Information Frictions and Market Power: A Laboratory Study

Anna Bayona[*], Jordi Brandts[†], and Xavier Vives[‡§]

July 2018

Abstract

Consider supply function competition with private information about correlated costs. We study whether information frictions lead to market power when agents potentially learn from the price. We find that if private signals are relatively noisy then interdependent costs lead to high market power, whereas if signals are relatively precise then interdependent costs do not raise market power over the case of uncorrelated costs. The results are broadly consistent with the comparative statics of Bayesian supply function equilibrium and, from an estimated mixture model, with an explanation in terms of the limited capacity of agents to process information signals.

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

JEL Codes: C92, D43, L13

1 Introduction

We present the results of a laboratory experiment in a setting representative of real-world markets characterised by competition in demand or supply schedules such as wholesale electricity

[*]Universitat Ramon Llull, ESADE Business School.

[†]Institut d'Anàlisi Econòmica (CSIC) and Barcelona GSE.

[‡]IESE Business School.

[§]We thank José Apesteguia, Maria Bigoni, Colin Camerer, Enrique Fatas, Dan Levin, Cristina Lopez-Mayan, Margaret Meyer, Rosemarie Nagel, Theo Offerman, Stanley Reynolds, Albert Satorra, Arthur Schram, Jack Stecher, and Chris Wallace for useful comments and discussions. Anna Bayona acknowledges the financial support from Banc de Sabadell and the Generalitat de Catalunya (AGAUR Grant 2017 SGR 640). 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.

markets, markets for pollution permits, or liquidity and Treasury auctions. In our experimental markets each seller has incomplete information about her costs, receives a private signal, and competes in supply functions (e.g., Klemperer and Meyer 1989). The aim of our experiment is to study the relationship between market power and information frictions. In particular, our focus is on the causal effect of the correlation of costs between sellers and the relative variation in the private signal noise with respect to the variation of the fundamental (hereafter the *variance ratio*) on market outcomes.[1] Our experiment is complex because we want to reflect some features of competition in actual markets. In conducting our experimental test, we take received theory about competition in supply functions seriously.

Our primary goal is to shed light on the extent to which information frictions lead to market power in supply function competition, such as in wholesale electricity markets. For example, according to Holmberg and Wolak (2016), cost uncertainty and information asymmetry are large in hydro-dominated markets. Understanding what factors contribute to the existence of market power is a central issue in industrial organisation and this is an area to which experimental economics can contribute. Our results may also yield insights into participants' decision processes in a specific market environment.

The issue is of policy importance. Indeed, a competition policy authority could mistakenly infer collusion from high margins in the wholesale electricity market when in fact generators were bidding non-cooperatively taking into account the information conveyed by the price. Or, similarly, the Treasury may suspect collusion in some bond auctions when in fact bidders are responding optimally to the incomplete information environment they face. It is therefore crucial to test whether the theoretical predictions are borne out in the laboratory, or alternatively, whether we find deviations towards either more competitive or less competitive behaviour than predicted by the equilibrium. In electricity markets, mitigating the market power is a primary concern of regulators.[2]

Our design is based on Vives (2011), which structures a wide range of competitive environments where agents compete in schedules with private information providing a very tractable model.[3] In this model, in a supply function equilibrium (hereafter SFE), noisy private information with cost correlation generates market power that exceeds the full-information benchmark. When costs are interdependent, the SFE predicts a higher degree of market power than when costs are uncorrelated. The mechanism that explains these comparative statics results can be stated as follows. With interdependent costs, a Bayesian-rational seller must also realise that a

[1]The 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).

[2]See, for example, Hortaçsu and Puller (2008) and Holmberg and Wolak (2015).

[3]For example, the models of Klemperer and Meyer (1989) and Kyle (1989) can be seen as limit cases of Vives (2011).

high price conveys the information that the average signal of her rivals is high, and therefore, the seller deduces that her own costs must also be high. Hence, she should compete less aggressively than if costs were uncorrelated in order to protect herself from adverse selection. As either costs become more correlated or private signals are relatively less precise then the price is more prominent in the inference of costs, and equilibrium supply functions become steeper, leading to greater market power. In equilibrium, the combination of incomplete information and strategic behaviour leads to greater market power when costs are interdependent than when they are uncorrelated. The SFE with uncorrelated costs coincides with the full-information equilibrium since sellers do not learn about costs from prices. 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 the original concept to multi-unit demand auctions.

The experimental design is as follows. We employ a between-subjects design with four treatments that differ in the correlation among costs and in the variance ratio.[4] Uncorrelated costs treatments are relevant in situations where cost shocks among sellers are purely idiosyncratic, while interdependent costs treatments are appropriate when cost shocks among sellers have a large systematic component due to events that commonly affect all sellers.

In each treatment, participants were randomly assigned to independent groups of twelve participants, each comprising four markets of three sellers. Sellers competed for several rounds, and 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 participants were assigned the role of sellers. Subjects received a private signal about the uncertain cost and 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.

Although we have simplified our design as much as possible, our experiment is rather complex. This is a reflection of the problem we are studying. The model we base our design on is an intricate one, and the markets the model represents are even more multifaceted. Naturally, simpler designs may make it possible to study participants' decision processes more in detail. However, we think that the insights gained by studying simple designs alone are not enough. In our view, complex and simple experiments are complements rather than substitutes. While simple experiments allow one to study a particular issue in detail, complex experiments are useful to obtain a more complete picture in environments where people face more than one

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

problem at the same time. More complex games are certainly required to understand situations of interest, such as how top managers decide to allocate company resources.[5]

The results that we wish to highlight pertain to the comparative statics of the SFE, which relate information frictions to market power, together with an explanation of observed behaviour. Our measure of market power is price impact, which reflects the ability of a seller to influence the market price. Price impact is given by the slope of residual demand facing a trader (a competitive trader would face an infinitely elastic demand). We find that the observed changes in market power and absences thereof are consistent with two of the three comparative statics predictions of the SFE. Consistent with the predictions of the SFE, we first find that average market power in markets with uncorrelated costs is the same regardless of the variance ratio, and that it is close to what the SFE prescribes. Second, with interdependent costs and a high variance ratio, average market power increases significantly. Inconsistent with the predictions of the SFE, our experiment shows that average market power does not increase significantly when the variance ratio is low and costs are interdependent. In other words, if the signal is relatively noisy then interdependent costs lead to high market power, whereas if the signal is relatively precise then interdependent costs do not lead to a difference in market power with respect to when costs are uncorrelated.

We also find that market power, as measured by price impact or the supply function slope, does not depend on participants' private signal. However, if the signal is relatively precise then the supply function's intercept is more responsive to the signal than when it is relatively noisy. This suggests an interpretation of our results in terms of the cognitive processes participants engage in. The relative attention participants pay to the private signal compared to the informational content of the price may depend on the noisiness of the signal. If costs are interdependent and the signal is rather precise then participants pay more attention to the signal and less to the price. By contrast, if the signal is relatively noisy then participants focus less on it and extract more information from the price leading to a larger increase in market power than in the previous case.

The results of our mixture model that are used to explain the observed heterogeneity in market power are consistent with this explanation. Changing from uncorrelated to interdependent costs and a constant low level of signal noise does not lead to higher market power. In both treatments, more than half of the participants deviate from perfect rationality either because they are price-takers or they ignore the information content of the market price, or both. By contrast, changing from uncorrelated to interdependent costs and a constant high level of signal noise does lead to higher market power. This change leads to a reduction in the proportion of price-takers and an increase in the proportion of participants who either play the SFE or best-respond to the empirical distribution of behaviour.

[5] A recent example of an interesting complex experiment is Selten, Pittnauer and Hohnisch (2012).

To the best of our knowledge, ours is the first laboratory experiment to test the relationship between information frictions and market power in the context of supply (or demand) function competition. We now briefly refer to some relevant literature.

Some 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 the supply function equilibrium concept, as well as Brandts et al. (2008), and Brandts et al. (2014), who study how the presence of forward markets and of pivotal suppliers affects prices. None of these supply function experiments incorporates informational frictions.[6]

Our experiment is also related to the literature of multi-unit uniform price auctions with incomplete information 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). 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. 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.

