International Journal of Industrial Organization 30 (2012) 508–517

Contents lists available at SciVerse ScienceDirect

International Journal of Industrial Organization

journal homepage: www.elsevier.com/locate/ijio

Collusion in experimental Bertrand duopolies with convex costs: The role of cost asymmetry[☆]

Cédric Argenton[a,*], Wieland Müller[a,b,1]

[a] CentER & TILEC, Tilburg University, The Netherlands
[b] VCEE, University of Vienna, Austria

ARTICLE INFO

Article history:
Received 20 May 2010
Received in revised form 27 March 2012
Accepted 30 May 2012
Available online 19 June 2012

JEL Classification:
L13
C72
C92

Keywords:
Bertrand competition
Convex costs
Collusion
Coordination
Experimental economics

ABSTRACT

Theory, experimental studies, as well as antitrust guidelines suggest that symmetry among firms is conducive to more collusive outcomes. We test this perception in a series of experimental repeated Bertrand duopolies where firms have convex costs. We implement symmetric as well as asymmetric markets that vary in their degree of cost asymmetry among firms. We find no evidence of symmetric markets being more prone to collusion than asymmetric markets. If anything, asymmetry *helps* firms coordinate on higher prices and achieve higher profits.

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, is a primary focus of economics. In market environments, the possibility for firms to collude and sustain high prices instead of harshly competing is not only of high theoretical interest but also of high policy relevance, as antitrust authorities throughout the world have come to see the fight against collusion as one of their main tasks. There is a large body of literature dealing with tacit collusion.[2] The general view about collusion-facilitating market characteristics is that firm symmetry (in product ranges, costs, demands, etc.) makes collusion easier. For instance, in his classical textbook, Scherer (1980; p. 205) writes "…the more cost functions differ from firm to firm, the more trouble firms will have maintaining a common price policy…" Motta (2004; p. 143) notes that "the more firms are asymmetric (in capacities, market shares, costs or product range) the less likely collusion will be" and proceeds to show in a technical section (4.2.5) that "symmetry helps collusion." Similarly, in their report to the European Commission, Ivaldi et al. (2003) elaborate on the finding that "cost asymmetries hinder collusion" (section III.9).

This consensus among economists has influenced policy-making. The US 'horizontal merger guidelines' state at section 2.11 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 (…)."[3] Similarly, the European Commission's 'horizontal merger guidelines' state at recital 48 that "[f]irms may find it easier to reach a common understanding on the terms of coordination if they are relatively symmetric, *especially in terms of cost structures*, market shares, capacity levels and levels of vertical integration."[4] As a matter

☆ We thank co-editor Yossi Spiegel and two anonymous referees for many constructive comments that greatly improved the paper. We are grateful to Ola Andersson, Charles Noussair, Ronald Peeters, Patrick Rey, Dirk Sliwka and Bert Willems as well as seminar and conference participants at Tilburg University, University of Amsterdam and NAKE for their helpful comments. Wieland Müller acknowledges financial support from the Netherlands Organisation for Scientific Research (NWO) through a VIDI grant.
* Corresponding author at: Department of Economics, Tilburg University, Warandelaan 2, Postbus 90153, 5000 LE Tilburg, The Netherlands.
E-mail addresses: c.argenton@uvt.nl (C. Argenton), wieland.mueller@univie.ac.at (W. Müller).
[1] Department of Economics, University of Vienna, Brünnerstrasse 72, 1210 Vienna, Austria.
[2] Ivaldi et al. (2003) summarize the relevant theoretical literature. For a state-of-the-art review of the theory of repeated games, see Mailath and Samuelson (2006). For a recent survey of experimental results on collusion, see Haan et al. (2009). A meta-study of experiments on collusion is performed by Engel (2007).

[3] Our italics. The guidelines can be found at http://www.usdoj.gov/atr/public/guidelines/horiz_book/toc.html (20 August 2011). In the following quotations "reaching terms of coordination" is to be taken as official parlance for "tacit collusion."
[4] Our italics. The guidelines can be found at <http://ec.europa.eu/competition/mergers/legislation/notices_on_substance.html> (20 August 2011).

0167-7187/$ – see front matter © 2012 Elsevier B.V. All rights reserved.
doi:10.1016/j.ijindorg.2012.05.006

of fact, the decision practice of the European Commission is remarkably in line with this guidance. Davies et al. (2011) show that in merger cases, concerns regarding 'coordinated effects' are raised only when the post-merger market structure is predicted to be a *symmetric* duopoly.

This general agreement on the role of cost asymmetry finds some support in the theoretical literature. A number of articles have argued that asymmetry makes tacit collusion harder in repeated games, in the sense of raising the critical discount rate needed for collusion to be sustainable as a subgame-perfect equilibrium. Those contributions typically argue that either the deviation profit or the punishment profit of a lower-cost firm is higher, in other terms, that such a firm has higher short-term incentives to cheat on a cartel agreement or that other cartel members have a reduced ability to retaliate. Bae (1987) and Harrington (1991) study repeated Bertrand competition under constant returns to scale. Rothschild (1999) does the same for Cournot competition and Davidson and Deneckere (1990) for Bertrand–Edgeworth competition. Those papers appeal to grim trigger strategies. Vasconcelos (2005) introduces harsher punishments and looks at a larger class of equilibria in the Cournot case; so do Compte et al. (2002) for Bertrand–Edgeworth competition. Miklós-Thal (2011) recently provided a full treatment of repeated Bertrand competition under different but constant unit costs by using optimal punishments à la Abreu (1986, 1988). The analysis is affected by the use of harsher or optimal punishments (and by the mode of competition) but the general conclusion remains that (efficient) collusion is harder to sustain under cost asymmetry.

The experimental results that are available are in line with this assertion. Mason et al. (1992) concern themselves with the impact of asymmetry on Cournot duopolists in a finitely repeated-game environment (fixed matching) and report that "asymmetric markets are less cooperative and take longer to reach equilibrium than symmetric markets."[5] Fonseca and Normann (2008) study experimental Bertrand–Edgeworth oligopolies and find that for a given number of firms (two or three) asymmetric markets exhibit lower prices than symmetric markets (under a fixed-matching, random-termination protocol). Keser (2000) investigates experimental price-setting duopolies with demand inertia, as introduced by Selten (1965), in environments where (constant) unit costs are symmetric. There are considerable variations in prices over time within markets, as well as across markets, but by comparing results to those of Keser (1993), which involved asymmetric costs, the author concludes that in the latter environment "we observed […] a significantly lesser degree of cooperation than in the symmetric cost situation."

The canonical model of Bertrand competition with constant unit costs has only recently been the subject of more intense experimental investigations.[6] Dugar and Mitra (2009) study the impact of asymmetry in constant unit costs on prices in experimental Bertrand duopolies under fixed matching and random assignment of roles across periods. They report that symmetric markets achieve higher prices than asymmetric markets.[7]

In this paper, we experimentally investigate the extent of tacit collusion in homogenous-product Bertrand duopolies under *convex* costs. Those markets are interesting from both a theoretical and a practical point of view. Theoretically, such games typically admit a whole interval of strict, Pareto-ranked Nash equilibria which may or may not contain the joint profit-maximizing strategy profile (Dastidar, 1995). When they don't, as in our experiment, there is a conflict between short-term private incentives and the joint interest of players but, contrary to standard prisoner's dilemma supergames (or, say, Bertrand competition under linear costs), the design of punishments or the implementation of trigger strategies in the repeated game is complicated by the multiplicity of equilibria in the one-shot game.

