RAND Journal of Economics
Vol. 00, No. 0, xxx 2014
pp. 1–20

Learning about common and private values in oligopoly

D. Bernhardt*

and

B. Taub**

We characterize a duopoly buffeted by demand and cost shocks. Firms learn about shocks from common observation, private observation, and noisy price signals. Firms internalize how outputs affect a rival's signal, and hence output. We distinguish how the nature of information —public versus private—and of what firms learn about—common versus private values—affect equilibrium outcomes. Firm outputs weigh private information about private values by more than common values. Thus, prices contain more information about private-value shocks.

1. Introduction

■ In the real world, firms are buffeted by shocks both to demand and to costs, some of which they learn about via direct common observation, some via private observation, and some of which they must attempt to extract from information in prices. Some shocks have a common-value nature—for example, a common demand shock that raises demand equally for each firm's product—entering directly into each firm's profit function. Other shocks take a private-value form—for example, an idiosyncratic cost shock specific to one firm's technology, or a firm-specific demand shock that raises demand for only one firm's product—and hence only indirectly affects a rival, affecting a rival's profits only to the extent that the shock alters the output of the affected firm.

In an oligopolistic industry, firms must account for the strategic behavior of rival firms when unraveling information from price signals, and they must account for how their own actions influence the price signals that rivals receive, and hence, their inferences and output choices. We distinguish how the strategic incentives to manipulate a rival's beliefs vary with the amount and extent of public and private information about the shocks, whether those values have a common value or private–value nature, and how these different types of information are aggregated and expressed in prices.

* University of Illinois and University of Warwick; danber@Illinois.edu.
** Glasgow University; bart.taub@gmail.com.
We thank Marco Ottaviani for helpful comments.

We start with a setting in which demand evolves from the beginning to the end of a production period, and firms compete in supply schedules, producing output before end-of-period demand is known. Firms receive private information about private-value and common-value shocks. They then submit supply schedules that detail for each beginning-of-period price an output that is optimal, conditional on their private information and the information contained in the fact that equilibrium prices are market clearing. After production occurs, demand evolves. Final prices are determined by market clearing, given end-of-period demand. We establish that the equilibrium of this economy is formally equivalent to a noisy rational expectations economy in which, rather than compete in supply schedules, firms observe a noisy signal of the final market-clearing price and base outputs on that information. In effect, the information in the fact that the beginning-of-period price is market clearing serves as a noisy signal of the final market clearing price—it is "as if" firms see a noisy signal of the end-of-period price. The key difference from standard noisy competitive rational expectations equilibrium (REE) models is that in an oligopoly, firms understand and internalize how their actions affect price signals, and hence the information of rivals, and thus their outputs. In contrast, competitive noisy REE models assume that agents are informationally small and ignore the impacts of their actions on the information contained in equilibrium prices.

In this economy, we characterize how and why firms' outputs weight private information about common-value shocks differently from private information about private-value shocks, which, in turn, they weight very differently from public information. We show how common and private *values* lead to very different impacts of public and private *information*.

We prove that firms always weight private-value shocks in their outputs more heavily than they weight common-value shocks. As a result, prices convey more information about private values than common values. The reason is that firms produce on opposite sides on private values, whereas they produce on the same side on common values. For example, via its price signal (equivalently, the information contained in the fact that the beginning-of-period price is market clearing), a firm positively weights common-value demand shocks observed by its rival. This causes a rival to cut its output weight on the common shock, which *reduces* the information in prices about the common shock. In contrast, when a firm positively weights a private-value, idiosyncratic shock to its demand, only the output consequences, and not the shock itself, matters to its rival. As a result, the rival partially offsets this output; and this partially offsetting position induces the firm receiving the shock to increase its weight on the private shock, which *raises* the information in price signals about the private demand shock. As a result, the information content of prices, and thus what rivals can learn from price, is always more sensitive to private signals about private-value cost or demand uncertainty than to private signals about common-value demand uncertainty.

The result that firms learn more from prices about private-value shocks than common-value shocks has important implications for equilibrium: whether an equilibrium exists, and how many equilibria exist, depends on the amount of private-value uncertainty relative to the noise in price signals. With noisy signals or little private-value uncertainty, a unique linear equilibrium exists regardless of the uncertainty about common values. However, as private-value uncertainty is increased, eventually we reach the point where two linear equilibria exist; and once there is "enough" private-value uncertainty, no equilibrium exists. Strategic complementarities in learning about private values underlie the multiple equilibria. In one equilibrium, a firm's output weights its private information about its private-value shocks by less; as a result, prices convey less information to its rival, so that a rival's output offsets little of the firm's output, confirming the optimality of not weighing shocks heavily. In the other equilibrium, a firm weights its private information heavily; as a result, prices contain more information, causing a rival to offset far more of a firm's output, validating decisions to heavily weight private information. In contrast, because firms produce in the same direction on common values, uncertainty about common values alone cannot support multiple equilibria.

In our linear-quadratic, Gaussian setting, learning is just the least squares projection of a rival's privately observed shocks on price signals. The relevant object of a firm's learning is the shift in the price due to the rival's privately observed information – the residual common value shock left after its rival's direct output on it plus the impact of private-value shocks to a rival on its output. We establish that what is learned is weighed just as if it were public information, reflecting that a rival can also forecast this learning.

We use this result to assess how firms weight private versus public information about private- and common-value shocks. The signal-jamming literature (Riordan, 1985; Aghion et al., 1991; Mirman, Samuelson, and Urbano, 1993; Caminal and Vives, 1996; or Harrington, 1995) highlights how firms overproduce on common-value public information – firm outputs weight public common demand by more than they would in a full-information setting – to try to persuade rivals via their price signals that the market is less profitable. However, what about privately observed shocks? Also, how do answers hinge on the private- and common-value natures of these demand and cost shocks?

We prove that firms collectively weight private information about private-value shocks by *more* than they weight public information about private values; but they weight private information about common-value shocks by *less* than they weight public information about common values. What underlies this result is that firms weight information learned from prices just like public information, reflecting that its rival can also forecast this learning. An informed firm acts like a monopolist on its residual private information – but a monopolist weights common-value information by *less* than a duopoly collectively would; whereas a monopolist weights private-value information by *more* than a duopoly, precisely because its rival partially offsets some of the price impact of its output.

Relatedly, firms always collectively weight private-value shocks by more than they would in a full-information setting. However, they only collectively weigh privately observed common demand shocks by more if prices convey enough information to a rival. Indeed, the two firms' *collective* output weight on a privately observed *private-value* shock always equals the monopoly output weight. This reflects that a firm hit by a shock can forecast what its rival learns via its price signal and take an offsetting position, leaving the firm in the position of a monopolist on a shock that only directly affects it. In contrast, a firm that privately observes a *common* demand shock always places a total (direct plus indirect) output weight equal to the monopoly output weight. These findings highlight how the output impacts of private information hinge fundamentally on whether that information concerns common or private values.

We next relate our article to various literatures. Section 2 presents our model. Given linear conjectures about a rival's output function, we solve a firm's optimization problem to obtain consistent linear outputs and pricing. Each firm projects its rival's private information about demand and cost on its noisy price signal. We solve for output strategies in terms of this pricing filter and develop a recursive mapping in the projection coefficient, the fixed points of which describe the equilibria. We characterize the correlation structure of information, the learning of firms via price signals, and the equilibrium output weights in terms of the economic primitives. We then describe how our static findings extend to a stationary dynamic setting. A conclusion follows. Proofs are in the Appendix.

Related literature. The analysis of oligopolistic competition in supply schedules with demand uncertainty dates back to Klemperer and Meyer (1989). They argue that competition in supply schedules better describes strategic competition between firms than competition in prices or quantities, because it more realistically allows firms to adjust to market conditions. In particular, in a supply schedule equilibrium, firms adjust to market conditions in an optimal manner given their rival's behavior – given knowledge of the market-clearing price, they would have no incentive to adjust their output level. In contrast, with stochastic Bertrand or Cournot competition, firms would want to alter their actions after learning something about demand.

Our model is most closely related to Vives (2011). His model also features private information and noisy signals: as in our model, the costs and benefits of manipulating beliefs via output are incurred simultaneously. In his model, firms receive private noisy signals about costs, and because costs are correlated across firms, one firm's signal is relevant for its rival. Firms compete in supply schedules and there is no demand uncertainty. As a result, the market-clearing price is privately fully revealing: in equilibrium, a firm's own cost signal and price yield the same forecast of its costs as when a firm sees the cost signals of each firm.[1] In contrast, in our model, firms know their own costs and they also have private information about demand, but they must estimate the shocks observed by rivals via the information in preliminary prices. Moreover, these prices are not privately fully revealing.

There is a large literature on learning and experimentation by firms when belief manipulation issues are absent. McLennan (1984), Aghion et al. (1991), and Harrington (1995) look at learning about a single parameter of demand, addressing whether and when firms eventually learn demand. Rustichini and Wolinsky (1995) look at experimentation and learning by a monopolist when unit demand moves according to a two-state Markov process, whereas Keller and Rady (1999) consider experimentation and learning by a monopolist about a two-state Markov demand process in a continuous time setting. Keller and Rady (2003) extend their analysis to a duopoly setting, assuming that firms observe each others' actions so that strategic manipulation considerations are absent. We extend this literature by incorporating feedback from prices, so that the learning process is entangled with the strategic efforts of firms to manipulate the beliefs of rivals; and we derive how the private/common value nature of the information being learned affects outcomes.

