WORKING PAPERS

Castellón (Spain)

UNIVERSITAT JAUME•I

On the evolution of monopoly pricing in Internet-assisted search markets

Economics Department

Aurora García-Gallego
Nikolaos Georgantzís
Ainhoa Jaramillo-Gutiérrez
Pedro Pereira
J. Carlos Pernías-Cerrillo

2013 / 05

On the evolution of monopoly pricing in Internet-assisted search markets

Aurora García Gallego
LEE-Universitat Jaume I
Department of Economics
mgarcia@eco.uji.es

Nikolaos Georgantzís
LEE-U. Jaume I, GLOBE-U. of Granada
Department of Economics
ngeorgantzis@ugr.es

Ainhoa Jaramillo-Gutiérrez
University of Valencia
EriCes & Dpt. of Applied Economics
Ainhoa.Jaramillo@uv.es

Pedro Pereira
Autoridade da Concorrência, Portugal
CEFAGE-UE, U. of Evora
pedro.br.pereira@gmail.com

J. Carlos Pernías-Cerrillo
Universitat Jaume I
Department of Economics
pernias@eco.uji.es

2013 / 05

Abstract

We study the evolution of prices in markets assisted by price-comparison engines. We use laboratory data obtained under two industry sizes and two conditions concerning the sample (complete, incomplete) of prices available to informed consumers. Distributions are typically bimodal. One of the two modes, corresponding to monopoly prices, tends to increasingly attract prices over time. The second one, corresponding to interior prices, presents a decreasing trend. Monopoly pricing can be used as an insurance against more competitive (but riskier) behavior. In fact, subjects earning low profits due to interior pricing in the past are more likely to choose monopoly pricing.

Keywords: Internet Economics, price-comparison search engines, mixed strategy equilibria, experimental economics

JEL classification: D0, D2, L1, L4

On the evolution of monopoly pricing in Internet-assisted search markets

Aurora García-Gallego*

LEE & Economics Department, U. Jaume I, Spain

Nikolaos Georgantzís

LEE, GLOBE & Economics Department, UJI and U. de Granada, Spain

Ainhoa Jaramillo-Gutiérrez

Eri-Ces & Department of Applied Economics, U. of Valencia, Spain

Pedro Pereira

Autoridade da Concorrência, Portugal

J. Carlos Pernías-Cerrillo

Economics Department, U. Jaume I, Spain

Abstract

We study the evolution of prices in markets assisted by price-comparison engines. We use laboratory data obtained under two industry sizes and two conditions concerning the sample (complete, incomplete) of prices available to informed consumers. Distributions are typically bimodal. One of the two modes, corresponding to monopoly prices, tends to increasingly attract prices over time. The second one, corresponding to interior prices, presents a decreasing trend. Monopoly pricing can be used as an insurance against more competitive (but riskier) behavior. In fact, subjects earning low profits due to interior pricing in the past are more likely to choose monopoly pricing.

Keywords: Internet Economics; Price-Comparison Search Engines; Mixed Strategy Equilibria; Experimental Economics.

*The authors gratefully acknowledge financial support by the *Spanish Ministry of Economics and Competitivity* (ECO2011-23634), and *Bancaixa* (P1-1B2010-17). N. Georgantzís wishes to acknowledge support by Proyectos de Excelencia de la Junta de Andalucía (P07-SEJ-03155). This paper benefitted from extremely helpful comments by E. Hopkins and M. Sefton.

1 Introduction

A *Search Engine* is a program that accesses and reads Internet pages, stores the results, and returns lists of pages, which match keywords in a query. *Price-Comparison search engines*, also known as shopping agents or shopping robots, are a class of search engines, which crawl commercial Internet sites. In addition to addresses from vendors, they also collect and display other information like prices, or return policies. From the consumers' perspective, one of the most promising aspects of e-commerce was that it would reduce search costs. With search engines, consumers could easily observe and compare the prices of a large number of vendors, and identify bargains. The consumers' enhanced ability of comparing prices would discipline vendors, and put downward pressure on prices. Presumably, the larger the number of vendors whose price a search engine lists on its site, and that thereby consumers can easily compare, the more competitive the market becomes.

The persistence of price dispersion in such markets is a fact which has been extensively confirmed by empirical findings and has been easily accommodated by the theoretical prediction of mixed strategy and, thus, dispersed price equilibria. However, the issue of price dynamics in search markets remains an open question. The existing empirical literature does not clarify whether disperse prices tend to converge towards some stable distribution or whether they follow some other systematic pattern. In this paper, we use a series of laboratory experiments to address this question. We identify two attractors of pricing strategies, which are systematically observed under different market conditions. The first of them is the monopoly price, whereas the second one corresponds to more competitive pricing. We find that the frequency of monopoly pricing systematically exhibits an increasing trend, whereas the mode around which lower prices are distributed presents a declining trend over time.

Since Diamond (1971), we know that the existence of uninformed consumers is crucial (and even sufficient) for monopoly pricing to be sustained by competing sellers in homogeneous good markets. This has been known as the *Diamond Paradox*, as it contradicts the intuitive principle that price competition in the absence of product differentiation should be expected to cause prices to fall close to the competitive level. Subsequently, Varian (1980), Rosenthal (1980), Burdett and Judd (1983) and Stahl (1989) have taught us the fundamental theory on price competition in markets in which informed and uninformed buyers coexist. Among the common properties of the resulting models, the most appealing one concerns the lack of pure strategy equilibria, leading to the theoretical prediction of equilibrium price dispersion. This property has linked these models with a growing theoretical and empirical literature on search markets assisted by Internet-based price-comparison search engines. Along this line, Brynjolfsson and

Smith (2000), Baye and Morgan (2001), Iyer and Pazgal (2000) and Kocas and Kiyak (2006) are representative examples of a recent second wave of papers on search markets, paying special attention to the role of the Internet on market competition. Apart from the novelty of focus on Internet-based search processes, a recurring endeavor of these papers is to explain the systematic empirical finding of persistent price dispersion even in homogeneous product markets.

Several experimental studies have addressed the comparative statics of price dispersion in markets with buyer search. Morgan et al. (2006) report experimental results confirming the basic comparative statics prediction of the Varian (1980) model. Also, Orzen (2008) offers some evidence for the conjecture by Janssen and Moraga-González (2004) concerning the collusive-pricing attractor which tends to reverse Varian (1980) comparative statics prediction regarding the size of an industry. The rotation of theoretical price distributions around the empirical ones reported in García-Gallego et al. (2004) may also be due to these coexisting equilibria. However, the dynamics of strategies in search markets with only mixed strategy equilibria has received less attention.

Diamond (1971) was the first to investigate the stability of price distributions in the presence of buyer and seller learning. In fact, that study goes beyond posterior analysis in that it proposes monopoly pricing as the only equilibrium which is stable under a specific type of learning dynamics. More recently, a small number of theoretical papers study the dynamic properties of price distributions in markets with consumer search. Hopkins and Seymour (2002) have shown that a broad family of learning dynamics may be stable under a relatively demanding condition on the proportion of uninformed consumers in the market ("sufficient ignorance"). In fact, these authors extend the existing learning models to allow for large (almost continuous) strategy spaces, making them applicable to models of price competition. More recently, Benaïm et al. (2009), recalling Shapley (1964) best response cycles in games with mixed strategy equilibria, have shown that the time average of prices set by adaptively learning subjects may converge close to the Nash equilibrium prediction even under conditions which do not guarantee the stability of the equilibrium price dispersion.

