Journal of Economic Behavior and Organization 233 (2025) 106968

Contents lists available at ScienceDirect

Journal of Economic Behavior and Organization

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

Research paper

Sharing rules in Bertrand duopolies with increasing returns ☆

Andreas Orland

Corvinus University of Budapest, Institute of Economics, Fővám tér 8, 1093 Budapest, Hungary

ARTICLE INFO

Dataset link: https://doi.org/10.17605/osf.io/jnufw

JEL classification:
C72
C90
L13

Keywords:
Sharing rules
Price competition
Tacit collusion
Increasing returns to scale
Experiment

ABSTRACT

Despite their empirical relevance, increasing returns to scale are understudied in experimental markets. We use Bertrand duopolies with increasing returns to examine the effects of two sharing rules on collusive behavior and prices in a pre-registered experiment: the symmetric rule (where each of the two firms that set the same price serves half of the market demand) and the winner-takes-all rule (where a fair randomization device decides which of the two firms serves the entire market). We hypothesized that market prices would be higher under the winner-takes-all rule because it provides a collusion mechanism that the symmetric rule does not. While we find that subjects under the winner-takes-all rule coordinate more often on one price than the symmetric sharing rule, this does not increase market prices. Coordination on high prices is rare. Additionally, the winner-takes-all rule facilitates the subjects' ability to coordinate on equal prices after sharing a market in the previous period.

1. Introduction

Despite their empirical relevance, increasing returns to scale (which result in decreasing average costs) and their effect on price-setting and tacit collusion are understudied in experimental markets. While the chemical and airline industries are examples, many other industries also produce under increasing returns to scale.[1] In this paper, we examine Bertrand competition in a laboratory experiment where returns to scale increase (where we use avoidable fixed costs in the cost function as a simple means of inducing increasing returns to scale).[2] In a treatment variation, we examine the effect of two different sharing rules on price-setting and collusive behavior using different coordination channels.

☆ The editor and two anonymous reviewers' comments helped improve the paper. The participants' helpful comments during the Potsdam Research Seminar in Economics, the Economic Science Association's 2022 World Meeting in Boston, and the Corvinus University of Budapest Economics Seminar are gratefully acknowledged. Lisa Bruttel, Wieland Müller, Holger Rau, Tetiana Sobolieva, Christopher Stapenhurst, Attila Tasnádi, and Bert Willems provided valuable feedback. Theo Schweisgut and Rahmatullah Yusefi provided excellent research assistance. The author wishes to thank the Corvinus Institute for Advanced Studies, Corvinus University of Budapest for their support of this research. This study was, before data collection, registered in the AEA RCT Registry (AEARCTR-0009056, https://doi.org/10.1257/rct.9056). Data and all scripts are available in a public OSF repository: https://doi.org/10.17605/osf.io/jnufw. This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

E-mail address: andreas.orland@uni-corvinus.hu.

[1] Blinder et al. (1998, chapter 12) and Robidoux and Lester (1992) highlight the empirical relevance of increasing returns to scale. Blinder et al. (1998) conducted an interview survey with the CEOs of 200 companies, a sample representative of the US economy. Asked about their company's cost structure, Blinder et al. report that about 40% of CEOs describe a cost structure with falling marginal costs (which resembles a situation with increasing returns to scale). Robidoux and Lester (1992) analyze cost-scale relationships and find evidence for increasing returns to scale for most of the 147 examined Canadian industries.

[2] In contrast to fixed costs, avoidable fixed costs are *only* incurred in the case of production—they are not sunk costs. Durham et al. (2004, p. 148) offer an illustrative example from the airline industry: adding passengers to an airline flight increases costs to some degree, but most of the costs are for fuel and crew and can be avoided by not flying (whereas the overhead costs for administration, etc., fall in the category of sunk fixed costs).

https://doi.org/10.1016/j.jebo.2025.106968
Received 19 September 2024; Received in revised form 22 February 2025; Accepted 1 March 2025

Available online 18 March 2025
0167-2681/

We use the well-known case of Bertrand competition with constant returns to highlight the differences that arise when returns are increasing. In the classic formulation of Bertrand competition, all firms are identical, face no fixed costs, and their marginal cost is constant. In theory, in a one-shot setting, firms should set their price equal to their marginal cost and make zero profits. However, firms can collude, especially in repeated interactions, by jointly setting a higher price. The question of how to share market demand in case of equal prices depends on the market's *sharing rule*. If market demand is split equally between all lowest-charging firms, one speaks of the *symmetric* sharing rule (labeled as such by Baye and Kovenock 2008). It is usually (implicitly) used in Bertrand competition. If, on the other hand, a fair randomization device determines a single firm (among all lowest-charging firms) to serve the entire market, this is called the *winner-takes-all* sharing rule (introduced by Baye and Morgan 2002). Under constant returns to scale, equilibrium prices are identical, and firms ensure the same expected profits when colluding under both sharing rules. Theory does not predict collusion to be more likely under either sharing rule.

When returns to scale are increasing, the firms' average cost decreases in output (thus, production becomes cheaper in output). Here, the sharing rule might affect collusive behavior. Assume that all firms consider colluding by setting the monopoly price. The consequences for each firm's expected profit depend on the market's sharing rule. The expected profit from collusion is higher for the winner-takes-all sharing rule than for the symmetric sharing rule, as, in the former, avoidable fixed costs are only incurred by the producing firm, and market demand is not split between firms. These higher expected profits make collusive behavior more plausible under the winner-takes-all sharing rule (and a comparison of critical discount factors assuming boundedly rational behavior supports this). We test this prediction and the underlying intertemporal mechanism in our experiment. It is informative for policymakers like competition and regulation authorities to know the sharing rule's effect on collusive behavior when human subjects take decisions in a controlled and incentivized setting. Public procurement auction designers might be interested in an experimental comparison of the two most frequently used sharing rules.

Our research contributes to two strands of literature: the literature on non-linear returns in Bertrand competition and the literature on sharing rules. There is extensive theoretical literature on *decreasing* returns in Bertrand competition. When returns decrease, i.e., due to convex costs, the Bertrand game might change fundamentally (serving the entire market can be considered a privilege under constant returns but might become an obligation under decreasing returns). Firms might not find it optimal to produce large quantities under the symmetric sharing rule—coordinating on a specific price might be more attractive for two firms than one firm slightly undercutting it. Dastidar (1995), Dixon (1984), and Hoernig (2002) examine the conditions for pure and mixed equilibria when costs are convex.

Besides the theoretical literature, convex costs have also been studied experimentally by Abbink and Brandts (2008) and Argenton and Müller (2012); both use only the symmetric sharing rule in their experiments. Abbink and Brandts (2008) vary the number of firms between two and four (who compete for 50 periods in fixed pairs). The observed frequency of collusive outcomes decreases with the number of firms, which is not predicted by theory. They also find imitation behavior and coordination on a focal point. Argenton and Müller (2012) examine the effect of cost asymmetries on collusion under convex costs. In one treatment, they conduct duopoly experiments with a symmetric cost structure; in two further treatments, they vary the level of cost asymmetry between firms (in all treatments, subjects interact for 40 periods in fixed pairs). Contrary to theory, the authors do not find that collusion decreases with cost asymmetry.

The case of *increasing* returns has also received some attention in the theoretical literature. Chaudhuri (1996) examines a symmetric Bertrand duopoly with avoidable fixed costs and constant marginal costs, and Chowdhury (2002) extends the model to asymmetric fixed costs. Both papers use a symmetric sharing rule and a discrete strategy space. The models' equilibria are characterized: there are no equilibria where both firms charge the same price. Edwards and Routledge (2023) provide necessary and sufficient conditions for the existence and uniqueness of a pure strategy Nash equilibrium in Bertrand competition with convex costs and a general sharing rule. There is no experimental test of Bertrand competition with increasing returns to scale.

Some contributions examine the effect of sharing rules on outcomes in Bertrand markets.[3] Baye and Morgan (2002) introduce winner-takes-all price competition and show how discontinuities in the (symmetric) cost function affect the undercutting argument in Bertrand competition. Hoernig (2007) examines asymmetric Bertrand games under different sharing rules. Baye and Kovenock (2008) analyze and compare symmetric Bertrand competition under the two sharing rules (we will refer to this paper in the theory section in more detail). The only experimental study concerned with sharing rules is due to Puzzello (2008). In a two-by-two experimental design, she varies the sharing rule (symmetric/winner-takes-all) and the divisibility of the price space (a fine treatment with 10,000,000 prices and a coarse treatment with 200 prices on the same interval) in a Bertrand competition with a capacity constraint and constant marginal costs. While price predictions do not differ between the symmetric and the winner-takes-all sharing rule, Puzzello finds that firms coordinate more frequently on the monopoly price under the symmetric rule, especially when the price grid is coarse.

