



Centre Interuniversitaire sur le Risque,
les Politiques Économiques et l'Emploi

Cahier de recherche/Working Paper **14-23**

Learning in a Perfectly Competitive Market

Leonard J. Mirman

Egas M. Salgueiro

Marc Santugini

Septembre/September 2014

Mirman: Department of Economics, University of Virginia

lm8h@virginia.edu

Salgueiro: DEGEI/GOVCOPP, Universidade de Aveiro

egas@ua.pt

Santugini: Department of Applied Economics, HEC Montréal; and CIRPÉE

marc.santugini@hec.ca

Abstract:

We study the informativeness of the price in a perfectly competitive market. A price-taking firm sells a good whose quality is unknown to some buyers. The uninformed buyers use the price to infer information about quality. The shape of the supply curve influences the amount of information contained in the equilibrium price.

Keywords: Asymmetric information, Learning, Perfect competition, Rational Expectations

JEL Classification: D2, D41, D8, L1

1 Introduction

In an economy with asymmetric information, learning agents understand that information is conveyed by the price system. Taking account of the information in the price system, they adjust their decisions accordingly. The amount of information contained in the price depends on the market structure as well as the decisions of all the agents. Hence, learning and decision-making cannot be separated, but are integrated through the price system. The relationship between the informativeness of the price and the market outcomes must be studied by taking account of the price system and its interaction with the learning activity and the decisions of all agents.

This paper studies the informativeness of the price in a perfectly competitive market. To that end, we embed learning in a model in which a perfectly competitive, representative firm, which has complete information about the market, sells a good whose quality is unknown to some buyers.¹ Demand is composed of both informed and uninformed buyers. The uninformed buyers use Bayesian methods to infer information from observing the price. On the supply side, the representative, price-taking firm produces and sells the good. The cost of production is assumed to be increasing in quality and quantity. There is also a demand shock, which is known to the firm but unknown to buyers, that prevents the market price from being perfectly informative about quality.

After characterizing the competitive equilibrium, the relationship between learning and market outcomes is addressed. Information flows and market outcomes are entwined because the uninformed buyers, who learn from prices, also participate in trading. In fact, the presence of uninformed buyers and their learning activity influence the informational content of the price. There is thus a two-way relation between trading and learning. Not only does learning from prices have an effect on decisions, but the agents' decisions impact the market price, thus influencing the informational content of the price and the learning process.

To study the impact of learning on the competitive outcome, we first

¹The assumption of *one* price-taking firm has no bearing on the results.

study the effect of the uninformed buyers' learning activity on demand. Since the uninformed buyers make decisions on the basis of both prior beliefs and the price, learning can be decomposed into two effects; a beliefs effect and a price effect. The beliefs effect reflects the change in behavior due to the asymmetry of information and the use of prior beliefs. The direction of the beliefs effect depends only on the bias of the prior beliefs. The price effect reflects the change in behavior due to updating beliefs. Unlike the beliefs effect, the sign of price effect depends on the bias of the prior beliefs and the demand shock.

After discussing the structure of demand in a learning environment, we study the informativeness of the price in the competitive equilibrium. Although it is the interaction of price-taking agents that determines the informativeness of the price, we identify two important aspects of the dissemination of information. First, the information conveyed by the price depends on the source of information. In our model, there is information present in both the demand and the supply sides. Specifically, on the demand side, there are informed buyers whose actions convey information. On the supply side, the information is due to the dependence of the marginal cost on quality. The second aspect that affects the information contained in the price-signal is the technology of the firm. Although a perfectly competitive firm has no control over prices and thus has no ability to influence directly the amount of information conveyed by prices, the technology and thus the shape of the cost function do affect the informativeness of the price. Indeed, in a competitive equilibrium, the price is equal to the marginal cost and thus the conveyance of information coming from either demand or supply depends on the functional form of the marginal cost. Hence, the distribution of the price-signal is defined by the distribution of the marginal cost.

To clarify the interaction of the information and the functional form of the marginal cost on the informativeness of the price, we split the discussion in two parts. We first discuss the competitive outcome in the case of horizontal supply curve, and then in the case of increasing supply curve.

When the marginal cost is constant implying a horizontal supply curve, the distribution of the price-signal is degenerate. The reason is that the

marginal cost is independent of output and thus independent of the demand shock. Hence, the information from the demand side is irrelevant, i.e., it is not conveyed by the price. In the case of a flat supply curve, two extreme outcomes are possible. On the one hand, when the constant marginal cost is independent of quality, the price has no informational content and is thus uninformative about quality.² In this case, the uninformed buyers revert to their prior beliefs. The lack of information in the price implies that there is no price effect. In other words, learning works only through the beliefs effect. When prior beliefs are bias upward, learning shifts the demand curve, which increases quantity.

Next, if the constant marginal cost depends on quality, then the price is informative and, in fact, it is fully-revealing about quality. That is, the uninformed buyers discard their prior beliefs and infer quality from observing the price. Note that the price is fully informative even when there is no information on the demand side. When the price is perfectly informative, learning has no effect on the competitive outcome since all buyers are informed in equilibrium. Indeed, beliefs and price effects work in opposite direction canceling each other out so that the intersection of demand and supply is unaffected by learning.

We then turn to the case of increasing supply curve. The dependence of the marginal cost on quantity implies that the price depends on the demand shock. Hence, the distribution of the price-signal is non-degenerate. Unless there are no informed buyers and the marginal cost is independent of quality, the price is partially informative about quality. Note that, unlike the case of constant supply, the presence of information on the supply side is not necessary for the price to be (partially) informative. Indeed, the information on the demand side is conveyed by the price system even when the supply side has no information. If there is information in the demand function as well as the supply function, it is the equilibrium interaction of supply and demand that determines the informational content of the equilibrium price. When the price is partially revealing (due to information coming from either demand or supply), the uninformed buyers use both prior information and

²Note that the price is uninformative even when there are informed buyers.

the information contained in the price to make decisions. Hence, both the beliefs and price effects are at work but do not, in general, cancel each other. In that case, the direction of the learning effect depends on both the bias of prior beliefs and on the demand shock.

Our work falls in the category of rational expectations models that study information flows in perfectly competitive markets (Kihlstrom and Mirman, 1975; Grossman, 1976, 1978; Grossman and Stiglitz, 1980).³ These studies on price transferring information are in endowment economies, i.e., with a vertical supply curve. We study the relationship between learning and market outcomes in a production economy. We make explicit the effect of the supply curve on the dissemination of information.⁴ Even though the perfectly competitive firm does not set the price, its decisions affect information flows through the equilibrium.⁵ For instance, in Grossman and Stiglitz (1980), the price is uninformative when there are no informed buyers. In our case, the price remains informative in the absence of informed buyers as long as the marginal cost depends on quality.⁶

Other studies consider the issue of information aggregation in large markets, i.e., with a continuum of firms. The firms have incomplete information and each receives a private signal about demand prior to setting quantity. This literature studies whether the market price aggregates the information efficiently. See Vives (1988) for instance. We consider price transferring information from informed to uninformed agents, i.e., there are learning agents who update beliefs upon observing the price. In their models, agents do not use the price to update beliefs, i.e., there is no learning.

The paper is organized as follows. Section 2 presents the model and equilibrium. Section 3 discusses how the uninformed buyers' learning activ-

³See also Grossman (1989) for a general treatment of this literature.

⁴Our demand setup is similar to Grossman and Stiglitz (1980), i.e., there are informed and uninformed buyers. As in Grossman and Stiglitz (1980), the market is noisy.

⁵Judd and Riordan (1994) and Mirman et al. (2014) study information flows in a noisy monopoly environment.

⁶Note that the relationship between the informativeness of the price and production has been studied in the context of insider trading. Indeed, Jain and Mirman (2000) explicitly links real and financial sectors through insider trading and studies the informativeness of the financial price when the insider is able to manipulate both real and financial prices.

ity influences demand. Section 4 derives and characterizes the competitive equilibrium. Section 5 provides final remarks.

2 Model

We embed learning in a model in which a perfectly competitive firm sells a good whose quality is unknown to some buyers. We first present the model. We then define the competitive equilibrium in which the price transmits information about quality to the uninformed buyers. In the next sections, we discuss how the uninformed buyers' learning activity influences demand. We finally derive and characterize the competitive equilibrium beginning with the case of constant marginal cost, followed by the case of increasing marginal cost.

