Collusion in experimental Bertrand duopolies with convex costs: the role of information and cost asymmetry*

Cédric Argenton†
CentER & TILEC, Tilburg University

Wieland Müller‡
CentER & TILEC, Tilburg University

November 6, 2009

Abstract

We report the results of a series of experimental Bertrand duopolies where firms have convex costs. Theoretically, these duopolies are 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 average 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 Ola Andersson, Dirk Sliwka and Bert Willems as well as seminar participants at Tilburg University 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

The ability of economic agents to reach outcomes that are collectively desirable, even if (and especially when) they conflict with short-term individual incentives, has been a primary focus of game theory and experimental economics since their beginnings. In market environments, the possibility for firms to collude and sustain high prices is of high theoretical interest as well as of high policy relevance, as antitrust authorities throughout the world see the fight against collusion as one of their main tasks. There is now a large body of literature dealing with collusion.[1] The received wisdom about collusion-facilitating market characteristics is that firm symmetry (in product range, costs, demands, etc.) and complete information (about cost and demand conditions, past actions, etc.) make collusion easier. For instance, Motta (2004) notes that "the more firms are asymmetric (in capacities, market shares, costs or product range) the less likely collusion will be" (p. 143) and proceeds to show in a technical section (4.2.5) that "symmetry helps collusion." The US "horizontal merger guidelines" (revised in 1997) state at section 2.11[2] that:

> [m]arket conditions may be conducive to or hinder reaching terms of coordination. For example, reaching terms of coordination may be facilitated by product or firm homogeneity and by existing practices among firms, practices not necessarily themselves antitrust violations, such as standardization of pricing or product variables on which firms could compete. Key information about rival firms and the market may also facilitate reaching terms of coordination. Conversely, reaching terms of coordination may be limited or impeded by product heterogeneity or by firms having substantially incomplete information about the conditions and prospects of their rivals' businesses, perhaps because of important differences among their current business operations. In addition, reaching terms of coordination may be limited or impeded by firm heterogeneity (...).

Results from experimental Cournot markets are remarkably in line with these assertions.[3]

In this paper, the focus is on tacit collusion in Bertrand duopolies with convex costs. Those markets are interesting from both a theoretical and a practical point of view. Indeed, such games

[1] For a state-of-the art review of the theory of repeated games, see Mailath and Samuelson (2006). For a survey of experimental results on collusion, see Haan, Schoobeek and Winkel (2009). A meta-study of experiments on collusion is performed by Engel (2007).

[2] The guidelines can be consulted at http://www.usdoj.gov/atr/public/guidelines/horiz_book/toc.html (28 July 2009).

[3] See Fouraker and Siegel (1963), Mason and Philips (1992, 1997).

typically admit a whole interval of Pareto-ranked Nash equilibria which may or may not contain the joint profit-maximizing strategy profile. When they don't, they share features of both large coordination games and prisoner's dilemma games. Besides, competition in certain sectors can be stylized as Bertrand pricing under convex costs. Indeed, utilities such as gas, water and electricity providers face rising marginal cost and are typically under the obligation to serve all demand addressed to them at posted prices.[4] In many countries, the markets on which they operate are dominated by two main firms.[5]

In our experiment, we examine whether repeatedly-interacting duopolists are able to coordinate on high prices under varying cost and information conditions. More precisely, we analyze price choices in symmetric and asymmetric cost 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. We want to know what the impact of differences in costs and (ex ante) payoff information is on *(i)* market prices, and *(ii)* the time it takes for markets to stabilize. Notice that from a welfare perspective there is a tension in the short run between (early) coordination (which is welfare-enhancing under convex costs as production is spread over two firms) and the ability of duopolies to charge collusive prices (which is obviously total welfare-reducing).

We obtain results that are at odds with those found in experimental Cournot markets. Indeed, we report some evidence that asymmetry *helps* collusion under complete information (which nevertheless lead to an overall positive welfare effect, in our environment), while incomplete information does *not* impair firms' ability to "agree" on high prices in symmetric markets (but leads to significant delays in coordination in asymmetric markets). Hence, we argue that in assessing the likelihood of collusion, it is necessary to pay more attention to the nature of market institutions, rather than rely on general rules which, although broadly valid for Cournot markets, may not apply to other environments.

The paper is organized as follows. Section 2 presents a short overview of previous experimental results on Bertrand competition in homogenous markets. Section 3 describes our experimental design as well as the theoretical predictions. In section 4 we report the experimental results, concentrating on the research questions outlined above. Finally, in section 5 we summarize our results, compare them to the Cournot experiment literature, and discuss their relevance with regards to

[4] More generally, "the Bertrand assumption is plausible when there are large costs of turning customers away" (Vives, 1999, p.118).

[5] 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).

competition policy.

2 Experimental literature on Bertrand competition in homogenous markets

Despite its central place in such fields as industrial organization, there is only a small experimental literature dealing with Bertrand competition in homogeneous markets with elastic demand. Next to their famous experiments with quantity-setting firms, Fouraker and Siegel (1963) report on Bertrand markets with either two or three symmetric firms under fixed matching. As we do in this paper, they consider markets in which firms have either complete or incomplete information about payoffs. Regarding the two dimensions of their design, the authors summarize their results as follows (p. 199): "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), Andersson and Wengström (2007), Boone, Müller, and Ray Chaudhuri (2008), Bruttel (2009) and Fatas, Haruvy, and Morales (2009) explicitly study Bertrand competition in their experiments with perfectly inelastic demand.[6]

Abbink and Brandts (2008, henceforth A&B) is the first study reporting on symmetric Bertrand competition experiments where firms have convex costs. They ran experiments with fixed groups 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. The authors 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.

Next to results from a Bertrand duopoly market with linear costs, Fatas, Haruvy, and

[6] There is a tradition of experiments on Bertrand pricing under product differentiation (an early reference is Dolbear et al. [1968]) but in that case the demand system is continuous. There is also a large body of literature on auctions, which can sometimes be interpreted as Bertrand games with perfectly inelastic demand. See Kagel (1995) for a survey and Dufwenberg and Gneezy (2002) for an example.

Morales (2009) also report results from Bertrand duopoly market with quadratic costs. In their experiment demand is inelastic, matching is fixed for 20 periods, and information is complete. The authors report market prices to converge to the payoff-dominant Nash equilibrium. Moreover, the authors check the relevance of alternative behavioral models.

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 automated buyers buy from the firm(s) offering the lowest price while sellers behave strategically. Although the experiment was described to the participants as a pricing game between firms, they were not given the details of the model. Instead, they were presented with payoff table(s).

Subjects were paired in one of four treatments. Treatments varied with respect to the difference in costs and the information that subjects were given about the profits to the other firm in their market. In treatments we call "symmetric," the profit tables of the two paired subjects were identical. In treatments we call "asymmetric", they were different. In some instances, subjects were given their profit 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 profit 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 payoff tables were generated from a linear demand curve ($Q = 100 - 1.5P$) and quadratic cost curves ($c(q) = c_i q^2$), with all numbers rounded to the nearest integer. 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, demand was split equally. In symmetric treatments, cost functions were identical ($c_1 = c_2 = 0.6$). In asymmetric treatments, one of the two subjects was endowed with a low cost parameter ($c_1 = 0.55$) while the other was endowed with a high cost parameter ($c_2 = 0.65$).[7] In all treatments, payoffs were expressed in a fictitious monetary unit, called "point."