In our experimental test, we use a Bertrand duopoly with symmetric firms. We implement increasing returns to scale using a cost structure with avoidable costs (fixed costs that are only incurred for positive output) and constant marginal costs, which leads to decreasing average costs. We use a finitely repeated setup and vary only the sharing rule between treatments. While we find that

[3] Sharing rules are also used in other fields. (i) In the literature on contests and tournaments, there are also experimental studies on sharing rules (how members of a winning team split up prizes); see Sheremeta (2018, p. 691) for an overview. A related study is by Brookins et al. (2021): they experimentally compare winner-take-all and proportional-prize indefinitely repeated contests. They report, under some conditions, evidence of less cooperation in the latter. (ii) In the auctions literature, there are also tie-breaking/rationing rules (which define what happens if two or more bidders place the same highest bid; the winner-takes-all rule is the standard in first-price auctions (Maskin and Riley, 2000, p. 440)). However, we do not know of any experimental comparison of different sharing rules in the context of auctions.

subjects under the winner-takes-all rule coordinate more than twice as often on one price compared to the symmetric sharing rule, we do not find that this increases market prices. This might be driven by the fact that subjects do *not* coordinate on sufficiently high prices. In further analyses, we report findings on alternation and intertemporal price adjustments. Here, compared to the symmetric rule, the winner-takes-all rule facilitates the subjects' ability to coordinate on equal prices after previously sharing a market.

The remainder of the paper is structured as follows: the next section describes the related theory. Section 3 explains and relates our experimental implementation to the theory. Section 4 describes the hypotheses we test with our experiment, and in Section 5, we present the results. In Section 6, we discuss our findings and conclude.

2. Theoretical considerations

2.1. The static game

Here, we provide an intuition for Bertrand pricing with increasing returns to scale under the two sharing rules for the one-shot game. We follow the original derivations for the symmetric sharing rule by Baye and Kovenock (2008, pp. 3–4) and the winner-takes-all sharing rule by Baye and Morgan (2002, pp. 278–280).

General framework Consider a market where two identical, risk-neutral firms compete to supply a homogeneous good. Both firms $i \in \{1, 2\}$ simultaneously set prices $p_i \in \mathbb{R}_{\geq 0}$. A linear function, $D(p) = (a - p)/b$, describes market demand. The firm with the lower price serves the entire market, and the firm with the higher price does not sell anything. The firms have an identical cost function:

$$C(q_i) = \begin{cases} 0 & \text{if } q_i = 0 \\ f + c\ q_i & \text{if } q_i > 0 \end{cases} \quad (1)$$

If a firm does not produce, it incurs no cost. If it does produce, it incurs an avoidable fixed cost f and a constant unit cost c. Thus, the average cost decreases as output q_i increases. For the case of a single firm serving the whole market, this results in monopoly quantity $q^M = (a - c)/(2\ b)$, monopoly price $p^M = (a + c)/2$, and monopoly profit $\pi^M = (a - c)^2/(4\ b) - f$. Depending on the parametrization of the cost and demand structure, there are two "breakeven" prices for the monopolist where profits are zero: $p^M_{0,-}$ (the lower breakeven point; below this price, profits are negative), and $p^M_{0,+}$ (the upper breakeven point; above this price, profits are negative). Between $p^M_{0,-}$ and $p^M_{0,+}$, monopoly profits are positive and peak at p^M. Illustrations of the monopoly profit π^M and the duopoly profit π^D (when two firms set the same price and share the market) for the parametrization used in our experiment are shown for each sharing rule in Fig. 1.

Symmetric sharing rule If the two firms charge the same price, they split the market evenly. However, since monopoly profits are negative for prices below $p^M_{0,-}$ (and above $p^M_{0,+}$), neither firm has an incentive to set prices in these ranges. If both firms were to charge this price at the lower breakeven price $p^M_{0,-}$, their duopoly profits π^D would be negative. This is because the costs associated with sharing the market demand are higher than sharing the cost, i.e., $C(D(p^M_{0,-})/2) > C(D(p^M_{0,-}))/2$, while the revenue is split equally. To avoid these negative profits, a firm could set a higher price p^H (up to p^M), which would improve its monopoly profit conditional on being the lower-pricing firm. However, this also reduces the probability of being that lower-pricing firm. Consequently, duopoly profits π^D are positive over a narrower range of prices compared to monopoly profits π^M. If one firm sets a price $p^D_{0,-}$ that ensures a non-negative duopoly profit, the other firm can benefit by undercutting it, thereby securing a positive monopoly profit. However, if both firms would charge $p^M_{0,-}$, they would each incur a negative duopoly profit (up until $p^D_{0,-}$, where their profits would be zero). If either firm chooses a probability distribution on the prices $p^M_{0,-}$ and above, the other firm could choose a different probability distribution to improve its position. As also noted by Chaudhuri (1996), there is no equilibrium in this setup where both firms set the same price (or probability distribution over prices) — whether using pure or mixed strategies — when the strategy space is continuous, as each firm has an incentive to continually adjust its price to improve its position (as there is always a smaller price that a firm can put weight on). This prediction changes for a discrete strategy space, and we will return to this argument when describing our experimental design; there, we present the pure and mixed strategies of the static game that emerge under the experiment's parametrization.

Winner-takes-all sharing rule A fair random draw decides which firm with the same price serves the entire market. Like the classic Bertrand competition with constant returns to scale, firms compete to serve the entire market and set a price equal to the lower breakeven price, $p^M_{0,-}$, where monopoly profits are zero. The random draw decides which firm makes zero profits because it does *not* serve the market and which firm serves the market at the breakeven price. Thus, the Bertrand paradox persists under this sharing rule.

2.2. The repeated game

To model firm behavior in the repeated game setting, we introduce a heuristic approach that accounts for bounded rationality through a form of myopia and firms not applying backward induction.[4] As often done in indefinitely repeated games, we also assume a grim trigger strategy: In the first period, both firms collude (set the monopoly price); if both firms colluded in all past periods,

[4] Spiegler (2011) summarizes research that incorporates bounded rationality in industrial organization. Orland and Roos (2013, 2019) provide examples of firms' myopic price-setting in macroeconomics experiments.

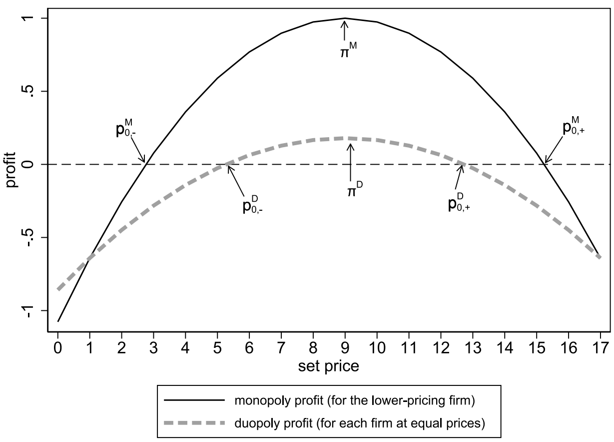

(a) Profits under the symmetric sharing rule

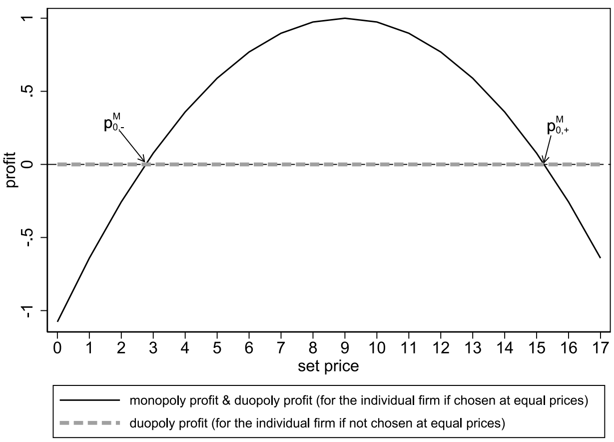

(b) Profits under the winner-takes-all sharing rule

Fig. 1. Profit functions (based on the experimental parametrization).