2.1 Preliminaries

Consider a market for a good of quality $\theta \geq 0$ sold at price $P \geq 0$. The demand side is composed of informed and uninformed price-taking buyers. Informed buyers know θ and have demand $Q_d^I = \theta - P$. Uninformed buyers do not know θ , and infer information about quality from observing the price. Given prior beliefs ξ and the price-signal P , the uninformed buyers use Bayes' rule to form posterior beliefs $\hat{\xi}(\cdot|P)$.⁷ That is, given a distribution of the price-signal conditional on any quality θ denoted as $\phi_P(P|\theta)$, posterior beliefs are $\hat{\xi}(\theta|P) \propto \xi(\theta)\phi_P(P|\theta)$. Hence, uninformed buyers have demand $Q_d^U = \hat{\mu}_\theta(P) - P$ where $\hat{\mu}_\theta(P) \equiv \int_{x \geq 0} x \hat{\xi}(x|P) dx$ is the posterior mean of quality upon observing P .⁸ The posterior mean $\hat{\mu}_\theta(P)$ is also referred to as the *updating rule*, which combines prior information and information contained in the price-signal. Note that posterior beliefs influence demand only through the posterior mean. The updating rule plays a role in determining the influence of learning on the competitive equilibrium.

⁷That is, for any $X \subset \mathbb{R}_+$, the uninformed buyer's prior and posterior probabilities that $\theta \in X$ are $\int_{x \in X} \xi(x) dx$ and $\int_{x \in X} \hat{\xi}(x|P) dx$, respectively.

⁸Note that x is used as a dummy variable for quality.

Normalizing the mass of buyers to one and letting $\lambda \in [0, 1]$ be the fraction of informed buyers, the market demand is $Q_d = \lambda Q_d^I + (1 - \lambda)Q_d^U + \varepsilon$ or

$$Q_d = \lambda(\theta - P) + (1 - \lambda)(\hat{\mu}_\theta(P) - P) + \varepsilon \quad (1)$$

where ε is a demand shock (e.g., noisy buyers), which is *unknown* to the buyers. We assume that ε is a realization of the random variable $\tilde{\varepsilon}$ with p.d.f. $\phi_\varepsilon(\varepsilon)$.⁹

On the supply side, a representative, perfectly competitive firm sells Q units of the good at total cost $C(Q, \theta) \geq 0$ such that $C_1(Q, \theta), C_2(Q, \theta) \geq 0$. The firm has complete information about demand and cost. In particular, θ and the realization ε are known. The objective of the price-taking firm is to set quantity so as to maximize profit

$$\pi = PQ - C(Q, \theta). \quad (2)$$

Note that the number of firms is inconsequential for the analysis.

2.2 Equilibrium Definition

We define the competitive equilibrium, which consists of the quantity Q^* , the market-clearing price P^* , the p.d.f of the price-signal conditional on $\theta \geq 0$, $\phi_P^*(\cdot|\theta)$, and the uninformed buyers' posterior beliefs about quality upon observing $P \geq 0$, $\hat{\xi}^*(\cdot|P)$. In terms of notation, x is a dummy variable for quality and the asterisk sign on a variable denotes the equilibrium value.

Definition 2.1. *The tuple $\{Q^*, P^*, \phi_P^*(\cdot|\theta), \hat{\xi}^*(\cdot|P)\}$ is a competitive equilibrium if,*

1. *Given P^* ,*

$$Q^* = \arg \max_{Q \geq 0} \{P^*Q - C(Q, \theta)\}. \quad (3)$$

⁹A tilde sign distinguishes a random variable from a realization.

2. Given Q^* and $\hat{\xi}^*(\cdot|P)$, P^* clears the market, i.e.,

$$\lambda(\theta - P^*) + (1 - \lambda)(\hat{\mu}_\theta^*(P)|_{P=P^*} - P^*) + \varepsilon = Q^* \quad (4)$$

where $\hat{\mu}_\theta^*(P) \equiv \int_{x \geq 0} x \hat{\xi}^*(x|P) dx$ is the updating rule.

3. The distribution of the price-signal conditional on quality $\theta \geq 0$, $\phi_P^*(P|\theta)$, is derived from (4).

4. Given $\phi_P^*(\cdot|\theta)$, posterior beliefs upon observing P are, for $\theta \geq 0$,

$$\hat{\xi}^*(\theta|P) \propto \xi(\theta) \phi_P^*(P|\theta) \quad (5)$$

by Bayes' rule.

From Statement 1 of the definition of equilibrium, the firm's conjecture about the price is correct. Moreover, from Statements 3 and 4, the uninformed buyers' conjecture of the distribution of the price-signal (conditional on θ) is correct. This correct conjecture is then used to form posterior beliefs. Finally, the market-clearing price and posterior beliefs are dependent on each other. On the one hand, the market-clearing condition and the distribution of the price-signal are influenced by the updating rule (Statements 2 and 3). On the other hand, from Statement 4, in equilibrium, posterior beliefs depend on the correct conditional distribution of the price-signal.

Although the pair $\{Q^*, P^*\}$ is defined by the usual intersection of demand and supply, learning alters the demand function through the presence of an updating rule. Moreover, in a learning environment, the informativeness of the price is linked to the supply function. Indeed, in a competitive equilibrium, the price is equal to the marginal cost, i.e., $P^* = C_1(Q, \theta)|_{Q=Q^*}$. Hence, the functional form of the marginal cost influences the distribution of the price-signal conditional on quality. To capture the effect of learning on a competitive equilibrium, we study the effect that an arbitrary updating rule has on the demand schedule, and then the effect of the marginal cost function on the informativeness of the price.

3 Demand and Supply

In this section, we derive demand and supply functions and we study the effect of an arbitrary updating rule on demand. Although the effect of the updating rule on demand can be studied in general, we focus on the class of *linear updating rules*, i.e., the updating rule is a linear combination of the prior mean and the price-signal. Formally, $\hat{\mu}_\theta(P) = a\mu_\theta + bP$ where $\mu_\theta \equiv \int_{x \geq 0} x\xi(x)dx$ is the prior mean and $a, b \geq 0$.¹⁰ When $a = 0$, prior beliefs have no influence on posterior beliefs whereas $a > 0$ means that the posterior mean takes account of prior beliefs. The parameter b determines the extent to which posterior beliefs are influenced by the price-signal. When $b = 0$, the price-signal is uninformative and thus has no effect on posterior beliefs. If $b > 0$, then the price is considered informative and is incorporated into posterior beliefs.

Demand. Plugging $\hat{\mu}_\theta(P) = a\mu_\theta + bP$ into (1) defines the demand schedule under learning (\mathcal{L}) so that the price-quantity pair $\{P_d^\mathcal{L}, Q_d^\mathcal{L}\}$ satisfies

$$P_d^\mathcal{L} = \frac{\lambda\theta + (1 - \lambda)a\mu_\theta + \varepsilon - Q_d^\mathcal{L}}{1 - (1 - \lambda)b} \quad (6)$$

when $(1 - \lambda)b \neq 1$. From (6), both the y -intercept and the slope of the learning demand depend on the updating rule through the parameter b . If $(1 - \lambda)b = 1$, then the demand schedule is vertical. That is, for all $P_d^\mathcal{L} \geq 0$, the price-quantity pair $\{P_d^\mathcal{L}, Q_d^\mathcal{L}\}$ satisfies $Q_d^\mathcal{L} = \lambda\theta + (1 - \lambda)a\mu_\theta + \varepsilon$. Note that the presence of uninformed buyers is necessary for the updating rule to have an impact on demand. Indeed, from (6), when $\lambda = 1$, the parameters a and b are absent from the demand schedule. In fact, evaluating (6) at $\lambda = 1$ defines the demand schedule under full-information (\mathcal{FI}) in which all buyers are informed. The full-information demand schedule regroups the

¹⁰It is pertinent to use the linearity of the updating rule because, in equilibrium, the updating rule is a linear combination of the prior mean and the price-signal, i.e., $\hat{\mu}_\theta^*(P) = a^*\mu_\theta + b^*P$ where the asterisk signs on a and b distinguish equilibrium from arbitrary updating rules.




