Bertrand competition with convex costs in symmetric and asymmetric markets: Results from a pilot study*

Cédric Argenton†
CentER & TILEC, Tilburg University

Wieland Müller‡
CentER & TILEC, Tilburg University

August 13, 2008

Abstract

We report the results of a series of experimental Bertrand duopolies where firms have convex costs. These duopolies are theoretically characterized by a multiplicity of Nash equilibria. Using a 2×2 design, we analyze price choices in symmetric and asymmetric markets under two information conditions (complete versus incomplete information about profits). We find that information has no effect in symmetric markets with respect to market prices and the time it takes for markets to stabilize. However, in asymmetric markets, complete information leads to higher market prices and quicker convergence of price choices.

JEL Classification numbers: L13, C72, C92.

Keywords: Bertrand competition, convex costs, collusion, coordination, experimental economics.

*We thank Bert Willems for helpful comments. Wieland Müller acknowledges financial support from the Netherlands Organisation for Scientific Research (NWO) through a VIDI grant.

†Department of Economics, Tilburg University, Warandelaan 2, Postbus 90153, 5000 LE Tilburg, The Netherlands. E-mail: *c.argenton@uvt.nl.*

‡Department of Economics, Tilburg University, Warandelaan 2, Postbus 90153, 5000 LE Tilburg, The Netherlands. E-mail: *w.mueller@uvt.nl.*

1 Introduction

Contrary to a widespread belief, Bertrand competition in a homogeneous market is not synonymous with perfect competition. As a matter of theory, the equilibrium price equals marginal cost in a one-shot Bertrand game only if returns to scale are constant and at least two firms share the efficient technology. When there are diseconomies of scale, Dastidar (1995) showed that there is a whole interval of pure-strategy Nash equilibrium prices. This begs the question as to which equilibrium can be predicted to emerge in such markets. The interval of Nash equilibria contains the competitive price, which involves marginal-cost pricing, and, often, the (Cournot) price that would emerge if firms competed in quantities under the same demand and costs conditions. Moreover, the lowest price in the set of Nash equilibria is the one that involves average-cost pricing (such that unilateral undercutting and serving the entire market involves sure losses). The Nash equilibria are Pareto-ranked with the highest Nash equilibrium being payoff-dominant as well as risk-dominant (in the sense of Harsanyi and Selten (1988)). Thus, this market game resembles a coordination game and firms might experience severe difficulties with coordinating, especially when their number is large. Hence, it appears hard to make a reasonably reliable prediction as to which price might prevail in such markets.

In this paper we consider *duopoly* markets in which firms repeatedly decide about prices in fixed pairs. Earlier experimental evidence on repeated market games with fixed pairs (e.g. Fouraker and Siegel (1963) or Huck, Müller and Normann (2001)) reports a high level of collusion, even in finitely-repeated game environments. Hence, Nash equilibria might be less relevant as theoretical benchmarks in this context. The focus may thus be shifted to the question as to whether or not duopolists are able to coordinate on perfectly collusive prices which, for the markets considered in this paper, are above the highest Nash equilibrium price. We analyze the question as to whether duopolists are able to coordinate on the joint-payoff-maximizing price under varying cost and information conditions. More precisely, we analyze price choices in symmetric and asymmetric duopoly markets under two information conditions: one in which firms have complete information about all payoff functions and one in which the competitor's payoff function is not known. The specific research questions we want to address in this paper are as follows. First, what is the effect of cost and information conditions on the ability of duopoly markets to reach collusive prices? If markets do not reach collusive prices, to which equilibrium price do they converge (if at all)? Second, what is the effect of cost and information conditions on the speed at which duopoly pricing stabilizes? Note that from a welfare perspective in the short run there is a tension between (early) coordination (which is welfare-enhancing under convex costs as production is spread over the two firms) and the ability of duopolies to charge collusive prices (which is obviously welfare-reducing).

A few comments on the motivation for our experiment are in order. First, Bertrand competition with convex costs is an underresearched topic both theoretically and experimentally. (With regard to theory, see the discussion in Weibull (2006), and with regard to experimental work, see the literature overview in the next section.) Second, although underresearched, Bertrand competition with convex costs seems a very relevant topic from a practical point of view as several industries, especially utilities such as gas and electricity provision, are characterized by convex costs (see, e.g., Wolfram (1999)) and the obligation to serve all demand addressed to them at the posted price. Third, considering the leading example of utilities, market structure is duopolistic in many countries. For instance, for five years after the privatization of the energy market in the UK in April 1990, the market for electricity generation was basically a duopoly consisting of the firms National Power and PowerGen (Wolfram (1998)). In Sweden, there are also two major firms competing throughout the country (Vattenfall and Sydkraft). Fourth, the effect of cost and information asymmetry on the behavior of markets are relevant from an antitrust policy point of view. Among the classical questions asked in antitrust economics are: Should symmetric or asymmetric markets be favored (as part of, e.g., merger control)? Should information exchange between firms about their cost structure be allowed? (See Kühn and Vives (1995) and Vives (2006) for overviews.)

To briefly summarize our main results, we find that information about payoffs has no effect in symmetric markets with respect to market prices and the time it takes until markets stabilize. However, in asymmetric markets, complete information leads to higher market prices and quicker convergence of price choices.

The paper is organized as follows. Section 2 presents a short overview of experimental results on Bertrand competition in homogenous markets. Section 3 describes our experimental design and the procedures followed. Section 4 provides the theoretical background and predictions. In section 5 we report the experimental results, concentrating on the research questions outlined above. Finally, in section 6 we discuss our results. In particular, we provide links to the experimental literature on coordination games, discuss the relevance of our results with respect to competition policy, and outline plans for future research.

2 Experimental literature on Bertrand competition

Despite its central place in fields such as industrial organization, there is only a small experimental literature dealing with Bertrand competition in homogeneous markets. Most of these papers explore price setting by firms with constant marginal costs. There is only one previous experimental study exploring price-setting behavior of firms with increasing marginal costs. We quickly review these papers.