From an empirical point of view, Cason and co-authors (Cason and Datta, 2006; Cason and Friedman, 2003; Cason et al., 2005) have paid attention to the dynamics of price dispersion using laboratory data. Especially, the Edgeworth cycles reported by Cason et al. (2005) can be seen as evidence for the instability of price dispersion predicted by Hopkins and Seymour (2002). Also, the serial correlation of individual strategies detected by Cason and Friedman (2003) can be interpreted as evidence against the hypothesis of mixed strategy play. Edgeworth cycles were also observed in experimental markets under Bertrand-Edgeworth competition in Kruse et al. (1994) and price-and-quantity competition in Guillén (2004). Furthermore, a number of empirical studies like Ross (1997) and Busse (2002) also provide evidence of Edgeworth price cycles in posted price markets.

The aforementioned empirical evidence of price cycles is compatible with the theoretical explanation offered by Hopkins and Seymour (2002) concerning the dynamics of learning in markets with only mixed strategy equilibria. In such markets, cycles of best responses emerge like in the well-known children's game of Rock-Paper-Scissors. Specifically, cycles emerge due to firms' switching from monopoly pricing, aimed at extracting maximal surplus from uninformed buyers, to competitive pricing, aimed at attracting the informed ones. Thus, the instability of these two strategies, monopoly and competitive pricing, is responsible for both the nonexistence of pure strategy equilibrium and the cyclical dynamics. Studying in depth the anatomy of the dynamic patterns of observed behavior requires focusing on the evolution of pricing with respect to these two poles of attraction. Furthermore, our interest in the evolution of monopoly pricing separately from other price observations is also dictated by a number of results which do not necessarily rely on the informed-uninformed consumers assumption. For example, while Varian (1980) predicts agglomeration on monopoly pricing as a result of rational behavior, Baye and Morgan (2004) directly relate the frequency of monopoly pricing to the level of bounded rationality in the market. Furthermore, although all the theoretical models discussed here contemplate risk neutral agents, it is straightforward to see why risk aversion would lead to more frequent monopoly pricing. This would happen because monopoly pricing yields a certain payoff while competitive pricing involves some probability of a higher profit but also the risk of a lower one. Therefore, for a variety of reasons, the agglomeration of strategies on monopoly pricing may affect both the comparative statics and the dynamics of observed behavior in markets with informed and uninformed consumers. Obviously, the evolution of monopoly pricing in such markets is not incompatible with other dynamic patterns like Edgeworth cycles.

In this paper, we report results from a laboratory experiment designed to study the evolution of monopoly pricing in the context of the Varian (1980) and Burdett and Judd (1983) models. We design four treatments which are inspired by the comparative statics of the former concerning the market size, and of the latter concerning the completeness of the sampling process available to informed consumers. Contrary to the majority of previous experimental studies[1], our experimental subjects are faced with the same history of rival prices, given that each subject is faced with a number of simulated rivals whose behavior is extracted from mixed strategy equilibrium distributions. Thus, deviations from equilibrium play cannot be explained as the result of noise

[1]Cason and Datta (2006) adopt a similar strategy, but they simulate optimal consumer search behavior.

due to simultaneously learning and strategically interacting agents. Furthermore, each subject's behavior can be treated as an independent observation. We identify two coexisting dynamic patterns occurring parallel to each other. The two patterns concern two alternative *peaks* of typically bimodal price distributions. The first peak is labelled as the *interior pricing* mode, whereas the second is referred to as the *monopoly pricing* one. We also show that the dynamics affecting the two modes occur in opposite directions and that each one of the corresponding patterns remains invariant across different experimental conditions.

The paper is organized in the following way: Section 2 outlines the theoretical framework. Section 3 describes the experimental design. Section 4 reports the results and section 5 concludes.

2 Theoretical framework

We briefly discuss the features and results of the theoretical setting.

Consider an electronic market for a homogeneous search good that opens for one period. There are: 1 price comparison search engine, $n \geq 3$ vendors, which we index through subscript $j = 1, \ldots, n$, and many consumers. The *Price-Comparison Search Engine*, lists the firms contained in its *Index*, and the prices they charge.

Denote by k the number of vendors indexed by the price-comparison search engine. We will refer to k as the *Size of the Index*. The search engine has *Complete Coverage* if it indexes all vendors present in the market: $k = n$. If $k < n$, the search engine has *Incomplete Coverage* and indexes each of the n vendors with the same probability k/n.[2]

There is a unit measure continuum of risk neutral consumers. Each consumer has a unit demand, and a reservation price of 1. There are 2 types of consumers, differing only with respect to whether they use the price-comparison search engine. *Non-Shoppers*, a proportion $\lambda \in (0, 1)$ of the consumer population, do not use the price-comparison search engine, perhaps because they are unaware of its existence, or perhaps because of the high opportunity cost of their time. The other consumers, a proportion $1 - \lambda$, are *Shoppers* and use the price-comparison search engine.

Consumers do not know the prices charged by individual vendors. Shoppers use the price-comparison search engine to learn the prices of vendors. Provided that the lowest price sampled by the price-comparison search engine is no higher than 1, all shoppers buy a unit of the product from the cheapest seller on the engine; in the case of a tie, they distribute themselves randomly among vendors; otherwise they reject the offer and exit the market. Non-shoppers distribute

[2]We focus on the case of unbiased incomplete coverage. A more complete version of the model accounting for biased incomplete coverage can be found in García-Gallego et al. (2004).

themselves evenly across vendors, i.e., each vendor has a share of non-shoppers of $1/n$. Given a price no higher than 1, each non-shopper buys a unit of the product from the firm to which he/she is randomly assigned; otherwise they reject the offer and exit the market.

Vendors are identical and risk neutral. Marginal costs are constant and equal to zero. Vendors know the functioning rules of the search engine. They know the probability with which they are indexed, but do not observe the identity of the indexed vendors, before choosing prices.

Denote by $\Pi_j(p)$, the expected profit of vendor j when it charges price p on $\mathbb{R}_0^+$. A vendor's *strategy* is a cumulative distribution function over prices, $F_j(\cdot)$. A vendor's *payoff* is its expected profit.[3]

A *Nash equilibrium* is a n-tuple $\{F_1(\cdot), \ldots, F_n(\cdot)\}$ of cumulative distribution functions over prices such that for some $\Pi_j^\star$ on $\mathbb{R}_0^+$, and $j = 1, \ldots, n$, $\Pi_j(p) = \Pi_j^\star$, for all p on the support of $F_j(\cdot)$, and $\Pi_j(p) \leq \Pi_j^\star$, for all p.

Denote by τ the type of the search engine, and let 'C', mean *Complete Coverage*, and 'I' mean *Incomplete Coverage*. Then, τ belongs to $\{C, I\}$. We will use superscripts 'C' and 'I' to denote variables or values associated with the cases where the search engine has that type. Denote by ϕ_j^τ the probability of firm j being indexed, such that $\phi_j^\tau = k/n$. Ignoring ties, the expected profit of a vendor that charges $p \leq 1$ is:

$$\Pi_j(p) = p\frac{\lambda}{n} + p(1-\lambda)\phi_j^\tau[1 - F_j^\tau(p)]^{k-1}. \tag{1}$$

Denote by l_j^τ the lowest price vendor j is willing to charge to sell to both types of consumers when the search engine has type τ, i.e., $l_j^\tau[\lambda/n + (1-\lambda)\phi_j^\tau] - \lambda/n \equiv 0$.

Then, it is shown that, in the symmetric equilibrium[4]:

$$p\frac{\lambda}{n} + p(1-\lambda)\phi_j^\tau[1 - F_j^\tau(p)]^{k-1} = \frac{\lambda}{n}. \tag{2}$$

All vendors are indexed with positive probability (including the extreme case of probability 1 when complete coverage is in place). Hence, they face the trade-off of charging a high price and selling only to non-shoppers, or charging a low price to try to sell also to shoppers, which leads them to randomize over prices. Then, the price distribution for the case in which the market consists of n vendors, and the price-comparison search engine has an unbiased index of size $k \leq n$, is identical to the price distribution for the case in which the price-comparison search engine has *Complete Coverage*, $k = n$, and the market consists of k vendors: $F^I(\cdot; n, k) = F^C(\cdot; k)$.

