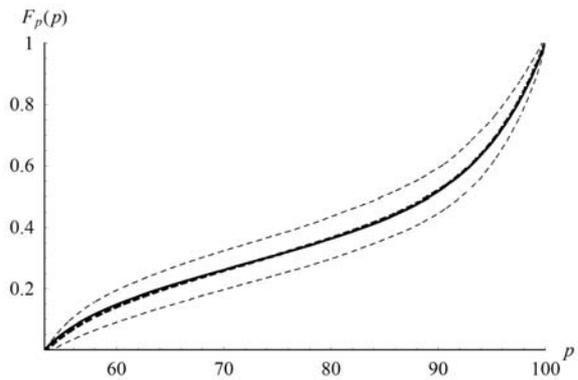
On the basis of this evidence, we conclude that our ML procedure outperforms the hybrid procedure described in Hong and Shum (2006), both from a numerical as well as from a goodness-of-fit point of view.



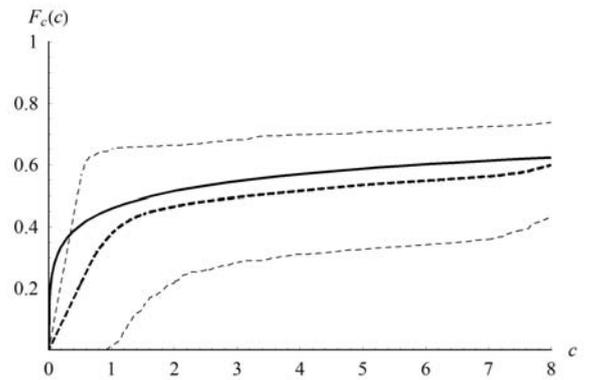
(a) Price cdf MEL method



(b) Search cost cdf Hong and Shum's method



(c) Price cdf ML method



(d) Search cost cdf ML method

Figure 5: True and estimated price and search cost cdf (ML vs Hong and Shum's (2006) method)

5 Empirical application

5.1 Data and empirical issues

In this section we apply our estimation procedure to a data set obtained from real-world markets. Before presenting the results, we discuss the data set and, following Lach (2002), check one by one the assumptions of the theoretical model. This also serves to identify the potential weaknesses and caveats of our empirical exercise.

The focus of the study is on on-line consumer markets for personal computer memory chips. At the time of data collection, the four memory chips we study were sold in stores advertising prices on the price comparison site www.shopper.com so we used this site to sample the prices of the firms over time.¹⁶

Shopper.com is one of the largest price comparison sites on the Internet.¹⁷ Internet shops get listed on shopper.com by subscribing to CNET, the owner of shopper.com. Stores can choose between three types of listing schemes, general, preferred or premier. Preferred or premier listing allows a shop to add a store logo. Shops can provide once or twice a day price information by uploading a so-called “price feed,” but it is not necessary to do so if a shop does not desire to alter its current price. The feed is collected four times a day and published on shopper.com. Shops are required to fill in eight fields in the feed: credit card price, manufacturer name, manufacturer Stock Keeping Unit (SKU), product URL, product name, availability, shipping and handling cost, and category.

By using a so-called “spider” computer code, we automatically collected this information for the four memory chips directly from shopper.com, from the beginning of August 2004 till the end of September 2004. Unfortunately we could not collect more data because the IP address of the computer we were using to download the data was blocked by the system managers of shopper.com at the end of September 2004.

¹⁶Since some consumers may proceed as we did and use the search engine to sample prices, an implicit assumption for our sampling method to be reasonable is that firms do not price discriminate between regular visitors to their web sites and visitors of search engines. We have manually checked this assumption and found overwhelming evidence that firms announce the same price in their web sites and in the search engines. Our estimate of q_N will of course include those individuals who use the search engine so our interpretation for these consumers is that they have search costs less than Δ_{N-1} . Under this interpretation, our estimates give the search costs of those consumers buying memory chips online.

¹⁷We are implicitly assuming that retailers, dealers, computer manufacturers, etc. buy from agents in the value chain other than the firms advertising in shopper.com, or directly from the memory chip manufacturers.

A first caveat of the study is that fitting the model in Section 2 to data from on-line markets assumes implicitly that consumers search for prices non-sequentially. Even though non-sequential search may be a good approximation of buyer behavior when consumers use web sites, web forums, and search engines to find price information, a caveat of the analysis is that sequential search might be a more adequate search protocol to model search activity on the World Wide Web.¹⁸

We selected four memory chips all manufactured by Kingston, which is by far the largest producer in the sector (the 2004 market share of Kingston was 27.0%, while the second biggest producer of memory chips, Smart Modular Technologies, had a 2004 market share of 8.1%). The details of these four products are given in Table 4.

Ellison and Ellison (2005) have pointed out that in these types of market firms often engage in “bait and switch” strategies. We selected the memory chips to avoid this potential problem: chosen chips were at the moment of data collection somewhat at the top of the product line, exhibiting relatively large storage capacity (512 megabytes) and fast speed of operation (above 266 MHz). Two of the memory chips are of the SO-DIMM (Small Outline Dual In-line Memory Modules) type, which are intended for notebooks only.

It may be argued that different memory chips are in the same relevant market so a differentiated products market model is more appropriate than our model with homogeneous products. To avoid this problem to the extent possible, we included in the analysis only memory chips intended for particular PC’s. More concretely, we chose two memory chips for notebooks, one intended for Toshiba notebooks and the other for Dell Inspiron notebooks. Arguably, consumers who own for example a Toshiba Satellite 5105 notebook and are contemplating to extend its memory by 512 MB would most likely consider to buy only the Kingston KTT3614 memory chip (see www.toshiba.com).¹⁹ The other two memory chips are intended for Dell desktop computers, in particular for the Dimension series.

Another form of heterogeneity we are ignoring is store differentiation. Like in Hortacısu and Syverson (2004), it would be reasonable to assume that consumers may sample the firms with unequal probability, simply because some firms are more popular than others, or

¹⁸For details on the optimality of non-sequential and sequential search see Morgan and Manning (1985).