they collude in the next period; otherwise (especially if one firm undercut the monopoly price), both firms punish each other until the end (by setting the stage-game equilibrium price). In the finitely repeated game, this threat of future punishment is not credible, assuming rationality as backward induction predicts that the stage-game equilibrium is also the subgame-perfect equilibrium of the repeated game, and (mutual) punishment should occur in every period from the start. For our behavioral modification, we assume that the boundedly rational firms ignore the mutual defection in the final period T (where the decision between the collusion payoff and the higher undercutting payoff is straightforward) and that in every starting period $t < T$, the firms compare the discounted expected collusion and defection payoffs over the *remaining* periods until the end. Thus, for each starting period $t \in \{1, \dots, T-1\}$, two equations describe the firms' discounted lifetime payoffs (where the discount factor δ_t is allowed to vary between the different starting periods t but is identical in all periods following starting period t):

$$\mathbb{E}[\Pi_t^{\text{Collusion}}] = \mathbb{E}[\sum_{k=t}^{T} \delta_t^{k-1} \pi^{\text{Collusion}}], \tag{2}$$

and

$$\mathbb{E}[\Pi_t^{\text{Defection}}] = \mathbb{E}[\pi^{\text{Undercutting}} + \sum_{k=t+1}^{T} \delta_t^{k-1} \pi^{\text{Punishment}}]. \tag{3}$$

The grim trigger strategy implies that firms compare the potential long-term benefits of collusion with the immediate gains from defection and following punishment payoffs, adjusted by their degree of bounded rationality as captured by the discount factor δ_t. Combining the two conditions leaves us with $\sum_{k=t}^{T} \delta_t^{k-1} \pi^D \geq \pi^M_{p-\epsilon} + \mathbb{E}[\sum_{k=t+1}^{T} \delta_t^{k-1} \pi^{\text{Punishment}}]$ (where $\pi^M_{p-\epsilon}$ stands for the monopoly profit from undercutting the collusive price by the smallest-possible amount) for Sym and $\sum_{k=t}^{T} \delta_t^{k-1} \pi^M/2 \geq \pi^M_{p-\epsilon} + \sum_{k=t}^{T} \delta_t^{k-1} \pi(p^M_{0,-})/2$

Table 1
Summary of the aggregated equilibria and the euro payoffs in the two treatments.

Treatment	# of equilibria		Mean	Median	$\mathbb{E}[\pi^{\text{Collusion}}]$	$\mathbb{E}[\pi^{\text{Defection}}]$	$\mathbb{E}[\pi^{\text{Punishment}}]$
	pure	mixed	Set price	Set price	In euros	In euros	In euros
WTA	1	–	3.00	3.00	0.50	1.00	0.04
Sym	2	19	3.06	3.00	0.18	1.00	−0.02

Note: The reported mean and median set prices and euro profits in Sym are derived by assigning equal weights to all 21 equilibria.

for WTA. For each of these $(T-1)$ inequalities per sharing rule, we can determine the smallest δ_t that describes the critical behavioral discount factor δ_t^*. This factor represents the threshold at which firms are indifferent between maintaining collusion and defecting and allows us to compare each starting period's discount factors between sharing rules to infer where collusion is more likely. In the next section, after showing the experimental parametrization, we determine the expected payoffs in both treatments and the implied critical behavioral discount factors.

3. Experimental design & procedures

Experimental design In our between-subject experimental design, we vary the sharing rule for an otherwise identical Bertrand duopoly. We label our treatments Sym (for the symmetric sharing rule) and WTA (for the winner-takes-all sharing rule).

We base the market environment on the following parametrization of the Bertrand competition: Two symmetric firms, $i \in \{1, 2\}$, play the stage game repeatedly. In each period, they compete by simultaneously setting prices from 0 to 9 (in 25-cent increments), $p_i \in \{0.00, 0.25, 0.50, \dots, 8.75, 9.00\}$. Each firm has the cost function $C(q_i) = \begin{cases} 0 & \text{if } q_i = 0 \\ 25 + q_i & \text{if } q_i > 0 \end{cases}$. The demand function is $D(p) = 17 - p$.

If firm i sets the lower price, it serves the whole market and makes a profit of $\pi_i(p_i, p_{-i}) = p_i \; D(p_i) - C(D(p_i))$ while the other firm, $-i$, receives nothing: $\pi_{-i}(p_i, p_{-i}) = 0$. If the two firms set the same price, the duopoly profits differ between the two treatments. Fig. 2(a) shows monopoly and duopoly euro profits in Sym.[5] Monopoly and duopoly profits increase in the set price and peak at 8.75 and 9.00, respectively. Monopoly profits are negative for prices between 0.00 and 2.50 (here, revenue is not high enough to cover costs). Aggregate duopoly profits are lower than monopoly profits (as firms share market demand evenly while each incurs avoidable costs), and profits are negative for a wider range of prices (between 0.00 and 5.00). Fig. 2(b) shows the profits in WTA. Monopoly prices are identical to Sym. A fair random draw decides which of the two firms serves the market. Thus, the duopoly profit is equivalent to the monopoly profit for the chosen firm and zero for the other firm. In contrast to Fig. 1(b), Fig. 2(b) shows *expected* duopoly profits for WTA.

In the following, we discuss some practical aspects of our choice of allowed prices and how they relate to equilibrium predictions. First, we focus on the price range up to the monopoly price, p^M. As the prices below p^M dominate the prices beyond it, we restrict the strategy space to the economically relevant part. This design choice has two advantages: For the subjects, it is intuitive to understand the incentives (a higher price comes with a higher (monopoly or duopoly) profit; beyond the monopoly price, profits would decrease again), and for us, it is easy to interpret the outcomes in the market (exercised market power increases in market price).

Second, as in any other price-setting experiment, we also have to discretize the strategy space and, thus, deviate from the assumption of a continuous strategy space.[6] For the WTA stage game, we find only one pure-strategy equilibrium where both firms set a price of 3.00, marginally above the zero-profit price of 2.75.[7] The use of a discrete strategy space affects predictions in the Sym stage game: As every game with a finite number of players and strategies has at least one equilibrium (Nash, 1951), in contrast to the theoretical prediction for the symmetric sharing rule with a continuous strategy space, for our experimental game, we identify 21 equilibria (in the range of prices where monopoly and duopoly profits equal zero, 2.75 and 5.25): two with pure strategies and 19 with mixed strategies. However, due to our choice to implement a strategy space with a relatively fine resolution, both treatments' mean and median prices are very close. Using a finer grid results in the Sym mean prediction being closer to the WTA mean prediction, but it also entails that coordinating on one price might become harder in the experiment. A coarser grid can yield clearer results in experiments, particularly when time and observations are limited.[8] As the only (payoff-relevant) difference between the two treatments occurs when both subjects choose an equal price, we intended (i) to observe periods with equal prices,

[5] We used a conversion rate of 39 "model monetary units" per euro to calculate the euro profits (where we rounded to the nearest cent). Then, in the experiment, we showed the subjects only the euro profits, avoiding the exchange from an experimental currency. Tables A.1 and A.2 in the Appendix show the exact euro profit values.

[6] Our price grid consists of 37 prices and is of comparable size to the two experiments testing price competition with decreasing returns to scale (Argenton and Müller 2012, 41 prices; Abbink and Brandts 2008, 40 prices).

[7] A Nash equilibrium price one step *above* the zero-profit price is a common result of a discrete strategy space. If both firms set a price of 2.75, each will make zero profit with certainty. If both set 3.00, one randomly determined firm will receive a positive profit, while the other will get nothing. Undercutting at 3.00 does not pay off.

[8] Both Brown-Kruse (1991) and Puzzello (2008) vary the resolution of the strategy space in experiments. Both find that the observed frequency of collusion is higher in their "coarse" treatments than in their "fine" treatments and that mean prices between treatments are not different (Brown-Kruse, 1991) or that the results regarding the differences of set prices are inconclusive (Puzzello, 2008).

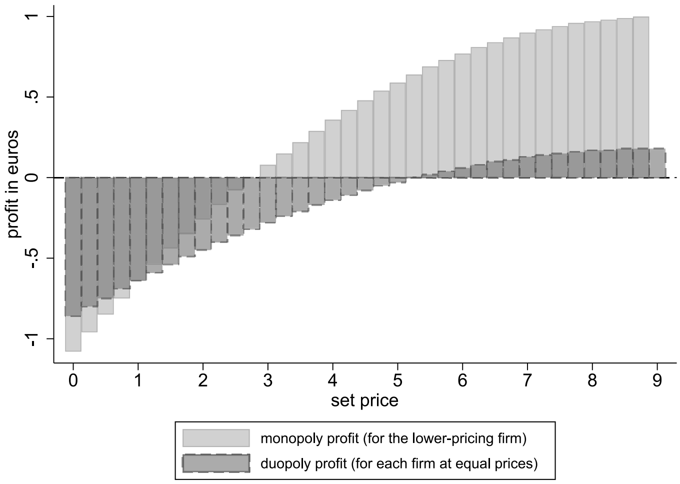

(a) Profits in Sym

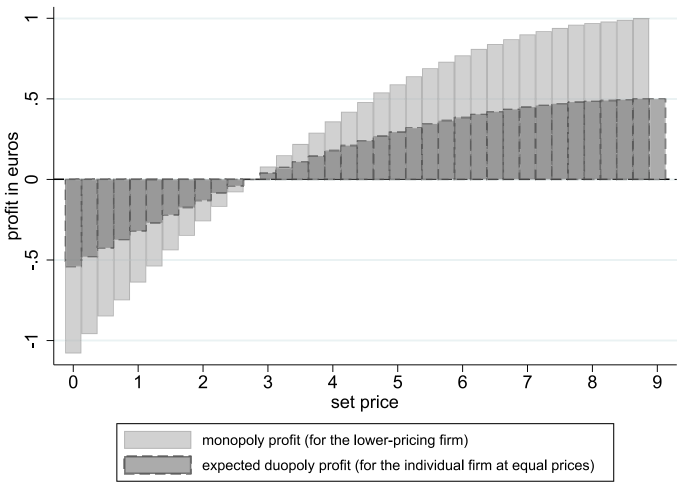

(b) Profits in WTA

Fig. 2. Profits in the experiment.

