Available online at www.sciencedirect.com

ScienceDirect

Journal of Economic Theory 167 (2017) 53–74

JOURNAL OF Economic Theory

www.elsevier.com/locate/jet

Information and coordination frictions in experimental posted offer markets ☆

Leif Helland [a,*], Espen R. Moen [a], Edgar Preugschat [b]

[a] *Department of Economics, BI Norwegian Business School, Norway*
[b] *Department of Economics, Technical University Dortmund, Germany*

Received 8 December 2014; final version received 8 May 2016; accepted 28 September 2016

Available online 17 October 2016

Abstract

We experimentally investigate buyer and seller behavior in small markets with two kinds of frictions. First, a subset of buyers may have (severely) limited information about prices, and choose a seller at random. Second, sellers may not be able to serve all potential customers. Such capacity constraints can lead to coordination frictions where some sellers and buyers may not be able to trade. Theory predicts very different equilibrium outcomes when we vary the set-up along these two dimensions. In particular, it implies that a higher number of informed buyers will lead to lower prices when sellers do not face capacity constraints, while prices may actually increase if sellers are capacity constrained, as shown by Lester (2011). In the experiment, the differences between the constrained and non-constrained case are confirmed; prices fall when sellers are not capacity constrained but either do not fall by much or even increase when they are not. We find that prices are quite close to the predicted equilibrium values except in treatments where unconstrained sellers face a large fraction of informed buyers. However, introducing noise into the theoretical decision making process produces a pattern of deviations that fits well with the observed ones.

☆ We are grateful for helpful comments from Kjell Arne Brekke, Urs Fischbacher, Knut-Eric Neset Joslin, Ola Kvaløy, Jean-Robert Tyran, Henrik Orzen, participants of the seminar of the Thurgau Institute of Economics, Kreuzlingen, November 2013, Samfunnsøkonomenes Forskermøte, Oslo, January 2014, the BI Workshop on Experimental Economics, Oslo, May 2014, the 9th NCBEE meeting, Aarhus, September 2014, the ESA European Meeting, Prague, September 2014, the Conference on Labor Market Models and their Applications, Sandbjerg Manor, October 2015, and the Search and Matching Conference, Amsterdam, May 2016. This research is financed by grant 212996/F10 from the Norwegian Research Council. Edgar Preugschat thanks the German Research Foundation for support through Priority Programme SPP 1764 and the Norwegian Research Council for grant No. 238159/F11.

* Corresponding author.
E-mail addresses: leif.helland@bi.no (L. Helland), espen.r.moen@bi.no (E.R. Moen), edgar.preugschat@tu-dortmund.de (E. Preugschat).

http://dx.doi.org/10.1016/j.jet.2016.09.007

Keywords: Posted offer markets; Coordination frictions; Information frictions; Experiment; Noisy decision making

1. Introduction

Many markets are affected by information frictions and capacity constraints. Information about prices or salaries is not always available before visiting a firm, or may be too costly to acquire. Furthermore, in some markets sellers can serve all customers, whereas in other markets sellers are constrained in their capacity. In labor markets, firms may or may not advertise wages, and may to a varying degree be capacity constrained depending on the number of equivalent open job slots they possess. In some retail markets, like the gasoline market, sellers are not capacity constrained, while buyers often have to visit the station to observe the price. In the customer-to-customer markets for used cars, by contrast, prices are advertised (although bargaining may occur), while the seller only has one car to sell and thus is capacity constrained.

The theoretical literature shows that even small changes in the capacity of sellers or the informedness of buyers can have a profound impact on market outcomes. Particularly, prices decrease sharply in the share of informed customers when capacity is not constrained, but might change little in the presence of such constraints. While the interactions between seller capacity and consumer information is well understood in theory, well-controlled empirical studies of such interactions are absent in the literature. This paper aims at filling the gap by setting up a laboratory experiment.

The results of our experiment confirm the different effects of higher informedness with and without capacity constraints. In particular, when sellers are capacity constrained prices fall little and may even increase as the number of informed buyers increases. However, our findings also indicate that as information increases, prices do not decrease as much as theory predicts when sellers do not face capacity constraints. This may imply that measures of consumer protection aiming at informing customers may be less effective in terms of reducing prices than previously thought. Our findings may also contribute to a more nuanced view regarding the potential consumer benefits of the rapid growth in essentially cost-free access to posted prices on the internet.

Our experiment is based on two strands of the theoretical literature on posted offer markets. One strand explores the effects of information frictions when sellers have unlimited capacity. Hence sellers can serve all buyers that show up, but some buyers are uninformed about prices. Varian (1980), Burdett and Judd (1983), Stahl (1989), and Janssen and Moraga-González (2004) analyze markets where only a fraction of buyers observe all the prices in the market. The remaining buyers are uninformed, and approach a seller at random. In the resulting equilibrium, sellers randomize over prices. As the fraction of informed buyers increases, the average price decreases, with the classic Bertrand equilibrium as the limiting case where price equals marginal cost.[1]

Another strand of this literature, starting with Montgomery (1991) and developed further by, among others, Burdett et al. (2001), explores the effects of search frictions when sellers have limited capacity to serve customers. Buyers have perfect information about prices, demand one unit of the good, and decide independently which seller to approach. Sellers only have a limited number of goods for sale, which can be normalized to one. Some sellers may get many and some

[1] An overview of this literature can be found in Baye et al. (2006).

sellers no customers, and when a queue forms, only one buyer will be served. Consequently, a coordination friction arises, as some market participants may end up without trading. The nature of the resulting equilibrium is in stark contrast to the equilibrium in which sellers are unconstrained. When sellers are capacity constrained, buyers trade off the price with the probability of obtaining the good, the price elasticity of demand is lower and the market price is strictly above the price when sellers are not capacity constrained (the Bertrand price). If the buyer–seller ratio is high, sellers' may even set prices close to the buyers' willingness to pay.

In a recent paper, Lester (2011) combines these two strands of the literature, by introducing information frictions into a market setting with capacity constraints. He demonstrates that increasing the fraction of informed buyers may produce effects that differ dramatically from those obtained in a setting where sellers are unconstrained. Generally, prices respond less when the fraction of informed customers increases compared with the unconstrained case, and they may even increase. This counter-intuitive result rests on the fact that a higher number of informed buyers stiffens the competition between informed buyers for the good. We refer to this as *Lester's paradox*.

In this paper we construct a unified model framework that allows us to study the interaction between limited information and capacity constraints in the lab. To this end we run six treatments, with a varying number of informed buyers and of units for sale. In each treatment there are three buyers with a unit demand, and the number of informed buyers ranges from 1 to 3. There are two sellers who in the "unconstrained" treatments can serve the entire market, and in the "constrained" treatments can serve at most one customer. Sellers simultaneously advertise a price, and buyers subsequently and simultaneously decide which seller to approach.

Our main contributions are the following. First, our experimental results show that when firms are capacity constrained, prices react substantially less to an increase in the share of informed buyers than they do in the unconstrained case. This is true for any fraction of informed buyers. Furthermore, when the number of informed buyers goes from two to three (all buyers) in the unconstrained case, Lester's paradox emerges. The average price goes up and the increase in transaction prices is significant at 10% confidence level. This indicates that the postulated relationship between prices and buyer information is not a mere theoretical curiosity. Second, we find that the model generally predicts prices better when sellers are capacity constrained than when they are not. Specifically, in the treatments with 2 and 3 informed buyers deviations are strong. Using the concept of Quantal Response equilibrium we analyze how noisy pricing behavior impacts on market prices. This allows us to explain the observed differences in deviations from Nash-equilibria well. We find that a little noise can push prices substantially above the Nash equilibrium when sellers are unconstrained in capacity while deviations are very small when they are constrained. Third, our experiment includes several important posted-offer market arrangements as special cases and thus makes them comparable. In particular, our evidence on the varying impact of noise on best responses shows the value of such cross-market comparisons. Further, some of the market structures contained in our experiment have been tested in isolation before. As we replicate their main findings our results appear to be robust to variations in the experimental setup. We discuss these related experiments in detail in the last part of Section 4.

The paper is organized as follows. In the next section we outline and explain the predictions from theory using a unified environment for all six market structures. In section 3 we present our design and hypotheses. Section 4 presents our main results, analyzes buyer and seller behavior, and relates our findings to existing experiments. Section 5 concludes. All additional material is gathered in an online appendix.

2. Theoretical predictions

In the following we briefly outline the theoretical framework on which our treatments are based and refer the reader to the supplementary online appendix 1 for the details. The framework encompasses the model of Lester (2011), the standard directed search model of Burdett et al. (2001), a version of Varian (1980),[2] and the classic Bertrand model as special cases.