More indirectly, 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. The winner's curse is a prevalent, consistent, and robust phenomenon in single-unit auctions featuring common (or interdependent) values (Kagel and Levin 1986; Goeree and Offerman 2003; Kagel and Levin 2015). Because the analogy between the winner's curse with competition in supply functions and with single-unit auctions is relevant with respect to adverse selection but not necessarily with respect to market power.[7] Our results are consistent with this literature when signal is

[6] As in most laboratory experiments, our subjects are university students. A natural question is how useful our data are to shed light on what happens in markets in which decisions are typically made by experience market traders. Fréchette (2015) surveys all the existing experimental studies in which the behaviour of students and experts are compared. His overall conclusion is that: "(...), overall much of the big picture seems the same whether one looks at professionals or students in laboratory experiments testing economic models."

[7] Our results are more closely related to the theoretical notion of 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

relatively precise but not when it is noisy.

Several other papers study behaviour in other contexts and show that individuals fail to extract information from other people's actions in various simpler contingent reasoning setups. Some examples of these contexts include bilateral negotiations (Samuelson and Bazerman 1985), in the acquiring a company game (Charness and Levin 2009), and voting (Esponda and Vespa 2014). Ngangoué and Weizsäcker (2018) find that traders have trouble deriving information from hypothetical values of prices. In addition, 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). Both of those papers report that (a) a large proportion of participants behave just as in the equilibrium where participants play simple but strategic private-information games yet (b) this proportion declines markedly with increasing strategic complexity of the game.

Our findings indicate that when there are two sources of information (the private signal and the price), the relatively quality of these two sources of information may matter for the agents' attentiveness when they have a limited capacity to process information (Kacperczyk et al. 2014). Specifically, we find that if costs are interdependent and the private signal is relatively noisy then participants pay more attention to the price leading to greater market power compared to when the private signal is relatively precise.

The rest of our paper is organised as follows. We explain the theoretical model in Section 2. Section 3 describes the experimental design and hypotheses, after which Section 4 explains the experimental procedures. In Section 5 we present the experimental results, and in section 6 we show the results of two robustness treatments. We conclude in Section 7. Further details of the equilibrium characterisation, experimental instructions, auxiliary empirical results, and an analysis of participants' characteristics are provided in Online Appendices.

2 Theoretical Background

There are n sellers who compete simultaneously in a uniform price auction, and each seller submits a supply function. Seller i's profits can be written as

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

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

aggressively.

The information structure is as follows. A seller does not know the value of the cost shock θ_i before submitting her supply schedule, and she receives a signal $s_i = \theta_i + \varepsilon_i$ for which the error term is distributed as $\varepsilon_i \sim N(0, \sigma_\varepsilon^2)$, and hereafter assume that $\sigma_\varepsilon^2 > 0$. Sellers' random cost parameters may be correlated, with $\mathrm{corr}(\theta_i, \theta_j) = \rho$ for $i \neq j$. Define the variance ratio of the noise in the private signal to the noise in the fundamental as $\phi \equiv \sigma_\varepsilon^2/\sigma_\theta^2$ (hereafter *variance ratio*). Hence, the private signal is relatively precise (imprecise) in relation to the noise in the fundamental if ϕ is low (high).

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.[8] Given the signal received, a strategy for seller i is to submit a price-contingent schedule, $X(s_i, p)$, of the form

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

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 the unique linear symmetric supply function equilibrium (SFE) such that the first order condition satisfies:

$$X(s_i, p) = \frac{p - E[\theta_i \mid s_i, p]}{d + \lambda}, \tag{3}$$

where (a, b, c) are functions of the information structure (μ, ϕ, ρ) and market structure (n, q, λ), and $d = ((n-1)c)^{-1}$ is the (endogenous) slope of the inverse residual demand for a seller. [9]

Let us focus on the equilibrium reasoning. Recall that $\sigma_\varepsilon^2 > 0$. If $\rho\sigma_\varepsilon^2 = 0$ then costs are uncorrelated and the seller does not learn about θ_i from the market price.[10]If $\rho\sigma_\varepsilon^2 > 0$ then costs are interdependent and the market price at the SFE has two roles as can be seen from (3): a) the market price is an index of scarcity, since a high equilibrium market price means that the seller has an incentive to increase her supply; b) it contains information about θ_i, since, if costs are interdependent, a rational seller should infer that a high price reveals that her cost is high (because $E[\theta_i \mid s_i, p]$ increases with p). Hence, if $\rho\sigma_\varepsilon^2 > 0$, then equilibrium supply functions are

[8]See, for example, Baldick et al. (2004).

[9]Throughout the paper, we shall often work with the inverse supply function, since it corresponds more closely to how supply functions were represented in the experiment in the $(Quantity, AskPrice)$ space: $p = \hat{b} + \hat{a}s_i + \hat{c}X(s_i, p)$, where the coefficients of the inverse supply function $(\hat{a}, \hat{b}, \hat{c})$ are related to (a, b, c) as follows: $\hat{b} = -\frac{b}{c}$, $\hat{a} = \frac{a}{c}$, and $\hat{c} = \frac{1}{c}$ for $c \neq 0$. Henceforth, we will omit the modifier "inverse" and refer simply to "the supply function". We shall use *InterceptPQ* as the empirical counterpart of $\hat{b} + \hat{a}s_i$ and *SlopePQ* as the empirical counterpart of $\hat{c}$. Our graphs plot the supply function in the usual $(Quantity; AskPrice)$ space.

[10]In the auction literature, the case when $\rho\sigma_\varepsilon^2 = 0$ is either referred to as independent values if values are uncorrelated, or to private values when signals are fully informative.

steeper in relation to when costs are uncorrelated. In this case, as either costs become more correlated (higher ρ) or private signals are relatively less precise (higher ϕ) then the price is more prominent in the inference of θ_i, and equilibrium supply functions become steeper, leading to greater market power.

The condition that supply equals demand in equilibrium, $q = \sum_{i=1}^{n} X(s_i, p)$, yields the expected equilibrium price. The model predicts that equilibrium market outcomes are less competitive in markets with interdependent costs compared to markets independent or private costs. As a consequence, the expected market price and profits of the equilibrium allocation can be shown to be higher in the former case than in the latter.

To summarise, the main hypotheses resulting from the SFE model is that information frictions associated with interdependent costs generate market power above markets with uncorrelated costs.

The slope of the inverse residual demand for a seller, d, is a measure of *price impact* and will be our main measure of market power, which reflects the ability of a seller to influence the market price.[11] At the SFE, price impact is equal to

$$d^{SFE} = \frac{\lambda(1+M)}{(n-2-M)}, \tag{4}$$

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. Equilibrium requires that $n-2 > M$. As M grows and approaches $n-2$, d tends to infinity and the equilibrium breaks down. Minimal market power obtains when $M=0$. M increases with ρ and with the variance ratio, $\phi \equiv \sigma_\varepsilon^2/\sigma_\theta^2$. Figure 1 graphically displays two inverse residual demand that a seller faces. The dashed (continuous) line represents a situation where a seller faces a flatter (steeper) inverse residual demand which means that her capacity to influence the price is low (high).

[11] Price impact as a measure of market power has been used extensively in the financial economics literature, such as for example, by Kyle (1989).

Figure 1: Inverse residual demands.

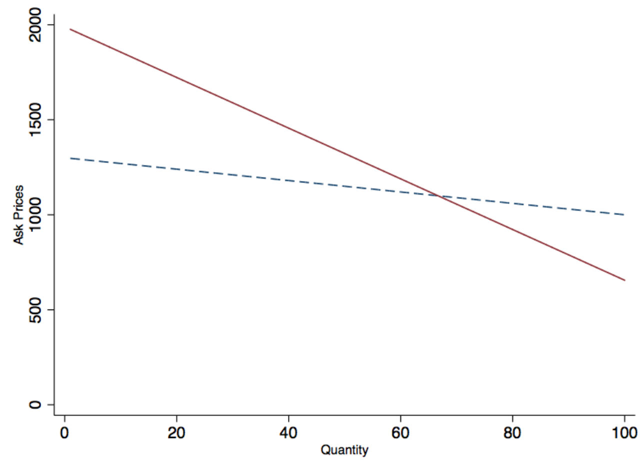

3 Experimental design and hypotheses

Experimental design

The experiment was based on the market described in the previous section. Each subject received a private signal and was subsequently asked to choose two *ask prices*: one for the first unit offered ($AskPrice1$) and one for the second ($AskPrice2$). We then used these two ask prices to construct a linear supply schedule, which was shown graphically on each subject's screen. The participant could revise the ask prices several times until she was satisfied with her decision. The buyer was simulated. Once all supply schedules had been submitted, the auctioneer (the computer) calculated the uniform market price, and participants obtained feedback on own and rivals' performance and costs. Treatments varied in the correlation among costs (ρ) and in the variance ratio (ϕ) in the design described in Table 1.

