7

Dynamic Information Aggregation

In chapters 1–5 we considered static models. However, information revelation has important dynamic implications as we saw in chapter 6 when considering social learning. In chapter 6 we considered models of pure informational externalities. A first issue is whether the results obtained are robust when considering models with payoff externalities, as most market models. A second issue to explore is whether there is dynamic information aggregation when in a corresponding static market there would not be. We saw in chapter 1 that simple market mechanisms, like Cournot, need not aggregate information. However, repeated market interaction may make the economy settle in a full- (or shared-) information equilibrium. In this case we may be interested in knowing how fast this convergence occurs and how learning and equilibrium behavior interact.

This chapter studies these issues in the context of dynamic markets which are repeated interaction extensions of the basic Cournot static model with a continuum of firms of chapter 1. Section 7.1 introduces the topic by putting the problems and literature in perspective. Section 7.2 presents a basic model with demand uncertainty and uninformed firms. Section 7.3 explores market dynamics with asymmetric information. Section 7.4 considers a variation of the model with cost uncertainty that displays slow learning and convergence to full-information equilibria. It provides an important link with social learning models as developed in chapter 6. Many proofs, given their technical nature, are gathered in the appendix to the chapter.

7.1 Rational Expectations, Full-Information Equilibria, and Learning

Simple (static) market mechanisms, like Cournot, need not aggregate information (chapter 1). Fully revealing rational expectations equilibria (FRREE) replicate the outcome of a competitive economy with (shared) full information (chapter 3). In the latter case the question is how agents form rational expectations. In some situations there is a game (typically in complex strategies like demand or supply functions) that implements the FRREE. More often, the FRREE is not implementable.

The issue then is whether FRREE are dynamically implementable with simple market mechanisms or how to attain a full-information equilibrium (FIE) with competitive dynamics. Repeated interaction in the market place provides an answer to both the lack of information aggregation with simple market mechanisms and the formation of (fully revealing) rational expectations. Agents process repeated observations of public market data and learn about the relevant uncertainty (adjusting their beliefs in response to observations).

The study of convergence to a full-information equilibrium with repeated interaction is important since it provides a foundation to competitive equilibria under private information. Indeed, the competitive model with full information will be approximately right even in private-information environments if repeated interaction in the market place resolves the uncertainty.

We consider in turn the following issues: The distinction between learning about an equilibrium, or how to form equilibrium expectations, and learning within an equilibrium, where expectations are equilibrium expectations always; the distinction between learning about an unknown parameter and convergence to a full-information equilibrium; and finally the characterization of the speed of learning and the rate of convergence to a full-information equilibrium.

7.1.1 Learning in Equilibrium and Learning about an Equilibrium

A distinction has to be made between learning *within* an (or *in*) equilibrium and learning *about* an equilibrium; in other words, between learning in rational expectations and learning about rational expectations. Furthermore, the learning procedure can be Bayesian or of another type. When learning in equilibrium, agents make the most efficient use of their information and therefore the procedure must be Bayesian. Expectations are "correct" all along, in the sense that they are based in a consistent and correct model of the economy, and the important issue is the revelation of information through time. When learning about an equilibrium the issue at stake is how agents come to have the beliefs associated with the equilibrium when at the start initial beliefs do not correspond to what the model predicts will occur. In this case agents may learn about an equilibrium in a Bayesian way or using other methods like least squares, for example.

To fix ideas consider an economy with an unknown parameter θ and agents with private information about θ. If the joint information reveals θ, then the shared-information equilibrium (SIE) is just the full-information equilibrium (FIE). In models of learning *in* rational expectations, agents have a correctly specified model of the economy and of the learning process, update their beliefs, and take actions accordingly. The issue studied is then information revelation (learning about θ) through prices and convergence to an SIE or an FIE. This limit equilibrium is sometimes called the (stationary) REE although expectations are rational all along. In order to prevent any misunderstanding the limit equilibrium will be refereed to as SIE or FIE.

Equilibrium learning tends to yield at least convergence of beliefs in the presence of i.i.d. shocks. Bayesian posterior beliefs converge with probability 1 (this follows from the martingale convergence theorem because posterior beliefs have the martingale property and are bounded). Convergence to the truth obtains under some regularity conditions as we will see below. (See section 10.3.3 for more details.) In economic models with equilibrium learning this fact is then used to show that the unknown parameter (or agents' types) are learned and that convergence to a limit equilibrium also obtains (see Townsend 1978; Frydman 1982; Blume and Easley 1984, 1998; Feldman 1987; Bray and Kreps 1988). However, "learning" does not imply that "prediction" is possible. That is, that the prediction of the future path of the process given its history up to date t converges to the correct conditional distribution as t grows. This is what matters for convergence to a limit equilibrium in a dynamic model. A distinction must therefore be made between learning the unknown parameter, or the types of the other agents, and the sequence of equilibria converging to a limit equilibrium (SIE or FIE). Blume and Easley (1998) point out that the learning of types is not sufficient to yield convergence to a limit equilibrium but they do not provide an example in an economic setting. We will study in sections 7.2 and 7.3 the relationship between learning θ and converging to an FIE in the context of equilibrium learning. We will see that learning θ is neither necessary nor sufficient to converge to an FIE.

In models of learning *about* a rational expectations equilibrium, agents maintain incorrect hypotheses in the face of the evolution of the economy and either are Bayesians or use "reasonable" behavioral updating procedures, like adaptive rules or least squares estimation, for example. It has been found that in this case convergence to REE is not guaranteed although positive results are available by assuming a certain degree of coordination of the forecasting strategies of agents (Blume and Easley 1982; Bray 1982; Frydman 1982; Jordan 1985; Bray and Savin 1986; Fourgeaud et al. 1986; Marcet and Sargent 1988, 1989; Woodford 1990; Evans and Honkapohja 2001). In Brock and Hommes (1997) global complicated dynamics arise in a model where agents adapt their beliefs by choosing among a limited set of predictors. What if agents who use behavioral rules and simple heuristics to form beliefs coexist with Bayesian agents in the market? A traditional argument has been that agents that do not predict as accurately as others will be driven out of the market (Alchian 1950; Friedman 1953). The reason is that the behavioral agents will end up losing money and dissipating their wealth. Several authors have questioned this idea pointing out that agents maximize expected utility and not wealth accumulation, and the result may be that agents with incorrect beliefs drive agents with correct beliefs out of the market (see, for example, De Long et al. (1989, 1990) for behavioral traders and Kyle and Wang (1997) and Hirshleifer and Luo (2001) for the effects of overconfidence). However, Sandroni (2000) has shown that this cannot happen in a dynamic general equilibrium model. The basic intuition is that agents allocate more wealth to events they think are more likely to occur and only agents with correct beliefs end up on average being right. Sandroni (2005) goes further and compares Bayesian with behavioral belief-updating agents in an environment where agents try to learn the structure of the economy. The economy is given by a standard dynamic asset pricing model (a Lucas tree (Lucas 1978)). In the end prices reflect the empirical models used by agents with predictive properties conducive to wealth accumulation. Sandroni distinguishes between a structure of the economy which is learnable from another which is not. If the structure is learnable (e.g., when the relevant data set is unlimited), then only Bayesian agents or agents who forecast very closely to them survive (the intuition is as before: only agents with Bayesian prediction end up allocating wealth to events that actually occur). If the structure is not learnable (e.g., when there are recurrent regime switches which make past data obsolete), Sandroni finds in fact that only Bayesian agents survive (that is, Bayesian agents drive other agents out of the market according to the probability distribution generated by the data). The conclusion is that asset prices are determined under the Bayesian paradigm even in the presence of non-Bayesian agents.

In the Bayesian setting agents have priors over possible sequences of market prices and update at each date. If the prior does not coincide with the objective distribution on price sequences generated by the behavior of agents, then there is learning about REE. In general, a crucial element to obtain convergence to a limit equilibrium with correct beliefs with respect to the underlying true economy is the a priori assumed coordination of expectations of agents. Indeed, Bayesian learning does not have a lot of content per se, all depends on the assumptions about the priors of the agents. Convergence to the truth requires that prior beliefs be consistent with the true model. A Bayesian agent will never learn the truth if he puts zero weight on the true value of the parameter or model to start with. In a Bayesian statistical setting a *sufficient* condition for conditional beliefs about the future given the past to converge to correct conditional beliefs is that the true process given any parameter be absolutely continuous with respect to the prior predicted distribution on observations (sample paths). This condition basically means that if the prior assigns probability zero to an event, then the true model must also assign probability zero to the event (see Blackwell and Dubins 1962; Blume and Easley 1998). The condition is very strong since it requires a countable parameter space. It is at the base, for example, of the convergence results to Nash equilibrium in games of Kalai and Lehrer (1993) and Nyarko (1997, 1998). Agents need not have the correct model but their priors (on the entire space of paths) need to have a "grain of truth" (the absolute continuity condition) for convergence to Nash equilibrium behavior to obtain.[1] Nyarko (1997) analyzes the model of Townsend (1978) and Feldman (1987), with which we deal in section 7.3, when players form a Bayesian hierarchy of beliefs about the behavior of other players and concludes that, apart from the mutual absolute continuity condition on the beliefs of the players, a contraction property of the best response maps is needed to obtain convergence to the Nash equilibrium of the model corresponding to the true fundamentals.

There is still another approach to building foundations for an REE which is sometimes called "eductive stability." For an equilibrium to be eductively stable it must be the outcome only of the rationality of agents and common knowledge

[1] See also the discussion at the end of section 10.3.3.

about payoffs. That is, the equilibrium itself need not be common knowledge. This is an insight provided by the rationalizability literature in games (see section 10.4.1). The question then is how the agents coordinate expectations in the equilibrium. Eductive stability basically means that the equilibrium must be the outcome of the iterated elimination of dominated strategies. If this process has a unique outcome, the game is then called dominance solvable (see section 10.4.1). This approach builds on introspection, equilibrium expectations being pinned down in a thought process in the minds of agents in virtual time. Not all REE are eductively stable. Indeed, a contraction condition on best-reply maps is also typically needed to show dominance solvability or eductive stability. (See Guesnerie (2006) for a collection of papers on the issue and section 8.4.2 for an example of a coordination game where conditions are found for the game to be dominance solvable.)

7.1.2 Speed of Learning and Rate of Convergence to Limit Equilibria

Taking for granted that a learning process converges to a limit (fully revealing/shared-information) equilibrium it is of the utmost importance to know how fast and what factors affect the speed of convergence. It is not of much use for practical purposes to show convergence if it is not known whether this will happen quickly or take a long time, when the underlying conditions of the economy will have changed and the parameters learned may well be irrelevant by then. In this sense "slow" convergence may mean in practice no convergence. We would also like to know how structural market conditions—like the nature of uncertainty and its relation to market observables (namely, prices), the potential persistence of shocks (autocorrelation), the degree of asymmetric information, and the precision of private signals—affect the speed of convergence and in what direction.

A relevant distinction is whether the situation is one of "learning from others," as in chapter 6, or whether the relevant uncertainty impacts directly the public statistics that agents observe. The first situation may arise, for example, with cost uncertainty while the second with demand uncertainty. We have already seen in chapter 6 a robust result of slow learning from others. A question is how the robust result of slow learning from others may be modified in a market environment with payoff externalities.

It is also worth exploring the role of persistence of shocks since it may be presumed that some forms of autocorrelation may slow the speed of learning. This is the case for the generalized least squares (GLS) estimator, for example, in classical econometric analysis when disturbances are positively autocorrelated according to a stationary AR(1) process. Obviously, even in the context of a linear model where GLS is optimal, a dynamic market will have feedback effects which may drive us away from the classical analysis and results. Will the presence of autocorrelation impair the capacity of agents to "learn" θ and predict market prices, and to agree on price estimates so that a (fully revealing or shared-information) equilibrium obtains?