The economy is populated by a number of $S = 2$ sellers (or "firms") and $B = 3$ buyers, all of which are risk neutral.[3] Buyers have a unit demand with a reservation price normalized to 100. The model consists of two stages: First, sellers simultaneously set and commit to prices $p_s \in [0, 100]$. In the second stage buyers simultaneously make buying decisions. A number $U \geq 0$ of uninformed buyers independently and randomly choose a seller, where each seller is visited with equal probability by a given buyer. Further, there are $N \geq 1$ (with $N + U = 3$) informed buyers who can costlessly observe all prices offered in the market and choose at which seller to buy.

Regarding the number of units each firm has for sale we distinguish between two cases. In the first case (denoted by index $z = c$), all firms are capacity constrained, and each firm has exactly one unit for sale. Hence, if two buyers show up, only one can be served. In the second case (denoted by $z = n$) firms are not capacity constrained, and each firm has B units for sale. In this case a seller can always serve all the customers that show up. We then denote a specific market setting by T_N^z, summarizing the parameter constellation of $z \in \{c, n\}$ and $N \in \{1, 2, 3\}$ which we will vary in the experiment.

For a given combination of z and N the expected payoff of a seller s is $\pi_s(p_s, p_{-s}) = \mu(p_s, p_{-s})p_s$, where $\mu(p_s, p_{-s})$ is the expected number of sales given the own price and the prices of other sellers. The expected payoff of a buyer i conditional on choosing a seller s is $v_i(\theta_{-i}^s) = \eta(\theta_{-i}^s)(1 - p_s)$, where $\eta(\theta_{-i}^s)$ is the probability of getting the good at seller s given that the other buyers go to the same seller with probability θ_{-i}^s in a symmetric equilibrium. If sellers are not capacity constrained, $z = n$, this probability is always equal to 1. If the sellers are capacity constrained the probability is typically strictly less than 1. If no seller is chosen the payoff is zero. It follows from the assumptions on uninformed buyers that $\theta_i^s = 1/S$ for all $i \in U$. We focus on sub-game perfect equilibria with symmetric (mixed) strategies. While this is the standard assumption in the theoretical literature, it is also justified in our experimental set-up since market participants are anonymous and new markets are formed randomly in each period, making coordination difficult.[4]

When varying the number of informed buyers N and the capacity z, different kinds of equilibria emerge which we summarize in Table 1. With no capacity constraints and all three buyers informed, Bertrand competition emerges, and the equilibrium price is zero (T_3^n). With uninformed buyers, the equilibrium price cannot be zero, as a seller can obtain a strictly positive profit by setting the price to 100 and rip off the uninformed buyers that come along. With two uninformed buyers (T_2^n), a seller that sets a price of 100 obtains an expected profit of 100 (two uninformed buyers arrive with probability 1/2 each). A seller that sets a price of 50 obtains the same expected profit if he attracts the informed seller with probability 1. It can be shown that in the resulting equilibrium, sellers randomize their prices on the interval [50, 100]. With only one informed buyer (T_1^n) sellers randomize on the interval [20, 100]. When comparing the treatments

[2] In contrast to Varian (1980), we set costs equal to a constant normalized to zero.

[3] The constraints on the number of agents are only imposed to focus on the more interesting cases.

[4] See also our discussion of collusion in section 4.

Table 1
Theoretical predictions: expectation, support and distribution of prices.

z	N		
	1	2	3
n	$E(p) = 69.3$	$E(p) = 40.2$	$E(p) = 0.0$
	$p \in [50, 100]$	$p \in [20, 100]$	
	$F(p) = \left(\frac{2p-100}{p}\right)$	$F(p) = \left(\frac{5p-100}{4p}\right)$	
	$E(p_T) = 66.7$	$E(p_T) = 33.3$	$E(p_T) = 0.0$
c	$E(p) = 86.3$	$E(p) = 66.7$	$E(p) = 72.7$
	$p \in [75, 100]$		
	$F(p) = \left(\frac{4p-300}{p}\right)$		
	$E(p_T) = 85.7$	$E(p_T) = 66.7$	$E(p_T) = 72.7$

p: posted prices; p_T: transaction prices

with no capacity constraints along the dimension of buyer informedness, the classic result that expected prices decline with N emerges, since competition for informed buyers becomes more intense.

Second, when there are capacity constraints and only one informed buyer (T_1^c), the informed buyer will always approach the seller with the lowest price. A seller that sets a price of 100 sells with probability 3/4, and hence gets an expected profit of 75. If a seller sets a price of 75, and attracts the informed buyer with certainty, he also gets an expected profit of 75. It can be shown that in equilibrium, sellers randomize over the price interval [75, 100]. If there are capacity constraints and more than one informed buyer (T_2^c and T_3^c) buyers play (symmetric) mixed strategies so that in equilibrium they are indifferent between sellers. While in general there can be equilibria with mixed strategies on the sellers' side, for the parameter constellations of our treatments there will be only symmetric equilibria where sellers play pure strategies.

Comparing to the equilibria without capacity constraints, we see that expected prices are higher for a given number of informed buyers, N (see Table 1). In stark contrast to the case of no capacity constraints the equilibrium price can increase in the number of informed buyers given S and B (Lester's paradox). In our setup this occurs when moving from T_2^c to T_3^c.

To gain intuition as to why more informed buyers' may lead to a higher price, we divide the effects of more informed buyers into two. First we have a rip-off effect, fewer uninformed buyers imply that there are fewer customers that are insensitive to prices, and this reduces the incentives to charge a high price. In contrast to the standard case, the presence of capacity constraints implies an additional competition effect that goes in the opposite direction. A higher number of informed buyers leads to stronger competition for sellers with a low price. As buyers not only care about the price but also about the probability of getting the good, higher competition makes low price sellers less attractive. The price elasticity will be lower if there is more competition on the buyer side, because congestion decreases the attractiveness of a low price seller. Thus, the competition effect tends to increase prices when the number of informed buyers goes up.

3. Parameters and procedures

To test the different predictions of an increase in informed buyers when there are capacity constraints compared to the case without such constraints we use a 2×3 design. The experiment consists of six treatments where we vary both the capacity constraint, i.e. $z = c$ or $z = n$ and the number (equivalently, the share) of informed buyers, i.e. $N = 1$, $N = 2$, and $N = 3$. Our design

Table 2
Treatments and blocks.

Treatment	# of blocks	
	Oslo	Konstanz
T_1^n	5	0
T_2^n	2	3
T_3^n	5	0
T_1^c	5	0
T_2^c	2	3
T_3^c	2	3

allows us to better isolate the effect of capacity constraints when varying buyer informedness. To have such an explicit comparison is particularly important as the theoretically predicted price increase from T_2^c to T_3^c is relatively small so that we initially did not expect to actually observe a price increase.

In the experiment uninformed buyers were computer programs flipping fair coins to determine where to purchase. All informed buyers and all sellers were human subjects.

In all treatments prices and payoffs were measured in experimental currency units (ECUs). Buyers valuations were set to 100 ECU, and sellers marginal costs to 0 ECU.

In each treatment, one market constellation is played. Each treatment consists of five blocks, and each block consists of three markets. Each of our sessions consisted of either two or three blocks. Each session investigated only one treatment. All treatments lasted 50 periods, and each period corresponded to a two-stage game. Subjects were randomly allocated a label prior to the start of trading, and kept this label for the 50 periods of play. Buyers were labeled "Blue", "Red" and "Green", and sellers were labeled "Circle" and "Square". Subjects were randomly assigned to the three markets within each block at the start of each period in such a way that all labels were present in all markets. No subject was ever allocated across blocks, and no information on behavior in other blocks was conveyed to subjects. Unique subjects were used in all blocks. Thus, observations at the block-level are independent.[5]

A total of 360 subjects were used for the experiment, and a total of 18000 individual decisions (by humans) were collected. Some sessions used students from the University of Konstanz, Germany, other sessions used students from the Norwegian Business School in Oslo, Norway. As we show in the online appendix, there are no significant differences between blocks collected in Oslo and Konstanz. We therefore pool data from the two locations. Data were collected between November 2012 and February 2014. Table 2 provides an overview of the experiment.

Subjects were recruited online using the ORSEE system (Greiner, 2004). The experiment was programmed in z-Tree (Fischbacher, 2007), and was contextualized as a market, using terms such

[5] Concerns that heterogeneous buyers can lead to a coordination equilibrium have been raised in the theoretical literature on directed search (Coles and Eeckhout, 2000). In our environment buyers have access to a minimal identification technology, since they play in fixed labels. We find that this is not sufficient to promote buyer coordination in treatments T_2^c and T_3^c: empirical visit probabilities of informed buyers match the theoretical visit probabilities of such buyers very closely (see Section 4 below). Another concern is that sellers may use labels (and set prices with decimals) to facilitate collusion on prices above equilibrium. However, only in treatments T_2^n and T_3^n prices are substantially above equilibrium levels. Moreover, in T_3^n the average price relative to the buyer valuation is very close to the ones found in the other studies discussed and where no fixed labels where used. We discuss the issue of collusion in detail in the online appendix.