¹⁹The information available at www.toshiba.com suggests that Kingston memory chips are original parts used by Toshiba. For many consumers buying the same part as the original part is important (see Delgado and Waterson’s (2003) study of the UK tyre market).

because they advertise more effectively than others. The main problem with this extension is that we would need quantity data to estimate the model, which we do not have.

We view the markets under study as consumer markets, where the typical buyer is an individual consumer. In this sense, the usual buyer is expected to buy a single chip to upgrade the memory capacity of his/her computer; indeed, often computers have just a single slot available for memory upgrades. As a result, the inelastic demand assumption of the model seems reasonable here.

Product name	KTT3614	KTDINSP8200	KTD4400	KTD8300
Manufacturer	Kingston	Kingston	Kingston	Kingston
PC	Toshiba notebooks	Dell Inspiron 8200	Dell Dimension 4400	Dell Dimension 8300
MB	512MB DDR SDRAM	512MB DDR SDRAM	512MB DDR SDRAM	512MB DDR SDRAM
Memory Speed	PC2100 (266 MHz)	PC2100 (266 MHz)	PC2100 (266 MHz)	PC3200 (333 MHz)
Type	SO-DIMM	SO-DIMM	DIMM	DIMM

Notes: SO-DIMM memory chips are for notebooks while DIMM memory chips are for desktop computers.

Table 4: List of products

Product name	KTT3614	KTDINSP8200	KTD4400	KTD8300
Total No. of Stores	25	24	24	23
Mean No. of Stores (Min, Max)	22.4 (20, 24)	21.8 (19, 23)	21.8 (19, 23)	20.3 (17, 22)
Mean Weeks in Sample (Std)	7.16 (1.72)	7.25 (1.70)	7.25 (1.73)	7.04 (1.74)
No. of Observations	179	174	174	162
Mean Price (Std)	142.96 (24.34)	142.09 (21.33)	117.56 (18.34)	126.02 (20.58)
Max. and Min. Prices	208.90, 115.00	200.50, 109.20	170.50, 96.00	182.50, 102.00
Coefficient of Variation (as %)	17.02	15.01	15.60	16.33

Notes: Prices are in US dollars. Pooled data is used for the estimates of the mean, max. and min. prices and the coefficient of variation.

Table 5: Summary statistics

The summary statistics of the data can be found in Table 5. We found distinct numbers of stores operating in different markets but in all cases the number was quite high. For the KTT3614 memory chip, 25 firms were seen quoting prices over the period under study; for the KTDINSP8200 and KTD4400 chips we collected prices from 24 different stores, and for the KTD8300 chip we found 23 stores.

In our study, we estimated N by the total number of firms which were listed in shopper.com. This number is based on the sample of firms that advertise in shopper.com and is probably lower than the true number of stores in the relevant market. Our Monte Carlo simulations above show the extent of the bias introduced by measuring N incorrectly. If the

true number of retailers is not dramatically different than our estimate of N , the results, though biased, will be economically meaningful.²⁰

Not every firm was quoting a price every week. For example, for the KTT3614 memory chip we saw an average of 22.4 stores quoting a price in a typical week. The lowest number of stores for this product was 20 and the highest number of stores was 24. Similar figures were found for the other products (see Table 5). There might be several reasons for this variation. For some stores there were missing values somewhere in the middle of the sampling period. This might have been due to technical problems when uploading the price feed to shopper.com. We also observed that some stores appeared in the sample only after some weeks had passed. In any case, on average, a typical firm was quoting prices during more than 88% of the sample period (7 weeks out of 8).

The estimations are conducted under the assumption that firms play a stationary repeated game of finite horizon so, in every period, the data should reflect the equilibrium of the static game analyzed in Section 2. This assumption has some testable implications. One, since the equilibrium is in mixed strategies, prices should be dispersed at any given moment in time. Two, since firms are supposed to draw prices from the same price cdf period after period, there should be variation in the position of a typical firm in the price ranking and prices should not exhibit serial correlation. Third, stationarity of the environment implies that price distributions should be similar across periods, i.e., that supply or demand shocks have been absent during the sample period. We now examine how these three features appear in the data.

Table 5 shows the mean price and corresponding standard deviation of prices, for each product. As expected, memory chips for notebooks are on average more expensive than those intended for desktop computers; moreover, the KTD8300 chip is more expensive than the KTD4400 chip due to its faster speed of operation. For all the products, we observe significant price dispersion as measured by the coefficient of variation. On average, relative to buying from one of the firms at random, the gains from being fully informed in this market are sizable, ranging from 21.56 to 32.89 US dollars.

A careful examination of the data reveals that most stores certainly change their price

²⁰For robustness purposes, we also estimated the model taking 5 more firms than those seen in the data. The qualitative nature of our results did not change significantly.

from time to time, but we observe that they do not do it synchronously, that is, the length of time between price revisions changes from firm to firm. For example, in the market for the KTT3614 memory chip, 20 stores out of 25 changed their price at least once during the period under study. On average, a typical firm selling the KTT3614 chip changed the price once every 5 weeks; however, while some firms did change their prices several times over the sample period (up to 5 times), other firms did not. For the other memory chips, we found similar patterns.²¹ The reason for this variation may be due to menu cost dispersion across firms.

We also observe some variation in the price ranking of a typical firm. For example, for the KTT3614 memory chip the standard deviation of the ranking of a firm ranges from 0 to 3.77. This is somewhat smaller than what we would expect on the basis of the theoretical model. One reason for these findings might be the short length of time of the sample period because some of the firms did not alter their prices.²² To check this hypothesis, we gathered prices at the time of writing this paper and compared the current ranking of a typical firm with that at the time of data collection (one year ago). For example, for the KTT3614 memory chip, we found 21 stores quoting prices so some stores seem to be no longer active in this market. This is not surprising since this market evolves very rapidly so after one year a product may be somewhat outdated. Of these 21 stores, 17 stores were either higher or lower in the ranking compared to one year ago. The difference in ranking ranged from 0 to 6 and was on average of 2.19. Finally, 9 stores out of 21 are now in a different quartile of the ranking distribution. Similar figures apply to the other memory chips.²³

To check for serial correlation, we calculated the autocorrelation function (ACF) for each

²¹A typical firm selling the KTDINSP8200 chip changed its price once every 6 weeks, once every 7 weeks for the KTD4400 chip and once every 6 weeks for the KTD8300 memory chip.

