

Bayona, Anna; Brandts, Jordi; Vives, Xavier

Working Paper

Supply Function Competition, Private Information, and Market Power: A Laboratory Study

CESifo Working Paper, No. 6005

Provided in Cooperation with:

Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

Suggested Citation: Bayona, Anna; Brandts, Jordi; Vives, Xavier (2016) : Supply Function Competition, Private Information, and Market Power: A Laboratory Study, CESifo Working Paper, No. 6005, Center for Economic Studies and ifo Institute (CESifo), Munich

This Version is available at:

<https://hdl.handle.net/10419/145040>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



Working Papers

www.cesifo.org/wp

Supply Function Competition, Private Information, and Market Power: A Laboratory Study

Anna Bayona
Jordi Brandts
Xavier Vives

CESIFO WORKING PAPER NO. 6005
CATEGORY 11: INDUSTRIAL ORGANISATION
JULY 2016

An electronic version of the paper may be downloaded

- *from the SSRN website:* www.SSRN.com
- *from the RePEc website:* www.RePEc.org
- *from the CESifo website:* www.CESifo-group.org/wp

ISSN 2364-1428

Supply Function Competition, Private Information, and Market Power: A Laboratory Study

Abstract

In the context of supply function competition with private information, we test in the laboratory whether—as predicted in Bayesian equilibrium—costs that are positively correlated lead to steeper supply functions and less competitive outcomes than do uncorrelated costs. We find that the majority of subjects bid in accordance with the equilibrium prediction when the environment is simple (uncorrelated costs treatment) but fail to do so in a more complex environment (positively correlated costs treatment). Although we find no statistically significant differences between treatments in average behaviour and outcomes, there are significant differences in the distribution of supply functions. Our results are consistent with the presence of sophisticated agents that on average best respond to a large proportion of subjects who ignore the correlation among costs. Experimental welfare losses in both treatments are higher than the equilibrium prediction owing to a substantial degree of productive inefficiency.

JEL-Codes: C920, D430, L130.

Keywords: divisible good auction, generalized winner's curse, correlation neglect, electricity market.

Anna Bayona
ESADE Business School
Spain - 08172 Sant Cugat
anna.bayona@esade.edu

Jordi Brandts
Institut d'Anàlisi Econòmica (CSIC)
& Barcelona GSE / Spain
jordi.brandts@uab.cat

Xavier Vives
IESE Business School
University of Navarra
Spain - 08034 Barcelona
xvives@iese.edu

July 2016

We thank José Apesteguía, Maria Bigoni, Colin Camerer, Enrique Fatas, Dan Levin, Cristina Lopez-Mayan, Margaret Meyer, Rosemarie Nagel, Stanley Reynolds, Albert Satorra, Arthur Schram, and Jack Stecher for useful comments and discussions. Anna Bayona acknowledges the financial support from Banc de Sabadell. Jordi Brandts acknowledges financial support from the Spanish Ministry of Economics and Competitiveness (Grant ECO2014-59302-P) and the Generalitat de Catalunya (AGAUR Grant 2014 SGR 510). Xavier Vives acknowledges financial support of the Spanish Ministry of Economics and Competitiveness (Grant ECO2015- 63711-P) and the Generalitat de Catalunya (AGAUR Grant 2014 SGR 1496). The usual disclaimers apply.

1 Introduction

We design a laboratory experiment which captures the complexity of the bidding and information environments which are representative of the real-world markets characterised by competition in schedules such as wholesale electricity markets, markets for pollution permits, as well as liquidity and Treasury auctions.¹ We provide experimental evidence of behaviour and outcomes in a market where each seller has incomplete information about her costs, receives a private signal, and competes in supply functions. The aim of our experiment is to study the relationship between information frictions and market power, and to examine the implications of the complexity of the environment: uncorrelated costs versus positively correlated costs.

We consider a market where firms compete in terms of supply functions (see Klemperer and Meyer 1989) and with incomplete cost information (Vives 2011). The latter paper finds that, in a unique linear Bayesian equilibrium, private information with cost correlation generates market power that exceeds the full-information benchmark.² When costs are positively correlated, the model predicts that the supply function's slope is steeper and the intercept is lower leading to higher expected market prices and profits than when costs are uncorrelated. The mechanism that explains these results can be stated as follows. A seller receives a private signal that is informative about her random costs. A fully rational seller who is strategic must also realise that, when costs are positively correlated, a high price conveys the information that costs are high; therefore, to protect herself from adverse selection, she should compete less aggressively than if costs were uncorrelated. So if all sellers are fully rational then the combination of private information and strategic behaviour leads to greater market power when costs are correlated than when they are uncorrelated. The mechanism that relates higher cost correlation to increased market power is connected to a generalised version of the winner's curse (Ausubel et al. 2014) that extends this concept to multi-unit demand auctions.³

The experimental design is as follows. We employ a between-subjects experimental design with two treatments that only differ in the correlation among costs.⁴ In each treatment, subjects were randomly assigned to independent groups of twelve subjects, each comprising four markets of three sellers. Within each group, we applied random matching between rounds in order to retain the theoretical model's one-shot nature. The buyer was simulated, and subjects were assigned the role of sellers. Subjects received a private signal about the uncertain cost and

¹The following papers argue for the importance of demand or cost uncertainty among bidders that compete in schedules: wholesale electricity markets (Holmberg and Wolak 2015); liquidity auctions (Cassola et al. 2013); Treasury auctions (Keloharju et al. 2005); carbon dioxide emission permits (Lopomo et al. 2011).

²The supply function equilibrium with uncorrelated costs coincides with the full-information equilibrium since sellers do not learn about cost uncertainty from prices.

³The connection is established in the Related Literature section.

⁴In a between-subjects design participants are either part of the control group or the treatment group but cannot participate in both.

were then asked to submit a (linear) supply function. As in the theoretical model, and in contrast to most of the experimental literature, we used a normally distributed information structure that well approximates the distribution of values in naturally occurring environments. After all decisions had been made, the uniform market price was calculated and each subject received detailed feedback about her own performance, the market price, and the behaviour and performance of rivals in the same market. At the end of the experiment, we administered a post-experiment questionnaire that asked about participant’s demographic information, bidding behaviour, and understanding of the game. Subjects were given incentives that incorporated both fixed and variable components, where the variable component reflected the individual participant’s performance during the game.

Our experimental data are in line with some of the theoretical predictions. First, we confirm that average behaviour in the uncorrelated costs treatment closely matches the theoretical prediction in the experiment’s early stages and that, over time, the average supply function tends even further toward the equilibrium supply function. This finding is important because the uncorrelated costs treatment yields a benchmark against which to compare behaviour in the positively correlated costs treatment. Second, we find that the features of the equilibrium that are common to both treatments are observed in the data. In particular, we observe that the supply function’s intercept is increasing in a bidder’s signal realisation. This result is consistent with subjects understanding that a higher signal implies a higher average intercept of the marginal cost, which means they should set a higher ask price for the first unit offered and this leads, in turn, to a higher supply function intercept. We observe also that the supply function’s slope is unrelated to the signal received.

Analysing the distribution of individual choices, we find that the cumulative distribution of supply function slopes in the positively correlated costs treatment first order stochastically dominates the cumulative distribution of slopes in the uncorrelated costs treatment, both in the first and last five rounds (in a stronger way in the latter). This result shows that differences in behaviour between treatments are consistent with the *direction* predicted by the theoretical model. However, we cannot reject the hypothesis that *average* supply functions are the same in both treatments. In the positively correlated costs treatment we observe that the average supply function is substantially flatter and has a higher intercept than predicted by the equilibrium, which is consistent with subjects being too strongly guided by the signal received. This divergence persists even as subjects gain bidding experience.

In terms of experimental outcomes, we also find differences (albeit insignificant ones) between treatments with regard to market prices, profits, and efficiency of the allocations. However, we do not find the predicted differences in market power between the two treatments. In the positively correlated costs treatment, bidders forgo a large percentage of profits—an outcome typical of auctions where bidders ignore the adverse effects of correlation among costs. Moreover,

experimental welfare losses in both treatments are larger than predicted by the equilibrium owing to the considerable extent of productive inefficiency. These results suggest that subjects in the positively correlated costs treatment fall prey to the generalised winner’s curse and, as a result, compete too aggressively in comparison with the equilibrium prediction. The greater heterogeneity in behaviour in the positively correlated (than in the uncorrelated) costs treatment merits further exploration.

We offer a detailed analysis of behaviour that proceeds in three parts. First, in each treatment we analyse theoretically subjects’ strategic incentives. We remark that the environment that we study is complex both in terms of the information setting (the combined uncertainty regarding cost and private signal) and the market structure (competition in supply functions and transaction costs which vary with quantity sold). The positively correlated cost treatment adds a further layer of complexity since the market price is informative about a seller’s cost; hence the equilibrium logic requires subjects to form not only correct beliefs about the economic environment but also correct higher-order beliefs. In other words, the equilibrium of the positively correlated costs treatment requires that sellers believe that other sellers are taking advantage of the correlation structure, that they believe that everyone believes this, that everyone believes that everyone believes this... We find that a (sophisticated) subject in our positively correlated costs treatment who best responds to the rivals’ average choices has an incentive to bid a supply function *between* the equilibrium of the uncorrelated costs treatment and the equilibrium of the positively correlated costs treatment. It follows that behaviour and outcomes between treatments (and for different subject types—namely, sophisticated types capable of best responding and naïve who fall prey to the generalised winner’s curse—) are less differentiated than predicted by the equilibrium (which assumes that all bidders are fully rational).

The second part of our analysis takes a descriptive approach to organising, via cluster analysis, the observed heterogeneity in individual-level behaviour; we then compare the clusters so derived to the various theoretical benchmarks. We use average behaviour in blocks of five rounds. In the *uncorrelated* costs treatment, we identify two clusters. One cluster assembles choices that are close to the equilibrium and theoretical best response to the average choice; this cluster includes 58% of the subjects in the first five rounds of bidding, a proportion that increases to 72% in the last five rounds. The other cluster groups subjects whose supply functions are steeper and have a lower intercept than the equilibrium of the uncorrelated costs treatment; this cluster includes 42% and 28% of the subjects in (respectively) the first and last five rounds. In the *positively correlated* costs treatment we identify three clusters. One cluster groups subjects whose supply functions are close to the benchmark for subjects who fall prey to the generalised winner’s curse, which includes 58% of the subjects in the first five rounds of bidding and 50% in the last five rounds. Another cluster gathers subjects whose bidding behaviour is not inconsistent with sophisticated behaviour in the sense that their supply functions are close to the

theoretical best response to the average supply function; this cluster includes 36% (respectively 42%) of subjects in the first (respectively last) five rounds of bidding. Given the complexity of the bidding environment, it is interesting to find a substantial number of subjects whose behaviour approximates the average best response. The third derived cluster groups the remaining subjects, who bid a steep supply function with a low intercept (6% and 8% of subjects in the first and last five rounds, respectively). Although there are few treatment differences in *average* behaviour there are substantial differences in the *distributions* of individual behaviour. These are driven by the two treatments' different levels of strategic complexity, which lead to the identification of distinct types of subjects (naïve and sophisticated) in each treatment. Notice, however, that even with a substantial percentage of subjects whose behaviour approximates the average best response only moves the aggregate behaviour slightly away from the uncorrelated costs treatment.

Finally, we analyse how behaviour changes across rounds and find treatment-based differences in the determinants of the evolution of behaviour. In particular, the *best-response dynamics* factor figures more prominently in the uncorrelated than in the positively correlated costs treatment, whereas *imitation of the best* plays a smaller role in the uncorrelated than in the positively correlated costs treatment. In each treatment, both *imitation of the average* and *reinforcement learning* are important factors in explaining the evolution of behaviour across rounds. The combination of these determinants, the initial conditions and strategic complexity explains why behaviour evolves toward the equilibrium in the uncorrelated costs treatment but not in the positively correlated costs treatment.

The rest of our paper is organised as follows. We review the related literature in Section 2 and explain the theoretical model in Section 3. Section 4 describes the experimental design, after which Section 5 presents our main results. In Section 6 we analyse in detail the behaviour observed during our experiment. We conclude in Section 7. (The experiment's instructions may be found in Appendix C.)

2 Related Literature

Our experimental paper studies competition in supply schedules within an information environment that includes both positively correlated costs and uncorrelated costs. This environment is reminiscent of the one described in Goeree and Offerman (2003), who also use normally distributed values and error terms. However, their paper compares behaviour in cases of common versus uncorrelated private values in a single-unit, second-price auction whereas ours compares behaviour in cases of correlated versus uncorrelated costs in a supply function, uniform price auction.

With respect to the competition environment, some early experiments used bid functions in

auctions with incomplete information (e.g., Selten and Buchta 1999), but few laboratory experiments have sought to analyse, as we do, competition in supply functions. Exceptions include the work of Bolle et al. (2013), who focus on testing predictions of Supply Function Equilibrium concept, as well as Brandts et al. (2014), whose paper compares testable predictions made by various models of how pivotal suppliers affect supply function bidding. In contrast, our experiment focuses on a framework in which market power is driven by a small number of firms, increasing marginal costs, and private information about costs. Outside the laboratory, Hortaçsu and Puller (2008) empirically evaluate strategic bidding behaviour in multi-unit auctions using data from the Texas electricity market. These authors find evidence that large firms bid according to the theoretical benchmark while smaller firms deviate significantly from that benchmark.

To the best of our knowledge, ours is the first laboratory experiment to test the relationship between informational frictions and market power in the context of supply function competition. Because of similarities in the information environment, our results are related to findings in the literature on the winner’s curse in single unit auctions where a savvy bidder avoids bidding aggressively because “winning” conveys the news that her signal was the highest in the market. A prevalent, consistent, and robust phenomenon in single-unit auctions featuring common (or correlated) values (Kagel and Levin 1986; Goeree and Offerman 2003; Kagel and Levin (forthcoming)) and where bidders ignore the adverse selection problem. However, the analogy between the winner’s curse with competition in supply functions and with single-unit auctions applies with respect to adverse selection but not necessarily with respect to market power. In essence, our results are more closely related to the *generalised winner’s curse* (Ausubel et al. 2014) which reflects that “winning” a larger quantity is worse news than “winning” a smaller quantity because the former implies a higher expected cost for the bidder (where bidders are sellers). In our environment a seller that faces a high price should think that it is likely that costs of her rivals are high and this is news that her own costs are also high because of the positive correlation. The result is that the seller should moderate her offer and this induces the supply function to be steeper. Therefore, rational bidders refrain from competing too aggressively. We find evidence of the generalised winner’s curse in a multi-unit, divisible-good auction with interdependent values.

The market structure in our setup is reminiscent of multi-unit uniform price auctions, for which there is evidence of demand reduction—in demand auctions characterised by independent private values and an indivisible good—both experimentally (Kagel and Levin 2001) and in the field (List and Lucking-Reiley 2000; Levin 2005; Engelbrecht-Wiggans et al. 2005; Engelbrecht-Wiggans et al. 2006). Unlike this literature, our paper addresses a uniform-price auction with interdependent values and a divisible good. The experiment we conduct is also related to that of Sade et al. (2006), who test the theoretical predictions of a divisible-good, multi-unit auction

model under different auction designs; they report some inconsistencies between the theoretical equilibrium strategies and actual experimental behaviour.

There is also an experimental literature on “correlation neglect” in various strategic contexts. Those contexts include bilateral negotiations (Samuelson and Bazerman 1985), trade with adverse selection (Holt and Sherman 1994), social learning (Weizsacker 2010), auctions with toeholds (Georganas and Nagel 2011), voting (Esponda and Vespa 2014; Levy and Razin 2015), and belief formation (Enke and Zimmermann 2013; Koch and Penczynski 2015). Our results are consistent also with these experimental findings in that a substantial proportion of our subjects ignore correlation among costs and hence its adverse effects.

The detailed analysis of our results echoes the analysis of other strategic games with private information; examples include Carrillo and Palfrey (2011) as well as Brocas et al. (2014), who conduct a similar cluster analysis. Both of those papers report that (a) a large proportion of subjects behave just as in the equilibrium where subjects play simple but strategic private-information games yet (b) this proportion declines markedly with increasing strategic complexity of the game. Charness and Levin (2009) attribute the origin of the winner’s curse to bounded rationality since individuals have difficulties thinking contingently about future events. This behavioural bias is also likely to apply to the subjects of our experiment since the equilibrium of the positively correlated costs treatment requires that a bidder thinks contingently about the relationship between the market price and unit costs. Furthermore, the evolution of choices across rounds that we observe is related to the work of Huck et al. (1999) and Bigoni and Fort (2013), who analyse learning in a Cournot setting. Much as in those papers, our own results indicate that learning is a composite process that involves elements of both adaptive and reinforcement learning—despite the greater complexity of our environment (in which subjects set supply functions with incomplete information rather than choosing quantities with full information).

The complexity of the our baseline model makes the comparison of our results to extant behavioural models somewhat difficult. Eyster and Rabin (2005) propose the *cursed equilibrium* concept, whereby players incorrectly assess the relationship between rivals’ strategies and their own private information. In our setting, then, a player in the positively correlated costs treatment would be “fully cursed” if she ignored the information that price conveys about costs. Except for the two extreme cases in which all players are fully cursed or all players are fully rational, there are several analytical difficulties associated with computing the cursed equilibrium in our complex setting; we therefore refrain from further relating our results to that notion. An alternative equilibrium concept that could explain our choices would be the *quantal response equilibrium* (QRE), which assumes that the probability of any strategy is increasing in the strategy’s expected payoff and rationalises the idea that players make mistakes. However, the calculation of the QRE for our baseline model is beyond the scope of this paper. Finally, a

nonequilibrium potential explanation of our results could be given by the level- k model of strategic thinking (Nagel 1995; Crawford and Iriberri 2007). We do not formally apply the level- k model to describe our experimental choices for the following reasons. First, if we define level-0 as the equilibrium behaviour of the uncorrelated costs treatment then we cannot explain the heterogeneity that we observe in the uncorrelated costs treatment. For other plausible definitions of level-0, the theoretical level- k predictions do not correspond to observed peaks that we observe in the distribution of choices. Second, the various levels are not sufficiently differentiated to enable clear identification of a subjects' degree of strategic thinking. Given these limitations, we focus on analysing the theoretical and empirical distribution of best replies, an approach that yields a categorisation of subjects into clusters: those who are naïve (subjects that fall prey to the generalised winner's curse) and sophisticated (responding best to the average strategy of rivals).

3 Theoretical Background

We use the framework of Vives (2011) to guide our experimental design. There are a finite number n of sellers who compete simultaneously in a uniform price auction, and each seller submits a supply function. Seller i 's profits can be written as

$$\pi_i = (p - \theta_i)x_i - \frac{\lambda}{2}x_i^2, \quad (1)$$

where x_i are the units sold, θ_i denotes a random cost parameter, p is a uniform market price, and $\lambda > 0$ represents a parameter that measures the level of transaction costs. The market-clearing condition allows us to find the uniform market price p . The (random) cost parameter θ_i is normally distributed as $\theta_i \sim N(\bar{\theta}, \sigma_\theta^2)$. The demand is inelastic and equal to q .

The information structure is as follows. A seller does not know the value of the cost shock θ_i before setting her supply schedule, and she receives a signal $s_i = \theta_i + \varepsilon_i$ for which the error term is distributed as $\varepsilon_i \sim N(0, \sigma_\varepsilon^2)$. Sellers' random cost parameters may be correlated, with $\text{corr}(\theta_i, \theta_j) = \rho$ for $i \neq j$. When $\rho = 1$ the model is equivalent to a common costs model, when $\rho = 0$ to an uncorrelated costs model, and when $0 < \rho < 1$ to a correlated costs model. Error terms are uncorrelated either among themselves or with the random cost shocks. In our experiment, the *treatment variable* is the correlation ρ among costs.

Since the payoff function is quadratic and since the information structure is normally distributed, we focus on linear supply schedules. Linear supply functions are a reasonable approximation of the types of supply functions submitted by bidders in real markets.⁵ Given the signal

⁵See, for example, Baldick et al. (2004).

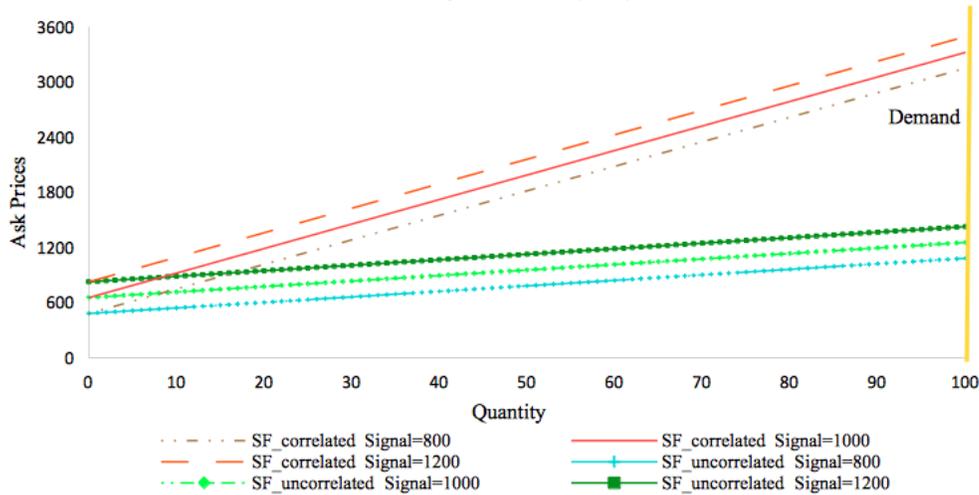
received, a strategy for seller i is to submit a price-contingent schedule, $X(s_i, p)$, of the form

$$X(s_i, p) = b - as_i + cp. \quad (2)$$

Thus the seller’s supply function is determined by the three coefficients (a, b, c) . We interpret these coefficients as follows: a is a bidder’s response to the private signal; b is the fixed part of the supply function’s intercept $f = b - as_i$; and c is the supply function’s slope.

Vives (2011) finds a unique supply function equilibrium and describes how the equilibrium parameters (a, b, c) depend on the information structure $(\bar{\theta}, \sigma_{\theta}^2, \sigma_{\epsilon}^2, \rho)$ and on the market structure (n, q, λ) ; see Appendix A for the formulae that characterise the equilibrium supply function and outcomes. Figure 1 illustrates the comparative statics of the unique, symmetric, and linear Bayesian Nash supply function equilibrium.

Figure 1: *Bayesian Nash equilibrium supply function predictions for the uncorrelated and positively correlated costs treatments when all agents are fully rational.*



Note. The theoretical predictions for each treatment are computed for three different signal realisations: a high signal with value 1,200, a signal equal to the ex ante value of θ_i with value 1,000, and a low signal with value 800.