Rate-of-convergence results are usually difficult to obtain. Work by Jordan (1992a) on a class of Bayesian myopic learning processes (without noise) establishes an exponential rate of convergence to Nash equilibria for finite normal form games. In models of learning about an equilibrium partial results, relying sometimes on simulations, have been obtained by Bray and Savin (1986), Fourgeaud et al. (1986), Jordan (1992b), and Marcet and Sargent (1992). We will provide in the rest of the chapter results on the speed of learning and rate of convergence to FIE in different scenarios in the context of equilibrium learning.

In this chapter, as in the rest of the book, we will concentrate attention on equilibrium learning. We start by considering the case where θ is the unknown demand intercept in a Cournot market similar to the one presented in chapter 1. Section 7.2 deals with learning and convergence to FIE when firms are uninformed and section 7.3 when there is asymmetric information about θ. Both sections are based on Jun and Vives (1996). Section 7.4 will consider the implications of uncertainty about costs and is based on Vives (1993). In all cases the speed of learning and the rate of convergence to full-information equilibria are characterized.[2]

7.2 Learning and Convergence to a Full-Information Equilibrium with Uninformed Firms

We consider here the classical linear partial equilibrium model first studied by Muth (1961) in his seminal work about rational expectations. This is an infinitely repeated version of the continuum-of-firms Cournot model of chapter 1 with a random demand intercept subject to potentially persistent shocks and uninformed firms. Asymmetric information is introduced in the next section.

Consider an infinite-horizon market, $t = 1, 2, \dots$, with a continuum of firms, indexed by $i \in [0, 1]$, endowed with the Lebesgue measure, and producing a homogeneous product. Firms are risk neutral and have identical quadratic cost functions: $C(x_i) = \frac{1}{2}\lambda x_i^2$, with $\lambda > 0$. The inverse demand for the product in period t is random and is given by

$$p_t = \theta + u_t - \beta x_t,$$

where θ is the unknown intercept of demand, u_t is a period-specific shock, $\beta > 0$, and x_t is the average (per capita) supply in period t, $x_t = \int_0^1 x_{it}\,\mathrm{d}i$.

The temporary shocks $\{u_t\}_{t=1}^{\infty}$ are potentially persistent and generated by an AR(1) process: $u_t = \varsigma u_{t-1} + \eta_t$, with $\{\eta_t\}_{t=1}^{\infty}$ a white noise normal process.[3] When $\varsigma = 1$ the process is a random walk. We do not impose stationarity of

[2] See Bisin et al. (2006) for the study of rational expectations equilibria in dynamic economies with local interactions (in such environments the interaction capacity of an individual depends on his position on the established network of relationships).

[3] Subject to the initial condition $u_t = 0$ for $t \leqslant 0$. The results that we will derive hold for a general AR(∞) process (Jun and Vives 1996).

any sort. When $|\varsigma| < 1$ the AR(1) process is asymptotically independent and asymptotically stationary (but not stationary). Whenever the stochastic process $\{u_t\}$ is asymptotically stationary the initial conditions will not affect the asymptotic results in our model. When $|\varsigma| \geqslant 1$, the process is neither asymptotically stationary nor asymptotically independent.[4]

Firms are uninformed about θ but have a common prior which is normally distributed: $\theta \sim N(\bar{\theta}, \sigma_\theta^2)$. The distributional assumptions, including the parameters of the distributions, as well as the parameters λ and β, are common knowledge among the agents in the economy.

Firms set output levels in every period to maximize expected discounted profits by taking into account the information that past prices carry about θ. Given the structure of the economy, the market equilibrium is just a competitive equilibrium and myopic behavior is optimal. Firms have no scope for collusion because they are negligible and only prices are observable. Discounting has no influence on the equilibrium. If firms were to maximize the long-run average of profits, then whatever would happen in a finite number of periods would be irrelevant. Profit maximization leads any firm in period t to choose the output level according to

$$x_t = \lambda^{-1} E[p_t \mid I_t],$$

where $I_t = \{p^{t-1}\}$, $p^{t-1} = \{p_1, \dots, p_{t-1}\}$, is the information the firm has in period t (past prices).

Given that firms have common expectations, and consequently all firms produce the same output x_t, the information content of the price p_t is equivalent to $p_t + \beta x_t = \theta + u_t$. Let $z_t = \theta + u_t$. The problem of a firm in period t is to predict the market price $p_t = \theta + u_t - \beta x_t$, or $z_t = \theta + u_t$, with the information z^{t-1}, $z^{t-1} = \{z_1, \dots, z_{t-1}\}$. The equilibrium output

$$x_t = (\lambda + \beta)^{-1} E[z_t \mid z^{t-1}]$$

follows from the profit maximization condition $x_t = \lambda^{-1} E[p_t \mid I_t]$.

Denote by θ_t the public belief or expectation of θ conditional on public information. That is, $\theta_t = E[\theta \mid z^t]$. Given normality, the random variable θ_t is a sufficient statistic of the information in prices about the unknown θ. Let $\Delta z_t = z_t - \varsigma z_{t-1}$ with $z_0 = 0$ for $t \geqslant 1$. The random variable Δz_t equals $(1 - \varsigma)\theta + \eta_t$ and represents the new information about θ in the current price p_t. It is then clear that $\theta_t = E[\theta \mid z^t] = E[\theta \mid \Delta z^t]$, where $\Delta z^t = \{\Delta z_1, \dots, \Delta z_t\}$. Given that $z_t = \Delta z_t + \varsigma z_{t-1}$ it follows that

$$x_t = (\lambda + \beta)^{-1} E[z_t \mid z^{t-1}] = (\lambda + \beta)^{-1}((1 - \varsigma)\theta_{t-1} + \varsigma z_{t-1}).$$

The equilibrium is unique.[5]

[4] In our context we say that the process $\{u_t\}$ is *asymptotically independent* if $\text{cov}[u_t, u_{t+h}] \to 0$ as $h \to \infty$. The process is *asymptotically stationary* if $\text{cov}[u_t, u_{t+h}]$ does not depend on t in the limit as $t \to \infty$ (see Spanos 1986, chapter 8, p. 153). Under our assumptions, $\text{cov}[u_t, u_{t+h}] = t\sigma_\eta^2$ if $|\varsigma| = 1$, and $\text{cov}[u_t, u_{t+h}] = \sigma_\eta^2 \varsigma^h (1 - \varsigma^{2t})/(1 - \varsigma^2)$ otherwise, $h \geqslant 0$.

[5] It is worth noting that if agents are risk averse in the Muth model we may have potential nonexistence, unique or multiple linear equilibria (see McCafferty and Driskill 1980). In section 9.1.3.1 we consider a related model with risk-averse agents where there may be three linear equilibria.

If firms were to know θ, then in period t they would have access to past demand disturbances u^{t-1} because they observe past prices p^{t-1}, or equivalently z^{t-1} with $z_t = \theta + u_t$. Furthermore, $E[z_t \mid \theta, u^{t-1}] = \theta + E[u_t \mid u^{t-1}] = \theta + \varsigma u_{t-1}$. It follows that the full-information equilibrium output (denoted by "f") is

$$x_t^{\mathrm{f}} = (\lambda + \beta)^{-1}(\theta + \varsigma u_{t-1}).$$

When $\varsigma = 0$ the FIE can be thought of as the REE of a static market. Indeed, in the market with inverse demand $p = \theta - \beta x$, with θ random, the competitive equilibrium $x^{\mathrm{f}} = (\lambda + \beta)^{-1}\theta$ and $p^{\mathrm{f}} = (\lambda + \beta)^{-1}\lambda\theta$ is an FRREE which is implementable in supply functions (just the competitive supplies, $X(p) = p/\lambda$, see chapter 3).

The difference of the FIE with the market equilibrium output x_t (noting that $z_{t-1} = \theta + u_{t-1}$) is easily seen to be

$$x_t - x_t^{\mathrm{f}} = (\lambda + \beta)^{-1}(1 - \varsigma)(\theta_{t-1} - \theta) \quad \text{for } t \geqslant 2.$$

We want to examine the asymptotic behavior of $x_t - x_t^{\mathrm{f}}$, that is, *convergence* of the market equilibrium to the FIE, and of $\theta_t - \theta$ or the difference between public information θ_t about the unknown parameter θ and the true value of θ, that is, *learning* by firms about the unknown θ.

We will see that although it may difficult, or impossible, to estimate the only unknown parameter θ in the economy, yet convergence to an FIE obtains. Learning θ and convergence to an FIE are thus not equivalent in the presence of persistent shocks.

Before stating our convergence results (the first in proposition 7.1) we will recall again some measures of speed of convergence. We will say that the sequence (of real numbers) $\{b_t\}$ is of the *order* t^{υ}, with υ a real number, whenever $t^{-\upsilon} b_t \xrightarrow[t]{} k$ for some nonzero constant k. Denote by $\xrightarrow{\mathrm{L}}$ convergence in law (distribution). We say that the sequence of random variables $\{x_t\}$ converges to x at the rate $t^{-\kappa}$, for $\kappa > 0$, if $x_t \xrightarrow[t]{} x$ (a.s. or in mean square) and $t^{\kappa}(x_t - x) \xrightarrow{\mathrm{L}} N(0, \mathrm{AV})$ for some positive constant AV. The asymptotic variance AV represents a refined measure of speed for a given convergence rate. A lower AV means faster convergence.[6]

Proposition 7.1 (Jun and Vives 1996).

(i) *If $\varsigma \neq 1$, then θ_t converges to θ (a.s. and in mean square) at the rate $1/\sqrt{t}$ with asymptotic variance $\sigma_\eta^2/(1-\varsigma)^2$, and $x_t - x_t^{\mathrm{f}} \to 0$ at the rate $1/\sqrt{t}$ with asymptotic variance $\sigma_\eta^2/(\lambda+\beta)^2$.*

(ii) *If $\varsigma = 1$, except for the first period, no information about θ can be inferred from prices (with the precision of θ_t constant at $\tau_\theta + \tau_\eta$) but the market equilibrium coincides with the FIE ($x_t - x_t^{\mathrm{f}} = 0$ for $t \geqslant 2$).*

---

[6] See sections 10.3.1 and 10.3.2 for a fuller development of the definitions and relationships between convergence concepts.

Proof. (i) Let $\Delta z_t = z_t - \varsigma z_{t-1}$ for $t \geqslant 1$ as before. Then $\Delta z_t = (1-\varsigma)\theta + \eta_t$ and $\theta_t = E[\theta \mid z^t] = E[\theta \mid \Delta z^t]$, where $\Delta z^t = \{\Delta z_1, \dots, \Delta z_t\}$. It follows from normal theory (see section 10.2.1) that $E[\theta \mid \Delta z^t] = (\tau_\theta \bar{\theta} + \tau_\eta z_1 + \tau_\eta (1-\varsigma)\sum_{k=2}^{t} \Delta z_k)/\tau_t$, where $\tau_t = \tau_\theta + \tau_\eta(1 + (1-\varsigma)^2(t-1))$ is the informativeness (precision) of the public statistic θ_t in the estimation of θ ($\tau_t = (\operatorname{var}[\theta \mid \theta_t])^{-1}$). The public precision τ_t is of the order of t since $\varsigma \neq 1$ and consequently tends to ∞ with t. It follows that $\theta_t \to \theta$ (a.s. and in mean square). Furthermore, $\sqrt{t}(\theta_t - \theta) \xrightarrow{\text{L}} N(0, \sigma_\eta^2/(1-\varsigma)^2)$ because $\operatorname{var}[\theta - \theta_t] = \operatorname{var}[\theta \mid \theta_t] = \tau_t^{-1}$ and $t^{-1}\tau_t \to \tau_\eta(1-\varsigma)^2 = (\sigma_\eta^2/(1-\varsigma)^2)^{-1}$. (This also shows that $\theta_t \to \theta$ in mean square.) The result for $x_t - x_t^{\text{f}}$ follows from $x_t - x_t^{\text{f}} = (\lambda + \beta)^{-1}(1-\varsigma)(\theta_{t-1} - \theta)$ for $t \geqslant 2$.