A key feature of Bertrand competition is the obligation for firms to serve all demand addressed to them at their posted price, even if rationing were more profitable.[8] Obviously, this characteristic makes Bertrand competition special, and different from most simple posted-offer markets where sellers have a fixed supply. However, competition in certain sectors can be stylized as Bertrand pricing under convex costs. For instance, utilities such as gas, water and electricity providers face rising marginal costs and are typically under the (legal or technical) obligation to adjust their supply to customers' demand. The same is true of businesses which operate under a subscription system, e.g., telecom services. In many countries, such markets are or have been dominated by two main firms.[9] More generally, "the Bertrand assumption is plausible when there are large costs of turning customers away" (Vives, 1999, p.118).[10] Hence, notwithstanding the usual and legitimate concerns about external validity, our experimental investigation can shed light on a common market structure, in which the strategic incentives are somewhat more complicated than usual.

In our experiment, we examine whether repeatedly interacting duopolists are able to coordinate tacitly on high prices under varying cost conditions. More precisely, we analyze price choices in three fixed-matching treatments featuring cost symmetry, a small asymmetry and a larger asymmetry. We want to know what the impact of differences in costs is on market prices, and more generally on the ability of firms to coordinate. Our results do not confirm the general view. Indeed, we find no evidence of symmetric markets being more collusive than asymmetric markets. With regards to some of our collusion measures, we even find that differences in costs actually *help* firms coordinate and come closer to cartel profits.

We are aware of only two other papers experimentally implementing Bertrand competition under convex costs. Abbink and Brandts (2008) study the impact of the number of firms (2, 3, or 4) on (long-run) outcomes when firms are symmetric. Moreover, one of the markets analyzed in Fatas et al. (2009) is a symmetric Bertrand duopoly market with quadratic costs and inelastic demand. To our knowledge, we are the first to investigate experimentally the effects of cost asymmetry under this market structure.

The paper is organized as follows. Section 2 describes our experimental design as well as the theoretical predictions. We report the experimental results in Section 3. Section 4 discusses the findings and concludes.

[5] See Mason and Phillips (1997) for a confirmation and the introduction of incomplete information (about the payoff structure) in this setting.

[6] See Fouraker and Siegel (1963), Dufwenberg and Gneezy (2000, 2002), and Bruttel (2008) for treatments of the symmetric cost case.

[7] However, that is also true of experimental designs that are more conducive to one-shot play. Indeed, Dugar and Mitra (2011) use a protocol involving random matching of subjects (but fixed assignment of roles), and reach the same conclusion. Boone et al. (2012) also analyze homogenous-product Bertrand markets in a within-subject random-matching design where firms have constant but different marginal costs. They find that while market prices converge to the Nash prediction in two of their three treatments where two or three firms all have different costs, market prices stay above the predicted level in the condition where two firms have the same low unit cost and a third firm has a higher one.

[8] This is not an issue under constant returns to scale. By contrast, under convex costs firms might find it suboptimal to produce large quantities.

[9] 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).

[10] Dixon (1990) shows that the combination of price competition with explicitly modelled costs of turning customers away delivers a range of Nash equilibrium prices. Thus, Dastidar's (1995) model can be seen as a reduced-form version of a more complicated game.

Table 1
Design details and static Nash equilibria.

Treatment name	"SYM"	"ASYM-L"	"ASYM-H"
Cost parameters	$c_1 = c_2 = 0.6$	$c_1 = 0.55$, $c_2 = 0.65$	$c_1 = 0.5$, $c_2 = 0.7$
Set of Nash equilibria	{21, 22, …, 39}	{22, 23, …, 38}	{23, 24, …, 36}
Number of subjects	$19 \times 2 = 38$	$23 \times 2 = 46$	$21 \times 2 = 42$

2. Experimental design and theoretical predictions

2.1. Experimental design

Since our main interest lies in the "comparative statics of collusion" when cost conditions vary, our concern in designing the experiment was to generate enough collusion in the first place. For this reason, we focussed attention on duopoly markets. Previous experiments have indeed shown that tacit collusion is rarely observed in markets with more than two firms.[11] 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. The experiment was described to the participants as a pricing game between firms but they were not given the details of the model. Instead, they were presented with payoff table(s). The use of payoff tables is common practice and can be traced back to as early as Fouraker and Siegel (1963).

Subjects were paired in one of three treatments, which varied with respect to the cost structure, and thus payoffs. In the treatment we call "SYM" the profit tables of the two paired subjects were identical. In treatments "ASYM-L" and "ASYM-H" they were different. L (H) stands for a low (high) degree of asymmetry regarding firms' costs.

More precisely, the payoff tables were generated from a linear demand curve ($D(p) = 100 - 1.5p$) and quadratic cost curves ($C(q) = c_i q^2$), with all numbers rounded to integer values. 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 the symmetric treatment, cost functions were identical with $c_1 = c_2 = c_0 = 0.6$. In the asymmetric treatment "ASYM-L", 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 ($c_2 = 0.65$). In the asymmetric treatment "ASYM-H", the low-cost parameters was $c_1 = 0.5$ and the high-cost parameter was $c_2 = 0.7$. This "symmetric-spread" design was chosen because, conditional on both firms charging the same price, total costs, joint profits and thus total welfare are the same in all treatments, which facilitates comparisons. In particular, the joint profit-maximizing (cartel) price is the same across symmetric and asymmetric treatments.

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.[12] 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 order not to foster imitation.[13]

In all treatments, payoffs were expressed in a fictitious monetary unit ("points"). Subjects were told that negative numbers stood for losses, which were indeed possible in the range of low prices. They started the experiment with an initial capital of 5000 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.[14] One Euro was exchanged for every 1800 points accumulated. Each treatment lasted between 30 and 45 min. The average monetary earnings across all treatments were €12.96.

All subjects were electronically recruited from the pool of participants registered with Tilburg University's CentERlab. At the time of the experiment, all were students enrolled in various programs 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).[15] Instructions were read, questions were taken and answered, after which the experiment started. Each participant took part in only one session. We analyze data from 8 lab sessions (3 for SYM, 3 for ASYM-L, 2 for ASYM-H). We have data on 19 pairs in treatment SYM, 23 pairs in treatment ASYM-L, and 21 pairs in treatment ASYM-H. Table 1 summarizes the design.

2.2. Theoretical predictions

Bertrand competition is not synonymous with perfect competition when firms face convex costs. In the symmetric case there is a whole interval of pure-strategy Nash equilibrium prices. The lower bound of this interval is determined by average-cost pricing. The upper bound is determined by the incentive to marginally undercut competitors. The interval contains the competitive price (which involves marginal-cost pricing). It may contain the price that maximizes 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. All general claims are proved by Dastidar (1995).[16]

We illustrate the intuition with Fig. 1. For the parameters of our symmetric treatment, it displays monopoly profits as a function of price (solid curve) as well as duopoly profits when firms charge the same price (dashed curve). Because of convexity in costs, industry profits are higher when production is split between two firms which then face lower marginal costs (compare monopoly profits with twice the duopoly profits). In the absence of fixed costs, a Nash equilibrium must be such that duopolists make nonnegative profits at the equilibrium price. Hence, only prices that lie to the right of the left-most vertical line are admissible. Incentives must also be such that one of the players does not want to undercut the other so as to reap higher monopoly profits. Hence, only prices that lie to the left of the right-most vertical line are admissible. Therefore, there is a whole interval of equilibrium prices between the two vertical lines. Note that the price that maximizes duopoly profits lies outside that interval.

