

Consumer Search and Oligopolistic Pricing:
a theoretical and empirical inquiry

ISBN 978 90 5170 556 0

Cover design: Crasborn Graphic Designers bno, Valkenburg a.d. Geul

This book is no. **408** of the Tinbergen Institute Research Series, established through cooperation between Thela Thesis and the Tinbergen Institute. A list of books which already appeared in the series can be found in the back.

Consumer Search and Oligopolistic Pricing: a theoretical and empirical inquiry

Zoekgedrag van consumenten en oligopolistische
prijsvorming:
een theoretische en empirische analyse

Proefschrift

ter verkrijging van de graad van doctor aan de
Erasmus Universiteit Rotterdam
op gezag van de
rector magnificus

Prof.dr. S.W.J. Lamberts

en volgens besluit van het College voor Promoties

De openbare verdediging zal plaatsvinden op
donderdag 21 juni 2007 om 13:30 uur

door

Matthijs Rudolf Wildenbeest

geboren te Alkmaar

Promotiecommissie

Promotoren: Prof.dr. J.L. Moraga-González
 Prof.dr. M.C.W. Janssen

Overige leden: Prof.dr. A. Hortaçsu
 Prof.dr. O.H. Swank
 Prof.dr. M. Waterson

Acknowledgement

This project started about five years ago by José Luis Moraga-González and Maarten Janssen hiring me as a student assistant to collect prices on the Internet. The data turned out to be very interesting, one thing led to another, and, after almost four years of work, it finally led to the completion of this thesis. As my main supervisor, José Luis had a great impact on this process with his valuable guidance, advice, comments and enthusiasm. He contributed to virtually all chapters of this thesis and taught me many things about economics and academia. We have been working together intensively, which was not only very rewarding, but also a lot fun. I could not have wished for a better supervisor. Maarten coauthored the fourth chapter of this thesis and had very useful remarks on other topics of this thesis as well. I am very much grateful to him for this. Another thanks goes to Zsolt Sándor. Zsolt got involved into the project, and coauthored the third chapter. I have learned a lot from his insights, and I very much enjoyed working with him. In addition I benefited from the help of many colleagues. Among them I mention Eddy Bekkers, Mariëlle Non, and Chris Wilson, who all gave comments on earlier versions of several chapters of my thesis.

I have very much enjoyed the time spent at the Tinbergen Institute, whether located on the 17th, 16th, or 9th floor of the H-building. I thank the staff of the Tinbergen Institute and my fellow PhD students for this. I very much enjoyed the atmosphere, the lunches, and the indoor soccer matches with our team H9 (formerly known as TI Knudde). Thanks to Bram, Chen, Diana, Eelco, Francesco, Gabriella, Haikun, Harris, Jan-Frederik, Josse, Karim, Marco, Nuno, Philippe, René, Robin, Sandra, Silvia, Simon, Stefano, Ward, and many others. Although I started (and ended) without roommates, over the years I have shared offices with Mariëlle, Michiel and Amy. Thanks to all three for the very good time. During the final phase of the completion of this thesis I was located at the department of economics on the 8th floor of the H-building. Thanks to all my colleagues there.

I am very much grateful to Ali Hortaçsu, Otto Swank and Mike Waterson for evaluating this thesis and for being part of the inner doctoral committee. In addition, I thank Saul Lach and Stefan Stremersch for their willingness to be part of the plenary doctoral committee. During my time as a PhD student I had the opportunity to present my work at several conferences abroad and to stay for a few months at Universidad Carlos III de Madrid. Thanks to Jordi Jau-mandreu for making the visit to Madrid possible. Financial support from Universidad Carlos III de Madrid, Stichting Dr. Hendrik Muller's Vaderlandsch Fonds, the Netherlands Organisation for Scientific Research (NWO), and the Trustfonds Erasmus Universiteit is gratefully acknowledged.

Finally, I wish to express my greatest thanks to my family and friends, who have supported me, especially to Phuong, for her unconditional support, patience and understanding of my frequent travels.

Rotterdam, May 2007

Contents

| | | |
|----------|---|-----------|
| 1 | Introduction | 1 |
| 1.1 | Price dispersion | 2 |
| 1.2 | Consumer search literature | 4 |
| 1.3 | Structure of this thesis | 7 |
| 2 | Estimation of a nonsequential search model | 11 |
| 2.1 | Introduction | 12 |
| 2.2 | The consumer search model | 15 |
| 2.3 | Maximum likelihood estimation | 17 |
| 2.4 | A Monte Carlo study | 19 |
| 2.4.1 | Performance of the estimates | 20 |
| 2.4.2 | Measurement error in the number of firms | 25 |
| 2.4.3 | Comparison of the ML estimation method and Hong and Shum (2006) method | 26 |
| 2.5 | Empirical application | 30 |
| 2.5.1 | Data and empirical issues | 30 |
| 2.5.2 | Estimation results | 37 |
| 2.5.3 | The effects of a sales tax | 40 |
| 2.6 | Conclusions | 43 |
| 3 | Estimation and identification of nonsequential search models | 45 |
| 3.1 | Introduction | 46 |
| 3.2 | The model | 48 |
| 3.3 | Theoretical analysis | 49 |

| | | |
|----------|--|------------|
| 3.4 | Mergers | 57 |
| 3.5 | Estimation | 61 |
| 3.5.1 | The simple model: heterogeneous search costs, homogeneous valuations | 62 |
| 3.5.2 | The general model: heterogeneous search costs and valuations | 73 |
| 3.6 | Conclusions | 75 |
| 4 | Truly costly sequential search and oligopolistic pricing | 83 |
| 4.1 | Introduction | 84 |
| 4.2 | The model | 86 |
| 4.3 | Equilibrium analysis | 87 |
| 4.4 | Comparative statics | 92 |
| 4.5 | Estimation | 99 |
| 4.6 | Empirical application | 101 |
| 4.7 | Conclusions | 103 |
| 5 | Estimation of a sequential search model | 105 |
| 5.1 | Introduction | 106 |
| 5.2 | The model | 108 |
| 5.3 | Equilibrium | 109 |
| 5.4 | Identification and estimation | 111 |
| 5.5 | Empirical application | 115 |
| 5.6 | Conclusions | 119 |
| 6 | Vertical product differentiation and search | 121 |
| 6.1 | Introduction | 122 |
| 6.2 | The model | 125 |
| 6.3 | Estimation | 130 |
| 6.4 | Empirical analysis | 132 |
| 6.4.1 | Description of the data | 135 |
| 6.4.2 | Estimation of search costs without time fixed effects | 140 |
| 6.4.3 | Estimation of search costs with time fixed effects | 143 |
| 6.4.4 | Comparison across periods and discussion of results | 148 |

| | | |
|----------|--|------------|
| 6.5 | Conclusions | 150 |
| 7 | Summary, achievements and directions for further research | 153 |
| 7.1 | Summary | 154 |
| 7.2 | Achievements | 156 |
| 7.3 | Directions for further research | 158 |
| | Samenvatting (Summary in Dutch) | 161 |
| | References | 167 |

1

Introduction

1.1 Price dispersion

According to the classical paper by Bertrand (1883) in a world with perfect information and without transaction costs, the price for a homogeneous good sold by firms with identical costs and no capacity constraints will converge to the Walrasian price, even if there is a finite number of firms. As explained by Bertrand (1883), if firms set prices, each firm can simply conquer the whole market by setting a slightly lower price than the prices of the rivals. As a result each rationally behaving firm will try to undercut the price of its competitors. This process will continue until the price is equal to marginal cost as a further lowering of the price would yield negative profits. The unique Nash Equilibrium is such that each firm sets a price equal to marginal cost.

The ease with which one can find two stores that charge different prices for similar products shows that this ‘law of one price’ does not hold in general. Instead, price dispersion seems to be the rule in real-world markets. Over the years economists have come up with several explanations for these observed differences in prices. A first explanation is that instead of what is assumed by Bertrand (1883) in reality it is often the case that there is spatial differentiation or that goods differ in terms of their characteristics. If consumers have different tastes for these characteristics, goods are no longer perfect substitutes and this will give firms market power. As a result there are less incentives for firms to undercut each other and goods might be offered at different prices.

What is intriguing is that also in markets where the goods and the firms are seemingly identical, price dispersion is observed. If we purely focus on the firm side, in such a situation there are clear incentives for firms to undercut their competitors. However, if instead of what is assumed by Bertrand (1883) consumers have imperfect information, the incentives to undercut each other weaken. In recent empirical work this line of thought is being followed to explain price differences for seemingly homogeneous goods. The idea is that, due to search frictions some consumers do not compare prices, which allows firms to set a higher price than in a frictionless economy. Firms might then set different prices to either maximize surplus from consumers who do not compare prices, by setting a high price, or to maximize surplus from price comparing consumers, by setting a low price. As a result, price dispersion emerges.

Notice that if goods are identical, but the firms are different in their marginal costs, price

dispersion might also emerge, but only in combination with search frictions. The idea is that without search frictions firms with higher marginal cost than their competitors cannot survive, because they can never set prices as low as their competitors without making a loss. If the goods are identical, consumers will simply buy from the lowest priced firm, which is the firm with the lowest marginal cost. However, if marginal cost differences coexist with search frictions also the firms with higher marginal cost might survive, by catering the consumers with high search costs.

Although product differentiation, uncertainty in demand and capacity constraints might also act as sources of price dispersion, this thesis will focus on search frictions as an important determinant of price variation. Nevertheless, the framework will be extended to include vertical product differentiation as well in the last essay of this book. Of course, this does not mean that uncertainty in demand and capacity constraints are not important sources of price dispersion. In some markets capacity constraints will be an important contributor or maybe even the most important source of price dispersion. This means that the models described in this thesis will be less applicable to those markets.

Existing consumer search models have showed that the way search frictions influence markets may depend on how search costs are distributed among consumers. For example, in the sequential search model of Stahl (1989), if search costs are high mean prices remain constant if a merger occurs, while otherwise mean prices will fall.¹ In addition, in the nonsequential search model of Janssen and Moraga-González (2004) whether a merger leads to higher or to lower mean prices depends on the height of search costs. More in general, this implies that for a sound assessment of the price and welfare effects of competition policy measures such as the increase of a sales tax, the softening of barriers to entry or the imposition of a price cap, one needs to have knowledge about supply and demand parameters, including the search cost distribution. For this, one could simply assume some search cost distribution, but given the dependence of the results on the exact shape of the search cost distribution, without having some good estimates this exercise might lead to misleading conclusions.

The purpose of this thesis is to come up with methods to estimate consumer search models, using a limited amount of data. This thesis is part of a relatively new strain of the consumer search literature that uses the structure of search models to identify and estimate search cost

¹See also Chapter 4 of this thesis.

distributions. It builds on existing theoretical and empirical work; the next section discusses this literature, focussing on the most relevant theoretical and empirical contributions.²

1.2 Consumer search literature

Starting from the seminal article of Stigler (1961), economists have put much effort in trying to understand the nature of price dispersion. In the theoretical literature the initial goal was to find a rationale for price dispersion as an equilibrium outcome. Although Stigler (1961) succeeded in giving this rationale using a search-theoretic approach, his model was criticized because it focussed only on the consumer side. The distribution of prices was taken as exogenous and not based on optimal firm price setting behavior. Diamond (1971) showed that when firms are also optimizing, positive search costs lead to a unique equilibrium where all firms charge the monopoly price. This result became known as the Diamond Paradox because even though there is a continuum of identical firms in his model, search frictions lead to monopoly pricing and not to the competitive outcome. From that moment on the theoretical literature was concerned with finding ways to overcome this striking result.

One of the first papers that shows that price dispersion can arise as an equilibrium outcome with both firms and consumers optimizing is Reinganum (1979). In this paper there is a continuum of firms with heterogeneous costs, selling a good to consumers who search sequentially for prices. Reinganum (1979) proves that an equilibrium with price dispersion arises as long as consumers hold elastic demands. Although consumers have to search for prices, because in equilibrium each firm will set a constant markup over its marginal cost, the source of price dispersion is marginal cost heterogeneity.

In a model with homogeneous firms, Stahl (1989) shows price dispersion can also be a result of search cost heterogeneity only. As in Reinganum (1979) consumers search sequentially, but in Stahl (1989) some consumers search costlessly, while others have positive search costs. The equilibrium is in mixed strategies: firms try to attract both types of consumers by randomizing their prices. Moreover, consumers with positive search costs do not search.

Burdett and Judd (1983) show that in order to have price dispersion as an equilibrium outcome, it is not necessary to have any ex-ante heterogeneity whatsoever. Burdett and Judd (1983)

²For a more extended overview of the existing literature on consumer search, see Baye et al. (2006).

assume that consumers search nonsequentially, that is, consumers determine before they start searching how many times to search. As in Stahl (1989), depending on the height of the search cost of consumers, it might be optimal for firms to randomize their prices. For consumers it might then be optimal to randomize between searching once and twice.

More recent theory papers focus on the relation between market characteristics, prices and price dispersion. For example, Janssen and Moraga-González (2004) study an oligopolistic version of Burdett and Judd (1983) with a two-type distribution of consumers. Comparative statics with respect to the number of firms show that equilibrium prices, price dispersion and welfare depend on the height of the search cost of consumers.

The empirical work focussed initially on documenting to what extent price dispersion is observed in real world markets. Stigler (1961) already gave some examples on price dispersion; he found that the coefficient of variation for Chevrolets in the Chicago area was 1.72, while that for anthracite coal delivered in Washington was around 6.8. The more systematic approach of Pratt et al. (1979) showed that there was also substantial price dispersion for 39 randomly chosen products from the Yellow Pages in the Boston Area.

While the earlier empirical work focused mainly on documenting price dispersion, later empirical work tried to test some of the comparative statics results of search models directly. Although Stigler (1961) and Pratt et al. (1979) already identified a relation between price dispersion and the benefits of search by comparing price dispersion for cheap versus expensive items, Sorensen (2000) focused more directly on this relationship by studying price dispersion in the market for prescription drugs. He finds substantial price differences and that price dispersion is inversely related to purchase frequency. The reason for this is that the consumers' benefit per search is higher for frequently purchased drugs, leading to more search and lower price dispersion.

With the onset of e-commerce in the nineties, the focus of empirical work shifted to the relation between the cost of searching and price dispersion. The common hypothesis studied in the large number of papers that appeared since then is that search cost of consumers shopping online is lower than the cost of searching for consumers shopping in traditional markets. This would imply that prices and price dispersion should be lower in online markets than in traditional ones; however, in general, the empirical results are quite mixed. Some papers find higher price dispersion and prices online than in traditional stores, while others find that they are the

same, or find an opposite relation. For instance, concerning price levels, Clay et al. (2001) report that prices in online and traditional stores are similar. Brown and Goolsbee (2002) find that there is no evidence that the Internet reduced prices before price comparison sites emerged and proliferated. Finally, Bailey (1998) finds that online prices are higher than off-line ones for books and CDs. On price dispersion, the effect of moving markets online seems to be ambiguous empirically as well.

Theory search models can be divided in models that generate price dispersion as a result of pure strategy equilibria (e.g., Reinganum, 1979) or as a result of mixed strategy equilibria (e.g., Burdett and Judd, 1983; Janssen and Moraga-González, 2004; Stahl, 1989). Lach (2002) explicitly studies whether price dispersion is in line with mixed strategy equilibria by studying the persistence and the nature of price dispersion over time. If firms play a mixed strategy equilibrium in prices, over time prices should fluctuate in a random fashion. Indeed, after correcting for observed and unobserved firm heterogeneities, Lach (2002) finds prices in his data set to be in line with random pricing strategies.

So far, surprisingly little attention has been given to quantifying search costs in real world markets. This is an important omission, since it is known from the existing theory that the predictions often depend on the shape of the search cost distribution. As a matter of fact, the mixed results found in empirical papers comparing price dispersion online and off line could be explained by the non-monotonic relation between search cost and price dispersion often found in the theoretical literature on consumer search. In addition, to study the effects of competition policy measures such as the increase of a sales tax, the softening of barriers to entry or the imposition of a price cap, one needs to have knowledge about the search cost distribution. Of course one could simply impose some search cost distribution to study the effects of such measures, but given the sensitiveness of the results to the exact shape of the search cost distribution, it is worth to study methods that give researchers and practitioners good estimates of these distributions.

This thesis tries fill the gap in the empirical literature by providing methods to estimate and identify consumer search costs. So far, in addition to this thesis, only two papers deal with the estimation of consumer search models. Hong and Shum (2006) were the first to show how to estimate search costs using only price data. Hong and Shum (2006) estimate both a mixed strategy sequential search model, and a mixed strategy nonsequential search model. The

estimation of the nonsequential search model is done nonparametrically while the alternative model is estimated parametrically, although Hong and Shum (2006) show that if one is also in the possession of data on the marginal cost of the firms, the sequential search model can also be estimated nonparametrically. Hortaçsu and Syverson (2004) show how to estimate search costs in a pure strategy sequential search model that allows for vertical product differentiation, using both price and quantity data. Given the importance of these two papers to the analysis of this thesis, both papers will be discussed in more detail in later chapters of this thesis.

1.3 Structure of this thesis

This thesis consists of five essays on the structural estimation of consumer search models and can be roughly divided into three parts. In the first two essays, I explain how to model, identify and estimate nonsequential search models. The first essay presents a structural methodology to estimate search cost distributions. The starting point is an oligopolistic version of the nonsequential search model presented in Hong and Shum (2006). The oligopoly assumption is useful because it helps the researcher separate the variation in prices caused by variation in the number of firms from that caused by variation in search costs. Using the equilibrium conditions derived from the model, it is shown how to estimate the model by a maximum likelihood procedure. The method is applied to a data set of online prices for different memory chips. The estimates suggest that online consumers have either quite high or quite low search costs so they either search exhaustively in the market or very little, for at most three prices. Search frictions confer a significant amount of market power to the firms: despite that more than 20 firms operate in each of the markets we study, price-cost margins are around 25%. Kolmogorov-Smirnov goodness-of-fit tests suggest that the null hypothesis that the observed prices are generated by the model cannot be rejected. This chapter also illustrates how the structural methodology can be employed to simulate the effects of policy interventions.

The second essay goes more into the details of the identification of search costs using a structural estimation approach. The model presented in the first essay is extended to a more general framework where consumers not only differ in their costs of searching but also in their valuations. It is shown that the search cost distribution cannot be identified nonparametrically at all the points of its support when the econometrician observes prices from only one market.

It is shown that this problem can be solved by studying a richer framework where the econometrician has price data from several markets with the same search cost distribution. To exploit the common feature of the markets, it is proposed to estimate the search cost density function by a semi-nonparametric (SNP) density estimator whose parameters maximize the joint likelihood corresponding to all the markets. This essentially nonparametric approach can be used when all consumers have identical valuations or when valuations differ but are not correlated with search costs. When there is correlation between search costs and valuations, only parametric identification obtains.

In the two subsequent essays I focus on the estimation of sequential search models. In the third essay the paper of Stahl (1989) on oligopolistic pricing and sequential consumer search is modified by relaxing the assumption that consumers obtain the first price quotation for free. In the theory part of this chapter it is shown that when all price quotations are costly to obtain, the unique symmetric equilibrium need not involve full consumer participation. The region of parameters for which non-shoppers do not fully participate in the market becomes larger as the number of shoppers decreases and/or the number of firms increases. The comparative statics properties of this new type of equilibrium are interesting. In particular, expected price increases as search cost decreases and is constant in the number of shoppers and in the number of firms. Welfare falls as firms enter the market. It is shown that monopoly pricing never obtains with truly costly search. In the empirical part of this chapter, using the equilibrium conditions derived from the model, it is shown how to estimate the model. The estimates show that the model does not do very good in explaining observed pricing patterns. This is most likely related to the outcome of the model that in equilibrium consumers either observe all prices or do not search, which has a big impact on the flexibility of the model.

The fourth essay studies the identification and estimation of a more general sequential search model. In order to make the setting empirically more meaningful, instead of having a two-type distribution, consumers now draw their search cost from an atomless distribution with the positive real line as support. It is shown that the search cost distribution can be identified nonparametrically using only price data, even with data from only one market. The method is applied to the same data set of prices for memory chips which was used in the first essay of the thesis. A comparison of the estimates obtained from the two different models reveals that estimated search costs are much higher with sequential search than with nonsequential search.

Moreover, compared to the nonparametric estimation method of Hong and Shum (2006), search costs are much higher.

Finally, in the last essay of this thesis, I extend the framework to include vertical product differentiation as a source of price dispersion as well. In the unique symmetric equilibrium firms with different characteristics draw utilities from a common utility distribution. Because the firms differ in their characteristics, they use different price distributions. The model therefore provides a theoretical rationale for explaining price dispersion as a result of quality differences and search frictions together. Using the equilibrium conditions derived from the model, it is shown how to separate the two effects from each other and how to estimate the model. A data set on prices from Dutch supermarkets reveals that the amount of search has decreased over the sampling period. Moreover, ignoring vertical product differentiation results in an overestimation of search costs. Although used in a nonsequential search setting, the ideas can also be applied to sequential search models.

The five essays provide a variety of models on consumer search behavior and the effects of search frictions on pricing strategies of firms. Which model is the most appropriate to use depends on the characteristics of the market under study and is up to the choice of the econometrician. For example, the nonsequential search model might be more appropriate in markets where the search outcome is observed with some delay, like in markets for labor, mortgages, refurbishing services, etc. Also, the choice which model to estimate depends on what kind of data is available. For example, compared to Hortaçsu and Syverson (2004), for estimation of the vertical product differentiation model presented in the last essay of this thesis less data is needed. However, the use of only price data comes at a cost, since the model presented here is less general than the model presented in Hortaçsu and Syverson (2004). Therefore, if one has both price and quantity data a model like the one presented in Hortaçsu and Syverson (2004) might be the best choice.

Estimation of a nonsequential search model

NOTE: This chapter is based on Moraga-González and Wildenbeest (2006).

2.1 Introduction

As argued in Chapter 1, 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 the predictions of the various theoretical models are quite sensitive to the height of the search cost in the market.¹ Since competition policy recommendations may depend on the amount of search in an industry, there is a need to develop methods to quantify search costs.

This essay presents methods to estimate a nonsequential search model. It builds on Hong and Shum (2006), who present structural methodologies to retrieve information on search costs in markets for homogeneous goods. Hong and Shum (2006) 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.

The nonsequential search model studied by Hong and Shum (2006) generalizes Burdett and Judd (1983) seminal paper by introducing search cost heterogeneity. The market is operated by a continuum of firms, which compete by setting prices. Consumers with heterogeneous search costs search to discover prices and buy from the cheapest firm they observe. In equilibrium, some consumers do not search whereas others do and this 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 exploit equilibrium behavior to obtain a (potentially infinitely large) number of moment conditions. The estimates of the parameters of the price cdf give the height of the search cost distribution evaluated at a series of cut-off points. In the second step, these cut-off points are estimated directly from the empirical cdf of prices. This method is interesting but it requires to solve a highly dimensional optimization problem, which is computationally quite demanding. Indeed, in practice, only a few parameters of the price distribution can be estimated; without a priori good information about search costs, discarding parameters has the problem of introducing biases in the estimates.²

¹See e.g. Janssen and Moraga-González (2004) for the influence of the height of the search cost on equilibrium search intensity and market performance.

²For 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.

In this chapter we present an alternative way, by *maximum likelihood* (ML), to estimate an oligopoly version of the nonsequential search model of Burdett and Judd (1983). 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 cut-off 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 cut-off points of the search cost distribution are also ML so the asymptotic theory for computing standard errors of and for conducting tests of hypotheses on the estimates of the search cost distribution remains standard. In addition, our method is relatively easy to implement in practice and we have not observed any numerical difficulties.

The model we study is an oligopolistic version of Burdett and Judd (1983). Vis-à-vis the competitive case studied by Hong and Shum (2006), the oligopoly model has the advantage that it captures variation in prices that is due to variation in the number of competitors; this makes our model useful for the study of competition policy issues. 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, independently of the number of prices the econometrician actually observes, we can estimate the relevant number of parameters of the price distribution. In this way we learn about the distribution of search costs at all relevant quantiles. In particular, we find out how much search is conducted by consumers who search thoroughly and how low their search costs are.³

To estimate the parameters of the price distribution, all we need is to observe the prices firms charge 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, ignoring low search cost consumers leads to significant biases in the estimates: search costs are substantially overestimated and price-cost margins largely exaggerated. These biases result in

³Brown 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.

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. The search cost distribution consistent with these estimates implies that consumers have either quite high or quite low search costs. The estimates suggest that the search cost of consumers who search thoroughly in the market is at most 17 US dollar cents.⁴

Consumer 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 be even 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.

This chapter also illustrates how the structural methodology 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. We find that sales taxes may affect the equilibrium in non-trivial ways. As a matter of fact, we observe that the tax shifts the price cdf to the right and, depending on the height of the tax rate, this may

⁴Using a different methodology, Sorensen (2000) finds that between 5% to 10% of the consumers conduct an exhaustive search for prices in the market for prescription drugs.

change significantly the search profile in the economy. For example, a 15% sales tax reduces search in such a way that the tax ends up being passed on to the consumers more than fully and after-tax firms profits increase.

The structure of this chapter is as follows. In the next section, we review and modify the nonsequential consumer search model studied in the paper of Hong and Shum (2006). In Section 2.3 we discuss our maximum likelihood estimation method. Section 2.4 presents a Monte Carlo study that, among other issues, compares our estimation method with that of Hong and Shum. In Section 2.5 we estimate the search cost distribution underlying price data obtained from some online markets for memory chips; in addition, this section shows how the market would be affected if a sales tax was introduced. Finally, Section 2.6 concludes.

2.2 The consumer search model

We study an oligopolistic version of the model proposed in Hong and Shum (2006); their model generalizes the nonsequential consumer search model of Burdett and Judd (1983) by adding consumer search cost heterogeneity.⁵ Assume 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 inelastically demands one unit of the good. Let \bar{p} be the consumer valuation. 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 ic .

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 .⁶ Given firm behavior, the number of prices $i(c)$ a consumer with search cost c observes must be optimal, i.e.,

$$i(c) = \arg \min_{i \geq 1} c(i-1) + \int_{\underline{p}}^{\bar{p}} ip(1 - F_p(p))^{i-1} f_p(p) dp. \quad (2.1)$$

⁵The oligopoly case is also studied in Janssen and Moraga-González (2004) but with a two-point search cost distribution.

⁶It will become clear later that the upper bound of the price distribution must be equal to the consumer valuation.

Since $i(c)$ must be an integer, the problem in equation (2.1) 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

$$\Delta_i = Ep_{1:i} - Ep_{1:i+1}, i = 1, 2, \dots, N - 1 \quad (2.2)$$

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, the fractions of consumers q_i sampling i prices are simply

$$q_1 = 1 - F_c(\Delta_1); \quad (2.3a)$$

$$q_i = F_c(\Delta_{i-1}) - F_c(\Delta_i), i = 2, 3, \dots, N - 1; \quad (2.3b)$$

$$q_N = F_c(\Delta_{N-1}). \quad (2.3c)$$

Given consumer search behavior it is indeed optimal for firms to mix in prices. The upper bound of the price distribution must be \bar{p} because a firm which charges the upper bound sells only to the consumers who do not compare prices (consumers in q_1), who 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.,

$$(p - r) \left[\sum_{i=1}^N \frac{i q_i}{N} (1 - F_p(p))^{i-1} \right] = \frac{q_1(\bar{p} - r)}{N}. \quad (2.4)$$

From equation (2.4) it follows that the minimum price charged in the market is

$$\underline{p} = \frac{q_1(\bar{p} - r)}{\sum_{i=1}^N i q_i} + r. \quad (2.5)$$

As shown in Hong and Shum (2006), equations (2.2) to (2.5) 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, this will also be the focus of our next section.

2.3 Maximum likelihood estimation

Assume the researcher observes the prices of the N firms operating in the market.⁷ 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 (2.4) cannot be solved for the equilibrium price distribution F_p and this makes it difficult to calculate the cut-off points

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

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 cut-off 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.⁸ We first rewrite the cut-off points as (by integration by parts)

$$\Delta_i = \int_p^{\bar{p}} F_p(p)(1 - F_p(p))^i dp, \quad i = 1, 2, \dots, N-1. \quad (2.6)$$

Since the distribution function $F_p(p)$ is monotonically increasing in p , its inverse exists. Using equation (2.4) we can find the inverse function:

$$p(z) = \frac{q_1(\bar{p} - r)}{\sum_{i=1}^N i q_i (1 - z)^{i-1}} + r. \quad (2.7)$$

⁷In practice, sometimes not all the firms are observed by the researcher; our Monte Carlo study in Section 2.4 examines the implication of this lack of data.

⁸Note that for our asymptotic arguments we 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.

Using this inverse function, a change of variables in equation (2.6) yields:

$$\Delta_i = \int_0^1 p(z)[(i+1)z - 1](1-z)^{i-1} dz, \quad i = 1, 2, \dots, N-1. \quad (2.8)$$

If we obtain ML estimates of $r, \bar{p}, \underline{p}$ and $q_i, i = 1, 2, \dots, N$, then, by the invariance property (see Greene, 1997), we can use equations (2.7) and (2.8) to calculate ML estimates of the cut-off 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, which will often be the situation. Since the price density cannot be obtained in closed form, we apply the implicit function theorem to equation (2.4), which yields

$$f_p(p) = \frac{\sum_{i=1}^N i q_i (1 - F_p(p))^{i-1}}{(p - r) \sum_{i=1}^N i(i-1) q_i (1 - F_p(p))^{i-2}}. \quad (2.9)$$

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 take the minimum price in the sample p_1 and the maximum one p_M to estimate 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.⁹ Using the estimates of \underline{p} and \bar{p} , equation (2.5) can be solved to obtain the marginal cost r as a function of the other parameters:

$$r = \frac{p_1 \sum_{i=1}^N i q_i - q_1 p_M}{\sum_{i=2}^N i q_i}. \quad (2.10)$$

Plugging this formula into equation (2.9) and using the fact that $q_N = 1 - \sum_{i=1}^{N-1} q_i$ we can

⁹See also Donald and Paarsch (1993) on using order statistics to estimate the lower and upper bound of bid distributions.

solve numerically the following maximum likelihood estimation problem:¹⁰

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

where

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

We note that in this formulation the estimate of r is obtained from equation (2.10) as a function of the estimates of the other parameters. This procedure introduces some dependence between the price observations. Our Monte Carlo study next 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 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.

2.4 A Monte Carlo study

The study in this section has various purposes. First, we investigate how precise the maximum likelihood estimates of the price distribution and of the search cost distribution 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 impact of some measurement error in the number of firms N that operate in the market, since this may be a

¹⁰The 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 (2.4), which in turn allows us to calculate $f_p(p_\ell)$ using (2.9). We use a trust region PCG method, which proceeds by changing the q_i 's until the log-likelihood function is maximized.

problem in real-world applications. Finally, we compare our estimation method to that of Hong and Shum (2006).

2.4.1 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 $\sigma = 5$. Further, 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 (2.3a)-(2.3c) and (2.8) hold simultaneously. Next, we use these parameters to construct the equilibrium price distribution, implicitly defined by equation (2.4). 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 2.1 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%.

In equilibrium a great deal of the consumers, about 38%, search for only one price; another important group of buyers searchers 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.¹¹

As discussed above, for the estimation of the model we need to assume that market interaction evolves over a finite number of $T \geq 2$ periods. We take the equilibrium of the static game described in Section 2.2 as the equilibrium of the repeated game with finite horizon. Our first

¹¹This 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%).

| | 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) |
| \underline{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 2.1: True and estimated parameter values

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 2.1 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 does not represent 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.

| | 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 2.2: True and estimated critical search cost values

The first column of Table 2.2 reports the true cut-off 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 cut-off points are highly significant, and quite close to the true ones.

The price and search cost cdf's as well as their mean estimates are plotted in Figure 2.1(a) and 2.1(b) respectively. In both graphs the thick curves are the true distributions while the thick dashed curves show the mean of the 1000 estimated distributions. The dashed curves




Figure 2.1: Estimated price and search cost cdf's (# periods $T = 4$; # obs. $M = 100$)

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.¹²

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 very well 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 bound 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 2.2(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 estimate is much closer to the true distribution than the previous one. The average estimate of q_1 goes down 0.04 points and this

¹²The last cut-off point of the search cost distribution we can estimate is Δ_1 and therefore we do not have information about search costs beyond that point.

results in an upward shift of the estimated search cost cdf.

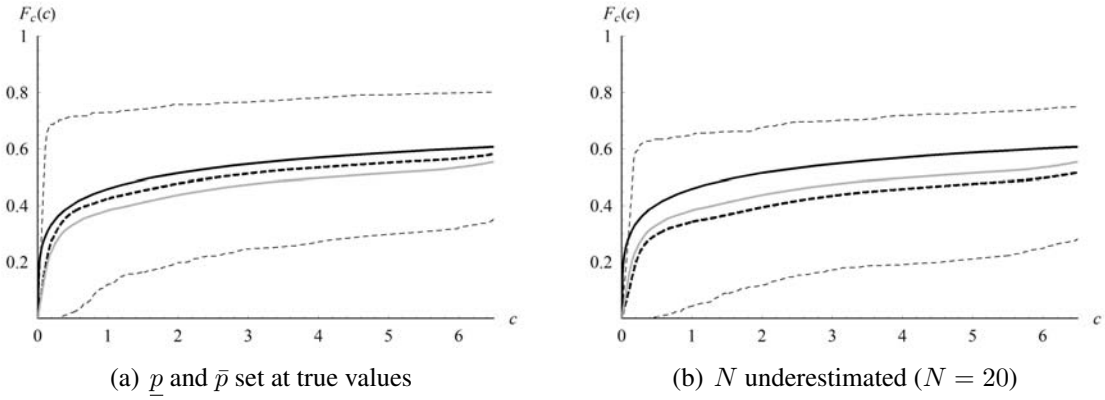


Figure 2.2: Estimated search cost cdf (# periods $T = 4$; # obs. $M = 100$)

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 cdf's becomes smaller as we increase the number of observations. This can be seen in columns 3 and 4 of Tables 2.1 and 2.2 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 bound) 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 2.3(a) and 2.3(b).

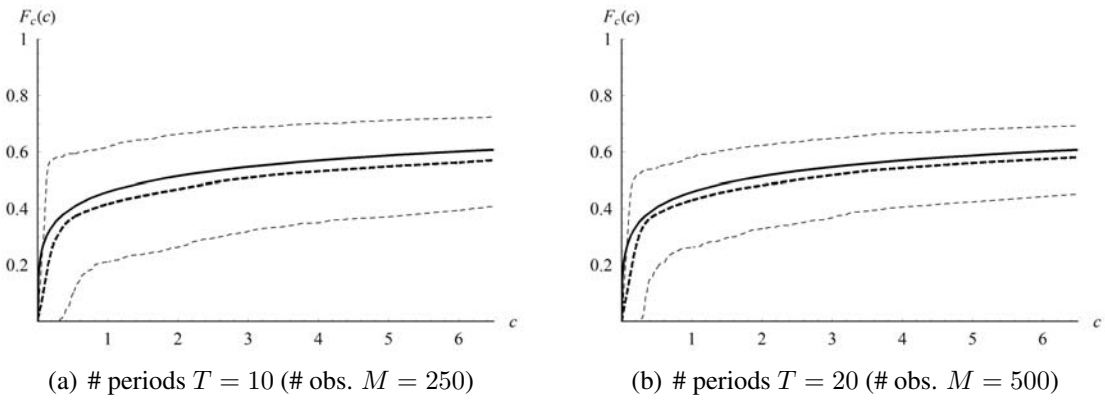


Figure 2.3: Estimated search cost cdf

2.4.2 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 captured a situation where the econometrician observes N with some but not very large error (20% fewer firms). The results can be seen in Figure 2.2(b). 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.




Figure 2.4: Estimated equilibrium price and search cost cdf ($N = 4$; # periods $T = 4$)

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 should be reflected in the estimate of the search cost distribution. Figure 2.4 gives the estimation results. In panel 2.4(a) we plot the estimated price cdf and the true empirical price distribution. Clearly, the fit is much worse than if we had estimated N more or less correctly (see Figures 2.1(a) and 2.2(b)). This has a large impact on the estimates of search cost and marginal cost parameters. As can be seen in Figure 2.4(b), the estimated search cost cdf is far from the true one. 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 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, and clarifies the nature of the bias introduced by this type of measurement error.

2.4.3 Comparison of the ML estimation method and Hong and Shum (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 distribution 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 revise briefly their approach.¹³ 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 charged with probability π_j . Using the equilibrium condition (2.4), each price $i = 1, 2, \dots, M - 1$ in the data set satisfies

$$(p_i - r) \left[\sum_{k=1}^N k q_k \left(1 - \left[\frac{1}{M} \sum_{j=1}^M \pi_j \mathbf{1}(p_j \leq p_i) \right] \right)^{k-1} \right] = (\bar{p} - r) q_1 \quad (2.11)$$

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 (2.11) 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

$$\frac{1}{M} \sum_{j=1}^M \pi_j \left[\mathbf{1} \left(p_j \leq r + \frac{(\bar{p} - r) q_1}{\sum_{k=1}^N k q_k (1 - s_\ell)^{k-1}} \right) - s_\ell \right] = 0 \quad (2.12)$$

¹³For details, we refer to the Appendix of Hong and Shum (2006).

The empirical likelihood problem consists of maximizing $\sum_{j=1}^M \log \pi_j$ subject to the constraints in (2.12) 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 (2.11). 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 (2.6) 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; 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 in case of no-convergence. The results of the simulations are reported in Tables 2.3(a) and 2.3(b).

Table 2.3(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

| | 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.

(a) True and estimated parameter values

| | 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.

(b) True and estimated Δ_i 's

Table 2.3: Estimation results (ML vs Hong and Shum (2006) method)

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 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 2.3(b) shows the true cutoff points of the search cost distribution, along with corresponding estimates based on the empirical cdf of prices and 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 2.5. Figure 2.5(a) shows the MEL estimate of the price cdf along with the true equilibrium price distribution. 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

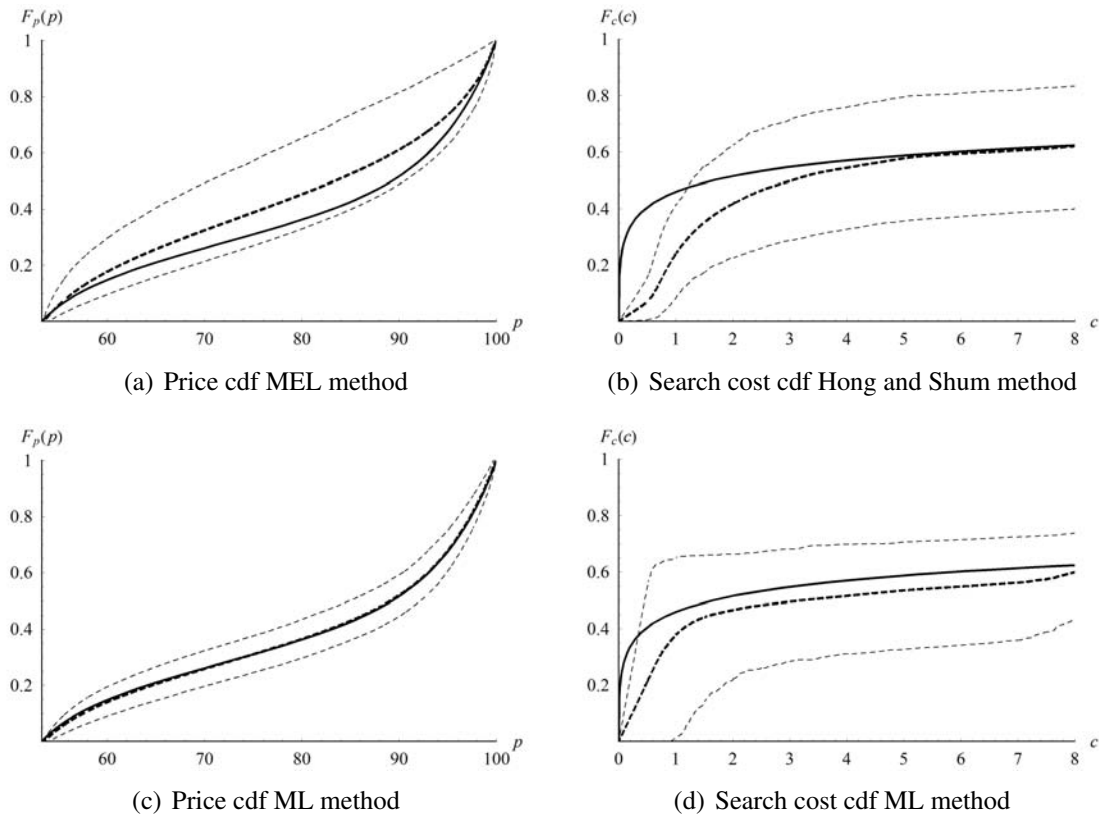


Figure 2.5: True and estimated price and search cost cdf (ML vs. Hong and Shum (2006) method)

Figure 2.5(c); the graph reveals that the fit is remarkably good and clearly outperforms the alternative MEL estimate.