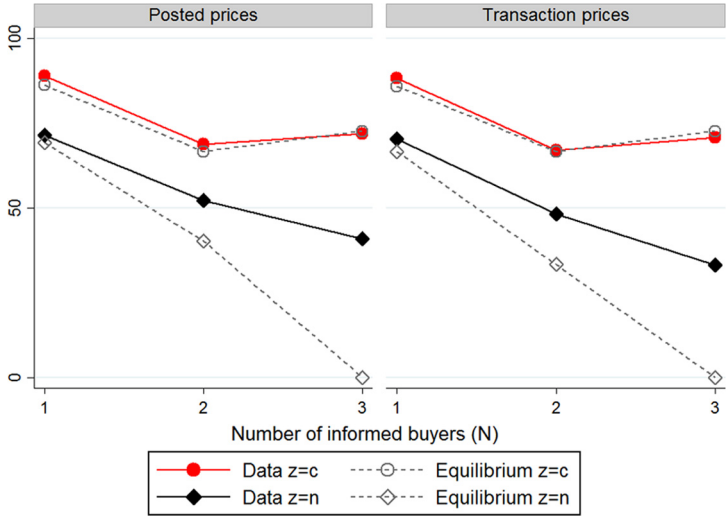

Fig. 1. Average posted prices and transaction prices for each treatment – data and theory.

as "sellers", "buyers", "prices" and "queues". Subjects were randomly allocated to numbered cubicles on entering the lab to break up social groups. After being seated, each subject was issued written instructions and these were read aloud by the administrator of the experiment to achieve public knowledge of the rules.[6] There were no test periods, and no control questions to check understanding. Sellers were allowed to post prices with two decimals. Strict anonymity was preserved throughout. Each period consisted of a posting stage, and a purchase stage. Sellers posted prices simultaneously, human buyers then observed the prices posted and simultaneously chose one seller to go to. In treatments with capacity constraints, if a queue formed at a seller the transacting buyer (human or computer program) was drawn with a uniform probability from the queue. At the end of each period all subjects got feedback on the whole history of posted prices, queues at each seller, transactions in the market he or she was operating, as well as own profit.

After period 50 was concluded, accumulated ECUs were converted to NOK or Euros (depending on the location) at a pre-announced exchange rate, and subjects were paid privately on leaving the lab. On average a session took 70 minutes. In the Oslo treatments average earnings were 54 US dollars. In the Bertrand treatment (T_3^n) all subjects got a (pre-announced) flat fee of 27 US dollars plus whatever they earned in the session. This was done in order to avoid sellers not earning money in the experiment. In all other treatments subjects got what they earned plus a show up fee. Earnings in the Konstanz treatments were adjusted to give the same consumer purchasing power as the Oslo treatments.

4. Results

Market behavior Fig. 1 provides a treatment-by-treatment comparison of observed prices and their theoretical counterparts, averaged over all periods and all blocks (see Table 3 for the actual numbers).

[6] A full set of instructions and a sample of screen-shots can be found in the online appendix.

Table 3
Observed average posted and average transaction prices.

Treatment	T_1^n	T_2^n	T_3^n	T_1^c	T_2^c	T_3^c
Posted prices	71.4	52.3	41.6	89.1	68.9	71.9
Transaction prices	68.8	44.3	32.2	88.0	66.5	70.3
Coefficient of variation (posted prices)	0.27	0.45	0.55	0.15	0.17	0.17

As can be seen, average posted prices are remarkably close to the theoretical equilibrium values in treatments T_1^n, T_1^c, T_2^c, and T_3^c, while they deviate substantially in treatments T_2^n and, especially T_3^n, the market with Bertrand competition. Transaction prices are similarly close to their respective equilibrium values, and also exhibit the strongest deviations for treatments T_2^n and T_3^n.[7] For both posted and transacted prices the predicted patterns are clearly visible: for a given number of informed buyers prices with capacity constrained sellers are always above prices with unconstrained sellers. Further, prices decrease with the number of informed buyers when there are no capacity constraints and either slightly fall or slightly increase otherwise. We summarize this in the following informal result.

Result 1 *(Average prices: data and theory).* Average posted prices are very close to the theoretically expected prices in treatments T_1^n, T_1^c, T_2^c, and T_3^c, while they deviate substantially in treatments T_2^n and T_3^n. Transaction prices are similarly close and exhibit the same pattern of deviations.

We test differences between treatments with one-sided Wilcoxon rank sum (WRS) tests using blocks as units of observation.

For posted prices the differences between treatment T_1^n and T_1^c ($W = -2.402$; $p = .008$), T_2^n and T_2^c ($W = -2.611$; $p = .005$), and T_3^n and T_3^c ($W = -2.611$; $p = .005$) are all significant at the 1% level. Furthermore, posted prices decrease when going from treatment T_1^n to T_2^n ($W = 2.611$; $p = .005$); and when going from T_2^n to T_3^n ($W = 2.193$; $p = .014$). These price decreases are significant at the 5% level or better. WRS tests also reveal that posted prices decrease significantly from treatment T_1^c to T_2^c ($W = 2.611$; $p = 0.005$). The increase in posted prices from treatment T_2^c to T_3^c, however, is not significant at conventional levels ($W = -1.149$; $p = 0.125$). Nonetheless it is close to being significant at the 10% level, and we find this quite remarkable, considering that theory predicts an increase in prices between T_2^c and T_3^c by a measly 6 ECUs, and that the WRS test uses only five observations in each treatment.

Result 2 *(Treatment differences for posted prices).* The differences in posted prices between the treatments with and without capacity constraints for a given number of informed buyers are all significant. Furthermore, the decrease in posted prices when going from treatment T_1^n to T_2^n, from T_2^n to T_3^n, and from T_1^c to T_2^c are significant.

Our results become stronger for transaction prices. The differences between treatment T_1^n and T_1^c ($W = -2.402$; $p = .008$), T_2^n and T_2^c ($W = -2.611$; $p = .005$), and T_3^n and T_3^c ($W =$

[7] In T_2^n the deviation in percent of the theoretical price is 30.1 for posted prices and 33.0 for transaction prices. In T_3^n this measure is not defined. For the other four treatments deviations in percent of theoretical posted prices are between 3.3 and 1.1, and between 3.3 and 0.3 for transaction prices.

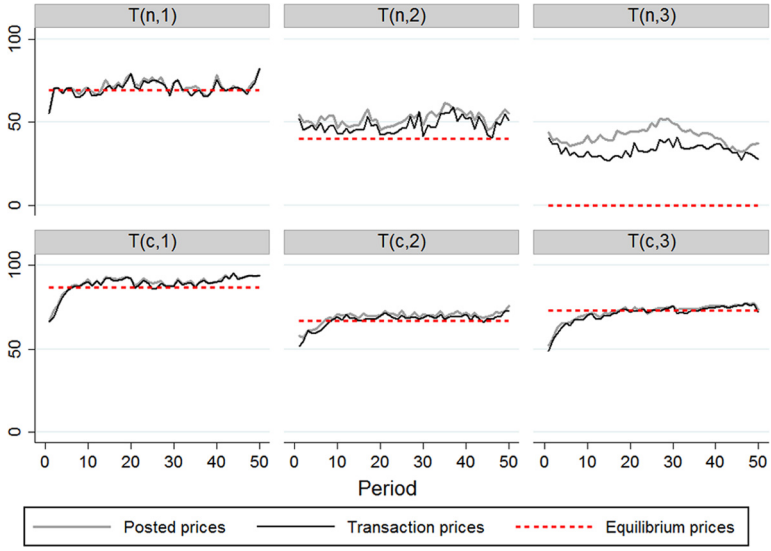

Fig. 2. Average posted prices and average transaction prices over periods.

-2.611; $p = .005$) are all significant at the 1% percent level. Transaction prices also decrease when going from treatment T_1^n to T_2^n ($W = 2.611$; $p = .005$), and when going from T_2^n to T_3^n ($W = 2.402$; $p = .008$). These reductions are significant at the 1% level or better. WRS tests also show that transaction prices decrease significantly from treatment T_1^c to T_2^c ($W = 2.611$; $p = .005$). Finally, the increase in transaction prices from treatment T_2^c to T_3^c is now significant at the 10% level, and almost significant at the 5% level ($W = -1.567$; $p = .059$).[8]

Result 3 *(Treatment differences for transaction prices).* The differences in transaction prices between the treatments with and without capacity constraints for a given number of informed buyers are all significant. Transaction prices also decrease significantly when going from treatment T_1^n to T_2^n, from T_2^n to T_3^n, and from T_1^c to T_2^c, while the increase in transaction prices is weakly significant when going from T_2^c to T_3^c.

