



Discussion Paper Series

Strategic interactions, incomplete information and learning

By

Michele Berardi

Centre for Growth and Business Cycle Research, Economic Studies,
University of Manchester, Manchester, M13 9PL, UK

April 2011
Number 157

Download paper from:

<http://www.socialsciences.manchester.ac.uk/cgbcr/discussionpapers/index.html>

Strategic interactions, incomplete information and learning

Michele Berardi
The University of Manchester

April 18, 2011

Abstract

In a model of incomplete, heterogeneous information, with externalities and strategic interactions, we analyse the possibility of adaptive learning to act as coordination device. We build on the framework introduced by Angeletos and Pavan (2007) and extend it to a setting where agents need to learn to coordinate. We analyse conditions under which learning obtains, and show that adaptive learning can solve the problem of socially inefficient coordination.

Key words: Learning, heterogeneity, interaction, coordination.

JEL classification: C62, C73, D83.

1 Introduction

In recent years, a growing literature has been studying macroeconomic models with learning dynamics (for an authoritative treatise, see Evans and Honkapohja, 2001). The majority of these works have analyzed learning at a macro aggregate level, either with homogeneous or heterogeneous agents, by replacing the expectational operators at the semi-reduced form level that arise after the aggregation and linearization of microfounded models of agents' optimal behavior with an explicit expectations formation mechanism based on adaptive learning.

While this practice is valid to a first approximation and has indeed delivered useful insights into the properties of economic models under learning, it neglects the fact that within a macro model there is often hidden, at the micro level, a component of coordination. This tension is usually resolved by the assumption of rational expectations, which delivers a fixed point in the coordination problem. But once agents are deprived of full rationality, as it happens in the learning literature, the coordination problem becomes relevant and it might interact in interesting ways with the learning activity of agents to generate belief dynamics that ultimately affect aggregate outcomes.

A typical example is Muth's price model, where firms need to coordinate their production decisions based on the information conveyed by prices. Carton and Guse (2010) study learning in a game theoretic setting of this model, and find that adaptive learning and replicator dynamics can give rise to rather different outcomes when firms have a discrete set of possible production levels.

More in general, there are a number of macroeconomic models that lay hidden underneath a coordination problem among agents and that are built on the assumption that such a problem has been somehow solved.

The aim of the present work is to investigate the conditions under which agents can learn to coordinate. To this end, we use a setting first introduced by Morris and Shin (2002), which neatly captures the need for agents to forecast other agents' actions in order to maximize their own utility.

In a model where individual utility depends not only on a fundamental of the economy but also on the aggregate action in the population, we want to see whether agents can learn to coordinate on the best strategy without having to engage in a mental process of guessing and outguessing the actions of others but simply relying on the information observable at the aggregate level and on statistical techniques to process such information.

The answer to this question is bound to depend on the amount and quality of information available to agents. We first assume that the fundamental is observable and focus on the problem of coordination among agents. In order to play their optimal action, agents need to have some expectations of what the average action in the economy will be, and therefore they estimate it using past observations. Agents are assumed to be able to observe past average actions with one period delay: after each agent has played his action and the economy has aggregated them into

aggregate outcomes, these are observable to agents.

We then build on the global games literature and assume that the fundamental itself is not observable but agents have access to noisy private and public information on the underlying fundamental. Given this information, each agent needs to choose his optimal action, which, given the structure of the economy, also depends on the actions implemented by other agents. This framework will allow us to investigate the interactions between the problem of learning, as usually addressed in the macro literature, and that of coordination. We will show how adaptive learning can in fact act as coordination device in a model with heterogeneous information and strategic interactions. The key parameter that governs learnability will turn out to be the private value of coordination: only if agents don't overreact to the expected actions of other agents, they will be able to coordinate on an equilibrium. Interestingly, the equilibrium that emerges under learning can differ from the one that emerges with fully rational agents that behave in a strategic way.

Lastly, we will consider the issue of coordination based on a sunspot variable, one that though unrelated to fundamentals could affect the economy simply because agents deem it relevant and use it in their forecasts. We will show, though, that in the present framework agents can not learn to coordinate based on a sunspot component.

1.1 Related literature

Our contribution is related and builds on a number of works, and it merges concepts from different strains of literature. The most directly related works, in terms of the basic framework used, are Morris and Shin (2002) and Angeletos and Pavan (2007). They both introduce a general setting in which agents' best actions depend on the aggregate action in the economy, and agents must solve a coordination problem in order to maximize their utility. They find that the value agents attach to coordination is crucial in driving the dynamics of the economy. Morris and Shin (2002) famously showed that some degree of uncertainty about the fundamentals can be beneficial as it solves the problem of multiple equilibria in the economy. Angeletos, Hellwig, Pavan (2007) then extended the static framework of global games, i.e., coordination games of incomplete information, to allow agents to take actions repeatedly over many periods and to learn about the underlying fundamentals: they show that in this dynamic setting multiplicity of equilibria can emerge under the same conditions that would guarantee uniqueness in the static benchmark. We will not touch upon this aspect in the present work and only focus on the fundamental symmetric equilibrium for the economy.

The spirit of the paper is close to several works. Marimon and McGrattan (1992), in a critical review of adaptive learning in repeatedly played strategic form games, show that if agents use adaptive learning rules with inertia and experimentation, the strategy played converges to a subset of rationalizable strategies. Beggs (2009) considers adaptive learning in Bayesian games with binary actions, a framework that includes many of the applications of the theory of global games, and shows the conditions under which convergence obtains. Crawford (1995) shows how agents can learn to coordinate using simple linear adjustment rules.

We also refer to concepts from the literature on rationalizable equilibria. Guesnerie (1992) first considered the problem of how a rational expectations equilibrium can emerge as the outcome

of the mental process of iterated deletion of dominated strategies by rational agents concerned with maximizing their own utility while recognizing that all other agents in the economy are doing the same. Evans and Guesnerie (1993) then examined the connection between expectational stability (adaptive learning) and strong rationality (eductive learning) by embedding a linear rational expectations model into a game-theoretic framework.

We also marginally touch on the problem of coordination and higher order beliefs, though we leave the explicit consideration of such a problem in a contest of adaptive learning to future research. Important and related works in this area are Townsend (1983) and Marcet and Sargent (1989): in the former, firms face the problem of forecasting the forecasts of others, and this gives rise to an infinite regress problem, which is solved by Marcet and Sargent (1989) by using adaptive learning to compute the relevant equilibrium for the model.

