Journal of Economic Behavior and Organization 192 (2021) 584–599

Contents lists available at ScienceDirect

Journal of Economic Behavior and Organization

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

Imperfect tacit collusion and asymmetric price transmission☆

Muhammed Bulutay[a], David Hales[b,*], Patrick Julius[c], Weiwei Tasch[d]

[a] Chair of Macroeconomics, Technische Universität Berlin, Germany
[b] Department of Economics, University of California at Santa Barbara, United States
[c] Department of Economics, University of California, Irvine, United States
[d] Philosophy, Politics, and Economics Program, University of Pennsylvania, United States

ARTICLE INFO

Article history:
Received 31 March 2021
Revised 26 August 2021
Accepted 19 October 2021
Available online 14 November 2021

JEL classification:
D43
L13
C92
C72
C73

Keywords:
Asymmetric price transmission
Tacit collusion
Oligopolistic competition
Market power
Rockets and feathers

ABSTRACT

We investigate asymmetric price transmission (APT) in laboratory experiments and find that imperfect tacit collusion is likely the cause in our otherwise frictionless markets. We vary the number of sellers across markets to evaluate the role competition plays in APT. We report similar magnitudes of asymmetry in markets with 3, 4, 6, and 10 sellers, but not in duopolies. Furthermore, sellers consistently set their prices above the best-response levels implied by their forecasts, particularly in periods following negative shocks. We interpret these pricing deviations as sellers' intentions to collude, and note that they mechanically drive the pricing asymmetries we observe.

1. Introduction

The phenomenon of Asymmetric Price Transmission (APT), that is, that supplier prices rise quickly after positive input cost shocks but fall relatively more slowly after similarly-sized negative shocks, has been repeatedly documented in the literature such that we can rightly describe it as a stylized fact.[1] However, while empirical evidence for the existence of

☆ This article is dedicated to the memory of Stephen P.A. Brown, who proposed this research effort and who provided tremendous mentorship and guidance. He is sorely missed. We also benefited enormously from the helpful advice of Frank Heinemann, Youssef Benzarti, Ted Bergstrom, Tim Cason, Gary Charness, John Duffy, Ted Frech, Uri Gneezy, Antoine Dib, Rod Garratt, Rosemarie Nagel, Gabriele Camera, Cheng-Zhong Qin, Shyam Sunder, Ryan Oprea, Emanuel Vespa, and Sevgi Yuksel. We are also indebted to the organizers and participants of the 11th BESLab Experimental Economics Summer School in Macroeconomics, the 7th Spring School in Behavioral Economics, the 2019 International ESA Meeting, 13th Maastricht Behavioral and Experimental Economics Symposium and the 11th International Workshop on Theoretical and Experimental Macroeconomics for insightful discussions. The authors gratefully acknowledge financial support from the Joachim Herz Foundation. Experimental research using human subjects was conducted under the authority of IRB protocol 3-18-0513, University of California, Santa Barbara, dated December 8, 2017.

* Corresponding author.
E-mail addresses: m.bulutay@tu-berlin.de (M. Bulutay), dhales@ucsb.edu (D. Hales), pjulius@uci.edu (P. Julius), wtasch@upenn.edu (W. Tasch).

[1] This is also sometimes termed as "positive APT" to distinguish it from the opposite phenomenon of "negative APT." In this paper we will simply refer to "APT" when we mean positive APT, except where doing so would create ambiguity. See Section 2.1 for an overview of the evidence.

https://doi.org/10.1016/j.jebo.2021.10.018

APT is ample, identification of its causal forces is not settled. Many theoretical explanations have been proposed, but the empirical literature has yet to conclusively determine which of these are valid or are most influential.

Empirical studies of APT predominantly examine aggregate-level variables (e.g. inflation, concentration) proposed to be relevant in the theory literature. The focus on such variables occurs because firm-level determinants are either not directly observable, or are not adequately measurable in panel data form. This approach yields helpful correlations between such variables, but the effort to identify causal relationships has met with only limited success, most notably in the context of firm-level underpinnings of the phenomenon. While the search for accurate firm-level data should certainly be continued, and where discovered used to further inform our understanding of pricing behavior, experimental methods offer a comparative advantage: testing theories that involve variables which are unobservable in the field (e.g. agents' information sets) lie outside the reach of empirical methods;[2] if however these same variables can be controlled through experimental design, we can overcome this obstacle to testing theory.

A question of primary interest is whether tacit collusion drives APT-like pricing behavior.[3] The field data does not convincingly exclude the possibility that market competitors secretly communicate, given the strong legal and even criminal incentives for firms to conceal – or avoid engaging in – such activities. This provides an obvious challenge for identification and motivates turning to the controlled setting of the laboratory, where we can directly observe competitor behavior and credibly prevent communication between sellers.[4]

An argument put forth by Borenstein et al. (1997) is that a variation of the "trigger price" model of oligopolistic coordination originally introduced by Green and Porter (1984) may explain the emergence of APT-type dynamics through tacit collusion. In their model, when positive shocks occur firms immediately raise prices in order to preserve profit margins; however, when negative shocks occur firms react adaptively, holding prices at pre-shock levels until they see convincing evidence that a rival has cut their prices. Rapidly lowering prices in response to a downward cost shock could be perceived as defection from a mutually beneficial regime of tacit collusion, thus inviting retaliation from other firms. In contrast, rapidly raising prices in response to an upward cost shock poses no such threat to one's competitors, and therefore incurs no corresponding risk of retaliation. Although their arguments are sound and are consistent with a deep empirical literature finding correlations suggestive of tacit collusion, Borenstein et al. (1997) conclude that they are unable to conclusively draw support for this hypothesis from their data. As no other empirical study of which we are aware has accomplished this either, we thus find motivation to turn to the laboratory to examine the role of tacit collusion in driving the APT dynamic.[5]

A second question of interest is whether the number of competing sellers in a given market plays a significant role in the realization of the APT phenomenon. Notably, in his broad study of U.S. wholesale and retail markets, Peltzman (2000) finds a negative relation between the number of competitors in a market and the magnitude of APT observed. As with any empirical study, however, this study does not exclude the possibility that explicit (but unobserved) communication between firms lies behind this result. Several non-APT focused studies of experimental oligopoly markets find that there is an inverse relation between the number of sellers in a market and the size of deviations from the Nash equilibrium (NE) outcomes (for example, see Huck et al., 2004; Dufwenberg and Gneezy, 2000; and Fonseca and Normann, 2012). However, we are unaware of any experimental study that specifically studies the role of the number of sellers in driving the APT phenomenon. We therefore incorporate the number of sellers in our markets as a treatment variable in our experimental design.

To our knowledge, Bayer and Ke (2018) is the only experimental study that directly targets the topic of APT. The authors' study employs a Bertrand duopoly setting in which sellers' costs either increase, decrease or stay constant at the halfway point of the experiment. With two extensions of this baseline condition, they further test the impacts of search costs and asymmetric information on APT. They find APT across all treatments, even in the absence of search and information frictions. They argue that the asymmetry can be explained with a backward-looking learning model: If a seller fails (manages) to sell the good in the period prior to the shock, it is more (less) likely that she will adjust her price downwards (upwards) in the following period. The authors' results support this regularity when the shock is negative, but not when it is positive. Hence, although this learning model may account for the downward rigidity, it falls short of explaining the asymmetry.[6]

While (Bayer and Ke, 2018)'s study provides a useful benchmark to our own, our design choices differ substantially from theirs, as we pursue different research questions. Whereas we aim to assess the roles of cooperative behavior and

[2] Meyer and von Cramon-Taubadel (2004) and Frey and Manera (2007) provide extensive discussions of methodological issues in econometric tests of APT.

[3] In this paper we will use the term "tacit collusion" to mean the phenomenon in which suppliers coordinate on prices above the competitive equilibrium level, through the channel of publicly visible pricing alone. Tacit collusion can also take the form of coordination on quantities below competitive equilibrium levels, but in this paper we will focus strictly on the role of coordination on prices.

[4] Furthermore, the laboratory may be the only environment in which we can reliably detect collusion, since the non-collusive prices or profits are unavailable without imposition of strong structural assumptions.

[5] There are some studies that regress the estimated asymmetry with measures of market concentration, e.g. Loy et al. (2016). Counter-intuitively, the authors find that asymmetry decreases with higher concentration in German milk markets. However, it is difficult to associate this estimate with the causal impact of collusion on APT as their observed higher concentration index may stem from higher efficiency or product differentiation rather than from competitor conduct.

[6] Bayer and Ke (2018) also reason that following positive cost shocks sellers will reason that other sellers will all immediately raise their prices, and so they do the same, while following a negative cost shock sellers do not see any reason to cut their prices unless and until they subsequently lose sales. They cite factors such as bounded rationality as explanations for this behavior, but do not offer a more precise explanation of the channels through which the observed behavior emerges.

tacit collusion on pricing asymmetries, they deliberately try to attenuate their impacts to isolate the role of learning.[7] In particular, in their experiment sellers whose stores are not visited by a buyer receive only limited information on the market price, due to the feedback structure. In our experiment, we inform sellers of the average market price of the other sellers, as we want to create the conditions in which price signalling can be studied more explicitly.

In our experimental setting, subjects play the role of sellers and a computer plays the role of buyers. Each seller faces demand that linearly decreases with one's own price and linearly increases with the average price of others. We vary the size of groups across sessions as 2, 3, 4, 6, and 10, while calibrating the demand function to hold the best-response functions of each seller identical, across all group sizes. Thus, we isolate and study the impact of group size on behavior, while holding the price-based incentives facing individual sellers constant across markets.[8] Throughout our experiment, sellers experience a series of input price shocks – either large or small – that shift the NE price either up or down. Through this design, we are able (i) to test whether APT emerges despite the absence of market frictions and information asymmetries that are often theorized to be the causal forces behind pricing asymmetries; and, (ii) if APT does occur, to assess the impact of number of sellers on the magnitude of the resulting asymmetries. To our knowledge, ours is the first experiment that study the role of number of sellers in shaping APT, but also the first to study the pure number effect in a price competition setting.