In appendix 2 we run treatment regressions. These regressions confirm the results from the non-parametric tests, but also indicate that results are neither driven by differences in lab population (Oslo versus Konstanz), the interaction of treatments and time, or idiosyncratic differences across blocks. We comment further on the coefficients of variation after looking at convergence over rounds.

Fig. 2 displays the average posted prices and transaction prices per period by treatment. Treatments T_2^n and T_3^n evidently deviate substantially from the theoretical predictions, and do not seem

[8] With one exception, results are unchanged if the WSR-tests use only data from periods 11–48 (after learning has taken place and before the onset of endgame effects). The one exception is that the drop in posted prices from T_2^n to T_3^n is no longer significant in a one sided test when data are restricted in this way ($W = 0.940$; $p = .174$).

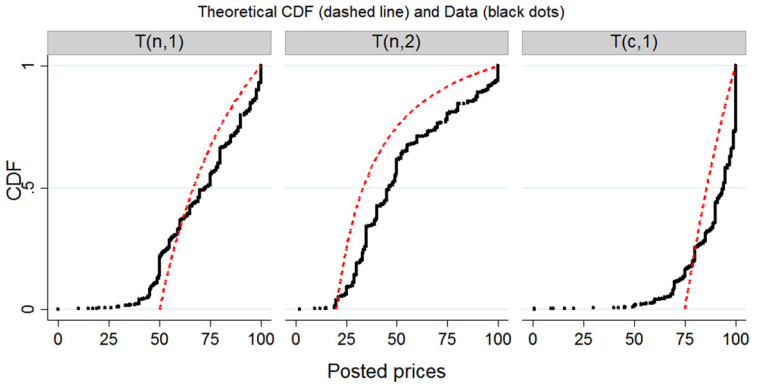

Fig. 3. Cumulative price distributions T_1^n, T_2^n, and T_1^c: data and theoretical prediction.

to converge to it. For the other treatments prices seem to approach the equilibrium value, or some value close to equilibrium, fairly rapidly, and then remain there.[9]

We now take a closer look at the distribution of prices. In the last row of Table 3 we report the coefficient of variation of posted prices for each treatment. These correspond reasonably well to the theoretical coefficients of variation in the treatments where prices are dispersed in equilibrium. For treatment T_1^n the predicted (observed) coefficient of variation is 0.20 (0.27), for T_2^n the predicted (observed) coefficient is 0.48 (0.45), while in treatment T_1^c the predicted coefficient of variation is 0.08, about half of the observed one (0.15). Unlike predicted by theory, there is also dispersion in prices for the remaining three treatments. Prices are to a limited degree dispersed in treatments T_2^c and T_3^c, whereas there is large variation in prices in the Bertrand treatment. We explore these deviations from theory further below. In Fig. 3 we compare the empirical distributions of treatments T_1^n, T_2^n and T_1^c to their theoretical counterparts.[10] In the figure dashed lines indicate theoretical price distributions, while solid lines are empirical posted prices.

First, data match the support of the equilibrium distributions in these treatments reasonably well. Using all periods, about 68% of the data lie within the support in treatment T_1^n. The corresponding numbers for treatments T_2^n and T_1^c are 85% and 80%, respectively.[11] While data track the theoretical distributions reasonably well, the empirical distributions do not have the convex shape of the theoretical distributions.

[9] In appendix 2 we run dynamic regressions to check formally for convergence. These regressions show that we can only be confident that posted prices weakly converge to the equilibrium value for treatment T_3^c. In the other treatments there is evidence of weak convergence, and with the exception of treatments T_2^n and T_3^n, these processes converge to a value close to the theoretical prediction. The precise definitions of strong and weak convergence are provided in the appendix.

[10] See Table 1 for the distribution functions.

[11] This finding is robust over rounds. Using only periods 39 to 48 (where behavior should have stabilized) improves the number only slightly T_1^n (to 72% of data lying within the support), while the numbers stay unchanged for the two other treatments. That a large share of the data lie within the support is perhaps not so surprising. A simple logic shows that prices below the lower bound of the support are dominated by the rip-off price of 100 ECU. As an example, consider T_2^n where there are two informed buyers. If the seller succeeds in posting the lower price she sells two units to the informed buyers, and has an equal chance of selling her third unit to the uninformed buyer. Thus, *given* that the seller has the lower price the expected profit equals her posted price times 2.5. Posting the rip-off price of 100 provides an expectation of 50. Thus any price below $50/2.5 = 20$, which is the lower bound of the support, is dominated by the rip-off price.

Result 4 *(Price distributions).* In treatments T_1^n, T_2^n, and T_1^c, where theory predicts price distributions, the empirical distributions of posted prices roughly match their predicted counterparts. While the shape is not always well matched, the support is matched quite closely.

Below we analyze how deviations from theoretical price distributions can be accounted for by noisy seller responses. Prior to that, however, we address the question of how consistent buyer responses are with theory.[12]

Buyer behavior For the theoretical pricing strategies to make sense, sellers need to believe that buyers will respond optimally to the prices they post. Do buyers respond optimally to posted prices? In treatments T_1^n to T_3^n and T_1^c the unconditional best response of an informed buyer is to (try to) purchase from the seller with the lower price. In these treatments a high fraction of purchase attempts follow the predicted best responses.

Result 5 *(Buyer behavior I).* When prices between sellers differ, the average percentage of buyers that go for the lower price is 92.4 in treatment T_1^n, 98.6 in treatment T_2^n, 97.1 in treatment T_3^n, and 88.4 in treatment T_1^c.[13]

In treatments T_2^c and T_3^c the equilibrium conditions require informed buyers to randomize over which seller to choose such as to make other informed buyers indifferent in their choice of a seller. To evaluate the optimality of buyer responses in these treatments we used the following procedure for each of these treatments. First we calculated for each informed buyer in every period the predicted equilibrium probability of choosing a fixed seller, given the pair of actual prices posted. Recall that sellers have fixed labels in our experiment; either square or circle. In our calculations the fixed seller is the one labeled square. We then estimate a logistic regression. The dependent variable in this regression is a dummy equal to one if the buyer in question went to seller square, and zero otherwise. This dummy was regressed on the equilibrium probability of choosing seller square. The regressions were estimated with buyer random effects. Table 4 reports the results.

The regression coefficients are precisely estimated, the fit of the models is good in each case, and the probability of choosing seller square, given a pair of prices, is positively and significantly related to the theoretical probability of making such a choice in both treatments. Taking exponents on both sides of the regressions and reorganizing, we obtain the estimated probabilities of choosing seller square for each observation (each buyer in each period) in each treatment. Averaging over the theoretical and the estimated probabilities (of choosing square) for each treatment, returns the results reported in Table 5.

Fig. 4 shows that these averages do not mask a weak buyer response to changes in the theoretical probability. In the figure circles provide the average fraction of buyers visiting seller square (y-axis) for brackets of length 0.025 on the theoretical probability of doing so (x-axis). If the

[12] Separate analysis of buyer and seller behavior is common in experiments where buyer reactions are not automated, see for instance Anbarci and Feltovich (2014) and Cason and Noussair (2007).

[13] The average payment in excess of the lower price paid by subjects in ECU (standard deviation) and by treatment was 13.5 (17.5) in T_1^n; 10.3 (16.7) in T_2^n; 11.1 (14.8) in T_3^n; 8.6 (10.1) in T_1^c; 6.2 (5.8) in T_2^c; and 8.2 (8.4) in T_3^c. In treatments T_1^n, T_2^n, T_3^n and T_1^c irrational buyer decisions are mainly due to one or two outlying subjects that make repeated – and often costly – mistakes. In T_2^c and T_3^c visiting the high price seller is more evenly distributed over buyers, as one would expect in equilibrium.

Table 4
Logistic regressions with random effects for buyers.

Treatment	T_2^c	T_3^c
Equilibrium probability of choosing seller "square" given posted prices	4.99***	3.69***
	(.271)	(.263)
Constant	−2.53***	−1.83***
	(.192)	(.143)
# of data points	1500	2250
# of buyers	30	45
Log likelihood	−737.1	−1436.5
χ^2 model	338.8***	198.0***

Dependent variable: choice of seller square. Standard errors in parentheses. Significant at level: *** 1%; ** 5%; * 10%.

Table 5
Average equilibrium- and estimated probability of choosing square.

