INFORMATION, MARKET POWER AND PRICE VOLATILITY

By

Dirk Bergemann, Tibor Heumann, and Stephen Morris

September 2019

COWLES FOUNDATION DISCUSSION PAPER NO. 2200

COWLES FOUNDATION FOR RESEARCH IN ECONOMICS
YALE UNIVERSITY
Box 208281
New Haven, Connecticut 06520-8281

http://cowles.yale.edu/

Information, Market Power and Price Volatility*

Dirk Bergemann† Tibor Heumann‡ Stephen Morris§

September 15, 2019

Abstract

We consider demand function competition with a finite number of agents and private information. We show that any degree of market power can arise in the unique equilibrium under an information structure that is arbitrarily close to complete information. In particular, regardless of the number of agents and the correlation of payoff shocks, market power may be arbitrarily close to zero (so we obtain the competitive outcome) or arbitrarily large (so there is no trade in equilibrium). By contrast, price volatility is always less than the variance of the aggregate shock across all information structures.

Jel Classification: C72, D43, D44, D83, G12.

Keywords: Demand Function Competition, Supply Function Competition, Price Impact, Market Power, Incomplete Information, Price Volatility .

*We gratefully acknowledge financial support from NSF SES 1459899. We would like to thank Piotr Dworczak, Drew Fudenberg, Meg Meyer, Alessandro Pavan, Marzena Rostek, Jean Tirole, Xavier Vives, and Marek Weretka, as well as many seminar participants, for informative discussions.

†Department of Economics, Yale University, New Haven, CT 06520, U.S.A., dirk.bergemann@yale.edu.

‡Department of Economics, HEC–Montreal, Montreal QC H3T 2A7, Canada; tibor.heumann@hec.ca.

§Department of Economics, MIT, Cambridge, MA 02139, USA, semorris@mit.edu

1 Introduction

1.1 Motivation and Results

Models of demand function competition (or equivalently, supply function competition) are a cornerstone to the analysis of markets in industrial organization and finance. Economic agents submit demand functions and an auctioneer chooses a price that clears the market. Demand function competition is an accurate description of many important economic markets, such as treasury auctions or electricity markets. In addition, it can be seen as a stylized representation of many other markets, where there may not be an actual auctioneer but agents can condition their bids on market prices and markets clear at equilibrium prices.

Under complete information, there is a well known multiplicity of equilibria under demand function competition (see Grossman (1981), Klemperer and Meyer (1989)). In particular, under demand function competition, the degree of market power – which measures the distortion of the allocation as a result of strategic withholding of demand – is indeterminate. This indeterminacy arises because, under complete information, an agent is indifferent about what demand to submit at prices that do not arise in equilibrium. Making the realistic assumption that there is incomplete information removes the indeterminacy because every price can arise with positive probability in equilibrium. We therefore analyze demand function competition under incomplete information (Vives (2011)). We consider a setting where a finite number of agents have linear-quadratic preferences over their holdings of a divisible good, and the marginal utility of an agent is determined by a payoff shock; we restrict attention to symmetric environments (in terms of payoff shocks and information structures) and symmetric linear Nash equilibria.

The outcome of demand function competition under incomplete information will depend on the fundamentals of the economic environment - the number of agents and the distribution of payoff shocks - but also on which information structure is assumed. However, it will rarely be clear what would be reasonable assumptions to make about the information structure. We therefore examine if it is possible to make predictions about outcomes under demand function competition in a given economic environment that are robust to the exact modelling of the information structure.

Our first main result establishes the impossibility of robust predictions about market power. We show that any degree of market power can arise in the unique equilibrium under an information structure that is arbitrarily close to complete information. In particular, regardless of the number of agents and the correlation of payoff shocks, market power may be arbitrarily close to zero (so we obtain the competitive outcome) or arbitrarily large (so there is no trade in equilibrium). The reason is that, when there is incomplete information, prices convey information to agents. The slope of the demand function an agent submits will then depend on what information is being revealed, and this will pin down market power in equilibrium.

Given the sharp indeterminacy in the level of market power induced by the information structure, it is natural to ask what predictions— if any— hold across *all* information structures.

Our second main result shows that – for any level of market power – price volatility is always (that is, regardless of the information structure) less than the price volatility that is achieved by an equilibrium under complete information. A direct corollary of our result is that price volatility is less than the variance of the average shock across agents across all information structures. Hence, we show that it is possible to provide sharp bounds on some equilibrium statistics, which hold across all information structures.

The information structures giving rise to extremal outcomes are special as they are constructed to simplify the Bayesian updating when solving for the Nash equilibrium, so they do not necessarily have an immediate interpretation. Thus, one could have expected that, within the class of "natural" information structures, market power is well behaved in the following sense: (i) only a large amount of asymmetric information can lead to a large increase in market power (that is, large increase with respect to the benchmark provided by Klemperer and Meyer (1989)) , and (ii) market power is related to the amount of interdependence in the payoff environment (that is, the correlation of the payoff shocks). However, neither of these conjectures is true, even when limiting attention to "natural" information structures. We discuss the basic insight–later formalized by the proof of Theorem 1– now in some detail.

We can always decompose agents' payoff shocks into idiosyncratic and common components. If there was common knowledge of the common component, but agents observed noisy signals of their idiosyncratic components, there would be a unique equilibrium and we can identify the market power as the noise goes to zero. If instead there was common knowledge of the idiosyncratic components, but each agent observed a different noisy signal of the common component, there will be a different unique equilibrium and a different market power in the limit as the noise goes to zero. In the latter case, unlike in the former case, higher prices will reveal positive information about the value of the good to agents and, as a result, agents will submit less price elastic demand functions and there will be high market power. More generally, if agents have distinct noisy but accurate signals of the idiosyncratic and common components of payoff shocks, then market power will be determined by the relative accuracy of the signals, even when all signals are very accurate.

We interpret our first main result and our parameterized information structures as establishing that the indeterminacy of market power is not an artifact of particular modelling choices, such as complete information, but rather is an intrinsic feature of the game. If economic agents interact in a market where demands can be conditioned on prices, then there can be an extreme sensitivity to the inferences that market participants draw from prices so that it will not be possible to make ex ante predictions about market power.

On the other hand, we interpret our second main result as showing that the same economic feature that gives rise to indeterminacy of market power – conditioning demand on market prices – puts tight bounds on price volatility that do not hold in other economic environments.

The tight bounds on price volatility and indeterminacy of market power are important features of demand function competition. Our methodology allows us to make an exact comparison of outcomes under demand function competition (under any information structure) with what could have arisen under alternative trading mechanisms. We illustrate this by showing that under Cournot competition, our qualitative results are reversed: market power is now completely determined by the number of firms (and independent of the information structure), by contrast the bounds on price volatility are now very weak.

1.2 Related Literature

The multiplicity of equilibria in demand function competition under complete information was identified by Wilson (1979), Grossman (1981) and Hart (1985), see also Vives (1999) for a more detailed account. Klemperer and Meyer (1989) emphasized that the complete information multiplicity was driven by the fact that agents' demand at non-equilibrium prices was indeterminate. They showed that introducing noise that pinned down best responses lead to a unique equilibrium and thus determinate market power. And they showed that the equilibrium selected was independent of the shape of the noise, as the noise became small. They were thus able to offer a compelling prediction about market power. Our results show that their results rely on a maintained private values assumption, implying that agents cannot learn from prices. We replicate the Klemperer and Meyer (1989) finding that small perturbations select a unique equilibrium but - by allowing for the possibility of a common value component of values - we can say nothing about market power in the perturbed equilibria.

Vives (2011) pioneered the study of asymmetric information under demand function competition, and we work in his setting of linear-quadratic payoffs and interdependent values. He studied a particular class of information structures where each trader observes a *one-dimensional normal* noisy signal of his own payoff type. The noise is represented by an idiosyncratic error term around his payoff-type. We study what happens for *all multi-dimensional normal* information structures. In particular, we allow each trader to observe signals about the other traders' payoff type. Moreover, each multivariate signal can be either noise-free or noisy, and the noise term can have idiosyncratic or common components.

We show that the impact of asymmetric information on the equilibrium market power can even be larger than the ones derived from the one-dimensional signals studied in Vives (2011). Our results reverse some of the comparative statics and bounds that are found using the specific class of one-dimensional signal structures. In particular, in this paper but not in Vives (2011) market power can be large even when the amount of asymmetric information is small; this holds regardless of the number of players, or the correlation of the payoff shocks. Rostek and Weretka (2012) and Rostek and Weretka (2015) relaxed the symmetry in the correlation of payoff shocks across agents, while maintaining the one-dimensional signal model of Vives (2011). This allows for a rich structure in the induced correlation of signals and large variation in market power. In our setting, the variation in market power arises through multidimensional signals despite maintaining symmetry in the correlation of payoff shocks across agents.

