



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

Cahier de recherche/Working Paper **15-02**

Noisy Learning in a Competitive Market with Risk Aversion

Leonard J. Mirman

Egas M. Salgueiro

Marc Santugini

Février/February 2015

Mirman: Department of Economics, University of Virginia

lm8h@virginia.edu

Salgueiro: DEGEI/GOVCOPP, Universidade de Aveiro

egas@ua.pt

Santugini: Department of Applied Economics and CIRPÉE, HEC Montréal

marc.santugini@hec.ca

Abstract:

We address the issue of risk aversion in a competitive equilibrium when some buyers engage in learning and information is conveyed through the price system. Specifically, since the learning process yields uncertainty, we study the effect of risk aversion on the equilibrium outcomes of the model, including the amount of information released by the market. We show that risk aversion has an effect on the market outcomes but not on the flow of information. In particular, an increase in risk aversion lowers the competitive price and quantity. However, an increase in risk aversion does not change the amount of information embedded in the equilibrium price.

Keywords: Learning, Risk aversion, Uncertainty

JEL Classification: D21, D42, D82, D83, D84, L12, L15

1 Introduction

One of the central questions in the field of economics of uncertainty is the effect of risk aversion on behavior. The question has been long studied. In general, agents interact in markets and thus risk aversion influences not only behavior directly, but also indirectly through the equilibrium or the market outcomes such as prices and quantities. This is especially relevant in markets with asymmetric information in which the agents face uncertainty due to incomplete information, but learn from observing prices.

It is well known that market prices are instrumental in disseminating information to market participants (Grossman, 1989). The informational role of prices is generally studied in a noisy environment in which there is asymmetric information about a characteristic of the good being traded. In much of the literature on the dissemination of information via prices, agents are assumed to be risk-neutral. In fact, little is known about the role risk aversion plays in the conveyance of information. It is our purpose to study the effect of risk aversion on the competitive equilibrium when agents in the markets are uninformed, but extract information from the price.

Specifically, in a noisy environment in which the price conveys information imperfectly, the learning process is itself a source of uncertainty, which must be taken account of by the risk-averse agents. In fact, in the interaction between uncertainty, learning, optimal behavior, and market equilibrium, risk aversion plays a central role not only in the decision-making process but also in the learning process. Since the learning process yields uncertainty, a natural question is the effect of risk aversion on the equilibrium outcomes of the model, including the amount of information released by the market.

In this paper, we consider a competitive market in which demand is composed of both informed and uninformed buyers. The informed buyers know the quality of the good. The uninformed buyers use Bayesian methods to infer information about quality 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 noise, a random noise component, which is known to the firm, but

unknown to buyers. The presence of noise in the market prevents complete learning by the uninformed consumers, which fundamentally affects the market equilibrium.¹

We provide two sets of results on the effect of risk aversion on the competitive equilibrium with learning. First, we study the effect of risk aversion on the competitive price and quantity. We show that an increase in risk aversion induces a downward movement of the demand curve, which decreases price and output. The effect of risk aversion on demand is both direct and indirect. The direct effect is to reduce the demand of the uninformed buyers, which shifts demand inward. The indirect effect of risk aversion is through the updating rule of the uninformed buyers. Indeed, the updating rule (as a function of price) depends on risk aversion.

Our second result concerns the influence of risk aversion on the amount of information transferred through the competitive equilibrium. We show that risk aversion has no effect on the conveyance of information through prices and thus on the updated beliefs (i.e., the posterior mean). To understand why risk aversion has no effect on what the uninformed buyers learn, note first that an increase in risk aversion does reduce the uninformed buyers demand, essentially reducing the influence of the uninformed buyers on the equilibrium and thus making the presence of informed buyers more important. However, this increase in the relative importance of informed buyers does not make the equilibrium price more informative. The reason is that the amount of information contained in the price can be summarized by a sufficient statistics that is independent of risk aversion. Hence, changes in risk aversion do not alter the conveyance of information through the price and thus the value of the posterior mean (which is relevant to the uninformed buyers demand) remains the same, i.e., is unaffected by changes in risk aversion. Note that since the level of the equilibrium price changes (although

¹Our 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; Diamond and Verrecchia, 1981; Hellwig, 1980). The majority of these papers considers the trading of a financial asset and, thus, there is no *firm* supplying the good. We take a different point of view by addressing the issue of learning in a market for a good or service in which the behavior of a price-taking firm is made explicit.

the information contained in it does not), the updating rule (as a function of price) adapts in such a way as to yield the same value of the posterior mean. It does not however mean that the updating rule puts more weight or less weight on the information contained in the price, i.e., this change in the updating rule is not due to changes in the informativeness of the signal since the price carries the same amount of information. Our result on information flows being independent of risk aversion complements the one stated in Grossman and Stiglitz (1980) in an endowment economy about the effect of the variance of the demand shock on the informativeness of the price. Indeed, from Grossman and Stiglitz (1980), “...an increase in noise reduces the informativeness of the price system: but it leads more individuals to become informed; the remarkable result obtained above establishes that *the two effects exactly offset each other* so that the equilibrium informativeness of the price system is unchanged” (their italic).

The paper is organized as follows. Section 2 characterizes the learning equilibrium with a competitive market, whereas Section 3 considers the effect of uncertainty and risk aversion on equilibrium, particularly on learning. We provide concluding remarks in Section 4.

2 Equilibrium with risk averse learning buyer in a competitive market

In this section we study the learning equilibrium of a competitive market with risk averse buyers.

On the supply side there is a competitive firm, which maximizes profit given price P . The firm has a cost function dependent on quantity Q , quality, $\theta \geq 0$, and a random noise $\tilde{\varepsilon}$:

$$C(Q) = C_F + (\gamma_\theta \theta + \tilde{\varepsilon})Q + \gamma_q Q^2, \gamma_\theta, \gamma_q > 0^2 \quad (1)$$

²Throughout this work we assume that γ_q is strictly positive. The case where $\gamma_q = 0$, i.e., the supply function is horizontal, was extensively studied in (Mirman et al., 2014).

where $\tilde{\varepsilon}$ is a normally-distributed noise³, with $\tilde{\varepsilon} \sim n(0, \sigma_{\tilde{\varepsilon}}^2)$ ⁴. It is assumed that the realization of $\tilde{\varepsilon}$ is known to the firm but not to the buyers.

From the cost function we derive the supply function. As the firm is a price taker, it will choose its output to maximize profit setting marginal cost equal to price, i.e.,

$$P = \gamma_{\theta}\theta + \gamma_q Q^S + \tilde{\varepsilon} \quad (2)$$

or

$$Q^S = \frac{P - \gamma_{\theta}\theta - \tilde{\varepsilon}}{\gamma_q} \quad (3)$$

On the demand side, there is a fraction $\lambda \in (0, 1)$ of consumers who are informed about the value of the θ . The remaining fraction, $(1 - \lambda)$, of consumers is not informed and must learn about the quality. These are, respectively, the informed consumers I and learning consumers L. The learning consumers extract information from observing the price. Specifically, a learning consumer has prior beliefs⁵ regarding quality that are given by

$$\tilde{\theta}_a \sim n(\mu_{\theta}, \sigma_{\theta}^2), \mu_{\theta} > 0 \quad (4)$$

represented by the p.d.f. $\hat{\xi}(\theta_a)$. Given P these prior beliefs are updated to the posterior beliefs $\tilde{\theta}_p|P \sim n(\hat{\mu}_{\theta}, \hat{\sigma}_{\theta}^2)$, with $\hat{\mu}_{\theta} = \int_{\mathbb{R}} x \hat{\xi}(x|P) dx > 0$. The corresponding posterior p.d.f. $\hat{\xi}(\theta_p|P)$ is computed according to Bayes' rule. We define $\chi(P)$ to be the updating rule for the posterior mean, i.e., $\hat{\mu}_{\theta} = \chi(P)$

³The tilde sign differentiates a random variable from its realization.

⁴In order to study noisy signaling in a competitive market, 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. A normal distribution combined with linear demand yields closed-form equilibrium values and makes the analysis tractable by focusing on the mean and variance of price and posterior beliefs (see Grossman and Stiglitz (1980), Kyle (1985), Judd and Riordan (1994) for the use of normal distributions to study the informational role of prices in single-agent problems). Hence, the normality assumption allows us to gain insight into information flows in a noisy environment. Although equilibrium price and posterior mean quality can admit negative values, restrictions on parameter values ensures that the probability of a negative price or a negative posterior mean be arbitrarily close to zero.