Table 1: *Summary of the experimental design*

Experimental Treatments	Low Φ	High Φ
Uncorrelated Costs (ρ=0)	TUL	TUH
Positively Correlated Costs (ρ>0)	TCL	TCH

Conducting the experiment required us to specify numerical values for the theoretical model's fixed parameters. For this purpose, we used three implementation criteria: (i) the existence of a unique equilibrium; (ii) sufficiently differentiated equilibrium behaviour and outcomes between the treatments with different correlation levels and equal or similar equilibrium behaviour and outcomes with the same correlation; and (iii) reduction of computational demands placed on participants.

Table 2: *Fixed parameters used in the experiment.*

	Symbol	Value
Number of Sellers	n	3
Inelastic demand	q	100
Transaction cost parameter	λ	3
Mean of random cost	μ	1,000
Mean of signal's error	$\bar{\varepsilon}$	0
Variance of signal's error	σ^2_ε	3,600

In addition, we told bidders that the simulated buyer would not purchase any unit at a price higher than 3,600 (price cap). The price cap was not part of the theoretical model, we chose a value high enough to preclude distortion of equilibrium behaviour.

Table 3 shows the price impact, the supply function, expected market prices and profits at the SFE given the experimental parameters. The numerical values of the treatment parameters in terms of the correlation among costs (ρ) and in the variance ratio (ϕ) of the experimental design are presented in the vector form (ρ, ϕ) next to the acronyms for the different treatments.

Table 3: *Treatment parameters and theoretical SFE predictions in each of the treatments.*

TUL≡(0, 0.36)	**TUH≡(0, 36)**	**Treatment name and parameters**
3.00	3.00	Price impact
p=1000+6X	p=1000+6X	Equilibrium inverse supply function when s_i=μ
1,200.00	1,200.00	Expected market price
5,612.75	5,000.23	Ex-ante expected profit at the SFE
TCL≡(0.6, 0.36)	**TCH≡(0.175, 36)**	**Treatment name and parameters**
13.34	12.52	Price impact
p=655.32+26.68X	p=682.72+25.04X	Equilibrium inverse supply function when s_i=μ
1,544.68	1,517.28	Expected market price
16,567.40	15,576.1	Ex-ante expected profit at the SFE

In the table, rows vary in terms of the correlation among costs: the first row corresponds to the two treatments with uncorrelated costs ($\rho = 0$); while the second row displays two treatments with interdependent costs, which have $\rho\sigma^2_\varepsilon > 0$. Columns vary in terms of the variance ratio (ϕ). Notice that treatments TUH and TCH have relatively imprecise private signals in relation to treatments TUL and TCL.

As mentioned above, the combinations (ρ, ϕ) were chosen so that the two treatments with $\rho > 0$ had similar degree of market power (d) to make them easily comparable. We 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 an equilibrium exists, and for $\phi = 0.36$, the highest correlation that satisfied our implementation criteria described earlier was $\rho = 0.6$. In addition, for $\phi = 36$ the equilibrium does not exist for any $\rho > 0.25$, and hence we chose $\rho = 0.175$ which, within the range of existence, has a very similar degree of market power as TCL.

Alternative Benchmarks

For the SFE predictions, we have assumed that all sellers are perfectly rational and that they have common knowledge of the rationality of other players, as well as coordinating expectations on the SFE. However, there are other behavioural possibilities, which include asymmetric strategies. Recall that our measure of market power, the *price impact*, corresponds to the slope of the inverse residual demand for a seller. We adapt the definition of price impact of the symmetric model to a situation where sellers' strategies are *not* symmetric. Suppose that the supply function of seller i can characterised by (a_i, b_i, c_i), with corresponding inverse residual demand $p = I_i - d_i x_i$, where its slope is

$$d_i = \left(\sum_{j \neq i} c_j \right)^{-1}, \tag{5}$$

and its intercept $I_i = d_i \left(q - \sum_{j \neq i} b_j + \sum_{j \neq i} a_j s_j \right)$. Note that d_i, our measure of market power, is the inverse of the sum of the rivals' supply function slopes in a market.

There are two effects that characterise the best-response slope of sophisticated seller i (BR_S). The 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 BR_S is increasing in the rivals' average supply function slope. When costs are positively correlated then there is –in addition to the strategic effect– also an *inference effect* that is related to information conveyed by the price. Seller i correctly thinks that a high price indicates 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.[12] We can also consider that some sellers best-respond without taking into account the informational content of the price - i.e. they are informationally naïve best-responders (BR_N). These sellers only consider the strategic effect and ignore the inference effect.[13]

Assuming that sellers' strategies are symmetric, three alternative benchmarks can be formu-

[12]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.

[13]See Online Appendix A.2 for further details of formulae that characterise the theoretical best-response strategy.

lated by relaxing the two key factors (strategic behaviour and optimal information processing) that lead to market power in our context. A first possibility is that subjects are price-takers since they do not perceive the influence of their supply function decisions on prices, but condition on their private signal and learn from the price (hereafter PT benchmark). This will be the case when all agents ignore the strategic effect. A second possibility is that sellers are informationally naïve since they do not extract information from the market price, but behave strategically in the market (IN). This will be the case when agents ignore the inference effect. A third possibility is that sellers are both price-takers and informationally naïve (PTIN). This is the case when agents ignore both the strategic and the inference effect.See Online Appendix A.1 for further details of formulae that characterise all the benchmarks. Furthermore, a sophisticated supplier may best respond to perceived behaviour of other players that does not conform to any of these benchmarks.

Table 4 presents the price impact predictions of these other benchmarks together with the prescription of the SFE. Notice that if costs are uncorrelated then the price is not informative about costs and, hence, the PT and PTIN benchmarks are identical, and IN coincides with the SFE benchmark. In addition, we can see that the degree of market power is greater when sellers are strategic (i.e. in the SFE benchmark) than when they are price takers (i.e. in the PT benchmark). When costs are interdependent, sellers have the greatest degree of market power when they are strategic and exploit the informational content of the price (SFE benchmark), followed by when sellers are strategic but informationally naïve (IN), followed by when sellers are price takers even if they exploit the informational content of the price (PT), and the smallest degree of market power occurs in markets with fully naïve sellers (PTIN). Notice also that in treatments with interdependent costs (TCL and TCH), the SFE predicts a substantially larger degree of market power than all the alternative benchmarks, while in treatments with uncorrelated costs (TUL and TUH), the predictions of the SFE and alternative benchmarks are more similar.

Table 4: *Price impact (d) predictions for the SFE and other benchmarks in each treatment.*

TUL	**TUH**	**Treatment** *Benchmarks*
3.00	3.00	SFE
3.00	3.00	IN
1.50	1.50	PT
1.50	1.50	PTIN
TCL	**TCH**	**Treatment** *Benchmarks*
13.34	12.52	SFE
3.00	3.00	IN
2.45	2.42	PT
1.50	1.50	PTIN

Hypotheses

In our analysis of information frictions and market power, we focus mostly on the comparative statics of the SFE, but also consider its point predictions.[14] Given the theoretical model, the SFE benchmark, and the parametrisation, we can formulate the following two hypotheses regarding market power and information frictions.

H1: *Regardless of the variance ratio, there is the same degree of market power in treatments with uncorrelated costs: market power is the same in TUL and TUH.*

H2: *Given a level of the variance ratio, there is greater market power in treatments with interdependent costs than with uncorrelated costs: a) market power is higher in TCL than in TUL; b) market power is higher in TCH than in TUH.*

Observe that the comparative statics predictions of the SFE and of the alternative benchmarks are not the same. Although all propose the same market power in TUL and TUH, they differ with respect to the comparison between the treatments with uncorrelated and interdependent costs. Whereas the SFE prescribes substantially higher market power with interdependent costs, the alternative benchmarks propose no effect (IN and PTIN) or a small increase (PT).

We next formulate a stricter hypothesis relating each treatment to the *point predictions* made by the SFE.[15]

H3: *Market power in each treatment is equal to the prediction made by the SFE benchmark.*

We also study whether the alternative benchmarks (PT, IN, PTIN) provide a good description of the relationship between market power and information frictions.

We next formulate an auxiliary hypothesis regarding the comparative statics relationship between market power and the private signal.