Lastly, we build on the literature on sunspot equilibria. The possibility of an economy being driven by sunspot variables, i.e., variables unrelated to fundamentals, has received a lot of attention in the literature, at least since the works of Azariadis (1981), Cass and Shell (1983) and Guesnerie (1986). In relation to learning, the possibility of sunspot equilibria to be stable under learning dynamics has been considered in Woodford (1990), Evans and Honkapohja (1994), Evans and Honkapohja (2003) and Evans and McGough (2005). The general message that can be taken from this literature is that, though sunspot equilibria can be learnable, this usually requires rather strict conditions, and the outcome depends on the representation used by agents.

1.2 Plan of the paper

The plan of the paper is as follows: Section 2 introduces the model and shows the symmetric equilibrium under full information and rationality; Section 3 introduces learning when there is full information about the fundamental but uncertainty about other agents' actions; Section 4 analyses learning when there is incomplete and private information about the fundamental; Section 5 considers the possibility of agents using a sunspot variable to coordinate; Section 6 discusses the main results of the paper; and Section 7 concludes.

2 The model

The basic framework is borrowed from Angeletos and Pavan (2007), though we introduce time and make it dynamic. There is a finite number of agents n , and each agent i needs to choose his action k_t^i in order to maximize his own utility, which depends on an exogenous fundamental θ_t and on the actions of other agents.

The utility of each agent i is given by

$$U_t^i = U(k_t^i, K_t, \sigma_{k,t}, \theta_t) \tag{1}$$

where

$$K_t = \frac{1}{n} \sum_{j=1}^n k_t^j \quad (2)$$

$$\sigma_{k,t} = \left[\frac{1}{n} \sum_{j=1}^n (k_t^j - K_t)^2 \right]^{\frac{1}{2}} \quad (3)$$

and U is quadratic with partial derivatives $U_{k\sigma} = U_{K\sigma} = U_{\theta\sigma} = 0$ and $U_\sigma(k, K, 0, \theta) = 0$ for all (k, K, θ) . This means that the dispersion of actions in the population has only a second order, non strategic effect on individual utility. Technically, it means that utility is separable in σ . Morris and Shin (2002) model is an instance of such setting.

Since each agents chooses k_t^i in order to maximize his own utility, given his expectations of other agents' actions and of the fundamental, we have

$$k_t^i = \arg \max_k E_t^i [U(k_t^i, K_t, \sigma_{k,t}, \theta_t)] . \quad (4)$$

This gives the optimality condition

$$k_t^i : E_t^i [U'_k(k_t^i, K_t, \sigma_{k,t}, \theta_t)] = 0, \quad (5)$$

which is linear for quadratic utility function.

In the course of this work, we will consider different assumptions about the expectations formation operator in a dynamic setting. To make things operational, we will adopt a specific utility function for agents, expressed in terms of loss, that gives a setting equivalent to the beauty contest framework used by Morris and Shin (2002):

$$U_t^i = -L_t^i = -E_t^i \left[\eta (k_t^i - K_t)^2 + \beta (k_t^i - \theta_t)^2 + \gamma \sigma_k^2 \right]. \quad (6)$$

By solving agent's maximization problem we get the optimal action

$$k_t^i = \frac{\eta}{\eta + \beta} E_t^i K_t + \frac{\beta}{\eta + \beta} E_t^i \theta_t$$

or, defining $\alpha = \frac{\eta}{\eta + \beta}$,

$$k_t^i = \alpha E_t^i K_t + (1 - \alpha) E_t^i \theta_t. \quad (7)$$

Note that if agents value diversity and want to be far apart from each other in their actions, $\alpha < 0$.

2.1 Equilibrium under full information and rationality

If agents are all the same, are rational and they observe θ_t , the problem reduces to

$$k_t^* = \arg \max_k E_t [U(k_t, k_t, 0, \theta_t)] \quad (8)$$

and assuming a quadratic utility function, the solution $k_t^* = k(\theta_t)$ is linear (see Angeletos and Pavan, 2007):

$$k(\theta_t) = k_0 + k_1\theta_t \tag{9}$$

with

$$k_0 = -\frac{U_k(0,0,0,0)}{U_{kk} + U_{kK}} \tag{10}$$

$$k_1 = -\frac{U_{k\theta}}{U_{kk} + U_{kK}}. \tag{11}$$

In this case agents have no uncertainty and their optimal action depends on their own preferences (through k_0 and k_1) and on an observable exogenous component (θ_t). They can therefore implement their optimal policy (9).

Following Angeletos and Pavan (2007, Supplement), we assume that $-U_{kK}/U_{kk} < 1$, which ensures uniqueness of equilibrium under complete information. Using our loss function (6) above, this restriction corresponds to $\alpha < 1$. We also have that solution (9), under utility function (6), gives $k_0 = 0$, $k_1 = 1$, and the optimal action is therefore given by

$$k_t^* = \theta_t, \tag{12}$$

which defines the unique symmetric equilibrium. Note that (12) is the only equilibrium under complete information and rationality, for any value of $\alpha < 1$. It therefore does not matter under rationality whether agents are homogeneous or heterogeneous in their preferences (i.e., in their specific α^i), since (12) is the optimal action for all of them as long as $\alpha^i < 1$, $\forall i$. Things could be different with a more generic utility function: in that case k_0 and k_1 in (9) could depend on α^i and so differ across agents. We will neglect this complication in this work and simply focus on results under loss function (6).

3 Complete information about fundamentals and learning

We have seen above the equilibrium fixed point of the model if agents are fully rational. In particular, this requires agents i) to have knowledge about the fundamental process θ_t and to be aware of the fact that everybody else in the economy does as well; and ii) to know that everybody has the same utility function and therefore will behave alike.

In this section we maintain the hypothesis about knowledge of the fundamentals, but relax the assumption about full knowledge of others' preferences. Agents therefore need to learn about each other's actions.