⁵We use $\tilde{\theta}_a$ to denote the prior (*ex ante*) beliefs of the learning buyers regarding quality, and $\tilde{\theta}_p$ for their posterior (*ex post*) beliefs. This avoids confusion between prior beliefs, posterior beliefs and the true quality θ .

Both types of consumers maximize the expected value of their CARA utility functions, given by $U_i(Q) = -\exp\left\{-a_i\left[\theta Q - \frac{Q^2}{2} + y\right]\right\}$, subject to their budget constraint ($I=PQ+y$). The parameter $a_i > 0$ is the Arrow-Pratt coefficient of absolute risk-aversion; a bigger a_i corresponds to a more risk averse agent. Here, y is an alternative good with no risk and price standardized to 1. While the informed buyers know θ , their certainty equivalent is $CE_I = \left[\theta Q - \frac{Q^2}{2} + y\right]$. The certainty equivalent of the learning buyers is

$$CE_L = \left[\hat{\mu}_\theta - \frac{1 + a_L \hat{\sigma}_\theta^2}{2} Q\right] Q + y \quad (5)$$

The demand functions of both types of buyers are then given, respectively, by :

$$Q_I^D = \theta - P \quad (6)$$

$$Q_L^D = \frac{\hat{\mu}_\theta - P}{1 + a_L \hat{\sigma}_\theta^2} \quad (7)$$

The demand of the informed buyers is the difference between the quality, θ , and price. The demand of the learning buyers is the difference between the expected value of their posterior beliefs about quality, $\hat{\mu}_\theta$, and price, divided by $(1 + a_L \hat{\sigma}_\theta^2)$. This last term reflects the effect of risk aversion and posterior uncertainty. When $a_L \rightarrow 0$, i.e. when we approach risk neutrality, the demand function of the learning buyers tends to the demand function of the informed buyers, but with θ replaced by the posterior beliefs. When either the coefficient of risk aversion or the posterior variance increases, demand by the learning buyers decreases.

Total demand is the sum of the demands of the informed and of the learning agents. This sum is given by

$$Q^D(P, \theta) = \lambda(\theta - P) + (1 - \lambda) \frac{\hat{\mu}_\theta - P}{1 + a_L \hat{\sigma}_\theta^2} \quad (8)$$

From (8), the inverse demand is given by

$$P = \frac{(1 + a_L \hat{\sigma}_\theta^2) \lambda \theta + (1 - \lambda) \hat{\mu}_\theta}{(1 + a_L \hat{\sigma}_\theta^2) \lambda + (1 - \lambda)} - \frac{(1 + a_L \hat{\sigma}_\theta^2) Q^D}{(1 + a_L \hat{\sigma}_\theta^2) \lambda + (1 - \lambda)} \quad (9)$$

Note that, while the inverse supply function has a random component, the inverse demand has not. However, the variability of the equilibrium price depends on demand through the slope of the inverse demand, which is

$$\frac{\partial P}{\partial Q^D} = - \frac{(1 + a_L \hat{\sigma}_\theta^2)}{(1 + a_L \hat{\sigma}_\theta^2) \lambda + (1 - \lambda)} \quad (10)$$

which becomes steeper, i.e. more negative, when risk-aversion increases:

$$\frac{\partial^2 P}{\partial Q^D \partial a_L} = - \frac{(1 - \lambda) \hat{\sigma}_\theta^2}{[(1 + a_L \hat{\sigma}_\theta^2) \lambda + (1 - \lambda)]^2} < 0 \quad (11)$$

This rotation of the demand function due to risk aversion and its effect on the variance of the equilibrium price is further discussed in section 2.2.

2.1 The equilibrium

Having described the market structure, we now define the equilibrium. The learning competitive equilibrium consists of the quantity of the firm, $Q^*(\theta, \varepsilon)$, the uninformed buyers' posterior beliefs about quality upon observing the realized price, $\hat{\xi}^*(\cdot|P)$, and the distribution of the market-clearing price $P^*(\theta, \varepsilon)$. 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 n -tuple $[Q^*(\theta, \varepsilon), \hat{\xi}^*(\cdot|P), P^*(\theta, \varepsilon)]$ is a competitive equilibrium with learning if,*

1. **Firm.** For all (θ, ε) , given $P^*(\theta, \varepsilon)$,

$$Q^*(\theta, \varepsilon) = \arg \max_{Q \geq 0} \{P^*(\theta, \varepsilon)Q - C(Q, \theta, \varepsilon)\}. \quad (12)$$

2. **Uninformed buyers.** For all θ , given the price-signal $P^*(\theta, \varepsilon)$ with corresponding p.d.f. $\phi_P^*(P|\theta)$, posterior beliefs upon observing the realization $P = P^*(\theta, \varepsilon)$ are

$$\hat{\xi}^*(\theta_p|P) \propto \xi(\theta_a)\phi_P^*(P|\theta) \quad (13)$$

by Bayes' rule.

3. **Market-clearing price.** For all (θ, ε) , given $Q^*(\theta, \varepsilon)$ and $\hat{\xi}^*(\cdot|P)$, $P^*(\theta, \varepsilon)$ clears the market, i.e.,

$$\lambda(\theta - P^*(\theta, \varepsilon)) + (1 - \lambda)\left(\frac{\hat{\mu}_\theta^*(P)|_{P=P^*(\theta, \varepsilon)} - P^*(\theta, \varepsilon)}{1 + a_L \hat{\sigma}_\theta^2}\right) = Q^*(\theta, \varepsilon) \quad (14)$$

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

From Statement 1 of the definition of equilibrium, the firm's conjecture about the price, after observing ε , is correct. Moreover, from Statements 2, the uninformed buyers' conjecture, not knowing ε , 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 2, in equilibrium, posterior beliefs depend on the correct conditional distribution of the price-signal.

Proposition 2.2 states that there exists a unique equilibrium in which the price retains the normal distribution and the updating rule is a linear function of the prior mean and the price signal.

Proposition 2.2. *Under the conditions of Definition 2.1, with cost function defined by (1), there exists a unique learning competitive equilibrium (LE) in which the updating rule for the posterior mean is a linear function of the prior beliefs and the price signal.*

1. In equilibrium, the posterior beliefs are given by $\hat{\xi}^*(\theta_p|P) \sim n(\hat{\mu}_\theta, \hat{\sigma}_\theta^2)$,

with

$$\hat{\mu}_\theta = a\mu_\theta + bP^*(\theta, \varepsilon) \quad (15)$$

and

$$\hat{\sigma}_\theta^2 = \frac{\sigma_\varepsilon^2 \sigma_\theta^2}{\sigma_\varepsilon^2 + \sigma_\theta^2 [\gamma_\theta + \gamma_q \lambda]^2} \quad (16)$$

with

$$a = \frac{\sigma_\varepsilon^2 (1 + a_L \hat{\sigma}_\theta^2)}{(1 + a_L \hat{\sigma}_\theta^2) [\sigma_\varepsilon^2 + \sigma_\theta^2 [\gamma_\theta + \gamma_q \lambda]^2] + \gamma_q (1 - \lambda) [\gamma_\theta + \gamma_q \lambda] \sigma_\theta^2} \quad (17)$$

and

$$b = \frac{(1 + \gamma_q \lambda) (1 + a_L \hat{\sigma}_\theta^2) [\gamma_\theta + \gamma_q \lambda] \sigma_\theta^2 + \gamma_q (1 - \lambda) [\gamma_\theta + \gamma_q \lambda] \sigma_\theta^2}{(1 + a_L \hat{\sigma}_\theta^2) [\sigma_\varepsilon^2 + \sigma_\theta^2 [\gamma_\theta + \gamma_q \lambda]^2] + \gamma_q (1 - \lambda) [\gamma_\theta + \gamma_q \lambda] \sigma_\theta^2} \quad (18)$$