Table 2 reproduces the payoff table we used in the symmetric treatment. As can be checked, all prices in {21, 22, …, 39} are Bertrand equilibria. The lowest Nash equilibrium price, 21, involves an

[11] This is a general conclusion drawn by Haan et al. (2009) or Engel (2007). See Abbink and Brandts (2008) for Bertrand competition under convex costs. For the case of Cournot markets, see Huck et al. (2004).

[12] It is known that collusion is not to be expected under random matching. See, for instance, Kübler and Müller (2002) in the case of Bertrand competition with differentiated products.

[13] Several papers have shown that imitation leads to competitive behavior in many market games and that the observation of rivals' payoffs is conducive to such imitation. See, e.g., Altavilla et al. (2006), Apesteguia et al. (2007), Huck et al. (1999), or Offerman et al. (2002).

[14] As expected, no participant depleted his or her entire capital at any point during the experiment.

[15] The instructions can be found in the Appendix A.

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

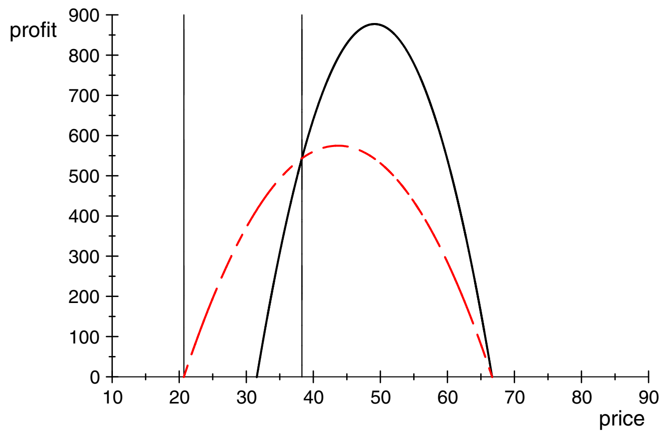

Fig. 1. Monopoly (solid curve) and duopoly profits (dashed curve) as a function of the (common) price.

equilibrium profit of 15 but a loss of 1377 in case of miscoordination. By contrast, the payoff-dominant equilibrium price, 39, involves 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.[17] 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.

In case of cost asymmetry, there is an equivalent to Fig. 1 for each firm. The highest Bertrand price is determined by the incentives for the lower-cost firm to undercut (because this firms is comparatively better at undercutting and producing large quantities). Conversely, the lowest Bertrand price is determined by the zero-profit condition for the higher-cost firm (because as price decreases, the profit of that firm turns negative first). In our first asymmetric treatment ASYM-L, the range of Bertrand equilibria ran from 22 to 38. The lowest equilibrium price involving no loss to either firm in case of miscoordination was 33. In our asymmetric treatment ASYM-H, the range of Bertrand equilibria ran from 23 to 36. The lowest equilibrium price involving no loss to either firm in case of miscoordination was 35. Table 1 summarizes those Nash predictions.

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 a number of Bertrand equilibrium prices at the bottom. Because of cost convexity, the potential losses from miscoordination keep on increasing with the size of demand so that low Bertrand prices are in this 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. In all treatments, the price that maximizes joint profits is 44 but because of cost differences, firms' interests are no longer perfectly aligned in asymmetric treatments. Conditional on both firms charging the same price, the profit to the low-cost firm is maximized at a price of 43, while the profit to the high-cost firm is maximized at a price of 45 in treatment ASYM-L. The numbers are 42 and 45, respectively, in treatment ASYM-H.

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

Table 2
Payoff table for symmetric treatments.

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[a]	−1377	15	0
22[a]	−1219	64	0
23[a]	−1068	110	0
24[a]	−922	154	0
25[a]	−781	195	0
26[a]	−647	235	0
27[a]	−518	272	0
28[a]	−394	307	0
29[a]	−277	340	0
30[a]	−165	371	0
31[a]	−59	400	0
32[a]	42	426	0
33[a]	136	451	0
34[a]	225	473	0
35[a]	309	493	0
36[a]	386	511	0
37[a]	458	526	0
38[a]	525	540	0
39[a]	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

[a] Static Nash equilibrium, **Perfect collusion**.

There is no unequivocal theoretical prediction for the outcome of play in such games. Shared expectations and common knowledge of rationality can give rise to the play of any Nash equilibrium in a one-shot context. Payoff dominance calls for the highest Nash equilibrium price to be played. Models of imitation (Abbink and Brandts, 2008 or Alós-Ferrer et al., 2000) predict convergence towards the competitive equilibrium under some conditions. We did not include the profit to the other firm in the feedback information received by participants. So, we did not expect them to follow that line of reasoning.

In any case, in our experimental setting the stage 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 is well-known from Benoît and Krishna (1985). There are of course a multiplicity of subgame-perfect equilibria.

Nonetheless, we now explain how the general view regarding cost asymmetry applies to our setting. To simplify the exposition, we make as if the action space of each firm were a continuum $[0, a]$ (where a is the price at which the demand curve intersects the vertical axis), and as if the interaction were perceived as infinitely

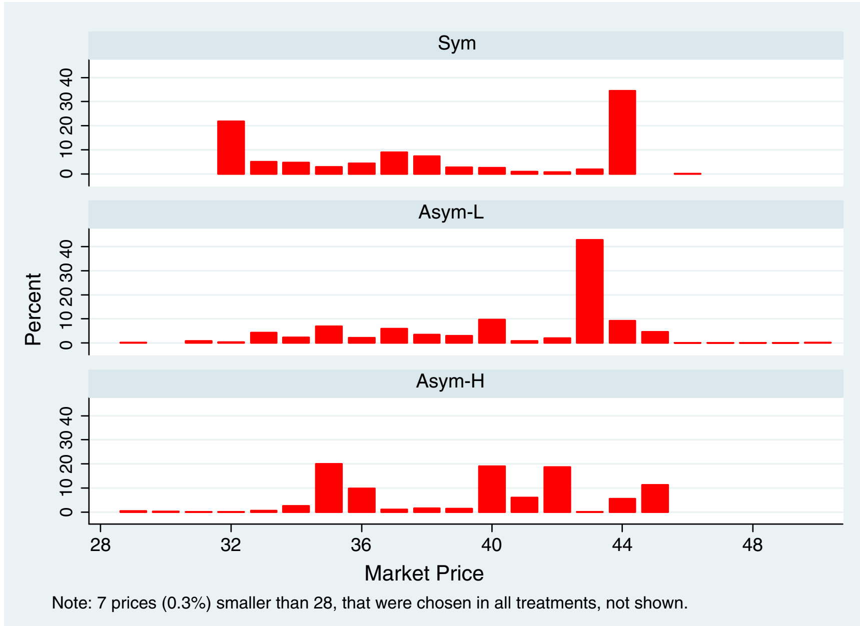

Fig. 2. Distribution of market prices in periods 1–37.

repeated under discount rates $\delta_i<1$.[18] In what follows, firm 0 will stand for the representative symmetric cost firm. Firm 1 will have a lower-cost parameter, while firm 2 will have a higher-cost parameter. That is, $c_1<c_0<c_2$. Symmetric treatments involve two type 0 firms, while asymmetric treatments oppose a type 1 firm to a type 2 firm. Notationwise, let $\pi_i^{(n)}(p)\equiv p\frac{D(p)}{n}-C_i\left(\frac{D(p)}{n}\right)$ stand for the one-period profit accruing to firm $i\in\{0,\ 1,\ 2\}$ when n firms charge the same price p (and any other firm charges a higher price). Thus, $\pi_i^{(1)}$ is firm i's monopoly profit while $\pi_i^{(2)}$ is her profit when the two firms share the market.