Figure 1: Demand Schedule under Arbitrary Linear Updating Rule

price-quantity pairs $\{P_d^{FI}, Q_d^{FI}\}$ such that

$$P_d^{FI} = \theta + \varepsilon - Q_d^{FI}. \quad (7)$$

In Figure 1 the learning demand is compared with the full-information demand for four different values of the parameters a and b . Although arbitrary values of a and b are considered, each case depicted in Figure 1 may occur in equilibrium. In each graph, the solid line is the demand schedule under learning (\mathcal{L}) whereas the dashed line represents the demand schedule under full-information (\mathcal{FI}). In Figure 1a, the condition $(1 - \lambda)b \neq 1$ implies that the learning demand is not vertical whereas $a > 0, b > 0$ ensures that both prior beliefs and new information influence demand. This case admits several possibilities in terms of slope and location with respect to the full-information demand. In Figure 1a, the learning demand is downward-

sloping and crosses the full-information demand from above because, from (6) and (7), the parameter values satisfy $a > 0, \theta/(\theta + \varepsilon) < b < 1/(1 - \lambda)$. However, for $(1 - \lambda)b \neq 1, a > 0, b > 0$, the learning demand is not necessarily downward-sloping. Indeed, if $b > 1/(1 - \lambda)$, then an increase in the price induces an increase in quantity through a higher posterior mean, which outweighs the standard negative effect of price on quantity demanded. Next, when $(1 - \lambda)b = 1, a > 0, b > 0$, the learning demand is a vertical line as depicted in Figure 1b. The last two figures consider two extreme cases of the updating rule. In Figure 1c, the uninformed buyers' updating rule ignores prior information, i.e., $a = 0, b > 0$. This case also admits several possibilities in terms of slope and position with respect to the full-information demand. In the case depicted in Figure 1c, the learning demand is below the full-information demand with a steeper slope because, from (6) and (7), the parameter values satisfy $a = 0, \theta/(\theta + \varepsilon) < b < 1/(1 - \lambda)$.¹¹ Finally, Figure 1d depicts the case in which the uninformed buyers do not engage in learning, i.e., purchases are made solely on the basis of prior beliefs. In this case, the learning and the full-information demand schedules are parallel. Indeed, from (6) evaluated at $a = 1, b = 0$ and (7), learning demand is to the right of the full-information demand if and only if prior beliefs are biased upward, i.e., $\mu_\theta > \theta$ as depicted in Figure 1d. Hence, if prior beliefs are unbiased, full-information and learning demand coincide.

The difference between learning and full-information demands is due to two distinct effects. The first is about asymmetry of information and prior beliefs. In a learning environment, the uninformed buyers use prior beliefs about quality. When prior beliefs are biased (i.e., $\mu_\theta \neq \theta$), learning influences the competitive equilibrium through *prior effect* of learning. The second component is about pure learning, i.e., the uninformed buyers update beliefs upon observing an informative price-signal. Thus, there is an informational role for the price through the updating rule. This is the *price effect* of

¹¹Note that when $(1 - \lambda)b = 1, a = 0, b > 0$, prior information is ignored and the learning demand is vertical.

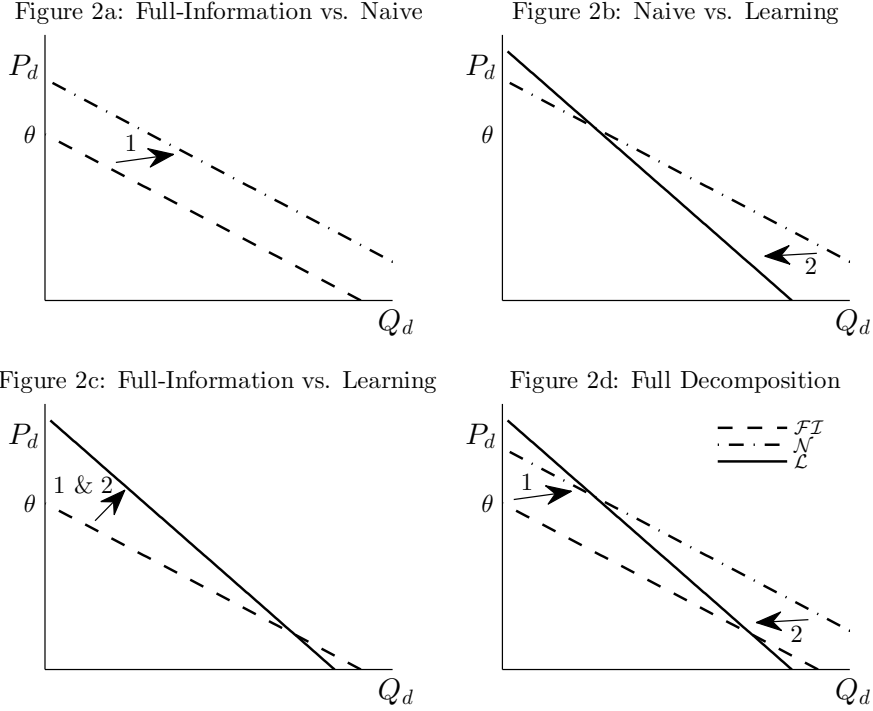


Figure 2: Demand Decomposition under Arbitrary Linear Updating Rule

learning.¹²

To decompose the effect of learning into these two components, we consider the intermediate case of a naive (\mathcal{N}) demand schedule in which the uninformed buyers do not learn, and thus, use only their prior beliefs. The naive demand schedule is defined by the price-quantity pair $\{P_d^{\mathcal{N}}, Q_d^{\mathcal{N}}\}$ such that

$$P_d^{\mathcal{N}} = \lambda\theta + (1 - \lambda)\mu_\theta + \varepsilon - Q_d^{\mathcal{N}} \quad (8)$$

where the prior mean μ_θ reflects the absence of updating. Hence, comparing (7) and (8) identifies the beliefs effect of learning whereas comparing (6) and (8) accounts for the price effect of learning.

Figure 2 provides a decomposition of the learning demand. Specifically,

¹²See Grossman (1989). In Koulovatianos, Mirman, and Santugini (2009), the effect of learning on optimal growth also depends on two components, beliefs and anticipation of learning.

Figure 2a highlights the prior effect of learning by depicting full-information and naive demand schedules when prior beliefs are upward-biased, i.e., $\mu_\theta > \theta$. Changing the composition of demand from informed buyers to a group of both informed and *naive* uninformed buyers causes a parallel shift of demand. From (7) and (8), the full-information demand schedule is to the right (left) of the naive demand schedule when $\theta > \mu_\theta$ ($\theta < \mu_\theta$). The prior effect is nil when $\theta = \mu_\theta$. Next, Figure 2b shows the price effect of learning by comparing the naive and the learning demand schedules. Given prior beliefs, the effect of the uninformed buyers' learning activity on demand is two-fold when $(1 - \lambda)b \neq 1$.¹³ Indeed, from (6) and (8), the price effect changes both the slope and the intercept of demand.¹⁴ Figure 2c aggregates the prior and price effects by depicting full-information and learning demands. Finally, Figure 2d summarizes the information from Figures 2a,b,c by providing all demand schedules.

Supply. Unlike the demand schedule, the supply schedule is unaltered by the uninformed buyers' learning activity. Given that the firm's cost function is $C(Q, \theta)$, the supply schedule is defined by the quantity-price pair $\{P_s, Q_s\}$ such that

$$P_s = C_1(Q_s, \theta). \quad (9)$$

The supply curve depends on θ when the marginal cost function depends on θ , i.e., when $C_{12}(Q_s, \theta) > 0$.

4 Equilibrium

Having discussed the structure of the demand and the supply functions in a learning environment, we next characterize the competitive equilibrium and study the informativeness of the price. The information contained in the market-clearing price emanates from both the demand side and the supply side. Specifically, the information from the demand comes from the presence

¹³For $(1 - \lambda)b = 1$, the updating component makes the learning demand vertical.

¹⁴From (6) and (8), the y -intercept between learning and naive demand curves is different when $\lambda\theta + (1 - \lambda)\mu_\theta + \varepsilon \neq \frac{\lambda\theta + (1 - \lambda)\alpha\mu_\theta + \varepsilon}{1 - (1 - \lambda)b}$.

of informed buyers who know θ whereas, from (9), the information from the supply is due to the dependence of the marginal cost function on θ .