H4: *Market power, as measured by price impact, does not depend on the private signal. However, the supply function intercept reacts more to the private signal when it is more precise.*

This hypothesis is important in order to check whether the features of this class of models are observed in our experimental data.

4 Experimental Procedures

We ran the experiment with 384 participants, 288 who participated in the main experiment grouped in 72 participants for each of the main treatments described in Table 1, and 96 participants who participated in the robustness sessions. Each main treatment had 6 independent

[14]For other studies in experimental industrial organisation that focus on the comparative statics of theoretical models, see Holt (1985) on quantity competition, Brandts et al. (2014) on supply function competition, Morgan et al. (2006) on price dispersion, Isaac and Reynolds (1988) on the relationship between spillovers and innovation.

[15]The comparison of market power between TCL and TCH belongs to the point predictions (H3) and not to comparative statics since its equivalence is specific to the chosen parametrisation and not to general shifts in market power due to information frictions.

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). Participants competed for 2 trial rounds followed by 25 (in TUL and TCL) or 18 (in TUH and TCL) incentivised rounds since it is an established fact that equilibrium does not appear instantaneously in experimental games.[16] 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. Thus, the composition of each of the four markets varied each round within a group.

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 units for sale. Second, we asked sellers to construct a nondecreasing and linear supply function.[17] Third, ask prices had to be nonnegative. Fourth, as mentioned above, we imposed a price cap in order to limit the potential gains of sellers in the experimental sessions.

At the end of each round, each participant 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).[18] 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.[19] 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.

At the end of the experiment, participants completed a questionnaire that requested personal information and asked questions about the subject's reflections after playing the game. Once

[16]In treatments TUH and TCH participants competed for 18 rounds since experimental sessions were too long with 25 rounds (more than 3 hours) and we reduced the number of rounds for practical purposes.

[17]Bolle et al. (2013) allowed participants to enter non-linear supply functions and found that participants often enter convex supply functions. We restricted supply functions to be linear since we want to exclusively focus on the relationship between market power and information frictions and compare our experimental results to the unique linear SFE.

[18]Subjects did *not* receive feedback on the signal received by others because that would not be expected to occur in reality.

[19]Note that this information feedback may have encouraged herding behaviour (among others, participants are informed of each others' supply functions). However, we do not have any specific indication that this occurred.

the questionnaire was completed, each participant was paid in private.

Sessions were conducted in the LINEEX laboratory of the University of Valencia (Spain). The participants were undergraduate students in the fields of economics, finance, business, mathematics, engineering, and natural sciences. All sessions were computerised.[20] Instructions were read aloud, questions were answered in private, and—throughout the sessions—no communication was allowed between participants. 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 with a comprehension test.

As for incentives, each participant started with 50,000 experimental points.[21] During the experiment, participants won or lost points. At the end of the experiment, the total number of points were exchanged for euros at the rate of 10,000 experimental points per euro.[22] The payments ultimately made to participants ranged from 10 to 32.1 euros and averaged 23.5 euros. Each session lasted an average of 2.5 hours. See Online Appendix B for the instructions, comprehension test, screenshots used for running the experiment, and the end-of-experiment questionnaire, and Online Appendix D for some demographic information of the participants.

5 Experimental Results

In Section 5.1 we provide an analysis of the experimental relationship between information frictions and average market power, while in section 5.2 we address the heterogeneity of market power in each of the treatments of our experiment. For comparability and consistency, we use the first 18 rounds in all treatments.

5.1 Information frictions and market power

In this part of the results section, we analyse all the available rounds of the experiment in each treatment. Further analysis on time trends and for the last 5 rounds can be found in the Online Appendix C.1 and C.2, respectively.

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

[21]These points were equivalent to 5 euros.

[22]An alternative to our payment scheme would have been to use the random-lottery incentive scheme and pay participants only the earnings of one period selected at random. This procedure has the potential of eliminating wealth effects that can arise when paying for all periods. At the same time, given that we have 25 periods, the probability of a particular period being chosen may be too small to have participants engage and focus their attention. In any case, note that we concentrate on a treatment comparison, which should not be affected by the specific features of the payment scheme. For an interesting discussion of different payment schemes see Bardsley (2009).

Comparative statics

We use price impact as our main measure for market power, and to get a more complete view of market power, we also analyse the supply function, characterised by the intercept and the slope. In addition, we report market prices and profits in each of the treatments.[23]

Figure 2 shows the average supply function in each treatment, when decisions are averaged across all rounds, together with the corresponding theoretical benchmarks, and Table 5 provides a numerical summary of the average experimental price impact and the average supply function in each treatment.

Figure 2: *Average experimental supply functions and theoretical benchmarks in each of the treatments.*

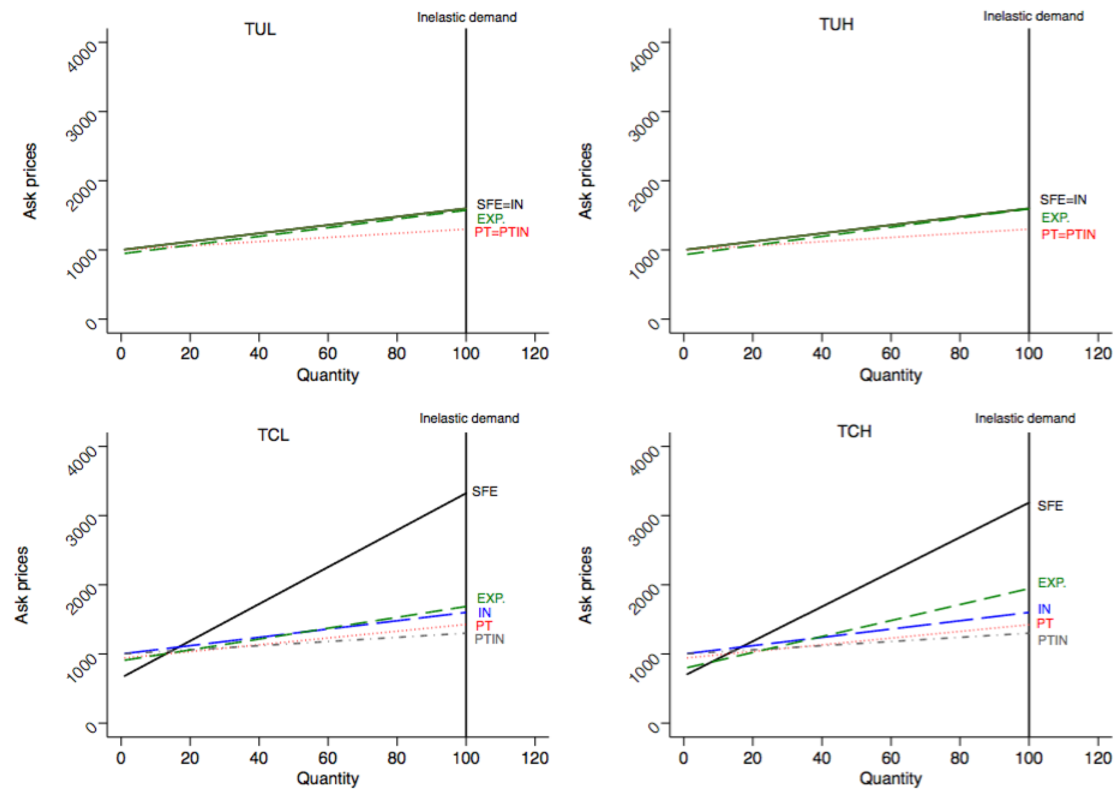

Note. *EXP. (in green) refers to experimental supply function; SFE (in black) to the supply function equilibrium benchmark; IN (in blue) to the informationally naïve benchmark; PT (in red) to the price-taking benchmark; and PTIN (in grey) to the informationally naïve and competitive benchmark.*

[23] We analyse bidding behaviour in terms of the inverse supply function since it corresponds to how participants made their decisions: $p = InterceptPQ + SlopePQ\, X(s_i, p)$, where $InterceptPQ$ and $SlopePQ$ are the intercept and slope of the inverse supply function, respectively. From any participant's two-dimensional choice, we can infer that $SlopePQ = AskPrice2 - AskPrice1$, and: $InterceptPQ = AskPrice1 - SlopePQ$.