In this case, therefore, agents do observe θ_t , but there is uncertainty about aggregate action K_t . It follows from (7) that the action of each agent i must satisfy the condition

$$k_t^i = (1 - \alpha)\theta_t + \alpha E_t^i K_t. \tag{13}$$

This requires agents to have expectations about K_t at each time t . Given (13), the aggregate

model for the economy is

$$K_t = \frac{\sum_{j=1}^n k_t^j}{n} = \frac{\sum_{j=1}^n (1 - \alpha) \theta_t}{n} + \frac{\sum_{j=1}^n \alpha E_t^j K_t}{n} = (1 - \alpha) \theta_t + \frac{\alpha \sum_{j=1}^n E_t^j K_t}{n}. \quad (14)$$

3.1 Adaptive expectations

Assume at first that agents have (naive) adaptive expectations, and simply take the past value as expectation for the current aggregate action: $E_t^i K_t = K_{t-1}$. Then, $\forall i$

$$k_t^i = (1 - \alpha) \theta_t + \alpha K_{t-1} \quad (15)$$

and aggregating across agents

$$K_t = (1 - \alpha) \theta_t + \alpha K_{t-1}, \quad (16)$$

which is stationary if and only if $|\alpha| < 1$, and in this case the economy converges to the full information equilibrium:

$$K_t = \theta_t.$$

This means that when agents have adaptive expectations on the aggregate action in the economy, in order for the dynamics to converge to the full information model we need to have $|\alpha| < 1$, i.e., the private value of coordination must not be too high or low when agents have adaptive expectations.

Proposition 1 *Under adaptive expectations, the economy converges to the full information symmetric equilibrium if and only if $|\alpha| < 1$.*

If agents assign great value to diversity ($\alpha < -1$), they will try to differentiate from each other and the economy will diverge away with oscillatory behavior. But also if α is high (> 1), so that agents assign great value to coordination, the economy will diverge, because agents overreact to other people's actions and thus destabilize the economy.

3.1.1 Heterogeneous preferences

We allow now agents to be heterogeneous in their preferences about the degree of coordination, i.e., agents have now different α^i . Then

$$k_t^i = (1 - \alpha^i) \theta_t + \alpha^i K_{t-1} \quad (17)$$

and the aggregate dynamics are given by

$$K_t = (1 - \bar{\alpha}) \theta_t + \bar{\alpha} K_{t-1} \quad (18)$$

where

$$\bar{\alpha} = \frac{\sum_{i=1}^n \alpha^i}{n}.$$

Stability of equilibrium therefore requires $|\bar{\alpha}| < 1$, in which case the economy converges towards the symmetric full information equilibrium $K_t = \theta_t$. This result shows that what matters for

aggregate dynamics is the average value of coordination in the population: even if some agents have extreme preferences, they do not destabilize the economy as long as average preferences satisfy the above restriction.

Proposition 2 *Under adaptive expectations and heterogeneous preferences, the economy converges to the full information symmetric equilibrium if and only if $|\frac{1}{n} \sum \alpha^i| < 1$.*

3.2 Adaptive learning

Assume now that agents form their expectations as adaptive learners, and in particular they use information about the exogenous fundamental and the past value of aggregate actions to infer information about current aggregate action according to the forecasting model, or perceived law of motion (PLM):

$$E_t^i K_t = a_t^i + b_t^i K_{t-1} + c_t^i \theta_t. \quad (19)$$

Parameters a, b, c are updated using econometric techniques such as recursive least squares (RLS), and agents use their most recent estimates to compute $E_t^i K_t$. Based on this value, they then choose k_t^i according to (13). Note that k_t^i is computed at each time t according to the anticipated utility model of Kreps (1998), i.e., taking the most recent parameter estimates as given and fixed.

Once k_t^i has been chosen, $\forall i$, the economy aggregates actions and K_t is determined. Parameters a, b, c can then be updated using standard statistical methods based on this new value for aggregate data. The question is: does $k_t^i \rightarrow k_t^*$ over time, i.e., can agents learn to coordinate on k_t^* ?

Since agents use model (19) to form expectations about K_t and then, on the basis of those expectations and the observed θ_t , decide their optimal action, k_t^i must have a (linear) representation of the form (obtained by plugging (19) into (13))

$$k_t^i = \phi_0^i + \phi_1^i \theta_t + \phi_2^i K_{t-1} \quad (20)$$

with

$$\begin{aligned} \phi_0^i &= \alpha a^i \\ \phi_1^i &= (1 - \alpha) + \alpha c^i \\ \phi_2^i &= \alpha b^i. \end{aligned}$$

Aggregating actions across agents in the economy we obtain:

$$K_t = \frac{\sum_{j=1}^n k_t^j}{n} = \alpha \frac{\sum a^j}{n} + \left((1 - \alpha) + \alpha \frac{\sum c^j}{n} \right) \theta_t + \alpha \frac{\sum b^j}{n} K_{t-1}. \quad (21)$$

Agents update parameters in their PLM (19) using forecast errors, according to the RLS

algorithm

$$\begin{bmatrix} a_t^i \\ b_t^i \\ c_t^i \end{bmatrix} = \begin{bmatrix} a_{t-1}^i \\ b_{t-1}^i \\ c_{t-1}^i \end{bmatrix} + t^{-1} R_{t-1}^{-1} w_t (K_t - E_t^i K_t) \quad (22)$$

$$R_t = R_{t-1} + t^{-1} (w_t w_t' - R_{t-1}) \quad (23)$$

with

$$w_t = \begin{bmatrix} 1 \\ K_{t-1} \\ \theta_t \end{bmatrix}$$

representing the vector of regressors and

$$K_t - E_t^i K_t = \left(\alpha \frac{\sum_j a^j}{n} - a^i \right) + \left(\alpha \frac{\sum_j b^j}{n} - b^i \right) K_{t-1} + \left((1 - \alpha) + \alpha \frac{\sum_j c^j}{n} - c^i \right) \theta_t$$

the forecast error.