2. The equilibrium price is

$$\begin{aligned} P^*(\theta, \varepsilon) &= \frac{(\sigma_\varepsilon^2 + \sigma_\theta^2 [\gamma_\theta + \lambda \gamma_q]^2) (1 + a_L \hat{\sigma}_\theta^2) + \gamma_q (1 - \lambda) [\gamma_\theta + \gamma_q \lambda] \sigma_\theta^2}{(\sigma_\varepsilon^2 + [\gamma_\theta + \lambda \gamma_q]^2 \sigma_\theta^2) [(1 + \gamma_q \lambda) (1 + a_L \hat{\sigma}_\theta^2) + \gamma_q (1 - \lambda)]} [\gamma_\theta + \lambda \gamma_q] \theta \\ &\quad + \frac{\sigma_\varepsilon^2 \gamma_q (1 - \lambda)}{(\sigma_\varepsilon^2 + [\gamma_\theta + \lambda \gamma_q]^2 \sigma_\theta^2) [(1 + \gamma_q \lambda) (1 + a_L \hat{\sigma}_\theta^2) + \gamma_q (1 - \lambda)]} \mu_\theta \\ &\quad + \frac{(\sigma_\varepsilon^2 + \sigma_\theta^2 [\gamma_\theta + \lambda \gamma_q]^2) (1 + a_L \hat{\sigma}_\theta^2) + \gamma_q (1 - \lambda) [\gamma_\theta + \gamma_q \lambda] \sigma_\theta^2}{(\sigma_\varepsilon^2 + [\gamma_\theta + \lambda \gamma_q]^2 \sigma_\theta^2) [(1 + \gamma_q \lambda) (1 + a_L \hat{\sigma}_\theta^2) + \gamma_q (1 - \lambda)]} \varepsilon \end{aligned} \quad (19)$$

with distribution given by $P(\theta, \tilde{\varepsilon}, \hat{\xi}^*(\cdot)) \sim n(\hat{\mu}_P, \hat{\sigma}_P^2)$ ⁶.

⁶Here $\hat{\mu}_P$ is written as

$$\begin{aligned} \hat{\mu}_P &= \lambda \frac{(\gamma_\theta + \gamma_q \lambda) (1 + a_L \hat{\sigma}_\theta^2)}{(1 + \gamma_q \lambda) (1 + a_L \hat{\sigma}_\theta^2) + \gamma_q (1 - \lambda)} \theta \\ &\quad + (1 - \lambda) \frac{(\gamma_\theta + \gamma_q \lambda) (1 + a_L \hat{\sigma}_\theta^2) + \gamma_q \sigma_\varepsilon^2 \mu_\theta + [\gamma_\theta + \gamma_q \lambda]^2 \sigma_\theta^2 \theta}{(\sigma_\varepsilon^2 + [\gamma_\theta + \lambda \gamma_q]^2 \sigma_\theta^2) [(1 + \gamma_q \lambda) (1 + a_L \hat{\sigma}_\theta^2) + \gamma_q (1 - \lambda)]} \end{aligned} \quad (20)$$

This formulation shows that the expected price depends on the behavior of the two types of consumers.

3. *Equilibrium output is*

$$Q^*(\theta, \varepsilon) = \frac{P^*(\theta, \varepsilon) - \gamma_\theta \theta - \varepsilon}{\gamma_q} \quad (21)$$

Proof. Substituting (1) into (12) and solving (12) yields (21). Then assume the posterior mean of the quality parameter is given by the updating rule in (15). Plugging (15) and (21) into (14) and solving for the price, yields the equilibrium price as a function of a and b

$$P(\theta, \varepsilon) = \frac{[\gamma_\theta + \gamma_q \lambda] (1 + a_L \hat{\sigma}_\theta^2) \theta + \gamma_q (1 - \lambda) a \mu_\theta + (1 + a_L \hat{\sigma}_\theta^2) \varepsilon}{(1 + \gamma_q \lambda) (1 + a_L \hat{\sigma}_\theta^2) + \gamma_q (1 - b) (1 - \lambda)} \quad (22)$$

Hence the conditional distribution

$$\hat{P}|\theta \sim n(\hat{\mu}_{P|\theta}; \hat{\sigma}_{P|\theta}^2) \quad (23)$$

with

$$\hat{\mu}_{P|\theta} = \frac{[\gamma_\theta + \gamma_q \lambda] (1 + a_L \hat{\sigma}_\theta^2) \theta + \gamma_q (1 - \lambda) a \mu_\theta}{(1 + \gamma_q \lambda) (1 + a_L \hat{\sigma}_\theta^2) + \gamma_q (1 - b) (1 - \lambda)} \quad (24)$$

$$\hat{\sigma}_{P|\theta}^2 = \left[\frac{(1 + a_L \hat{\sigma}_\theta^2)}{(1 + \gamma_q \lambda) (1 + a_L \hat{\sigma}_\theta^2) + \gamma_q (1 - b) (1 - \lambda)} \right]^2 \sigma_\varepsilon^2 \quad (25)$$

We define

$$z = \frac{[(1 + \gamma_q \lambda) (1 + a_L \hat{\sigma}_\theta^2) + \gamma_q (1 - b) (1 - \lambda)]}{[\gamma_\theta + \gamma_q \lambda] (1 + a_L \hat{\sigma}_\theta^2)} P(\theta, \varepsilon) - \frac{\gamma_q (1 - \lambda) a}{[\gamma_\theta + \gamma_q \lambda] (1 + a_L \hat{\sigma}_\theta^2)} = \theta + \frac{\varepsilon}{[\gamma_\theta + \gamma_q \lambda]} \quad (26)$$

with conditional distribution $z|\theta$

$$z|\theta \sim n\left(\theta; \frac{\sigma_\varepsilon^2}{[\gamma_\theta + \gamma_q \lambda]^2}\right) \quad (27)$$

From (27) and from the prior distribution of beliefs (4) the posterior distribution $\tilde{\theta}_p|z$ is obtained,

$$\hat{\theta}_p|z \sim n \left(\frac{z [\gamma_\theta + \gamma_q \lambda]^2 \sigma_\theta^2 + \sigma_\varepsilon^2 \mu_\theta}{\sigma_\varepsilon^2 + \sigma_\theta^2 [\gamma_\theta + \gamma_q \lambda]^2}; \frac{\sigma_\varepsilon^2 \sigma_\theta^2}{\sigma_\varepsilon^2 + \sigma_\theta^2 [\gamma_\theta + \gamma_q \lambda]^2} \right) \quad (28)$$

Substituting in (28) z for its expression in order of P (given in (26)), we obtain the posterior pdf for $\hat{\theta}_p|P$, i.e.,

$$\hat{\theta}_p|P \sim n(\hat{\mu}_\theta; \hat{\sigma}_\theta^2) \quad (29)$$

with

$$\hat{\mu}_\theta = \frac{(1 + a_L \hat{\sigma}_\theta^2) \sigma_\varepsilon^2 \mu_\theta - \gamma_q (1 - \lambda) (\gamma_\theta + \gamma_q \lambda) \sigma_\theta^2 a}{[\sigma_\varepsilon^2 + \sigma_\theta^2 (\gamma_\theta + \gamma_q \lambda)^2] (1 + a_L \hat{\sigma}_\theta^2)} + \frac{(1 + a_L \hat{\sigma}_\theta^2) (1 + \gamma_q \lambda) + \gamma_q (1 - b) (1 - \lambda) (\gamma_\theta + \gamma_q \lambda) \sigma_\theta^2}{[\sigma_\varepsilon^2 + \sigma_\theta^2 (\gamma_\theta + \gamma_q \lambda)^2] (1 + a_L \hat{\sigma}_\theta^2)} P(\theta, \varepsilon) \quad (30)$$

$$\hat{\sigma}_\theta^2 = \frac{\sigma_\varepsilon^2 \sigma_\theta^2}{(\gamma_\theta + \gamma_q \lambda)^2 \sigma_\theta^2 + \sigma_\varepsilon^2} \quad (31)$$

Equating $\hat{\mu}_\theta = a\mu_\theta + bP$ and solving for a and b we arrive at (17) and (18), confirming the existence of a linear updating rule. Substituting (17) and (18) into (22) yields (19). \square

Note that the distribution of the posterior beliefs has a nonrandom variance (see (16)) but a random mean. By (15), the posterior mean is a linear function of the random price, and so also has a normal distribution.