Our "anything goes" result for market power has the same flavor as abstract game theory results establishing that fine details of the information structure can be chosen to select among multiple rationalizable or equilibrium outcomes of complete information games (Rubinstein (1989) and Weinstein and Yildiz (2007)). But this work relies on extremal information structures and, in particular, a "richness" assumption in Weinstein and Yildiz (2007), which in our context would require the strong assumption that there exist "types" with a dominant strategy to submit particular demand functions.[1] Our results do not require richness and do exploit the structure of the demand function competition game. And we show that while "anything goes" for market power, we also show that the complete information equilibria are "extremal" in the sense that they generate the maximum price volatility; there is no counterpart for this result in the class of general games studied by Weinstein and Yildiz (2007).

2 Model

2.1 Payoff Environment

There are N agents who have demand for a divisible good. The utility of agent $i \in \{1, ..., N\}$ who buys $q \in \mathbb{R}$ units of the good at price $p \in \mathbb{R}$ is given by:

$$u_i(\theta_i, q_i, p) \triangleq \theta_i q_i - p q_i - \frac{1}{2} q_i^2, \tag{1}$$

where $\theta_i \in \mathbb{R}$ is the payoff shock of agent i. The payoff shock θ_i describes the marginal willingness to pay of agent i for the good at $q = 0$. The payoff shocks are symmetrically and normally distributed across the agents, and for any i, j:

$$\begin{pmatrix} \theta_i \\ \theta_j \end{pmatrix} \sim \mathcal{N}\left(\begin{pmatrix} \mu_\theta \\ \mu_\theta \end{pmatrix}, \begin{pmatrix} \sigma_\theta^2 & \rho_{\theta\theta}\sigma_\theta^2 \\ \rho_{\theta\theta}\sigma_\theta^2 & \sigma_\theta^2 \end{pmatrix} \right),$$

where $\rho_{\theta\theta}$ is the correlation coefficient between the payoff shocks θ_i and θ_j.

[1] Weinstein and Yildiz (2011) provides a similar result without requiring a richness condition. However, their results apply only for games with one-dimensional strategies, and continuous and concave payoffs.

The realized average payoff shock among all the agents is denoted by:

$$\overline{\theta} = \frac{1}{N} \sum_{i \in N} \theta_i, \tag{2}$$

and the corresponding joint distribution of θ_i and $\overline{\theta}$ is given by

$$\begin{pmatrix} \theta_i \\ \overline{\theta} \end{pmatrix} \sim \mathcal{N}\left(\begin{pmatrix} \mu_\theta \\ \mu_\theta \end{pmatrix}, \begin{pmatrix} \sigma_\theta^2 & \frac{1+(N-1)\rho_{\theta\theta}}{N}\sigma_\theta^2 \\ \frac{1+(N-1)\rho_{\theta\theta}}{N}\sigma_\theta^2 & \frac{1+(N-1)\rho_{\theta\theta}}{N}\sigma_\theta^2 \end{pmatrix} \right).$$

The supply of the good is given by an exogenous supply function $S(p)$ as represented by a linear inverse supply function with $\alpha, \beta \in \mathbb{R}_+$:

$$p(q) = \alpha + \beta q. \tag{3}$$

For notational simplicity, we normalize the intercept α of the affine supply function to zero.[2]

2.2 Information Structure

Each agent i observes a multidimensional signal $s_i \in \mathbb{R}^J$ about the payoff shocks:

$$s_i \triangleq (s_{i1}, ..., s_{ij}, ..., s_{iJ}).$$

The joint distribution of signals and payoff shocks $(s_1, ..., s_N, \theta_1, ..., \theta_N)$ is symmetrically and normally distributed. We discuss specific examples of information structures in the following sections.

2.3 Demand Function Competition

The agents compete via demand functions. Each agent i submits a demand function $x_i : \mathbb{R} \times \mathbb{R}^J \to \mathbb{R}$ that specifies the demanded quantity as a function of the market price $p \in \mathbb{R}$

[2] The general affice case with $\alpha \neq 0$ is equivalent to a market with a different mean payoff shock. More precisely, considering $\alpha \neq 0$ is mathematically equivalent to considering a model in which $\tilde{\alpha} = 0$ and $\tilde{\theta}_i = \theta_i - \alpha$.

and the private signal $s_i \in \mathbb{R}^J$, denoted by $x_i(p, s_i)$. The Walrasian auctioneer sets a price p^* such that the market clears for every realization of signal profiles s:

$$p^* = \beta \sum_{i \in N} x_i(p^*, s_i). \tag{4}$$

If a market clearing price satisfying (4) does not exist, then we assume that there is a market shutdown, thus $q_1 = \cdots = q_N = 0$. We note that in the class of linear equilibria we study, a market clearing price will always exist, so the aforementioned rule is mentioned only for completeness. We denote by $x_i(p, s_i)$ the quantity demand by agent i as a function of the received signal s_i and the market price p.

We study the Nash equilibrium of the demand function competition game. The strategy profile $(x_1^*, ..., x_N^*)$ forms a Nash equilibrium if:

$$x_i^* \in \underset{\{x_i:\mathbb{R}^{J+1}\to\mathbb{R}\}}{\arg\max}\ \mathbb{E}\left[\theta_i x_i(p^*, s_i) - p^* x_i(p^*, s_i) - \frac{x_i(p^*, s_i)^2}{2}\right], \tag{5}$$

where

$$p^* = \beta(x_i(p^*, s_i) + \sum_{j \neq i} x_j^*(p^*, s_j)).$$

We say a Nash equilibrium $(x_1^*, ..., x_N^*)$ is linear and symmetric if there exists a vector of coefficients $(c_0, ..., c_J, m) \in \mathbb{R}^{J+2}$ such that for all $i \in N$:

$$x_i(p, s_i) = c_0 + \sum_{j \in J} c_j s_{ij} - mp.$$

Thus, the private information s_i of agent i determines the intercept of the demand curve, whereas the slope m of the demand curve, determined in equilibrium, is invariant with respect to the signal s_i. Throughout the paper we focus on symmetric linear Nash equilibria and so hereafter we drop the qualifications "symmetric" and "linear". When we say an equilibrium is unique, we refer to uniqueness within this class of equilibria.

2.4 Equilibrium Statistics: Market Power and Price Volatility

We analyze the set of equilibrium outcomes in demand function competition under incomplete information. We frequently describe the equilibrium outcome through two central statistics of the equilibrium: *market power* and *price volatility*.

The marginal utility of agent i from consuming the q_i-th unit of the good is $\theta_i - q_i$. We define the market power of agent i as the agent's gross marginal utility minus the price divided by the equilibrium price:

$$l_i \triangleq \frac{\theta_i - q_i - p}{p}.$$

This is the natural demand side analogue of the supply side price markup defined by Lerner (1934), commonly referred to as the "Lerner's index". We define the (expected) *equilibrium market power* by:

$$l \triangleq \mathbb{E}\left[\frac{1}{N}\sum_{i\in N} l_i\right] = \frac{1}{N}\mathbb{E}\left[\frac{\sum_{i\in N}(\theta_i - q_i - p)}{p}\right]. \tag{6}$$

The market power l is defined as the expected average of the Lerner index across all agents. If the agents were price takers, then the market power would be $l = 0$.

A second equilibrium statistic of interest is *price volatility*, the variance of the equilibrium price, which we denote by:

$$\sigma_p^2 \triangleq \mathrm{var}(p). \tag{7}$$

Price volatility measures the ex ante uncertainty about the equilibrium price.

These two statistics of the equilibrium outcome, market power and price volatility, will completely describe the first and second moments of aggregate market outcomes. Thus, the equilibrium market power will determine the expected aggregate demand as well as the expected equilibrium price. Likewise, the price volatility, is the variance of the equilibrium price, and will determine the variance of the aggregate demand as well. Thus, within the linear-quadratic normal environment, these two statistics completely describe the aggregate equilibrium outcomes.

While most of our paper focuses on price volatility and market power, we will explain how our results extend to other statistics of an equilibrium outcome, such as mean and variance of the individual demand.

3 The Case of Complete Information

As previously discussed, the existence of multiple equilibria in demand function competition has been long-established. Here, we focus on the implications that the multiplicity has for the induced market power and price volatility. With complete information every agent i observes the entire vector of payoff shocks $(\theta_1, ..., \theta_N)$ before submitting his demand $x_i(p, \theta)$. The complete information setting allows us to introduce some key ideas.