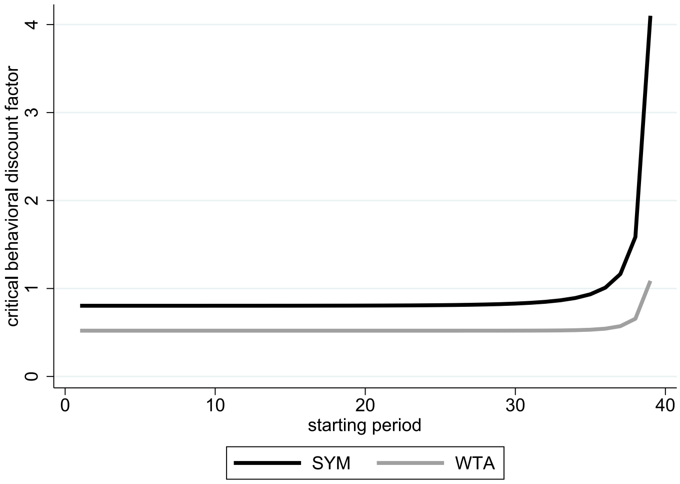

Fig. 3. The time-varying critical behavioral discount factors in both treatments.

and (ii) to be able to observe a difference in these numbers between the two treatments by choosing the strategy space. Table 1 shows the aggregated set prices according to the stage game equilibria for the two treatments and the expected per-period euro payoffs for collusion, undercutting, and punishment.[9]

Based on the reasoning in Section 2.2 and the expected payoffs in Table 1, we calculate the time-varying critical behavioral discount factors over periods 1–39 and show them in Fig. 3. The discount factors in Sym start at a higher level and increase more sharply towards the end than in WTA.[10] That means that subjects in Sym might have a stronger preference for early payoffs, which potentially leads to more or earlier undercutting in Sym compared to WTA.[11]

Experimental procedures Before each session, we randomly determined it to be either a Sym or a WTA session (up to the pre-determined number of subjects in each treatment). Then, at the beginning of the experiment, each subject was randomly matched with another subject and interacted anonymously with this subject for the entire experiment. The experiment lasted 40 periods, and this and all other design features of their treatment were known to the subjects.[12]

We conducted a finitely repeated game in our experiment instead of an indefinitely repeated game (where a random draw with a known probability determines whether the game ends in the next period; proposed by Roth and Murnighan 1978) for four reasons. First, we intended to replicate the main design features of Abbink and Brandts 2008, Puzzello 2008, Argenton and Müller 2012, which all use a fixed and known number of periods. Second, a finitely repeated game is transparent and easy to communicate to the subjects (which is crucial for ensuring that subjects understand the game they are participating in and making it easier for them to strategize given the incentives of the environment); the uniformity of the data facilitates the analysis and gives us more control over the duration of the sessions. Third, we collected the experimental data in a small number of seven sessions; thus, a random draw terminating the indefinitely repeated game for all duopolies in the same session at once can result in a very uneven number of periods' observations between the two treatments. Fourth, despite its advantages, like avoiding endgame effects, indefinitely repeated games in experiments also show some problems. E.g., Lugovskyy et al. (2017) do not find consistent evidence that overall cooperation rates in Public Goods Game experiments depend on whether the number of periods is finite or determined probabilistically. Roth (1995, pp. 26–28), Dal Bó (2005, pp. 1591–1594), Normann and Wallace (2012, pp. 708–709) and Dal Bó and Fréchette (2018) discuss the pros and cons of finitely and indefinitely repeated games in the lab.

We framed the experiment in terms of price competition. Instead of informing subjects about all the details of the underlying model, we explained the experimental game using payoff tables (which was also done in the experiments by Abbink and Brandts 2008, Argenton and Müller 2012). Our payoff tables show the profit consequences (in euros) for all possible prices, depending on

[9] We used the software package Gambit (by McKelvey et al. 2016) to calculate these equilibria (based on the euro profits in the payoff tables). The Gambit data files for both experimental games are available in the public repository. Table A.3 in the Appendix displays the equilibria in Sym.

[10] In both treatments, the calculated δ_t^*s become larger than 1 in the final starting periods. This is not in line with non-behavioral discounting. Discount factors larger than 1 imply that firms place disproportionate importance on future outcomes in the final starting periods. We use the behavioral discount factors to explain market prices and coordination behavior in regressions pooling both treatments; thus, we do not normalize them. The standard critical discount factors of the hypothetical indefinitely repeated game are $\delta^* \geq \frac{1-0.5}{1-0.04} \cong 0.520833$ in Sym and $\delta^* \geq \frac{1-0.18}{1+0.02} \cong 0.520833$ in WTA; thus, the standard discount factors have the same ordering as the behavioral discount factors and the hypotheses in Section 4 do not change for an infinitely repeated experiment.

[11] Bruttel (2009) shows experimentally that the degree of stability of a cartel is lower with a higher critical discount factor (in an indefinitely repeated Bertrand experiment).

[12] Partner matching ensures that we can expect more collusive behavior than when applying stranger matching. The relatively high number of 40 periods limits endgame effects (i.e., the breakdown of collusion in the last periods of the experiment when duration is known; see Selten and Stoecker 1986). Normann and Wallace (2012) compared different termination rules and found that finitely repeated Prisoner Dilemma games do not hinder cooperation.

the other firm's decision. These tables were part of the experimental instructions and could be used by the subjects when taking decisions. Subjects were told that negative numbers stand for losses. We gave the subjects an initial endowment of 1.50 euros to cover possible losses. After reading the instructions on-screen, all subjects had to answer a set of six control questions correctly to proceed (and this was made known to all subjects, too).[13]

When entering their decision on the computer, subjects used an on-screen slider. Subjects chose a price between 0 and 9 (in steps of 0.25) using the slider and were shown, below the slider, the profit consequences of the chosen price (and this information was updated in real-time when sliding to a different price). After confirming their decision, subjects received feedback: their price, the other firm's price, their period profit, and their cumulated profit over all periods. At the end of the experiment, we summed up all periods' profits (and losses) and the initial endowment. If subjects incurred overall losses, their profits were capped at zero euros.

After the main part of the experiment, we conducted a questionnaire. We asked for the subjects' gender (male/female/non-binary), elicited their cognitive abilities by asking for their final math grade in school (on a range from 1 to 5, best to worst grade in the German system), asked if they study or have studied a field with compulsory economics courses (yes/no), and asked if subjects have experience with game theory (yes/no). We also conducted a simple, monetarily incentivized task to elicit the subjects' risk aversion (which is due to Gneezy and Potters 1997): each subject was endowed with one euro and decided about the amount (in 10-cent increments) to invest in a risky asset—with 50% probability the investment was lost, with 50% probability the investment paid out 2.5-fold.

The experiment was conducted in seven sessions at PLEx, the Potsdam Laboratory for Economic Experiments, between April 11 and April 19, 2022. Economic experiments are not subject to IRB approval at the University of Potsdam, Germany. We invited subjects using a pre-existing database (based on ORSEE; Greiner 2015). The subjects were exclusively students from various fields of study from the University of Potsdam and other nearby universities. Every subject took part in only one session. Upon arrival at the laboratory, all subjects were seated at computer workstations with privacy walls. Communication between subjects was prohibited. The computer programs were coded in z-Tree (Fischbacher, 2007). After the experiment, we paid all subjects in cash and in private. On top of the payoff from the incentivized parts of the experiment, subjects received a show-up fee of five euros (which was not offset against potential losses from the main part of the experiment).

Before collecting the data, we pre-registered the number of observations to collect, our hypotheses, and the data analysis plan. We collected 25 independent observations (duopolies) in each of the two treatments (100 subjects in total).[14] We conducted seven sessions, each lasting between 60 and 80 min. On average, the subjects earned 12.59 euros (median 12.40, minimum 1.80, maximum 27.70), not including the show-up fee. Table A.4 in the Appendix shows subjects' sociodemographic characteristics from the post-experimental questionnaire. According to the elicited variables, the two treatments are balanced.