Other research takes the opposite tack of exploring how actions are affected when firms have incentives to manipulate beliefs of rivals in symmetric information settings. The signal-jamming literature (Riordan, 1985; Aghion et al., 1991; Mirman, Samuelson, and Urbano, 1993; Caminal and Vives, 1996; Alepuz and Urbano, 2005; Harrington, 1995) explores belief manipulation incentives when firms learn about the level of demand from prices. In these two-date models, firms are symmetrically uninformed about demand or costs. At date two, firms observe prices and condition outputs accordingly. At date one, firms internalize how outputs affect price and thus the inferences of rivals, and hence their date-two outputs. Due to the absence of private information, in equilibrium at date two, firms perfectly infer the demand realization. Separating the benefits from signal jamming (reduced date-two output by rivals) from the costs (excessive date-one output) raises incentives to signal jam. In contrast, in our model, costs and benefits are incurred simultaneously, so that strategic output choices cannot be divorced from the filtering of price. We also distinguish how incentives to manipulate beliefs vary with public and private information, and with the common/private-value nature of that information.

There is a large literature in which firms have private information about demand or costs. In the two-period limit pricing models of Harrington (1986, 1987), Caminal (1990), or Bagwell and Ramey (1991), an incumbent (or duopoly) have private information, for example, about a potential entrant's costs, and choose prices that signal this information, influencing entry decisions. Mailath (1989) considers a differentiated good duopoly, where firms simultaneously signal costs via date one price announcements, and determines when a separating equilibrium exists.[2] Athey and Bagwell (2008) introduce private cost uncertainty to the literature on collusion with imperfect monitoring, analyzing a stationary procurement auction game in which a firm's costs evolve according to a two-state Markov process, and firms make cheap-talk announcements about costs prior to making bids. Histories matter for incentives, but, with cheap talk, are never used to obtain information about fundamentals.

Methodologically and philosophically, our model is closest to the literature on strategic speculation in financial markets (Kyle, 1989; Bernhardt, Seiler, and Taub, 2010) where noise enters prices due to "liquidity" traders, and speculators condition demands on prices (submitting

[1] Mirman, Salgueiro, and Santugini (2014) analyze a monopolist whose price provides consumers a noisy signal of quality.

[2] A more distantly related literature identifies conditions under which firms with private information earn higher profits if they first share information at a cheap-talk stage and then make output decisions than if they conceal their information (Gal-or, 1985, 1986, 1988; Vives, 1984; Malueg and Tsutsui, 1996).

demand schedules), internalizing how their demands influence price. However, oligopoly and speculative incentives are diametrically opposed – speculators seek to conceal private information; whereas firms seek to convince rivals that demand and costs are low.

□ **Discussion.** Our analysis describes the strategic interactions of firms in settings with three distinguishing features. First, the firms face substantial uncertainty about demand and/or costs, and there is significant learning. Second, the industry has an oligopolistic structure – the actions of individual firms affect prices and price signals, and firms internalize the direct effects in their decision making, and the indirect effects on the other firm's actions via their learning from the information contained in prices. Third, production takes place over time with important decisions (e.g., regarding inventory or scale) being made before final demand and/or final unit production costs are known, being based on private information and the information in preliminary prices.

For example, our analysis describes strategic interactions in the large passenger-jet industry. The two main rival firms, Airbus and Boeing, make production decisions long before they know the prices they will receive: contracts with airlines are signed years ahead of production, but orders can be cancelled, prices can rise or fall reflecting individual situations of airlines, inflation can occur, and so on. Not only does new information about "final" demand come to firms as production proceeds, but Boeing learns about Airbus's production imperfectly, and with a lag. For example, Boeing very imperfectly learns about the attributes that enter the quality of Airbus's aircraft – properties such as avionics, interior design, fuel economy stemming from aerodynamic properties – and the costs of these new features.

One can view Boeing and Airbus as determining supply schedules, physically realized in part by the inventories of parts built up in anticipation of production, decisions about speed of production, and the confidential price schedules they provide their sales representatives. These details are largely private in nature. As a result, information about demand, and how that demand will translate into future prices is observed with noise. This means that the firms may use current and past prices as signals from which they extract information about market conditions and a rival's costs. In turn, this means that a rival has incentives to consciously manipulate those signals via production decisions – again, choices about stockpiles of inputs, speed of manufacturing, and so on.

Our analysis also describes strategic interactions in the cell phone industry. The major players, Apple and Samsung, are large enough to take into account how their production and pricing decisions interact and influence decisions by rivals. Cell phone manufacturers periodically create new models and face substantial uncertainty about how consumers will value different attributes, how technological advances will affect new applications and their value, the costs of producing new models, and so on. Full-scale production of new models begins months before they are taken to the market, and firms update about what that final demand will be from intermim information arrival that is influenced by the actions of rival cell phone manufacturers. Further, most phones are sold via service providers like Sprint, T-Mobile, AT&T, and outlets like Best Buy, and prices are determined via negotiation between the providers and manufacturers, implicitly aggregating preliminary demand, just as we model.

It is tempting to use a two-period model to capture such strategic interactions. In such a model, Airbus, given its initial private information about demand and costs, might make period 1 production decisions based on that information, but with those decisions modified by the knowledge that Boeing will, in part, base period 2 output on information extracted from period 1 prices. The problem with this modelling approach is that the period 2 payoffs from manipulating Boeing's inferences from period 1 prices are not synchronized with the period 1 manipulation costs, distorting manipulation incentives upward. This also complicates interpretations of results. This leads us to model the strategic interplay in a setting where firms incur the costs and benefits of manipulating a rival's information at the same time, leaving incentives undistorted. One could alternatively model this strategic interplay in a dynamic stationary setting in which firms strategically manipulate a rival's information in each period, incurring the costs of current manipulation, but also gleaning the benefits from past manipulation as the rival reacts. Section 3 describes how such intertemporal considerations alter firm behavior.

2. The model

■ We begin by considering firm learning about private- and common-value components of profits in a duopoly setting in which firms compete in supply schedules and demand evolves after production decisions are made. The supply schedules detail for each beginning-of-period price an output level. Firms recognize that demand will evolve after output is produced, so that realized end-of-period sale prices will be different. We establish an equivalence between the equilibrium of this supply schedule game, where firms do not directly see the market-clearing beginning-of-period price, and a noisy rational expectations setting in which firms directly observe a noisy signal of the market-clearing end-of-period price, and internalize how their output choices affect the information and output of a rival.

At the start of the period, firms 1 and 2, receive information, some of it private in nature, about common- and private-value shocks to demand and to costs. Beginning-of-period demand is given by

$$p_i^P(q_1, q_2) = \hat{a} + b^c(\hat{a}_1^c + \hat{a}_2^c) + b^p\hat{a}_i^p - (q_1 + q_2) + e^P,$$

where (i) $\hat{a} > 0$ is a public information component of demand that is common to both firms; (ii) $\hat{a}_j^c$, $j = 1, 2$ is a common demand shock that firm j privately observes; (iii) $\hat{a}_i^p$ is a private demand shock that is specific to firm i's product that firm i privately observes; (iv) $b^p > 0$ and $b^c > 0$ are the associated weights on these shocks in demand; (v) q_j is output by firm $j = 1, 2$; and (vi) e^P is a common, but unobservable, component of demand that disappears before final production.[3] The demand shocks are independently and normally distributed, with $\hat{a}_j^p \sim N(0, \hat{\sigma}_a^{2p})$, $\hat{a}_j^c \sim N(0, \hat{\sigma}_a^{2c})$, for $j = 1, 2$, and $e^P \sim N(0, \sigma_e^2)$.

As we describe below, firms compete in supply schedules that detail for each preliminary price an output level. What complicates matters for firms is that after outputs are determined, some components of preliminary demand disappear, for example, some sources of expected customer demand may drop out, and there may be other unanticipated shocks to demand. What firms $i = 1, 2$ care about is the final end-of-period demand,

$$p_i^F(q_1, q_2) = \hat{a} + b^c\left(\hat{a}_1^c + \hat{a}_2^c\right) + b^p\hat{a}_i^p - (q_1 + q_2) + e^F, \quad \text{where } e^F \sim N\left(0, \sigma_{eF}^3\right),$$

which determines the prices that their outputs will actually receive. This formulation captures the fact that some potential consumers will appear or disappear from demand both prior to output choices and afterward, and that firm output choices must account for both of these possibilities.

Each firm faces constant marginal costs of production. To produce q_i, firm i incurs costs

$$c_i(q_i) = (\hat{c} + \hat{c}_i^p)q_i,$$

where $\hat{c}$ is a publicly observed component of costs that is common to the two firms, and $\hat{c}_i^p$ is a privately observed shock that affects firm i's costs, but not its rival's, where $\hat{c}_i^p$ is independently and normally distributed, $\hat{c}_i^p \sim N(0, \hat{\sigma}_c^2)$. The normalization of the means of the shocks to zero is without loss of generality.

---

[3] One can allow for a more general formulation of firm demand, in which a firm's own output has a greater weight on its price than that of its rival, so that

$$p_1^P(q_1, q_2) = \hat{a} + b^c(\hat{a}_1^c + \hat{a}_2^c) + b^p\hat{a}_1^p - (q_1 + yq_2) + e^P,$$

where $0 < y \leq 1$. The approach to solving for the equilibrium mirrors that here. However, $y < 1$ introduces quadratic terms in the output weights that firms place on their different sources of information. As a result, closed-form analytical solutions do not obtain. One can solve for equilibrium outcomes numerically and verify that the qualitative properties are unchanged as long as y is large enough (so that firm 1 cares enough about the output choices of firm 2).