Figure 2.5(b) shows the estimate of the search cost distribution using Hong and Shum procedure, along with the true lognormal distribution with parameters 0.5 and 5. 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 2.5(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, both from a numerical as well as from a goodness-of-fit point of view.

2.5 Empirical application

2.5.1 Data and empirical issues

In this section we apply our estimation procedure to data 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 this 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.¹⁴

Shopper.com is one of the largest price comparison sites on the Internet.¹⁵ 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 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.

¹⁴Since 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.

¹⁵We 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 2.2 to data from on-line markets assumes implicitly that consumers search for prices nonsequentially. Even though nonsequential 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.¹⁶

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 2.4.

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).¹⁷ 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 Hortaçsu 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 because they

¹⁶For details on the optimality of nonsequential and sequential search see Morgan and Manning (1985).

¹⁷The 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 (2003) study of the UK tyre market).

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 2.4: 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, 105.00 | 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 2.5: Summary statistics

The summary statistics of the data can be found in Table 2.5. 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 measurement errors in N . If the true number of retailers is not dramatically different than our estimate of N the results, though biased, will be economically meaningful.¹⁸

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 2.5). 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 be 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 the data in every period should reflect the equilibrium of the static game analyzed in Section 2.2. 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 cdf's 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 2.5 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 from

¹⁸For 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.

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 but 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.¹⁹ 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.²⁰ To check this hypothesis, we gathered prices at the time of writing this chapter 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 are 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, 16 stores were either higher or lower in the ranking compared to one year ago. The difference in ranking ranged from 0 to 7 and was on average of 2.48. Finally, 9 stores out of 21 are now in a different quartile of the ranking distribution. Similar figures apply to the other memory chips.²¹

To check for serial correlation, we calculated autocorrelations for each product at each store, i.e.,

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

¹⁹A 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.

²⁰Lach (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.

²¹For 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 3.39. Finally, for the KTD8300 chip 11 out of 22 stores were in a different quartile, while the average difference in ranking is 3.55.

where p_{av} denotes the store's average price for the product. Excluding the store-product pairs for which we had fewer than eight observations, we found that for 71% of the store-product pairs autocorrelations were not significantly different from zero at a 5% significance level. Although the number of observations for each store-product pair is limited, this evidence suggests that serial correlation is not a serious issue in our data set.²²

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.²³ Another reason is that shipping costs and sales taxes depend on the 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.²⁴

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.

²²In 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 (2004) for a related finding.

²³Actually 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.

²⁴Tables 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.

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:

$$PRICE_{jt} = \beta_0 + \beta_1 \cdot RATING_{jt} + \beta_2 \cdot SHIP_{jt} + \beta_3 \cdot STOCK_{jt} + \varepsilon_{jt}, \quad (2.13)$$

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 (2.13) 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.²⁵ 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 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. However, given the short period of data collection and given the observation 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. Actually, in a related study of online prices, using a similar kind of data set which extends over a much longer period of time (monthly data from November 1999 to May 2001), Baye et al. (2007) find that unobserved firm heterogeneities (in costs, branding, trust, etc.) can only explain up to 72% of the variation in prices.²⁶ Still, a caveat

²⁵For 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 on a 1% level for the KTDINSP8200 chip, significant on 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).

²⁶Lach (2002) presents a similar finding for a data set of prices for goods sold in physical stores located in Israel.

of the current model is that it cannot control for this unobserved firm heterogeneity and that it treats all variation in the price data as variation due to search frictions.

2.5.2 Estimation results

The estimation results for the four different memory chips are presented in Table 2.5. 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 2.6: Estimation results

The estimates of the share of consumers who search once, q_1 , range from 22% to 30% and are all highly significant.²⁷ 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 all approximately zero. Finally, the estimates of the fraction of consumers comparing all the prices in the market, q_N , range

²⁷To be able to calculate the standard errors, we deleted the columns and rows of the Hessian for which the corresponding parameter estimates were zero.

from 4% to 13% and are, except for the KTD3614 memory chip, significant at a 5% level.²⁸

These results suggest a clear picture of consumer 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 2.7: Estimated critical search cost values

The estimated cut-off points of the search cost distribution, Δ_i , with corresponding standard errors are presented in Table 2.7. All the cut-off 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 2.6 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

²⁸In a study of the consumer click-through behavior online, Johnson et al. (2004) also point out that many consumers search quite little.

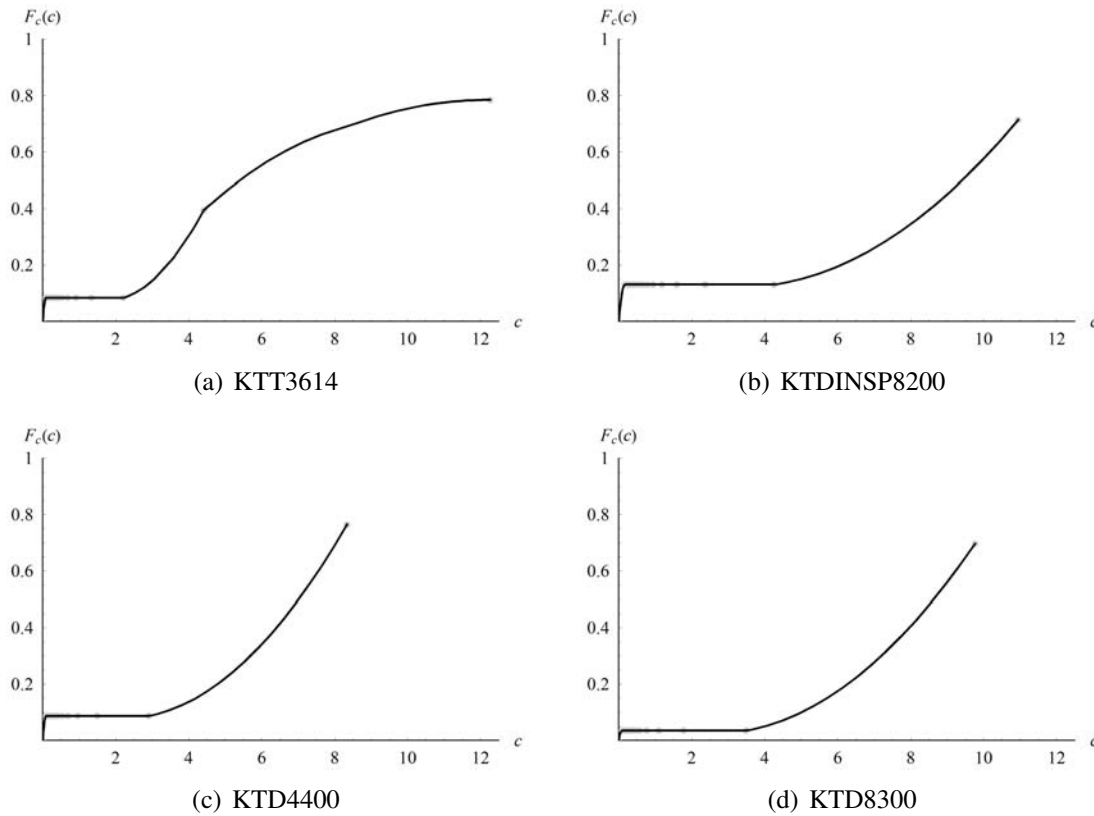


Figure 2.6: Estimated search cost cdf

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 den-

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

sity function fits the data, we use the Kolmogorov-Smirnov test (KS-test) to compare the actual distribution to the fitted distribution.³⁰ 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 2.6 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.³¹ The goodness of fit of the model to the data can be visualized in Figure 2.7. A solid curve represents an empirical price distribution, while a dashed curve represents an estimated one.³²

2.5.3 The effects of a sales tax

The value of estimating a structural model of demand and supply is that the aggregate implications of policy changes can be computed by generating 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 (2.4) 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 \bar{p} while the lower bound of

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.

³¹Because some of the parameters that enter the test are estimated we also calculated the Rao-Robson Statistic, which is a 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.

³²We 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 2.4.3, 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.

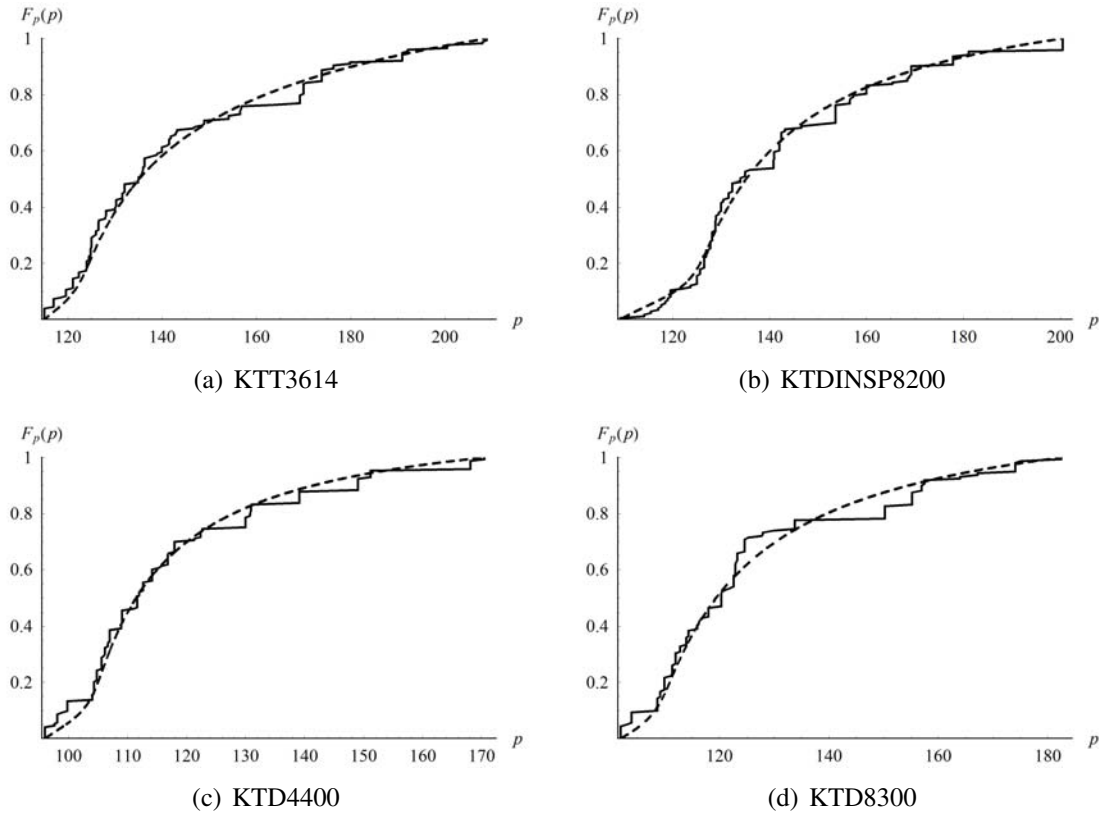


Figure 2.7: Estimated and empirical price cdf

the price cdf changes to

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

Using these two equations, we can simulate the implications of sales taxes on equilibrium. 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 2.6(b)). The fitted search cost density we obtain is

$$\hat{f}_c(c) = 0.36 \cdot \text{lognormal}(c/8.64, 0.27, 9.76) + 0.64 \cdot \text{lognormal}(c/8.64, 0, 0.24)$$

The fitted curve and the estimated points of the search cost distribution can be seen in Figure 2.8(a).

Using this estimate, we simulate the effects of a 5%, 10%, 15% and 20% sales tax. The




Figure 2.8: The effects of a sales tax

results are reported in Table 2.8 while the original and the different post-tax equilibrium price distributions are drawn in Figure 2.8(b). As the graph shows, the introduction of a tax results in a rightward shift of the price distribution so all consumers end-up paying higher final prices. What is interesting is that a tax, by compressing the price distribution, changes the incentives to search in the economy. Inspection of columns 3 to 6 in Table 2.8 reveals that as the tax increases, the number of consumers who do not exercise price comparisons increases. For example, when the tax is 20% (last column of the table) this number is about 83% which implies that firms can charge prices quite close to monopoly (average price is indeed 194.28 US dollars, quite close to valuation, which is 200.50 US dollars).

How much of the tax is passed along to the consumers turns out to depend on the height 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; after-tax firm profits are lower than in the case of no taxation (see columns 2 and 3 of Table 2.8). 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 profits that are higher than in the absence of taxes (see columns 2 and 5 of Table 2.8).

| | KTDINSP8200 estimated | KTDINSP8200 fitted | KTDINSP8200 5% tax | KTDINSP8200 10% tax | KTDINSP8200 15% tax | KTDINSP8200 20% tax |
|--------------|--------------------------|-----------------------|-----------------------|------------------------|------------------------|------------------------|
| p | 109.20 | 108.82 | 114.48 (+5.2%) | 120.92 (+11.1%) | 128.49 (+15.3%) | 142.27 (+30.7%) |
| t | 0 | 0 | 0.05 (+5.0%) | 0.10 (+10.0%) | 0.15 (+15.0%) | 0.20 (+20.0%) |
| v | 200.50 | 200.50 | 200.50 | 200.50 | 200.50 | 200.50 |
| N | 24 | 24 | 24 | 24 | 24 | 24 |
| q_1 | 0.29 (0.04) | 0.27 | 0.30 | 0.34 | 0.42 | 0.83 |
| q_2 | 0.58 (0.02) | 0.56 | 0.53 | 0.49 | 0.42 | 0.01 |
| q_3 | 0.00 | 0.01 | 0.01 | 0.01 | 0.01 | 0.00 |
| q_4 | 0.00 | 0.01 | 0.01 | 0.01 | 0.00 | 0.00 |
| \vdots | \vdots | \vdots | \vdots | \vdots | \vdots | \vdots |
| q_{N-1} | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 |
| q_N | 0.13 (0.05) | 0.12 | 0.12 | 0.12 | 0.12 | 0.14 |
| r | 103.15 (0.84) | 103.15 | 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%) | 194.28 (+78.5%) |
| $Ep_{1:2}$ | 130.01 | 129.01 | 135.37 (+4.9%) | 143.03 (+10.9%) | 153.25 (+18.8%) | 189.42 (+46.8%) |
| $Ep_{1:N}$ | 113.54 | 112.71 | 118.59 (+5.2%) | 125.40 (+11.3%) | 133.77 (+18.7%) | 154.89 (+37.4%) |
| $E\pi$ (net) | 1.16 | 1.10 | 1.09 (-1.3%) | 1.09 (-0.5%) | 1.16 (+5.6%) | 1.97 (+78.9%) |

Notes: Estimated standard errors in parenthesis.

Table 2.8: The effects of a sales tax: estimated and simulated parameter values

2.6 Conclusions

Consumer search models have shown for example that, depending on the nature of search costs in the market, an increase in the number of firms can increase or decrease the price levels and price dispersion. Since competition policy recommendations may depend on the nature of the search cost distribution, there is a need to develop methods to identify 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 up on this line of work, this chapter 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 variation the number of competitors from variation 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 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 the 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 consumer 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 operate in the markets we study, search frictions confer significant market power to the firms. 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.

This chapter also illustrates how the structural methodology 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 one of the markets studied above. We find that a sales tax reduces the gains from search and this may affect the equilibrium in unexpected ways.

Estimation and identification of nonsequential search models

NOTE: This chapter is based on Moraga-González, Sándor, and Wildenbeest (2007).

3.1 Introduction

In this chapter we study the identification and estimation of the costs of search structurally in a more general setting than in the previous chapter. More specifically, we take an oligopoly model where firms are symmetric and sell homogeneous products; on the other side of the market, consumers, who are heterogeneous in regard to their valuations and their search costs, search nonsequentially (Burdett and Judd, 1983).¹ The firm symmetry assumption along with the homogeneous products assumption ensure that the price equilibrium is in mixed strategies; an equilibrium where firms mix provides the econometrician with a first-order condition and a constancy-of-profits condition, which can both be exploited to estimate the model using only price data. Relaxing either of these two assumptions leads to another class of search models where equilibria are in pure strategies so price variation across firms is not only due to search frictions; as shown in Hortaçsu and Syverson (2004), the estimation of these types of model require quantity data.

We start our discussion on identification and estimation of search costs by studying a simpler version of our model where we assume away consumer valuation heterogeneity.² Given that prices reflect the search behavior of groups of consumers and not the behavior of individual buyers, it turns out that the search cost distribution can only be identified at a series of critical points that are determined by consumers' optimal search. In fact, if there are N firms operating in the market then only N points of the search cost distribution can be identified, each point corresponding to the search cost of the marginal consumer who is indifferent between searching k and $k + 1$ times.

The model studied in Hong and Shum (2006) is a particular case of this simpler model where the number of firms operating in the market is set to infinity. Hong and Shum (2006, p.262) suggest that when the number of firms goes to infinity, identification of the search cost distribution in its full support should obtain. We show that this presumption is wrong because the (infinite) series of critical search costs turns out to be convergent to zero. This property, which stems from the fact that the marginal gains from an extra search are declining in the

¹Nonsequential (or fixed-sample-size) search is appealing when consumers find it more advantageous to gather price information quickly (Morgan and Manning, 1985). This occurs when the search outcome is observed with some delay, like in markets for labor (Fershtman and Fishman, 1992), for mortgages, for refurbishing services, for specialized inputs, etc.

²The simpler model is also studied in Chapter 2.

number of searches, implies that the set of search cost values the econometrician can retrieve from the price data is not dense in the support of the search cost distribution so nonparametric identification of the search cost cdf in its entire support fails.

To overcome this identification problem we propose to consider a different framework where the econometrician has price data from several oligopolistic markets, each observed over some period of time. In particular, we consider markets with the common feature that the search cost distribution is the same, while consumer valuations differ across markets. We show that the search cost cdf can be identified fully in such setting; the reason is that every market generates a distinctive set of search cost values for which the econometrician can retrieve the density of search costs, and this forces the search cost distribution to be uniquely determined for a larger set of points.

Given that we need to pool price data from multiple markets to identify and quantify search costs, the spline approximation methods employed earlier in the literature (cf. Hong and Shum, 2006; Hortaçsu and Syverson, 2004) are not valid in our setting. The reason is that spline approximations use procedures in which distinct markets are not linked via the same underlying search cost distribution. To exploit the linkage between markets, we propose to estimate the search cost density function directly by a flexible polynomial-type parametric function, namely, a semi-nonparametric (SNP) density estimator (Gallant and Nychka, 1987). In this way, we obtain an essentially nonparametric estimator of the search cost distribution common to all the markets, because the SNP density estimators approximate arbitrarily closely any sufficiently smooth density function (Gallant and Nychka, 1987). Since the parameters of the SNP density estimator are chosen directly to maximize the joint likelihood corresponding to all the markets, our SNP maximum likelihood estimator does not have the problem of spline approximations.

The identification problem of the simpler model without valuation heterogeneity of course carries over to the general model. In this case we also need to pool price data from several markets if we want to make some progress in the nonparametric identification issue. The natural proposal is to consider a setting where the econometrician has price data from several oligopolistic markets all of them with the same joint density of valuations and search costs, but with variation in the number firms and marginal costs between markets. It turns out that this is not sufficient for identifying a bivariate density nonparametrically. If search costs and valuations are uncorrelated, we can identify nonparametrically the search cost distribution but for the

identification of the valuation density we need to make distributional assumptions. By contrast, if there is correlation between valuations and search costs, only parametric identification of the joint-density obtains.

The structure of this essay is as follows. In the next section, we present our consumer search model with valuation and search cost heterogeneity. In Section 3.3 we study the characterization of a price dispersed symmetric equilibrium. Section 3.4 is dedicated to the simulation of the welfare effects of mergers and serves to motivate the identification and estimation of search costs. Our first set of identification results and our SNP estimation method are presented in Section 3.5, where we study the simpler version of the model with no valuations heterogeneity. In this section we also present some simulation results to illustrate the identification and estimation issues. Section 3.5 also discusses the estimation and identification of the full model. Finally, Section 3.6 concludes.

3.2 The model

We examine a model of firm competition in the presence of consumer search. The model is an oligopolistic version of Burdett and Judd (1983) with heterogeneity in consumer search costs and in consumer valuations. On the supply side of the market there are N firms producing a homogeneous good at constant returns to scale. Their identical unit cost is equal to r and they compete in prices. On the demand side of the market there is a unit mass of buyers. Each consumer wishes to purchase a single unit of the good at most. A consumer is characterized by his/her valuation and his/her search cost, i.e., a pair (v, c) , which is assumed to be a random draw from a joint distribution function $G(v, c)$, defined over the set $\mathbb{R}^+ \times \mathbb{R}^+$. Let $g(v, c)$ be the joint density function and assume that $g(v, c) > 0$ everywhere. We shall denote the marginal valuation and search cost distributions as $G_v(v)$ and $G_c(c)$, respectively, with corresponding densities $g_v(v)$ and $g_c(c)$.

Consumers must engage in costly search to observe prices. Assume they search non-sequentially, i.e., they decide before hand how many prices to observe. Once a consumer has observed the desired number of prices, he/she chooses to buy from the store in his/her sample charging the lowest price, or not to buy at all. A consumer with search cost c sampling k firms

incurs a total search cost kc .³

Firms and buyers play a simultaneous moves game. An individual firm chooses its price taking rivals' prices as well as consumers' search behavior as given. A firm i 's strategy is denoted by a distribution of prices $F_i(p)$. Let $F_{-i}(p)$ denote the vector of prices charged by firms other than i . The (expected) profit to firm i from charging price p_i given rivals' strategies is denoted $\pi_i(p_i, F_{-i}(p))$. Likewise, an individual buyer takes as given firm pricing and decides on his/her optimal search strategy to maximize his/her expected utility. The strategy of a consumer with search cost c is then a number k of firm prices to sample. Let the fraction of consumers sampling k firms be denoted by μ_k , with $\sum_{k=0}^N \mu_k = 1$.

We shall concentrate on symmetric Nash equilibria. A symmetric equilibrium is a distribution of prices $F(p)$ and a grouping of consumers $\{\mu_0, \mu_1, \dots, \mu_N\}$ such that (a) $\pi_i(p, F_{-i}(p))$ is equal to a constant $\bar{\pi}$ for all p in the support of $F(p)$, $\forall i$; (b) $\pi_i(p, F_{-i}(p)) \leq \bar{\pi}$ for all $p, \forall i$; and (c) a consumer sampling k firms obtains no lower expected utility than by sampling any other number of firms. Let us denote the equilibrium density of prices by $f(p)$, with maximum price \bar{p} and minimum price \underline{p} .

3.3 Theoretical analysis

Our first result indicates that, for an equilibrium to exist, there must be some consumers who search just once and others who search more than once. The intuition is simple. If all consumers compared two or more prices in the market then firms would have no other choice than pricing competitively, in which case consumers would not be willing to search that much. Alternatively, if no consumer compared prices then firms would charge the same price in equilibrium. In that case consumers would continue to enter the market until the expected utility from buying $v - p$ equals search cost, in which case a firm would gain by deviating by increasing its price.

Lemma 3.1 *If a symmetric equilibrium exists, then $1 > \mu_1 > 0$ and $\mu_k > 0$ for some $k = 2, 3, \dots, N$.*

³We note that we are assuming that the first search is also costly, in contrast to most of the literature on consumer search. This assumption alone implies that not all consumers will participate in the market, and that the participation rate is endogenous.

Proof. First, suppose, on the contrary, that $\mu_1 = 0$. Then we have two possibilities: (i) either $\mu_0 = 1$ in which case the market does not open, or (ii) $\mu_k > 0$ for some $k = 0, 2, 3, \dots, N$ in which case all firms would charge a price equal to the marginal cost r . But if this were so, consumers would gain by deviating and searching for fewer prices.

Second, suppose, on the contrary, that $\mu_1 = 1$. Since equilibrium prices must be above the marginal cost r and since search is costly, it is obvious that consumers with valuations below r would not enter the market so all consumers searching once cannot be part of an equilibrium.

Finally, suppose, on the contrary, that $1 > \mu_1 > 0$ and that $\mu_k = 0$ for all $k = 2, 3, \dots, N$, i.e., $\mu_0 + \mu_1 = 1$. Given that consumers observe a single price at most, all firms must charge the same price in equilibrium. Let p be the equilibrium price. Then, if consumers expect to see a price p when they venture a store, all consumers with a valuation larger than p and search cost of at most $v - p$, i.e., a fraction

$$\int_p^\infty \int_0^{v-p} g(v, c) dc dv$$

would indeed wish to enter the market. Consider a deviant firm charging $\tilde{p} > p$; this firm would make a profit

$$\pi_i(\tilde{p} > p) = \frac{1}{N}(\tilde{p} - r) \int_{\tilde{p}}^\infty \int_0^{v-p} g(v, c) dc dv. \quad (3.1)$$

Taking the derivative of (3.1) with respect to \tilde{p} we obtain

$$\frac{d\pi_i(\cdot)}{d\tilde{p}} = \frac{1}{N} \int_{\tilde{p}}^\infty \int_0^{v-p} g(v, c) dc dv - \frac{1}{N}(\tilde{p} - r) \int_0^{\tilde{p}-p} g(\tilde{p}, c) dc$$

and evaluating it at $\tilde{p} = p$ we get

$$\left. \frac{d\pi_i(\cdot)}{d\tilde{p}} \right|_{\tilde{p}=p} = \frac{1}{N} \int_p^\infty \int_0^{v-p} g(v, c) dc dv > 0,$$

which implies that p cannot be an equilibrium price. ■

Our next observation is that, given that, conditional on entering the market, some consumers will search only once while others will search for two prices or more, for an equilibrium to exist it must be the case that firm pricing is characterized by mixed strategies.

Lemma 3.2 *If a symmetric equilibrium exists, $F(p)$ must be atomless.*

Proof. Suppose, on the contrary, that firms did charge a price $\hat{p} > r$ with strictly positive probability in equilibrium. Consider a firm i charging \hat{p} . The probability that \hat{p} is the lowest price in the market is strictly positive. This occurs, for example, when all other firms are also charging \hat{p} . From Lemma 3.1 we know that in equilibrium there must exist some $\hat{k} \in \{2, 3, \dots, N\}$ for which $\mu_{\hat{k}} > 0$. Consider the fraction of consumers sampling \hat{k} firms. The probability that these consumers are sampling firm i is strictly positive; as a result, firm i would gain by deviating and charging $\hat{p} - \varepsilon$ since in that case the firm would attract all consumers in $\mu_{\hat{k}}$ who happened to sample firm i and were willing to buy at price \hat{p} . This deviation would give firm i a discrete increase in its profits and thus rules out atoms at any price $\hat{p} > r$.

It remains to be proven that an atom at the marginal cost r cannot be part of an equilibrium either. Consider a firm charging r . From Lemma 3.1 we know that in equilibrium $1 > \mu_1 > 0$. This implies that some consumers will sample the firm in question with strictly positive probability and will proceed buying provided their valuation is not lower than r . The firm would thus obtain positive demand with a strictly positive probability and would however make zero profits. It is then clear that the firm would have an incentive to deviate by increasing its price. ■

Lemma 3.2 implies that the measure of consumers searching for k prices must be strictly positive, $k = 0, 1, 2, \dots, N$. We now calculate the fractions of consumers $\mu_k, k = 0, 1, 2, \dots, N$ who sample k prices. We start by calculating the set of consumers who, owing to their high search costs or to their low valuations, get out of the market and do not search at all.⁴ Denoting the expectation operator by E , the expected utility of a consumer (v, c) who searches k times is $v - E[\min\{p_1, p_2, \dots, p_k\}] - kc$. Therefore, the set of consumers who do not participate in the market is given by $S_0 = \{(v, c) \in \mathbb{R}^+ \times \mathbb{R}^+ \text{ such that } v - E[\min\{p_1, p_2, \dots, p_k\}] - kc \leq 0, \forall k\}$, so the fraction of consumers who get out of the market is

$$\mu_0 = \int_{S_0} g(v, c) dc dv.$$

⁴Notice that in this model there will be consumers who find it optimal to search $k = 1, 2, \dots, N$ times but end up not buying at all simply because their draw of prices is a bad one.

We now calculate the fraction of consumers who search k times. For a consumer (v, c) to search for k prices, the following conditions must hold:

$$(i) \ v - E[\min\{p_1, p_2, \dots, p_k\}] - kc > 0. \quad (3.2a)$$

$$(ii) \ v - E[\min\{p_1, p_2, \dots, p_k\}] - kc > v - E[\min\{p_1, p_2, \dots, p_{k+1}\}] - (k+1)c \quad (3.2b)$$

$$(iii) \ v - E[\min\{p_1, p_2, \dots, p_k\}] - kc > v - E[\min\{p_1, p_2, \dots, p_{k-1}\}] - (k-1)c \quad (3.2c)$$

Condition (i) says that a consumer who decides to search for k prices must obtain a positive expected utility from doing so and conditions (ii) and (iii) guarantee that the consumer prefers to search for k prices rather than for $k+1$ or $k-1$, respectively. Notice that conditions (ii) and (iii) are independent of the valuation v so if a consumer prefers k over $k+1$ or $k-1$ searches, this also holds for all consumers with the same search cost but different valuations.

Since $F(p)$ is atomless, the order statistic $E[\min\{p_1, p_2, \dots, p_k\}]$ is a decreasing and convex function of k . Therefore there exists a unique search cost value, denoted c_k , corresponding to the consumer indifferent between searching k times and searching $k+1$ times, which satisfies $v - E[\min\{p_1, p_2, \dots, p_k\}] - kc_k = v - E[\min\{p_1, p_2, \dots, p_{k+1}\}] - (k+1)c_k$. Isolating c_k from this expression leads to

$$c_k = E[\min\{p_1, p_2, \dots, p_k\}] - E[\min\{p_1, p_2, \dots, p_{k+1}\}], \ k = 1, 2, \dots, N-1. \quad (3.3)$$

Combining equations (3.2a)-(3.2c) and equation (3.3) leads to the following result.

Lemma 3.3 *For any given atomless price distribution $F(p)$, optimal consumer search behavior leads to a unique partition of the set of consumers where the fraction of buyers searching for k prices is given by:*

$$\mu_1 = \int_{c_1}^{\infty} \int_{c+E[p]}^{\infty} g(v, c) dv dc \quad (3.4a)$$

$$\mu_k = \int_{c_k}^{c_{k-1}} \int_{kc+E[\min\{p_1, p_2, \dots, p_k\}]}^{\infty} g(v, c) dv dc; \ k = 2, \dots, N-1 \quad (3.4b)$$

$$\mu_N = \int_0^{c_{N-1}} \int_{Nc+E[\min\{p_1, p_2, \dots, p_N\}]}^{\infty} g(v, c) dv dc, \quad (3.4c)$$

with $\mu_0 = 1 - \sum_{k=1}^N \mu_k$ and where c_k is given by equation (3.3).

The calculation of the fractions of consumers μ_k 's who search k times is illustrated in Figure 3.3 for a market where the number of firms $N = 5$ and where the distribution of search cost and valuations is uniform on $\mathbb{R}^+ \times \mathbb{R}^+$.




Figure 3.1: The μ_k 's graphically

We now move to examine firms' pricing decisions. Taking consumer search behavior μ_k , $k = 0, 1, 2, \dots, N$ as given, an individual firm quotes a price to maximize expected profits. Consider a consumer who searches for k prices. To write out the payoff of a firm i charging p_i we ask when a consumer buys from firm i . Three events must be realized: one, the consumer must sample the price of firm i ; two, the price of firm i must be lower than the price of the other $k - 1$ firms sampled by the consumer; and three, the valuation of the consumer v should not be lower than p_i . As a result, the profit to a firm i charging p_i given the strategies of the rivals is:

$$\pi_i(p_i; F(\cdot)) = (p_i - r) \left[\sum_{k=1}^N \frac{k d_k(p_i)}{N} (1 - F(p_i))^{k-1} \right],$$

where $d_k(p_i)$ represents the probability of the joint event that a consumer (v, c) samples firm i plus $k - 1$ other firms and has a valuation v greater than p_i . These probabilities are given by the following expressions:

$$d_1(p) = \begin{cases} \int_p^\infty \int_{c_1}^{v-E[p]} g(v, c) dc dv & \text{if } p > c_1 + E[p] \\ \mu_1 & \text{otherwise,} \end{cases}$$

and for $k = 2, 3, \dots, N$

$$d_k(p) = \begin{cases} \int_p^\infty \int_{c_k}^{\min\{\frac{1}{k}(v-E[\min\{p_1, p_2, \dots, p_k\}]), c_{k-1}\}} g(v, c) dc dv & \text{if } p > kc_k + E[\min\{p_1, p_2, \dots, p_k\}] \\ \mu_k & \text{otherwise.} \end{cases}$$

In a mixed strategy equilibrium, a firm must be indifferent between charging any price p_i in the support of $F(\cdot)$. Consider, in particular, a firm charging the upper bound of the price distribution \bar{p} . Since such a firm will only sell to those consumers who search once, this upper bound must be equal to the maximizer of

$$\pi_i(\bar{p}; F(\cdot)) = \frac{(\bar{p} - r)}{N} d_1(\bar{p})$$

Or

$$\pi_i(\bar{p}; F(p)) = \begin{cases} \frac{(\bar{p} - r)}{N} \int_{\bar{p}}^\infty \int_{c_1}^{v-E[p]} g(v, c) dc dv & \text{if } \bar{p} > c_1 + E[p] \\ \frac{(\bar{p} - r)}{N} \mu_1 & \text{if } \bar{p} \leq c_1 + E[p] \end{cases} \quad (3.5)$$

Inspection of this last equation reveals that the upper bound \bar{p} must be no lower than $c_1 + E[p]$ just because the profit of the firm is strictly increasing in \bar{p} for all $\bar{p} < c_1 + E[p]$. Therefore, the upper bound \bar{p} is the solution to the following first order condition:

$$\int_{\bar{p}}^\infty \int_{c_1}^{v-E[p]} g(v, c) dc dv - (\bar{p} - r) \int_{c_1}^{\bar{p}-E[p]} g(\bar{p}, c) dc = 0 \quad (3.6)$$

For a finite upper bound to exist, it is sufficient to assume that $\lim_{p \rightarrow \infty} p(1 - G_v(p)) = 0$ (recall $G_v(\cdot)$ is the marginal cdf of valuations). This assumption is satisfied by distributions G_v with tails as thick as the lognormal distribution. To see that the assumption is sufficient, notice

that

$$\int_{\bar{p}}^{\infty} \int_{c_1}^{v-E[p]} g(v, c) dc dv < \int_{\bar{p}}^{\infty} \int_0^{\infty} g(v, c) dc dv = \int_{\bar{p}}^{\infty} g_v(v) dv = 1 - G_v(\bar{p}).$$

This implies that the profit expression in equation (3.5) goes to 0 as $\bar{p} \rightarrow \infty$, and since it is strictly positive for $\bar{p} = c_1 + E[p]$ for which $d_1(\bar{p}) = \mu_1$ and it is continuous in \bar{p} , we conclude that it has a (finite) global maximum.

Given $\bar{p} = \arg \max_p \pi_i(p; F(\cdot))$, the symmetric Nash equilibrium is characterized by the constancy of profits condition: $\pi_i(p_i; F(\cdot)) = \pi_i(\bar{p}; F(\cdot))$, i.e.,

$$(p_i - r) \left[\sum_{k=1}^N k d_k(p_i) (1 - F(p_i))^{k-1} \right] = (\bar{p} - r) d_1(\bar{p}). \quad (3.7)$$

Equation (3.7) cannot be solved in closed-form for $F(p_i)$. However, the minimum price charged in the market can be found by setting $F(\underline{p}) = 0$ in equation (3.7) and solving it for \underline{p} . This yields:

$$\underline{p} = r + (\bar{p} - r) \frac{d_1(\bar{p})}{\sum_{k=1}^N k \mu_k}. \quad (3.8)$$

To prove uniqueness of $F(p_i)$ in equation (3.7), which is useful for estimating the model, let us rewrite equation (3.7) as

$$\sum_{k=1}^N k d_k(p) (1 - z)^{k-1} = \frac{(\bar{p} - r) d_1(\bar{p})}{p - r}. \quad (3.9)$$

Then we need to prove that for any price p , this equation cannot have more than one solution in $z \in [0, 1]$. Note that the RHS of equation (3.9) is positive and does not depend on z , while the LHS of equation (3.9) is a positive-valued function strictly decreasing in z . Therefore, for any price p , there can exist at most one $F(p)$ that satisfies equation (3.7). These arguments lead to the following result.

Lemma 3.4 *Given consumer search behavior $\{\mu_k\}_{k=0}^N$, with $\mu_k > 0$ for all $k \leq N$, a distribution of consumer valuations and search costs $G(v, c)$, as well as a marginal cost r , if there exists an equilibrium price cdf $F(\cdot)$ with upper bound \bar{p} then it is unique and it is given by the simultaneous solution to equations (3.6) and (3.7).*

Lemma 3.4 shows that for any given grouping of consumers, the equilibrium price distribution is unique. An equilibrium of the game requires the conjectured grouping of consumers to be the outcome of optimal consumer search given correct conjectures about the price distribution.

It will be useful to rewrite the critical search cost values c_k 's given in equation (3.3). Using the distributions of the order statistics, and after successively integrating by parts, we can rewrite them as follows:

$$c_0 = \int_{\underline{p}}^{\bar{p}} F(p) dp; \quad (3.10a)$$

$$c_k = \int_{\underline{p}}^{\bar{p}} F(p)(1 - F(p))^k dp, \quad k = 1, 2, \dots, N - 1. \quad (3.10b)$$

Since $F(p)$ is monotonically increasing in p , its inverse exists. Let $p(z)$ denote the inverse of $F(p)$. Using this inverse function, a change of variables in equations (3.10a) and (3.10b) yields:

$$c_0 = \bar{p} - \int_0^1 p(z) dz; \quad (3.11a)$$

$$c_k = \int_0^1 p(z)[(k+1)z - 1](1 - z)^{k-1} dz, \quad k = 1, 2, \dots, N - 1. \quad (3.11b)$$

Therefore:

Proposition 3.1 *If a symmetric equilibrium of the game described above exists then firms set prices according to Lemma 3.4, consumers search according to Lemma 3.3, and the numbers c_k 's are given by the solution to the system of equations (3.11a) and (3.11b).*

An equilibrium of the full game is given by the simultaneous solution of equations (3.7), (3.4a)-(3.4c), and (3.11a) and (3.11b). Proving that such a solution exists and is unique turns out to be very difficult. Our numerical simulations (some of which we present below), however, suggest that the symmetric equilibrium exists and is unique.

3.4 Mergers

Economists and practitioners need to take sound public policy decisions and this requires to have a reasonable knowledge about the market. In this section we illustrate the importance of estimating models of consumer search; we show how absence of good knowledge about consumer search costs may lead to wrong (or at least debatable) recommendations. We shall focus on the impact of mergers on prices and overall welfare but the analysis can be extended to study other public policies such as price controls, taxes and subsidies, etc.

The price and welfare effects of mergers in our model are difficult to derive analytically since the equilibrium price distribution cannot be obtained in closed-form. We then proceed by solving the model numerically. For simplicity, we consider a market where consumer valuations are identical and we set $v = 100$; in addition, we assume that the first price quotation is obtained at no cost. These two assumptions together imply that all consumers buy in equilibrium. This is convenient since it allows us to isolate the single mechanism through which mergers influence the aggregate outcome: the amount of search.⁵ The firms' marginal cost r is set equal to 50.

Let us assume that search costs follow a log-normal distribution, with parameters (ν_c, σ_c) . In what follows, we fix the mean search cost to 50 and compare how the market works for two different levels of search cost dispersion. In particular, we focus on the effects of mergers on prices and surplus and study how these effects depend on the amount of search cost dispersion. We start with a market where search cost dispersion is relatively low. For this we set $(\nu_c, \sigma_c) = (2.63, 1.6)$.⁶ Given the other data, we solve for the equilibrium of the model for different number of firms. The results are reported in Table 3.1.

Table 3.1 shows how consumer search intensities change as we increase the number of firms. Remarkably, in this market a large majority of the consumers searches only once. For example, when there are just two firms in the market the fraction of consumers who do not compare prices is almost 100%. This number remains high but decreases as we increase the number of firms. A second important feature is that very few consumers make an exhaustive search in the market;

⁵The results do not depend on this assumption. Adding consumer valuation heterogeneity obscures the welfare results because there are several effects at work; in particular, the fraction of consumers not buying at all is endogenous and depends on the number of firms operating in the market so not only the amount of search influences the aggregate results but also the rate of consumer participation.