The residual supply faced by agent i, denoted by $r_i(p, \theta)$, is determined by the demand functions of all the agents other than i:

$$r_i(p, \theta) \triangleq S(p) - \sum_{j \neq i} x_j(p, \theta).$$

Agent i can then be viewed as a monopsonist over his residual supply. That is, if agent i submits demand $x_i(p, \theta)$, then the equilibrium price p^* satisfies $x_i(p^*, \theta) = r_i(p^*, \theta)$ for every i. Hence, agent i only needs to determine what is the optimal point along the curve $r_i(\theta, p)$; this will determine the quantity that agent i purchases and the equilibrium price.

To compute the first order condition for agent i's demand, it is useful to define the *price impact* λ_i of agent i:

$$\frac{1}{\lambda_i} \triangleq \frac{\partial r_i(p)}{\partial p}. \tag{8}$$

This parameter is frequently referred to as "Kyle's lambda" (see Kyle (1985)). The price impact determines the rate at which the price increases when the quantity bought by agent i increases:

$$\lambda_i = \frac{\partial p}{\partial r_i(p)}$$

The first-order condition of agent i determines the equilibrium demand of agent i:

$$x_i(p^*) = \frac{\theta_i - p^*}{1 + \lambda_i}.$$

It is easy to check that λ_i determines how much demand agent i withholds to decrease the price at which he purchases the good. For example, if $\lambda_i = 0$, then agent i behaves as a price taker. As λ_i increases, agent i withholds more demand to decrease the equilibrium price. Hence, λ_i determines the incentive of agent i to withhold demand to decrease the price.

The following proposition established what happens with complete information (i.e., when every agent observes $(\theta_1, ..., \theta_N)$) in our linear-quadratic setting with interdependent values.

Proposition 1 (Equilibrium Statistics with Complete Information)

The market power and the price volatility (l, σ_p^2) are induced by some Nash equilibrium under complete information if and only if they satisfy:

$$l \geq -\frac{1}{2\beta N} \quad \textit{and} \quad \sigma_p^2 = \frac{(\beta N)^2}{(1 + \beta N + \beta N l)^2} \sigma_{\overline{\theta}}^2, \tag{9}$$

Proposition 1 can be established using the general arguments provided for Theorem 1. We therefore relegated the proof to the Appendix. The lower bound on the market power follows from the fact that a small amount of negative market power can occur only when an agent faces a downward sloping residual supply. However, the slope of the residual supply cannot be too inelastic as otherwise the agent would be able to achieve infinite utility by buying an arbitrarily large quantity at an arbitrarily low price. The relation between price volatility and market power is intuitive. As market power increases, every agent withholds more demand to lower the price. This leads to a smaller response to the payoff shocks and consequently, lower price volatility.

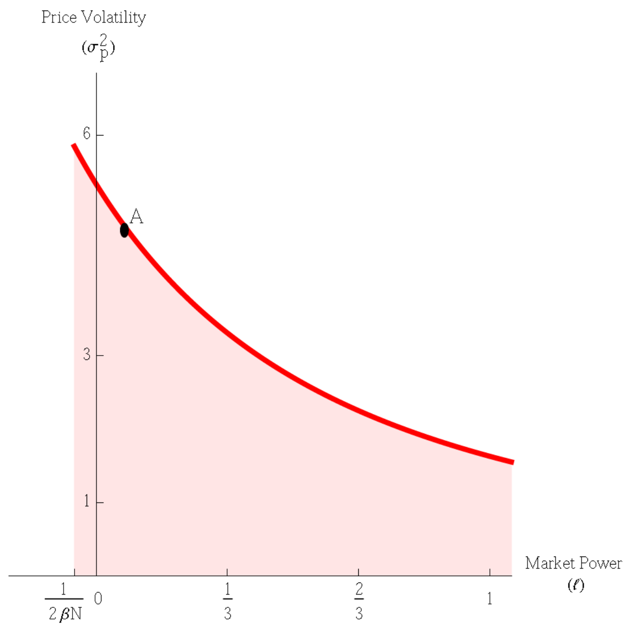

Figure 1: Set of equilibrium pairs $\left(l, \sigma_p^2\right)$ of market power and price volatility with complete information ($\beta = 3, N = 3$).

The reason for multiple equilibria is that each agent has multiple best responses. In particular, there are multiple affine demand functions agent i can submit which would intercept his residual supply at the same point (hence, inducing the same equilibrium price and quantities). Thus, agent i is indifferent between the multiple demand functions that intercept with his residual supply at the same point. Yet, the slope of the demand function submitted by agent i determines the slope of the residual supply of agent j. If agent j submits a sufficiently elastic demand, then the price impact of agent i will be close to 0; any increase in the quantity bought by agent i will be offset by a decrease in the quantity bought by agent j, keeping the equilibrium price unchanged. If agent j submits a sufficiently inelastic demand, then the price impact of agent i may be arbitrarily large; any increase in the quantity bought by agent i will be reinforced by an increase in the quantity bought by agent j, leading to arbitrarily large changes in the equilibrium price.

In Figure 1 the bold red curve plots all feasible equilibrium pairs of market power and price volatility that can be a attained under complete information. The point labelled A depicts the equilibrium outcome that would be attained under complete information if we selected the outcome using the equilibrium selection proposed by Klemperer and Meyer (1989). The results in the next Section will establish that the set of all possible pairs of market power and price volatility is the set of pairs under this red curve established by the complete information equilibrium, thus the area in light red under the boundary curve in bold red.

4 Robust Prediction of Market Power and Price Volatility

With incomplete information, market power and price volatility will be uniquely pinned down given a specific information structure. What robust predictions can be made then that do not depend on the fine details of the information structure? We will show that we cannot make any robust predictions about market power: any positive level of market power can arise as the unique equilibrium even when we restrict attention to arbitrarily small amounts of incomplete information. But we can offer a sharp prediction about the price volatility: no matter the amount of incomplete information, it cannot be higher than what happens under complete information.

We say that an information structure is $\varepsilon-$close to complete information if the conditional variance of the estimate of each payoff shock θ_j is small given the signal s_i received by agent i:

$$\forall i, j \in N, \qquad \text{var}(\theta_j | s_i) < \varepsilon. \tag{10}$$

In an information structure that is ε-close to complete information an agent can observe his own payoff shock and the payoff shock of the other agents with a residual uncertainty of at most ε. If an information structure is $\varepsilon-$close to complete information for a sufficiently small ε, then the information structure will effectively be a perturbation of complete information. We now show that any equilibrium under complete information can be selected as the unique equilibrium in a perturbation of complete information.

Theorem 1 (Equilibrium Selection)

For every $\varepsilon > 0$, and every market power and price volatility (l, σ_p^2) satisfying (9) there exists an information structure that is $\varepsilon-$close to complete information and induces (l, σ_p^2) as the unique equilibrium.

Proof. We now construct a symmetric linear Nash equilibrium using a "guess-and-verify" method. A linear demand function x^* is a symmetric Nash equilibrium if and only if it solves (4) and (5) for all i. In a linear Nash equilibrium $x^*(s_i, p^*)$ is linear in p^* and so the first order condition of (5) is given by:

$$x^*(s_i, p^*) = \frac{\mathbb{E}[\theta_i | p^*, s_i] - p^*}{1+\lambda}, \tag{11}$$

where λ is the derivative of the inverse residual supply defined earlier in (8). We now write λ explicitly:

$$\lambda \triangleq \left(\frac{\partial r_i(p)}{\partial p}\right)^{-1} = \frac{\beta}{1+\beta m(N-1)}. \tag{12}$$

The objective function of (5) is a quadratic function of $x(s_i, p)$ and the coefficient on the quadratic component is equal to $-(\lambda + 1/2)$. Thus, the second order conditions is satisfied if and only if $\lambda \geq -1/2$. It is clear that, if $\lambda < -1/2$, then the objective function is strictly convex and hence (5) does not have a solution. Therefore, there is no equilibrium satisfying $\lambda < -1/2$.