The presence of informed buyers or a marginal cost function dependent on quality (or both) are thus necessary for the price to convey information about quality. Indeed, if demand is composed only of uninformed buyers (i.e., $\lambda = 0$) and the marginal cost is independent of quality (i.e., $C_{12}(Q, \theta) = 0$), then there is no source of information in the market, i.e., the market-clearing condition defined by (4) is independent of θ .

Remark 4.1. *In a competitive equilibrium, if $\lambda = 0$ and $C_{12}(Q, \theta) = 0$, then P^* is uninformative about θ .*

On the other hand, the presence of informed buyers is not sufficient for the price to be informative. Indeed, in equilibrium, $P^* = C_1(Q, \theta)|_{Q=Q^*(\theta, \varepsilon)}$. Hence, the distribution of the price-signal conditional on θ depends not only on the information contained in demand and supply, but also on the functional form of the marginal cost. Specifically, the informativeness of the price depends on the signs of the derivatives of the marginal cost. If $C_{11}(Q, \theta) = 0$, then the price is independent of output, and so the demand schedule has no bearing on the equilibrium price. This means that the information incorporated in the demand is not observable through price. If $C_{12}(Q, \theta) = 0$, then the supply function is independent of quality and no information from the supply side is conveyed through the price.

To understand the role played by the shape of the supply function on the informativeness of the price, we characterize the competitive equilibrium first in the case of constant supply (i.e., constant marginal cost) and then in the case of increasing supply (i.e., increasing marginal cost).

4.1 Constant Marginal Cost

Suppose that marginal cost is constant, i.e., $C_{11}(Q, \theta) = 0$. A horizontal supply curve implies that P^* is independent of output, and thus, independent of the demand shock ε . The distribution of the price-signal conditional on quality is, therefore, degenerate. It follows that, in the case of constant

marginal cost, the price is either uninformative or perfectly informative about quality. Specifically, when $C_{12}(Q, \theta) = 0$, the market-clearing condition yields an uninformative price. On the other hand, when $C_{12}(Q, \theta) > 0$, the market-clearing condition yields a price that fully reveals quality. We now examine each case separately.

Marginal Cost Independent of Quality. When the marginal cost is constant and independent of quality, i.e., $C_{11}(Q, \theta) = 0 = C_{12}(Q, \theta) = 0$, $C_1(Q, \theta) \equiv \kappa$ where $\kappa \geq 0$ is unrelated to θ . Then, $P^* = \kappa$ is uninformative about θ and the uninformed buyers revert to their prior information, i.e., $\hat{\mu}_\theta^*(P) = \mu_\theta$ where $\mu_\theta \equiv \int_{x \geq 0} x\xi(x)dx$.¹⁵

Proposition 4.2. *Suppose that the marginal cost is constant (i.e., $C_{11}(Q, \theta) = 0$) such that $C_1(Q, \theta) \equiv \kappa \geq 0$. Then, in the competitive equilibrium,*

1. *The firm produces*

$$Q^* = \lambda\theta + (1 - \lambda)\mu_\theta + \varepsilon - \kappa \quad (10)$$

at the price

$$P^* = \kappa. \quad (11)$$

2. *Posterior beliefs do not depend on P and are equal to prior beliefs, i.e., $\hat{\mu}_\theta^*(P) = \mu_\theta$ and $\hat{\sigma}_\theta^{*2} = \sigma_\theta^2$.*

Proof. Suppose that $C_{11}(Q, \theta) = 0$ such that $C_1(Q, \theta) \equiv \kappa \geq 0$. Then, $P^* = \kappa$. Since the price-signal is uninformative, posterior beliefs satisfy $\hat{\mu}_\theta^*(P) = \mu_\theta$ and $\hat{\sigma}_\theta^{*2} = \sigma_\theta^2$. Finally, plugging (11) into (4) yields (10), which satisfies the firm's maximization defined in (3). \square

In the case of marginal cost independent of quality, the absence of information on the supply side prevents the price from being informative even when there is information on the demand side.

¹⁵In other words, in equilibrium, the updating rule $\hat{\mu}_\theta^*(P) = a^*\mu_\theta + b^*P$ is such that $a^* = 1$ and $b^* = 0$.

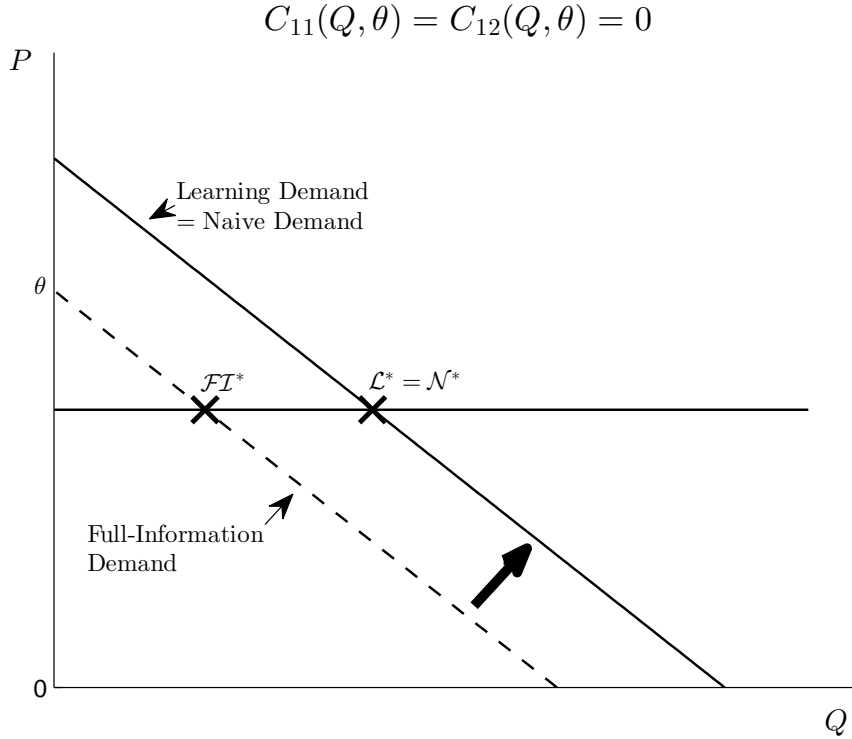


Figure 3: Constant Marginal Cost Independent of Quality

Remark 4.3. *Suppose that the marginal cost is constant (i.e., $C_{11}(Q, \theta) = 0$) such that $C_1(Q, \theta) \equiv \kappa \geq 0$. Then, in a competitive equilibrium, the supply side contains no information and, for any $\lambda \in (0, 1)$, the information present on the demand side is not revealed by the price.*

Figure 3 provides a depiction of the competitive equilibrium under a constant marginal cost which is independent of quality when prior beliefs are upward-biased, i.e., $\mu_\theta > \theta$. An uninformative price implies that the updating rule is equal to the prior. Hence, the learning and naive demand schedules coincide since the price effect of learning is absent. Moreover, the learning demand has the same slope as the full-information demand, i.e., learning and full-information demands are parallel. The distance between learning and full-information demands depends on the bias of the prior beliefs through the beliefs effect. In Figure 3, learning demand (and thus naive demand) is to the right of the full-information demand because $\mu_\theta > \theta$.

With a horizontal supply curve, the equilibrium quantity is defined by the position of the demand curve. The effect of learning on output is shown in Figure 3 for the case in which prior beliefs is upward-biased, i.e., $\mu_\theta > \theta$. Remark 4.4 states the effect of learning in the case of a constant marginal cost independent of quality.

Remark 4.4. *Suppose that the marginal cost is constant (i.e., $C_{11}(Q, \theta) = 0$) such that $C_1(Q, \theta) \equiv \kappa \geq 0$. Then, in a competitive equilibrium, from (10), $Q^*|_{\lambda=1} > Q^*|_{\lambda \in [0,1)} \iff \theta > \mu_\theta$.*

Constant Marginal Cost Dependent of Quality. Next, consider the case in which the constant marginal cost depends on quality, i.e., $C_{12}(Q, \theta) > 0$. Specifically, let $C_1(Q, \theta) \equiv K(\theta) \geq 0$ such that $K'(\theta) > 0$. Then, $P^* = K(\theta)$ is informative about θ , which implies that, in equilibrium, $\hat{\mu}_\theta^*(P) = K^{-1}(P)$.