[7] This design allows us to meaningfully compare welfare across treatments later on: conditional on both firms charging the same price, total welfare is the same in symmetric and asymmetric markets.

Subjects were told that negative numbers stood for losses, which were indeed possible in the range of low prices.

Bertrand competition is not synonymous with perfect competition when firms face convex costs. Dastidar (1995) proved that in the symmetric case there is a whole interval of pure-strategy Nash equilibrium prices.[8] The lower bound of this interval is determined by average-cost pricing. The upper bound is determined by the incentive marginally 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, a pure-strategy Nash equilibrium always exists. It may be unique or not, symmetric or not. In the linear-quadratic specification we implement, it is still the case that there is a continuum of symmetric equilibria.

Figure 1 reproduces the payoff table we used in symmetric treatments. As can be checked, all prices in $\{21, 22, ..., 39\}$ are 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 is 32.[9] The monopoly price is 49 but, due to decreasing returns to scale, the price maximizing joint profits (and thus an obvious candidate for tacit collusion) is 44.[10]

In our asymmetric treatments, the range of symmetric 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.[11]

Several features are of interest. First, 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

[8] See also Weibull (2006). There are also continua of nonzero-profit mixed-strategy equilibria, as demonstrated by Hoernig (2002).

[9] This is the equivalent in our specification of the "near-magic" number 24 in A&B.

[10] 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).

[11] When prices and quantities are continuous variables, it is possible to solve for the competitive equilibrium price (31.5) or the Cournot equilibrium price (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 1: Payoff table for symmetric treatments

so that low Bertrand prices are in a sense riskier. Second, from the point of view of firms, Nash equilibria are Pareto-ranked: the higher the equilibrium price, the higher the equilibrium profits. (The ordering is of course reversed when one considers consumer surplus.) Third, 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. Under cost symmetry the perfectly collusive price is 44. Fourth, 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 firm's "preferred" collusive prices (43 for the low-cost and 45 for the high-cost firm) are separated by two integers.[12]

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).[13] 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 at the start of the experiment and interacted with him or her in all 40 rounds. Subjects were made aware of this feature in the instructions. 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. 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 not to foster imitation (see, e.g., Apesteguia *et al.*, 2007).

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. Each treatment lasted between 30 and 45 minutes. The average

[12] 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.

[13] The instructions for all treatments are available at:

http://center.uvt.nl/staff/muller/InstructionsBetrandWithConvexCosts.pdf. The instructions for two representative treatments are reprinted in the Appendix.

	Complete Information	Incomplete Information
Symmetric Costs	SymC $c_1 = c_2 = 0.6$ $(19 \times 2 = 38$ subjects)	SymI $c_1 = c_2 = 0.6$ $(23 \times 2 = 46$ subjects)
Asymmetric Costs	AsymC $c_1 = 0.55, \quad c_2 = 0.65$ $(23 \times 2 = 46$ subjects)	AsymI $c_1 = 0.55, \quad c_2 = 0.65$ $(23 \times 2 = 46$ subjects)

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

monetary earnings across all treatments were 12.75 Euros.

We analyze data from 12 lab sessions, 3 for each treatment. We have data on 19 pairs in the SymC treatment, 23 pairs in the SymI, AsymC, and AsymI treatments. Table 2 summarizes the design.

There is no univocal theoretical prediction for the outcome of the game. Shared expectations and common knowledge of rationality can give rise to the play of any Nash equilibrium in a one-shot context. However, payoff dominance calls for the highest Nash equilibrium price to be played.[14] This disregards the fact that in our experimental setting the game was in fact repeatedly played by the same players. Given the multiplicity of static Nash equilibria, tacit collusion on higher prices can theoretically arise, even under a finite horizon, as well-known from Benoît and Krishna (1985). Models of imitation (A&B or Alós-Ferrer et al., 2000) predict convergence towards the competitive equilibrium under some conditions. Since we did not include the profit to the other firm in the feedback information received by participants, we did not expect them to follow that line of reasoning. On the contrary, on the basis of existing literature on collusion in duopolies (see footnote 1), we were expecting players, after a series of trials and errors, to settle on prices close to the collusive level, and the closer, the more information and symmetry there was in the market.

[14] No Nash equilibrium risk-dominates all other ones in pairwise comparisons.

	SymC	SymI	AsymC	AsymI
Rounds 1-17	37.3	37.3	40.0	38.7
	(0.83)	(0.68)	(0.65)	(0.76)
Rounds 18-37	39.0	39.7	41.2	40.1
	(1.07)	(0.82)	(0.63)	(0.91)
Rounds 1-37	38.3	38.6	40.6	39.5
	(0.87)	(0.71)	(0.59)	(0.79)

Table 3: Mean market prices (standard errors of the mean in parentheses)

4 Experimental results

We report the results of the experiment in two subsections.[15] In a first subsection, we give a quick overview of the results and then test for differences in market prices, consumer surplus, producer surplus as well as total surplus across treatments. In a second subsection, we focus on convergence patterns.

4.1 Prices and welfare

Summary statistics of the experimental results are given in Table 3, which displays the mean of market prices and the standard error of the mean in the four treatments (separately for various time intervals).[16] The market price is defined as the minimum of the prices posted by the two firms in a market. This is the price at which consumers would obtain the good in a market characterized by Bertrand competition. Observe in Table 3 that market prices are highest on average (and noticeably less dispersed) in treatment AsymC, followed by prices in treatment AsymI. In contrast, prices are lowest and very similar in the two symmetric treatments.[17]

Figure 1 shows histograms of market prices in rounds 1-37 for all four treatments. In all treatments, pricing above the highest Nash equilibrium was quite common. The distribution of market prices is quite similar in the two symmetric treatments. In fact, in both of them, the mode is the perfectly collusive price of 44 and the second most often chosen price is 32 (which is the lowest price in the range of Nash equilibria which involves no loss in case of miscoordination).

[15] Preliminary results from a few pilot sessions were reported in Argenton and Müller (2009).

[16] Due to a clear endgame effect we excluded the last three rounds. We divided the remaining periods in two, refering to rounds 1-17 as the first half and to rounds 18-37 as the second half.

[17] Statistics regarding individual prices instead of market prices display the same features.

In contrast, the distribution in treatment AsymC is single-peaked at 43, which is the price that maximizes the profit of the low-cost firm, conditional on both firms charging the same price. Finally, the distribution of market prices in treatment AsymI appears to be more evenly distributed in comparison to all other treatments. Note that in treatment AsymI there are two markets that converge relatively early (one in round 3, the other in round 10) to a price of 50 which is much higher than the collusive prices of 43 (low-cost firm) or 45 (high-cost firm). In no other market did we observe convergence to such a high price.

It is instructive to quickly compare the histograms in Figure 1 with the distribution of market prices in the duopoly treatment in A&B. The distribution in our treatment SymC (as well as the one in treatment SymI) is very similar to the one reported in Figure 3 in A&B for the case $n = 2$. We also find that the lowest price in the range of Nash equilibria which involves no loss in case of miscoordination (32 in our case and 24 in A&B) is chosen quite often in our symmetric duopoly markets, despite the difference in feedback information.[18] However, in our asymmetric duopoly markets, the corresponding price of 33 is only chosen in very few cases: the focality of this Nash equilibrium thus seems to be the product of special circumstances, although it is impossible to tell at this stage whether symmetry or full feedback information (or both) is (are) crucial.