2.1 Constant marginal costs

Next to experiments with quantity-setting firms, Fouraker and Siegel (1963) report on Bertrand markets with either two or three symmetric firms. As we do in this paper, they also consider markets in which firms have either complete or incomplete information. Regarding the two dimensions of their design, Fouraker and Siegel (1963, p. 199) summarize their results as follows: "As the number of bargainers in the oligopolistic structure increases, oligopolies under complete information show a stronger tendency, and faster approach, to the Bertrand price." Furthermore: "As the amount of relevant information available to the bargainers increases, duopolies decrease their tendency to the Bertrand competitive price."

Dufwenberg and Gneezy (2000) analyze the effect of the number of competitors in a market with constant marginal cost and inelastic demand. In their experiments, groups of two, three, or four subjects (out of a pool of 12 subjects) were randomly matched to play the Bertrand game. All bids and subject identities were made publicly known after each period. While market prices clearly stay above the competitive level in duopoly, they quickly converge to the competitive level when there are three or four firms in the market.

In Dufwenberg and Gneezy (2002) the authors use the same setup but consider only duopoly markets while varying the feedback. Subjects were informed either about all bids, all winning bids, or no bids (of the 12 subjects in the pool) at the end of each period.[1] The authors find that, as before, market prices stay consistently above the competitive level when all bids are announced. However, in both of the other treatments, market prices converge quickly to the competitive price level. The authors conclude that in the treatment where all bids are revealed, players can signal their willingness to cooperate more easily by choosing very high prices.

[1] The treatments where all bids are announced at the end of a period is therefore the same as the one used in Dufwenberg and Gneezy (2000).

Boone, Müller, and Ray Chaudhuri (2008) analyze Bertrand markets where firms have constant but different marginal costs. They consider three market conditions in a within-subject design. In condition 1, there are two firms that have unit costs of either 10 or 20. In condition 2, there are three firms with unit costs of 10, 10, and 30, respectively. In condition 3, there are again three firms that have unit costs of 10, 20, and 30, respectively. In a random matching scheme, players are informed about the prices chosen in their own market at the end of each period. The authors find that while market prices converge to the predicted price level of 19 in conditions 1 and 3, market prices stay above the predicted price of 10 in condition 2.[2]

2.2 Increasing marginal costs

Abbink and Brandts (2008) is the first and, as far as we know, the only study reporting on symmetric Bertrand competition experiments where firms have convex costs. They ran experiments with fixed pairs of two, three, and four symmetric firms that have complete information about competitors' payoff functions. Their results indicate that duopolists are often able to collude on the joint profit-maximizing price. However, with more than two firms in a market, the predominant market price is the lowest price in the range of Nash equilibria which involves no loss in case of miscoordination, a much smaller number than the collusive price. Abbink and Brandts (2008) also develop a learning model based on imitation which predicts long-run convergence toward the competitive outcome and report further experimental evidence in its support.

3 Experimental design

Subjects in our experimental design repeatedly made price choices out of the set $\{10, 11, ..., 50\}$. The design aimed at reproducing the conditions of the model of Bertrand pricing under convex costs, in which buyers passively buy from the firm(s) offering the lowest price while sellers behave strategically. Although the experiment was described to the subjects as a pricing game between firms, they were not given the details of the model. Instead, they were presented with a payoff table describing their own payoff as a function of the price they choose and the one chosen by the other firm in their market, along with another, similar table displaying this other firm's payoff in

[2] Note that condition 2 resembles the duopoly situation considered in Dufwenberg and Gneezy (2000), although the feedback was more limited. The results are in line with what was found in Dufwenberg and Gneezy (2000) and in the duopoly market with feedback of all prices in Dufwenberg and Gneezy (2002).

some treatments. The use of payoff tables in experiments is widespread and old, starting as early as Fouraker and Siegel (1963).

Subjects were paired in one of four treatments. Treatments varied with respect to the difference, if any, in payoffs and the information that subjects were given about the payoffs to the other firm in their market. In treatments we call "symmetric," the payoff tables of the two paired subjects were identical. In treatments we call "asymmetric", they were different. In some instances, subjects were given their payoff table along with the one of their rival. Those treatments we describe as "complete information." In some other instances, subjects were given only their own payoff table and were told that the one of their rival might or might not be identical. Those treatments we describe as "incomplete information." Crossing the two criteria, we obtained the following treatments: symmetric costs and complete information (SymC), symmetric costs and incomplete information (SymI), asymmetric costs and complete information (AsymC) and asymmetric costs and incomplete information (AsymI).

The payoffs were generated from a linear demand curve and quadratic cost curves. The subject posting the lowest price was assumed to serve all the demand addressed to him or her at this price. In case both subjects chose the same price, sales were split equally. In symmetric treatments, both payoffs were computed from the demand curve $Q = 100 - 1.5P$ and individual cost functions: $c(q) = 0.6q^2$. In asymmetric treatments, one of the two paired subjects was endowed with a high cost parameter ($c(q) = 0.65q^2$) while the other was endowed with a low cost parameter ($c(q) = 0.55q^2$). To produce the payoff tables, all numbers were rounded. In addition, in the asymmetric treatments, some numbers were substituted with a neighboring integer in order to produce unique profit-maximizers in all cases. This linear-quadratic specification was chosen not only because it made computations easy but also because *(i)* it provided for a clear separation between the set of Nash equilibria and the best collusive price(s) and *(ii)* it provided for a direct comparison of consumer, producer and total surpluses across treatments (since, conditional on both firms charging the same price, consumer surplus and producer surplus are the same in symmetric and asymmetric treatments). In all treatments, payoffs were presented in a fictitious monetary unit, called "point." Subjects were told that negative numbers stood for losses, which were indeed possible in the low range of prices.