Treatment	T_2^c	T_3^c
Equilibrium probability	.542	.508
Estimated probability	.541	.511

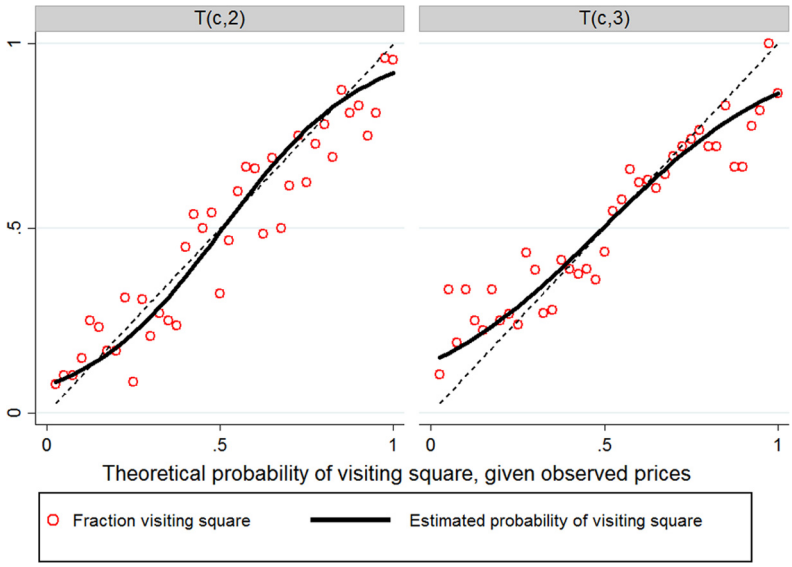

Fig. 4. Estimated and actual buyer reactions in treatments T_2^c and T_3^c.

theoretical probability is a perfect predictor of the actual choices, all circles will be located on the (dashed) 45-degree line.

From Fig. 4 we conclude that the theoretical probability has substantial predictive power over its entire range. The black lines are estimated probability curves, using the regressions in Table 4.

We appreciate that these curves are close to linear over the range of the theoretical probability, indicating the absence of threshold effects. The slope of the estimated probability curve is closer to unity for T_2^c than for T_3^c, where buyers overshoot somewhat for low theoretical probabilities, and undershoot somewhat for high theoretical probabilities. Still, the general impression is that theoretical choice probabilities are remarkably close to the actual ones also in T_3^c.

Result 6 *(Buyer behavior II).* For treatments T_2^c and for T_3^c the average probability of buying at a specific seller is almost identical to the average predicted probability given prices. The estimated probabilities follow the predicted probabilities very closely.

While we cannot easily compare them, it seems that informed buyer responses correspond better with theory when responses are more complicated to work out (i.e. when mixed strategies are required) than when they are not (i.e. where buyers have dominant pure strategies). In appendix 4 we investigate the confidence interval around the standard errors for the mean reactions of buyers in treatments T_2^c and T_3^c. In the vast majority of cases we are unable to reject the null of a perfect match between theoretical and empirical choice probabilities using a 95% confidence interval.

Given that buyer responses are very close to the theoretical predictions for treatments T_2^n, T_3^n, T_2^c, and T_3^c, and fairly close for treatments T_1^n and T_1^c, we investigate the sources of deviations coming from seller behavior given optimal buyer behavior.

Seller behavior As buyers' decisions are by and large consistent with theory, the observed deviations from theory should primarily be caused by sellers' behavior. Fig. 1 is suggestive about the pattern of deviations. First, the capacity constrained treatments are on average much closer to the theoretical predictions than the non-constrained treatments. Furthermore, in the absence of capacity constraints the deviations become stronger as the share of informed buyers increases. The first observation may appear surprising as the capacity constrained treatments involve a seemingly more complex reasoning for both sellers and buyers.

We think that it is reasonable to assume that subjects in a laboratory setting make mistakes relative to the behavioral requirements of the equilibria. Errors seem to be more likely if they are associated with a smaller loss in profits. Furthermore, the effects of mistakes may depend on how the opponent firm reacts to them, i.e. the derivative of the best response function to a change in the opponent's price. To illustrate this we compare the (pure strategy) best-response functions across treatments. In the treatments without capacity constraints, T_1^n–T_3^n, as well as in treatment T_1^c, the best response to a price in the equilibrium support is to slightly undercut by setting the price ε below the competitor's price.[14] Hence the best response to a deviation by the other seller in these treatments is to increase the price with the same amount. The reaction is much more muted in treatments T_2^c and T_3^c, where a unit price increase is only followed by a raise of 0.61 units.

The first row of Table 6 shows the best response to a unit increase of the opponent's price, starting from the equilibrium expected transaction price (where we approximate $\varepsilon \approx 0$). In all treatments prices move in the same direction, but the reactions are much weaker in the last two treatments. The two last rows of the table show the corresponding absolute and relative increases

[14] Recall that the equilibrium support is $p \in [50, 100]$ for T_1^n, $p \in [20, 100]$ for T_2^n, and $p \in [75, 100]$ for T_1^c. If the competitor sets a price below the minimum of the support the best response is to set a price of 100.

Table 6
Deviations from equilibrium.

Treatment	T_1^n	T_2^n	T_3^n	T_1^c	T_2^c	T_3^c
Best response price increase	1.00	1.00	1.00	1.00	0.61	0.61
Change in profits	2.00	2.50	3.00	1.00	0.88	0.88
% change in profits	1.50	3.00	∞	1.17	1.51	1.39

All measures are relative to the (expected) equilibrium transaction price reported in Table 1. "Best response price increase" refers to an optimal price increase when the opponent increases the price by one unit. The amount of undercutting is set to $\varepsilon = 0$. Both initial and resulting profits are calculated assuming that the player (marginally) undercuts the price.

in profits, respectively. The changes in profits are stronger on average for the treatments without capacity constraints. Moreover, for these treatments the profit changes are more pronounced the more informed buyers there are. Thus, upward deviations are more likely to be reinforced in treatments T_2^n and T_3^n where profit changes are large.

In the Bertrand case, the Nash equilibrium price is zero. Hence the only possible deviations are upwards. Furthermore, the loss associated with a deviation if the opponent plays Nash is zero, and the equilibrium strategy is indeed a weakly dominated strategy. Hence, in the presence of noise, a rational player will not play zero, but set a strictly higher price. Furthermore, as there is a strong strategic complementarity in the price setting behavior of the sellers, this rationalizes why prices may spiral away from zero. For treatment T_2^n, where we also observe strong upward deviations, this argument is less clear, as the expected Nash price is below but close to $1/2$. Thus it is not obvious how a rational player will react to noisy play by the opponent. In order to shed more light on the observed deviations from Nash equilibrium we analyze noisy play using the concept of quantal response equilibrium (QRE).

QRE has been successfully applied in the experimental literature to rationalize deviations from Nash outcomes in various games.[15] We isolate the effect of noise on seller's strategies by taking optimal play in the buyers' sub-game as given. Our aim is twofold: First, we investigate to what extent QRE can capture the deviations from theory with respect to average prices and price distributions. Second, following Goeree and Holt (2001), we use the QRE concept to measure the sensitivity of the Nash equilibrium with respect to noise. Our conjecture with respect to this second goal is that market settings with a steeper best response function and larger associated relative profit gains are more likely to be sensitive to the introduction of noise.

The structural approach to QRE, first introduced by McKelvey and Palfrey (1995), is based on a random payoff model, where the profit π of a seller i, given the other seller's cumulative distribution function for pricing strategies, F_{-i}, is perturbed by a random error: $\hat{\pi}_i(p, F_{-i}) = E_{p_{-i}}\pi_i(p, F_{-i}) + \epsilon_{i,p}$.[16] Each player assigns a probability to a given action equal to the probability that this action is a best response given the error. The resulting quantal responses can thus be interpreted as noisy best responses. Equilibrium requires players' beliefs about the opponents

[15] See Goeree et al. (2008) for a brief review. An application of QRE to a Bertrand market is given in Baye and Morgan (2004). Two alternative approaches are the ε-Equilibrium concept by Radner (1980) and the introduction of "noise traders" (De Long et al. (1990)) who set prices according to a given (exogenous) distribution. We have utilized the latter concept but found that it does not explain the data as well as the QRE (our results are available on request). See also the partial approach to noise trading in the context of the Bertrand model in Dufwenberg and Gneezy (2000).

[16] See McKelvey and Palfrey (1995) for the details of the definition and the derivation of the logit specification. Alternatively, QRE can be defined in an axiomatic way, see Goeree et al. (2005).

Table 7
QRE estimates with common noise parameter.