We prove the result by decomposing the payoff shock θ_i into two independent payoff shocks:

$$\theta_i \triangleq \eta_i + \phi_i, \tag{13}$$

where the sets of payoff shocks $\{\eta_i\}_{i \in N}$ are independent of the shocks $\{\phi_i\}_{i \in N}$. The shocks are jointly normally distributed:

$$\mu_\eta = \mu_\phi = \mu_\theta / 2 \quad \text{and} \quad \operatorname{corr}(\eta_i, \eta_j) = \operatorname{corr}(\phi_i, \phi_j) = \operatorname{corr}(\theta_i, \theta_j), \tag{14}$$

and the variance of the shocks are:

$$\text{var}(\phi_i) = \varepsilon \quad \text{and} \quad \text{var}(\eta_i) = \sigma_\theta^2 - \varepsilon. \tag{15}$$

It follows from (14) and (15) that:

$$\text{var}(\phi_i + \eta_i) = \sigma_\theta^2 \quad \text{and} \quad \text{cov}(\phi_i + \eta_i, \phi_j + \eta_j) = \text{cov}(\theta_i, \theta_j).$$

Thus, the joint distribution of the random variables $\{\eta_i + \phi_i\}_{i\in N}$ is equal to the the joint distribution of the original payoff shocks $\{\theta_i\}_{i\in N}$.

We first establish the argument for $l \in \mathbb{R}\backslash\{0, 1\}$; we address the special cases $l = 0$ and $l = 1$ at the end of the proof. We assume that every agent observes the realization of all shocks $\{\eta_i\}_{i\in N}$. Additionally, agent i observes a signal that is equal to a weighted difference between his shock ϕ_i and the average of the shocks $\{\phi_j\}_{j\in N}$:

$$t_i = \phi_i - (1-\gamma)\frac{1}{N}\sum_{j\in N}\phi_j, \tag{16}$$

where $\gamma \in \mathbb{R}\backslash\{0\}$ confounds the payoff shocks ϕ_i with ϕ_j for all $j \neq i$. Thus, the signals observed by agent i are summarized by $s_i = (t_i, \eta_1, ..., \eta_N)$. We remark that under this information structure:

$$\forall i, j \in N, \quad \text{var}(\theta_i|\eta_1, ..., \eta_N, s_j) = \text{var}(\phi_i|s_j) \leq \text{var}(\phi_i) = \varepsilon.$$

It follows that under this information structure (10) is satisfied.

In any linear Nash equilibrium, the equilibrium price must be a linear function of the shocks $\{\eta_i\}_{i\in N}$ and the signals $\{t_i\}_{i\in N}$. The symmetry of the conjectured equilibrium, implies that there exists constants $\hat{c}_0, \hat{c}_1, \hat{c}_2$ such that the equilibrium price satisfies:

$$p^* = \hat{c}_0 + \hat{c}_1\bar{\phi} + \hat{c}_2\bar{\eta}.^3 \tag{17}$$

[3] Recall that according to the notation already introduced $\bar{\eta} = \sum_{i\in N}\eta_i/N$ and $\bar{\phi} = \sum_{i\in N}\phi_i/N$.

Regardless of the values of $\hat{c}_0, \hat{c}_1, \hat{c}_2$, as long as $\hat{c}_1 \neq 0$, the following equation is satisfied:

$$\mathbb{E}[\theta_i | \{\eta_i\}_{i \in N}, t_i, p^*] = \theta_i. \tag{18}$$

That is, agent i can infer perfectly θ_i using the realization of the shocks $\{\eta_i\}_{i \in N}$, the signal t_i and the equilibrium price. This is because agent i can infer $\bar{\phi}$ from p^*, which in addition to t_i, allows agent i to perfectly infer ϕ_i (note that $\bar{\eta}$ is common knowledge). We now verify that there is no equilibrium in which $\hat{c}_1 = 0$. If $\hat{c}_1 = 0$, then $\mathbb{E}[\theta_i | \{\eta_i\}_{i \in N}, t_i, p^*] = \eta_i + \mathbb{E}[\phi_i | t_i]$ and so each agent will submit a demand function:

$$x^*(s_i, p^*) = \frac{\eta_i + \mathbb{E}[\phi_i | t_i] - p^*}{1 + \lambda},$$

for some λ. Therefore, market clearing implies that:

$$\beta \sum_{i \in N} \frac{\eta_i + \mathbb{E}[\phi_i | t_i] - p^*}{1 + \lambda} = p^*.$$

Thus,

$$p^* = \frac{1}{1 + \beta N + \lambda} \beta \sum_{i \in N} \eta_i + \mathbb{E}[\phi_i | t_i].$$

However, $\sum_{i \in N} \mathbb{E}[\phi_i | t_i] \propto \bar{\phi}$ (recall that $\gamma \neq 0$, so $\sum_{i \in N} t_i \neq 0$). Thus, the market clearing price p^* must depend on $\bar{\phi}$, which contradicts $\hat{c}_1 = 0$.

Using (11) and (18) we conclude that, in equilibrium, agent i buys a quantity equal to:

$$x^*(s_i, p) = q_i^* = \frac{\theta_i - p^*}{1 + \lambda}, \tag{19}$$

for some $\lambda \geq -1/2$. The market clearing condition implies that $p^* = \beta \sum q_i^*$, and so the equilibrium price is given by:

$$p^* = \frac{\beta N \bar{\theta}}{1 + \lambda + \beta N}, \tag{20}$$

for some $\lambda \geq -1/2$.

Given the expression for the equilibrium price in (20), we note that:

$$\mathbb{E}[\theta_i | p^*, t_i, \{\eta_i\}_{i\in N}] = t_i + \eta_i + (1-\gamma)\left(\frac{p^*}{\beta N}(1+\lambda+\beta N) - \bar{\eta}\right) = \theta_i.$$

Recall that in equilibrium agent i submits demand function (11), so the slope of the demand submitted by agent i is given by:

$$m = -\frac{\partial x_i(p)}{\partial p} = \frac{-1}{1+\lambda}\left(\frac{\partial \mathbb{E}[\theta_i | p^*, t_i, \{\eta_i\}_{i\in N}]}{\partial p^*} - 1\right) = \frac{1-(1-\gamma)\frac{1}{\beta N}(1+\lambda+\beta N)}{1+\lambda}. \tag{21}$$

This gives a relation between agent i's price impact (i.e. λ) and the slope of the demand function submitted by agent i (i.e. m). Using these equations and (12) we find λ in terms of the confounding parameter γ:

$$\lambda = \frac{1}{2}\left(-1 - N\beta\frac{\gamma(N-1)-1}{\gamma(N-1)+1} \pm \sqrt{\left(N\beta\frac{\gamma(N-1)-1}{\gamma(N-1)+1}\right)^2 + 2N\beta + 1}\right). \tag{22}$$

Only the positive root is a valid solution as the negative root yields λ less than $-1/2$. Hence, for every γ, there is a unique linear Nash equilibrium. In this equilibrium the price impact is equal to the positive root of (22) and each agent i submits a demand function:

$$x(s_i, p) = \frac{1}{1+\lambda}\left(t_i + \eta_i + (1-\gamma)\left(\frac{p}{\beta N}(1+\lambda+\beta N) - \bar{\eta}\right) - p\right).$$

We note that this demand function is equal to (11), so the demand of every agent satisfies the first-order condition by construction and $\lambda \geq -1/2$ so the second-order condition is also satisfied. This shows that this is a Nash equilibrium.

We note that for all $\lambda \geq -1/2$, there exists a γ that satisfies (22) (with the positive root). To verify this, note that (22) is continuous as function of γ except at $\gamma = -1/(N-1)$. Moreover, at $\gamma = -1/(N-1)$, the right limit is $+\infty$ while the left limit it $-1/2$. Since (22) is equal to 0 in the limits $\gamma \to \pm\infty$, we have that every $\lambda \in [-1/2, \infty)$ is achieved by some γ.

From (20) it follows directly that in equilibrium the price volatility is:

$$\sigma_p^2 = \left(\frac{\beta N}{1+\lambda+\beta N}\right)^2 \sigma_{\bar{\theta}}^2.$$

Using the definition of market power (6), the expression for the price (20) and the expression for q_i (19) we get that:

$$l = \frac{1}{N}\mathbb{E}\left[\frac{\sum_{i\in N}\left(\theta_i - \frac{\theta_i - \frac{\beta N\bar{\theta}}{1+\lambda+\beta N}}{1+\lambda} - \frac{\beta N\bar{\theta}}{1+\lambda+\beta N}\right)}{\frac{\beta N\bar{\theta}}{1+\lambda+\beta N}}\right] = \mathbb{E}\left[\frac{\left(\bar{\theta} - \frac{\bar{\theta} - \frac{\beta N\bar{\theta}}{1+\lambda+\beta N}}{1+\lambda} - \frac{\beta N\bar{\theta}}{1+\lambda+\beta N}\right)}{\frac{\beta N\bar{\theta}}{1+\lambda+\beta N}}\right].$$

Simplifying terms, we get that:

$$l = \frac{\lambda}{\beta N}.$$

Thus, we have that $l \geq -1/2\beta N$ and price volatility can be written as a function of market power as in Theorem 1. Recall that we have in fact constructed a linear Nash equilibria in which every agent i submits a demand function given by (11), which by construction satisfied the agent's first-order condition and we have also shown that the second-order condition is satisfied.