Our contributions to the literature are two-fold: First, we document the prevalence of the APT phenomenon through experiments in which we possess strict control over the environment. In particular, our results indicate that the APT may emerge even in the absence of market frictions and information asymmetries that are often theorized to be the causes of pricing asymmetries. This suggests that in markets with three or more sellers, the presence of agents who attempt to coordinate on prices via price signaling may suffice for APT pricing dynamics to emerge. In our duopoly markets, however, our results suggest that coordination can be so successful that rather than the (positive) APT phenomenon, persistent pricing at collusive price levels, or even negative APT, may instead result.

Second, keeping incentives the same across differing group size, we are able to isolate and perform hypothesis tests on the pure effect of increased group size on APT. For markets with three or more sellers, we do not find significant differences in the magnitude of observed APT. Together, the results of our study support theories that highlight the role of tacit collusion on APT. We conclude that APT may be the product environments in which collusion is significant, but imperfect (i.e., unstable).

2. Related literature

2.1. Field evidence

Bacon (1991) provides an early empirical study suggesting that retail gasoline prices in the United Kingdom experience faster and more concentrated responses to crude oil price increases than they do to similar crude oil price decreases. Bacon termed this phenomenon "Rockets and Feathers", and since this paper was published dozens of other researchers have detected the presence of this sort of asymmetry in a variety of consumer and intermediate goods markets.

Peltzman (2000) provides one of the most comprehensive empirical examinations of APT. He conducts a broad study of pricing behavior of 77 consumer and 165 producer goods markets in the U.S., and he concludes that in more than two-thirds of these markets prices rise faster than they fall, in response to input cost changes. Peltzman also seeks correlations between various features of markets and industries, and the degree to which evidence of APT is present. Most notably, he finds that markets with fewer competitors tend to exhibit more pricing asymmetry, while on the other hand markets with higher levels of concentration tend to be less likely to exhibit pricing asymmetry, as in Loy et al. (2016). Peltzman's study, however, does not provide an explanation for these correlations.

In an early survey of field evidence, Meyer and von Cramon-Taubadel (2004) find that (excluding Peltzman, 2000's study), symmetry in price response is rejected in almost one-half of all cases in the literature. Their survey also shows that different test methods yield highly varying rejection rates (between 6% and 80%). Frey and Manera (2007) and Perdiguero-García (2013) provide meta-regression analyses with more comprehensive and recent data sets. Both studies confirm that APT is very likely to occur but also emphasize the variation of reported outcomes. Their results show that this heterogeneity can be explained by certain characteristics of the data (e.g., data frequency) and of the employed econometric model. Most notably, Perdiguero-García (2013) reports that the asymmetry tends to decrease in more competitive segments of the industry.

[7] Although Bayer and Ke (2018) exert effort to minimize the role played by tacit collusion with their study, their typed-stranger matching protocol significantly reduces but does not completely eliminate the possibility that subjects might repeatedly interact, and thus have the opportunity to establish reputation over time. By contrast, the perfect-stranger matching protocol, in which a subject is assured they will be matched with another only once in a session, does more credibly eliminate this possibility. Moreover, the duopoly setting of their study makes collusion presumably more achievable, since coordination is easier when there is only one other market participant. As a result, it is hard to assess the extent of the role to which cooperative behavior played in their study.

[8] This is akin to the "pure number effect" studied by Isaac and Walker (1988) in the context of public goods game. See Hanaki and Masiliūnas (2021) for a similar approach in Cournot oligopolies.

2.2. Theoretical explanations

There is a growing body of literature on the theoretical accounts of APT, an unsurprising fact given that pricing asymmetries are not predicted by standard price competition models.[9] These studies propose explanations of APT mainly by introducing market frictions, information asymmetries or boundedly rational agents into the underlying models. One reason there is such a variety in the way different studies explain the APT is because these studies typically focus on specific market structures (e.g., wholesale petroleum markets) and their idiosyncrasies. In this subsection, we review some of these studies in an attempt to categorize them as well as to highlight discrepancies.[10]

Borenstein et al. (1997) consider the role of search costs in facilitating APT. They hypothesize that negative cost shocks in the presence of costly search provide firms temporary pricing power, which they then use to delay reductions in prices, yielding temporarily superior profits. Benabou and Gertner (1993) and Yang and Ye (2008) also develop explanations based on consumer search costs, but also on the volatility of input costs. They reason that volatility should reduce search incentives for consumers; producers, realizing that they face demand that is temporarily more inelastic, yielding them increased pricing power over the short term, respond by reducing prices more slowly. Reagan and Weitzman (1982) and Borenstein and Shepard (1996) propose explanations based on inventory costs, reasoning that it is relatively more costly for manufacturers and suppliers facing capacity constraints or sharply rising short-term production costs to deal with unanticipated increases in demand resulting from price drops, than it is to respond to corresponding drops in demand due to price increases. Ball and Mankiw (1994) consider a menu-cost model in conjunction with positive trend inflation as an explanation of APT. In another study, Ahrens et al. (2017) show that the presence of consumers with loss aversion may explain why prices are more sluggish to adjust downwards than upwards in response to permanent demand shocks.

The various explanations and models described above provide differing implications for government policy: if APT occurs due to collusion, there may be room for regulation to improve economic efficiency; if however APT is primarily caused by the presence of inventory costs, asymmetric menu costs, or search costs, then regulation that controls pricing behavior may actually induce inefficiency rather than attenuate it. Given the robust evidence of the widespread existence of APT and its non-trivial magnitude and impact on consumer outcomes, identifying which theories best describe the asymmetric pricing behavior is key to determining effective public policy.

2.3. APT and experiments

Despite the many possible explanations that have been proposed, the empirical literature yields only mixed evidence that is often inconclusive due to identification issues. This suggests there is room for further research to shed light on the phenomenon. We consider the advantages of experimental methods in isolating and studying causal determinants of APT.[11] In this subsection, we summarize the most relevant literature to our study.

There are two studies of which we are aware – in addition to Bayer and Ke (2018) – that conduct market experiments with APT-related results. Deck and Wilson (2008) investigate gasoline markets and find that retail prices adjust asymmetrically to changes in station costs in zones with clustered stations, but not in zones with stations that are relatively isolated from competitors. Cason and Friedman (2002) find weak evidence of APT in posted offer markets where customers incur switching costs. While these studies examine their findings on APT, their experimental designs are optimized to investigate questions regarding the structure of gasoline markets (e.g., zone pricing, divorcement) and of consumer markets (e.g., switching costs), not to identify causes of APT. In particular, sellers' costs in both experiments follow random-walk shocks, which may not be salient enough to detect APT. Our study distinguishes itself from this string of literature by examining APT with larger, persistent shocks.[12]

Apart from studies that directly target APT, price competition experiments that study the impact of group size on tacit collusion are also relevant to the current paper. Dufwenberg and Gneezy (2000) provide an early evidence for such a relation through an oligopoly game that corresponds to a discrete version of the Bertrand model. They find that winning prices tend to converge to NE levels in groups of three or four competitors, but stay consistently high in duopolies.

[9] A notable exception is the case of Markov-perfect equilibria, and in particular the case of the Edgeworth cycle. In this phenomenon, firms undercut each others' prices successively until prices approach marginal cost; at this point, one of the firms decides with some positive probability to spike its price, and once this occurs the cycle is repeated, yielding each firm positive economic profits. Maskin and Tirole (1988) further show that these cycles provide a case where asymmetric pricing can be sustained in equilibrium. However, the Edgeworth cycle model requires that firms make price decisions alternately; the model does not support an equilibrium when price decisions are made simultaneously or continuously. Moreover, the emergence of the phenomenon seems in practice to be limited to environments in which competitors rapidly and publicly change prices (see for example Byrne and De Roos, 2019 for an interesting case in Perth, Australia petrol markets, in which a government mandate for retail suppliers to publish their prices daily seems to have facilitated the emergence of a weekly cycle of Edgeworth-like pricing dynamics that persisted for many years.). Thus the Edgeworth cycle model arguably applies only to a relatively narrow range of market contexts.

[10] For more exhaustive surveys of theoretical explanations, see Meyer and von Cramon-Taubadel (2004) and Brown and Yucel (2000).

[11] The usage of experimental methods in macroeconomic research is becoming more and more prevalent. See Duffy (2016) and Cornand and Heinemann (2019) for recent surveys.

[12] Fehr and Tyran (2001) also employ large positive and negative shocks and report APT-like behavior in a price-setting game. However, the authors do not analyze the phenomenon, nor do they probe its implications. In another related experimental study, Duersch and Eife (2019) consider Bertrand duopolies with zero marginal cost in either inflationary, deflationary or constant price environments. They find that real prices are significantly lower in the inflationary environment compared to non-inflationary environments.

Morgan et al. (2006) find that increasing the number of sellers from 2 to 4 decreases the prices paid by some consumers (the ones informed about the entire distribution of prices) but not for others (the ones who buy with motives other than prices). Abbink and Brandts (2008) also find that there is a negative relationship between the number of competing firms and price levels.[13] Nevertheless, as in Dufwenberg and Gneezy (2000), they find that collusive pricing is the modal outcome in duopolies. Fonseca and Normann (2012), Orzen (2008), Davis (2009) and Horstmann et al. (2018) provide further evidence that collusive prices are very likely to be observed in duopolies. Average prices approach considerably close to the NE in the baseline condition of these studies (fixed matching, no communication, symmetric sellers etc.) when the number of sellers is 3 or greater.