Suppose that firm 1's cost-advantage is not too drastic, so that any price $p\in[\dot{p},\ddot{p}]$ is a Nash equilibrium of the stage game under asymmetry, where $\dot{p}$ is defined by $\pi_2^{(2)}(\dot{p})=0$ and $\ddot{p}$ is defined by $\pi_1^{(2)}(\ddot{p})=\pi_1^{(1)}(\ddot{p})$. It is easy to show that for either firm 1 or firm 2, the minmax payoff in the stage game is 0. (If the price charged by i is high enough, then j will either charge the same price or undercut, and thus make positive profits. If the price charged by i is instead low, then j will charge any higher price, avoid making any sale, and earn zero profit.) For the construction of subgame-perfect equilibria, a crucial question is: can firms be "minmaxed" as part of a punishment equilibrium? The answer is straightforward for firm 2. Following any deviation by the latter, we can prescribe firms to play $\dot{p}$ forever, which by definition of $\dot{p}$ guarantees her zero profit and constitutes a subgame-perfect equilibrium path. For firm 1, things are slightly more complicated. Consider the following punishment strategies. In the period immediately following a deviation by firm 1, that firm prices at $\tilde{p}$ defined by

$$\pi_1^{(1)}(\tilde{p})=-\frac{\delta}{1-\delta}\pi_1^{(2)}(\dot{p}),$$

while firm 2 charges any strictly higher price. In all remaining periods, both firms charge $\dot{p}$. That is, firm 1 is forced to make a loss in the first period by serving the whole market at a low price. That loss is then exactly recovered in all subsequent periods. (As firm 1 has lower costs than firm 2, $\pi_1^{(2)}(\dot{p})>0$.)[19] In the punishment phase, any unilateral deviation is prescribed to lead to the (new) start of the (relevant) punishment. One therefore sees that both firms can be brought down to zero profits as part of a subgame-perfect punishment phase. Hence, cost asymmetry does not necessarily weaken retaliation possibilities if firms use optimal punishments, a point stressed by Miklós-Thal (2011). However, the deviation incentives (the extra payoff a firm makes in the period where she deviates and undercuts the other firm) are always more than proportionally higher for firm 1, because of cost convexity. That is,

$$\pi_1^{(1)}(p)/\pi_1^{(2)}(p)>\pi_0^{(1)}(p)/\pi_0^{(2)}(p)>\pi_2^{(1)}(p)/\pi_2^{(2)}(p).$$

Given the availability of minmax punishments, the condition for a price p to be sustained in equilibrium can then be written as:

$$\frac{1}{1-\delta_i}\pi_i^{(2)}(p)\geq\pi_i^{(1)}(p)+\delta_i\cdot 0.$$

The left-hand side stands for the discounted payoff along the equilibrium path while the right-hand side stands for the discounted sum of the deviation payoff and the punishment payoff. The conclusion that $\delta_1>\delta_0>\delta_2$ then immediately follows. Hence, the minimum discount rate required for the maintenance of collusion is higher in asymmetric treatments (δ_1) than in symmetric treatments (δ_0).

Therefore, on the basis of the existing literature on collusion and the general view about the role of cost asymmetry, as applied to our Bertrand supergame with convex costs, we expect that players in symmetric markets will be able to coordinate more easily on prices that are closer to the collusive level than players in asymmetric markets. In the next section we will define and analyze various measures of collusion. These measures will have the same property as market prices, namely, that higher values correspond to higher levels of collusion. Hence, the main hypothesis which we formulate and test in this study can be summarized (somewhat vaguely, at this stage) as follows:

[18] Note that in the experimental economics literature it is known that play in finitely repeated interactions might be more cooperative even if the stage-game equilibrium is unique (see, e.g., Selten and Stoecker, 1986, or Andreoni and Miller, 1993). This result is not inconsistent with the idea that subjects in finitely repeated games behave as if they perceived the time horizon as indefinite.

[19] If firm 1's cost advantage is strong, a price $\tilde{p}$ so defined may not exist but it is easy to go for punishments with a 'stick' phase of more than one period: find $\tilde{p}$ and $T\in\mathbb{N}^*$ such that $\pi_1^{(1)}(\tilde{p})+\delta\pi_1^{(1)}(\tilde{p})+\ldots\delta^{T-1}\pi_1^{(1)}(\tilde{p})=-\frac{\delta^T}{1-\delta}\pi_1^{(2)}(\dot{p})$, which is always possible.

Table 3
Collusion and coordination measures and test results. (Data from all markets.).

Collusion and coordination measures

	Market price			Supra Nash Price Count[a]			Supra Price Nash Index[b]			Collusion Index[c]			Price Coordination Count[d]		
	SymC	Asym-L	Asym-H	SymC	Asym-L	Asym-H	SymC	Asym-L	Asym-H	SymC	Asym-L	Asym-H	SymC	Asym-L	Asym-H
Rounds	37.3	40.0	39.2	6.0	11.4	11.5	−0.042	0.052	0.090	−9.200	0.416	0.719	7.8	11.2	10.4
1–17	(0.84)	(0.65)	(0.66)	(1.38)	(1.28)	(1.25)	(0.02)	(0.02)	(0.02)	(1.26)	(0.05)	(0.05)	(1.28)	(0.92)	(1.06)
Rounds	39.0	41.2	39.5	9.3	15.6	12.6	0.0008	0.085	0.097	−3.392	0.428	0.874	16.5	17.9	17.8
18–37	(1.07)	(0.63)	(0.76)	(2.13)	(1.65)	(2.01)	(0.03)	(0.02)	(0.02)	(1.14)	(0.06)	(0.04)	(1.33)	(0.94)	(1.07)
Rounds	38.3	40.6	39.4	15.3	27.0	24.1	−0.0190	0.070	0.094	−6.061	0.432	0.803	24.3	29.1	28.1
1–37	(0.88)	(0.59)	(0.70)	(3.22)	(2.74)	(3.12)	(0.02)	(0.02)	(0.02)	(1.06)	(0.04)	(0.04)	(2.22)	(1.75)	(1.98)

Results (p-values) of 2-tailed Nonparametric Tests (Mann–Whitney U)

		Sym	Asym-L		Sym	Asym-L		Sym	Asym-L		Sym	Asym-L		Sym	Asym-L
Rounds	Asym-L	0.0487	–	Asym-L	0.0153	–	Asym-L	0.0071	–	Asym-L	0.0000	–	Asym-L	0.0797	–
1–17	Asym-H	0.0695	0.3178	Asym-H	0.0110	0.9622	Asym-H	0.0001	0.2544	Asym-H	0.0000	0.0002	Asym-H	0.1613	0.6789
Rounds	Asym-L	0.3506	–	Asym-L	0.0258	–	Asym-L	0.0015	–	Asym-L	0.3342	–	Asym-L	0.3924	–
18–37	Asym-H	1.0000	0.0631	Asym-H	0.1918	0.3525	Asym-H	0.0080	0.5946	Asym-H	0.0058	0.0009	Asym-H	0.3818	0.8957
Rounds	Asym-L	0.0748	–	Asym-L	0.0190	–	Asym-L	0.0035	–	Asym-L	0.0002	–	Asym-L	0.1595	–
1_37	Asym-H	0.3501	0.1266	Asym-H	0.0627	0.8130	Asym-H	0.0021	0.5181	Asym-H	0.0000	0.0000	Asym-H	0.2322	0.7590

Notes: Averages of individual market averages are reported in the upper part of this table with standard errors of the mean in parentheses.
[a] "Supra Nash Price Count" denotes the number of market prices that are larger than the highest Nash price.
[b] "Supra Price Nash Index" is defined as $(p\text{-}NE^H)/NE^H$, where p denotes the market price and "NE^H" denotes the highest Nash price, which is equal to 39 in SymC, 38 in Asym-L, and 36 in Asym-H, respectively.
[c] "Collusion Index" is defined as $(\pi^{Observed} - \pi^{HighestNash})/(\pi^{Cartel} - \pi^{HighestNash})$.
[d] "Price Coordination Count" denotes the number of cases in which the two firms of a market chose the same price.