Proposition 4.5 provides the competitive equilibrium under constant marginal cost dependent on quality. Comparing Propositions 4.2 and 4.5 highlights the influence of the supply function on information flows, through the sign of $C_{12}(Q, \theta)$. Specifically, when $C_{12}(Q, \theta) = 0$, $\hat{\mu}_\theta(P^*) = \mu_\theta$ whereas $C_{12}(Q, \theta) > 0$ implies that $\hat{\mu}_\theta(P^*) = \theta$.

Proposition 4.5. *Suppose that the marginal cost is constant (i.e., $C_{11}(Q, \theta) = 0$) such that $C_1(Q, \theta) \equiv K(\theta) \geq 0$, $K'(\theta) > 0$. Then, in the competitive equilibrium,*

1. *The firm produces*

$$Q^* = \theta + \varepsilon - K(\theta) \tag{12}$$

at the price

$$P^* = K(\theta). \tag{13}$$

2. *Posterior beliefs are such that $\hat{\mu}_\theta^*(P) = K^{-1}(P)$ and $\hat{\sigma}_\theta^{*2} = 0$.*

Proof. Suppose that $C_{11}(Q, \theta) = 0$ such that $C_1(Q, \theta) \equiv K(\theta) \geq 0$, $K'(\theta) > 0$. Then, $P^* = K(\theta)$. Since the price-signal is perfectly informative, posterior beliefs satisfy $\hat{\mu}_\theta^*(P) = K^{-1}(P)$ and $\hat{\sigma}_\theta^{*2} = 0$. Finally, plugging (13) into (4) yields (12), which satisfies the firm's maximization defined in (3). \square




Figure 4: Constant Marginal Cost Dependent on Quality

In the case of a constant marginal cost dependent on quality, the price is fully-revealing regardless of the information contained in the demand side.

Remark 4.6. *Suppose that the marginal cost is constant (i.e., $C_{11}(Q, \theta) = 0$) such that $C_1(Q, \theta) \equiv K(\theta) \geq 0, K'(\theta) > 0$. Then, in a competitive equilibrium, the price conveys perfect information about quality even when there is no information present on the demand side (i.e., $\lambda = 0$).*

Figure 4 provides a general depiction of the competitive equilibrium under a constant marginal cost dependent on quality when prior beliefs are upward-biased, i.e., $\mu_\theta > \theta$. A perfectly informative price implies that prior beliefs have no effect on the updating rule. Moreover, the uninformed buyers become informed upon observing the price so that the learning demand crosses the full-information demand at the intersection with supply. Since the price effect of learning is present, full-information and learning demands

have different slopes. Moreover, the prior effect and the price effect work in opposite directions in equal strength. Specifically, in Figure 4, $\mu_\theta > \theta$ implies that the naive demand is to the right of the full-information demand, but since there is full-revelation, the price effect works in the opposite direction so that learning demand shifts back to cross the supply curve at the full-information point of intersection. Comparing Remark 4.7 with Remark 4.4 highlights the influence the shape of the supply curve has on the effect of learning for equilibrium quantity.

Remark 4.7. *Suppose that the marginal cost is constant (i.e., $C_{11}(Q, \theta) = 0$) such that $C_1(Q, \theta) \equiv K(\theta) \geq 0, K'(\theta) > 0$. Then, in a competitive equilibrium, from (12), $Q^*|_{\lambda=1} = Q^*|_{\lambda \in [0,1]}$.*

4.2 Increasing Marginal Cost

In this section, we consider the case of increasing marginal cost. Unlike the constant marginal cost case, both sources of information (demand and supply) influence the learning process and the price conveys imperfect or perfect information about quality.

To study learning with an increasing marginal cost, we rely on the fact that the family of normal distributions with an unknown mean is a conjugate family for samples from a normal distribution.¹⁶ We also assume that the cost function is quadratic. Assumptions 4.8, 4.9 and 4.10 yield closed-form equilibrium values, which allows us to gain insight on information flows in a noisy environment.¹⁷

Assumption 4.8. *Cost is $C(Q, \theta) = \gamma_\theta \theta Q + \gamma_Q Q^2/2$ where $\gamma_\theta \in [0, 1)$ and $\gamma_Q \geq 0$.*

¹⁶For the case of constant marginal cost, there is no need to make distributional assumptions about prior beliefs and the demand shock.

¹⁷See Grossman and Stiglitz (1980), Kyle (1985), Judd and Riordan (1994), and Mirman et al. (2014) for the use of normal distributions to study the informational role of prices in endowment and single-agent (e.g., monopoly) problems. See also Vives (2011) for the use of normal distribution in a rational expectations environment with supply function competition. Although price and quantity can be negative, restrictions on parameter values ensures that the probability of a negative price or a negative quantity be arbitrarily close to zero.

Assumption 4.9. *The distribution of the demand shock is $\tilde{\varepsilon} \sim N(0, \sigma_\varepsilon^2)$ where $\sigma_\varepsilon^2 > 0$.*

Assumption 4.10. *Prior beliefs about quality are $\tilde{\theta} \sim N(\mu_\theta, \sigma_\theta^2)$ where $\mu_\theta > 0$ and $\sigma_\theta^2 > 0$.*

Proposition 4.11 characterizes the competitive equilibrium for the case of an increasing marginal cost. Note that in equilibrium the updating rule is a linear combination of the prior mean and the price-signal.

Proposition 4.11. *Suppose that Assumptions 4.8, 4.9, and 4.10 hold. Then, there exists a unique linear competitive equilibrium. In equilibrium,*

1. *The firm produces*

$$Q^* = \frac{(1 - \gamma_\theta)\theta + \varepsilon}{1 + \gamma_Q} + \frac{(1 - \lambda)\gamma_Q}{1 + \gamma_Q} \frac{(\mu_\theta - \theta)\gamma_Q\sigma_\varepsilon^2 + (\gamma_\theta + \lambda\gamma_Q)\varepsilon\sigma_\theta^2}{\gamma_Q^2\sigma_\varepsilon^2 + (\gamma_\theta + \lambda\gamma_Q)^2\sigma_\theta^2} \quad (14)$$

at the price

$$P^* = \frac{(\gamma_\theta + \gamma_Q)\theta + \gamma_Q\varepsilon}{1 + \gamma_Q} + \frac{(1 - \lambda)\gamma_Q^2}{1 + \gamma_Q} \frac{(\mu_\theta - \theta)\gamma_Q\sigma_\varepsilon^2 + (\gamma_\theta + \lambda\gamma_Q)\varepsilon\sigma_\theta^2}{\gamma_Q^2\sigma_\varepsilon^2 + (\gamma_\theta + \lambda\gamma_Q)^2\sigma_\theta^2}. \quad (15)$$

2. *The price-signal conditional on θ is distributed as $\tilde{P}^*|\theta \sim N(\hat{\mu}_P^*(\theta), \hat{\sigma}_P^{*2})$ where*

$$\hat{\mu}_P^*(\theta) = \frac{(\gamma_\theta + \gamma_Q)\theta}{1 + \gamma_Q} + \frac{(1 - \lambda)\gamma_Q^2}{1 + \gamma_Q} \frac{(\mu_\theta - \theta)\gamma_Q\sigma_\varepsilon^2}{\gamma_Q^2\sigma_\varepsilon^2 + (\gamma_\theta + \lambda\gamma_Q)^2\sigma_\theta^2}, \quad (16)$$

$$\hat{\sigma}_P^{*2} = \frac{\gamma_Q\sigma_\varepsilon^2}{1 + \gamma_Q} + \frac{(1 - \lambda)\gamma_Q^2}{1 + \gamma_Q} \frac{(\gamma_\theta + \lambda\gamma_Q)\sigma_\theta^2\sigma_\varepsilon^2}{\gamma_Q^2\sigma_\varepsilon^2 + (\gamma_\theta + \lambda\gamma_Q)^2\sigma_\theta^2}. \quad (17)$$