⁶Mean search cost is equal to $\exp(\nu + \sigma^2/2) \simeq 50$ and standard deviation $\sqrt{(\exp(\sigma^2) - 1) \exp(2\nu + \sigma^2)} \simeq 172.74$. The coefficient of variation (CV) is equal to $172.74/50 = 3.45$.

| | $N = 2$ | $N = 3$ | $N = 4$ | $N = 5$ | $N = 6$ | $N = 7$ | $N = 8$ | $N = 9$ | $N = 10$ | $N = 11$ | $N = 12$ |
|------------|---------|---------|---------|---------|---------|---------|---------|---------|----------|----------|----------|
| μ_1 | 1.00 | 0.86 | 0.81 | 0.79 | 0.78 | 0.78 | 0.78 | 0.78 | 0.77 | 0.77 | 0.77 |
| μ_2 | 0.00 | 0.07 | 0.09 | 0.09 | 0.09 | 0.09 | 0.09 | 0.09 | 0.09 | 0.09 | 0.09 |
| μ_3 | - | 0.07 | 0.04 | 0.05 | 0.05 | 0.05 | 0.05 | 0.05 | 0.05 | 0.05 | 0.05 |
| μ_4 | - | - | 0.06 | 0.03 | 0.03 | 0.03 | 0.03 | 0.03 | 0.03 | 0.03 | 0.03 |
| μ_5 | - | - | - | 0.05 | 0.02 | 0.02 | 0.02 | 0.02 | 0.02 | 0.02 | 0.02 |
| μ_6 | - | - | - | - | 0.04 | 0.01 | 0.01 | 0.01 | 0.01 | 0.01 | 0.01 |
| μ_7 | - | - | - | - | - | 0.03 | 0.01 | 0.01 | 0.01 | 0.01 | 0.01 |
| μ_8 | - | - | - | - | - | - | 0.02 | 0.01 | 0.01 | 0.01 | 0.01 |
| μ_9 | - | - | - | - | - | - | - | 0.02 | 0.00 | 0.00 | 0.00 |
| μ_{10} | - | - | - | - | - | - | - | - | 0.01 | 0.00 | 0.00 |
| μ_{11} | - | - | - | - | - | - | - | - | - | 0.01 | 0.00 |
| μ_{12} | - | - | - | - | - | - | - | - | - | - | 0.01 |

Table 3.1: Equilibrium search intensities for $(\nu_c, \sigma_c) = (2.63, 1.6)$ (mean is 50; CV is 3.45)

in fact for example if there are 10 firms in the industry about 94% of the consumers searches for a maximum of 4 firms.

The fact that most consumers do not compare prices is reflected in equilibrium prices. Figure 3.2(a) shows how mean prices change with the number of firms. The average price under

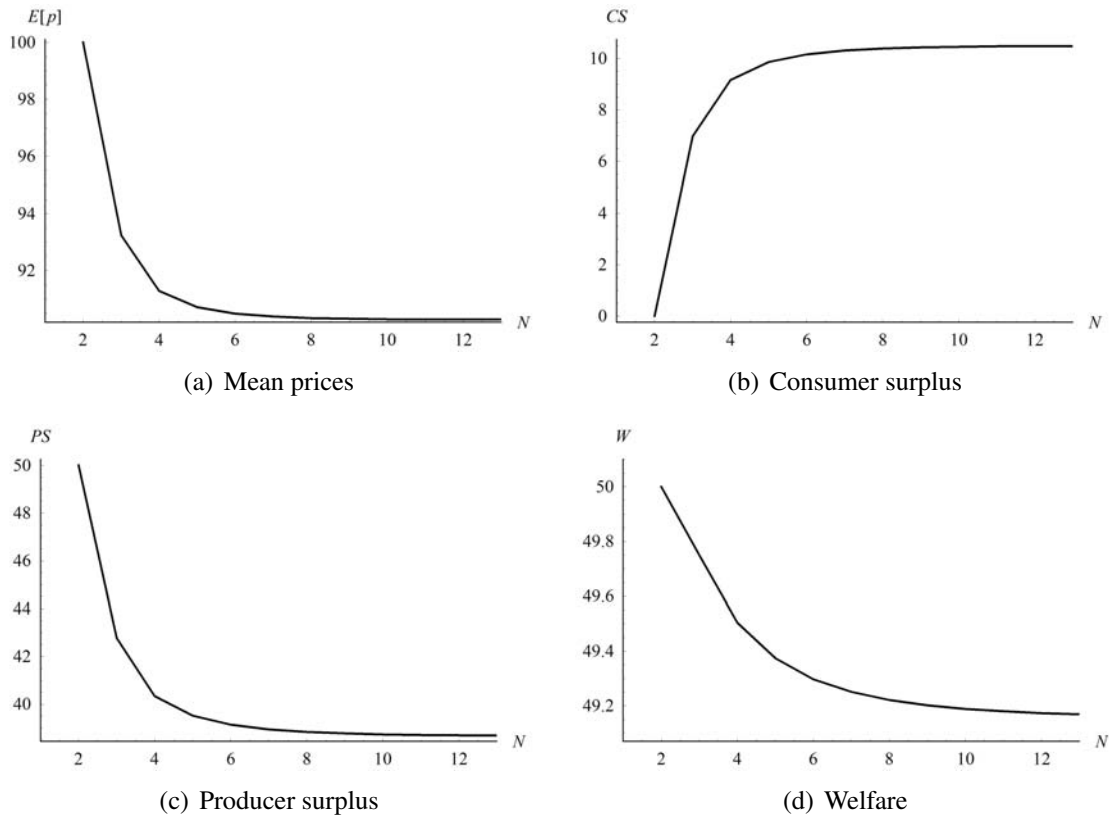


Figure 3.2: Comparative statics of an increase in the number of firms $((\nu_c, \sigma_c) = (2.63, 1.6))$

duopoly is very high and it decreases as the number of firms rises. The decrease of the mean price is due to the fact that the share of consumers comparing two or three prices increases in the number of firms. The average price is what is important for consumers who do not exercise price comparisons so consumers benefit from the resulting average price decreases. These gains are also reflected in that consumer surplus, plotted in Figure 3.2(b), increases in N . Figures 3.2(c) and 3.2(d) show the behavior of aggregate profits and social welfare. The welfare result is perhaps surprising and it deserves an explanation. Note from Table 1 that the amount of search increases in N and that, since search costs are wasteful, more search generates a welfare loss. In sum, mergers in this case of low search cost dispersion would lead to higher average prices, lower consumer surplus, greater industry profits and greater welfare.

The situation is quite different when search costs are much more dispersed, holding everything else equal. Let us set $(\nu_c, \sigma_c) = (0.79, 2.5)$, which implies the new search cost distribution is a mean-preserving spread of the previous one. The new equilibrium search intensities are reported in Table 3.2. What is different in this case of high search cost dispersion is that a great deal of consumers conduct an exhaustive search; as before, the extent of price comparison in the market increases as the number of firms rises.

| | $N = 2$ | $N = 3$ | $N = 4$ | $N = 5$ | $N = 6$ | $N = 7$ | $N = 8$ | $N = 9$ | $N = 10$ | $N = 11$ | $N = 12$ |
|------------|---------|---------|---------|---------|---------|---------|---------|---------|----------|----------|----------|
| μ_1 | 0.37 | 0.32 | 0.30 | 0.30 | 0.29 | 0.29 | 0.29 | 0.29 | 0.29 | 0.29 | 0.29 |
| μ_2 | 0.63 | 0.14 | 0.12 | 0.11 | 0.11 | 0.10 | 0.10 | 0.10 | 0.10 | 0.10 | 0.10 |
| μ_3 | - | 0.54 | 0.10 | 0.10 | 0.09 | 0.09 | 0.08 | 0.08 | 0.08 | 0.08 | 0.08 |
| μ_4 | - | - | 0.47 | 0.08 | 0.08 | 0.07 | 0.07 | 0.07 | 0.07 | 0.07 | 0.07 |
| μ_5 | - | - | - | 0.42 | 0.06 | 0.06 | 0.06 | 0.06 | 0.06 | 0.06 | 0.06 |
| μ_6 | - | - | - | - | 0.37 | 0.05 | 0.05 | 0.05 | 0.05 | 0.05 | 0.05 |
| μ_7 | - | - | - | - | - | 0.33 | 0.04 | 0.04 | 0.04 | 0.04 | 0.04 |
| μ_8 | - | - | - | - | - | - | 0.30 | 0.04 | 0.04 | 0.04 | 0.04 |
| μ_9 | - | - | - | - | - | - | - | 0.27 | 0.03 | 0.03 | 0.03 |
| μ_{10} | - | - | - | - | - | - | - | - | 0.25 | 0.03 | 0.03 |
| μ_{11} | - | - | - | - | - | - | - | - | - | 0.22 | 0.02 |
| μ_{12} | - | - | - | - | - | - | - | - | - | - | 0.21 |

Table 3.2: Equilibrium search intensities for $(\nu_c, \sigma_c) = (0.79, 2.5)$ (mean is 50; CV is 22.73)

Figure 3.3 plots the equilibrium mean price against the number of competitors in the industry. Under duopoly, the average price is relatively low compared to the previous case. What is remarkably different is that the mean price increases as more firms enter the industry. Moreover, we see that consumer surplus can decrease and profits increase as we the number of competitors goes up. The crucial distinction between these two examples is the equilibrium consumer

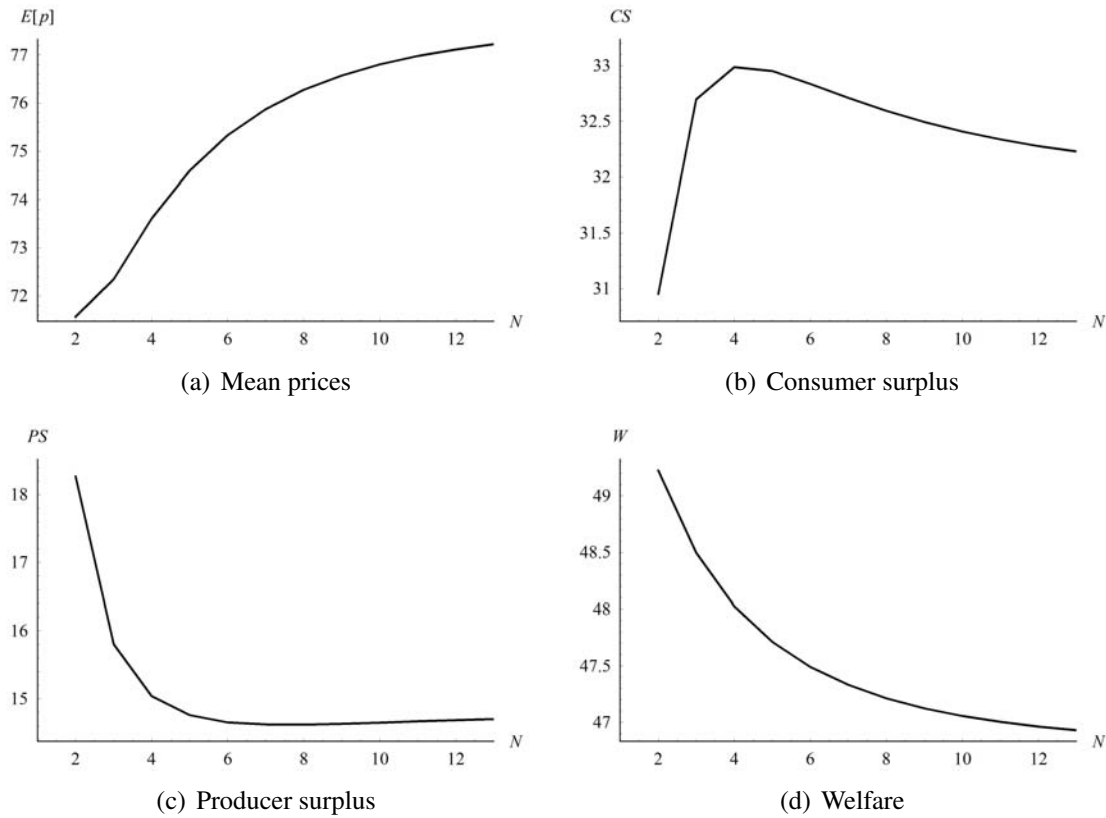


Figure 3.3: Comparative statics of an increase in the number of firms ($\nu_c = 0.79, \sigma_c = 2.5$)

search intensity. Table 3.2 shows that most of the consumers (more than 63%) exercise price comparisons in this case while Table 3.1 showed the opposite evidence. Consumers who conduct an exhaustive search in the market become disproportionately less attractive for a firm as more competitors are around. This effect, which leads to higher prices, has here a dominating influence and results in lower consumer surplus and greater industry profits. Welfare is again decreasing in N due to the rise of actually incurred search costs.

In summary, this section shows that mergers can lead to an increase or to a decrease in average prices and that the direction of the effect depends on the extent of search frictions in the market. For the econometrician, the question is how to elicit such information on the basis of observable market variables. This is the question we address in the remaining of the paper.

3.5 Estimation

In this section we discuss how to estimate the unknown distribution function $G(\cdot, \cdot)$ of search costs and consumer valuations if we observe the prices of the products in several markets. As mentioned in the Introduction, the estimation method employed is maximum likelihood. This method is well suited for the model presented in Section 3.2 because the model predicts that the prices in a market are the outcome of a mixed strategy price equilibrium. Therefore, observed prices are draws from the same equilibrium price distribution. In this section we maintain the assumption that such an equilibrium price distribution exists in each market. This is not a restrictive assumption because, without it, any method of estimation of this model fails. The validity of this assumption can be verified posterior to estimation by comparing the empirical price distribution with the estimated price distribution in a given market.

Taking these considerations into account, we can write the log-likelihood function corresponding to a market as $LL(\mathbf{p}) = \sum_i \log f(p_i)$, where subscript i stands for firm i in the market and f is the density function of the prices. Since we assume that prices in different markets are independent, the log-likelihood corresponding to all markets will be the sum of the log-likelihoods for the markets.

The density of prices in a market can be found by taking the implicit derivative of equation (3.7) with respect to p . This gives

$$f(p) = \frac{\sum_{k=1}^N k d_k(p) (1 - F(p))^{k-1} + (p - r) \sum_{k=1}^N k \frac{\partial d_k(p)}{\partial p} (1 - F(p))^{k-1}}{(p - r) \sum_{k=2}^N k(k-1) d_k(p) (1 - F(p))^{k-2}},$$

where

$$\frac{\partial d_1(p)}{\partial p} = \begin{cases} - \int_{c_1}^{p-E[p]} g(p, c) dc & \text{if } p > c_1 + E[p] \\ 0 & \text{otherwise,} \end{cases}$$

and for $k = 2, 3, \dots, N$

$$\frac{\partial d_k(p)}{\partial p} = \begin{cases} - \int_{c_k}^{\min\{\frac{1}{k}(p-E[\min\{p_1, p_2, \dots, p_k\}]), c_{k-1}\}} g(p, c) dc & \text{if } p > k c_k + E[\min\{p_1, p_2, \dots, p_k\}] \\ 0 & \text{otherwise.} \end{cases}$$

We note that in the expression of the density we will substitute the distribution function $F(\cdot)$

computed from equation (3.7), which by Lemma 3.4 is determined uniquely. Regarding the estimation of the firms' marginal cost in a market, we use equation (3.8), which yields

$$r = \frac{p \sum_{k=1}^N k \mu_k - \bar{p} d_1(\bar{p})}{\sum_{k=1}^N k \mu_k - d_1(\bar{p})}.$$

We now make some remarks regarding the consistency of the maximum likelihood estimator. Since we use price data from several markets, it is likely that the distributions of these prices are different. Therefore, the conditions for consistency of the estimator are slightly different from the iid case. Hoadley (1971) studies the maximum likelihood estimator for non-identically distributed observations, so one option is to verify the conditions considered there.

The most challenging condition is identification. Since, as we will see, the bivariate valuation and search cost distribution function $G(\cdot, \cdot)$ cannot be identified nonparametrically, we discuss first the estimation and identification of a special case of the model, namely the model without valuation heterogeneity. Then we propose some flexible parametric assumptions for our general model presented above, and discuss the estimation of the valuation and search cost distribution function in that case.

3.5.1 The simple model: heterogeneous search costs, homogeneous valuations

It is useful to start the discussion on identification and estimation of the unknown parameters by considering a simpler model where all consumers have identical valuations. Such a model has been studied by Hong and Shum (2006) and is also studied in Chapter 2. In what follows we argue that the distribution of search costs cannot be identified in its full support using data from only one market. This problem, which of course carries over to the more general model, is dealt with in this paper by pooling data from multiple markets with similar search costs, which requires the use of a novel estimation technique in this literature.

Let $v > 0$ denote the common consumer valuation for the good. It is straightforward to see that the upper bound of the equilibrium price distribution must be equal to v and the equilibrium

equation (3.7) simplifies to

$$(p_i - r) \left[\sum_{k=1}^N \frac{k\mu_k}{N} (1 - F(p_i))^{k-1} \right] = \frac{\mu_1(v - r)}{N} \quad (3.12)$$

with lower bound

$$\underline{p} = \frac{\mu_1(v - r)}{\sum_{k=1}^N k\mu_k} + r. \quad (3.13)$$

Equation (3.12) can easily be inverted to obtain

$$p(z) = \frac{\mu_1(v - r)}{\sum_{k=1}^N k\mu_k(1 - z)^{k-1}} + r.$$

so that the critical search cost parameters can be computed from the following system of equations:

$$c_0 = v - \int_0^1 \left(\frac{(v - r)\mu_1}{\sum_{k=1}^N k\mu_k(1 - z)^{k-1}} + r \right) dz; \quad (3.14a)$$

$$c_k = \int_0^1 \left(\frac{(v - r)\mu_1}{\sum_{k=1}^N k\mu_k(1 - z)^{k-1}} + r \right) [(k + 1)z - 1](1 - z)^{k-1} dz, \quad k = 1, 2, \dots, N. \quad (3.14b)$$

The price density function is in this case

$$f(p) = \frac{\sum_{k=1}^N \frac{k\mu_k}{N} (1 - F(p))^{k-1}}{(p - r) \sum_{k=1}^N \frac{k(k-1)\mu_k}{N} (1 - F(p))^{k-2}}. \quad (3.15)$$

the log-likelihood function is $LL(\mathbf{p}) = \sum_i \log f(p_i)$.

Before discussing how to determine the density function explicitly we make some observations on it. First, the fraction of consumers μ_N can be eliminated from equation (3.15) by substituting it by $1 - \sum_{k=0}^{N-1} \mu_k$; here we assume that the econometrician observes μ_0 , i.e., the fraction of consumers who do not participate in the market. Second, we can solve equation

(3.13) for

$$r = \frac{p \sum_{k=1}^N k \mu_k - \mu_1 v}{\sum_{k=2}^N k \mu_k},$$

and eliminate it from the price density. Third, we can estimate \underline{p} and v superconsistently by taking the minimum price and the maximum price observed in the data, respectively. As a result, the only parameters unknown in the density function are $\{\mu_k\}_{k=1}^{N-1}$.

One possibility is to estimate these fractions of consumers μ_k directly by maximum likelihood. Using the ML estimates of these fractions of consumers, and the fact that

$$\mu_k = \int_{c_k}^{c_{k-1}} dG_c(c) = G_c(c_{k-1}) - G_c(c_k), \text{ for all } k = 1, 2, \dots, N-1; \quad (3.16a)$$

$$\mu_N = \int_0^{c_{N-1}} dG_c(c) = G_c(c_{N-1}). \quad (3.16b)$$

we can compute ML estimates of $\{c_k\}_{k=0}^{N-1}$ from equations (3.14a) and (3.14b). Once we have obtained ML estimates for the $\{\mu_k\}_{k=1}^N$ and $\{c_k\}_{k=0}^{N-1}$, we can construct a spline approximation estimate of the search cost distribution $G_c(c)$. This approach can only be applied if the econometrician observes the prices of the firms over an extended period of time, as shown in Chapter 2.⁷ Here we do not follow this approach because, as we show below, it cannot identify the search cost distribution in its entire support, not even if we let $N \rightarrow \infty$.

In this paper we employ a method that is different in essence and which, by pooling price data from different markets, allows us to identify the whole search cost distribution (for details we refer to the subsection on identification). More precisely, we estimate the search cost density function directly by a flexible polynomial-type parametric function. Then we compute $\{\mu_k\}_{k=1}^{N-1}$ and $\{c_k\}_{k=0}^{N-1}$ from the system of equations in (3.14a), (3.14b), (3.16a) and (3.16b) in terms of the parameters of this function, and using equation (3.12) and then equation (3.15) we express the $F(p_i)$'s and the $f(p_i)$'s, respectively, in terms of these parameters. This way we obtain the log-likelihood as a function of the parameters of the polynomial-type parametric function used

⁷Since there are $N-1$ parameters to estimate, observation of a single period cross-section of prices is not sufficient because after estimating the upper and the lower bound of the price distribution only $N-2$ data points are left. Hong and Shum (2006) propose to estimate only some of the μ'_k 's but this has the problem that the choice of μ'_k 's becomes arbitrary. Alternatively, in Chapter 2 we propose to study a finitely repeated version of the game. Since the static equilibrium is also the equilibrium of the repeated game, prices observed in a market during certain time horizon are iid draws from the same distribution.

to estimate the search cost density.

For the polynomial-type parametric function that estimates the search cost density we employ the so-called semi-nonparametric (SNP) density estimator (Gallant and Nychka, 1987). This SNP estimator is based upon a Hermite polynomial expansion. The idea of this is that any reasonable density can be mimicked by such a Hermite polynomial series. SNP density estimators are essentially nonparametric, just like the spline approximation method described above, because the set of all Hermite polynomial expansions is dense in the set of density functions that are relevant (Gallant and Nychka, 1987).

To apply the SNP estimation in our problem, we specify the search cost density as

$$g_c(c; \gamma, \sigma, \theta) = \frac{\left[\sum_{i=0}^{p_n} \theta_i u_i(c) \right]^2}{\sum_{i=0}^{p_n} \theta_i^2}, \theta \in \Theta_{p_n}, \Theta_{p_n} = \{\theta : \theta = (\theta_0, \theta_1, \dots, \theta_{p_n}), \theta_0 = 1\},$$

where p_n is the number of polynomial terms,

$$\begin{aligned} u_0(c) &= (c\sigma\sqrt{2\pi})^{-1/2} \exp(-((\log c - \gamma)/\sigma)^2/4), \\ u_1(c) &= (c\sigma\sqrt{2\pi})^{-1/2} ((\log c - \gamma)/\sigma) \exp(-((\log c - \gamma)/\sigma)^2/4), \\ u_i(c) &= \left[((\log c - \gamma)/\sigma) u_{i-1}(c) - \sqrt{i-1} u_{i-2}(c) \right] / \sqrt{i} \text{ for } i \geq 2. \end{aligned}$$

This parametric form corresponds to the univariate SNP estimator studied extensively by Fenton and Gallant (1996). Our expressions are obtained by transforming their random variable x with the density defined in their Section 4.3 to $c = \exp^{\gamma + \sigma x}$. This transformation is useful in our case since search costs are positive.

The vector of parameters to be estimated by maximum likelihood is $\{\gamma, \sigma, \theta_0, \theta_1, \dots, \theta_{p_n}\}$. The consistency of the maximum likelihood estimator can be established by verifying the conditions provided by Gallant and Nychka (1987) combined with the conditions from Hoadley (1971). A condition required by Gallant and Nychka (1987) is that the search cost density is differentiable and strictly positive on its support.

Identification of the simple model

In order to study the question whether the model can be identified, we ask whether the model provides sufficient information to recover the unknowns of interest given that we have full knowledge of the price distribution. This kind of treatment of the identification problem is in the spirit of Koopmans and Reiersøl (1950). It should be obvious by now that, since prices reflect only the behavior of the groups of consumers μ_k and not the behavior of individual consumers in this model, considering a market operated by a finite number of firms is not sufficient to identify the search cost distribution in its entire support. Because of this, in this subsection we consider markets with infinitely (but countably) many firms, and therefore, we have infinite sequences $(\mu_k)_{k \geq 1}$ of fractions of consumers who search for k firms and $(c_k)_{k \geq 0}$ of cutoff points.

We prove three results here. Our first result provides conditions for identifying the search cost distribution at the cutoff points c_k when we observe the price distribution from one market. More precisely, it says that if we know the price distribution $F(p)$, r and μ_0 then we can identify the height of the search cost distribution corresponding to the cut-off points c_k 's. This result is auxiliary to the next two results. The second result shows that we cannot identify the whole search cost distribution when we observe prices from only one market. The third result shows that, when we observe prices from several markets, we can identify the search cost distribution on the interval $[0, \sup c_0]$, where $\sup c_0$ is the supremum of the set of c_0 -cutoff points from all markets.⁸ We place the proofs in the Appendix to ease the reading.

Proposition 3.2 *Consider the simple model where all consumers have identical valuations v . Suppose that the triple of endogenous variables $(F, (\mu_k)_{k \geq 1}, (c_k)_{k \geq 0})$ is generated by the triple of exogenous variables (G_c, v, r) and that $(F', (\mu'_k)_{k \geq 1}, (c'_k)_{k \geq 0})$ is generated by (G'_c, v', r') , where G_c, G'_c are distribution functions with support $(0, \infty)$ and positive density on this support. Suppose also that F is a distribution function with support (\underline{p}, \bar{p}) and that $F' = F$. In addition, assume the conditions*

1. $r' = r$,
2. $\mu'_0 = \mu_0$.

⁸We note that, since the main reason for estimating a search cost distribution is to perform policy analyses, estimating the search cost distribution on the interval $[0, \widehat{\sup c_0}]$ is sufficient in most practical applications; here $\widehat{\sup c_0}$ denotes the estimate of $\sup c_0$ that one obtains from our estimation procedure.

Then $\mu'_k = \mu_k$, $c'_k = c_k$ and $G'_c(c_k) = G_c(c_k)$ for any $k \geq 0$.

Regarding the conditions we note that, in our belief, Condition 1 is not necessary. We adopt it here in order to make the problem analytically tractable (see the remark after the proof in the Appendix for some intuition). With respect to Condition 2 we note that it reflects the need, already mentioned above, that the econometrician must observe μ_0 .

Even though the search cost distribution can be identified at the cutoff points, our second result shows that identification of the whole search cost distribution using data from only one market is not possible. Intuitively, the reason is that the sequence of critical points $(c_k)_{k \geq 0}$ is convergent to zero so we cannot get identification at higher quantiles.

Proposition 3.3 *Consider the simple model where all consumers have identical valuations v . Suppose that $(F, (\mu_k)_{k \geq 1}, (c_k)_{k \geq 0})$ and $(F', (\mu'_k)_{k \geq 1}, (c'_k)_{k \geq 0})$ are generated by (G_c, v, r) and (G'_c, v', r') , respectively, where G_c, G'_c are distribution functions with support $(0, \infty)$ and positive density on this support. Suppose also that F and F' are distribution functions with supports (\underline{p}, \bar{p}) and $(\underline{p}', \bar{p}')$, respectively. In addition, assume the conditions*

1. $r' = r$ and $v' = v$,
2. $c'_k = c_k$ for any $k \geq 0$,
3. $G'_c(c_k) = G_c(c_k)$ for any $k \geq 0$.

Then $F' = F$.

Proposition 3.3 implies that the search cost distribution cannot be fully identified using price information from only one market. This is because conditions 2 and 3 allow for two search cost distributions G_c and G'_c , to be different for a non-negligible set of points. This is due to the fact that the model predicts that the sequence of cut-off points $(c_k)_k$ converges monotonically to zero and therefore it is not dense in the support $[0, \infty)$ of the search cost distribution.

This observation can be seen in Figure 3.4 where we plot the critical cutoff points c_k for different number of firms ($N = 10, 15, 50$ and 100). In these plots we set $v = 500$, $r = 50$ and assume consumer search costs follow a log-normal distribution with parameters $(\nu_c, \sigma_c) = (0.5, 5)$. The graphs illustrate how the sequence of critical search costs c_k is convergent to zero




Figure 3.4: Non-identification from only one market

so increasing the number of firms does not help much to get information of the height of search costs at high quantiles.

To overcome this problem we propose to consider a richer framework where the econometrician has price data from several markets. In particular, we consider markets where the consumer valuations are different but the search cost distribution is the same. Intuitively, this solves the problem of identification because every market generates a distinctive set of cutoff points, and this forces the search cost distribution to be uniquely determined for a larger set of search cost values.

More formally, let m be a market index. Assume that in each market m there are infinitely (countably) many firms. Assume also that there are infinitely (countably) many such markets.

Proposition 3.4 *Consider the simple model where all consumers in a market m have identical valuations v_m . Suppose that search costs have a distribution function G_c with support $(0, \infty)$ and that the conditions in Proposition 3.2 are satisfied so in each market m the values $G_c(c_{m,k})$ of the search cost distribution corresponding to the cut-off points $c_{m,k}$, $k \geq 0$, are identified. In*

addition assume that

1. *the valuations $\{v_m\}_{m \geq 1}$ are random variables drawn independently from a distribution with support $(0, \infty)$ and*
2. *the marginal costs $\{r_m\}_{m \geq 1}$ are random variables drawn independently from a distribution with support denoted S_r .*

Then G_c is identified on the interval $[0, \sup c_0]$, where $\sup c_0$ is the supremum of the set of c_0 -cutoff points from all markets.

The marginal cost of the firms will naturally vary across markets, but our identification result needs only variation in consumer valuation. This can be seen explicitly from Condition 2 on the marginal costs, because the support of their distribution S_r is not specified; it can be an interval, an infinitely countable set, a finite set, or even a set with a single element. The only requirement regarding the marginal costs is that they are independently drawn from a distribution with this support.

We note that in practice the situation that there are many firms in each market is not the only possibility to identify the search cost distribution nonparametrically. Another situation is when the markets have only a few firms, whose prices are independent realizations of their distributions and observed over many periods. This provides sufficient information to estimate the valuations in each market, and Proposition 3.4 ensures nonparametric identification of the search cost distribution.

It is also important to note that, in both situations, the number of firms in the market or the number of time periods over which we observe the firms need not be very large. This is because these observations serve the estimation of the valuation in the respective market, and this is done by taking the maximum of the prices observed. This estimator is known to converge in probability at a rate proportional to the number of observations, so for achieving a reasonable precision it is not necessary to have a very large number of observations.

The ideas put forward in Proposition 3.4 are illustrated in Figure 3.5, where we plot all the critical cutoff points c_k obtained from different markets, each of them operated by 10 firms. In these plots we set $r = 50$ and again assume consumer search costs follow a log-normal distribution with parameters $(\nu_c, \sigma_c) = (0.5, 5)$. We assume that the market with the highest consumer




Figure 3.5: Identification from multiple markets

valuation has a v_m which is 500 and that the market with the lowest consumer valuation has a v_m equal to 100. Furthermore, when there is only one market, we set $v_m = 500$, when there are two markets $v_m = \{100, 500\}$, when there are five markets $v_m = \{100, 200, 300, 400, 500\}$, and so on. The graphs illustrate how, when we increase the number of markets M from 1 to 5, 25 and 50, the set of critical search cost points c_k becomes denser and denser in the full support of the search cost distribution.

Given that we need to pool price data from multiple markets to identify and quantify search costs, the spline approximation method used earlier in the literature (and described above) cannot be used. The reason is that the spline approximation method uses a two-step procedure such that markets are not linked via the same underlying search cost distribution. In particular the set of μ_k 's in a given market m maximizes the likelihood of prices in that market and not necessarily the likelihood of all prices. The SNP estimator does not have this problem, because the parameters of the search cost distribution are chosen directly to maximize the log-likelihood of all markets together.

Simulations

The purpose of this section is to illustrate the identification and estimation of search costs in the simple model using price data. In particular we focus on the effects of using multiple markets rather than only one market when estimating search costs. To obtain price samples, we assume that the search cost density is a mixture of a log-normal and a gamma distribution, in particular

$$g_c(c) = 0.3 \cdot \text{lognormal}(c, 2, 5) + 0.7 \cdot \text{gamma}(c, 5, 2).$$

Table 3.3 describes the parameters chosen for the simulations. Our first experiment (EXP1) compares the estimated search cost distribution using data from a single market (EXP1A, $M = 1$) with that using data from six markets (EXP1B, $M = 6$). In every market there are five firms active ($N = 5$) and the marginal cost is $r = 50$. In the estimation with one market (EXP1A) we set $v = 70$, while in the estimation with six markets (EXP1B) v takes on values 70, 75, 80, 85, 90 and 95. The structure of experiments 2 (EXP2A versus EXP2B) and 3 (EXP3A versus EXP3B) is similar, but we take markets operated by a larger number of firms (15 and 25 respectively) and at the same time we increase the number of markets (16 and 26 respectively). Note that, to avoid biases related to the sample size, in each experiment the number of observations in the single-market case is the same as in the multiple-markets case. The performance of the estimates is evaluated by computing a measure of distance between the estimated search cost densities and the true density: $L_1 = \int_0^\infty |\hat{g}_c - g_c| dx$.

| | EXP1A | EXP1B | EXP2A | EXP2B | EXP3A | EXP3B |
|----------|----------|----------|----------|----------|----------|----------|
| M | 1 | 6 | 1 | 16 | 1 | 26 |
| N | 5 | 5 | 15 | 15 | 25 | 25 |
| $\#obs$ | 30 | 30 | 240 | 240 | 650 | 650 |
| r | 50 | 50 | 50 | 50 | 50 | 50 |
| v_1 | 70 | 70 | 70 | 70 | 70 | 70 |
| v_2 | - | 75 | - | 75 | - | 75 |
| v_3 | - | 80 | - | 80 | - | 80 |
| \vdots | \vdots | \vdots | \vdots | \vdots | \vdots | \vdots |
| v_M | - | 95 | - | 145 | - | 195 |

Table 3.3: Values for r and v

For our estimations, we set the number of polynomial terms in equation (3.12) equal to the integer closest to the fifth root of the total number of observations (Fenton and Gallant,

| | EXP1A | EXP1B | EXP2A | EXP2B | EXP3A | EXP3B |
|------------------|--------|--------|---------|---------|---------|----------|
| M | 1 | 6 | 1 | 16 | 1 | 26 |
| N | 5 | 5 | 15 | 15 | 25 | 25 |
| $\#obs$ | 30 | 30 | 240 | 240 | 650 | 650 |
| γ | 0.00 | 1.12 | 2.23 | 1.88 | 14.82 | 2.50 |
| σ | 0.57 | 1.32 | 1.76 | 1.15 | 9.06 | 1.25 |
| θ_0 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 |
| θ_1 | 0.81 | 0.58 | 0.16 | 0.15 | 7.97 | -0.33 |
| θ_2 | 0.83 | 0.38 | -0.03 | -0.21 | -0.93 | -0.28 |
| θ_3 | - | - | -0.38 | -0.43 | -3.47 | -0.11 |
| θ_4 | - | - | - | - | -1.21 | 0.43 |
| LL | 38.664 | 35.976 | 197.845 | 596.332 | 386.207 | 2016.254 |
| $L_1(0, \infty)$ | 1.552 | 0.801 | 0.932 | 0.453 | 1.362 | 0.409 |

Table 3.4: Parameter estimates

1996). For example, in our first experiment (EXP1) we estimate two additional polynomial terms, while experiment 3 has four extra polynomial parameters. Table 3.4 gives the estimates of the parameters of the SNP distribution for the different experiments. The table shows that the measure of fit is always lower when we use multiple markets than when we use a single market, which means that the fit with multiple markets is superior.

The goodness of fit can be seen in Figure 3.6, where we have plotted the estimated search cost cdf's and pdf's together with the true ones. In Figures 3.6(a) and 3.6(b), we present the cdf's and pdf's obtained with data from one market only. It can readily be seen that the fit of the estimated search cost cdf and pdf gets better as we increase the number of firms and the number of observations, but only for search costs less than 2, approximately. This is because the search cost distribution cannot be identified for larger search cost values.

In Figures 3.6(c) and 3.6(d), we present the cdf's and pdf's obtained with data from multiple markets. Like in the case of data from one market only, the fit of the estimated search cost cdf and pdf gets better as we increase the number of firms and markets. However, the crucial difference is that multiple markets allow us to identify the search cost distribution not only at low search cost values but also at high ones. This is reflected by the fact that the estimates of the unknown distributions are closer to the true ones in the entire support.

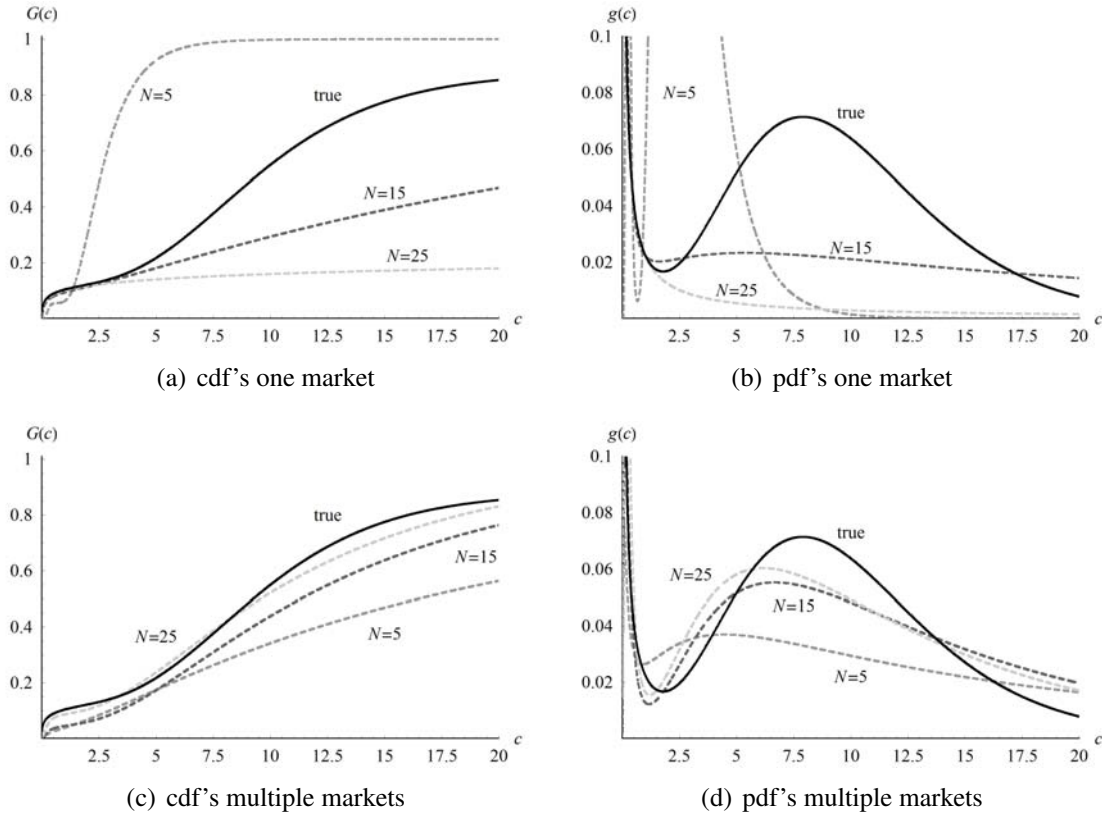


Figure 3.6: Estimated and true search cost cdf's and pdf's

3.5.2 The general model: heterogeneous search costs and valuations

In this section we turn our attention to the general model where there is search cost and valuation heterogeneity. We argue here that nonparametric identification of the joint valuation and search cost distribution function is not possible. The argument builds on the discussion of the identification of the simple model without valuation heterogeneity. For this simple model the information from the price distribution, which can be regarded as a continuum, is mapped to a discrete non-dense set of critical search cost cutoffs c_k 's. In order to identify the model we have shown we have to make up for the information lost by using price observations from many markets, which enables us to make the set of the critical search cost values dense. For the general model the situation is similar so by using price observations from many markets we can only make up for the information lost needed to estimate the search cost distribution, which is not sufficient for estimating the joint distribution of valuations and search costs.

If the econometrician suspects there is significant variation in consumer valuations, we pro-

pose to estimate the general model using a flexible parametric joint distribution for valuations and search costs. The parametric form is a bivariate SNP density function that corresponds to random variables transformed, as for the univariate nonparametric case described in Section 3.5.1, by the exponential so that

$$g(v, c) = \left[\sum_{i=0}^K \sum_{j=0}^K \frac{\alpha_{ij}}{\sqrt{\pi}} (\log v - \gamma_v)^i (\log c - \gamma_c)^j \right]^2 \frac{\exp \left[-\frac{(\log v - \gamma_v)^2}{2\sigma_v^2} - \frac{(\log c - \gamma_c)^2}{2\sigma_c^2} \right]}{2cv\sigma_c\sigma_v}.$$

This transformed SNP distribution inherits some computational advantages from the SNP distribution proposed by Gallant and Nychka (1987) because the corresponding probabilities can be computed as

$$\begin{aligned} \int_{a_v}^{b_v} \int_{a_c}^{b_c} g(v, c) dc dv &= \sum_{i,j,k,l=0}^K \frac{\alpha_{ij}\alpha_{kl}}{2\pi\sigma_c\sigma_v} \int_{a_v}^{b_v} \frac{(\log v - \gamma_v)^{i+k}}{v} \exp \left[-\frac{(\log v - \gamma_v)^2}{2\sigma_v^2} \right] dv \times \\ &\times \int_{a_c}^{b_c} \frac{(\log c - \gamma_c)^{j+l}}{c} \exp \left[-\frac{(\log c - \gamma_c)^2}{2\sigma_c^2} \right] dc, \end{aligned}$$

and the involved univariate integrals can be transformed to univariate Hermite integrals. Since it is an SNP density, the fact that we can increase the number of its parameters by taking a larger K allows for the estimation of a large parametric class of joint valuation and search cost densities. To ensure integration to 1 we scale the density by S , which is defined as

$$S = \int_0^\infty \int_0^\infty g(v, c) dc dv.$$

In this case, identification of the parameters of the joint density of valuations and search costs is a question of parametric identification of nonlinear models. As such, due to the complications caused by the nonlinear structure, it is very difficult to give an answer to this question. The standard intuitive reasoning, however, applies in our setting as well. If we observe a large number of prices, then equation (3.7) is satisfied for each of them, so we have a large number of equations with as many unknowns as the number of parameters, which appear in the expression of the d_k 's. Due to nonlinearity, the equations contain sufficient new information about the parameters, so these equations determine the parameters uniquely. The simulations presented below illustrate, among others, the parametric identification of the joint density of valuations and search costs.