All subjects were electronically recruited from the pool of participants registered with Tilburg University's CentERlab. At the time of the experiment, they were all students enrolled in various programmes of the university. They reported to the experimental laboratory, where they were assigned to a computer workstation and given a set of instructions and payoff table(s).[3] The section of the instructions explaining the way to read the payoff table(s) was read aloud as they followed along on their own copy. Questions were taken and answered, after which the experiment started.

The experiment consisted of 40 decision rounds. Subjects were randomly matched with an anonymous counterpart and interacted with him or her in all 40 rounds. Subjects were made aware of this feature. In each round, each subject had to make only one decision, namely to set the price at which he or she was willing to sell the fictitious product of the firm he or she represented. All integers from 10 to 50 were possible choices. After each round, each subject was presented with a summary screen displaying the price chosen by this subject, the price chosen by his or her rival as well as his own payoff. The rival's payoff was *not* displayed (although it could have been recovered from the payoff tables in the complete information treatments) in order no to foster imitation.

In all treatments, subjects started the experiment with an initial capital of 5,000 points to cover possible losses. At the end of the experiment, their monetary earnings were determined by the sum of this capital and the profits (or losses) in all rounds. One euro was exchanged for every 1,800 points accumulated. Subjects were not given the details of the procedure to be followed in case their accumulated earnings turned negative during the course of the experiment. As expected, this event did not occur. Each treatment lasted between 30 and 45 minutes. Subjects had been recruited for a length of time 45 to 60 minutes longer. The average monetary earnings across all treatments were 12.75 Euros.

We analyze data from 4 experimental sessions, one for each treatment. We have data on 8 pairs in the SymC treatment; 9 pairs in the SymI treatment; 6 pairs in the AsymC treatment; and 9 pairs in the AsymI treatment. Table 1 summarizes the design.

4 Theoretical predictions

Bertrand competition is not synonymous with perfect competition when the technology exhibits diseconomies of scale (i.e. when firms face convex cost functions). Dastidar (1995) proved that in the symmetric case (where all firms have the same cost function) there is a whole interval of pure-strategy Nash equilibrium prices.[4] The lower bound of this interval is determined by average-cost

[3] The instructions for all treatments are available at:
http://center.uvt.nl/staff/muller/InstructionsBetrandWithConvexCosts.pdf.

[4] There are also continua of non-zero profit mixed-strategy equilibria, as demonstrated by Hoernig (2002).

	Complete Information	Incomplete Information
Symmetric Costs	SymC $c_1 = c_2 = 0.6$ $(8 \times 2 = 16)$	SymI $c_1 = c_2 = 0.6$ $(9 \times 2 = 18)$
Asymmetric Costs	AsymC $c_1 = 0.55, \quad c_2 = 0.65$ $(6 \times 2 = 12)$	AsymI $c_1 = 0.55, \quad c_2 = 0.65$ $(9 \times 2 = 18)$

Table 1: The 2 by 2 factorial design of cost and information conditions and the numbers of subjects participating in the four treatments.

pricing. The upper bound is determined by the incentives to undercut competitors. The interval contains the competitive price, which involves marginal-cost pricing. It may contain the price that maximize joint profits or not, but in the linear-quadratic specification we implement, it doesn't. In the asymmetric case (where firms have different cost functions), a pure-strategy Nash equilibrium always exists. It may be unique or non unique. In the linear-quadratic specification we implement, it is still the case that there is a continuum of equilibria.

As a practical matter, equilibria are found by checking whether at a given candidate price *(i)* both firms make non-negative profits and *(ii)* no firm wants to marginally undercut the other one so as to reap the benefits of serving the entire demand. In our specification, equilibria have several interesting features. Firstly, the lowest equilibrium is determined by a zero-profit condition. Because costs are convex, this means that a player who posts the corresponding price runs the risk of making a loss if it happens that the other player chooses a higher price. This is in fact true for all the prices at the bottom of the interval of Bertrand equilibria. Because of cost convexity, the potential losses from miscoordination keep on increasing with the size of demand so that low Bertrand prices are riskier. Secondly, from the point of view of firms, Nash equilibria are Pareto-ranked: the higher the equilibrium price, the higher the equilibrium profits. The ranking is reversed when one considers total surplus or consumer surplus. Thirdly, the price which maximizes players' joint profits lies outside the interval of Nash equilibria, so that there is room for collusion in a repeated-game environment. Fourthly, under cost asymmetry, players disagree about the best course of action under collusion. Naturally, a high-cost firm maximizes its own profit at a higher price than a low-cost firm. In our specification, it is the case that each firm's "preferred" price is separated from the other by two integers.[5]

Figure 1 reproduces the payoff table we used in symmetric treatments. As can be checked, all prices in $\{21, 22, ..., 39\}$ were Bertrand equilibria. The lowest Nash equilibrium price, 21, involved an equilibrium profit of 15 but a loss of 1377 in case of miscoordination. By contrast, the payoff-dominant equilibrium price, 39, involved an equilibrium profit of 551 and a gain of 585 in case of miscoordination. The lowest Bertrand equilibrium price involving no loss in case of miscoordination was 32.[6] The monopoly price was 49 but, due to decreasing returns to scale, the price maximizing joint profits (and thus an obvious candidate for tacit collusion) was 44.[7]

In our asymmetric treatments, the range of Bertrand equilibria ran from 22 to 38. The lowest equilibrium price involving no loss in case of miscoordination was 33. Conditional on having both firms charging the same price, the profit to the low-cost firm was maximized at a price of 43, while the profit to the high-cost firm was maximized at a price of 45.[8]

Thus, the duopoly setting we investigate is akin to a large, finitely-repeated coordination game with many Nash equilibria. There is no univocal theoretical prediction. Shared expectations and common knowledge of rationality can give rise to the play of any Nash equilibrium in a one-shot context. However, both payoff dominance and risk dominance in the sense of Harsanyi and Selten (1988) call for the highest Nash equilibrium price to be played. This disregards the fact that in our experimental setting the game is in fact repeatedly played by the same players, which can give rise to tacit collusion on higher prices, as well-known from Benoît and Krishna (1985).