3. *Posterior beliefs $\hat{\xi}^*(\cdot|P)$ conditional on P are normally distributed with posterior mean $\hat{\mu}_\theta^*(P) = a^*\mu_\theta + b^*P$,*

$$a^* = \frac{\gamma_Q^2\sigma_\varepsilon^2}{\gamma_Q^2\sigma_\varepsilon^2 + (\gamma_\theta + \gamma_Q)(\gamma_\theta + \lambda\gamma_Q)\sigma_\theta^2}, \quad (18)$$

$$b^* = \frac{(1 + \gamma_Q)(\gamma_\theta + \lambda\gamma_Q)\sigma_\theta^2}{\gamma_Q^2\sigma_\varepsilon^2 + (\gamma_\theta + \gamma_Q)(\gamma_\theta + \lambda\gamma_Q)\sigma_\theta^2}. \quad (19)$$

and posterior variance

$$\hat{\sigma}_\theta^{*2} = \frac{\gamma_Q^2 \sigma_\varepsilon^2 \sigma_\theta^2}{\gamma_Q^2 \sigma_\varepsilon^2 + (\gamma_\theta + \lambda \gamma_Q)^2 \sigma_\theta^2}. \quad (20)$$

Proof. Suppose that $\gamma_Q > 0$. Using Assumption 4.8, given P^* , the first-order condition corresponding to (3) yields

$$Q^* = \frac{P^* - \gamma_\theta \theta}{\gamma_Q}. \quad (21)$$

Plugging (21) and the posterior mean $\hat{\mu}_\theta^*(P) = \int_{\mathbb{R}} x \hat{\xi}^*(x|P) dx = a^* + b^* P$ into (4) and solving for the market-clearing price yields

$$P^* = \frac{(\gamma_\theta + \lambda \gamma_Q) \theta + \gamma_Q (1 - \lambda) a^* + \gamma_Q \varepsilon}{1 + \gamma_Q (1 - b^* (1 - \lambda))} \quad (22)$$

where it remains to solve for a^* and b^* . From (22) and Assumption 4.9, the price-signal conditional on θ is normally distributed with mean and variance

$$\hat{\mu}_P^*(\theta) = \frac{(\gamma_\theta + \lambda \gamma_Q) \theta + \gamma_Q (1 - \lambda) a^*}{1 + \gamma_Q (1 - b^* (1 - \lambda))} \quad (23)$$

$$\hat{\sigma}_P^{*2} = \frac{\gamma_Q^2 \sigma_\varepsilon^2}{(1 + \gamma_Q (1 - b^* (1 - \lambda)))^2}. \quad (24)$$

Given that prior beliefs are normally distributed (Assumption 4.10) and that the price-signal is normally distributed,¹⁸ posterior beliefs conditional on P are normally distributed with mean

$$\begin{aligned} \hat{\mu}_\theta^*(P) &= \frac{\sigma_\varepsilon^2 \gamma_Q^2 \mu_\theta - \sigma_\theta^2 (\gamma_\theta + \lambda \gamma_Q) (1 - \lambda) a^*}{\gamma_Q^2 \sigma_\varepsilon^2 + (\gamma_\theta + \lambda \gamma_Q)^2 \sigma_\theta^2} \\ &\quad + \frac{\sigma_\theta^2 (1 + \gamma_Q (1 - b^* (1 - \lambda))) (\gamma_\theta + \lambda \gamma_Q)}{\gamma_Q^2 \sigma_\varepsilon^2 + (\gamma_\theta + \lambda \gamma_Q)^2 \sigma_\theta^2} P \end{aligned} \quad (25)$$

¹⁸Using (22), let $z = \frac{(1 - \gamma_Q (1 - b^* (1 - \lambda))) P - \gamma_Q (1 - \lambda) a^*}{\gamma_\theta + \lambda \gamma_Q} = \theta + \frac{\gamma_Q \varepsilon}{\gamma_\theta + \lambda \gamma_Q}$ such that $\tilde{z}|\theta \sim N\left(\theta, \frac{\gamma_Q^2 \sigma_\varepsilon^2}{(\gamma_\theta + \lambda \gamma_Q)^2}\right)$. Hence, $\tilde{\theta}|z \sim N\left(\frac{\sigma_\theta^2 z + \frac{\gamma_Q^2 \sigma_\varepsilon^2}{(\gamma_\theta + \lambda \gamma_Q)^2} \mu_\theta}{\sigma_\theta^2 + \frac{\gamma_Q^2 \sigma_\varepsilon^2}{(\gamma_\theta + \lambda \gamma_Q)^2}}, \frac{\gamma_Q^2 \sigma_\theta^2 \sigma_\varepsilon^2}{\gamma_Q^2 \sigma_\varepsilon^2 + \sigma_\theta^2 (\gamma_\theta + \lambda \gamma_Q)^2}\right)$.

and variance given by (20). It remains to determine the values of a^* and b^* . Equating (25) to the conjecture $\hat{\mu}_\theta^*(P) = a^* + b^*P$ and solving for a^* and b^* yields (18) and (19). Finally, plugging (18) and (19) into (21), (22), (23), and (24), and rearranging yields (14), (15), (16), and (17), respectively. \square

Having characterized the competitive equilibrium when the supply curve is upward-sloping, we study the informativeness of the price. We begin by showing that in the case of an increasing marginal cost, the market-clearing price corresponding to an arbitrary linear updating rule depends on both demand and supply parameters. Unlike the case of a constant marginal cost, an increasing marginal cost makes the market-clearing price dependent on the demand side and in particular, the fraction of informed buyers, the uninformed buyers's updating rule, and the demand shock.

Formally, given P and using (3), $Q^* = \frac{P - \gamma_\theta \theta}{\gamma_Q}$. Plugging $\hat{\mu}_\theta(P) = a\mu_\theta + bP$ and $Q = \frac{P - \gamma_\theta \theta}{\gamma_Q}$ into (4) yields

$$\lambda(\theta - P) + (1 - \lambda)(a\mu_\theta + bP - P) + \varepsilon = \frac{P - \gamma_\theta \theta}{\gamma_Q} \quad (26)$$

so that the market-clearing price is

$$P = \frac{(\gamma_\theta + \lambda\gamma_Q)\theta + (1 - \lambda)\gamma_Q a\mu_\theta + \gamma_Q \varepsilon}{1 + \gamma_Q(1 - (1 - \lambda)b)}. \quad (27)$$

Expression (27) implies that, unlike the case of constant marginal cost, the information from the demand side is reflected in the market clearing price through the parameter λ . In addition, an increasing marginal cost links the price to the demand shock, which implies that the distribution of the price-signal is non-degenerate.¹⁹ Hence, the price is either uninformative or imperfectly informative, but never perfectly informative about quality.

Propositions 4.12 and 4.13 provide conditions for which the price reveals

¹⁹When demand is noiseless, the distribution of the price-signal is degenerate (as in the case of constant marginal constant is constant) and fully-revealing about quality. Evaluating (15) at $\sigma_\varepsilon^2 = \varepsilon = 0$ yields $P^*|_{\sigma_\varepsilon^2=\varepsilon=0} = (\gamma_\theta + \gamma_Q)\theta / (1 + \gamma_Q)$, which is independent of ε and strictly increasing in θ . Hence, plugging (15), (18), (19) evaluated at $\sigma_\varepsilon^2 = 0, \varepsilon = 0$ into $\hat{\mu}_\theta^*(P) = a^* + b^*P$ yields $\hat{\mu}_\theta^*(P^*) = \theta$.

either no information or some information about quality. Specifically, unless there are some informed buyers or the cost depends on quality, the price cannot provide any information to the uninformed buyers. Proposition 4.12 is consistent with Remark 4.1.

Proposition 4.12. *Suppose that $\gamma_Q > 0$. If $\lambda = 0$ and $\gamma_\theta = 0$, then, in the competitive equilibrium, P^* is independent of θ and $\hat{\mu}_\theta^*(P^*) = \mu_\theta$ and $\hat{\sigma}_\theta^{*2} = \sigma_\theta^2$.*

Proof. Evaluating (15) at $\lambda = 0$ and $\gamma_\theta = 0$ implies that the price is independent of θ , which yields non-revelation. \square

Proposition 4.13 states that the price is imperfectly informative as long as λ and γ_θ are not simultaneously zero. Although quality is not fully revealed, the posterior variance is less than the prior variance. Here, \mathbb{E} is the expectation operator with respect to the p.d.f. $\phi_P^*(\cdot|\theta)$.