Finally, we address the case of $l = 0$ and $l = 1$. An equilibrium with market power $l = 1$ and price volatility given by (9) is attained by the following information structure. Every agent observes $\{\eta_i\}_{i\in N}$ and $\{(\theta_i - \bar{\phi})\}_{i\in N}$, and additionally agent i privately observes signal:

$$t_i = \bar{\phi} + \varepsilon_i,$$

where ε_i is an error term normally distributed with variance one and correlation across agents equal to $-1/(N-1)$. Since the errors are perfectly negatively correlated, this implies that $\sum_{i\in N}\varepsilon_i = 0$. Note that agent i knows η_i and $(\phi_i - \bar{\phi})$, so by virtue of observing only his own signals he does not know $\bar{\phi}$.[4] We also note that:

$$\mathbb{E}[\bar{\phi} \mid t_i] = \frac{1}{N}\sum_{j\in N} t_j. \tag{23}$$

That is, $\bar{\phi}$ can be inferred perfectly by averaging our the signals $\{t_i\}_{i\in N}$.

[4] Agent i cannot infer $\bar{\phi}$ by observing $\{(\theta_i - \bar{\phi})\}_{i\in N}$, as by construction the sum of all the signals $\{(\theta_i - \bar{\phi})\}_{i\in N}$ is equal to zero. Thus, knowing $\{(\theta_i - \bar{\phi})\}_{i\in N}$ is equivalent to knowing only $(\theta_i - \bar{\phi})$.

An equilibrium with $l = 0$ and price volatility given by (9) is attained by considering the following information structure. Every agent knows $\{\eta_i\}_{i\in N}$ and also $\bar{\phi}$, and additionally agent i privately observes a signal:

$$t_i = \phi_i - \bar{\phi} + \varepsilon,$$

where ε is a common error term, normally distributed with variance one. Note that agent i knows η_i and $\bar{\phi}$, so by virtue of observing only his own signals he only does not know $\phi_i - \bar{\phi}$. We also note that:

$$\mathbb{E}[\phi_i - \bar{\phi} \mid t_i] = t_i - \frac{1}{N}\sum_{j\in N} t_j. \tag{24}$$

That is, $\phi_i - \bar{\phi}$ can be inferred perfectly by subtracting from i the average of the signals $\{t_i\}_{i\in N}$.

Using the two aforementioned information structures one can construct a unique linear equilibrium in which $l = 1$ and $l = 0$. The equilibrium construction is the same as before, so we will not repeat the steps. However, we note that, for the noise-free signals we constructed in (16), we have that:

$$\mathbb{E}[\phi_i \mid \{t_i\}_{i\in N}] = t_i + \frac{(1-\gamma)}{\gamma}\frac{1}{N}\sum_{j\in N} t_j. \tag{25}$$

Here the two cases that are not well defined are the cases $\gamma = 0$ and $\gamma = \infty$, which intuitively corresponds to the cases in which the expectations are given by (23) and (24).[5] ■

Theorem 1 shows that all combinations of market power and price volatility that can be achieved as an equilibrium under complete information can also be achieved as a unique equilibrium in an information structure that is close to complete information. In fact, the result is stronger, every equilibrium outcome under complete information is the unique equilibrium outcome of an information structure that is close to complete information.

[5] The right-hand-side of (25) diverges in the limit $\gamma \to 0$. However, what is relevant for the analysis is that in this limit $\mathbb{E}[\phi_i \mid \{t_i\}_{i\in N}] \propto \frac{1}{N}\sum_{j\in N} t_j$, which is satisfied in (23).

The proof of Theorem 1 uses a class of information structures that we refer to as *noise-free signals* as they represent the payoff shocks in a linear combination without adding any extraneous noise. Our construction uses a (N+1)-dimensional signal. Given that there are N different shocks in our model, in any information structure that is close to complete information, agents must observe at least N different signals.

To give an intuition for the proof of Theorem 1 it is useful to consider a one-dimensional version of the noise-free signals described in the proof of Theorem 1. Consider a situation in which every agent i observes a one-dimensional signal:[6]

$$s_i = \theta_i - (1-\gamma)\frac{1}{N}\sum_{j\in N}\theta_j.$$

Every market power can arise even when agents observe one-dimensional signals.[7]

The key insight is that by changing the weights that the signal places on an agent's own payoff shock relative to the other agents' payoff shocks we can change the (perceived) degree of payoff interdependence between the agents. To see this, note that the expected payoff of agent i conditional on all signals $\{s_i\}_{i\in N}$ is given by:

$$\mathbb{E}[\phi_i \mid \{t_i\}_{i\in N}] = t_i + \frac{(1-\gamma)}{\gamma}\frac{1}{N}\sum_{j\in N} t_j.$$

By confounding the information of the agent, we can change the weight an agent places on the price to predict his own payoff shock, which in turn changes the equilibrium degree of market power. We illustrate the relation between γ and the equilibrium degree of market power in Figure 2.

[6] This is a particular case of the information structure used in the proof of Theorem 1. It corresponds to $\text{var}(\phi_i) = \sigma_\theta^2$. In this case, each agent effectively observes just a one-dimensional signal.

[7] In this case the agents will not be close to complete information since they observe only one-dimensional signals.

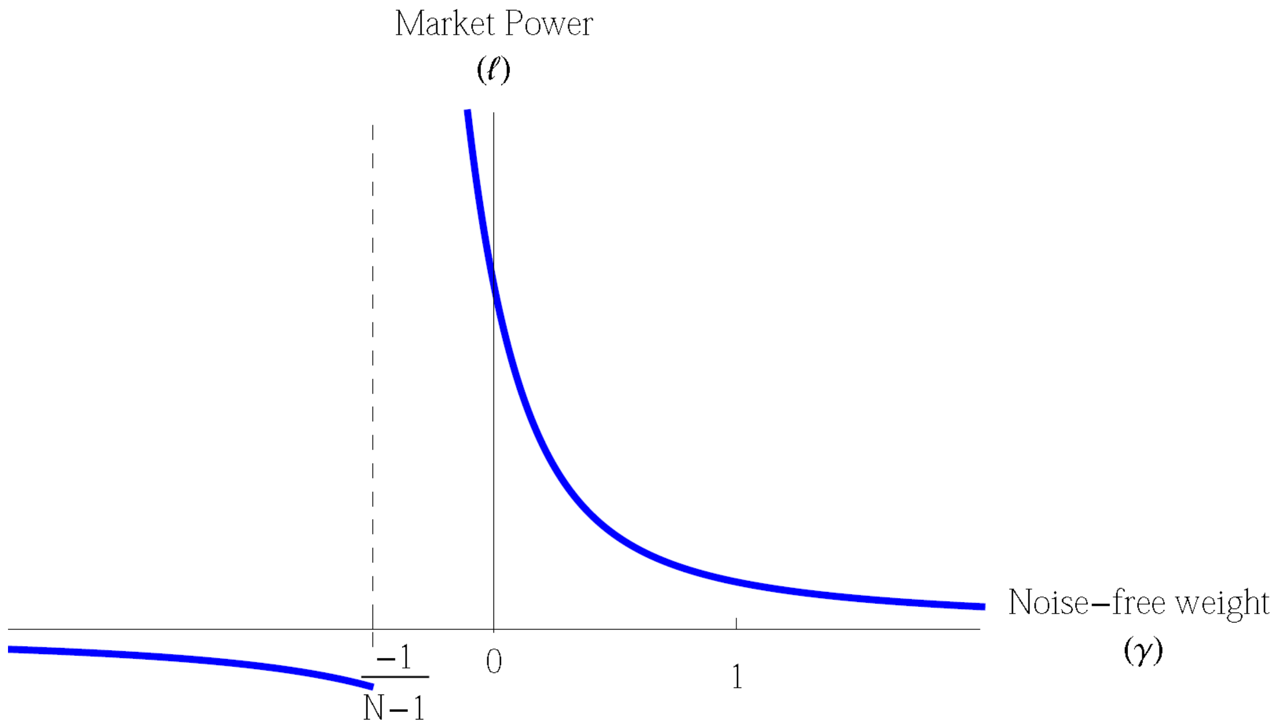

Figure 2: Set of equilibrium market power l with noise-free signal parametrized by γ.