The updating rule $\chi(P)$ combines the prior beliefs of the learning consumers (given by μ_θ) with the information inferred from the price (see (15), (17) and (18)). The weights given to the prior beliefs and the price, respectively a and b , vary with: i) the ratio between the variance of the prior beliefs and the variance of the random supply noise ($\sigma_\theta^2/\sigma_\varepsilon^2$); ii) the coefficient of risk aversion (a_L) and the posterior variance ($\hat{\sigma}_\theta^2$, which itself depends on the parameters in i) and iii)); iii) the other parameters (λ , γ_θ and γ_q), which influence the slopes of the demand and supply curves. The expected equilibrium price and quantity depend on all these parameters.

2.2 Discussion

In this section we study the impact that the learning risk averse buyers have on the equilibrium. For this purpose we compare the LE in Proposition 2.2 with the full information equilibrium (FIE), i.e., the equilibrium when all buyers know θ .

We calculate the FIE, substituting $\lambda = 1$ in equations (15) to (21).

Remark 2.3. *With $\lambda = 1$, FIE is:*

$$P = \frac{\gamma_\theta + \gamma_q}{1 + \gamma_q} \theta + \frac{\varepsilon}{1 + \gamma_q} \quad (32)$$

$$\hat{\mu}_P = \frac{\gamma_\theta + \gamma_q}{1 + \gamma_q} \theta \text{ and } \hat{\sigma}_P^2 = \frac{\sigma_\varepsilon^2}{(1 + \gamma_q)^2}$$

$$E[Q|\theta] = \frac{1 - \gamma_\theta}{1 + \gamma_q} \theta \quad (33)$$

For the expected output to be positive we assume that $\gamma_\theta < 1$.

The difference between this FIE and the LE from Proposition 2.2 is due to the impact of the presence of uninformed buyers. We depict the FIE and the LE in Fig.1, assuming $\mu_\theta < \theta^7$.

To understand the impact of the presence of risk-averse learning buyers, the change from the FIE to the LE can be decomposed into two components. The first is due to the presence of uninformed but nonlearning (*naive*) buyers, i.e., buyers that do not update their prior beliefs. The second component reflects the change in equilibrium due to the learning process, i.e., the updating of prior beliefs upon observing the price. We call the first component the *prior beliefs effect* and the second component the *price effect*.

To identify the *prior beliefs effect* we calculate the equilibrium of the intermediate case with *naive* uninformed buyers⁸. In this intermediate case,

⁷We could also assume that $\mu_\theta > \theta$. In this case the LE demand function would start above the FIE demand function and would have a steeper slope

⁸This *naive* equilibrium is important in itself. We show in Section 3.2 that, when uncertainty increases, the LE tends to the *naive* equilibrium

demand by the learning buyers depends on the prior mean and variance. Total demand is given by

$$Q^D(P, \theta) = \lambda(\theta - P) + (1 - \lambda) \frac{\mu_\theta - P}{1 + a_L \sigma_\theta^2} \quad (34)$$

Proposition 2.4. *Assuming naive (nonlearning) uninformed buyers, with supply and demand defined by (1) and (34), the competitive equilibrium is given by, :*

1.

$$P(\theta, \varepsilon) = \frac{[\gamma_\theta + \gamma_q \lambda] (1 + a_L \sigma_\theta^2) \theta + \gamma_q (1 - \lambda) \mu_\theta + (1 + a_L \sigma_\theta^2) \varepsilon}{(1 + \gamma_q \lambda) (1 + a_L \sigma_\theta^2) + \gamma_q (1 - \lambda)} \quad (35)$$

and

$$\tilde{P}(\theta, \varepsilon) \sim n(\hat{\mu}_P, \hat{\sigma}_P^2) \quad (36)$$

2.

$$Q(\theta, \varepsilon) = \frac{P(\theta, \varepsilon) - \gamma_\theta \theta - \varepsilon}{\gamma_q} \quad (37)$$

Proof. Equation (37) derives immediately from (1), while (35) results from (22), with $a = 1$ and $b = 0$. Finally, from (35) P is a linear function of ε , hence it is normally distributed. \square

Note that, from (34), the *naive* inverse demand is given by

$$P(Q^D, \theta) = \frac{(1 + a_L \sigma_\theta^2) \lambda \theta + (1 - \lambda) \mu_\theta - (1 + a_L \sigma_\theta^2) Q^D}{(1 + a_L \sigma_\theta^2) \lambda + (1 - \lambda)} \quad (38)$$

From (38), $P(Q^D|Q^D = 0)$, the intercept on the y-axis, is equal to (larger than, smaller than) the equivalent intercept in the FIE expected demand if $\mu_\theta = \theta$ ($\mu_\theta > \theta$, $\mu_\theta < \theta$). If the uninformed buyers are *naive* then the demand function shifts up or down (relative to the FIE), depending on the prior being over or undervalued. The presence of uninformed risk-averse *naive* buyers also rotates the demand function, increasing its slope⁹. So,

⁹From (34), $Q^D(P|P = 0)$, the x-axis intercept, is smaller than the FI solution if $\mu_\theta \leq \theta$, and can go either way if $\mu_\theta > \theta$.

even if prior beliefs are unbiased, there is always a difference between *naive* and FI demand, due to the risk aversion of the uninformed buyers. Note that in the case of risk neutral¹⁰ *naive* buyers the demand curve does not rotate. This rotation of the *naive* demand curve is due to the effect of risk aversion and has an impact on the equilibrium quantity and price.

In fact, even though the supply function is not affected by risk aversion, the presence of uninformed *naive* buyers changes the equilibrium price and output, again *even if prior beliefs are unbiased*. The increase in the slope of the demand curve (compared with the FIE) due to the presence of uninformed *naive* buyers means that: a) if $\mu_\theta \leq \theta$, the *naive* price and output are always below the FI solutions; b) if $\mu_\theta > \theta$, the *naive* equilibrium price and output may be above, below or coincide with the FI solutions, depending on the values of all parameters. We depict these two situation in Fig.2 and Fig.3.

Next consider the *price effect*. The change in equilibrium due to the *price effect* is the difference between the naive equilibrium of Proposition 2.4 and the LE of Proposition 2.2. This *price effect* shifts the LE demand function back towards the FI demand while also rotating it. This is similar to the risk neutral case. However, the presence of risk aversion ensures that the LE demand never coincides with the FIE demand even if prior beliefs are unbiased, as in the risk neutral case. Fig.4 and Fig.5 illustrate the cases of under and overvalued prior.

3 The effects of risk aversion and uncertainty

In this section we study the effect of uncertainty and risk aversion on the learning process and the learning equilibrium. When referring to uncertainty, we must distinguish between prior uncertainty (introduced by the prior variance, σ_θ^2 , and the variance of noisy demand, σ_ε^2) and posterior uncertainty, characterized by the posterior variance, $\hat{\sigma}_\theta^2$. We show that risk aversion affects the learning process changing the distribution of $\hat{\theta}_p|P$, i.e.,

¹⁰Henceforth, when we refer the risk neutral case we are referring to the limit case when $a_L \rightarrow 0$ and the demand of the learning buyers tends to $Q_L^D = \hat{\mu}_\theta - P$. This limit case corresponds to the risk-neutral case studied in (Mirman et al., 2014).

the conditional distribution of posterior beliefs for an observed P . However, it has no effect on $\mathbb{E}\hat{\mu}_\theta^*(P^*(\theta, \tilde{\varepsilon}))$, the average value of the posterior mean¹¹. Increasing risk aversion shifts the price distribution, lowering expected price and expected output. On the other hand, prior uncertainty affects both the learning process and the marginal posterior distribution. Finally, the effects of risk aversion and posterior uncertainty cannot be studied separately.

Starting with this last comment, notice that the coefficient of risk aversion always appears connected with posterior uncertainty, in $(1 + a_L \hat{\sigma}_\theta^2)$. As noted in Section 2, the term $(1 + a_L \hat{\sigma}_\theta^2)$ reflects the effect of posterior uncertainty about quality on the demand function of the uninformed buyers. Without posterior uncertainty ($\hat{\sigma}_\theta^2 = 0$) there is no risk aversion effect and without risk aversion ($a_L \rightarrow 0$) there is no effect of posterior uncertainty.

The posterior variance, $\hat{\sigma}_\theta^2 = \frac{\sigma_\varepsilon^2 \sigma_\theta^2}{\sigma_\varepsilon^2 + \sigma_\theta^2 [\gamma_\theta + \gamma_q \lambda]^2}$, increases with the prior variance, σ_θ^2 , and the variance of the supply noise, σ_ε^2 . Moreover, $\hat{\sigma}_\theta^2$ decreases with γ_θ , the impact of quality on the cost, with the slope of the supply function, γ_q , and with λ , the fraction of informed buyers.