The equilibrium prediction for a fully rational bidder’s behaviour is described as follows. When costs are positively correlated, a high price conveys that the bidder’s costs are high; in equilibrium, then, bidders submit steeper schedules (than when costs are uncorrelated) to protect themselves from adverse selection. Also, for the same signal realisation, the equilibrium supply function’s intercept is lower when costs are positively correlated than when they are uncorrelated; hence the model also predicts that equilibrium market outcomes are less competitive when costs are positively correlated than when they are uncorrelated because both the expected market price and profits are higher in the former case than in the latter. Note that outcomes in both treatments lie between the Cournot and competitive benchmarks, since each treatment features supply functions with positive and finite slopes. In terms of welfare, we observe that

the equilibrium allocation is inefficient due to distributive inefficiency. Because demand is inelastic, there is no aggregate inefficiency at the equilibrium allocation. At this allocation, sellers supply quantities that exhibit too little dispersion vis-à-vis the efficient benchmark, which is reflected in the ex-ante expected deadweight loss at the equilibrium allocation (it is defined as the difference between expected total surplus at the efficient and equilibrium allocations).

The comparative statics of the unique Bayesian Nash equilibrium—which assumes that all sellers are fully rational and ex ante symmetric—allow us to derive the testable predictions encapsulated by the six hypotheses listed here as (A)–(F). These hypotheses focus on the model’s general predictions and on the comparative statics with respect to the correlation among costs, since ρ is our treatment variable.

- (A) *In each treatment, the supply function slope is positive and unrelated to a bidder’s signal realisation.*
- (B) *In each treatment, the supply function intercept is nonzero and increasing in a bidder’s signal realisation.*
- (C) *The supply function is steeper in the positively correlated costs treatment than in the uncorrelated costs treatment.*
- (D) *For a given signal realisation, the supply function’s intercept is lower in the positive correlated costs treatment than in the uncorrelated costs treatment; therefore, the expected supply function intercept is lower in the positive correlated costs treatment than in the uncorrelated costs treatment.*
- (E) *The expected market price and profits are larger in the positively correlated costs treatment than in the uncorrelated costs treatment.*
- (F) *The expected deadweight loss is larger in the uncorrelated costs treatment than in the positively correlated costs treatment.*⁶

If subjects ignore the correlation among costs and thus do not understand that winning a larger quantity is worse news (when costs are positively correlated) than winning a smaller quantity, then those subjects fall prey to the winner’s curse in the context of a multi-unit auction with interdependent values. (This phenomenon was termed the generalised winner’s curse by Ausubel et al. 2014.) Therefore, the benchmark for subjects who fall prey to the generalised winner’s curse is the equilibrium of the uncorrelated costs treatment. If all subjects were to fall prey to that curse then we should expect behaviour and market outcomes in both of our

⁶This claim follows from the chosen constellation of experimental parameters (see Table 2 in Section 4), although the theoretical prediction asserts that the expected deadweight loss can either increase or decrease with the correlation among costs.

experimental treatments to be indistinguishable, in which case hypotheses (C)–(F) might not be hold.⁷ The hypotheses just listed pertain to point estimates and expected behaviour, yet experimental variation in behaviour and outcomes may also be of interest. We explore these issues in Section 5.

4 Experimental Design

Sessions were conducted in the LINEEX laboratory of the University of Valencia. The participants were undergraduate students in the fields of economics, finance, business, engineering, and natural sciences. All sessions were computerised.⁸ Instructions were read aloud, questions were answered in private, and—throughout the sessions—no communication was allowed between subjects. Instructions explained all details of the market rules, distributional assumptions on the random costs, the nature of signals, and the correlation among costs (the meaning of correlation was explained both with a definition and graphically). Before starting the experiment, we tested participants’ understanding. See Appendix C for the instructions and the first part of Appendix D for the comprehension test.

We ran the experiment with 144 participants, half of whom participated in the uncorrelated costs treatment and half in the positively correlated costs treatment.⁹ Each treatment had 6 independent groups of 12 members each, which consisted of 4 markets with 3 sellers in each market. We chose a market size of 3 because this is the minimum market size that does not lead to collusion in other, similar environments—for example, a Bertrand game (Dufwenberg and Gneezy 2000) and a Cournot market (Huck et al. 2004). Subjects competed for 2 trial rounds followed by 25 live rounds (since it is an established fact that equilibrium does not appear instantaneously in experimental games). In all of these rounds, in order to keep the spirit of the theoretical model’s one-shot nature, we employed random matching between rounds.¹⁰ Thus, the composition of each of the four markets varied each round within a group. Table 1 summarises the structure of our experimental design.

⁷Because costs and signals follow a normal distribution, a subject could fall prey to the so-called news curse whereby she ignores prior information and takes a signal at face value (Goeree and Offerman 2003). We can show that, in our model, a bidder who falls prey to the news curse sets the same supply function slope as one who falls prey to the generalised winner’s curse. As a result, these two curses are not easily distinguished.

⁸For this purpose we used the *z-tree* software (Fischbacher 1999).

⁹Before running the experiment, we conducted two pilot versions (one for each treatment).

¹⁰Although this random matching within a group results in fewer independent observations than if we had chosen fixed markets, we considered it more important that our treatments accord with the theoretical model.

Table 1: *Experimental design.*

Treatment		Independent Groups	Markets/ Group	Subjects/ Market	Number of Subjects	Number of Rounds
Uncorrelated Costs ($\rho=0$)	Control T0	6	4	3	72	25
Positively Correlated Costs ($\rho=0.6$)	Treatment T1	6	4	3	72	25

In the second part of Appendix D we reproduce the screenshots used for running the experiment. In each round, all subjects received a private signal and were subsequently asked to choose two *ask prices*: one for the first unit offered and one for the second. We then used these two ask prices to construct a linear supply schedule, which was shown in the form of a graph on each subject’s screen. The participant could then revise the ask prices several times until she was satisfied with her decision. The buyer was simulated.

Once all supply schedules had been submitted, each bidder received feedback on the uniform market price, her own performance (with regard to revenues, production costs, transaction costs, units sold, and profits), the performance of the other two market participants (units sold, profits, and supply functions), and the values of the random variables drawn (her own cost and the costs of the other two participants in the same market).¹¹ Participants were allowed to consult the history of their own performance. Other experiments have shown that feedback affects behaviour in the laboratory. In a Cournot game, for example, Offerman et al. (2002) report that different feedback rules can result in outcomes that range from competitive to collusive. Given the complexity of our experiment, we maximised the feedback given after each round in order to maximise the potential learning of participants. After each participant had checked her feedback, a new round of the game would start. Note that, in each market and for each round, we generated three random unit costs from a multivariate normal distribution. Also, in each round and for each participant, the unit costs and signals were independent draws from previous and future rounds.

Conducting the experiment required us to specify numerical values for the theoretical model’s parameters; see Table 2. For this purpose we used three criteria: (i) the existence of a unique equilibrium; (ii) sufficiently differentiated behaviour and outcomes between the two treatments; and (iii) reduction of computational demands placed on participants. It is important to bear in mind that $\rho = 0$ for the uncorrelated costs (control) treatment whereas $\rho = 0.6$ for the positively correlated costs treatment.¹² Refer to Appendix B for the equilibrium supply function and

¹¹Subjects did *not* receive feedback on the signal received because that would not be expected to occur in reality. That is: after trading, firms may observe the actual costs of competitors but are unlikely to observe the private signals competitors had received at the time of their decisions.

¹²We would have preferred to set a higher correlation among unit costs so that our predictions would be maximally differentiated. Yet inelastic demand reduces the range for which a unique equilibrium exists, and

outcomes (based on Table 2’s experimental parameters) and for a statistical description of the distribution of random costs and errors used in the experiment.

Table 2: *Experimental parameters.*

	Symbol	Value/s
Number of Sellers	n	3
Inelastic demand	q	100
Transaction cost parameter	λ	3
Mean of random cost	Θ	1,000
Variance of random cost	σ_{Θ}^2	10,000
Mean of signal's error	$\bar{\varepsilon}$	0
Variance of signal's error	σ_{ε}^2	3,600
Correlation	ρ	0 and 0.6

We imposed certain market rules, which were inspired by the theoretical model and facilitated implementation of the experiment. First, we asked each seller to offer all 100 units for sale. Second, we asked sellers to construct a nondecreasing and linear supply function. Third, ask prices had to be nonnegative. Fourth, we told bidders that the simulated buyer would not purchase any unit at a price higher than 3,600; this price cap was imposed in order to limit the potential gains of sellers in the experimental sessions. Although the price cap was not part of the theoretical model, we chose a value high enough to preclude distortion of equilibrium behaviour. The only difference between treatments was the correlation among costs and hence the distribution of random costs and signals.

At the end of the experiment, participants completed a questionnaire (see Appendix E) that requested personal information and asked questions about the subject’s reflections after playing the game. Once the questionnaire was completed, each participant was paid in private.

As for incentives, each participant started with 50,000 experimental points.¹³ During the experiment, subjects won or lost points. At the end of the experiment, points were exchanged for euros at the rate of 10,000 experimental points per euro. In addition, each subject received a 10 euros show-up fee. The payments ultimately made to subjects ranged from 10 to 27.8 euros and averaged 20.8 euros. Each session lasted between two and three hours.

5 Experimental Results

We present our results in three sections. In Section 5.1 we provide an analysis of experimental behaviour (supply functions) and in Section 5.2 we analyse experimental outcomes (market price and profits) and efficiency of allocations. Section 5.3 addresses trends of behaviour and

$\rho = 0.6$ was the highest correlation that satisfied our implementation criteria. Inelastic demand was stipulated in order to simplify the participants’ computations.

¹³These points were equivalent to 5 euros.

outcomes across rounds. Throughout the section, we evaluate Hypotheses (A)–(F) formulated in Section 3. Appendix G conducts a robustness test of our main experimental results using a panel data approach.

We shall discuss the experimental results in terms of the inverse supply function, since it corresponds to how participants made their decisions. From any participant’s two-dimensional decision, $(AskPrice1, AskPrice2)$, we can infer the slope and intercept of each participant’s inverse supply function, $p = \hat{f} + \hat{c}X(s_i, p)$. The inverse supply function slope is \hat{c} , where $\hat{c} = AskPrice2 - AskPrice1$ and the intercept is \hat{f} , defined as $\hat{f} = AskPrice1 - \hat{c}$. The coefficients of the inverse supply function are related to the coefficients of equation (2)’s supply function as follows: $\hat{b} = \frac{-b}{c}$; $\hat{a} = \frac{a}{c}$; $\hat{c} = \frac{1}{c}$ for $c \neq 0$; the inverse supply function’s intercept is $\hat{f} = \hat{b} + \hat{a}s_i$. We will omit the modifier “inverse” and refer simply to “the supply function”. We shall use *InterceptPQ* as the empirical counterpart of \hat{f} and *SlopePQ* as the empirical counterpart of \hat{c} . Our graphs plot the supply function in the usual $(Quantity, AskPrice)$ space.¹⁴

5.1 Analysis of experimental behaviour: supply functions

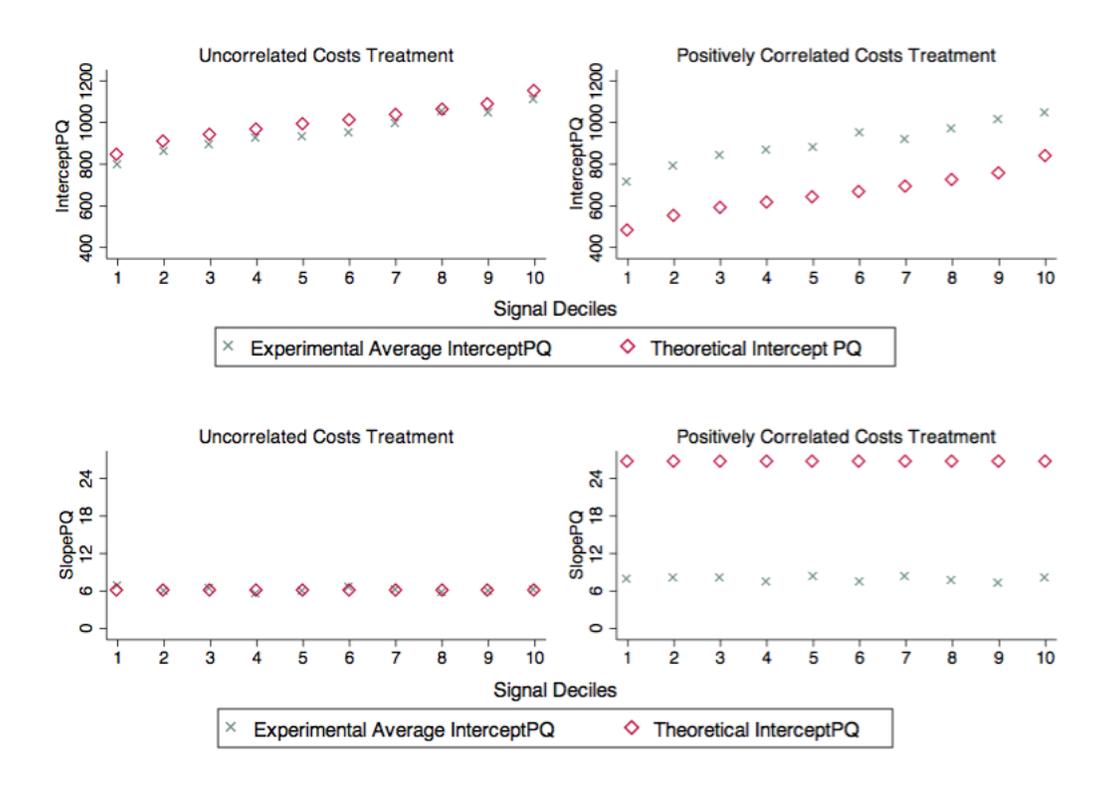
We first present evidence of the most general testable predictions of the theoretical model that are common in both treatments. In Section 3 we saw the theoretical framework predicting that, in both treatments, the supply function *slope* should be independent of the signal received whereas the supply function *intercept* should increase with the signal realisation. The latter prediction reflects that a higher signal implies a higher average intercept of the marginal cost and so the bidder should set a higher ask price for the first unit offered, leading to a higher supply function intercept.

In each treatment, we find that the average supply function intercept increases with the signal in each decile. In contrast, the supply function slope remains approximately constant in each signal decile in either treatment. Furthermore, a regression of the supply function intercept on signal yields a coefficient of 0.856 with $p = 0.000$ in the uncorrelated costs treatment and a coefficient of 0.926 with $p = 0.000$ in the positively correlated costs treatment; while a regression of the supply function slope on signal yields an insignificant coefficient (with $p = 0.703$ in the uncorrelated costs treatment and with $p = 0.655$ in the positively correlated costs treatment).¹⁵ Figure 2 illustrates the average supply function slope and intercept for each signal decile using all the choices from both treatments.

¹⁴Note that a steep supply function in the $(Quantity, AskPrice)$ space has a high \hat{c} and a low c .

¹⁵The unit of observation for the regressions reported is the group *across rounds*. In each treatment there are 150 observations: 6 groups over the course of 25 rounds.

Figure 2: *Experimental supply function intercept and slope in each signal decile and the corresponding theoretical prediction in each treatment.*



Note. SF refers to supply function. The categories in the horizontal axis represent signal deciles.

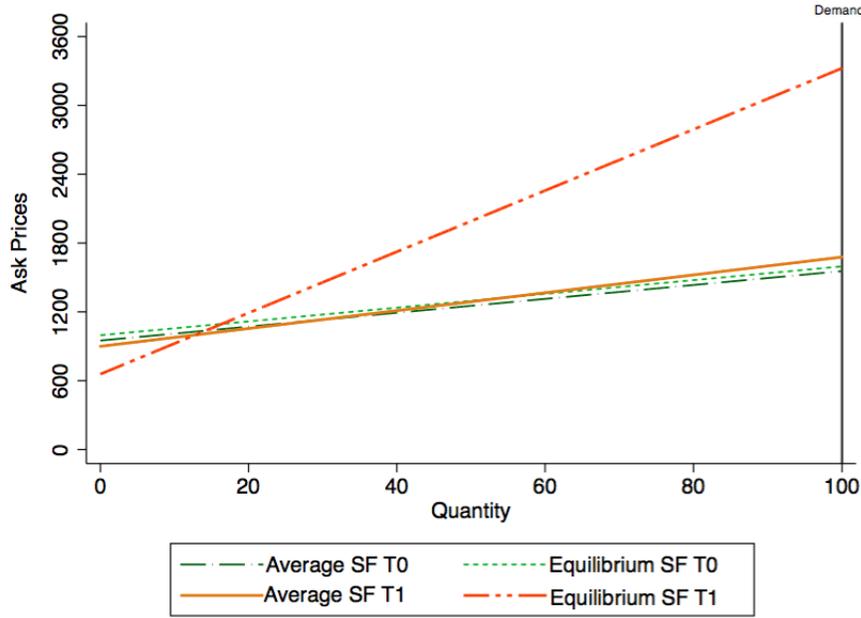
This evidence suggests that bidding behaviour is consistent, qualitatively, with the most general testable predictions of the theoretical model that are common to both treatments. Our first result pertains to Hypotheses (A) and (B).

Result 1 (Predictions common to both treatments). As posited by the theoretical model, in each treatment: **(i)** the average supply function’s intercept is increasing in a bidder’s signal realisation; and **(ii)** the average supply function’s slope is positive and unrelated to a bidder’s signal realisation.

We interpret this result to mean that our data supports Hypothesis (A) and also Hypothesis (B).

We next present bidders’ behaviour in each treatment, when decisions are aggregated across all rounds. For each treatment, Figure 3 plots the observed average supply function as well as the corresponding equilibrium prediction; Table 3 gives the slope and intercept of the average experimental supply function. We employ both parametric and nonparametric tests. Here our unit of analysis is the group average, which aggregates individual choices or outcomes within the group and across rounds.

Figure 3: *Average experimental and equilibrium supply functions in each treatment.*



Note. SF refers to supply function; T0 to uncorrelated costs treatment; T1 to the positively correlated costs treatment.

Table 3: *Average behaviour, and their corresponding theoretical predictions, by treatment.*

Behaviour: Supply Function	Number of observations/Treatment	Uncorrelated Costs Treatment		Positively Correlated Costs Treatment	
		Mean (s.d.)	Theoretical Prediction	Mean (s.d.)	Theoretical Prediction
Intercept PQ	1,800	950.11 (207.73)	1,000	899.25 (291.11)	655.32
Slope PQ	1,800	6.05 (5.32)	6.00	7.79 (6.98)	26.68

Note. Theoretical predictions refer to the equilibrium prediction for the supply function slope and expected intercept. The standard deviation (s.d.) is given in parentheses below the reported average. For all variables, reported standard deviations are at the individual level.

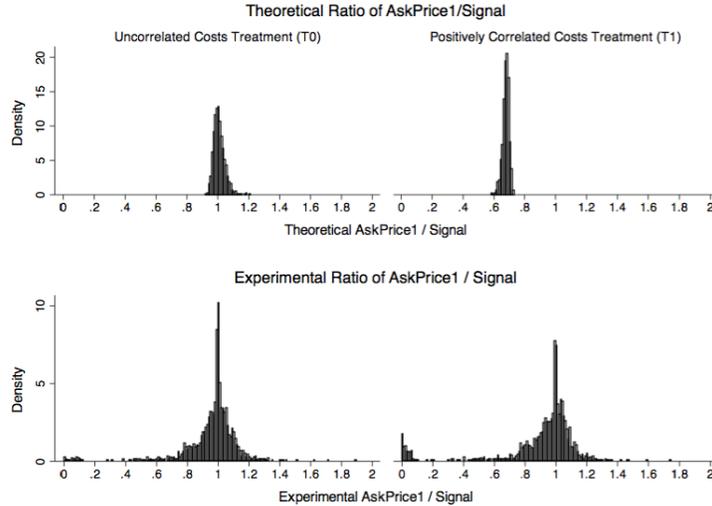
Figure 3 and Table 3 show that, in the uncorrelated costs treatment, the average supply function is close to the theoretical prediction. In fact, we cannot reject the hypothesis that the supply function slope is the *same* as the theoretical prediction (two-sided Wilcoxon signed-rank test: $n_1 = 6$, $n_2 = 6$, $p = 0.600$; two-sided t -test: $n_1 = 6$, $n_2 = 6$, $p = 0.914$). The average intercept is lower than expected (950.11 vs. 1,000). The evaluation test of the hypothesis that the average intercept is equal to its expected value yields mixed results at the 5% significance level (two-sided Wilcoxon signed-rank test: $n_1 = 6$, $n_2 = 6$, $p = 0.046$; two-sided t -test: $n_1 = 6$, $n_2 = 6$, $p = 0.054$). Finding that behaviour in the uncorrelated costs treatment is, on average, close to the theoretical prediction is important because it allows us to use the control treatment

as the benchmark for our analysis.

Comparing the average supply function of the two treatments reveals that, as predicted by the Bayesian Nash equilibrium, the average supply function in the positively correlated costs treatment has a higher slope (7.79 vs. 6.05) and lower intercept (899.25 vs. 950.11) than in the uncorrelated costs treatment. This difference between the average supply functions in the two treatments accords with the theoretical model qualitatively but not quantitatively, as the differences between treatments are substantially smaller than predicted.¹⁶ In fact, this difference is not statistically significant in terms of either the supply function intercept (one-sided Mann–Whitney U-test: $n_1 = 6$, $n_2 = 6$, $p = 0.556$; one-sided t -test with unequal variance: $n_1 = 6$, $n_2 = 6$, $p = 0.198$) or the slope (one-sided Mann–Whitney U-test: $n_1 = 6$, $n_2 = 6$, $p = 0.667$; one-sided t -test with unequal variance: $n_1 = 6$, $n_2 = 6$, $p = 0.120$).

Next, we want to explore whether the behavioural response to the private signal is quantitatively consistent with the theoretical model (the analysis reported in Figure 2 has only explored the direction and not the magnitude of the response to the signal). We further disaggregate choices conditional on the private signal and analyse individual ask prices for the first unit offered—which is equivalent, in essence to the supply function intercept—in each treatment. Figure 4 displays both the equilibrium and experimental ratio $AskPrice1 / Signal$ in each treatment using disaggregated individual choices during all rounds of the experiment.

Figure 4: *Histogram of the theoretical and experimental ratio of AskPrice1 to Signal received for individual choices in all rounds of the experiment.*



Note. This histogram omits two observations for which the experimental $AskPrice1 / Signal$ ratio was greater than 2.

¹⁶The tests reported are as follows. The null hypothesis for all of them is $H_0 : \mu_0 - \mu_1 = 0$, where μ_1 (resp. μ_0) is the mean for the positively correlated (resp. uncorrelated) costs treatment. The alternative hypothesis for the supply function slope ($SlopePQ$) is $H_1 : \mu_0 - \mu_1 < 0$. For the supply function intercept ($InterceptPQ$) the alternative hypothesis is $H_1 : \mu_0 - \mu_1 > 0$. The reported p -values are for the corresponding one-sided tests.

In the uncorrelated costs treatment, the Bayesian Nash equilibrium predicts that the ratio of *AskPrice1* to *Signal* is between 0.92 and 1.2, with both the mean and the median equal to 1.01. In other words, in this treatment the Bayesian Nash equilibrium predicts that *AskPrice1* will be (on average) equal to the signal received.¹⁷ Experimental choices in the uncorrelated costs treatment, as illustrated in the lower left graph of Figure 4, reveal that the mean of the *AskPrice1/Signal* ratio is similar to the mean (0.96) of its theoretical counterpart—but with a larger standard deviation owing to the heterogeneity in individual choices.