Table 5: *Average market power in each treatment.*

TUL	TUH	Treatment
2.74	2.90	Price impact
(2.05)	*(1.39)*	*(s.d. Price impact)*
p = 943.01 + 6.30X	p = 929.87 + 6.64X	Average inverse supply function
(222.66) (5.48)	*(160.14), (4.44)*	*(s.d. InterceptPQ), (s.d. SlopePQ)*
TCL	**TCH**	**Treatment**
3.38	4.89	Price impact
(2.82)	*(2.81)*	*(s.d. Price impact)*
p = 901.35 + 7.85X	p = 790.74 + 11.54X	Average inverse supply function
(281.81) (6.89)	*(237.54) (7.38)*	*(s.d. InterceptPQ), (s.d. SlopePQ)*

Note. The first row shows the average, while the second in brackets and italics is the standard deviation (s.d.). The number of observations in each treatments is equal to 1,296. In addition, price impact is not defined for the 63 observations that have an undefined denominator in d_i.

We observe from Figure 2 that the average experimental supply functions in the uncorrelated costs treatments (TUL, TUH) are close to the corresponding SFE benchmark, and less competitive than the two price-taking (PT and PTIN) benchmarks. In TCL and TCH observed behaviour appears to be far from the SFE, and close to the IN alternative benchmark. However, this closeness needs to be evaluated statistically; we do this in Table 8.

Table 5 shows that market power is similar in the uncorrelated costs treatments (TUL and TUH), as reflected by both the average price impact and the slope of the supply function. However, when costs are interdependent (TCL and TCH treatments), the average supply functions are far from the SFE prediction. In particular, we observe that in the TCL treatment, the experimental average supply function is close to the informationally naïve (IN) benchmark, which displays a substantially lower degree of market power than in the SFE. This result bears some resemblance to the result presented in Goeree and Offerman (2003) in the context of a single-unit and a second-price auction comparing common versus uncorrelated private values. Observe also that the interdependent costs treatments (TCL and TCH) have a higher price impact and steeper supply functions compared to when costs are uncorrelated, with the highest market power in the treatment with a high variance ratio (TCH).

In Table 6 shows test results for treatment differences in market power using the Mann-Whitney-Wilcoxon rank-sum test (the unit of observation is the average per group). These results are the basis for our statistical evaluation of H1 and H2.

Table 6: *Tests of hypotheses pertaining to market power.*

Null Hypotheses	Price Impact p-values	SlopePQ p-values	InterceptPQ p-values
H_0: $MP_{TUL} = MP_{TUH}$	0.337	0.423	0.262
H_0: $MP_{TUL} = MP_{TCL}$	0.423	0.423	0.749
H_0: $MP_{TUH} = MP_{TCH}$	0.0065***	0.0039***	0.0039***

Note: *MP* denotes the variable related market power, *i.e. MP={Price Impact, SlopePQ, InterceptPQ}*. The table reports p-values of two-sided tests. Denote *, **, *** significance at the 10%, 5% and 1% levels, respectively. There are 6 observations per treatment.

We find that, as predicted by the SFE, and the alternative benchmarks, there are no significant difference between market power in the two treatments with uncorrelated costs (TUL and TUH). Increasing the correlation, while keeping the variance ratio fixed at a low level, does not lead to greater market power as the SFE would predict (since we cannot reject the null hypothesis that market power in TUL and TCL are the same). By contrast, we find that an increase in correlation, while keeping the variance ratio fixed at a high level, leads to a significant difference in market power, according to all three measures, in relation to the equivalent treatment with uncorrelated costs.

Table 7 shows the summary statistics of additional indicators of market power: market prices and profits.[24]

Table 7: *Average market prices and profits in each treatment.*

TUL	**TUH**	**Treatment**
1,109.14	1,122.37	Average market price
(125.79)	*(60.96)*	*(s.d. Market price)*
2,090.08	2,090.08	Average profit
(2,762.75)	*(2,762.76)*	*(s.d. Profits)*
TCL	**TCH**	**Treatment**
1,123.10	1,149.32	Average market price
(136.09)	*(99.15)*	*(s.d. Market price)*
1,875.44	2,983.52	Average profit
(6,136.85)	*(3,974.92)*	*(s.d. Profits)*

Note. The first row shows the average, while the second in brackets and italics is the standard deviation (s.d.). For market prices, the unit of analysis is the market with 432 in each treatment, while for profits the unity of analysis is the individual choice with 1,800 observations.

The results on market prices and profits corroborate the previous findings. Using rank-sum tests for differences in the distribution of market prices and profits when the unit of observation is the group, we find that, there are no differences in average market outcomes between treatments TUL and TUH (p-value=0.055 for market prices, p-value=0.87 for profits) and between

[24] In addition, Appendix C.4 displays results on the efficiency of the experimental allocations.

treatments TUL and TCL (p-value=1.00 for market prices, p-value=0.42 for profits). However, average market prices are higher in treatment TCH than in TUH, while the evidence for profits is weaker (p-value=0.016 for market prices, p-value=0.055 for profits). In addition, we find that average experimental market prices and profits are lower than theoretically predicted by the SFE. This result bears some resemblance to a similar fact that typically occurs in common value auctions where bidders ignore the correlation among costs (Kagel and Levin 1986).

We can summarise our findings with respect to H1 and H2 as follows:

> **Result 1** [Comparative statics related to market power and information frictions]:
>
> *(i) There is no significant difference in average market power between TUL and TUH. Average market power in markets with uncorrelated costs is the same regardless of the variance ratio. This result is consistent with H1.*
>
> *(ii) Average market power is significantly higher in TCH than in TUH, but not significantly different between TCL and TUL. It increases with cost correlation when the variance ratio is high but it does not increase significantly when the variance ratio is low. This result is consistent with our H2.b but not with H2.a.*

This is our central result since our main aim is to study the relationship between market power and information frictions.

Point Predictions

We next examine the specific point predictions made in H3 regarding price impact in each of the treatments. We also test for the alternative benchmarks given by PT, IN, PTIN. Note that this is a very stringent test of the benchmarks, and is not the primary focus of our investigation. In order to test them more precisely, we use individual level data, which allows us to use all the available information, while taking into account the dependence between observations.

We first run a random-effects panel regression of price impact on the treatment dummies, and then use these regressions to carry out post-estimation tests of the equality of the regression's coefficients to the specific point prediction made by each model. In this regression, the unit of analysis is the individual across rounds and we cluster standard errors at the independent group level. The dependent variable is *Price impact* (di), while the explanatory variables are treatment dummies: *D_correlation* is 1 in treatments with positive correlation, and 0 in treatments with uncorrelated costs; *D_phi* is 1 in treatments with high variance ratio, and 0 otherwise; *D_correlationphi* is an interaction term that is 1 in the treatment with positive correlation and high variance ratio, i.e. in treatment TCH, and 0 otherwise. The results of the estimated coefficients are in Table 8. A. In Table 8. B, we conduct a formal test of each of the specific point predictions made by each of the theoretical benchmarks regarding price impact in each of the treatments (see Table 4), where the null hypothesis is that the relevant coefficient in the panel regressions presented in Table 8. A is equal to the specific point prediction made by each of the theoretical benchmarks made in each treatment, using t-tests.

Table 8: *Post-estimation t-tests regarding price impact of the point predictions made by each of the theoretical benchmarks.*

A. Panel Regression

	Price impact
Constant	2.73***
	(0.12)
Treatment variables	
D_correlation	0.64
	(0.48)
D_phi	0.17
	(0.16)
D_correlationphi	1.34**
	(0.62)
Observations	5,089
R^2	0.12

B. Post Estimation t-tests

Benchmarks/ Treatments	TUL	TUH	TCL	TCH
SFE	0.0261**	0.388	0.000***	0.000***
PT	0.000***	0.000***	0.0462**	0.000***
IN	-	-	0.421	0.000***
PTIN	-	-	0.000***	0.000***

Notes. The null hypotheses in Table 8.B is that the relevant regression coefficient in the panel regressions presented in Table 8.A is equal to the specific point prediction made by each of the theoretical benchmark considered in each of the treatments. The p-values for testing the IN and PTIN models of the two uncorrelated costs treatments are identical to the SFE and PT models, respectively, since they lead to identical predictions. **, *** denote significance at the 5%, and 1% levels, respectively.

Results of Table 8.A complement Result 1, and give a more formal perspective. We do not find that differences in market power due to the correlation when the variance ratio is low, however, market power can arise due to information frictions when the variance ratio is high. The results in Table 8.B show that we can reject the point predictions of the SFE for all treatments except for TUH. At the same time, we can see that none of the alternative benchmarks does statistically clearly better than the SFE. PTIN and PT are rejected in all treatments. The IN benchmark does well for the TCL but not for TCH. The next result sums up the evidence relating to H3.[25]

Result 2 [Tests of the point predictions made by each theoretical benchmark in each of the treatments]:

(i) Average market power is significantly different from the SFE point predictions in TUL, TCL and TCH, but not in TUH. Therefore, we reject H3 in all treatments but TUH.