Proposition 4.13. *Suppose that $\gamma_Q > 0$. If either $\lambda \in (0, 1)$ or $\gamma_\theta > 0$, then, in the competitive equilibrium,*

1. $\mathbb{E}\hat{\mu}_\theta^*(\tilde{P}^*) = \alpha\mu_\theta + (1 - \alpha)\theta$ where

$$\alpha = \frac{\gamma_Q^2 \sigma_\varepsilon^2}{\gamma_Q^2 \sigma_\varepsilon^2 + (\gamma_\theta + \lambda\gamma_Q)^2 \sigma_\theta^2}. \quad (28)$$

2. $\hat{\sigma}_\theta^{*2} < \sigma_\theta^2$.

Proof. Plugging (15), (18), and (19) into $\hat{\mu}_\theta^*(P) = a^* + b^*P^*$ and taking the expectation with respect to \tilde{P}^* yields

$$\begin{aligned} \mathbb{E}\hat{\mu}_\theta^*(\tilde{P}^*) &= \frac{\gamma_Q^2 \sigma_\varepsilon^2 \mu_\theta}{\gamma_Q^2 \sigma_\varepsilon^2 + (\gamma_\theta + \gamma_Q)(\gamma_\theta + \lambda\gamma_Q)\sigma_\theta^2} \\ &\quad + \frac{(1 + \gamma_Q)(\gamma_\theta + \lambda\gamma_Q)\sigma_\theta^2}{\gamma_Q^2 \sigma_\varepsilon^2 + (\gamma_\theta + \gamma_Q)(\gamma_\theta + \lambda\gamma_Q)\sigma_\theta^2} \\ &\quad \cdot \left(\frac{(\gamma_\theta + \gamma_Q)\theta}{1 + \gamma_Q} + \frac{(1 - \lambda)\gamma_Q^2}{1 + \gamma_Q} \frac{(\mu_\theta - \theta)\gamma_Q\sigma_\varepsilon^2}{\gamma_Q^2 \sigma_\varepsilon^2 + (\gamma_\theta + \lambda\gamma_Q)^2 \sigma_\theta^2} \right). \end{aligned} \quad (29)$$

Rearranging (29) yields $\mathbb{E}\hat{\mu}_\theta^*(\tilde{P}^*) = \alpha\mu_\theta + (1-\alpha)\theta$ where α is defined by (28). Next, from (20), the posterior variance is smaller than the prior variance, i.e.,

$$\hat{\sigma}_\theta^2 = \frac{\sigma_\varepsilon^2 \sigma_\theta^2 \gamma_Q^2}{\sigma_\varepsilon^2 \gamma_Q^2 + \sigma_\theta^2 (\gamma_\theta + \lambda \gamma_Q)^2} = \frac{1}{1 + \frac{\sigma_\theta^2 (\gamma_\theta + \lambda \gamma_Q)^2}{\sigma_\varepsilon^2 \gamma_Q^2}} \sigma_\theta^2 < \sigma_\theta^2. \quad (30)$$

□

Note that when the price is informative about quality, the information contained in the price comes from both demand and supply. In particular, if cost is unrelated to quality (i.e., $\gamma_\theta = 0$), learning still occurs because information about quality is conveyed through demand when there are informed buyers (i.e., $\lambda \in (0, 1]$). The parameter λ controls the amount of information released by the demand side. An increase in the number of informed buyers reveals more information. That is, the expected posterior mean is closer to the value of quality and the posterior variance is decreased as a result of an increase in the fraction of informed buyers.

Proposition 4.14. *Suppose that $\gamma_Q > 0$. Then, in a competitive equilibrium,*

$$\frac{\partial \mathbb{E}\hat{\mu}_\theta^*(\tilde{P}^*)}{\partial \lambda} > 0 \iff \theta > \mu_\theta \quad (31)$$

and

$$\frac{\partial \hat{\sigma}_\theta^{*2}}{\partial \lambda} < 0. \quad (32)$$

Proof. From Proposition 4.13,

$$\frac{\partial \mathbb{E}\hat{\mu}_\theta^*(\tilde{P}^*)}{\partial \alpha} = (\mu_\theta - \theta) < 0 \iff \theta > \mu_\theta, \quad (33)$$

and from (28),

$$\frac{\partial \alpha}{\partial \lambda} < 0. \quad (34)$$

The signs of (33) and (34) imply (31). Taking the derivative of (20) with respect to λ yields (32). □

On the supply side, the marginal effect of quality on cost, γ_θ , governs the

amount of information emanating from the supply side. Indeed, an increase in cost through γ_θ provides more information to the uninformed buyers.

Proposition 4.15. *Suppose that $\gamma_Q > 0$. Then, in a competitive equilibrium,*

$$\frac{\partial \mathbb{E} \hat{\mu}_\theta^*(\tilde{P}^*)}{\partial \gamma_\theta} > 0 \iff \theta > \mu_\theta \quad (35)$$

and

$$\frac{\partial \hat{\sigma}_\theta^{*2}}{\partial \gamma_\theta} < 0. \quad (36)$$

Proof. From Proposition 4.13,

$$\frac{\partial \mathbb{E} \hat{\mu}_\theta^*(\tilde{P}^*)}{\partial \alpha} = (\mu_\theta - \theta) < 0 \iff \theta > \mu_\theta, \quad (37)$$

and from (28),

$$\frac{\partial \alpha}{\partial \gamma_\theta} < 0. \quad (38)$$

The signs of (37) and (38) imply (35). Taking the derivative of (20) with respect to γ_θ yields (36). \square

Having studied the informativeness of the price with an increasing marginal cost, we next turn to the effect of learning on quantity and price. As noted, the effect of learning on output and price is reflected in the difference between the full-information case (i.e., $\lambda = 1$) and the learning case (i.e., $\lambda \in (0, 1)$). From (14) and (15), the firm produces

$$Q^* = \frac{(1 - \gamma_\theta)\theta + \varepsilon}{1 + \gamma_Q} + \underbrace{\frac{(1 - \lambda)\gamma_Q}{1 + \gamma_Q} \frac{(\mu_\theta - \theta)\gamma_Q\sigma_\varepsilon^2 + (\gamma_\theta + \lambda\gamma_Q)\varepsilon\sigma_\theta^2}{\gamma_Q^2\sigma_\varepsilon^2 + (\gamma_\theta + \lambda\gamma_Q)^2\sigma_\theta^2}}_{\text{Learning Effect}} \quad (39)$$

at the market-clearing price

$$P^* = \frac{(\gamma_\theta + \gamma_Q)\theta + \gamma_Q\varepsilon}{1 + \gamma_Q} + \underbrace{\frac{(1 - \lambda)\gamma_Q^2}{1 + \gamma_Q} \frac{(\mu_\theta - \theta)\gamma_Q\sigma_\varepsilon^2 + (\gamma_\theta + \lambda\gamma_Q)\varepsilon\sigma_\theta^2}{\gamma_Q^2\sigma_\varepsilon^2 + (\gamma_\theta + \lambda\gamma_Q)^2\sigma_\theta^2}}_{\text{Learning Effect}}. \quad (40)$$

To highlight the effect of learning, expressions (39) and (40) are decomposed into two terms. The first term in each expression is the equilibrium value when all buyers are informed. If $\lambda = 1$, then all buyers are informed and there is no learning activity, which reduces the second terms in (39) and (40) to zero.²⁰ The second term represents the effect of learning due to the presence of uninformed buyers in the market.

Proposition 4.16 states that the sign of the *learning effect* term depends on the parameter values, i.e., fraction of informed buyers, demand shock, bias in prior beliefs, and cost parameters.

Proposition 4.16. *Suppose that $\sigma_\varepsilon^2 > 0$ and $\gamma_Q > 0$. Then, from (39) and (40),*

$$\begin{aligned} Q^*|_{\lambda=1} > Q^*|_{\lambda \in [0,1)} &\iff \theta - \mu_\theta > \frac{(\gamma_\theta + \lambda\gamma_Q)\varepsilon\sigma_\theta^2}{\gamma_Q\sigma_\varepsilon^2}. \\ P^*|_{\lambda=1} > P^*|_{\lambda \in [0,1)} & \end{aligned} \quad (41)$$

Proposition 4.16 is in contrast to Remark 4.4. That is, in the case of a positive constant marginal cost, the prior bias governs the direction of the effect of learning on output. With an increasing supply curve, the prior bias has only a partial influence. The reason is that with an increasing marginal cost, both the beliefs effect and the price effect are present and do not cancel each other. If prior beliefs are unbiased and the demand shock is positive (negative), then learning increases (decreases) both quantity and price. However, if prior beliefs are downward-biased (i.e., $\theta > \mu_\theta$) and demand shock is negative, learning unambiguously increases both quantity and price. Figure 5 depicts the competitive equilibrium with an upward-sloping supply curve when $\theta - \mu_\theta > \frac{(\gamma_\theta + \lambda\gamma_Q)\varepsilon\sigma_\theta^2}{\gamma_Q\sigma_\varepsilon^2}$.