In the positively correlated costs treatment, the theoretically predicted *AskPrice1/Signal* ratio is between 0.58 and 0.73, with both the mean and the median equal to 0.68; thus the equilibrium predicts that *AskPrice1* will be lower than the signal received. The experimental distribution has a mean of 0.91 and median of 0.99, which means that: (i) for a given signal, *AskPrice1* is (on average) larger than predicted by the Bayesian Nash equilibrium in this treatment; and (ii) subjects in the positively correlated costs treatment are strongly guided, when choosing *AskPrice1*, by the signal received (lower right graph of Figure 4) and so may choose a number close to the signal received (since it acts as a focal point).

The foregoing analysis illustrates systematic divergences between behaviour—that is, how subjects respond (through the supply function intercept) to the private signal—and Bayesian Nash equilibrium prediction in the positively correlated costs treatment. We interpret this result as offering additional evidence that subjects, when setting the supply function intercept, do not account for the effects of correlation among costs and consequently respond too strongly to the signal received.

The next result concerns the evaluation of Hypotheses (C) and (D).

Result 2 (Differences between experimental behaviour and theoretical predictions; differences in the average supply function between treatments). **(i)** The average supply function in the uncorrelated costs treatment is close to the corresponding equilibrium prediction; however, we reject the hypothesis that the average supply function in the positively correlated costs treatment is the same as the corresponding equilibrium prediction. **(ii)** Differences in the average supply function between treatments are not statistically significant.

Thus we do not find empirical support for Hypothesis (C) or Hypothesis (D). In addition, Result 2 suggests that the generalised winner’s curse is a prevalent phenomenon in the positively correlated costs treatment. The explanation is that: (a) behaviour in the uncorrelated costs treatment is consistent with the theoretical prediction and so can serve as a benchmark for subjects that fall prey to the generalised winner’s curse; and (b) *average* behaviour in the

¹⁷The variables *AskPrice1* and *InterceptPQ* are related as follows: $AskPrice1 = InterceptPQ + SlopePQ$. In each round, each subject was asked to specify (*AskPrice1*, *AskPrice2*); therefore, *Ask Price1* more clearly reflects how participants made their decisions.

positively correlated costs treatment is *not* different than average behaviour in the uncorrelated costs treatment.

We now address the question posed at the end of Section 3 about the variation and distribution of behaviour. In Table 4 we present an analysis of variance for the different levels of aggregation (subject level; group level; treatment level).

Table 4: *Analysis of variance of the experimental supply function slope and intercept.*

Variable	Mean	s.d. Overall Subject level ¹	s.d. Between Subjects ¹	s.d. Within Subjects (period)	s.d. Overall Group level	s.d. Between Groups	s.d. Within Groups (period)	s.d. Overall Treatment level (period)
Observations / Treatment		1800	72	25	150	6	25	25
Uncorrelated Costs Treatment								
<i>Behaviour: Supply Function</i>								
Intercept PQ	950.11	207.73	133.31	160.05	65.04	48.88	47.19	22.07
Slope PQ	6.05	5.32	3.71	3.83	1.7	1.17	1.32	0.83
Positively Correlated Costs Treatment								
<i>Behaviour: Supply Function</i>								
Intercept PQ	899.25	291.11	233.65	175.73	129.56	128.25	54.66	21.11
Slope PQ	7.79	6.98	5.35	4.53	3.18	3.06	1.50	0.71

Note. s.d. refers to standard deviation. The explanation of the variables reported in the table is as follows: s.d. Overall subject level: the variation of individual choice across all rounds; s.d. Between subjects: the variation of the individual choice averaged across all rounds; s.d. Within Subjects: the variation within subject over time; s.d. Overall group level: the variation the average choice of the group across all rounds; s.d. Between groups: the variation of the average choice of the group averaged across all rounds; s.d. Within groups: the variation within groups over time; and s.d. Overall Treatment level: the variation of the average choice of the treatment across all rounds.

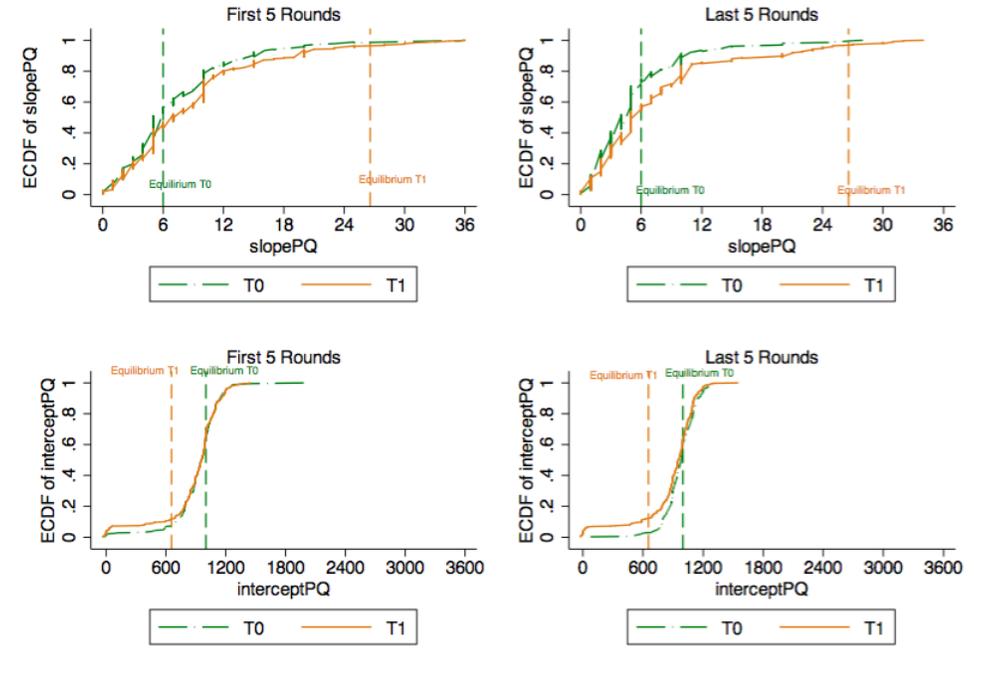
First, the table shows that, at the different levels of aggregation, there is a substantial amount of heterogeneity in behaviour in both treatments. Second, we can observe that the variances of both the supply function slope and intercept are relatively larger in the positively correlated than in the uncorrelated costs treatment, a difference that is driven mainly by the heterogeneity across subjects and groups and not by heterogeneity across rounds. So even though there are no differences in average behaviour between treatments, there are differences in the *variance* of behaviour between treatments. We provide a detailed explanation in Section 6.2. We can show that the variance of the supply function slope is larger in the positively correlated than in the uncorrelated treatment at the different levels of aggregation: between groups (variance ratio test, one-sided, $n_1 = 6$, $n_2 = 6$, $p = 0.027$) and between subjects (variance ratio test, one-sided, $n_1 = 72$, $n_2 = 72$, $p = 0.0012$). The same finding is also true for the supply function intercept: between groups (variance ratio test, one-sided $n_1 = 6$, $n_2 = 6$, $p = 0.027$) and between subjects (variance ratio test, one-sided, $n_1 = 72$, $n_2 = 72$, $p = 0.000$). The differences across periods at different levels of aggregation are not substantial.

The analysis just presented has established that—even when differences in the average supply

function of each treatment are insignificant—the variance in the supply function slope and intercept is greater in the positively correlated than in the uncorrelated costs treatment.

Figure 5 plots the empirical cumulative distribution function (ECDF) for the supply function slope and intercept in each treatment, where the source data for these graphs are the disaggregated individual choices made by subjects during the first and last five rounds of bidding. These two time periods are important because they allow us to summarise behaviour at the beginning of the experiment, when subjects have no experience, and at the end of the experiment, when they have bid for 20 rounds.

Figure 5: *Empirical cumulative distribution function (ECDF) for the supply function slope and intercept in the first five and last five rounds of bidding.*



There is a significant difference between treatments as regards the ECDFs of supply function slopes in both the first and last five rounds of bidding. The ECDF of supply function slopes in the positively correlated costs treatment first order stochastically dominates the ECDF of slopes in the uncorrelated costs treatment both in the first and last five rounds (in a stronger way in the latter). In fact, we can reject the hypothesis that the distribution of slopes is the same in the two treatments during the first five rounds of bidding (Kolmogorov–Smirnov equality of distributions test: $n_1 = 360$, $n_2 = 360$, $p = 0.023$) and also during the last five rounds (Kolmogorov–Smirnov equality of distributions test: $n_1 = 360$, $n_2 = 360$, $p = 0.000$).¹⁸

¹⁸Repeating the same Kolmogorov–Smirnov equality of distributions test for the intermediate periods, in groups of five rounds, we find that the distribution of slopes in the two treatments is significantly different in all intermediate time periods considered; this result is significant at the 5% level.

With regard to the supply function intercepts, we can see that the differences in ECDFs between treatments are quite small in the first five rounds of bidding, for which they are not statistically significant (Kolmogorov–Smirnov equality of distributions test: $n_1 = 360$, $n_2 = 360$, $p = 0.689$). In the last five rounds of bidding, however, differences in the distribution of supply function intercepts are more substantial—albeit mainly driven by a few subjects, in the positively correlated costs treatment, who bid a supply function with a very low intercept. Even so, we can reject the hypothesis that the ECDFs of supply function intercepts between treatments are the same in the last five rounds of bidding (Kolmogorov–Smirnov equality of distributions test: $n_1 = 360$, $n_2 = 360$, $p = 0.009$).¹⁹

We now address the question posed at the end of Section 3 about the variation and distribution of behaviour.

Result 3 (Difference in the distributions of the supply function slope and intercept between treatments). **(i)** There is a substantial amount of heterogeneity in behaviour in both treatments. Variances in the supply function slope and intercept are relatively greater in the positively correlated than in the uncorrelated costs treatment because of the heterogeneity in subjects’ and groups’ behaviour, not because of heterogeneity across rounds. **(ii)** Using individual choices, we reject the hypothesis that the distribution of supply function *slopes* is the same in both treatments for the first five and last five rounds of bidding. The cumulative distribution function of supply function slopes in the positively correlated costs treatment first order stochastically dominates the cumulative distribution function of slopes in the uncorrelated costs treatment both in the first and last five rounds. We cannot reject the hypothesis that the distribution of supply function *intercepts* is the same in both treatments for the first five rounds of bidding, but we do reject the hypothesis that the distribution of supply function intercepts is the same in both treatments for the last five rounds of bidding.

Result 3 indicates that there are differences in the distribution of supply functions between treatments. The differences in behaviour between treatments are consistent with the direction predicted by the theoretical model. We shall examine this issue more closely in Section 6.2.

In sum, the results reported in this section have shown that *average* differences in behaviour between treatments are not statistically significant. However, there are significant differences in the *distribution* of supply functions between treatments, suggesting that that some subjects in the positively correlated costs treatment—because they ignore the correlation among costs and

¹⁹Repeating the same Kolmogorov–Smirnov equality of distributions test for the intermediate periods, in groups of five rounds, we find that the distribution of intercepts in the two treatments is not significantly different in the second and third groups of five rounds (Rounds [6, 10] and Rounds [11, 15]) but that the distribution becomes significantly different after the 15th round.

its adverse effects—fall prey to the generalised winner’s curse while others do not.²⁰ The more detailed examination in Section 6 seeks to determine whether these results are driven by a small proportion of subjects or whether the generalised winner’s curse is prevalent in the positively correlated costs treatment.

5.2 Analysis of experimental outcomes (market price and profits) and efficiency of allocations

Table 5 reports experimental outcomes and the corresponding predictions in terms of the market price, profits, and efficiency levels of the allocations (as captured by deadweight losses) together with the corresponding theoretical predictions.²¹ Our unit of analysis is the group average, which aggregates outcomes within the group and across rounds.²²

Table 5: *Average outcomes, and their corresponding theoretical predictions, by treatment.*

Variable	Number of observations /Treatment	Uncorrelated Costs Treatment		Positively Correlated Costs Treatment	
		Mean (s.d.)	<i>Theoretical Prediction</i>	Mean (s.d.)	<i>Theoretical Prediction</i>
<i>Outcomes</i>					
Market Price	600	1,110.38 (117.36)	1,200.00	1,123.73 (130.27)	1,544.68
Profits	1,800	2,255.35 (7,191.77)	5,612.75	1,937.81 (5,620.92)	16,567.37
<i>Efficiency</i>					
Deadweight Loss	600	2,222.51 (3,142.66)	612.75	2,011.01 (2,290.29)	467.73

Note. The theoretical predictions for outcomes and efficiency refer to ex ante expected market prices, ex ante expected profits, and ex ante expected deadweight loss at the equilibrium allocation.

The standard deviation (s.d.) is given in parentheses below the reported average. For profits reported standard deviations are at the individual level; for market price and deadweight loss, standard deviations are at the market level.

Turning now to average market prices, we see that the difference in these prices between treatments is not statistically significant (one-sided Mann–Whitney U-test: $n_1 = 6$, $n_2 = 6$, $p = 0.556$; one-sided t -test with unequal variance: $n_1 = 6$, $n_2 = 6$, $p = 0.210$). In both treatments, average market prices are lower than theoretically predicted: 92.7% and 72.6% of the theoretically predicted values in the uncorrelated and positively correlated costs treatment, respectively.

²⁰Refer to Section 3 for an explanation of the generalised winner’s curse.

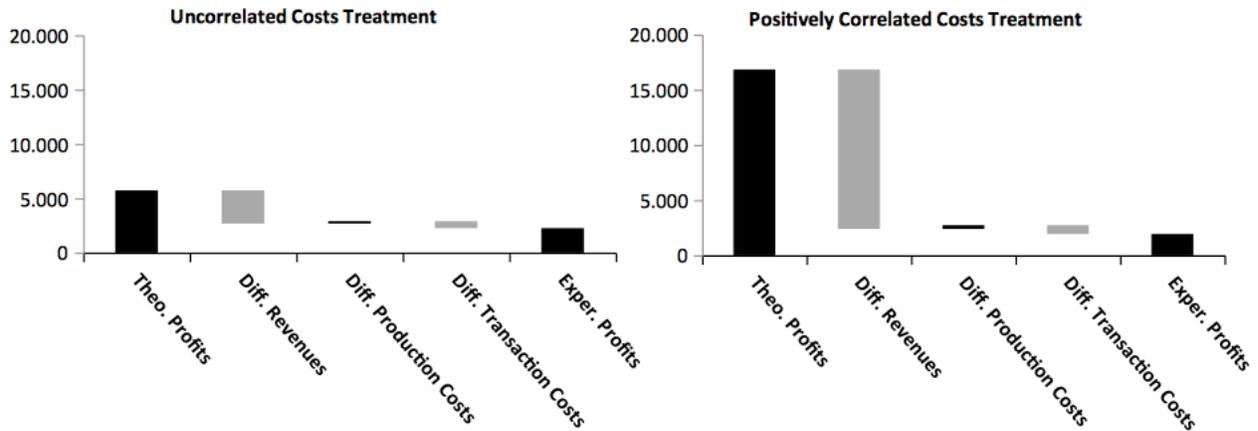
²¹See Appendix A for the formulae which have been used to calculate the outcome and efficiency variables.

²²The tests reported for the rest of the section are as follows. The null hypothesis for all of them is $H_0 : \mu_0 - \mu_1 = 0$, where μ_1 (resp. μ_0) is the mean for the positively correlated (resp. uncorrelated) costs treatment. The alternative hypothesis for market price and profits is $H_1 : \mu_0 - \mu_1 < 0$. For deadweight loss, the alternative hypothesis is $H_1 : \mu_0 - \mu_1 > 0$. The reported p -values are for the corresponding one-sided tests.

Average profits in each treatment are substantially lower than their corresponding ex ante equilibrium predictions. In particular, average profits are 40.2% (resp., 11.7%) of the theoretically predicted values in the uncorrelated (resp., positively correlated) costs treatment.²³ We also observe that, in the correlated costs treatment, bidders forgo a large percentage of ex ante expected profits—as typically occurs in auctions where bidders ignore the correlation among costs (Kagel and Levin 1986). The average difference in profits between the two treatments is not statistically significant (one-sided Mann–Whitney U-test: $n_1 = 6$, $n_2 = 6$, $p = 0.389$; one-sided t -test with unequal variance: $n_1 = 6$, $n_2 = 6$, $p = 0.716$).

In order to further understand the very large divergence between theoretical and experimental profits, we decompose this difference into the various components of profits (revenues, production costs and transaction costs) and we illustrate it by a waterfall chart in Figure 6.

Figure 6: *Decomposition of the differences between theoretical and experimental profits into the various components in each treatment.*



Note. In each graph, the first bar of each graph corresponds to ex-ante expected profits at the equilibrium allocation. The second bar to the difference between theoretical and experimental revenues; the third bar to the difference between theoretical and experimental production costs; The fourth bar to the difference between theoretical and experimental transaction costs. The last bar corresponds to average experimental profits. A waterfall chart shows how an initial value (whole column) increases or decreases by a sequence of intermediate positive or negative values (floating columns) to reach a final value (whole column). A colour-code is used for distinguishing positive (grey) from negative (black) intermediate values.

The figure shows that, in both treatments, the largest difference between theoretical and experimental profits is driven by the revenues component.²⁴ Therefore, differences in experimental

²³The theoretical predictions assume that all subjects play the Bayesian Nash equilibrium in the corresponding treatment. Note that if one subject realises that her opponents are *not* playing the Bayesian Nash equilibrium then it is not optimal for her to play the Bayesian Nash equilibrium, either; therefore, ex ante expected profits are different. See Section 6 for further discussion.

²⁴Recall that profits of seller i at time t are $\pi_{it} = p_i x_{it} - \theta_{it} x_{it} - \frac{\lambda}{2} x_{it}^2$. The first term corresponds to revenues; the second to production costs and the third to transaction costs.

and theoretical profits are primarily driven by market prices being lower than theoretically predicted, since the supply functions submitted by subjects are flatter than the equilibrium supply function and exhibit a large heterogeneity in both treatments; the difference and heterogeneity here are more pronounced in the positively correlated costs treatment, as seen in Table 3. These differences in market prices are amplified when translated to revenues since the market price is multiplied by units sold, the average of which is 33.33.²⁵

With regard to the efficiency of experimental allocations, there is a difference between treatments—in the direction predicted by Hypothesis (F)—but it is not statistically significant (one-sided Mann–Whitney U-test: $n_1 = 6$, $n_2 = 6$, $p = 0.583$; one-sided t -test with unequal variance: $n_1 = 6$, $n_2 = 6$, $p = 0.242$). So that we can better understand why the more heterogeneous behaviour in the positively correlated (than in the uncorrelated) costs treatment does *not* lead to differences in the efficiency of allocations, we decompose the experimental deadweight loss (dwl) into three components as follows:

$$dwl = \left(\frac{\lambda}{2}\right) \left(\sum_{i=1}^n \left(x_i^e - \frac{q}{n}\right)^2 + \left(x_i^o - \frac{q}{n}\right)^2 - 2\left(x_i^e - \frac{q}{n}\right)\left(x_i^o - \frac{q}{n}\right) \right), \quad (3)$$

The first term of (3) captures the variance of the experimental allocation, the second term captures the variance of the efficient allocation, and the third term captures the *covariance* of the experimental and efficient allocations.

After comparing these three components of deadweight loss across treatments, we can make several observations. First, variances in the experimental allocations do not differ between treatments. Second, the variance in efficient allocations is greater in the uncorrelated than in the positively correlated costs treatment; this finding corresponds to the theoretical prediction that follows from the experiment’s configuration of parameters. Third, the covariance of the experimental and efficient allocations is substantially larger in the uncorrelated than in the positively correlated costs treatment. When we consider both the three dwl components and their corresponding coefficients, we find that the difference between treatments with regard to variance in the efficient allocations is offset by the difference in the covariance term. For a graph displaying of these components, see Figure 10 (and Table 14) in Appendix F.

We also compare the deadweight losses of the experimental and equilibrium allocations. The values reported in Table 5 show that, in both treatments, the average experimental deadweight losses are considerably larger than those of the equilibrium allocations. This result is explained as follows. In both treatments, the deadweight losses at the equilibrium allocations are due to insufficient dispersion and covariation with respect to the corresponding efficient allocations. In

²⁵Our analysis of the variance in experimental outcomes (see Appendix F) establishes that variances in market price, profits, and deadweight loss are driven by neither subject nor group heterogeneity but rather, for the most part, by heterogeneity across rounds.

the *uncorrelated* costs treatment, the deadweight loss at the experimental allocation is larger than at the equilibrium because the covariation between the experimental and efficient allocations is not enough to compensate for the sum of the variances in the experimental and efficient allocations (these two variances are very similar). In the *positively correlated* costs treatment, however, the difference between the deadweight losses of the experimental and equilibrium allocations is driven mainly by the former’s comparatively greater variance.

We now evaluate Hypotheses (E) and (F), which address differences in experimental outcomes between treatments.

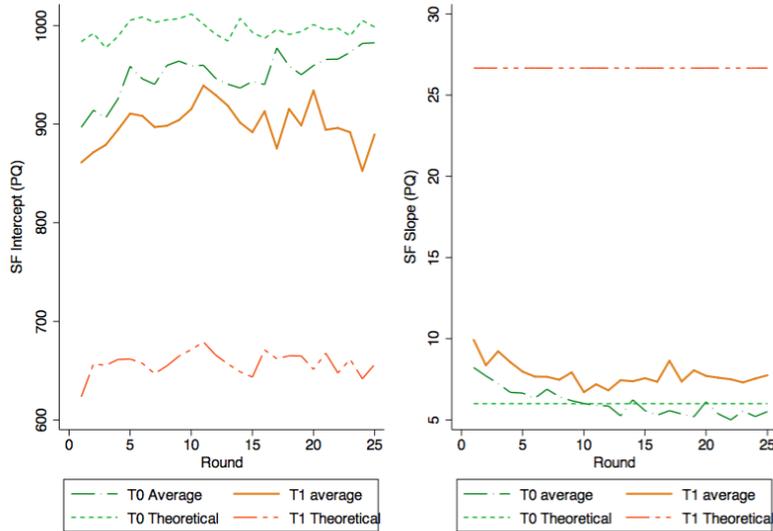
Result 4 (Differences in average market prices, profits, and deadweight losses between treatments). There is no average difference between treatments with respect to market prices, profits, or deadweight losses.

Result 4 implies that market power is *not* greater in the positively correlated than in the uncorrelated costs treatment. Hence we find no supporting evidence for Hypothesis (E) or Hypothesis (F).

5.3 Time trends in behaviour and outcomes

Figure 7 plots, for each treatment, the evolution across rounds of the supply function intercept and slope as well as the corresponding theoretical predictions.

Figure 7: *Evolution across rounds of the supply function slope and intercept in each treatment.*



Note. The theoretical supply function intercept is calculated using the average signal received in each round for each treatment; the theoretical predictions are calculated using the actual draws of the signal received. See Appendix B for a comparison between actual draws and the theoretical distribution. SF = supply function; T0 = uncorrelated costs treatment; T1 = positively correlated costs treatment.