(ii) When $\varsigma = 1$, $\Delta z_t = \eta_t$ for $t \geqslant 2$ and prices contain no information about θ. For $t = 1$, $\Delta z_1 = \theta + \eta_1$ and $\tau_1 = \tau_\theta + \tau_\eta$. Furthermore, for $t \geqslant 2$, $x_t - x_t^{\text{f}} = (\lambda + \beta)^{-1}(1-\varsigma)(\theta_{t-1} - \theta) = 0$ when $\varsigma = 1$. □

Learning θ and converging to an FIE are different phenomena. While learning θ may be very slow (because ς is close to 1, or even impossible if $\varsigma = 1$), convergence to the FIE is unaffected. Indeed, the rate of convergence to the FIE is independent of ς if $\varsigma \neq 1$. If $\varsigma = 1$, then convergence is immediate. This highlights in a very stark form the fact that learning θ is not necessary for convergence to an FIE. It is worth noting that with i.i.d. shocks ($\varsigma = 0$), $x_t - x_t^{\text{f}} = (\lambda + \beta)^{-1}(\theta_{t-1} - \theta)$, and learning θ is equivalent to converging to the FIE.

When $\varsigma \neq 1$ the new information in p_t about θ is $\Delta z_t = (1-\varsigma)\theta + \eta_t$ and firms face a classical ordinary least squares (OLS) estimation problem of θ. However, the asymptotic variance of the public belief $\theta_t \to \infty$ as ς tends to 1. Indeed, when $\varsigma = 1$, θ cannot be estimated since prices give no information on θ (the price innovation is pure noise, $\Delta z_t = \eta_t$ for $t \geqslant 2$). Nevertheless, firms are not interested in estimating θ but the market price p_t, or, equivalently, the statistic $z_t = \theta + u_t$. When $\varsigma = 1$, $\{z_t\}$ follows a martingale and the estimation of θ is irrelevant to the prediction of z_t: $E[z_t \mid z^{t-1}] = z_{t-1}$. Therefore, the market and the FIE outcomes coincide. When $\varsigma \neq 1$, the market equilibrium will converge to the FIE if (and only if) the expectations $E[z_t \mid z^{t-1}] - E[z_t \mid \theta, u^{t-1}]$ merge as t grows. This will happen if firms learn θ. As ς approaches 1 it becomes more and more difficult to estimate θ (the asymptotic variance of $\theta_t - \theta$ grows) but it is also less and less important to do so to predict the current price (since $E[z_t \mid z^{t-1}] = (1-\varsigma)\theta_{t-1} + \varsigma z_{t-1}$). This is why the asymptotic variance of $x_t - x_t^{\text{f}}$ is independent of ς.

It is worth remarking that the asymptotic variance of $x_t - x_t^{\text{f}}$ increases with σ_η^2 and decreases with λ and β. We have that the smaller the slopes of marginal cost or inverse demand, the larger the response of agents to public information and the slower the convergence rate. This means that convergence in a market with uniformly lower marginal costs (λ lower) or of larger size (β lower) will be slower.

In summary, we have characterized the speed of learning and convergence to FIE in the classical Muth rational expectations model. Convergence to the FIE happens fast (at the rate of $1/\sqrt{t}$) independently of whether learning the unknown parameter is slow or fast. We have seen in this model with uninformed agents that learning the unknown parameter and converging to a full-information equilibrium are not equivalent. We will see in section 7.3 that provided there is no positive mass of perfectly informed agents the results obtained so far will go through. That is, private information will have no asymptotic effect in the limit although equilibrium behavior will have a real feedback in learning. When there is a positive mass of perfectly informed agents, private information will have an asymptotic effect and results may change. In particular, learning the unknown parameter θ may not be *sufficient* for convergence to the FIE to obtain.

7.3 Market Dynamics with Asymmetric Information

Townsend (1978) and Feldman (1987) studied convergence to FIE in a classical linear partial equilibrium model with asymmetric information and i.i.d. shocks to demand. Their analysis is extended here to persistent shocks.[7]

We consider the same model as in section 7.2 but with firms asymmetrically informed about θ. Firm i, $i \in [0,1]$, is endowed with a private signal about θ, $s_i = \theta + \varepsilon_i$, where ε_i is an error term. We assume that θ, $\{\varepsilon_i\}$, and $\{\eta_t\}$ are independently (across agents and across time) normally distributed random variables: $\theta \sim N(\bar{\theta}, \sigma_\theta^2)$, $\eta_t \sim N(0, \sigma_\eta^2)$, and $\varepsilon_i \sim N(\bar{\theta}, \sigma_{\varepsilon_i}^2)$. The precision of the signals is given by a (measurable) function $T_\varepsilon : [0,1] \to \mathbb{R}_+ \cup \{\infty\}$, where $T_\varepsilon(i) = \sigma_{\varepsilon_i}^{-2} (\equiv \tau_{\varepsilon_i})$ is the value of this function for $i \in [0,1]$. Let $\mu \geqslant 0$ be the mass of agents with perfectly informative signals (that is, signals of infinite precision).

As in the previous chapters we adopt the "strong law of large numbers" convention about the average of a continuum of independent random variables. In particular, we write $\int_0^1 s_i \, \mathrm{d}i = \int_0^1 (\theta + \varepsilon_i) \, \mathrm{d}i = \theta + \int_0^1 \varepsilon_i \, \mathrm{d}i = \theta$ (a.s.) when signals have uniformly bounded variances.[8] If there is a positive mass of firms with signals with precisions uniformly bounded away from zero, θ is revealed by forming an average of the signals of agents with precision bounded away from zero and ignoring the signals of other agents. In this case the shared-information equilibrium coincides with the full-information equilibrium.[9] Firms know all the parameters of the model, including the parameters of the distributions, except for θ.

---

[7] See Nyarko (1997) for an analysis of the model and the conditions for convergence to an FIE to obtain when Bayesian firms are free to form beliefs about the behavior of others.

[8] See section 10.3.1 for a justification of the convention.

[9] As before, when $\varsigma = 0$ the FIE can be thought as the FRREE of a static market with inverse demand $p = \theta - \beta x$ with θ random. This FRREE is implementable in supply functions with the competitive supplies, $X(p, s_i) = p/\lambda$ for any s_i.

The dynamic game is as in section 7.2 but now firms start with their private estimates of the uncertain intercept of demand. With this information each firm sets an output level, and a market-clearing price obtains. Upon observing the first-period price, firms set the period 2 output and so on. The equilibrium of the market is just the (perfect) Bayesian equilibrium (PBE) of the dynamic game.[10] Given that agents are negligible and that the action of a single firm does not affect public information (the market price) the equilibrium has a very simple structure. It just involves a sequence of one-shot Bayes–Nash equilibria with evolving information structures according to the refinement of public information.

Profit maximization leads firm i in period t to choose the output level $x_{it} = \lambda^{-1}E[p_t \mid I_{it}]$, where $I_{it} = \{s_i, p^{t-1}\}$ is the information set of the firm. In order to estimate the market price $p_t = \theta + u_t - \beta x_t$, a firm in period t needs to estimate $\theta + u_t$ and x_t. Under private information a firm does not know the average market output x_t since it depends on the aggregate private information received by all the firms (which in our model with a continuum of firms is θ).

It is possible to show, by following a standard procedure (detailed in the appendix to this chapter, see proposition 7.6), that there is a unique linear equilibrium in the dynamic game (in fact, the equilibrium is unique in the class of strategies with bounded means and uniformly bounded variances across firms; see exercise 7.3 for a closely related result). In equilibrium

$$p_t = z_t - \beta c_t \theta_{t-1} - \beta(\lambda + \beta)^{-1} \varsigma z_{t-1},$$

with $z_t = \hat{a}_t \theta + u_t$, $\hat{a}_t = 1 - \beta a_t$, where a_t (c_t) is the average responsiveness of firms to private (public) information. The equilibrium parameters are determined in a recursive way.

It is clear that p^t is observationally equivalent to z^t. Let $\Delta z_t = z_t - \varsigma z_{t-1}$ and $\theta_t = E[\theta \mid \Delta z^t]$, where $\Delta z^t = \{\Delta z_1, \dots, \Delta z_t\}$. The random variable Δz_t represents the *new information* about θ in the price p_t. Its informativeness depends on the coefficient $\Delta \hat{a}_t = \hat{a}_t - \varsigma \hat{a}_{t-1}$. We have that $\Delta z_t = \Delta \hat{a}_t \theta + \eta_t$, where, by definition, $z_0 = 0$ and $\hat{a}_0 = 0$.

The discrepancy between the FIE output $x_t^{\mathrm{f}} = (\lambda + \beta)^{-1}(\theta + \varsigma u_{t-1})$ (see section 7.2) and the average output with private information $x_t = a_t \theta + c_t \theta_{t-1} + (\lambda + \beta)^{-1} \varsigma z_{t-1}$ is easily seen to be equal to the difference between the public predictor of θ_{t-1} and θ, weighted by the aggregate responsiveness of agents to public information:

$$x_t - x_t^{\mathrm{f}} = c_t(\theta_{t-1} - \theta).$$

This difference corresponds to the difference between average market expectations about the price, $\int_0^1 E[p_t \mid I_{it}]\,\mathrm{d}i$, and full-information expectations $E[p_t \mid I_t^{\mathrm{f}}]$, $I_t^{\mathrm{f}} = \{\theta, u^{t-1}\}$. Indeed, $\int_0^1 E[p_t \mid I_{it}]\,\mathrm{d}i - E[p_t \mid I_t^{\mathrm{f}}] = \lambda(x_t - x_t^{\mathrm{f}})$.

We now examine whether, and if so how fast, the market equilibrium with private information converges to the FIE benchmark. We show that if there is

[10] See section 10.4.4 for a short introduction to PBE.

no positive mass of *perfectly* informed firms (i.e., $\mu = 0$), then the presence of private information does not affect the convergence results obtained with uninformed firms (proposition 7.1(i) holds and (ii) also holds in a modified way). Private information has no asymptotic effect in the limit although equilibrium behavior will have a real feedback in learning. When $\mu > 0$, private information has an asymptotic effect, although convergence results and rates of convergence are not affected if the process of the shocks is not too explosive ($|\varsigma| < 1 + \lambda/\beta\mu$).

Proposition 7.2 states the result. We will say that the sequence of random variables $\{y_t\}$ converges to 0 *exponentially* if there is a constant $\kappa > 1$ such that $\kappa^t y_t \to 0$ (a.s.).[11] Obviously, this implies that $y_t \to 0$ a.s.

Proposition 7.2 (Jun and Vives 1996). *Let $|\varsigma| < 1 + \lambda/\beta\mu$, then as t grows:*

(i) *If $\varsigma \neq 1$, θ_t converges to θ at the rate $1/\sqrt{t}$ with asymptotic variance $((1-\varsigma)^{-1} + \beta\mu\lambda^{-1})^2\sigma_\eta^2$, and $x_t - x_t^{\mathrm{f}} \to 0$ at the rate $1/\sqrt{t}$ with asymptotic variance $(1-\mu)^2\sigma_\eta^*/(\lambda+\beta)^2$.*

(ii) *If $\varsigma = 1$, $\theta_t - \theta$ converges to a nondegenerate normal distribution with precision no larger than $\tau_\theta + \tau_\eta$, and $x_t - x_t^{\mathrm{f}} \to 0$ exponentially.*

Proof. See the appendix (section 7.6). □

When $\mu = 0$ and $\varsigma \neq 1$ private information does not matter asymptotically. We then obtain the same results as with uninformed firms (for $\mu = 0$, the upper bound on $|\varsigma|$, $1 + \lambda/\beta\mu$, is infinite). With $\varsigma \neq 1$ prices end up revealing θ and therefore the weight imperfectly informed firms put on their private information about θ, a_t, vanishes over time.[12] Consequently, average output is eventually known by the firms and need not be estimated. Therefore, asymptotically (and approximately) to predict p_t only $\theta + u_t$ needs to be estimated, as in the benchmark case with no private information.

As in section 7.2 the degree of serial correlation ς neither affects the rate nor the asymptotic variance of convergence to the FIE. Changes in autocorrelation are asymptotically neutral since they produce two effects that exactly offset each other. An increase in ς toward 1 makes the price statistic θ_t less informative about θ but it also makes it (via an equilibrium mechanism) less important for agents to estimate θ (and consequently firms put less weight on θ_t, a lower c_t, when choosing their production).