Associated with the stochastic recursive algorithm (22), there is the system of differential equations

$$\dot{a} = [\alpha n^{-1} [1]_{n,n} - I_{n,n}] a \quad (24)$$

$$\dot{b} = [\alpha n^{-1} [1]_{n,n} - I_{n,n}] b \quad (25)$$

$$\dot{c} = [1 - \alpha]_{n,1} + [\alpha n^{-1} [1]_{n,n} - I_{n,n}] c \quad (26)$$

where $a = [a^1 \dots a^n]'$ is the vector of individual estimated parameters (and similarly for b and c), $I_{n,n}$ is an n by n identity matrix and $[1]_{n,n}$ is an n by n matrix of ones. The stability of each of the three systems depends on the eigenvalues of the matrix $[\alpha n^{-1} [1]_{n,n} - I_{n,n}]$ which are $\alpha - 1$ (with multiplicity 1) and -1 (with multiplicity $n - 1$). Therefore, stability of the system requires $\alpha < 1$. Remember that α is the private value of coordination: this condition says that such value must not be too high. It also implies that when agents give negative value to coordination (i.e., $\alpha < 0$), the system is stable: agents, trying to move away from each other, induce stability under adaptive learning dynamics.

Proposition 3 *Under adaptive learning, the fundamental symmetric equilibrium is learnable if and only if $\alpha < 1$.*¹

Note that α represents the private value of coordination. Our Proposition 2 says that this value has to be small: if agents value coordination too much, they overreact to their expectations of other agents' actions and the economy does not converge to the fundamental symmetric equilibrium.

Solution values for parameters are, after learning has converged and agents all have the same

¹Note that if we let the number of agents in the system grow without bound, i.e., we let $n \rightarrow \infty$, the system (24)-(26) would still have an eigenvalue equal to $\alpha - 1$ (from $\frac{\alpha}{n}n - 1$) and an infinite number of eigenvalues equal to -1 , so its stability properties would not change.

expectations:

$$\begin{aligned} a^{eq} &= 0 \\ b^{eq} &= 0 \\ c^{eq} &= 1, \end{aligned}$$

which imply that the economy converges to the fundamental symmetric equilibrium

$$K_t = \theta_t$$

since all agents implement the action $k_t^i = k_y^* = \theta_t$.

3.2.1 Heterogeneous preferences

Assume now that agents are heterogeneous in their preferences, so each agent has his own α^i . Then the system (24)-(26) becomes

$$\dot{a} = \left[n^{-1} \begin{bmatrix} \alpha^1 & \alpha^2 & \dots & \alpha^n \\ \alpha^1 & \ddots & & \\ \vdots & & & \\ \alpha^1 & \alpha^2 & \dots & \alpha^n \end{bmatrix}_{n,n} - I_{n,n} \right] a \quad (27)$$

$$\dot{b} = \left[n^{-1} \begin{bmatrix} \alpha^1 & \alpha^2 & \dots & \alpha^n \\ \alpha^1 & \ddots & & \\ \vdots & & & \\ \alpha^1 & \alpha^2 & \dots & \alpha^n \end{bmatrix}_{n,n} - I_{n,n} \right] b \quad (28)$$

$$\dot{c} = \begin{bmatrix} 1 - \alpha^1 \\ 1 - \alpha^2 \\ \vdots \\ 1 - \alpha^n \end{bmatrix}_{n,1} + \left[n^{-1} \begin{bmatrix} \alpha^1 & \alpha^2 & \dots & \alpha^n \\ \alpha^1 & \ddots & & \\ \vdots & & & \\ \alpha^1 & \alpha^2 & \dots & \alpha^n \end{bmatrix}_{n,n} - I_{n,n} \right] c. \quad (29)$$

Now each system has $n - 1$ eigenvalues equal to -1 , and one eigenvalue equal to $\frac{1}{n} \sum \alpha^i - 1$. For it to be stable therefore we need $\frac{1}{n} \sum \alpha^i < 1$. This means that we do not need all agents to value coordination in the same way, but only that on average the value they attach to coordination is small enough.

Proposition 4 *With heterogeneous α^i , adaptive learning converges if and only if $\frac{1}{n} \sum \alpha^i < 1$, i.e., the average value of coordination in the population must be less than one.*

This result shows that when preferences are heterogeneous, as long as the average value of coordination is less than one, the learning process of *all* agents converges, even for those agents that have $\alpha^i > 1$, since the evolution of other agents' expectations acts as stabilizer. This result is very important and must be stressed: learning conditions for each agent depend not on individual

preferences, but on the average value in the population, since it is this average value that governs the dynamics of the underlying variables agents are trying to learn about.

3.3 Eductive learning

Eductive learning was first introduced by Guesnerie (1992) as a way to investigate whether rational and fully informed agents could coordinate on the rational expectations equilibrium with a process of mental reasoning, that would lead them to exclude alternative outcomes thanks to the notion of rationalizable strategies. Evans and Guesnerie (1993) showed the connection between eductive learning and adaptive learning in a cobweb model.

We compare now the condition for adaptive learning with that for eductive learning in our setting. In a cobweb model it is well known that the condition for adaptive learning is $\beta < 1$, where β measures the feedback from expectations to prices, while that for eductive learning is $|\beta| < 1$. The two conditions therefore differ from each other. Gaballo (2010) recently showed how to bridge this gap by modifying the condition for rationalizability usually employed to define eductive learning.

In our setting, eductive learning requires agents to be able to coordinate on a strategy by reasoning about what would be best for other agents to do and then implement their best response to such behavior. Suppose agent i thinks that everybody else is implementing the aggregate action K_0 , then his best reply, according to (13), would be

$$k_0^i = \alpha K_0 + (1 - \alpha) \theta_t.$$

Now, since this holds for any agent i , the aggregate action that follows, K^1 , would be

$$K_1 = \alpha K_0 + (1 - \alpha) \theta_t$$

which in turn would imply a best response from each agent that would give rise to aggregate action K_2

$$K_2 = \alpha K_1 + (1 - \alpha) \theta_t.$$

This mental process defines a difference equation for the aggregate action

$$K_n = \alpha K_{n-1} + (1 - \alpha) \theta_t \tag{30}$$

which is stable for $|\alpha| < 1$, and in this case it converges to the symmetric full information equilibrium $K_t = \theta_t$. Note that the difference equation that arises under eductive learning is the same (though in "notional time", instead of real time) as the one that emerges with adaptive expectations.

Proposition 5 *Under eductive learning, the economy converges to the symmetric full information equilibrium if and only if $|\alpha| < 1$.*

In the model under consideration, therefore, eductive and adaptive learning conditions differ from each other, similarly to what happens for the cobweb model.

3.3.1 Heterogeneous preferences

Suppose now that agents are heterogeneous in their α^i . It is easy to verify that in this case educative learning would require $|\frac{1}{n} \sum \alpha^i| < 1$, i.e., the average private value of coordination must be less than one in absolute value.