The evolution of the average market price in all treatments is shown in Figure 2. 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. This increase happens at a very fast rate in treatment AsymC, followed by AsymI and then, at a clearly slower rate, in the two symmetric treatments. Second, treatment AsymC stands out as the average market price in this treatment is clearly higher than in all other treatments. Average market prices are lowest in the two symmetric treatments SymC and SymI (with not much of a difference between the two) and they are somewhat in-between in treatment AsymI. Third, in all treatments we observe a typical endgame effect with average prices sharply decreasing in the last two or three periods.[19]

We now turn to testing for differences across treatments in market prices and surpluses. For this purpose, we run the following GLS panel model

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

where p_{jt} is the market price in market j in round t and ε_{jt} is a market-specific error term. D_{SymC},

[18] In A&B, after each round, players get to learn not only the prices posted by firms in their market but also the payoffs to all of them.

[19] See Selten and Stoecker (1986) for a classical investigation of this phenomenon.

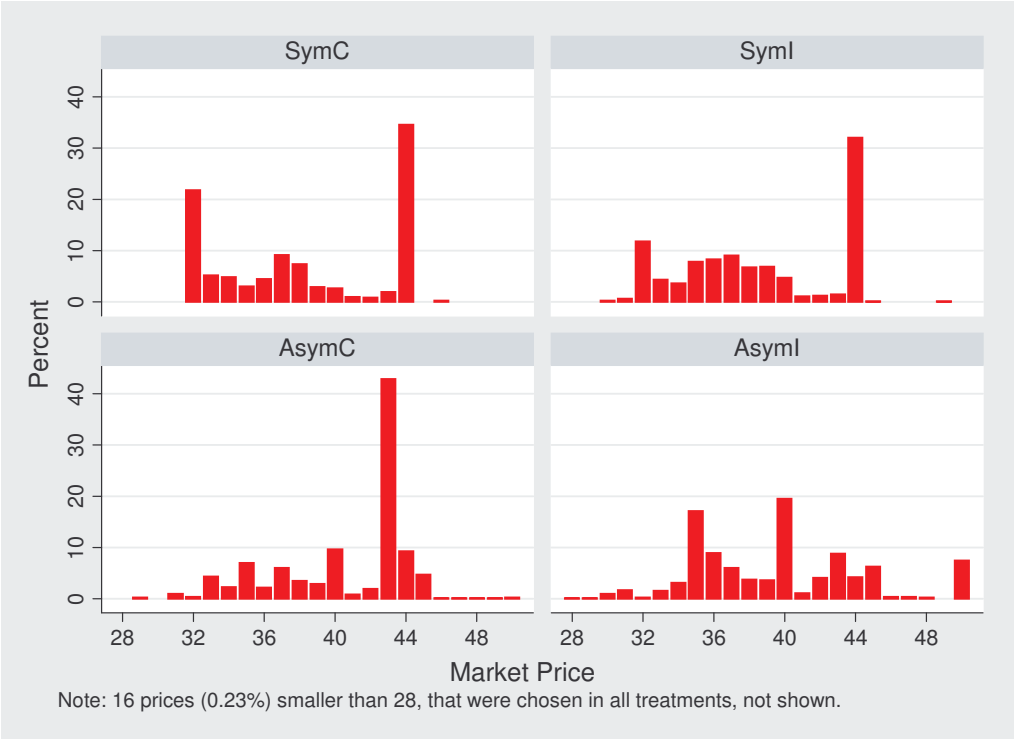

Figure 1: Distribution of market prices in periods 1-37.

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 relatively to this benchmark. Similar regressions were run to test for across-treatment differences in consumer, 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. We also clustered errors by market to account for heterogeneity. (For a similar approach see e.g. Mason and Phillips 1992, 1997). We ran regressions separately for the first half, the second half as well as for all rounds (see footnote 16).

The regression results for market prices and total surplus are reported in Table 4. While the treatment variables in this table indicate whether or not market prices 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 treatment dummy coefficients in regression 1.

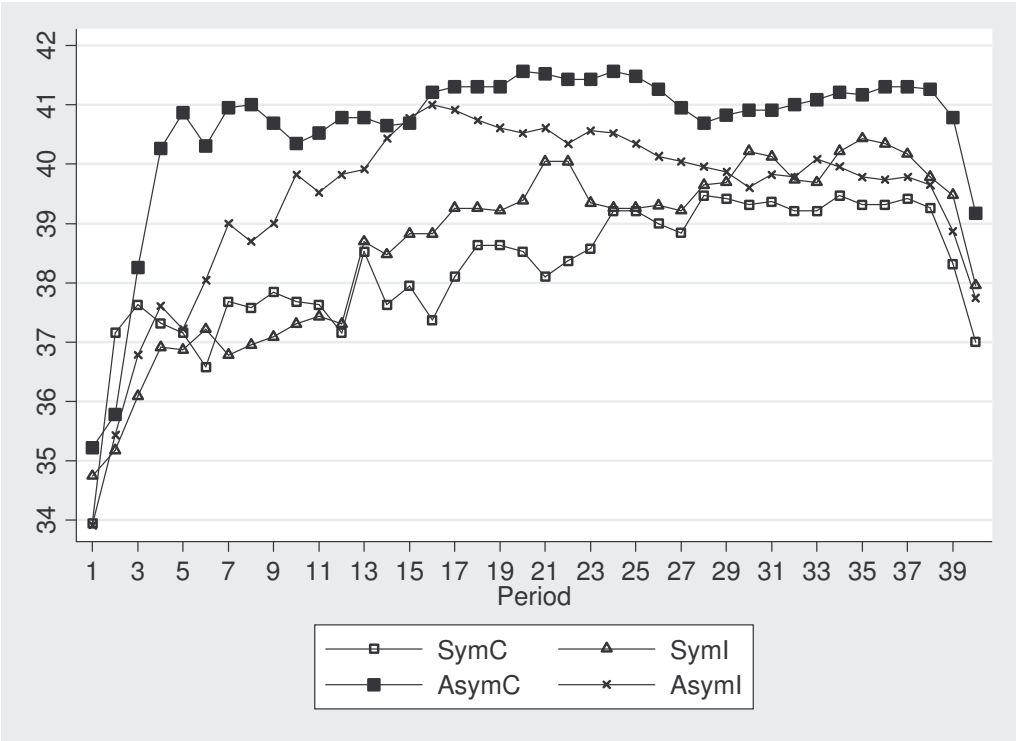

Figure 2: The evolution of market prices over time.