As discussed in detail in García-Gallego et al. (2004), an unbiased decrease in the size of the index has two impacts. First, for indexed vendors, the decrease in the size of the index reduces

[3] It is well known that this game has no equilibrium in pure strategies.

[4] According to Baye et al. (1992) there is also a continuum of asymmetric equilibria.

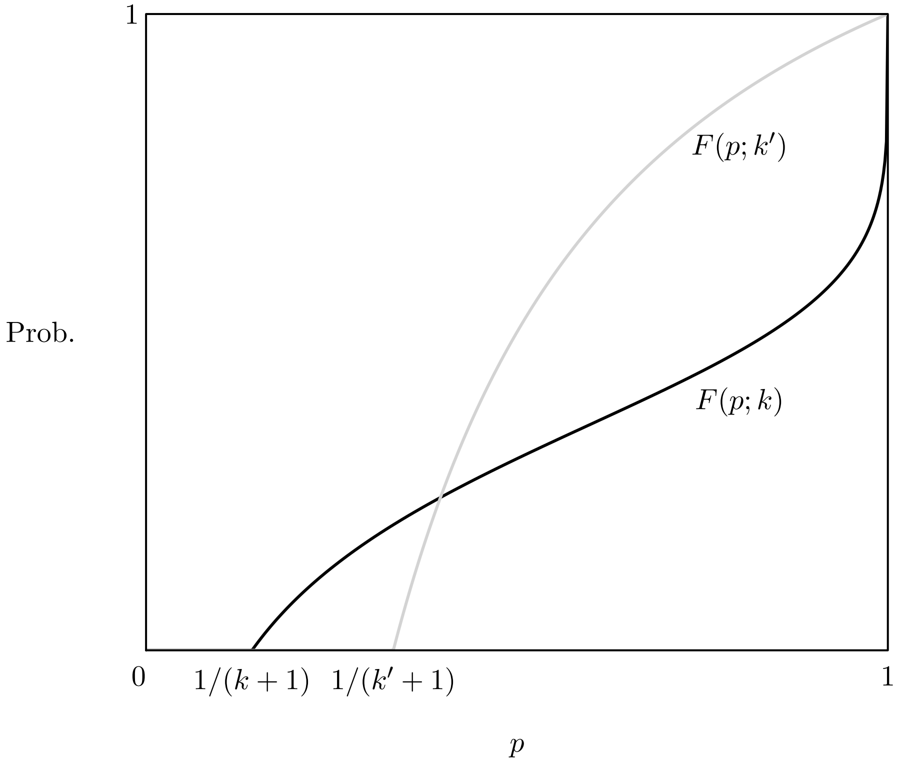

Figure 1: A reduction of the index size

The symmetric equilibrium cumulative distribution function of prices rotates counterclockwise when the size of the index k decreases to k'.

the number of rivals with which a vendor has to compete to sell to shoppers from $k-1$ to $k-2$. This increases the probability that an indexed vendor will have the lowest price, $(1-F^I)^{k-1}$, which increases the *Volume of Sales effect.* The first impact leads vendors to shift probability mass from higher to lower prices. As a consequence, the price distribution shifts to the left. Second, the decrease in the size of the index reduces the probability that a given vendor is indexed from k/n to $(k-1)/n$, which reduces the *Volume of Sales effect.* The second impact leads vendors to raise the lower bound of the support, and to shift probability mass from lower to higher prices. As a consequence, the price distribution rotates. The total impact of an unbiased decrease in the size of the index is to cause the price distribution to rotate counter clock-wise, see Figure 1.

The increase in the lower bound of the support, $l^I(k) < l^I(k-1)$, raises the expected price paid by shoppers. However, the average price paid in the market remains constant and equal to λ. This implies that the expected price by non-shoppers decreases. Recall that vendors now charge lower prices with a higher probability. Shoppers and non-shoppers have conflicting interests with respect to *Incomplete Coverage*, as compared with *Complete Coverage.* Shoppers prefer a large to a small unbiased index, and non-shoppers prefer a small to a large unbiased index.

Under *Incomplete Coverage*, the equilibrium price distribution does not depend on the number of vendors in the market, $F^I(\cdot;n,k) = F^I(\cdot;n{+}1,k)$. The probability with which a vendor is indexed, k/n, depends on the number of vendors. Besides, each vendor's share of non-shoppers, λ/n, also depends on the number of vendors.

Given that $F^I(\cdot;n,k) = F^C(\cdot;k)$, comparing the price distributions under *Incomplete Coverage* and under *Complete Coverage* is equivalent to comparing $F^I(\cdot;n,k)$ and $F^I(\cdot;n,n)$, i.e., is equivalent to analyzing the impact of an increase in the size of the index, under *Incomplete Coverage.* Thus, compared with *Complete Coverage*, *Incomplete Coverage* causes the price distribution to rotate counter-clockwise, which increases the expected price paid by shoppers and decreases the expected price paid by non-shoppers.

3 Experimental design

We present data from 4 experimental conditions implemented in a market environment like the one described above. The 4 treatments (C3, I3, C6, I6) correspond to the combination of two different industry sizes, $n \in \{3,6\}$, with two sizes of the sample (*Complete*, *Incomplete*) of prices available to shoppers searching on the price-comparison search engine. In terms of the theoretical model, the design can be used to test the model's comparative statics predictions

Table 1: Design parameters

Treatment	n	k	ϕ	λ	Price Mean (s. d.)	Min. Price	$\Pr(p = 1)$
C3	3	3	1	1/2	0.60 (0.25)	0.25	1.83%
I3	3	2	2/3	1/2	0.55 (0.18)	0.33	0.05%
C6	6	6	1	1/2	0.70 (0.31)	0.42	17.56%
I6	6	4	2/3	1/2	0.65 (0.28)	0.20	6.30%

concerning the size of the market under complete coverage by comparing a *complete coverage triopoly* (C3) *to a complete coverage hexapoly* (C6). Furthermore, by comparison of C3 to I3 and C6 to I6, we can test the model's hypotheses concerning the completeness of the index. In fact, both incomplete coverage treatments are run with a 2/3 probability of a seller being indexed by the search engine.[5] In all treatments, the λ parameter was set to 1/2.

Table 1 presents the details of our design and the statistics corresponding to the theoretical price distributions. Apart from the hypothesis of equality between the observed and the theoretical means, these statistics help us formulate the following qualitative testable hypotheses emerging from the theoretical framework:

H1: An increase in the number of firms in the market (from C3 to C6 and from I3 to I6) leads to a higher average price.

H2: An increase in the number of firms indexed by the search engine (from I3 to C3 and I6 to C6) leads to a higher average price.

These two hypotheses provide a test for the main prediction of the Varian (1980) and Burdett and Judd (1983) models respectively. We turn now to the main objective of our paper which is the frequency and evolution of monopoly pricing. We state the main comparative statics predictions:

H3: An increase in the number of firms in the market increases the frequency of monopoly pricing.

H4: An increase in the number of firms indexed by the search engine increases the frequency of monopoly pricing.

As an alternative to these four hypotheses, the conjecture by Janssen and Moraga-González (2004) concerning the collusion facilitating role of a decrease in the size of the industry tends to reverse the effects predicted here.

[5] Thus, in treatment I3 we use $k = 2$ and in treatment I6 we set $k = 4$.

Regarding the behavior of individual prices over time, the hypothesis of mixed strategy play implies that a firm's pricing strategy in period t in independent from its strategy in period $t+1$. The most important testable implication of this hypothesis is:

H5: A firm's probability of setting a given price (including monopoly pricing) remains invariant over time.