(ii) None of the alternative benchmarks explains market power better than the SFE. PTIN and PT are rejected in all treatments. The IN benchmark can explain average

[25]Note, however that, in TUL average market power is between the SFE(=IN) and PT(=PTIN) benchmarks, while in TCH average market power is between the SFE and IN benchmarks.

market power for the TCL but not for TCH.

The picture that emerges from our examination of average market power in the four treatments is that the experimental data is consistent with two of the three comparative statics predictions of the SFE, whereas most of its point predictions are rejected. At the same time, none of the alternative benchmarks yields overall more reliable predictions than the SFE.

Market power and the private signal

We next study how responsive market power and supply functions are to the private signal in each treatment. We run random effects panel regressions with the dependent variables market power (d_i), supply function slope (*SlopePQ*) and intercept (*InterceptPQ*) on the independent variable *Signal*. These regressions cluster standard errors at the independent group level. We find that the private signal is not significant in the regressions of market power (p-value=0.853) and supply function slope (p-value=0.261). However, the private signal does have a significant effect on the supply function intercept (p-value=0.000). In order to get a more complete view of how the supply function intercept depends on the private signal, we report the results of random effects panel regressions with the supply function's intercept on private signal separately for each of the treatments in Table 9.

Table 9: *Results from panel regressions of InterceptPQ on the private signal (Signal) in each of the treatments.*

	InterceptPQ TUL	InterceptPQ TCL	InterceptPQ TUH	InterceptPQ TCH
Constant	206.68***	203.19***	501.62***	355.16***
	(67.24)	(64.20)	(136.93)	(60.19)
Signal	0.740***	0.695***	0.428***	0.435***
	(0.055)	(0.033)	(0.140)	(0.068)
Observations	1,296	1,296	1,296	1,296
R^2	0.16	0.085	0.029	0.014

Notes. *, **, *** denote significance at the 10%, 5%, and 1% levels, respectively. Results of the panel regressions correspond to random effects with robust standard errors are clustered at the independent group level and are given in parentheses.

The results of the regressions show that in treatments where the private signal is relatively precise (TUL and TCL), then the supply function's intercept is more responsive to the private signal, whereas in treatments with a relatively noisy private signal (TUH and TCH), then supply function's intercept is less responsive to the signal.[26] This empirical observation confirms the comparative statics of the SFE benchmark. The next result summarises our findings with respect to H4.

[26] Formal tests of these results can be found in Appendix C.3, and show that the response to the private signal is the same in TUL and TCL, while the response to the private signal is lower and the same in TUH and TCH.

Result 3 [The responsiveness of supply functions to the private signal]:

(i) Market power, as measured by price impact, is not influenced by the private signal.

(ii) With regard to the supply function intercept, we find that if the private signal is relatively precise, then the supply function's intercept is more responsive to the private signal than when the private signal is relatively noisy.

Both (i) and (ii) are consistent with H4.

This evidence suggests a possible explanation of the main result of our experiment. When costs are interdependent, two signals are informative: the private signal and the price. When the signal is relatively informative (TCL), then participants pay most of their attention to the private signal and do not focus on the informational content of the price. The signal is readily available, whereas using the informational value of the cost correlation requires more attention. However, when the signal is very noisy (in TCH), then participants focus more on the informativeness of the price to estimate the marginal cost. This leads to greater market power and can potentially explain our main result: that if the private signal is relatively noisy then interdependent costs lead to market power (TUH vs. TCH), whereas if the signal is relatively precise then information frictions do not lead to greater market power (TUL vs. TCL). This observation can be related to the theoretical literature of rational inattention (i.e., Kacperczyk et al. 2014), and to the experimental results found in Ngangoué and Weizsäcker (2018).

5.2 Heterogeneity of behaviour and mixture models

In this sub-section, we use price impact and its related supply function variable – the supply function slope – as measures of market power for analysing the heterogeneity in individual choices. Here we use the last five rounds since in our analysis we include a best-response benchmark, which are more relevant during the last five rounds when participants have learnt the most during the experiment.

The assessment of market power for each subject reveals substantial heterogeneity at the individual level in all treatments during the last 5 rounds, as can be seen in Figure 3.

Figure 3: *Histograms for price impact averaged over the last 5 rounds for each subject in each treatment, and benchmarks.*

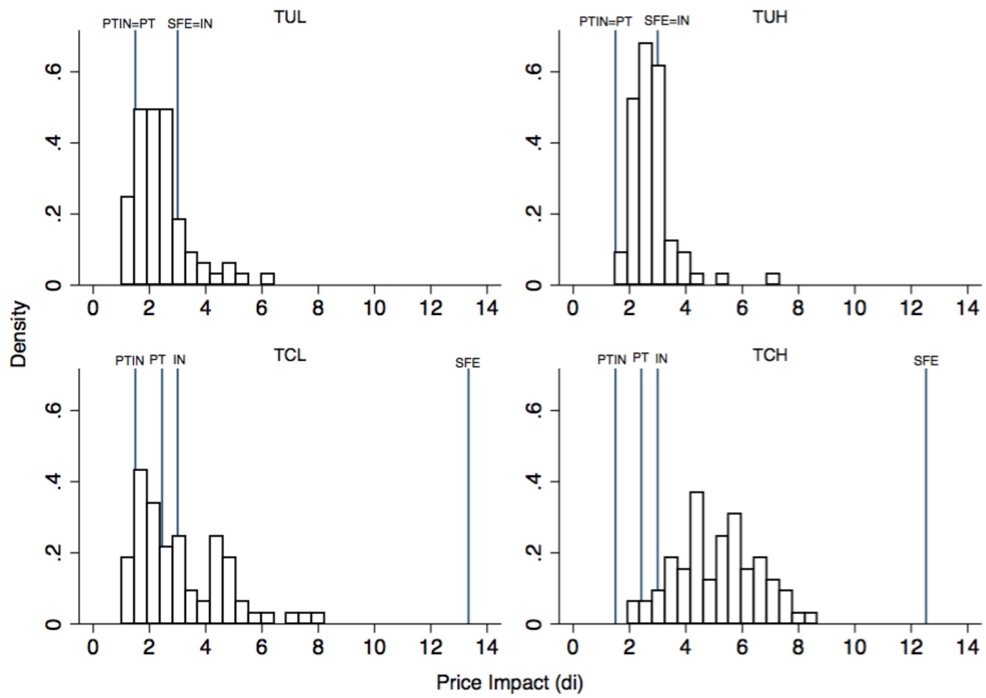

Note. The unit of analysis is the price impact averaged over the last 5 rounds for each subject. Hence, there are 72 observations in each treatment. Note that each observation averages the price impact of a subject in the last 5 rounds, and therefore, there may be considerable variance in individual behaviour over these last 5 rounds.

The figure shows that market power in the uncorrelated costs treatments is more homogeneous than in the interdependent costs treatments. We test this formally using Levene's test of homogeneity of variances where the unit of analysis is the same as in Figure 3, and obtain that there are no significant differences in the variance between TUL and TUH (p-value=0.104), while we reject that variances are the same in TUL and TCL (p-value=0.000), and TUH and TCH (p-value=0.000). In addition, we observe that several of the benchmarks combined might explain the observed heterogeneity in individual choices during the last 5 rounds.

To go one step further in our analysis of heterogeneity, we estimate a finite mixture model. As a measure of behaviour, we use the supply function slopes submitted by each subject in the last 5 rounds in each treatment, and assume that they are i.i.d. and generated by a normal distribution with mean γ_{jt} which is equal to the prediction of a particular benchmark j at round t, and variance σ_j, which is a free parameter.[27]

[27] We use the supply function slope as our measure of market power, instead of price impact, since it is more directly related to how participants made their decisions in each round and what participants observe about rivals in the experimental feedback. Given that we also compute the best-response benchmark, it seems more natural to analyse the supply function slope rather than price impact as a measure of market power.