Theorem 1 shows that, (i) all equilibrium outcomes under complete information can turn into unique equilibrium outcomes under incomplete information, and (ii) restricting attention to information structures close to complete information does not allow us to provide sharper predictions about market power and price volatility. The large indeterminacy in the set of possible outcomes suggests that it is difficult to offer robust predictions for market power under demand function competition. By contrast, it is possible to provide sharp predictions regarding price volatility with demand function competition.

Theorem 2 (Equilibria Under All Information Structures)

There exists an information structure that induces a pair of market power and price volatility (l, σ_p^2) *if and only if:*

$$l \geq -\frac{1}{2}\frac{1}{\beta N} \quad \text{and } \sigma_p^2 \leq \frac{(\beta N)^2}{(1+\beta N(1+l))^2}\sigma_{\bar{\theta}}^2. \tag{26}$$

Moreover, all feasible pairs (l, σ_p^2) *are induced by a unique equilibrium for some information structure.*

Proof. We prove necessity and sufficiency separately, and start with necessity.

We established in Theorem 1 that in any linear Nash equilibrium an agent's demand is given by (11). Adding (11) over all agents and multiplying times β we get:

$$\beta \sum_{i \in N} x^*(s_i, p^*) = \beta \sum_{i \in N} \frac{\mathbb{E}[\theta_i | s_i, p^*] - p^*}{1+\lambda}.$$

Market clearing implies that $\beta \sum_{i \in N} x^*(s_i, p^*) = p^*$. It follows that

$$p^* = \beta \sum_{i \in N} \frac{\mathbb{E}[\theta_i | s_i, p^*] - p^*}{1+\lambda}.$$

Rearranging terms we obtain

$$p^* = \frac{\beta N}{1+\lambda+\beta N}\frac{1}{N}\sum_{i \in N} \mathbb{E}[\theta_i | s_i, p^*].$$

Taking the expectation of the previous equation conditional on p^* (i.e., taking the expectation $\mathbb{E}[\cdot | p^*]$) and using the law of iterated expectations we get:

$$p^* = \frac{\beta N}{1+\lambda+\beta N}\frac{1}{N}\sum_{i \in N} \mathbb{E}[\theta_i | p^*] = \frac{\beta N}{1+\lambda+\beta N}\mathbb{E}[\frac{1}{N}\sum_{i \in N} \theta_i | p^*]. \tag{27}$$

It follows that:

$$\sigma_p^2 = \left(\frac{\beta N}{1+\lambda+\beta N}\right) \text{cov}(p, \bar{\theta}) = \left(\frac{\beta N}{1+\lambda+\beta N}\right) \rho_{p\bar{\theta}} \sigma_p \sigma_{\bar{\theta}}.$$

Thus, we have that:

$$\sigma_p^2 = \left(\frac{\beta N}{1+\lambda+\beta N}\right)^2 \rho_{p\bar{\theta}}^2 \sigma_{\bar{\theta}}^2. \tag{28}$$

We now prove that $\lambda = l\beta N$ in every linear equilibrium. We write the market power (6) as follows:

$$l = \frac{1}{N}\mathbb{E}\left[\frac{1}{p}\mathbb{E}\left[\sum_{i \in N}(\theta_i - q_i - p) \mid p\right]\right], \tag{29}$$

where we have used that the law of iterated expectations implies that $\mathbb{E}[\cdot] = \mathbb{E}[\mathbb{E}[\cdot \mid p]]$.The first-order condition (11) implies that:

$$q_i = \frac{\mathbb{E}[\theta_i|p, s_i] - p}{1+\lambda}.$$

Replacing q_i back into the equation (29) for market power, we get:

$$l = \frac{1}{N}\mathbb{E}\left[\frac{1}{p}\mathbb{E}\left[\sum_{i \in N}\left(\theta_i - \frac{\mathbb{E}[\theta_i|p, s_i] - p}{1+\lambda} - p\right) \mid p\right]\right].$$

Using the law of iterated expectations we get that $\mathbb{E}[\mathbb{E}[\theta_i|p, s_i] \mid p] = \mathbb{E}[\theta_i|p]$. Simplifying terms, we get:

$$l = \frac{1}{N}\mathbb{E}\left[\frac{1}{p}\sum_{i \in N}\frac{\lambda}{1+\lambda}\left(\mathbb{E}\left[\theta_i \mid p\right] - p\right)\right]. \tag{30}$$

Using (27), we get that:

$$Np\frac{1+\lambda+\beta N}{\beta N} = \mathbb{E}[\sum_{i \in N}\theta_i|p].$$

Replacing the conditional expectation of the payoff shock back into (30) we have that:

$$l = \mathbb{E}\left[\frac{1}{p}\frac{\lambda}{1+\lambda}\left(\frac{1+\lambda+\beta N}{\beta N}p - p\right)\right].$$

Here p cancels out, so we can omit the expectation on the right-hand-side of this equation. Simplifying terms, we get that $\lambda = l\beta N$.

Since $\rho^2_{p\bar{\theta}} \leq 1$, we have that (28) implies that:

$$\sigma_p^2 \leq \frac{(\beta N)^2}{(1+\beta N(1+l))^2}\sigma_{\bar{\theta}}^2.$$

Moreover, in the proof of Theorem 1 we also proved that in any linear Nash equilibrium $\lambda \geq -1/2$ and thus $l \geq -1/(2\beta N)$.

Sufficiency. Let (l, σ_p^2) be such that (26) is satisfied. We show that there exists an information structure that induces this market power and price volatility as a unique equilibrium. Suppose the payoff shocks are decomposed as in (13) and (14). Additionally, assume that the variances are given by:

$$\mathrm{var}(\phi_i) = \sigma_p^2 \frac{N}{\rho_{\theta\theta}(N-1)+1} \frac{(1+\beta N(1+l))^2}{(\beta N)^2} \quad \text{and} \quad \mathrm{var}(\eta_i) = (\sigma_\theta^2 - \mathrm{var}(\phi_i)). \tag{31}$$

Note that $\mathrm{var}(\eta_i)$ (as defined in (31)) is always positive because the theorem states that:

$$\sigma_p^2 \leq \frac{(\beta N)^2}{(1+\beta N(1+l))^2} \sigma_\theta^2,$$

and so $\mathrm{var}(\phi_i)$ (as defined in (31)) is less or equal than $\mathrm{var}(\theta_i)$. We assume that agents have no information on the realization of the shocks $\{\eta_i\}_{i\in N}$ and they observe a one-dimensional signals about the payoff shocks $\{\phi_i\}_{i\in N}$:

$$t_i = \phi_i - \left(1 - \frac{(l\beta N+1)(\beta N - l\beta N)}{\beta N l(N-1)(\beta N + l\beta N + 1)}\right)\frac{1}{N}\sum_{j\in N}\phi_j.$$

Note that this signal is the same as the one-dimensional signal used in the proof of Theorem 1, where we replace:

$$\gamma = \frac{(l\beta N+1)(\beta N - l\beta N)}{\lambda(N-1)(\beta N + l\beta N + 1)}.$$

We can now construct a linear equilibrium in the same way as in the proof of Theorem 1, so we get that the price volatility is:[8]

$$\frac{(\beta N)^2}{(1+\beta N+\beta N l)^2}\mathrm{var}\left(\frac{1}{N}\sum_{i\in N}\phi_i\right) = \frac{(\beta N)^2}{(1+\beta N+\beta N l)^2}\frac{\rho_{\phi\phi}(N-1)+1}{N}\mathrm{var}(\phi_i).$$

Since $\mathrm{var}(\phi_i)$ is defined as in (31) and $\rho_{\phi\phi} = \rho_{\theta\theta}$, the previous equation implies that the price volatility is given by σ_p^2. Moreover, the market power is l. Thus, (l, σ_p^2) is induced by an

[8] The only difference with the construction in the proof of Theorem 1 is that agents have no information on the realization of the shocks $\{\eta_i\}_{i\in N}$. However, the equilibrium can be constructed essentially the same way.

equilibrium outcome. Moreover, this is the unique equilibrium, as we already established in the proof of Theorem 1. ■

Theorem 2 provides a sharp bound on all possible equilibrium outcomes. It shows that the equilibrium outcome is bounded by the outcomes that are achieved under complete information. Thus the outcomes that arise under complete information can be seen as the "upper boundary" of the set of outcomes that can arise under all information structures, as illustrated earlier in Figure 1.

The "if" part of the statement closely resembles the proof of Theorem 1. In particular, the set of market power and price volatility that satisfy (26) would be achieved under complete information if one could reduce the variance of the aggregate shocks (i.e. by making $\mathrm{var}(\overline{\theta})$ smaller). By decomposing the payoff shocks into an observable and a non-observable component, we can effectively achieve the same outcomes as if there was complete information but the variance of the shocks was smaller.