Noting that cost and demand components enter profits linearly as differences, we can consolidate these terms and write firm profits in reduced form

$$\Pi_i = (\bar{a} + a_1 + a_2 + e^F - (q_1 + q_2))q_i - c_i q_i, \tag{1}$$

where $\bar{a} = \hat{a} - \hat{c}$ is the net public-information demand-cost difference, $a_j = b^c \hat{a}_j^c$ is a common-value shock to demand privately observed by firm j, where $a_j \sim N(0, \sigma_a^2 = (b^c)^2 \hat{\sigma}^{2c})$ and $c^j = (\hat{c}_j^p - b^p \hat{a}_j^p)$ is a net private-value demand-cost difference shock, so that $c^j \sim N(0, \sigma_c^2 = (b^p)^2 \hat{\sigma}^{2p} + \hat{\sigma}^{2c})$.

By rewriting profits so that it is "as if" private values only enter costs and common values only enter demand, exposition is eased, and we obtain the simple profit formulation in (1). Moreover, by loading all private–value components into costs, it is "as if" firms see the same price signals,

$$p_1^{P*}(a_1, c_1, a_2, c_2, e^P) = p_2^{P*}(a_1, c_1, a_2, c_2, e^P) = p^{P*}(a_1, c_1, a_2, c_2, e^P),$$

easing presentation. We work with this in what follows.

Firms compete in supply schedules. At the beginning of the period, after observing public and private information about demand and costs, firms choose supply schedules that detail for each possible preliminary price an output choice. Firms optimize, choosing supply schedules that maximize expected profits given public and private information and the information contained in the preliminary price being market clearing (i.e., consistent with their output choices).

Definition 1. A strategy for firm i is a function Q_i that maps its private information, a_i and c_i, into a supply schedule $q_i(p^P; a_i, c_i)$ that details for each possible market-clearing preliminary price, $p^P \in R$, the quantity that firm i supplies to the market. Letting S denote the space of differentiable supply schedules $q: R \to R$, a strategy specifies for each $(a_i, c_i) \in R \times R$ a supply schedule $q_i(p^P; a_i, c_i) \in S$.

Definition 2. A preliminary price function $p^P : (\overline{a}, a_1, a_2, q_1, q_2, e^P) \mapsto R$ is market clearing if it is consistent with the supply schedules $q_i : (p^P; a_i, c_i) \mapsto R$. That is, for each $(a_1, a_2, c_1, c_2, e^P)$,

$$\begin{aligned} p^P(a_1, c_1, a_2, c_2, e^P) = \bar{a} + a_1 + a_2 + e^P - q_1(p^P(a_1, c_1, a_2, c_2, e^P); a_1, c_1) \\ - q_2(p^P(a_1, c_1, a_2, c_2, e^P); a_2, c_2). \end{aligned} \tag{2}$$

Timing. The sequence of events and actions within each period as follows:

(i) $(a_1, a_2, c_1, c_2, e^P)$ are realized.
(ii) Given (a_i, c_i), firms $i = 1, 2$ submit supply schedules $q_i(\cdot; a_i, c_i)$.
(iii) The market-clearing beginning-of-period price $p^P(a_1, a_2, c_1, c_2, e^P)$ is determined, which, in turn, determines outputs $q_i(p^P; a_i, c_i)$, $i = 1, 2$.
(iv) End-of-period price, $p(a_1, a_2, c_1, c_2, e^P, e^F) = p^P(a_1, a_2, c_1, c_2, e^P) - e^P + e^F$ is realized and firms realize profits.

When submitting supply schedules, firms understand and internalize the facts that (i) the beginning-of-period price clears the market, and (ii) its rival's output will hinge on this market-clearing price. Each firm i seeks to maximize expected realized profits conditional on its information, $\{a_i, c_i\}$, and the information contained in the fact that this price is market clearing. Thus, for each possible beginning-of-period price p^P, firm i's supply schedule specifies a q_i that solves

$$\max_{q_i \in R} E[(\bar{a} + a_1 + a_2 - (q_i + q_{-i}(p^P; a_{-i}, c_{-i}) + e^F)q_{-i} - c_i q_i | a_i, c_i, p^P(a_1, a_2, c_1, c_2, e^P)], \tag{3}$$

where $q_{-i}(p^P; a_{-i}, c_{-i})$ is the supply schedule of its rival, evaluated at the market-clearing beginning-of-period price, $p^P(a_1, a_2, c_1, c_2, e^P)$. In particular, firm i internalizes that as it varies

q_i, it affects the market-clearing beginning-of-period price, and hence its rival's output, because its rival's supply schedule is evaluated at a different price. Of course, in equilibrium, a rival's supply schedule is evaluated at the market-clearing beginning-of-period price.

Definition 3. An equilibrium is a pair of strategies, Q_1^*, Q_2^* and a market-clearing price function $p^{P*}(a_1, a_2, c_1, c_2, e^P)$ such that for each $(a_1, a_2, c_1, c_2, e^P)$, (a) $q_i^*(p^{P*}(a_1, a_2, c_1, c_2, e^P); a_i, c_i)$ maximizes expected profits, solving (3), for $i = 1, 2$, and (b) preliminary prices are market clearing,

$$p^{P*}(a_1, c_1, a_2, c_2, e^P) = \bar{a} + a_1 + a_2 + e^P - q_1^*(p^{P*}(a_1, c_1, a_2, c_2, e^P); a_1, c_1) \\ -q_2^*(p^{P*}(a_1, c_1, a_2, c_2, e^P); a_2, c_2).$$

In general, solving for equilibrium in such a game is extremely difficult – the sensitivity of the rival's supply response to an increase in firm i's output will hinge both on the curvature of the pricing function, which determines the change in price, and on the sensitivity of the rival's supply schedule to that price change, which both enters that price change, and can conceivably vary in complicated ways with its private information, $\{a_{-i}, c_{-i}\}$. Equilibrium is the fixed point of such a problem, and, typically, it is completely intractable. Here, we exploit the Gaussian nature of the demand and supply shocks, and conjecture that the rival's equilibrium supply schedule is a linear function of its information. This means that from firm i's perspective, price is a linear function of the fundamental primitives, which means that its conditional expectations of the unknown primitives given its own information are linear. In turn, this means that the solution to its resulting quadratic optimization problem is linear, confirming the linear structure of the equilibrium.

This leads us to conjecture that the rival's supply schedule is linear, taking the form

$$q_{-i}(p^P; a_{-i}, c_{-i}) = \alpha_{-i}a_{-i} + \beta_{-i}\bar{a} - \gamma_{-i}c_{-i} + \delta_{-i}p^P. \tag{4}$$

Given the linear conjecture about the rival firm's supply schedule, we can solve for the market-clearing preliminary price from the perspective of firm i:

$$p^P = \bar{a} + a_1 + a_2 + e^P - q_i - (\alpha_{-i}a_{-i} + \beta_{-i}\bar{a} - \gamma_{-i}c_{-i} + \delta_{-i}p^P).$$

Firm i takes into account that as it increases its output q^i, it alters the price at which the rival's supply schedule is evaluated, causing the rival to supply less, which ameliorates the direct impact of firm i's output on price. Taking this indirect impact into account, we solve for how price is affected by firm i's output from its perspective:

$$p^P = \frac{a_i + (1 - \alpha_{-i})a_{-i} + (1 - \beta_{-i})\bar{a} + \gamma_{-i}c_{-i} - \delta_{-i}e^P - q_i}{1 + \delta_{-i}}.$$

Thus, from firm i's perspective, the marginal impact of an increase in q_i on price is only $\frac{1}{1+\delta_{-i}} < 1$.

Firm i chooses $q_i(p^P; c_i, a_i)$ to maximize expected profits, solving

$$\max_{q_i \in R} E\left[\left(\frac{a_i + (1 - \alpha_{-i})a_{-i} + (1 - \beta_{-i})\bar{a} + \gamma_{-i}c_{-i} - \delta_{-i}e^P - q_i}{1 + \delta_{-i}} + e^F - c_i\right)\right. \\ \left. \times q_i \middle| c_i, a_i, \frac{a_i + (1 - \alpha_{-i})a_{-i} + (1 - \beta_{-i})\bar{a} + \gamma_{-i}c_{-i} - \delta_{-i}e^P - q_i}{1 + \delta_{-i}}\right]. \tag{5}$$

Because firm i knows $a_i, \bar{a}$, and c_i, we can reformulate its information set in terms of net information:

$$\{a_i, c_i, (1 - \alpha_{-i})a_{-i} + \gamma_{-i}c_{-i} + e^P\}.$$

In turn, because firm i knows a_i and c_i, the net information, $(1-\alpha_{-i})a_{-i}+\gamma_{-i}c_{-i}+e^P$ is informationally equivalent from firm i's perspective to the market-clearing preliminary price $p^P(a_1,a_2,c_1,c_2,e^P)$. We can therefore write firm i's objective (5) as:

$$\max_{q_i\in R} E\left[\left(\frac{(a_i+(1-\alpha_{-i})a_{-i}+(1-\beta_{-i})\bar{a}+\gamma_{-i}c_{-i}-\delta_{-i}e^P-q_i)}{1+\delta_{-i}}+e^F-c_i\right)q_i \middle| c_i, a_i, (1-\alpha_{-i})a_{-i}+\gamma_{-i}c_{-i}+e^P\right].$$

By construction, the final demand shock e^F is orthogonal to all other shocks. Because it is not realized *ex ante*, its conditional expectation is zero, so it has no influence on the optimization. The remaining optimization problem is deterministic. The first-order condition describing firm i's best response to the rival firm $-i$'s conjectured strategy is

$$2q_i = a_i-(1+\delta)c_i+(1-\beta)\bar{a}+E[(1-\alpha)a_{-i}+\gamma c_{-i}-\delta e^P|(1-\alpha)a_{-i}+\gamma c_{-i}+e^P].$$