5 Experimental results

We report the results of the experiment in two subsections. In a first subsection, we give a quick overview of the results and then test for differences in market prices, producer surplus and total welfare across treatments. In a second subsection, we analyze convergence patterns in the individual

[5] *Joint* profit maximization is achieved at a unique, well-defined price, which is the same in our symmetric and asymmetric treatments. The disagreement between firms in the asymmetric case arises from the impossibility of making side payments to compensate for the difference in costs. The payoff implications of coordinating on 43, 44 or 45 are arguably very small.

[6] This is the equivalent in our specification of the "near-magic" number 24 in Abbink and Brandts (2008).

[7] When prices and quantities are continuous variables, it is possible to solve for the competitive equilibrium price (about 31.6) or the Cournot equilibrium price (about 38.9).

[8] When prices and quantities are continuous variables, it is possible to solve for the competitive equilibrium price (about 31.5) or the Cournot equilibrium price (about 38.8).

Your price	Your profit when you have the lowest price	Your profit when you are tied for the lowest price	Your profit when you don't have the lowest price
10	-3485	-659	0
11	-3265	-587	0
12	-3050	-517	0
13	-2842	-449	0
14	-2639	-383	0
15	-2441	-320	0
16	-2250	-258	0
17	-2064	-199	0
18	-1883	-142	0
19	-1709	-88	0
20	-1540	-35	0
21*	-1377	15	0
22*	-1219	64	0
23*	-1068	110	0
24*	-922	154	0
25*	-781	195	0
26*	-647	235	0
27*	-518	272	0
28*	-394	307	0
29*	-277	340	0
30*	-165	371	0
31*	-59	400	0
32*	42	426	0
33*	136	451	0
34*	225	473	0
35*	309	493	0
36*	386	511	0
37*	458	526	0
38*	525	540	0
39*	585	551	0
40	640	560	0
41	689	567	0
42	733	572	0
43	770	574	0
44	802	575	0
45	829	573	0
46	849	569	0
47	864	563	0
48	874	554	0
49	877	544	0
50	875	531	0

Note: * Static Nash equilibrium, **Perfect collusion**

Table 2: Payoff table for symmetric treatments

Treatment	Mean of market averages	Standard error of the mean	N
SymC	38.56	1.31	8
SymI	39.25	1.32	9
AsymC	42.65	0.64	6
AsymI	38.32	1.10	9

Table 3: Descriptive statistics regarding market prices (all 40 periods considered)

markets. We here concentrate on two questions. *(i)* To which price(s) do markets converge? *(ii)* If they do, how long does it take for markets to stabilize?

5.1 Market prices and welfare

A summary of the experimental results is given in Table 2, which displays the mean of market prices and the standard error of the mean in the various treatments (taking all periods into account). The market price is defined as the minimum of the two prices posted by the firms. It is the price at which consumers would obtain the good on a market characterized by Bertrand competition. Observe that market prices seem noticeably higher on average (by about 4 points or 10%) and noticeably less dispersed in treatment AsymC. In contrast, the three other treatments look remarkably similar.

The evolution of the average market price in each treatment is shown in Figure 1. Inspecting this figure, we make a number of observations. First, in all treatments there is a noticeable increase of the average market price at the beginning of the experiment. Whereas average market prices increase initially during the first 10 to 15 periods in the treatments with incomplete information, this increase takes place quicker in the treatments with complete information. Second, after an initial increase, the average market price is rather stable in all treatments. This stabilization is most remarkable in treatment AsymC where the average market price hardly changes over time after the initial increase and stays at a level of about 43 until shortly before the end of the experiment. Third, in all treatments we observe a noticeable endgame effect with average prices sharply decreasing in the last two or three periods. Fourth, and most importantly, treatment AsymC clearly stands out: the average market price in this treatment is distinctively higher than in all other treatments. Fifth, the average market prices in treatments SymC, SymI, and AsymI are quite close to one another throughout the experiment (and perhaps start diverging only towards the end of the experiment).

Next, we turn to test for cross-treatment differences in market prices, producer surplus, and

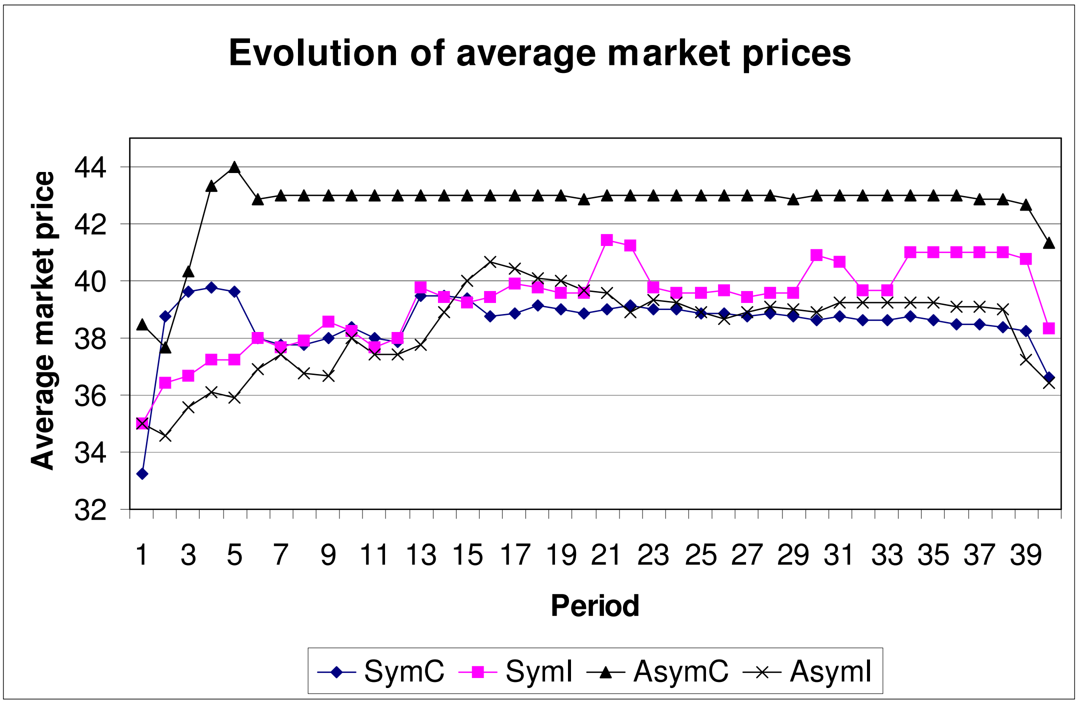

Figure 1: The evolution of prices over time.