²²Lach (2002) examined the Israeli markets for chicken, coffee, flour and refrigerators during 48 months. The median duration of a store's ranking in a given quartile ranged from 1 month for coffee and chicken to 2 to 3 months for flour and refrigerators; in that period most of the firms were seen quoting prices in all quartiles of the price distribution.

²³For the KTDINSP8200 memory chip 11 out of 22 stores were in a different quartile, with an average difference in ranking of 3.35. For the KTD4400 chip 14 out of 23 stores were in a different quartile, with an average difference in ranking of 3.48. Finally, for the KTD8300 chip 11 out of 22 stores were in a different quartile, while the average difference in ranking is 3.45.

product at each store, i.e.,

$$\text{ACF} = \frac{\sum_{t=2}^T (p_t - p_{av})(p_{t-1} - p_{av})}{\sum_{t=1}^T (p_t - p_{av})^2},$$

where p_{av} denotes the store’s average price for the product. The results are summarized in Table 6. Although the number of observations for each store-product pair is too small for the autocorrelations to be estimated precisely, this evidence suggests that serial correlation is not a serious issue in our data set.²⁴

	KTT3614	KTDINSP8200	KTD4400	KTD8300
Mean ACF (Std)	0.41 (0.19)	0.42 (0.15)	0.05 (0.31)	-0.02 (0.34)
Min. ACF	-0.06	-0.13	-0.50	-0.48
Max. ACF	0.58	0.59	0.46	0.46
Number of stores included	16	17	10	8

Notes: Store-product pairs for which we had fewer than eight observations are excluded, as well as pairs for which we observed no variation over time. If the autocorrelation is within $\pm 2/\sqrt{T}$, where $T = 8$ is the number of observations over time, it is not significantly different from zero at (approximately) the 5% significance level. This turns out to be the case for all individual autocorrelations.

Table 6: Summary autocorrelations store-product pairs

To check whether absence of demand and supply shocks is a reasonable assumption in our data set we tested the null hypothesis that price distributions in two different periods were equal using two-sample Kolmogorov-Smirnov tests. The results indicate that for the KTD4400 and KTD8300 memory chips, the null hypothesis that the distributions are the same cannot be rejected for any possible pair of periods, at a 5% significance level. For the other two memory chips, the KTT3614 and the KTDINSP8200, the null hypothesis was rejected only for pairs of periods that included the last period, which suggests that for these memory chips the last period is somewhat different than earlier periods.

The prices used for our estimations include neither shipping costs nor sales taxes. One reason for not including shipping costs in the main analysis is that we do not have the data for all the stores.²⁵ Another reason is that shipping costs and sales taxes depend on the

²⁴In a recent study of retail price variation, Hosken and Reiffen (2004) find that prices of most grocery products are at their annual mode more than 55% of the time and that temporary discounts account for 20% to 50% of the annual variation in retail prices, which suggests a large degree of serial correlation in their data set. See also Pesendorfer (2002) for a related finding.

²⁵Actually stores may choose to report blank in the shipping and handling cost field of the price feed form. As a result, shopper.com reports “See Site” in the shipping and handling column for that particular store.

state in which the consumer lives, which makes it difficult to compare total prices. In spite of these considerations, for robustness purposes, we also estimated the model neglecting sales taxes but using the shipping costs as if we were living in New York. Since a store not providing shipping cost information cannot be considered to ship for free (otherwise they would announce it as a promotional strategy), either we visited the web sites to discover shipping costs or we attributed average shipping costs to the missing values. The qualitative nature of the results did not change in these two cases.²⁶

Some of the variation in prices may be due to store differentiation. Consumers might view some stores more appealing than others and base this view on observable store characteristics like firm reputation, return policies, stock availability, order fulfillment, payment methods, etc. Unfortunately, we do not have information on all these indicators. But we do have information on whether the item was in stock or not, on whether firms disclosed shipping cost on shopper.com or not, and on the CNET certified ranking of a store, which is a store quality index computed by CNET on the basis of consumer feedback. To see the impact of these (observable) variables on the prices of each memory chip in our data set, we estimated the following model:

$$PRICE_{jt} = \beta_0 + \beta_1 \cdot RATING_{jt} + \beta_2 \cdot SHIP_{jt} + \beta_3 \cdot STOCK_{jt} + \varepsilon_{jt}, \quad (13)$$

where, for each product, $PRICE_{jt}$ is the list price of store j in week t , $RATING_{jt}$ is the CNET certified ranking of store j in week t , $SHIP_{jt}$ is a dummy for whether shop j disclosed shipping cost in week t , and $STOCK_{jt}$ is a dummy for whether shop j had the item in stock in week t . We estimated equation (13) by OLS. The resulting R-squared values indicate that only between 6% and 17% of the total variation in prices can be attributed to observable differences in store characteristics.²⁷ This suggests that the rest of the price variation can be due to strategic price setting in the presence of search costs or to unobserved

²⁶Tables containing the estimates using the data including shipping costs, as well as plots of the resulting search cost distributions and fitted price cdf's can be obtained from the authors upon request.

²⁷For all memory chips, the OLS estimates of the coefficient of $SHIP_{jt}$ are negative and highly significant. The estimates of the coefficient of $RATING_{jt}$ are positive and significant at a 1% level for the KTDINSP8200 chip, significant at a 10% level for the KTT3614 and KTD4400 chips, and not significant for the KTD8300 chip. The coefficient of $STOCK_{jt}$ was not significant for any of the products, but this could be due to the lack of variation of this variable in our data (upon reporting on shopper.com, almost all stores had the product in stock).

firm heterogeneity.