The asymptotic variance of $\theta_t - \theta$ is increasing with μ and with β. It is decreasing with λ. The fact that uninformed agents find prices less precise as an estimator of θ for larger μ may seem surprising. Nevertheless, perfectly informed firms by reacting to their information make the price less sensitive to θ (and the less sensitive to θ, the larger is β and the smaller is λ). The asymptotic variance of $x_t - x_t^{\mathrm{f}}$ decreases with μ. When μ approaches 1 the asymptotic variance of the discrepancy between market output and the FIE tends to 0.

[11] Note that this implies that convergence is at least at an exponential rate.

[12] Whenever $\mu = 0$ or $\varsigma = 1$, $a_t \to (1-\varsigma)\mu/[\lambda + (1-\varsigma)\beta\mu] = 0$.

When $\varsigma = 1$ with uninformed firms prices cannot reveal any information about θ but $x_t - x_t^{\mathrm{f}} = 0$ after the second period. With private information, prices reveal information over time (but less in the aggregate than in the first period with uninformed firms) and there is some discrepancy between x_t and x_t^{f} that vanishes fast. Now, for any $\mu \geqslant 0$ the weight a firm puts on its private information vanishes over time. When $\varsigma = 1$ agents put less and less weight on their private information (and on the price statistic) very fast. The consequence is that the new information about θ in the current price, $\Delta z_t = \Delta\hat{a}_t\theta + \eta_t$, looks like noise (since $\Delta\hat{a}_t \to 0$ very fast) and θ cannot be learned. Nevertheless, this also implies (since $\hat{a}_t \to 1$) that to predict the price in period t very quickly firms need only predict $\theta + u_t$ (that is, θ no longer matters in the estimation of $z_t = \hat{a}_t\theta + u_t$, only $\theta + u_t$) and firms therefore rely on the new information of the last-period price, z_{t-1}. Even though θ is not revealed by prices there is consensus on price expectations given the information of firms. More precisely, $\int_0^1 E[p_t \mid I_{it}]\,\mathrm{d}i - E[p_t \mid I_{it}] \to 0$ (a.s.) and $x_t - x_t^{\mathrm{f}} \to 0$ (a.s.) at an exponential rate.

We now see that when the process of shocks is sufficiently explosive, we then have fast learning and no convergence to the FIE!

Proposition 7.3 (Jun and Vives 1996). *If $\varsigma > 1 + \lambda/\beta\mu$, then as $t \to \infty$:*

(i) *$\theta_t - \theta$ tends exponentially to 0, and*

(ii) *$x_t - x_t^{\mathrm{f}}$ tends to a nondegenerate normal random variable with variance $\sigma_\eta^2(\lambda + \beta)^{-2}(1 - \mu)^2(\hat{\varsigma}^2 - 1)$, where $\hat{\varsigma} = \varsigma(1 + \lambda/\beta\mu)^{-1}$.*

Proof. See the appendix (section 7.6). □

When $\varsigma > 1 + \lambda/\beta\mu$ prices reveal θ very fast (at an exponential rate) but the uninformed firms' reaction to the public information contained in prices (c_t) is so strong and increasing (also growing exponentially) that the discrepancy with the FIE, $c_t(\theta_{t-1} - \theta)$, does not disappear as $t \to \infty$. Even though public information reveals θ very fast, price expectations do not converge to the full-information level because the sensitivity of prices to public information is also growing very fast.

The combination of $\mu > 0$ with a sufficiently explosive $\{u_t\}$ process yields the nonconvergence outcome. Suppose for simplicity that a mass $1 - \mu$ of agents is uninformed. We have that $\Delta\hat{a}_t$ (and a_t, c_t) tend exponentially to $-\infty$ and $\hat{a}_t$ tends exponentially to ∞. This explosive equilibrium is sustained as follows. The price process depends more and more on θ since $\hat{a}_t$ tends exponentially to ∞ (and, in consequence, $\Delta\hat{a}_t$ tends exponentially to $-\infty$). Informed firms put an increasing negative weight on θ ($a_t \to -\infty$) since the prediction of the price innovation z_t depends more and more on θ: $E[z_t \mid \theta, z^{t-1}] = \Delta\hat{a}_t\theta + \varsigma z_{t-1}$. This makes it more and more important for uninformed firms to predict θ in order to predict the market price and consequently they put an increasing negative weight on public information θ_t: $E[z_t \mid z^{t-1}] = \Delta\hat{a}_t\theta_{t-1} + \varsigma z_{t-1}$.[13]

In summary, agents may not learn θ but converge to the FIE, and agents may learn θ but not converge to the FIE. That is, learning the unknown parameter θ is neither a necessary nor a sufficient condition for convergence to the FIE.

It must be pointed out that the phenomena mentioned above happen when usual stability assumptions imposed on the AR process are violated (although, as we have emphasized, the distinction between learning θ and converging to the FIE remains under the stability assumptions). Indeed, when $|\varsigma| \geqslant 1$ the AR(1) process is not even asymptotically stationary and the variance of u_t is unbounded.[14]

The normality assumption implies in general that prices and quantities, as well as θ, may take negative values. Nevertheless, choosing appropriately the mean and variance of θ the probability of such events can be made small provided $|\varsigma| < 1$. A problem with the variance of $\{u_t\}$ exploding is that the same happens with the variance of prices and quantities, and therefore prices and quantities become negative with increasing probability as time grows. Nevertheless it is possible to show, at least for the case $\mu = 0$, that the results in this section hold for an isoelastic-lognormal model in which prices and quantities are always positive. (See exercise 7.2.)

7.4 Slow Learning and Convergence

Consider a variation of the model in section 7.3 with no autocorrelation in demand ($\varsigma = 0$) where firms know the mean demand parameter α,

$$p_t = \alpha + u_t - \beta x_t,$$

but are uncertain about the intercept θ of their linear marginal cost function. Firm i has costs[15]

$$C(x_i) = \theta x_{it} + \tfrac{1}{2}\lambda x_{it}^2.$$

As in section 7.3 firm i is endowed with a private signal s_i about the unknown θ and the joint distribution of the random variables, and the values of the parameters β, α, and λ are common knowledge with $\alpha > \bar{\theta} > 0$ and $\lambda > 0$.

Consider the following scenario. Suppose that the firms' output pollutes or is toxic but that the level of toxicity or long-term pollution of a unit of output produced is unknown (the product is a new chemical, for example). Let θ be

---

[13] Now as t grows it does not become easier to predict the next-period price. The price prediction error, $\operatorname{var}[z_t - E[z_t \mid z^{t-1}]] = (\Delta\hat{a}_t)^2 \operatorname{var}[\theta - \theta_{t-1}] + \sigma_\eta^2$, is bounded away from the full-information irreducible error σ_η^2. Indeed, as $t \to \infty$, $\operatorname{var}[z_t - E[z_t \mid z^{t-1}]] \to (1 + \hat{\varsigma}^2) = \sigma_\eta^2$. This means that uninformed firms can never catch up with informed firms and predict prices as accurately as them.

[14] $\operatorname{var}[u_t]$ grows linearly with t for $|\varsigma| = 1$ and exponentially (at the rate ς^{2t}) for $|\varsigma| > 1$. When $|\varsigma| < 1$, then $\sigma_\eta^2/(1 - \varsigma^2)$ provides an upper bound for $\operatorname{var}[u_t]$.

[15] The model is inspired by an example provided by McKelvey and Page (1986). This section is based on Vives (1993).

the assessed toxicity or pollution damage of a unit of output produced. Firms have private assessments of θ and are allowed to produce but they will have to pay the corresponding damages in proportion to total production once the value of θ is realized. At each period there is an independent (small) probability $1-\delta > 0$ that θ is realized (that is, that the tests on the toxicity of the product are definitive). In this case the probability of θ not being realized in period t, δ^t, tends to 0 as $t \to \infty$. When θ is realized the firm has to pay the damage of the accumulated production. In any period a firm, given θ, expects a damage cost per unit of production[16] of exactly θ, and the firm maximizes the (expected) discounted profits with discount factor δ.

In period t firm i produces x_{it} obtaining a net revenue $p_t x_{it} - \frac{1}{2}\lambda x_{it}^2$. If θ is realized after t periods of production, the firm has to pay $\theta(\sum_{k=1}^{t} x_{ik})$. Profits of firm i corresponding to period t are

$$\pi_{it} = (p_t - \theta)x_{it} - \tfrac{1}{2}\lambda x_{it}^2,$$

although only the net revenue $p_t x_{it} - \frac{1}{2}\lambda x_{it}^2$ is observable in period t. Firm i tries to maximize the (expectation of the) discounted sum of period profits $\sum_{k=1}^{t} \delta^k \pi_{ik}$.

The model could also be restated in terms of competition among buyers of an asset of unknown *ex post* return θ. Consider, for example, firms purchasing labor of unknown productivity θ because of technological uncertainty, and facing a random inverse linear labor supply and adjustment costs in the labor stock. Firm i would buy x_{it} in period t. The monetary return to the final labor stock $\sum_{k=1}^{t} x_{ik}$ would then be $\theta(\sum_{k=1}^{t} x_{ik})$. Buyer i would face a quadratic cost of adjustment of his position in period t, x_{it}, equal to $\frac{1}{2}\lambda x_{it}^2$, $\lambda > 0$. Inverse supply in period t would be given by $p_t = u_t + \beta x_t$, with x_t the average (per capita) quantity demanded in period t. Buyer i would obtain benefit $\theta x_{it} - \frac{1}{2}\lambda x_{it}^2$ (with profits $\pi_{it} = (\theta - p_t)x_{it} - \frac{1}{2}\lambda x_{it}^2$) from the quantity x_{it} demanded in period t.[17]

In the exposition we will keep the convention that agents are sellers. In order to save on notation set $\beta = 1$.

Firm i in period t has to decide how much to produce by estimating the period price p_t on the basis of the information it has available: the private signal s_i plus the (public) information contained in past prices $p^{t-1} = \{p_1, \dots, p_{t-1}\}$. That is, in period t the information set of firm i is $\{s_i, p^{t-1}\}$. The distributional assumptions are as in section 7.3. It is assumed that there is a positive measure set of firms who receive signals of precision bounded away from zero.

It is worth remarking that the present model is one of learning from others. That is, what agents can learn about θ is given by their joint information. This is so because prices depend on θ only through the average market action. Indeed, if agents were to receive perfectly correlated signals, then they would have the same information and there would be nothing to learn. The assumption

---

[16] That is, $(1-\delta+\delta(1-\delta)+\delta^2(1-\delta)+\cdots)\theta = (1-\delta)(1-\delta)^{-1}\theta = \theta$.

[17] See section 9.1.2 for a financial market example.

of (conditionally) independent signals, plus the fact that a positive measure of agents receive signals of precision bounded away from zero, together with our convention on the "strong law of large numbers" for i.i.d. processes, ensures that if agents were to share information they would obtain θ, and agents can therefore, at least potentially, learn θ.

The perfect Bayesian equilibria of the dynamic game will be investigated. As in section 7.3 the equilibria will necessarily involve a sequence of Bayesian equilibria of the one-shot game. Otherwise, there would be a (positive measure) subset of agents that at some stage could, individually, improve their expected payoffs by reacting optimally to the average market action.

As in section 7.3 after t periods of trading with linear strategies the information contained in the price sequence $\{p_t\}$ will be summarized by a public belief statistic $\theta_t = E[\theta \mid z^t]$, where $z^t = \{z_1, \dots, z_t\}$ and $z_t = a_t\theta + u_t$, with a_t the average responsiveness of firms to their private signals. As before the random variable z_t represents the new information about θ in the price p_t. Profit maximization leads firm i in period t to choose the output level

$$x_{it} = \lambda^{-1} E[p_t - \theta \mid s_i, \theta_{t-1}].$$

It can be shown that there is a unique linear equilibrium (again unique in the class of strategies with bounded means and uniformly bounded variances across firms). Proposition 7.7 in the appendix provides the characterization of the equilibrium.