Proposition 6 *Under educative learning with heterogeneity, the economy converges to the symmetric full information equilibrium if and only if $|\frac{1}{n} \sum \alpha^i| < 1$.*

This result states that also in the case of educative learning, it is sufficient that the condition for stability holds on average in the population.

4 Learning with incomplete and private information

We are now interested in understanding the problem of coordination when agents do not directly observe the fundamental process driving the economy but have to learn about it from imperfect signals. In order to decide their best strategy, agents therefore need now to form expectations about a fundamental exogenous component and about other agents' actions.

Following the literature on global games (see, e.g., Morris and Shin (2001)), we assume that agents do not observe the fundamental process θ_t but receive instead noisy private (x_t^i) and public (y_t) signals. The stochastic processes involved are:

$$\theta_t = \theta + \varepsilon_t \tag{31}$$

$$y_t = \theta_t + u_t \tag{32}$$

$$x_t^i = \theta_t + v_t^i \tag{33}$$

where ε, u, v^i are i.i.d. shocks, normally distributed with mean zero and variances $\sigma_\varepsilon^2, \sigma_u^2$ and $\sigma_{v^i}^2$ respectively. The first is a noise in the drawn made by nature at the beginning of each period to determine the fundamental, while u and v^i are observational noise in the public and private signals.

Angeletos and Pavan (2007) show in their static setting that in the case of agents not observing θ , but instead receiving a private signal x and a public signal y , agents' optimal action has a linear representation of the form

$$k(x, y) = k_0 + k_1 [(1 - \gamma)x + \gamma z] \tag{34}$$

with

$$\gamma = \delta + \frac{\alpha\delta(1-\delta)}{1-\alpha(1-\delta)}$$

where

$$z = E[\theta | y]$$

and

$$\alpha = -\frac{U_{kK}}{U_{kk}}$$

$$\delta = \frac{\sigma_y^{-2} + \sigma_\theta^{-2}}{\sigma_x^{-2} + \sigma_y^{-2} + \sigma_\theta^{-2}}.$$

Would this strategy be learnable by agents in a repeated game? In other words, can agents learn k_0 , k_1 , and γ (or at least the combination of these parameters needed to implement their optimal action)? Note that while α is a behavioral parameter, that depends on the preferences of agents, δ represents some characteristics of the economy (the variances of the various shocks), and it is rather farfetched to assume that agents know exactly these values.

In Angeletos and Pavan (2007), agents are heterogeneous in terms of their private information, but the functional form of the best response is the same for everyone. Given that private and public signals are distributed normally around θ , once learning has converged there will be a distribution of actions $k \sim N(\mu_k, \sigma_k^2)$, where

$$\mu_k = k_0 + k_1\theta$$

$$\sigma_k^2 = [k_1(1-\gamma)]^2 \sigma_x^2 + [k_1\gamma]^2 \sigma_y^2.$$

The question is: suppose agents start with a prior on the parameters in their decision rule (34), and they update it using statistical learning. Will actions converge towards a distribution $\sim N(\mu_k, \sigma_k^2)$? Here there is uncertainty about the fundamental, so agents' actions will never collapse on a point as in the previous section, but they might converge to a stationary ergodic distribution. For this to happen, agents must solve a coordination problem, i.e., agents need to learn to coordinate.

Starting from the optimality condition (13), and assuming now that also θ_t is unknown, we have that agents' optimal action can be represented as

$$k_t^i = (1-\alpha) E_t^i \theta_t + \alpha E_t^i K_t. \quad (35)$$

Agents therefore use their private (x_t^i) and a public (y_t) signals to learn about $E_t^i \theta_t$ and $E_t^i K_t$, according to the PLMs:

$$E_t^i K_t = E^i(K_t | x_t^i, y_t) = a_K^i + b_K^i x_t^i + c_K^i y_t \quad (36)$$

$$E_t^i \theta_t = E^i(\theta_t | x_t^i, y_t) = a_\theta^i + b_\theta^i x_t^i + c_\theta^i y_t \quad (37)$$

which imply, from (35)

$$k_t^i = \alpha E_t^i K_t + (1-\alpha) E_t^i \theta_t =$$

$$= \alpha (a_K^i + b_K^i x_t^i + c_K^i y_t) + (1-\alpha) (a_\theta^i + b_\theta^i x_t^i + c_\theta^i y_t).$$

Aggregating, we then obtain

$$\begin{aligned}
 K_t &= \frac{\sum_{j=1}^n k_t^j}{n} = \alpha \left(\frac{\sum_{j=1}^n a_K^j}{n} + \frac{\sum_{j=1}^n b_K^j x_t^j}{n} + \frac{\sum_{j=1}^n c_K^j}{n} y_t \right) + \\
 &+ (1 - \alpha) \left(\frac{\sum_{j=1}^n a_\theta^j}{n} + \frac{\sum_{j=1}^n b_\theta^j x_t^j}{n} + \frac{\sum_{j=1}^n c_\theta^j}{n} y_t \right) \\
 &= [\alpha \bar{a}_K + (1 - \alpha) \bar{a}_\theta] + [\alpha \bar{c}_K + (1 - \alpha) \bar{c}_\theta] y_t + \alpha \frac{\sum_{j=1}^n b_K^j x_t^j}{n} + (1 - \alpha) \frac{\sum_{j=1}^n b_\theta^j x_t^j}{n},
 \end{aligned}$$

where \bar{a} and \bar{c} indicate population averages. Since agents have private information, learning is heterogeneous and the last two terms can not be reduced down to averages. We therefore have

$$K_t = [\alpha \bar{a}_K + (1 - \alpha) \bar{a}_\theta] + [\alpha \bar{c}_K + (1 - \alpha) \bar{c}_\theta] y_t + \frac{1}{n} \sum_{j=1}^n [\alpha b_K^j + (1 - \alpha) b_\theta^j] x_t^j. \quad (38)$$