	Market Price			Total Surplus		
	(MP1)	**(MP2)**	**(MP3)**	**(TS1)**	**(TS2)**	**(TS3)**
Rounds →	1-17	18-37	1-37	1-17	18-37	1-37
Constant	36.88***	39.03***	37.75***	1,309.04***	1,525.30***	1,415.72***
	(51.50)	(46.77)	(57.61)	(40.93)	(65.64)	(61.40)
SymI	0.30	0.69	0.44	50.60	43.27	41.93
	(0.31)	(0.61)	(0.50)	(1.17)	(1.38)	(1.35)
AsymC	2.42**	2.25**	2.02**	91.62**	27.98	56.24*
	(2.50)	(2.00)	(2.28)	(2.12)	(0.89)	(1.80)
AsymI	1.34	1.20	0.78	30.96	6.95	14.21
	(1.39)	(1.07)	(0.88)	(0.72)	(0.22)	(0.46)
N	1496	1760	3256	1496	1760	3256

Note: * significant at 10%; ** significant at 5%; *** significant at 1%. z statistics in parentheses

Table 4: Regression results

(MP1)	SymC	SymI	AsymC
SymI	0.759	—	—
AsymC	**0.012**	**0.0210**	—
AsymI	0.165	0.2551	0.2420

(TS1)	SymC	SymI	AsymC
SymI	0.242	—	—
AsymC	**0.034**	0.3184	—
AsymI	0.474	0.6327	0.1400

(MP2)	SymC	SymI	AsymC
SymI	0.539	—	—
AsymC	**0.046**	0.1456	—
AsymI	0.286	0.6340	0.3276

(TS2)	SymC	SymI	AsymC
SymI	0.168	—	—
AsymC	0.373	0.6087	—
AsymI	0.825	0.2240	0.4815

(MP3)	SymC	SymI	AsymC
SymI	0.615	—	—
AsymC	**0.023**	**0.0622**	—
AsymI	0.378	0.6911	0.1422

(TS3)	SymC	SymI	AsymC
SymI	0.178	—	—
AsymC	**0.071**	0.6293	—
AsymI	0.648	0.3497	0.1562

Note: Statistically significant results are highlighted in bold font.

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

According to Tables 4 and 5, market prices in treatment AsymC are always significantly higher than in treatment SymC (and usually higher than in treatment SymI). On the other hand, there are no statistically significant differences in market prices between the symmetric treatments SymC and SymI and between the two asymmetric treatments AsymC and AsymI. The finding of statistically higher market prices in treatment AsymC compared with market prices in treatment SymC has direct effects on consumer surplus (which is significantly lower in treatment AsymC than in treatment SymC) and on producer surplus (which is significantly higher in treatment AsymC than in treatment SymC).[20]

Since most of our observations are at or above the competitive price, higher prices should lead to lower total welfare in treatment AsymC as compared to treatment SymC. However, the results on the right-hand side in Tables 4 and 5 indicate that total welfare in treatment AsymC is higher than in treatment SymC in the first half of the experiment and overall. The only candidate explanation for this pattern is that a positive early-coordination effect outweighs a negative price effect in treatment AsymC as compared to treatment SymC. 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

[20] Results regarding across-treatment differences in consumer surplus and producer surplus are reported in Appendix A.

to scale. Therefore, coordination, be it on higher prices, leads to efficiency gains. Indeed, Figure 2 suggests quicker coordination in treatment AsymC, which we document and discuss in the next section. The results thus suggest that in treatment AsymC, although prices are initially higher, the overall effect on total surplus is positive, for the efficiency gains from coordination are reaped earlier.

4.2 Convergence patterns

In this section we analyze convergence patterns. We try to see whether markets "stabilize" or "converge" (i.e. whether firms end up charging the same price). If so, we want to know to which price and in which period they converge.

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. To some extent, this definition is arbitrary: it aims at capturing the idea that players have eventually reached an understanding, without excluding the possibility of one (failed) attempt at switching to another price.

We find that markets typically converge. In fact, the percentage of markets converging in treatments SymC is 73.3% (14 out of 19 markets) whereas this number is 78.3% (18 out of 23 markets) in all other treatments. Hence, the percentage of converging markets is, perhaps surprisingly, lowest in the treatment with symmetric firms and full information. However, the difference is not statistically significant.

Figure 3 shows the evolution of average market prices in all four treatments for those markets that eventually converged to a stable price. This figure suggests that average prices are highest and convergence to stable prices fastest in treatment AsymC, while average prices (and their evolution over time) are very similar in the two symmetric treatments.

Which prices do markets converge to? Detailed information regarding the prices to which markets converge is graphically represented in Figure 4, which shows a histogram of prices to which markets converged for each treatment separately (conditional on convergence). We observe that if markets converge, they usually converge either to a Nash equilibrium or to a (quasi-)perfectly collusive price of either 43, 44, or 45. (Perfect collusion is particularly common in treatment AsymC.) However, there are also markets that converge to a price between the highest Nash price and the perfectly collusive price. Moreover, as mentioned earlier, there are two markets in AsymI that converges to the price of 50, distinctively above the collusive prices. Prices to which

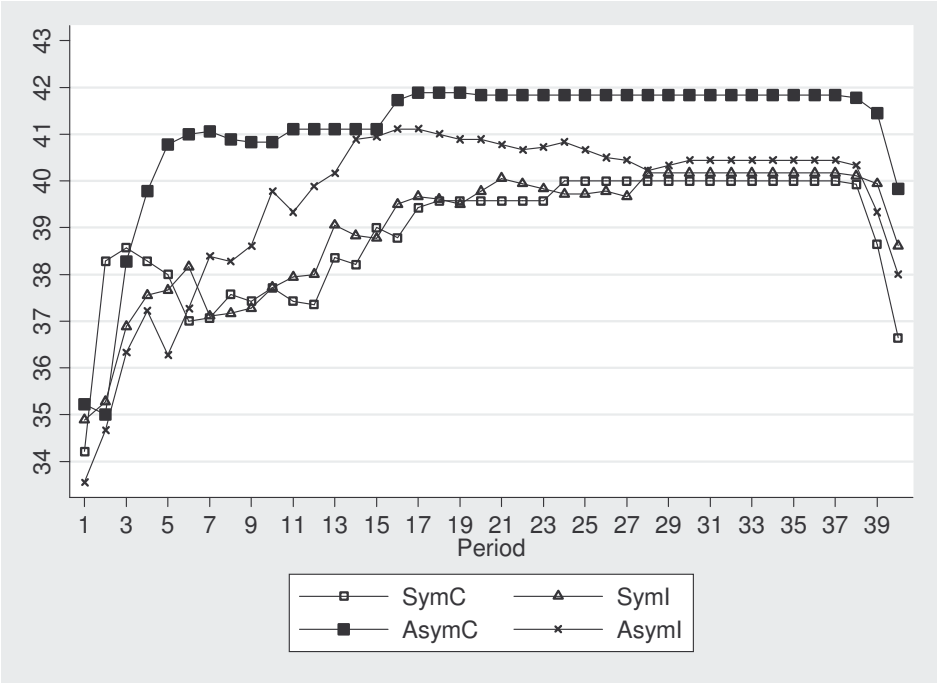

Figure 3: The evolution of market prices over time for markets that eventually converged to a stable price.

markets converge appear most concentrated in treatment AsymC. The average prices to which markets converge in treatments SymC, SymI, AsymC, and AsymI are 40.0, 40.2, 41.9, and 40.8, respectively. This said, using a two-tailed Mann-Whitney U 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. This suggests that the higher average market prices in treatment AsymC are attributable to markets that do not converge (in addition to differences earlier in the game). Indeed, regressions run separately on the subsamples of markets that have converged and markets that have not (available on request from the authors) corroborate this hypothesis.