Given the Gaussian model structure, the conditional expectation is linear:

$$E[(1-\alpha_{-i})a_{-i}+\gamma_{-i}c_{-i}-\delta_{-i}e|(1-\alpha_{-i})a_{-i}+\gamma_{-i}c_{-i}+e^P]=\lambda_i[(1-\alpha_{-i})a_{-i}+\gamma_{-i}c_{-i}+e^P],$$

where

$$\lambda_i=\frac{(1-\alpha_{-i})^2\sigma_a^2+\gamma_{-i}^2\sigma_c^2-\delta_{-i}\sigma_e^2}{(1-\alpha_{-i})^2\sigma_a^2+\gamma_{-i}^2\sigma_c^2+\sigma_e^2} \tag{6}$$

is a projection coefficient. That is, firm i forecasts the residual component of demand (after subtracting off components that are predictable given its own private information or public information) by projecting it onto the residual information in the fact that the preliminary price is market clearing. We then solve for firm i's supply schedule,

$$q_i(p^P;a_i,c_i)=\frac{a_i-(1+\delta_{-i})c_i+(1-\beta_{-i})\bar{a}+\lambda_i[(1-\alpha_{-i})a_{-i}+\gamma_{-i}c_{-i}+e^P]}{2}, \tag{7}$$

which is a linear function of the net information. We then have the following lemma:

Lemma 1. If firm $-i$'s supply schedule takes the linear form of (4), then firm i's optimal supply schedule takes a linear form,

$$q_i(p^P;a_i,c_i)=\alpha_i a_i+\beta_i\bar{a}-\gamma_i c_i+\delta_i p^P,$$

as does the market-clearing preliminary price.

Proof. All proofs are in the Appendix. □

Observe that the end-of-period demand shock e^F does not enter the solution: because it is orthogonal to the information that firms have when they determine their supply functions, it disappears from the calculations. Conversely, even though demand, and therefore profit, is ultimately not influenced by the beginning-of-period demand shock e^P, it necessarily influences output via its influence on the market-clearing, beginning-of-period price.

Mechanically, the impact of the beginning-of-period demand shock and the consequent structure of beginning-of-period price results in an isomorphism between this setting and one where firms can base output on a noisy observation of end-of-period price. That is, e^P can be viewed as noise interfering with firms' observations of end-of-period prices. This can be seen explicitly in the signal extraction expressed in the projection coefficient in equation (6). As we set out in the Introduction, the model can therefore be viewed as a noisy REE model, but with the critical added element that firms strategically influence prices, and hence the net information of rivals.

With the linear structure of equilibrium strategies and prices established, it follows that the equilibria of the two frameworks coincide. In what follows, we pose the language and analysis within the noisy REE price framework. In this setting, the analogue of the volatility of the beginning-of-period demand shock e^P shrinking, is that the noisy signal of end-of-period price becomes more informative.

We next solve for all symmetric linear equilibria. To do this, we exploit symmetry of the two firms' equilibrium strategies, substituting $\alpha_1 = \alpha_2 = \alpha$, $\beta_1 = \beta_2 = \beta$, $\gamma_1 = \gamma_2 = \gamma$, and $\delta_1 = \delta_2 = \delta$. Recalling that λ_i is the projection on the firm's *net* information in price $(1-\alpha_{-i})a_{-i} + \gamma_{-i}c_{-i} + e^P$, we solve for α, β, γ, and δ in terms of λ.

To do this, we exploit symmetry, substituting for q_1 and q_2 into preliminary price to obtain

$$p^P = \frac{1-\lambda(1-\alpha)}{2}(a_1+a_2) + \beta\overline{a} + \frac{(1+\delta)-\lambda\gamma}{2}(c_1+c_2) - \lambda e^P. \tag{8}$$

We next solve for the equilibrium parameters of a firm's output strategy. Equation (7) reveals that firm 1's output weight on a_1 is $\frac{1}{2}$, which equals the sum of its direct output weight on a_1 plus its indirect output weight on a_1 via price. The coefficient on a_1 in price is $\frac{1-\lambda(1-\alpha)}{2}$. Substituting this into the conjectured form of output strategies reveals that its direct output weight α on a_1 solves

$$a_1: \quad \frac{1}{2} = \alpha + \delta\left(\frac{1-\lambda(1-\alpha)}{2}\right) \Rightarrow \alpha = \frac{1-\delta(1-\lambda)}{2+\delta\lambda}.$$

An analogous exercise yields the solutions for the other parameters:

$$\begin{aligned}
a_2: &\quad \lambda\frac{1-\alpha}{2} = \delta\left(\frac{1-\lambda(1-\alpha)}{2}\right) &&\Rightarrow \alpha = 1 - \frac{\delta}{\lambda(1+\delta)}\\
\bar{a}: &\quad \frac{1-\beta}{2} = \beta(1+\delta) &&\Rightarrow \beta = \frac{1}{3+2\delta}\\
c_1: &\quad -\frac{(1+\delta)}{2} = -\gamma + \delta\frac{(1+\delta)-\lambda\gamma}{2} &&\Rightarrow \gamma = \frac{(1+\delta)^2}{2+\lambda\delta}\\
c_2: &\quad \frac{\gamma\lambda}{2} = \delta\frac{(1+\delta)-\lambda\gamma}{2} &&\Rightarrow \gamma = \frac{\delta}{\lambda}\\
e: &\quad \delta(1-\lambda) = \frac{\lambda}{2} &&\Rightarrow \delta = \frac{\lambda}{2(1-\lambda)}.
\end{aligned}$$

One can show that after substituting for δ, that the two equations for γ (or α) yield the same solution. Substituting for δ, yields the primitive parameters α, β, and γ solely in terms of the price filter, λ:

$$\alpha = \frac{1-\lambda}{2-\lambda}, \quad \beta = \frac{1-\lambda}{3-2\lambda}, \quad \gamma = \frac{1}{2(1-\lambda)}, \quad \text{and} \quad \delta = \frac{\lambda}{2(1-\lambda)}. \tag{9}$$

The direct output weights, β and α, on common-value demand primitives *fall* with the information content of prices, λ, reflecting that more informative prices lead a rival to produce more aggressively in the *same* direction on this information. In contrast, the direct output weight on privately observed private-value costs *rises* when prices convey more information because firms produce in opposite directions on private values, as the unaffected firm responds to the output of the affected firm, and not the shock itself.

To solve for the equilibrium, we substitute the solutions for γ, β, α, and δ into the expression for the projection coefficient λ to obtain an equation that implicitly characterizes the equilibrium,

$$\lambda = \frac{\frac{1}{(2-\lambda)^2}\frac{\sigma_a^2}{\sigma_e^2} + \frac{1}{4(1-\lambda)^2}\frac{\sigma_c^2}{\sigma_e^2} - \frac{\lambda}{2(1-\lambda)}}{\frac{1}{(2-\lambda)^2}\frac{\sigma_a^2}{\sigma_e^2} + \frac{1}{4(1-\lambda)^2}\frac{\sigma_c^2}{\sigma_e^2} + 1}. \tag{10}$$

We divide through by σ_e^2 to emphasize that the equilibrium projection coefficient is a function of the signal-to-noise ratios. Observe that public information ($\bar{a}$) does not affect the weight firms place on price signals, because $\bar{a}$ is perfectly forecastable. The converse is *not* true: when prices contain more information – when λ is higher – firms reduce direct output weights on $\bar{a}$, that is, β falls.

We now provide necessary and sufficient conditions describing the existence of symmetric linear equilibria of this oligopoly game.

Proposition 1.

(i) If $\frac{\sigma_c^2}{\sigma_e^2} \leq 2$, then there is a unique symmetric linear equilibrium, and $\lambda < \frac{3}{4}$.

(ii) There exists an $s(\frac{\sigma_a^2}{\sigma_e^2}) > 2$, where $s' < 0$ with $s(0) = \frac{9}{4}$ and $s(\infty) = 2$ such that for $2 < \frac{\sigma_c^2}{\sigma_e^2} < s(\frac{\sigma_a^2}{\sigma_e^2})$, two symmetric linear equilibria exist, and the larger equilibrium λ exceeds $\frac{3}{4}$.

(iii) For $\frac{\sigma_c^2}{\sigma_e^2} \geq s(\frac{\sigma_a^2}{\sigma_e^2})$, no symmetric linear equilibrium exists.

Proposition 2. The equilibrium price projection λ rises with the signal-to-noise ratios $\frac{\sigma_c^2}{\sigma_e^2}$ and $\frac{\sigma_a^2}{\sigma_e^2}$.

- The information content of prices is more sensitive to privately observed private-value (cost) uncertainty than to privately observed common-value (demand) uncertainty: $\frac{\partial \lambda}{\partial \sigma_c^2} - \frac{\partial \lambda}{\partial \sigma_a^2} > 0$.
- The higher sensitivity of λ to privately observed private-value uncertainty than to privately observed common-value uncertainty rises with σ_c^2 and σ_a^2: $\frac{\partial}{\partial \sigma_c^2}(\frac{\partial \lambda}{\partial \sigma_c^2} - \frac{\partial \lambda}{\partial \sigma_a^2}) > 0$ and $\frac{\partial}{\partial \sigma_a^2}(\frac{\partial \lambda}{\partial \sigma_c^2} - \frac{\partial \lambda}{\partial \sigma_a^2}) > 0$.