Hypothesis. The higher the level of cost asymmetry, the lower the measure of collusion.

3. Experimental results

We now turn to the experimental evidence regarding the ability of firms to collude under various cost conditions.[20] Different measures of "success" in reaching terms of coordination can be thought of. In this section, we will first comment on the general pattern of play in the various treatments. We will then analyze market prices in some detail. However, as the set of Nash equilibria is not the same in all treatments, the comparison of absolute prices can be criticized. That is why we subsequently introduce and discuss various other measures of collusion. Those are the frequency of prices in excess of the highest Nash equilibrium price ("Supra Nash Price Count"), the (normalized) deviation from the highest Nash equilibrium price ("Supra Nash Price Index"), and the extent to which firms improve on the highest Nash equilibrium profits and come closer to the highest joint profits ("Collusion Index"). Measures of the ability of subjects to choose one and the same price are useful in interpreting the results. We consider the eventual stabilization of play on a single price ("Convergence") along with the frequency with which subjects manage to charge the same price ("Price Coordination Count").

Let us first comment on the general pattern of play. Fig. 2 shows histograms of market prices in rounds 1–37 for all three treatments.[21] In all treatments, pricing above the highest Nash equilibrium was quite common. In treatment SYM, the distribution is double-peaked, with the mode at the perfectly collusive price of 44 and the second most often chosen price being 32 (which is the lowest price in the range of Nash equilibria which involves no loss in case of miscoordination). In contrast, the distribution in treatment ASYM-L 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 ASYM-H appears to be more evenly distributed in comparison to the other treatments. Thus, the comparison between the symmetric and the asymmetric treatments (especially treatment ASYM-L) does not suggest that symmetry helps players choose higher prices.[22]

We now turn to our various measures of collusion. Averages (of individual market averages) for those measures are presented in the upper part of Table 3, separately for various time horizons. The lower part of this table displays the results of Mann–Whitney U tests for pair-wise treatment comparisons, again separately for various time horizons. The unit of observation for those tests is the average for each individual market.

3.1. Market prices

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. Contrary to our Hypothesis, we observe in Table 3 that market prices are highest on average (and less dispersed) in treatment ASYM-L, followed by prices in treatment ASYM-H and then treatment SYM.[23]

The evolution of the average market price in all treatments is shown in Fig. 3. Inspecting this figure, we make a number of observations. First, in treatment SYM, the average price in period 1 is about 34, then increases sharply during the first three periods and more slowly in the periods that follow. The average market price then stabilizes at a price slightly above 39 in later periods. Importantly, average prices in treatment SYM are the lowest of all treatments. Second,

[20] We also ran a few sessions for treatments with incomplete information about the existence of a cost asymmetry. Results were reported in Argenton and Müller (2009).
[21] Due to a clear endgame effect we excluded the last three rounds. In order to account for players' learning, we divided the remaining periods in two, referring to rounds 1–17 as the first half and to rounds 18–37 as the second half.

[22] Abbink and Brandts (2008) report that in their symmetric duopoly markets the lowest price in the range of Nash equilibria which involves no loss in case of miscoordination (24 in their markets, 32 in our treatment SYM) had special attraction for players. This is also the case in our symmetric markets. However, in our asymmetric duopoly markets, the corresponding price of 33 (ASYM-L) and 35 (ASYM-H) is only chosen in relatively few cases, suggesting that the focality of this price seems to be the product of special circumstances, e.g., symmetry. Note that in a one-shot game, Argenton et al. (2010) submit that strategic uncertainty makes 32 an attractor of play. Testing this theory would require eliciting the beliefs that players hold about their rival's actions in the (early rounds of the) experiment.
[23] Statistics regarding individual prices instead of market prices display the same features.

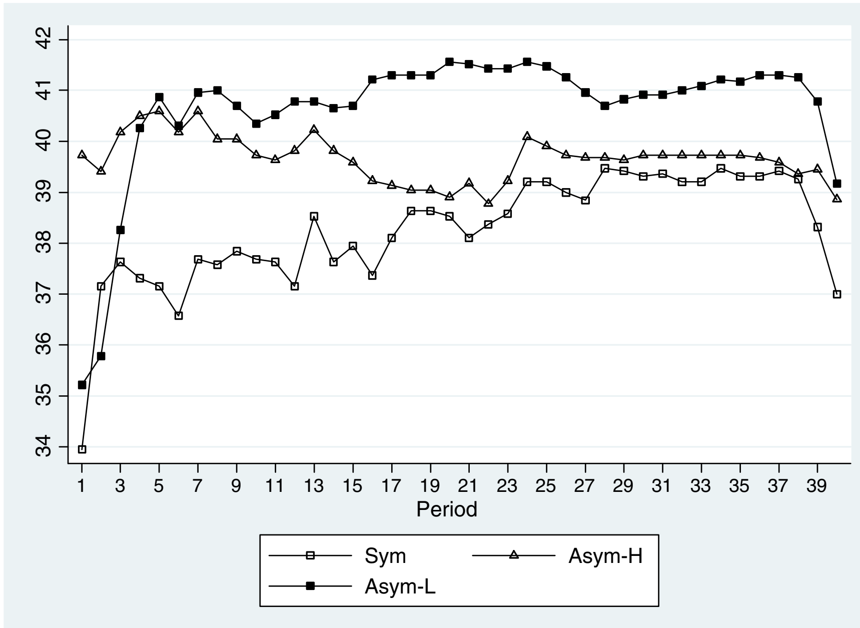

Fig. 3. The evolution of market prices over time.

in treatment ASYM-L the average price in period 1 is slightly higher than 35 and then shoots up to a level of around 41 during the first 5 periods and then more or less stabilizes at this high level. Average prices in treatment ASYM-L are the highest of all treatments. Third, average prices in treatment ASYM-H start at a high level of about 40 in period 1, then almost reach the level of prices in treatment ASYM-L in the next few rounds, but then converge from above to the average prices in treatment SYM. Fourth, whereas in the first half there is a clear gap in average prices between the asymmetric treatments and the symmetric treatment, in the second half there is clear gap in averages prices between treatment ASYM-L and the other two treatments. Fifth, in all treatments we observe a typical endgame effect with average prices sharply decreasing in the last two or three periods.[24] Thus, the information contained in Fig. 3 is, again, not suggestive of the validity of our Hypothesis, according to which symmetry should help players reach higher prices.