If either a_L or $\hat{\sigma}_\theta^2$ tend to zero we approach the risk neutral case and the same results as in Mirman et al. (2014) are applicable.

Remark 3.1. *If $a_L \rightarrow 0$ or $\hat{\sigma}_\theta^2 = 0^{12}$, then the LE tends to the risk neutral results:*

$$a = \frac{\sigma_\varepsilon^2}{(\gamma_\theta + \gamma_q)(\gamma_\theta + \gamma_q \lambda)\sigma_\theta^2 + \sigma_\varepsilon^2} \quad (39)$$

$$b = \frac{(1 + \gamma_q)(\gamma_\theta + \gamma_q \lambda)\sigma_\theta^2}{(\gamma_\theta + \gamma_q)(\gamma_\theta + \gamma_q \lambda)\sigma_\theta^2 + \sigma_\varepsilon^2} \quad (40)$$

¹¹Here, \mathbb{E} is the expectation operator with respect to the p.d.f. $\phi_P^*(\cdot|\theta)$.

¹² $\hat{\sigma}_\theta^2 = 0$ may result from two different situations: if $\gamma_q = 0$ or $\sigma_\varepsilon^2 = 0$ we have perfect learning and non random price; if $\sigma_\theta^2 = 0$ we have no learning at all.

$$\hat{\mu}_P = \frac{\gamma_\theta + \gamma_q \lambda}{1 + \gamma_q} \left\{ \frac{[\gamma_\theta + \gamma_q] [\gamma_\theta + \gamma_q \lambda] \sigma_\theta^2 + \sigma_\varepsilon^2}{[\gamma_\theta + \gamma_q \lambda]^2 \sigma_\theta^2 + \sigma_\varepsilon^2} \right\} \theta + \frac{(1 - \lambda) \gamma_q}{1 + \gamma_q} \left\{ \frac{\sigma_\varepsilon^2}{[\gamma_\theta + \gamma_q \lambda]^2 \sigma_\theta^2 + \sigma_\varepsilon^2} \right\} \mu_\theta \quad (41)$$

$$E [Q|\theta] = \frac{\hat{\mu}_P - \gamma_\theta \theta - \varepsilon}{\gamma_q} \quad (42)$$

3.1 The effect of risk aversion

This section studies the impact that risk aversion has on the learning process and on the LE. Starting with the learning process, we analyze the effect of risk aversion on the updating rule, $\hat{\mu}_\theta = a\mu_\theta + bP$. Note that $\hat{\mu}_\theta$ is not a convex combination of μ_θ and P , i.e., $a + b \neq 1$. Define

$$b' = \frac{[(\gamma_\theta + \gamma_q \lambda)(1 + a_L \hat{\sigma}_\theta^2) + \gamma_q(1 - \lambda)] [\gamma_\theta + \gamma_q \lambda] \sigma_\theta^2}{\sigma_\varepsilon^2(1 + a_L \hat{\sigma}_\theta^2) + [(\gamma_\theta + \gamma_q \lambda)(1 + a_L \hat{\sigma}_\theta^2) + \gamma_q(1 - \lambda)] [\gamma_\theta + \gamma_q \lambda] \sigma_\theta^2} \quad (43)$$

Then $a + b' = 1$, and the posterior mean is given by the convex combination $\hat{\mu}_\theta = a\mu_\theta + b'P'$, where P' , the revised price¹³, is

$$P' = \frac{(1 + \gamma_q \lambda)(1 + a_L \hat{\sigma}_\theta^2) + \gamma_q(1 - \lambda)}{(\gamma_\theta + \gamma_q \lambda)(1 + a_L \hat{\sigma}_\theta^2) + \gamma_q(1 - \lambda)} P > P \quad (44)$$

We now look to the effect of risk aversion on the weights, a and b' . Dividing both the numerator and the denominator of a by $(1 + a_L \hat{\sigma}_\theta^2)$ yields:

$$a = \frac{\sigma_\varepsilon^2}{\sigma_\varepsilon^2 + \left[[\gamma_\theta + \gamma_q \lambda] + \frac{\gamma_q(1 - \lambda)}{(1 + a_L \hat{\sigma}_\theta^2)} \right] [\gamma_\theta + \gamma_q \lambda] \sigma_\theta^2} \quad (45)$$

Compared with the coefficient a in the risk neutral case from (39), note that the denominator decreases with a_L , thus increasing a . As $a + b' = 1$, when a increases b' decreases. Hence, the presence of risk aversion changes

¹³With Full Information, $\lambda = 1$, the expected value of this revised price is equal to the true quality, θ .

the weights of the prior beliefs and the revised observed price in the learning process, giving more weight to prior beliefs and giving less to the information conveyed by the revised observed price¹⁴.

Despite these changes in the coefficients of the updating rule, from Proposition 3.2 the effect of risk aversion on $\mathbb{E}\hat{\mu}_\theta^*(P^*(\theta, \varepsilon))$ is null. In fact, risk aversion changes the posterior distribution of $\tilde{\theta}_p|P$ but not the conditional distribution of $\tilde{\theta}_p|\varepsilon$ ¹⁵. This means the price signal adjusts to compensate for the changes in a and b' , in a way that ensures that, for any realization of the supply noise ε the posterior beliefs in the risk aversion case are the same as in the risk neutral case.

Proposition 3.2. *Under the conditions of Definition 2.1, with supply and demand defined by (1), (12) and (8), the coefficient of risk aversion has no effect on the conditional distribution of the posterior beliefs for a given noise, $\hat{\xi}^*(\theta_p|\varepsilon) \sim n\left(\mu_{\theta_p}, \sigma_{\theta_p}^2\right)$, with μ_{θ_p} and $\sigma_{\theta_p}^2$ independent from a_L .*

Moreover, learning always occurs on average, i.e.:

1.

$$\mathbb{E}\hat{\mu}_\theta^*(\tilde{P}^*) = \alpha\mu_\theta + (1 - \alpha)\theta, 0 < \alpha \leq 1 \quad (46)$$

2.

$$\hat{\sigma}_\theta^2 = \frac{\sigma_\varepsilon^2 \sigma_\theta^2}{\sigma_\varepsilon^2 + \sigma_\theta^2 [\gamma_\theta + \gamma_q \lambda]^2} < \sigma_\theta^2 \quad (47)$$

Proof. Substituting (26) in (28) we obtain

$$\hat{\xi}^*(\theta_p|\varepsilon) \sim n\left(\frac{[\gamma_\theta + \gamma_q \lambda]^2 \sigma_\theta^2 \left[\theta + \frac{\varepsilon}{\gamma_\theta + \gamma_q \lambda}\right] + \sigma_\varepsilon^2 \mu_\theta}{\sigma_\varepsilon^2 + \sigma_\theta^2 [\gamma_\theta + \gamma_q \lambda]^2}; \frac{\sigma_\varepsilon^2 \sigma_\theta^2}{\sigma_\varepsilon^2 + \sigma_\theta^2 [\gamma_\theta + \gamma_q \lambda]^2}\right) \quad (48)$$

¹⁴Note that P' itself changes with risk aversion, increasing away from P . However the revised price always lays in the interval $\left[\frac{1+\gamma_q}{\gamma_\theta+\gamma_q}P; \frac{1+\gamma_q\lambda}{\gamma_\theta+\gamma_q\lambda}P\right]$, as a_L goes from zero to infinity.