When do markets converge? Figure 5 shows a histogram of the periods in which markets converged for each treatment separately. Conditional on convergence, such a period is defined as the beginning of the time window during which both firms uninterruptedly posted the price they converged to. Consistent with Figure 2, showing the evolution of average market prices, it appears that markets in treatment AsymC converge earlier than markets in the other treatments on average.[21] In fact, the average periods in which markets converge in treatments SymC, SymI,

[21] By the nature of the markets we consider, the cardinality of the set of Nash equilibria is smaller in the asymmetric

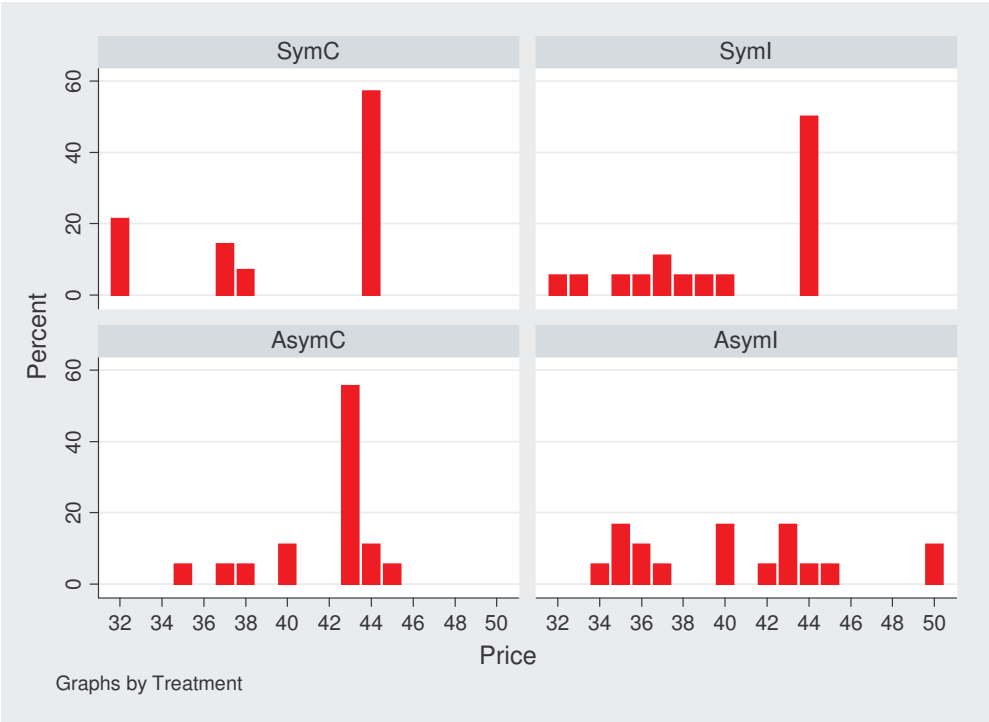

Figure 4: Histograms of prices to which markets converged

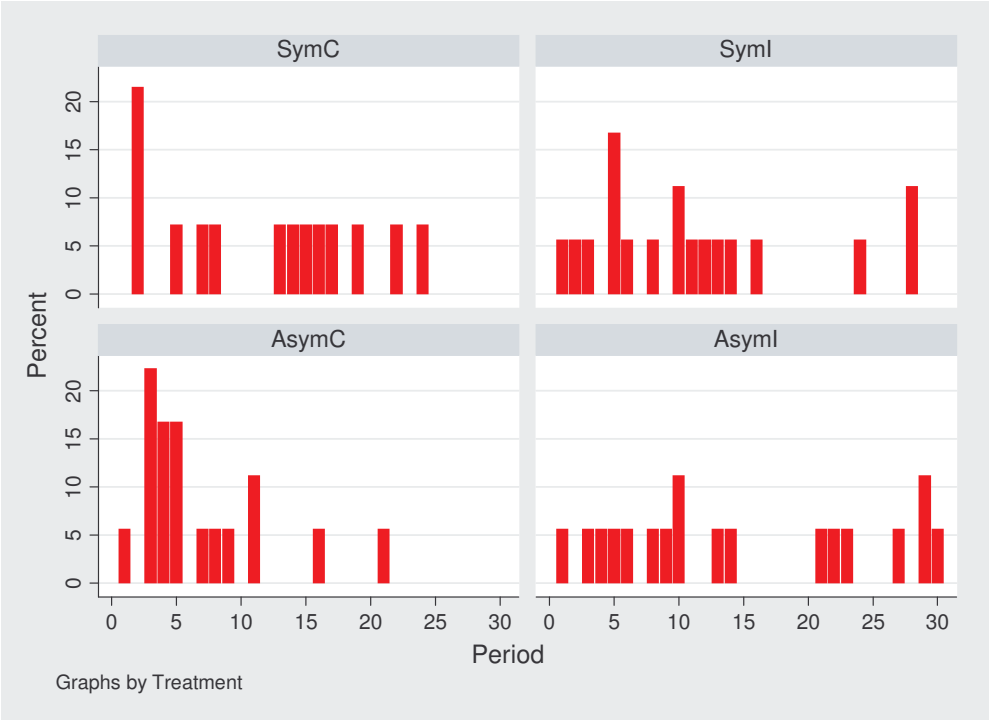

Figure 5: Histograms of periods in which markets stabilized

	SymC	SymI	AsymC
SymI	0.6341	—	—
AsymC	**0.0864**	**0.0751**	—
AsymI	0.4818	0.3498	**0.0111**

Note: Statistically significant results are highlighted in bold font.

Table 6: p-levels of pairwise Mann-Whitney U tests on differences of periods at which markets converged

AsymC, and AsymI are 11.9, 11.2, 6.8, and 14.6, respectively. The results of two-tailed Mann-Whitney U tests of pair-wise across-treatment differences in the period of convergence (where each market counts as an independent observation) are given in Table 6. We find that markets in treatment AsymC converge significantly earlier than in other treatments.[22]

However, the analysis above does not account for markets that have not converged. To account for this form of right-censoring, we next turn to a Kaplan and Meier (1958) time-to-failure analysis, where "failure" in our context refers to "convergence" of a market. The Kaplan-Meier non-parametric estimate of the probability of failure (i.e., of convergence) after time t is given by $1 - \pi_{t_i<t} = \frac{n_i - d_i}{n_i}$, where n_i is the number of markets that have not yet converged at t_i and d_i is the number of markets that converged at time t_i. The probabilities of failure are plotted for all treatments in Figure 6. Recall that while 73.3% of the markets in treatment SymC converge, 78.3% of markets in all other treatments do so. This is why the graphs in Figure 6 do not approach 1. It is apparent from the Figure that markets in treatment AsymC converge much earlier that those in all other treatments. Also, there is quite a number of markets in treatment AsymI that converge relatively late, which is consistent with the information given in Figure 5. Using two-tailed pairwise Wilcoxon-Breslow-Gehan tests for equality of the failure functions shown in Figure 6 indicate that only the speed of convergence in treatments AsymC and AsymI are significantly different ($p = 0.0745$).

What can explain earlier convergence in AsymC*?* All in all, there is some evidence that asymmetry fosters early coordination under complete information (while it leads to delays in "agreeing" under incomplete information). The first finding is somewhat unexpected. Which feature in

treatments, but recall that the set of actions is the same in all treatments.