Since θ_t is exogenous, parameters in equation (37) will converge over time to their ordinary least squares estimates (i.e., conditions $E(\theta_t - E_t^i \theta_t) = 0$, $E[x_t^i (\theta_t - E_t^i \theta_t)] = 0$ and $E[y_t (\theta_t - E_t^i \theta_t)] = 0$ will hold in equilibrium):

$$a_\theta^i \rightarrow \frac{\sigma_\varepsilon^{-2}}{\sigma_\varepsilon^{-2} + \sigma_u^{-2} + \sigma_v^{-2}} E(\theta_t) := a_\theta^{eq} \quad (39)$$

$$b_\theta^i \rightarrow \frac{\sigma_v^{-2}}{\sigma_\varepsilon^{-2} + \sigma_u^{-2} + \sigma_v^{-2}} := b_\theta^{eq} \quad (40)$$

$$c_\theta^i \rightarrow \frac{\sigma_u^{-2}}{\sigma_\varepsilon^{-2} + \sigma_u^{-2} + \sigma_v^{-2}} := c_\theta^{eq}. \quad (41)$$

As for parameters in the PLM for K_t , if agents update their estimates using RLS, the evolution of parameters over time is represented by the stochastic recursive algorithm:

$$\phi_{t+1}^i = \phi_t^i + t^{-1} R_t^{-1} w_t (K_t - E_t^i K_t) \quad (42)$$

$$R_t = R_{t-1} + t^{-1} (w_t w_t' - R_{t-1}), \quad (43)$$

where

$$\phi^i = \begin{bmatrix} a_k^i \\ b_k^i \\ c_k^i \end{bmatrix}, \quad w_t = \begin{bmatrix} 1 \\ x_t^i \\ y_t \end{bmatrix}.$$

Since the PLM turns out to be misspecified with respect to the ALM, as the former depends on individual x_t^i and the latter on their population average θ_t , we need to project the PLM onto the ALM to find equilibrium values for the parameters. Using stochastic approximation theory we have

$$\begin{aligned}
 \frac{d\phi}{d\tau} &= \lim_{t \rightarrow \infty} EQ(t, \phi, z_t) \\
 Q(t, \phi, z_t) &= R_t^{-1} w_t (K_t - E_t^i K_t),
 \end{aligned}$$

where $z_t = [w_t' \theta_t]'$ and expectations are taken over the invariant distribution of z_t for fixed ϕ .

Since

$$\begin{aligned} K_t - E_t^i K_t &= [\alpha \bar{a}_K + (1 - \alpha) \bar{a}_\theta] + [\alpha \bar{c}_K + (1 - \alpha) \bar{c}_\theta] y_t + \\ &\quad + \frac{1}{n} \sum_{j=1}^n [\alpha b_K^j + (1 - \alpha) b_\theta^j] x_t^j - a_K^i - b_K^i x_t^i - c_K^i y_t, \end{aligned}$$

we have

$$\lim_{t \rightarrow \infty} EQ(\cdot) = \lim E \left[R_t^{-1} w_t \left(\begin{bmatrix} 1 & x_t^i & y_t \end{bmatrix} \begin{bmatrix} \alpha \bar{a}_K + (1 - \alpha) \bar{a}_\theta - a_K^i \\ -b_K^i \\ \alpha \bar{c}_K + (1 - \alpha) \bar{c}_\theta - c_K^i \end{bmatrix} \right) + R_t^{-1} w_t \frac{1}{n} \sum_{j=1}^n [\alpha b_K^j + (1 - \alpha) b_\theta^j] x_t^j \right].$$

By denoting

$$R^{-1} := \lim_{t \rightarrow \infty} E R_t^{-1} = \begin{bmatrix} 1 & \theta & \theta \\ \theta & \theta^2 + \sigma_\varepsilon^2 + \sigma_v^2 & \theta^2 + \sigma_\varepsilon^2 \\ \theta & \theta^2 + \sigma_\varepsilon^2 & \theta^2 + \sigma_\varepsilon^2 + \sigma_u^2 \end{bmatrix}^{-1}$$

and noting (from (42)) that b_K^j and b_θ^j are independent of w_t , we then obtain

$$\frac{d\phi^i}{d\tau} = \begin{bmatrix} \alpha \bar{a}_K + (1 - \alpha) \bar{a}_\theta - a_K^i \\ -b_K^i \\ \alpha \bar{c}_K + (1 - \alpha) \bar{c}_\theta - c_K^i \end{bmatrix} + R^{-1} [\alpha \bar{b}_K + (1 - \alpha) \bar{b}_\theta] E w_t x_t^i,$$

which, denoting $B := \lim_{t \rightarrow \infty} [\alpha \bar{b}_K + (1 - \alpha) \bar{b}_\theta]$, leads to

$$\begin{aligned} \dot{a}_K^i &= \alpha \bar{a}_K + (1 - \alpha) \bar{a}_\theta - a_K^i + BR_{11}^{-1} E x_t^i + BR_{12}^{-1} E [x_t^i x_t^i] + BR_{13}^{-1} E [y_t x_t^i] \\ \dot{b}_K^i &= -b_K^i + BR_{21}^{-1} E x_t^i + BR_{22}^{-1} E [x_t^i x_t^i] + BR_{23}^{-1} E [y_t x_t^i] \\ \dot{c}_K^i &= \alpha \bar{c}_K + (1 - \alpha) \bar{c}_\theta - c_K^i + BR_{31}^{-1} E x_t^i + BR_{32}^{-1} E [x_t^i x_t^i] + BR_{33}^{-1} E [y_t x_t^i]. \end{aligned}$$

By defining $a_K := [a_K^1 \dots a_K^n]'$ the column vector of agents' estimated parameters and $\dot{a}_K := [\dot{a}_K^1 \dots \dot{a}_K^n]'$ the column vector of the corresponding differential equations (and similarly for b and c), we then have

$$\dot{a}_K = [\alpha \bar{a}_K + (1 - \alpha) \bar{a}_\theta + \Delta_a] I_{n,n} - a_K \quad (44)$$

$$\dot{b}_K = [\Delta_b] I_{n,n} - b_K \quad (45)$$

$$\dot{c}_K = [\alpha \bar{c}_K + (1 - \alpha) \bar{c}_\theta + \Delta_c] I_{n,n} - c_K, \quad (46)$$

where I_n is the n by n identity matrix and

$$\begin{aligned} \Delta_a &= BR_{11}^{-1} \theta + BR_{12}^{-1} (\theta^2 + \sigma_v^2 + \sigma_\varepsilon^2) + BR_{13}^{-1} (\theta^2 + \sigma_\varepsilon^2) \\ \Delta_b &= BR_{21}^{-1} \theta + BR_{22}^{-1} (\theta^2 + \sigma_v^2 + \sigma_\varepsilon^2) + BR_{23}^{-1} (\theta^2 + \sigma_\varepsilon^2) \\ \Delta_c &= BR_{31}^{-1} \theta + BR_{32}^{-1} (\theta^2 + \sigma_v^2 + \sigma_\varepsilon^2) + BR_{33}^{-1} (\theta^2 + \sigma_\varepsilon^2). \end{aligned}$$