Treatment	T_1^n	T_2^n	T_3^n	T_1^c	T_2^c	T_3^c
μ			0.062			
Expected QRE price	71.5	53.4	36.9	85.2	70.6	64.0
Average distance QRE and data			0.714			
Distance QRE and data by treatment	0.489	0.593	0.479	0.966	0.329	1.429
Expected Nash price	69.3	40.2	0.0	86.3	66.7	72.7
Distance Nash and data by treatment	0.799	1.409	5.358	1.083	1.691	1.621

Minimized distances between the c.d.f.s of QRE and data (square root of the sum of squared deviations) estimated on a grid of integer prices. When distance is measured as an average over treatments, each treatment receives the same weight. Payoffs in the treatments without capacity constraints (n) are scaled by factor $1/3$.

mixing probabilities to be correct. While this equilibrium requirement puts high demands on the rationality of the players if taken literally, the resulting rule for the mixing probabilities is very intuitive: the probability of choosing an action increases with its expected payoff.

We assume a Gumbel distribution for the error, i.i.d. across actions and players, leading to the logistic form of the quantal response. In our symmetric case the quantal responses are given by the (identical) distribution function over strategies for each seller, $F^Q(p)$, that solves the following functional fixed point:

$$F^Q(p) = \int_0^p \exp(\frac{1}{\mu} E\pi_i(p, F^Q(x)))dx \Big/ \int_0^{100} \exp(\frac{1}{\mu} E\pi_i(p, F^Q(x)))dx \ , \forall p \in [0, 100],$$

where $\mu > 0$ is the parameter governing noisiness.[17] With the logistic specification we follow the majority of the experimental literature, making our findings comparable. This specific choice of the distributional form, together with the restriction that all treatments are estimated with the same noise, puts discipline on the resulting QRE which depends only on one free parameter, μ.[18] When μ approaches infinity, all prices are equally likely, which can be interpreted as completely noisy strategies. On the other extreme, if μ goes to zero, the quantal response approaches the best response of the underlying pricing game, and behavior converges to a Nash equilibrium.

To judge whether QRE can rationalize the data we fit the QRE cumulative distribution functions of posted prices to the corresponding CDFs in the data by choosing a common μ to minimize the sum of squared deviations.[19] We follow Goeree et al. (2003) and normalize the maximum possible payoffs in treatments T_1^n–T_3^n (which is 300) to 100 to obtain the same payoff range in all treatments. Table 7 reports the parameter estimate and the implied expected prices for simultaneously fitting μ, the average distance, the implied distances treatment-by-treatment, as well as the corresponding distances between the theoretical distributions and the data. Fig. 5 displays the implied CDFs together with distributions from data and theory.

[17] We use the reciprocal value of the noise parameter λ often used in the literature, i.e $\lambda \equiv \frac{1}{\mu}$.

[18] See Haile et al. (2008) for a discussion of the issue of falsifiability of QRE.

[19] We follow the approach by Baye and Morgan (2004). The more standard maximum likelihood estimation is not suitable in our case as in some treatments the density of the QRE distribution approaches zero for part of the support of the empirical distribution. For our purposes there is otherwise no difference between the procedures.

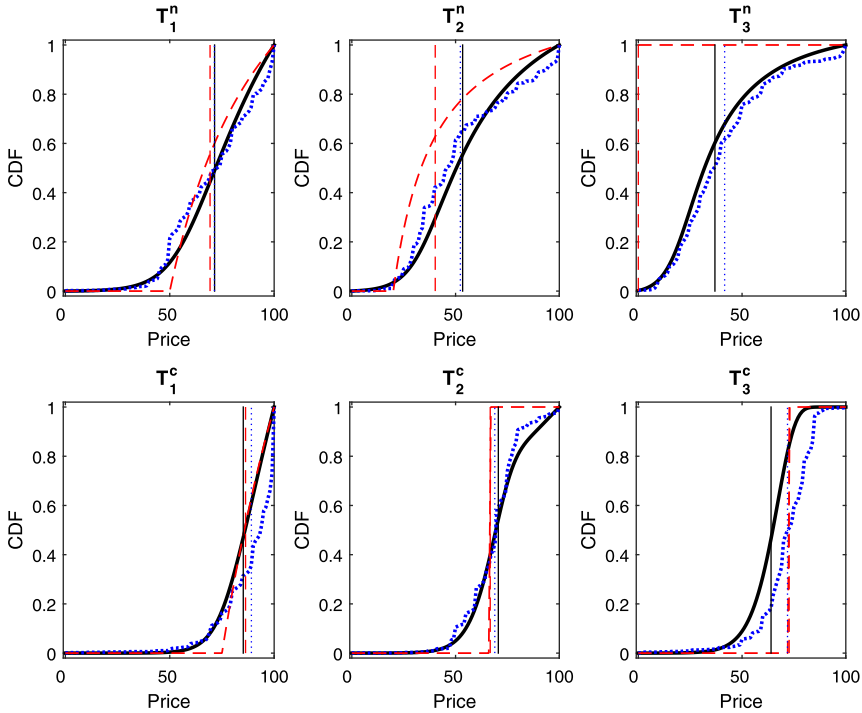

Fig. 5. QRE distributions (solid black), theoretical distributions (dashed red) and actual posted price distributions (blue dots) and corresponding average prices indicated by the vertical lines. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

As can be seen, fitting the QRE distributions allows us to match both the expected prices and price distributions of the data quite well, with the exception of treatment T_3^c where the expected price is not well matched. However, in all treatments, the CDF of the QRE fits the data (individually and on average) better than the CDF implied by Nash-equilibrium. As a consequence, the QRE estimates capture well the large deviations of the average prices in treatments T_2^n and T_3^n, while staying close to the equilibrium distributions in the remaining treatments. Furthermore, QRE rationalizes the observed price distributions in the treatments where the Nash equilibrium predicts only point prices (i.e. T_3^n, T_2^c, and T_3^c).[20]

Another way to see this is by comparing the observed CVs and the CVs implied by the QRE. Such a comparison reveals a remarkably tight fit, taking into account that the QRE is estimated with a common noise parameter for all six treatments. In particular, the CVs implied by the QRE indicates that the QRE capture the observed prices well both in treatments with Nash point prices and in treatments with Nash price distributions. Details are provided in the online appendix 5.1.

Result 7 *(Seller behavior I).* The QRE distributions and expected prices match well the empirical distributions of T_1^n to T_2^c and roughly match the distribution for T_3^c and can thereby rationalize the observed price dispersion in treatments T_3^n, T_2^c and for T_3^c. Furthermore, the QRE estimates account well for the substantial deviations from Nash equilibrium in treatments T_2^n and T_3^n.

[20] In section 5 of the appendix we also display the CDFs of the QRE estimated treatment by treatment.

Table 8
Absolute and relative (percentage) changes in QRE expected price.

Treatment	T_1^n	T_2^n	T_3^n	T_1^c	T_2^c	T_3^c
Relative change	2.48	10,0	∞	0.81	1.42	1.96
Absolute change	1.719	4.025	16.010	0.695	0.944	1.423

Relative (percentage) and absolute changes in expected price when moving from Nash equilibrium ($\mu = 0$) to QRE equilibrium ($\mu = 0.01$). All numbers are absolute values.

Finally, the coefficients of variation of the QRE distributions closely match the observed coefficients of variation.

The QRE estimates are the result both of the added noise to individual best responses and equilibrium interaction. To better see the direct effect of noise across treatments we consider two additional exercises. First we analyze a scenario where one player noisily best responds to a player that plays the Nash equilibrium strategy. Second, we study the effect of noise close to the Nash equilibrium.

First, we consider the noisy response of a seller if the opponent plays his Nash equilibrium strategy, given the estimated noise parameter $\mu = 0.062$. That is, for each treatment we characterize the distribution

$$\tilde{F}(p) = \int_0^p \exp(\frac{1}{\mu} E\pi_i(p, F(x)))dx \Big/ \int_0^{100} \exp(\frac{1}{\mu} E\pi_i(p, F(x)))dx \ , \forall p \in [0, 100],$$

where $F(x)$ is the Nash equilibrium strategy in a given treatment. $\tilde{F}(p)$ captures the direct effect (or first round effect) of noise given that the opponent plays his Nash equilibrium strategy.[21] The density of $\tilde{F}(p)$, together with its expected value and the expected Nash price, is given in the online appendix 5.2 for all six treatments. The deviation is enormous in the Bertrand case, where the expected noisy price response is 50 while the Nash price is 0. The reason is that when playing against the Nash strategy ($p = 0$), the pay-off is zero for all choices of p. Furthermore, also in T_2^n, the expected price with noise is substantially higher than the expected Nash price. For the other treatments however, the differences between the expected Nash prices and the expected noisy responses are small. Thus, the deviations in the "first round" carry over to the ones we find for the estimated QREs.