In order to avoid a noisy learning environment[6] we face each subject with a number of robot players whose behavior is randomly extracted from the equilibrium mixed strategy distributions described in the previous section. This allows us to face all subjects with the same history of rival prices, creating a more controlled and statistically robust sample of individual reactions to the same history as far as other players' strategies are concerned. Specifically, each subject participating in treatments C3 and I3 is faced in each period with 2 simulated players, whereas each subject in treatments C6 and I6 is faced with another 5 simulated players. A total of 45 independent observations (a series of 50 period strategies per subject) are obtained under each treatment. We ran 10 sessions at the Laboratori d'Economia Experimental of the Universitat Jaume I in Castellón (Spain). In each session, 18 subjects were randomly assigned to two different treatments under the same price-sampling condition but with different industry sizes. A total of 180 subjects participated in the experiments, earning an average of approximately 20 Euros each. They were recruited among the students attending business-related courses at the Universitat Jaume I.

The prices set by robot players were randomly extracted from the theoretical distributions described above, using the inverse transformation method (e.g. Cameron and Trivedi, 2005).[7] In order to guarantee that observed heterogeneity can be genuinely attributed to differences across individuals, we have maintained the actual realizations of the simulated players fixed within each treatment. Subjects were told that they were faced with computer- simulated rivals whose actions were programmed to pursue exactly the same objectives as those induced to them.[8] In that sense, the noise due to strategic interaction among simultaneously learning human agents is eliminated creating a controlled environment in which subjects are exposed to similar rival price histories extracted from exactly the same mixed strategy equilibrium play.

[6]Warnick and Hopkins (2006) and Cason and Friedman (2003) warn us on the difficulties entailed in experimental tests of mixed strategy equilibria due to the fact that an individual agent is learning in the presence of noise created by rivals learning strategies.

[7]The experiment was run using a specific software programmed in Java.

[8]No mention was made to equilibrium mixed strategy play by the computer- simulated agents because it could be mistakenly interpreted as correct play, yielding undesirable myopic imitation. See the instructions in the appendix.

4 Data analysis and main results

We present here the results obtained from the experimental sessions. Table 2 provides descriptive statistics for observed period average prices and profits. The evolution of these variables is plotted in Figures 2 to 4. We express profits as a proportion of the corresponding monopoly profits:

$$\pi_{it} = \frac{p_{it} q_{it}}{\lambda / n_i}, \tag{3}$$

where i indexes subjects, t indexes time periods, p_{it} and q_{it} are the prices and sales of firm i in period t, and λ/n_i is the profit level that a firm obtains by setting $p = 1$ and selling only to its captive consumers. Then, $\pi_{it} = 1$ means that subject i earned the monopoly profits in period t. This monopoly profit level can be thought of as a reference point for subjects, because it is the level of earnings that can be reached independently of others' strategies. This strategy can be used by subjects as an insurance device against the riskier but potentially more profitable alternative of interior prices aimed at capturing clients on the search engine in case the firm is the minimum-price seller on the index.

To further analyze our experimental data, we focus on individual pricing behavior. Table 3 reports estimates of the dynamic regression equation

$$\begin{aligned} \ln p_{it} = \beta_0 + \beta_1 \ln p_{i,t-1} + \beta_2 \ln \pi_{i,t-1} + \beta_3 \ln rp_{i,t-1} \\ + \beta_4 t + \beta_5 i3_i + \beta_6 c6_i + \beta_7 i6_i + u_i + \epsilon_{it} \end{aligned} \tag{4}$$

$rp_{i,t-1}$ is the previous period logarithm of average rival prices, $\pi_{i,t-1}$ is the past period profits of subject i, $i3_i$, $c6_i$ and $i6_i$ are treatment dummy variables, u_i is an unobserved variable capturing unobserved time-invariant heterogeneity and ϵ_{it} is a regression disturbance. We also present in Table 4 estimates of the dynamic linear probability model:

$$\begin{aligned} \ln mp_{it} = \alpha_0 + \alpha_1 \ln mp_{i,t-1} + \alpha_2 \ln \pi_{i,t-1} + \alpha_3 \ln rp_{i,t-1} \\ + \alpha_4 t + \alpha_5 i3_i + \alpha_6 c6_i + \alpha_7 i6_i + v_i + \omega_{it} \end{aligned} \tag{5}$$

where mp_{it} is an indicator which take the value 1 if subject i set the monopoly price in period t, and 0 otherwise, v_i stands for unobserved time-invariant subject heterogeneity, and ω_{it} is a regression disturbance. Thus, our approach accounts for the substantial individual heterogeneity present in our data and the underlying dynamics.

The estimates reported in Tables 3 and 4 were obtained by the GMM dynamic panel estimators developed by Arellano and Bond (1991), Arellano and Bover (1995), and Blundell and Bond (1998). These estimators are valid under very general situations. In particular, we do not need strong distributional assumptions on the disturbances, nor ruling out correlation among the individual effects and the right hand side variables. The estimators allow for regressors that

Table 2: Descriptive statistics: Observed prices and profits

Treatment	All prices	Interior Pricing	Mon. Pricing	Profits
C3	0.56 (0.04)	0.47 (0.07)	0.17 (0.06)	0.89 (0.24)
I3	0.55 (0.05)	0.48 (0.05)	0.14 (0.05)	0.86 (0.36)
C6	0.47 (0.05)	0.35 (0.08)	0.18 (0.05)	0.95 (0.35)
I6	0.49 (0.07)	0.41 (0.08)	0.13 (0.05)	1.02 (0.56)

Price and profit averages; standard deviations in parentheses; *Mon. Pricing* corresponds to the observed frequency with which a subject sets $p = 1$. Statistics provided under the heading *Interior pricing* refer to prices satisfying $p \neq 1$.

are not strictly exogenous and are robust to arbitrary conditional heteroscedasticity, and serial correlations among the observations of a given individual, provided that observations of different individuals are independent, a condition that our experimental setup guarantees. These estimators are designed for a sampling scheme of small T-large N but our sample consists of 180 individuals observed during 50 consecutive periods. So, we have used "stacked" instruments as suggested by Arellano (2003, section 8.7). Also, in Tables 3 and 4 we present estimates of extended models that allow for two lags of the regressors. We do not find that the additional lags are significant. Also the specification tests do not detect problems of serial correlation, suggesting that our simple specifications (4) and (5) adequately capture the dynamics in our data.

From the statistics of Table 2, it is interesting to observe that 6-firm markets yield profits which very close to monopoly pricing. On the contrary, 3-firm markets yield profits which are lower than the profit which could be earned by subjects adopting the *safe* option of monopoly pricing. Anyway, we cannot reject the hypothesis that the average profit equals 1 in any of the treatments. The dispersion of profits is increasing in the number of firms present in the market and is greater under incomplete than under complete sampling. This pattern can be explained as a consequence of the increased uncertainty as we move from the complete to the incomplete sampling case. Figure 4 shows the evolution of average profit per treatment. We do not observe any systematic pattern in the temporal evolution of profits. We cannot reject the hypothesis of serial independence according to the results of run tests. Moreover, there are not obvious trends in these data.