4. Hypotheses

While our stage-game prediction is that mean set prices in Sym are slightly higher than in WTA (see Table 1), we expect market prices (the minimum of prices set by both firms) to be higher in WTA than in Sym. We base this expectation on two arguments: (i) The calculated critical behavioral discount factors in Sym are higher than the ones in WTA in all periods, indicating a lower stability of collusion in Sym, and (ii) the fact that the incentive structure in WTA inherently encourages aiming for the maximum possible profit that coordination can achieve, given the absence of a loss in efficiency when sharing market demand.[15] The latter argument opens the following channel in WTA: Coordination on one price is in expectation more profitable regardless of the price, and both firms simultaneously set the same price that ensures the maximum profit, either 8.75 or 9.00, and let the random draw decide about the receiver of the entailed profit. This is because, in WTA, the entire market demand is served by one firm (chosen randomly if prices are equal), making the highest prices the most attractive targets for coordination. Repeating this ensures equal (expected) profits for both firms. Hypothesis 1 describes the expected market outcome:

Hypothesis 1. Market prices in WTA are *higher* than in Sym.

Hypothesises 2a and 2b describe the underlying mechanism (where Hypothesis 2a serves as the necessary condition, and Hypothesis 2b as the sufficient condition; only coordination on high enough prices increases market prices):

Hypothesis 2a. Subjects in WTA more often set equal prices than subjects in Sym.

[13] Section B in the Appendix provides the instructions, the quizzes, example screenshots of the decision stage of both treatments, their translations, and the payoff tables as handed out to the subjects.

[14] We based this number of observations on previous experiments in the related literature: (i) Abbink and Brandts (2008) collected between 8 and 10 independent observations per treatment with 50 periods of interaction per independent observation; (ii) Puzzello (2008) between 7 and 17 independent observations per treatment (60 periods of interaction); (iii) Argenton and Müller (2012) between 19 and 23 independent observations (40 periods of interaction).

[15] Assume that the two firms tacitly coordinate on a price $\bar{p}$. In Sym, a firm's expected profit is $\mathbb{E}[\pi_i(p_i = \bar{p}, p_{-i} = \bar{p})] = \frac{1}{2}\ \bar{p}\ D(\bar{p}) - C(\frac{1}{2}D(\bar{p}))$ (half the revenue minus the cost for half the market demand with certainty). In WTA, it is $\mathbb{E}[\pi_i(p_i = \bar{p}, p_{-i} = \bar{p})] = \frac{1}{2}[\bar{p}\ D(\bar{p}) - C(D(\bar{p}))]$ (serving the entire market with 50% probability, a profit of zero otherwise). In our experiment, the expected profit in each period in WTA is 0.32 euros larger than in Sym for all prices. Our line of reasoning is the following: (i) When, all other things being equal, coordination is more profitable in WTA, we expect to observe more coordination in this treatment. (ii) When coordination is more profitable at higher prices, we should observe more coordination at these prices.

Table 2
Summary statistics of key variables in the two treatments.

	Sym		WTA	
	Mean	Median	Mean	Median
Market prices	5.34 (0.28)	5.00	4.89 (0.30)	4.25
Periods with equal prices (per duopoly)	5.52 (0.93)	4.00	11.36 (2.18)	8.00
Periods with equal prices of 8.75 or 9.00 (per duopoly)	0.64 (0.36)	0.00	2.00 (1.36)	0.00
Periods with alternation between 8.75 and 9.00 (per duopoly)	3.44 (1.82)	0.00	1.88 (1.31)	0.00

Note: Data is aggregated on the duopoly level. Standard errors in parentheses (clustered at the duopoly level).

Hypothesis 2b. Subjects in WTA more often set prices of 8.75 or 9.00 than subjects in Sym.

Another strategy that might counteract the effect of coordination on market prices is alternation.[16] Alternation in our experiment means that one firm (intentionally) sets a price of 9.00 while the other firm slightly undercuts this price. Subjects take turns between periods and thus share the monopoly profit equally; this can be applied in both treatments. Therefore, alternation could serve as a viable *intertemporal* strategy for firms to achieve similar outcomes to the *intratemporal* strategy of direct coordination (as sequential rather than simultaneous decisions might be easier to coordinate). Alternation can be considered a strategic response to the experimental conditions that might emerge when direct coordination is challenging. We consider this strategy in our exploratory analysis as examining the role of alternation alongside the main hypotheses might help us better understand market dynamics under different sharing rules.

5. Results

5.1. Data overview and hypotheses tests

Table 2 shows summary statistics for the key variables. For our hypotheses tests, we aggregate our data at the duopoly level and calculate the mean market prices and the sum of periods with equal prices for each of the 50 duopolies. First, we compare the mean market prices between the two treatments. In contrast to our hypothesis, market prices in Sym are higher than in WTA. A one-sided two-sample Wilcoxon-Mann–Whitney test suggests that the evidence supporting a positive shift of average market prices in WTA compared to Sym is not strong enough (exact p-value, $p = 0.8927$).[17]

Result 1. Market prices, in contrast to our prediction, are not higher under the winner-takes-all sharing rule compared to the symmetric sharing rule.

Next, we compare the number of periods in which the two firms set an equal price. The mean shows that subjects set an equal price in 28.4% of periods in WTA and in 13.8% of periods in Sym.[18] In WTA, the mean and the median are at least twice as large as in Sym. A one-sided two-sample Wilcoxon-Mann–Whitney test finds strong evidence supporting a positive shift of coordination on one price in WTA, compared to Sym (exact p-value, $p < 0.01$).

Result 2a. As predicted, we observe a higher incidence of tied prices under the winner-takes-all sharing rule than under the symmetric sharing rule.

Finally, we compare the number of periods in which the firms coordinate on a price of 8.75 or 9.00, the prices where duopoly profits are highest. The mean and median in both treatments are much smaller than the unrestricted variable. While the median is zero in both treatments, the mean in WTA is thrice as large as in Sym (coordination drops to 1.6% of periods in WTA and 5.0% in Sym). However, a one-sided two-sample Wilcoxon-Mann–Whitney test suggests that the evidence supporting a positive shift of coordination on high prices in WTA compared to Sym is not strong enough (exact p-value, $p = 0.2328$).[19] Puzzello (2008, p. 172)

[16] See Amelio and Biancini (2010), Lau and Mui (2008, 2012), Duffy et al. (2017), Sibly and Tisdell (2018), and Riyanto and Roy (2019) for real-world examples of this strategy and evidence from the lab.

[17] To complement the one-sided test with aggregated data, we also show the result of a clustered two-sided two-sample Wilcoxon-Mann–Whitney test (using the `clusrank` package by Jiang et al. 2020). This test is not part of the pre-registration. It uses the full set of observations (2000 round observations in 50 clusters instead of 50 aggregated duopoly observations) and is highly significant ($p < 2.2e{-}16$, Datta and Satten method, exact p-value).

[18] We compare these numbers to the share of same-price periods in other studies. In Argenton and Müller's experiment, depending on their treatment, the authors report between 65.7% and 75.9% of same-price periods in the first 37 periods (excluding the final three periods to account for endgame effects). This high level of coordination is not surprising, as convex costs can make serving the market alone costly compared to sharing it.

[19] We also compare mean prices, conditioned on both firms setting equal prices. In Sym, this is 5.24 (with a standard error of 0.17); in WTA 5.03 (0.14). A two-sided, two-sample Wilcoxon-Mann–Whitney test is not significant (exact p-value, $p = 0.7639$). In Figure A.1 in the Appendix, we show frequency distributions of the number of occurrences of equal prices in both treatments. We observe that in Sym, equal prices are relatively evenly distributed over the price range. In contrast to that, we see that in WTA, there is one peak around the stage-game Nash equilibrium (at 2.75–3.00) and another, lower peak, around the monopoly price (at 8.75–9.00). The coordination on competitive prices offsets the effect of coordination on high prices on the mean.