The left-hand side of this figure shows that the supply function intercept increases across rounds in the uncorrelated costs treatment (a regression of the supply function intercept on round yields a coefficient of 2.28, with $p = 0.002$); in the positively correlated costs treatment, the intercept increases during the first ten rounds of play and then decreases across rounds (a regression of the supply function intercept on round yields an insignificant coefficient with $p = 0.927$). The right-hand side of Figure 7 shows that the average supply function slope becomes smaller across rounds in the uncorrelated costs treatment (a linear regression of the supply slope on round yields a coefficient of -0.096 with $p = 0.000$), but there is no significant time trend in the positively correlated costs treatment (a linear regression of supply function slope on round yields an insignificant coefficient with $p = 0.212$).²⁶ Thus we find that the evolution of behaviour across rounds is different in each treatment.

Furthermore, we observe that the change in behaviour (with regard to both the intercept and slope of the supply function) is more pronounced in the first ten rounds of play and especially in the first five rounds. In the last five rounds of play, the average supply function slope is stable; however, during these rounds the supply function intercept increases sharply (resp., decreases) in the uncorrelated (resp., positively correlated) costs treatment.

In the uncorrelated costs treatment, we find that both the intercept and slope of the supply function tend toward the theoretical prediction as the number of rounds increases. Yet in the positively correlated costs treatment we observe no decline across rounds in the difference between the average supply function and the theoretical prediction, which indicates that naïve behaviour persists.

Next we investigate the evolution of market outcomes across rounds. Table 6 shows the evolution of the average market price and profits in blocks of five rounds.

Table 6: *Evolution across rounds of average market price and profits in each treatment.*

Blocks of 5 rounds	Uncorrelated Costs Treatment			Positively Correlated Costs Treatment		
	Average Market Price	Average Profit	Deadweight loss	Average Market Price	Average Profit	Deadweight loss
Rounds [1,5]	1,106.05	1,960.58	3,133.28	1,122.86	1,847.73	2,549.45
Rounds [6,10]	1,111.44	1,911.60	2,115.95	1,115.48	1,641.48	2,467.70
Rounds [11,15]	1,112.27	2,400.06	1,850.10	1,119.46	1,994.96	1,937.91
Rounds [16,20]	1,110.58	2,570.13	2,277.78	1,144.07	2,347.59	1,571.81
Rounds [21,25]	1,111.58	2,434.41	1,735.43	1,116.79	1,857.32	1,528.18
Theoretical Prediction	1,200.00	5,612.75	612.75	1,544.68	16,567.37	467.73

The table reveals that, irrespective of treatment, there is no time trend in the average market price (a regression of market price on round yields an insignificant coefficient in both cases, with $p = 0.561$ in the uncorrelated costs treatment and $p = 0.817$ in the positively correlated

²⁶Our unit of observation for the regressions reported in Section 5.3 is again the group across rounds; as before, there are 150 observations in each treatment.

costs treatment). In addition, we find that profits increase after the tenth round in the uncorrelated costs treatment but not in the positively correlated costs treatment (regressing profits on round yields an insignificant coefficient, with $p = 0.103$ in the uncorrelated costs treatment and $p = 0.577$ in the positively correlated costs treatment).²⁷ Deadweight losses decrease over time in both treatments (a regression of deadweight loss on round gives a coefficient of -56.40 with $p = 0.003$ in the uncorrelated costs treatment and of -56.29 with $p = 0.000$ in the positively correlated costs treatment). In the uncorrelated costs treatment, deadweight losses are decreasing in the number of rounds because the covariance between the experimental and efficient allocations increases substantially. In the positively correlated costs treatment, deadweight losses decrease for a different reason: variance in the experimental allocations decreases significantly with the number of rounds, thus reducing the difference between the variance of the experimental and efficient allocations.

Finally, we report an additional finding about the evolution of behaviour across rounds. This result emerged from our analysis of the experimental data.

Result 5 (Evolution of supply functions and outcomes across rounds). **(i)** The evolution of supply functions across rounds is different in the two treatments: in the uncorrelated costs treatment, average behaviour starts close to the equilibrium prediction and, across rounds, moves even closer to that prediction; in the positively correlated costs treatment, average behaviour starts far from from the equilibrium prediction and, across rounds, does not move much closer to that prediction. **(ii)** In both treatments, deadweight losses decrease as subjects gain bidding experience; however, we find no evidence that market prices evolve across rounds. In the uncorrelated costs treatment, there is some evidence that profits increase as subjects gain bidding experience. In the positively correlated costs treatment, we do not find that profits evolve across rounds.

Result 5 suggests that the learning process may be different in the two treatments, a possibility that we explore further in Section 6.3. The finding that, as the number of rounds increases, behaviour in the positively correlated costs treatment does not move much closer to the equilibrium prediction suggests that naïve behaviour persists in this treatment.

²⁷Although we observe a time trend in profits (after the tenth round) in the uncorrelated costs treatment, that trend is not statistically significant when data are aggregated at the group level. However, in Section 5.4 we establish that this time trend is statistically significant when choices are considered at the individual level across rounds.

6 Examining the Data More Closely

In this section, we first study a subjects' strategic incentives; we then perform cluster analysis to descriptively study how close subjects' choices are to the various theoretical benchmarks. Third, we provide a description of the determinants of the evolution of behaviour across rounds, and finally we analyse subjects' responses to our post-experiment questionnaire. Overall, all the parts provide an explanation of the results we observe.

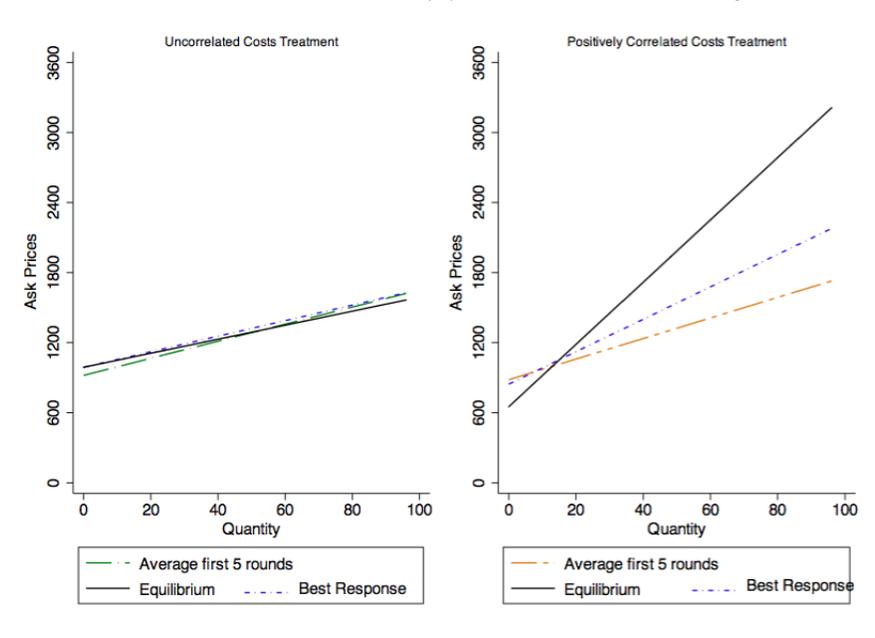
6.1 Best-response analysis

Bidding in the positively correlated costs treatment involves a higher degree of strategic complexity than bidding in the uncorrelated costs treatment. This is because the market price is not informative about costs in the uncorrelated costs treatment whereas, in the positively correlated costs treatment, equilibrium reasoning requires that a subject correctly understand how the market price *is* informative about the average cost, which also involves having the correct higher-order beliefs. To increase our understanding of a bidder's strategic incentives, we derive theoretically the best-response strategy of seller i while assuming that she knows the average strategy of rivals, which determines her residual demand. We then compute the comparative statics of the best-response with respect to the average strategy of those rivals.

In each treatment, we first analyse the best-response supply function to the rivals' average strategy during the first five rounds of bidding.²⁸ In the uncorrelated costs treatment, the average supply function for the first five rounds is steeper and has a lower intercept than the equilibrium. If we assume that rivals bid as the "representative seller" during the first five rounds, then seller i 's best-response is to bid a supply function that is flatter and has a higher intercept than that of her rivals but that is still steeper (and with a higher intercept) than the equilibrium supply function. In the positively correlated costs treatment, however, the average supply function for the first five rounds deviates substantially from the equilibrium supply function: the former is much flatter and has a higher intercept than predicted. In this case, seller i 's best-response supply function is steeper than that of her rivals' yet flatter than the equilibrium supply function. The intercept of seller i 's best-reply supply function is also between these two benchmarks. We illustrate these features of the best-response supply function in each treatment and compare them with the corresponding equilibrium supply functions in Figure 8.

²⁸Since behaviour does not evolve much across rounds, it follows that the features of the best-response supply function would be similar had we considered alternative definitions of the rivals' average strategy.

Figure 8: *Equilibrium supply function versus average supply function for the first five rounds of bidding, by treatment, and best-response supply function to the average.*



We shall identify the drivers of these features, focusing on the supply function slope because it is the most relevant variable from a strategic standpoint.²⁹ There are two effects that characterise the best-response slope of seller i . First is the *strategic effect*, which occurs regardless of the correlation among costs. If rivals bid a steep supply function, then the slope of the inverse residual demand increases and so seller i also has an incentive to bid a steep supply function. Yet if rivals bid a flat supply function, then the slope of the inverse residual demand decreases and so seller i likewise has an incentive to bid a flat supply function. Because of this strategic effect, the best-response slope of seller i is increasing in the rivals' average supply function slope.

In addition to the strategic effect, when costs are positively correlated there is also an *inference effect* that is related to information conveyed by the price. Seller i correctly thinks that a high price that the average signal of her rivals is high; therefore, if costs are positively correlated then seller i deduces that her own costs must be high. Hence seller i has an incentive to bid a steeper supply function than if costs were uncorrelated. When costs are positively correlated, the inference effect causes seller i 's best-response slope to be decreasing in her rivals' average supply function slope. The reason is that, when costs are positively correlated, the inference effect moderates the reaction to the price, the more so, the more rivals react to the price. This is because a higher reaction to the price by rivals induces a trader to also give a higher weight to the price in the estimation of her cost and hence it increases the magnitude of the inference effect.

²⁹The mathematical derivations are given in Appendix H.

Suppose seller i bids in the positively correlated costs treatment and that seller i 's rivals fall prey to the generalised winner's curse, thus bidding as in the equilibrium of the uncorrelated costs treatment. Then the slope of seller i 's best-response supply function is increasing in the supply function slope of rivals, which means that the strategic effect dominates the inference effect. However, the inference effect does *moderate* the magnitude of the increase in the best-response supply function's slope as a result of an increase in the slope of her rivals' supply function—when costs are uncorrelated. It follows that the optimal response for seller i is to bid a supply function whose slope is between the slope of naïve sellers' (average) supply function and the slope of the equilibrium supply function. In other words: a sophisticated seller who is best responding to naïve rivals has an incentive to bid a flatter supply function than the equilibrium would predict, which leads to the behaviour of naïve and sophisticated sellers being less distinct. An equivalent result was first noted by Camerer and Fehr (2006) in the context of games characterised by strategic complementarities and by sophisticated and boundedly rational subjects.³⁰ Figure 11 (in Appendix H) plots, for each treatment, the best-response supply function slope as a function of rivals' average supply function slope.

In sum, the positively correlated costs treatment presents a higher degree of strategic complexity than the uncorrelated costs treatment. This explains why average choices in the positively correlated costs treatment are farther from the equilibrium than in the uncorrelated costs treatment. In the former, the best-response strategy of a sophisticated seller who best responds to her rivals' actual choices is one that falls *between* the equilibrium of the positively correlated costs treatment and the benchmark of the generalised winner's curse (i.e., the equilibrium of the uncorrelated costs treatment). Therefore, we see less difference in behaviour—between treatments and types of subjects (naïve and sophisticated)—than is predicted by the equilibrium.

6.2 Cluster analysis

We can use cluster analysis to conduct a descriptive study of experimental choices; this approach enables us to “organise” the heterogeneity in behaviour and relate it to the various theoretical benchmarks. We shall present the results of analysing subjects' choices in the first five rounds of bidding, when strategic thinking is most relevant, and also in the last five rounds, when subjects are most experienced.

We use a model-based clustering technique, which is based on mixture models, for endogenously and simultaneously determining (a) the number of clusters and (b) the type of model characterised by properties of the underlying probability distributions (orientation, volume, and

³⁰We can also relate best-response and equilibrium strategies to models of strategic thinking. In the terminology of a level- k model (Nagel 1995), a level-1 subject would best respond to the strategies of rivals believing that rivals are level-0 subjects (e.g., subjects who fall prey to the generalised winner's curse). So for $k \geq 1$, a level- k subject would best respond to rivals believing they are all of level $k - 1$. The Bayesian Nash equilibrium would then be obtained as the limiting behaviour when all subjects perform an infinite number of level- k iterations.

shape). The Bayesian information criterion (BIC) allows us to select the most appropriate clustering model. We implement this procedure using the *Mclust* package in R (developed by Fraley and Raftery 2002), which uses the expectation-maximisation (EM) algorithm.

The data used for the cluster analysis are as follows. We first compute the average subject choice for each group of five rounds. Each choice is two-dimensional in that they each consist of the supply function slope and intercept. We then compute the deviation of each individual’s average choice with respect to the corresponding equilibrium benchmark, and the resulting deviations form the basis of our cluster analysis. We conduct the cluster analysis separately for each treatment and each group of five rounds. To facilitate comparisons within a given treatment—and our interpretation thereof—we choose the same number of clusters and type of model in each group of five rounds.³¹

Table 7 summarises the results of our cluster analysis by reporting the type of model, the number of components, the frequency of subjects in each cluster, and the average supply slope and intercept of subjects belonging to a particular cluster (in Appendix I we provide results from cluster analysis based on the intermediate time periods in five-round blocks). The findings reported in Table 7 are plotted in Figure 9.

In the uncorrelated costs treatment, the BIC yields two clusters with a model of VEV (varying volume, equal shape, and varying orientation with ellipsoidal distribution) type. In both time periods, cluster 1 groups subjects whose choices are close to both the Bayesian Nash equilibrium choice and the best-response to the average choice. Cluster 1 contains 58% of the subjects in the first five rounds, and this percentage rises to 72% in the last five rounds. Subjects in cluster 2 bid a steeper supply function with a lower intercept than the equilibrium prediction. These results suggest that most subjects’ behaviour in the uncorrelated costs treatment accords fairly well with the theoretical predictions (as already indicated by the average results reported in Section 5). Furthermore, we observe that 14% of the subjects move from cluster 2 to cluster 1 from the beginning to the end of the experiment.

³¹In the *uncorrelated* costs treatment, the five clustering models presented in the text (and in Appendix I) have the highest possible BIC or there is a BIC difference of less than 2 with respect to the optimal model. (Fraley and Raftery 2002 argue that, if the difference in BIC between two models is less than 2, then this constitutes at least weak evidence that one model is better than the other.) In the *positively correlated* costs treatment, no single model has the highest BIC for all five sets of data. Hence we consider the second-best model for all five time periods, which is a three-component model of VEI (varying volume, equal shape, and coordinate axes orientation with diagonal distribution) type.

Table 7: *Descriptive statistics characterising each cluster.*

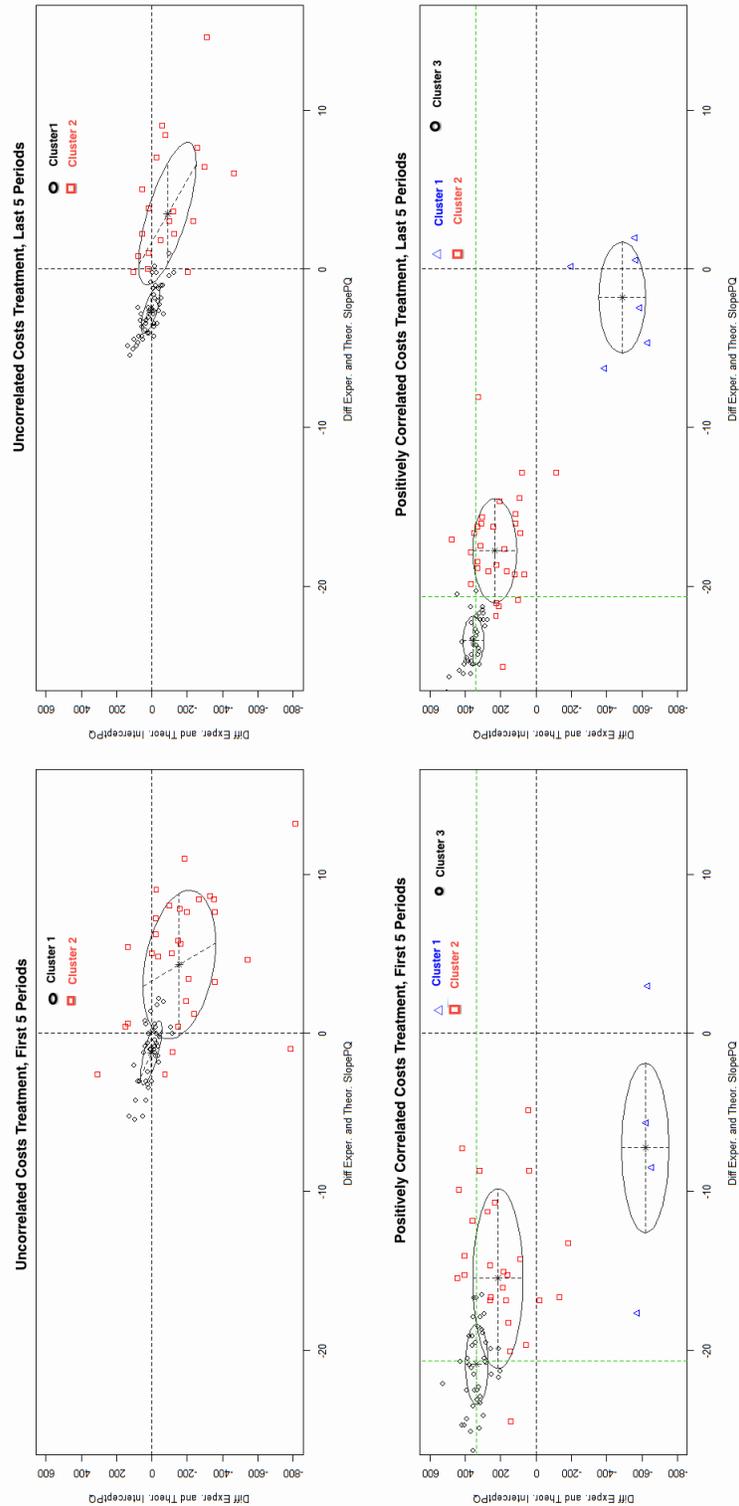
Period	Model & Components	Log-likelihood & BIC	Cluster	Number of subjects	Relative frequency	Average SF slope (s.d.)	Average SF intercept (s.d.)
Uncorrelated costs treatment							
First 5 rounds	VEV 2	-636.45 -1,315.66	Cluster 1	42	58%	4.83 (2.00)	995.59 (69.39)
			Cluster 2	30	42%	10.77 (4.04)	814.91 (238.64)
Last 5 rounds	VEV 2	-576.34 -1195.44	Cluster 1	52	72%	3.44 (1.57)	1,002.99 (74.16)
			Cluster 2	20	28%	10.25 (3.78)	897.9 (145.21)
Equilibrium SF						6.00	1,000.00
Positively correlated costs treatment							
First 5 rounds	VEI 3	-664.98 -1,381.27	Cluster 1	4	6%	19.45 (8.49)	12.95 (16.55)
			Cluster 2	26	36%	12.32 (4.34)	844.82 (169.04)
			Cluster 3	42	58%	5.62 (2.51)	990.07 (79.19)
Last 5 rounds	VEI 3	-640.22 -1,331.77	Cluster 1	6	8%	24.87 (3.22)	126.83 (194.50)
			Cluster 2	30	42%	9.17 (3.22)	880.37 (120.61)
			Cluster 3	36	50%	3.29 (1.61)	1,014.94 (56.71)
Equilibrium SF						26.68	655.32

Note. The intercept of the equilibrium supply function is calculated using the average signal realisation. Standard deviations (s.d.) are reported in parentheses.

In the positively correlated costs treatment, the BIC yields three clusters and a model of VEI type (see note 30). In the first five rounds, no subject bids as in the equilibrium. In these rounds, we see that: 58% of the subjects are in cluster 3, which approximately corresponds to the benchmark of the generalised winner's curse (equilibrium of the uncorrelated costs treatment); 36% of the subjects are in cluster 2, which means their behaviour is not inconsistent with that of a sophisticated bidder (i.e., they bid a supply function close to the average theoretical best reply); and 6% of the subjects are in cluster 1, meaning that they choose a supply function with an extremely steep slope and a very low intercept (close to zero). The subjects in cluster 1 are attracted to the zero-intercept focal point.³²

³²These cluster results can be loosely related to the level- k model of strategic thinking. Subjects in cluster 3 can be thought to be level-0; subjects in cluster 2 to be level-1; and subjects in cluster 1 cannot be easily related to a level of strategic thinking since one dimension (the supply function slope) is close to the equilibrium but the other dimension (the supply function intercept) is far from all benchmarks.

Figure 9: *Results of cluster analysis.*



Note. Each data point corresponds to a subject's two-dimensional (slope and intercept) average choice over the first five or last five rounds of bidding. The x -axes (resp., y -axes) represent the difference between theoretical and experimental supply function slopes (resp., intercepts). The black dashed lines plot the equilibrium supply function slope and intercept in each treatment, where $(0,0)$ corresponds to the equilibrium. In the two graphs corresponding to the positively correlated treatment, there are green dashed lines that plot the generalised winner's curse benchmark (i.e., the equilibrium of the uncorrelated costs treatment).

To explore further the relationship between clusters and optimal behaviour, in Table 18 of Appendix I we summarise—for each cluster—the average difference between the experimental supply function and the theoretical best-response supply function. In the uncorrelated costs treatment, cluster 1 groups subjects whose supply functions are close to the theoretical best-response in both time periods; in the positively correlated costs treatment, cluster 2 groups subjects whose choices are close to the theoretical best-response slope and intercept in both time periods. In the positively correlated costs treatment, subjects in cluster 1 bid supply functions that are too steep and have too low an intercept vis-à-vis the theoretical best reply while subjects in cluster 3 bid supply functions that are too flat and have too high an intercept.

In the positively correlated costs treatment, there is an exodus of subjects from cluster 3 to cluster 2 as the experiment progresses from the first five to the last five rounds of bidding; that movement might lead us to infer that most subjects bid a supply function with a steeper slope. However, Table 7 shows that the average supply function slope of subjects in clusters 2 and 3 actually decreases from the first five to the last five rounds of bidding, which means that—during those last five rounds—the benchmark of the generalised winner’s curse is located between clusters 2 and 3. In light of this evidence, we conclude that there are no important changes in the evolution of clusters between the beginning and the end of the experiment. We can summarise these findings as follows.