Simulations

To illustrate the parametric identification of the joint density of valuations and search costs we generate price data using the SNP distribution function given above with $K = 1$, $\gamma_v = 5$, $\gamma_c = 2$, $\sigma_v = 0.3$, $\sigma_c = 1.25$, $\alpha_{00} = 1$, $\alpha_{01} = 1$, $\alpha_{10} = 0.1$ and $\alpha_{11} = 0.1$. Next we estimate the model using either 100 price observations (EXP1), 200 price observations (EXP2), or 1000 price observations (EXP3). Because the scale of the α 's is not determined we set $\alpha_{00} = \alpha_{01} = 1$ so the only α 's estimated are α_{10} and α_{11} . The parameter estimates are given in Table 3.5. In Figure 3.7(a) the true marginal search cost distribution is plotted together with the estimated marginal search cost distributions for the several experiments. Similarly, in Figure 3.7(b) the true marginal valuation distribution is shown together with the estimated marginal valuation distributions.

| | TRUE | EXP1 | EXP2 | EXP3 |
|---------------|-------|--------|-------|-------|
| N | 5 | 5 | 5 | 5 |
| M | - | 100 | 200 | 1000 |
| γ_v | 5.000 | 5.049 | 5.053 | 5.053 |
| γ_c | 2.000 | 2.044 | 1.964 | 1.955 |
| σ_v | 0.300 | 0.124 | 0.163 | 0.186 |
| σ_c | 1.250 | 1.172 | 1.260 | 1.245 |
| α_{00} | 1.000 | 1.000 | 1.000 | 1.000 |
| α_{01} | 1.000 | 1.000 | 1.000 | 1.000 |
| α_{10} | 0.100 | -0.622 | 0.090 | 0.483 |
| α_{11} | 0.100 | 0.742 | 0.080 | 0.012 |

(a) SNP parameters

| | TRUE | EXP1 | EXP2 | EXP3 |
|-----------------|--------|--------|--------|--------|
| N | 5 | 5 | 5 | 5 |
| M | - | 100 | 200 | 1000 |
| r | 20.00 | 7.92 | 16.91 | 15.84 |
| \underline{p} | 68.82 | 68.82 | 68.82 | 68.82 |
| \bar{p} | 135.51 | 135.51 | 135.51 | 135.51 |
| μ_0 | 0.414 | 0.286 | 0.319 | 0.340 |
| μ_1 | 0.440 | 0.549 | 0.511 | 0.522 |
| μ_2 | 0.066 | 0.091 | 0.085 | 0.083 |
| μ_3 | 0.011 | 0.010 | 0.012 | 0.015 |
| μ_4 | 0.001 | 0.002 | 0.001 | 0.001 |
| μ_5 | 0.068 | 0.062 | 0.073 | 0.071 |

(b) μ_k 's

Table 3.5: Estimation results

As can be seen from Figure 3.7(a), using 1000 price observations gives the best results, although all the estimated marginal cost distributions are quite close to true marginal distribution, even when only 100 prices are used for the estimation. The same is true for the marginal valuation distribution, although the differences are now somewhat bigger.

3.6 Conclusions

Since the seminal contribution of Stigler (1961), economists have dedicated a significant amount of effort to understand the nature of competition in markets where price information is not readily available to consumers. One of the lessons learned is that consumer search models may




Figure 3.7: True and estimated marginal search cost and valuation cdf

lead to predictions different than those from conventional economic theory. For example, in the model presented in this paper mergers of firms may lead to higher or to lower average prices in the market, the particular direction depending on the moments of the search cost distribution. This motivates the development of methods to identify and estimate search costs.

This paper studies the estimation of a class of nonsequential consumer search models, where symmetric firms sell homogenous products and consumers differ in their valuations and in their search costs. We find that this model does not allow for the nonparametric identification of the search cost distribution at all the points of its support when the econometrician observes prices from only one market. We propose a way to solve this pitfall, namely by considering a richer framework where the econometrician has price data from several markets with the same search cost distribution. Pooling price data from multiple markets enables us to identify nonparametrically the search cost density fully. To take advantage of the common search cost distribution underlying all the markets, we estimate the search cost density function by a semi-nonparametric (SNP) density estimator whose parameters maximize the joint likelihood corresponding to all the markets. This approach can be used when all consumers have identical valuations. If consumers differ in their willingness to pay but valuations are not correlated with search costs, still search costs can be identified nonparametrically. When there is correlation, only parametric identification obtains.

Our model has featured a market where symmetric firms sell homogeneous products. These two assumptions together imply that the equilibrium is in mixed strategies, with all the competitors drawing prices from the same price distribution. The fact that firms mix in equilibrium gives the econometrician sufficient information to estimate the model using only price data.

Relaxing these assumptions is an interesting avenue for further research. The drawback in practical applications is that firm heterogeneity should lead to pure-strategy equilibria, which would require the use of market share data for estimation.

APPENDIX: Proofs identification simple model

Here we summarize the problem of identification in the simple model and give the proofs of the propositions in Section 3.5.1. Identification seeks an answer to the question whether the model provides sufficient information to recover the unknowns of interest when we have full knowledge of the distribution of the observed variables. In the case of our model we study whether we can determine the search cost distribution G_c , the consumer valuation v and the firms' marginal cost r , when we know the price distribution.

For this, we consider infinitely many firms in the model. Then the model can be described as follows. The exogenous variables are the triplet (G_c, v, r) that generate the endogenous variables $(F, \{\mu_k\}_{k \geq 1}, \{c_k\}_{k \geq 0})$. In this section we maintain the assumption that these latter variables exist. They satisfy

$$\sum_{k \geq 1} k \mu_k (1 - F(p))^{k-1} = \mu_1 \frac{\bar{p} - r}{p - r} \quad \text{for any } p \in (\underline{p}, \bar{p}] , \quad (\text{A3.17a})$$

$$\bar{p} = v, \quad (\text{A3.17b})$$

$$\mu_k = G_c(c_{k-1}) - G_c(c_k) \quad \text{for any } k \geq 1, \quad (\text{A3.17c})$$

$$c_k = \int_{\underline{p}}^{\bar{p}} F(p) (1 - F(p))^k dp \quad \text{for any } k \geq 0. \quad (\text{A3.17d})$$

As before in the paper, here we also use the notation $\mu_0 = 1 - \sum_{k \geq 1} \mu_k$.

Proof of Proposition 3.2. As argued in the text, the upper bound of F must be equal to the consumer valuation, i.e., $\bar{p} = v$, so $v' = v$; by Condition 1, $r' = r$. First we show that $\mu'_k = \mu_k$ for any k . For this we note first that neither μ_1 nor μ'_1 can be equal to zero. If $\mu_1 = 0$ then by equation (A3.17a) $\sum_{k \geq 2} k \mu_k (1 - F(p))^{k-1} = 0$ for any $p \in (\underline{p}, \bar{p}]$, which, due to the fact that F is continuous, can only happen if $\mu_k = 0$ for any $k \geq 2$. This further implies by equation (A3.17c) that $G_c(c_k) = G_c(c_0)$ for any $k \geq 1$. By equation (3.4b), $c_n = e_n - e_{n+1}$, where $e_n = E[\min\{p_1, \dots, p_n\}]$. Since $e_n \geq e_{n+1}$ and $e_n \geq \underline{p}$ for any n , the series $(e_n)_n$ is convergent. Hence $c_n \rightarrow 0$ as $n \rightarrow \infty$. Because G_c is continuous in 0, $G_c(c_n) \rightarrow G_c(0) = 0$, so $G_c(c_0) = 0$. Because the density function corresponding to G_c is positive on $(0, \infty)$, this can only happen if $c_0 = 0$. But by equation (A3.17d) $c_0 = \int_{\underline{p}}^{\bar{p}} F(p) dp$, which is positive because F is a continuous cdf with support (\underline{p}, \bar{p}) , so we arrive at a contradiction. Since exactly the same arguments apply to μ'_1 , we have shown that μ_1 and μ'_1 are strictly positive.

From equation (A3.17a) we obtain

$$\sum_{k \geq 1} k \frac{\mu_k}{\mu_1} (1 - F(p))^{k-1} = \frac{v - r}{p - r} = \sum_{k \geq 1} k \frac{\mu'_k}{\mu'_1} (1 - F(p))^{k-1} \quad \text{for any } p \in (\underline{p}, \bar{p}].$$

This is equivalent to

$$\sum_{k \geq 2} \lambda_k t^{k-1} = 0 \quad \text{for any } t \in (0, \alpha), \quad (\text{A3.18})$$

where $\lambda_k = k \left(\frac{\mu_k}{\mu_1} - \frac{\mu'_k}{\mu'_1} \right)$ for $k \geq 1$ and $t = 1 - F(p)$. This latter transformation is possible because F is strictly increasing on some interval (\tilde{p}, \bar{p}) , where $1 - F(\tilde{p}) = \alpha$. We now refer to Lemma 3.5 below. This lemma implies that equation (A3.18) can only hold if $\lambda_k = 0$ for any $k \geq 2$. Therefore $\frac{\mu_k}{\mu_1} = \frac{\mu'_k}{\mu'_1}$. On the other hand, by Condition 2, $\mu_1 + \sum_{k \geq 2} \mu_k = \mu'_1 + \sum_{k \geq 2} \mu'_k = 1 - \mu_0$, which implies $\frac{1 - \mu_0}{\mu_1} = \frac{1 - \mu_0}{\mu'_1}$. Therefore $\mu'_k = \mu_k$ for any $k \geq 1$.

The equalities $c'_k = c_k$ follow from equation (A3.17d). It remains to show that $G'_c(c_k) = G_c(c_k)$ for any $k \geq 0$. We do so by showing that the $G_c(c_k)$'s for $k \geq 0$ are uniquely determined by the μ_k 's. By equation (A3.17c) $G_c(c_{k-1}) - G_c(c_k) = \mu_k$ for any $k \geq 1$. This implies that $G_c(c_0) - G_c(c_n) = \sum_{k=1}^n \mu_k$. The limit of the right hand side, when $n \rightarrow \infty$, exists and is $1 - \mu_0$. Therefore $G_c(c_0) - \lim_{n \rightarrow \infty} G_c(c_n) = 1 - \mu_0$. Because G_c is continuous in 0, $G_c(c_n) \rightarrow G_c(0) = 0$. Therefore $G_c(c_0) = 1 - \mu_0$ and $G_c(c_n) = 1 - \sum_{k=0}^n \mu_k$ for any $n \geq 1$. The result then follows from the equality of the μ_k 's established above. ■

Remark 1 Condition 1 is probably not necessary but we adopt it here for the simplicity of the proof. If this condition does not hold then equation (A3.17a) implies

$$\frac{p - r}{v - r} \sum_{k \geq 1} k \frac{\mu_k}{\mu_1} (1 - F(p))^{k-1} = \frac{p - r'}{v - r'} \sum_{k \geq 1} k \frac{\mu'_k}{\mu'_1} (1 - F(p))^{k-1} \quad \text{for any } p \in (\underline{p}, \bar{p}],$$

and this cannot be simplified to a power series identity as equation (A3.18). Still intuition suggests that the equalities $\mu'_k = \mu_k$ and $r' = r$ follow, since we may view this as a system of a continuum of equations with countably many unknowns r, r', μ_k, μ'_k for $k \geq 1$. ■

Proof of Proposition 3.3. Conditions 2 and 3 together with equation (A3.17c) imply that $\mu'_k = \mu_k$ for any $k \geq 1$. From the proof of Proposition 3.2 we know that $\mu_1 > 0$, so by equations

(A3.17a) and (A3.17b), and the first condition we have

$$\sum_{k \geq 1} k \mu_k (1 - F'(p))^{k-1} = \mu_1 \frac{v-r}{p-r} \quad \text{for any } p \in (\underline{p}', \bar{p}'] ,$$

without the right hand side being identically 0. Therefore $F'(p) = F(p)$ for any $p \in (\underline{p}', \bar{p}']$. This implies, by using the continuity of F , that $F(\underline{p}') = 0$ and $F(\bar{p}') = 1$ so $\underline{p}' \leq \underline{p}$ and $\bar{p} \leq \bar{p}'$. If now we interchange F and F' in this argument then we obtain $\underline{p} \leq \underline{p}'$ and $\bar{p}' \leq \bar{p}$, so $\underline{p}' = \underline{p}$ and $\bar{p} = \bar{p}'$. This implies that $F' = F$. ■

Proof of Proposition 3.4. In the proof we write $c_0(v, r)$ to make explicit the dependence of c_0 on v and r . The dependence of $c_0(v, r)$ on v and r is continuous because the dependence of the price distribution F on v and r is continuous and c_0 is given by equation (A3.17d). Take an arbitrary interval $(a, b) \subset \left(0, \sup_{\substack{v \in (0, \infty) \\ r \in S_r}} c_0(v, r)\right)$. Then the pre-image set defined as $c_0^{-1}(a, b) = \{(v, r) | c_0(v, r) \in (a, b)\}$ is a nonempty set, open in $(0, \infty) \times S_r$ because c_0 is a continuous function of v and r . Therefore, with probability 1, there exists an m such that $(v_m, r_m) \in c_0^{-1}(a, b)$, which means that $c_0(v_m, r_m) \in (a, b)$. Because the interval (a, b) has been chosen arbitrarily, we have proved that for any interval, with probability 1 we can find an m such that $c_0(v_m, r_m)$ is included in the interval. Hence the set $\{c_0(v_m, r_m) | m \geq 1\}$ is dense in $(0, \infty)$ with probability 1. We know that $G_c(c_{m,0}) = G_c(c_0(v_m, r_m))$, $m \geq 1$, are identified. Then, since it is continuous, G_c is identified on $\left(0, \sup_{\substack{v \in (0, \infty) \\ r \in S_r}} c_0(v, r)\right)$. ■

Lemma 3.5 (Power Series) Suppose that $(a_n)_{n \geq 1} \subset \mathbb{R}$ and $\sum_{n \geq 1} a_n x^n = 0 \quad \forall x \in (0, \alpha)$ for some $\alpha > 0$. Then $a_n = 0$ for any $n \geq 1$.

Proof. $\sum_{n \geq 1} a_n x^n = 0$ implies $a_1 + x \sum_{n \geq 0} a_{n+2} x^n = 0 \quad \forall x \in (0, \alpha)$. This can also be written as $\sum_{n \geq 0} a_{n+2} x^n = -\frac{a_1}{x} \quad \forall x \in (0, \alpha)$, which means that the power series $\sum_{n \geq 0} a_{n+2} x^n$ converges $\forall x \in (0, \alpha)$. Then by Lemma 3.6 below there exists $\rho \in (0, \alpha)$ such that $\sum_{n \geq 0} a_{n+2} x^n$ is uniformly convergent on $[-\rho, \rho]$. Let $p_1(x)$ be its limit, where $p_1 : [-\rho, \rho] \rightarrow \mathbb{R}$, that is, $\sum_{n \geq 0} a_{n+2} x^n = p_1(x) \quad \forall x \in [-\rho, \rho]$. Therefore

$$a_1 = -x p_1(x) \quad \forall x \in [-\rho, \rho]. \quad (\text{A3.19})$$

The function p_1 is continuous because it is the uniform limit of a sequence of continuous func-

tions, so $\lim_{x \rightarrow 0} p_1(x) = p_1(0) = a_2$. This further implies that $\lim_{x \rightarrow 0} x p_1(x) = 0$, so based on equation (A3.19), for any $\varepsilon > 0$ there is $\delta(\varepsilon) > 0$ such that $|a_1| = |x p_1(x)| < \varepsilon$ for any x with $|x| < \delta(\varepsilon)$. This implies that $a_1 = 0$.

So we have obtained that $\sum_{n \geq 2} a_n x^n = 0 \forall x \in (0, \alpha)$, which implies $\sum_{n \geq 2} a_n x^{n-1} = 0 \forall x \in (0, \alpha)$. By renaming the sequence $(a_n)_{n \geq 2}$ as $(b_n)_{n \geq 1}$ with $b_n = a_{n+1}$ we have $\sum_{n \geq 1} b_n x^n = 0 \forall x \in (0, \alpha)$. The arguments of the previous paragraph imply that $b_1 = 0$, that is, $a_2 = 0$. Going on this way we can show that $a_n = 0$ for any $n \geq 1$. ■

The following lemma is a version of a result also known as Abel's Uniform Convergence Test.

Lemma 3.6 (Abel) *Suppose that the series $\sum_{n \geq 0} a_n x_0^n$ is convergent. Then $\forall \rho$ with $0 < \rho < |x_0|$ the series $\sum_{n \geq 0} a_n x^n$ is uniformly convergent $\forall x \in [-\rho, \rho]$.*

Proof. Let y be arbitrary with $0 < |y| < |x_0|$. First we note that the convergence of the series $\sum_{n \geq 0} a_n x_0^n$ implies that $\lim_{n \rightarrow \infty} a_n x_0^n = 0$ and therefore there exists M with $|a_n x_0^n| < M \forall n$. The sequence $b_n = \sum_{k=0}^n |a_k| |y|^k$ is convergent because it is increasing and

$$\sum_{k=0}^n |a_k| |y|^k = \sum_{k=0}^n |a_k| |x_0|^k \frac{|y|^k}{|x_0|^k} < M \sum_{k=0}^n \left| \frac{y}{x_0} \right|^k \leq \frac{M}{1 - \left| \frac{y}{x_0} \right|} \quad \forall n,$$

that is, $(b_n)_n$ is bounded above. Let $b = \lim_{n \rightarrow \infty} b_n = \sum_{k \geq 0} |a_k| |y|^k$. Then the sequence $\sum_{k \geq n+1} |a_k| |y|^k = b - b_n$, and hence it converges to 0.

In particular, by taking $y = \rho$ we have obtained that $\sum_{k \geq n+1} |a_k| \rho^k$ converges to 0 for arbitrary ρ with $0 < \rho < |x_0|$ and by taking $y = |x|$ we have obtained that $\sum_{k \geq 0} |a_k| |x|^k$ is convergent for $\forall x \in [-\rho, \rho]$. This latter statement means that the series $\sum_{k \geq 0} a_k x^k$ is absolutely convergent and hence convergent for $\forall x \in [-\rho, \rho]$. So we can write

$$\sup_{x \in [-\rho, \rho]} \left| \sum_{k \geq 0} a_k x^k - \sum_{k=0}^n a_k x^k \right| = \sup_{x \in [-\rho, \rho]} \left| \sum_{k \geq n+1} a_k x^k \right| \leq \sup_{x \in [-\rho, \rho]} \sum_{k \geq n+1} |a_k| |x|^k \leq \sum_{k \geq n+1} |a_k| \rho^k.$$

Since the right hand side goes to 0 as $n \rightarrow \infty$, we have obtained that $\sum_{k=0}^n a_k x^k$ converges to $\sum_{k \geq 0} a_k x^k$ uniformly for $x \in [-\rho, \rho]$. ■

4

Truly costly sequential search and oligopolistic pricing

NOTE: This chapter is based on Janssen, Moraga-González, and Wildenbeest (2004, 2005).

4.1 Introduction

In this chapter and the next chapter we move from nonsequential search to sequential search. The main difference between the two search protocols is that with sequential search consumers decide *after* each search whether to continue searching or not, while with nonsequential search consumers determine how many times to search *before* they start searching. As shown in Morgan and Manning (1985), which of the two is optimal depends on the market studied.

This chapter is more theory oriented than the other chapters in the thesis, although in the end of this chapter it will be briefly discussed how to estimate the model presented. The starting point of this chapter is a celebrated article by Stahl (1989). In this chapter oligopolistic pricing in the presence of sequential consumer search is being studied. There are two types of consumers in the market. Fully informed consumers (referred to as *shoppers* in his article) have no opportunity cost of time and thus search for all prices at no cost; non-shoppers search sequentially, i.e., they first observe one price and then decide whether or not to observe a second price, and so on.¹ Stahl (1989) assumes that consumers observe the first price quotation for free, as do many other papers in the search literature, which implies that every buyer makes at least one search. In the first part of this chapter, we study the implications of relaxing this assumption.

The optimal sequential search rule implies that a consumer with a price at hand continues searching if, and only if, the observed price is higher than a certain reservation price. Knowing this, no firm will charge prices above consumers' reservation price. Therefore, under the assumption that obtaining the first price observation is costless, buyers 'search' exactly once in equilibrium and buy at the observed price. In this chapter, we refer to this type of equilibrium as one with *full consumer participation*. This equilibrium is one of the two possible equilibrium configurations when the first price quotation is *not* for free. The new type of equilibrium that arises with truly costly search is one with *partial consumer participation*, where some buyers decide not to search at all as they rationally expect prices to be so high that they are indifferent between searching and not searching. The existence and characterization of this new type of equilibrium is one main contribution of this chapter.

Another main contribution is to provide the comparative statics properties of the equilibrium with partial consumer participation. These comparative statics effects differ from those under

¹The functioning of markets in the presence of sequential consumer search is also examined in Anderson and Renault (1999), Reinganum (1979), Rob (1985) and Stahl (1996).

full consumer participation in interesting ways. First, the equilibrium distribution of prices with a given search cost dominates in a first-order stochastic sense the price distribution with a lower search cost; as a result expected price increases as search cost decreases. This is due to the fact that a decrease in search cost raises participation of non-shoppers, who happen to search only once in equilibrium. As firms have monopoly power over these consumers, they raise their prices. A second result is that an increase in the number of shoppers does not influence the equilibrium price distribution. This is because more shoppers foster the participation of non-shoppers in such a way that prices remain the same. Finally, we find that firm entry results in a mean-preserving spread of prices and in a decrease in welfare because the market participation rate of non-shoppers falls. The last two results imply that, unlike in Stahl's model, expected price does not tend to the monopoly price when the number of shoppers converges to zero, nor when the number of firms goes to infinity. This is because when the fraction of shoppers becomes very small, or the number of firms very large, the economy turns into an equilibrium with partial consumer participation and in such an equilibrium expected price is insensitive to changes in those parameters.

A third contribution is to provide a maximum likelihood method to estimate the model. The estimation procedure exploits the structure of the model as in previous chapters. We show that one cannot easily distinguish between the full and the partial consumer participation equilibrium empirically. As a matter of fact, to tell whether one equilibrium is more plausible than the other, one would need to have a priori knowledge of the share of consumers not participating in the market. We apply the estimation method to the same database of prices as in Chapter 2. The estimates indicate that the model does not do very well in explaining observed pricing patterns. Seen together with the results in Chapter 2, this is somewhat surprising because sequential search seems to be a more usual form of search online than nonsequential. This result, though striking, should be seen cautiously because the model of this chapter is quite restrictive in the sense that the search cost distribution is discrete and with only two atoms. Generalizing this model to the case of a continuous search cost distribution is the subject matter of the next chapter.

The rest of this chapter is organized as follows. Section 4.2 presents the model. A full characterization and an overview of the two types of equilibrium are given in Section 4.3. Section 4.4 presents the different comparative statics results. Section 4.5 presents the estimation pro-

cedure. In Section 4.6 the method is applied to a data set of online prices. Finally, Section 4.7 concludes.

4.2 The model

We examine the model of oligopolistic competition and sequential consumer search presented in Stahl (1989), but we assume that *all* price quotations are costly to obtain for non-shoppers. The features of the model are as follows. There are N firms that produce a homogeneous good at constant returns to scale. Their identical unit cost can be normalized to zero and prices can be interpreted as price-to-cost margins.² There is a unit mass of buyers and we assume that buyers hold inelastic demands.³ A consumer wishes to purchase at most a single unit of the good and his/her valuation for the item is $v > 0$. A proportion $\mu \in (0, 1)$ of the consumers has zero opportunity cost of time and therefore searches for prices costlessly. These consumers are referred to as *shoppers*. The other $1 - \mu$ percent of the buyers, referred to as *non-shoppers*, must pay search cost $c > 0$ to observe every price quotation they get, including the first one. Non-shoppers search sequentially, i.e., a buyer first decides whether to sample a first firm or not and then, upon observation of the price of the first firm, decides to search for a second price or not, and so on. We assume that $v > c$.

Firms and buyers play the following game. An individual firm chooses its price taking price choices of the rivals as well as consumers' search behavior as given. Likewise, an individual buyer forms conjectures about the distribution of prices in the market and decides on his/her optimal search strategy. We restrict the analysis to symmetric Nash equilibria. The distribution of prices charged by a firm is denoted by $F(p)$, its density by $f(p)$ and the lower and the upper bound of its support by \underline{p} and \bar{p} , respectively.

²In Section 4.5 we relax this assumption by introducing a identical unit cost r .

³Stahl (1989) considers a more general specification of the demand function. The assumption of inelastic demand allows us to compute explicitly the reservation price and give a full characterization of which type of equilibrium exists for which configurations of parameters. Provided that consumer surplus at the monopoly price does not fully cover the search cost, the main qualitative results of this chapter do not depend on the assumption of inelastic demand.

4.3 Equilibrium analysis

We first derive some auxiliary results.

Lemma 4.1 *An equilibrium where non-shoppers do not search at all does not exist.*

Proof. Suppose non-shoppers did not search. Then, the only consumers left in the market would be the shoppers. Therefore, competition between stores would drive prices down to marginal cost. But then, as $v - c > 0$, the non-shoppers would gain by deviating and searching once. ■

Lemma 4.1 reveals that existence of equilibrium requires the non-shoppers to be active in the market with strictly positive probability. The next result is provided by Stahl (1989).

Lemma 4.2 *In equilibrium non-shoppers will not search beyond the first firm.*

Proof. See Lemma 2 of Stahl (1989). ■

The idea behind Lemma 4.2 is that pricing above consumers' reservation price is never optimal for firms since buyers would continue searching if that were the case; as a result, the price buyers find at the first store they encounter is always accepted and no further search takes place.

Let us introduce the following notation. Let θ_1 be the probability with which a non-shopper searches once. Lemmas 4.1 and 4.2 together imply that only two candidates for equilibrium exist: either (a) $\theta_1 = 1$, or (b) $0 < \theta_1 < 1$. The first case is similar to Stahl (1989). We shall refer to this equilibrium as one with *full consumer participation*. This is because if all consumers search once, they will all buy the good. This contrasts with case (b) where consumers mix between not searching at all and searching once so not all consumers enter the market. In this case, we will speak of *partial consumer participation*.

The next remark is that, since $\theta_1 > 0$ in any equilibrium, the equilibrium price distribution must be atomless.

Lemma 4.3 *Irrespective of the search behavior of non-shoppers, if $F(p)$ is an equilibrium price distribution, then it is atomless. Hence, there is no pure strategy equilibrium.*

Proof. See Lemma 1 of Stahl (1989). The proof extends straightforwardly to the case of partial consumer participation. ■

We note that firms have an incentive to charge low prices in order to attract all the shoppers but at the same time they also have an incentive to charge high prices to extract income from the consumers who do not compare prices. These two forces are balanced when firms randomize their prices. Lemma 4.3 shows that equilibria must necessarily exhibit price dispersion, and that firm pricing is always characterized by atomless price distributions. In what follows we shall examine the characterization and the existence of the different types of equilibrium.

Case a: Equilibrium with full consumer participation

Suppose that non-shoppers search for one price with probability 1, i.e., $\theta_1 = 1$. This is the case analyzed by Stahl (1989) with two modifications. First, as Stahl considers a more general demand structure, an explicit expression for the reservation price cannot be obtained. Second, as Stahl (1989) assumes the first price quotation to be for free, this full participation equilibrium exists for all values of the parameters in his model, but not in ours. We will explicitly define the parameter space for which an equilibrium with full consumer participation exists when the first price quotation is costly. These two modifications deserve a slightly extended analysis.

Under full consumer participation, the expected payoff to firm i from charging price p_i when its rivals choose a random pricing strategy according to the cumulative distribution $F(\cdot)$ is

$$\pi_i(p_i, F(p_i)) = p_i \left[\frac{1 - \mu}{N} + \mu(1 - F(p_i))^{N-1} \right]. \quad (4.1)$$

This profit expression is easily interpreted. Firm i attracts the μ shoppers when it charges a price that is lower than its rivals' prices, which happens with probability $(1 - F(p_i))^{N-1}$. The firm also serves the $1 - \mu$ non-shoppers whenever they visit its store, which occurs with probability $1/N$.

In equilibrium, a firm must be indifferent between charging any price in the support of $F(\cdot)$. Let us denote the upper bound of $F(\cdot)$ by \bar{p} . Any price in the support of $F(\cdot)$ must then satisfy $\pi_i(p_i, F(\cdot)) = \pi_i(\bar{p})$, i.e.,

$$p_i \left[\frac{1 - \mu}{N} + \mu(1 - F(p_i))^{N-1} \right] = \frac{(1 - \mu)\bar{p}}{N}. \quad (4.2)$$

Solving this equation for the price distribution yields

$$F(p) = 1 - \left(\frac{(1 - \mu)(\bar{p} - p)}{N\mu p} \right)^{\frac{1}{N-1}}. \quad (4.3)$$

Since $F(\cdot)$ is a distribution function there must be some \underline{p} for which $F(\underline{p}) = 0$. Solving for \underline{p} one obtains the lower bound of the price distribution $\underline{p} = (1 - \mu)\bar{p}/(\mu N + (1 - \mu))$.

The cumulative distribution (4.3) represents optimal firm pricing. We now turn to discuss optimal consumer behavior. Consider a buyer who has observed a given price p . This consumer will continue to search if the expected benefits from searching further exceed the search cost. We can define the reservation price ρ as the price that makes a consumer indifferent between searching once more and accepting the price at hand; this price satisfies:

$$\int_{\underline{p}}^{\rho} (\rho - p)f(p)dp = c. \quad (4.4)$$

No firm will charge a price above ρ since this will lead to continued search (Stahl, 1989). As a result the upper bound $\bar{p} = \rho$. We now derive an expression for ρ . First rewrite equation (4.4) as

$$\rho - E[p] - c = 0. \quad (4.5)$$

To calculate $E[p]$ we solve equation (4.3) for p , which gives

$$p = \frac{\rho}{1 + bN(1 - F)^{N-1}}, \quad (4.6)$$

where $b = \mu/(1 - \mu) > 0$. Note that $E[p] = \rho - \int_{\underline{p}}^{\rho} F(p)dp$. By changing variables we can write $E[p] = \int_0^1 p dy$. Plugging p from equation (4.6) gives, after rearranging,

$$E[p] = \rho \int_0^1 \frac{dy}{1 + bNy^{N-1}}. \quad (4.7)$$

Equation (4.7) can be plugged into equation (4.5) to solve for ρ :

$$\rho = \frac{c}{1 - \int_0^1 \frac{dy}{1 + bNy^{N-1}}}, \quad (4.8)$$

It can be shown that the reservation price ρ increases in c and in N , decreases in μ and is insensitive to v .

It must be the case that $\rho \leq v$. In addition, non-shoppers must find it profitable to search once, rather than not searching at all, i.e.,

$$v - E[p] - c \geq 0. \quad (4.9)$$

Inspection of (4.5) and (4.9) implies that $\rho \leq v$. It is useful to rewrite condition (4.9) as

$$1 - \int_0^1 \frac{dy}{1 + bNy^{N-1}} \geq \frac{c}{v}. \quad (4.10)$$

This equation gives the set of parameters for which an equilibrium where buyers search once for sure exists. For future reference, let us denote the left-hand-side of equation (4.10) as $\Phi(1; \mu; N)$.⁴ We note that $0 < \Phi(1; \mu; N) < 1$ for all values of the parameters.

Proposition 4.1 *Let $0 < \frac{c}{v} \leq \Phi(1; \mu; N)$. Then a market equilibrium with full consumer participation exists where firms prices are distributed according to equation (4.3) on the set $[(1 - \mu)\rho/(\mu N + (1 - \mu)), \rho]$ and all non-shoppers search once, where $\rho = \rho$ and ρ solves equation (4.8).*

From equation (4.8) it follows that the reservation price ρ converges to v when c/v approaches $\Phi(1; \mu; N)$. The question that arises is: What happens when search cost is high, in particular when $\frac{c}{v} > \Phi(1; \mu; N)$? In what follows we show that the symmetric equilibrium involves partial consumer participation. This type of equilibrium is new in the sequential search literature and, as we shall see later, its properties are quite interesting.

Case b: Equilibrium with partial consumer participation

Now suppose that non-shoppers randomize between searching once and not searching at all, i.e., $0 < \theta_1 < 1$. The expected payoff to firm i is

$$\pi_i(p_i, F(p_i)) = p_i \left[\frac{(1 - \mu)\theta_1}{N} + \mu(1 - F(p_i))^{N-1} \right]. \quad (4.11)$$

⁴The number 1 in the arguments of $\Phi(\cdot)$ stands for $\theta_1=1$.

The economic interpretation of this profit function is analogous to that of equation (4.1), except that now there are only $(1 - \mu)\theta_1$ non-shoppers active, rather than $1 - \mu$. A similar analysis as above yields the following equilibrium price distribution:

$$F(p) = 1 - \left(\frac{\theta_1(1 - \mu)(\bar{p} - p)}{N\mu p} \right)^{\frac{1}{N-1}}, \quad (4.12)$$

with support $[\underline{p}, \bar{p}]$ where $\underline{p} = (1 - \mu)\theta_1\bar{p}/(\mu N + (1 - \mu)\theta_1)$. We now notice that the upper bound is no longer equal to ρ , but equal to v . To see this, note that non-shoppers optimal behavior requires that they are indifferent between searching once and not searching all, i.e., it must be the case that

$$v - E[p] - c = 0 \quad (4.13)$$

It is obvious that conditions (4.5) and (4.13) can only hold together if $\bar{p} = v$. Condition (4.13) can be rewritten as:

$$1 - \int_0^1 \frac{dy}{1 + \frac{1}{\theta_1} b N y^{N-1}} = \frac{c}{v}. \quad (4.14)$$

where $b = \mu/(1 - \mu)$. For future reference, denote the left-hand-side of equation (4.14) as $\Phi(\theta_1, \mu, N)$. Inspection of this function reveals that $\Phi(0, \mu, N) = 1$ and that $\Phi(\theta_1, \mu, N)$ is monotonically decreasing in θ_1 and increasing in μ . It can also be shown that $\Phi(\theta_1, \mu, N)$ decreases in N .

Proposition 4.2 *Let $\Phi(1, \mu, N) < \frac{c}{v} < 1$. Then a market equilibrium with partial consumer participation exists where firms prices are distributed according to equation (4.12) on the set $[(1 - \mu)\theta_1 v/(\mu N + (1 - \mu)\theta_1), v]$ and non-shoppers search once with probability θ_1 , which is the solution to equation (4.14) (with the remaining probability, non-shoppers stay out of the market).*

In summary, the game outlined above has a unique symmetric equilibrium where consumers may either search once surely, or mix between searching and not searching, depending on parameters. Inspection of the equations above immediately reveals that whether consumers participate in the market fully or partially depends on three critical parameters: (i) the value of the

purchase compared to the search cost c/v , (ii) the number of consumers with zero opportunity cost of time μ , and (iii) the number of firms N . To illustrate this issue we have represented the regions of parameters for which the equilibrium exhibits full or partial consumer participation in Figure 4.1. In these graphs we set $N = 2$ and vary μ . The left graph exhibits a market with many shoppers while the right one illustrates a market with just a few of them. The decreasing curve represents $\Phi(\theta_1; \mu, N)$ as a function of θ_1 . For large search cost parameters, say c_1 , non-shoppers participate in the market with probability less than one. This probability is given by the point at which the curve $\Phi(\cdot)$ and the line c_1/v intersect. When search cost is low enough, e.g. c_2 , non-shoppers search for one price with probability one.

The region of parameters for which there is partial consumer participation is larger the lower the parameter μ . This is because, as mentioned above, the function $\Phi(\cdot)$ falls as μ decreases (see equation (4.14)). A similar remark can be made when N increases. Indeed, when μ approaches 0 or N becomes very large, $\Phi(\theta_1; \mu, N)$ approaches 0 for all values of θ_1 . In that case the region for which consumers participate partially in equilibrium covers almost the entire parameter space. This indicates that partial consumer participation is relevant when the number of firms in the market is large and/or there are few shoppers.




Figure 4.1: Parameter regions for which distinct types of equilibrium exist ($N = 2$)

4.4 Comparative statics

In this section we study the influence of changes in the parameters of the model on search intensity θ_1 , on the equilibrium price distribution, on the average price charged in the market $E[p]$, and on welfare, denoted W . The main results are summarized in Table 4.1. The results

for the case of full consumer participation are similar to Stahl (1989); the others are new. An upwards (downwards) arrow means that the variable under consideration increases (falls); the symbol ‘—’ means that the variable remains constant. Our discussion shall concentrate on the most striking and interesting observations concerning the equilibrium with partial consumer participation.

| | Partial consumer participation | | | Full consumer participation | | |
|----------------|--------------------------------|------------|--------------|-----------------------------|--------------|------------|
| | θ_1 | $E[p]$ | W | ρ | $E[p]$ | W |
| $\downarrow c$ | \uparrow | \uparrow | \uparrow | \downarrow | \downarrow | \uparrow |
| $\uparrow \mu$ | \uparrow | — | \uparrow | \downarrow | \downarrow | \uparrow |
| $\uparrow v$ | \uparrow | \uparrow | \uparrow | — | — | \uparrow |
| $\uparrow N$ | \downarrow | — | \downarrow | \uparrow | \uparrow | — |

Table 4.1: Summary of comparative statics results

a. The effects of a reduction in search cost c

The first result we want to emphasize is that, under partial consumer participation, a reduction in search cost leads to an increase in expected price. This follows immediately from the equilibrium condition $v - E[p] - c = 0$. The intuition behind this result is simple. As Figure 4.1 shows, the intensity with which non-shoppers search in this type of equilibrium rises as c falls. Note further that these consumers are precisely those who do not exercise price comparisons, and thus they are prepared to accept higher prices. Consequently, a fall in c increases sellers’ incentives to charge higher prices. Indeed, inspection of equation (4.12) reveals that the distribution of prices with a high θ_1 (lower c) dominates in a first-order stochastic sense the price distribution with a low θ_1 (higher c).

As c decreases further, non-shoppers eventually start participating fully. In such a case, a decline in c results in a fall in the reservation price ρ (equation (4.8)). Inspection of the equilibrium price distribution in Proposition 4.1 reveals that F increases as ρ decreases. As a result, expected price decreases in c under full consumer participation.

These observations are illustrated in Figure 4.2. In Figure 4.2(a) we have simulated an economy where the number of informed consumers is large ($\mu = 0.8$). This graph shows that expected price is non-monotonic in relative search cost c/v . When there are few shoppers in the market ($\mu = 0.1$) expected price decreases in search cost for almost the entire parameter region

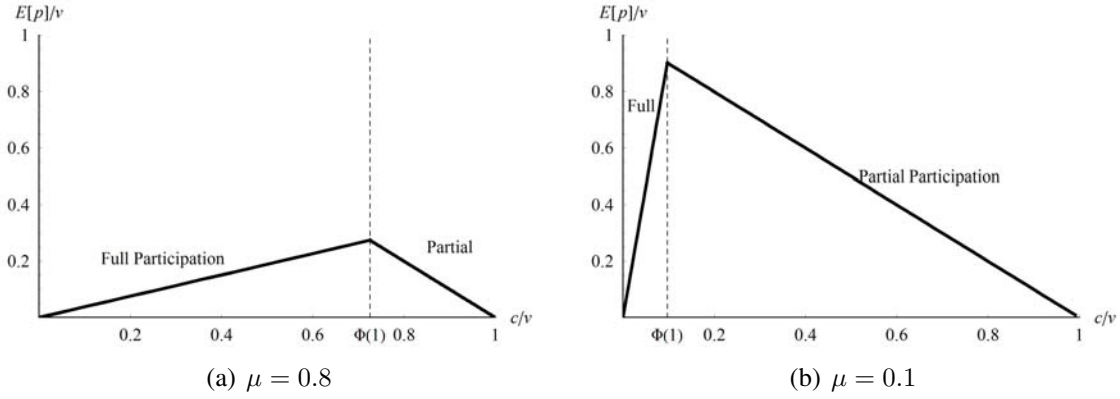


Figure 4.2: The influence of lower search cost c on expected price ($N = 2$)

(cf., Figure 4.2(b)). This is because for this parameter constellation the economy is most likely in an equilibrium with partial consumer participation.