¹⁵The conditional distribution of $\tilde{\theta}_p|P$ is the subjective distribution of the posterior beliefs of the learning buyers, resulting from their updating process for a given equilibrium price. The conditional distribution of $\tilde{\theta}_p|\varepsilon$ is not relevant for the learning buyers, because they cannot observe the random noise ε . However, as the informed observers (us) know the function $P^*(\theta, \varepsilon)$, we can predict how the posterior beliefs react for each realization of the random noise, obtaining the conditional distribution of $\tilde{\theta}_p|\varepsilon$.

showing that μ_{θ_p} and $\sigma_{\theta_p}^2$ are independent from a_L , the coefficient of risk aversion. Moreover, μ_{θ_p} is random (linear in ε) while $\sigma_{\theta_p}^2$ is nonrandom (independent of ε). Taking expectations of μ_{θ_p} with respect to the distribution of ε , we obtain

$$E\mu_{\theta_p} = \mathbb{E}\hat{\mu}_{\theta}^*(\tilde{P}^*) = \frac{[\gamma_{\theta} + \gamma_q\lambda]^2 \sigma_{\theta}^2 \theta + \sigma_{\varepsilon}^2 \mu_{\theta}}{\sigma_{\varepsilon}^2 + \sigma_{\theta}^2 [\gamma_{\theta} + \gamma_q\lambda]^2} \quad (49)$$

From (49) we derive (46), with $\alpha = \frac{\sigma_{\varepsilon}^2}{\sigma_{\varepsilon}^2 + (\gamma_{\theta} + \gamma_q\lambda)^2 \sigma_{\theta}^2}$. Equation (47) comes from (16). \square

From Proposition 3.2 the conditional distribution $\tilde{\theta}_p|\varepsilon$ does not depend on risk aversion, i.e., for a given realization of the random noise, ε , the posterior beliefs are exactly the same with or without risk aversion. However, for a given observed price, posterior beliefs depend on risk aversion, because the weights a and b depend on risk aversion.

Risk aversion changes the likelihood function, i.e., the distribution of price for a given θ , so that the same realization of the supply noise ε translates into different observed prices in the risk aversion and the risk neutral case. The effect of risk aversion is to diminish the demand of the learning buyers, so the demand of the informed buyers has a bigger fraction of total demand. Then the observed price becomes depends more on the informed buyers than with no risk aversion. In other words, the price in the risk averse case is more dependent (comparing with the risk neutral case) on θ and less on the prior beliefs of the uninformed buyers. This change in the price completely compensate the changes in a and b .

Analytically, equation (19) shows that the expected mean of the equilibrium price is a linear combination of θ and μ_{θ} . After some manipulation, the expected equilibrium price is written as:

$$\hat{\mu}_{P|\theta} = \frac{1}{b} \left\{ \frac{[\gamma_{\theta} + \gamma_q\lambda]^2 \sigma_{\theta}^2}{(\gamma_{\theta} + \gamma_q\lambda)^2 \sigma_{\theta}^2 + \sigma_{\varepsilon}^2} \theta + \left[\frac{\sigma_{\varepsilon}^2}{(\gamma_{\theta} + \gamma_q\lambda)^2 \sigma_{\theta}^2 + \sigma_{\varepsilon}^2} - a \right] \mu_{\theta} \right\} \quad (50)$$

Hence, as the presence of risk aversion changes the weights a and b , the

expected price is inversely proportional to b and annuls the effect of a , maintaining the expected posterior mean unchanged. This complete compensation happens because the posterior distribution can be written as a function of a sufficient statistic, z (used in the proof of Proposition 2.2), which does not depend on risk aversion¹⁶.

From Proposition 3.2, the effect of risk aversion on the expected price and output is determined. Compare the risk averse and the risk neutral cases. From equation (14), for the same random noise, if the equilibrium price was equal in both cases, the demand at the market equilibrium with positive risk aversion would be smaller than without risk aversion, creating excess supply. So, for the same random noise, with risk aversion the equilibrium price must be smaller than in the risk neutral case. Note that, in this argument, the supply function and the posterior beliefs are the same in both cases, as was shown in Proposition 3.2. The supply function remaining unchanged, output also decreases with risk aversion. If, for any random noise, equilibrium price and output are smaller with risk aversion, then expected price and output decrease with risk aversion.

Remark 3.3. *Expected price and output diminishes with risk aversion.*

3.2 The effect of prior uncertainty

In this section we study the effect of both components of prior uncertainty: the variance of the supply noise, σ_ε^2 (henceforth called supply uncertainty) and the prior variance, σ_θ^2 .

¹⁶(Grossman and Stiglitz, 1980) arrive at similar results, but concerning the effect of the variance of the demand shock on the informativeness of the price. They say: “..an increase in noise reduces the informativeness of the price system: but it...leads more individuals to become informed; the remarkable result obtained above establishes that *the two effects exactly offset each other* so that the equilibrium informativeness of the price system is unchanged” (their italic). In their model it is the percentage of informed traders, λ , that is endogenous (in our model it is exogenous) and so it is λ that adjusts. Again, in (Grossman and Stiglitz, 1980), what happens is that there is a sufficient statistic for the unknown parameter θ ($w_\lambda(\theta, x)$) which does not depend on the demand uncertainty, hence the latter has no influence on posterior beliefs.

3.2.1 The effect on learning

When $\sigma_\varepsilon^2 = 0$ there is no posterior uncertainty, learning becomes perfect and risk aversion irrelevant. Hence, as in the FI equilibrium.

$$\begin{aligned}\hat{\mu}_\theta &= \theta \text{ and } \hat{\sigma}_\theta^2 = 0 \\ a &= 0 \text{ and } b = \frac{1+\gamma_q}{\gamma_\theta+\gamma_q} \\ \hat{\sigma}_P^2 &= 0 \text{ and } P = \hat{\mu}_P = \frac{\gamma_\theta+\gamma_q}{1+\gamma_q}\theta \\ Q &= E[Q|\theta] = \frac{1-\gamma_\theta}{1+\gamma_q}\theta\end{aligned}$$

On the other hand, when $\sigma_\theta^2 = 0$ there is no posterior uncertainty and risk aversion becomes irrelevant, but there is no learning¹⁷. Hence,

$$\begin{aligned}\hat{\mu}_\theta &= \mu_\theta \text{ and } \hat{\sigma}_\theta^2 = 0 \\ a &= 1 \text{ and } b = 0 \\ \hat{\mu}_P &= \frac{\gamma_\theta\theta+\gamma_q[\lambda\theta+(1-\lambda)\mu_\theta]}{1+\gamma_q} \text{ and } \hat{\sigma}_P^2 = \frac{\sigma_\varepsilon^2}{(1+\gamma_q)^2} \\ Q &= \frac{\lambda\theta+(1-\lambda)\mu_\theta-\gamma_\theta\theta-\varepsilon}{1+\gamma_q}\end{aligned}$$

Next consider the case of $\sigma_\varepsilon^2 > 0$ and $\sigma_\theta^2 > 0$. Contrary to the effect of risk aversion, prior uncertainty does influence the posterior distribution, through both the posterior mean and the posterior variance. Prior uncertainty affects the posterior mean through the updating rule. The weight a (b' is $1-a$), depends directly on the ratio of the prior and the supply noise variance, $\frac{\sigma_\theta^2}{\sigma_\varepsilon^2}$:

$$a = \frac{1}{1 + \left[[\gamma_\theta + \gamma_q\lambda]^2 + \frac{\gamma_q(1-\lambda)[\gamma_\theta+\gamma_q\lambda]}{(1+a_L\hat{\sigma}_\theta^2)} \right] \frac{\sigma_\theta^2}{\sigma_\varepsilon^2}} \quad (51)$$

This direct effect implies that, when $\frac{\sigma_\theta^2}{\sigma_\varepsilon^2}$ increases (decreases) a decreases (increases) and b' increases (decreases). Hence, if the variability of the prior

¹⁷If $\sigma_\varepsilon^2 = \sigma_\theta^2 = 0$ we have an indeterminacy. The learning buyers are completely sure about their prior beliefs and, simultaneously, price is perfectly informative. What should they believe?

beliefs increases relative to the supply noise variability, learning buyers put more weight on the observed information (b') and less weight on the prior beliefs (a).

But there is also an indirect effect, acting through the posterior variance $\hat{\sigma}_\theta^2$. When either σ_ε^2 or σ_θ^2 increases so does the posterior variance $\hat{\sigma}_\theta^2$. An increase in $\hat{\sigma}_\theta^2$ has the same effect as risk aversion: increases a , decreases b' . Hence:

1. If σ_ε^2 increases, $\frac{\sigma_\theta^2}{\sigma_\varepsilon^2}$ decreases and $\hat{\sigma}_\theta^2$ increases, implying that a increases, while b' decreases. More supply uncertainty increases the weight of prior beliefs and diminishes the weight of the revised observed price.
2. If σ_ε^2 decreases, $\frac{\sigma_\theta^2}{\sigma_\varepsilon^2}$ increases and $\hat{\sigma}_\theta^2$ decreases, implying that a decreases and b' increases. Less supply uncertainty decreases the weight of prior beliefs and increases the weight of the revised observed price.
3. If σ_θ^2 increases, both $\frac{\sigma_\theta^2}{\sigma_\varepsilon^2}$ and $\hat{\sigma}_\theta^2$ increase, and the effect in a and b' is ambiguous.
4. If σ_θ^2 decreases, both $\frac{\sigma_\theta^2}{\sigma_\varepsilon^2}$ and $\hat{\sigma}_\theta^2$ decrease, and the effect in a and b' is ambiguous.