Result 6 (Cluster analysis). **(i)** In the uncorrelated costs treatment, 58% of the subjects are in the same cluster as the equilibrium in the first five rounds—a percentage that increases to 72% in the last five rounds. This cluster groups subjects whose supply function is closest to the theoretical best-response. **(ii)** In the positively correlated costs treatment, no subject bids as predicted by the equilibrium in either the first five rounds or the last five rounds of bidding. The proportion of subjects who are in the same cluster as the “generalised winner’s curse” benchmark is 58% in the first five rounds of bidding and 50% in the last five rounds; the proportion of subjects whose bidding behaviour is not inconsistent with sophisticated behaviour (i.e., best responding to the average supply function) is 36% in the first five rounds and 42% in the last five rounds of bidding; and the proportion of subjects who bid a steep supply function with a low intercept is 6% (resp. 8%) in the first (resp. last) five rounds of bidding.

Observe that, for a subject who is best responding to actual choices in the positively correlated costs treatment, it is individually optimal to bid a supply function with a flatter slope and with a higher intercept than the equilibrium would predict. This fact explains why average behaviour and outcomes are not sufficiently differentiated between treatments and hence why it is difficult empirically to distinguish the behaviour of naïve and sophisticated subjects. Furthermore, if

some subjects never learn (or learn very slowly) and remain in cluster 3, then a subject who is best responding to actual choices may have no incentive to bid increasingly steeper supply functions during the course of the experiment. This may explain why slopes do not evolve toward the equilibrium prediction in the positively correlated costs treatment. The rest of Section 6 is dedicated to uncovering what determines the evolution of choices across rounds in the two treatments. Section 6.3 conducts a more detailed analysis of behaviour since it considers the each individual choice across all rounds of the experiment (rather than in blocks of 5 periods as has been done in this section).

6.3 Evolution of choices across rounds

Our goal is to provide a description of the evolution behaviour observed experimentally and thereby to understand what drives differences in the changes of behaviour (across rounds) in the two treatments. The experimental design provided subjects with complete feedback after each round, feedback that included each subjects' own choice and profits as well as the choices and profits of the other two market participants. Subjects might be therefore be influenced by the average choices of rivals in the same group. We designed this high-information environment, and stipulated 25 rounds per session, so as to foster participants' learning about the equilibrium. A subject's use of a particular learning model depends on both the available information and her cognitive abilities. Given the experiment's complexity, we considered only those learning models in which subjects had the necessary information to learn.

We therefore consider that the evolution of behaviour across rounds can be explained in terms of three main categories of learning models (for an excellent review, see Camerer 2003): experiential learning, in which subjects learn from their own experience; imitation-based learning, in which subjects choose—after the first round—a strategy chosen by other players in the previous round; and belief learning, an adaptive style whereby subjects update their beliefs based on the history of play. Unlike the case of Cournot competition, there are very few theoretical results on the convergence properties of learning models in the context of supply function competition with private information.³³ In this section, then, we aim only to provide an informative description of actual behaviour across rounds.

Our empirical strategy follows the approach of Huck et al. (1999) and Bigoni and Fort (2013); more specifically, we consider a representative model (or models) from each of the three categories just described. Moreover, we presume that the supply function slope is the strategic variable best able to summarise how a subject adjusts her behaviour across rounds. We forgo analysing the dynamics of the supply function intercept because, in our view, the subjects do not have enough information to learn from it; that is, we provided no feedback on the signals

³³For some general considerations, see Vives (2008, Chapter 7).

received by rivals. As a result, the supply function intercept could be affected by confounding factors, such as how subjects respond to the private signal. Note also that supply function intercepts are nearly identical in the two treatments.

The type of experiential learning on which we focus is *reinforcement learning* (RL), whereby a player increases her probability of playing a strategy (i.e., the strategy is reinforced) as a function of the profits previously earned via that strategy. A strategy’s “propensity” is the cumulative sum of the previous profits that it generated with such strategy, and a strategy’s “relative propensity” is its propensity divided by the aggregate propensities of all strategies employed in a given round (cf. Erev and Roth 1998). Our implementation of reinforcement learning assumes that a player bids the supply function slope exhibiting the highest relative propensity.

We then consider two models based on learning by imitation. Of these, the first is a payoff-independent model—*imitation of the average* (IA)—under which subjects imitate their rivals’ average strategy in the previous round. The second is a payoff-dependent model, *imitation of the best* (IB), whereby subjects imitate the strategy of the subject who garnered the highest profits in the previous round (and where that most successful player imitates herself).³⁴

As a representative of belief learning models we consider *best-response dynamics* (BR), whereby subjects best respond to their rivals’ previous-round strategies.³⁵ Recall that the features of a best-response strategy were described, for each treatment, in Section 6.1.³⁶

We estimate a dynamic panel regression to identify the determinants of each subject’s supply function slope across rounds with respect to the various learning models described previously.³⁷ Thus, we estimate the following equation:

$$\begin{aligned} SlopePQ_{it} = & \gamma SlopePQ_{it-1} + \beta_{RL}RL_{it-1} + \beta_{IA}IA_{it-1} \\ & + \beta_{IB}IB_{it-1} + \beta_{BR}BR_{it-1} + \nu_i + \omega_{it}. \end{aligned} \tag{4}$$

Here $SlopePQ_{it}$ is subject i ’s supply function slope in round t ; $SlopePQ_{it-1}$ is the supply function

³⁴We have not considered learning models such as *imitate the exemplary firm* (Offerman et al. 2002), since this would require subjects to calculate the joint maximisation of profits for three “firms” producing a given level of output. That task is far from trivial in our experimental environment of supply functions and private information.

³⁵The difference between the best-response analysis in this section and the version analysed elsewhere (in Section 6.1, Section 6.2, Appendix H, and Appendix I) is that here we consider the best-response to rivals’ strategy of the previous rather than the current round.

³⁶We do not consider the *fictitious play* model because it is cognitively more demanding than best-response dynamics. In particular, it requires subjects to remember the average strategies of all previous rounds of play; best-response dynamics, in contrast, requires subjects to remember the strategies only of the other two players in the previous round.

³⁷As compared with Huck et al. (1999) we have added the reinforcement learning term. As compared with Bigoni and Fort (2013), we have replaced the “trial and error” experiential learning model with the reinforcement learning model, which has gained wider acceptance in the literature.

slope of subject i in the previous round (i.e., round $t-1$), which gives an indication of behavioural persistence; RL_{it-1} is the supply function slope corresponding to the highest relative propensity based on player i 's profits in round $t-1$ (reinforcement learning); IA_{it-1} is the average supply function slope of i 's rivals in the previous round (imitation of the average); IB_{it-1} is the supply function slope that corresponds to the highest profits of i 's rivals in round $t-1$ and if i 's profits were the highest in period $t-1$ then subject i imitates herself (imitation of the best); BR_{it-1} is subject i 's theoretical best reply to her rivals' supply function slope in the previous round (best-response dynamics); ω_{it} is the error term; and v_i represents panel-level random effects.³⁸

If a given learning model fully explained how a subject adjusts her supply function slope across rounds, then its regression coefficient would be 1 and the other coefficients would be insignificant. In the more likely event that the model's explanation is only partial, we expect its coefficient to be positive and significantly different from 0.

Table 8 presents the results from estimating equation (4)—for individual choices across rounds in each treatment—using the approach proposed by Arellano and Bover (1995) and Blundell and Bond (1998).³⁹

Table 8: *Dynamic panel estimation: Evolution of behaviour across rounds.*

	Uncorrelated Costs Treatment	Positively Correlated Costs Treatment	
	Dependent variable: Slope _{it} (sl _{it})		
	(1)	(2)	(3)
Lag of Dependent Variable (sl _{it-1})	0.12*** (0.041)	0.11** (0.047)	0.31*** (0.041)
Reinforcement learning	0.30*** (0.069)	0.35*** (0.057)	0.25*** (0.063)
Imitation of the average	-	0.11*** (0.040)	0.12*** (0.038)
Imitation of the best	0.031 (0.044)	0.061 (0.038)	0.11*** (0.033)
Best response	0.34*** (0.063)	-	0.02 (0.016)
Number of observations	1,727	1,728	1,663
Number of instruments	184	184	230

Note. Robust standard errors are given in parentheses.

*, **, and *** denote significance at (respectively) the 10%, 5%, and 1% levels.

Columns [1] and [2] of Table 8 present the results of estimating equation (4) for the un-

³⁸Details of how these learning variables have been created are available from the authors upon request.

³⁹The estimation was conducted via Stata's *xtdpdsys* command using the two-step estimator. The regressors corresponding to each of the learning models were considered to be predetermined because, although possibly correlated to past errors, they are not correlated with future errors. The two restrictions on these instruments are that (i) we use at most one lag of the dependent variable as an instrument and (ii) each regressor appears as contemporaneous and at most one lag is used as an instrument.

In all the equations presented, the post-estimation specification test for autocorrelation of the error terms gives no evidence of autocorrelation of order higher than 1 at the 10% significance level. In addition, we cannot reject the null hypothesis that the overidentifying restrictions are valid (Sargan test) for all the regressions displayed.

correlated costs treatment. We cannot estimate the full learning model of equation (4) because the best-response dynamics variable (BR_{it-1}) is perfectly collinear with the imitation of the average variable (IA_{it-1}) owing to the functional relationship between them.⁴⁰ Hence we cannot disentangle the effects of these two learning models in the uncorrelated costs treatment.

The results of estimating equation (4) for the uncorrelated costs treatment are presented in Table 8[1], where IA is excluded, and in Table 8[2], where BR is excluded. In both equations, all significant coefficients have the expected signs. We find that RL is an important component of the adjustment dynamics whereas IB is not significant. In column [1], BR is significant and has the largest coefficient; in column [2], IA has a positive and significant coefficient. In summary: we find that reinforcement learning, imitation of the average, and best-response dynamics jointly contribute to explain the evolution of behaviour across rounds in the uncorrelated costs treatment. Behaviour is fairly persistent.

Column [3] of Table 8 reports regression results for the positively correlated costs treatment; here the estimation *excludes* 64 observations—namely, those for which a theoretical best-response was unfeasible given our experimental design.⁴¹ In this treatment, too, behaviour exhibits some persistence over time. Also, the terms for RL, IB, and IA jointly provide a partial description of how subjects adjust their supply function slope across rounds; in contrast, the BR term does not contribute to explaining how subjects' behaviour evolves.

These findings may be summarised as follows.

Result 7 (Determinants of evolution of behaviour over rounds). There are significant differences between how subjects learn in the two treatments. Best-response dynamics is a more prominent determinant of the evolution of behaviour across rounds in the uncorrelated than in the positively correlated costs treatment. Imitation of the best is a significant factor in the positively correlated costs treatment but not in the uncorrelated costs treatment. In both treatments, reinforcement learning and imitation of the average help to explain the evolution of behaviour across rounds.

Result 7 shows that the evolution of behaviour across rounds is a composite process whose elements include the various learning models considered but with different weights in each treatment. That best-response dynamics has a significant effect on learning when costs are uncorrelated but not when they are positively correlated is consistent with the notion that the former treatment environment is cognitively simpler for experimental subjects.

⁴⁰For details, see Appendix H and equation (17) therein.

⁴¹More specifically, the theoretical best-response would have required a negative supply function slope, which was not allowed. The average supply function slope of these 64 excluded observations does not differ significantly from the average of those that were included. If we instead include the 64 observations when estimating equation (4), the coefficients are strongly similar to those reported in Table 8[3].

Notice that this result may appear to contrast with Result 6 which shows that, in the positively correlated costs treatment, the proportion of subjects whose bidding behaviour is not inconsistent with sophisticated behaviour (i.e., best responding to the average supply function) is 36% in the first five rounds and 42% in the last five rounds of bidding. However, the analysis of Section 6.2 used the average subject choice for each group of five rounds while the dynamic analysis of this section considers the individual choice for each round. Therefore, we conclude that on a round-to-round basis, subjects in the positively correlated costs treatment do not appear to be best responding to the choices of rivals in the previous round. However, there is evidence that when aggregating rounds in groups of five, some subjects in the positively correlated costs treatment (36% in the first five rounds and 42% in the last five rounds of bidding) are close to the average best-response.

6.4 Post-experiment questionnaire: Analysis

Our experiment’s main result is that most subjects in the positively correlated treatment behave *as if* they did not understand that the market price is informative about costs. Analysing the responses given by subjects in our post-experiment questionnaire offers further evidence of this interpretation. We asked two questions: “Do you think that a high market price generally means good, mixed, or bad news about the level of your costs?” and then “Explain your answer.” The aim of these questions was to see whether, after making decisions in 25 rounds, the subjects understood that price is informative in the positively correlated costs treatment; we also wanted to see if there were differences between the two treatments. We analyse the explanations given (responses to the second question) because they reveal more about the reasoning processes used by participants. Table 9 groups the questionnaire responses into types and reports a typical answer given in each category.

Table 9: *Classification of responses to post-experiment questionnaire.*

Relationship to equilibrium reasoning	Classification	Sub-Classification	Typical Answer
Equilibrium logic of the positively correlated costs treatment	<i>Market price is informative about costs</i>	<i>Market price is informative about costs</i>	“It can give you an approximation in terms of the cost, since it is related, and it would be high”
Equilibrium logic of the uncorrelated costs treatment	<i>Market price is not informative about costs</i>	<i>Market price is not informative about costs</i>	“The market price and the level of unit cost are not related. Each seller has its own unit cost independently of the market price”
Ambiguous	<i>High costs imply a high market price</i>	<i>High costs imply a high market price</i>	“If all sellers have a high unit cost, it is logical that the market price rises”
Ambiguous	<i>Other factors</i>	<i>Strategic uncertainty</i>	“It will depend on the unit cost of each seller/company and the prices set by each of them”
Ambiguous	<i>Other factors</i>	<i>Depends on other factors</i>	“It cannot be said precisely since it depends on other factors”
Ambiguous	<i>Other factors</i>	<i>Experience in the game</i>	“This has been my experience in the experiment”
Ambiguous	<i>Other factors</i>	<i>Unclassified</i>	Various answers
Ambiguous	<i>Not answered</i>	<i>Not answered</i>	Not answered

The answers classified as “market price is informative about costs” reflect the logic of the

equilibrium reasoning applicable to the positively correlated costs treatment, while answers classified as “market price is not informative about costs” reflect the equilibrium reasoning of the uncorrelated costs treatment. The answers classified as “high costs imply a high market price” are ambiguous, since they might reflect either (i) an understanding of the equilibrium reasoning of the positively correlated costs treatment or (ii) an understanding of the prediction, common to both treatments, that a high signal implies high costs and so subjects should bid a supply function with a higher intercept.⁴² Hence we shall not focus on interpreting either this third category of answers or on responses classified as “other factors”. Table 10 reports, for each treatment, the percentage of answers that correspond to each of the categories analysed.

Table 10: *Responses to post-experiment questionnaire by group and treatment.*

Answer Classification	<i>Market price is informative about costs</i>	<i>Market price is not informative about costs</i>	<i>High costs imply a high market price</i>	<i>Other factors</i>	<i>Not answered</i>
Uncorrelated Costs Treatment	8%	42%	17%	31%	3%
Positively Correlated Costs Treatment	17%	44%	18%	17%	4%

This table reveals that nearly twice the percentage of subjects fall in the “market price is informative about costs” category under the positively correlated treatment than under the uncorrelated treatment (17% vs. 8%), a result indicating that only a few subjects in the positively correlated costs treatment understood that the market price is informative about the level of costs. We observe also that “market price is not informative about costs” is the most prevalent answer in both treatments, with no major differences—as a function of treatment—in subjects making this response. Thus the responses to our post-experiment questionnaire are consistent with the interpretation that, in the positively correlated costs treatment, a large percentage of participants that fall prey to the generalised winner’s curse and a much smaller percentage of subjects are sophisticated bidders whose best-response is more similar to those that fall prey to the generalised winner’s curse than the equilibrium would predict.

Furthermore, we conducted two robustness sessions to disentangle whether the fact that subjects in the positively correlated costs treatment do not understand that the market price is informative about costs can be attributed to a failure to engage in Bayesian updating. We report the results of these additional sessions in Appendix H. We find that our results in this treatment are *not* due to a bias related to simple Bayesian updating.

⁴²We remark, however, that the latter is the first step toward understanding that the market price *is* informative about costs.

7 Concluding Remarks

We have analysed bidding behaviour in an experiment which reflects the complexity of real-world markets where bidders compete in supply functions, have incomplete information about their costs, and receive a private signal. We used the unique Bayesian Nash equilibrium prediction and its comparative statics as benchmarks when evaluating the experimental results. Our experiment employed a between-subjects design with two treatments: uncorrelated costs and positively correlated costs; the former served as our control treatment. In conducting the experiment, we chose numerical values for the model's parameters that allowed theoretical predictions about behaviour and outcomes to be sufficiently differentiated between the two treatments.

We find that most subjects bid in accordance with the equilibrium predictions in a simple strategic environment (uncorrelated costs treatment), yet most subjects fail to do so in a more complex strategic environment (positively correlated costs treatment) in which they are required to use the market price to extract information about the level of costs. On average, subjects in the positively correlated costs treatment bid flatter supply functions than predicted by the equilibrium. Our findings are consistent with the experimental literature on correlation neglect (see Section 2), a phenomenon that in our setting is intertwined with the generalised winner's curse. To the best of our knowledge, this paper is the first to document that phenomenon in an environment characterised by interdependent values and the auctioning of multiple units of a divisible good.

In electricity markets, mitigating the market power is a primary concern of regulators.⁴³ If costs are positively correlated then our experiment shows that, when a large proportion of subjects who are competing in supply functions neglect the correlation among costs, market outcomes are more competitive than in the equilibrium. Yet despite being more competitive, the resulting allocations present a higher productive inefficiency than the corresponding equilibrium allocations.

Our experiment suggests a few open questions for future research, both experimental and theoretical. Future work could explore mechanisms by which subjects learn to overcome the generalised winner's curse (e.g., asking subjects to come back to the laboratory a few days later (experienced bidders); replicating the experiment with professional traders; extending the number of rounds). The experimental findings reported here also call for the development of theoretical models that analyse market competition among participants who exhibit various degrees of strategic sophistication—as well as for examination of the convergence properties of different types of learning models in markets characterised by supply function competition and private information.

⁴³See, for example, Hortaçsu and Puller (2008) and Holmberg and Wolak (2015).

References

- [1] Arellano, M., and Bover, O., (1995), “Another Look at the Instrumental Variable Estimation of Error-Components Models,” *Journal of Econometrics*, 68: 29-51.
- [2] Ausubel, L. M., Cramton, P., Pycia, M., Rostek, M., and Weretka, M., (2014), “Demand Reduction and Inefficiency in Multi-unit Auctions,” *The Review of Economic Studies*, 81(4), 1366-1400.
- [3] Baldick, R., Grant, R., and Kahn, E., (2004), “Theory and Application of Linear Supply Function Equilibrium in Electricity Markets,” *Journal of Regulatory Economics*, 25(2), 143-167.
- [4] Baltagi, B.H., (2008), *Econometric Analysis of Panel Data*, 4th ed., Chichester, UK: Wiley.
- [5] Bigoni, M., and Fort, M., (2013), “Information and Learning in Oligopoly: An Experiment,” *Games and Economic Behavior*, 81: 192-214.
- [6] Blundell, R., and Bond, S., (1998), “Initial Conditions and Moment Restrictions in Dynamic Panel Data Models,” *Journal of Econometrics*, 87: 115-143.
- [7] Bolle, F., Grimm, V., Ockenfels, A., and del Pozo, X., (2013), “An Experiment on Supply Function Competition,” *European Economic Review*, 63(C): 170-185.
- [8] Brandts, J., Reynolds, S.S., and Schram, A., (2014), “Pivotal Suppliers and Market Power in Experimental Supply Function Competition,” *The Economic Journal*, 124: 887–916.
- [9] Brocas, I., Carrillo, J. D., Wang, S.W., and Camerer, C.F., (2014), “Imperfect Choice or Imperfect Attention? Understanding Strategic Thinking in Private Information Games,” *Review of Economic Studies*, 81: 944-970.
- [10] Camerer, C. F., (2003), *Behavioral Game Theory: Experiments in Strategic Interaction*, Princeton University Press.
- [11] Camerer, C.F., and Fehr, E., (2006), “When Does “Economic Man” Dominate Social Behavior?,” *Science*, 311: 47-52.
- [12] Carrillo, J. D., and Palfrey, T., (2011), “No Trade,” *Games and Economic Behavior*, 71(1), 66-87.
- [13] Cassola, N., Hortaçsu, A. and Kastl, J., (2013), “The 2007 Subprime Market Crisis Through the Lens of European Central Bank Auctions for Short-Term Funds,” *Econometrica*, 81: 1309–1345.

- [14] Charness, G., and Levin, D., (2009), “The origin of the winner’s curse: a laboratory study,” *American Economic Journal: Microeconomics*, 207-236.
- [15] Crawford, V. P., and Iriberry, N., (2007), “Level-k Auctions: Can a Nonequilibrium Model of Strategic Thinking Explain the Winner’s Curse and Overbidding in Private-Value Auctions?,” *Econometrica*, 75(6): 1721-1770.
- [16] Dufwenberg, M., and Gneezy, U., (2000), “Price Competition and Market Concentration: An Experimental Study,” *International Journal of Industrial Organization*, 18, 7-22.
- [17] Engelbrecht-Wiggans, R., List, J. A., & Reiley, D. H., (2005), “Demand Reduction in Multi-unit Auctions: Evidence from a Sportscard Field Experiment: Reply,” *The American Economic Review*, 95(1), 472-476.
- [18] Engelbrecht-Wiggans, R., List, J. A., and Reiley, D. H., (2006)., “Demand Reduction in Multi-Unit Auctions With Varying Number of Bidders: Theory and Evidence from a Field Experiment,” *International Economic Review*, 47(1), 203-231.
- [19] Enke, B., and Zimmermann, F., (2013), “Correlation Neglect in Belief Formation,” *IZA Discussion Paper No. 7372*.
- [20] Erev, I., and Roth, A., (1998), “Predicting how People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria,” *American Economic Review*, 88: 848–881.
- [21] Esponda, I., and Vespa, E., (2014), “Hypothetical Thinking and Information Extraction in the Laboratory,” *American Economic Journal: Microeconomics*, 6(4): 180-202.
- [22] Eyster, E., and Rabin, M., (2005), “Cursed Equilibrium,” *Econometrica*, 73(5), 1623-1672.
- [23] Fischbacher, U., (1999), “z-Tree 1.1.0: Experimenter’s Manual,” University of Zurich, Institute for Empirical Research in Economics.
- [24] Fraley, C., and Raftery, A.E. (2002), “Model-based Clustering, Discriminant Analysis, and Density Estimation,” *Journal of the American Statistical Association*, 97, 611-631.
- [25] Georganas, S., and Nagel, R., (2011), “Auctions with Toeholds: An Experimental Study of Company Takeovers,” *International Journal of Industrial Organization*, 29(1), 34-45.
- [26] Goeree, J., and Offerman, T., (2003), “Winner’s Curse Without Overbidding,” *European Economic Review*, 47(4), 625-644.