From simple inspection of Table 2 and Figures 2 and 3, we can see that the means, standard deviations and the distributions of observed prices significantly deviate from the corresponding theoretical ones. On the contrary, from Figure 5 we see that the counter-clockwise rotation of price distributions predicted as the result of a decrease in the number of firms and the size of

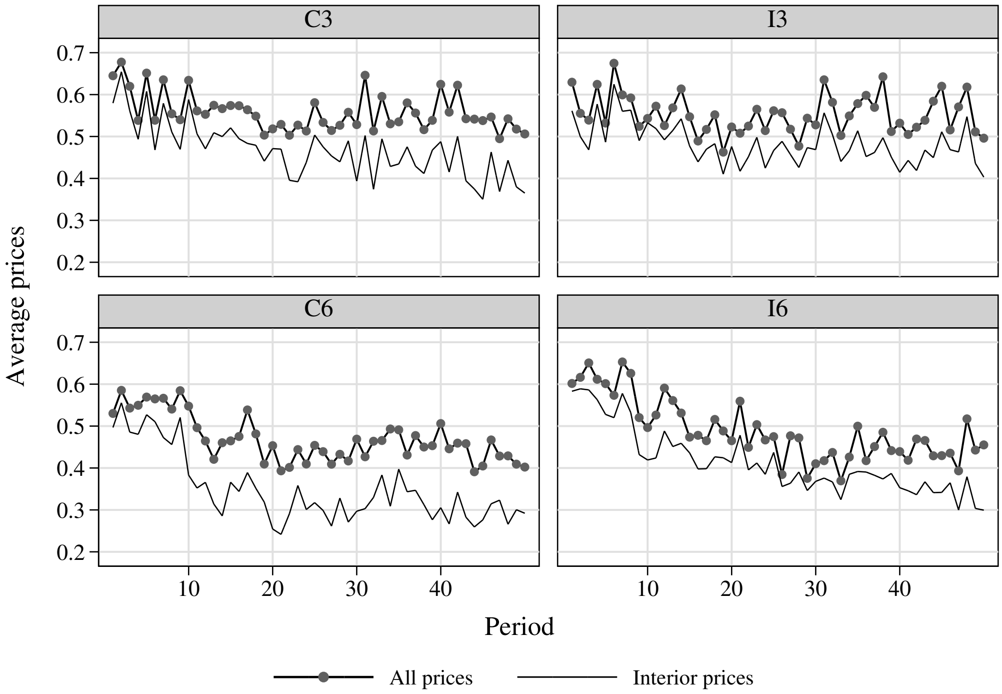

Figure 2: Evolution of price average and interior price ($p < 1$) average.

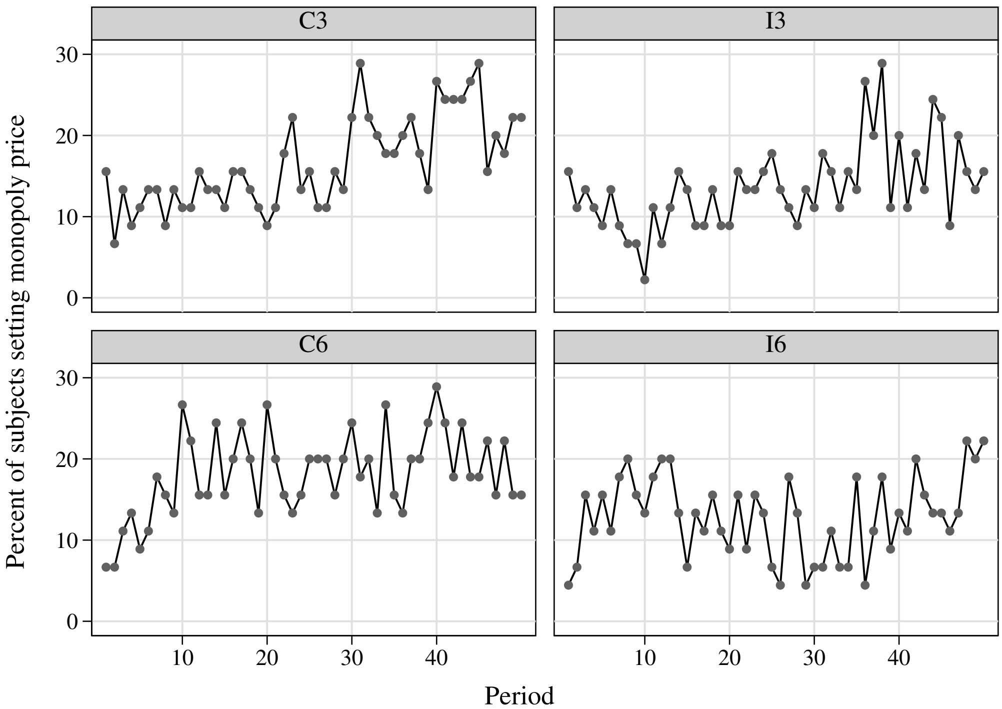

Figure 3: Evolution of the frequency of monopoly pricing

Table 3: Price regressions

	GMM (1 lag)	GMM (2 lags)
$\ln p_{i,t-1}$	0.1111 (0.0316)***	0.0823 (0.0477)*
$\ln p_{i,t-2}$		−0.0212 (0.0292)
$\ln rp_{i,t-1}$	0.1646 (0.0353)***	0.1696 (0.0363)***
$\ln rp_{i,t-2}$		0.0148 (0.0243)
$\ln \pi_{i,t-1}$	−0.0293 (0.0152)*	−0.0471 (0.0201)**
$\ln \pi_{i,t-2}$		−0.0161 (0.0142)
t	−0.0044 (0.0008)***	−0.0044 (0.0009)***
$i3_i$	0.0457 (0.0560)	0.0572 (0.0598)
$c6_i$	−0.3010 (0.0859)***	−0.3226 (0.0942)***
$i6_i$	−0.1805 (0.0606)***	−0.1904 (0.0666)***
constant	−0.4294 (0.0536)***	−0.4667 (0.0684)***
Specification tests (p-values)		
m_1	0.000	0.000
m_2	0.423	0.856
J	0.198	0.099
F 2^{nd} lag		0.447

Two-step system GMM estimates of equation (4). The dependent variable is $\ln p_{it}$. Number of observations is 8820 (180 individuals, 49 periods). Standard error of the estimates between parentheses. Significant estimates at the 10%, 5%, and 1% significance level are marked with *, **, and ***, respectively. Three lags of $\ln rp_{it}$, $\ln p_{i,t-1}$, $\ln \pi_{i,t-1}$ used as "stacked" instruments (see Arellano, 2003, section 8.7) in the differences equation. For the levels equation the instruments are a constant, t, $i3_i$, $c6_i$, and $i6_i$. Total number of instruments: 14. The specification tests m_1, and m_2 are the Arellano and Bond (1991) tests for autocorrelation of order 1 and 2. If ϵ_{it} are not serially correlated, m_1 should reject the null, but m_2 should not reject. J is the Hansen's test of overidentifying restrictions. The F test in the last row test the joint significance of the second lag of the regressors.

Table 4: Monopoly pricing regressions

	GMM (1 lag)	GMM (2 lags)
$\ln mp_{i,t-1}$	0.0726 (0.0351)**	0.0495 (0.0529)
$\ln mp_{i,t-2}$		-0.0117 (0.0381)
$\ln rp_{i,t-1}$	-0.0326 (0.0173)*	-0.0385 (0.0182)**
$\ln rp_{i,t-2}$		-0.0204 (0.0126)
$\ln \pi_{i,t-1}$	-0.0251 (0.0086)***	-0.0313 (0.0120)***
$\ln \pi_{i,t-2}$		-0.0062 (0.0080)
t	0.0014 (0.0004)***	0.0013 (0.0005)***
$i3_i$	-0.0176 (0.0499)	-0.0266 (0.0513)
$c6_i$	0.0347 (0.0556)	0.0389 (0.0569)
$i6_i$	-0.0217 (0.0469)	-0.0243 (0.0477)
constant	0.0788 (0.0411)*	0.0748 (0.0427)*
Specification tests (p-values)		
m_1	0.000	0.000
m_2	0.707	0.674
J	0.112	0.160
F 2$^{\text{nd}}$ lag		0.135

Two-step system GMM estimates of equation (5). The dependent variable is $\ln mp_{it}$. Number of observations is 8820 (180 individuals, 49 periods). Standard error of the estimates between parentheses. Significant estimates at the 10%, 5%, and 1% significance level are marked with *, **, and ***, respectively. Three lags of $\ln rp_{it}$, $\ln mp_{i,t-1}$, $\ln \pi_{i,t-1}$ used as "stacked" instruments (see Arellano, 2003, section 8.7) in the differences equation. For the levels equation the instruments are a constant, t, $i3_i$, $c6_i$, and $i6_i$. Total number of instruments: 14. The specification tests m_1, and m_2 are the Arellano and Bond (1991) tests for autocorrelation of order 1 and 2. If ω_{it} are not serially correlated, m_1 should reject the null, but m_2 should not reject. J is the Hansen's test of overidentifying restrictions. The F test in the last row test the joint significance of the second lag of the regressors.