total welfare more formally.[9] For the purpose of testing for price differences across treatments, we run the following GLS model

$$p_{jt} = \alpha_0 + \alpha_1 D_{SymI} + \alpha_2 D_{AsymC} + \alpha_3 D_{AsymI} + \varepsilon_{jt} \quad (1)$$

where p_{jt} is the market price in market j in round t and ε_{jt} is a market-specific error term. D_{SymC}, D_{AsymC}, and D_{AsymI} are treatment dummies equal to 1 if a market price stems from treatment SymI, AsymC, AsymI, respectively, and equal to 0 otherwise. Thus, treatment SymC serves as the reference group, and the treatment dummies measure the effect of the various treatments relative to this benchmark. Similar regressions were run to test for across-treatment differences in producer and total surpluses. The Arellano and Bond (1991) test for autocorrelation indicates serial autocorrelation, as is expected in any learning process in which players act upon the feedback they receive. Hence, in all regressions we corrected for serial autocorrelation described by an AR(2) process. (For a similar approach see e.g. Mason and Phillips (1992, 1997)). Furthermore, we used the Huber-White/sandwich estimator of variance. We ran regressions separately for the first half, the second half as well as for all rounds.[10]

[9] Note that differences in market prices translate directly into differences in consumer welfare.

[10] Due to the clear endgame effect we excluded the last three rounds. We divided the remaining periods in two,

	Market Price			Producer Surplus			Total Surplus		
	(MP1)	(MP2)	(MP3)	(PS1)	(PS2)	(PS3)	(TS1)	(TS2)	(TS3)
Rounds →	1-17	18-37	1-37	1-17	18-37	1-37	1-17	18-37	1-37
Constant	37.68**	38.81**	37.81**	678.23**	990.37**	831.02**	1,317.63**	1,588.35**	1,454.90**
	(33.15)	(24.76)	(30.54)	(6.98)	(17.10)	(12.29)	(22.14)	(46.15)	(53.45)
SymI	0.16	1.51	1.01	27.24	27.89	20.82	26.46	−29.57	−10.18
	(−0.1)	(−0.72)	(−0.59)	(−0.21)	(−0.33)	(−0.22)	(−0.36)	(−0.64)	(−0.25)
AsymC	4.11**	4.12*	4.18**	303.18**	136.77*	209.22**	134.64*	−37.57	37.4
	(3.19)	(2.43)	(3.06)	(2.98)	(2.29)	(2.92)	(2.17)	(−0.9)	(−1.11)
AsymI	−0.18	0.68	0.12	−33.86	−45.7	−49.76	−13.19	−66.98	−44.69
	(−0.13)	(−0.32)	(−0.08)	(−0.27)	(−0.64)	(−0.59)	(−0.16)	(−1.16)	(−0.81)
N	544	640	1184	544	640	1184	544	640	1184

Note: * significant at 5%; ** significant at 1%. z statistics in parentheses

Table 4: Regression results

The regression results are reported in Table 4. While the treatment variables in this table indicate whether or not market prices, producer and total welfare are different when compared to the reference treatment SymC, we are interested in all pairwise cross-treatment differences. For this purpose we present in Table 5 the (two-tailed) p-values associated with Wald tests for equalities between pairs of the treatment dummy coefficients in the regressions reported in Table 4.

Table 5 contains an important piece of information: with one major and one minor exception, all pairwise treatment comparisons including treatment AsymC indicate significant differences. The major exception is that total surplus in the second half of the experiment and overall (see subtables TS2 and TS3 in Table 5) is not significantly different across treatments. The minor exception is that the difference in prices and producer surpluses between treatments AsymC and SymI when only later rounds are taken into account is significant only at the 10% level (see subtables MP2 and PS2 in Table 5)). All other pairwise comparisons are insignificant.

In particular, Tables 4 and 5 indicate that in treatment AsymC market prices are markedly higher than in all other treatments, and this is especially true for the first half of the experiment. This has a direct effect on consumer surplus (which is significantly lower in treatment AsymC than in all other treatments) and on producer surplus (which is significantly higher in treatment AsymC

calling rounds 1-17 the first half and rounds 18-37 the second half.

(MP1)	SymC	SymI	AsymC	**(PS1)**	SymC	SymI	AsymC	**(TS1)**	SymC	SymI	AsymC
SymI	0.924	—	—	SymI	0.833	—	—	SymI	0.718	—	—
AsymC	0.001	0.0040	—	AsymC	0.003	0.0023	—	AsymC	0.030	0.0191	—
AsymI	0.896	0.8163	0.0000	AsymI	0.786	0.5978	0.0001	AsymI	0.876	0.5914	0.0186
(MP2)	SymC	SymI	AsymC	**(PS2)**	SymC	SymI	AsymC	**(TS2)**	SymC	SymI	AsymC
SymI	0.474	—	—	SymI	0.743	—	—	SymI	0.525	—	—
AsymC	0.015	0.0947	—	AsymC	0.022	0.0880	—	AsymC	0.366	0.8375	—
AsymI	0.750	0.6845	0.0319	AsymI	0.525	0.3284	0.0001	AsymI	0.247	0.5046	0.5721
(MP3)	SymC	SymI	AsymC	**(PS3)**	SymC	SymI	AsymC	**(TS3)**	SymC	SymI	AsymC
SymI	0.558	—	—	SymI	0.826	—	—	SymI	0.805	—	—
AsymC	0.002	0.0177	—	AsymC	0.004	0.0074	—	AsymC	0.267	0.1947	—
AsymI	0.940	0.5707	0.0005	AsymI	0.556	0.3971	0.0000	AsymI	0.420	0.5473	0.1159