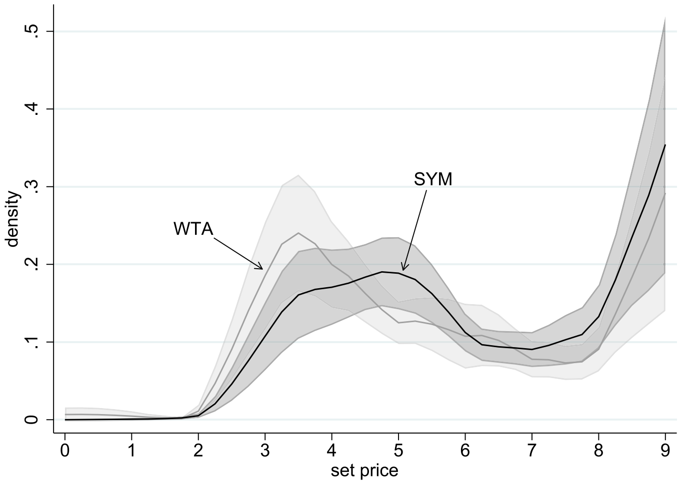

Fig. 4. Kernel estimates of set prices in both treatments (with confidence intervals).
Note: The kernel estimates use an Epanechnikov kernel, adjusted for a lower boundary of 0, an upper boundary of 9, and an evaluation grid size of 37. The 95% confidence intervals use bootstrapped standard errors (with 200 repetitions), clustered at the duopoly level.

reports a higher incidence of tacit collusion under the symmetric than under the winner-takes-all rule for both the fine and the coarse price grid; this might be driven by constant returns to scale (and identical expected profits) in her experiment.

Result 2b. In contrast to our prediction, under the winner-takes-all sharing rule, we do not observe significantly more periods where firms jointly set a price that maximizes profits (compared to the symmetric sharing rule).

As the initial tests used aggregated data, we also visualize the distributions of both firms' set prices and supply kernel estimates in the two treatments in Fig. 4. We observe the following: (i) There is considerable price dispersion. (ii) The kernel estimate in WTA seems to have a peak around the predicted price of 3.00, whereas in Sym, there seems to be a plateau between circa 3 and 5. (iii) Both treatments' set prices peak at the end of the price range, indicating (attempts of) collusive behavior (where one firm intentionally sets a very high price so the other firm can undercut slightly below). (iv) However, the 95% confidence intervals of the two treatments overlap over the entire price range.[20]

We show market prices' evolution over the 40 periods in Fig. 5. The mean market prices in Sym and WTA are above the mean Nash predictions, as also reported in Puzzello (2008, p. 171). In both treatments, we observe an inverse U-shaped curve of the quadratic fit of market prices.[21] Other studies also report this form (e.g., in Argenton and Müller 2012): After an initial increase in collusive behavior, this form of cooperation between firms gradually wears off because of end game effects; however, Puzzello (2008, p. 171) reports an unstable pattern of mean and median posted prices under both sharing rules (which can be explained by the underlying model's capacity constraint and the resulting Edgeworth cycling). When we compare the mean market prices of the first and second half of the experiment between the treatments, we detect weak evidence that the divergence of mean prices between treatments increases over time: a two-sided two-sample Wilcoxon-Mann–Whitney test detects a difference significant at the 10%-level for the second half (for periods 1–20: $p = 0.6581$, for periods 21–40: $p = 0.0701$; exact p-values). We observe significantly more equal prices in WTA than in Sym in the first and second half of the experiment, though at different levels of significance (using a one-sided two-sample Wilcoxon-Mann–Whitney test, $p = 0.0971$ in the first half, $p = 0.0099$ in the second half, exact p-values). Given the treatment effect of time-dependent differences in prices and coordination, we later report robustness checks that control for the time dimension.

Can the critical behavioral discount factors explain the observed endgame effect/price divergence? In Table 3, we regress market prices and periods with an equal price-dummy on the calculated critical discount factors. We observe that, as predicted, a higher critical discount factor is significantly related with lower market prices after accounting for the WTA-dummy and the number of periods. However, we do not find a significant effect on setting equal prices.

[20] To complement the kernel estimates, we also supply histograms that allow us to compare the fractions of observed set prices with their one-shot Nash predictions. Figure A.2 in the Appendix shows histograms of the set prices in both treatments, 95% confidence intervals (based on linear probability models and clustered standard errors), and weights predicted by Nash equilibria (in Sym). In Sym, shown in Figure A.2a, we can observe that only for two prices (3.75 and 4.25), fractions of set prices are indistinguishable from predictions. In WTA, shown in Figure A.2b, less than 10% of observed set prices lie on the predicted price of 3.00.

[21] Table A.5 in the Appendix supplies the OLS regressions of the quadratic fit shown in Fig. 5. All regression coefficients differ significantly from zero but not significantly between the two treatments.

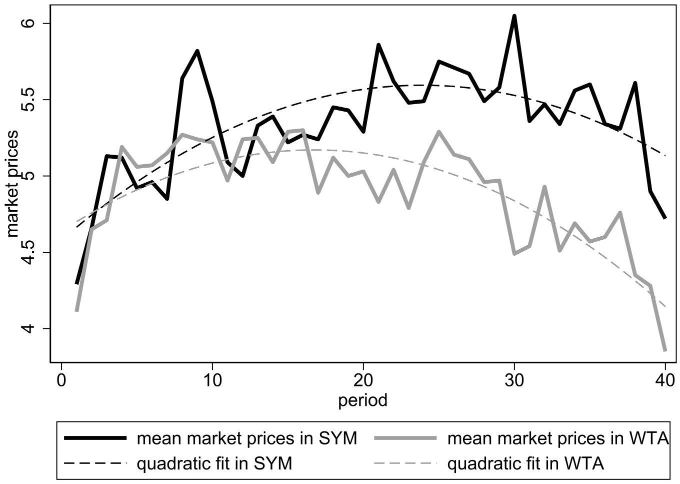

Fig. 5. Market prices' evolution over time in both treatments.

Table 3
The effect of the behavioral discount factor on market prices and periods with equal prices.

	DV: market price			DV: equal price-*dummy*		
	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6
Discount factor	0.116	−0.148	−0.187**	−0.673	0.051	0.070
	(0.248)	(0.100)	(0.084)	(0.637)	(0.236)	(0.241)
WTA-*dummy*		−0.494	−0.509		0.783**	0.791**
		(0.408)	(0.413)		(0.326)	(0.320)
Period			0.004			−0.002
			(0.009)			(0.007)
Constant	5.050***	5.492***	5.445***	−1.679***	−2.639***	−2.618***
	(0.282)	(0.277)	(0.287)	(0.533)	(0.337)	(0.347)
# obs.	1,950	1,950	1,950	1,950	1,950	1,950
# clusters	50	50	50	50	50	50
adj. R^2/pseudo R^2	0.000	0.012	0.012	0.005	0.019	0.019

Note: Models 1–3 show OLS regressions (with adj. R^2), Models 4–6 logit regressions (with pseudo R^2). ***, **, and * show difference from zero at the 1%-, 5%-, and 10%-level, respectively. Standard errors clustered at the duopoly level.

5.2. Price alternation as a substitute collusion strategy

This section investigates the role of price alternation as a potential collusion strategy in duopolies and its prevalence across the two treatments. Price alternation, defined as duopolists taking turns to set lower prices (specifically, one firm undercuts at 8.75 while the other firm sets a price of 9.00), is explored as a substitute for direct price coordination.

The final row of Table 2 shows the number of alternation periods. While the median is zero in both treatments, the mean in Sym is almost twice as large as in WTA (8.6% vs. 4.7% of periods). A one-sided two-sample Wilcoxon-Mann–Whitney test suggests evidence that the difference in alternation periods between the two treatments is not strong enough (exact p-value, $p = 0.26985$).

Figure A.3 in the Appendix presents a scatter plot where, for each duopoly, we show the number of periods with equal prices and the number of alternation periods. In WTA, there are only two duopolies with a positive number of alternations (8% of independent observations); in Sym, there are four (16%). The rare usage of alternation contrasts with the frequency of equal prices. The figure also shows the estimated linear relationship between the two variables. The relationship between the two strategies in both treatments is negative: the correlation coefficient between the two strategies is $\rho = -0.371$ (the difference from zero is significant at $p = 0.0682$) in Sym and $\rho = -0.184$ ($p = 0.3800$) in WTA.

A comparison with the related literature, such as the study by Sibly and Tisdell (2018), highlights the challenges of using alternation in environments with a large strategy space, like price-setting experiments, compared to binary decision contexts such as the Prisoner's Dilemma. Our results indicate a negative correlation between alternation and price coordination, suggesting these strategies are used as substitutes in our experimental setting but only infrequently.

5.3. Past outcomes as determinants of pricing behavior

Coordination on one price is difficult and depends on the sharing rule. This section examines intertemporal competition and price coordination, exploring how subjects adjust their prices in response to past market outcomes and the differences between the two

Table 4
Contingency table of pricing as a response to past period's market outcome in Sym.

Price in period t is...	Market in period $t-1$ was...			
	Lost	Shared	Won	Total
Lower	639	55	226	920
	70.45%	40.44%	24.92%	47.18%
Equal	108	31	110	249
	11.91%	22.79%	12.13%	12.77%
Higher	160	50	571	781
	17.64%	36.76%	62.95%	40.05%
Total	907	136	907	1,950
	100.00%	100.00%	100.00%	100.00%

Note: Percentages report relative column frequency.