For formal test results regarding market prices, we turn to the lower part of Table 3. For the first half of the experiment we find that average market prices in the two asymmetric treatments are significantly higher than in the symmetric treatment while there is no difference in average market prices across the two asymmetric treatments. In the second half, however, the differences in market prices between the two asymmetric and the symmetric treatment cease to be significant (while the difference between the two asymmetric treatments becomes significant). In sum, we don't find evidence that market prices decrease with the level of cost asymmetry.[25]

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

[25] Imposing more structure does not help validate the Hypothesis. Regressions available from the authors use GLS panel regression techniques, allowing for autocorrelation of the error terms and clustering at the market level, to compare average market prices across treatments. Market prices in treatment ASYM-L are always significantly higher than in treatment SYM , while in treatment ASYM-H they are only weakly significantly higher than in treatment SYM in the first half of the experiment, which clearly rejects our Hypothesis. Furthermore, average prices in the two asymmetric treatments differ (weakly) only in the second half of the experiment.

3.2. Supra Nash Price Count

As the set of Nash equilibria changes from one treatment to the other, the comparison of absolute market prices can be criticized for not capturing subject's ability to depart from Nash play. Our first attempt at addressing this criticism consists in measuring the number of times market prices happen to be strictly higher than the highest Nash equilibrium price. With one exception, in Table 3 we observe that this measure is statistically significantly larger for the two asymmetric treatments than for the symmetric treatment (while there is no statistical difference between the two asymmetric treatments). Thus, we find evidence against our Hypothesis that asymmetry impairs subjects' ability to price above static Nash equilibrium prices.

3.3. Supra Nash Price Index

Another way to avoid problems with comparisons of absolute market prices consists in normalizing the deviation from Nash prices. The next measure, which we call the "Supra Nash Price Index", is defined as $(p-\mathrm{NE}^H)/\mathrm{NE}^H$, where p denotes the market price and NE^H denotes the highest Nash equilibrium price in the one-shot game (which is equal to 39 in SYM, 38 in ASYM-L, and 36 in ASYM-H, respectively). In comparison to the market prices we analyzed above, this index effectively controls for the fact that the highest Nash equilibrium price is different across the three treatments. Clearly, this index is negative (positive) if the market price is usually lower (higher) than the highest Nash price. The (absolute value of this) index measures the difference of the market price and the highest Nash equilibrium price as a percentage of the highest Nash equilibrium price. Inspecting the averages of this index across treatments (see Table 3), we find that it is negative in treatment SYM during the first half of the experiment and overall, while it is positive in all other cases. Furthermore, it is monotonic in the degree of asymmetry. The test results in the lower part of Table 3 deliver a clear-cut result: The Supra Nash Price Index is significantly larger in both asymmetric treatments than in the symmetric treatment, but there is no statistical difference between the two asymmetric treatments. Hence, with respect to this index, we find that asymmetry helps subjects deviate from Nash-pricing but we cannot reject the hypothesis that deviations are of the same order of magnitude in the two asymmetric treatments.

3.4. Collusion Index

So far, we focussed only on prices or related measures. A typical measure of collusion consists in measuring the extent to which firms manage to increase their profits above the (highest) Nash equilibrium profits and come closer to cartel profits (see, e.g., Holt, 1995). Therefore, we define the Collusion Index as follows:

$$\text{Collusion Index} = \frac{\pi^{\text{Observed}} - \pi^{\text{HighestNash}}}{\pi^{\text{Cartel}} - \pi^{\text{HighestNash}}},$$

where π^{Observed} stands for the joint profits actually achieved, $\pi^{\text{HighestNash}}$ is the joint profit at the highest Nash equilibrium, and π^{Cartel} stands for the maximum possible joint profits (which are achieved at a common price of 44 in all treatments). The Collusion Index is equal to 1 at the maximal joint profit, and equal to 0 if both sellers choose the highest Nash equilibrium price. Averages of the Collusion Index and related test results are again presented in Table 3. The Collusion Index turns out to be consistently negative across all time horizons in treatment SYM, indicating that subjects do not fare better than under Nash play on average. By contrast, the index is positive in the two asymmetric treatments, independently of the time horizon. Moreover, the Collusion Index is monotonic in the degree of asymmetry. Test results indicate that pair-wise across-treatment differences are (with one exception) statistically significantly different. Hence, there is evidence against our Hypothesis that symmetry helps subjects achieve higher levels of collusion.

3.5. Convergence

The fact the asymmetry seems to be conducive to higher profits is potentially the result of two different effects. On the one hand, because of cost convexity, having both firms charge one and the same price increases profits as compared to the situation where only one firm serves all demand at the market price. On the other hand, conditional on charging the same price, firms have an interest in getting as close as possible to the cartel price of 44. The evidence on market prices does not suggest that they decrease with the level of cost asymmetry. Therefore, the significant differences in the values of the Collusion Index must be explained by a higher ability of subjects to coordinate on the same price, independently of its level. To first look at this issue, we investigate whether subjects in a typical market manage eventually to “agree” on charging the same price, in which case we say that this market “converges”, and the time it takes in case they do. For this purpose, 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. Although somewhat arbitrary, this definition aims at capturing the idea that players have eventually reached a common 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 SYM is 73.3% (14 out of 19 markets) whereas this number is 78.3% (18 out of 23 markets) in treatment ASYM-L and 85.7% (18 out of 21) in treatment ASYM-H. Hence, the percentage of converging markets is lowest in the treatment with symmetric firms. Two-tailed χ^2 tests indicate, however, that only the difference between treatments SYM and ASYM-H is statistically significant ($p = 0.030$). Furthermore, conditional on convergence, the period of convergence is defined as the beginning of the time interval during which both firms uninterruptedly posted the price they converged to. The average period in which markets converge in treatments SYM, ASYM-L, and ASYM-H is, respectively, 11.9, 6.8, and 9.4. 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) indicate that markets in treatment ASYM-L converge significantly earlier than markets in treatment SYM ($p = 0.0864$, two-tailed Mann–Whitney U test), all other differences are insignificant.[26] Thus, there is some evidence that convergence is more easily achieved in asymmetric treatments.

3.6. Price Coordination Count

Another way to look into the issue of coordination is to ask how often firms manage to charge one and the same price, independently of any convergence.[27] For this purpose, we consider the average number of periods in which the two firms in a market chose the same price. Indeed, symmetry could help firms along that dimension as well. We thus briefly analyze the effect of symmetry on the ability of firms to post the same price (see Table 3). Overall, we find that, if anything, it is again asymmetry that helps subjects coordinate on the same price, especially in the early periods of the experiment. For instance, let us consider the incidence of same price choices in the first half of the experiment. We find that firms manage to choose the same price in 45.9% (on average 7.8 out of 17 rounds in treatment SYM), 65.9% (on average in 11.2 out of 17 cases in treatment ASYM-L), and 61.2% (on average 10.4 out of 17 periods in treatment ASYM-H) of the cases. The corresponding test results in the lower part of Table 3 indicate that the difference between treatment SYM and ASYM-L is significant. Note, however, that during the second half of the experiment the statistical differences in this measure across treatments disappear. Hence, there is no evidence that symmetric markets are more conducive to coordination.

3.7. Robustness check

The important result derived so far is that symmetry does not help firms coordinate on higher prices and achieve higher profits. However, the analysis up to now is based on the data for all markets, independently of whether or not they finally converged on a common price. To check the robustness of the absence of evidence in favor of our Hypothesis that symmetry is conducive to achieving higher prices and profits, we repeat the analysis presented in Table 3 by taking into account only those markets that eventually converged. The results of the analysis with data from converged markets only are presented in Table 4. The results are clear-cut. Although the positive effect of cost asymmetry on market prices is usually not statistically significant, we again find no evidence in favor of the claim that symmetry helps achieving more collusion than asymmetry. Again, if anything, asymmetry helps firms achieve higher levels of collusion (according to the measures “Supra Nash Price Count”, “Supra Nash Price Index” and the “Collusion Index” in Table 4).