Table 5: p-values of (two-tailed) pairwise cross-treatment differences in market prices

than in all other treatments). Everything else being equal, higher prices should lead to lower total welfare in treatment AsymC. However, the results in Tables 4 and 5 indicate that total welfare in treatment AsymC is the same overall as in all other treatments. The only candidate explanation for this pattern is that a positive early-coordination effect outweighs a negative price effect in treatment AsymC. Conditional upon both firms charging the same price, only the level of price determines the level of total welfare. However, under convex costs, it is always preferable to split production between the two firms to fight the decreasing returns to scale. Therefore, coordination, be it on higher prices, leads to efficiency gains. Indeed, there is evidence for quicker coordination in treatment AsymC, which we document in the next section (see also Figure 1).

The results thus suggest that in treatment AsymC, although prices are higher, the overall effect on total surplus is nul, for the efficiency gains from coordination are reaped earlier. However, an important caveat applies to this finding: it obviously depends on the duration of the experiment. If the experiment had lasted for 80 periods and stabilized markets had been left undisturbed, the benefits of low prices in treatments SymC, SymI and AsymI would have eventually prevailed. Conversely, if markets had been shocked after only 20 periods, then the costs of miscoordinating would have been primordial. That is, the net effect depends on how long markets are in operation without disturbance.

5.2 Convergence patterns

In this section we analyze convergence patterns. Guiding questions are: Do markets converge? If they do, to which price and in which period? Also, we ask whether markets typically converge from above or from below? To answer these questions, we collect all relevant information in Table 6.

Do markets converge? We classify a market as having converged if both firms charge one and the same price in periods 31-37 where we allow for one exception in which one firm charges a price one unit higher or lower than the price the market converges to. The columns labeled "Convergence?" in Table 6 reveal that markets typically converge. Perhaps not surprisingly, whereas 7 out of 8 markets (87.5%) converge in treatment SymC (where firms are symmetric and this is known), in treatment AsymI (where firms are asymmetric and this is not known) "only" 6 out of 9 markets (66.7%) converge. The percentage of markets converging in the treatments SymI and AsymC is 77.7% and 83.3%, respectively.

Which price do markets converge to? The columns labeled "Where to?" in Table 6 show the prices markets converge to (if, in fact, they do). We observe that if markets converge, they either converge to a Nash equilibrium or to a collusive price of either 43, 44, or 45. There is only one exception in treatment AsymI where one market converges to the price of 50, distinctively above the collusive prices. The information regarding the prices markets converge to is also graphically represented in the upper half of Figure 2 that shows a histogram of prices to which markets converged for each treatment separately (conditional on convergence). We note that prices to which markets converge are quite dispersed in treatments SymC, SymI, and AsymI, typically ranging from a low Nash equilibrium to the price of joint-profit maximization of 44. However, the range of prices to which markets converge in treatment AsymC is quite concentrated ranging from 40 to 45. Thus, if markets do converge in treatment AsymC, they converge to a price higher than the highest Nash equilibrium price of 38. The average prices markets converge to in treatments SymC, SymI, AsymC, and AsymI are 39.4, 41.3, 43, and 39.8, respectively. This said, using a robust rank-order test (where each market counts as an independent observation), we find no significant differences in the prices markets converge to in any pairwise treatment comparison.

When do markets converge? The columns labeled "Which period?" in Table 6 show the period during which the price to which the market converged (if they did) started being charged by both firms. Again, perhaps it is easier to understand this information using graphical means. The lower half of Figure 2 shows the periods in which markets converge for each treatment separately.

	Treatment SymC					**Treatment** SymI			
Market	Conver-gence?	Where to?	Which period?	From where?	Market	Conver-gence?	Where to?	Which period?	From where?
1	yes	32	14	above	1	yes	33	28	above
2	yes	37	13	below	2	yes	36	5	below
3	yes	37	22	above	3	yes	44	1	constant
4	yes	38	19	below	4	yes	44	2	below
5	yes	44	2	below	5	yes	44	6	below
6	yes	44	2	below	6	yes	44	12	below
7	yes	44	16	below	7	yes	44	16	below
8	no			above	8	no			below
					9	no			
[Per cent] (Mean)	[87.5]	(39.4)	(12.6)			[77.7]	(41.3)	(10)	
%(perfectly collud. markets)	42.9					71.4			

	Treatment AsymC					**Treatment** AsymI			
Market	Conver-gence?	Where to?	Which period?	From where?	Market	Conver-gence?	Where to?	Which period?	From where?
1	yes	40	7	below	1	yes	35	6	above
2	yes	43	3	constant	2	yes	35	23	?
3	yes	43	5	below	3	yes	36	30	?
4	yes	44	3	constant	4	yes	40	1	constant
5	yes	45	3	below	5	yes	43	27	below
6	no			?	6	yes	50	10	below
					7	no			?
					8	no			?
					9	no			above
[Per cent] (Mean)	[83.3]	(43)	(4.2)			[66.7]	(39.8)	(16.2)	
%(perfectly collud. markets)	80.0					16.7			

Notes: In column "Convergence?" a "yes" ("no") indicates that a market has (has not) converged. In column "From where?" the entry "above" ("below") indicates that a market converged from above (below). The entry "constant" indicates that the same market price was charged throughout the experiment. The entry "?" in this column indicates that the direction of convergence cannot be determined.

Table 6: Convergence patterns.