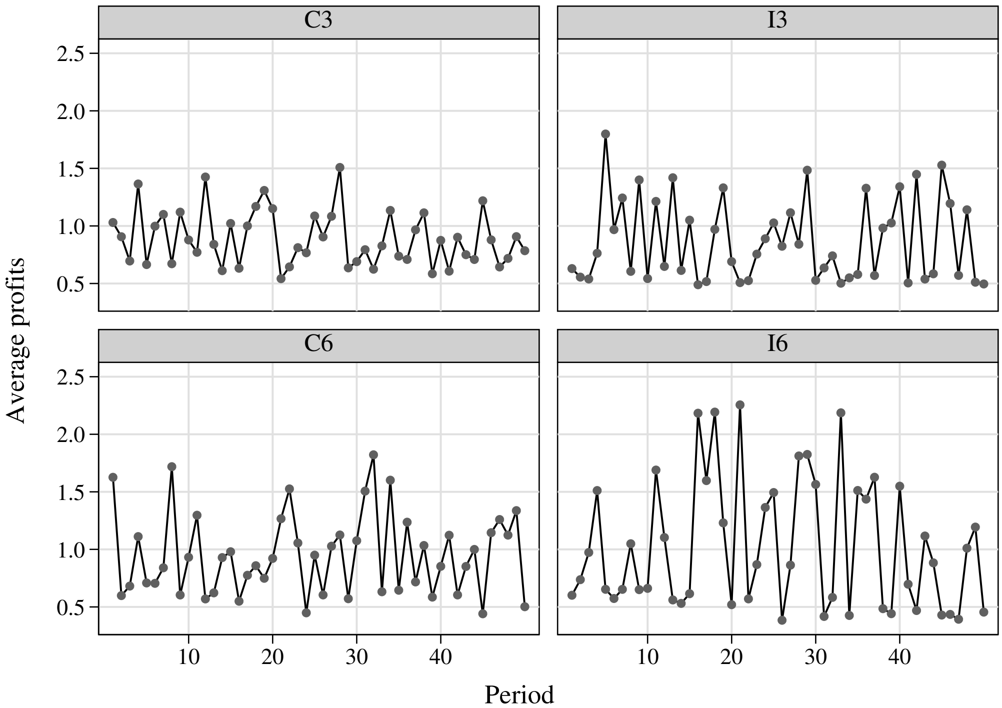

Figure 4: Evolution of average profit

the index is compatible with the observed intersections between the cumulative distributions of treatment pairs C3-I3 and C6-I6. Although this observation can be interpreted as a moderate confirmation of the Varian (1980) and Burdett and Judd (1983) models' predictions, hypotheses H1 and H2 are rejected through a number of different tests. Among them, the estimates of treatment dummy coefficients obtained from the model of individual prices in Table 3 indicate that the ranking obtained for mean prices corresponding to treatments with different firm numbers and index sizes is the contrary to that predicted by the theory. That is, a decrease in the size of the index of prices sampled by the search engine and a decrease in the total number of firms in the industry would both lead to higher average prices, providing evidence for the collusive effect conjectured by Janssen and Moraga-González (2004).

We move now to a more detailed discussion of the observed behavior regarding the main issue addressed in this paper, which is the frequency of monopoly pricing and its evolution. Figure 6 contains the histograms of price distributions obtained under each treatment. From simple inspection of the graphs, we see that, in all cases, price distributions are bimodal as predicted by the theory. One mode corresponds to *monopoly* pricing, whereas a second mode corresponds to *interior* prices. The frequency of monopoly pricing in treatments C3, I3 and I6 (17%, 14% and 13% respectively) is much larger than the model predicts (1.83%, 0.05% and 6.3% respectively). Monopoly pricing obtained under treatment C6 is approximately the same as

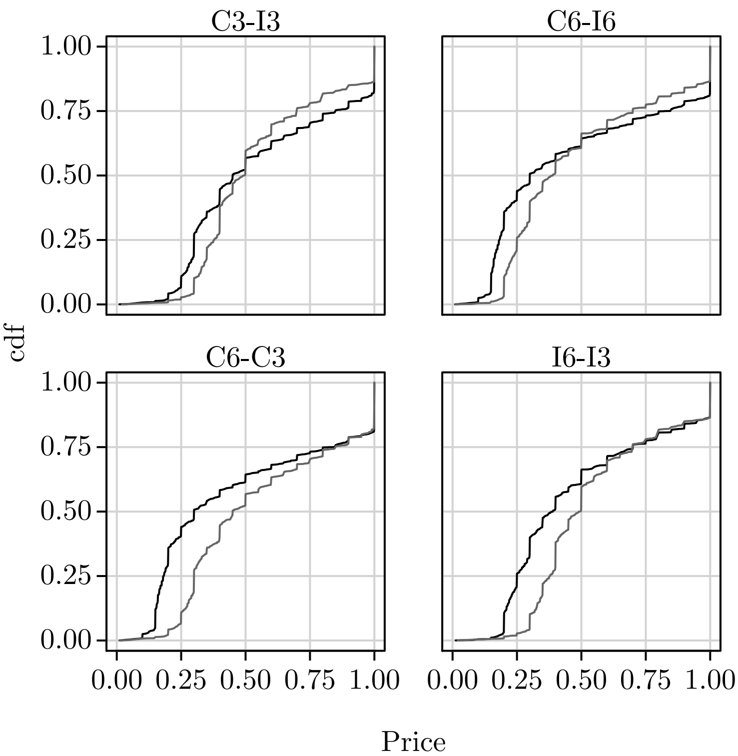

Figure 5: Observed price distributions
Comparison of empirical cumulative distributions of prices among pairs of treatments. In each panel, the darker lines corresponds to the treatment with a larger size of the index.