Inspection of equation (A1) reveals that the left-hand side rises in both $\frac{\sigma_a^2}{\sigma_e^2}$ and $\frac{\sigma_c^2}{\sigma_e^2}$ (more quickly in $\frac{\sigma_c^2}{\sigma_e^2}$), implying the comparative statics. Intuitively, with more noise, the signal-to-noise ratio shrinks in equilibrium, so the projection coefficient λ falls. The reason why private-value uncertainty contributes more than common-value uncertainty to the information content of prices reflects their opposing impacts on a rival's actions. Firm 2's output weights a_1 positively (via price) because a_1 enters firm 2's payoff directly; but this causes firm 1 to reduce its output weight on a_1 – firms produce in the "same" direction on common demand. In turn, this reduces the information about a_1 in firm 2's price signal. In contrast, firms produce in "opposite" directions on private-value shocks, and when firm 2 takes an offsetting position against firm 1's output on c_1, this causes firm 1 to weight c_1 even more heavily. As a result, firm 2's price signal provides it even more information about c_1.

Thus, there are strategic complementarities in learning about private values, and strategic substitution in learning about common values. The learning complementarities about private values not only can cause equilibrium to break down when there is too much uncertainty about private values, but they can also lead to *multiple* equilibria in linear strategies. That is, for $2 < \frac{\sigma_c^2}{\sigma_e^2} < s(\frac{\sigma_a^2}{\sigma_e^2})$, two equilibria exist. In one equilibrium, a firm produces less aggressively on its private information about private values; as a result, price signals contain less information, so that its rival's output does not aggressively offset a firm's production on its costs, confirming the optimality of not weighing cost shocks heavily in output. In the other equilibrium, a firm produces aggressively on its private information; as a result, price signals contain more information, causing its rival to offset far more of a firm's production on its costs, validating a decision to heavily weight costs in output.

The driving force for equilibrium breakdown when there is substantial private-value uncertainty can also be seen in the stability properties of best responses. The best-response dynamics $\frac{d\lambda_1}{d\lambda_2}$ in the neighborhood of an equilibrium switch from being stable to unstable at $s(\frac{\sigma_a^2}{\sigma_e^2})$. Intuitively, $\frac{d\lambda_1}{d\lambda_2}$ rises with the information in prices about c_1 due to the strategic learning complementarities, and best responses become unstable once prices become sufficiently informative.

When $\sigma_c^2 \geq \frac{9}{4}\sigma_e^2$, an equilibrium does not exist because price signals convey enough information to firms about costs that the low cost firm goes infinitely long on output, whereas the high-cost firm takes the offsetting position, "producing" a negative infinity of output. One can show that existence of a linear equilibrium is always retrieved if there is an additional quadratic component to costs, so that $c_j(q_j) = \hat{c}q_j + c_j q_j + f q_j^2$, where $f > 0$ is public information. Increasing marginal costs ensure that when σ_c^2 is large, "trading" unbounded production roles by communicating via price is not optimal. Intuitively, with increasing marginal costs, the costs of large outputs eventually swamp the "gains from trade" between firms with different costs, ensuring bounded outputs. In turn, bounded outputs mean that noise-to-signal-ratios in price signals are bounded away from zero, so that λ, although increasing σ_c^2, remains strictly bounded away from one.

□ **Learning.** A firm learns about its rival's private information by forecasting via least-squares projection on its price signal. Firm 1's supply function in equation (7) expresses this via the term λ_1 that multiplies the elements of the net information in price. Firm 1 seeks to forecast the net shift in the price that it will receive due to the shocks privately observed by its rival, $(1-\alpha_2)a_2 + \gamma_2 c_2$. That is, it seeks to forecast the sum of the residual common-value demand shock left after accounting for its rival's *direct* production on that information, $(1-\alpha_2)a_2$, plus the *direct* price impact of its rival's cost shock via its output, $\gamma_2 c_2$. By expressing λ_1 in terms of a more fundamental projection, we now show that a firm treats information learned from this projection as if it were a publicly-observable common-value shock. This reflects that its rival knows the true shocks and hence can forecast this learning.

The net information in firm 1's price signal is $(1-\alpha_2)a_2 + \gamma_2 c_2 + e$. Firm 1 projects the grouped private signal $(1-\alpha_2)a_2 + \gamma_2 c_2 - \delta_2 e$ onto this net information using the projection coefficient λ_1. Our decomposition begins by projecting the *net* price shock $(1-\alpha_2)a_2 + \gamma_2 c_2$ onto the net information in firm 1's price signal; the associated projection coefficient is

$$\pi_1 = \frac{\mathrm{cov}((1-\alpha_2)a_2, (1-\alpha)a_2 + \gamma c_2 + e)}{\mathrm{var}((1-\alpha)a_2 + \gamma c_2 + e)} = \frac{(1-\alpha_2)^2\sigma_a^2 + \gamma_2^2\sigma_c^2}{(1-\alpha_2)^2\sigma_a^2 + \gamma_2^2\sigma_c^2 + \sigma_e^2}.$$

The projection coefficient λ_1 can be decomposed in terms of the direct projection π_1:

$$\lambda_1 = \frac{\mathrm{cov}((1-\alpha_2)a_2 + \gamma c_2 - \delta_2 e, (1-\alpha_2)a_2 + e)}{\mathrm{var}((1-\alpha_2)a_2 + \gamma_2 c_2 + e)} = \frac{(1-\alpha_2)^2\sigma_a^2 + \gamma_2^2\sigma_c^2 - \delta_2\sigma_e^2}{(1-\alpha_2)^2\sigma_a^2 + \gamma_2^2\sigma_c^2 + \sigma_e^2}.$$

Thus,

$$\lambda_1 = \pi_1 - \delta_2(1-\pi_1). \tag{11}$$

The decomposition reflects *firm* 2's awareness of both firm 1's learning, and of the errors that firm 1 makes in its learning. Because firm 2 sees a_2 and c_2, it knows what firm 1 learns and it knows the forecast error that firm 1 makes, and it weights them differently; in turn, firm 1 understands this. The term $-\delta_2(1-\pi_2)$ is the *reduction* in the weight on learned information that reflects firm 2's strategic exploitation of its private knowledge about the forecast error in firm 1's projection, $1-\pi_1$; and firm 1's understanding of this impact.

Using the decomposition in equation (11), we now establish:

Proposition 3. Information that firm 1 learns about $(1-\alpha)a_2 + \gamma c_2$ from its price signal is weighted as if it were publicly observed common-value information. That is,

$$\lambda = (1-\beta)\pi. \tag{12}$$

□ **Correlation in private information.** In this economy, the relevant private information that firms hold is the error in the other firm's forecast of its private information conditional on the information in prices. We next derive the correlation structure of this residual private information:

Proposition 4. The residual private information of firms, that is, the forecast errors in the projection of private information on price signals, is positively correlated between firms.

Defining ρ_c to be the projection of the common-value demand shock observed by a rival on the net information in a firm's price signal, and ρ_p to be the analogous projection of the rival's private-value cost shock, the correlation in the errors of the two firms in their forecasts of common demand shocks is $\rho_c^2\sigma_e^2$, the correlation of their forecasts of each other's cost shocks is $\rho_p^2\sigma_e^2$, and the correlation of one firm's forecast error of demand and the other firm's error in its forecast of costs is $\rho_c\rho_p\sigma_e^2$. It follows that more informative prices raise forecast error correlations. This result stands in contrast to the correlation of private information in models of strategic speculation in financial markets. In financial markets, the relevant private information of privately informed speculators – the differences between their forecasts of an asset's value and the market's forecast – is *negatively* correlated. This reflects that when speculators receive uncorrelated signals, their net private information tends to be on "opposite sides" of the asset's expected value (the price), that is, the errors in the market makers' forecasts of their signals are negatively correlated (Bernhardt, Seiler, and Taub, 2010; Back, Cao, and Willard, 2000; Foster and Viswanathan, 1996). What underlies the opposing correlation structures in speculation and duopoly frameworks is that the relevant private information of firms concerns *levels* (e.g., of demand) rather than *differences*.

□ **Output weights.** From an economic perspective, what matters is not just a firm's direct output weight on a fundamental, but rather its *total* direct plus indirect (via its price signal) output weight. We now derive both how an individual firm's output weights compare with their full-information counterparts *and* how aggregate (across firms) output weights compare. In a full-information setting, firm 1's output is

$$q_1 = \frac{\bar{a} + a_1 + a_2 - c_1 - q_2}{2} = \frac{\bar{a} + a_1 + a_2 - 2c_1 + c_2}{3}. \tag{13}$$

The second equality shows that because firms produce in the same direction on common components, they scale back weights relative to monopoly levels; but because firms produce in opposite directions on private-value components, a firm hit by a shock raises (the absolute value of) its weight.

We next show how private information affects how firms' outputs weight differently public and private information about common and private values.

Proposition 5. In the private information economy,

- Firms overproduce on the publicly-known, common $\bar{a}$ relative to a full-information setting.
- Firms overproduce by the same amounts on common-value learned and public information.
- Aggregate (across firms) output weights on publicly-known, common demand exceed those on privately observed common demand by the product $\beta\alpha$ of the direct output weights on privately observed and publicly observed demand.
- Aggregate (across firms) output weights on a firm's privately observed private-value cost shocks exceed those on publicly-known, private-value costs by $\frac{\beta}{2}$.

A firm's weight on the publicly-known, common value $\bar{a} = \hat{a} - \hat{c}$ of $\frac{2-\lambda}{2(3-2\lambda)} = \frac{1-\beta}{2}$ exceeds its full-information economy weight of $\frac{1}{3}$. Similarly, reflecting Proposition 3, firm 1 also weights its forecast of its rival's net privately observed demand shock by $\frac{1-\beta}{2}$; and it weights its forecast of how its rival's output is affected by its cost shock by $\frac{1-\beta}{2}$. That firms overproduce on both public and this learned information relative to a full-information setting reflects classical signal-jamming incentives. Each firm overproduces on $\overline{a}$, which captures the market's *ex ante* expected profitability, to try to convince its rival that the market is less profitable. This overproduction grows with λ, reflecting that when prices are more informative, firms weight price signals by more, so that a rival's output becomes more sensitive to signal jamming. Thus, there are *synergies* between the informativeness of price signals and signal-jamming incentives.