A finding common to each of these studies is that persistent coordination over collusive prices is unlikely in markets other than duopolies. This, however, does not preclude the possibility that players might manage to coordinate temporarily on high prices following negative shocks. Experiments also indicate that increasing the number of sellers often leads to more competitive outcomes (in terms of price and output), which in turn should make APT less likely. Although, the meta-analyses of Fiala and Suetens (2017) and Horstmann et al. (2018) on oligopoly experiments indicate that there may not be a linear relationship between the number of competing firms and the degree of tacit collusion. On the one hand, Horstmann et al. (2018) argue that this result stems from the relatively small number of studies that provide pairwise comparisons and the lack of statistical power in these studies. On the other hand, Hanaki and Masiliūnas (2021) consider, as we do, the fact that a change in group size simultaneously influences the difficulty of coordination (aka the "pure number effect") and the incentives provided to collude. They find that the pure effect of group size is small if exists at all. Our study contributes to the literature through improvements of these axis, in particular, by changing the group size without changing the incentives provided to individual sellers in different markets in a price competition game.

3. Method

3.1. Pricing game

We develop a price competition game with a linear demand model and employ this in our experimental markets.[14] In this setting, the demand facing seller $i \in N$ in period $t \in T$ is equal to

$$q_{i,t}(p_{i,t}, p_{-i,t}; \delta, \gamma) = \begin{cases} \delta - \gamma \cdot (p_{i,t} - p_{-i,t}), & p_{i,t} \in \left[p^{min}, \overline{p} \right] \\ 0, & \text{otherwise} \end{cases} \tag{1}$$

where δ and γ are parameters of demand, $p_{i,t}$ is the price set by seller i and $p_{-i,t}$ is the average price chosen by the rival sellers in the same market (i.e. $p_{-i,t} \equiv \frac{1}{N-1} \sum_{j \neq i}^{N-1} p_{j,t}$) at period t and $\overline{p} = \min\left(p^{\max}, p_{-i,t} + \frac{\delta}{\gamma}\right)$. Parameters $p^{\min}$ and $p^{\max}$ refer to the price floor and the reservation price of the representative consumer, respectively. Variable $\overline{p}$ regulates the maximum price level below which (conditional on the average price of other sellers) $q_{i,t}$ takes on strictly positive values.[15]

The model of linear demand that we use, with own- and cross-price parameters equal in magnitude, is micro-founded by the Spokes models of Chen and Riordan (2007) and Bos and Vermeulen (2021), which are themselves adaptations of Salop's canonical Circular City address model (Salop, 1979). In addition, this specification of demand can also be motivated by a limiting case of quasi-linear quadratic utility (see Online Appendix A for an exposition). Note that any demand specification with linear parameters on the price of each competing good can be equivalently represented in terms of own-price and the (weighted) average price of other sellers, as in (1). Thus the presence of average price figure $p_{-i,t}$ in this equation does not imply that when contemplating their appetite for good i that consumers consider the *average* price set by firms competing with firm i; rather, it implies that the aggregate effect of all consumers' individual demand for good i is based on the entire vector of prices, but can nevertheless be mathematically represented concisely in terms of own price p_i and and the average price of all other sellers, p_{-i}. This fact will be helpful in keeping both our analysis and our experimental design tractable, as we will shortly see.

Given the own-demand specification in (1), seller profits are calculated as

$$\pi_{i,t} = (p_{i,t} - mc_t) \cdot q_{i,t} - f, \tag{2}$$

where $q_{i,t}$ is quantity demanded from seller i as defined in (1), mc_t is marginal cost that shifts every $\overline{T}$ periods that comprise a round (denoted $r \in R$) and f is fixed cost. Sellers set their prices in each period simultaneously from a discrete set that is bounded as $p_{i,t} \in [mc_t, p^{\max}]$, such that the price floor is equal to the marginal cost of that round.[16]

[13] Their results are particularly interesting since in their price competition setting, there exist multiple equilibria.

[14] Linear demand systems have a number of advantages over alternatives and long been applied in the industrial organization literature. Such systems lead to closed-form best-response functions and Nash Equilibrium specifications, greatly enabling interpretation of empirical or experimental results. To the extent that non-linear systems of demand (e.g. the Almost Ideal Demand System of Deaton and Muellbauer, 1980 or the Relative Love of Variation model of Zhelobodko et al., 2012) may be preferred on theoretic or other grounds, linear systems provide approximations with first-order accuracy.

[15] In practice subjects chose prices that were revealed to be above this value in fewer than 0.2% of all cases.

[16] As a practical matter we needed to set a price floor of $p_{i,t} \geq mc_t$ to ensure subjects would not complete the experiment with a negative payoff.

Table 1
Experimental design parameters.

General parameters	
Number of periods per round	$\overline{T} = 15$
Number of rounds per session	$R = 5$
Demand parameters	$\delta = 8.50$, $\gamma = 7.275$
Fixed cost	$f = 1$
Maximal/reservation price	$p^{\max} = 3$
Varying parameters	
Group size across treatments	$N \in \{2, 3, 4, 6, 10\}$
Marginal cost across rounds	$mc : (0.90, 0.50, 1.30, 0.50, 0.90)$
Cost shock sequence	$\Delta mc \equiv \eta : (-0.40, +0.80, -0.80, +0.40)$
NE price across rounds	$p^{NE} : (2.07, 1.67, 2.47, 1.67, 2.07)$

In the described finitely repeated game, we can express the maximization problem as

$$\max_{\mathbf{p_i}} \sum_{h=0}^{T} \beta^h \mathbb{E}_{i,t-1} \pi_{i.t+h} \quad \text{subject to} \quad p_{i,t+h} \in [mc_{t+h}, p^{\max}]. \tag{3}$$

Sellers thus maximize the expected discounted sum of profits over T periods by choosing a vector of prices $\mathbf{p_i}$.[17] This in turn leads to best-response function

$$p_{i,t}^{BR} = \frac{1}{2}\left(mc_t + \frac{\delta}{\gamma} + \mathbb{E}_{i,t-1}[p_{-i,t}]\right). \tag{4}$$

The current period's marginal cost mc is revealed to the sellers prior to their pricing decisions thus is outside the expectation operator. Sellers form expectations of the prices their competitors will set during the current period $\left(\mathbb{E}_{i,t-1}[p_{-i,t}]\right)$ by conditioning on all available information. The system of best responses for all sellers implied by (4) solves for steady-state prices as:

$$p_{i,t}^{*} = p_{i,t}^{BR}\left(\mathbb{E}_{i,t-1}[p_{-i,t}^{*}]\right) = mc_t + \frac{\delta}{\gamma} \equiv p_t^{NE}. \tag{5}$$

This price level also corresponds to the unique stage-game NE (p_t^{NE}) and the unique subgame perfect Nash equilibrium (SPNE). Sellers may achieve the joint profit maximum (JPM) if they each set their prices to the maximum price $p^{\max}$.

In this pricing game, neither own-demand nor own-profit depend on the number of sellers. These only depend on own-price and the average price of rival sellers. The best-response action is also independent of N for a wide range of expectation models, including rational expectations and adaptive expectations (Evans and Honkapohja, 2001). This feature assures that the incentives given to the sellers of different group sizes are matched and the market power of each seller, measured by the size of markup over marginal cost, is *ex-ante* equal. We consider this as necessary for ensuring a *ceteris paribus* comparison between the treatment conditions and to capture the pure number effect.

3.2. Experimental design

Sellers interact repeatedly in the described pricing game for R rounds, which are each composed of $\overline{T}$ periods. Marginal cost mc_t fluctuates across rounds, modeling large exogenous cost shocks, but remains invariant during a round. Our experimental manipulations consist of varying the size of markets across sessions in a between-subjects design, and of varying the size and direction of shocks across rounds in a within-subjects design. We implement a fixed-matching protocol during a session.

The calibration of the experimental game is summarized in Table 1. The experiment consists of 5 rounds of 15 periods each, with a new marginal cost announced at the beginning of each round. The sequence of shocks is identical across all treatments: marginal cost starts at \$0.90 in Round 1, drops to \$0.50 in Round 2, rises to \$1.30 in Round 3, falls again to \$0.50 in Round 4, then rises to \$0.90 for Round 5.

3.3. Procedures

Experimental sessions were conducted at the University of California, Santa Barbara's Experimental and Behavioral Economics Laboratory (EBEL) using the z-Tree platform (Fischbacher, 2007), between September and December of 2018. A total of 245 subjects were recruited from the experimental economics subject pool of the same university, using the ORSEE tool

[17] Although the model assumes that sellers choose a vector of prices for all periods, subjects only submit a price decision for the current period in the experiment.

(Greiner, 2015). Subjects were allocated to markets of size 2, 3, 4, 6 and 10, with a total of 36, 39, 52, 48 and 70 subjects assigned to each group size condition, respectively. This setup yields 59 independent markets for the analysis.[18]

At the beginning of each experiment, subjects are provided written instructions which are also read to them aloud by an experimenter. Subjects then proceed to take a short comprehension quiz.[19] In the main part of the experiment, each subject plays the role of sellers and makes a series of 75 pricing decisions, whereas consumer behavior is simulated by computer. We also elicit subjects' one-period-ahead expectations about the average price chosen by rival sellers (i.e., $\mathbb{E}_{i,t-1}[p_{-i,t}]$). These expectations are not rewarded separately, to avoid creating hedging issues. Subjects are able to set a price between the marginal cost and the maximum price (of \$3.00), in increments of \$0.01. Once all subjects set their prices and expectations, they are individually notified by the computer of the average price established by the others in their market, reminded of their own price, and shown their own resulting payoff for that period. Subjects are able to track the previous values of these outcomes through a history box that is available in their screen (see Online Appendix B.3).