- [27] Holmberg, P., and Wolak, F.A., (2015), “Electricity Markets: Designing Auctions Where Suppliers Have Uncertain Costs,” *Cambridge Working Papers in Economics*, No. 1541.
- [28] Holt, C. A., and Sherman, R., (1994), “The Loser’s Curse,” *American Economic Review*: 642-652.
- [29] Hortaçsu, A., and Puller, S. L., (2008), “Understanding Strategic Bidding in Multi-unit Auctions: A Case Study of the Texas Electricity Spot Market,” *The RAND Journal of Economics*, 39: 86–114.
- [30] Huck, S., Normann, H.T., and Oechssler, J., (1999), “Learning in Cournot Oligopoly- An Experiment,” *Economic Journal*, 109, C-80-90.
- [31] Huck, S., Normann, H.T., and Oechssler, J., (2004), “Two are Few and Four are Many: Number Effects in Experimental Oligopolies,” *Journal of Economic Behavior and Organization*, 53(4): 435-446.
- [32] Kagel, J. H., and Levin, D., (1986), “The Winner’s Curse and Public Information in Common Value Auctions,” *American Economic Review*, 76(5): 894-920.
- [33] Kagel, J. H., and Levin, D., (2001), “Behavior in Multi-Unit Demand Auctions: Experiments with Uniform Price and Dynamic Vickrey Auctions,” *Econometrica*, 69(2): 413-454.
- [34] Kagel, J. H. and Levin, D., (forthcoming) “Auctions: A Survey of Experimental Research, 1995 - 2008,” in *The Handbook of Experimental Economics*, Volume II , J. H. Kagel and A. E. Roth (eds). Princeton: Princeton University Press.
- [35] Keloharju, M., Nyborg, K. G., and Rydqvist, K., (2005), “Strategic Behavior and Underpricing in Uniform Price Auctions: Evidence from Finnish Treasury Auctions,” *The Journal of Finance*, 60(4), 1865-1902.
- [36] Klemperer, P., and Meyer, M., (1989), “Supply Function Equilibria in Oligopoly under Uncertainty,” *Econometrica*, 57(6): 1243–1277.
- [37] Koch, C., and Penczynski, S., (2015), “The Winner’s Curse: Conditional Reasoning and Belief Formation,” *Working paper SSRN* No. 2563438.
- [38] Levin, D., (2005), “Demand Reduction in Multi-Unit Auctions: Evidence from a Sportscard Field Experiment: Comment,” *The American Economic Review*, 95(1): 467-471.
- [39] Levy, G., and Razin, R., (2015), “Correlation Neglect, Voting Behavior, and Information Aggregation,” *American Economic Review*, 105(4), 1634-1645.

- [40] List, J.A., and Lucking-Reiley, D., (2000), "Demand Reduction in Multi-Unit Auctions Evidence from a Sportcard Field Experiment," *American Economic Review*, 90(4): 961-972.
- [41] Lopomo, G., Marx, L. M., McAdams, D., and Murray, B., (2011), "Carbon Allowance Auction Design: An assessment of Options for the United States," *Review of Environmental Economics and Policy*, 5(1), 25-43.
- [42] Nagel, R., (1995), "Unraveling in Guessing Games: An Experimental Study," *American Economic Review*, 85(5): 1313-1326.
- [43] Nguyen, M., (2010), "XTSUR: Stata Module to Estimate Seemingly Unrelated Regression Model on Unbalanced Panel Data," *Statistical Software Component No. S456953*, Boston College Department of Economics.
- [44] Offerman, T., Potters, J., and Sonnemans, J., (2002), "Imitation and Belief Learning in an Oligopoly Experiment," *The Review of Economic Studies*, 69(4): 973-997.
- [45] Sade, O., Schnitzlein, C., and Zender, J.F., (2006), "Competition and Cooperation in Divisible Good Auction: An Experimental Examination," *The Review of Financial Studies*, 19(1): 195-235.
- [46] Samuelson, W. F., and Bazerman, M. H., (1985), "Negotiation under the Winner's Curse," *Research in Experimental Economics*, 3, 105-38.
- [47] Selten, R., and Buchta, J., (1999), "Experimental Sealed Bid First Price Auctions with Directly Observed Bid Functions," *In Games and Human Behavior*, Mahwah NJ, L. Erlbaum Associates Inc., 79-102.
- [48] Vives, X., (2008), *Information and Learning in Markets: The Impact of Market Microstructure*. Princeton University Press.
- [49] Vives, X., (2011), "Strategic Supply Function Competition with Private Information," *Econometrica*, 76(3): 1919-1966.
- [50] Weizsacker, G., (2010), "Do We Follow Others when We Should? A Simple Test of Rational Expectations," *American Economic Review*, 100, 2340-2360.

Appendices

Appendix A: Theoretical considerations

This section uses results from Vives (2011). When demand is inelastic and equal q and if $\frac{-1}{n-1} < \rho < 1$, $\frac{\sigma_\varepsilon^2}{\sigma_\theta^2} < \infty$ and $\lambda > 0$, the model described in Section 3 has a unique linear Supply Function Equilibrium (SFE) given as

$$X(s_i, p) = b - as_i + cp, \quad (5)$$

where

$$a = \frac{(1 - \rho)\sigma_\theta^2}{(1 - \rho)\sigma_\theta^2 + \sigma_\varepsilon^2} (d + \lambda)^{-1}, \quad (6)$$

$$b = \frac{1}{1 + M} \left(\frac{qM}{n} - \frac{\sigma_\varepsilon^2 \bar{\theta} (d + \lambda)^{-1}}{(1 + (n - 1)\rho)\sigma_\theta^2 + \sigma_\varepsilon^2} \right), \quad (7)$$

and

$$c = \frac{n - 2 - M}{\lambda(n - 1)(1 + M)}, \quad (8)$$

where $M = \frac{\rho n \sigma_\varepsilon^2}{(1 - \rho)((1 + (n - 1)\rho)\sigma_\theta^2 + \sigma_\varepsilon^2)}$ represents an index of adverse selection and $d = \frac{1}{(n - 1)c}$ is the slope of the inverse residual demand. The expected intercept is equal to $E[f] = b - a\bar{\theta}$.

The expected market price is equal to

$$E[p] = \bar{\theta} + \frac{(d + \lambda)q}{n}. \quad (9)$$

Ex-ante expected profits of seller i at the SFE given the predicted values with full information are equal to:⁴⁴

$$E[\tilde{\pi}(t; d)] = (d + \frac{\lambda}{2}) \left(\frac{q^2}{n^2} + \frac{(1 - \rho)^2 (n - 1) \sigma_\theta^4}{n(\sigma_\varepsilon^2 + \sigma_\theta^2(1 - \rho))} \frac{1}{(d + \lambda)^2} \right), \quad (10)$$

where $t = (E[\theta_1 | s], E[\theta_2 | s], \dots, E[\theta_n | s])$, $s = (s_1, s_2, \dots, s_n)$ and $\tilde{\pi}(t; d) = \frac{1}{n} \sum_i \pi_i(t; d)$. We note that $x_i(t; d) = \tilde{x}(t; d) + \frac{\tilde{t} - t_i}{(d + \lambda)}$. Therefore, the first term of expected profits corresponds to expected profits at the average quantity since they are equal to $(\bar{p} - \bar{\theta})\frac{q}{n} - \frac{\lambda}{2}(\frac{q}{n})^2 = (d + \frac{\lambda}{2})\frac{q^2}{n^2}$. The second term is related to the dispersion of the predicted values. For each subject and round, we compute profits conditional on the private signal. We then calculate the average profits in

⁴⁴At the SFE the market price and the signal provide sufficient information on the joint information of the market.

each treatment, which give us an estimate for ex-ante expected profits.

In terms of the efficiency at the equilibrium allocation, we note the following. First, there is no aggregate inefficiency since the demand is inelastic. Second, the ex-ante expected deadweight loss, $E[DWL]$, at the equilibrium allocation is the difference between expected total surplus at the efficient allocation and at the equilibrium allocation. We can write it as

$$E[DWL] = \frac{n\lambda}{2} E[(x_i - x_i^o)^2]. \quad (11)$$

Vives (2011) shows that the efficient allocation is equal to

$$x_i^o = \frac{q}{n} + \frac{(1-\rho)\sigma_\theta^2}{\lambda((1-\rho)\sigma_\theta^2 + \sigma_\varepsilon^2)} (\tilde{s} - s_i), \quad (12)$$

and that expected deadweight loss at the equilibrium allocation can be then written as

$$E[DWL] = \frac{\lambda}{2} \left(\frac{1}{\lambda} - \frac{1}{\lambda + d} \right)^2 \frac{(1-\rho)^2 (n-1) \sigma_\theta^4}{(\sigma_\varepsilon^2 + \sigma_\theta^2 (1-\rho))}. \quad (13)$$

For each market we compute the empirical counterpart of the (interim) deadweight loss at the experimental allocation as follows

$$dwl = \left(\frac{n\lambda}{2} \right) \left(\frac{1}{n} \sum_{i=1}^n (x_i^e - x_i^o)^2 \right), \quad (14)$$

where x_i^e is the experimental allocation (superscript e is to differentiate it from the equilibrium allocation, x_i) and x_i^o is the efficient allocation as defined above. In order to calculate the average deadweight loss in each treatment, we average the interim deadweight losses for the 600 markets in each treatment, which gives us an estimate for the ex-ante expected deadweight loss.⁴⁵

Appendix B: Equilibrium predictions, ex-ante expected outcomes, and statistical distribution of random costs and errors used in the experiment

In our experiment, given the experimental parameters of Table 2, the numerical equilibrium supply function and ex-ante expected outcomes in the two treatments can be summarised in the table below.

⁴⁵There are 600 markets in each treatment since we have 6 groups, which consist of 4 markets each for 25 rounds.

Table 11: *Equilibrium supply function & ex-ante expected outcomes given the experimental parameters.*

	Uncorrelated costs treatment ($\rho=0$)	Positively correlated costs treatment ($\rho=0.6$)
Supply Function		
a	0.12	0.032
b	-44.12	7.65
c	0.17	0.037
Supply function (PQ)		
Intercept	1,000.00	655.32
Slope	6.00	26.68
Outcomes		
Expected market price	1,200.00	1,544.68
Ex-ante expected profit at the SFE	5,612.75	16,567.37
Efficiency		
Ex-ante expected inefficiency at the SFE	612.75	467.73

Notes. The numerical value for the intercept corresponds to the expected intercept in each treatment. Supply function (PQ) refers to the supply function viewed in the usual (Quantity, Ask Price) space. SFE refers to supply function equilibrium.

The next table reports the summary statistics of the draws of the random variables used in the experiment. The theoretical model assumes that in each market we draw an n-dimensional multivariate normal distribution with $\boldsymbol{\theta} \sim N(\boldsymbol{\mu}, \boldsymbol{\Sigma}_{\boldsymbol{\theta}})$, where the covariance matrix $\boldsymbol{\Sigma}_{\boldsymbol{\theta}}$ is such that the variance of each component, θ_i , is equal to σ_{θ}^2 and the covariance between any $i \neq j$ is equal to $\rho\sigma_{\theta}^2$, where ρ is the correlation coefficient. In Table 2 we have specified the following parameter configuration: $n = 3$, $\bar{\theta} = 1,000$, $\sigma_{\theta}^2 = 10,000$; $\rho = 0$ for the uncorrelated costs treatment and $\rho = 0.6$ in the positively correlated costs treatment. The first part of the table below reports the statistics of the actual draws for random costs for each of the three sellers, where subscripts 1, 2 and 3 correspond to seller 1, seller 2 and seller 3, respectively. Because of our random matching protocol, in each round we have a different subject to take the role of sellers 1, 2 and 3. The model also assumes that in each market we draw an n-dimensional multivariate normal distribution of signals such that $\mathbf{s} = \boldsymbol{\theta} + \boldsymbol{\varepsilon}$, where the signals' errors are distributed $\boldsymbol{\varepsilon} \sim N(\mathbf{0}, \boldsymbol{\Sigma}_{\boldsymbol{\varepsilon}})$. The variance of each component, ε_i , is equal to σ_{ε}^2 and the covariance between any $i \neq j$ is equal to zero. In Table 2, we have specified that $\sigma_{\varepsilon}^2 = 3,600$. The second part of the table below reports the statistics of the actual draws for the signals' errors for each of the three sellers. The theoretical model also specifies that error terms in the signals are uncorrelated with the θ_i 's. We check this in the third part of the table below. There are no statistically significant differences between the draws used in the experiment and the theoretically specified distributions.

Table 12: *Statistical Distribution of Random Costs and Errors used in the Experiment.*

	Uncorrelated costs treatment ($\rho=0$)	Positively correlated costs treatment ($\rho=0.6$)
Distribution of Random Costs θ_i for $i=1,2,3$		
Mean ($\theta_1 / \theta_2 / \theta_3$)	997.06/ 998.76/ 998.57	1,000.89/1,000.15/1,005.22
Variance ($\theta_1 / \theta_2 / \theta_3$)	10,484.45/9,927.97/11,476.69	9,567.90/9,710.48/9621.23
Correlations (θ_1, θ_2) / (θ_1, θ_3) / (θ_2, θ_3)	-0.04/-0.02/0.01	0.60/0.61/0.59
Distribution of errors ϵ_i for $i=1,2,3$		
Mean ($\epsilon_1 / \epsilon_2 / \epsilon_3$)	-2.30/-3.09/-2.56	2.59/2.86/-2.30
Variance ($\epsilon_1 / \epsilon_2 / \epsilon_3$)	3,875.56/3,146.13/3,780.09	3,517.05/3,814.36/3,571.94
Correlations (ϵ_1, ϵ_2) / (ϵ_1, ϵ_3) / (ϵ_2, ϵ_3)	-0.05/-0.03/-0.04	0.01/-0.006/0.04
Correlations errors and signals		
Correlations (ϵ_1, θ_1) / (ϵ_2, θ_2) / (ϵ_3, θ_3)	0.05/-0.03/0.01	-0.01/-0.04/-0.04
Correlations (ϵ_1, θ_2) / (ϵ_1, θ_3) / (ϵ_2, θ_3)	-0.01/0.005/-0.03	-0.004/-0.03/-0.06
Number of draws	600/600/600	600/600/600

Appendix C: Instructions for the experiment

These instructions are for the treatment with positively correlated costs and have been translated from Spanish (except from figures, which are exactly as presented to participants).

INSTRUCTIONS

You are about to participate in an economic experiment. Your profits depend on your decisions and on the decisions of other participants. Read the instructions carefully. You can click on the links at the bottom of each page to move forward or backward. Before starting the experiment, we will give a summary of the instructions and there will be two trial rounds.

THE EXPERIMENT

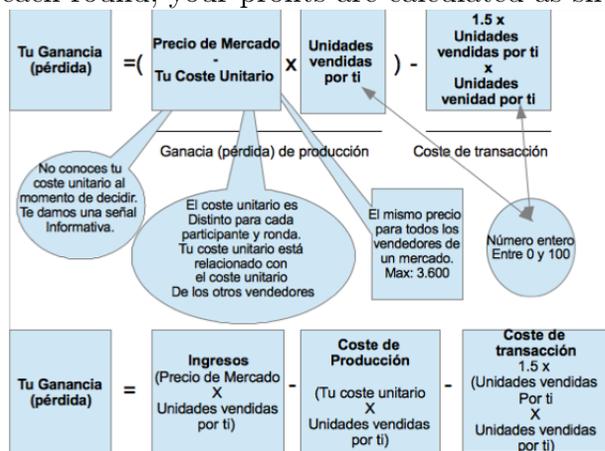
You will earn 10 Euros for participating in the experiment regardless of your performance in the game. You will gain (lose) points during the experiment. At the end of the experiment, points are exchanged for euros. 10,000 points are equivalent to 1 Euro. Each player will start with an initial capital of 50,000 points. Gains (losses) that you accumulate during the experiment will be added (subtracted) to the initial capital. Players who have accumulated losses at the end of the experiment will receive 10 Euros for participating. Players with gains will receive their gains converted to Euros plus the 10 Euro participation fee.

The experiment will last 25 rounds. In the experiment you will participate in a market. You will be a seller of a fictitious good. Each market will have 3 sellers. Market participants will change randomly from round to round. At any given time, no one knows who she is matched with. We guarantee anonymity. The buying decisions will be made by the computer and not

a participant of the experiment. In each round and market, the computer will buy exactly 100 units of the good.

YOUR PROFITS

In each round, your profits are calculated as shown in the figure below:⁴⁶



Your profits are equal to the income you receive from selling units minus total costs (consisting of production and transaction costs).

Some details to keep in mind: you only pay the total costs of the units that you sell. If you sell zero units in a round, your profits will also be zero in this round. You can make losses when your income is less than the total costs (production and transaction). The cumulative profits are the sum of the profits (losses) on each round. Losses will be deducted from the accumulated profits. Throughout the experiment, a window in the upper left corner of your screen will show the current round and accumulated profits.

YOUR DECISION

In each round, you have to decide the minimum price that you are willing to sell each unit for. We call these Ask Prices.

THE MARKET PRICE

Once the three sellers in a given market have entered and confirmed their decisions, the computer calculates the market price as follows.

1. In each market, the computer observes the 300 Ask Prices introduced by the sellers of your market.
2. The computer ranks the 300 Ask Prices from the lowest to the highest.

⁴⁶English explanation: the top part of the figure explains the formula for profits given in equation (1). The bottom part of the figure expresses $\text{Profits} = \text{Revenues} - \text{Production Costs} - \text{Transaction Costs}$.

3. The computer starts buying the cheapest unit, then it buys the next unit, etc. until it has purchased exactly 100 units. At this time the computer stops.

4. The Ask Price of the 100th unit purchased by the computer is the market price (the price of the last unit purchased by the computer).

The market price *is the same for all units* sold in a market. In other words, a seller receives a payment, which is equal to the market price for each unit she sells. If more than one unit is offered at the market price, the computer calculates the difference:

Units Remaining = 100 - Units that are offered at prices below the market price.

The Units Remaining are then split proportionally among the sellers that have offered them at an Ask Price equal to the market price.

UNITS SOLD

In each round and market, the three sellers offer a total of 300 units. The computer purchases the 100 cheapest units. Each seller sells those units that are offered at lower Ask Prices than the market price. Note that those units that are offered at higher Ask Prices than the market price are not sold. Those units offered at an Ask Price which is equal to the market price will be divided proportionally among the sellers that have offered them.

MARKET RULES

In each round and market, the computer buys exactly 100 units of the good at a price not exceeding 3,600. In order to simplify the task of entering all Ask Prices in each round, we request that you to enter:

- Ask Price for Unit 1
- Ask Price for Unit 2

Ask Prices can be different for different units. To find Ask Prices for the other units, we will join the Ask Price for Unit 1 and the Ask Price for Unit 2 by a straight line. In this way, we find the Ask Prices for all the 100 units. In the experiment, you will be able to see this graphically and try different values until you are satisfied with your decision.

We apply the following five market rules.

1. You must offer all the 100 units for sale.
2. Your Ask Price for one unit must always be greater than or equal to the Ask Price of the previous unit. Therefore, the Ask Price for the second unit cannot be less than the Ask Price for the first unit. You can only enter integers for your decisions.
3. Both Ask Prices must be zero or positive.
4. The buyer will not purchase any unit at a price above the price cap of 3,600.

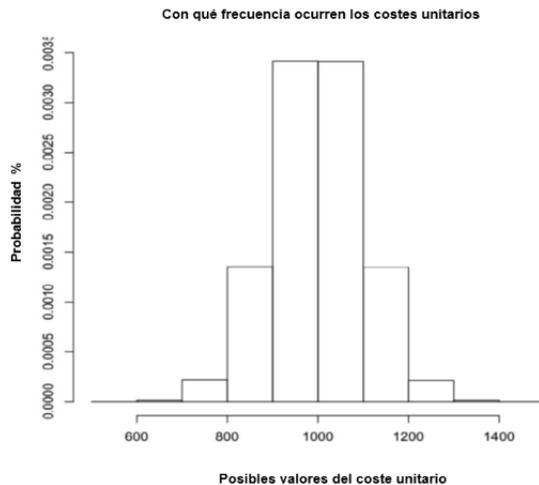
5. The Ask Price for some units may be lower your unit cost, since unit costs are unknown at the time when you decide the Ask Prices. You may have losses.

EXAMPLE

This example is illustrative and irrelevant to the experiment itself. We give the example on paper. Here you can see how the computer determines the market price and units sold by each seller in a market.

UNIT COST

In each round the unit cost is random and unknown to you at the time of the decision. The unit cost is independent of previous and future round. Your unit cost is different from the unit cost of other participants. However, your unit cost is related to the unit costs of the other market participants. Below we explain how unit costs are related and we give a figure and explanation of the possible values of unit costs and their associated frequencies. This figure is the same for all sellers and all round.⁴⁷



The horizontal displays the unit cost while the vertical axis shows the frequency with which each unit cost occurs (probability). This frequency is indicated by the length of the corresponding bar.

In the figure you can see that the most frequent unit cost is 1,000. We obtain 1,000 as unit cost with a frequency of 0.35%. In general terms, we would obtain a unit cost of 1,000 in 35 of 1,000 cases.

In 50% of the cases (50 of 100 cases), the unit cost will be between 933 and 1,067.

In 75% of the cases (75 of 100 cases), the unit cost will be between 885 and 1,115 .

⁴⁷ *English explanation:* the figure shows the possible values of costs θ_i (horizontal axis) and their corresponding probabilities (vertical axis).

In 95% of the cases (95 of 100 cases), the unit cost will be between 804 and 1,196 .

There is a very small chance that the unit cost is less than 700. This can occur in 1 of 1,000 cases approximately. Similarly, there is a very small chance that the unit cost is greater than 1,300. This occurs can occur in 1 of 1,000 cases, approximately.

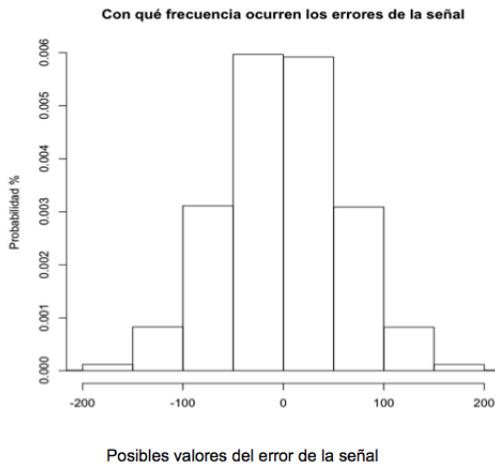
For participants with knowledge of statistics: the unit cost is normally distributed with mean 1,000 and standard deviation 100.

INFORMATION ABOUT YOUR UNIT COST (YOUR SIGNAL)

In each round, each participant receives information on her unit costs. This information is not fully precise. The signal that you receive is equal to:

$$Signal = UnitCost + Error$$

The error is independent of your unit cost, it is also independent from the unit costs of other participants and it is independent from past and future errors. The following figure describes the possible values of the error term and an indication of how likely each error is likely to occur. This graph is the same for all sellers and rounds.⁴⁸



On the horizontal axis you can observe the possible values of the error terms. On the vertical axis, you can observe the frequency with which each error occurs (probability). This frequency is indicated by the length of the corresponding bar.