The strategy of a firm is a linear function of s_i and θ_{t-1} and in equilibrium

$$p_t = z_t + c_t\theta_{t-1} + \alpha\lambda(\lambda + 1)^{-1},$$

where as before c_t is the average responsiveness of firms to public information. The parameters are determined in a recursive way.

The informativeness of prices $\tau_t = \tau_\theta + \tau_u \sum_{k=1}^{t} a_k^2$ always increases and tends to ∞ and the price statistic θ_t eventually becomes fully revealing. Indeed, the only possibility for τ_t to be bounded above is for a_t to converge to zero (and fast) but this is self-contradictory: if τ_t does not tend to ∞, a_t will not tend to zero since agents will keep putting some weight on their private information. Consequently, the average responsiveness to private information a_t will (weakly) decrease with t. In fact, the average weight given to private information in $E[\theta \mid s_i, \theta_{t-1}]$, $\xi_{t-1} = \int_0^1 (\tau_{\varepsilon_i}/(\tau_{\varepsilon_i} + \tau_{t-1}))\,\mathrm{d}i$, converges to the proportion of perfectly informed agents μ, and the average weight given to public information converges to $1 - \mu$. Since $a_t = \xi_{t-1}(\lambda + \xi_{t-1})^{-1}$ we have that $a_t \xrightarrow[t]{} \mu(\lambda + \mu)^{-1}$.

If the informativeness of the price statistic θ_t, τ_t, is of the order t^{υ}, define the *asymptotic precision* as $A\tau_\infty = \lim_{t\to\infty} t^{-\upsilon}\tau_t$.

If there is a positive mass of perfectly informed agents $\mu > 0$, then a_t will be bounded away from zero, $a_t \to a_\infty > 0$, because these agents put constant weight on their private perfect information since they have nothing to learn from prices. The order of magnitude of τ_t depends on the information content of current prices (the new information) for t large. This is given by the random variable $z_t = a_t\theta + u_t$ and asymptotically z_t looks like $a_\infty\theta + u_t$. This means that the information content of z_t is asymptotically constant and the order of magnitude of τ_t is t, as in the standard linear regression model or i.i.d. noisy observations of θ. For t large τ_t can be approximated by a linear function of t with slope equal to $A\tau_\infty = \tau_u(\mu/(\lambda+\mu))^2$, $\tau_t \approx \tau_\theta + A\tau_\infty t$. The asymptotic precision of τ_t (the "slope" of convergence) is larger the larger is the proportion of perfectly informed agents μ and smaller the larger is the slope of the adjustment cost λ.

If there is no positive mass of perfectly informed agents, $\mu = 0$, then $a_t \to 0$ as the number of trading periods increases and prices become more informative, and asymptotically the new information is pure noise: z_t looks like u_t for t large. This will not preclude convergence but it will slow it down. This is so since agents then put less and less weight on their private signal and more and more on the price statistic, which becomes more informative and eventually fully revealing. The consequence is that τ_t will tend to ∞ more slowly. Indeed, agents, by reacting on average less to their private information, will incorporate less of it in the current price, slowing down the convergence of the price statistic to θ. This self-correcting property of learning from others has been explored in section 6.3, where a heuristic explanation of the slow learning result was given.

When the average precision of private information in the market τ_ε is finite, the order of τ_t is $t^{-1/3}$. For t large τ_t is approximated by a strictly concave function of t: $\tau_t \approx A\tau_\infty t^{1/3}$, where $A\tau_\infty = (3\tau_u(\tau_\varepsilon/\lambda)^2)^{1/3}$. The results follow from the fact that $a_t\tau_t$ converges to τ_ε/λ as $t \to \infty$. The asymptotic precision $A\tau_\infty$ is increasing with τ_u and with τ_ε, and decreasing with the slope of the adjustment cost λ. In all cases an increase in the adjustment cost λ decreases the asymptotic precision $A\tau_\infty$. An increase in λ has a direct effect of restricting the response of agents to information, for a given price precision, and a countervailing indirect effect of decreasing the price precision, inducing agents to put more weight on their signals. The direct effect dominates eventually and the asymptotic price precision diminishes.

The following lemma summarizes the results thus far.

Lemma 7.1. *In equilibrium, as $t \to \infty$:*

(i) $\tau_t \to \infty$, $a_t \to \mu(\lambda+\mu)^{-1}$ *and* $c_t \to \lambda(1-\mu)(\lambda+\mu)^{-1}(\lambda+1)^{-1}$.

(ii) *The informativeness of the price statistic θ_t, τ_t, is of order t^{υ}, and a_t is of order $t^{-\kappa}$, where $2\kappa + \upsilon = 1$, $\upsilon \in [\frac{1}{3}, 1]$, $\kappa \geqslant 0$. Furthermore:*

(iia) *If $\tau_\varepsilon \equiv \int_0^1 \tau_{\varepsilon_i}\,\mathrm{d}i < \infty$, then $\upsilon = \frac{1}{3}$ and $A\tau_\infty = (3\tau_u(\tau_\varepsilon/\lambda)^2)^{1/3}$.*

(iib) *Otherwise, if $\mu > 0$, then $\upsilon = 1$ and $A\tau_\infty = \tau_u(\mu/(\lambda+\mu))^2$; if $\mu = 0$, then $1 > \upsilon > \frac{1}{3} > \kappa > 0$.*

Proof. See the appendix (section 7.6). □

This result generalizes the slow learning result of the prediction model of section 6.3 to a market context. More precisely, the model is closely related to the model of learning from others with long-lived agents of section 6.3.3. The difference is that here there is a payoff externality and that we allow for agents to receive signals of different precisions. The payoff externality does not change the slow learning result and the differences in the precisions of the signals of the agents do not change the result as long as the average precision τ_ε is finite. If it is infinite but $\mu = 0$, then the precision of public information τ_t still grows more slowly than the benchmark rate of t but faster than $t^{1/3}$. If there is a positive mass of informed agents ($\mu > 0$), then learning is at the benchmark rate. (Exercise 7.4 states the results for the general prediction model.)

If agents were to know θ (i.e., $\mu = 1$), the equilibrium action would be $x^{\mathrm{f}} = (\alpha - \theta)/(1 + \lambda)$ in any period. This corresponds to an FRREE of a static market with inverse demand $p = \alpha - x$. Note, however, that this FRREE is not implementable in supply functions (see section 3.2).

Once the order of magnitude and asymptotic value of the precision of the price statistic θ_t is known its asymptotic distribution follows immediately. The following propositions establish the asymptotic results for learning θ and convergence to FIE. Learning follows from τ_t tending to ∞ with t. Convergence to equilibrium then follows from lemma 7.1(i) and the fact that $x_t - x^{\mathrm{f}} = c_t(\theta_{t-1} - \theta)$. The speed of learning and the rates of convergence follow from lemma 7.1(iii) immediately.

Proposition 7.4 (learning (Vives 1993)). *As $t \to \infty$:*

(i) *The public belief θ_t converges (almost surely and in mean square) to θ.*

(ii) *If τ_ε is finite, $\sqrt{t^{1/3}}(\theta_t - \theta) \xrightarrow{\mathrm{L}} N(0, (3\tau_u(\lambda/\tau_\varepsilon)^2)^{-1/3})$;
if $\mu > 0$, $\sqrt{t}(\theta_t - \theta) \xrightarrow{\mathrm{L}} N(0, (\mu/(\lambda + \mu))^{-2}\tau_u^{-1})$; and
if $\mu = 0$ and τ_ε is infinite, $\sqrt{t^{\upsilon}}(\theta_t - \theta) \xrightarrow{\mathrm{L}} N(0, (A\tau_\infty)^{-1})$,
for some $\upsilon \in (\frac{1}{3}, 1)$ and appropriate positive constant $A\tau_\infty$.*

Proposition 7.5 (convergence to FIE (Vives 1993)). *As $t \to \infty$:*

(i) *$x_t - x^{\mathrm{f}}$ converges (almost surely and in mean square) to zero.*

(ii) *If τ_ε is finite, $\sqrt{t^{1/3}}(x_t - x^{\mathrm{f}}) \xrightarrow{\mathrm{L}} N(0, (1 + \lambda)^{-2}(3\tau_u(\tau_\varepsilon/\lambda)^2)^{-1/3})$;
if $\mu > 0$, $\sqrt{t}(x_t - x^{\mathrm{f}}) \xrightarrow{\mathrm{L}} N(0, (\lambda/(\lambda + 1))^2\tau_u^{-1})$; and
if $\mu = 0$ and τ_ε is infinite, $\sqrt{t^{\upsilon}}(x_t - x^{\mathrm{f}}) \xrightarrow{\mathrm{L}} N(0, (1 + \lambda)^{-2}(A\tau_\infty)^{-1})$,
for some $\upsilon \in (\frac{1}{3}, 1)$ and appropriate positive constant $A\tau_\infty$.*

Agents eventually learn the unknown parameter θ since the price statistic θ_t is a (strongly) consistent estimator of θ. Nevertheless the speed of learning depends crucially on the distribution of private information in the market. If a positive mass of agents is perfectly informed, then convergence is at the standard rate $1/\sqrt{t}$. Otherwise, convergence is at a lower rate, reaching the lower bound $1/\sqrt{t^{1/3}}$ when the average precision of private information is finite.

The asymptotic variance (AV) of the departure from the FIE is larger the larger the asymptotic variance of τ_t and the response of agents to public information. If τ_ε is finite, an increase in λ may increase or decrease AV according to whether

λ is small or large. An increase in λ increases $(A\tau_\infty)^{-1}$ but decreases c_∞ and the overall effect depends on the size of λ. If μ is positive, then AV always increases in λ, and if all agents are perfectly informed, obviously, AV equals zero.

7.5 Summary

This chapter has introduced rational learning dynamics in versions of the static Cournot market of chapter 1 with an unknown payoff relevant parameter θ. Each period firms face a demand with a period-specific shock, receive a private signal about θ in period 0, and learn from past prices. The aim has been to see how and when repeated interaction makes the economy settle in a full-information equilibrium (interpretable as a fully revealing rational expectations equilibrium) and what determines the rate of convergence to this outcome. The most important results are:

- When the unknown parameter affects directly the public statistic (θ is a parameter of the demand function):
  - If period-specific shocks are stationary, then learning θ and convergence to the FIE happen at the same rate $1/\sqrt{t}$, where t is the number of periods.
  - If period-specific shocks are nonstationary, then learning θ is neither necessary nor sufficient for convergence to the FIE:
    - When shocks follow a random walk, θ is never fully learned but convergence to the FIE obtains very fast (exponentially).
    - When the process of shocks is sufficiently explosive learning is very fast but there is no convergence.
- When the unknown parameter affects indirectly the public statistic (θ is a parameter of the cost function), the model is one of learning from others—since firms, by observing prices, at most can hope to learn the pooled information of all agents in the market—and shocks are i.i.d.:
  - If there is no positive mass of agents perfectly informed, learning θ and converging to the FIE is slow, typically of order $1/\sqrt{t^{1/3}}$.

Thus we see that when agents learn in equilibrium, learning and convergence to an FIE obtains under standard assumptions (i.e., stationary shocks) as the number of market interactions increases. However, learning and convergence may be very slow when learning is from others. The reason is that Bayesian learning from others has a self-correcting property that does not preclude convergence but slows it down in the presence of noisy observation. The slow learning result generalizes the results obtained in the prediction model of section 6.3 to a general distribution of precisions of private signals and a market environment with payoff externalities.