We notify subjects that a new cost shock will occur at the beginning of each new round, either an increase or decrease, of either \$0.40 or \$0.80. We reveal the magnitude and direction of each shock immediately prior to the first period of each respective round. At that time, we also hand out copies of a printed payoff table corresponding to the new marginal cost. These tables assist subjects in estimating the profits they will receive, conditional on the hypothetical prices they and others may set in each period of that round (see Online Appendix B.3).

Sessions lasted a total of 90 to 125 min. Subjects were paid \$18.66 on average (a minimum of \$10.89 and a maximum of \$28.50), which includes the \$5.00 show-up fee and \$3.00 for the completion of the optional survey (no subject declined this offer). The remaining payoff is determined as the average payoff of a randomly chosen round of the game.

4. Hypotheses

As the experimental design specifically avoids any of the features outlined in Section 2.2 (e.g., frictions, information asymmetries), standard theory suggests that prices react symmetrically to shocks. However, we may observe asymmetry if certain conditions are satisfied even in the absence of these frictions. To unravel these conditions, we need to investigate the strategic tension underlying our pricing game. On the one hand, individual incentives promote undercutting opponents' prices until the NE price is reached. On the other hand, sellers may generate higher profits if they successfully sustain coordination at a price level above NE.

We first focus on the incentive each individual seller has to undercut other sellers. Consider the following variable:

$$Inc2Dev(p_{-i,t}) \equiv \pi_{i,t}(p_{i,t}^{BR}(p_{-i,t-1}), p_{-i,t-1}) - \pi_{i,t}(p_{-i,t-1}, p_{-i,t-1}), \tag{6}$$

which expresses the incentive of a myopic seller to deviate from the coordinated price level from the previous period, $p_{-i,t-1}$.[20] This seller is myopic as s/he ignores the possibility that the competitors might also engage in similar reasoning. By substitution of (1), (2), and (4) we can rewrite (6) as

$$Inc2Dev(p_{-i,t}) = \frac{\gamma}{4}(p_{-i,t-1} - \delta/\gamma - mc_t)^2 = \frac{\gamma}{4}\left(p_{-i,t-1} - p_t^{NE}\right)^2 \tag{7}$$

The incentive to deviate from price $p_{-i,t-1}$ during period t thus increases quadratically with the difference between the previous period market price and the current period NE price. Thus, individual incentives promote convergence to the NE.

We now separately analyze *Inc2Dev* in reaction to positive and negative shocks. Consider the case where a shock of magnitude $\eta \equiv |mc_t - mc_{t-1}|$ occurs in period t. Then, application of (5) allows us to express $p_t^{NE} = p_{t-1}^{NE} + \eta$ and $p_t^{NE} = p_{t-1}^{NE} - \eta$ for positive and negative shocks, respectively, and we thus define:

$$\begin{aligned} Inc2Dev_{i,t}^{+} &\equiv \tfrac{\gamma}{4}(p_{-i,t-1} - (p_{t-1}^{NE} + \eta))^2 \\ Inc2Dev_{i,t}^{-} &\equiv \tfrac{\gamma}{4}(p_{-i,t-1} - (p_{t-1}^{NE} - \eta))^2. \end{aligned} \tag{8}$$

Straightforward manipulation of (8) allows us to express the difference in the incentives to deviate as:

$$\Delta Inc2Dev_{i,t} \equiv Inc2Dev_{i,t}^{+} - Inc2Dev_{i,t}^{-} = -\gamma\eta \cdot (p_{-i,t-1} - p_{t-1}^{NE}). \tag{9}$$

This equation implies that when market prices are above (below) NE immediately prior to a cost shock, the incentive to deviate *following the shock* will be greater (lesser) if the shock is negative than if it is positive. When however $p_{-i,t-1} \to p_{t-1}^{NE}$, or when sellers are not myopic, the incentive to deviate converges to zero. We consequently consider the absence of APT as our first null hypothesis:

Hypothesis 1. *Prices respond symmetrically to (equally sized) positive and negative shocks.*

We can test this hypothesis by exploiting the exogenous within-subjects treatment variations in marginal cost.

[18] In one session (20 subjects), the data from the final period is lost due to technical reasons. All the analysis in the results section is performed based on all the available data.

[19] We reviewed answers for each subject and provided explanations where needed. See Online Appendix B for all experimental material.

[20] We are indebted to an anonymous reviewer for proposing analysis using this approach.

An interesting implication of Eq. (9) is that if prices are already collusive prior to the shock, the relatively greater temptation to deviate from the coordinated price following negative shocks suggests we should observe negative APT as opposed to the positive APT that (Borenstein et al., 1997) and many others have observed. The arguments of myopic best-response and of tacit collusion – explained in the next hypothesis – therefore go in the opposite directions.

Our second hypothesis concerns the second force underlying this strategic tension: the prospect of sustained higher profits through tacit collusion. We use market power, measured through the size of markups over marginal cost, to study collusion (see Section 5.2. for the description of the exact measure). Markups should be invariant to the number of rival sellers, given that both the profit and best-response functions are independent of group size. Moreover, in the absence of frictions and the ability of competitors to communicate, the theory predicts a constant markup for all levels of marginal cost. However, if tacit collusion occurs, we expect to observe higher market power (i) in markets with fewer sellers, and (ii) in the periods occurring soon after negative shocks. For (i), we expect to observe persistent coordination more often where there are fewer sellers to dampen the strength of price signals. Moreover, it is arguably easier to sustain coordination above NE pricing when one has fewer competitors, as any one of them can undermine joint coordination if they individually fail to cooperate. For (ii), we conjecture that negative shocks boost the market power of sellers (at least temporarily), as such shocks may play the role of a coordination device for attempts of collusion. We can test this hypothesis by using the between-subjects treatment variations in group size, and within-subject treatment variations in marginal cost.

Hypothesis 2. *Sellers' market power is invariant to the number of sellers in the market, and is unaffected by the existence of periodic shocks.*

Finally, our third hypothesis concerns individual pricing strategies. The Rational Expectations Hypothesis (REH) of (Muth, 1961) admits the possibility of expectation errors at the individual level, but which should tend to cancel out in aggregate. Also, after observing $t-1$ periods of price history, a seller may learn that the others do not behave consistently

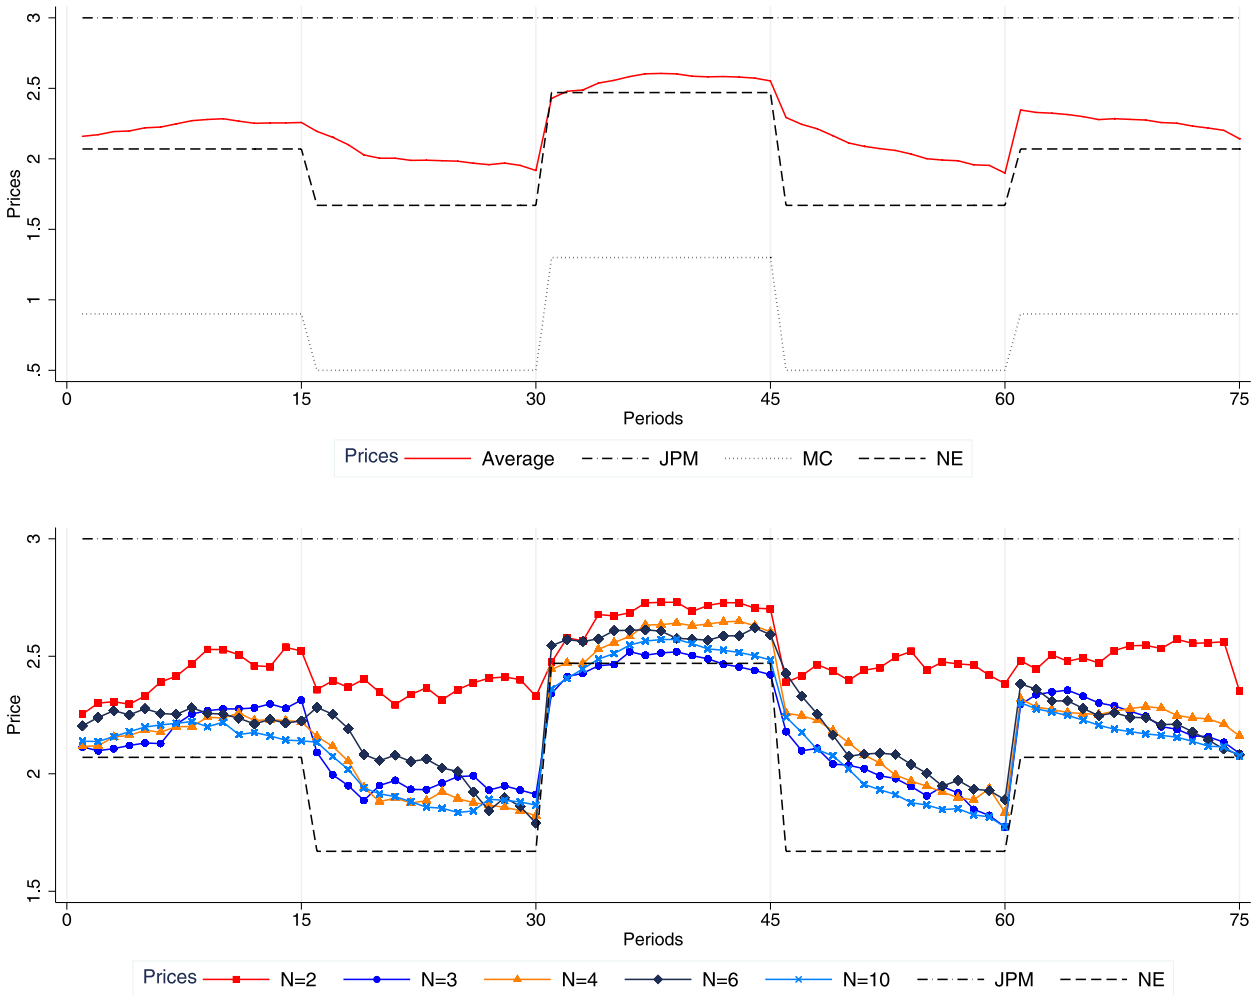

Fig. 1. Average pricing behavior across periods and group sizes.