4. Conclusion

One of the classical questions in antitrust is whether symmetric or asymmetric market structures should be favored (as part of, e.g., merger control when coordinated effects are assessed) or monitored (e.g., when deciding which sectors to target for cartel detection)? Textbooks and policy guidelines, informed by some theoretical

[26] For a detailed analysis of the time it takes for markets to stabilize in the context of Cournot markets, see Mason et al. (1992) and Mason and Phillips (1997).

[27] As pointed out by a referee, this statistic is not a good measure of collusion in markets which, unlike ours, are characterized by constant returns to scale. For instance, players taking turns in capturing the entire market at the monopoly price is highly collusive but would deliver a price coordination count of zero.

Table 4
Collusion and coordination measures and test results. (Data from converged markets only.).

Collusion and Coordination Measures

	Market Price			Supra Nash Price Count[a]			Supra Nash Price Index[b]			Collusion Index[c]			Price Coordination Count[d]		
	SymC	Asym-L	Asym-H	SymC	Asym-L	Asym-H	SymC	Asym-L	Asym-H	SymC	Asym-L	Asym-H	SymC	Asym-L	Asym-H
Rounds	37.8	40.2	39.3	6.7	11.9	11.4	−0.030	0.057	0.091	−7.828	0.411	0.745	9.2	12.8	11.0
1–17	(1.05)	(0.78)	(0.75)	(1.71)	(1.47)	(1.42)	(0.03)	(0.02)	(0.02)	(1.50)	(0.06)	(0.04)	(1.56)	(0.62)	(1.03)
Rounds	39.9	41.8	39.8	11.0	16.8	13.1	0.022	0.101	0.104	−1.022	0.437	0.913	19.6	19.8	19.1
18–37	(1.35)	(0.64)	(0.82)	(2.68)	(1.75)	(2.22)	(0.03)	(0.02)	(0.02)	(0.69)	(0.10)	(0.03)	(0.29)	(0.17)	(0.63)
Rounds	38.9	41.1	39.5	17.7	28.7	24.5	−0.002	0.081	0.098	−4.150	0.426	0.836	28.8	32.6	30.2
1–37	(1.09)	(0.66)	(0.77)	(3.98)	(3.08)	(3.49)	(0.03)	(0.02)	(0.02)	(0.91)	(0.06)	(0.03)	(1.72)	(0.74)	(1.53)

Results (p-values) of 2-tailed Nonparametric Tests (Mann Whitney U)

		Sym	Asym-L		Sym	Asym-L		Sym	Asym-L		Sym	Asym-L		Sym	Asym-L
Rounds	Asym-L	0.2030	–	Asym-L	0.0639	–	Asym-L	0.0440	–	Asym-L	0.0000	–	Asym-L	0.1745	–
1–17	Asym-H	0.2313	0.3188	Asym-H	0.0860	0.6896	Asym-H	0.0018	0.3756	Asym-H	0.0000	0.0003	Asym-H	0.4334	0.990
Rounds	Asym-L	0.8913	–	Asym-L	0.0420	–	Asym-L	0.0013	–	Asym-L	1.0000	–	Asym-L	0.4342	–
18–37	Asym-H	0.5521	0.0370	Asym-H	0.3932	0.1534	Asym-H	0.0663	0.8728	Asym-H	0.0608	0.0311	Asym-H	0.8452	0.5806
Rounds	Asym-L	0.2388	–	Asym-L	0.0538	–	Asym-L	0.0226	–	Asym-L	0.0044	–	Asym-L	0.2142	–
1–37	Asym-H	0.7040	0.0905	Asym-H	0.2013	0.5439	Asym-H	0.0226	0.7042	Asym-H	0.0000	0.0000	Asym-H	0.5669	0.3721

Notes: Averages of individual market averages are reported in the upper part of this table with standard errors of the mean in parentheses.
[a] "Supra Nash Price Count" denotes the number of market prices that are larger than the highest Nash price.
[b] "Supra Price Nash Index" is defined as (p-NE^H)/NE^H, where p denotes the market price and "NE^H" denotes the highest Nash price, which is equal to 39 in SymC, 38 in Asym-L, and 36 in Asym-H, respectively.
[c] "Collusion Index" is defined as $(\pi^{\text{Observed}} - \pi^{\text{HighestNash}})/(\pi^{\text{Cartel}} - \pi^{\text{HighestNash}})$.
[d] "Price Coordination Count" denotes the number of cases in which the two firms of a market chose the same price.

literature and the available experimental studies, suggest that asymmetry is to be favored, as symmetry among firms is thought to be conducive to collusive outcomes. We test this perception in a series of experimental repeated Bertrand duopolies where firms have convex costs. We implement symmetric as well as asymmetric markets that vary in their degree of cost asymmetry among firms. For our lab markets, we never find evidence of symmetric markets being more collusive than asymmetric markets. In fact, for some measures of collusion, we have the opposite result. Firms in our asymmetric treatments come closer to the cartel profits not so much because of clearly higher prices but because they appear to be able to coordinate more often, and to converge earlier, on the same price. This suggests that asymmetric cost conditions may have stabilizing properties in markets with convex costs, where the incentives to split production are high and require firms to charge the same price.

Although the evidence is not overwhelming, we feel that some remarks are in order on this topic. In particular, what could explain earlier convergence in treatment ASYM-L as compared to treatment SYM? We venture that "leadership" may explain this pattern. First, in the asymmetric treatments, players know 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 ASYM-L, 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, in treatment ASYM-L 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 the Online Appendix. However, for reasons that are not entirely clear to us, in treatment ASYM-H we do not find a clear "leadership" tendency of low-cost firms. Instead, in this treatment we find that (especially high-cost) firms have a somewhat higher tendency of not changing their prices from one period to the next in comparison to firms in treatment SYM. This could explain why markets in treatment ASYM-H converge somewhat earlier than those in treatment SYM. On the other hand, a higher extent of more stable play of high-cost firms in treatment ASYM-H in comparison to high-cost firms in treatment ASYM-L also means that high-cost firms in ASYM-H make fewer attempts to pull prices upward. This could be one explanation for why prices in treatment ASYM-H do not stay as high as they are in the first few periods of the experiment. Again, we provide more details in the Online Appendix but acknowledge that more research is to be conducted to account for such puzzling differences.

In future work, we plan on testing the possibly stronger stabilizing properties of asymmetric markets by periodically shocking markets and studying oligopolists' adaptation to changing conditions. In general, we think that the properties of Bertrand markets with convex costs and the results presented here and elsewhere warrant further (experimental) study.

Appendix A *Instructions*

We here display the instructions for treatment SYM. *The changes in the instructions for asymmetric treatments* ASYM-L *and* ASYM-H *are displayed between brackets after the corresponding passages in the instructions for* SYM.

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 is identical to the one you have.

[Attached to the instructions are two tables. In the first 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. In the second table you can see the profits (or losses) that the other firm will make depending on the prices chosen by itself and by yourself. Notice that the two tables are different.]

Down the first column of the table [of your profit table (Table 1)] 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.”

[The profit table for the other firm in your market (Table 2) can be read in an analogous manner.]

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.

Appendix B. Supplementary data

Supplementary data to this article can be found online at http://dx.doi.org/10.1016/j.ijindorg.2012.05.006.