Welfare is given by $W = \mu v + \theta_1(1 - \mu)(v - c)$ in this market, with $\theta_1 = 1$ in the full consumer participation case. A decrease in search cost c increases the surplus of the non-shoppers as well as their participation rate; as a result, welfare increases as search cost falls. Proposition 4.3 summarizes these findings. For this purpose, let $\varepsilon > 0$ be a small enough number so that if non-shoppers participate fully when the search cost is c , they also do it when the search cost increases to $c + \varepsilon$.

Proposition 4.3 *If non-shoppers participate partially (fully), the equilibrium distribution of prices with search cost c dominates (is dominated by) the price distribution with a higher search cost $c + \varepsilon$ in a first-order stochastic sense. As a result, expected price is non-monotonic in search cost. Moreover, welfare increases as search cost decreases.*

b. The effects of an increase in the fraction of shoppers μ

We next consider the effects of an increase in μ . Under partial consumer participation, a change in the number of shoppers does not influence expected price-to-cost margins as nothing changes in the equilibrium condition $v - E[p] - c = 0$. To understand the economic forces underlying this result, we first note that an increase in μ has in principle a pro-competitive effect. Keeping the search intensity of non-shoppers constant, firms would tend to charge lower prices as the number of shoppers in the market becomes higher. However, a change in μ also affects θ_1 . To

see how, one can apply the implicit function theorem to equation (4.14) to obtain

$$\frac{d\theta_1}{d\mu} = -\frac{\Phi'_\mu}{\Phi'_{\theta_1}} = \frac{\theta_1}{\mu(1-\mu)} > 0, \quad (4.15)$$

which means that an increase in μ results in an increase in the search intensity of the non-shoppers. This is because more informed consumers in the market makes searching more attractive for the non-shoppers, as the former buyers put pressure on firms to cut prices. A higher participation rate of non-shoppers in turn gives firms incentives to increase prices, since non-shoppers do not compare prices. Interestingly, these two opposite forces offset each other so that expected price remains constant. Indeed, it is easy to see that the entire distribution of prices (4.12) does not change in μ . The reason is that, using equation (4.15), the ratio of non-shoppers to shoppers $\theta_1(1-\mu)/\mu$ is constant in μ .

If non-shoppers search for one price for sure, the pro-competitive effects of an increase in μ mentioned above are strengthened by the fact that the reservation price ρ decreases in μ and thus expected prices fall (cf., Stahl (1989)). These remarks are illustrated in Figure 4.3. Figure 4.3(a) simulates an economy where product's valuation is relatively low compared to search cost ($c/v = 0.5$). The figure depicts expected price-to-cost margins as a function of μ . As $\Phi(1; \mu; N)$ is decreasing in μ , it easily follows that for a given N and c/v , there is a unique $\hat{\mu}$ such that condition (4.10) holds with equality. Therefore, the market equilibrium exhibits partial consumer participation when μ lies in the interval $(0, \hat{\mu})$ while full participation arises when μ lies in the interval $[\hat{\mu}, 1)$. Starting from full consumer participation, expected price increases as μ falls. As μ decreases further, the economy eventually moves into an equilibrium with only partial consumer participation. Even if $\mu \rightarrow 0$, expected price remains below v , so a Diamond type of result does not arise in a setting with truly costly search. In Figure 4.3(b) we have simulated an economy where search cost is relatively low ($c/v = 0.05$). The only difference is that the region of parameters for which consumers search once surely is much larger than before.

Welfare increases in μ in the full consumer participation case simply because consumers who incur search cost are replaced by zero search cost consumers. If non-shoppers participate only partially, an increase in μ raises the participation rate of non-shoppers, which implies that welfare also goes up in this case. These findings are summarized in the following result. As

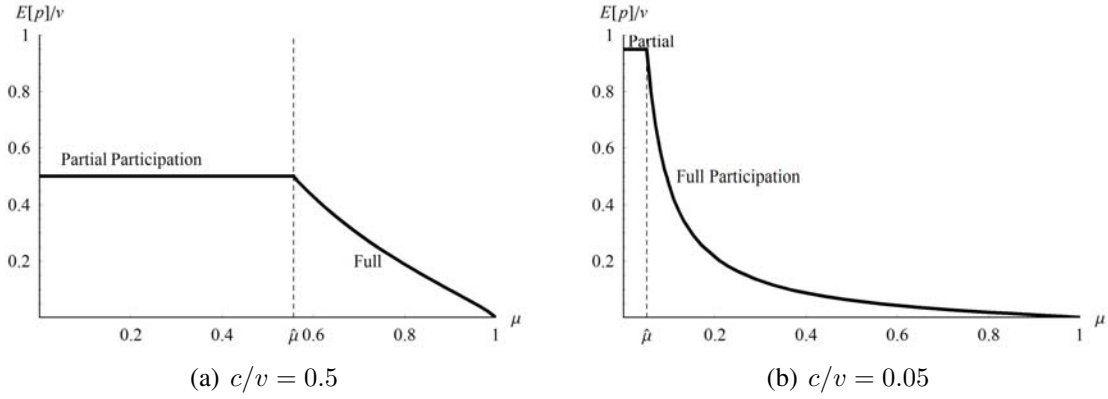


Figure 4.3: The influence of μ on expected price ($N = 2$).

above, let $\varepsilon > 0$ be a small enough number so that if non-shoppers participate partially when the fraction of shoppers is μ they also do it when the fraction of shoppers increases to $\mu + \varepsilon$.

Proposition 4.4 *If non-shoppers participate partially, an increase in the proportion of shoppers from μ to $\mu + \varepsilon$ leaves the equilibrium price distribution unchanged; by contrast, if non-shoppers participate fully, the price distribution with a higher proportion of shoppers $\mu + \varepsilon$ is dominated by the distribution with a fraction of shoppers μ in a first-order stochastic sense. As a result, expected price is weakly decreasing in μ . Moreover, welfare increases in the proportion of non-shoppers for all parameters.*

c. The effects of an increase in the number of firms N

Under partial consumer participation, it immediately follows from the equilibrium condition $v - E[p] - c = 0$ that an increase in N does not affect expected price. The economic forces underlying this result are, however, less straightforward. We have noted before that $\Phi(\theta_1; \mu; N)$ is decreasing in N . This means that without a change in θ_1 expected price would rise: the idea is that as N increases it becomes more and more unlikely that an individual firm sells to the shoppers and, consequently, it concentrates more and more on selling to the non-shoppers. A higher expected price means, however, that a larger fraction of non-shoppers prefers not to participate in the market and this effect exactly offsets the first effect.

We now notice that, under partial consumer participation the distribution of prices with $N + 1$ firms is a mean-preserving spread of the price distribution with N firms. This follows from the following two remarks. First, since the equilibrium θ_1 falls as a result of an increase

in N , the lower bound of the price distribution in Proposition 4.2 decreases; the upper bound does not change. Second, it is easy to see that there exists a unique value of p , denoted by \tilde{p} , such that $F(\tilde{p}; N, \theta_1^*(N)) = F(\tilde{p}; N+1, \theta_1^*(N+1))$, where $\theta_1^*(N)$ and $\theta_1^*(N+1)$ denote the non-shopper participation rates when there are N and $N+1$ firms in the market, respectively. From equation (4.12), this value is defined by the equality

$$\left(\frac{1 - \mu \frac{v - \tilde{p}}{\tilde{p}}}{\mu} \right)^{\frac{1}{N(N-1)}} = \frac{\left(\frac{\theta_1^*(N+1)}{N+1} \right)^{\frac{1}{N}}}{\left(\frac{\theta_1^*(N)}{N} \right)^{\frac{1}{N-1}}}. \quad (4.16)$$

As expected price remains constant, these two observations imply that $F(p; N+1, \theta_1^*(N+1))$ is a mean-preserving spread of $F(p; N, \theta_1^*(N))$.

Under full consumer participation Stahl (1989) shows that expected price rises in N . Intuitively, only the first effect discussed above is relevant here as θ_1 is fixed to be equal to 1. Stahl (1989) also shows that expected price converges to the monopoly price as $N \rightarrow \infty$ (Diamond type of result).

It is then interesting to see which type of equilibrium arises for which values of N . We find that for any given c/v and μ , the equilibrium is characterized by partial consumer participation when N is sufficiently large. What happens as N increases is that, if the non-shoppers keep searching once with probability one, expected price tends to the monopoly price and eventually, the condition $v - E[p] - c > 0$ is violated. This is easily seen upon inspection of condition (4.10) and noting that $\Phi(1; \mu, N)$ declines monotonically in N and converges to zero as N approaches infinity. This implies that starting from an equilibrium with full consumer participation, our model does not yield the Diamond result in the limit when $N \rightarrow \infty$ since at some point the economy turns to a situation of partial consumer participation. Figure 4.4 below illustrates these comparative statics results.

The effects of firm entry on welfare are straightforward. If non-shoppers participate fully, firm entry has no bearing on welfare. If non-shoppers participate only partially, an increase in the number of firms reduces the participation rate of these consumers, which reduces welfare. These findings are summarized in the following proposition.

Proposition 4.5 *An increase in the number of firms results in a mean-preserving spread of the price distribution if non-shoppers participate partially, while if they participate fully the price*

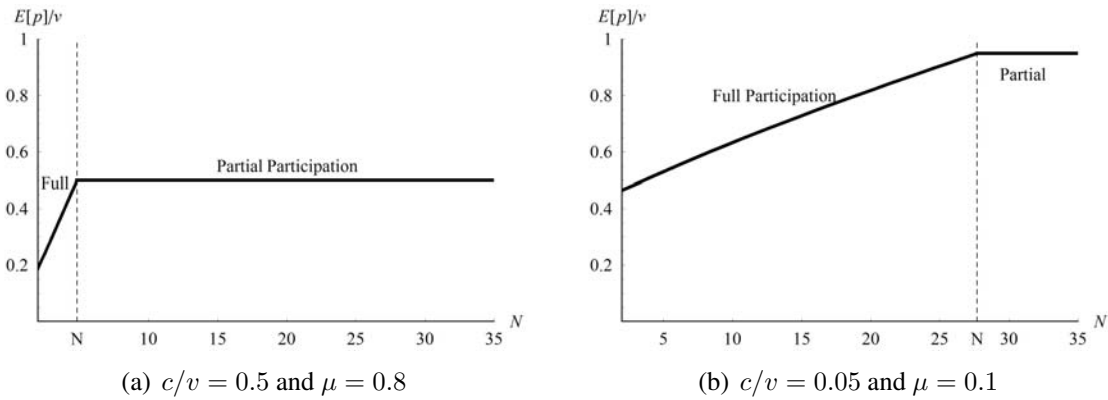


Figure 4.4: The impact of entry of firms on expected price

distribution shifts downwards; as a result, expected price is weakly increasing in N . Welfare decreases with firm entry under partial consumer participation while it remains constant if non-shoppers search once surely; as a result, welfare is weakly decreasing in N .

d. The effects of an increase in the value of the purchase v

We next briefly discuss the effects of changes in v . The main difference with the effects of a change in c is that, under full participation, c affects ρ whereas v does not affect ρ . As a result, the only difference with the discussion on the comparative statics effects of changes in search cost c is that now when buyers search for one price for sure, an increase in v does not alter price-to-cost margins.

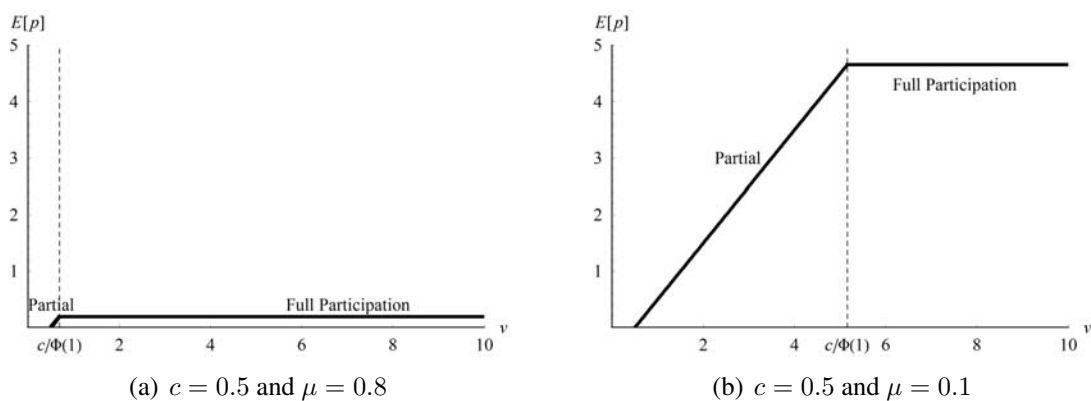


Figure 4.5: The impact of v on expected price ($N = 2$)

Figure 4.5 shows the influence of an increase in v on expected prices. In Figure 4.5(a) the number of informed consumers is relatively high ($c = 0.5$ and $\mu = 0.8$). When v lies in the

interval $(c, c/\Phi(1))$, non-shoppers participate in the market with probability less than one. In this parameter area, expected price rises as v increases. When v is above $c/\Phi(1)$ non-shoppers search once surely. For this region of parameters expected price is unaffected by a change in v . Figure 4.5(b) shows the case of an economy with relatively few shoppers.

4.5 Estimation

In this section we discuss how to estimate the model by maximum likelihood. To make the exercise empirically more meaningful, we introduce a marginal cost r by subtracting r from v , ρ and p , respectively. This changes the price distribution of the partial consumer participation equilibrium to

$$F(p) = 1 - \left(\frac{\theta_1(1-\mu)(\bar{p}-p)}{N\mu(p-r)} \right)^{\frac{1}{N-1}}. \quad (4.17)$$

Notice that setting $\theta_1 = 1$ gives the equilibrium with full consumer participation. Similarly, the lower bound becomes $\underline{p} = (1-\mu)\theta_1(\bar{p}-r)/(\mu N + (1-\mu)\theta_1) + r$, which can be solved for r to characterize r as a function of the parameters of the model, i.e.,

$$r = \underline{p} - \frac{(\bar{p}-\underline{p})\theta_1(1-\mu)}{\mu N}. \quad (4.18)$$

In order to estimate the different search equilibria, we derive the density functions associated with the equilibrium price distributions given in equations (4.3) and (4.12) by taking the derivative with respect to p . The density function is then

$$f(p) = \frac{1}{N-1} \left(\frac{\bar{p}-p}{p-r} \right)^{\frac{2-N}{N-1}} \left(\frac{\theta_1(1-\mu)}{N\mu} \right)^{\frac{1}{N-1}} \frac{\bar{p}-r}{(p-r)^2}, \quad (4.19)$$

where again setting $\theta_1 = 1$ gives the equilibrium with full consumer participation. Notice that although the price densities are different for each of the two equilibrium it is empirically impossible to distinguish between the two. This is easily seen by inspection of equation (4.19): θ_1 and μ can only be identified jointly as $\theta_1(1-\mu)/\mu$. This also means that empirically one cannot distinguish between the two equilibria without additional information on the share of consumers not entering the market. Therefore, in the remainder of this chapter we assume that

$\theta_1 = 1$, allowing us to focus only on the full consumer participation equilibrium. As shown in Section 4.3, in the full consumer participation equilibrium $\bar{p} = \rho$.

We set the number of firms N equal to the number of price observations found in the data. The objective is to maximize the log-likelihood function, which is given by $LL = \sum_{i=1}^N \log f(p_i)$. Because the support of the log-likelihood function depends upon unknown parameters, maximum likelihood estimation is not standard. Following Kiefer and Neumann (1993) we take the sampling minimum and maximum to estimate \underline{p} and ρ , respectively.⁵ These super-consistent estimates allow us to proceed as if \underline{p} and ρ were known. In addition, equation (4.18) can be plugged in the density function to eliminate r so that the only parameter left for estimation is μ .

We use the Broyden, Fletcher, Goldfarb and Shanno (BFGS) quasi-Newton method with a mixed quadratic and cubic line search procedure to find the μ that maximizes the log-likelihood function given N , the consistent estimates of \underline{p} and ρ and the vector of remaining prices, i.e., all the prices apart from the maximum and the minimum price. The BFGS quasi-Newton method uses the observed behavior of the log-likelihood function and its gradient to make an approximation to the Hessian matrix using the BFGS iterative updating technique. This method is implemented via a Matlab routine. We pick a random value for μ as a starting point and estimate the model several times to make sure that the outcome does not depend on the starting values.⁶ Given these estimates of the parameters we can derive the level of search cost c that is consistent with optimal consumer search using equation (4.8). The standard error for μ is estimated by calculating the analytic Hessian matrix, given the estimate of μ . Standard errors for r and c are calculated using the delta method.⁷

⁵More on using order statistics to estimate the lower and upper bound of distributions can be found in Donald and Paarsch (1993).

⁶These random starting value for μ , together with N , \underline{p} , ρ and a price p_i are plugged into the equilibrium condition (4.2) to find the value of $F(p_i)$ that solves this equation; this is done for all prices in the cross-section p_i . Subsequently, the calculated values of $F(p)$, together with N , \underline{p} , ρ , the vector of remaining prices and the starting value of μ are plugged into equation (4.19) to calculate the vector $f(p)$, from which we can calculate the log-likelihood value LL . The BFGS quasi-Newton method then comes with a new μ , from which we calculate a new value for LL . This process continues until the optimal μ is found, i.e., LL is maximized.

⁷If r is a function of the estimated μ , then the variance of r is given by $\frac{\partial r}{\partial \mu} \cdot (-H)^{-1} \cdot \frac{\partial r}{\partial \mu}'$, where H is the Hessian matrix.

4.6 Empirical application

A market most suitable in terms of the model of the previous section would consist of firms offering a homogeneous product to consumers who demand at most a single unit of the product. Keeping this in mind, we have chosen to focus on the online markets for memory chips. To make the empirical application comparable to previous chapters, the estimation method is applied to the same data set as in Chapter 2. Therefore, for an overview of the data the reader is referred to the discussion in Section 2.5.

The estimation results are presented in Table 4.2. As the table shows, estimates of the share of consumers having zero cost μ range between 0.75 and 0.85 and are highly significant. These consumers can be interpreted as shoppers, or as consumers using a price comparison site to find the lowest price around. We note that these estimates seem quite high on the basis of findings by other authors (although somewhat outdated, Whelan (2001) finds that around 36 percent of on-line shoppers can be classified as bargain hunters). Using the estimates of μ , we can calculate estimates of unit costs r and consumer search cost c . As Table 4.2 shows, in most cases unit cost r is close to the lower bound of the price distribution \underline{p} , which means that shops quoting low prices have quite small margins. The search cost that is consistent with the sequential search model, denoted c in the Table, varies between \$11.16 and \$13.89 and seems to be increasing in the value of the product.

The Kolmogorov-Smirnov test results show that for all four product the null hypothesis that the observed prices are generated from the sequential search model can be rejected. This is a striking result, especially given that the nonsequential search model presented in Chapter 2 is able to explain the data in a satisfactory manner. However, in this chapter we assume that consumers either have positive search cost, or have zero search cost, which is a lot more restrictive than the assumption made in Chapter 2 that consumers draw their search cost from some distribution. As we will show in the next chapter, if the sequential search model presented here is extended to a setting in which consumers draw their search cost from some distribution, the model performs much better.

Figure 4.6 gives an illustration of how the estimated theoretical distribution functions match observed price data. The dotted curves in the graphs represent the empirical distribution function of prices for the four different memory chips, while the solid curves give the estimated

| | KTT3614 | KTDINSP8200 | KTD4400 | KTD8300 |
|-----------------|---------------|---------------|--------------|---------------|
| \underline{p} | 115.00 | 109.20 | 96.00 | 102.00 |
| ρ | 208.90 | 200.50 | 170.50 | 182.50 |
| N | 25 | 24 | 24 | 23 |
| M | 179 | 174 | 174 | 162 |
| μ | 0.85 (0.05) | 0.75 (0.07) | 0.84 (0.05) | 0.84 (0.05) |
| r | 114.34 (0.26) | 107.93 (0.50) | 95.42 (0.23) | 101.34 (0.27) |
| c | 13.89 (1.61) | 12.53 (1.81) | 11.16 (1.34) | 12.33 (1.54) |
| LL | 993.53 | 1004.23 | 925.30 | 867.92 |
| KS | 9.77 | 10.29 | 9.76 | 9.20 |

Notes: Estimated standard errors in parenthesis. The 95%-critical value of the KS-statistic is 1.36.

Table 4.2: Estimation results

theoretical distribution functions. In line with the Kolmogorov-Smirnov test results, the graphs reveal that the model does not do very well in explaining observed pricing behavior.

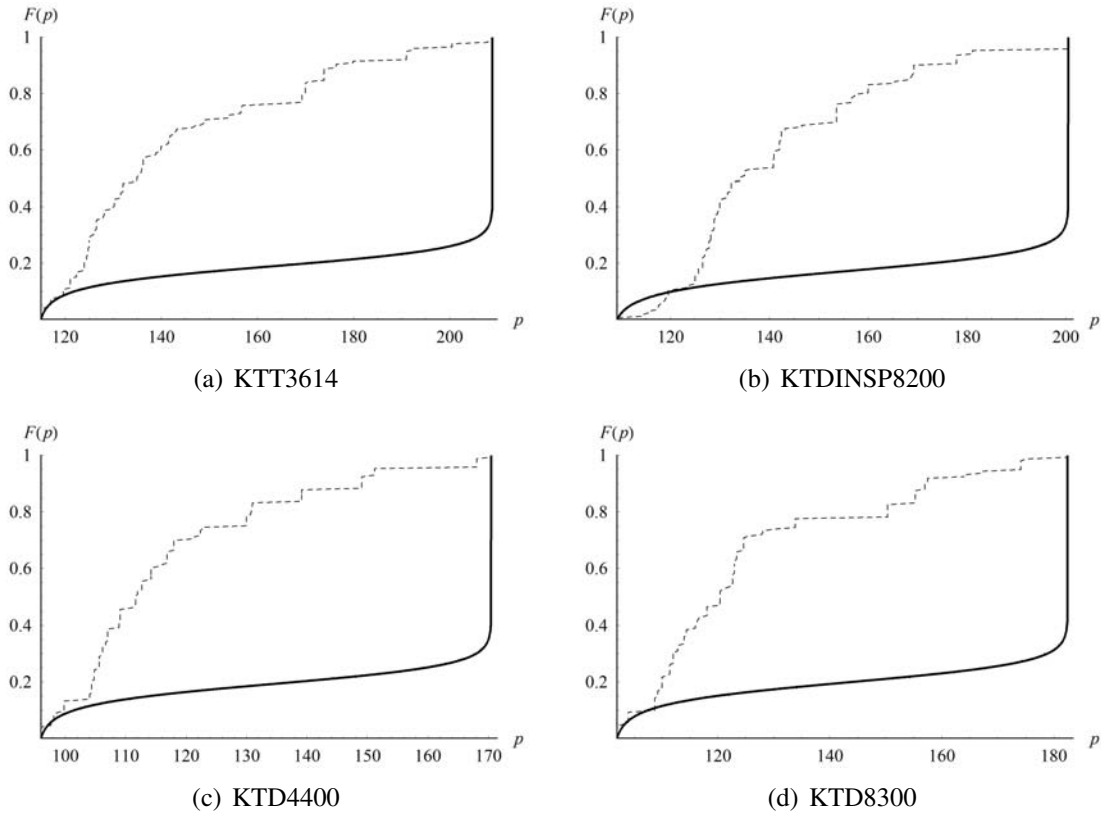


Figure 4.6: Estimated search cost cdf

4.7 Conclusions

In this chapter we have taken the seminal model of Stahl (1989) on sequential consumer search and oligopolistic pricing and studied the implications of relaxing the assumption that consumers obtain the first price quotation for free. When also the first price quotation is costly, the Nash equilibrium need not entail full consumer participation. Partial consumer participation arises when search cost is above a certain threshold value which depends on the other parameters. This threshold value becomes arbitrarily low as the number of firms becomes large enough and/or the number of shoppers sufficiently small. Therefore, especially in markets with many firms and/or with few shoppers, this situation of partial consumer participation should be seriously considered.

This new type of equilibrium exhibits interesting comparative statics properties. In particular, expected price increases as search cost decreases, and is constant in the number of shoppers and in the number of firms; welfare decreases as firms enter the market. This chapter also showed that, starting from an equilibrium with full consumer participation, monopolistic pricing never obtains when the number of shoppers goes to zero and/or the number of firms goes to infinity because with truly costly search the economy eventually turns into an equilibrium with partial consumer participation.

We have also shown how to estimate the model using maximum likelihood. Estimates indicate that the model does not very well in explaining empirical observed prices of memory chips sold online. A likely explanation for this result is that the assumption that there are only two types of consumers is too restrictive. The next chapter generalizes the model presented here to a setting in which there is a continuum of consumer types.

Estimation of a sequential search model

NOTE: This chapter is based on Wildenbeest (2007b).

5.1 Introduction

The search theoretic literature has studied a variety of models where consumers use different search protocols. For example, in the model of Varian (1980) consumers use newspapers to search for prices, in the model of Stahl (1989) consumers visit the shops sequentially and, finally, in the model of Burdett and Judd (1983) consumers search nonsequentially to find out the prices the firms charge. The most commonly studied search protocols are the sequential search protocol and the nonsequential search protocol. As demonstrated by Morgan and Manning (1985) neither of these protocols is superior universally; as a matter of fact, sequential search is optimal when price information can be gathered rather quickly because in that case the decision whether to search further at a given point in time can be conditioned on the observed prices. If offers are observed with some delay, as it would happen in the labor market for example, nonsequential search is more advantageous to the decision-makers. The main results of the theory, however, are robust: no matter whether consumers search sequentially or nonsequentially, prices are typically dispersed in equilibrium and the prices different consumers pay may behave in non-standard ways as a response to entry.

When it comes to empirical work, however, a sensible choice for an econometrician is to employ a model whose assumptions are most in line with the real-world practice. As mentioned in previous chapters, the literature on estimation of full-fledged search models is scant. The only papers we are aware of that estimate a sequential search model are Hortaçsu and Syverson (2004) and Hong and Shum (2006). Hortaçsu and Syverson (2004) model and estimate a framework in which search frictions coexist with vertical product differentiation as an explanation for price dispersion. Firms differ in marginal cost and quality level and play pure strategies in prices. Hortaçsu and Syverson (2004) show that search costs of consumers can be identified from aggregate price and quantity data.¹ In a related paper Hong and Shum (2006) show how to estimate a sequential search model using only price data. Hong and Shum (2006) present a method that relies on parametric assumptions on the shape of the search cost distribution, but show in addition that if one has marginal cost data it is possible to estimate search costs nonparametrically. A potential drawback of these two papers is that they assume that the distribution of realized prices is known by the consumer before searching. This assumption has

¹In the next chapter we show that in a slightly more restrictive setting it is possible to estimate search costs in the presence of vertical product differentiation using only price data.

been criticized (see e.g. Stahl, 1996) because it implies that consumers have a lot of information prior to searching. In markets with many firms the assumption may be innocuous because the theoretical and the empirical distribution of prices will be quite similar. However, in oligopolistic markets the market distribution of realized prices may be quite different than the theoretical distribution of prices.

The purpose of this chapter is to contribute to this recent literature by presenting a consumer sequential search model which can be estimated nonparametrically using a limited amount of data and which does not need to assume that realized prices are known in advance of search. Our paper is thus closely related to Hong and Shum (2006), but we study a more general model for which it is not necessary to assume that firms operate in an environment with infinitely many firms. The fact that we model an oligopoly game allows us to separate variation in prices due to changes in the number of firms to variation due to changes in search costs. Moreover, contrarily to what is assumed in both Hortaçsu and Syverson (2004) and Hong and Shum (2006), we do not assume consumers know realized prices before they start searching.

In our model a finite number of firms compete for consumers that have to pay a search cost to obtain price information. We derive a condition that implicitly defines the equilibrium price distribution and show how to use it to characterize the search cost distribution as a function of the price distribution, marginal cost and the number of firms. Then using the empirical price distribution search costs of consumers with reservation price equal to each observed price can be estimated as well as the corresponding quantiles of the search cost distribution.

We apply the estimation method to the same price data for several memory chips sold online as in Chapter 2. Our estimation results point towards relatively high search costs among consumers. More specifically, median search costs are estimated to be between 7.78 and 19.72 US dollars.

A question that arises is what sort of mistake would the researcher commit if she used the wrong type of model, i.e., if she used a sequential search model to estimate the relevant parameters when the data are actually generated from a nonsequential search strategy. To address this question, we compare the estimates obtained with the sequential search assumption to the estimates we obtained from the nonsequential search model of Chapter 2. For all memory chips, we find that estimated search costs are much higher in the sequential search model than in the nonsequential search model.

The chapter also presents a comparison between our estimation method and the nonparametric estimation method of the competitive sequential search model of Hong and Shum (2006). Also compared to that model, search costs are much higher using our method. A possible explanation for this is that in a competitive model, firms put more mass on high prices since the probability of serving the consumers with low reservation prices is low. Therefore, to rationalize the same set of prices, an oligopolistic model needs higher search costs than a competitive model.

The structure of the rest of this chapter is as follows. In the next section, we present our oligopolistic sequential search model. In Section 5.3 we study the characterization of a price dispersed symmetric equilibrium. In Section 5.4 we show how to estimate search costs nonparametrically using only price data. Moreover, we present some simulation results to illustrate the identification and estimation issues. In Section 5.5 we apply the estimation method to the market for online memory chips. Finally, Section 5.6 concludes.

5.2 The model

The model we study extends the standard sequential search model of Stahl (1989) by adding search cost heterogeneity. The model is similar in nature to Stahl (1996). The features of the model are as follows. On the supply side of the market, there are N firms each producing a homogeneous good at constant returns to scale. Let $r > 0$ be the common marginal cost. On the demand side of the market, there is a unit mass of buyers each demanding at most one unit of the good. Let v be the common valuation for the good. Consumers differ in their search costs. We assume that the search cost of a consumer is a random draw from a distribution $G(c)$, with density $g(c) > 0$, $c \in \mathbb{R}^+$. Consumers search sequentially. A consumer first decides whether to sample a first firm or not and then, upon observation of the price of the first firm, decides whether to search for a second price or not, and so on. For simplicity we assume that the first price quotation is obtained for free so if a consumer with search cost c searches k times she incurs a total costs $(k - 1)c$.

Firms and buyers play the following game. An individual firm chooses its price taking price choices of the rivals as well as consumers' search behavior as given. An individual buyer forms conjectures about the distribution of prices in the market and decides on her optimal

sequential search strategy. We confine ourselves to the analysis of symmetric Nash equilibria. The distribution of prices charged by a firm is denoted by $F(p)$, its density by $f(p)$ and the lower and the upper bound of its support by \bar{p} and \underline{p} , respectively.

5.3 Equilibrium

We start by discussing consumer search behavior. Consumers take the price distribution as given and decide whether to continue searching or not using a reservation price strategy. Let $F(p)$ be a (correct) conjecture about the price distribution. If a consumer has observed a price \hat{p} , the gains from searching one more time are

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

Integrating by parts, we can rewrite these gains from search as

$$H(\hat{p}; F) = \int_{\underline{p}}^{\hat{p}} F(p)dp. \quad (5.1)$$

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

$$H(\rho; F) - c = 0. \quad (5.2)$$

As it is well-known (see e.g. Weitzman, 1979), a consumer with search cost c will continue searching as long as $p > \rho(c; F)$, and will stop otherwise. Inspection of equation (5.1) reveals that the reservation price $\rho(c; F)$ is increasing in c .

We now move to calculate the profits of a firm i charging a price p_i . For this we follow Stahl (1996). Take a consumer at random whose search cost is c . There are two kinds of situation in which firm i sells to this consumer. The first situation is when the price of firm i induces no further search by this consumer, i.e., $p_i < \rho(c; F)$. In that case, the firm will sell to this consumer at any moment she encounters firm i during the course of her search. This may be in her first search, or in her second search, and so on till her last search. The other situation is when the consumer prefers to continue searching after observing the price of firm i , i.e., when

$p_i > \rho(c; F)$, but firm i happens to have the lowest price in the market so the consumer comes back to buy from firm i after visiting all the rival firms.

Let \tilde{c} be the search cost of the consumer who is indifferent between accepting the price of firm i and continue searching, i.e., $p_i = \rho(\tilde{c}; F)$. From equation (5.2), it follows that $\tilde{c} = H(p_i; F)$. Therefore, the profit to a firm i charging price p_i is:

$$\pi_i(p_i; F) = (p_i - r) \left[G(H(p_i; F)) (1 - F(p_i))^{N-1} + y(p_i; F) \right],$$

where the demand from the consumers with reservation values satisfying $\rho(c; F) < p_i$ is given by the first term between brackets and the demand from the consumers with $\rho(c; F) > p_i$ is given by

$$y(p_i; F) = \int_{H(p_i; F)}^{\infty} \frac{1}{N} \sum_{k=0}^{N-1} [1 - F(\rho(c; F))]^k g(c) dc. \quad (5.3)$$

We shall study only symmetric equilibria. Let \bar{p} denote the upper bound of the price distribution. Because it can never be optimal for a firm to set a price higher than the valuation v of the good, the optimal upper bound is given by $\bar{p} = \min\{v, p^M\}$, where the monopoly price $p^M = \arg \max_p (p - r)y(p; F)$, if it exists.² Furthermore, in equilibrium we require that $\pi_i(p_i; F) = \pi_i(\bar{p}; F)$ for all $p_i \in [\underline{p}, \bar{p}]$. This is the usual constancy-of-profits condition stating that a firm is indifferent between all the prices in the support of the price distribution. Because a firm setting a price equal to the upper bound only serves consumers with a reservation value higher than \bar{p} , this means that

$$(\bar{p} - r) \frac{1 - G(H(\bar{p}; F))}{N} = (p_i - r) \left[G(H(p_i; F)) (1 - F(p_i))^{N-1} + y(p_i; F) \right]. \quad (5.4)$$

The equilibrium price distribution is implicitly defined by this condition. Although the price distribution cannot be computed in closed-form, in the next section we show that if we evaluate the equilibrium condition together with equation (5.2) only at observed prices, we can express each search cost cdf value as a function of the observed prices, the number of firms and the

²The monopoly price p^M exists if, for an infinitely high valuation, a firm changing the highest price in the market does not gain by increasing its price. This is the case if we assume that $\lim_{\bar{p} \rightarrow \infty} \bar{p}(1 - G(H(\bar{p}; F))) = 0$. This is because $(\bar{p} - r)(1 - G(H(\bar{p}; F))) > (\bar{p} - r)y(\bar{p}; F)$ so if profits go to zero as $\bar{p} \rightarrow \infty$, by continuity, there must exist a finite \bar{p} which maximizes the profits at \bar{p} .

marginal cost. As a result, an important advantage of our method is that even though we study sequential search in oligopolistic setting, we do not need to evaluate any integral to estimate search costs nonparametrically.

5.4 Identification and estimation

The idea of our estimation procedure is as follows. If one did know the price distribution $F(\cdot)$, then one could use equation (5.2) to calculate all search cost values c that correspond to reservation values ρ in the support of $F(\cdot)$. Then the latent search cost density $g(c)$ in equation (5.3) should be chosen in such a way that the equilibrium condition in equation (5.4) holds. Unfortunately, $F(\cdot)$ is unknown, but we can proceed by estimating it from observed prices. The estimation of the search cost distribution then consists of two steps. In the first step we use the empirical price distribution in equation (5.2) to construct search cost values. In the second step we use the empirical price cdf to estimate the height of the search cost distribution corresponding to the search cost values obtained in the first step.

Assume we have a data set with M different prices. Let us collect them by ascending order, i.e., $p_1 < p_2 < \dots < p_M$. We can use the observed prices to estimate $F(p)$ nonparametrically by the empirical distribution, i.e., by

$$\tilde{F}(p) = \frac{1}{M} \sum_{i=1}^M \mathbb{1}(p_i < p),$$

where M is the total number of prices.

As mentioned above, since $F(\rho(c; F))$ depends on c there is no closed-form solution for equation (5.3). For the purpose of estimation it is therefore useful to rewrite $y(p; F)$ as

$$y(p_i; \tilde{F}) = \frac{1}{N} \sum_{j=i}^M \int_{H(p_j; \tilde{F})}^{H(p_{j+1}; \tilde{F})} \sum_{k=0}^{N-1} [1 - \tilde{F}(\rho)]^k g(c) dc, \quad (5.5)$$

where $H(p_{M+1}; \tilde{F}) = \infty$ and $H(p_1; \tilde{F}) = 0$.

From equation (5.2) it follows that there is a one-to-one mapping from prices to search costs via the reservation values. As a matter of fact, if we define c_i as the search cost of a consumer with reservation value $\rho(c_i, F)$ and this reservation value is equal to a price p_i , we

have $H(p_i, F) = c_i$. Therefore, we can calculate c_i directly from the data by using the empirical price cdf, i.e., $c_i = \frac{1}{M} \sum_{k=1}^i (p_i - p_k)$. Moreover, it follows from the empirical 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 $\tilde{F}(\rho) = \tilde{F}(p_j)$. This means that we can rewrite the integral in equation (5.5) as

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

where $\gamma_j = G(H(p_{j+1})) - G(H(p_j))$ is the share of consumers with a reservation value ρ_i in between p_i and p_{i+1} . Using this, equation (5.5) can be written as

$$y(p_i; \tilde{F}) = \frac{1}{N} \sum_{j=i}^M \gamma_j \sum_{k=1}^N [1 - \tilde{F}(p_j)]^{k-1}. \quad (5.6)$$

so for estimation we do not need to evaluate any integral.

We can use the expression for $y(p_i; \tilde{F})$ in equation (5.6) to write the equilibrium condition as

$$(\bar{p}-r) \frac{\gamma_M}{N} = (p_i-r) \cdot \left[\left(1 - \sum_{j=i}^M \gamma_j \right) (1 - \tilde{F}(p_i))^{N-1} + \frac{1}{N} \sum_{j=i}^M \gamma_j \sum_{k=1}^N (1 - \tilde{F}(p_j))^{k-1} \right]. \quad (5.7)$$

For estimation we rewrite equation (5.7) as

$$\gamma_i = \frac{\frac{\bar{p}-r}{p_i-r} \gamma_M - \sum_{j=i+1}^M \gamma_j \sum_{k=1}^N (1 - \tilde{F}(p_j))^{k-1} - N \left(1 - \sum_{j=i+1}^M \gamma_j \right) (1 - \tilde{F}(p_i))^{N-1}}{\sum_{k=1}^N (1 - \tilde{F}(p_i))^{k-1} - N(1 - \tilde{F}(p_i))^{N-1}}. \quad (5.8)$$

Equation (5.8) expresses the differences between subsequent values of the search cost cdf, denoted by γ_i , as a function of the price p_i , the distribution of prices $F(\cdot)$, the number of firms N , the maximum price \bar{p} , and the marginal cost r . More specifically, equation (5.8) states that if $F(\cdot)$ is the equilibrium distribution, then the share of consumers with a reservation value ρ_i in between p_i and p_{i+1} must satisfy equation (5.8).

Given estimates of $F(p)$, N , r , \bar{p} and γ_M it is possible to calculate first γ_{M-1} , then γ_{M-2} , and so on. The estimated γ_i 's in combination with the calculated c_i 's then form a nonparametric

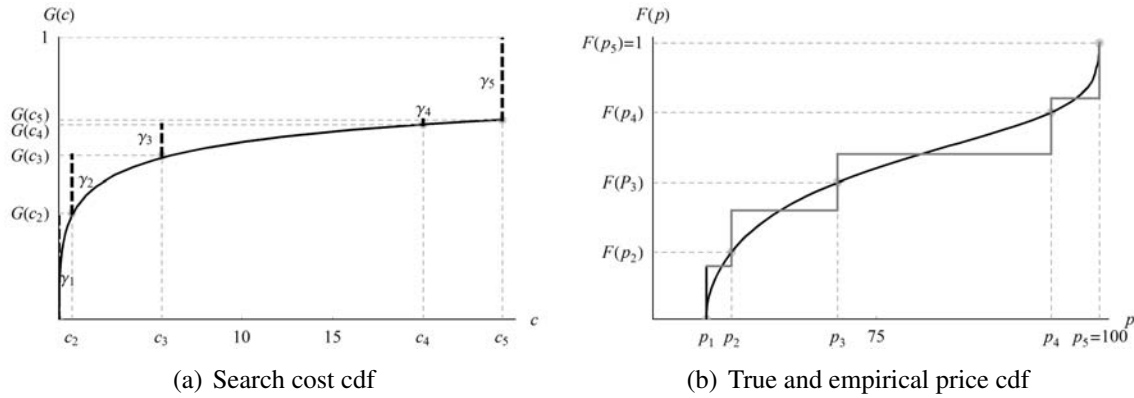


Figure 5.1: Example equilibrium search model ($M = 5$)

estimate of the search cost distribution $G(c)$. Notice that \bar{p} is estimated by the highest observed price. The estimate for unit cost r should be taken from another data source. The sum of all γ_i 's should be one, so γ_M should be chosen in such a way that this indeed holds.