Stability of this system of ODEs is governed by the Jacobian

$$J = \begin{bmatrix} J_a & \frac{\delta \dot{a}}{\delta b} & 0 \\ 0 & J_b & 0 \\ 0 & \frac{\delta \dot{c}}{\delta b} & J_c \end{bmatrix}, \quad (47)$$

whose eigenvalues are those of the three matrices on the main diagonal

$$\begin{aligned} J_a &= [\alpha n^{-1} [1]_{n,n} - I_{n,n}] \\ J_b &= \left[\alpha \frac{\delta \Delta_b}{\delta B} n^{-1} [1]_{n,n} - I_{n,n} \right] \\ J_c &= [\alpha n^{-1} [1]_{n,n} - I_{n,n}]. \end{aligned}$$

The relevant eigenvalues for stability are therefore $\alpha - 1$ (with multiplicity 2) and $\alpha \left(\frac{\delta \Delta_b}{\delta B} \right) - 1$ with multiplicity 1.

Proposition 7 *Under incomplete private information and adaptive learning, learning dynamics converge if and only if $\alpha < 1$ and $\alpha \left(\frac{\delta \Delta_b}{\delta B} \right) < 1$.*

Since

$$\frac{\delta \Delta_b}{\delta B} = (R_{21}^{-1} \theta + R_{22}^{-1} (\theta^2 + \sigma_v^2 + \sigma_\varepsilon^2) + R_{23}^{-1} (\theta^2 + \sigma_\varepsilon^2)) > 0$$

but could be ≤ 1 , learnability thus requires $\alpha < 1$ if $\frac{\delta \Delta_b}{\delta B} < 1$, and $\alpha < \left(\frac{\delta \Delta_b}{\delta B} \right)^{-1} < 1$ if $\frac{\delta \Delta_b}{\delta B} > 1$.

Note that once learning has converged, since agents receive idiosyncratic information x_t^i , at any time t actions will be distributed normally with mean μ_k and variance σ_k^2 .

4.1 Special case: $\theta = 0$

By assuming that $\theta = 0$ in (31), which reduces the matrix R^{-1} to

$$R^{-1} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \sigma_\varepsilon^2 + \sigma_v^2 & \sigma_\varepsilon^2 \\ 0 & \sigma_\varepsilon^2 & \sigma_\varepsilon^2 + \sigma_u^2 \end{bmatrix}^{-1} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \frac{\sigma_\varepsilon^2 + \sigma_u^2}{\sigma_\varepsilon^2 \sigma_u^2 + \sigma_\varepsilon^2 \sigma_v^2 + \sigma_u^2 \sigma_v^2} & -\frac{\sigma_\varepsilon^2}{\sigma_\varepsilon^2 \sigma_u^2 + \sigma_\varepsilon^2 \sigma_v^2 + \sigma_u^2 \sigma_v^2} \\ 0 & -\frac{\sigma_\varepsilon^2}{\sigma_\varepsilon^2 \sigma_u^2 + \sigma_\varepsilon^2 \sigma_v^2 + \sigma_u^2 \sigma_v^2} & \frac{\sigma_\varepsilon^2 + \sigma_v^2}{\sigma_\varepsilon^2 \sigma_u^2 + \sigma_\varepsilon^2 \sigma_v^2 + \sigma_u^2 \sigma_v^2} \end{bmatrix},$$

we are able to obtain a new set of differential equations that replaces (44)-(46):

$$\dot{a}_K^i = \alpha \bar{a}_K + (1 - \alpha) \bar{a}_\theta - a_K^i \quad (48)$$

$$\dot{b}_K^i = \alpha \bar{b}_K + (1 - \alpha) \bar{b}_\theta - b_K^i \quad (49)$$

$$\dot{c}_K^i = \alpha \bar{c}_K + (1 - \alpha) \bar{c}_\theta - c_K^i, \quad (50)$$

and again considering the system with all agents together, we get that the relevant stability conditions depend on the three Jacobians

$$\begin{aligned} J_a &= \left[\alpha n^{-1} [1]_{n,n} - I_{n,n} \right] \\ J_b &= \left[\alpha n^{-1} [1]_{n,n} - I_{n,n} \right] \\ J_c &= \left[\alpha n^{-1} [1]_{n,n} - I_{n,n} \right]. \end{aligned}$$

The relevant eigenvalues for stability are $\alpha - 1$ (with multiplicity 3) and -1 (with multiplicity $3n - 3$), and therefore learnability requires now $\alpha < 1$.

Proposition 8 *Under incomplete private information and adaptive learning, with $\theta = 0$, learning dynamics converge if and only if $\alpha < 1$.*

We can therefore see that the condition for adaptive learning to converge is in this case the same as the one we derived under full information about the fundamental.

4.1.1 Heterogeneous preferences

Suppose now that agents are heterogeneous in their α^i . Going through the previous reasoning, only now with heterogeneous α^i , we can show that stability under learning depends on $\frac{1}{n} \sum \alpha^i$. Again, the average value of coordination has to be less than one.

Proposition 9 *Under incomplete private information and adaptive learning, with $\theta = 0$ and heterogeneous preferences, learning dynamics converge if and only if $\frac{1}{n} \sum \alpha^i < 1$, i.e., the average value of coordination is less than one.*

4.1.2 Equilibrium under learning

We are now going to show that under adaptive learning and incomplete information, parameter estimates in agents' PLMs do not converge towards the optimal values as implied by Angeletos and Pavan (2007), but towards values that would be optimal only for $\alpha = 0$, i.e., only if there was no strategic component in agents' interactions. By learning statistically, it seems therefore that agents are not able to take into account the strategic component of their interactions, and fail to coordinate on the full information equilibrium. They are nevertheless able to coordinate on a different outcome that represents an equilibrium in their learning activity.