The "only if" part of the statement is economically more interesting because it uses the restrictions that arise from agents' first order condition. By aggregating the agents' demands and using the market clearing condition, we can establish that the equilibrium price satisfies:

$$p^* = \frac{\beta}{1+\beta N(1+l)} \sum_{i \in N} \mathbb{E}[\theta_i | s_i, p^*]. \tag{32}$$

That is, the equilibrium price is proportional to the average of the agents' expected payoff shocks. Taking expectations of (32) conditional on p^* and using the law of iterated expectations, we can write the equilibrium price as follows:

$$p^* = \frac{\beta N}{1+\beta N(1+l)} \mathbb{E}[\bar{\theta} | p^*]. \tag{33}$$

Since (33) relates p^* with the expectation of $\bar{\theta}$ conditional on p^*, it follows that the variance of

p^* is directly related to the variance of $\bar{\theta}$ and the correlation between p^* and $\bar{\theta}$.[9]

It is crucial for the argument that the expected payoff shock of agent i is computed conditional on the equilibrium price — this is an implication of the fact that agents compete in demand functions and hence agent i can condition the quantity he buys on the equilibrium price. The fact that an agent can condition on the equilibrium price disciplines beliefs, which ultimately allows us to bound the price volatility. This allows us to relate p^* to the average payoff shock $\bar{\theta}$ (as in (33)), instead of p^* being related only to the average of the agents' expected payoff shocks (as in (32)).

Theorem 2 allows to give predictions about market power and price volatility that hold across every information structure. We can compare these outcomes with the set of outcomes that can arise when agents observe one-dimensional noisy signals as in Vives (2011). Suppose every agent i observes a one-dimensional signal:

$$s_i = \theta_i + \varepsilon_i, \tag{34}$$

where ε_i is a normally distributed noise term and independently distributed across agents. This set of information structures is parametrized by the variance of the noise term σ_ε^2. The yellow line in Figure 3 illustrates the set of market power and price volatility and that can arise for all $\sigma_\varepsilon^2 \in [0, \infty)$. The arrows indicate the direction in which the variance increases. We can see that increasing the variance of the noise term increases the equilibrium market power and decreases the price volatility. However, not every market power is attained by some σ_ε^2 and also the price volatility decreases at a faster rate than under the complete-information equilibria (illustrated by a red line). The reason price volatility decreases faster is that the signals are noisy and so the price co-moves less with the payoff shocks.The noise-free signals and the class of noisy one-dimensional signals (34) have the common feature that in equilibrium prices are privately revealing. This

[9]For any two random variables (y, z), if $y = \mathbb{E}[z|y]$, then $\sigma_y = \rho_{yz}\sigma_z$.

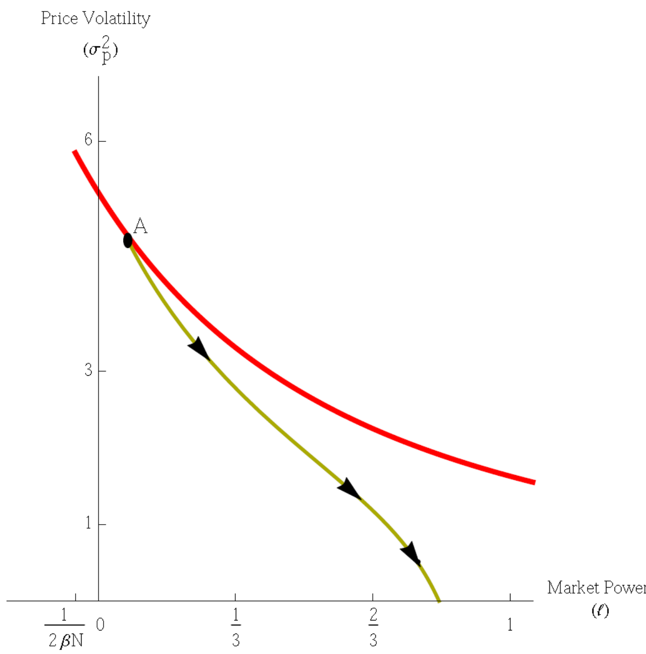

Figure 3: Set of equilibrium pairs of market power and price volatility $\left(l, \sigma_p^2\right)$ under noisy one-dimensional signals.

means that the information contained in the price plus the private signals observed by agent i are a sufficient statistic for all the signals observed by every agent to predict θ_i. When prices are privately revealing constructing a linear Nash equilibrium is analytically tractable because it is easy to guess the informational content of the price. However, we note that the characterization in Theorem 2 also holds for equilibria in which the price is not privately revealing.

An interesting conclusion that one can derive from the proof of Theorem 2 is that the expected price is uniquely determined by the equilibrium degree of market power. More precisely, (27) implies that:

$$\mu_p = \frac{\beta N}{1 + \lambda + \beta N} \mu_\theta .$$

Thus, the equilibrium degree of market power l and the price volatility describe the first and second moments of the equilibrium price. By extension, they also describe the first and second moment of the aggregate demand. This in turn allows us to describe other statistics of an equilibrium outcome. For example, the expected revenue is given by:

$$\mathbb{E}[p \sum_{i \in N} q_i] = \mathbb{E}[\frac{p^2}{\beta}] = \frac{1}{\beta}(\mu_p^2 + \sigma_p^2).$$

Here we have used that $\sum_{i \in N} q_i = p/\beta$ by the market clearing condition. Since Theorem 2 gives bounds on the price volatility and market power, and market power determines the mean price, it is easy to see how Theorem 2 also allows to give bounds on the equilibrium revenue of the seller.

5 Comparing Market Mechanisms

Our paper focuses on characterizing the set of possible pairs of market power and price volatility that can be attained by a Nash equilibrium of a particular mechanism, the demand function competition game. A methodological contribution of the paper is that we provide a characterization of key statistics of the equilibrium outcome independent of the specific equilibrium strategies that generates these outcomes. This approach has the advantage that it permits an easy and insightful comparison of *different* mechanisms or game forms.

We now illustrate this by means of the Cournot competition. Thus, we consider the outcome of the economic environment as described in Section 2 in terms of payoff functions and payoff shocks, but where agents now compete by submitting quantities (i.e., Cournot competition) instead of submitting demand schedules. Bergemann, Heumann, and Morris (2015) characterizes all the equilibrium outcomes of the quantity competition game. We can use their results (and the results we have given here) to compare both forms of market competition.[10]

[10] The explicit calculations and comparison can be found in an earlier version of this paper (see Bergemann, Heumann, and Morris (2018)).

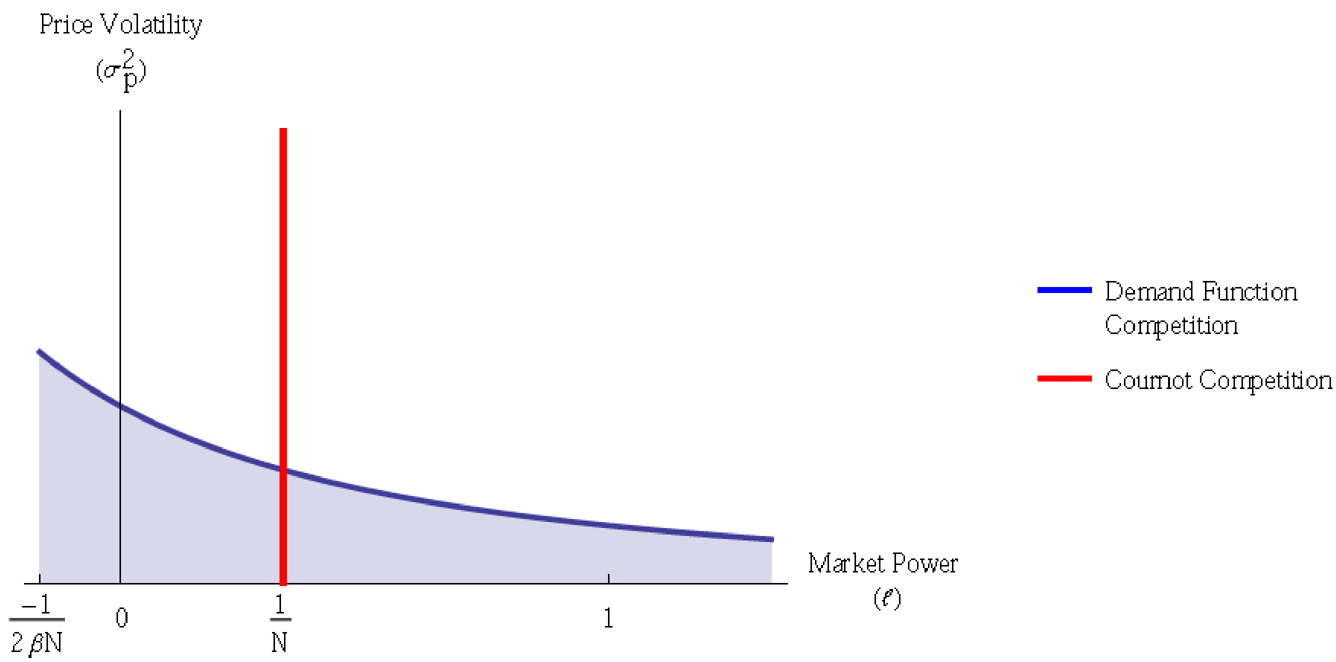

Figure 4: Comparison of the first and second equilibrium moments under demand function competition and Cournot competition.