[22] Regressing the period of convergence on treatment dummies in a GLS model delivers similar results with the exception that the comparison SymC vs AsymC becomes significant at the 5% level.

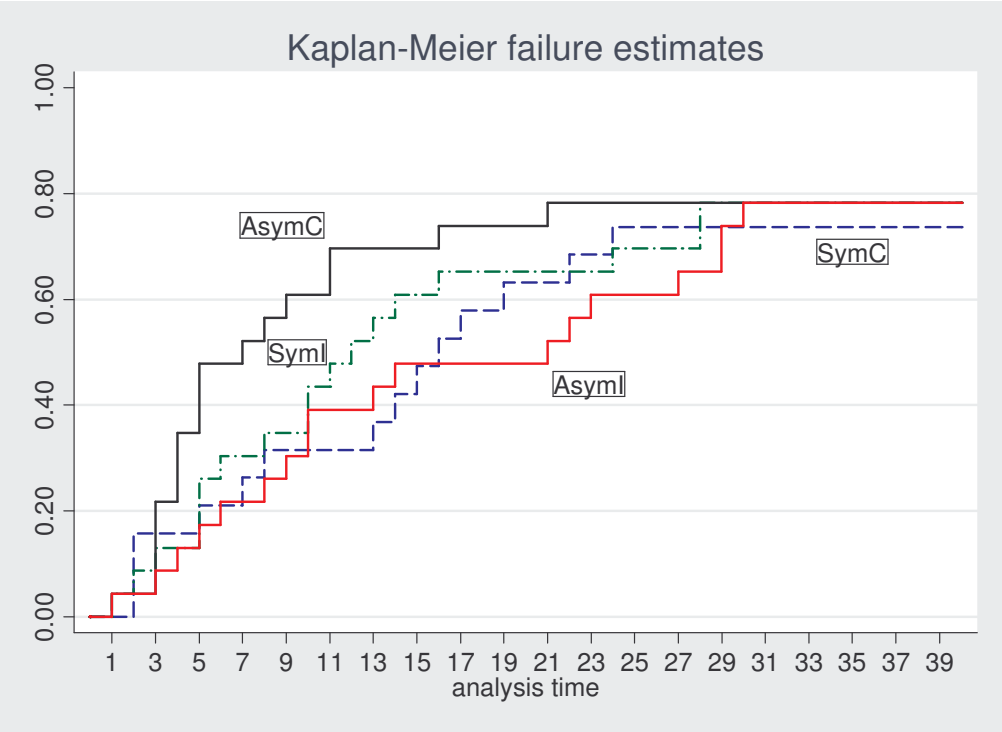

Figure 6: Kaplan-Meier failure (convergence) estimates

treatment AsymC could be conducive to early coordination when compared to treatment SymC? We venture that "leadership" may explain this pattern. Indeed, AsymC is the only treatment in which players know for sure that one of the two firms has got an advantage, which may influence the ability of this firm to "propose prices" in the adjustment process leading to stabilization. In support of this hypothesis, we first note that in treatment AsymC, a very high fraction of the posted prices happen to be 43, which is the price maximizing the profits of the low-cost firm, rather than 44 or 45. Second, we observe that the pattern of reactions to miscoordination systematically favors the low-cost firm, in the sense that it is the high-cost firm that more often adjusts its price in response to a gap in the posted prices. We elaborate on the latter observation and present relevant figures in Appendix B.

5 Summary and discussion

We find that in repeated Bertrand duopolies with convex costs average market prices are overall high, typically above the level of the highest Nash equilibrium. Regarding our treatment variables, we find that price levels in both symmetric and asymmetric markets are unaffected by information

differences about the rival's profit. However, under complete information, prices are on average higher in asymmetric markets (in spite of the arguably small cost differences between the two treatments), whereas under incomplete information, asymmetry does not make a difference. Furthermore, the lowest Nash equilibrium involving no loss in case of miscoordination, which is an attractor of play in symmetric markets, plays no role in asymmetric markets.

Neither information nor cost conditions seem to have an influence on the ability of subjects to coordinate eventually. About three fourths of all markets have converged to some common price after 30 periods of play. Besides, although asymmetric markets under complete information converge significantly more often to collusive prices than asymmetric markets under incomplete information, the difference in the prices to which markets converge on average is not statistically significant in any pairwise cross-treatment comparison. Hence, the higher market prices observed in the asymmetric treatment under complete information are attributable to pricing in the first periods of the game and to markets that do not converge.

In the end, the speed of convergence is the main variable affected by variations in the information and cost conditions, in a way that is not expected. In asymmetric markets, complete information induces early coordination while delays are observed under incomplete information. Symmetric markets, whose ability to stabilize is unaffected by information conditions, stand between those two extremes with respect to the time it takes to converge. We find some evidence that in asymmetric markets under complete information, low-cost firms are able to influence the price formation process to a greater extent, thus leading to early convergence.

These findings are not in line with those that have been reported about Cournot competition under similar designs. Mason and Phillips (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, in markets with complete information, we find the opposite.

In the same environment, Mason and Phillips (1997) tested the impact of information on duopoly outcomes and reported that "symmetric markets are more cooperative when profitability is common knowledge. Asymmetric markets are unaffected by information differences." Fouraker and Siegel (1963) had reached similar conclusions with regards to symmetric markets. Again, we find the opposite, as our symmetric markets are unaffected by changes in information, while asymmetric markets react strongly to completeness of information by achieving cooperation faster.

Among the classical questions asked in antitrust economics are: Should symmetric or asym-

metric market structures be favored (as part of, e.g., merger control when coordinated effects are assessed)? Should information exchange between firms about their cost structure be allowed (e.g. when a trade association offers to perform a benchmarking exercise)? (See Kühn and Vives (1995) and Vives (2006) for overviews of the latter question.) Our claim is that the answer to those questions depend on market institutions. In Bertrand markets, our findings do not accord well with widely-held views about symmetry and complete information as collusion-facilitating market characteristics, which seem to be mostly relevant for Cournot markets.[23] Hence, we believe that in assessing the likelihood of tacit collusion, closer attention should be paid to the mode of competition between the firms in the market at hand.

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] Alós-Ferrer, C, A. B. Ania, and K. R. Schenk-Hoppé (2000): An evolutionary model of Bertrand oligopoly, *Games and Economic Behavior*, 33, 1-19.

[3] Andersson, O. and E. Wengström (2007): Do antitrust laws facilitate collusion? Experimental evidence on costly communication in duopolies, *Scandinavian Journal of Economics*, 109(2), 312-339.

[4] Apesteguia, J., S. Huck, and J. Oechssler (2007): Imitation: Theory and experimental evidence, *Journal of Economic Theory*, 136, 217-235.

[5] 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.

[6] Argenton, C. and W. Müller (2009): Bertrand competition with convex costs in symmetric and asymmetric markets: Results from a pilot study, in: Wolfgang Franz, Werner Güth, Hans Jürgen Ramser and Manfred Stadler (eds.), *Experimentelle Wirtschaftsforschung*, Mohr Siebeck, Tübingen.

[23] We already knew that "Bertrand colludes more than Cournot" (Suetens and Potters, 2007), although the authors compared price outcomes in three (series of) experiments with *differentiated* Bertrand competition treatments and Cournot competition treatments.