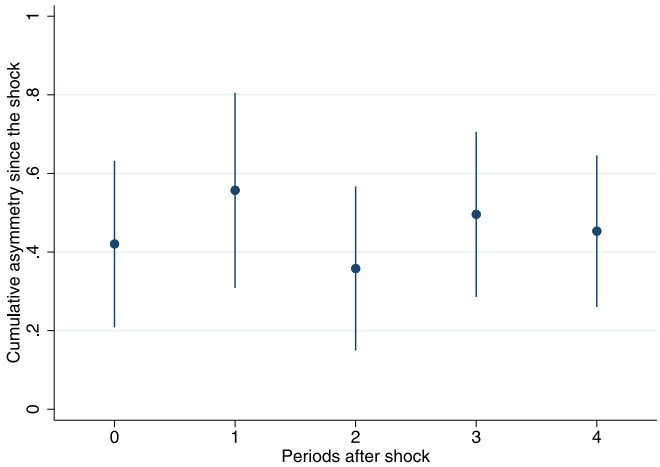

Fig. 2. Cumulative response after K periods. Dots refer to $\sum_{k=0}^{K} c_{t-k}$. Lines represent 95% confidence intervals.

with the predictions of REH. Nevertheless, conditional on expectations, sellers should select the best-response action as this maximizes their profit. As we elicit subjects' guesses on the average price set by others, we can test this hypothesis without assuming a specific expectation model. We can interpret any intentional deviation from the profit-maximazing action as an attempt to reach collusive prices.

Hypothesis 3. *Conditional on expectations, pricing behavior follows the best-response function.*

5. Results

Figure 1 provides a depiction of the average price per period, as the average of all market prices and as broken out by group size. Here, market price refers to the average of all prices in market m (i.e., $p_{m,t} = \frac{1}{N}\sum_{i=1}^{N} p_{i,t}$). The reader can readily discern that for groups of size 3 and greater, average prices rise rapidly after positive cost shocks, while they fall more slowly after negative cost shocks. By contrast, for groups of size 2, it is not immediately obvious whether average pricing behavior is affected by cost shocks. A second observation that is immediately clear is that average prices are generally above the NE price, with deviations being higher following negative shocks than when following positive shocks. Overall, the visual inspection of the data suggests the presence of market behavior consistent with APT.

5.1. Estimation of the asymmetry

We follow Peltzman (2000) and estimate the coefficients of the distributed lag model (DLM) to measure the magnitude of APT. This model can be expressed as:

$$\Delta p_{i,t} = \sum_{k=0}^{K} b_{t-k} \cdot \Delta mc_{t-k} + \sum_{k=0}^{K} c_{t-k} \cdot (\mathbb{1}[\Delta mc_{t-k} > 0] \cdot \Delta mc_{t-k}) + \epsilon_{i,t} \tag{10}$$

where the change in output price (i.e., $\Delta p_{i,t} = p_{i,t} - p_{i,t-1}$) is modelled as a function of the lagged changes in marginal cost (i.e., Δmc_{t-k}). The indicator variable $\mathbb{1}[\Delta mc_{t-k} > 0]$ takes the value 1 if the change in marginal cost in period $t-k$ is positive and equal to 0 otherwise. The sum of interaction coefficients $\sum_{k=0}^{K} c_{t-k}$ reflects the magnitude of asymmetry and its persistence over K periods.

We estimate model (10) with Ordinary Least Squares (OLS) regressions in a step-wise manner. Figure 2 reports the estimated asymmetry for $K = 4$.[21] Estimates indicate that the APT is both strong and persistent. Immediate price reactions are 32.9 cents greater in magnitude for positive than for negative 80-cent shocks.

We now assess the reaction of prices to equally sized shocks between our treatment groups with non-parametric tests. We compare immediate pass-through rates of shocks that are β_0^+ and β_0^- calculated as:

$$\begin{aligned} p_{i,t+\tau}^{+} &= p_{i,t-1}^{+} + \beta_\tau^{+}\eta^{+} \\ p_{i,t+\tau}^{-} &= p_{i,t-1}^{-} + \beta_\tau^{-}\eta^{-} \end{aligned} \tag{11}$$

[21] We report the full set of results in Online Appendix C.1. All estimations employ robust standard errors that are clustered at market level. We also include a set of indicator variables that are specific to each group size (i.e., $\mathbb{1}[N = s]$), the lagged change in the average price of rival sellers (i.e., $\Delta p_{-i,t-1}$), a three-way interaction term (between $\mathbb{1}[\Delta mc_{t-k} > 0]$, Δmc_{t-k} and group size specific indicator variables) and autoregressive terms amongst the set of regressors to check the robustness of estimates. The significance of asymmetry coefficients as well as their magnitude are robust to the inclusion of these variables.

Table 2
Asymmetry in the immediate pass-through rates.

	All	$N > 2$	$N = 2$	$N = 3$	$N = 4$	$N = 6$	$N = 10$
Small shocks							
β_0^-	0.159	0.115	0.411	0.558	0.158	−0.144	0.0154
	(0.0676)	(0.0739)	(0.162)	(0.247)	(0.138)	(0.130)	(0.0985)
β_0^+	1.119	1.270	0.244	1.305	1.209	1.233	1.322
	(0.0672)	(0.0659)	(0.196)	(0.172)	(0.121)	(0.121)	(0.123)
p-value	0.000	0.000	0.411	0.028	0.000	0.000	0.000
Large shocks							
β_0^-	0.324	0.313	0.391	0.303	0.432	0.206	0.304
	(0.0362)	(0.0405)	(0.0734)	(0.107)	(0.0811)	(0.0706)	(0.0712)
β_0^+	0.639	0.718	0.181	0.537	0.779	0.945	0.619
	(0.0396)	(0.0400)	(0.111)	(0.116)	(0.0636)	(0.0796)	(0.0643)
p-value	0.000	0.000	0.035	0.202	0.013	0.000	0.006
Observations	245	209	36	39	52	48	70

The averages of pass-through rates by differing group sizes are reported. Below averages, standard errors are reported in parentheses. p-values correspond to the result of the Wilcoxon signed-rank test on the equality of pass-through rates for small or large shocks (i.e. $H_0 : \beta_0^+ = \beta_0^-$).

where η^+ (η^-) reflects either the large or small positive (negative) shock and $t-1$ corresponds to the period just before the shock. Note that aggregate market demand implied by (1) is perfectly inelastic, so that shifts in the NE price following cost shocks are exactly equal to the magnitude of the cost shock itself. Accordingly, we can denote the cases $\beta_0^+ = 1$ and/or $\beta_0^- = 1$ as incidences of "full pass-through" of cost shocks. The ratio of β_0^+ and β_0^- thus conveys information on the degree of APT in immediate cost-shock responses. A ratio of 1 would indicate the absence of APT.

Table 2 provides average value of pass-through rates for different aggregation levels.[22] First, we note that the hypothesis of full pass-through can generally be rejected. Exceptions consist of the small positive shock (i.e., $\eta^+ = 0.40$) and $N = 6$ for the large positive shock. Second, we test APT in the immediate post-shock responses by testing the equality of immediate pass-through rates for equally sized shocks as $H_0 : \beta_0^+ = \beta_0^-$ via Wilcoxon signed-rank tests. The pooled data and the data for groups of size greater than 2 suggest rejecting the null. For groups of size 3, we reject symmetry for the smaller but not for the larger shock.

For duopolies, we see that the asymmetry is reversed; the average price response following the larger cost shock is significantly greater for the negative than for the positive cost shock. Reflecting on the differences in incentives to deviate calculated in (9) from Section 4, and noting from Fig. 1 that pre-shock prices were generally much higher than the NE price immediately prior to shocks in duopoly markets, we see that the relative incentive to deviate after negative shocks was much stronger in duopoly markets than in other markets. We explore the possibility that this relatively higher incentive to deviate in duopoly markets may have led to this divergent outcome.

Taken together with the estimates of the DLM, we reach the first two results of our paper:

Result 1.1: *Prices do not react symmetrically to equally sized positive and negative shocks.*

Result 1.2: *Price reactions in triopoly and larger markets are consistent with positive APT. Price reactions in duopoly markets are consistent with negative APT.*

5.2. Market power

We now turn to our second hypothesis. We follow the literature in applying the Lerner index as the relevant measure of market power: $L_{i,t} = \frac{p_{i,t} - mc_t}{p_{i,t}}$ (Lerner, 1934). We propose that the difference between the observed Lerner index (i.e., $L_{i,t}$) and the"theoretical" Lerner index, that is the index that would be relevant if behavior was consistent with NE predictions (i.e., $L_t^{NE} = \frac{p_t^{NE} - mc_t}{p_t^{NE}}$), provides a measure of "excess" market power due to collusion. We further propose this as an appropriate measure of tacit collusion, as our price competition structure incorporates homogeneous goods and we control marginal costs. Thus, we do not suffer the identification problem of observational studies. Our measure of excess market power can then be expressed as:

$$L_{i,t}^x = L_{i,t} - L_t^{NE} = mc_t \left(\frac{1}{p_t^{NE}} - \frac{1}{p_{i,t}} \right). \tag{12}$$

Figure 3 depicts the average of our measure of excess market power, by period and treatment. Upon visual examination one can immediately see that excess market power generally lies above the theoretical "Nash" level, consistent with an environment in which tacit collusion exists. Also, this average measure reaches its highest levels during the second and fourth rounds, the two rounds that immediately follow negative shocks. Following the large positive shock at the beginning

[22] We report the results of an identical analysis for $\tau = 14$ in the Online Appendix C.2. Consequently, asymmetry in pass-through rates reduces but does not disappear entirely even 14 periods after the shock.