The finding that quite a few stores do change their price often and also that store rankings change from week to week gives an indication that store heterogeneity cannot be the only factor in explaining price setting behavior. To check to what extent unobserved heterogeneity across shops (e.g. based on brand recognition, or on marginal cost) plays an important role in explaining price setting behavior in our data set, we regressed prices on a constant and a set of store dummies. In this case the R-squared was very high, ranging from 0.93 to 0.99. Given the short period of data collection and given the fact that within this 8 week period quite a few firms either did not change their price at all, or changed it only once, these high R-squared values are not very surprising. Still, a caveat of the current model is that it cannot control for unobserved firm heterogeneity and therefore it treats all variation in the price data as variation due to search frictions.

5.2 Estimation results

The estimation results for the four different memory chips are presented in Table 7. An interesting observation is that even though the products differ in their characteristics, the estimates are quite similar across memory chips. This suggests that the consumers acquiring these products have similar search cost distributions.

	KTT3614	KTDINSP8200	KTD4400	KTD8300
\underline{p}	115.00	109.20	96.00	102.00
v	208.90	200.50	170.50	182.50
N	25	24	24	23
M	179	174	174	162
q_1	0.22 (0.05)	0.29 (0.04)	0.24 (0.05)	0.30 (0.04)
q_2	0.39 (0.15)	0.58 (0.02)	0.68 (0.01)	0.66 (0.03)
q_3	0.31 (0.14)	0.00	0.00	0.00
q_4	0.00	0.00	0.00	0.00
\vdots	\vdots	\vdots	\vdots	\vdots
q_{N-1}	0.00	0.00	0.00	0.00
q_N	0.08 (0.07)	0.13 (0.05)	0.09 (0.05)	0.04 (0.01)
r	109.69 (1.43)	103.15 (0.84)	90.91 (1.16)	90.55 (1.91)
LL	715.42	677.81	644.64	616.39
KS	1.07	1.01	1.11	1.26

Notes: Estimated standard errors in parenthesis.

Table 7: Estimation results

The estimates of the share of consumers who search once, q_1 , range from 22% to 30%

and are all highly significant.²⁸ These consumers do not compare prices and thus confer monopoly power to the firms. Firms compete for the rest of the consumers, who happen to search for 2 or 3 prices or for all the prices in the market. In particular, the estimates of q_2 range from 39% for the KTT3614 memory chip to 68% for the KTD4400 chip and are highly significant as well. The KTT3614 chip has also a sizable share of consumers comparing three prices, about 31%. For all the products, the estimates of parameters q_4 till q_{N-1} are approximately zero. Finally, the estimates of the fraction of consumers comparing all the prices in the market, q_N , range from 4% to 13% and are, except for the KTD3614 memory chip, significant at a 5% level.

These results suggest a clear picture of consumers' search costs. The entire consumer population can roughly be grouped into three subsets: buyers who do not search, buyers who compare at most three prices and buyers who compare all the prices in the market. This is consistent with the view that consumers have either quite high search costs or quite low search costs.

	KTT3614	KTDINSP8200	KTD4400	KTD8300
Δ_1	12.26 (1.42)	10.94 (0.28)	8.34 (0.47)	9.76 (0.28)
Δ_2	4.41 (1.19)	4.25 (0.27)	2.91 (0.32)	3.49 (0.17)
Δ_3	2.21 (0.81)	2.37 (0.20)	1.50 (0.20)	1.78 (0.10)
Δ_4	1.34 (0.58)	1.59 (0.16)	0.96 (0.14)	1.10 (0.07)
Δ_5	0.92 (0.44)	1.19 (0.13)	0.69 (0.11)	0.77 (0.05)
Δ_6	0.68 (0.35)	0.94 (0.11)	0.54 (0.09)	0.58 (0.04)
Δ_7	0.53 (0.29)	0.78 (0.09)	0.44 (0.07)	0.46 (0.03)
Δ_8	0.43 (0.24)	0.67 (0.08)	0.37 (0.06)	0.38 (0.02)
Δ_9	0.36 (0.21)	0.58 (0.07)	0.32 (0.06)	0.32 (0.02)
Δ_{10}	0.31 (0.18)	0.51 (0.06)	0.28 (0.05)	0.28 (0.02)
Δ_{11}	0.27 (0.16)	0.45 (0.06)	0.24 (0.04)	0.24 (0.02)
Δ_{12}	0.24 (0.15)	0.41 (0.05)	0.22 (0.04)	0.22 (0.01)
Δ_{13}	0.21 (0.13)	0.37 (0.05)	0.20 (0.04)	0.19 (0.01)
Δ_{14}	0.19 (0.12)	0.33 (0.04)	0.18 (0.03)	0.17 (0.01)
Δ_{15}	0.17 (0.11)	0.30 (0.04)	0.16 (0.03)	0.16 (0.01)
Δ_{16}	0.16 (0.10)	0.28 (0.04)	0.15 (0.03)	0.15 (0.01)
Δ_{17}	0.14 (0.09)	0.26 (0.03)	0.14 (0.03)	0.13 (0.01)
Δ_{18}	0.13 (0.08)	0.24 (0.03)	0.13 (0.02)	0.12 (0.01)
Δ_{19}	0.12 (0.08)	0.22 (0.03)	0.12 (0.02)	0.11 (0.01)
Δ_{20}	0.11 (0.07)	0.20 (0.03)	0.11 (0.02)	0.11 (0.01)
Δ_{21}	0.11 (0.07)	0.19 (0.02)	0.10 (0.02)	0.10 (0.01)
Δ_{22}	0.10 (0.06)	0.18 (0.02)	0.10 (0.02)	0.09 (0.01)
Δ_{23}	0.09 (0.06)	0.17 (0.02)	0.09 (0.02)	-
Δ_{24}	0.09 (0.06)	-	-	-

Notes: Estimated standard errors in parenthesis.