Second, we conduct a reversed exercise in which we measure the sensitivity to noise when starting at the Nash equilibrium and then move to QRE with a low value of μ. Table 8 reports the relative and absolute change in the expected price when changing μ from 0 to .01. It reveals that the expected prices diverge from the Nash equilibrium at a substantially lower rate when sellers are capacity constrained than when they are not. Moreover, the absolute values of the price elasticities with respect to noise are increasing in the number of informed buyers for both the unconstrained and the constrained treatments. Note that both the relative and the absolute price changes are highest in T_3^n and T_2^n where observed deviations from the Nash equilibrium are most pronounced.

[21] Let Γ denote the QRE mapping so that $\tilde{F}(p) = \Gamma F(x)$. Then the QRE distribution F^Q is a fixed-point of Γ, $F^Q = \Gamma F^Q$. Let $\Gamma^k(F(x))$ denote the mapping performed k times. Then $F^Q = \lim_{k\to\infty} \Gamma^k F(x)$ provided that the limit exists.

Result 8 *(Seller behavior II).* 1. The "first round" effects of noise on best responses are strong in treatments T_2^n and T_3^n. Thus, the direct effects of noise seem to carry over to the deviations found in the estimated QREs. 2. The unconstrained treatments react more strongly to an increase in noise (μ) when starting at a level of μ close to zero. For both the constrained and the unconstrained treatments the expected QRE price responds more strongly to an increase in μ when the number of informed buyers is higher.

In summary, the QRE analysis confirms our intuition coming from the best responses which suggests that the capacity constrained treatments should be less sensitive to noisy play, despite being computationally more complex.

Finally, an alternative possibility to explain the large deviations is collusive behavior. Evidence of collusive behavior has been documented in a number of market experiments.[22] One may think that the complementarity in pricing would give strong incentives to cooperate.[23] However, there are two countervailing forces at play in our experiment. First, sellers are constantly re-matched within a block, which makes it very difficult to establish and maintain a tacit agreement to collude. In line with this Orzen (2008) finds that prices are close to the Nash equilibrium in duopolies resembling the setup in T_1^n and T_2^n if subjects are randomly re-matched from period to period, while prices deviate substantially upwards if subjects stay in the same markets. A similar point is made in Ochs (1990) for a more general setting in which capacity constraints create coordination problems. Second, in treatments T_2^n and T_3^n where we observe the largest deviations from equilibrium there are also strong incentives to undercut the opponent's price, making successful collusion unlikely.

In appendix 5.4 we analyze our data with respect to collusion. In particular we look for simple arrangements, such as constant prices and uncomplicated rotation schemes, by investigating descriptive statistics of levels and variation in prices and profits. We also follow Friedman et al. (2015) in plotting the probability of a price change in the current period against the price in the previous period. Finally, we perform analysis of end-game effects. We do not find support for collusive behavior in any of these analyses. Furthermore, our indicators suggest that, if anything, coordinating behavior seems to be the least pronounced in treatments T_2^n and T_3^n.

Comparison to previous experiments How do our results compare to previous studies?[24] In our treatments with capacity constraints we test for equilibria in which buyer-coordination is

[22] E.g. Davis and Holt (1994) find strong evidence of supra-competitive pricing in Bertrand competition with heterogeneous sellers, and Friedman et al. (2015) find strong evidence of long run collusive behavior for both Cournot duopolies and triopolies. In Davis and Holt (1994) subjects stay in the same group over the course of the experiment. In Friedman et al. (2015) subjects play for 1200 periods, and are re-matched into new groups after each block of 400 periods.

[23] Potters and Suetens (2009) find that collusion in the sense of supra-competitive prices is more often observed when there are strategic complements than when there are strategic substitutes. They compare treatments with complements and substitutes rather than changing the degree of complementarity in a setting with continuous payoff functions and fixed pairs of subjects over many rounds.

[24] The following is an incomplete list of experimental studies of market environments that relate to our study. Anbarci and Feltovich (2013) examine the two-price model of Coles and Eeckhout (2000); Cason and Friedman (2003) examine the noisy sequential-search model of Burdett and Judd (1983); Cason and Datta (2006) examine a model due to Robert and Stahl (1993), in which sellers can advertise at a cost; Otto and Bolle (2011) and Abrams et al. (2000) examine markets in which there is bargaining after matching. Deck and Wilson (2006) investigate an environment in which sellers have the ability to track customers and offer discriminatory prices based on observed history. There is also a large literature on posted offer markets in which buyers enter the market in a random order, see Ketcham et al. (1984), Davis and Holt (1994) and the surveys in chapter 6-8 of Plott and Smith (2008). Since buyers enter in a random order and make their

not permitted. Cason and Noussair (2007) (hereafter CN) investigate small experimental markets (a two-seller, three-buyer treatment; and a three-seller, two-buyer treatment) to test the Burdett et al. (2001) model against the large market model of Montgomery (1991). Their study is based on a design very close to our T_3^c. In their two-seller, three buyer treatment CN find average posted prices of 83.7 for periods 39–48.[25] In comparison, average posted prices in periods 39–48 is 75.2 in our T_3^c treatment. So, while CN overshoot the equilibrium value by 11 percentage points in these 10 periods, we overshoot by only 2.2 percentage points. Finally, our data converge more rapidly on a value closer to equilibrium in the T_3^c treatment than the CN data does.[26]

Anbarci and Feltovich (2014) (hereafter AF) also run a T_3^c treatment. Their design differs from ours (and that of CN) in important ways.[27] They run their T_3^c treatment for 20 periods. Averaging posted prices over all periods, AF undershoot the equilibrium value by 13.5 percentage points.[28] Averaging posted prices only over the last 5 periods reduces this undershooting to 7.9 percentage points.

In general, buyer reactions in our T_3^c treatment are substantially more in line with theory than those of CN and AF. While we observe the same *qualitative* biases in buyer reactions as AF and CN, these biases are far weaker in our T_3^c treatment than in theirs.[29]

AFs study is concurrent and independent to ours. They also examine a treatment similar to our T_2^c. In contrast to us they fail to find support for Lester's paradox. Their interpretation of this deviation from theory is based on fair pricing. Our design differs in important ways from theirs. In particular, we have more independent observations for each treatment and more than twice as many rounds within each observation. The latter difference might be important as behavior is converging slowly within the first 10 rounds of play.[30] A further difference is that our design also allows us to benchmark the impact of information frictions against the case where sellers do not face capacity constraints.

Morgan et al. (2006) (hereafter MOS) test the Varian (1980) model, in which sellers are not constrained. Their design differs from ours in a number of ways.[31] In their two-seller treatments, they test for the change in posted prices as the fraction of informed buyers is increased from $\frac{1}{2}$ to $\frac{5}{6}$. Their qualitative results are in line the results we obtain for T_2^n and T_1^n. Increasing the share of informed buyers reduces posted prices, as it should do in equilibrium. As in our T_2^n

purchases one at a time, there can be no coordination frictions in these markets. Also, in much of this literature sellers are heterogeneous. In a Bertrand setting, seller heterogeneity seems to drive prices down towards the competitive level.

[25] After behavior has stabilized, but before the end game effects set in.

[26] To see this, compare the dynamic regressions in the appendix of this paper with those in CN.

[27] AF run the same subjects in various treatments, using only three separate matching blocks. The design combines within – and between subjects comparisons, controlling for order effects.

[28] Average transaction prices over all rounds undershoots the equilibrium value of the T_3^c treatment by a full 14.4 percentage points in AF. In an identical design by Anbarci and Feltovich (2013), undershooting in the T_3^c treatment is reduced to 7.1 percentage points (compare Table 2 in AF with Table 5 in Anbarci and Feltovich, 2013).

[29] Compare our Fig. 4 with Figure 4 in CN and Figure 4 in AF.

[30] This is in line with the finding of Anbarci and Feltovich (2014) that the contradictory result is alleviated when only the last 5 rounds are taken into account. See Cason and Noussair (2007) for related finding regarding convergence behavior.

[31] Among other things they ran the same subjects in various treatments, using six separate matching blocks, combining a within – and between subjects design with control for order effects. In contrast to our design, MOS also used robots to mimic the responses of informed as well as uninformed buyers.

and T_1^n treatments, the overshooting of posted prices compared to equilibrium values increases substantially with the fraction of informed buyers.[32]

As in our T_2^n and T_1^n treatments, the support of the empirical price distributions match the support of the theoretical price distributions well in the two-seller treatments of MOS. Furthermore, and again as in our T_2^n and T_1^n treatments, empirical price distributions in MOS are somewhat closer to theoretical distributions the larger the fraction of informed buyers is.

Several tests of Bertrand duopoly competition exist. In Dufwenberg and Gneezy (2000) (hereafter DG), and in Dufwenberg et al. (2007) (hereafter DGGN) buyer reactions are automated, and treatments are conducted with pen and paper. In DG marginal costs are 1 and in DGGN 2, while buyer valuations are 100 in both experiments. Sellers compete for 10 periods. Average posted (transaction) prices over these 10 periods are 34.5 (27.1) in DG and 28.7 (21.9) in DGGN.