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

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

[9] Bruttel, L.V. (2009): Group Dynamics in Experimental Studies - The Bertrand Paradox Revisited, *Journal of Economic Behavior and Organization*, 69, 51-63.

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

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

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

[13] Engel, C. (2007): How Much Collusion? A Meta-Analysis On Oligopoly Experiments, *Journal of Competition Law and Economics* 3, 491-549.

[14] Fatas, E., E. Haruvy, and A. J. Morales (2009): An experimental re-examination of the Bertrand paradox, mimeo.

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

[16] Haan, M..A., L. Schoonbeek, and B.M. Winkel (2009): Experimental results on collusion; in Jeroen Hinloppen and Hans-Theo Normann, *Experiments and competition policy*, New York: Cambridge University Press, 9-33.

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

[18] Kagel, J.H. (1995): Auctions: a survey of experimental research, in: John H. Kagel and Alvin A. Roth (eds.), *The handbook of experimental economics*, Princeton, NJ: Princeton University Press.

[19] Kaplan, E.L. and Meier, P. (1958): Nonparametric estimation from incomplete observations, *Journal of the American Statistical Association*, 53, 457-481.

[20] 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.

[21] Mailath, G.J. and L. Samuelson (2006). *Repeated games and reputations: Long-run relationships*, New York: Oxford University Press.

[22] 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.

[23] 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.

[24] Motta, M. (2004): *Competition policy: theory and practice*, Cambridge: Cambridge University Press.

[25] Selten, R. and R. Stoecker (1986): End behavior in finite prisonner's dilemma supergames, *Journal of Economic Behavior and Organization*, 1, 47-70.

[26] Suetens, S. and J. Potters (2007): Bertrand colludes more than Cournot, *Experimental Economics*, 10, 71-77.

[27] Vives, X. (1999): *Oligopoly Pricing: Old Ideas and New Tools*, Cambridge, Mass.: MIT Press.

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

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

[30] 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.

A Regression and test results for consumer and producer surplus

	Consumer Surplus			Producer Surplus		
	(CS1)	**(CS2)**	**(CS3)**	**(PS1)**	**(PS2)**	**(PS3)**
Rounds →	1-17	18-37	1-37	1-17	18-37	1-37
Constant	683.0637***	590.8985***	1415.723***	632.5129***	933.4687***	779.8765***
	(23.13)	(17.68)	(61.40)	(13.33)	(30.63)	(23.54)
SymI	-15.98112	-33.6193	41.92996	60.73585	77.77934*	61.86704
	(-0.40)	(-0.74)	(1.35)	(0.95)	(1.89)	(1.38)
AsymC	-106.1099***	-100.7274**	56.23723*	198.4769***	127.9323***	153.4718***
	(-2.66)	(-2.23)	(1.80)	(3.10)	(3.11)	(3.43)
AsymI	-59.26518	-51.79762	14.21312	90.28444	58.01539	59.74634
	(-1.49)	(-1.15)	(0.46)	(1.41)	(1.41)	(1.33)
N	1496	1760	3256	1496	1760	3256

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

Table 7: Regression results for consumer and producer surplus

B Price adjustment dynamics in treatments SymC and AsymC (not for publication)

To account for a possibly different adjustment behavior of asymmetric firms, we analyze the adjustment behavior of the two firms in AsymC separately. One way to analyze dynamics in the play of the repeated duopoly games is to study how a firm i adjusted its own price p_t^i in the current period, t , (in relation to its own price p_{t-1}^i in the previous period $t-1$) in response to the difference $p_{t-1}^i - p_{t-1}^{-i}$ between its own price, p_{t-1}^i, and the price of the other firm, p_{t-1}^{-i}, in the previous period. We distinguish three cases (i) $p_{t-1}^i - p_{t-1}^{-i} < 0$, (ii) $p_{t-1}^i - p_{t-1}^{-i} = 0$, and (iii) $p_{t-1}^i - p_{t-1}^{-i} > 0$, depending on whether firm i's price in the previous period was smaller, equal, or larger than the other firm's price, p_{t-1}^{-i}. Likewise, a firm's reaction to this price difference can either be to decrease, keep, or increase its own price in the current period relatively to its price in the previous period. Hence, we will distinguish between the three cases (i) $p_t^i - p_{t-1}^i < 0$, (ii) $p_t^i - p_{t-1}^i = 0$, and (iii)

(CS1)	SymC	SymI	AsymC
SymI	0.689	—	—
AsymC	**0.008**	**0.0176**	—
AsymI	0.137	0.2541	0.2171

(PS1)	SymC	SymI	AsymC
SymI	0.344	—	—
AsymC	**0.002**	**0.0239**	—
AsymI	0.159	0.6281	**0.0761**

(CS2)	SymC	SymI	AsymC
SymI	0.457	—	—
AsymC	**0.026**	0.1182	—
AsymI	0.251	0.6721	0.2546

(PS2)	SymC	SymI	AsymC
SymI	**0.059**	—	—
AsymC	**0.002**	0.2004	—
AsymI	0.159	0.6139	**0.0743**

(CS3)	SymC	SymI	AsymC
SymI	0.178	—	—
AsymC	**0.071**	0.6293	—
AsymI	0.648	0.3497	0.1562

(PS3)	SymC	SymI	AsymC
SymI	0.167	—	—
AsymC	**0.001**	**0.0314**	—
AsymI	0.182	0.9603	**0.0277**

Note: Statistically significant results are highlighted in bold font.

Table 8: p-values of (two-tailed) pairwise cross-treatment differences w.r.t. to estimates reported in Table 7

$p_t^i - p_{t-1}^i > 0$. Table 9 shows cross tables of the two variables $p_{t-1}^i - p_{t-1}^{-i}$ and $p_t^i - p_{t-1}^i$ for the symmetric firms in treatment SymC (top) and the two asymmetric firms in treatment AsymC (bottom). Recall that in treatment SymC markets stabilized on average in period 11.9 whereas markets converged on average in period 6.8 in treatment AsymC. In order to capture the average time interval needed for markets to stabilize, in Table 9 we only include data from period 1-12 for treatment SymC and period 1-7 for treatment AsymC. Moreover, we report percentages for ease of comparison. Comparing adjustments made in the two treatments and conditional on the sign of $p_{t-1}^i - p_{t-1}^{-i}$, we make the following observations:

- $p_{t-1}^i - p_{t-1}^{-i} < 0$: In this case, high-cost firms in AsymC increase their own price in 73% of the cases, while firms in SymC do so in only 52.8% of the cases. Hence, it appears as if high-cost firms in AsymC that are lagging behind "catch up" more often with the other firm in the market than firms in treatment SymC that are in the same situation.[24]

- $p_{t-1}^i - p_{t-1}^{-i} = 0$: In this case, low-cost firms in AsymC do not change their own price in 93.8% of the cases, whereas this happens in only 81.1% of the cases in SymC. Hence, it appears as

[24] A χ^2-test reveals that the adjustment patterns of symmetric firms in SymC and high-cost firms in AsymC in case of $p_{t-1}^i - p_{t-1}^{-i} < 0$ are statistically different at the 10% level, while adjustment patterns of symmetric firms in SymC and low-cost firms in AsymC are not statistically different.