As $\hat{\sigma}_\theta^2$ changes, so changes $\frac{(1+\gamma_q\lambda)(1+a_L\hat{\sigma}_\theta^2)+\gamma_q(1-\lambda)}{(\gamma_\theta+\gamma_q\lambda)(1+a_L\hat{\sigma}_\theta^2)+\gamma_q(1-\lambda)}$, the proportion between P' and P , adding to the ambiguity of the total effects. However, the change in P' is bounded within the interval $\left[\frac{1+\gamma_q}{\gamma_\theta+\gamma_q}P; \frac{1+\gamma_q\lambda}{\gamma_\theta+\gamma_q\lambda}P\right]$. Asymptotically, all ambiguities disappear.

In the limit, if $\sigma_\varepsilon^2 \rightarrow \infty$ then $a \rightarrow 1$ and $b' \rightarrow 0$. On the other hand, if $\sigma_\theta^2 \rightarrow \infty$ then $a \rightarrow 0$ and $b' \rightarrow 1$. Hence, when the variance of the supply noise increases to infinity, more weight is put on the prior beliefs and less weight on the observed price, reducing learning. In fact, from (15) and (16), in the limit there is no learning if $\sigma_\varepsilon^2 \rightarrow \infty$:

$$\hat{\mu}_\theta \rightarrow \mu_\theta \text{ and } \hat{\sigma}_\theta^2 \rightarrow \sigma_\theta^2$$

When the prior variance increases to infinity, i.e., $\sigma_\theta^2 \rightarrow \infty$, learning depends only on the observed information:

$$\hat{\mu}_\theta \rightarrow P' \text{ and } \hat{\sigma}_\theta^2 \rightarrow \frac{\sigma_\varepsilon^2}{[\gamma_\theta+\gamma_q\lambda]^2}$$

3.2.2 The effect on price and output

The market equilibrium (see (14)) is affected by σ_ε^2 through the posterior mean and the posterior variance. The posterior mean tends to the prior mean as market uncertainty increases. If the prior mean is overvalued, $\mu_\theta > \theta$, demand increases; if the prior mean is undervalued, $\mu_\theta < \theta$, demand decreases. The posterior variance increases with σ_ε^2 , thus decreasing demand.

If the prior is undervalued demand unequivocally decreases when σ_ε^2 increases. If the prior is overvalued, the joint effect of the posterior mean and the posterior variance can go either way, depending on the values of the parameters. Again, from Proposition 2.2, the limits of the price distribution as $\sigma_\varepsilon^2 \rightarrow \infty$ are:

$$\hat{\mu}_P \rightarrow \frac{[\gamma_\theta + \gamma_q \lambda](1 + a_L \sigma_\theta^2) \theta + \gamma_q (1 - \lambda) \mu_\theta}{(1 + \gamma_q \lambda)(1 + a_L \sigma_\theta^2) + \gamma_q (1 - \lambda)} \text{ and } \hat{\sigma}_P^2 \rightarrow \infty$$

In the limit, expected price is equal to (35) and the equilibrium tends to the naive equilibrium defined in Proposition 2.4. Because learning disappears in the limit, uninformed buyers act asymptotically naive.

Market equilibrium is also affected by the prior variance, σ_θ^2 , through the posterior mean and the posterior variance. As the prior variance increases the posterior mean tends to the revised price, while the posterior variance increases. If the prior mean is undervalued, $\mu_\theta < \theta$, demand increases; if the prior mean is overvalued, $\mu_\theta > \theta$, demand decreases. The posterior variance increases with σ_θ^2 , thus decreasing demand.

If the prior is overvalued demand unequivocally decreases when σ_θ^2 increases. If the prior is undervalued, the joint effect of the posterior mean and the posterior variance can go either way, depending on the values of the parameters. Again, from Proposition 2.2 and equation (44), the limits of the price distribution as $\sigma_\theta^2 \rightarrow \infty$ are:

$$\hat{\mu}_P \rightarrow \frac{\gamma_\theta + \gamma_q \lambda}{1 + \gamma_q \lambda} \theta \text{ and } \hat{\sigma}_P^2 \rightarrow \frac{\sigma_\varepsilon^2}{[\gamma_\theta + \gamma_q \lambda]^2}$$

3.2.3 The effect with no risk aversion

When studying the impact of prior uncertainty on learning and on the market equilibrium, we remarked that there is a direct and an indirect effect,

the latter through the posterior variance. We also noted that the indirect effect acts through $(1 + a_L \hat{\sigma}_\theta^2)$. When there is no risk aversion, i.e. $a_L \rightarrow 0$, this indirect effect disappears and the impact of prior uncertainty on learning and on the equilibrium becomes more straightforward. Proposition 3.4 states that, in the risk neutral case, increases in supply uncertainty makes the posterior mean approach the prior mean while diverging from the true value. This implies that the mean equilibrium price increases if the prior is overvalued, decreases if it is undervalued and stays the same if the prior is unbiased.

Proposition 3.4. *If $a_L \rightarrow 0$ then in the limit case*

$$\frac{\partial a}{\partial \sigma_\varepsilon^2} > 0 ; \frac{\partial b}{\partial \sigma_\varepsilon^2} < 0$$

$$\frac{\partial \hat{\mu}_\theta}{\partial \sigma_\varepsilon^2} > 0 \Leftrightarrow \mu_\theta > \theta ; \frac{\partial \hat{\mu}_\theta}{\partial \sigma_\varepsilon^2} = 0 \Leftrightarrow \mu_\theta = \theta ; \frac{\partial \hat{\mu}_\theta}{\partial \sigma_\varepsilon^2} < 0 \Leftrightarrow \mu_\theta < \theta$$

$$\frac{\partial \hat{\mu}_P}{\partial \sigma_\varepsilon^2} > 0 \Leftrightarrow \mu_\theta > \theta ; \frac{\partial \hat{\mu}_P}{\partial \sigma_\varepsilon^2} = 0 \Leftrightarrow \mu_\theta = \theta ; \frac{\partial \hat{\mu}_P}{\partial \sigma_\varepsilon^2} < 0 \Leftrightarrow \mu_\theta < \theta$$

Proof. When $a_L \rightarrow 0$, a and b and $\hat{\mu}_P$ are given by (39), (40) and (41). Deriving these expressions in order of σ_ε^2 the results above are obtained. \square

From Proposition 3.4 we conclude that supply uncertainty impairs learning, because the posterior variance increases and the posterior mean moves away from the true quality value.

Proposition 3.5 states that, in the risk neutral case, increases in prior variance makes the posterior mean approach the true value while diverging from the prior mean. This implies that the mean equilibrium price decreases if the prior is overvalued, increases if it is undervalued and stays the same if the prior is unbiased.

Proposition 3.5. *If $a_L \rightarrow 0$ then in the limit case*

$$\frac{\partial a}{\partial \sigma_\theta^2} < 0 ; \frac{\partial b}{\partial \sigma_\theta^2} > 0$$

$$\frac{\partial \hat{\mu}_\theta}{\partial \sigma_\theta^2} > 0 \Leftrightarrow \mu_\theta < \theta ; \frac{\partial \hat{\mu}_\theta}{\partial \sigma_\theta^2} = 0 \Leftrightarrow \mu_\theta = \theta ; \frac{\partial \hat{\mu}_\theta}{\partial \sigma_\theta^2} < 0 \Leftrightarrow \mu_\theta > \theta$$

$$\frac{\partial \hat{\mu}_P}{\partial \sigma_\theta^2} > 0 \Leftrightarrow \mu_\theta < \theta ; \frac{\partial \hat{\mu}_P}{\partial \sigma_\theta^2} = 0 \Leftrightarrow \mu_\theta = \theta ; \frac{\partial \hat{\mu}_P}{\partial \sigma_\theta^2} < 0 \Leftrightarrow \mu_\theta > \theta$$

Proof. With $a_L \rightarrow 0$, a and b and $\hat{\mu}_P$ are given by (39), (40) and (41). Deriving these expressions in order of σ_θ^2 the results above are obtained. \square

From Proposition 3.5 we conclude that prior variance has contradictory effects upon learning. On one hand, the posterior mean approaches the true quality value. But on the other hand the posterior variance increases.