The incentives to over weight publicly-known common values are clear. In contrast, it is unclear whether firms should overweight or underweight *privately* observed shocks, and how this should vary with their common- and private-value natures. The proposition shows that in our private information economy, firms collectively weight publicly observed common-value shocks by exactly $\beta\alpha$ more than when those shocks are privately observed. This is *reversed* for private-value shocks: firms collectively weight privately observed private-value cost shocks by $\frac{\beta}{2}$ more than when they are publicly observed. Thus, the combined output weights by the two firms on the privately observed private-value shocks exceed the combined weights on publicly-known common demand, which, in turn, exceed the combined weights on privately observed common demand. This result reflects that firms weight information learned from price signals about privately observed shocks just like public information, so the net effects depend on how the informed firm's output weights its residual private information. An informed firm acts like a monopolist on this information – but a monopolist weights common-value information by *less* than a duopoly collectively would; whereas a monopolist weights private-value information by *more* than a duopoly precisely because its rival partially offsets some of the price impact of its output.

Proposition 6. Each firm acts as a monopolist on the common demand shock that it privately observes, whereas firms collectively act as monopolists on a privately observed private cost shock to one firm: regardless of the informativeness of prices, firm i's total (direct plus indirect via price) output weight on a_i is always $\frac{1}{2}$; whereas the aggregate (across firms) output weight on c_i is $\frac{1}{2}$.

An individual firm always overweights a privately observed common demand shock relative to full-information settings, trying to convince its rival that the *absolute* magnitude of the shock is smaller. Indeed, firm i overweights a_i by exactly $\frac{\beta}{2}$ more than it overweights publicly-known common demand, $\overline{a}$, where $\frac{\beta}{2}$ is half of its direct output weight on $\overline{a}$. In contrast, the two firms collectively weight a privately observed private-value cost shock by (minus) $\frac{1}{2}$, as the firms take offsetting positions on what the rival learns about the cost shock. That is, a firm places a weight $-[\frac{1}{2}+\frac{\lambda}{4(1-\lambda)}]$ on its own marginal cost shock, but its rival offsets $\frac{\lambda}{4(1-\lambda)}$. The offsetting positions reflect that firm i can forecast what its rival learns. It follows that relative to a full information setting, collective firm output always weighs a firm's privately observed private-value costs by more (by $\frac{1}{2}-\frac{1}{3}=\frac{1}{6}$). In contrast, firms collectively overweight privately observed common demand shocks relative to a full-information setting *only* if price signals are informative enough that $\lambda>\frac{1}{2}$.

This contrast highlights the very different strategic consequences of privately observed common-value and private-value shocks. In both cases, a firm has an incentive to signal jam on its private information. However, with privately observed common-value demand shocks, firms produce in the "same direction," so enough information must be transmitted via prices for the signal-jamming strategic effect to dominate the reduced information that a rival has about a demand shock in the private information economy in terms of collective firm output. As price signals become less noisy, the weight λ that firms place on their price signals rises, and hence so do both the signal-jamming strategic effects and the information conveyed to a rival via its price signal. As a result, output weights rise, and they eventually exceed their full-information levels.

Proposition 7. Relative to a full-information setting, a firm overweights its own private-value shock c_i, and its rival's private-value and common-value shocks c_{-i} and a_{-i}, if and only if prices are sufficiently informative. The required informativeness of prices λ is higher for common-value shocks than private-value shocks.

TABLE 1 Output Weights
Column 1 has the full-information economy weights; column 2 has the private-information economy weights; column 3 has the difference in private-information economy weights on private and public information. The first row shows weights on publicly-known common values. The middle cell shows weights on privately observed common value shocks a_i: the top row is the direct weight after netting out a firm's own indirect weight via price; the second row is the indirect weight on information observed by the rival via price; and the third row is the sum. The bottom cell has a similar decomposition for privately observed private-value shocks. See Propositions 5 through 7

	Full Info.	Noisy	Difference Noisy Private v. Noisy $\overline{a}$
$\overline{a}$	$\frac{1}{3}$	$\frac{2-\lambda}{2(3-2\lambda)} = \frac{1-\beta}{2}$	
a_i own	$\frac{1}{3}$	$\frac{1}{2}$	$\frac{1-\lambda}{2(3-2\lambda)} = \frac{\beta}{2} > 0$
a_{-i} indirect	$\frac{1}{3}$	$\frac{\lambda}{2(2-\lambda)}$	$-\frac{4-7\lambda+3\lambda^2}{2(2-\lambda)(3-2\lambda)} < 0$
$a_i + a_{-i}$	$\frac{2}{3}$	$\frac{1}{2} + \frac{\lambda}{2(2-\lambda)}$	(v. $2\bar{a}$) $-\frac{(1-\lambda)^2}{(2-\lambda)(3-2\lambda)} = -\beta\alpha < 0$
c_i own	$-\frac{2}{3}$	$-\left(\frac{1}{2} + \frac{\lambda}{4(1-\lambda)}\right)$	$-\frac{2-\lambda}{4(1-\lambda)(3-2\lambda)} < 0$
c_{-i} indirect	$\frac{1}{3}$	$\frac{\lambda}{4(1-\lambda)}$	$-\frac{1}{4}\frac{4\lambda^2-9\lambda+4}{(1-\lambda)(3-2\lambda)} > 0$ if $\lambda > 0.61$
$c_i + c_{-i}$	$-\frac{1}{3}$	$-\frac{1}{2}$	(v. $-\overline{a}$) $\frac{1-\lambda}{2(3-2\lambda)} = \frac{\beta}{2} > 0$

To understand the intuition, first suppose that price signals are extremely noisy. Then, a firm acts as a monopolist on the private-value shocks that it privately observes, and its output minimally weighs the shocks observed by its rival. Thus, the full-information output weights on these shocks are greater (in particular, because a rival does not offset any output on a private-value shock that it does not learn about). Now reduce the noise in price signals. Then λ rises. As a result, the weights firm i places on c_i, c_{-i}, and a_{-i} rise, reflecting the increased learning and signal-jamming incentives. Eventually, the synergies between these two forces dominate, reflecting that firms always overweight public/learned information relative to a full-information setting. That prices must be more informative for the weights on common demand shocks observed by a rival to exceed their public-information levels than for the weights on a rival's private cost shocks to do so reflects the strategic learning complementarities about private values and strategic learning substitutes about common values. Table 1 presents the calculations for the results in Propositions 5–7.

3 Dynamics

■ Our static analysis can be extended to a dynamic setting. In Bernhardt and Taub (2013), we develop a stationary dynamic version of our static model and analyze it numerically. In the dynamic model, firms are buffeted each period by new, long-lived shocks to demand and costs that evolve according to persistent autoregressive processes driven by Gaussian innovations. Now firms learn about rivals' private shocks via current and past price signals. In turn, each firm chooses output strategically, knowing that its current output also influences a rival's future actions – each firm combines the information it extracts from the history of price signals with that in the history of its privately- and publicly observed innovations to determine how much to produce.

We establish that many key structural properties of our static model also hold in the steady state of the dynamic model. For example, the economics of how firms weigh privately observed private- and common-value shocks carry over: in the steady state of the dynamic model, at any given lag, the output weight placed by one firm on a common-value shock that it privately observes equals the collective output weight that the *two* firms together have on a privately observed private-value shock. So, too, in the dynamic model, just as in the static model, (i) the level of public information does not affect how firms filter information in price signals or how they weigh information that is privately observed, (ii) information learned from the history of price signals is treated in output decisions as if it is public information, and (iii) indeed, only the current value of public information components and not the timing of their arrival enter firm output decisions. More generally, (iv) the result in the static setting that firms learn more about privately observed private-value shocks than privately observed common-value shocks carries over to the dynamic economy for exactly the same reasons. That is, as a firm learns more about a rival's private cost shocks, it produces more intensively in the *opposite* direction, leading its rival to raise its output weight further. In contrast, as a rival learns more about a common value shock observed by a firm, the rival produces more intensively on that information, causing the firm seeing the shock to cut back.

However, a conjecture that with autoregressive shocks, the firms' output choices in the dynamic setting simply echo those processes, with the same fixed autocorrelations – essentially replicating the static model iteratively – is misplaced for two reasons. First, firms understand that their strategic actions affect a rival's inferences not just in the current period, but also in future periods. As a result, the incentives to manipulate a rival's beliefs (e.g., a firm's incentives to convince a rival that its costs are lower than they actually are) rise in a dynamic setting because a firm gains not only in the current period from this manipulation, but also in future periods. Second, the correlation structure of the firms' information evolves – each firm learns more and more over time via new price signals about older shocks observed by its rival, so that less of that information remains private information to the rival. Output weights on public and private components differ, so it follows directly that the dynamics are far more complicated. As a firm learns more about a common value shock observed by its rival, its output weight rises, causing its rival to cut back. In contrast, weights on privately observed, private-value shocks rise with lags: as a firm learns more about a rival's private cost shocks, it produces more intensively in the *opposite* direction, leading its rival to further raise its output weight. Consequently, price signals convey more information about a rival's privately observed older private-value shocks than privately observed older common-value shocks. Moreover, the induced patterns on outputs mean that firm profits move together with privately observed private-value shocks and publicly observed common value shocks, but firm profits move in opposite directions on privately observed common-value and publicly observed private-value shocks.