Table 5
Contingency table of pricing as a response to past period's market outcome in WTA.

Price in period t is...	Market in period $t-1$ was...			
	Lost	Shared	Won	Total
Lower	576	88	216	880
	68.57%	32.59%	25.71%	45.13%
Equal	135	114	157	406
	16.07%	42.22%	18.69%	20.82%
Higher	129	68	467	664
	15.36%	25.19%	55.60%	34.05%
Total	840	270	840	1,950
	100.00%	100.00%	100.00%	100.00%

Note: Percentages report relative column frequency.

sharing rules. This analysis, in general, can help us better understand tacit collusion and, more specifically, can inform regulators and market designers about how sharing rules might affect market outcomes intertemporally.

In the first approach, we categorize the price adjustment and determine if a subject sets a lower, equal, or higher price than in the previous period (thereby neglecting the size of the price adjustment). Then, we categorize market outcomes and determine if a subject won the market (by having set the lower price), shared it with the other subject (by having set the same price), or lost the market. Tables 4 and 5 show three-by-three contingency tables, where the categorical pricing behavior in one period is correlated with the categorical market outcome of the previous period, separately for each treatment.[22] We observe the following:

(i) Dynamic price competition—in both treatments, subjects set a lower (higher) price after losing (winning) the market in the majority of decisions: 70.45% (68.57%) of prices are lowered after losing a market in Sym (WTA); 62.95% (55.6%) of prices are increased after winning a market in Sym (WTA).
(ii) Increased market sharing in WTA (a reformulation of Result 2a)—in periods 1–39, the number of shared markets is almost twice as high in WTA as in Sym (270 vs. 136 prices); in periods 2–40, the number of periods in which both firms set the same price is higher in WTA as in Sym (406 vs. 249 prices).
(iii) Enhanced coordination in WTA—the probability of setting the same price as a response to a previously shared market is almost twice as high in WTA as in Sym (42.22% vs. 22.79%).[23] Regressions confirm this finding, as shown in Table 6.[24] Similarly, Puzzello (2008, pp. 173–174) reports that the effect of coordinating in the previous period on current coordination is positive, as well, but stronger under the symmetric sharing rule.

Next, we consider a richer model where we link the *size* of the price adjustment between two periods to previous market outcomes. Besides the main outcome, whether a market was won or lost (sharing the market serves as a baseline outcome now), we also consider the difference between the two firms' prices in the same period. Table 7 shows the results from OLS regressions of the two treatments. According to Model 1, subjects in both treatments react significantly stronger to lost markets than to won markets with their price adjustment (Wald tests reject the equality of the two coefficients with $p < 0.0001$ in both regressions). Model

[22] In the Appendix, we supply the results of this categorization exercise separately for the first and second half of the experiment and both treatments, in Tables A.6–A.9. These results do not differ qualitatively from the ones shown here.

[23] This ignores the result of the random draw that determines who serves the market in WTA. In Table A.10 in the Appendix, we show pricing behavior as a response to the random draw when prices are equal in WTA. More lowered (increased) prices follow a favorable (unfavorable) random draw (67.05% and 57.35%). However, Table A.10 only shows 270 pricing decisions.

[24] Here, we use OLS regressions because interaction terms in logit and probit regressions have problems (Ai and Norton, 2003). Two further tables in the Appendix show regressions with a dummy variable for the second half of the experiment (Table A.11) and with period variables (Table A.12) as robustness checks. They confirm our findings (the coefficients controlling for the time dimension are not significantly different from zero, and the coefficients of the variables of interest are not different from the ones in Table 6).

Table 6
Coordination on equal prices between treatments (OLS regressions).

DV: equal price-*dummy*	Model 1	Model 2	Model 3
L.market shared-*dummy*	0.211***	0.199***	0.108*
	(0.051)	(0.048)	(0.057)
WTA-*dummy*		0.067**	0.054*
		(0.031)	(0.032)
L.market shared-*dummy* * WTA-*dummy*			0.141*
			(0.083)
constant	0.146***	0.047	0.120***
	(0.016)	(0.024)	(0.021)
# obs.	3,900	3,900	3,900
# clusters	50	50	50
adj. R^2	0.030	0.037	0.041

Note: ***, **, and * show difference from zero at the 1%-, 5%-, and 10%-level, respectively. Standard errors clustered at the duopoly level.

Table 7
Set prices' adjustment reaction to past outcomes (OLS regressions).

DV: set price difference	Model 1			Model 2		
	Sym	WTA	Difference	Sym	WTA	Difference
L.lost	−0.998***	−0.912***	$p = 0.645$	−0.513***	−0.350***	$p = 0.325$
	(0.136)	(0.130)		(0.110)	(0.123)	
L.won	0.409**	0.406***	$p = 0.986$	−0.193	−0.145	$p = 0.794$
	(0.148)	(0.098)		(0.132)	(0.132)	
L.lost * L.abs(price span)				−0.326***	−0.375***	$p = 0.599$
				(0.063)	(0.068)	
L.won * L.abs(price span)				0.405***	0.367***	$p = 0.687$
				(0.069)	(0.063)	
constant	0.272**	0.194**	$p = 0.564$	0.272**	0.194**	$p = 0.564$
	(0.104)	(0.085)		(0.104)	(0.85)	
# obs.	1,950	1,950		1,950	1,950	
# clusters	25	25		25	25	
adj. R^2	0.164	0.134		0.243	0.227	

Note: *** and ** show difference from zero at the 1%- and 5%-, respectively. Standard errors clustered at the duopoly level.

2 also accounts for the previous period's price span (the difference between the two firms' set prices).[25] We see that the subjects' reaction to a one-point difference between previously set prices is of comparable size for won and lost markets. Still, the response to a lost market, corrected for the price span, is significantly stronger than the reaction to a won market (again, Wald tests reject the equality of the two coefficients with $p = 0.0005$ in both regressions).[26] However, we cannot detect significant differences between the coefficients of the two treatments in both models.

Our analysis draws parallels with Fonseca and Normann (2012), who investigate price adjustment dynamics in experimental Bertrand markets with pre-competition chat communication. While their study emphasizes the role of communication in establishing and maintaining (explicit) collusion, our focus on a non-communicative setting shows that subjects can collude tacitly by coordinating on equal prices and that the sharing rule makes a difference in how they coordinate. When we examine the size of price adjustments, we observe an asymmetry in the subjects' reactions to won and lost markets—losing in the previous period triggers a larger response than winning.

6. Discussion & conclusion

We conducted an experiment where we varied the sharing rule in a Bertrand market with increasing returns to scale. As in previous research on Bertrand competition with decreasing returns to scale (Abbink and Brandts 2008, Argenton and Müller 2012), we also find pricing above (mean) Nash predictions regardless of the sharing rule and an inverse U-shaped curve of market prices over time. Our findings differ in some aspects from the ones reported in Puzzello (2008), who reported more coordination on the monopoly price under the symmetric sharing rule. While we find that subjects under the winner-takes-all rule coordinate more than twice as often on one price compared to the symmetric sharing rule, we do not find that this increases market prices (the

[25] Figure A.4 in the Appendix shows, separately for each treatment, that the price adjustment in period t is closely related to the price span in period $t-1$.

[26] In the Appendix, we show regressions extended by a second-half dummy (Table A.13) and by period variables (Table A.14). Again, the results are not qualitatively different from the ones presented here. As an additional robustness check, in regressions not shown here, we also included the past period's ***profit*** in the two models in Table 7. The coefficients in all specifications are negative and significantly different from zero. After including the lagged profit variable, we still observe that subjects in both treatments react more strongly to lost markets than to won markets when adjusting their prices.

increasing divergence of market prices over time, and a rank-sum test using non-aggregated observations from all periods instead of the pre-registered aggregated test suggest that the market price under WTA might be significantly larger than under Sym, especially when firms interact longer). We attribute this to two factors: (i) Subjects under the winner-takes-all sharing rule do *not* coordinate more frequently on the prices that grant them monopoly payoffs. (ii) Although not definitive, some evidence suggests that alternation and coordination on the same price are substitute strategies under the symmetric sharing rule. So, while subjects in the winner-takes-all treatment do not use coordination on the same price to the full extent, at least some subjects in the symmetric treatment use alternation to increase profits and equalize payoffs. These two effects might offset each other. However, as we aggregate data at a high level and work with few observations here, this conclusion has to be taken with a grain of salt.