Example equilibrium

To simulate the equilibrium it is useful to write equation (5.7) as

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

Now, given unit cost r , maximum price \bar{p} , number of firms N , search cost distribution $G(c)$, and M initial search cost values c_i , we can calculate prices p_i .³ It turns out that the equilibrium is a contraction. This means that with a set of computed prices we can calculate a new set of c_i 's, which will be used to generate a new set of prices. This process continues until the change in c_i 's compared to the previous round is approximately zero.

Figure 5.1 gives an example of an equilibrium. Here we take a lognormal search cost distribution with parameters 1 and 4, $\bar{p} = 100$, $r = 50$ and $N = 5$. Graph 5.1(a) gives the cdf of the search cost distribution together with the γ_i 's and the c_i 's that are identified in case of 5 price observations ($M = 5$). Graph 5.1(b) gives the constructed equilibrium price cdf and the empirical price cdf when there are 5 price observations.

³Remember that \bar{p} is endogenous, so this should be kept in mind when generating data. Depending on the chosen search cost cdf, a way would be to set $\bar{p} = v$ and check afterwards whether v is indeed smaller than p^M .




Figure 5.2: Identification

We note now that the identification of the search cost distribution works differently in the sequential search model than in the nonsequential search model. In Chapter 3 we have shown that in the nonsequential search model prices reflect only behavior of groups of consumers. For a fixed number of firms, having more observations does not lead to more groups of consumers, so this does not help to get more points on the search cost distribution. On the other hand, increasing the number of firms does not give more information of the height of search costs at high quantiles because the sequence of critical search costs c_k is convergent to zero. However, we show in Chapter 3 that a richer framework with price data from several markets overcomes the problem.

In the sequential search model one does not need price data from several markets to identify the search cost distribution in its full support, since increasing the number of observations is already enough to get more points on the support. Intuitively, this is because each extra price p_i corresponds via the definition of the reservation price in equation (5.2) to an additional search cost value c_i . Figure 5.2 shows how the identification exactly works. The graph shows the points

identified when we have 10, 25, 100 and 1000 price observations. Clearly, with 1000 price observations the search cost distribution is fully identified. Notice that as in the nonsequential search model, search costs are only identified up to some maximum search cost value. Using data from markets with the same search cost distribution but a higher valuation v would help us increase this maximum search cost value.

5.5 Empirical application

The purpose of this section is twofold. A first objective is to apply the estimation procedure to real world data to investigate what level of search costs can rationalize empirically observed prices. Secondly, a question that arises is what sort of mistake would the researcher commit if she used a sequential search model to estimate the relevant parameters when the data are actually generated from a nonsequential search strategy. The second question can easily be addressed by using the same data set as the one we used to estimate the nonsequential search model in Chapter 2. Once we get estimates of search costs for the sequential search model, we can compare them to the estimates we obtained in the earlier chapter with nonsequential search.

As shown in the previous section, to estimate the model we need to observe prices, the number of firms and the marginal cost. The prices used in this section come from the same data set as in Chapter 2 (prices of four memory chips sold on the Internet from the beginning of August 2004 till the end of September 2004).⁴ As in Chapter 2 we estimate the number of firms by the maximum number of firms found selling the memory chips in the period of data collection. Since we do not have data on the marginal cost of the memory chips, we propose to take the estimated values from Chapter 2 and check afterwards how robust the estimates are to changes in these values.

Figure 5.3 gives the estimated search cost distributions for all four memory chips. What is striking is that for all memory chips estimated search costs are quite high. For example, about 30% percent of the online consumers of the KTT3614 memory chip have search costs higher than 30 dollars. To compare the estimates of search costs under the two different search protocols, for each memory chip we plot in Figure 5.4 the search cost cdfs obtained from estimating the two models. It can be seen that estimated search costs are much higher in the case

⁴The data is discussed in more detail in Section 2.5.

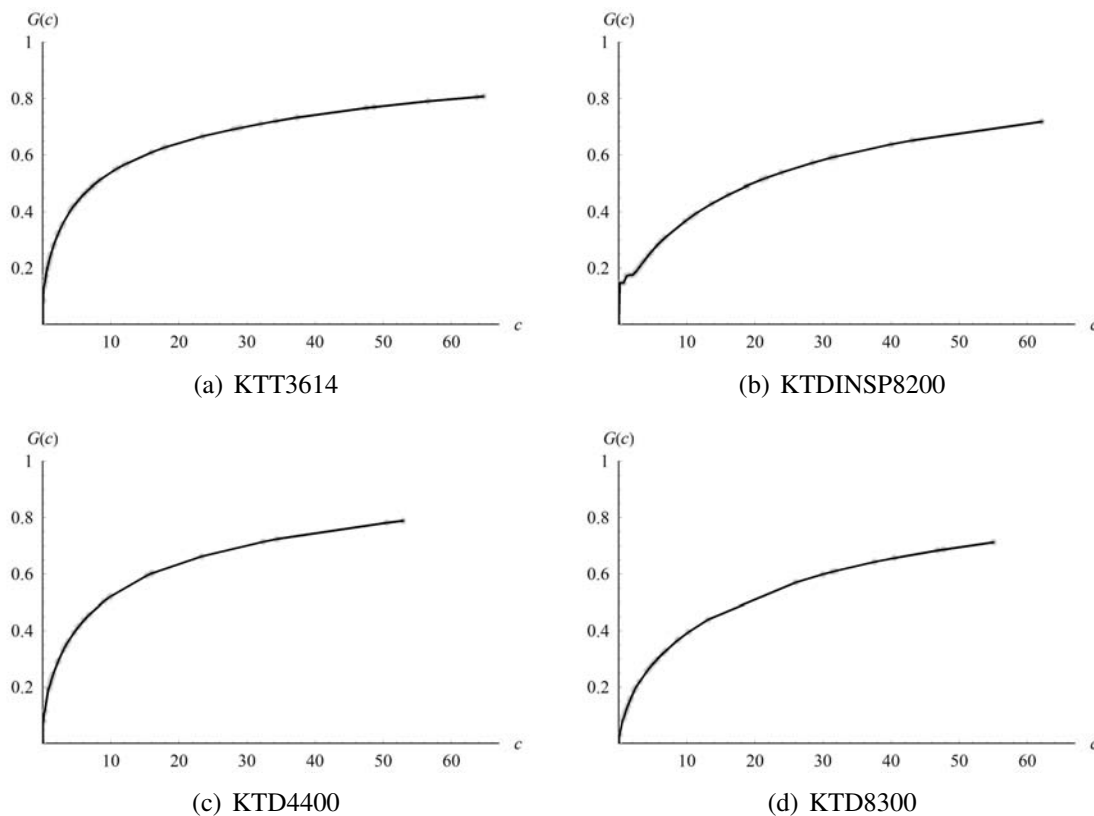


Figure 5.3: Estimated search cost cdf

of sequential search (solid curve) than in the case of nonsequential search (dashed curve).

Although, as can be seen from the graphs, the share of consumers with search costs close to zero is in both models roughly the same, in general estimated search costs according to the nonsequential search model are much lower. For example, median search costs of consumers buying the KTDINSP8200 memory chip are 19.72 US dollars according to the sequential search model, while only 9.41 US dollars according to the nonsequential search model. As can be seen from Table 5.1 a similar pattern arises for the other memory chips.

| | KTT3614 | KTDINSP8200 | KTD4400 | KTD8300 |
|----------------------------------|---------|-------------|---------|---------|
| Median search costs | | | | |
| Sequential | 7.78 | 19.72 | 8.84 | 19.05 |
| Nonsequential (Chapter 2) | 5.39 | 9.41 | 7.01 | 8.63 |
| Sequential (Hong and Shum, 2006) | 0.28 | 0.25 | 0.50 | 1.50 |

Notes: Marginal cost r is fixed at the marginal cost estimated by the nonsequential search model.

Table 5.1: Median search costs sequential and nonsequential search

The shapes of the estimated search cost distributions seems to be to in line with the shapes

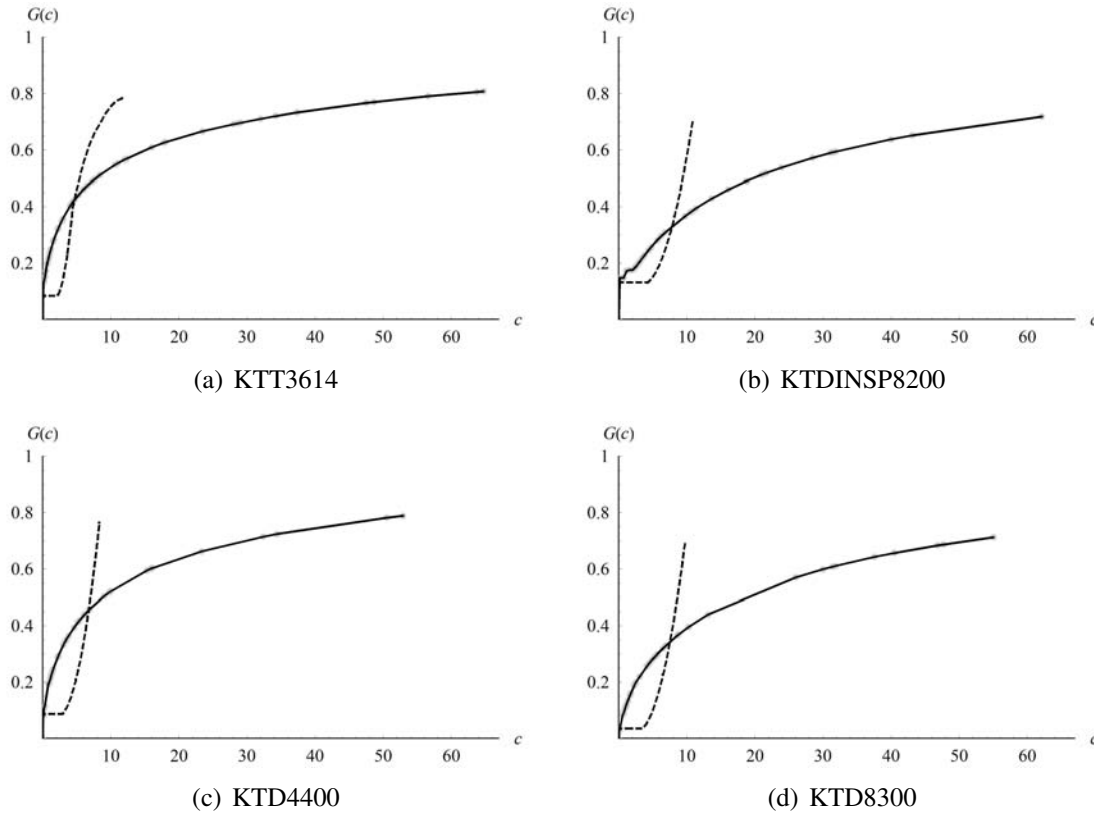


Figure 5.4: Comparison estimated search cost cdf with nonsequential search model

of the estimated search cost cdfs in the market for mutual funds as studied by Hortaçsu and Syverson (2004). In a related paper Hong and Shum (2006) estimate a sequential search model in a setting with infinitely many firms. Although their main focus is to estimate the model parametrically, as a specification check they also estimate a nonparametric version of their model. As in this chapter, they need marginal cost data to be able to do that. Hong and Shum (2006) show that in their model the relation to be estimated is simply

$$G(H(p)) = \frac{p - \underline{p}}{p - r},$$

so if one evaluates this relation only at the search cost values corresponding to reservation prices that equal observed prices, the search cost distribution can be estimated nonparametrically. Figure 5.5 presents a comparison of the results obtained from the two different estimation methods. Estimated search costs using the method of Hong and Shum (2006) are shown by the dashed curves, while the estimated search cost distribution using our method is given by the solid curves. As also indicated by the low median search cost values in Table 5.1 search costs

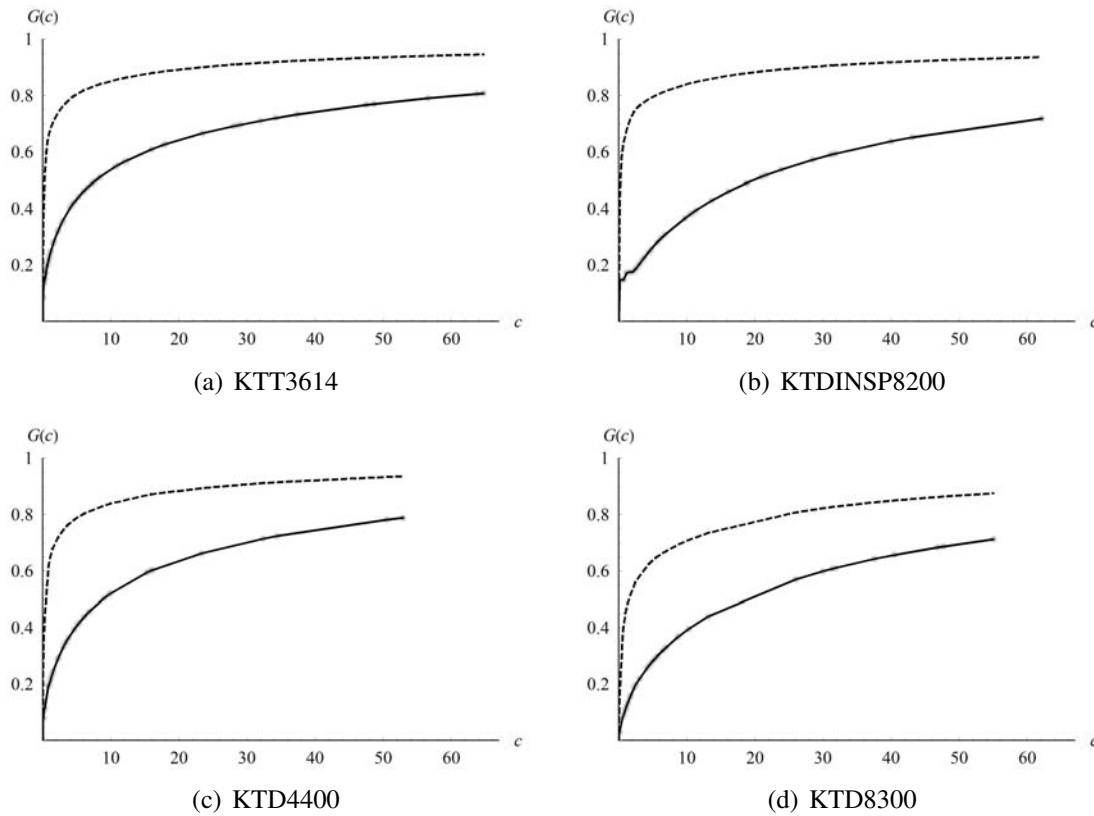


Figure 5.5: Comparison estimated search cost cdf with Hong and Shum (2006)

are much lower in the competitive model of Hong and Shum (2006) than in the N firm model presented here. The intuition behind this result is as follows. To rationalize the empirically observed pricing patterns, the competitive model needs lower overall search costs than the N firm model. This seems counterintuitive since this would imply that the same search cost distribution would lead to more competitive pricing in the N firm model than in the competitive model. A mechanism that could potentially explain this surprising result is that the probability that one has the lowest price decreases in the number of firms. In a setting with infinitely many firms this probability goes to zero, so if the number of competitors increases firms will put more weight on the consumers with high search costs and will allocate more probability mass to higher prices. As illustrated by Figure 5.5 the difference in estimated search costs is substantial, so this makes the choice which model to use an important one.

5.6 Conclusions

In this chapter we have shown how to estimate a sequential search model nonparametrically. The model has N firms competing for consumers searching for price information. Using a condition that implicitly defines the equilibrium price distribution, we have shown that each search cost cdf value can be expressed as a function of the price distribution, the marginal cost and the number of firms. As a result, using the empirical price distribution search costs of consumers having reservation prices equal to observed prices can be computed, as well as the corresponding quantiles of the search cost distribution.

The model has been estimated using prices of memory chips sold online. Estimates indicate that search costs are quite high; in fact, mean search costs are up to almost 20 US dollars. Compared to the estimates obtained in Chapter 2 where we estimated a nonsequential search model, search costs appear to be much higher under sequential search than under nonsequential search.

We have also compared our estimation method to the nonparametric estimation method of the sequential search model presented in Hong and Shum (2006). We have found that our estimates of search costs are much higher. A possible explanation for this striking result is that Hong and Shum (2006) model competition in a fully competitive setting, while we have only N firms competing. Everything else equal, firms put more mass at higher prices in a competitive setting, which implies that in order to rationalize the same set of prices, in the oligopolistic model one needs higher search costs.

6

Vertical product differentiation and search

NOTE: This chapter is based on Wildenbeest (2007a).

6.1 Introduction

Almost all theoretical search models discussed in Chapter 1 of this thesis deal with homogeneous goods, despite the fact that even goods that are seemingly homogeneous differ in aspects indirectly related to the good. For example, books are typically homogeneous; books are uniquely identified by an ISBN number and if there are two books with the same ISBN number there are no clear reasons to prefer one over the other. However, when the same book is sold by different retailers, it is a different story. Then the characteristics of the stores play a role as well. Although the book is the same, the retailers might differ in things like offered service or location, which will have an impact on pricing strategies.

This chapter studies a search model which allows for (vertical) product differentiation. This product differentiation can either be because the goods themselves are differentiated, or, as in the example above, because there are differences between the stores selling the goods. By letting both valuation and unit cost vary across firms, it is shown that a symmetric equilibrium exists in which firms offer deals to consumers that are randomly drawn from a common utility distribution. Because of the firm differences in quality and unit cost, the equilibrium entails different price distributions across firms. Consumers, on the other hand, know the valuation of the goods, but have to pay a search cost in order to observe prices. Given the utility distribution, consumers decide how much to search to maximize expected utility.

The search model in this essay can explain random pricing patterns, and unlike other existing models it also offers an explanation why some firms have persistently higher or lower prices than other firms. In consumer search models, the latter is typically explained by pure pricing strategies, while the former is usually explained by mixed pricing strategies. The search model in this paper has firms mixing in utilities. As a result, firms randomly draw prices from price distributions with supports that depend on the quality level of the product. To the best of my knowledge this feature is novel in the literature.

By using the structure of the equilibrium of the model, I show that it is possible to estimate both search costs and the impact of firm characteristics on prices using price data alone. Utilities can be estimated from prices by taking the negative of the residuals of a regression of prices on a firm dummy. The utilities are then used to estimate the model by maximum likelihood. The estimation method is similar to the method described in Chapter 3, although the likelihood

function is now written in terms of utility instead of prices.

I apply the estimation method to price data from supermarkets in the Netherlands. The data covers the period between January 2004 and June 2006. This period is interesting because it was characterized by a declining trend in prices.¹ The estimation results show that the model does quite well in explaining the observed prices. Besides that, I find that the share of consumers who search thoroughly has decreased over time. Since real monthly wages went down during the sampling period, this shift in the search cost distribution cannot be explained by changes in real income. Alternative explanations are changes in the utility of shopping itself or changes in consumer confidence.

In a related empirical paper, Lach (2002) studies existence and persistence of price dispersion using price data of four different products in Israel. Several predictions from search models are tested and he finds the patterns in the price data to be in line with these predictions. Lach (2002) controls for differences between firms in a similar way as I do here. In that sense, the analysis presented here shows that vertical product differentiation can be captured in a theoretical model in such a way that Lach's approach is theoretically justified. Moreover, this paper goes one step further by using the structure of the theoretical search model to estimate the underlying search cost distribution. In a recent paper Hong and Shum (2006) show how to estimate search cost distributions using the restrictions of an equilibrium search model. Using a maximum empirical likelihood approach they show that in their homogeneous good model only price data is needed to estimate search costs. In Chapter 2 this approach is extended to an oligopoly search model and it is shown how to estimate search costs using a maximum likelihood approach instead. Hortaçsu and Syverson (2004) show how to estimate search costs in a model of vertical product differentiation. The model presented here is similar in nature to Hortaçsu and Syverson (2004) since both allow for vertical product differentiation. Nevertheless, there are some important differences. First of all, Hortaçsu and Syverson (2004) assume consumers observe all prices, but do not know which firm charges which price, while here I assume consumers only know the distributions from which prices are drawn. When the market is populated by just a few firms, the assumption that consumers already know all prices around

¹In the Dutch press this process was called a 'price war'. The meaning of 'price war' in the economics literature is somewhat different because it refers to a deviation from a collusive equilibrium (see, e.g., Rotemberg and Saloner, 1986). It is not likely that supermarkets in the Netherlands were colluding at the beginning of the sampling period, so in this context using the term price war would be misleading.

is a particularly strong one because the realized prices might differ a lot from the distribution from which they are drawn. A second difference is that Hortaçsu and Syverson (2004) study a sequential search model, while I assume consumers search nonsequentially. Nonsequential search is optimal in some circumstances, for example when several price quotations arrive at a time.² Finally, and most importantly, another difference is that price dispersion in Hortaçsu and Syverson (2004) is a result of firms playing pure strategies, while in the model presented here it is a result of mixed strategies. In a mixed strategy equilibrium profits need to be the same across firms, which gives an extra condition that can be used for the estimation of the model. This extra condition makes that here only price data is needed to estimate the model, while Hortaçsu and Syverson (2004) need both price and quantity data.

The empirical section is connected to the literature on retail pricing dynamics in supermarkets. In recent work it has been observed that a typical grocery product is sold at a regular price for a number of time periods, whereas only once in a while the product is sold at a discount price (see e.g. Hosken and Reiffen, 2004; Pesendorfer, 2004). An implication of this is that current prices depend on past prices. This seems to be at odds with the implication of the search model in this paper that prices are random draws from some distribution, since this means that prices are not predictable and that there is no correlation over time. However, an important difference between the papers mentioned above and this paper is that while they focus on single products, the focus here is on baskets of goods. Consumers typically buy more than one grocery product at a time when visiting a supermarket, so the sum of prices of all products being bought together is of greater importance to the consumer than the price of individual products. It turns out that considering a basket of goods one cannot reject the null hypothesis that prices are random draws from some distribution.

The structure of this chapter is as follows. In the next section I discuss the theoretical model. Section 6.3 continues with a method to estimate search costs using maximum likelihood. In Section 6.4 I apply the estimation method to price data from supermarkets. Finally, the last section concludes.

²See Morgan and Manning (1985) for the optimality of sequential versus nonsequential search.

6.2 The model

I study a model of firms offering a differentiated good competing for incompletely informed consumers. On the supply side there are N firms producing goods at a unit cost r . The goods are vertically differentiated, i.e., the goods can be ranked according to their characteristics and consumers all agree in the ranking. The model can be used to address two sources of product differentiation. The goods can be differentiated either because the goods themselves are heterogeneous, for example because the products have different features. It could also be that the good is homogeneous, but that the firms selling the product are differentiated. An example of this would be stores selling the same products but offering different service levels.

On the demand side there is a continuum of consumers demanding at most one good, all deriving the same utility from consumption of good j :

$$u_j(X_j, p_j; \beta) = \beta X_j - p_j + \xi_j, \quad (6.1)$$

where X_j are the observable characteristics of the good, p_j is the price of firm j 's good, ξ_j are (by the econometrician) unobservable characteristics of good j and where the parameter β describes the relation between X_j and u_j . Since the coefficient of price is normalized to -1 , utilities are measured in the same unit as prices. In what follows, let the valuation v_j for the good produced by firm j be the sum of the contribution of the observable and unobservable characteristics to utility, so that equation (6.1) can be rewritten as

$$u_j = v_j - p_j. \quad (6.2)$$

Consumers know their valuation for the good produced at the different firms. However, consumers have no information about which firm provides the good at what utility level, since prices are only observed after searching. By engaging in costly search the consumers can gain information about the prices of the good at a subset of the firms. Consumers are characterized by their search cost c , which is a random draw from the distribution function $G(c)$, with density function $g(c)$. I assume consumers search nonsequentially, i.e., consumers determine before entering the market how many times to search.³ Consumers then buy the product from the firm

³The way consumers search is similar to the nonsequential search model of Burdett and Judd (1983).

in their sample providing the highest utility level.

Firms and consumers play a simultaneous move game. The store characteristics are assumed to be drawn from some distribution and are fixed in the short run. I assume the difference between valuation and unit cost $v_j - r_j$ to be the same across firms, i.e., $x = v_j - r_j$, where x is the maximum possible margin that can be attained by the firms. This means that more favorable characteristics come at a higher cost. Moreover, by restricting v_j and r_j to be related in this way, as I will show below, firms are symmetric in the margin received at each offered utility level. In this way incentives for the firms are identical, which gives rise to the existence of a symmetric equilibrium.

Valuations and unit costs are common knowledge. Therefore, an individual firm takes the expected utilities of the other firms and the search behavior of consumers as given while setting its own price. An individual consumer takes the firm pricing strategies as given and decides on a number k of firms to visit in order to maximize utility. The fraction of consumers sampling k firms is denoted by μ_k .

I only focus on symmetric equilibria. A condition for a symmetric equilibrium to exist is that some consumers should search once, while others should search more than once. The intuition for this is that if all consumers did compare prices, all firms would set a price equal to their unit cost, which implies that all firms would be offering the same utility level x . As a result, there is no reason to search. On the other hand, if no consumers were willing to compare prices, firms would set their price equal to their valuation, which means that all firms would offer a utility level of zero. Consumers would not participate, because they have to pay a search cost c to enter the market.⁴

A second condition for a symmetric equilibrium to exist is that the firms must play mixed strategies in setting their utility level. The proof for this is similar to the proof of Lemma 3.2 of Chapter 3 and can be explained by the idea that offering slightly more utility gives a discrete jump in profits when dealing with consumers that compare utilities. Hence there are no atoms in the utility distribution. On the other hand, at utility levels close to zero only consumers searching once will buy, so in that case offering a lower utility increases profits. As a result, firms draw utilities from a common atomless utility distribution, which I denote $L(u)$, with a lower bound equal to zero.

⁴See also Lemma 3.1 of Chapter 3.

Consumer search behavior should be optimal. This means that for a consumer searching k times, the expected utility should be higher than the expected cost of searching kc . Moreover, the net benefit of searching k times should be higher than the net benefit of searching $k - 1$ or $k + 1$ times. Now define c_k as the search cost of the consumer indifferent between searching k and $k + 1$ times. For this consumer $E[\max\{u_1, u_2, \dots, u_k\}] - kc = E[\max\{u_1, u_2, \dots, u_{k+1}\}] - (k + 1)c$, or

$$c_k = E[\max\{u_1, u_2, \dots, u_{k+1}\}] - E[\max\{u_1, u_2, \dots, u_k\}]. \quad (6.3)$$

The share of consumers who search k times is then given by

$$\mu_k = \int_{c_k}^{c_{k-1}} g(c)dc = G(c_{k-1}) - G(c_k). \quad (6.4)$$

Now consider optimal firm behavior. Given expected consumer behavior μ_k and expectations on $L(u)$, the profit of firm j offering utility u_j is given by

$$\pi_j(u_j; L(u)) = (x - u_j) \sum_{k=1}^N \frac{k\mu_k}{N} L(u_j)^{k-1}.$$

Since $x - u_j = p_j - r_j$, the first part of this equation is the margin the store makes on its product. The second part represents the expected quantities sold, and is explained as the summation over all N consumer groups of the share of consumers searching k times multiplied by the probability that these μ_k consumers visit the firm (which is k/N) and by the probability that a firm selling the product at a utility level of u_j offers the highest utility out of k firms, which is $L(u_j)^{k-1}$.

Given the mixed strategies, in equilibrium a store should be indifferent between setting any utility in the support of $L(u)$. In addition, the lower bound of $L(u)$ should be equal to zero. This is because a firm offering a utility of zero will only sell to the consumers searching once, and surplus extracted from these consumers is maximized by setting $\bar{p}_j = v_j$ so that $\underline{u} = 0$. In this case the profit equation simplifies to $\pi(\underline{u}) = x\mu_1/N$. Setting this equal to the equilibrium profits in general gives the equilibrium condition for this model:

$$(x - u) \sum_{k=1}^N \frac{k\mu_k}{N} L(u)^{k-1} = x \cdot \frac{\mu_1}{N}. \quad (6.5)$$

Unfortunately, this equation cannot be solved for $L(u)$, so the equilibrium distribution of utilities is only implicitly defined. Solving equation (6.5) for u gives

$$u = x \cdot \frac{\sum_{k=2}^N k \mu_k L(u)^{k-1}}{\sum_{k=1}^N k \mu_k L(u)^{k-1}}. \quad (6.6)$$

Although the utility distribution is the same for each firm, since $u = v_j - p_j$, the price distribution is different across firms:

$$F_j(p) = \Pr[p_j \leq p] = \Pr[p \geq v_j - u_j] = \Pr[u_j \geq v_j - p] = 1 - L(v_j - p).$$

The maximum utility in the market can be found by setting $L(u) = 1$, which gives

$$\bar{u} = x \cdot \frac{\sum_{k=2}^N k \mu_k}{\sum_{k=1}^N k \mu_k}. \quad (6.7)$$

Individual firms choose a utility level to maximize expected profits given expected search behavior of the consumers and given the expected utility distribution function, so in equilibrium the first order condition with respect to u should be zero, i.e.,

$$\frac{\partial \pi}{\partial u} = \sum_{k=1}^N \frac{k \mu_k}{N} L(u)^{k-1} - (x - u) \sum_{k=1}^N \frac{k(k-1) \mu_k}{N} L(u)^{k-2} l(u) = 0.$$

Solving this expression for $l(u)$ gives the density function of utility

$$l(u) = \frac{\sum_{k=1}^N k \mu_k L(u)^{k-1}}{(x - u) \sum_{k=1}^N k(k-1) \mu_k L(u)^{k-2}}. \quad (6.8)$$

Using the characterization of the utility distribution equation (6.3) can be rewritten as a function of the utility distribution:

$$c_k = \int_{\underline{u}}^{\bar{u}} (k+1) u L(u)^k l(u) du - \int_{\underline{u}}^{\bar{u}} k u L(u)^{k-1} l(u) du.$$

By using the change of variable $y = L(u)$, we obtain $dy = l(u) du$. Plugging this into the equation above, transforming the lower limit into $y = L(\underline{u}) = 0$ and the upper limit into

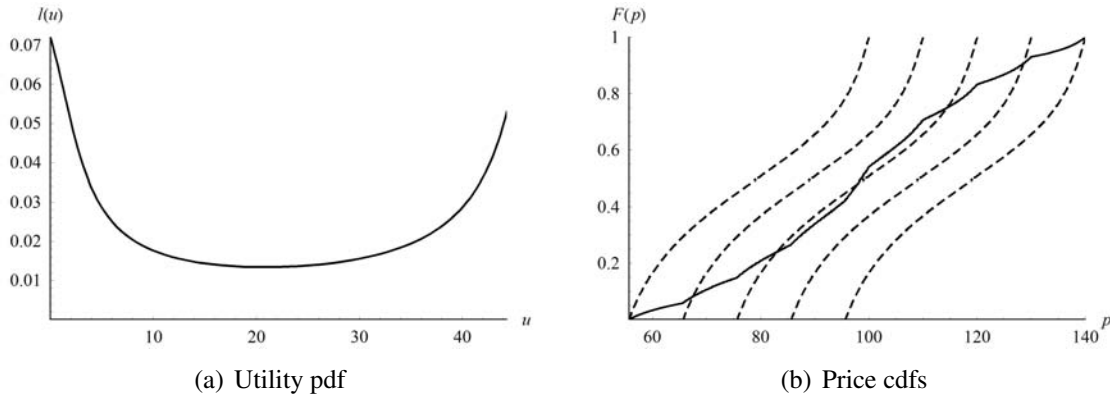


Figure 6.1: Example equilibrium search model

$y = L(\bar{u}) = 1$ and solving gives

$$c_k = \int_0^1 u(y)[(k+1)y - k]y^{k-1} dy. \quad (6.9)$$

Then using the same change of variable in equation (6.6) we can get rid of $u(y)$ in equation (6.9).

As an example, I calculate equilibrium when consumers search costs are drawn from a log-normal distribution with parameters 0.5 and 5. Figure 6.1 gives plots of the equilibrium for 5 firms with valuations ranging from 100 to 140 and $x = 50$ so that marginal cost range from 50 to 90. In Figure 6.1(a) the equilibrium utility density is plotted. Most mass is at the extremes of the distribution, with slightly more mass at lower utilities than at higher utilities. This shows the tradeoff firms face: set a high utility to attract consumers who compare several offerings or set a low utility in order to maximize surplus extracted from consumers who do not search. In Figure 6.1(b) the equilibrium price cdfs are drawn; the dashed lines are the firms' individual price cdfs and the solid line is the price cdf for all the firms together. What is interesting to note is that the shape of the individual price cdfs is quite different from the shape of the price cdf of all firms together. This means that assuming all firms are selling the same homogeneous product when in fact they are not might likely lead to wrong estimates of the underlying search cost distribution. I will come back to this issue in the empirical section.

Note that price dispersion in Hortaçsu and Syverson (2004) is a result of firms playing pure strategies, while in the model presented here it is a result of mixed strategies. In a mixed strategy equilibrium profits need to be the same across firms, which gives an extra condition that can be

used for the estimation of the model. As we will see in the next section, this extra condition makes that here only price data is needed to estimate the model, while Hortaçsu and Syverson (2004) need both price and quantity data.

6.3 Estimation

The goal of this section is to present a method to estimate the model presented in the previous section using only price data. Assume the prices N firms charge for the same good are observed for a certain period of time, the latter being indicated by the subscript t . There are two methods to calculate utilities from observed prices. In the first method v_j is (superconsistently) estimated by taking the maximum observed price \bar{p}_j for each firm j during the sampling period. Then using equation (6.2) we get $u_{jt} = v_j - p_{jt} = \bar{p}_j - p_{jt}$, so corresponding utilities for all observed prices can be calculated. The second method follows from rewriting equation (6.2) to $p_{jt} = v_j - u_{jt}$. This equation can be estimated by carrying out a fixed effects regression of prices on a constant, i.e.,

$$p_{jt} = \alpha + \delta_j + \epsilon_{jt},$$

where α is a constant, δ_j are the firm fixed effects and ϵ_{jt} are the residuals. Note that with this specification, valuations v_j are estimated by $\alpha + \delta_j$ and utilities are calculated by taking the negative of the residuals ϵ_{jt} . Moreover, ϵ_{jt} is simply the price at time t for firm j minus the average price of firm j within the period, which means that $u_{jt} = -\epsilon_{jt} = p_j - p_{jt}$, where p_j is the average price for store j .⁵ In both methods utilities are calculated by restricting the shape of the price distribution to be the same across firms (although they might have different means), but instead of using the maximum observed prices across firms the second method uses the average observed prices across firms to serve as a proxy for differences in valuations. If the underlying utility distribution is the same, asymptotically both methods give the same results. However, in small samples estimates using the maximum price to estimate valuation leads to an overestimation of the density at zero utility. For each firm by definition the highest observed

⁵For the analysis the height of the utilities is not important, all what matters are the differences between the utilities. This means that our estimate of the search cost distribution does not change when the same constant is added to all valuations v_j .

price during the sampling period leads to a zero utility observation, so if the sample consists of 10 firms competing over 10 periods, 10 percent of the observations lead to zero utility, while if the firms compete over 100 periods, only 1 percent of the observations lead to zero utility. Such an overestimation of zero utility in the utility distribution seriously affects the estimates of the underlying search cost distribution, so in small samples it is better to estimate the valuations using the mean of the prices.

The density function in equation (6.8) can be used to estimate the search cost distribution by maximum likelihood. Since all firms are assumed to draw utilities from the same distribution, all the calculated utilities can simply be pooled. The log-likelihood function is then $LL = \sum_{i=1}^M \log l(u_i)$, where M is the total number of observations. The likelihood function can be concentrated by solving the calculated upper bound of the utility distribution in equation (6.7) for x as a function of the rest of the parameters, i.e.,

$$x = \bar{u} \cdot \frac{\sum_{k=1}^N k\mu_k}{\sum_{k=2}^N k\mu_k},$$

and by plugging this into equation (6.8). I then proceed as in Chapter 3 to estimate the parameters of the model using the semi-nonparametric (SNP) density estimator. The idea is to mimic the underlying search cost distribution by the flexible, Hermite polynomial based form of the SNP estimator. As in Chapter 3 I take the parametric form of the univariate SNP estimator of Fenton and Gallant (1996) but I use the log-normal instead of the normal distribution as the base case. The search cost density is then specified as

$$g(c; \gamma, \sigma, \theta) = \frac{\left[\sum_{i=0}^{p_n} \theta_i w_i(c) \right]^2}{\sum_{i=0}^{p_n} \theta_i^2}, \theta \in \Theta_p, \Theta_p = \{\theta : \theta = (\theta_0, \theta_1, \dots, \theta_p), \theta_0 = 1\}, \quad (6.10)$$

where p_n is the number of polynomial terms,

$$\begin{aligned} w_0(c) &= (c\sigma\sqrt{2\pi})^{-1/2} \exp(-((\log c - \gamma)/\sigma)^2/4), \\ w_1(c) &= (c\sigma\sqrt{2\pi})^{-1/2} ((\log c - \gamma)/\sigma) \exp(-((\log c - \gamma)/\sigma)^2/4), \\ w_i(c) &= \left[((\log c - \gamma)/\sigma) w_{i-1}(c) - \sqrt{i-1} w_{i-2}(c) \right] / \sqrt{i} \text{ for } i \geq 2. \end{aligned}$$

The search cost density is then, together with the specification of c_k as in equation (6.9), plugged into the specification of μ_k as in equation (6.4), which is subsequently plugged into the log-likelihood function to get rid of the μ_k 's and the c_k 's. The parameters left for estimation by maximum likelihood are the parameters of the SNP density function, i.e., $\{\gamma, \sigma, \theta_0, \theta_1, \dots, \theta_{p_n}\}$.⁶

6.4 Empirical analysis

In this section I apply the estimation method to prices collected from Dutch supermarkets between January 2004 and June 2006. The application of the search model described in Section 6.3 to grocery products might need some justification. The supermarket sector is typically a sector in which vertical product differentiation plays an important role. Everyone who visits a supermarket once in a while knows that things like shopping location, parking possibilities, product variety, queue length, etc. do matter a lot. However, although favorable characteristics increase the utility level of the typical visitor of a supermarket, they usually come at a cost. Full-service supermarkets, focussing on quality, are in general more expensive than for example discounters, whose primary focus is on low prices and not on service.

Compared to shops selling memory chips online as studied in previous chapters, in the offline world physical locations become important as well. Although this implies that horizontal product differentiation issues might be relevant, allowing for horizontal product differentiation in addition to vertical product differentiation would complicate the analysis too much. We therefore ignore horizontal characteristics, so location should be interpreted as a vertical characteristic, which can be justified by the idea that some supermarkets in general have better locations than others.

Another assumption made is that consumers search nonsequentially for the highest utility around. Nonsequential search implies that consumers determine before they start searching how many times to search. To justify this assumption one could think of consumers using advertisements in for example newspapers to collect information about prices at different supermarkets, the use of price comparison sites on the Internet, or a situation where there are a lot of shops at the same distant place in town.⁷

⁶Note that the identification results of Chapter 3 carry over to this model as well.

⁷At this point I would like to notice that the sequential search model of Chapter 5 could also be extended in the

In the analysis we do not explicitly take advertising into account, although it is clear that advertising is more of an issue for grocery products than for online memory chips. Through advertising consumers essentially get some price information for free, so ignoring advertising puts a lower bound on the estimated search costs. On the other hand, as I will explain in more detail below, my focus will be on a basket of goods, so ignoring advertising could be justified on the basis of the argument that consumers are not so much interested in the prices of only a few advertised products, but only in the price of a basket of grocery products. Our focus on a basket of goods also helps to justify the inelastic demand assumption, since the usual buyer is expected to buy a single basket at a time.

The model is applied to a data set of prices that are collected over time, so the implicit assumption is that supermarkets play a stationary repeated game of finite horizon. This means that we are ignoring dynamic effects caused by for example loyalty cards, advertising, and switching costs. However, since part of the share of consumers searching only one time can also be interpreted as consumers being loyal to some supermarket, to some extent loyalty can also be accommodated in the current setting.

The focus of this study will be on homogeneous goods, but I allow for the possibility that supermarkets are differentiated in terms of the service they offer. Most theory models explain price dispersion by either random pricing strategies of homogeneous firms or by pure strategies of heterogeneous firms. The model described in Section 6.2 combines the two: heterogeneous firms mix over price distributions with different support. As is shown below in some detail, in the data set average prices across stores are persistently different over time, but at the same time, stores randomize their prices. These observations make the model presented here a suitable theoretical framework to study price setting behavior of supermarkets in relation to search behavior of consumers, as traditional search models cannot explain both things at the same time.