The static setting also informs about the comparative statics of the dynamic model. For example, the partial learning that takes place in the static model is mirrored in the dynamic model. However, when shocks are larger or more persistent, learning is faster, which increases the positive correlation in the private information of the two firms. Accordingly, this alters the evolution of output weights at longer lags.

4 Conclusion

■ In this article, we characterize a duopoly buffeted by demand and cost shocks. Firms learn about these shocks via common observation, private observation, and the information conveyed by noisy price signals. We derive how firms extract information from noisy price signals and strategically choose output to manipulate a rival's beliefs and hence output. We distinguish how the nature of information – public versus private – and of what firms learn about – common versus private values – affect equilibrium outcomes.

Firms treat private information about common values differently from private information about private values. In particular, firm outputs weigh private information about private value by more than common values. As a result, price signals contain more information about private-value shocks than common-value shocks. This reflects strategic learning complements in output on private values, and strategic learning substitutes in output on common values, due to the fact that firms produce in the same direction on common values, but in opposite directions on private values. In turn, total firm output weights on privately observed private-value shocks exceed those on publicly-known private values; but total weights on privately observed common value shocks are less than those on publicly-known common values. A further consequence is that when most private information concerns common values, equilibrium is unique; but when private information about private values predominates, there can be multiple equilibria, and, indeed, with extensive private information about private values, equilibrium can break down.

Appendix

In this Appendix, we provide proofs for Lemma 1 and Propositions 1-4.

Proof of Lemma 1. To complete the proof, we rewrite output in terms of the beginning-of-period price rather than net information. To begin, write p^P in terms of net information and known terms:

$$p^P = \frac{[a_i + (1-\beta_{-i})\overline{a} - q_i] + [(1-\alpha_{-i})a_{-i} + \gamma_{-i}c_{-i} + e^P]}{1+\delta_{-i}}.$$

Thus, we can write net information as

$$[(1-\alpha_{-i})a_{-i} + \gamma_{-i}c_{-i} + e^P] = (1+\delta_{-i})p^P - [a_i + (1-\beta_{-i})\overline{a} - q_i].$$

Substituting this into the solution for q_i in equation (7) yields

$$q_i(p^P; a_i, c_i) = \frac{a_i - (1+\delta_{-i})c_i + (1-\beta_{-i})\overline{a} + \lambda_i[(1+\delta_{-i})p^P - [a_i + (1-\beta_{-i})\overline{a} - q_i]]}{2}.$$

Solving for q_i yields

$$q_i = \frac{a_i - (1+\delta_{-i})c_i + (1-\beta_{-i})\overline{a} + \lambda_i[(1+\delta_{-i})p^P - [a_i + (1-\beta_{-i})\overline{a}]]}{2-\lambda_i}.$$

This is linear in preliminary price as asserted, with

$$\delta_i \equiv \frac{\lambda_i}{2-\lambda_i}(1+\delta_{-i}).$$

□

Proof of Proposition 1. Our geometric approach rearranges the equation (10) describing the equilibrium λ as

$$\frac{(1-\lambda)^2}{(2-\lambda)^2}\frac{\sigma_a^2}{\sigma_e^2} + \frac{1}{4}\frac{\sigma_c^2}{\sigma_e^2} = \lambda\frac{3-2\lambda}{2}. \tag{A1}$$

The LHS decreases in λ and exceeds the RHS at $\lambda = 0$, but is less than the RHS at $\lambda = 1$ if $\frac{\sigma_c^2}{\sigma_e^2} < 2$, implying an intersection in [0, 1]. The RHS achieves a maximum of $\frac{9}{16}$ at $\lambda = \frac{3}{4}$. Hence, the LHS necessarily exceeds the RHS if $\frac{\sigma_c^2}{\sigma_e^2} > \frac{9}{4}$. For $\frac{\sigma_c^2}{\sigma_e^2} \in (2, \frac{9}{4})$, the LHS intersects the RHS only if $\frac{\sigma_a^2}{\sigma_e^2}$ is sufficiently small, where "sufficiently small" is decreasing in $\frac{\sigma_c^2}{\sigma_e^2}$. To complete the proof, we show that (i) if there is an intersection, then there are two, but (ii) only intersections with $\lambda < 1$ are associated with equilibria – fixed points with $\lambda > 1$ feature negative expected profits.

Substitute $z(\lambda) = \frac{1}{(2-\lambda)^2}\frac{\sigma_a^2}{\sigma_e^2} + \frac{1}{4(1-\lambda)^2}\frac{\sigma_c^2}{\sigma_e^2}$ to write

$$\lambda = \frac{z(\lambda) - \frac{\lambda}{2(1-\lambda)}}{z(\lambda)+1} \Leftrightarrow \lambda z(\lambda) + \lambda = z(\lambda) - \frac{\lambda}{2(1-\lambda)}. \tag{A2}$$

Solve this for

$$(1-\lambda)z(\lambda) = \lambda\left(1 + \frac{1}{2(1-\lambda)}\right) = \lambda\frac{3-2\lambda}{2(1-\lambda)}.$$

Cross-multiply by $1-\lambda$, and substitute for $z(\lambda)$ to obtain

$$\frac{(1-\lambda)^2}{(2-\lambda)^2}\frac{\sigma_a^2}{\sigma_e^2} + \frac{1}{4}\frac{\sigma_c^2}{\sigma_e^2} = \lambda\frac{3-2\lambda}{2}. \tag{A3}$$

Note that *LHS*(*0*) > *RHS*(*0*) and $LHS(1) < RHS(1)$, for $\frac{\sigma_c^2}{\sigma_e^2} < 2$, implying at least one solution with $\lambda \in (0, 1)$. The RHS achieves a maximum of $\frac{9}{16}$ at $\lambda = \frac{3}{4}$, so that the LHS necessarily exceeds the RHS if $\frac{\sigma_c^2}{\sigma_e^2} > 9/4$. The LHS is increasing in

$\frac{\sigma_a^2}{\sigma_e^2}$, so that the maximal $\frac{\sigma_c^2}{\sigma_e^2} \in (2, \frac{9}{4})$ for which the LHS ever intersects the RHS is decreasing in $\frac{\sigma_a^2}{\sigma_e^2}$. Hence, there exists a $s(\frac{\sigma_a^2}{\sigma_e^2}) \in (2, \frac{9}{4}]$ such that the LHS and RHS intersect if and only if $\frac{\sigma_c^2}{\sigma_e^2} \le s(\frac{\sigma_a^2}{\sigma_e^2})$.

The RHS is a concave (quadratic) function of λ, that reaches a maximum of $\frac{9}{16}$ at $\lambda = \frac{3}{4}$, and it equals $\frac{1}{2}$ at $\lambda = 1$. The first derivative of the LHS is $\frac{-2(1-\lambda)}{(2-\lambda)^3}\frac{\sigma_a^2}{\sigma_e^2} < 0$ for $\lambda < 1$ (and going to zero at 1). The second derivative of the LHS is $2\frac{2\lambda-1}{(2-\lambda)^4}\frac{\sigma_a^2}{\sigma_e^2}$, which is negative for $\lambda < \frac{1}{2}$, and positive for $\lambda > \frac{1}{2}$. Hence, when the LHS and RHS intersect, there are two intersections. It remains first to determine when there are multiple intersections for $\lambda \in [0, 1)$, and then to show that $\lambda \ge 1$ cannot be part of an equilibrium. If there are multiple intersections for $\lambda \in [0, 1)$, the larger one exceeds $\frac{3}{4}$, and multiple intersections occur for $\lambda \in [0, 1)$ if and only if LHS and RHS intersect (and are not just tangent at $\frac{3}{4}$) and LHS(1)> RHS(1) (as the concavity of the RHS and convexity of LHS for $\lambda > \frac{1}{2}$ imply that the LHS crosses the RHS from below and there is not a third intersection). Finally, LHS(1)>RHS(1) if and only if $\frac{\sigma_c^2}{\sigma_e^2} > 2$.

It remains to show that a fixed point with $\lambda > 1$ cannot be part of an equilibrium. Inspection immediately reveals that $\lambda \ge 1.5$ cannot be a solution to equation (A3). Inspection of (9) reveals that for $\lambda \in (1, 1.5)$ that γ, β, α, and δ are all negative. We now show that this implies that the solutions for output weights would mean that firms expect to lose money, a contradiction. Profits are:

$$p - c_1 = (1 - \lambda(1-\alpha))(a_1 + a_2) + \beta\bar{a} + ((1 + \delta - \lambda\gamma)(1/2))(c_1 + c_2) - c_1 - \lambda e$$

times

$$q_1 = \alpha a_1 + \beta\bar{a} - \gamma c_1 + \delta(p + e).$$

The coefficients on $a_1, a_2, \bar{a}, e$ in $p - c_1$ are negative, but they are positive in q_1, implying that they contribute expected losses to the firm. In particular, the coefficient in price of (i) a_i is $1 - \lambda(1-\alpha) = 1 - \frac{\lambda}{2-\lambda} = 2\frac{1-\lambda}{2-\lambda} < 0$, (ii) $\bar{a}$ is $\beta < 0$, (iii) e is $-\lambda < 0$. Conversely, the coefficient in q_1 of (i) a_1 is positive because $|\alpha| = |\frac{1-\lambda}{2-\lambda}| < |2\delta\frac{1-\lambda}{2-\lambda}|$, that is, $|\frac{\lambda}{1-\lambda}| > 1$ for $\lambda \in (1, 1.5)$, (ii) a_2 is positive because $\delta < 0$, (iii) $\bar{a}$ is positive because $\beta(1+\delta) > 0$, and (iv) e is $\frac{\lambda}{2} > 0$. Finally, substituting γ and δ, reveals that the coefficients on c_1 and c_2 in $p - c_1$ are $-\frac{1}{2}$ and $\frac{1}{2}$, respectively. The coefficient on c_1 in q_1 is $\frac{\lambda-2}{4(1-\lambda)} > 0$, and that on c_2 is $\frac{\lambda}{4(1-\lambda)} < 0$. Hence, the signs on c_1 in $p - c_1$ and q_1 differ, as do those on c_2. Hence, they also contribute expected losses. Thus, $\lambda \in [1, 1.5)$ cannot be part of an equilibrium: the solution implies that firms necessarily earn negative expected profits, and they can earn zero profits by not producing. □