References

Abbink, K., Brandts, J., 2008. 24. Pricing in Bertrand competition with increasing marginal costs. Games and Economic Behavior 63, 1–31.
Abreu, D., 1986. Extremal equilibria of oligopolistic supergames. Journal of Economic Theory 39, 191–225.
Abreu, D., 1988. Towards a theory of discounted repeated games. Econometrica 56, 383–396.
Alós-Ferrer, C., Ania, A.B., Schenk-Hoppé, K.R., 2000. An evolutionary model of Bertrand oligopoly. Games and Economic Behavior 33, 1–19.
Altavilla, C., Luini, L., Sbriglia, P., 2006. Social learning in market games. Journal of Economic Behavior & Organization 61, 632–652.
Andreoni, J., Miller, J.H., 1993. Rational cooperation in the finitely repeated prisoner's dilemma: experimental evidence. The Economic Journal 103, 570–585.
Apesteguia, J., Huck, S., Oechssler, J., 2007. Imitation: theory and experimental evidence. Journal of Economic Theory 136, 217–235.
Argenton, C., Müller, W., 2009. Bertrand competition with convex costs in symmetric and asymmetric markets: Results from a pilot study. In: Franz, W., Güth, W., Ramser, H.J., Stadler, M. (Eds.), Experimentelle Wirtschaftsforschung. Tübingen, Mohr Siebeck, pp. 147–165.
Argenton, C., Andersson, O., Weibull, J.W., 2010. Robustness to Strategic Uncertainty in Price Competition. SSE/EFI working paper no. 726.
Bae, H., 1987. A price-setting supergame between two heterogeneous firms. European Economic Review 31, 1159–1171.
Benoît, J.-P., Krishna, V., 1985. Finitely-repeated games. Econometrica 53, 905–922.
Boone, J., Larraín Aylwin, J.L., Müller, W., Ray Chaudhuri, A., 2012. Bertrand competition with asymmetric costs: experimental evidence. Economics Letters 117, 134–137.
Bruttel, L., 2008. Group dynamics in experimental studies—The Bertrand Paradox revisited. Journal of Economic Behavior & Organization 69, 51–63.
Compte, O., Jenny, F., Rey, P., 2002. Capacity constraints, mergers and collusion. European Economic Review 46, 1–29.
Dastidar, K.G., 1995. On the existence of pure strategy Bertrand equilibrium. Economic Theory 5, 19–32.
Davidson, C., Deneckere, R., 1990. Excess capacity and collusion. International Economic Review 31, 521–541.
Davies, S., Olczak, M., Coles, H., 2011. Tacit collusion, firm asymmetries and numbers: evidence from EC merger cases. International Journal of Industrial Organization 29, 221–231.
Dixon, H.D., 1990. Bertrand–Edgeworth equilibria when firms avoid turning customers away. The Journal of Industrial Economics 39, 131–146.
Dufwenberg, M., Gneezy, U., 2000. Price competition and market concentration: an experimental study. International Journal of Industrial Organization 18, 7–22.
Dufwenberg, M., Gneezy, U., 2002. Information disclosure in auctions: an experiment. Journal of Economic Behavior and Organization 48, 431–444.
Dugar, S., Mitra, A., 2009. The size of the cost asymmetry and Bertrand competition: experimental evidence. Mimeo.
Dugar, S., Mitra, A., 2011. Bertrand competition with different marginal costs: the role of the cost asymmetry. Mimeo.
Engel, C., 2007. How much collusion? A meta-analysis of oligopoly experiments. Journal of Competition Law and Economics 3, 491–549.
Fatas, E., Haruvy, E., Morales, A.J., 2009. An experimental re-examination of the Bertrand paradox. Mimeo.
Fonseca, M.A., Normann, H.-T., 2008. Mergers, asymmetries and collusion: experimental evidence. The Economic Journal 118, 387–400.
Fouraker, L., Siegel, S., 1963. Bargaining Behavior. McGraw–Hill, New York.
Haan, M.A., Schoonbeek, L., Winkel, B.M., 2009. Experimental results on collusion. In: Hinloppen, J., Normann, H.-T. (Eds.), Experiments and competition policy. Cambridge University Press, New York, pp. 9–33.
Harrington, J.E., 1991. The determination of price and output quotas in a heterogeneous cartel. International Economic Review 32, 767–792.
Hoernig, S., 2002. Mixed Bertrand equilibria under decreasing returns to scale: an embarrassment of riches. Economics Letters 74, 359–362.
Holt, C.A., 1995. Industrial Organization: A Survey of Laboratory Research. In: Kagel, J., Roth, A. (Eds.), Handbook of Experimental Economics. Princeton University Press, Princeton, pp. 349–443.
Huck, S., Normann, H.-T., Oechssler, J., 1999. Learning in Cournot oligopoly — an experiment. The Economic Journal 109, C80–C95.
Huck, S., Normann, H.-T., Oechssler, J., 2004. Two are few and four are many: number effects in experimental oligopolies. Journal of Economic Behavior & Organization 53, 435–446.
Ivaldi, M., Jullien, B., Rey, P., Seabright, P., Tirole, J., 2003. The economics of tacit collusionfinal report for DG Competition, European Commission (available at , http://idei.fr/doc/wp/2003/tacit_collusion.pdf20036 May 2010).
Keser, C., 1993. Some results of experimental duopoly markets with demand inertia. The Journal of Industrial Economics 41, 133–151.
Keser, C., 2000. Cooperation in symmetric duopolies with demand inertia. International Journal of Industrial Organization 18, 23–38.
Kübler, D., Müller, W., 2002. Simultaneous and sequential price competition in heterogeneous duopoly markets: experimental evidence. International Journal of Industrial Organization 20, 1437–1460.
Mailath, G.J., Samuelson, L., 2006. Repeated games and reputations: long-run relationships. Oxford University Press, New York.
Mason, C.F., Phillips, O.R., 1997. Information and cost asymmetry in experimental duopoly markets. The Review of Economics and Statistics 79, 290–299.
Mason, C.F., Phillips, O.R., Nowell, C., 1992. Duopoly behavior in asymmetric markets: an experimental evaluation. The Review of Economics and Statistics 74, 662–670.
Miklós-Thal, J., 2011. Optimal collusion under cost asymmetry. Economic Theory 46, 99–125.
Motta, M., 2004. Competition policy: theory and practice. Cambridge University Press, Cambridge.
Offerman, T., Potters, J., Sonnemans, J., 2002. Imitation and belief learning in an oligopoly experiment. Review of Economic Studies 69, 973–997.
Rothschild, R., 1999. Cartel stability when costs are heterogeneous. International Journal of Industrial Organization 17, 717–734.
Scherer, F.M., 1980. Industrial market structure and economic performance. Houghton Mifflin Company, Boston.
Selten, R., 1965. Spieltheorische Behandlung eines Oligopolmodells mit Nachfrageträgheit. Zeitschrift für die Gesamte Staatswissenschaft 121, 301–324.
Selten, R., Stoecker, R., 1986. End behavior in finite prisonner's dilemma supergames. Journal of Economic Behavior & Organization 1, 47–70.
Vasconcelos, H., 2005. Tacit collusion, cost asymmetries, and mergers. The RAND Journal of Economics 36, 39–62.
Vives, X., 1999. Oligopoly Pricing: Old Ideas and New Tools. MIT Press, Cambridge, Mass.
Weibull, J.W., 2006. Price competition and convex costs. SSE/EFI Working Paper No 622.
Wolfram, C.D., 1998. Strategic bidding in a multiunit auction: an empirical analysis of bids to supply electricity in England and Wales. The RAND Journal of Economics 29, 703–725.