The learning-effect term identified in (39) and (40) includes both the be-

²⁰As noted, the presence of uninformed buyers (i.e., $\lambda \in [0, 1)$) is not a sufficient condition for learning to have an effect on the competitive equilibrium. Indeed, for all $\lambda \in [0, 1)$, the *learning effect* terms in (39) and (40) are equal to zero when demand is noiseless ($\sigma_\varepsilon^2 = 0$ so that $\tilde{\varepsilon}$ is a degenerate random variable at $\varepsilon = 0$) or marginal cost is constant ($\gamma_Q = 0$). The reason is that a noiseless demand or a constant marginal cost implies that the price-signal is strictly increasing in quality and does not depend on the demand shock. Hence, quality is perfectly inferred from observing the price. It follows that the competitive equilibrium is unaffected by learning when θ is fully revealed.

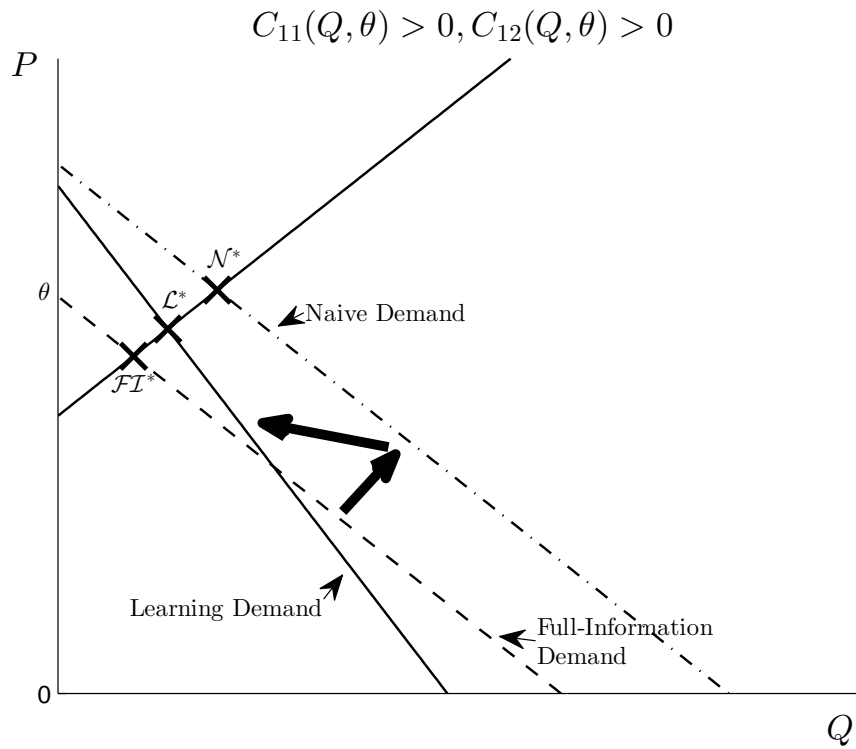


Figure 5: Increasing Marginal Cost Dependent on Quality

beliefs effect and the price effect. The price effect reflects the pure informational role of the price on the equilibrium. In Figure 5, the beliefs effect is the difference between the equilibrium points $\mathcal{F}\mathcal{I}^*$ and \mathcal{N}^* whereas the price effect is the difference between the equilibrium points \mathcal{N}^* and \mathcal{L}^* . That is, for the

price, the learning-effect term in (40) is decomposed as

$$\begin{aligned}
& \underbrace{\frac{(1-\lambda)\gamma_Q^2(\mu_\theta - \theta)\gamma_Q\sigma_\varepsilon^2 + (\gamma_\theta + \lambda\gamma_Q)\varepsilon\sigma_\theta^2}{1 + \gamma_Q \quad \gamma_Q^2\sigma_\varepsilon^2 + (\gamma_\theta + \lambda\gamma_Q)^2\sigma_\theta^2}}_{\text{Learning Effect}} \\
&= \underbrace{\frac{(1-\lambda)\gamma_Q(\mu_\theta - \theta)}{1 + \gamma_Q}}_{\text{Beliefs Effect of Learning}} \\
& \quad + \underbrace{\frac{(1-\lambda)(\gamma_\theta + \lambda\gamma_Q)\gamma_Q\sigma_\theta^2((\gamma_\theta + \lambda\gamma_Q)(\theta - \mu_\theta) + \gamma_Q\varepsilon)}{(1 + \gamma_Q)(\gamma_Q^2\sigma_\varepsilon^2 + (\gamma_\theta + \lambda\gamma_Q)^2\sigma_\theta^2)}}_{\text{Price Effect of Learning}}. \tag{42}
\end{aligned}$$

From (42), the sign of the beliefs effect of learning depends only on the bias of the prior whereas the sign of the price effect depends on the bias of the prior as well as the demand shock. Hence, the noise in demand influences the learning effect only when the price is informative.

5 Final Remarks

In our model, the firm is assumed to be risk-neutral. Moreover, our demand specification does not make explicit the role played by the buyers' risk aversion on the competitive equilibrium. Little is known about the role of risk aversion in an environment in which agents face uncertainty, but engage in learning. For instance, the behavior of risk-averse perfectly competitive firms maximizing the expected utility of profit has only been studied in models of uncertainty without learning (Baron, 1970; Sandmo, 1971; Leland, 1972). Combining learning and risk aversion would further our understanding of the role risk aversion has not only on decisions and market outcomes, but also on information flows. Future work should consider how buyers' and sellers' risk preferences influence the conveyance of information through the equilibrium.

References

- D.P. Baron. Price Uncertainty, Utility, and Industry Equilibrium in Pure Competition. *Int. Econ. Rev.*, 11(3):463–480, 1970.
- S.J. Grossman. On the Efficiency of Competitive Stock Markets where Traders Have Different Information. *J. Finance*, 31(2):573–585, 1976.
- S.J. Grossman. Further Results on the Informational Efficiency of Competitive Stock Markets. *J. Econ. Theory*, 18(1):81–101, 1978.
- S.J. Grossman. *The Informational Role of Prices*. MIT Press, 1989.
- S.J. Grossman and J.E. Stiglitz. On the Impossibility of Informationally Efficient Markets. *Amer. Econ. Rev.*, 70(3):393–408, 1980.
- N. Jain and L.J. Mirman. Real and Financial Effects of Insider Trading with Correlated Signals. *Econ. Theory*, 16(2):333–353, 2000.
- K.L. Judd and M.H. Riordan. Price and Quality in a New Product Monopoly. *Rev. Econ. Stud.*, 61(4):773–789, 1994.
- R.E. Kihlstrom and L.J. Mirman. Information and Market Equilibrium. *Bell. J. Econ.*, 6(1):357–376, 1975.
- C. Koulovatianos, L.J. Mirman, and M. Santugini. Optimal Growth and Uncertainty: Learning. *J. Econ. Theory*, 144(1):280–295, 2009.
- A.S. Kyle. Continuous Auctions and Insider Trading. *Econometrica*, 53(6):1315–1335, 1985.
- H.E. Leland. Theory of the Firm Facing Uncertain Demand. *Amer. Econ. Rev.*, 62(3):278–291, 1972.
- L.J. Mirman, E.M. Salgueiro, and M. Santugini. Noisy Signaling in Monopoly. *Int. Rev. Econ. Finance*, 29(C):504–511, 2014.
- A. Sandmo. On the Theory of the Competitive Firm under Price Uncertainty. *Amer. Econ. Rev.*, 61(1):65–73, 1971.

- X. Vives. Aggregation of Information in Large Cournot Markets. *Econometrica*, 56(4):851–876, 1988.
- X. Vives. Strategic Supply Function Competition With Private Information. *Econometrica*, 79(6):1919–1966, 2011.