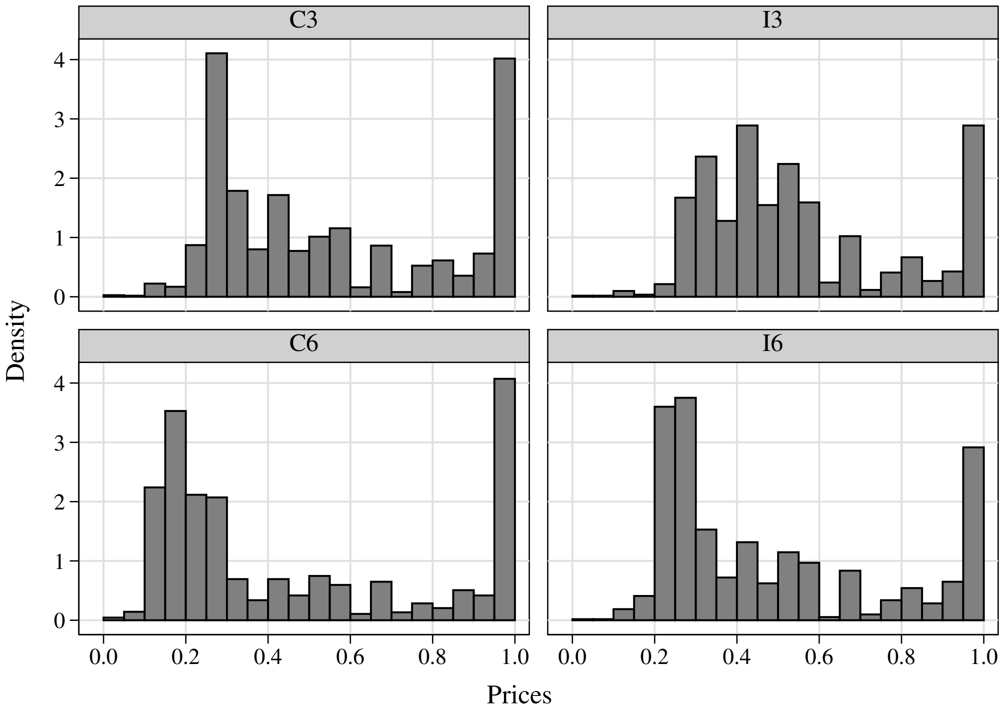

Figure 6: Histograms of price distributions per treatment

predicted by the theoretical model (17.56% against the theoretical frequency of 18%). Complete coverage treatments, C3 and C6, yield higher frequencies of monopoly pricing (17% and 18% respectively) than the corresponding incomplete coverage treatments, which is compatible with H4. On the contrary, changes in the total number of firms (from C3 to C6 and from I3 to I6) produces insignificant effects on the frequency of monopoly pricing, which contradicts H3. However a more rigorous test of H3 and H4 can be performed looking at the significance of treatment dummies in the Probit model of monopoly pricing. The estimates presented in s Table 4 reject both H3 and H4 against the alternative that the frequency of monopoly pricing does not vary with the total number of firms in the market and the number of firms indexed by the search engine.

Therefore, we see that monopoly pricing is a rather systematic pattern of behavior which does not depend on the parameters of the two models which have inspired our design. Following the last hypothesis, H5, we want to investigate whether individual pricing strategies are compatible with mixed strategy play, or alternatively, whether a firm's price in one period is correlated with its price in the previous period. In the framework of this hypothesis, we pay special attention to monopoly pricing, asking whether the persistence of this pattern across different design parameters implies that it is also invariant over time. Figure 2 shows period average prices. The general pattern is that average prices steadily decrease in the first twenty periods and then tend to stabilize. The initial decrease is more pronounced in treatments C6 and I6. Indeed in treatment C6, average prices do not get stable before period 30. In fact, it is not clear whether prices in this treatment stabilize at all. Figure 2 also shows the price series excluding monopoly prices ($p = 1$). In all treatments, the evolution of interior prices is characterized by a pronounced decreasing trend. On Figure 3, we observe some interesting patterns with respect to the frequency with which subjects set $p = 1$. In general, the frequency of monopoly pricing steadily increases in treatments C3 and I3. There are also increases in treatments C6 and I6, but the pattern is less clear. So, generally speaking, as the experiment evolves people adopt monopoly prices more frequently and, at the same time, set lower interior prices.

The estimates of the linear probability model for monopoly pricing, equation (5), presented in Table 4 reveal that there is an increasing trend in the frequency with which monopoly pricing is adopted along the 50 periods of a session. At the same time, the estimates of the individual pricing model (4) presented in Table 3 show that, overall, prices present a declining trend. This, together with the increasing trend in the frequency of monopoly pricing implies that interior prices present a strongly decreasing tendency. Also, given the significance of the lagged own strategy coefficients, both the overall and the monopoly pricing models reject the hypothesis of independence of individual strategies over time, which is a rejection of H5.

The overall individual pricing and the monopoly pricing models allow us to further investigate the underlying behavioral dynamics which can explain the aforementioned trends. First, the overall pricing model confirms that individual prices are compatible with Cournot best responses, given that a firm's price in a given period positively relates to the average price of rival firms in the previous period. The reaction to rival prices in periods before the last is rejected by the corresponding model with two lags. Also, a period's price is lower following a period of low earnings, while further lags of a firm's own profit were also found to be nonsignificant. Intuitively, this finding implies that a firm which has failed to be the most competitive of those included in the index is more likely to abandon competitive pricing in favor of a higher price in the following period. In fact, in that case, a rational player should then revert to monopoly pricing. The estimated monopoly pricing model confirms exactly this conjecture. The lower the average of rival prices are, the more likely it is that a firm adopts monopoly pricing in the next period.

Summarizing the findings obtained from the estimation of the two individual pricing models, we can describe a situation which is perfectly compatible with the price cycles described in several theoretical, empirical and experimental studies reviewed above. A firm engages in competitive pricing trying to be the one with the lowest price among those sampled by the search engine.

In that effort, their reaction to rival past prices resembles Cournot best responses, because a fall (increase) in rival firms' prices is followed by a fall (increase) in the firm's own price in the following period. Firms will more frequently fail to be the one with the lowest price, thus experiencing a sharp fall in their profits. In that case, firms will abandon competitive (interior) prices in order to adopt monopoly pricing. However, such cycles coexist with an increasing trend in the frequency with which monopoly pricing is adopted and a decreasing trend of interior prices. Another important novelty of our analysis is that the aforementioned dynamics concern individual behavior obtained in the absence of any noise from rival firms' learning dynamics.

5 Conclusions

The rapid increase in the volume of Internet-based consumer search processes has led to a re-emergence of economists' interest in the functioning of markets with informed and uninformed consumers.

In this paper we have presented experimental results which are relevant for the evolution of prices in Internet-assisted search markets. Our design implements conditions which are inspired by the comparative statics of the models by Varian (1980) and Burdett and Judd (1983). The data offer limited support for the static predictions of the two models. However, our focus on the dynamics underlying the observed price distributions reveals some interesting features which, to our knowledge, have not been studied so far.

First, we observe that price distributions are bimodal. One mode corresponds to monopoly pricing and the other to interior dispersed equilibria. While overall prices tend to react to the treatments implemented here, the frequency of monopoly pricing remains invariant across treatments. We show that interior prices exhibit a decreasing trend, whereas the frequency with which monopoly pricing is adopted increases over time. This finding can be considered as a confirmation of the instability prediction by Hopkins and Seymour (2002) for this setting. Firms set higher prices in response to higher rival prices in the past and are more likely to adopt monopoly pricing following a period of low profits. On the contrary, they are more likely to undercut, setting interior prices in response to rivals adopting monopoly pricing in the previous period. These patterns are compatible with Cournot best responses and price cycles caused by firms' switching from interior to monopoly pricing and viceversa.

We cannot propose a unique behavioral explanation for the systematic and excessive frequency with which monopoly pricing is adopted. Nevertheless, it seems that the conjecture by Janssen and Moraga-González (2004) is only confirmed by interior prices, because monopoly pricing is not affected by the number of firms in the market and the number of those sampled by the search engine. Therefore, the adoption of monopoly pricing can be explained as the result of risk averse behavior aimed at guaranteing the safe monopoly profit against the potentially more profitable but riskier alternative of competitive pricing. Alternatively, according to Baye and Morgan (2004), excessive monopoly pricing may be a symptom of bounded rationality.

Our results go beyond testing the stability of price dispersions or a set of specific predictions based on comparative statics. According to the main finding of our experiment, in search markets like the ones assisted by Internet-based price comparison search engines, some prices should be expected to decrease over time, while, at the same time, the frequency of monopoly pricing should be expected to increase. Hopefully, this clear cut prediction will inspire future empirical research on the divergent dynamics affecting the two attractors of pricing behavior identified here.

References

Arellano, M. (2003): *Panel Data Econometrics.* Oxford University Press.

Arellano, M., and S. Bond (1991): "Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Equations." *Review of Economic Studies*, 58, 277–297.

Arellano, M., and O. Bover (1995): "Another Look at the Instrumental Variables Estimation of Error-Components Models." *Journal of Econometrics*, 88, 29–51.