In Abrams et al. (2000) (hereafter ASY) sellers and buyers were randomly matched, and buyers had the opportunity to search at a cost after a match was formed and posted prices had been observed.[33] Buyers and sellers were humans, the experiment lasted for 25 periods, and was computerized. Buyer valuations were set at 120 and marginal costs at 0. Re-scaled to a valuation of 100 the average posted prices over the 25 periods were 40.5, while the average transaction prices were 24.2. Common to DG, DGGN and ASY is that prices are volatile and do not drop monotonically over time towards the equilibrium in which prices equate marginal costs.

In our T_3^n treatment average posted prices over all periods were 41.0, while average transaction prices were 33.2. The time paths of average posted and transaction prices are displayed in Fig. 2. As is evident, neither price measure falls monotonically over time. Thus, our T_3^n results are comparable to those of DG, DGGN and ASY in the sense that prices deviate substantially from equilibrium and stay in the same broad range as in existing experiments; that transaction prices are substantially below posted prices; and that prices do not fall monotonically over time.

Summing up, we succeed in replicating the behavioral patterns of existing experiments for treatments T_1^n, T_2^n, T_3^n, and T_3^c.

5. Conclusion

In this paper we have tested the effects of information and coordination frictions due to capacity constraints in small posted offer markets. Our experiments have confirmed the theoretical predictions that the presence of capacity constraints dramatically changes the effect of an increased share of informed buyers. In the absence of capacity constraints prices clearly fall as more buyers become informed. In the presence of capacity constraints prices fall only slightly, or even increase as more buyers become informed. Our experiment confirms the counter-intuitive prediction of Lester's paradox. In addition, our experiment demonstrates that different market settings lead to differently strong deviations from theory. Our results indicate that these deviations are mostly due to deviating seller behavior and not so much due to buyers' choices. Noisy price setting can rationalize the observed price choices.

[32] In the MOS treatment with 1/2 of buyers informed overshooting is 5.3 percentage points, while it increases to 15.4 percentage points in the treatment with 5/6 of buyers informed. In our T_1^n treatment (1/3 of buyers informed) the overshooting is 1.8 percentage points, compared to 11.9 percentage points in our T_2^n treatment (2/3 of buyers informed).

[33] The search opportunity should be of no consequence in the Bertrand treatments, and is generally not used.

Appendix. Supplementary material

Supplementary material related to this article can be found online at http://dx.doi.org/10.1016/j.jet.2016.09.007.

References

Abrams, E., Sefton, M., Yavas, A., 2000. An experimental comparison of two search models. Econ. Theory 16 (3), 735–749.

Anbarci, N., Feltovich, N., 2013. Directed search, coordination failure, and seller profits: an experimental comparison of posted pricing with single and multiple prices. Int. Econ. Rev. 54 (3), 873–884.

Anbarci, N., Feltovich, N., 2014. Pricing in competitive search markets: experimental evidence of the roles of price information and fairness perceptions. Mimeo.

Baye, M.R., Morgan, J., 2004. Price dispersion in the lab and on the Internet: theory and evidence. Rand J. Econ., 449–466.

Baye, M.R., Morgan, J., Scholten, P., 2006. Information, search, and price dispersion. In: Hendershott, T. (Ed.), Handbook on Economics and Information Systems, vol. 1. Elsevier, Amsterdam.

Burdett, K., Judd, K.L., 1983. Equilibrium price dispersion. Econometrica 51 (4), 955–969.

Burdett, K., Shi, S., Wright, R., 2001. Pricing and matching with frictions. J. Polit. Econ. 109 (5), 1060–1085.

Cason, T., Friedman, D., 2003. Buyer search and price dispersion: a laboratory study. J. Econ. Theory 112 (2), 232–260.

Cason, T., Noussair, C., 2007. A market with frictions in the matching process: an experimental study. Int. Econ. Rev. 48 (2), 665–691.

Cason, T.N., Datta, S., 2006. An experimental study of price dispersion in an optimal search model with advertising. Int. J. Ind. Organ. 24 (3), 639–665.

Coles, M.G., Eeckhout, J., 2000. Heterogeneity as a coordination device. UPF Economics & Business Working Paper (510).

De Long, J.B., Shleifer, A., Summers, L.H., Waldmann, R.J., 1990. Noise trader risk in financial markets. J. Polit. Econ., 703–738.

Deck, C.A., Wilson, B.J., 2006. Tracking customer search to price discriminate. Econ. Inq. 44 (2), 280–295.

Davis, Douglas D., Holt, C.A., 1994. Market power and mergers in laboratory markets with posted prices. Rand J. Econ. 25 (3), 467–487.

Dufwenberg, M., Gneezy, U., 2000. Price competition and market concentration: an experimental study. Int. J. Ind. Organ. 18 (1), 7–22.

Dufwenberg, M., Gneezy, U., Goeree, J.K., Nagel, R., 2007. Price floors and competition. Econ. Theory 33 (1), 211–224.

Fischbacher, U., 2007. z-Tree: Zurich toolbox for ready-made economic experiments. Exp. Econ. 10 (2), 171–178.

Friedman, D., Huck, S., Oprea, R., Weidenholzer, S., 2015. From imitation to collusion: long-run learning in a low-information environment. J. Econ. Theory 155, 185–205.

Goeree, J.K., Holt, C.A., 2001. Ten little treasures of game theory and ten intuitive contradictions. Am. Econ. Rev., 1402–1422.

Goeree, J.K., Holt, C.A., Palfrey, T.R., 2003. Risk averse behavior in generalized matching pennies games. Games Econ. Behav. 45 (1), 97–113.

Goeree, J.K., Holt, C.A., Palfrey, T.R., 2005. Regular quantal response equilibrium. Exp. Econ. 8 (4), 347–367.

Goeree, J.K., Holt, C.A., Palfrey, T.R., 2008. Quantal response equilibrium. In: Durlauf, S.N., Blume, L. (Eds.), The New Palgrave Dictionary of Economics. Palgrave Macmillan, Basingstoke.

Greiner, B., 2004. An online recruitment system for economic experiments. In: Kremer, K., Macho, V. (Eds.), Forschung und wissenschaftliches Rechnen. In: GWDG-Ber., vol. 63. Ges. für Wiss. Datenverarbeitung, Göttingen, pp. 79–93.

Haile, P.A., Hortaçsu, A., Kosenok, G., 2008. On the empirical content of quantal response equilibrium. Am. Econ. Rev. 98 (1), 180–200.

Janssen, M.C., Moraga-González, J.L., 2004. Strategic pricing, consumer search and the number of firms. Rev. Econ. Stud. 71 (4), 1089–1118.

Ketcham, J., Smith, V.L., Williams, A.W., 1984. A comparison of posted-offer and double-auction pricing institutions. Rev. Econ. Stud. 51 (4), 595–614.

Lester, B., 2011. Information and prices with capacity constraints. Am. Econ. Rev. 101 (4), 1591–1600.

McKelvey, R.D., Palfrey, T.R., 1995. Quantal response equilibria for normal form games. Games Econ. Behav. 10 (1), 6–38.

Montgomery, J.D., 1991. Equilibrium wage dispersion and interindustry wage differentials. Q. J. Econ. 106 (1), 163–179.
Morgan, J., Orzen, H., Sefton, M., 2006. An experimental study of price dispersion. Games Econ. Behav. 54 (1), 134–158.
Ochs, J., 1990. The coordination problem in decentralized markets: an experiment. Q. J. Econ., 545–559.
Orzen, H., 2008. Counterintuitive number effects in experimental oligopolies. Exp. Econ. 11 (4), 390–401.
Otto, P.E., Bolle, F., 2011. Matching markets with price bargaining. Exp. Econ. 14 (3), 322–348.
Plott, C.R., Smith, V.L., 2008. Handbook of Experimental Economics Results, vol. 1. Elsevier.
Potters, J., Suetens, S., 2009. Cooperation in experimental games of strategic complements and substitutes. Rev. Econ. Stud. 76 (3), 1125–1147.
Radner, R., 1980. Collusive behavior in noncooperative epsilon-equilibria of oligopolies with long but finite lives. J. Econ. Theory 22 (2), 136–154.
Robert, J., Stahl, D.O., 1993. Informative price advertising in a sequential search model. Econometrica, 657–686.
Stahl, D.O., 1989. Oligopolistic pricing with sequential consumer search. Am. Econ. Rev., 700–712.
Varian, H.R., 1980. A model of sales. Am. Econ. Rev. 70 (4), 651–659.