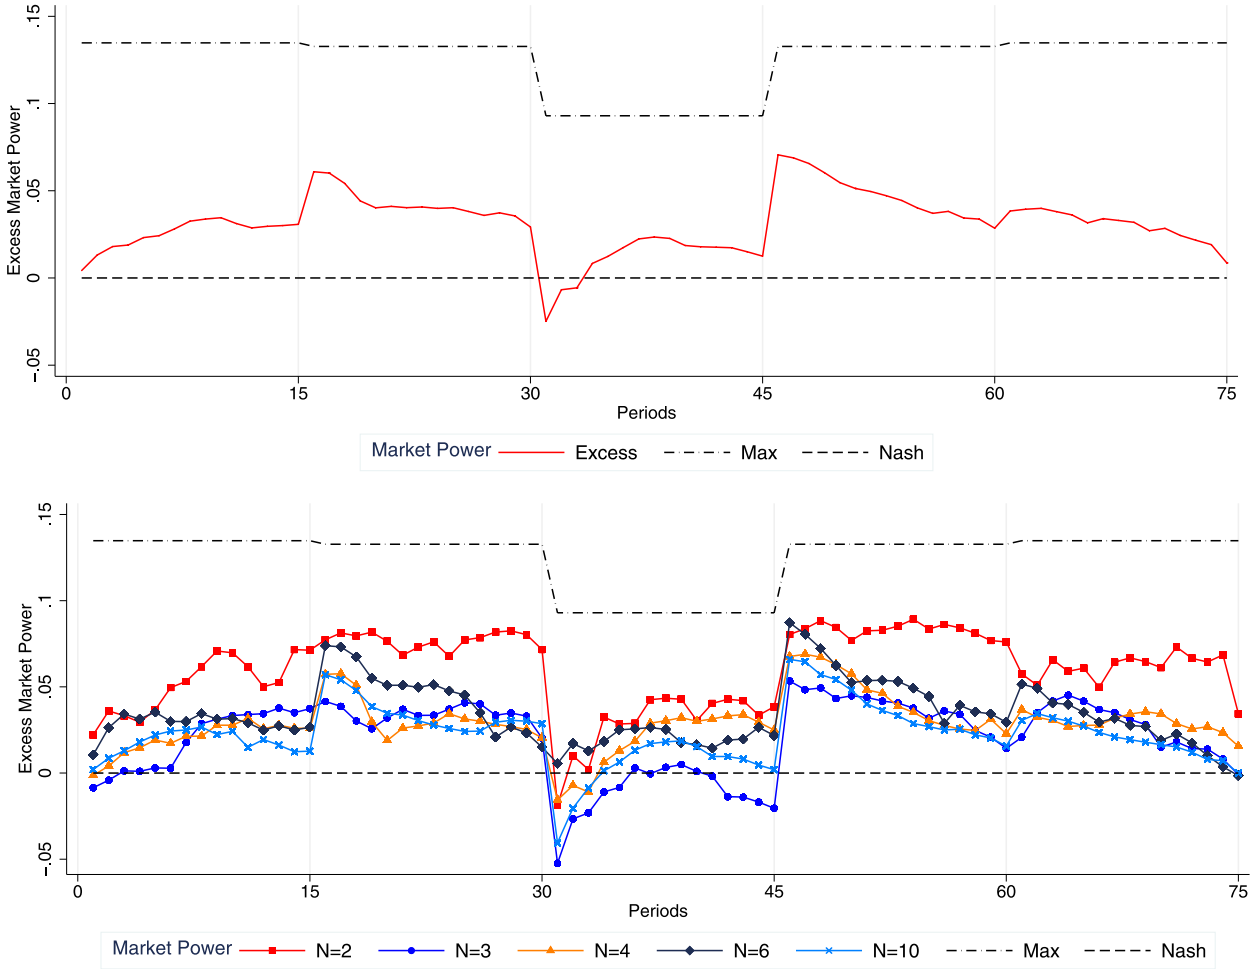

Fig. 3. Excess market power across periods and group sizes. In both subfigures, "Max" refers to the maximum excess market power that can be observed (i.e., when $L_{i,t} = L_t^{\max} = \frac{p^{\max} - MC_r}{p^{\max}}$) and "Nash" refers to the case $L_{i,t} = L_t^{NE}$.

of the third round, excess market power falls so much that it turns negative for several periods. Following the small positive shock at the beginning of the fifth round, excess market power does not notably react.

We test the veracity of these observations by performing OLS regressions.[23] We consider the following specification:

$$L_{i,t}^{x} = \alpha + \sum_{s \neq 2} \delta_s \cdot \mathbb{1}[N = s] + \sum_{e \neq 1} \gamma_e \cdot \mathbb{1}[r = e] + \epsilon_{i,t}, \tag{13}$$

where the excess market power of seller i in period t is modeled as a function of group size- and round-specific indicator variables. According to our null hypothesis, the model with no independent variables should fit the data as well as this model.

Table 3 reports the estimates in a step-wise manner. In model (5), we truncate the data to the periods where shocks shift the marginal cost (i.e., periods 16, 31, 46 and 61) and replace the dependent variable with the change in excess market power as $\Delta L_{i,t}^{x}$. This allows us to interpret the estimates of round specific indicator variables as the immediate effect of cost shocks on the tacit collusion in model (5).

First, we reject the null hypothesis in all specifications except (2) at a confidence level of 0.01 with the F-test. The fact that the constant α is positive and significant in model (1) indicate the overall presence of tacit collusion. Second, the coefficients of round-specific indicator variables in model (3) indicate that tacit collusion is higher during the second and fourth rounds, and lower during the third round relative to the first round. In model (5) where we truncate the data, the coefficient of rounds 3 and 5 are negative and that of round 4 is positive. Furthermore, we reject the hypothesis $H_0 : \alpha + \delta_s + \gamma_e = 0$ at a confidence level of 0.05 (i.e., $\alpha + \delta_{\{N=3,4,10\}} + \gamma_5 = 0$). We can thus say that immediately after a negative

[23] The non-parametric counterpart of this test is reported in Online Appendix C.3

Table 3
Excess market power.

	(1)	(2)	(3)	(4)	(5)
Constant	0.032***	0.060***	0.025***	0.054***	0.004
	(0.004)	(0.011)	(0.004)	(0.011)	(0.007)
$N = 3$		−0.039*		−0.039*	0.018
		(0.015)		(0.015)	(0.020)
$N = 4$		−0.031*		−0.031*	0.028**
		(0.014)		(0.014)	(0.010)
$N = 6$		−0.026		−0.026	0.047***
		(0.014)		(0.014)	(0.010)
$N = 10$		−0.038**		−0.038**	0.029*
		(0.012)		(0.012)	(0.011)
$r = 2$			0.017***	0.017***	
			(0.003)	(0.003)	
$r = 3$			−0.014***	−0.014***	−0.084***
			(0.004)	(0.004)	(0.009)
$r = 4$			0.023***	0.023***	0.028***
			(0.004)	(0.004)	(0.006)
$r = 5$			0.005	0.005	−0.020**
			(0.004)	(0.004)	(0.007)
Observations	18,355	18,355	18,355	18,355	980
Adjusted R^2	-	0.052	0.054	0.106	0.225
F-statistic	-	2.607	31.289	19.288	27.003

Results of OLS regressions on specification (13) are reported. In model (5), the dependent variable is the change in excess market power, $\Delta L^x_{i,t}$. Below estimates, robust standard errors that are clustered at the market level are reported in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

(the large positive) shock, the excess market power increases (decreases). Third, coefficients of group size specific indicator variables are negative and significant, although at marginal level for $N = 6$ (p-value$= 0.071$) in model (2). Here, we also reject the hypothesis $H_0 : \alpha + \delta_s = 0$ for all s (p-value< 0.01). This suggests that tacit collusion is present in all markets but its magnitude is smaller when $N > 2$. The sign of these coefficients in model (5) suggests that markets larger than size 3 increase their market power in response to the first negative shock. Lastly, we generally reject the hypothesis $H_0 : \alpha + \delta_s + \gamma_e = 0$ in the most unrestricted model (4) (11 times out of 15 tests at p-value< 0.05). The overall interpretation of these tests provide the basis of our second result:

Result 2: *Excess market power (i.e., tacit collusion) is not invariant to shock direction and group size. It is persistently higher in duopolies, and in larger-sized markets it rises following negative cost shocks.*

5.3. Deviations from best-response

Finally, we assess the deviation of subjects' prices from the profit-maximizing best-response action that is computed by using the submitted expectations/guesses (i.e., $p_{i,t} - p^{BR|\mathbb{E}}_{i,t}$). Conditional on the subjects accurately reporting their expectations, they risk lower profits if they fail to best-respond to these reported expectations. Consequently, the observed deviations can be attributed either to error or alternatively to strategic motives (i.e. attempts to collude). To argue that the deviations we observe in our experiments are not entirely due to erroneous behavior, we compare the magnitudes and directions of such deviations to the average magnitude of expectation errors (i.e., $\mathbb{E}_{i,t-1}[p_{-i,t}] - p_{-i,t}$) and the average of absolute expectation errors.[24] If subjects deviate from their best-response action in a way that is different from their expectation errors, this suggests that deviations reflect intentional behavior.