7.6 Appendix

7.6.1 Proofs

Section 7.3

Derivation of the linear equilibrium. Posit the following form for the candidate equilibrium price function: $p_t = z_t + L(p^{t-1})$, with L a linear function. It follows that in a linear equilibrium p^t is observationally equivalent to z^t. Firm i in period t is interested in predicting the current price p_t with the information $I_{it} = \{s_i, p^{t-1}\}$. In equilibrium $p_t = z_t - \beta c_t \theta_{t-1} - \beta(\lambda + \beta)^{-1} \varsigma z_{t-1}$, with $z_t = \hat{a}_t \theta + u_t$, $\hat{a}_t = 1 - \beta a_t$, where a_t (c_t) is the average responsiveness of firms to private (public) information. The best predictor of the current price is $E[p_t \mid I_{it}]$. This is equivalent to predicting z_t with the information $\{s_i, z^{t-1}\}$: $E[z_t \mid s_i, z^{t-1}]$. Firm i is Bayesian and will compute $E[z_t \mid s_i, z^{t-1}]$ with its knowledge of the structure of the model and the (equilibrium) knowledge of the coefficients $\hat{a}_t$. In general, to predict z_t (compute $E[z_t \mid I_{it}]$) it will be necessary to predict θ (compute $E[\theta \mid I_{it}]$).

Recall the definitions: $\Delta z_t = z_t - \varsigma z_{t-1}$ and $\Delta \hat{a}_t = \hat{a}_t - \varsigma \hat{a}_{t-1}$. Then $\Delta z_t = \Delta \hat{a}_t \theta + \eta_t$ (where, by definition, $z_0 = 0$ and $\hat{a}_0 = 0$) and $\theta_t = E[\theta \mid \Delta z^t]$. Also let $\theta_0 \equiv \bar{\theta}$ and $\tau_0 \equiv \tau_\theta$. Bayesian prediction of $z_t = \Delta z_t + \varsigma z_{t-1}$ with information $\{s_i, z^{t-1}\}$ follows from normal theory:

$$E[z_t \mid s_i, z^{t-1}] = \Delta \hat{a}_t E[\theta \mid s_i, \Delta z^{t-1}] + \varsigma z_{t-1},$$
$$E[\theta \mid s_i, \Delta z^{t-1}] = \xi_{it-1} s_i + (1 - \xi_{it-1}) \theta_{t-1},$$

where $\xi_{it} = \tau_{\varepsilon_i} / (\tau_{\varepsilon_i} + \tau_t)$ and $\tau_t = \tau_\theta + \tau_\eta \sum_{k=1}^{t} (\Delta \hat{a}_k)^2$. The estimator of θ is a convex combination of private and public information with weights according to their relative precisions. In order to predict z_t the estimation of θ matters provided $\Delta \hat{a}_t$ does not equal 0. The coefficient $\Delta \hat{a}_t = 0$ (for $t \geqslant 2$) when $\varsigma = 1$ and $\hat{a}_t = \hat{a}_{t-1}$. That is, in order to predict z_t the estimation of θ does not matter when the AR process follows a random walk *and* the sensitivity of the price to θ is stationary.

Because the triple $(s_i, \theta_{t-1}, z_{t-1})$ is sufficient in the estimation of z_t, we may restrict our attention to strategies of the type $X_{it}(s_i, \theta_{t-1}, z_{t-1})$. As usual we use notation in such a way that a parameter with subscript "i" represents the individual coefficient of the corresponding aggregate coefficient, for example, $\int_0^1 a_{it} \, \mathrm{d}i = a_t$. The following proposition provides a characterization of linear equilibria.

Proposition 7.6 (Jun and Vives 1996). *There is a unique equilibrium in which the price in period t is a linear function of θ, u_t, θ_{t-1}, and z_{t-1}:*

$$p_t = \hat{a}_t \theta + u_t - \beta c_t \theta_{t-1} - \beta(\lambda + \beta)^{-1} \varsigma z_{t-1},$$

with $\hat{a}_t = 1 - \beta a_t$, $c_t = (\lambda + \beta)^{-1} \Delta \hat{a}_t (1 - \xi_{t-1})$, and the coefficients a_t are recursively defined $a_t = (\lambda + \beta \xi_{t-1})^{-1} \xi_{t-1} (1 - \varsigma(1 - \beta a_{t-1}))$, with $a_1 = (\lambda + \beta \xi_0)^{-1} \xi_0$.

Corollary 7.1. *Firms' strategies are given by*

$$X_{it}(s_i, \theta_{t-1}, z_{t-1}) = a_{it} s_i + c_{it} \theta_{t-1} + (\lambda + \beta)^{-1} \varsigma z_{t-1},$$

with $a_{it} = \lambda^{-1} \Delta \hat{a}_t \xi_{it-1}$ *and* $c_{it} = \lambda^{-1} \Delta \hat{a}_t [(1 - \xi_{it-1}) - \beta(\lambda + \beta)^{-1}(1 - \xi_{t-1})]$.

Outline of the proof. Consider period t. Posit a linear price function as in the proposition. Profit maximization leads each firm to choose the output level according to $x_{it} = \lambda^{-1} E[p_t \mid I_{it}]$, where $I_{it} = \{s_i, p^{t-1}\}$, or, equivalently, $I_{it} = \{s_i, \theta_{t-1}, z_{t-1}\}$. Aggregate output can then be obtained by integrating the individual outputs. We now have a price function and aggregate output expressed in terms of the parameters that we want to determine. By substituting these expressions into both sides of the inverse demand and identifying coefficients, we can solve (uniquely) for these parameters. Individual and aggregate outputs follow. Firms are maximizing profit and the market clears at the expected price level. The linear structure preserves normality of distributions.

Proofs of propositions 7.2 and 7.3. Convergence of the market equilibrium x_t to the FIE x_t^{f} depends on the asymptotic behavior of the parameter c_t and the price statistic θ_t. The information contained in θ_t is determined by the sequence $\{\Delta \hat{a}_t\}$. The precision incorporated in θ_t is given by $\tau_t = \tau_\theta + \tau_\eta \sum_{k=1}^{t} (\Delta a_k)^2$. The asymptotic behavior of c_t also depends on the properties of $\{\Delta \hat{a}_t\}$ since $c_t = (\lambda + \beta)^{-1} \Delta \hat{a}_t (1 - \xi_{t-1})$. We characterize the asymptotic behavior of equilibrium parameters in lemma 7.2, and learning about θ and convergence to the FIE in propositions 7.2 and 7.3.

Lemma 7.2. *As* $t \to \infty$*:*

(i) *If* $\varsigma \neq 1$*, then* $\tau_t \to \infty$.

(ii) *If* $|\varsigma| < 1 + \lambda/\beta\mu$*, then* $\Delta \hat{a}_t \to \Delta \hat{a}_\infty$ *and* $c_t \to (\lambda + \beta)^{-1} \Delta \hat{a}_\infty$*, with* $\Delta \hat{a}_\infty = (1 - \varsigma)\lambda / [\lambda + (1 - \varsigma)\beta\mu]$.

(iii) *If* $\varsigma \neq 1$ *and* $|\varsigma| < 1 + \lambda/\beta\mu$*, then* $t^{-1}\tau_t \to (\Delta \hat{a}_\infty)^2 \tau_\eta$.

Proof. (Sketch.) (i) The precision τ_t is monotone increasing. If it converges to a finite number, then $\Delta \hat{a}_t$ must converge to 0. From the definition of $\hat{a}_t$, we can obtain that $\hat{a}_t = 1 - \beta\lambda^{-1} \Delta \hat{a}_t \xi_{t-1}$. Hence, $\hat{a}_t \to 1$, but this implies $\Delta \hat{a}_t \to 1 - \varsigma \neq 0$, which is a contradiction.

(ii) An unbounded public precision means that agents who do not have perfect information asymptotically disregard private information, and therefore the aggregate weight given to private information tends to the proportion μ of perfectly informed agents ($\xi_t \to \mu$). This implies that for t large the dynamic system for a_t is approximately $a_t = (1 - \varsigma(1 - \beta a_{t-1}))(\lambda + \beta\mu)^{-1}\mu$. The system is stable, with a_t tending to $(1 - \varsigma)\mu / [\lambda + (1 - \varsigma)\beta\mu]$ as $t \to \infty$, if $|\varsigma| < 1 + \lambda/\beta\mu$. The limits for $\Delta \hat{a}_t$, c_t, and $\hat{a}_t$ ($\hat{a}_t \to \lambda/[\lambda + (1 - \varsigma)\beta\mu]$) then follow from $a_t = \lambda^{-1} \xi_{t-1} \Delta \hat{a}_t$, $\hat{a}_t = 1 - \beta a_t$, and $c_t = (\lambda + \beta)^{-1}(1 - \xi_{t-1})\Delta \hat{a}_t$. If $\varsigma = 1$ and $\tau_t \to \infty$, the same argument applies. Otherwise, $\Delta \hat{a}_t$ must converge to 0 and the result also follows.

(iii) follows from (i) and (ii) since $\tau_t = \tau_\theta + \tau_\eta \sum_{k=1}^t (\Delta\hat{a}_k)^2$, $\Delta\hat{a}_t \to \Delta\hat{a}_\infty$, and, therefore, $t^{-1}\tau_t \to (\Delta\hat{a}_\infty)^2\tau_\eta$. □

Proof of proposition 7.2 (see p. 259). (i) First, the result for θ_t follows as in the proof of proposition 7.1(i) by noting that now $\Delta z_t = \Delta\hat{a}_t\theta + \eta_t$, $\Delta\hat{a}_t \to \Delta\hat{a}_\infty = (1-\varsigma)\lambda/[\lambda + (1-\varsigma)\beta\mu]$. We have that $\mathrm{var}[\theta_t - \theta] = \tau_t^{-1}$ and the result follows because $t^{-1}\tau_t \to A\tau_\infty = (\Delta\hat{a}_\infty)^2\tau_\eta$. Second, we have that $x_t - x_t^{\mathrm{f}} = c_t(\theta_{t-1}-\theta) = (\lambda+\beta)^{-1}\Delta\hat{a}_t(1-\xi_{t-1})(\theta_{t-1}-\theta)$. If $\varsigma \neq 1$, then $\Delta\hat{a}_t \to \Delta\hat{a}_\infty$, $\tau_t \to \infty$, and $\xi_t \to \mu$. Hence, convergence obtains from (i). Now, $\mathrm{var}[x_t - x_t^{\mathrm{f}}] = (\lambda+\beta)^{-2}(\Delta\hat{a}_t)^2((1-\xi_{t-1})^2\,\mathrm{var}[\theta_{t-1}-\theta])$. Furthermore, from (i) $t\,\mathrm{var}[\theta_{t-1}-\theta] \to ((\Delta\hat{a}_\infty)^2\tau_\eta)^{-1}$ and therefore $t((\Delta\hat{a}_t)^2\,\mathrm{var}[\theta_{t-1}-\theta]) \to \sigma_\eta^2$. The result follows since $\xi_t \to \mu$.

(ii) (Sketch.) When $\varsigma = 1$ the precision of the public statistic is bounded. Indeed, it is easily seen that $\hat{a}_t < 1$ increases with t (and therefore $\Delta\hat{a}_t = \hat{a}_t - \hat{a}_{t-1} \geqslant 0$). It then follows that

$$\sum_{k=1}^t (\Delta\hat{a}_k)^2 \leqslant \left(\sum_{k=1}^t |\Delta\hat{a}_k|\right)^2 = \left(\sum_{k=1}^t \Delta\hat{a}_k\right)^2 = (\hat{a}_t)^2 \leqslant 1,$$

and therefore $\tau_t = \tau_\theta + \tau_\eta\sum_{k=1}^t \Delta\hat{a}_k \leqslant \tau_\theta + \tau_\eta$. We conclude that $\tau_t \to \infty$, since τ_t is increasing, with $\tau_\theta \leqslant \tau_\infty \leqslant \tau_\theta + \tau_\eta$. The increment $\Delta\hat{a}_t$ is of the same order as a_t, which converges to zero exponentially. It then follows that $x_t - x_t^{\mathrm{f}} = (\lambda+\beta)^{-1}\Delta\hat{a}_t(1-\xi_{t-1})(\theta_{t-1}-\theta)$ also converges to zero exponentially with t. □

Proof of proposition 7.3 (see p. 260). (Sketch.) (i) As in the proof of the lemma it is clear that $\tau_t \to \infty$ and that the aggregate weight given to private information tends to the proportion μ of perfectly informed agents ($\xi_t \to \mu$) as $t \to \infty$. Again this implies that for t large the dynamic system for a_t is approximately $a_t = (1-\varsigma(1-\beta a_{t-1}))(\lambda+\beta\mu)^{-1}\mu$ but this system is unstable with a_t tending to $-\infty$ exponentially with t if $\varsigma > 1+\lambda/\beta\mu$. From the expressions $a_t = \lambda^{-1}\xi_{t-1}\Delta\hat{a}_t$, $\hat{a}_t = 1-\beta a_t$, and $c_t = (\lambda+\beta)^{-1}(1-\xi_{t-1})\Delta\hat{a}_t$ we have that $\{|\hat{a}_t|\}$ and $\{|\Delta\hat{a}_t|\}$, $\{|c_t|\}$ also tend to ∞ exponentially ($\hat{a}_t$ to ∞, $\Delta\hat{a}_t$ and c_t to $-\infty$). Given that $\tau_t = \tau_\theta + \tau_\eta\sum_{k=1}^t(\Delta\hat{a}_k)^2$, it is possible to show that

$$\lim_{t\to\infty}(\Delta\hat{a}_t)^{-2}\tau_t = \hat{\varsigma}^2(\hat{\varsigma}^2-1)^{-1}\tau_\eta.$$

(ii) This follows because τ_t grows exponentially with t.