Table 8: Estimated critical search cost values

²⁸To be able to calculate the standard errors, we deleted the columns and rows of the Hessian for which the corresponding parameter estimates were zero.

The estimated cutoff points of the search cost distribution, Δ_i , with corresponding standard errors are presented in Table 8. All the cutoff points are highly significant and notice again that there is very little variation in the estimates across products. The estimated critical search cost values in combination with the estimated shares of consumers searching i times allow us to construct estimates of the search cost distributions underlying firm and consumer behavior.

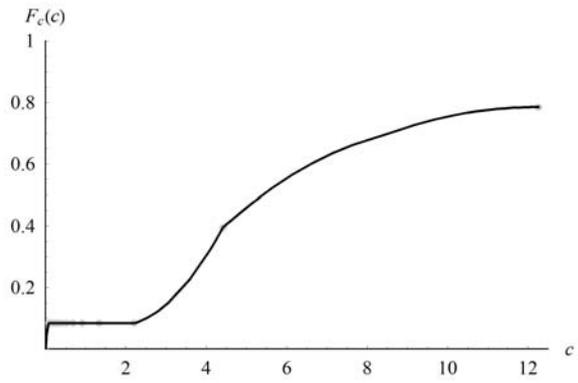
Figure 6 gives the estimated cumulative search cost distributions for the four memory chips. For example, for the KTT3614 memory chip we see that around 22% of the consumers have search costs higher than 12.26 US dollars; these costs are so high that these consumers only search once in equilibrium. Around 70% of the consumers have search costs in between 2.21 and 12.26 US dollars and for these consumers it is worth to search 2 or 3 times. Finally, around 8% of the buyers have search costs that are at most 9 dollar cents; these costs are so low that these buyers check the prices of all vendors. In sum, these estimates imply that typical on-line consumers have either very high search costs or very low search costs.

In spite of having more than 20 stores operating in each of the markets, we observe that market power is substantial. The estimates of r indicate that unit costs are between 50% and 53% of the value of the product so the average price-cost margins range between 23% and 28%.²⁹ This is of course the consequence of search costs, suggesting that demand side characteristics might be even more important than supply side ones to assess market competitiveness (Waterson, 2003).

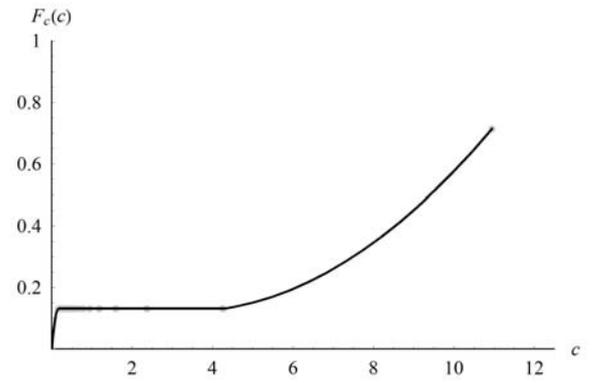
We finally test the goodness of fit of the model. To see how well the estimated price density function fits the data, we use the Kolmogorov-Smirnov test (KS-test) to compare the actual distribution to the fitted distribution. The KS-test is based on the maximum difference between the empirical cdf and the hypothesized estimated cdf. The null hypothesis for this test is that they have the same distribution, the alternative hypothesis is that they have different distributions. As Table 7 shows, since all KS values are below the 95%-critical value of the KS-statistic, which is 1.36, for all four memory chips we cannot reject that the prices are drawn from the estimated price distribution.³⁰ The goodness of fit of the

²⁹These margins are similar to those found in the book industry (Clay et al., 2001).

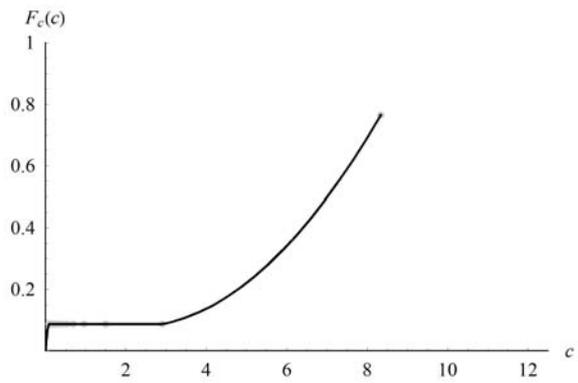
³⁰In this table KS is calculated as $\sqrt{M} \cdot \tau_M$, where M is the number of observations and τ_M is the maximum absolute difference over all prices between the estimated price cdf and the empirical price cdf. Because some of the parameters that enter the test are estimated, we also calculated the Rao-Robson Statistic, which is a