Treatment SymC

		$p_t^i - p_{t-1}^i$			
		< 0	$= 0$	> 0	Total
$p_{t-1}^i - p_{t-1}^{-i}$	< 0	12.6	34.6	52.8	100
	$= 0$	4.9	81.1	14.0	100
	> 0	59.8	33.1	7.1	100
	Total	23.9	52.4	23.7	

Treatment AsymC

Low-cost firm

		$p_t^i - p_{t-1}^i$			
		< 0	$= 0$	> 0	Total
$p_{t-1}^i - p_{t-1}^{-i}$	< 0	7.5	35.8	56.6	100
	$= 0$	4.2	93.8	2.1	100
	> 0	32.4	62.2	5.4	100
	Total	13.0	63.0	23.9	

High-cost firm

		$p_t^i - p_{t-1}^i$			
		< 0	$= 0$	$>$	Total
	< 0	5.4	21.6	73.0	100
	$= 0$	2.1	89.6	8.3	100
	$>$	54.7	32.1	13.2	100
	Total	23.2	49.3	27.5	

Note: All frequencies are expressed in percentages.

Table 9: Adjustment dynamics of symmetric firms in treatment SymC (top) and of low- and high-cost firms in treatment AsymC (bottom)

if low-cost firms in AsymC keep play stable more often than firms in treatment SymC that are in the same situation.[25]

- $p^i_{t-1} - p^{-i}_{t-1} > 0$: In this case, low-cost firms in AsymC do not change their own price in 62.2% of the cases, whereas this occurs in only 33.1% of the cases in SymC. Hence, it appears as if low-cost firms in AsymC that are ahead, more often give the other firm in the market the chance to catch up than firms in treatment SymC that are in the same situation.[26]

These observations are consistent with the hypothesis that low-cost firms act as price "leaders" in the asymmetric treatment. Upon miscoordination, they stick to their choice more often and the burden of the adjustment falls on the high-cost firms.

C Instructions (not for publication)

C.1 Symmetric treatments, complete information

INSTRUCTIONS

Welcome to this experiment!

Please read these instructions carefully! Do not speak to your neighbours and keep quiet during the entire experiment! If you have a question, please raise your hand. We will then come to your seat.

In this experiment you will repeatedly make decisions. By doing so you can earn money. How much you earn depends on your decisions and on the decisions of another participant in the experiment. All participants receive the same instructions.

YOUR TASK IN THE EXPERIMENT

In this experiment, you represent a firm which, along with one other firm, produces and sells a fictitious product in a market. In each of the 40 rounds of this experiment, you and the other firm will always have to make one decision, namely, to set the price at which you are willing

[25] A χ^2-test reveals that the adjustment patterns of symmetric firms in SymC and low-cost firms in AsymC in case of $p^i_{t-1} - p^{-i}_{t-1} = 0$ are statistically different at the 10% level, while adjustment patterns of symmetric firms in SymC and high-cost firms in AsymC in this case are not statistically different.

[26] A χ^2-test reveals that the adjustment patterns of symmetric firms in SymC and low-cost firms in AsymC in case of $p^i_{t-1} - p^{-i}_{t-1} > 0$ are statistically different at the 1% level, while adjustment patterns of symmetric firms in SymC and high-cost firms in AsymC in this case are not statistically different.

to sell the fictitious product. Prices can be chosen from the set {10, 11, 12, ..., 50}. That is, all integer numbers from 10 to 50 are possible choices.

YOUR PROFIT

The profits are denoted in a fictitious unit of money which we call "Points". Negative numbers stand for losses.

In the attached table you can see the profits (or losses) that you will make depending on the prices chosen by yourself and the other firm in your market. The participant who represents the other firm in your market has a profit table that is identical to the one you have.

Down the first column of the table are listed the prices that you may choose in any given round. Columns 2, 3, and 4 show your profit depending on the prices chosen by yourself and the other firm in your market in the three cases that can be distinguished.

- If the price that you have chosen is the lowest price of all prices chosen in your market, you will receive the profit shown in Column 2 entitled "Your profit when you have the lowest price."
- If the price that you have chosen is the same as the price chosen by the other firm in your market, you will receive the profit shown in Column 3 entitled "Your profit when you are tied for the lowest price."
- If the price that you have chosen is higher than the price of the other firm in your market, you will receive the profit shown in Column 4 entitled "Your profit when you don't have the lowest price."

MATCHING

The experiment consists of 40 decision rounds. In all rounds, you will interact with the same participant, who will be randomly selected at the beginning of the experiment. The identity of this participant will never be revealed to you.

FEEDBACK

At the end of each round, you will learn the price chosen by the other firm in your market and your own profit (or loss).

YOUR MONETARY EARNINGS

You will start the experiment with an initial capital of 5000 Points.

At the end of the experiment, your monetary earnings will be determined by the sum of your initial capital and your profits (or losses) in all rounds. You will receive 1 Euro for every 1800 Points you have accumulated.

C.2 Asymmetric treatment, incomplete information

INSTRUCTIONS

Welcome to this experiment!

Please read these instructions carefully! Do not speak to your neighbours and keep quiet during the entire experiment! If you have a question, please raise your hand. We will then come to your seat.

In this experiment you will repeatedly make decisions. By doing so you can earn money. How much you earn depends on your decisions and on the decisions of another participant in the experiment. All participants receive the same instructions.

YOUR TASK IN THE EXPERIMENT

In this experiment, you represent a firm which, along with one other firm, produces and sells a fictitious product in a market. In each of the 40 rounds of this experiment, you and the other firm will always have to make one decision, namely, to set the price at which you are willing to sell the fictitious product. Prices can be chosen from the set {10, 11, 12, ..., 50}. That is, all integer numbers from 10 to 50 are possible choices.

YOUR PROFIT

The profits are denoted in a fictitious unit of money which we call "Points". Negative numbers stand for losses.

In the attached table you can see the profits (or losses) that you will make depending on the prices chosen by yourself and the other firm in your market. The participant who represents the other firm in your market has a profit table that may or may not be identical to the one you have.

Down the first column of the table are listed the prices that you may choose in any given round. Columns 2, 3, and 4 show your profit depending on the prices chosen by yourself and the other firm in your market in the three cases that can be distinguished.

- If the price that you have chosen is the lowest price of all prices chosen in your market, you will receive the profit shown in Column 2 entitled "Your profit when you have the lowest price."
- If the price that you have chosen is the same as the price chosen by the other firm in your market, you will receive the profit shown in Column 3 entitled "Your profit when you are tied for the lowest price."
- If the price that you have chosen is higher than the price of the other firm in your

market, you will receive the profit shown in Column 4 entitled "Your profit when you don't have the lowest price."

MATCHING

The experiment consists of 40 decision rounds. In all rounds, you will interact with the same participant, who will be randomly selected at the beginning of the experiment. The identity of this participant will never be revealed to you.

FEEDBACK

At the end of each round, you will learn the price chosen by the other firm in your market and your own profit (or loss).

YOUR MONETARY EARNINGS

You will start the experiment with an initial capital of 5000 Points.

At the end of the experiment, your monetary earnings will be determined by the sum of your initial capital and your profits (or losses) in all rounds. You will receive 1 Euro for every 1800 Points you have accumulated.