Besides estimating search cost parameters and investigating whether the model does well in explaining observed pricing behavior in the supermarket sector, an additional purpose of this empirical section is to investigate how search costs have evolved over the sampling period. This is an especially interesting exercise because during the period of data gathering the Dutch supermarkets were engaged in a process of lowering their prices. This process was initiated on

same way as I did in this chapter for the nonsequential search model of Chapter 2. As discussed in Chapter 5, one could then evaluate the extent of the bias introduced by estimating a model with the “wrong” search protocol.

October 20th 2003 by market leader Albert Heijn by cutting prices of around 1000 products. All other competitors reacted swiftly on this move and from that moment on Albert Heijn cut prices several times.

Albert Heijn is a typical full-service supermarket, offering high service at high prices, and, according to the management of Albert Heijn, the reason to lower their prices was to improve its price-quality ratio by offering high service at more competitive prices.⁸ In terms of the model described in Section 6.2, improving the price-quality ratio while at the same time cutting prices means that a firm with unit cost r_i has to decrease more than the maximum price consumers are willing to pay for its basket v_i (i.e., supermarkets need to become more efficient) and this is exactly what happened: supermarket chains tried to renegotiate contracts with their suppliers in order to lower their unit costs.⁹

I divide the entire sampling period into three periods and I estimate the model for each period separately. This not only allows me to estimate the search cost distribution for each period separately, but is also necessary to capture the effects of the price cutting process on the firm side of the model. I assume that within a period supermarket chains are in equilibrium, with exogenous parameters fixed, but that between periods the exogenous parameters of the firm side of the model are allowed to change. This means that the stores' margins in periods 2 and 3 are not necessarily the same as in period 1, which is necessary to explain why Albert Heijn initiated the price cutting process in the first place.

The setup of the empirical analysis is as follows. In the next subsection, I start by giving a description of the entire data set. Here I check some of the implications of the model, like random pricing. I then divide the data set in three periods of equal length. Next, because the analysis of the first two periods is complicated by a declining monthly average price, for expositional reasons I first estimate the model for the last period, without correcting for any trend. Thereafter, I estimate the model for all periods, but this time I do take changes in the average monthly price into account. I end the study with a comparison of the estimates.

⁸See <http://www.ah.nl/albertheijn/persberichten/article.jsp?id=219162> for the press release of Albert Heijn announcing price cuts on more than 1000 of its products on October 20, 2003.

⁹In a document of the Dutch Competition Authority (NMa) on supermarkets and their relationship with suppliers, the NMa concludes that the supermarkets passed on the benefits of lower supply prices to the consumers (see point 6 of <http://www.nmanet.nl/images/veelgestelde%20vragen%20inkoopmacht.tcm16-55025.pdf>).

6.4.1 Description of the data

The data was collected using the web site Supers.nl, an independent price monitor of supermarket prices in the Netherlands.¹⁰ Each month, Supers.nl compares prices for around 25-30 products sold at several Dutch supermarkets chains. In the Netherlands, all supermarkets price uniformly across locations. This means that even though supermarkets serve local customers, prices are set nationally, so pricing strategies of the stores can be investigated at an aggregate level.¹¹ As already said, besides seeing whether the model could actually explain observed behavior in this sector, an additional purpose of this empirical analysis is to investigate how search costs evolve over time. Given that supermarkets initiated a process of lowering prices in October 2003, ideally I would have liked to have data from before and after the start of this process. However, since price comparison on Supers.nl started on January 2004 only data from during the process is available.

The data set consists each month of around 25-30 products, but the set of products is not necessarily the same each month. In this analysis I focus on a basket of goods. In supermarkets, consumers usually buy many products at a time, so most consumers are interested in the price of a basket of goods and not so much in the price of individual products.¹²

In total there are 40 different products in the data set. Figure 6.2 gives a kernel estimate of the price density of the entire basket, that is, the price density of all the 40 products which appear at least once during the period of data collection together, where prices are in logarithmic form and as deviations from the monthly average to control for differences in products within the sample each month. Moreover, to correct for changes in the CPI I divide each price observation by the monthly CPI with January 2004 as the base case, so that all prices are in terms of January 2004 prices. As a result, all the variation depicted in Figure 6.2 is solely due to differences between stores. As the graph shows, there is substantial price dispersion for the

¹⁰See <http://www.supers.nl>.

¹¹In a press release of the Dutch Competition Authority (NMa) on the acquisition of Groenwoudt by Laurus, the NMa speaks about a so-called chain effect in the way supermarket compete in the Netherlands; because areas from where supermarkets attract their customers overlap, in the end supermarkets compete at a national level (see http://www.nmanet.nl/nederlands/home/Actueel/Nieuws.Persberichten/NMa_Persberichten/2000/00_14.asp).

¹²In addition, a supermarket is a multi-product firm so a single-product model as described in Section 6.2 is probably not the right model when investigating individual products. A drawback of looking at baskets of products instead of individual products is that behavior of consumers who go to different supermarkets for each different product is not captured.




Figure 6.2: Price density (logarithm of real price in deviation from monthly average)

basket.¹³ Whether this is due to quality differences or due to strategic price differences will be the focus of the remainder of this section.

Supers.nl compares prices of products at all the big supermarket chains in the Netherlands selling branded products. In total there are 18 different supermarket chains in the data set.¹⁴ Not all chains are in each others proximity at the local level, but since prices are set nationally, I assume all stores are competing with each other and set the number of stores N equal to 18. The goods within a basket are all homogeneous brands, but because they are sold by different sellers, the heterogeneity of the supermarkets is transferred to the products. In terms of the theoretical model presented in Section 6.2, consumers valuation v is different across supermarket chains, even though the good is the same. Therefore, supermarkets with better valued characteristics can ask higher prices on average.

To investigate whether the right model is used to study search behavior of consumers in this setting, I will first check if some of the implications of the model, like random pricing and persistent differences in average price across stores, are observed in the data. As in Lach (2002), I concentrate on the relative position of the stores in terms of price rankings, and how these positions evolve over time, but in contrast to his approach, because my model allows

¹³Table A6.6 in the Appendix gives the names and some simple statistics for the 40 products which appear at least once during the period of data collection. The coefficient of variation fluctuates across products, but as the table shows, also for most products there seems to be substantial price dispersion as well.

¹⁴I left two supermarkets out of the analysis because there were only price observations for 8 months or less.




Figure 6.3: Price rankings (real price)

for vertical product differentiation the prices I use to calculate rankings are not corrected for differences between stores.

According to the model, prices are randomly drawn from a distribution, so as a result of this, price rankings should change over time, but, as different stores have different supports of the price distribution, it is unlikely that a store's variation in ranking is very large. Figure 6.3 shows how the price rankings of the stores evolve over time for the basket of goods. Clearly, rankings do change over time, but it does not happen that a store with an initial low price for the basket, and thus a low ranking, moves up all the way in the ranking distribution. The same holds for stores with an initial high ranking.

To study this issue in more detail, Table 6.1(a) gives information about the time a firm spends in each quartile of every month's empirical price distribution. Stores change their relative position in the price rankings, but usually not every month. For example, 6.19% of the price observations in the first quartile were for stores that had a price in this quartile for one successive month. Likewise, 3.54% of prices in the first quartile belong to stores that were in this quartile for two successive months. Especially the stores that have a price within the first or fourth quartile stay there for many months: at least 65% of stores for more than 6 months. Among supermarkets pricing in the second and third quartile there is more fluctuation; prices are on average between 4 and 5 successive months in one of these quartiles and 50% of stores keep their prices in these quartiles for at most 5 months.

| Duration | q_1 | q_2 | q_3 | q_4 |
|-----------|-------|-------|-------|-------|
| 1 month | 6.19 | 14.29 | 24.79 | 12.77 |
| 2 months | 3.54 | 12.03 | 11.57 | 4.26 |
| 3 months | 7.96 | 11.28 | 17.36 | 6.38 |
| 4 months | 3.54 | 12.03 | 16.53 | 11.35 |
| 5 months | 8.85 | 11.28 | 12.40 | 0.00 |
| 6+ months | 69.91 | 39.10 | 17.36 | 65.25 |
| mean | 11.19 | 4.76 | 4.14 | 12.01 |
| median | 13 | 5 | 4 | 6 |
| max | 20 | 9 | 12 | 27 |

Notes: In percentages.

(a) Durations by quartile

| Quartile | Quartile spend most time in | | | |
|----------|-----------------------------|-------|-------|-------|
| | q_1 | q_2 | q_3 | q_4 |
| q_1 | 71.67 | 19.01 | 0.00 | 0.00 |
| q_2 | 25.83 | 59.15 | 10.83 | 3.97 |
| q_3 | 2.50 | 20.42 | 61.67 | 11.90 |
| q_4 | 0.00 | 1.41 | 27.50 | 84.13 |
| N | 4 | 7 | 4 | 5 |
| M | 120 | 142 | 120 | 126 |

Notes: In percentages.

(b) Quartiles by quartile spend most time in

Table 6.1: Stores' positioning

Table 6.1(b) groups supermarkets by the quartile they spent the most time in and gives the percentage of prices in each quartile. In line with Table 6.1(a), supermarkets that have most observations in either quartile 1 or 4 are the least found in other quartiles. On average 69% of prices of a supermarket are observed in one quartile only. Moreover, on average 95% of prices within a supermarket are in two successive quartiles only. There are no supermarkets which have prices in all four quartiles. We can conclude from all this evidence that supermarkets change their relative positions, but stick to two successive quartiles of the empirical price distribution only. Together with Figure 6.3 and the duration results this shows that supermarkets randomize their prices, but over different supports. This is consistent with the theoretical model of Section 6.2.

To investigate if search costs have changed over time, and to capture the dynamics of the firm side of the model, I divide the sampling period into three distinct subperiods. More specifically, I divide the 30 months into three periods of 10 months, so that all periods are of equal length. This means that within a period it is assumed that the maximum possible margin x and the search cost distribution $G(c)$ do not change, but that between periods both x and $G(c)$ are free to fluctuate. In this way I am able to estimate the search cost distribution for each period separately, while controlling for possible changes in x as a result of the supermarkets' attempts to improve their price-quality ratio or to become more efficient.

To be able to estimate valuations within a period, the products in the basket should not change. However, between periods it does not matter if the composition of the basket changes, as the underlying search cost distribution to be estimated does not depend in any way on the




Figure 6.4: Monthly average price of the basket per period

basket or the products in the basket. Selecting only the products for which I had price observations for a complete period, I end up with a basket containing 20 products in period 1, 25 products in period 2 and 22 products in period 3.¹⁵

Figure 6.4 gives the average monthly price for the basket for each month in the sample. Note that each period has different products in the basket so the average monthly price cannot be compared between periods. Especially in the first period there seems to be a downward trend in the average monthly price. Even though in terms of the search model each average price is the result of the mean of only around 18 random draws from a particular distribution it is unlikely that declining price pattern is due to randomness. Instead it is more likely that the declining pattern is caused by the ongoing price cutting process. In the third period prices are more or less stabilized, which suggests that the price cutting process has probably ended.¹⁶

The downward trend in average monthly price in period 1 and 2 complicates the estimation of the model. For expositional reasons I will therefore estimate the model first for the third period without correcting for differences in average price across time; in subsection 6.4.3 the search cost distributions will be estimated for all periods, but this time by allowing for time

¹⁵In period 2 there is one product for which there are 9 out of 10 prices in the sample. To increase the size of this periods basket I use the price of the week before to replace the missing price.

¹⁶Although there is a downward trend in the average monthly price in the first months, this is not necessarily at odds with the assumption that x is constant within a period. To see this, note that in terms of the model of Section 6.2, when both v and r decline by the same amount, x remains constant while the average price goes down.

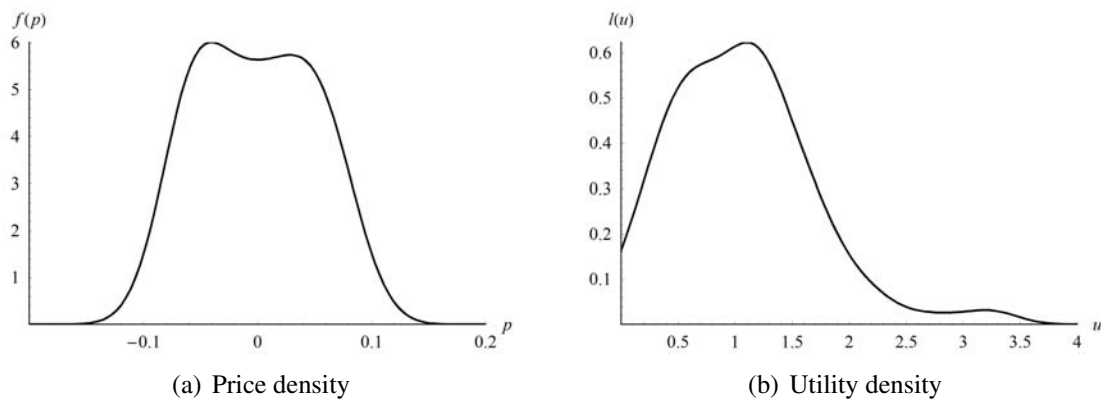


Figure 6.5: Price and utility density period 3

fixed effects as well. In subsection 6.4.4 the results will be compared and discussed.

6.4.2 Estimation of search costs without time fixed effects

Figure 6.5(a) shows a kernel estimate of the price density for the basket in the third period, where prices are the logarithm of real prices in deviation from the period average. According to the search model presented in Section 6.2, the substantial price dispersion shown in this graph is explained as a combination of quality differences between stores and random pricing strategies. Because of the way utilities are defined in the model, utilities are essentially prices controlled for quality differences between stores. As described in the previous section, utilities can be calculated by taking the difference between the maximum price for store j over time and the price at time t , i.e., $u_{jt} = \bar{p}_j - p_{jt}$, or by a fixed effects regression, where the utilities are given by $u_{jt} = -\epsilon_{jt} = p_j - p_{jt}$. Given that there are only 10 monthly observations per store, the second approach will be used here. Utility u_j at time t is then calculated as the average price for store j over the 10 months minus the price at time t for store j .

Table 6.2 gives the results of the fixed effects regression. Since the R^2 is quite high, shop characteristics explain a large part of the variation in the data. To see whether the store fixed effects are jointly significant, an F -test is performed. As can be seen in Table 6.2, the p -value for the F -test is equal to zero, which suggests that store fixed effects indeed matter.

Figure 6.5(b) graphs a kernel estimate of the utility density function for period 3, where the utilities are the negatives of the residuals of the fixed effects regression. Figure 6.5(b) shows that there is substantial utility dispersion. The utility density is right-skewed, which tells us that

| Variables | price p |
|----------------------|--------------|
| Constant | 32.74 (0.05) |
| Albert Heijn | 1.64 |
| Bas Dirk Digros | -1.62 |
| Bonimarkt | -1.88 |
| C1000 | 1.21 |
| Coop | 2.07 |
| Deen | -0.50 |
| Dekamarkt | -0.50 |
| Edah | - |
| Foodfactory | -1.62 |
| Golff | 1.62 |
| Hoogvliet | -1.62 |
| Jan Linders | -0.23 |
| Jumbo | -1.37 |
| Konmar | 0.86 |
| Nettorama | -1.91 |
| Plus | 1.38 |
| Super de Boer | 2.19 |
| Vomar | 0.29 |
| N | 170 |
| R^2 | 0.85 |
| \bar{R}^2 | 0.83 |
| p -value F -test | 0.00 |

Notes: Standard errors in parenthesis.

Table 6.2: Fixed Effects regression results period 3

although it is possible to encounter utilities of more than 2.5, it happens with small probability. This already gives some indication that the share of consumers searching intensively will not be very large in this market.

The calculated utilities are used in the maximum likelihood procedure described in Section 6.3. The estimation results are presented in Table 6.3(a). The estimates of the parameters of the log-normal part of the SNP function, γ and σ , are highly significant, while only some of the parameters of the polynomial part of the SNP function, the θ 's, are significant at a 5%-level. Table 6.3(b) gives the search intensities derived from the estimated parameters of the SNP distribution. What strikes is that consumers either search for prices at one or two, or at all chains. The estimated share of consumers searching once or twice is around 97%, while only 3% of consumers compare all prices. A similar picture arises when the estimated search cost cdf and pdf are graphed, as in Figure 6.6. The flat part in the lower part of the search cost distribution indicates that consumers either have search cost of more than 0.22 euro, or they have search cost almost equal to zero. Finally, Table 6.3(a) shows that the estimated maximum


| Period 3 | | Period 3 | |
|------------|---------------|------------|------|
| p^n | 5 | μ_1 | 0.81 |
| # obs | 170 | μ_2 | 0.16 |
| γ | 2.882 (0.32) | μ_3 | 0.00 |
| σ | 0.620 (0.18) | μ_4 | 0.00 |
| θ_0 | 1.000 | μ_5 | 0.00 |
| θ_1 | 1.317 (0.81) | μ_6 | 0.00 |
| θ_2 | 0.845 (1.62) | μ_7 | 0.00 |
| θ_3 | -0.484 (1.87) | μ_8 | 0.00 |
| θ_4 | -0.710 (1.39) | μ_9 | 0.00 |
| θ_5 | -0.931 (0.83) | μ_{10} | 0.00 |
| | | μ_{11} | 0.00 |
| x | 6.13 | μ_{12} | 0.00 |
| LL | 914.200 | μ_{13} | 0.00 |
| $KS F(p)$ | 1.00 | μ_{14} | 0.00 |
| $KS L(u)$ | 0.86 | μ_{15} | 0.00 |
| | | μ_{16} | 0.00 |
| | | μ_{17} | 0.00 |
| | | μ_{18} | 0.03 |

Notes: Standard errors in parenthesis.

(a) Parameter estimates
SNP function

Notes: The μ_k 's are calculated using equation (6.4)

(b) Implied μ_k 's

Table 6.3: Estimation results period 3**Figure 6.6:** Estimated search cost distribution period 3

price-cost margin $x = v_j - r_j$ is 6.13 euro.

As can be seen in Figure 6.7(a) the model does quite well in explaining the data since the estimated price cumulative distribution function, as indicated by the solid line, is quite close to the empirical one. The results of a more formal Kolmogorov-Smirnov test are put in Table 6.3(a).¹⁷ Since all $KS F(p)$ values are below the 95%-critical value of the KS -statistic, which is

¹⁷In this table $KS F(p)$ 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.




Figure 6.7: Fit price and utility distributions period 3

1.36, we cannot reject that the prices are drawn from the estimated price distribution. Of course, given that around 85% of the variation can be explained by store fixed effects, a substantial part of the fit in Figure 6.7(a) is due to non-search related causes. Since the utility distribution is derived by controlling for store fixed effects, in principle the fit of the utility distribution is a better indicator for determining to what extent search matters. Figure 6.7(b) shows the estimated utility cdf compared to the calculated utility cdf. As can be seen in this graph, the estimated utility distribution is close to the calculated utility distribution. That the model does quite well in explaining utilities can also be concluded from the corresponding $KS L(u)$ value in Table 6.3(a), which is well below the critical value of 1.36.

6.4.3 Estimation of search costs with time fixed effects

In this subsection search cost distributions for all three periods will be estimated. Because of the declining trend in average price in period 1 and to a lesser extent in period 2, I will allow for time fixed effects in the calculation of utilities. To be able to compare the estimates across periods, the data in period 3 will be corrected for differences in average price across time as well.

Figure 6.8(a) shows kernel estimates of the price density for each period. Because of the declining trend in the first months I take the logarithm of real prices in deviation from the monthly average. It cannot be concluded from this graph whether pricing was more competitive in one period versus another one, but the graph does show that the shape of the price distribution,

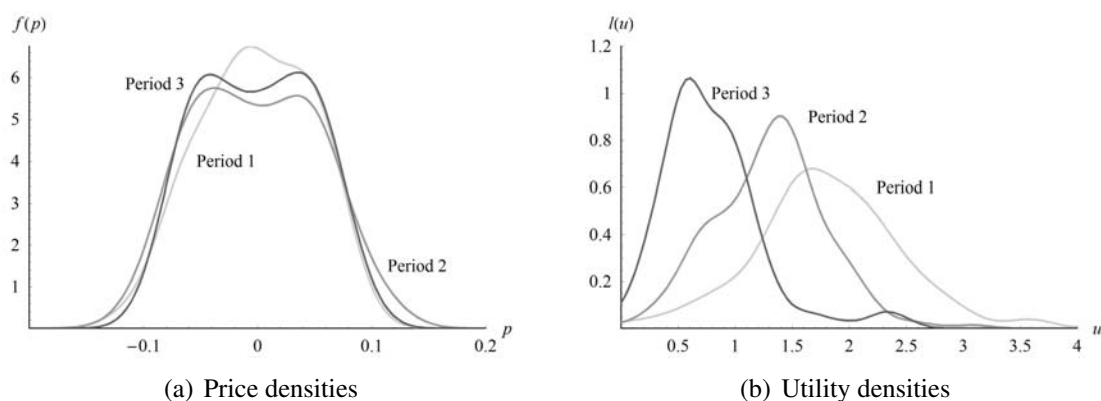


Figure 6.8: Price and utility densities all periods

corrected for the differences in the basket and the declining trend, is different between periods. What causes this remains to be seen.

To correct for the declining average price in the first periods in the calculation of the utilities, in this section I allow for time fixed effects as well. Utility u_j at time t is then calculated as the price at time t for store j minus the average price of the basket at time t minus the average of the time corrected price for store j in the specific period, i.e., $u_{jt} = p_t + p_j - p_{jt}$, where p_j is now the average of $p_t - p_{jt}$ across stores. Notice that adding time fixed effects effectively corrects for the declining average price, but it also eliminates part of the variation in prices caused by the search process, so this should be taken into account when evaluating the results.

Table 6.4 gives the results of the fixed effects regression, including the store fixed effects. Since for all periods the R^2 is quite high, shop characteristics and time effects explain a large part of the variation in the data. Moreover, for each period the results of an F -test to test whether time and store fixed effects are jointly significant suggest that both store and period fixed effects indeed matter.

Figure 6.8(b) graphs a kernel estimate of the utility density function for each period, where the utilities are the negatives of the residuals of the fixed effects regression. Figure 6.8(b) shows that there is substantial utility dispersion. On average, supermarkets have offered the highest utility in period 1 and the lowest in period 3. A likely reason for the difference in average utility between period 1 and period 2 is that the basket of period 1 is bigger in size. However, the basket of period 2 is smaller in size than the basket of period 3 and still average utility is lower in the last period. As will be shown later, this might have to do with changes in the underlying search cost distribution.

| Variables | Period 1 | Period 2 | Period 3 |
|----------------------|--------------|--------------|--------------|
| Constant | 36.27 (0.06) | 31.62 (0.04) | 32.74 (0.04) |
| Albert Heijn | 1.52 | 1.71 | 1.64 |
| Bas Dirk Digros | -2.48 | -2.29 | -1.62 |
| Bonimarkt | -0.33 | -1.30 | -1.88 |
| C1000 | 0.82 | 1.29 | 1.21 |
| Coop | 2.33 | 3.12 | 2.07 |
| Deen | -1.19 | -0.57 | -0.50 |
| Dekamarkt | -0.34 | -0.20 | -0.50 |
| Edah | 1.02 | 1.71 | - |
| Foodfactory | - | -1.93 | -1.62 |
| Golff | - | - | 1.62 |
| Hoogvliet | -1.04 | -1.71 | -1.62 |
| Jan Linders | - | 0.08 | -0.23 |
| Jumbo | -1.80 | -1.42 | -1.37 |
| Konmar | 0.47 | -0.15 | 0.86 |
| Nettorama | -2.81 | -2.08 | -1.91 |
| Plus | 1.24 | 1.30 | 1.38 |
| Super de Boer | 2.39 | 1.92 | 2.19 |
| Vomar | 0.22 | 0.53 | 0.29 |
| N | 150 | 170 | 170 |
| R^2 | 0.89 | 0.92 | 0.92 |
| \bar{R}^2 | 0.87 | 0.91 | 0.90 |
| p -value F -test | 0.00 | 0.00 | 0.00 |

Notes: Standard errors in parenthesis.

Table 6.4: Fixed Effects Regression Results

The estimation results are presented in Table 6.5. The parameter estimates of γ and σ in Table 6.5(a) are all highly significant, while only some of the θ 's are significant at a 5%-level. Table 6.5(b) gives the search intensities derived from the estimated parameters of the SNP distribution. In line with the results from the previous section, consumers either search for prices at one or two, or at all chains. Moreover, the estimates indicate that the share of consumers searching once or twice increases over time, from 86% in period 1 to 89% in period 2 and even 94% in period 3. In contrast, the share of consumers comparing prices from all supermarkets decreases from 12% in period 1 to 8% in period 2 and 6% in period 3. A similar picture arises when the estimated search cost cdf and pdf are graphed, as in Figure 6.9. For all three periods there is a flat part in the lower part of the search cost distribution, indicating that consumers either have search cost of more than 0.20 euro, or they have search cost almost equal to zero.

According to the estimates the share of consumers searching for all prices declines over time, which is also shown in Figure 6.9. However, although the shares of consumers searching

| | period 1 | period 2 | period 3 |
|-------------|---------------|---------------|---------------|
| p^n | 5 | 5 | 5 |
| # obs | 150 | 170 | 170 |
| γ | 2.642 (0.28) | 3.334 (0.31) | 2.572 (0.21) |
| σ | 0.652 (0.11) | 0.782 (0.09) | 0.608 (0.10) |
| θ_0 | 1.000 | 1.000 | 1.000 |
| θ_1 | 1.510 (0.75) | 0.525 (0.68) | 1.311 (0.68) |
| θ_2 | 1.662 (1.61) | -0.568 (0.85) | 0.951 (1.58) |
| θ_3 | -0.094 (1.71) | -0.356 (0.63) | -0.559 (2.04) |
| θ_4 | -0.082 (1.67) | -0.119 (0.95) | -0.678 (1.61) |
| θ_5 | -1.378 (0.78) | 0.716 (0.64) | -1.103 (0.89) |
| x | 4.21 | 3.91 | 3.83 |
| x/\bar{p} | 10.38% | 11.05% | 10.73% |
| LL | 831.438 | 906.433 | 857.248 |
| $KS F(p)$ | 0.94 | 0.57 | 0.98 |
| $KS L(u)$ | 1.80 | 1.68 | 1.39 |

Notes: Estimated standard errors in parenthesis.

(a) Parameter estimates SNP function

| | period 1 | period 2 | period 3 |
|------------|----------|----------|----------|
| μ_1 | 0.50 | 0.61 | 0.76 |
| μ_2 | 0.36 | 0.28 | 0.18 |
| μ_3 | 0.00 | 0.01 | 0.00 |
| μ_4 | 0.00 | 0.00 | 0.00 |
| μ_5 | 0.00 | 0.00 | 0.00 |
| μ_6 | 0.00 | 0.00 | 0.00 |
| μ_7 | 0.00 | 0.00 | 0.00 |
| μ_8 | 0.00 | 0.00 | 0.00 |
| μ_9 | 0.00 | 0.00 | 0.00 |
| μ_{10} | 0.00 | 0.00 | 0.00 |
| μ_{11} | 0.00 | 0.00 | 0.00 |
| μ_{12} | 0.00 | 0.00 | 0.00 |
| μ_{13} | 0.00 | 0.00 | 0.00 |
| μ_{14} | 0.00 | 0.00 | 0.00 |
| μ_{15} | 0.00 | 0.00 | 0.00 |
| μ_{16} | 0.00 | 0.00 | 0.00 |
| μ_{17} | 0.01 | 0.01 | 0.00 |
| μ_{18} | 0.12 | 0.08 | 0.06 |

Notes: The μ_k 's are calculated using equation (6.4).

(b) Implied μ_k 's

Table 6.5: Estimation results




Figure 6.9: Estimated Search Cost Distribution

once and twice increases when moving from period 1 to period 2 and 3, the estimated search cost distributions for higher search costs show a less straightforward pattern: according to Figure 6.9 the search costs for people with search costs above 0.25 actually went down over time.

Table 6.5(a) also gives the estimated maximum price-cost margin $x = v_j - r_j$. Because the baskets are different over time these estimates cannot be directly compared between periods. Therefore, in the table I also give x as a percentage of the maximum price found in each period, so that it can be compared across periods. Clearly, compared to period 1, the maximum price-cost margin went up in period 2 and 3, which gives some indication that supermarkets have

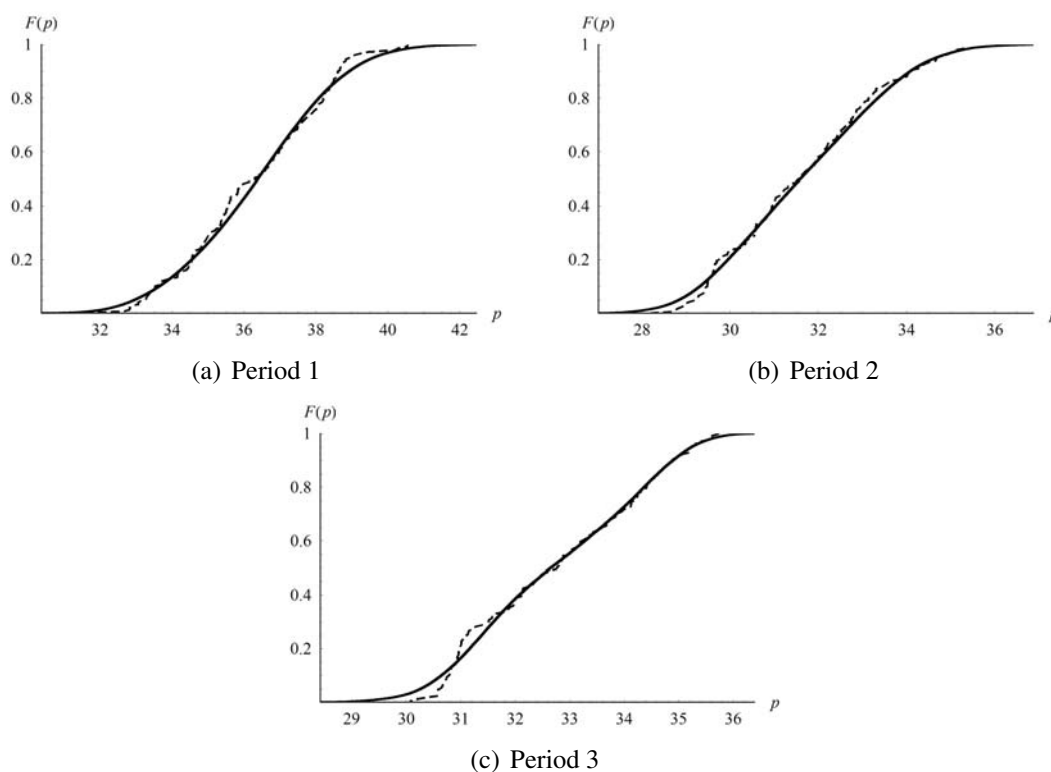


Figure 6.10: Fit price distributions

become more efficient.

As can be seen in Figure 6.10 the model does quite well in explaining the data since all three estimated price cumulative distribution functions, as indicated by the solid lines, are quite close to the empirical ones. The results of a Kolmogorov-Smirnov test in Table 6.5(a) are in line with these observations, since all $KS F(p)$ values are below the 95%-critical value of 1.36. In contrast, the fit of the estimated utility distribution to the calculated utility distribution is less good. Figure 6.11 shows the estimated utility cdfs compared to the calculated utility cdfs. Especially in period 1 and 2 the fit is not so good, which can also be concluded from the corresponding $KS L(u)$ values in Table 6.5(a). The fit in period 3 is already much better. Prices stabilized in this period, so the relatively poor fit of the utility distribution in the first two periods could be due to a non-equilibrium situation caused by the price cutting process. Another reason for the relatively poor fit is that part of the variation in prices is eliminated by controlling for time fixed effects. From Subsection 6.4.2 it is already known that estimating the model for period 3 without controlling for time fixed effects gives a much better fit.

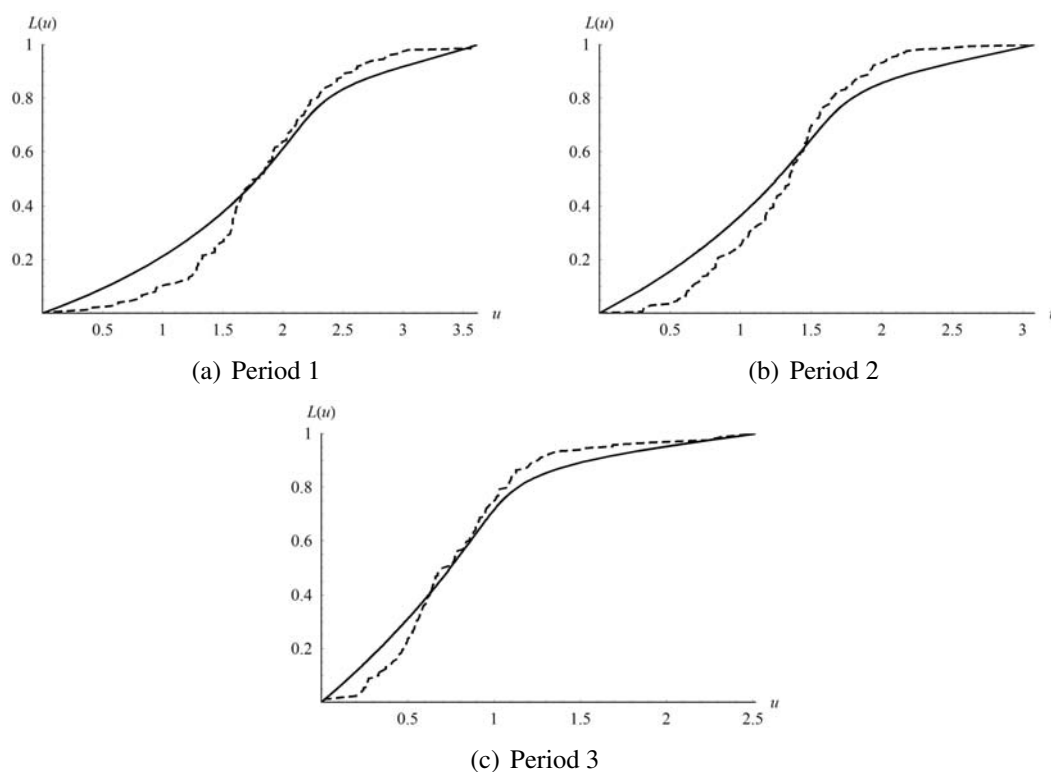


Figure 6.11: Fit utility distributions

6.4.4 Comparison across periods and discussion of results

The estimated search cost distributions in Subsection 6.4.3 point out that the amount of search has decreased over time. In addition, according to Table 6.5(b) the share of consumers with higher search cost went up, although Figure 6.9 points out that this did not lead to a general shift to the right of the search cost cdf. The most intuitive way to interpret search cost is as an opportunity cost of time. A change in search costs over time is then naturally related to a change in income, as time becomes more valuable as people earn more. During the sampling period, changes in income were modest in the sampling period and changes in union monthly wages were even negative in real terms, so although this might explain the decrease in search cost of part of the population, it is not very likely that a shift in real income is driving the changes in search intensity.¹⁸

The changes in search intensity could instead be explained by changes in utility of the

¹⁸According to the Dutch Central Bureau of Statistics, union monthly wages increased 2.5% between January 2004 and June 2006. In real terms there was a decline of 1.5%.

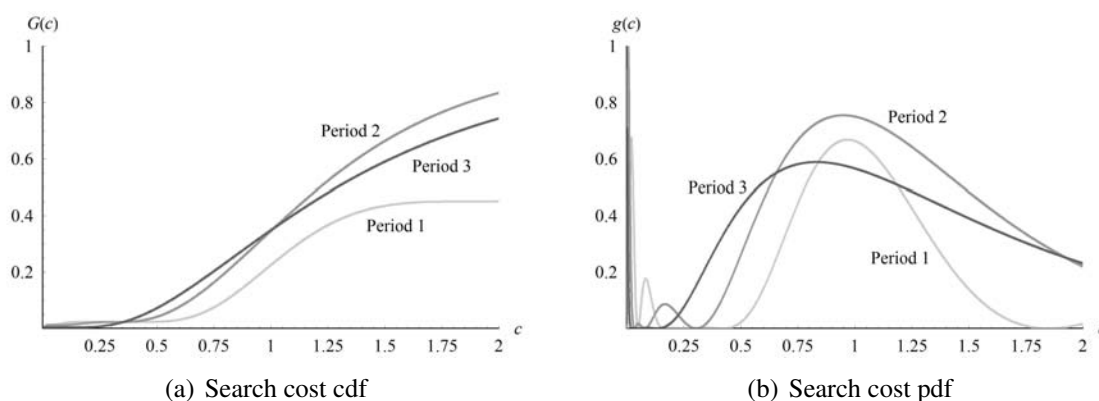


Figure 6.12: Estimated search cost distribution without product differentiation

shopping itself. As the so-called ‘shoppers’ in the search model of Stahl (1989), people may like shopping, even though they have a high opportunity cost of time. These shoppers have essentially search costs equal to zero. Apart from the products being bought, the shopping itself provides utility to these consumers as well, for example, because finding the lowest price gives satisfaction. If people dislike shopping more and more over time, this would decrease the share of shoppers. As a result people become less sensitive to price changes. Such a change in utility of shopping might even be caused by the process of price cutting itself, as a psychological reaction. Even though there are reasons to search, consumers are simply tired of looking for low prices. Indeed, in a report published in June 2004 about the Dutch supermarket price cutting process, consultancy firm Deloitte concludes that since its start in October 2003, less and less consumers change supermarkets because of low prices.¹⁹

An alternative, but related explanation for the changes in search intensity over time would be related to changes in the level of consumer confidence. People care less about finding the lowest price when they believe the economy is doing well. Instead of linking search cost to actual income, in this interpretation it is linked to expected income. Indeed, the consumer confidence level in the Netherlands has increased sharply in the last period compared to the other periods (from on average around -24 in period 1 and 2 to -13 in period 3), so this could explain that over time, less people compare all prices.

An interesting question would be how the estimates of the search cost model with vertical product differentiation compare to estimates from a model without vertical product differenti-

¹⁹See http://www.deloitte.com/dtt/press_release/0,1014,sid%253D13354%2526cid%253D52455,00.html for a press release of Deloitte published on June 23, 2004.

ation. Figure 6.12 gives the estimated search cost cdf and pdf if it is assumed the stores are homogeneous instead of vertically differentiated. What strikes is that estimated search costs are now much higher. Given that around 90% of the variation in prices can be attributed to differences between stores and that this is no longer captured in different valuations across stores but in the prices itself, the gains from searching are much higher in the homogeneous search model. To be able to explain observed prices, the population of consumers should have higher search cost on average and should search less than in the search model with vertical product differentiation. As Figure 6.12(a) shows, the share of consumers comparing prices from all stores is in general much lower now, although in period 1 it is still the highest and in period 3 the lowest. The homogenous search model does only slightly worse in explaining the observed prices, but as reported earlier, it fails to explain patterns at a more detailed level, so on those grounds the homogeneous products model can be rejected for this data set.

6.5 Conclusions

This chapter presented a nonsequential search model that allows for vertical product differentiation. Firms offering distinct products at different prices can be seen as competing in terms of utilities. It is shown that a symmetric mixed strategy equilibrium in utility exists. Firms draw utilities from a common utility distribution, but because valuations and unit costs are different across firms, firms have different price distributions. A result of this is that firms randomize their prices, but over different supports, so that prices are persistently different across firms over time, something which so far could not be explained by existing search models.

It is shown how to estimate the model using price data only. Utilities are calculated by taking the negative of the residuals of a fixed effects regression of prices on store and time dummies. The calculated utilities then serve as an input to a maximum likelihood estimation procedure in order to estimate the underlying search cost distribution. For this a semi-nonparametric estimator is used.

The method is applied to data from Dutch supermarkets in the period January 2004 till June 2006, a period which was characterized by a process of price cutting. The model does reasonably well in explaining observed prices. Moreover, estimates indicate that the amount of search has decreased over time. Since real monthly wages have decreased over time this shift

cannot be explained by a change in real income. Alternatively, the changes can be attributed to a change in the utility shopping gives to consumers, possibly as a result of events related to the price cutting process. An increase in the level of consumer confidence in the final period might explain the decreasing share of consumers comparing all prices as well.