When we examine the intertemporal coordination process in further analyses, we find that coordinating on a price after a period with equal prices is more likely in the winner-takes-all treatment than in the symmetric treatment. Furthermore, we document a plausible competition process: If subjects lose (win) a market, they react by price decreases (increases) in both treatments. When we examine a model that includes the previous period's price span, we see that the reaction to a lost market corrected for the price span is stronger than the reaction to a won market. It would be interesting to examine the intertemporal coordination process in Bertrand markets with decreasing costs, where coordination on one price can be profit-increasing for firms in the same period (and is not used to even out profits between periods).

Future research could examine the role of sharing rules in Bertrand markets with different returns to scale in a unified experimental framework (where the instructions, user interface, etc., are constant). Possible experimental design modifications include: (i) Subjects participate in more than one supergame, each against a different opponent. This could speed up learning, and coordination on equal prices might increase in later supergames. However, Mengel et al. 2022 showed that the realized number of interactions in indefinitely repeated games can influence behavior in later interactions, leading to potential experimental design challenges; a finitely repeated game eliminates similar variations due to differing game lengths in indefinitely repeated game designs. (ii) One could consider an indefinitely repeated version of the experiment with stronger predictions. For example, the "alternating" strategy might require a higher discount factor than coordinating on the same price (because alternation evens out profits intertemporally while coordination on the same price has an intratemporal effect).

Declaration of Generative AI and AI-assisted technologies in the writing process

During the preparation of this work, the author used Grammarly and ChatGPT to improve the readability of the manuscript. After using these tools/services, the author reviewed and edited the content as needed and takes full responsibility for the content of the published article.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.jebo.2025.106968.

Data availability

Data and scripts are available in a public OSF repository: https://doi.org/10.17605/osf.io/jnufw.

References

Abbink, K., Brandts, J., 2008. 24 — Pricing in Bertrand competition with increasing marginal costs. Games Econom. Behav. 63 (1), 1–31.
Ai, C., Norton, E.C., 2003. Interaction terms in logit and probit models. Econom. Lett. 80 (1), 123–129.
Amelio, A., Biancini, S., 2010. Alternating monopoly and tacit collusion. J. Ind. Econ. 58 (2), 402–423.
Argenton, C., Müller, W., 2012. Collusion in experimental Bertrand duopolies with convex costs: the role of cost asymmetry. Int. J. Ind. Organ. 30 (6), 508–517.
Baye, M.R., Kovenock, D., 2008. Bertrand competition. In: Durlauf, S.N., Blume, L.E. (Eds.), The New Palgrave Dictionary of Economics, second ed. Palgrave Macmillan, pp. 1–7.
Baye, M.R., Morgan, J., 2002. Winner-take-all price competition. Econom. Theory 19 (2), 271–282.
Blinder, A., Canetti, E.R., Lebow, D.E., Rudd, J.B., 1998. Asking About Prices: A New Approach to Understanding Price Stickiness. Russell Sage Foundation.
Brookins, P., Ryvkin, D., Smyth, A., 2021. Indefinitely repeated contests: an experimental study. Exp. Econ. 24 (4), 1390–1419.
Brown-Kruse, J.L., 1991. Contestability in the presence of an alternate market: an experimental examination. Rand J. Econ. 22 (1), 136–147.
Bruttel, L.V., 2009. The critical discount factor as a measure for cartel stability? J. Econ. 96, 113–136.
Chaudhuri, P.R., 1996. The contestable outcome as a Bertrand equilibrium. Econom. Lett. 50 (2), 237–242.
Chowdhury, P.R., 2002. Limit-pricing as Bertrand equilibrium. Econom. Theory 19 (4), 811–822.
Dal Bó, P., 2005. Cooperation under the shadow of the future: experimental evidence from infinitely repeated games. Am. Econ. Rev. 95 (5), 1591–1604.
Dal Bó, P., Fréchette, G.R., 2018. On the determinants of cooperation in infinitely repeated games: a survey. J. Econ. Lit. 56 (1), 60–114.
Dastidar, K.G., 1995. On the existence of pure strategy Bertrand equilibrium. Econom. Theory 5 (1), 19–32.
Dixon, H., 1984. The existence of mixed-strategy equilibria in a price-setting oligopoly with convex costs. Econom. Lett. 16 (3–4), 205–212.
Duffy, J., Lai, E.K., Lim, W., 2017. Coordination via correlation: an experimental study. Econom. Theory 64 (2), 265–304.

Durham, Y., McCabe, K., Olson, M.A., Rassenti, S., Smith, V., 2004. Oligopoly competition in fixed cost environments. Int. J. Ind. Organ. 22 (2), 147–162.
Edwards, R., Routledge, R., 2023. Existence and uniqueness of Nash equilibrium in discontinuous Bertrand games: a complete characterization. Int. J. Game Theory 52, 569–586.
Fischbacher, U., 2007. Z-tree: Zurich toolbox for ready-made economic experiments. Exp. Econ. 10 (2), 171–178.
Fonseca, M.A., Normann, H.T., 2012. Explicit vs. tacit collusion—the impact of communication in oligopoly experiments. Eur. Econ. Rev. 56 (8), 1759–1772.
Gneezy, U., Potters, J., 1997. An experiment on risk taking and evaluation periods. Q. J. Econ. 112 (2), 631–645.
Greiner, B., 2015. Subject pool recruitment procedures: organizing experiments with ORSEE. J. Econ. Sci. Assoc. 1 (1), 114–125.
Hoernig, S.H., 2002. Mixed Bertrand equilibria under decreasing returns to scale: an embarrassment of riches. Econom. Lett. 74 (3), 359–362.
Hoernig, S.H., 2007. Bertrand games and sharing rules. Econom. Theory 31 (3), 573–585.
Jiang, Y., Lee, M.L.T., He, X., Rosner, B., Yan, J., 2020. Wilcoxon rank-based tests for clustered data with R package clusrank. J. Stat. Softw. 96 (6), 1–26.
Lau, S.H.P., Mui, V.L., 2008. Using turn taking to mitigate coordination and conflict problems in the repeated battle of the sexes game. Theory and Decision 65, 153–183.
Lau, S.H.P., Mui, V.L., 2012. Using turn taking to achieve intertemporal cooperation and symmetry in infinitely repeated 2×2 games. Theory and Decision 72, 167–188.
Lugovskyy, V., Puzzello, D., Sorensen, A., Walker, J., Williams, A., 2017. An experimental study of finitely and infinitely repeated linear public goods games. Games Econom. Behav. 102, 286–302.
Maskin, E., Riley, J., 2000. Equilibrium in sealed high bid auctions. Rev. Econ. Stud. 67 (3), 439–454.
McKelvey, R.D., McLennan, A.M., Turocy, T.L., 2016. Gambit: software tools for game theory, version 16.0.1. URL http://www.gambit-project.org.
Mengel, F., Orlandi, L., Weidenholzer, S., 2022. Match length realization and cooperation in indefinitely repeated games. J. Econom. Theory 200, 105416.
Nash, J., 1951. Non-cooperative games. Ann. Math. 54 (2), 286–295.
Normann, H.T., Wallace, B., 2012. The impact of the termination rule on cooperation in a Prisoner's Dilemma experiment. Int. J. Game Theory 41 (3), 707–718.
Orland, A., Roos, M.W., 2013. The New Keynesian Phillips curve with myopic agents. J. Econom. Dynam. Control 37 (11), 2270–2286.
Orland, A., Roos, M.W., 2019. Price-setting with quadratic adjustment costs: Experimental evidence. J. Econ. Behav. Organ. 163, 88–116.
Puzzello, D., 2008. Tie-breaking rules and divisibility in experimental duopoly markets. J. Econ. Behav. Organ. 67 (1), 164–179.
Riyanto, Y.E., Roy, N., 2019. Path of intertemporal cooperation and limits to turn-taking behavior. J. Econ. Behav. Organ. 165, 21–36.
Robidoux, B., Lester, J., 1992. Econometric estimates of scale economies in Canadian manufacturing. Appl. Econ. 24 (1), 113–122.
Roth, A., 1995. Introduction to experimental economics. In: Kagel, J., Roth, A. (Eds.), Handbook of Experimental Economics. Princeton University Press, pp. 3–109.
Roth, A.E., Murnighan, J.K., 1978. Equilibrium behavior and repeated play of the Prisoner's Dilemma. J. Math. Psych. 17 (2), 189–198.
Selten, R., Stoecker, R., 1986. End behavior in sequences of finite Prisoner's Dilemma supergames: a learning theory approach. J. Econ. Behav. Organ. 7 (1), 47–70.
Sheremeta, R.M., 2018. Behavior in group contests: a review of experimental research. J. Econ. Surv. 32 (3), 683–704.
Sibly, H., Tisdell, J., 2018. Cooperation and turn taking in finitely-repeated Prisoners' Dilemmas: an experimental analysis. J. Econ. Psychol. 64, 49–56.
Spiegler, R., 2011. Bounded Rationality and Industrial Organization. Oxford University Press.