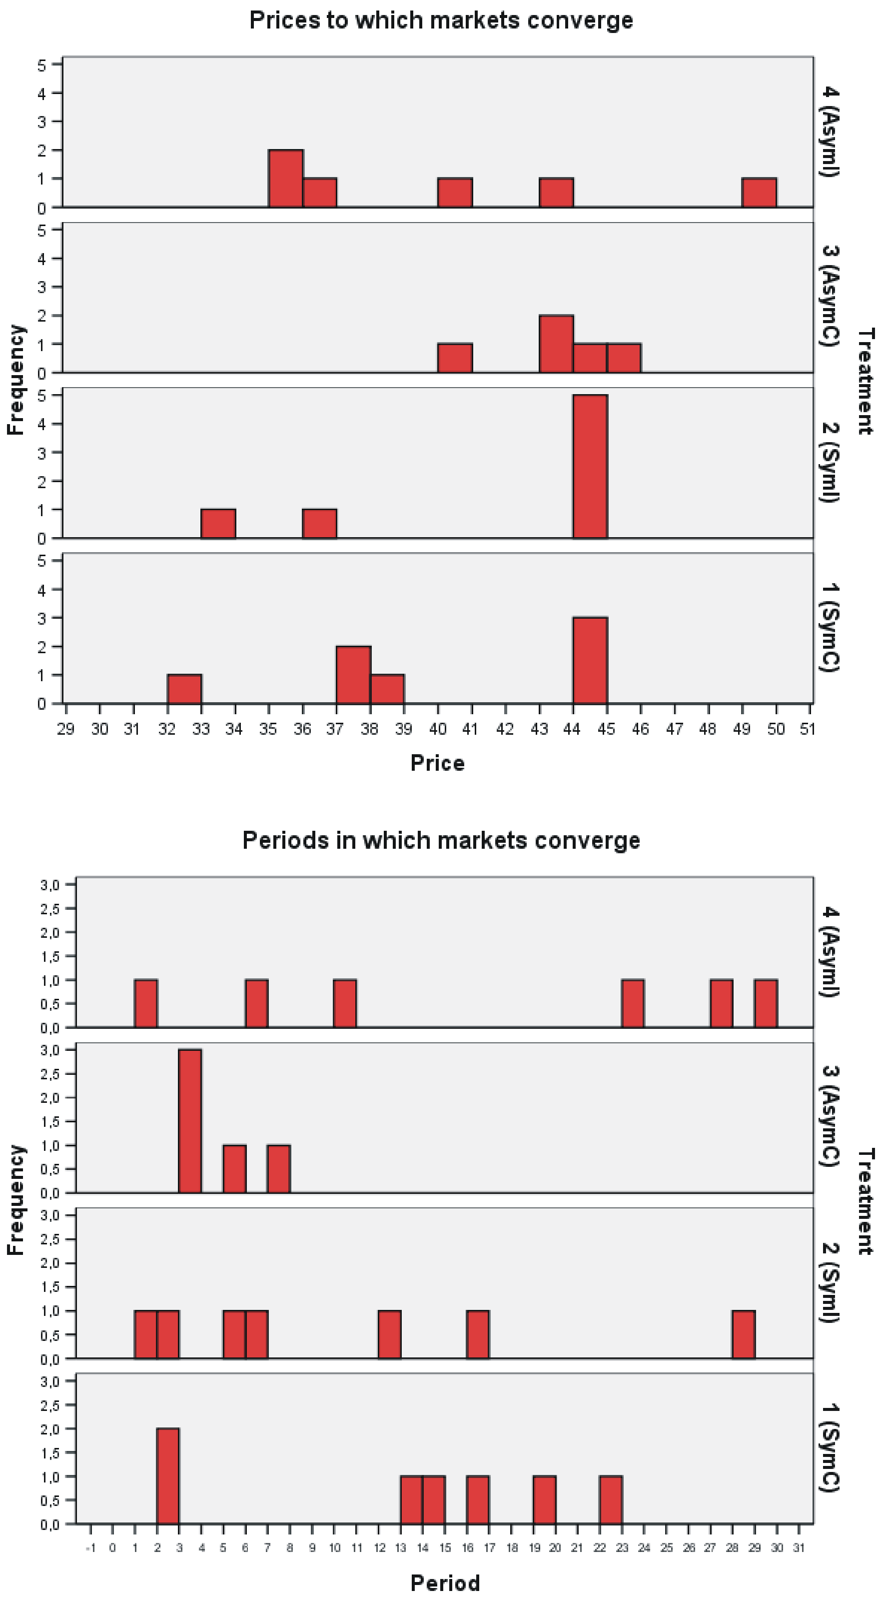

Figure 2: Histograms of prices markets converged to (top) and histograms of periods in which markets stabilized (bottom).

A pattern similar to the one above emerges. The ranges of periods in which markets converge are quite large in all treatments except for treatment AsymC where, conditional on convergence, markets converge within the first seven periods. Not surprisingly, it can take up to 28-30 periods until markets stabilize in the treatments with incomplete information. The average periods markets converge to in treatments SymC, SymI, AsymC, and AsymI are 12.6, 10, 4.2, and 16.2, respectively. Applying robust rank-order tests (where each market counts as an independent observation), we only find a weakly significant difference in the periods in which markets converge when comparing the two asymmetric markets ($p = 0.075$, two-tailed).

Finally and briefly, regarding the question which direction markets converge from, we note that those duopolies that are able to collude perfectly converge from below or right from the start. However, for those duopolies that converge to a Nash equilibrium price, the pattern is not clear.

6 Discussion

On the basis of our pilot experiment, we report that in the case of Bertrand competition with convex costs, asymmetry in payoffs, when coupled with perfect information about those payoffs on the part of subjects, seems associated with quicker convergence to higher prices. Those findings are not in line with those that have been reported about Cournot competition under similar designs.

Mason and Philips (1992) studied the impact of asymmetry on Cournot duopolists in a repeated-game environment and reported that "asymmetric markets are less cooperative and take longer to reach equilibrium than symmetric markets," whereas we find the opposite.

In the same environment, Mason and Philips (1997) tested the impact of complete information on duopoly outcomes and reported that "symmetric markets are more cooperative when profitability is common knowledge. Asymmetric markets are unaffected by information differences." Again, we find the opposite, as symmetric markets are unaffected by our information condition, while asymmetric markets react strongly to the completeness of information by reaching higher levels of cooperation quicker.