Figure 4 (a) depicts the average value of these deviations over time. The average expectation errors are remarkably close to zero, with no obvious trend across periods. Although this suggests that beliefs are on average correct, it does not imply the complete absence of errors: the average measure of absolute expectation errors lies well above zero throughout the experiment. The latter peaks following both positive and negative cost shocks but subsequently trends downward. The patterns of high initial absolute expectation errors and slow convergence are consistent with those of prior experiments in which prices are strategic complements (e.g., Hommes et al. 2005; Cooper et al. 2021). However, deviations from the best-response action reveal a different and rather interesting pattern: they peak sharply following negative shocks and remain high during these rounds, but do not peak similarly following positive shocks. The second graph in Fig. 4 depicts the average of deviations from best-response action by group size. The same pattern can be traced across our treatment groups.

We perform OLS regressions to study deviations from best-response. Consider the following specification:

$$p_{i,t} - p^{BR|\mathbb{E}}_{i,t} = \alpha + \sum_{s \neq 2} \delta_s \cdot \mathbb{1}[N = s] + \sum_{e \neq 1} \gamma_e \cdot \mathbb{1}[r = e] + \epsilon_{i,t} \tag{14}$$

[24] We label these latter two as "errors" rather than as "deviations" as there is no strategic benefit to knowingly submitting inaccurate guesses/expectations in our experiment.

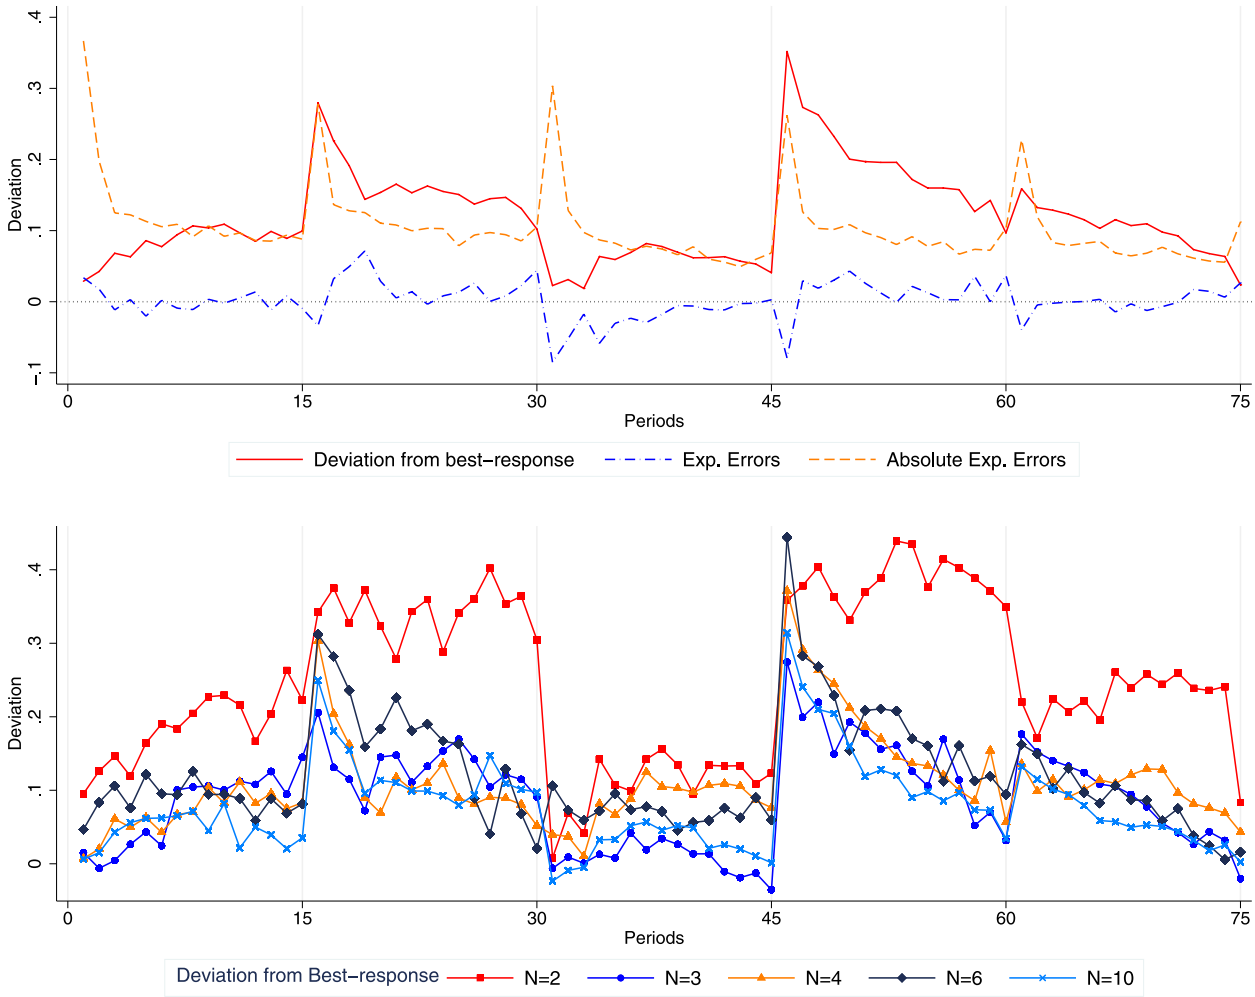

Fig. 4. Deviations from best-response action and errors in expectations.

where the deviation of subject i's price from the best-response action conditional on the submitted guess is modeled as a function of group size- and round-specific indicator variables. Our null hypothesis concerns the overall significance of this model.

Table 4 reports the estimates in a step-wise manner. In model (5), we truncate the data to the periods where shocks shift marginal cost the same way as we did in the previous section, and replace the dependent variable with the change in deviation from best-response action following a cost shock.

We reject the null hypothesis in all specifications. The hypotheses $H_0 : \alpha + \gamma_s = 0$ in model (2) and $H_0 : \alpha + \delta_s = 0$ in model (3) can be rejected at a significance of $p < 0.01$. This points out to the following two results: (i) Sellers deviate more (less) from the associated best-response action following negative (positive) shocks and (ii) deviations are lower when $N > 2$. We see that the sign of group size indicator coefficients in models (4) and (5) are flipped. In model (4), they reflect the fact that groups of size 3 and larger deviate less, on average, relative to duopolies. In model (5), they correspond to the immediate reaction of these groups to the first negative shock. These deviations rise further when a large negative shock shifts the marginal cost down ($\hat{\gamma}_4 = 0.131$) while they drop significantly in response to the large positive shock ($\hat{\gamma}_3 = -0.260$). In consequence, we reach to the following results:

Result 3.1: *Sellers deviate on average above their best-response action.*

Result 3.2: *Deviations from the best-response action grow (shrink) following negative (positive) shocks.*

6. Discussion

Our results point to the co-appearance of asymmetric price transmission and tacit collusion. The latter seems to be the result of strategic behavior, as our analysis of deviations from best-response action reveals. These findings are consistent with theories that cast tacit collusion as having a significant role in the emergence of APT, such as the trigger price model

Table 4
Deviations from best-response.

	(1)	(2)	(3)	(4)	(5)
Constant	0.120***	0.083***	0.248***	0.212***	0.044
	(0.013)	(0.012)	(0.041)	(0.040)	(0.028)
$N = 3$			−0.159**	−0.159**	0.122*
			(0.049)	(0.049)	(0.055)
$N = 4$			−0.138**	−0.138**	0.163**
			(0.050)	(0.050)	(0.051)
$N = 6$			−0.128*	−0.128*	0.210***
			(0.051)	(0.051)	(0.041)
$N = 10$			−0.170***	−0.170***	0.144**
			(0.044)	(0.044)	(0.046)
$r = 2$		0.080***		0.080***	
		(0.010)		(0.010)	
$r = 3$		−0.028**		−0.028**	−0.260***
		(0.010)		(0.010)	(0.031)
$r = 4$		0.112***		0.112***	0.131***
		(0.015)		(0.015)	(0.025)
$r = 5$		0.018		0.018	−0.118***
		(0.013)		(0.013)	(0.024)
Observations	18,355	18,355	18,355	18,355	980
Adjusted R^2	-	0.044	0.051	0.095	0.170
F-statistic	-	37.840	4.013	19.546	40.442

Results of OLS regressions on specification (14) are reported. In model (5), the dependent variable is the change in deviation from the best-response action following a cost shock, $\Delta(p_{i,t} - p_{i,t}^{BR|\mathbb{E}})$. Robust standard errors are clustered at the market level and are reported in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

in Borenstein et al. (1997). Most of the other theoretical explanations of APT in the literature cannot account for the pricing behavior observed in our results. We can reasonably exclude, for example, the influence of explicit collusion (i.e., involving direct communication), capacity constraints, inventory limitations, (a)symmetric menu costs, consumer loss aversion, (a)symmetric search costs, contexts of alternating price moves and price lockup periods, and so forth, as being necessary conditions for APT, since these features are excluded by our design.

We cannot, however, claim a monotonic relation between the magnitude of APT and tacit collusion: the pricing behavior of duopolies in our experiment is revealed to be consistent with 1) a negative APT response in the immediate aftermath of shocks, and 2) elevated and relatively stable pricing behavior between shocks. We propose an explanation for the exceptionality of the duopoly result follows, tackling the latter result first: in duopolies collusion is so strong that sellers are, by and large, able to maintain cooperative (tacitly collusive) pricing over a sustained period of time, with pricing showing no reversion to Nash. We therefore argue that APT requires significant, but imperfect, collusion.[25]

The success of duopoly markets in maintaining average prices well above equilibrium levels between shocks, and well above average price levels of triopoly and larger markets, provides a plausible explanation for the negative APT result we observed with duopolies for large shocks: as is apparent in an examination of (9), the high price deviations above NE prices in the lead-up to shocks in duopoly markets implies a correspondingly greater incentive for firms to deviate downward following negative than positive shocks. Thus we have the interesting result that while the dynamics of our duopoly markets appears to have enabled more success in attempts to tacitly collude between shocks, this success (in the form of elevated prices) also made these markets relatively more prone to respond strongly to downward than upward cost shocks.[26]

If tacit collusion is indeed a significant causal force behind APT, then our work has important implications for antitrust enforcement policy against collusion and price-fixing. In particular, regulators may consider APT in a market as a signal for the presence of collusion between firms in that market. Since many real-world interactions between competitive firms are repeated indefinitely, such collusion may even be sustainable as a NE. Further research is needed to determine whether collusion is an important cause of APT behavior in field settings, and if so, whether suitable forms of regulatory intervention might exist to reduce such collusion without increases in inefficiency.