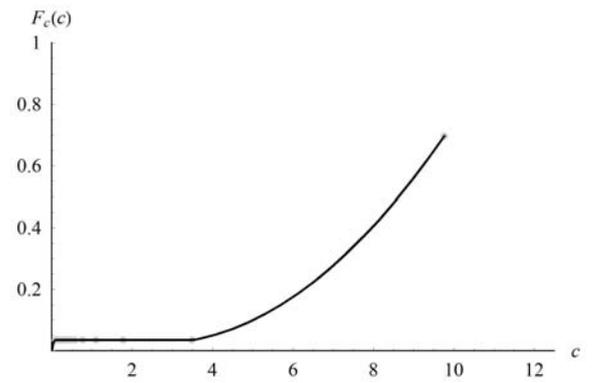
(a) KTT3614



(b) KTDINSP8200



(c) KTD4400



(d) KTD8300

Figure 6: Estimated search cost distributions

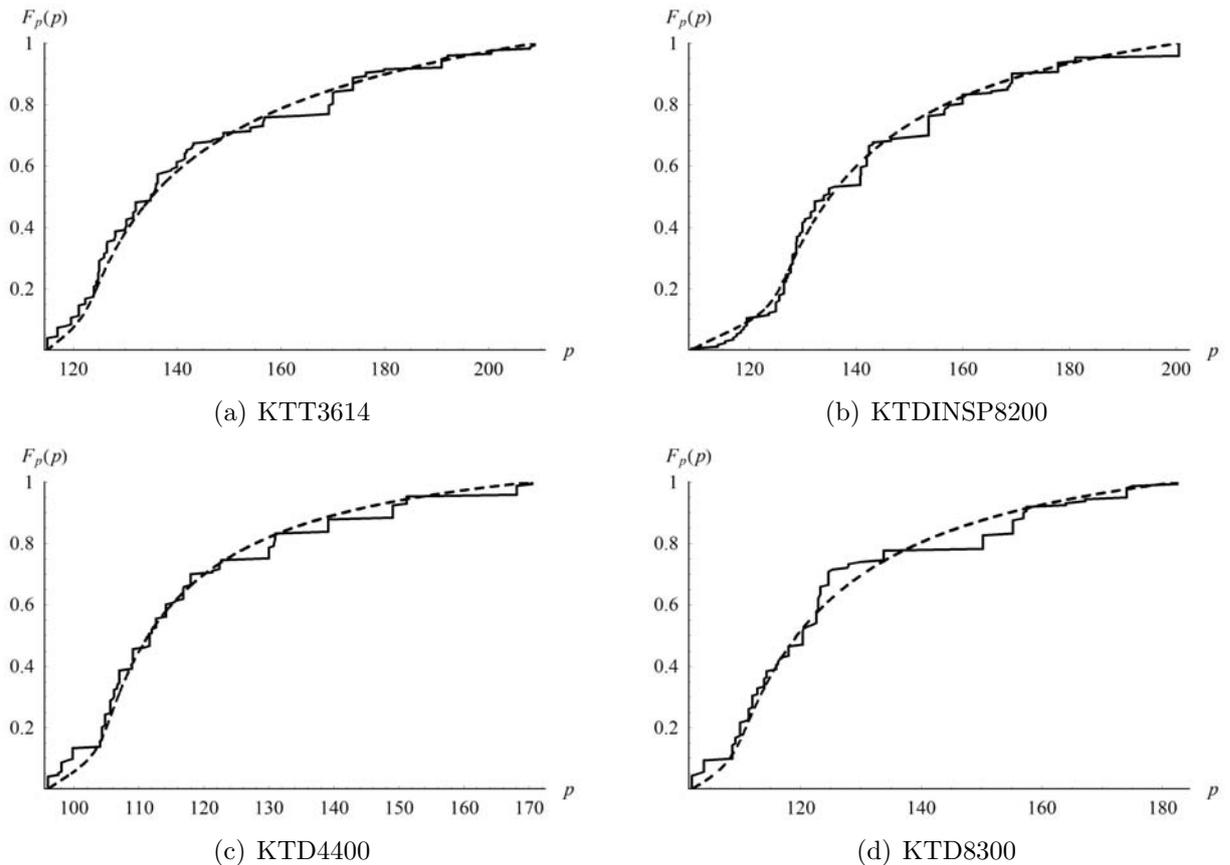


Figure 7: Estimated and empirical price distributions

model to the data can be visualized in Figure 7. A solid curve represents an empirical price distribution, while a dashed curve represents an estimated one.³¹

5.3 The effects of a sales tax

In consumer search markets the assessment of public policy may be difficult because policy changes affect not only firm pricing but also search intensity. In addition, since consumers pay different prices in equilibrium, it is not clear *a priori* how different consumers may be

kind of chi-squared test corrected for the uncertainty involved in estimating some of the parameters of the distribution that has to be fitted (for more details see Moore, 1986). The Rao-Robson statistics for two of the four products are below their corresponding critical values (KTT3614 and KTD4400), which means that for these products we cannot reject the null hypothesis that the estimated and empirical price cdf are the same.

³¹We also tried to estimate the model using the method of Hong and Shum (2006). Unfortunately, we were unable to obtain meaningful estimates. We encountered exactly the same problems as those reported in Section 4.3, i.e., the algorithm either did not move away from the starting values or did not converge. The reason is that the number of stores we observe in the data is quite high.

affected by a policy change. Since an economist can hardly observe search intensity directly, it seems difficult for him/her to come to a sensible conclusion as to how the market will perform after an intervention. The value of estimating a structural model of demand and supply is that the aggregate implications of policy changes can be assessed by computing what would be the after-policy equilibrium. To illustrate this feature, in this section we study the effects of a sales tax in the market for the KTDINSP8200 memory chip.

Denoting by t the *ad valorem* tax rate, a firm charging p receives a price $\hat{p} = (1 - t)p$. Therefore, in the presence of a sales tax, the equilibrium equation (4) is rewritten as

$$((1 - t)p - r) \left[\sum_{i=1}^N \frac{iq_i}{N} (1 - F_p(p))^{i-1} \right] = \frac{q_1((1 - t)\bar{p} - r)}{N}.$$

The upper bound of the price distribution continues to be equal to v , while the lower bound of the price cdf changes to

$$\underline{p} = \frac{1}{1 - t} \left(\frac{q_1((1 - t)\bar{p} - r)}{\sum_{i=1}^N iq_i} + r \right)$$

Using these equilibrium conditions, it is easy to see that if consumers did not change their search behavior the tax would result in a rightward shift of the price distribution. What is then interesting, is that a tax, by compressing the price distribution, lowers the incentives to search in the economy. This, in turn, gives the firms incentives to increase prices even further. Our next simulations show the final effect of a 5%, 10% and 15% sales tax.

Before moving to the results, we note that we first need to have a suitable (smooth) estimate of the search cost distribution. For this purpose, we fit a mixture of lognormals to the search cost points obtained in the estimation section (see Figure 6(b)). The fitted search cost density we obtain is

$$\hat{f}_c(c) = 0.36 \cdot \text{lognormal}(c, 2.43, 9.76) + 0.64 \cdot \text{lognormal}(c, 2.16, 0.24).$$

The fitted curve and the estimated points of the search cost distribution can be seen in Figure 8(a).