(iii) This follows from the expression of $x_t - x_t^{\mathrm{f}} = c_t(\theta_{t-1}-\theta)$, $c_t = (\lambda+\beta)^{-1}(1-\xi_{t-1})\Delta\hat{a}_t$, $\lim_{t\to\infty}(\Delta\hat{a}_t)^{-2}\tau_{t-1} = (\hat{\varsigma}^2-1)^{-1}\tau_\eta$. □

Section 7.4

Proposition 7.7 (Vives 1993). *In the model of section 7.4, there is a unique equilibrium (in the class of strategies with bounded means and uniformly—across agents—bounded variances). The equilibrium is linear and the price function is given by*

$$p_t = a_t\theta + u_t + c_t\theta_{t-1} + \lambda\alpha(1+\lambda)^{-1},$$

where $a_t = \xi_{t-1}(\lambda+\xi_{t-1})^{-1}$ *and* $c_t = (\lambda+\xi_{t-1})^{-1}(\lambda+1)^{-1}\lambda(1-\xi_{t-1})$.

Strategies are given by

$$X_{it}(s_i, \theta_{t-1}) = \alpha(1+\lambda)^{-1} - a_{it}s_i - c_{it}\theta_{t-1},$$

where $a_{it} = \xi_{it-1}(\lambda + \xi_{t-1})^{-1}$, $c_{it} = (\lambda + \xi_{t-1})^{-1}(\lambda+1)^{-1}((\lambda+1)(1-\xi_{t-1}) - (1-\xi_{t-1}))$, *with* $\xi_{it-1} = \tau_{\varepsilon_i}/(\tau_{\varepsilon_i} + \tau_{t-1})$ *and* $\xi_{t-1} = \int_0^1 \xi_{it-1}\,\mathrm{d}i$.

Proof. Existence and characterization of a linear equilibrium are similar to the proof of proposition 7.6. Uniqueness follows from exercise 7.3. See the proofs of propositions 4.1 and 4.2 in Vives (1993). □

Proof of lemma 7.1 (see p. 264). The proof of (i) follows from proposition 7.7 and the fact that $\xi_t \to \mu$ as $t \to \infty$ (see the proof of lemma 5.1 in Vives 1993).

(ii) Let a_t be of order $t^{-\kappa}$ and τ_t be of order t^{υ}. Since $a_t \underset{t}{\longrightarrow} \mu(\lambda+\mu)^{-1}$ and $\tau_t \underset{t}{\longrightarrow} \infty$, necessarily $\kappa \geqslant 0$ and $\upsilon > 0$. If $\mu > 0$, then $a_\infty = \mu(\lambda+\mu)^{-1} > 0$ and therefore $\kappa = 0$ and $\upsilon = 1$. It then follows that $t^{-1}\tau_t \underset{t}{\longrightarrow} \tau_\theta + \tau_u(\mu/(\lambda+\mu))^2$. If $\mu = 0$, the result follows from the two claims below. □

Claim 1. $a_t\tau_t$ converges to τ_ε/λ (τ_ε can be infinite) as $t \to \infty$.

Proof. Let $A_t \equiv \sum_{k=0}^{t} a_k^2$. Then $\tau_t = \tau_\theta + \tau_u A_t$. We show the equivalent result that $a_t\tau_{t-1}$ or $a_t A_{t-1}$ converges to τ_ε/λ. From $a_t = \xi_{t-1}(\lambda+\xi_{t-1})^{-1}$ we have that $a_t A_{t-1} = ((A_{t-1})^{-1} + \lambda(A_{t-1}\xi_{t-1})^{-1})^{-1}$. Now,

$$A_{t-1}\xi_{t-1} = \int_0^1 A_{t-1}\xi_{it-1}\,\mathrm{d}i = \int_0^1 \tau_{\varepsilon_i}((\tau_\theta + \tau_{\varepsilon_i})(A_{t-1})^{-1} + \tau_u)^{-1}\,\mathrm{d}i.$$

Observing that $A_{t-1}\xi_{it-1}$ is a monotone increasing sequence of nonnegative measurable functions of i converging almost everywhere to $\tau_{\varepsilon_i}/\tau_u$, it is possible to conclude (Lebesgue Monotone Convergence Theorem (see Royden 1968, p. 227)) that $\int_0^1 A_{t-1}\xi_{it-1}\,\mathrm{d}i$ converges to $\tau_u^{-1}\int_0^1 \tau_{\varepsilon_i}\,\mathrm{d}i$ and the result follows (note that $A_{t-1} \to \infty$ with t). □

Claim 2. If τ_ε is finite, then $t^{-1/3}A_t \to 3^{1/3}(\tau_\varepsilon/\lambda\tau_u)^{2/3}$ as $t \to \infty$. If τ_ε is infinite (and $\mu = 0$), then $1 > \upsilon > \frac{1}{3} > \kappa > 0$.

Proof. If $\mu = 0$, then $a_t \underset{t}{\longrightarrow} 0$. If τ_ε is finite, the result follows from claim 1 and lemma 7.3. If τ_ε is infinite, then $a_t A_t \underset{t}{\longrightarrow} \infty$ with t. Further, $t^{-1}A_t \underset{t}{\longrightarrow} 0$ (since $a_t \to 0$) and $t^{-1/3}A_t \underset{t}{\longrightarrow} \infty$ (using the fact that $A_t \geqslant t a_t^2$, $(a_t A_t)^{-1} \geqslant (t^{-1/3}A_t)^{-3/2}$, from which the result follows since $a_t A_t \underset{t}{\longrightarrow} \infty$). It then follows from lemma 7.3 (with $\upsilon > \kappa > 0$) that $2\kappa + \upsilon = 1$ and $1 > \upsilon > \frac{1}{3}$. Note that in all cases $2\kappa + \upsilon = 1$ and $1 \geqslant \upsilon \geqslant \frac{1}{3}$. □

Lemma 7.3. *Assume that* $a_t A_t \underset{t}{\longrightarrow} k > 0$, *then* $t^{-1/3}A_t \underset{t}{\longrightarrow} 3(k/3)^{2/3}$ *and* $t^{1/3}a_t \underset{t}{\longrightarrow} (k/3)^{1/3}$.

Proof. See the heuristic argument in section 6.3.1 and the proof of lemma A.1 in Vives (1993). An alternative proof is provided in the appendix to chapter 6. □

Lemma 7.4. *Assume that* $t^{-\upsilon}\sum_{k=1}^{t} a_k^2 \xrightarrow[t]{} k$*, where* $k > 0$*, and that* a_t *is of the order* $t^{-\kappa}$*, where* $\upsilon \geqslant \kappa > 0$*. Then* $\kappa < \frac{1}{2}$ *and* $\sum_{k=1}^{t} a_k^2$ *is of the order of* $t^{1-2\kappa}/(1-2\kappa)$*. It follows that* $\upsilon + 2\kappa = 1$ *and* $1 > \upsilon \geqslant \frac{1}{3}$*.*

Proof. See the proof of lemma A.2 in Vives (1993). □

7.7 Exercises

7.1 (*information acquisition and fully revealing equilibria*). Consider the Cournot game in the continuum economy with increasing marginal costs of section 1.6 and add another market period. At a first stage each firm has the opportunity to purchase the precision of its signal (like in section 1.6). Firms compete in quantities contingent on the received signals at a first market period, observe the market price, and compete again in a second market period. Show that each firm purchasing "precision" ξ^* at the first stage, producing $X_i(s_i) = a(s_i - \mu) + b\mu$ at the first market period (second stage), where $a = \xi^*/(\lambda + \beta\xi^*)$ (with ξ^*, a, and b as in the two-stage game), inferring θ from the market price, and producing $x = b\theta$ at the second market period (third stage) is an equilibrium path for the three-stage game. Bearing in mind the Grossman–Stiglitz paradox (see section 4.2.2), what do you conclude about the incentives to purchase information in the presence of fully revealing equilibria?

Solution. Suppose that firms have already chosen their precision of the information, ξ_i, $i \in [0,1]$. Firm i at stage two will follow the strategy $X_i(s_i) = a_i(s_i - \bar{\theta}) + b\bar{\theta}$ as in the two-stage game (where $a_i = \xi_i/(\lambda + \beta\xi)$ and $\xi = \int_0^1 \xi_i \,\mathrm{d}i$) since in the continuum economy the firm cannot affect the market price through its quantity choice. The market price conditional on θ will then be $p = \theta - \beta\tilde{X}(\theta)$, where $\tilde{X}(\theta)$ is the average output, $\tilde{X}(\theta) = a(\theta - \bar{\theta}) + b\bar{\theta}$ with $a = \int_0^1 a_i \,\mathrm{d}i = \xi/(\lambda + \beta\xi)$. Therefore, $p = \theta(1 - \beta a) - \beta(b - a)\bar{\theta}$ and provided that $1 - \beta a \neq 0$, θ can be obtained by observing the market price: $\theta = (p + \beta(b - a)\mu)/(1 - \beta a)$ (note that $a\beta < 1$ except if $\lambda = 0$). When θ is revealed a full-information competitive equilibrium obtains at the last stage, $x = b\theta$ and $p = \lambda b\theta$. At the last stage a fully revealing rational expectations equilibrium prevails but nevertheless firms have an incentive to purchase information since this affects expected profits at the second stage before the market price is observed (see Dubey et al. (1982) for an elaboration of this idea in a strategic market game context where they show that Nash equilibria of the continuum economy are fully revealing generically).

***7.2** (*isoelastic-lognormal model*). Similarly to section 1.2.4.2, consider a market with constant elasticity inverse demand and costs. In period t inverse demand is given by $p_t = \mathrm{e}^{\theta + u_t} x_t^{-\beta}$, $\beta > 0$, where the information structure is exactly as in section 7.3 with the simplifying assumption that the signals received by agents $s_i = \theta + \varepsilon_i$ have equal (finite) precision ($\tau_{\varepsilon_i} = \sigma_\varepsilon^{-2}$ for all i). Firms have constant elasticity cost functions given by $C(x_{it}) = (1 + \lambda)^{-1} x_{it}^{1+\lambda}$, $\lambda > 0$. Find an equilibrium in log-linear strategies (it looks similar to the linear model, but in log form, with the addition of a constant term which depends on the various variances). Show that the FIE output is given by

$$x_t^{\mathrm{f}} = \mathrm{e}^{(\lambda+\beta)^{-1}(\theta+\varsigma u_{t-1}+\sigma_\eta^2/2)}$$

and that convergence to the FIE is analogous to the result in the linear-normal model replacing $x_t - x_t^{\mathrm{f}}$ by $\log(x_t/x_t^{\mathrm{f}})$. (The departure from the FIE depends as before on the term $c_t(\theta_{t-1} - \theta)$ but there is an additional variance term.)