Baye, M., D. Kovenock, and C. G. de Vries (1992): "It Takes Two to Tango: Equilibria in a Model of Sales." *Games and Economic Behavior*, 4, 493–510.

Baye, M. and J. Morgan (2001): "Information Gatekeepers on the Internet and the Competitiveness of Homogeneous Product Markets." *American Economic Review*, 91, 3, 454–474.

Baye, M. and J. Morgan (2004): "Price Dispersion in the Lab and on the Internet: Theory and Evidence." *The RAND Journal of Economics*, 35, 448–466.

Benaïm, M., J. Hofbauer, and E. Hopkins (2009): "Learning in Games with Unstable Equilibria." *Journal of Economic Theory*, 144, 1694–1709.

Blundell, R., and S. Bond (1998): "Initial Conditions and Moment Restrictions in Dynamic Panel Data Models." *Journal of Econometrics*, 87, 115–143.

Brynjolfsson, E. and M. Smith (2000): "Frictionless Commerce? A Comparison of Internet and Conventional Retailers." *Management Science*, 46, 4, 563–585.

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

Busse, M. (2002): "Firm Financial Condition and Airline Price Wars." *The RAND Journal of Economics*, 33, 298–318.

Cameron, A. C. and P. K. Trivedi (2005): *Microeconometrics. Methods and Applications.* Cambridge University Press.

Cason, T. and S. Datta (2004): "An Experimental Study of Price Dispersion in an Optimal Search Model with Advertising." *International Journal of Industrial Organization*, 24, 639–665.

Cason, T. N. and D. Friedman (2003): "Buyer Search and Price Dispersion: a Laboratory Study." *Journal of Economic Theory*, 112, 232–260.

Cason, T. N., D. Friedman, and F. Wagener (2005): "The Dynamics of Price Dispersion or Edgeworth Variations." *Journal of Economic Dynamics and Control*, 29, 801–822.

Diamond, P. (1971): "A Model of Price Adjustment." *Review of Economic Theory*, 3, 156–178.

García-Gallego, A., N. Georgantzís, P. Pereira, and J. C. Pernías (2004): "Risk Attitudes and Internet Search Engines: Theory and Experimental Evidence." Working Paper 04-03, NET Institute.

Guillén, P. (2004): "Price-Quantity Competition and Edgeworth Cycles." Mimeo.

Hofbauer, J. and E. Hopkins (2005): "Learning in Perturbed Asymmetric Games." *Games and Economic Behavior*, 52, 133–152.

Hopkins, E. and R. Seymour (2002): "The Stability of Price Dispersion under Seller and Consumer Learning." *International Economic Review*, 43, 1157–1190.

Iyer, G. and A. Pazgal (2000): "Internet Shopping Agents: Virtual Co-Location and Competition." *Marketing Science*, 22, 85–106.

Janssen, M. C. W. and J. L. Moraga-González (2004): "Strategic Pricing, Consumer Search and the Number of Firms." *Review of Economic Studies*, 71, 1089–1118.

Kocas, C. and T. Kiyak (2006): "Theory and Evidence on Pricing by Asymmetric Oligopolies." *International Journal of Industrial Organization*, 24, 83–105.

Kruse, J. B, S. Rassenti, S. R. Reynolds, and V. L. Smith (1994): "Bertrand-Edgeworth Competition in Experimental Markets." *Econometrica*, 62, 343–371.

Morgan, J., H. Orzen, and M. Sefton (2006): "An Experimental Study of Price Dispersion." *Games and Economic Behavior*, 54, 134–158.

Orzen, H. (2005): "Counterintuitive Number Effects in Experimental Oligopolies." *Experimental Economics*, 11, 390-401.

Rosenthal, R. (1980): "A Model in Which an Increase in the Number of Sellers Leads to a Higher Price." *Econometrica*, 48, 1575–1579.

Ross, L. (1997): "When Will an Airline Stand its Ground? An Analysis of Fare Wars." *International Journal of Economics and Business*, 4, 109–127.

Shapley, L. (1964): "Some Topics in Two Person Games." In Drescher (ed.): *Advances in Game Theory*. Princeton: Princeton University Press.

Stahl, D. (1989): "Oligopolistic Pricing with Sequential Consumer Search." *American Economic Review*, 79, 700–712.

Varian, H. (1980): "A Model of Sales." *American Economic Review*, 70, 651–659.

Warnick, J. E. and E. Hopkins (2006): "A Simple Test of Learning Theory." University of Edinburgh, ESE Discussion Paper 153.

Appendix

Instructions (translated from Spanish)

- The purpose of this experiment is to study human behavior in specific economic contexts. The session will last (approximately) an hour. This experiment is part of a research project which has received support by public and private institutions. Your decisions during this session are going to be of great importance for the success of our research. Follow the instructions carefully. At the end of the session you will receive an amount of money in cash which will depend on your performance during the session. Next, we describe the economic scenario in which the decision making will take place.

- The environment is an industry characterized by:

(a) A number of firms (3 or 6) which will remain constant through the session. Each firm produces a product which is the same for all firms,

(b) A price comparison search engine like the ones on the Internet,

(c) 1200 consumers.

- You are *one* of the firms in the industry. *At the beginning of the session, the exact number of firms in the industry will appear on your screen.*
- The session will consist of *50 rounds*. In each round, you have to decide *the price* at which you want to sell the product. Price is your only decision variable.
- For simplicity, your production costs are zero, so that, each period, your profits will be equal to your revenue (PRICE × number of consumers you sell to).
- Each period, the *Price Search Engine* lists the price of *all* (*Treatments I3 and I6*:several but *not all*) firms in the industry. (*Treatments I3 and I6*:*The exact number of firms whose prices will be indexed by the search engine will appear on your screen at the beginning of the session.*)
- *Treatments I3 and I6*): In particular, the search engine chooses randomly the firms whose prices are going to be indexed in each period.
- Transactions will take place in *UMEX*, our lab's *Experimental Monetary Unit.*
- Each consumer buys one unit of the product per round. His maximum willingness to pay for a unit of the product is 1000 UMEX. That is, if the price you fix is higher than 1000 UMEX, the demand for your product will be zero.
- There are two types of consumers. Half of them (600) are consumers who read the list of prices created by the price search engine. The other half do not actually read the list of prices of the search engine (maybe because they are not able to do so). The behavior of consumers is simulated by the computer.
- The consumers who read the price list of the search engine will buy from the firm whose price for that period is the lowest *among all prices included in the price list*, if such price does not exceed 1000 UMEX. In case of a "tie" (i.e. several firms fix the same minimum price) the consumers are distributed evenly among the firms with the same minimum price.
- The consumers who do not read the search engine's price list will buy randomly from any vendor, so that this group of consumers will be distributed evenly among all firms in the industry.

- In each round, the computer randomly assigns the other firms you compete with among computer-simulated agents who are programmed to pursue identical objectives to those induced to you by these instructions. Therefore, the probability of competing with a firm which adopts the same price over several periods is very low.

Feedback

- After each round's prices are posted you will receive feedback on the outcome of the period, like (*Treatments I3 and I6:* whether you were part of the search engine,) your own sales and your earnings for that round and the prices fixed by the rest of the firms in your market. (*Treatments I3 and I6:* The prices of firms indexed by the search engine, will appear in bold characters.)
- During the session, you will be able to see the complete history of the aforementioned feedback on your screen.

Payoffs

- At the end of the session you will be paid in cash. Your reward will be determined taking into account the earnings you accumulate over 15 (randomly selected) periods. The exchange rate will be: $1,000,000$ UMEX = 12 € (if you are part of a market with 3 firms), or $1,000,000$ UMEX = 24 € (if you are part of a market with 6 firms).

Thank you very much for your collaboration. Good luck!