There are several possible explanations for these findings. First, our pilot experiment involves a relatively small number of pairs in each treatment, and particularly so in our asymmetric costs, complete information treatment. Obviously, if design is not an issue, accumulating more data by holding additional sessions is the next step to take so as to assess the results. Second, there may be a fundamental difference between experimental behavior under Cournot competition and under

Bertrand competition with convex costs. In other words, something in the asymmetric-complete information treatment may be conducive to more collusion.

In that last case, it is tempting to turn to the voluminous literature on coordination games to look for coordination-facilitating practices. This literature, recently surveyed by e.g. Devetag and Ortmann (2007), has focused on two classes of games: order-statistic games (e.g. minimum effort game) or stag-hunt games (typically, 2×2 or 3×3). Both classes are coordination games with several Pareto-ranked Nash equilibria. However, we note that the results are not directly comparable to ours. For one thing, in order-statistic games, all possible actions are Nash equilibria, and payoff-dominance and risk-dominance call for different outcomes, two features that are absent from Bertrand competition with convex costs. In stag-hunt games, embedment in a larger payoff matrix can ensure that not all actions are equilibrium actions. Yet, with the exception of one particular experiment (Rankin, Van Huyck and Battalio, 2000), there is again a tension between the safe choice and the high-payoff choice. Thus, most researchers have focused on testing the predictive power of risk versus payoff dominance. In addition, most of the experiments on order-statistic games have used fixed matching among many players (at least 6), while most of the experiments on stag-hunt games have used random-matching protocols (or related protocols), some designs that are more conducive to the play of a static Nash equilibrium than to cooperation.

This calls for more research to link our results to this literature. In that respect, a switch to random matching may be an obvious first step. Under fixed matching, in any case, our plans include an increase in the number of firms to assess the robustness of our findings. A more thorough investigation of the convergence patterns is also called for.

From the point of view of welfare, our main message is the existence of a trade-off between higher prices (a demand-side negative effect) and earlier coordination (a supply-side positive effect) under asymmetry and complete information. Indeed, by a total welfare standard, our results suggest that, provided the time span during which markets stabilize is not too long, one should favor asymmetric market structures and exchanges of information about costs *because* they are more conducive to cooperation, a claim that runs contrary to received wisdom in competition policy circles. Whether this finding is robust awaits further research. We plan to introduce demand shocks (or entry and exit of players) to test for the net benefits of cooperation independently of the duration of the interaction between firms.

References

[1] Abbink, K. and J. Brandts (2008): 24. Pricing in Bertrand competition with increasing marginal costs, *Games and Economic Behavior* 63, 1-31.

[2] Arellano, M. and S. Bond (1991): Some tests of specification for panel data: Monte Carlo evidence and an application to employment equations, *The Review of Economic Studies* 58: 277-97.

[3] Benoît, J.-P. and V. Krishna (1985): Finitely-Repeated Games, *Econometrica*, 53(4), 905-922.

[4] Boone, J., W. Müller and A. Ray Chaudhuri (2008): Bertrand competition with asymmetric costs: Experimental evidence, mimeo.

[5] Dastidar, K. G. (1995): On the existence of pure strategy Bertrand equilibrium, *Economic Theory* 5, 19-32.

[6] Devetag G. and A. Ortmann (2007); When and why? A critical survey on coordination failure in the laboratory, *Experimental Economics*, 10, 331-344.

[7] Dufwenberg, M. and Gneezy, U. (2000): Price competition and market concentration: an experimental study, *International Journal of Industrial Organization* 18, 7-22.

[8] Dufwenberg, M. and Gneezy, U. (2002): Information disclosure in auctions: an experiment, *Journal of Economic Behavior & Organization* 48, 431–444.

[9] Fouraker, L. and S. Siegel (1963): *Bargaining Behavior*, New York: McGraw–Hill.

[10] Harsanyi, J. C. and Selten, R. (1988): *A General Theory of Equilibrium Selection in Games*, Cambridge, Mass. and London: MIT Press.

[11] Hoernig, S. (2002): Mixed Bertrand Equilibria under Decreasing Returns to Scale: An Embarrassment of Riches, *Economics Letters*, 74, 359-362.

[12] Huck, S., Müller,W., Normann, H.-T. (2001): Stackelberg beats Cournot: On collusion and efficiency in experimental markets, *Economic Journal* 111 (474), 749–765.

[13] Kühn, K.-U. and X. Vives (1995): Information Exchanges among Firms and their Impact on Competition, Office for Official Publications for the European Community, Luxemburg.

[14] Mason, C. F. and O. R. Phillips (1992): Duopoly Behavior in Asymmetric Markets: An Experimental Evaluation, *Review of Economics and Statistics* 74(4), 662-670.

[15] Mason, C. F. and O. R. Phillips (1997): Information And Cost Asymmetry In Experimental Duopoly Markets, *Review of Economics and Statistics* 79(2), 290-299.

[16] Rankin, F., J. B. Van Huynck and R. C. Battalio (2000): Strategic similarity and emergence of conventions: Evidence from payoff perturbed stag hunt games, *Games and Economic Behavior*, 32, 315-337.

[17] Vives, X. (2006): Information sharing and antitrust, in: *The pros and cons of information sharing*, Mats Bergman (ed.), Swedish Competition Authority, 83-100.

[18] Weibull, J. W. (2006): Price competition and convex costs, *SSE/EFI Working Paper Series in Economics and Finance* No 622.

[19] Wolfram, C. D. (1998): Strategic Bidding in a Multiunit Auction: An Empirical Analysis of Bids to Supply Electricity in England and Wales, *RAND Journal of Economics* 29(4), 703-725.

[20] Wolfram, C. D. (1999): Measuring Duopoly Power in the British Electricity Spot Market, *American Economic Review* 89(4), 805-826.