Solution (from Jun and Vives 1995). The procedure is similar to the one outlined for the proofs of propositions 7.2 and 7.6 (with $\bar{\theta} = 0$) using our convention about the average of a continuum of random variables (which implies that $\int_0^1 \mathrm{e}^{\varepsilon_i}\,\mathrm{d}i = \int_0^1 E[\mathrm{e}^{\varepsilon_i}]\,\mathrm{d}i$) and the properties of lognormal distributions (namely that if $\varepsilon_i \sim N(0, \sigma_\varepsilon^2)$, then $E[\mathrm{e}^{\varepsilon_i}] = \mathrm{e}^{\sigma_\varepsilon^2/2}$). The result is that

$$\log(X_t(s_i, \theta_{t-1}, z_{t-1})) = a_t s_i + c_t\theta_{t-1} + (\lambda+\beta)^{-1}\varsigma z_{t-1} + k_t$$

and

$$\log(p_t) = \hat{a}_t\theta + u_t - \beta c_t\theta_{t-1} - \beta(\lambda+\beta)^{-1}\varsigma z_{t-1} - \beta(k_t + \tfrac{1}{2}a_t^2\sigma_\varepsilon^2),$$

where all parameters are as in proposition 7.6 and

$$k_t = (\lambda+\beta)^{-1}[\tfrac{1}{2}\sigma_\eta^2 + \lambda^2(a_t/\xi_{t-1})^2/2\tau_{t-1} - \tfrac{1}{2}\beta a_t^2\sigma_\varepsilon^2].$$

With full information a firm at time t knows $\{\theta, u^{t-1}\}$. From the FOC of profit maximization we get

$$x_t^{\mathrm{f}} = (E[p_t \mid \theta, u^{t-1}])^{1/\lambda} = \exp\{\lambda^{-1}(\theta + \varsigma u_{t-1} + \tfrac{1}{2}\sigma_\eta^2)\}(x_t^{\mathrm{f}})^{-\beta/\lambda}.$$

The expression for x_t^{f} follows. We then have that $\ln(x_t/x_t^{\mathrm{f}}) = c_t(\theta_{t-1} - \theta) + [\lambda^2\xi_{t-1}^{-2}\tau_{t-1}^{-1} - \phi\sigma_\varepsilon^2]a_t^2/2(\alpha+\beta)$ and the asymptotic results follow with a little bit of work. As $t \to \infty$:

(i) $x_t/x_t^{\mathrm{f}} \to 1$ (a.s. and in mean square), and

(ii) if $\varsigma \neq 1$, then $\sqrt{t}\log(x_t/x_t^{\mathrm{f}}) \xrightarrow{\mathrm{L}} N(0, \sigma_\eta^2/(\alpha+\beta)^2)$.

*__7.3__ (*uniqueness of dynamic equilibrium*). Prove the uniqueness part of proposition 7.7.

Solution. The uniqueness argument proceeds by showing that at any stage only linear equilibria are possible. Consider a generic stage, drop the period subscript, and denote by $\hat{\theta}$ the public statistic. Assume that the information available to agent i is $\{s_i, \hat{\theta}\}$, where the triple $(\theta, s_i, \hat{\theta})$ is jointly normally distributed. Then using FOCs show that for any candidate equilibrium average output $X(\theta, \hat{\theta})$, $E[(X(\theta,\hat{\theta}) - (\alpha(1+\lambda)^{-1} - a\theta - c\hat{\theta}))^2] = 0$, where $\alpha(1+\lambda)^{-1} - a\theta - c\hat{\theta}$ is the linear equilibrium. This implies that $X(\theta, \hat{\theta}) = \alpha(1+\lambda)^{-1} - a\theta - c\hat{\theta}$ (a.s.) and therefore necessarily $X_i(s_i, \hat{\theta}) = \alpha(1+\lambda)^{-1} - a_i s_i - c_i\hat{\theta}$ (a.s.). Check finally by induction that at any stage $(\theta, s_i, \hat{\theta})$ is jointly normally distributed. (See the proof of proposition 4.2 in Vives (1993) for the details.)

** **7.4** (*slow learning in the generalized prediction model*). Consider the prediction model of section 6.3.1 with no correlation in signals of agents but where the distribution of precisions of private signals for every generation is given by a (measurable) function $T_\varepsilon : [0,1] \to \mathbb{R}_+ \cup \{\infty\}$, where $T_\varepsilon(i) = \sigma_{\varepsilon_i}^{-2} (\equiv \tau_{\varepsilon_i})$ is the value of this function for $i \in [0,1]$. Denote by $\mu \geqslant 0$ the mass of agents with perfectly informative signals. Revise proposition 6.1 and show that as $t \to \infty$:

(i) $a_t \to \mu$ and $\tau_t \to \infty$.

(ii) $\theta_t \to \theta$ almost surely and in mean square;

(iii) τ_t is of order t^{υ} and a_t is of order $t^{-\kappa}$, where $2\kappa + \upsilon = 1$, $\upsilon \in [\frac{1}{3}, 1]$, $\kappa \geqslant 0$. Let $A\tau_\infty = \lim_{t\to\infty} t^{-\upsilon}\tau_t$. If $\tau_\varepsilon \equiv \int_0^1 \tau_{\varepsilon_i}\,\mathrm{d}i < \infty$, then $\upsilon = \frac{1}{3}$ and $A\tau_\infty = (3\tau_u)^{1/3}(\tau_\varepsilon)^{2/3}$. Otherwise, if $\mu > 0$, then $\upsilon = 1$ and $A\tau_\infty = \tau_u\mu^2$; if $\mu = 0$, then $1 > \upsilon > \frac{1}{3} > \kappa > 0$.

(iv) If $\tau_\varepsilon < \infty$, $\sqrt{t^{1/3}}(\theta_t - \theta) \xrightarrow{\mathrm{L}} N(0, (3\tau_u\tau_\varepsilon^2)^{-1/3})$;
if $\mu > 0$, $\sqrt{t}(\theta_t - \theta) \xrightarrow{\mathrm{L}} N(0, (\tau_u\mu^2)^{-1})$; and
if $\mu = 0$ and τ_ε is infinite, $\sqrt{t^{\upsilon}}(\theta_t - \theta) \xrightarrow{\mathrm{L}} N(0, (A\tau_\infty)^{-1})$ for some $\upsilon \in (\frac{1}{3}, 1)$ and appropriate positive constant $A\tau_\infty$.

Solution. Adapt the arguments in section 7.3 and the appendix.

References

Alchian, A. 1950. Uncertainty, evolution and economic theory. *Journal of Political Economy* 58:211–21.

Bisin, A., U. Horst, and O. Özgur. 2006. Rational expectations equilibria of economies with local interactions. *Journal of Economic Theory* 127:74–116.

Blackwell, D., and L. Dubins. 1962. Merging of opinions with increasing information. *Annals of Mathematical Statistics* 33:882–86.

Blume, L. E., and D. Easley. 1982. Learning to be rational. *Journal of Economic Theory* 26:340–51.

——. 1984. Rational expectations equilibrium: an alternative approach. *Journal of Economic Theory* 34:116–29.

——. 1998. Rational expectations and rational learning. In *Organizations with Incomplete Information: Essays in Economic Analysis: A Tribute to Roy Radner*, pp. 61–109. Cambridge University Press.

Bray, M. 1982. Learning, estimation, and the stability of rational expectations. *Journal of Economic Theory* 26:318–39.

Bray, M., and D. M. Kreps. 1988. Rational learning and rational expectations. In *Arrow and the Ascent of Modern Economic Theory* (ed. G. Feiwel), pp. 597–625. London: Macmillan.

Bray, M., and N. E. Savin. 1986. Rational expectations equilibria, learning, and model specification. *Econometrica* 54:1129–60.

Brock, W., and C. Hommes. 1997. A rational route to randomness. *Econometrica* 65: 1059–95.

De Long, J., A. Shleifer, L. Summers, and R. Waldmann. 1989. The size and incidence of the losses from noise trading. *Journal of Finance* 44:681–96.

De Long, J., A. Shleifer, L. Summers, and R. Waldmann. 1990. Noise trader risk in financial markets. *Journal of Political Economy* 98:703–38.

Dubey, P., J. Geanakoplos, and M. Shubik. 1982. Revelation of information in strategic market games: a critique of rational expectations. Cowles Foundation Discussion Papers 634R, Yale University.

Evans, D., and S. Honkapohja. 2001. *Learning and Expectations in Macroeconomics.* Princeton University Press.

Feldman, M. 1987. An example of convergence to rational expectations with heterogeneous beliefs. *International Economic Review* 28:635–50.

Fourgeaud, C., C. Gourieroux, and J. Pradel. 1986. Learning procedures and convergence to rationality. *Econometrica* 54:845–68.

Friedman, M. 1953. *Essays in Positive Economics.* University of Chicago Press.

Frydman, R. 1982. Towards an understanding of market processes: individual expectations, learning, and convergence to rational expectations equilibrium. *American Economic Review* 72:652–68.

Guesnerie, R. 2006. *Assessing Expectations*, volume 2: *Eductive Stability in Economics.* Cambridge, MA: MIT Press.

Hirshleifer, D., and G. Luo. 2001. On the survival of overconfident traders in a competitive securities market. *Journal of Financial Markets* 4:73–84.

Jordan, J. S. 1985. Learning rational expectations: the finite state case. *Journal of Economic Theory* 36:257–76.

——. 1992a. The exponential convergence of Bayesian learning in normal form games. *Games and Economic Behavior* 4:202–17.

——. 1992b. Convergence to rational expectations in a stationary linear game. *Review of Economic Studies* 59:109–23.

Jun, B., and X. Vives. 1995. Learning and convergence to REE with persistent shocks. Mimeo.

——. 1996. Learning and convergence to a full-information equilibrium are not equivalent. *Review of Economic Studies* 63:653–74.

Kalai, E., and E. Lehrer. 1993. Rational learning leads to Nash equilibrium. *Econometrica* 61:1019–45.

Kyle, A. S., and F. Wang. 1997. Speculation duopoly with agreement to disagree: can overconfidence survive the market test? *Journal of Finance* 52:2073–90.

Lucas, R. 1978. Asset prices in an exchange economy. *Econometrica* 46:1429–45.

Marcet, A., and T. J. Sargent. 1988. The fate of systems with "adaptive" expectations. *American Economic Review* 78:168–72.

——. 1989. Convergence of least squares learning mechanisms in self-referential linear stochastic models. *Journal of Economic Theory* 48:337–68.

——. 1992. The convergence of vector autoregressions to rational expectations equilibria. In *Macroeconomics: A Survey of Research Strategies* (ed. A. Vercelli and N. Dimitri), pp. 139–64. Oxford University Press.

McCafferty, S., and R. Driskill. 1980. Problems of existence and uniqueness in nonlinear rational expectations models. *Econometrica* 48:1313–17.

McKelvey, R. D., and T. Page. 1986. Common knowledge, consensus, and aggregate information. *Econometrica* 54:109–27.

Muth, J. F. 1961. Rational expectations and the theory of price movements. *Econometrica* 29:315–35.

Nyarko, Y. 1997. Convergence in economic models with Bayesian hierarchies of beliefs. *Journal of Economic Theory* 74:266–96.

——. 1998. Bayesian learning without common priors and convergence to Nash equilibrium. *Economic Theory* 11:643–56.

Royden, H. L. 1968. *Real Analysis*. New York: Macmillan.

Sandroni, A. 2000. Do markets favor agents able to make accurate predictions? *Econometrica* 68:1303–41.

——. 2005. Efficient markets and Bayes' rule. *Economic Theory* 26:741–54.

Spanos, A. 1986. *Statistical Foundations of Econometric Modeling*. Cambridge University Press.

Townsend, R. M. 1978. Market anticipations, rational expectations, and Bayesian analysis. *International Economic Review* 19:481–94.

Vives, X. 1993. How fast do rational agents learn? *Review of Economic Studies* 60:329–47.

Woodford, M. 1990. Learning to believe in sunspots. *Econometrica* 58:277–307.