The theoretical benchmarks that we use to estimate the mixture model are the SFE, IN, PT and PTIN, which have been discussed in Section 3. In our mixture model, we also consider the empirical best-response strategy, which may be seen as a relevant way to behave. Note that the naïve and sophisticated best-responses are identical in the uncorrelated costs treatments. However, in the complex environment of the interdependent costs treatments, it makes more sense to use the naïve best-response strategy (which ignores the informational content included in the market price) than the sophisticated best-response.[28]

For each at a given round, the naïve best-response is obtained by first computing the average supply function conditional on the private signal of all rivals that i has faced in previous rounds. These are then used to calculate the BR_N of each subject in a given round. [29] Hence, $j \in$ *{SFE, IN, PT, PTIN, BR_N}*. Recall that for the two uncorrelated costs treatments, the PT and SFE are identical to the $PTIN$ and IN respectively.

A finite mixture model allows each seller to use one of the benchmarks just described. Denote P_j the probability that any seller uses the j benchmark, such that $\sum_j P_j = 1$. The likelihood contribution for subject i *is*

$$L_i = \sum_j P_j \prod_{t=14}^{18} f(SlopePQ_{it}|j), \tag{6}$$

where $f(SlopePQ_{it}|j)$ is the conditional probability that subject i in round t chooses $SlopePQ_{it}$ given that it follows benchmark j. The sample log-likelihood in each treatment is then $logL = \sum_i^n log(L_i)$. Table 10 reports the maximum likelihood results for the general mixture model in each treatment. In addition to estimating the general model for each treatment, we have tested whether any of the possible nested models might provide a more parsimonious representation of the heterogeneity of our choices. Results of likelihood ratio tests (at the 1% level) favour the most general model which allows all benchmarks to be present in each treatment.

[28]The computational requirements that are needed to calculate the sophisticated best-response in the interdependent costs treatments are harder than for the equilibrium (refer to Online Appendix A.2, Lemma 1). This is because a subject must consider the asymmetric strategies of rivals, estimate the rivals' response to the private signal, use them optimally to calculate the inference effect, and combined it with the strategic effect, as explained in Section 3. In addition, in TCH we find that the naïve and sophisticated best-response strategies are practically identical.

[29]For further details of the empirical best-response strategies in each treatment, refer to Online Appendix C.5.

Table 10: *Maximum likelihood estimation for the general mixture models in each treatment.*

Treatments	TUL	TUH
Model	PT; BR; SFE	PT; BR; SFE
-log L	2,872.93	2,640.47
σ_{PT}	1.65 (0.054)	1.12 (0.079)
σ_{BR}	2.81 (0.25)	10.48 (0.63)
σ_{SFE}	10.60 (0.65)	1.92 (0.073)
P_{PT}	0.54 (0.031)	0.30 (0.029)
P_{BR}	0.26 (0.034)	0.17 (0.023)
P_{SFE}	0.20 (0.028)	0.53 (0.032)
Treatments	**TCL**	**TCH**
Model	PTIN; PT; IN; BR_N:BR_S;SFE	PTIN; PT; IN; BR_N:BR_S;SFE
-log L	2,890.19	3,386.41
σ_{PTIN}	1.40 (0.059)	1.30 (0.21)
σ_{PT}	0.67 (0.063)	0.16 (0.020)
σ_{IN}	9.76 (1.71)	0.60 (0.068)
σ_{BR_N}	3.54 (0.17)	4.45 (0.18)
σ_{SFE}	8.17 (0.90)	8.90 (0.42)
P_{PTIN}	0.096 (0.020)	0.031 (0.010)
P_{PT}	0.043 (0.018)	0.048 (0.014)
P_{IN}	0.33 (0.029)	0.018 (0.0084)
P_{BR_N}	0.43 (0.033)	0.64 (0.031)
P_{SFE}	0.097 (0.018)	0.26 (0.028)

Notes. The likelihood is computed using the supply function slope choices for periods [14,18] for all participants in each treatment, thus there are 360 observations in each treatment. The standard error of each parameter is reported in parentheses.

The results of the previous table exhibit the following patterns. Increasing the correlation while keeping the variance ratio fixed leads to both a lower proportion of the population to be price-takers and a lower proportion of participants being consistent with SFE. At the same time, the proportion of the participants who best-respond rises. Increasing the variance ratio for both uncorrelated and interdependent costs yields a lower proportion of participants being consistent with price-taking behaviour, whereas the percentage of SFE-consistent behaviour rises. In addition, when costs are interdependent the proportion of informationally naïve participants is greater when the variance ratio is low.

This finding means that, even though there should be a similar increase in market power due to a combination of either a high correlation and low variance ratio, or non-zero but low correlation and high variance ratio, these two scenarios lead to different experimental results. In contrast to what is predicted by the SFE, a high correlation and low variance ratio (TCL) does not lead to a significant difference in market power in relation to TUL since in both treatments more than half of the participants deviate from perfect rationality, either because they are price-takers (in TUL), or because they ignore the informational content of the market price (in TCL), or both. However, a non-zero but low correlation and a high variance ratio (TCH) brings about a significant degree of market power in relation to when costs are uncorrelated (TUH) since the low quality of their private signal makes participants pay more attention to rivals' strategies, are more strategic, and place more weight on the price than when the signal more precise. Consistent with this we find that a large percentage of participants who either play the SFE or best-respond.

6 Robustness

In this section, we report the results of additional sessions that we conducted in order to further understand the relationship between information frictions and market power.

6.1 Bayesian updating

One could argue that participants fail to engage in Bayesian updating and that this is why, in treatment TCL, they fail to understand that the market price is informative about costs. To explore this possibility further and to assist participants in the decision-making process, we conducted an additional treatment with 24 participants with two independent groups in the modified TCL treatment, hereafter TCL-Bayesian. The experimental design features were 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 $E[\theta_i \mid s_i]$ and $E[\theta_j \mid s_i]$ for $i \neq j$).[30] Second, we explicitly asked each subject to think about what her rivals would do and provided a simulation tool that participants 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 rounds.[31] The following table presents the summary statistics of behaviour and outcomes in the TCL-Bayesian treatment.

Table 11: *Market power and information frictions in the TCL-Bayesian treatment, and comparisons.*

TCL	**TCL-Bayesian**	**Treatment**
3.35	3.13	Price impact
(2.76)	*(2.13)*	*(s.d. Price impact)*
p=899.25+7.79X	P=898.46+7.38X	Average inverse supply function
(291.11), (6.98)	*(212.94), (6.50)*	*(s.d. InterceptPQ), (s.d. SlopePQ)*

Note. The first row shows the average, while the second in brackets and italics is the standard deviation *(s.d.)*. The number of observations in treatment TCL is 1,800, while in treatments TCH-Bayesian there are 360 observations, respectively.

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

[31]We reduced the number of rounds because our three modifications increased the experiment's duration.

We find that the average supply function and outcomes in the robustness session are similar to the averages of the baseline treatments, presented in Table 5, that correspond to the TCL treatment. Our results in this treatment suggest that they are *not* simply due to a bias related to Bayesian updating.

6.2 Increasing the noisiness of the private signal

We could ask how increasing the noisiness of the private signal (or the variance ratio) affects the competitiveness of markets. To answer this question, we conducted a treatment with correlation $\rho = 0.175$ and variance ratio $\phi = 3600$ (hereafter TCHH), in which the private signal is practically uninformative. Treatment TCHH can be then easily compared with treatment TCH (that has $\rho = 0.175$ and $\phi = 36$). The SFE predicts that the TCHH is slightly less competitive than treatment TCH, with a theoretical price impact of $d = 13.49$. The SFE predicts only a small increase in market power (an increase of 0.97 in price impact due to an increase in the variance ratio by a factor of 100). We ran the TCHH treatment with 72 participants, which had 6 independent groups of 12 sellers each. The experimental procedures were identical to those of the main experiment, described in Section 4. Table 12 displays the market power and information friction results in the TCHH treatment.

Table 12: *Market power and information frictions in the TCHH treatment, and comparisons.*

TCH	**TCHH**	**Treatment**
4.87	5.44	Price impact
(2.80)	*(3.05)*	*(s.d. Price impact)*
p=790.74+11.54X	P=769.59+12.85X	Average inverse supply function
(237.54), (7.38)	*(341.90), (8.04)*	*(s.d. InterceptPQ), (s.d. SlopePQ)*

Note. The first row shows the average, while the second in brackets and italics is the standard deviation *(s.d.)*. The number of observations in each treatment is 1,800.

We find that average market power is slightly larger in the TCHH treatment than in the TCH treatment, however, when tested formally, we find that the two treatments have the same degree of market power ($p = 0.200$, Mann-Whitney rank-sum test where the unit of observation is the independent group). Thus, we find that increasing the variance ratio from high to very high does not increase market power in a significant way, and this is broadly consistent with the qualitative predictions of the SFE model.