[25] The fact that our duopolies reached almost stable collusion, while larger markets did not, is consistent with the literature we review in Section 2.3. This can be attributable to a combination of two factors: first, coordination between market participants becomes increasingly difficult with each new seller, and three may well be the number from which the difficulties and costs involved in maintaining coordination start to exceed the marginal benefits; second, our duopolies are unique in that each participant can deduce the choices made by the other participant by observing aggregate market outcomes. In a triopoly or larger market, by contrast, it is not possible for sellers to detect whether an aggregate market outcome is due to the defection of a single competitor, or from a broader but shallower defection by multiple competitors.

[26] As we noted in the Results section, our dupoloy markets exhibit pricing behavior consistent with *negative* APT, while larger markets' behavior was consistent with *positive* APT. We see these apparently contradictory results as indicating a tension between two phenomena operating in opposite directions: on the one hand, collusive dynamics encourage competitors to respond more sharply to upward than to downward cost shocks, as we have argued throughout the paper; however, as prices rise above NE levels the differential incentives to deviate become stronger following negative shocks than positive shocks. At some point the success of tacit collusion may be so great as to reverse the direction of the resulting APT.

We propose that follow-on research may yield further insights into the mechanisms through which tacit collusion leads to APT, as well as potential policy responses that might diminish its frequency and magnitude. In particular, future experiments should address the impact of different levels of information transparency. Most notably, testing the effects of providing feedback on individual prices and/or payoffs of rivals on APT may provide particularly helpful insights. The latter is shown to lead to more rivalistic outcomes in experimental oligopoly studies, as it initiates imitation dynamics (Fiala and Suetens, 2017), while the former can lead to more collusive levels. Nevertheless, both may reduce the degree of asymmetry in price transmission. Another area of needed research is to explore the roles of market power and market concentration in shaping APT pricing behavior. In our experimental design we explicitly kept the incentives provided to sellers the same across markets of varying sizes to study the pure number effect similar to Hanaki and Masiliūnas (2021). Finally, future studies may benefit from testing the robustness of our findings to alternative demand specifications. The parameters of demand in our experimental markets are consistent with goods such as retail gasoline, which are demanded inelastically but for which suppliers face elastic demand. Our results are also more relevant to markets in which there are close substitutes.

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.

Supplementary material

Supplementary material associated with this article can be found, in the online version, at doi:10.1016/j.jebo.2021.10.018.

References

Abbink, K., Brandts, J., 2008. Pricing in Bertrand competition with increasing marginal costs. Games Econ. Behav. 63 (1), 1–31.
Ahrens, S., Pirschel, I., Snower, D.J., 2017. A theory of price adjustment under loss aversion. J. Econ. Behav. Organ. 134, 78–95.
Bacon, R.W., 1991. Rockets and feathers: the asymmetric speed of adjustment of uk retail gasoline prices to cost changes. Energy Econ. 13 (3), 211–218.
Ball, L., Mankiw, N.G., 1994. Asymmetric price adjustment and economic fluctuations. Econ. J. 104 (423), 247–261.
Bayer, R.-C., Ke, C., 2018. What causes rockets and feathers? An experimental investigation. J. Econ. Behav. Organ. 153, 223–237.
Benabou, R., Gertner, R., 1993. Search with learning from prices: does increased inflationary uncertainty lead to higher markups? Rev. Econ. Stud. 60 (1), 69–93.
Borenstein, S., Cameron, A.C., Gilbert, R., 1997. Do gasoline prices respond asymmetrically to crude oil price changes? Q. J. Econ. 112 (1), 305–339.
Borenstein, S., Shepard, A., 1996. Sticky Prices, Inventories, and Market Power in Wholesale Gasoline Markets. Technical Report. National Bureau of Economic Research.
Bos, I., Vermeulen, D., 2021. On the microfoundation of linear oligopoly demand. B.E. J. Theor. Econ.. forthcoming
Brown, S.P., Yucel, M.K., 2000. Gasoline and crude oil prices: why the asymmetry? Econ. Financ. Rev. (Third Quarter) 23–29.
Byrne, D.P., De Roos, N., 2019. Learning to coordinate: a study in retail gasoline. Am. Econ. Rev. 109 (2), 591–619.
Cason, T.N., Friedman, D., 2002. A laboratory study of customer markets. BE J. Econ. Anal. Policy 2 (1), 1–43.
Chen, Y., Riordan, M., 2007. Price and variety in the spokes model. Econ. J. 117, 897–921.
Cooper, K., Schneider, H., Waldman, M., 2021. Limited rationality and the strategic environment: further evidence from a pricing game. J. Behav. Exp. Econ. 90, 101632.
Cornand, C., Heinemann, F., 2019. Experiments in macroeconomics: methods and applications. Handbook of Research Methods and Applications in Experimental Economics. Edward Elgar Publishing.
Davis, D., 2009. Pure numbers effects, market power, and tacit collusion in posted offer markets. J. Econ. Behav. Organ. 72 (1), 475–488.
Deaton, A., Muellbauer, J., 1980. An almost ideal demand system. Am. Econ. Rev. 70 (3), 312–326.
Deck, C.A., Wilson, B.J., 2008. Experimental gasoline markets. J. Econ. Behav. Organ. 67 (1), 134–149.
Duersch, P., Eife, T.A., 2019. Price competition in an inflationary environment. J. Monet. Econ. 104, 48–66.
Duffy, J., 2016. Macroeconomics: a survey of laboratory research. Handb. Exp. Econ. 2, 1–90.
Dufwenberg, M., Gneezy, U., 2000. Price competition and market concentration: an experimental study. Int. J. Ind. Organiz. 18 (1), 7–22.
Evans, G.W., Honkapohja, S., 2001. Learning and Expectations in Macroeconomics. Princeton University Press.
Fehr, E., Tyran, J.-R., 2001. Does money illusion matter? Am. Econ. Rev. 91 (5), 1239–1262.
Fiala, L., Suetens, S., 2017. Transparency and cooperation in repeated dilemma games: a meta study. Exp. Econ. 20 (4), 755–771.
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.
Frey, G., Manera, M., 2007. Econometric models of asymmetric price transmission. J. Econ. Surv. 21 (2), 349–415.
Green, E.J., Porter, R.H., 1984. Noncooperative collusion under imperfect price information. Econometrica 52 (1), 87–100.
Greiner, B., 2015. Subject pool recruitment procedures: organizing experiments with ORSEE. J. Econ. Sci. Assoc. 1 (1), 114–125.
Hanaki, N., Masiliūnas, A., 2021. Market Concentration and Incentives to Collude in Cournot Oligopoly Experiments. ISER DP 1131.
Hommes, C., Sonnemans, J., Tuinstra, J., Van de Velden, H., 2005. Coordination of expectations in asset pricing experiments. Rev. Financ. Stud. 18 (3), 955–980.
Horstmann, N., Krämer, J., Schnurr, D., 2018. Number effects and tacit collusion in experimental oligopolies. J. Ind. Econ. 66 (3), 650–700.
Huck, S., Normann, H.-T., Oechssler, J., 2004. Two are few and four are many: number effects in experimental oligopolies. J. Econ. Behav. Organ. 53 (4), 435–446.
Isaac, R.M., Walker, J.M., 1988. Group size effects in public goods provision: the voluntary contributions mechanism. Q. J. Econ. 103 (1), 179–199.
Lerner, A.P., 1934. The concept of monopoly and the measurement of monopoly power. Rev. Econ. Stud. 1 (3), 157–175.
Loy, J.-P., Weiss, C.R., Glauben, T., 2016. Asymmetric cost pass-through? Empirical evidence on the role of market power, search and menu costs. J. Econ. Behav. Organ. 123, 184–192.
Maskin, E., Tirole, J., 1988. A theory of dynamic oligopoly, II: price competition, kinked demand curves, and edgeworth cycles. Econometrica 56 (3), 571–599.
Meyer, J., von Cramon-Taubadel, S., 2004. Asymmetric price transmission: a survey. J. Agric. Econ. 55 (3), 581–611.
Morgan, J., Orzen, H., Sefton, M., 2006. An experimental study of price dispersion. Games Econ. Behav. 54 (1), 134–158.
Muth, J.F., 1961. Rational expectations and the theory of price movements. Econometrica 315–335.
Orzen, H., 2008. Counterintuitive number effects in experimental oligopolies. Exp. Econ. 11 (4), 390–401.
Peltzman, S., 2000. Prices rise faster than they fall. J. Polit. Economy 108 (3), 466–502.

Perdiguero-García, J., 2013. Symmetric or asymmetric oil prices? a meta-analysis approach. Energy Policy 57, 389–397.
Reagan, P.B., Weitzman, M.L., 1982. Asymmetries in price and quantity adjustments by the competitive firm. J. Econ. Theory 27 (2), 410–420.
Salop, S.C., 1979. Monopolistic competition with outside goods. Bell J. Econ. 10 (1), 141–156.
Yang, H., Ye, L., 2008. Search with learning: understanding asymmetric price adjustments. RAND J. Econ. 39 (2), 547–564.
Zhelobodko, E., Kokovin, S., Parenti, M., Thisse, J.-F., 2012. Monopolistic competition: beyond the constant elasticity of substitution. Econometrica 80 (6), 2765–2784.