We first give a characterization of market power and price volatility in quantity competition that is analogous to Theorem 2. There exists an information structure that induces a pair of market power and price volatility, (l, σ_p^2), under quantity competition if and only if:

$$l = \frac{1}{N} \quad \text{and} \quad \sigma_p^2 \leq \frac{1}{4}\left(\frac{\sqrt{1+\beta}\sigma_{\bar{\theta}} + \sqrt{(\beta+\beta N+1)\sigma_{\theta}^2 + \beta N \sigma_{\bar{\theta}}^2}}{\sqrt{1+\beta}(\beta+\beta N+1)}\right)^2. \tag{35}$$

The restrictions that quantity competition imposes on market power and price volatility are strikingly different to those imposed under demand function competition. We illustrate the possible market power and price volatility that can be attained by some information structure under both forms of market competition in Figure 4; this figure illustrates the discussion that we give next.

Under quantity competition, market power is constant instead of being (almost) completely indeterminate. The expression for price volatility is slightly more cumbersome, but there is one important feature that is worth highlighting. Even if the aggregate shock is close to zero (i.e., $\sigma_{\theta}^2 \approx 0$), there may be non-negligible price volatility. In contrast, under demand function competition, price volatility is always negligible if the variance of the aggregate shock is negligible (see Theorem 2). The difference between quantity competition and demand function competition is explained as follows.

Under quantity competition an agent's price impact is equal to the slope of the exogenous supply; this explains why market power is constant across all information structure. In contrast, when agents submit price-contingent demands, an increase in the demand of an agent has an effect on the price, which in turn, affects the demand of other agents. Thus, an agent's price impact is not determined solely by the exogenous supply.

For price volatility, the intuition is as follows. In quantity competition the equilibrium quantities may be correlated because signals are correlated, which may lead to "excess" aggregate volatility. In contrast, when agents submit price-contingent demands, their beliefs are disciplined by the equilibrium price. More precisely, because agents choose their demand contingent on the price, the price serves as a public signal about the average quantity purchased by all agents. This additional signal disciplines beliefs in such a way that the quantities they purchase cannot be "overly correlated" which ultimately disciplines price volatility.

6 Discussion

In this paper we study demand function competition. Our results provide positive and negative results regarding our ability to make predictions in this empirically important market microstructure. We showed that any market power is possible— from $-1/(2\beta N)$ to infinity. Considering small amounts of incomplete information does not allow us to provide any sharper predictions, unless one is able to make additional restrictive assumptions regarding the nature of the incomplete information. Yet, we showed that we can provide many substantive predictions regarding the demand function competition that are robust to weak informational assumptions.

While our analysis focused on studying market power and price volatility the conclusions can be extended to other equilibrium statistics. For example, an analyst may be interested in the variance of the quantities bought by each agent, that is $\operatorname{var}(q_i)$. In an earlier version of this paper (see Bergemann, Heumann, and Morris (2018)), we explored more broadly how information may determine *any* given equilibrium statistic. Our conclusions there extended the current results in the sense that the equilibria under complete information are "extremal" equilibria. For example, the variance of the quantities is bounded the variance attained under a complete-information equilibrium. Thus, while it is difficult to rule out any of the equilibria that arise under complete information, these equilibria can be used to provide bounds of what can happen across all information structures.

References

Bergemann, D., T. Heumann, and S. Morris (2015): "Information and Volatility," *Journal of Economic Theory*, 158, 427–465.

——— (2018): "Information and Market Power," Discussion paper, https://ssrn.com/abstract=3276348.

Grossman, S. (1981): "Nash Equilibrium and the Industrial Organization of Markets with Large Fixed Costs," *Econometrica*, 49, 1149–1172.

Hart, O. (1985): "Imperfect Competition in General Equilibrium: An Overview of Recent Work," in *Frontiers of Economics*, ed. by K. Arrow, and S. Honkapohja. Basil Blackwell.

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

Kyle, A. (1985): "Continuous Auctions and Insider Trading," *Econometrica*, 53, 1315–1355.

Lerner, A. (1934): "The Concept of Monopoly and the Measurement of Monopoly Power," *The Review of Economic Studies*, 1, 157–175.

Rostek, M., and M. Weretka (2012): "Price Inference in Small Markets," *Econometrica*, 80, 687–711.

——— (2015): "Information and Strategic Behavior," *Journal of Economic Theory*, 158, 536–557.

Rubinstein, A. (1989): "The Electronic Mail Game: Strategic Behavior under 'Almost Common Knowledge'," *American Economic Review*, 79, 385–391.

Vives, X. (1999): *Oligopoly Pricing.* MIT Press, Cambridge.

——— (2011): "Strategic Supply Function Competition with Private Information," *Econometrica*, 79, 1919–1966.

Weinstein, J., and M. Yildiz (2007): "A Structure Theorem for Rationalizability with Applications to Robust Predictions of Refinements," *Econometrica*, 75, 365–400.

——— (2011): "Sensitivity of Equilibrium Behavior to Higher-Order Beliefs in Nice Games," *Games and Economic Behavior*, 72, 288–300.

Wilson, R. (1979): "Auction of Shares," *Quarterly Journal of Economics*, 93, 675–689.

7 Appendix

Proof of Proposition 1. Using the equilibrium construction in the proof of Theorem 1, in particular (11) we find that in any symmetric Nash equilibrium x^*:

$$x^*(s_i, p) = \frac{\theta_i - p}{1 + \lambda}, \tag{36}$$

where λ is the derivative of an agent's residual supply. This equation is the complete information counterpart to (11). The market clearing condition implies that the equilibrium price is given by:

$$p^* = \frac{\beta N \bar{\theta}}{1 + \lambda + \beta N}. \tag{37}$$

The arguments established in the proof of Theorem 1 imply that market power satisfies $l = \lambda/\beta N$ and that the second-order condition of an agent's maximization problem implies that $\lambda \geq -1/2$. Thus, (37) implies that the equilibrium relation between price volatility and market power is as in (9).

To check that the conditions are sufficient we consider the following family of demand functions parametrized by $\gamma \in \mathbb{R}$:

$$x_i(p) = \frac{1}{1 + \lambda}\left(\theta_i - (1 - \gamma)\bar{\theta}\right) - \frac{1}{N - 1}\left(\frac{1}{\lambda} - \frac{1}{\beta}\right)p, \tag{38}$$

where λ is determined as a function of γ by the positive root of (22). We first observe that this is the same demand function as the one that the agents submit in the Nash equilibrium when they observe the information structure constructed in Theorem 1 (there γ parametrizes the one-dimensional signal (16)). Thus, if (38) constitutes a Nash equilibrium, then the equilibrium market power and price volatility is given by:

$$l = \lambda/\beta N \quad \text{and} \quad \sigma_p^2 = \frac{(\beta N)^2}{(1 + \beta N + \lambda)^2}\sigma_{\bar{\theta}}^2,$$

where $\lambda \in [-1/2, \infty)$ is the positive root of (22). To check that (38) is an equilibrium, we note that when λ is determined by the positive root of (22) we have that:

$$x_i(p) = \frac{1}{1+\lambda}\left(\theta_i - (1-\gamma)\bar{\theta}\right) - \frac{1}{N-1}(\frac{1}{\lambda} - \frac{1}{\beta})p = \frac{\theta_i - p}{1+\lambda}.$$

Thus, the demand satisfies the first-order condition (36) and it also satisfies the second-order condition because $\lambda \geq -1/2$. Thus, (38) satisfies the optimality conditions and thus constitutes a symmetric Nash equilibrium. ■