Using this estimate, we simulate the effects of the sales tax. The results are reported in Table 9; the original and the different post-tax equilibrium price distributions are drawn in Figure 8(b). As the graph shows, the introduction of a tax results in a rightward shift of the

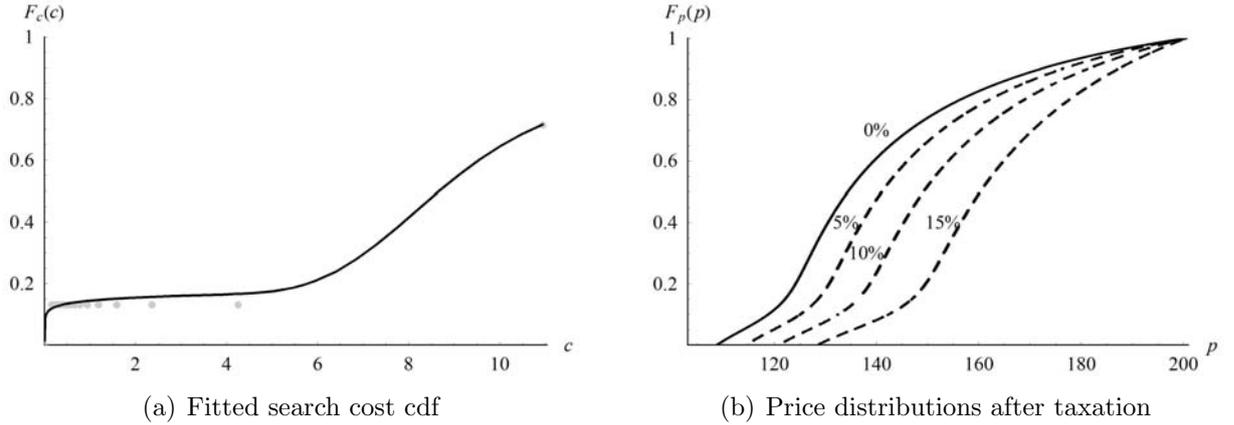


Figure 8: The effects of a sales tax

price distribution so all consumers end-up paying higher final prices. Inspection of columns 3 to 5 in Table 9 reveals that as the tax increases, the number of consumers who do not exercise price comparisons increases. For example, when the tax is 15% (last column of the table) this number is about 42%, which implies that firms can charge higher prices.

How much of the tax is passed on to consumers turns out to depend on the magnitude of the sales tax. A relatively small tax (5%) does not alter significantly the search profile in the economy and, though prices increase for all consumers, only about 5% is passed on to them; firms' profits after the tax are lower than in the case of no taxation (see columns 2 and 3 of Table 9). By contrast, a higher tax, for example 15%, ends-up reducing the gains from search substantially; in that case, the average price paid by the consumers who only search once increases by 16%, while the price the consumers who search thoroughly expect to pay increases by 18.7%. In sum, a consumer at random pays a price about 18% higher and this leads to after-tax firms' profits higher than in the absence of taxes (see columns 2 and 5 of Table 9).

6 Conclusions

Consumer search models have shown that the effects of an increase in the number of firms on the level and dispersion of market prices can depend upon the nature of search costs. Results of this kind imply that competition policy recommendations may depend on the nature of the search cost distribution and so there is a need to develop methods to identify

	KTDINSP8200 estimated	KTDINSP8200 fitted	KTDINSP8200 5% tax	KTDINSP8200 10% tax	KTDINSP8200 15% tax
p	109.20	108.82	114.48 (+5.2%)	120.92 (+11.1%)	128.49 (+15.3%)
t	0	0	0.05 (+5.0%)	0.10 (+10.0%)	0.15 (+15.0%)
v	200.50	200.50	200.50	200.50	200.50
N	24	24	24	24	24
q_1	0.29 (0.04)	0.27	0.30	0.34	0.42
q_2	0.58 (0.02)	0.56	0.53	0.49	0.42
q_3	0.00	0.01	0.01	0.01	0.01
q_4	0.00	0.01	0.01	0.01	0.00
\vdots	\vdots	\vdots	\vdots	\vdots	\vdots
q_{N-1}	0.00	0.00	0.00	0.00	0.00
q_N	0.13 (0.05)	0.12	0.12	0.12	0.12
r	103.15 (0.84)	103.15	103.15	103.15	103.15
Ep	140.96	140.17	146.08 (+4.2%)	153.20 (+9.3%)	162.65 (+16.0%)
$Ep_{1:2}$	130.01	129.01	135.37 (+4.9%)	143.03 (+10.9%)	153.25 (+18.8%)
$Ep_{1:N}$	113.54	112.71	118.59 (+5.2%)	125.40 (+11.3%)	133.77 (+18.7%)
$E\pi$ (net)	1.16	1.10	1.09 (-1.3%)	1.10 (-0.5%)	1.16 (+5.6%)

Notes: Column 1: estimated standard errors in parenthesis. Columns 3 to 5: percent changes relative to the fitted equilibrium in parenthesis.

Table 9: The effects of a sales tax: estimated and simulated parameter values

and quantify search costs.

Hong and Shum (2006) were the first to exploit the restrictions equilibrium search models place on the joint distribution of prices and search costs to structurally estimate unobserved search cost parameters. Following this research path, this paper has presented a new method to estimate search costs. Our method has three important features. First, we use a model with a finite number of firms, which helps separate variation in prices caused by changes in the number of competitors from variation in prices due to changes in search frictions. Second, our method yields *maximum likelihood* estimates of search cost parameters, which allows for standard asymptotic theory and hypothesis testing. Finally, our method is relatively easy to implement, in practice the estimation algorithm converges very rapidly and we have not observed numerical problems.