7 Concluding Remarks

We have analysed bidding behaviour in an experiment that reflects some of the complexity of real-world markets where bidders compete in supply functions, have incomplete information about their costs, and receive a noisy private signal. We used the supply function equilibrium (SFE) prediction and its comparative statics as benchmark when evaluating the experimental results. Alternative benchmarks resulting from relaxing the perfect rationality assumption are also considered. Our experiment employed a between-subjects design with four treatments that vary in the correlation among costs and the relative precision of the private signals in relation to the fundamental.

Our experimental results document certain causal relations between information frictions and market power in the context of supply function competition. Average market power in markets with uncorrelated costs is the same regardless of the precision of private signals. However, with interdependent costs and relatively noisy private signals, average market power rises significantly, whereas average market power does not increase if costs are interdependent and private signals are relatively precise.

Our analysis of observed individual heterogeneity in experimental market power suggests an explanation of our findings in terms of the reasoning processes of our participants. Indeed, our results are consistent with rational inattention models where agents have a limited capacity to interpret information signals. In both treatments with relatively precise signals, we find that more than half of participants deviate from perfect rationality either because they are price-takers or because they ignore the information content of the market price, or both. With relatively noisy signals, we find that, in the absence of cost correlation, there is a higher percentage of the participants that are price-takers. With interdependent costs, which is the case where agents should infer information from prices, there is a larger number of participants who exhibit a behaviour consistent either with the SFE or with a best response to the empirical distribution of observed actions.

If our experiment reflects behaviour in real markets, characterised by schedule (supply or demand) competition, then competition authorities should be weary of situations where agents have imprecise signals about the fundamentals and the variance of the fundamentals is not too large. In this case, which may apply to some wholesale electricity markets (Holmberg and Wolak 2016), traders may infer information from the price and market power may be increased, and mitigation measures may be appropriate. However, this behaviour is not necessarily indicative of collusion since it is the outcome of a non-cooperative equilibrium. Therefore, competition authorities have to be alert to the information structure of the market and be careful not to confound collusion with an environment conducive to substantial market power. Indeed, authorities when designing transparency measures have to check their impact on the variance ratio and the correlation of values/costs. If more transparency translates into lower noise in the private signals then this will have a pro-competitive effect since agents will tend to ignore the informational content of the price. Otherwise, an increase in transparency may backfire by inducing a higher level of market power.

Our experiment suggests a few open questions for future research, both experimental and theoretical. Future work could explore mechanisms by which participants learn how to improve information extraction from the price (e.g., asking participants to come back to the laboratory a few days later -experienced bidders; replicating the experiment with professional traders; extending the number of rounds). Increased capacity to learn from the price would imply a potential higher exercise of market power. 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 in markets characterised by supply function competition and private information.

Universitat Ramon Llull, ESADE Business School.

Institut d'Anàlisi Econòmica (CSIC) and Barcelona GSE.

IESE Business School.

Additional Supporting Information may be found in the online version of this article.

Online Appendix A. Theoretical benchmarks: equilibrium characterisation for the various benchmarks and theoretical best-response strategy.

Online Appendix B. Experimental instructions, comprehension test, screenshots, and post-experimental questionnaire.

Online Appendix C. Additional empirical results: analysis of the last 5 rounds of the experiment, time trends, efficiency of the allocations, and details of the best-response strategy calculations.

Online Appendix D. Analysis of participants' demographic, cognitive information, and post-experimental questionnaire.

References

[1] 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.

[2] 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.

[3] Bardsley, N., Cubitt, R., Loomes, G., Moffatt, P., Starmer, C., and Sugden, R., (2009), *Experimental Economics: Rethinking the Rules*, Princeton University Press, Princeton.

[4] Bolle, F., Grimm, V., Ockenfels, A., and del Pozo, X., (2013), "An Experiment on Supply Function Competition," *European Economic Review*, 63(C): 170-185.

[5] Brandts, J., Pezanis-Christou, P., and Schram, A., (2008), "Competition with Forward Contracts: A Laboratory Analysis Motivated by Electricity Market Design," *The Economic Journal*, 118(525), 192-214.

[6] 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.

[7] 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.

[8] Camerer, C.F., and Fehr, E., (2006), "When Does "Economic Man" Dominate Social Behavior?," *Science*, 311: 47-52.

[9] Carrillo, J. D., and Palfrey, T., (2011), "No Trade," *Games and Economic Behavior*, 71(1), 66-87.

[10] 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.

[11] Charness, G., and Levin, D., (2009), "The Origin of the Winner's Curse: A Laboratory Study," *American Economic Journal: Microeconomics*, 207-236.

[12] Crawford, V. P., and Iriberri, 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.

[13] Dufwenberg, M., and Gneezy, U., (2000), "Price Competition and Market Concentration: An Experimental Study," *International Journal of Industrial Organization,* 18, 7-22.

[14] 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.

[15] 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.

[16] Esponda, I., and Vespa, E., (2014), "Hypothetical Thinking and Information Extraction in the Laboratory," *American Economic Journal: Microeconomics,* 6(4): 180-202.

[17] Eyster, E., and Rabin, M., (2005), "Cursed Equilibrium," *Econometrica*, 73(5), 1623-1672.

[18] Fischbacher, U., (1999), "z-Tree 1.1.0: Experimenter's Manual," University of Zurich, Institute for Empirical Research in Economics.

[19] Fraley, C., and Raftery, A.E., (2002), "Model-based Clustering, Discriminant Analysis, and Density Estimation," *Journal of the American Statistical Association*, 97, 611-631.

[20] Fréchette, G. R., (2015), "Laboratory Experiments: Professionals versus Students," in *Handbook of Experimental Economic Methodology*, Guillaume R. Fréchette and Andrew Schotter (eds.), Oxford University Press, 360-390.

[21] Goeree, J., and Offerman, T., (2003), "Winner's Curse Without Overbidding," *European Economic Review,* 47(4), 625-644.

[22] Holmberg, P., and Wolak, F.A., (2015), "Electricity Markets: Designing Auctions Where Suppliers Have Uncertain Costs," *Cambridge Working Papers in Economics,* No. 1541.

[23] Holt, C. A., (1985), "An experimental test of the consistent-conjectures hypothesis," *The American Economic Review*, 75(3), 314-325.

[24] 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.

[25] 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.

[26] Isaac, R. M., & Reynolds, S. S., (1988), "Appropriability and market structure in a stochastic invention model," *The Quarterly Journal of Economics*, 103(4), 647-671.

[27] 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.

[28] 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.

[29] Kagel, J. H. and Levin, D., (2015) "Auctions: A Survey of Experimental Research," in *The Handbook of Experimental Economics*, Volume 2, J. H. Kagel and A. E. Roth (eds). Princeton: Princeton University Press.

[30] 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.

[31] Klemperer, P., and Meyer, M., (1989), "Supply Function Equilibria in Oligopoly under Uncertainty," *Econometrica*, 57(6): 1243–1277.

[32] Kyle, A. S., (1985), "Continuous Auctions and Insider Trading," *Econometrica,* 53(6): 1315-1335.

[33] 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.

[34] 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.

[35] Morgan, J., Orzen, H., & Sefton, M., (2006), "An experimental study of price dispersion," *Games and Economic Behavior*, 54(1), 134-158.

[36] Nagel, R., (1995), "Unraveling in Guessing Games: An Experimental Study," *American Economic Review*, 85(5): 1313-1326.

[37] 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.

[38] Potters, J., and Suetens, S., (2013), "Oligopoly experiments in the current millennium," *Journal of Economic Surveys*, 27(3), 439-460.

[39] 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.

[40] Samuelson, W. F., and Bazerman, M. H., (1985), "Negotiation under the Winner's Curse," *Research in Experimental Economics*, 3, 105-38.

[41] 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.

[42] Selten, R., Pittnauer, S., and Hohnisch, M., (2012), "Dealing with Dynamic Decision Problems when Knowledge of the Environment is Limited: An Approach Based on Goal Systems," *Journal of Behavioral Decision Making*, 25(5), 443-457.

[43] Vives, X., (2008*), Information and Learning in Markets: The Impact of Market Microstructure*. Princeton University Press.

[44] Vives, X., (2011), "Strategic Supply Function Competition with Private Information," *Econometrica,* 76(3): 1919-1966.