In the figure you can see that the most common error is 0. The frequency of error 0 is 0.66%. In general terms, this means that in approximately 66 of 10,000 cases you would get an error equal to 0.

In 50% of the cases (50 of 100 cases), the error term is between -40 and 40.

In 75% of the cases (75 of 100 cases), the error is between -69 and 69.

⁴⁸ *English explanation:* the figure shows the possible values of the signal's error (horizontal axis) and their corresponding probabilities (vertical axis).

In 95% of the cases (95 of 100 cases), the error is between -118 and 118.

There is a very small chance that the error is less than -200. This occurs in 4 out of 10,000 cases. Similarly, there is a very small probability that the error is greater than 200. This occurs in 4 out of 10,000 cases.

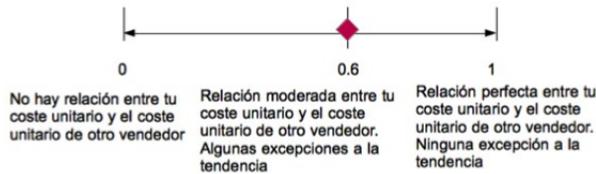
For participants with knowledge of statistics: the error has a normal distribution with mean 0 and standard deviation 60.

HOW YOUR COST IS RELATED TO THE COSTS OF THE OTHER SELLERS

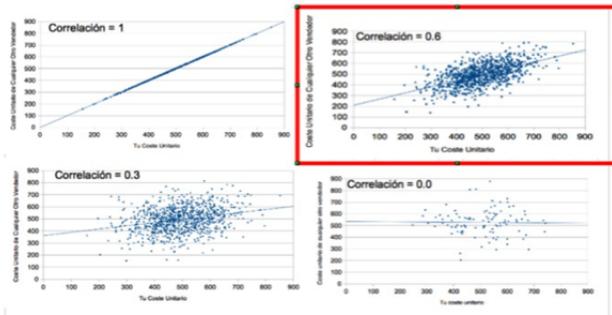
The unit cost is different for each seller and your unit cost is related to the unit cost of the other sellers in your market. The association between your unit cost and unit cost of another seller in your market follows the trend:

- The higher your unit cost, the higher will be the unit cost of the other sellers.
- The lower your unit cost, the lower the unit cost of the other sellers.

The strength of the association between your unit cost and unit cost of another seller is measured on a 0 to 1 scale. The strength of the association between your unit cost and unit cost of the other seller is $+0.6$.⁴⁹



Graphically we can see the relationship between your unit cost (horizontal axis) and the unit cost of another seller (vertical axis) for some strengths of association. The figure that has a red frame corresponds to an intensity of association of $+0.6$.



For participants with knowledge of statistics: the correlation between your unit cost and unit cost of any other player is $+0.6$.

⁴⁹ English explanation: the figure explains the meaning of correlation equal to 0, 0.6 and 1.

END OF ROUND FEEDBACK

At the end of each round, we will give you information about:

- Your profits (losses) and its components (Revenue-Cost of Production - Cost of transaction)
- Market price
- Your units sold
- Other market participants feedback: decisions; profits and unit costs.

You can also check your historical performance in a window in the upper right corner of your screen. During the experiment the computer performs mathematical operations to calculate the market price, units sold, Ask Prices for intermediate units, etc. For these calculation we use all available decimals. However, we show all the variables rounded to whole numbers, except from the market price.

THE END

This brings us to the end of the instructions. You can take your time to re-read the instructions by pressing the BACK button. When you understand the instructions you can indicate it to us by pressing the OK button at the bottom of the screen. Next you have to answer a questionnaire about the instructions, unit cost distributions and signals. When all participants have taken the questionnaire and indicated OK, we will start the practice rounds. Your profits or losses of the practice rounds will not be added or subtracted to your earnings during the experiment.

Appendix D: Comprehension test and experimental screenshots

The first part of this appendix reports the comprehension questionnaire which was administered before the trial rounds.

Comprehension Test

Questions. Answer True or False.

1. The unit cost has the same value for each of the participants in your market.
2. The unit cost has the same value for each of the participants in your market.

3. If my unit cost is high, it is rather likely that the unit cost of another seller is high.
4. Unit costs between 1000 and 1200 occur with the same frequency than unit costs between 1000 and 700.
5. Unit costs larger than 1000 occur with the same frequency as unit costs smaller than 1000.
6. Errors larger than 0 occur more frequently than errors smaller than 0.
7. An error of 5 is the most frequent error.
8. The seller who sells most units will always have the highest profit.
9. If my unit cost is low, it is rather likely that the unit cost of another seller is high.
10. The market price is the same for all units and sellers.

Answers (True (T) and False (F)): Q1. F Q2. T Q3. T (treatment 0.6); F (treatment 0) Q4. F Q5. T Q6. F Q7. F Q8. F Q9. F Q10. T

Notes: These notes appeared on the screen if a participant answered wrongly any of the previous questions.

Q1. *Treatment 0.6:* Your unit cost is different from the unit cost of other participants but it is related. *Treatment 0:* Your unit cost is different from the unit cost of other participants. There is no relation between your unit cost and that of other participants.

Q2. In each round, the unit cost is random and independent from the unit cost of past and future rounds.

Q3. *Treatment 0.6:* The higher your unit cost, the higher the unit cost of the other sellers will tend to be. *Treatment 0:* There is no relation between your unit cost and that of other participants. Therefore, if my unit cost is high, I can not deduce anything from the unit cost of the other participants.

Q4. Unit costs between 1000 and 1200 occur with higher frequency than unit costs between 1000 and 700.

Q5. The unit cost of 1000 is the most frequent one. Unit costs larger than 1000 occur with the same frequency as unit costs smaller than 1000.

Q6. Errors larger than 0 occur with the same frequency as errors smaller than 0.

Q7. An error of 0 is the most frequent error.

Q8. Profit does not only depend on the number of units sold. Remember that: $Profit = (MarketPrice - UnitCost)UnitsSold - 1.5UnitsSold^2$.

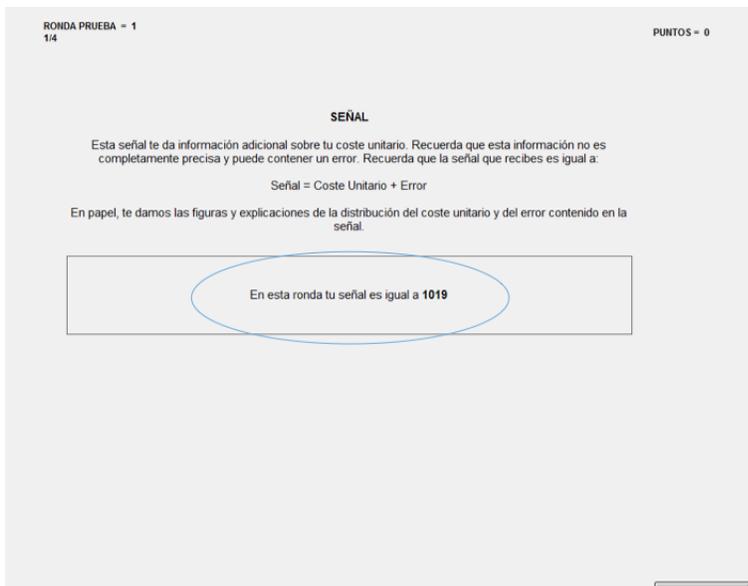
Q9. *Treatment 0.6:* The lower your unit cost, the lower the unit cost of the other sellers will tend to be. *Treatment 0:* There is no relation between your unit cost and that of other

participants. Therefore, if my unit cost is high, I can not deduce anything about the unit cost of the other participants.

Q10. The market price is the same for all units and sellers in a market.

The second part of this appendix reports the screenshots used during the experiment.

Screen 1: signal screen⁵⁰



⁵⁰ *English explanation:* the first part of the screen reminds participants of the definition of the signal. The second part gives the participant the realisation of her signal for the current round.

Screen 2: decision screen⁵¹

RONDA = 1
2/4

PUNTOS = 50000

INTRODUCCIÓN Y CONFIRMACIÓN de los PRECIOS DE OFERTA

Recuerda, en esta ronda tu señal es igual a: 1000

Precio de oferta: Precio mínimo al que estás dispuesto a vender una unidad.

Precio de Oferta de la Unidad 1:

Precio de Oferta de la Unidad 2:

Representación gráfica de los Precios de Oferta para todas las unidades entre 0 y 100.

Puedes probar diferentes Precios de Oferta hasta que estés satisfecho con tu decisión. Cuando estés satisfecho, deberás confirmar tu elección. El experimento no continuará hasta que todos los vendedores hayan confirmado su decisión.

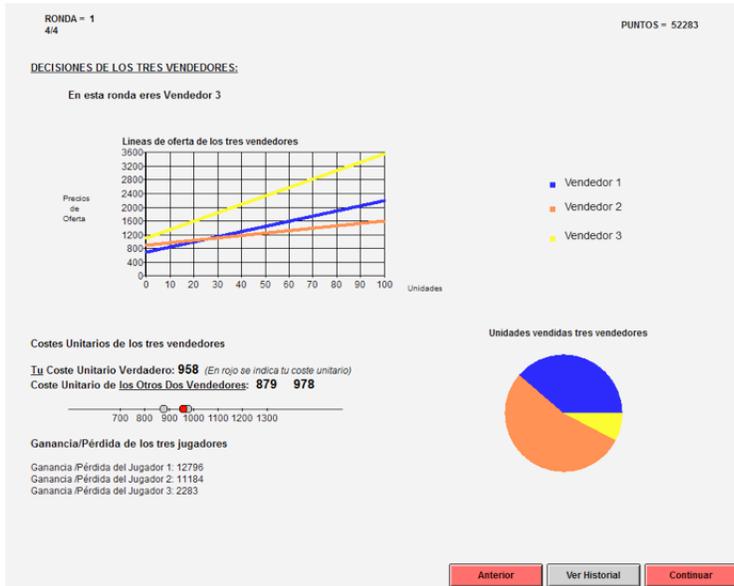
Screen 3: feedback about a seller's own performance⁵²



⁵¹English explanation: First, the screen reminds the participant of the realisation of her signal for the current round. Second, it asks the participant to enter the *ask price* for the first unit and the *ask price* for the second unit. The participant has access to a calculator and can see a graphical representation of the supply schedule. The participant can try a few decisions before it presses the “Confirm” button.

⁵²English explanation: this screen reports the participants’ profits/losses in the current round and splits them into revenues (price multiplied by units sold), production costs (cost multiplied by units sold) and transaction costs (1.5 multiplied by the square of units sold). The screen displays them both graphically and numerically.

Screen 4: feedback about market performance and other sellers' in the same market⁵³



Appendix E: Post-experiment questionnaire

After the rounds were completed, we asked for 3 demographic questions: *age*, *gender* and *degree studying*.

We then asked the following additional questions regarding understanding of the game.

1. Do you think that a high market price generally means GOOD/MIXED/BAD news about the level of your costs?
2. Explain your answer.
3. Do you think that the other sellers have answered the same as you to the previous question?
4. Explain your answer.

⁵³ *English explanation*: this screen gives the participant feedback about: own supply function and supply functions of the rivals (graphically); units sold by each of the players; the own value of θ_i and that of the two rivals; profits/losses for each of the players in the market.

Appendix F: Analysis of variance of experimental outcomes and decomposition of deadweight loss

The following table shows the analysis of variance for experimental outcomes and deadweight losses in each treatment.

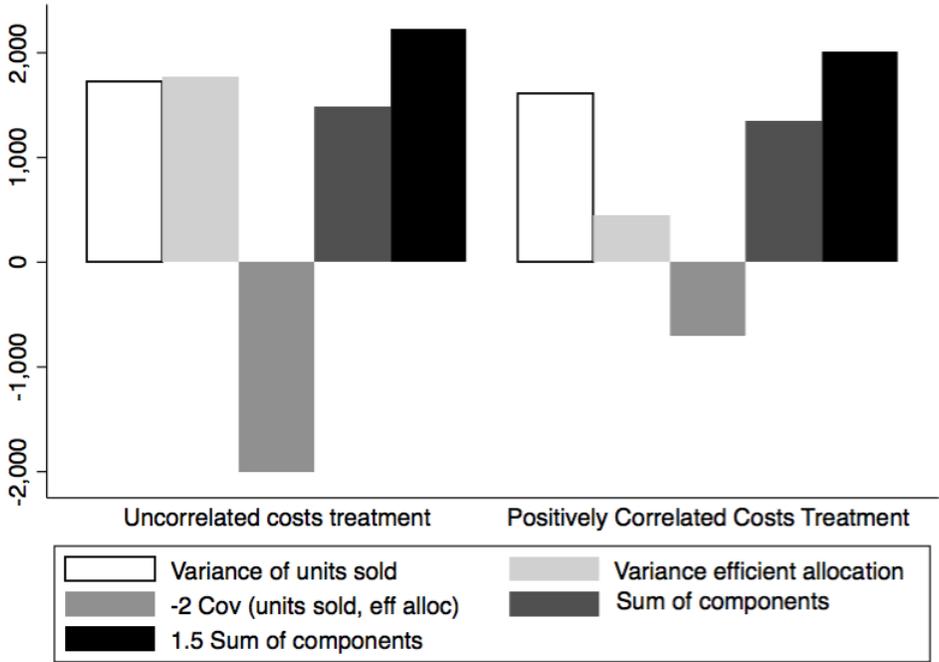
Table 13: *Analysis of variance of experimental outcomes and deadweight losses.*

Variable	Mean	s.d. Overall Subject level ¹	s.d. Between Subjects ¹	s.d. Within Subjects (period)	s.d. Overall Group level	s.d. Between Groups	s.d. Within Groups (period)	s.d. Overall Treatment level (period)
Observations / Treatment		1800	72	25	150	6	25	25
Uncorrelated Costs Treatment								
<i>Outcomes</i>								
Market Price	1,110.38	117.36	25.38	114.69	55.16	15.76	53.24	16.45
Profits	2,255.35	7,191.77	2,150.21	6,867.30	2,043.79	683.89	1,945.43	688.66
<i>Efficiency</i>								
Deadweight loss	2,222.51	3,142.66	619.20	3,083.55	1,667.18	472.55	1,610.01	768.85
Positively Correlated Costs Treatment								
<i>Outcomes</i>								
Market Price	1,123.73	130.27	40.46	124.09	62.83	34.70	54.20	20.20
Profits	1,877.88	5,816.32	2,679.11	5,171.82	2,004.70	1,140.88	1,710.81	701.28
<i>Efficiency</i>								
Deadweight loss	2,011.01	2,290.29	578.85	2,218.96	1,166.74	534.53	1,059.05	585.88

Note. s.d. refers to standard deviation. The explanation of the variables reported in the table is as follows: s.d. Overall subject level: the variation of individual choice across all rounds; s.d. Between subjects: the variation of the individual choice averaged across all rounds; s.d. Within Subjects: the variation within subject over time; s.d. Overall group level: the variation the average choice of the group across all rounds; s.d. Between groups: the variation of the average choice of the group averaged across all rounds; s.d. Within groups: the variation within groups over time; and s.d. Overall Treatment level: the variation of the average choice of the treatment across all rounds.

The following graph illustrates how the components of deadweight loss combine in each treatment.

Figure 10: *Components of deadweight loss at the experimental allocation in each treatment.*



Notes. This graph corresponds to the decomposition of deadweight loss, $dwl = (\frac{\lambda}{2})(\sum_{i=1}^n (x_i^e - \frac{q}{n})^2 + (x_i^o - \frac{q}{n})^2 - 2(x_i^e - \frac{q}{n})(x_i^o - \frac{q}{n}))$. The first term is variance of units sold, the second is the variance of the efficient allocation, and the third is -2 multiplied by the covariance of units sold and the efficient allocation. The bars associated with “Sum of components” refer to the sum $(\sum_{i=1}^n (x_i^e - \frac{q}{n})^2 + (x_i^o - \frac{q}{n})^2 - 2(x_i^e - \frac{q}{n})(x_i^o - \frac{q}{n}))$. *dwl* refers to deadweight loss.

The next table displays the components of deadweight loss at both the equilibrium and experimental allocations in each treatment.

Table 14: *Components of Deadweight Loss at the equilibrium and experimental allocations in each treatment.*

			Uncorrelated Costs Treatment	Positively Correlated Costs Treatment
	Variable	Number of observations/ Treatment	Mean	Mean
(1)	Variance of Efficient Alloc.	600	1,767.69 (1,706.84)	446.28 (453.06)
(2)	Variance of equilibrium alloc.	600	441.92 (426.71)	15.04 (15.27)
(3)	Covariance (equilibrium alloc., efficient alloc.)	600	883.84 (853.42)	81.94 (83.18)
(4)	Sum of components (1)+(2)-2*(3)	600	441.92 (426.71)	297.45 (301.97)
(5)	dwl equilibrium alloc.	600	662.88 (640.06)	446.17 (452.95)
(6)	Variance of experimental alloc.	600	1,724.29 (1,647.34)	1,605.65 (1,612.68)
(7)	Covariance (experimental alloc., efficient alloc.)	600	1,005.15 (1,343.87)	355.63 (637.39)
(8)	Sum of components (1)+(6)-2*(7)	600	1,481.67 (2,095.11)	1,340.67 (1,526.86)
(9)	dwl experimental alloc.	600	2,222.51 (3,142.66)	2,011.01 (2,290.29)

Notes. Alloc. refers to allocation. In order to calculate *dwl* in rows (5) and (9), we have used the actual distribution of draws of the random variables used in the experiment.

Appendix G: Panel data analysis

In this section we present the results of panel data analysis, which is based on the individual choice or outcome across rounds. For that purpose, we use a random effects panel data approach.⁵⁴

First, we evaluate whether there is a difference—between the positively correlated costs treatment and the (control) uncorrelated costs treatment—with regard to either market behaviour or outcomes. Table 15 reports random effects regression results of jointly estimating, for our panel data, the equation for the supply function slope and the intercept. The so-called seemingly unrelated regression (SUR) method we use, which assumes that disturbances across equations may be correlated, yields efficient estimates.⁵⁵ In Table 15[1], we jointly estimate the

⁵⁴For market price and deadweight loss, the unit of analysis is the market across rounds.

⁵⁵See Baltagi (2008) for the theoretical background on this estimation.

following set of equations:

$$\begin{aligned}
 SlopePQ_{it} &= \beta_{0S} + \beta_{1S}D_Treatment_i + \sum_j \beta_j D_Group_{ij} + \nu_i + u_{it}; \\
 InterceptPQ_{it} &= \beta_{0I} + \beta_{1I}D_Treatment_i + \sum_j \beta_j D_Group_j + w_i + v_{it}.
 \end{aligned}$$

Here the subscripts S and I signify (respectively) slope and intercept, the subscript it denotes the choice of subject i at round t , and j is an index for the group to which a subject belongs. The term $D_Treatment$ is an indicator variable that takes the value 0 (resp., 1) for the uncorrelated (resp., positively correlated) costs treatment, and the D_Group_{ij} are a set of dummies that take the value 1 if subject i belongs to group j (and value 0 otherwise). The terms u_{it} and v_{it} are observation-specific errors; ν_i and w_i are unobserved individual effects.

In Table 15[2] we report results from the following set of augmented regressions:

$$\begin{aligned}
 SlopePQ_{it} &= \beta_{0S} + \beta_{1S}D_Treatment_i + \sum_j D_Group_{ij} + \beta_{2S}t \\
 &\quad + \beta_{3S}(D_Treatment_i \times t) + \beta_{4S}Signal_{it} + \nu_i + u_{it}; \\
 InterceptPQ_{it} &= \beta_{0I} + \beta_{1I}D_Treatment_i + \sum_j D_Group_{ij} + \beta_{2I}t \\
 &\quad + \beta_{3I}(D_Treatment_i \times t) + \beta_{4I}Signal_{it} + w_i + v_{it}.
 \end{aligned}$$

In these expressions, t is a variable for the round number, $D_Treatment \times t$ is the interaction term between the round number and treatment, and $Signal$ is the signal received by subject i at round t . The errors u_{it} and v_{it} are correlated across equations.

Table 15: *Panel data: SUR of random effects.*

	Dependent variables		Dependent variables	
	SlopePQ	InterceptPQ	SlopePQ	InterceptPQ
	(1)		(2)	
Constant	4.29*** (1.17)	869.19*** (48.16)	6.38*** (1.34)	207.08*** (52.40)
D_Treatment	0.52 (1.65)	51.44 (68.11)	-0.04 (1.67)	57.35 (67.75)
Round	-	-	-0.096*** (0.014)	2.17*** (0.48)
D_Treatment * Round	-	-	0.050*** (0.02)	-2.19*** (0.68)
Signal	-	-	0.0002 (0.0006)	0.74*** (0.021)
Number of observations	3,600	3,600	3,600	3,600
Group Dummies	Yes	Yes	Yes	Yes

Note. Each equation was estimated with group dummies; standard errors are given in parentheses. This estimation was performed via the *xtsur* command in Stata (cf. Nguyen 2010), which is used for one-way random effects estimation of seemingly unrelated regressions (SURs) in a panel data set. Note that the *xtsur* command does not allow for *clustered* standard errors.

*, **, and *** denote significance at (respectively) the 10%, 5%, and 1% levels.

In both Table 15[1] and Table 15[2] we see that $D_Treatment$ is not statistically significant regardless of whether it is the only regressor or whether we also control for additional variables. With regard to the regression for the supply function slope, from Table 15[2] we can draw the following conclusions: round has a significant and negative coefficient, which means that supply functions become flatter as the number of rounds increases; and $D_Treatment \times t$ has a significant and positive coefficient, which means that the decrease in the slope of the supply function is less pronounced in the positively correlated costs treatment. We also find that *Signal* is not statistically significant in determining the supply function slope. With regard to the regression for the supply function intercept, we may conclude as follows: round has a significant and positive coefficient, which means that the supply function intercept increases as the number of rounds increases; $D_Treatment \times t$ has a significant and negative coefficient that is nearly identical to the coefficient for round, which means that the supply function intercept does not vary as a function of round in the positively correlated costs treatment; and *Signal*, as predicted by the theoretical model, has a significant and positive coefficient.

Second, we estimate a panel with random effects and where standard errors are clustered at the group level for the dependent variables (market prices, profits, and deadweight losses). In Table 16[1] we estimate the following equation for market price:

$$MarketPrice_{mt} = \beta_0 + \beta_1 D_Treatment_t + \nu_m + \omega_{mt}.$$

Here the subscript mt denotes market m at round t ; ϵ_{mt} is the error term; ν_m is the random

effect, which is uncorrelated with the regressor; and $D_Treatment$ is the treatment dummy defined previously. In Table 16[2], we augment the regression with additional controls and estimate the following equation:

$$MarketPrice_{mt} = \beta_0 + \beta_1 D_Treatment_m + \beta_2 t + \beta_3 (D_Treatment_m \times t) + \beta_4 Signal_{mt} + \nu_m + \omega_{mt},$$

where the unit of observation is the market across time and where all the variables are as defined before.