Using a data set of prices for four memory chips, we find that between 4% and 13% of consumers search for all prices in the market. These consumers have a search cost of at most 17 US dollar cents and obtain sizable gains relative to buying from one of the firms at random, namely, from 21 to 33 US dollars. Our estimates of the consumers' search cost density underlying the price observations for the memory chips suggest that consumers have either quite low or quite high search costs. Even though quite a few firms are active in the markets, search frictions confer significant market power. The estimates reveal that average

price-cost margins range from 23% to 28%. Finally, according to the Kolmogorov-Smirnov goodness-of-fit test, we cannot reject the null hypothesis that the price observations were drawn from the distribution functions specified by the theoretical search model.

In consumer search markets the assessment of public policy is difficult because policy changes affect not only firm pricing but also search intensity. The paper also illustrates how the structural method can be employed to simulate the effects of policy interventions. In particular, we study how the introduction of a sales tax would affect the market equilibrium for one of the memory chips. We find that a sales tax would reduce consumers' search intensity which would lead to some interesting equilibrium implications.

References

- [1] Baye, M., Morgan, J., Scholten, P., 2004. Price dispersion in the small and in the large: Evidence from an Internet price comparison site. *Journal of Industrial Economics* 52, 463-496.
- [2] Van den Berg, G.J., Ridder, G., 1998. An empirical equilibrium search model of the labor market. *Econometrica* 66, 1183-1221.
- [3] Brown, J.R., Goolsbee, A., 2002. Does the Internet make markets more competitive? Evidence from the life insurance industry. *Journal of Political Economy* 110, 481-507.
- [4] Burdett, K., Judd, K.L., 1983. Equilibrium price dispersion. *Econometrica* 51, 955-969.
- [5] Clay, K., Krishnan, R., Wolff, E., 2001. Prices and price dispersion on the web: Evidence from the on-line book industry. *Journal of Industrial Economics* 49, 521-539.
- [6] Dahlby, B., West, D.S., 1986. Price dispersion in an automobile insurance market. *Journal of Political Economy* 94, 418-438.
- [7] Delgado, J., Waterson, M., 2003. Tyre price dispersion across retail outlets in the UK. *Journal of Industrial Economics* 51, 491-509
- [8] Donald, S.G., Paarsch, H.J., 1993. Piecewise pseudo-maximum likelihood estimation in empirical models of auctions. *International Economic Review* 34, 121-148.
- [9] Eckstein, Z., Wolpin, K.I., 1990. Estimating an equilibrium search model using longitudinal data on individuals. *Econometrica* 58, 783-808.
- [10] Eckstein, Z., Van den Berg, G.J., 2007. Empirical labor search: A survey. *Journal of Econometrics* 136, 531-564.
- [11] Ellison, G., Ellison, S.F., 2005. Search, obfuscation, and price elasticities on the Internet. Mimeo.
- [12] Gautier, P.A., Moraga-González, J.L., Wolthoff, R., 2007. Structural estimation of job search intensity: Do non-employed workers search enough? Mimeo.

- [13] Greene, W.H., 1997. *Econometric Analysis*, third edition. Prentice Hall, New Jersey.
- [14] Hong, H., Shum, M., 2006. Using price distributions to estimate search costs. *Rand Journal of Economics* 37, 257-275.
- [15] Hortaçsu, A., Syverson, C., 2004. 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, 403-456.
- [16] Hosken, D., Reiffen, D., 2004. Patterns of retail price variation. *RAND Journal of Economics* 35, 128-146.
- [17] Janssen, M.C.W., Moraga-González, J.L., 2004. Strategic pricing, consumer search and the number of firms. *Review of Economic Studies* 71, 1089-1118.
- [18] Johnson, E.J., Moe, W.M., Fader, P.S., Bellman, S., Lohse, G.L., 2004. On the depth and dynamics of on-line search behavior. *Management Science* 50, 299-308.
- [19] Kiefer, N.M., Neumann, G.R., 1993. Wage dispersion with homogeneity: The empirical equilibrium search model. In: Bunzel, H., Jensen, P., Westergård-Nielsen, N. (Eds.), *Panel Data and Labour Market Dynamics*. North-Holland, Amsterdam.
- [20] Lach, S., 2002. Existence and persistence of price dispersion: An empirical analysis. *Review of Economics and Statistics* 84, 433-444.
- [21] Moore, D.S., 1986. Tests of chi-squared type. In: D'Agostino, R.B., Stephens, M.A. (Eds.), *Goodness-of-Fit Techniques*. Marcel Dekker, New York.
- [22] Morgan, P., Manning, R., 1985. Optimal search. *Econometrica* 53, 923-944.
- [23] Owen, A.B., 1990. Empirical likelihood confidence regions. *Annals of Statistics* 18, 90-120.
- [24] Pesendorfer, M., 2002. Retail sales: A study of pricing behavior in supermarkets. *Journal of Business* 75, 33-66.
- [25] Pratt, J.W., Wise, D.A., Zeckhauser, R.J., 1979. Price differences in almost competitive markets. *Quarterly Journal of Economics* 93, 189-211.

- [26] Qin, J., Lawless, J., 1994. Empirical likelihood and general estimating equations. *Annals of Statistics* 22, 300-325.
- [27] Rob, R., 1985. Equilibrium price distributions. *Review of Economic Studies* 52, 487-504.
- [28] Stahl, D.O., 1989. Oligopolistic pricing with sequential consumer search. *American Economic Review* 79, 700-712.
- [29] Sorensen, A.T., 2000. Equilibrium price dispersion in retail markets for prescription drugs. *Journal of Political Economy* 108, 833-850.
- [30] Sorensen, A.T., 2001. An empirical model of heterogeneous consumer search for retail prescription drugs. NBER Working Paper 8548.
- [31] Stigler, G., 1961. The economics of information. *Journal of Political Economy* 69, 213-225.
- [32] Varian, H.R., 1980. A model of sales. *American Economic Review* 70, 651-659.
- [33] Waterson, M., 2003. The role of consumers in competition and competition policy. *International Journal of Industrial Organization* 21, 129-150.