APPENDIX

| Product | Mean Price (Std) | Minimum Price | Maximum Price | Coefficient of Variation ($\times 100$) | Number of Observations |
|--|---------------------|------------------|------------------|--|---------------------------|
| Andrelon Iedere dag shampoo | 2.02 (0.18) | 1.16 | 2.45 | 8.88 | 508 |
| Ariel Color 1.35 kg | 4.76 (0.38) | 3.24 | 5.83 | 8.07 | 471 |
| Axe deodorant | 3.01 (0.43) | 1.97 | 3.99 | 14.14 | 272 |
| Bolletje beschuit blauw 135 g | 0.52 (0.05) | 0.41 | 0.87 | 8.83 | 508 |
| Calve whisky cocktail saus 250 ml | 0.87 (0.06) | 0.57 | 1.19 | 6.99 | 490 |
| Chicken Tonight Hawaii 1 pot | 1.36 (0.14) | 0.86 | 1.65 | 10.21 | 508 |
| Coca Cola 1,5 l | 0.94 (0.13) | 0.64 | 1.25 | 13.75 | 508 |
| Conimex mix voor bahmi speciaal | 0.90 (0.07) | 0.57 | 1.09 | 8.18 | 508 |
| Douwe Egberts Roodmerk Snelfiltermaling 500 g | 2.52 (0.26) | 0.96 | 2.91 | 10.14 | 508 |
| Dubbelfriss sinaasappel mandarijn 1 liter | 0.58 (0.04) | 0.38 | 0.64 | 6.07 | 444 |
| Dubbelfriss sinaasappel mandarijn 1,5 liter | 0.93 (0.03) | 0.83 | 0.97 | 3.53 | 19 |
| Fristi roze normaal 1 liter | 0.92 (0.07) | 0.39 | 1.04 | 7.95 | 317 |
| Goudkuipje klein 100g | 0.55 (0.06) | 0.46 | 0.75 | 10.26 | 508 |
| Hak Bruine Bonen klein 210 g | 0.47 (0.05) | 0.24 | 0.64 | 10.98 | 384 |
| Heineken Bier 24 x 0,3 l krat | 9.00 (0.47) | 6.78 | 9.72 | 5.26 | 253 |
| Hero aardbeienjam 340 g | 1.43 (0.12) | 0.96 | 1.67 | 8.52 | 355 |
| Hertog Slagroomijs 1 liter | 2.43 (0.18) | 1.54 | 2.90 | 7.58 | 463 |
| Honig chinese kippensoep pakje | 0.85 (0.07) | 0.57 | 1.16 | 8.07 | 355 |
| Honig Macaroni vlugkokend (groen) 500 g | 0.55 (0.16) | 0.38 | 0.95 | 29.21 | 107 |
| Honig macaroni vlugkokend kiloprijs | 0.94 (0.04) | 0.78 | 1.23 | 4.55 | 148 |
| Honig mix voor macaroni en spaghetti 62 g | 0.72 (0.07) | 0.28 | 0.95 | 10.03 | 508 |
| Honig Spaghetti vlugkokend (groen) | 0.68 (0.06) | 0.37 | 0.78 | 8.63 | 253 |
| Kellog's Smacks 375 g | 2.09 (0.29) | 1.68 | 2.68 | 13.69 | 508 |
| Knorr Wereldgerechten | 1.74 (0.19) | 1.06 | 1.93 | 10.99 | 17 |
| Krat Amstel bier 24 x 0,3 l | 8.31 (0.51) | 5.93 | 9.19 | 6.14 | 346 |
| Lassie toverrijst 400 g | 0.74 (0.04) | 0.67 | 0.84 | 5.07 | 102 |
| Lays chips naturel | 0.70 (0.05) | 0.57 | 0.93 | 7.57 | 238 |
| Lays chips naturel kiloprijs | 3.30 (0.19) | 2.88 | 3.64 | 5.81 | 17 |
| Optimel drink framboos 1 l | 1.13 (0.10) | 0.66 | 1.34 | 8.55 | 355 |
| Peijnenburg ontbijtkoek 500 g | 0.82 (0.04) | 0.79 | 0.89 | 5.06 | 15 |
| Peijnenburg ontbijtkoek 600 g | 0.95 (0.02) | 0.87 | 0.98 | 1.98 | 76 |
| Pickwick thee Engelse Melange (groen) 20 x 4 g | 0.75 (0.04) | 0.64 | 1.04 | 5.96 | 508 |
| Pringles Naturel 200 gram | 1.44 (0.14) | 0.72 | 1.67 | 9.51 | 403 |
| Rexona deodorant crystal | 2.37 (0.40) | 1.55 | 2.99 | 17.09 | 344 |
| Riedel Appelsientje Sinaasappelsap 1 liter | 0.78 (0.07) | 0.36 | 0.99 | 8.58 | 508 |
| Robijn Black Velvet 750 ml | 3.79 (0.17) | 3.34 | 4.12 | 4.52 | 55 |
| Spa Blauw 1,5 l | 0.42 (0.04) | 0.34 | 0.56 | 10.33 | 508 |
| Sportlife Peppermint 4 stuk | 1.49 (0.09) | 1.33 | 1.79 | 6.18 | 90 |
| Unox Cup-a-soup kip 3 zakjes | 0.66 (0.04) | 0.48 | 0.82 | 6.53 | 508 |
| Verkade knappertjes 220 g | 0.90 (0.06) | 0.64 | 0.99 | 6.73 | 508 |

Notes: Prices are deflated to January 2004.

Table A6.6: Simple statistics for products in the sample

**Summary, achievements and directions
for further research**

7.1 Summary

This thesis focused on the structural estimation of consumer search models. As argued in Chapter 1, because the way search frictions influence markets may depend on how search costs are distributed among consumers, meaningful competition policy requires knowledge of these distributions. Using the structure of several existing and newly developed consumer search models, the various essays of this thesis gave methods to identify and estimate search costs of consumers. An important finding is that in most settings price data is enough to identify how search costs are distributed among consumers, even in a setting where search frictions coexist with quality differences among firms.

The common approach in the different essays of this thesis is to take a particular consumer search model as the starting point for identification of the search cost distribution of consumers. Then using the structure of this consumer search model an equilibrium price distribution can be derived. This equilibrium price distribution characterizes how search costs are related to the pricing strategies of the firms. By taking the inverse of this relation, the quantiles of the search cost distribution can be expressed as a function of the structural parameters of the consumer search model. This inverse relation serves as the starting point for the estimation of search costs. How exactly the estimation works depends on the type of model. For example, we have shown that for the nonsequential search models of Chapters 2, 3 and 6, as well the sequential search model of Chapter 4 maximum likelihood can be used, while the estimation of the sequential search model of Chapter 5 relies on the empirical price distribution.

We showed how to estimate search costs within several theoretical settings. Chapter 2 and 3 focused on nonsequential consumer search models. Chapter 2 discussed how to estimate a nonsequential search model with symmetric firms and heterogeneous consumers in terms of search costs. The starting point of the analysis was an oligopolistic version of Hong and Shum (2006). We explained how to estimate the model using a maximum likelihood procedure. Furthermore, we applied the estimation method to data collected from an online price comparison site on memory chips. According to the estimates, consumers either search for at most three prices or they compare prices of all stores around. Goodness-of-fit statistics indicated that the model does quite well in explaining empirical pricing patterns.

Chapter 3 extended the framework of Chapter 2 to allow for valuation heterogeneity on the

consumer side of the model. In addition, this essay went more into the details of how exactly the identification of nonsequential search models works. We showed that in order to identify the model, one needs data from several markets with the same underlying search cost distribution. In order to connect the different markets, a spline approximation cannot be used. Instead, to exploit the linkage between the different markets via the search cost distribution, we proposed a semi-nonparametric (SNP) density estimator. The parameters of this SNP density function maximize the log-likelihood function corresponding to all the markets together.

The focus of Chapter 4 and 5 was on sequential search. In Chapter 4 we modified the sequential search model of Stahl (1989) by relaxing the assumption that the first price observation is for free. In the first part of this essay we showed that the so-called Diamond result never obtains when search is truly costly, that is, when the first price observation is also costly, not even as a limiting result. In the second part of this essay, the estimation of this model was studied. The estimation results indicated that the model does not do well in explaining empirically observed pricing patterns for memory chips sold online. This result seems to be intimately linked to the assumption that all searchers have the same search cost.

The framework of Chapter 4 was extended in Chapter 5 by allowing consumers to draw their search cost from an atomless search cost distribution. We showed how to identify and estimate the search cost distribution using price data. Interestingly, compared to the nonsequential search model, in this sequential search setting data from one market is enough to identify the search cost distribution. As in Chapter 2, we estimated the model using price data for memory chips obtained from a price comparison site. The estimates indicated that search costs are much higher when it is assumed that consumers search sequentially than when consumers search non-sequentially. Also compared to the nonparametric estimation method of the sequential search model of Hong and Shum (2006) estimated search costs are relatively high.

Finally, Chapter 6 extended the basic framework with homogeneous goods by allowing for vertical product differentiation. We showed that a unique mixed strategy equilibrium exists in terms of utility levels offered to consumers. Because firms differ in their characteristics, the price equilibrium is asymmetric. This model therefore provided a theoretical rationale for explaining price dispersion as a result of both quality differences between firms and search frictions. Using the equilibrium conditions derived from the model, we showed how to estimate the model using only price data. We estimated the model using a data set of products sold at

several grocery stores in the Netherlands. The estimates indicated that most consumers search at most twice, and that only a small fraction of consumers sampled all stores.

7.2 Achievements

The main contribution of this thesis is methodological: the several chapters of this thesis provide methods to retrieve the costs consumers make to obtain price observations using a limited amount of data. Even when part of the price differences seen between firms are due to quality differences, in most cases the only information needed to identify search costs is a data set of prices. As explained in this thesis, this is true for a wide range of models, even without making a priori assumptions about the way search costs are distributed.¹ As shown in Chapter 3, a previously overlooked condition for identification is that in a nonsequential search setting one needs to have price data from several markets, all with the same underlying search cost distribution. In addition, relative to the existing literature on the estimation of consumer search models, our estimation methods do not rely on the assumption that there are infinitely many firms in the market. Instead, we start from an oligopolistic setting, which allows us to study the impact of the number of firms on prices. Moreover, we assume consumers only know the distribution from which prices are drawn and not realized prices. As pointed out by Stahl (1996), since the purpose of our models is to investigate the information acquisition process of consumers, consumers should have no information regarding actual prices.

Although the main focus of this thesis is on the identification and estimation of consumer search costs, the thesis delivers a number of useful byproducts. The starting point for all described estimation procedures is consumer search theory. Several existing theoretical models have been adapted to make them empirically more useful. For example, the nonsequential search model of Burdett and Judd (1983) is extended in Chapter 3 to an oligopolistic setting with consumers that are heterogeneous in their search costs as well as in their valuations. Compared to a model where consumers have one common valuation for the good, after having observed a sample of prices, consumers who observe that their valuation is lower than the prices in their sample now leave the market. This puts some extra competitive pressure on the pricing of the

¹This is true for nonsequential search models. To identify search costs in a sequential search setting in addition one needs to have an estimate of the (common) marginal cost of the firms. Otherwise, only parametric identification is possible.

firms. Another lesson that can be learned from the several search models described in this thesis is that the exact shape of the search cost distribution matters a lot for comparative static results. This can be quite important when doing policy experiments. In Chapter 2 we have shown how to use the estimated search cost distribution to see what happens when there is a change in the tax regime of a particular market. Interestingly, a substantial increase in taxes might be passed on to consumers more than fully, resulting in greater profits for the firms.

When it comes to the actual pricing of firms, another thing that can be learned from this thesis is that most models discussed in this thesis do quite well in explaining observed pricing patterns. Estimated price distributions match very well empirical price distributions, which means that the empirically observed prices could very well be drawn from the estimated theoretical models. Also in other dimensions the models described in this thesis seem to be supported by the data. For example, like Lach (2002) we find that our price data on memory chips and grocery products largely follow a random pattern, which is in line with the mixed strategies we derived as equilibria in all of our search models.

Finally, we can learn from Chapter 5 that the choice between a sequential or a nonsequential search model can have a large impact on the estimates of the search cost distribution. When consumers search sequentially they always have the option to continue searching if an encountered price is higher than their reservation price. This puts much more competitive pressure on prices than in the nonsequential search model, since in the nonsequential search setting consumers will never continue searching after they have gathered the number of prices they committed to before the start of the search process. As a result, to explain the same price distribution estimated consumer search costs are lower in case of sequential search than in case of nonsequential search. Although one could estimate both models to get an idea of the upper and lower bounds of the search cost distribution, to get more specific estimates this means that the researcher should be very careful in deciding which of the two search protocols to start the analysis with. Knowledge of the market being studied will provide natural guidelines, i.e., for markets where the search outcome is observed with some delay, like in markets for labor, mortgages, refurbishing services, etc., nonsequential search is a natural choice, while for other markets sequential search is more appropriate.

7.3 Directions for further research

This thesis showed how to estimate search cost distributions within several theoretical frameworks. The focus was mostly on search frictions, even though price dispersion might also be related to capacity constraints, uncertainty in demand, quality differences between firms, marginal cost differences between firms, or differences in tastes among consumers. Chapter 6 is an attempt to combine search frictions as a source of price dispersion with product differentiation. Consumers are assumed to rank firms in terms of quality and all consumers agree in this ranking. A first direction for future research would be to introduce some heterogeneity in how consumers value quality, since it is likely that different groups of consumers have different tastes for quality. Also, an interesting extension would be to model horizontal product differentiation. In both cases, it is very likely that in addition to price data, sales data are needed to be able to identify the models. A useful starting point for a horizontal product differentiation model with consumer search would be a model along the lines of Sovinsky Goeree (2005), a paper that presents a discrete choice model of product differentiation that allows for heterogeneity in the information different consumers have due to their relative exposure to advertising. A nontrivial difference would be that in a consumer search setting the choice of information set is determined by the consumer, while in an advertising setting the firms decide how much information the consumer will get. Another way to capture firm heterogeneities would be to allow for unequal sampling probabilities. As a result of for example advertising, some firms have more visibility than others, so that they are more likely to be visited by consumers. Hortaçsu and Syverson (2004) allow for unequal sampling probabilities in their model, and show how to estimate these probabilities. However, to be able to identify sampling probabilities, quantity data is needed as well.

Another avenue for future research could evolve around the optimality of the different search protocols. As shown in Chapter 5 estimated search costs according to the sequential search model are in general higher than those according to the nonsequential search model. Without a priori knowledge of the way consumers search it is impossible to distinguish empirically between the two frameworks within the current set of models. To deal with this issue one could think of a model that makes the choice between sequential and nonsequential search endogenous. Another possibility would be a model that combines the two protocols. As shown

by Morgan and Manning (1985) the optimal search strategy is a combination of sequential and nonsequential search. It would be interesting to see the extent to which such a hybrid model outperforms the sequential and nonsequential search models.

A third direction for further research would be to abandon the static framework and focus on the dynamics of the search process. As shown by Fershtman and Fishman (1992), a dynamic environment might lead to results very different than those from the static environment. Also, when consumers face switching costs in addition to search costs, because of for example loyalty cards, advertisement, etc., dynamic effects should be taken into account. Combining switching costs and search costs in a model is nontrivial. Although it is often assumed that switching costs are similar to search costs, Wilson (2006) shows that the anticompetitive effects of search costs are larger than those stemming from switching costs.

Finally, an interesting extension would be to investigate the estimation of a model that allows for marginal cost heterogeneity. Emre, Hortaçsu, and Syverson (2005) show that a model with marginal cost heterogeneity results in firms playing pure strategies in prices. As argued in this thesis, in general models with pure strategy equilibria put less structure on the data than models with mixed strategy equilibria, so to estimate such a model quantity data is needed as well.

Samenvatting

(Summary in Dutch)

Volgens het beroemde artikel van Bertrand (1883) zal zelfs wanneer we te maken hebben met een beperkt aantal bedrijven, de prijs van een homogeen product naar de Walrasiaanse prijs convergeren, mits we in een wereld leven waarin informatie volledig is, waarin bedrijven dezelfde kostenstructuur hebben en waarin transactiekosten geen rol spelen. Zoals wordt uitgelegd in Bertrand (1883), wanneer bedrijven hun prijs bepalen kan ieder bedrijf de hele markt veroveren door simpelweg het product voor een iets lagere prijs te verkopen dan de concurrent. Een uitvloeisel van dit mechanisme is dat alle bedrijven door middel van prijsverlagingen de laagste prijs proberen te krijgen. Dit proces gaat net zolang door totdat de prijs gelijk is aan de marginale kosten.

Het gemak waarmee we twee winkels kunnen vinden die verschillende prijzen voor vergelijkbare producten vragen, laat zien dat de “wet van één prijs” niet altijd opgaat. In plaats daarvan lijkt het alsof prijs spreiding de regel is. Economen zijn door de jaren heen met verschillende verklaringen gekomen voor dit fenomeen. Een eerste verklaring is dat in tegenstelling tot wat wordt verondersteld door Bertrand (1883), in werkelijkheid het vaak het geval is dat producten verschillen in hun kenmerken. Wanneer consumenten een verschillende smaak hebben voor deze kenmerken zijn producten niet langer substituten, wat leidt tot marktmacht aan de kant van de bedrijven. Hierdoor zullen er minder prikkels zijn voor bedrijven om net onder de prijs van de anderen te gaan zitten, waardoor de producten voor verschillende prijzen aangeboden kunnen worden.

Het intrigerende is dat prijs spreiding ook wordt geobserveerd in markten waarin de goederen en de bedrijven vrijwel identiek zijn. Wanneer we ons puur op de kant van de bedrijven richten zijn er in zo’n situatie duidelijk prikkels voor bedrijven om net onder de prijs van

de concurrenten te gaan zitten. Echter, wanneer in tegenstelling tot Bertrand (1883) wordt verondersteld dat consumenten geen volledige informatie hebben over bijvoorbeeld de prijs van het goed bij de verschillende winkels, zal de prikkel om net een lagere prijs te vragen dan de concurrent minder worden. In recent empirisch werk wordt deze gedachte uitgewerkt om te verklaren waarom op het oog identieke goederen verschillende prijzen hebben. Het idee is dat vanwege zoekfricties sommige consumenten geen prijzen vergelijken, waardoor bedrijven een hogere prijs kunnen vragen voor hun product. Immers, deze consumenten zullen iedere prijs accepteren, mits lager dan hun waardering voor het goed. Aan de andere kant zijn er consumenten die wel prijzen vergelijken, wat weer een drukkend effect heeft op de ideale prijs. Om de juiste balans te vinden tussen deze twee prikkels, kan het zo zijn dat bedrijven prijzen aselekt bepalen en veranderen. Dit om aan de ene kant de winst te maximaliseren behaald op consumenten die geen prijzen vergelijken, door een hoge prijs te vragen. En aan de andere kant om de winst te maximaliseren behaald op consumenten die wel degelijk prijzen vergelijken, door een lage prijs te vragen. Dit proces leidt tot prijs spreiding.

Ook als de producten identiek zijn, maar bedrijven een verschillende kostenstructuur hebben, kan er prijs spreiding ontstaan. Echter, dit kan alleen in combinatie met zoekfricties. Het idee hierachter is dat in het geval van volledige informatie, bedrijven met hogere marginale kosten dan hun concurrenten niet kunnen overleven om de reden dat deze bedrijven niet mee kunnen gaan met de lage prijzen van bedrijven met lage marginale kosten zonder verlies te maken. Consumenten kopen immers bij het bedrijf of de winkel met de laagste prijzen. Wanneer deze verschillen in marginale kosten evenwel samengaan met zoekfricties, kunnen de bedrijven met hoge marginale kosten overleven door vooral te verkopen aan consumenten met hoge zoekkosten.

Alhoewel productdifferentiatie, onzekerheid wat betreft de vraag en capaciteitsrestricties eveneens een bron van prijs spreiding kunnen zijn, zal dit proefschrift zich vooral richten op zoekfricties als verklaring voor prijs spreiding. Desondanks zal het raamwerk ook uitgebreid worden naar verticale productdifferentiatie. Dit wil niet zeggen dat onzekerheid wat betreft de vraag en capaciteitsbeperkingen niet belangrijk zijn als veroorzaker van prijs spreiding. In sommige markten zullen capaciteitsrestricties een belangrijke bijdrage leveren aan het ontstaan van prijs spreiding of misschien zelfs de belangrijkste bijdrager zijn. Dit betekent dat de modellen ontwikkeld in dit proefschrift minder van toepassing zijn op die markten.

Bestaande zoekmodellen hebben laten zien dat de manier waarop zoekfricties markten beïnvloeden afhangt van hoe zoekkosten verdeeld zijn over consumenten. Zo kunnen in het sequentiële zoekmodel van Stahl (1989) bij hoge zoekkosten gemiddelde prijzen constant blijven na een fusie, terwijl bij lage zoekkosten het tegenovergestelde gebeurt. Iets soortgelijks zien we gebeuren in het nonsequentiële zoekmodel van Janssen and Moraga-González (2004): of een fusie leidt tot een daling dan wel stijging van de gemiddelde prijs hangt af van de hoogte van de zoekkosten. In meer algemene zin betekent dit dat voor een degelijke analyse van de prijs- en welvaartseffecten van beleidsmaatregelen zoals het verhogen van de btw, het verminderen van toegangsbarrières of het instellen van een maximum prijs, men kennis moet hebben van vraag- en aanbodparameters, waaronder de zoekkostendistributie. Men kan simpelweg een bepaalde distributie veronderstellen, maar gegeven het belang van de exacte vorm van de verdeling kan dit tot de verkeerde conclusies leiden.

Het doel van dit proefschrift is om met methodes te komen om zoekmodellen te schatten. Dit proefschrift bestaat uit vijf essays en kan grofweg verdeeld worden in drie delen. In de eerste twee essays leg ik uit hoe nonsequentiële zoekmodellen te modelleren, te identificeren en te schatten. Het eerste essay beschrijft een methode om zoekkostendistributies te schatten. Het startpunt is een oligopolistische versie van het nonsequentiële zoekmodel beschreven in Hong and Shum (2006). De oligopolieveronderstelling is nuttig omdat het de onderzoeker helpt een onderscheid te maken tussen prijsvariatie veroorzaakt door variatie in het aantal bedrijven en variatie in zoekkosten. Gebruikmakend van de evenwichtscondities afgeleid van het model wordt getoond hoe het model met behulp van *maximum likelihood* geschat kan worden. De schattingsmethode wordt toegepast op een dataset bestaande uit prijzen van geheugenchips die online verkocht worden. De schattingsresultaten suggereren dat consumenten of vrij hoge of juist vrij lage zoekkosten hebben, wat betekent dat consumenten of vrij uitgebreid zoeken of maximaal drie keer. Zoekfricties geven bedrijven een significante hoeveelheid marktmacht: ondanks dat er 20 bedrijven actief zijn in de onderzochte markten vinden we prijs-kostenmarges van ongeveer 25%. Kolmogorov-Smirnov testresultaten suggereren dat de nulhypothese dat de geobserveerde prijzen door het model gegenereerd worden niet verworpen kan worden. In het eerste essay wordt ook geïllustreerd hoe de structurele methodologie gebruikt kan worden om de effecten van beleidsinterventies te simuleren.

Het tweede essay gaat dieper in op hoe de identificatie van zoekkosten precies werkt in

nonsequentiële zoekmodellen. Het model beschreven in het eerste essay wordt uitgebreid naar een meer algemener raamwerk waarin consumenten niet alleen verschillen in zoekkosten, maar ook in waarderingen voor het goed. In het essay laten we zien dat de zoekkostendistributie niet volledig nonparametrisch geïdentificeerd kan worden indien slechts één markt in isolatie wordt bestudeerd. Dit nonidentificatieprobleem kan opgelost worden door een volledig raamwerk te nemen waarin de econometrist observaties heeft van meerdere markten met dezelfde achterliggende zoekkostendistributie. Om de relatie tussen de verschillende markten optimaal te kunnen gebruiken stellen we in het essay voor om de zoekkostenkansdichtheidsfunctie te schatten met behulp van een semi-nonparametrische (SNP) kansdichtheidsschatter. De parameters van deze kansdichtheidsschatter maximaliseren de gezamenlijke aannemelijkheid van alle markten tezamen. Deze in essentie nonparametrische benadering kan gebruikt worden wanneer alle consumenten een identieke waardering hebben voor het goed of wanneer waarderingen verschillen maar niet gecorreleerd zijn met zoekkosten. Indien er wel correlatie is tussen waarderingen en zoekkosten is het alleen mogelijk om het model parametrisch te identificeren.

In de twee daarop volgende essays richt ik mij op het schatten van sequentiële zoekmodellen. In het derde essay veranderen we het artikel van Stahl (1989) over oligopolistische prijsvorming en sequentieel consumenten zoekgedrag door de veronderstelling dat consumenten de eerste prijsobservatie voor niks ontvangen te veranderen. In het theoretische gedeelte van dit hoofdstuk laten we zien dat wanneer alle prijsobservaties kostbaar zijn het zo kan zijn dat niet alle consumenten deelnemen. De verzameling parameters waarvoor *non-shoppers* niet volledig meedoen wordt groter naarmate het aantal *shoppers* afneemt en/of het aantal bedrijven toeneemt. Welvaart neemt af naarmate meer bedrijven de markt betreden. In het essay laten we zien dat de prijs nooit gelijk zal zijn aan de monopolieprijs wanneer zoeken daadwerkelijk kostbaar is. Gebruikmakend van de evenwichtscondities afgeleid uit het model, laten we in het empirische gedeelte van dit hoofdstuk zien hoe het model geschat kan worden. Met behulp van dezelfde dataset als in voorgaande hoofdstukken, schatten we het model. De schattingsresultaten laten zien dat het model niet goed de geobserveerde prijspatronen kan verklaren. Waarschijnlijk is dit gerelateerd aan het feit dat in het beschreven model consumenten of alle prijzen waarnemen of helemaal niet zoeken, wat van grote invloed is op de flexibiliteit van het model.

In het vierde essay wordt een algemener sequentieel zoekmodel bestudeerd. Om het raam-

werk empirisch relevanter te maken, wordt in dit essay verondersteld dat consumenten hun zoekkosten uit een distributiefunctie trekken. In dit hoofdstuk laat ik zien dat zoekkosten met gebruik van alleen prijsdata nonparametrisch geïdentificeerd kunnen worden, zelfs wanneer er data van slechts één markt beschikbaar is. Een vergelijking met het nonsequentiële model laat zien dat om dezelfde prijsdistributie te verklaren hogere zoekkosten nodig zijn in het sequentiële zoekmodel dan in het nonsequentiële zoekmodel. Daarnaast zijn de zoekkosten geschat volgens de methode beschreven in dit essay hoger dan zoekkosten geschat volgens de nonparametrische methode van Hong and Shum (2006).

In het laatste essay van het proefschrift breid ik het raamwerk uit zodat verticale productdifferentiatie ook als een bron van prijsspreiding kan dienen. In het unieke symmetrische evenwicht trekken bedrijven met verschillende kenmerken nutseenheden uit een gemeenschappelijke nutsdistributie. Omdat de bedrijven verschillen in hun kenmerken leidt dit tot verschillende prijsverdelingen. Het model levert daarom een theoretische verklaring voor het verklaren van prijsspreiding als een gevolg van zowel kwaliteitsverschillen als zoekfricties. Gebruikmakend van de evenwichtsrestricties afgeleid uit het model, wordt er in het essay uitgelegd hoe de twee van elkaar te scheiden en hoe het model te schatten. Het model wordt toegepast op prijzen van Nederlandse supermarktproducten. Schattingen van de zoekkosten laten zien dat consumenten in de onderzoeksperiode minder zijn gaan zoeken. Daarnaast leidt het negeren van kwaliteitsverschillen tussen aanbieders tot een te hoge schatting van de zoekkosten van consumenten. Hoewel de ideeën toegepast worden op een non-sequentieel zoekmodel, kan een sequentieel zoekmodel op dezelfde manier uitgebreid worden.

De vijf essays geven een verscheidenheid aan zoekmodellen. Welk model het meest geschikt is om te gebruiken in de praktijk hangt af van de kenmerken van de markt die bestudeerd wordt. Het nonsequentiële zoekmodel is meer geschikt voor markten waarin de uitkomsten van het zoekproces met enige vertraging worden waargenomen, zoals bijvoorbeeld bij het zoeken naar een baan, hypotheek of verhuisbedrijf. Daarnaast hangt de keuze van het model af van de beschikbaarheid van data. Zo is voor het schatten van het model beschreven in het vijfde essay van dit proefschrift alleen prijsdata nodig, terwijl voor een vergelijkbaar model als Hortaçsu and Syverson (2004) daarnaast ook marktaandelen nodig zijn. Echter, het gebruik van alleen prijsdata gaat ten koste van de nauwkeurigheid van het model, omdat het model beschreven in dit proefschrift minder algemeen is dan het model van Hortaçsu and Syverson (2004). In een

situatie waarin de econometrist zowel marktaandelen als prijzen tot zijn beschikking heeft, is het model beschreven in Hortaçsu and Syverson (2004) daarom de beste keus.

References

Simon P. Anderson and Régis Renault (1999). Pricing, product diversity, and search costs: A Bertrand-Chamberlin-Diamond model. *RAND Journal of Economics* 30, 719–735.

Joseph Bailey (1998). Electronic commerce: Prices and consumer issues for three products: Books, compact discs, and software. *Organization for Economic Co-operation and Development OCDE/GD(98)4*.

Michael R. Baye, John Morgan, and Patrick Scholten (2006). Information, search and price dispersion. In T. Hendershott (Ed.), *Handbook on Economics and Information Systems*. Elsevier.

Michael R. Baye, John Morgan, and Patrick Scholten (2007). Persistent price dispersion in online markets. In D. Jansen (Ed.), *The New Economy and Beyond: Past, Present and Future*. Edward Elgar.

Joseph Bertrand (1883). Théorie mathématique de la richesse sociale. *Journal des Savants* 67, 499–508.

Jeffrey Brown and Austan Goolsbee (2002). Does the Internet make markets more competitive? Evidence from the life insurance industry. *Journal of Political Economy* 110, 481–507.

Kenneth Burdett and Kenneth L. Judd (1983). Equilibrium price dispersion. *Econometrica* 51, 955–969.

Karen Clay, Ramayya Krishnan, and Eric Wolff (2001). Prices and price dispersion on the web: Evidence from the online book industry. *Journal of Industrial Economics* 49, 521–539.

Juan Delgado and Michael Waterson (2003). Tyre price dispersion across retail outlets in the UK. *Journal of Industrial Economics* 51, 491–509.

- Peter A. Diamond (1971). A model of price adjustment. *Journal of Economic Theory* 3, 156–168.
- Stephen G. Donald and Harry J. Paarsch (1993). Piecewise pseudo-maximum likelihood estimation in empirical models of auctions. *International Economic Review* 34, 121–48.
- Glenn Ellison and Sara Fisher Ellison (2005). Search, obfuscation, and price elasticities on the Internet. *Mimeo*.
- Onsel Emre, Ali Hortaçsu, and Chad Syverson (2005). E-commerce and the market structure of retail industries. *NET Institute Working Paper* 05-24.
- Victor M. Fenton and A. Roland Gallant (1996). Qualitative and asymptotic performance of SNP density estimators. *Journal of Econometrics* 74, 77–118.
- Chaim Fershtman and Arthur Fishman (1992). Cycles and booms: Dynamic search equilibrium. *American Economic Review* 82, 1221–1233.
- A. Ronald Gallant and Douglas W. Nychka (1987). Semi-nonparametric maximum likelihood estimation. *Econometrica* 55, 363–390.
- William H. Greene (1997). *Econometric Analysis* (3rd ed.). Prentice Hall.
- Bruce Hoadley (1971). Asymptotic properties of maximum likelihood estimators for the independent not identically distributed case. *Annals of Mathematical Statistics* 42, 1977–1991.
- Han Hong and Matthew Shum (2006). Using price distributions to estimate search costs. *RAND Journal of Economics* 37, 257–275.
- Ali Hortaçsu and Chad Syverson (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.
- Daniel Hosken and David Reiffen (2004). Patterns of retail price variation. *RAND Journal of Economics* 35, 128–146.
- Maarten C. W. Janssen and José Luis Moraga-González (2004). Strategic pricing, consumer search and the number of firms. *Review of Economic Studies* 71, 1089–1118.
- Maarten C. W. Janssen, José Luis Moraga-González, and Matthijs R. Wildenbeest (2004). Consumer search and oligopolistic pricing: An empirical investigation. *Tinbergen Institute Discussion Paper* TI 2004-071/1.

- Maarten C. W. Janssen, José Luis Moraga-González, and Matthijs R. Wildenbeest (2005). Truly costly sequential search and oligopolistic pricing. *International Journal of Industrial Organization* 23, 451–466.
- Eric J. Johnson, Wendy W. Moe, Peter S. Fader, Steven Bellman, and Gerald L. Lohset (2004). On the depth and dynamics of on-line search behavior. *Management Science* 50, 299–308.
- Nicholas M. Kiefer and George R. Neumann (1993). Wage dispersion with homogeneity: The empirical equilibrium search model. In H. Bunzel, P. Jensen, and N. Westergård-Nielsen (Eds.), *Panel Data and Labour Market Dynamics*. Amsterdam: North-Holland.
- Tjalling C. Koopmans and Olav Reiersøl (1950). The identification of structural characteristics. *Annals of Mathematical Statistics* 21, 165–181.
- Saul Lach (2002). Existence and persistence of price dispersion: An empirical analysis. *Review of Economics and Statistics* 84, 433–444.
- David S. Moore (1986). Tests of chi-squared type. In R. B. D’Agostino and M. A. Stephens (Eds.), *Goodness-of-Fit Techniques*. New York: Marcel Dekker.
- José Luis Moraga-González, Zsolt Sándor, and Matthijs R. Wildenbeest (2007). Estimation of nonsequential search models. *Mimeo*.
- José Luis Moraga-González and Matthijs R. Wildenbeest (2006). Maximum likelihood estimation of search costs. *Tinbergen Institute Discussion Paper* TI 2006-019/1.
- Peter Morgan and Richard Manning (1985). Optimal search. *Econometrica* 53, 923–944.
- Art B. Owen (1990). Empirical likelihood confidence regions. *Annals of Statistics* 18, 90–120.
- Martin Pesendorfer (2004). Retail sales: A study of pricing behavior in supermarkets. *Journal of Business* 75, 33–66.
- John W. Pratt, David A. Wise, and Richard Zeckhauser (1979). Price differences in almost competitive markets. *Quarterly Journal of Economics* 93, 189–211.
- Jin Qin and Jerry Lawless (1994). Empirical likelihood and general estimating equations. *Annals of Statistics* 22, 300–325.
- Jennifer F. Reinganum (1979). A simple model of equilibrium price dispersion. *Journal of Political Economy* 87, 851–858.

- Rafael Rob (1985). Equilibrium price distributions. *Review of Economic Studies* 52, 487–504.
- Julio J. Rotemberg and Garth Saloner (1986). A supergame-theoretic model of price wars during booms. *American Economic Review* 76, 390–407.
- Alan T. Sorensen (2000). Equilibrium price dispersion in retail markets for prescription drugs. *Journal of Political Economy* 108, 833–850.
- Michelle Sovinsky Goeree (2005). Advertising in the US personal computer industry. *Mimeo*.
- Dale O. Stahl (1989). Oligopolistic pricing with sequential consumer search. *American Economic Review* 79, 700–712.
- Dale O. Stahl (1996). Oligopolistic pricing with heterogeneous consumer search. *International Journal of Industrial Organization* 14, 243–268.
- George Stigler (1961). The economics of information. *Journal of Political Economy* 69, 213–225.
- Hal R. Varian (1980). A model of sales. *American Economic Review* 70(4), 651–659.
- Michael Waterson (2003). The role of consumers in competition and competition policy. *International Journal of Industrial Organization* 21, 129–150.
- Martin L. Weitzman (1979). Optimal search for the best alternative. *Econometrica* 47, 641–654.
- David Whelan (2001). A tale of two consumers. *American Demographics* 22, 54–57.
- Matthijs R. Wildenbeest (2007a). An empirical model of search with vertically differentiated products. *Mimeo*.
- Matthijs R. Wildenbeest (2007b). Structural estimation of a sequential search model. *Mimeo*.
- Chris M. Wilson (2006). Markets with search and switching costs. *ESRC Centre for Competition Policy Working Paper* 06-10.

The Tinbergen Institute is the Institute for Economic Research, which was founded in 1987 by the Faculties of Economics and Econometrics of the Erasmus Universiteit Rotterdam, Universiteit van Amsterdam and Vrije Universiteit Amsterdam. The Institute is named after the late Professor Jan Tinbergen, Dutch Nobel Prize laureate in economics in 1969. The Tinbergen Institute is located in Amsterdam and Rotterdam. The following books recently appeared in the Tinbergen Institute Research Series:

358. S. LONGHI, *Open regional labour markets and socio-economic developments: Studies on adjustment and spatial interaction.*
359. K.J. BENIERS, *The quality of political decision making: Information and motivation.*
360. R.J.A. LAEVEN, *Essays on risk measures and stochastic dependence: With applications to insurance and finance.*
361. N.VAN HOREN, *Economic effects of financial integration for developing countries.*
362. J.J.A. KAMPHORST, *Networks and learning.*
363. E. PORRAS MUSALEM, *theory in practice: Joint replenishments and spare parts control.*
364. M. ABREU, *Spatial determinants of economic growth and technology diffusion.*
365. S.M. BAJDECHI-RAITA, *The risk of investment in human capital.*
366. A.P.C. VAN DER PLOEG, *Stochastic volatility and the pricing of financial derivatives.*
367. R. VAN DER KRUK, *Hedonic valuation of Dutch Wetlands.*
368. P. WRASAI, *Agency problems in political decision making.*
369. B.K. BIERUT, *Essays on the making and implementation of monetary policy decisions.*
370. E. REUBEN, *Fairness in the lab: The effects of norm enforcement in economic decisions.*
371. G.J.M. LINDERS, *Intangible barriers to trade: The impact of institutions, culture, and distance on patterns of trade.*
372. A. HOPFENSITZ, *The role of affect in reciprocity and risk taking: Experimental studies of economic behavior.*
373. R.A. SPARROW, *Health, education and economic crisis: Protecting the poor in Indonesia.*
374. M.J. KOETSE, *Determinants of investment behaviour: Methods and applications of meta-analysis.*
375. G. MÜLLER, *On the role of personality traits and social skills in adult economic attainment.*
376. E.H.B. FEIJEN, *The influence of powerful firms on financial markets.*
377. J.W. GROSSER, *Voting in the laboratory.*

378. M.R.E. BRONS, *Meta-analytical studies in transport economics: Methodology and applications.*
379. L.F. HOOGERHEIDE, *Essays on neural network sampling methods and instrumental variables.*
380. M. DE GRAAF-ZIJL, *Economic and social consequences of temporary employment.*
381. O.A.C. VAN HEMERT, *Dynamic investor decisions.*
382. Z. ŠAŠOVÁ, *Liking and disliking: The dynamic effects of social networks during a large-scale information system implementation.*
383. P. RODENBURG, *The construction of instruments for measuring unemployment.*
384. M.J. VAN DER LEIJ, *The economics of networks: Theory and empirics.*
385. R. VAN DER NOLL, *Essays on internet and information economics.*
386. V. PANCHENKO, *Nonparametric methods in economics and finance: dependence, causality and prediction.*
387. C.A.S.P. SÁ, *Higher education choice in The Netherlands: The economics of where to go.*
388. J. DELFGAAUW, *Wonderful and woeful work: Incentives, selection, turnover, and workers' motivation.*
389. G. DEBREZION, *Railway impacts on real estate prices.*
390. A.V. HARDIYANTO, *Time series studies on Indonesian rupiah/USD rate 1995 - 2005.*
391. M.I.S.H. MUNANDAR, *Essays on economic integration.*
392. K.G. BERDEN, *On technology, uncertainty and economic growth.*
393. G. VAN DE KUILEN, *The economic measurement of psychological risk attitudes.*
394. E.A. MOOI, *Inter-organizational cooperation, conflict, and change.*
395. A. LLENA NOZAL, *On the dynamics of health, work and socioeconomic status.*
396. P.D.E. DINDO, *Bounded rationality and heterogeneity in economic dynamic models.*
397. D.F. SCHRAGER, *Essays on asset liability modeling.*
398. R. HUANG, *Three essays on the effects of banking regulations.*
399. C.M. VAN MOURIK, *Globalisation and the role of financial accounting information in Japan.*
400. S.M.S.N. MAXIMIANO, *Essays in organizational economics.*
401. W. JANSSENS, *Social capital and cooperation: An impact evaluation of a women's empowerment programme in rural India.*
402. J. VAN DER SLUIS, *Successful entrepreneurship and human capital.*
403. S. DOMINGUEZ MARTINEZ, *Decision making with asymmetric information.*
404. H. SUNARTO, *Understanding the role of bank relationships, relationship marketing, and organizational learning in the performance of people's credit bank.*
405. M.Â. DOS REIS PORTELA, *Four essays on education, growth and labour economics.*
406. S.S. FICCO, *Essays on imperfect information-processing in economics.*
407. P.J.P.M. VERSIJP, *Advances in the use of stochastic dominance in asset pricing.*