The results reported in Table 16[1] show that $D_Treatment$ is not statistically significant—either when the treatment dummy is used as the only regressor or when we include the additional controls. In Table 16[2] we observe that, with respect to determining market price, the signal received is the only variable with a positive and significant coefficient.

Table 16: *Panel data: Regression of random effects for market outcomes and deadweight loss.*

	Dependent variables					
	Market Price (1)	Market Price (2)	Profits (3)	Profits (4)	dwl (5)	dwl (6)
Constant	1,110.38*** (6.14)	377.31*** (59.20)	2,255.35*** (266.24)	15,115.29*** (2,153.69)	2,222.51*** (184.02)	2,143.46 (1,335.17)
D_Treatment	13.35 (14.84)	10.91 (26.64)	-377.47 (517.84)	30.82 (962.98)	-211.50 (277.83)	-218.54 (445.33)
Round		0.26 (0.72)		39.66* (22.25)		-56.52*** (15.72)
D_Treatment * Round		-0.24 (1.26)		-23.49 (43.01)		0.061 (22.01)
Signal		0.73*** (0.061)		-13.44*** (2.05)		0.82 (1.12)
Number of observations	1,200	1,200	3,600	3,600	1,200	1,200
s.e. Cluster (Group level)	Yes	Yes	Yes	Yes	Yes	Yes

Note. This estimation was performed using the *xtreg, re* command in Stata. Standard errors (s.e.), reported in parentheses, are clustered at the group level.

*, **, and *** denote significance at (respectively) the 10%, 5%, and 1% levels.

Third, for profits we use the same panel data approach as for market price—but now the unit of observation is the subject across rounds.⁵⁶ We find that $D_Treatment$ is not significant in the results of either Table 16[3] or Table 16[4]. The signal received has a negative and significant coefficient, reflecting the negative correlation between signal (and also unit costs) and profits. In addition, we find that profits increase slightly across rounds in the uncorrelated costs treatment; in the positively correlated costs treatment, however, there is no evolution of profits (in either direction) across rounds.

⁵⁶In Table 16[3], we estimate the equation $Profit_{it} = \beta_0 + \beta_1 D_Treatment_i + \nu_i + \omega_{it}$, where the subscript it signifies subject i at round t and where ϵ_{it} is the error term. In Table 16[4], we augment that regression as follows: $Profit_{it} = \beta_0 + \beta_1 D_Treatment_i + \beta_2 t + \beta_3 (D_Treatment_i \times t) + \beta_4 Signal_{it} + \nu_i + \omega_{it}$.

Fourth, for deadweight loss we use the same panel data approach as for market price.⁵⁷ The results presented in Table 16[5] and Table 16[6] show that $D_Treatment$ is not statistically significant. Yet we do find that, in both treatments, deadweight losses decline across rounds; this result indicates that allocations become more efficient as the subjects gain bidding experience.

In short: all panel data results are consistent with the stylised facts and with the tests presented in the Section 5.

Appendix H: Best-response analysis

This section uses the notation of the theoretical background presented in Section 3 and computes a seller's best-response strategy given arbitrary strategies of rivals. We assume that an agent knows the strategies of the rivals, and that she forms correct beliefs about events in the competitive and information environments. The best-response strategy of seller i can be written as $X_i(s_i, p) = b_i - a_i s_i + c_i p$, and the actual strategies of the rivals as $X_j(s_j, p) = b_j - a_j s_j + c_j p$, where $j \neq i$. The rivals' average supply function slope is $c_{-i} = \frac{1}{n-1} \sum_{i \neq j} c_j$, and the rivals' average fixed part of the intercept is $b_{-i} = \frac{1}{n-1} \sum_{j \neq i} b_j$. Furthermore, for simplicity, we assume that all rivals set the same response to private information, i.e. $a_j = a_{-i}$ for all $j \neq i$.

Lemma 1 determines agent i 's best-response strategy given arbitrary strategies of the rivals. The assumptions required for Lemma 1 to hold are consistent with the features of our experimental design.

Lemma 1: Best-response strategy

Assume that $\rho \in [0, 1)$ and that the rivals' average supply function is characterised by (a_{-i}, b_{-i}, c_{-i}) . Suppose that rivals set a supply function such that $c_{-i} > 0$ and that all rivals set the same response to private information such that $a_j = a_{-i} > 0$ for all $j \neq i$. The best-response strategy for seller i is then given as:

$$a_i = \frac{R}{d_i + \lambda + T_i}, \quad (15)$$

$$b_i = \frac{(\bar{\theta}(R + T_i(n-1)a_{-i} - 1) + T_i(q - (n-1)b_{-i}))}{d_i + \lambda + T_i}, \quad (16)$$

⁵⁷In Table 16[5] we estimate the following equation for deadweight loss: $dwl_{mt} = \beta_0 + \beta_1 D_Treatment_m + \nu_m + \omega_{mt}$, where again the subscript mt denotes market m at round t and ϵ_{mt} is the error term. In Table 16[6], we augment the regression with additional and estimate the following equation: $dwl_{mt} = \beta_0 + \beta_1 D_Treatment_m + \beta_2 t + \beta_3 (D_Treatment_m \times t) + \beta_4 Signal_{mt} + \nu_m + \omega_{mt}$.

$$c_i = \frac{1 - T_i(n-1)c_{-i}}{d_i + \lambda + T_i}, \quad (17)$$

where $d_i = \frac{1}{(n-1)c_{-i}}$, $R = \frac{\sigma_\theta^2(\sigma_\theta^2(1-\rho)(1+(n-1)\rho) + \sigma_\epsilon^2)}{(\sigma_\theta^2(1-\rho) + \sigma_\epsilon^2)(\sigma_\theta^2(1+(n-1)\rho) + \sigma_\epsilon^2)}$, and

$$T_i = \frac{\rho\sigma_\theta^2\sigma_\epsilon^2}{(\sigma_\theta^2(1-\rho) + \sigma_\epsilon^2)(\sigma_\theta^2(1+(n-1)\rho) + \sigma_\epsilon^2)a_{-i}}.$$

Proof: Given the strategies of the rivals, seller i maximises her profits and the first order condition is given as: $X_i(s_i, p) = \frac{p - E[\theta | s_i, p]}{(d_i + \lambda)}$. Market clearing implies that $q = \sum_{j=1}^n X(s_j, p)$, and given the definitions of b_{-i} , c_{-i} and the assumption that $a_j = a_{-i}$ for all $j \neq i$, we can re-write the market clearing expression as: $p(n-1)c_{-i} = q - (n-1)b_{-i} + a_{-i} \sum_{i \neq j} s_j - x_i$. Then define $d_i = \frac{1}{(n-1)c_{-i}}$ and $I_i = q - (n-1)b_{-i} + a_{-i} \sum_{i \neq j} s_j$ so that we can write $p = I_i - d_i x_i$. All the information contained in p is also contained in h_i , where $h_i = a_{-i} \sum_{j \neq i} s_j$ and can be shown to be equal to $h_i = (n-1)b_{-i} - q + (n-1)c_{-i}p + x_i$. The second order condition is satisfied if $2d_i + \lambda > 0$, which is always satisfied if $c_{-i} > 0$.

We can now find an expression for $E[\theta_i | s_i, p] = E[\theta_i | s_i, h_i]$. The mean of the vector $\begin{pmatrix} \theta_i \\ s_i \\ h_i \end{pmatrix}$

is equal to $\begin{pmatrix} \bar{\theta} \\ \bar{\theta} \\ a_{-i}(n-1)\bar{\theta} \end{pmatrix}$ and the variance-covariance matrix is:

$$\begin{pmatrix} \sigma_\theta^2 & \sigma_\theta^2 & \rho\sigma_\theta^2 a_{-i}(n-1) \\ \sigma_\theta^2 & \sigma_\theta^2 + \sigma_\epsilon^2 & \rho\sigma_\theta^2 a_{-i}(n-1) \\ \rho\sigma_\theta^2 a_{-i}(n-1) & \rho\sigma_\theta^2 a_{-i}(n-1) & a_{-i}^2(n-1)((\sigma_\epsilon^2 + \sigma_\theta^2) + (n-2)\rho\sigma_\theta^2) \end{pmatrix}$$

Using the expressions for conditional expectations of normally distributed random variables, we obtain:

$$E[\theta_i | s_i, h_i] = \bar{\theta} + R(s_i - \bar{\theta}) + T_i(h_i - (n-1)a_{-i}\bar{\theta}), \quad (18)$$

where

$$R = \frac{\sigma_\theta^2(\sigma_\theta^2(1-\rho)(1+(n-1)\rho) + \sigma_\epsilon^2)}{(\sigma_\theta^2(1-\rho) + \sigma_\epsilon^2)(\sigma_\theta^2(1+(n-1)\rho) + \sigma_\epsilon^2)}, \quad (19)$$

and

$$T_i = \frac{\rho\sigma_\theta^2\sigma_\epsilon^2}{(\sigma_\theta^2(1-\rho) + \sigma_\epsilon^2)(\sigma_\theta^2(1+(n-1)\rho) + \sigma_\epsilon^2)a_{-i}}. \quad (20)$$

In order to obtain the best-response strategy for seller i , we first equate the coefficient on the signal, s_i and obtain: $a_i = \frac{R - T_i a_i}{(d_i + \lambda)}$, or equivalently: $a_i = \frac{R}{d_i + \lambda + T_i}$. Second, we equate

the coefficient on the price and obtain: $c_i = \frac{1-T_i(n-1)c_{-i}-T_i c_i}{(d_i+\lambda)}$. Grouping terms with c_i , we obtain: $c_i = \frac{1-T_i(n-1)c_{-i}}{(d_i+\lambda+T_i)}$. Third, we equate the coefficient on the constant and obtain: $b_i = \frac{\bar{\theta}(R+T_i(n-1)a_{-i}-1)+T_i(q-(n-1)b_{-i})}{(d_i+\lambda+T_i)}$. ■

The next Lemma describes the comparative statics of agent i 's best-response function.

Lemma 2: Strategic incentives

Assume that $\rho \in [0, 1)$ and that the rivals' average supply function is characterised by (a_{-i}, b_{-i}, c_{-i}) . Suppose that rivals bid a supply function such that $c_{-i} > 0$ and that all rivals set the same response to private information such that $a_j = a_{-i} > 0$ for all $j \neq i$.

If $\rho = 0$ then: $\frac{\partial a_i}{\partial a_{-i}} \Big|_{c_{-i}=\text{const}} = 0$, $\frac{\partial b_i}{\partial b_{-i}} \Big|_{a_{-i}=\text{const}, c_{-i}=\text{const}} = 0$ and $\frac{\partial c_i}{\partial c_{-i}} \Big|_{a_{-i}=\text{const}} > 0$.

If $\rho > 0$ then: $\frac{\partial a_i}{\partial a_{-i}} \Big|_{c_{-i}=\text{const}} > 0$, $\frac{\partial b_i}{\partial b_{-i}} \Big|_{a_{-i}=\text{const}, c_{-i}=\text{const}} < 0$,

$\frac{\partial c_i}{\partial c_{-i}} \Big|_{a_{-i}=\text{const}} > 0$ for $0 < c_{-i} < c_{-i}^*$, where c_{-i}^* is a positive number

and $\frac{\partial c_i}{\partial c_{-i}} \Big|_{a_{-i}=\text{const}} < 0$ for $c_{-i} > c_{-i}^*$.

Proof: We first evaluate:

$$\frac{\partial c_i}{\partial c_{-i}} \Big|_{a_{-i}=\text{const}} = \frac{-T_i(n-1)^2(\lambda+T_i)(c_{-i})^2 - 2T_i(n-1)c_{-i} + 1}{(n-1)(c_{-i})^2(d_i+\lambda+T_i)^2}. \quad (21)$$

since $\frac{\partial d_i}{\partial c_{-i}} = \frac{-1}{(n-1)(c_{-i})^2}$. When $\rho = 0$ then $T_i = 0$ and $\frac{\partial c_i}{\partial c_{-i}} \Big|_{a_{-i}=\text{const}} > 0$. When $\rho > 0$ then $T_i > 0$ and the numerator is a polynomial of degree two in c_{-i} , which can be written as:

$$f(c_{-i}) = -T_i(n-1)^2(\lambda+T_i)(c_{-i})^2 - 2T_i(n-1)c_{-i} + 1 \quad (22)$$

The maximum of this polynomial occurs at $c_{-i}^{max} = \frac{-1}{(n-1)(\lambda+T_i)} < 0$ and the parabola opens downwards since the coefficient on $(c_{-i})^2$ is negative. The polynomial has a positive (c_{-i}^*) and a negative root. Since we have assumed that $c_{-i} > 0$, we note that $\frac{\partial c_i}{\partial c_{-i}} \Big|_{a_{-i}=\text{const}} > 0$ for $0 < c_{-i} < c_{-i}^*$ and $\frac{\partial c_i}{\partial c_{-i}} \Big|_{a_{-i}=\text{const}} < 0$ for $c_{-i} > c_{-i}^*$.

Second, we take the derivative of a_i with respect to a_{-i} . When $\rho = 0$ then $\frac{\partial a_i}{\partial a_{-i}} \Big|_{b_{-i}=\text{const}, c_{-i}=\text{const}} = 0$. When $\rho > 0$ then

$$\frac{\partial a_i}{\partial a_{-i}} \Big|_{c_{-i}=\text{const}} = \frac{-R}{(d_i+\lambda+T_i)^2} \frac{\partial T_i}{\partial a_{-i}} > 0 \quad (23)$$

since $\frac{\partial T_i}{\partial a_{-i}} < 0$.

Third, we take the derivate of b_i with respect to b_{-i} . When $\rho = 0$ then $\frac{\partial b_i}{\partial b_{-i}} \Big|_{a_{-i}=\text{const}, c_{-i}=\text{const}} = 0$. When $\rho > 0$ then

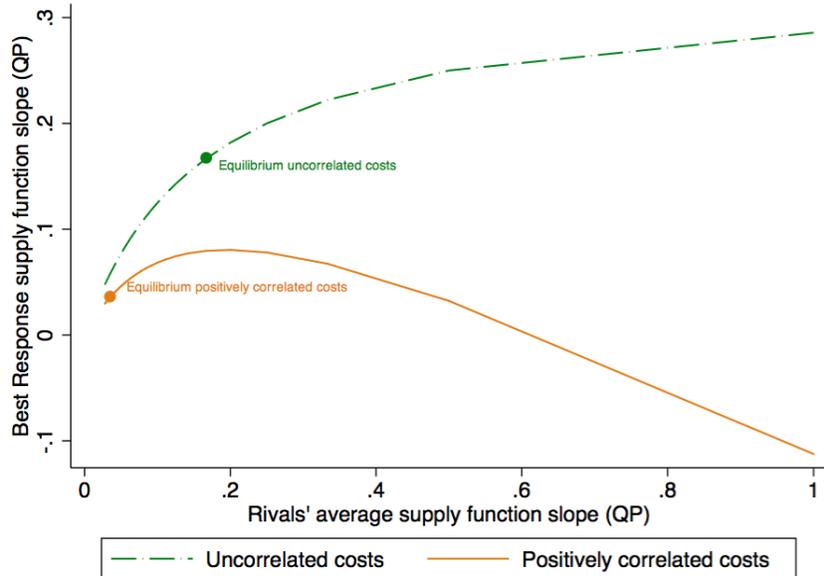
$$\frac{\partial b_i}{\partial b_{-i}} \Big|_{a_{-i}=\text{const}, c_{-i}=\text{const}} = \frac{-T_i(n-1)}{d_i + \lambda + T_i} < 0 \quad (24)$$

■

In the experiment, for each two dimensional choice, we observe a subject's supply function slope and intercept. Concerning the intercept, we cannot separately observe the two components of the supply function's intercept: response to the private signal and fixed part. This issue has consequences for the computation of the best-response to the rivals' supply function in the positively correlated costs treatment, where we need to have an estimate of the rivals' response to the private signal, a_{-i} . We make the assumption that all sellers equally and optimally respond to the private signal, i.e. $a_j = a_{-i}$ for all $j \neq i$. Introducing this assumption into the best-response strategy, we obtain: $a_j = a_{-i} = \frac{(1-\rho)\sigma_\theta^2}{(1-\rho)\sigma_\theta^2 + \sigma_\epsilon^2} (d_i + \lambda)^{-1}$.

The figure below illustrates the best-response slope, c_i , as a function of the rivals' average slope, c_{-i} , in each treatment (with the supply function viewed in (*Ask Price, Quantity*) space).

Figure 11: *Best-response supply function slope as a function of the rivals' average supply function slope in each treatment and the corresponding equilibrium predictions in the (Ask Price, Quantity) space.*



Note. The supply function is viewed in the (*Ask Price, Quantity*) space, which inverts the axes in relation to most of the figures shown throughout the text. Note that a steep supply function has a low c when represented in the (*Ask Price, Quantity*).

Appendix I: Cluster analysis

The next table presents the results of cluster analysis in the intermediate time periods in groups of five rounds.

Table 17: *Cluster analysis for the intermediate time periods.*

Period	Model & Components	Log-likelihood & BIC	Cluster	Number of subjects	Relative frequency	Average SF slope (s.d.)	Average SF intercept (s.d.)
Uncorrelated costs treatment							
Rounds [6,10]	VEV 2	-615.47 -1273.7	Cluster 1	49	68%	4.13 (1.67)	1,001.73 (79.14)
			Cluster 2	23	32%	11.14 (4.41)	851.73 (212.81)
Rounds [11,15]	VEV 2	-597.62 -1,238.02	Cluster 1	56	78%	4.14 (1.82)	994.97 (78.47)
			Cluster 2	16	22%	11.45 (5.14)	771.83 (242.75)
Rounds [16,20]	VEV 2	-590.66 -1,224.1	Cluster 1	57	79%	3.83 (1.82)	993.34 (73.94)
			Cluster 2	15	21%	11.87 (4.33)	820.09 (293.07)
Equilibrium SF						6.00	1,000.00
Positively correlated costs treatment							
Rounds [6,10]	VEI 3	-642.43 -1,336.17	Cluster 1	6	8%	21.53 (7.87)	189.8 (239.16)
			Cluster 2	22	31%	9.59 (3.10)	856.82 (110.29)
			Cluster 3	44	61%	4.52 (2.05)	1,026.03 (60.03)
Rounds [11,15]	VEI 3	-637.61 -1,326.54	Cluster 1	7	10%	20.31 (5.35)	363.26 (394.17)
			Cluster 2	32	44%	8.18 (2.82)	906.98 (107.64)
			Cluster 3	33	46%	3.64 (1.53)	1,042.26 (56.22)
Rounds [16,20]	VEI 3	-656,98 -1.365,27	Cluster 1	7	10%	21.00 (10.14)	415.29 (599.65)
			Cluster 2	34	47%	9.22 (3.03)	890.97 (109.11)
			Cluster 3	31	43%	3.31 (1.48)	1,036.34 (81.56)
Equilibrium SF						26.68	655.32

Note. The numbers in parenthesis below the average correspond to standard deviations (s.d.). The equilibrium SF has been calculated using the average signal realisation.

The following table summarises the average difference between the experimental supply function and the theoretical best-response supply function in each cluster.

Table 18: *Difference between the average experimental and best-response supply function in each cluster.*

Period	Cluster	Frequency of subjects	SF Slope PQ: Difference Experimental and Best Response	SF Intercept PQ: Difference Experimental and Best Response
Uncorrelated Costs Treatment				
First 5 rounds	Cluster 1	58%	-1.88 (2.12)	5.19 (51.47)
	Cluster 2	42%	4.19 (4.04)	-173.42 (243.77)
Last 5 rounds	Cluster 1	72%	-2.25 (1.76)	5.33 (57.54)
	Cluster 2	28%	4.64 (3.83)	-98.28 (153.18)
Equilibrium SF			6.00	1,000.00
Positively Correlated Costs Treatment				
First 5 rounds	Cluster 1	6%	3.67 (8.41)	-744.61 (93.68)
	Cluster 2	36%	-3.19 (4.94)	51.63 (171.09)
	Cluster 3	58%	-9.26 (2.75)	197.88 (105.12)
Last 5 rounds	Cluster 1	8%	5.41 (5.93)	-562.76 (196.65)
	Cluster 2	42%	-6.20 (3.53)	72.10 (159.18)
	Cluster 3	50%	-12.73 (2.78)	271.14 (133.94)
Equilibrium SF			26.68	655.32

Note. The numbers in parenthesis correspond to the standard deviations. The theoretical best-response has been calculated for each individual supply function choice and then averaged for subjects in a particular cluster and time period (blocks of 5 rounds). In order to calculate these averages, we have excluded the choices which had a best-response which was unfeasible for subjects in the experiment (73 observations).

Appendix H: Robustness check. Bayesian updating

One could argue that subjects fail to engage in Bayesian updating and that this is why, in the positively correlated costs treatment, they fail to understand that the market price is informative about costs. To explore this possibility further and to assist subjects in the decision-making process, we conducted an additional session with two groups in the positively correlated costs treatment. These sessions had the same experimental design features as in the baseline treatment but with three exceptions as follows. First, in addition to the signal received, each subject received the expected value of her own costs—and of her rivals' costs—conditional on the signal

received (thus subject i received a signal s_i and was also given $E[\theta_i | s_i]$ and $E[\theta_j | s_i]$ for $i \neq j$).⁵⁸ Second, we explicitly asked each subject to think about what her rivals would do and provided a simulation tool that subjects could use to make a provisional decision, based on those beliefs, and then visualise the resulting market price; the participant could then revise her decision. Third, the experiment lasted for 15 (rather than 25) rounds.⁵⁹ The following table presents the summary statistics of outcomes and choices of these two groups in the positively correlated costs treatment.

Table 19: *Robustness session: behaviour and outcomes (positively correlated costs treatment).*

Positively Correlated Costs Treatment				
Variable	Number of observations	Mean in Robustness Sessions	Mean in Baseline Treatment	<i>Theoretical Prediction</i> ³
<i>Behaviour: Supply Function</i>				
Intercept PQ	360	898.46	899.25	655.32
Slope PQ	360	7.26	7.79	26.68
<i>Outcomes</i>				
Market Price	120	1,115.66	1,123.73	1,544.68
Profits	360	2,007.11	1,937.81	16,567.37
<i>Efficiency</i>				
Deadweight Loss	120	2,011.01	1,654.35	467.74

We find that the average supply function and outcomes in the robustness session are similar to the averages of the baseline treatments, presented in Tables 3 and 5, that correspond to the positively correlated costs treatment. In other words: assisting subjects in their decision-making process seemed not to have a significant effect, in the positively correlated costs treatment, on either their behaviour or the experimental outcomes.

Our interpretation—namely, subjects in the positively correlated costs treatment fail to understand that the market price is related to the level of their costs and hence that they should bid less aggressively—is robust.

⁵⁸The participant instructions for these additional treatments are available upon request. We explained conditional expectations by telling subjects that, in each round, an expert would give them the expected value of both their and their rivals' unit cost.

⁵⁹We reduced the number of rounds because our three modifications increased the experiment's duration and because (as described in Section 5.3) subjects do not adjust much their behaviour in the last few rounds of bidding.