4 Final remarks

This model, with production and asymmetric information in a competitive market, explored the role of risk aversion and uncertainty on the learning process and on the equilibrium price and output. However, risk aversion was only relevant to the behavior of the learning buyer, not to the decisions on the informed buyer or the firm.

It would be interesting to explore a model where risk aversion affects the decisions of the firm, not only in a competitive market but also in a monopolistic market. The monopolistic case may also provide the framework to study strategic dynamic behavior, involving experimentation by some set of agents.

References

- W.D. Diamond and R.E. Verrecchia. Information Aggregation in a Noisy Rational Expectations Economy. *J. Financ. Econ.*, 9(3):221–235, 1981.
- 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.
- M.F. Hellwig. On the aggregation of Information in Competitive Markets. *J. Econ. Theory*, 22(3):477–498, 1980.
- 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.
- A.S. Kyle. Continuous Auctions and Insider Trading. *Econometrica*, 53(6):1315–1335, 1985.
- L.J. Mirman, E.M. Salgueiro, and M. Santugini. Learning in a Perfectly Competitive Market. CIRPEE working paper 1423, 2014.

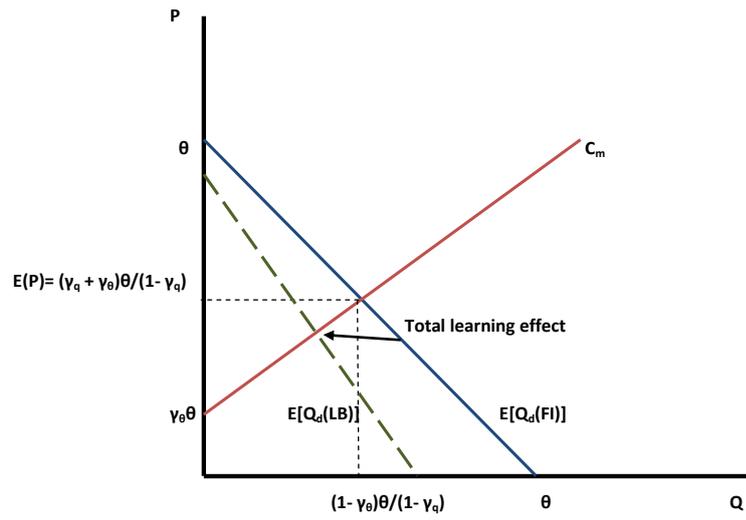


Figure 1: Full information and Bayesian Learning equilibrium

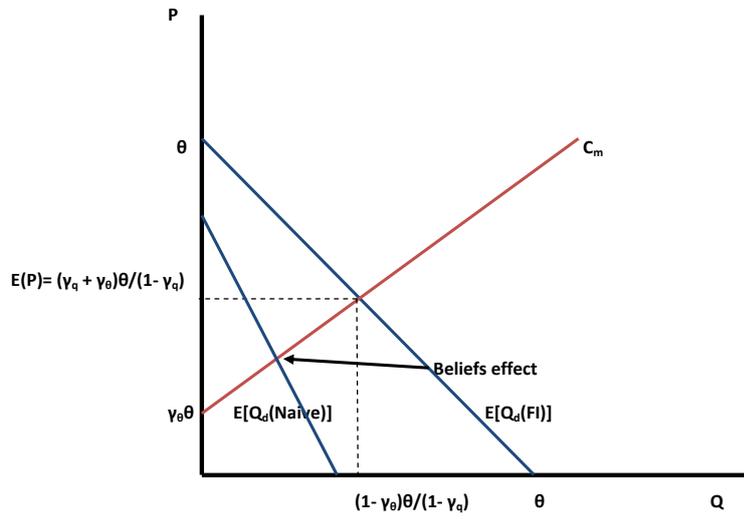


Figure 2: Full information and Naive (undervalued) equilibrium

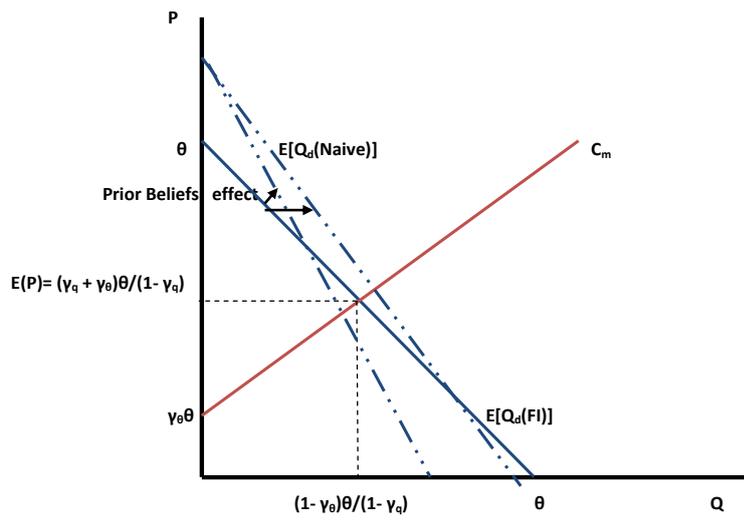


Figure 3: Full information and Naive (overvalued) equilibrium

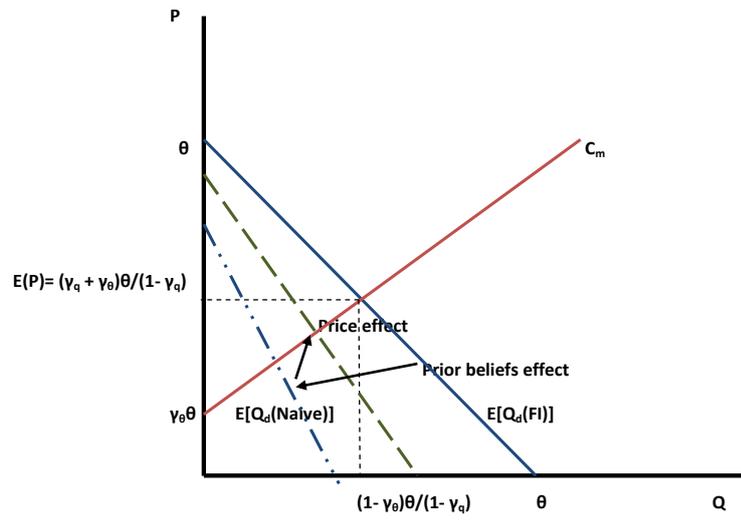


Figure 4: Full information, Naive (undervalued) and Bayesian Learning equilibrium

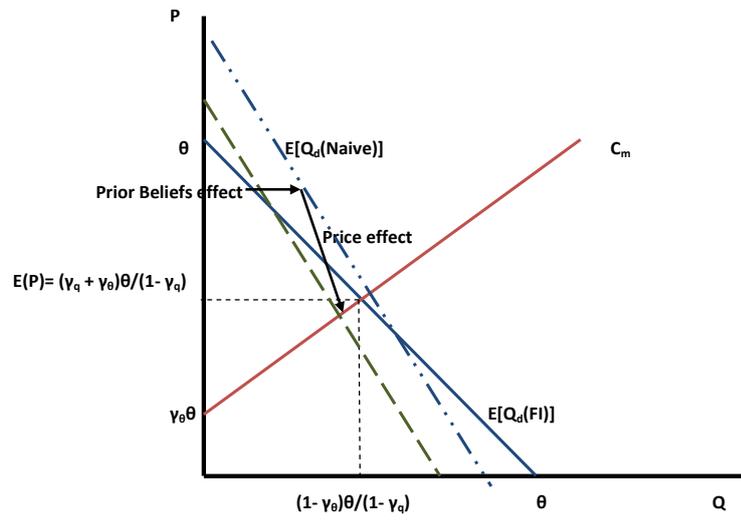


Figure 5: Full information, Naive (overvalued) and Bayesian Learning equilibrium