Proof of Proposition 2. Examining equation (A3) reveals that variances enter only according to their ratios $\frac{\sigma_a}{\sigma_e^2}$ and $\frac{\sigma_c}{\sigma_e^2}$ on the left-hand side and do not enter the right-hand side. Increasing σ_e^2 shrinks the left-hand side and hence the intersection point, whereas increasing σ_a or σ_c does the opposite. To see that λ is more sensitive to σ_c^2 than σ_a^2, observe that the coefficient of σ_c^2 in (A3) is $\frac{1}{4}$, whereas that on σ_a^2 is $\frac{(1-\lambda)^2}{(2-\lambda)^2}$, which is less than $\frac{1}{4}$ when there is private information (so $\lambda > 0$), and decreasing in λ. □

Proof of Proposition 3. Imposing symmetry, begin by solving for π in terms of λ:

$$\pi = \frac{\lambda + \delta}{1 + \delta} = \frac{\lambda + \frac{\lambda}{2(1-\lambda)}}{1 + \frac{\lambda}{2(1-\lambda)}} = \frac{3\lambda - 2\lambda^2}{2 - \lambda} = \frac{3 - 2\lambda}{2 - \lambda}\lambda,$$

where the first equality follows from substituting for δ. Next, observe that

$$1 - \beta = 1 - \frac{1-\lambda}{3-2\lambda} = \frac{2-\lambda}{3-2\lambda}.$$

Thus,

$$(1-\beta)\pi = \frac{2-\lambda}{3-2\lambda}\frac{3-2\lambda}{2-\lambda}\lambda = \lambda.$$

The supply schedule equation (7) shows that the firm's output weights this projection by one half. The weight $\frac{(1-\beta_2)}{2}$ on the learned component, πa_2, equals that on publicly known common demand—the projected or learned part is weighed just like public information. □

Proof of Proposition 4. Firm 1's projection of the common value demand shock observed by firm 2, a_2, on the net information in firm 1's price signal is:

$$\begin{aligned} P[a_2|(1-\alpha)a_2 + \gamma c_2 + e] &= \frac{\mathrm{cov}(a_2, (1-\alpha)a_2 + \gamma c_2 + e)}{\mathrm{var}((1-\alpha)a_2 + \gamma c_2 + e)} \\ &= \frac{(1-\alpha)\sigma_{a_2}^2}{(1-\alpha)^2\sigma_a^2 + \gamma^2\sigma_{c_2}^2 + \sigma_e^2} \equiv \rho_c. \end{aligned} \tag{A4}$$

Firm 2's residual private information is the forecast error:

$$a_2 - \rho_c((1-\alpha)a_2 + \gamma c_2 + e).$$

Hence, the correlation of firm 1 and firm 2 residual private information about demand is:

$$\mathrm{cov}(a_1 - \rho_c((1-\alpha)a_1 + \gamma c_1 + e), a_2 - \rho_c((1-\alpha)a_2 + \gamma c_2 + e)) = \rho_c^2\sigma_e^2 > 0.$$

Similarly, Firm 1's projection of the private-value cost shock to firm c_2 on the net information in firm 1's price signal is:

$$P[c_2|(1-\alpha)a_2 + \gamma c_2 + e] = \frac{\mathrm{cov}(c_2, (1-\alpha)a_2 + \gamma c_2 + e)}{\mathrm{var}((1-\alpha)a_2 + \gamma c_2 + e)}$$

$$= \frac{\gamma\sigma_{c_2}^2}{(1-\alpha)^2\sigma_a^2 + \gamma^2\sigma_{c_2}^2 + \sigma_e^2} \equiv \rho_p. \tag{A5}$$

Hence, the correlation in the residual private information of the two firms about costs is:

$$\mathrm{cov}(c_1 - \rho_p((1-\alpha)a_1 + \gamma c_1 + e), c_2 - \rho_p((1-\alpha)a_2 + \gamma c_2 + e)) = \rho_p^2\sigma_e^2 > 0,$$

and the correlation in their residual private information about demand and costs is $\rho_c\rho_p\sigma_e^2 > 0$. □

References

AGHION, P., BOLTON, P., HARRIS, C., AND JULLIEN, B. "Optimal Learning By Experimentation." *Review of Economic Studies*, Vol. 58 (1991), pp. 621–654.

ALEPUZ, D. AND URBANO, A. "Learning in Asymmetric Duopoly Markets: Competition in Information and Market Correlation." *Spanish Economic Review*, Vol. 7 (2005), pp. 1435–5469.

ATHEY, S. AND BAGWELL, K. "Collusion with Persistent Cost Shocks." *Econometrica*, Vol. 76 (2008), pp. 493–540.

BACK, K., CAO, C. H., AND WILLARD, G. A. "Imperfect Competition Among Informed Traders." *Journal of Finance*, Vol. 55 (2000), pp. 2117–2155.

BAGWELL, K. AND RAMEY, G. "Oligopoly Limit Pricing." *RAND Journal of Economics*, Vol. 22 (1991), pp. 155–172.

BERNHARDT, D., SEILER, P., AND TAUB, B. "Speculative Dynamics." *Economic Theory*, Vol. 44 (2010), pp. 1–52.

CAMINAL, R. "A Dynamic Duopoly Model with Asymmetric Information." *Journal of Industrial Economics*, Vol. 38 (1990) pp. 315–333.

CAMINAL, R. AND VIVES, X. "Why Market Shares Matter: An Information-Based Theory." *RAND Journal of Economics*, Vol. 27 (1996), pp. 221–239.

FOSTER, D. AND VISWANATHAN, S. "Strategic Trading When Agents Forecast the Forecasts of Others." *Journal of Finance*, Vol. 48 (1996), pp. 1437–1478.

GAL-OR, E. "Information Sharing in Oligopoly." *Econometrica*, Vol. 53 (1985), pp. 329–343.

———."Information Transmission—Cournot and Bertrand Equilibria." *Review of Economic Studies*, Vol. 53 (1986), pp. 85–92.

———. "The Advantages of Imprecise Information." *RAND Journal of Economics*, Vol. 19 (1988), pp. 266–275.

HARRINGTON, J. "Limit Pricing When the Potential Entrant Is Uncertain of Its Cost Function." *Econometrica*, Vol. 54 (1986), pp. 429–437.

———. "Oligopolistic Entry Deterrence under Incomplete Information." *RAND Journal of Economics*, Vol. 18 (1987), pp. 211–231.

———. "Experimentation and Learning in a Differentiated-Products Duopoly." *Journal of Economic Theory*, Vol. 66 (1995), pp. 275–288.

KELLER, G. AND RADY, S. "Optimal Experimentation in a Changing Environment." *Review of Economic Studies*, Vol. 66 (1999), pp. 475–507.

———. "Price Dispersion and Learning in a Dynamic Differentiated-Goods Duopoly." *RAND Journal of Economics*, Vol. 34 (2003), pp. 138–165.

KLEMPERER, P.D. AND MEYER, M.A. "Supply Function Equilibria in Oligopoly under Uncertainty." *Econometrica*, Vol. 57 (1989), pp. 1243–1277.

KYLE, A. "Informed Speculation with Imperfect Competition." *Review of Economic Studies*, Vol. 56 (1989), pp. 317–355.

MAILATH, G. "Simultaneous Signaling in an Oligopoly Model." *Quarterly Journal of Economics*, Vol. 104 (1989), pp. 417–427.

MALUEG, D. AND TSUTSUI, S. "Duopoly Information Exchange: The Case of Unknown Slope." *International Journal of Industrial Organization*, Vol. 14 (1996), pp. 119–136.

MCLENNAN, A. "Price Dispersion and Incomplete Learning in the Long Run." *Journal of Economic Dynamics and Control*, Vol. 7 (1984), pp. 331–347.

MIRMAN, L., SALGUEIRO, E., AND SANTUGINI, M. "Noisy Signaling in Monopoly." *International Review of Economics and Finance*, Vol. 29 (2014), pp. 504–511.

MIRMAN, L., SAMUELSON, L. AND URBANO, A. "Duopoly Signal Jamming." *Economic Theory*, Vol. 3 (1993), pp. 129–149.

RIORDAN, M. "Imperfect Information and Dynamic Conjectural Variations." *RAND Journal of Economics*, Vol. 16 (1985), pp. 41–50.

RUSTICHINI, A. AND WOLINSKY, A. "Learning About Variable Demand in the Long Run." *Journal of Economic Dynamics and Control*, Vol. 19 (1995), pp. 1283–1292.

SEILER, P. AND TAUB, B. "The Dynamics of Strategic Information Flows in Stock Markets." *Finance and Stochastics*, Vol. 12 (2008), pp. 43–82.

VIVES, X. "Duopoly Information Equilibria: Cournot and Bertrand." *Journal of Economic Theory*, Vol. 34 (1984), pp. 71–94.

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