This is an important result: note that we are allowing agents to perform their best action (as given by equation (35)) according to their utility function, and the only restriction we impose is that they don't know the exact values of the arguments in their best action equation and need to learn them from data.

Equilibrium points for the learning algorithm are resting points of the system (48)-(50). This

leads to

$$a_K^{eq} = a_\theta^{eq} = \frac{\sigma_\varepsilon^{-2}}{\sigma_\varepsilon^{-2} + \sigma_u^{-2} + \sigma_v^{-2}} E(\theta) \quad (51)$$

$$b_K^{eq} = b_\theta^{eq} = \frac{\sigma_v^{-2}}{\sigma_\varepsilon^{-2} + \sigma_u^{-2} + \sigma_v^{-2}} \quad (52)$$

$$c_K^{eq} = c_\theta^{eq} = \frac{\sigma_u^{-2}}{\sigma_\varepsilon^{-2} + \sigma_u^{-2} + \sigma_v^{-2}}. \quad (53)$$

From solution (34) for the Angeletos and Pavan (2007) model, we can work out what equilibrium parameters in agents PLMs should be in order to be consistent with that solution:

$$a_K^{AP} = a_\theta^{AP} = \frac{\sigma_\varepsilon^{-2}}{\sigma_\varepsilon^{-2} + \sigma_u^{-2} + (1-\alpha)\sigma_v^{-2}} E(\theta) \quad (54)$$

$$b_K^{AP} = b_\theta^{AP} = \frac{(1-\alpha)\sigma_v^{-2}}{\sigma_\varepsilon^{-2} + \sigma_u^{-2} + (1-\alpha)\sigma_v^{-2}} \quad (55)$$

$$c_K^{AP} = c_\theta^{AP} = \frac{\sigma_u^{-2}}{(1-\alpha)(\sigma_\varepsilon^{-2} + \sigma_u^{-2}) + \sigma_v^{-2}}. \quad (56)$$

Comparing (51-53) with (54-56), we can see that by learning adaptively from data agents neglect the source of strategic interactions and learning converges to the equilibrium that would emerge in Angeletos and Pavan (2007) with $\alpha = 0$. Under adaptive learning and incomplete information, therefore, while the private value of coordination α still determines whether convergence obtains or not, it does not affect the equilibrium for the economy.

Proposition 10 *Under incomplete private information and adaptive learning, if learning dynamics converge, the economy converges to the equilibrium characterized by $\alpha = 0$.*

5 Sunspot coordination

We now investigate whether in the framework under consideration it could be possible for agents to use a sunspot variable to coordinate their actions. Building on the literature on sunspot equilibria, we consider the possibility of agents using a sunspot variable, one that is uncorrelated with fundamentals, to gain information and help coordinate their actions.

In particular, we assume that agents still know their own preferences and so are able to determine that their optimal action is given by (35), but they now believe that an additional variable ξ_t is relevant for forecasting the fundamental θ_t and/or other agents' actions. Our aim is therefore to assess stability under learning, in a framework with strategic interactions and incomplete information, of forecasting rules that condition on an extraneous sunspot component.

If agents condition their forecasts on a sunspot component ξ_t , which is i.i.d. and independent from x_t^i , y_t and θ_t , PLMs (36)-(37) are modified as follows:

$$E_t^i K_t = E^i(K_t | x_t^i, y_t, \xi_t) = a_K^i + b_K^i x_t^i + c_K^i y_t + d_K^i \xi_t \quad (57)$$

$$E_t^i \theta_t = E^i(\theta_t | x_t^i, y_t, \xi_t) = a_\theta^i + b_\theta^i x_t^i + c_\theta^i y_t + d_\theta^i \xi_t. \quad (58)$$

Under these expectations, the temporary equilibrium for the economy would be

$$\begin{aligned}
 K_t = & [\alpha \bar{a}_K + (1 - \alpha) \bar{a}_\theta] + [\alpha \bar{c}_K + (1 - \alpha) \bar{c}_\theta] y_t + \frac{1}{n} \sum_{j=1}^n [\alpha b_K^j + (1 - \alpha) b_\theta^j] x_t^j + \\
 & + [\alpha \bar{d}_K + (1 - \alpha) \bar{d}_\theta] \xi_t
 \end{aligned} \tag{59}$$

Since θ_t is exogenous and independent of ξ_t , and the sunspot component is independent from the other regressors, it is immediate to show that over time estimates for d_θ^i in (58) would converge to zero. As for the sunspot parameter in PLM (57) for aggregate action K_t , it is easy to verify that the T-map from PLM (57) to ALM (59) for this parameter gives rise to the ODE

$$\dot{d}_K^i = \alpha \bar{d}_K + (1 - \alpha) \bar{d}_\theta - d_K^i,$$

for each agent i , where \bar{d} represents population averages. Since in equilibrium $\bar{d}_\theta = 0$, it follows that the only symmetric solution, for generic α , is $d_K^i = 0, \forall i$, and its stability requires $\alpha < 1$. This means that even if agents allow for aggregate actions to depend on an extraneous component and use such component in deciding their optimal action, they will learn over time to discard it under the same condition that ensures stability of the fundamental equilibrium.

Note that this result would carry over to a setting with heterogeneous preferences: even if agents hold different α^i , equilibrium under learning implies $d_K^i = 0, \forall i$: in this case the condition for stability under learning would be $\frac{1}{n} \sum \alpha^i < 1$.

Proposition 11 *Under incomplete information and adaptive learning, agents can not coordinate on an equilibrium with sunspots. The economy converges to the fundamental equilibrium if $\alpha < 1$ or, under heterogeneous preferences, if $\frac{1}{n} \sum \alpha^i < 1$.*

6 Discussion

The framework used to analyze the issues of learning and coordination in this paper can be interpreted as representing a number of specific economic models. For example, it could be interpreted as a model of investment complementarities, or a beauty contest economy where financial investors try to outbid each other.

Our results show that if agents have perfect information about the fundamental process, they would coordinate on the equilibrium provided a certain condition on their preferences holds. The specific condition required, though, would depend on whether they engage in higher order thinking or if instead they rely on the gathering and processing of external information in order to predict other agents' actions.

If instead agents have incomplete and private information about the fundamental, they can not rely only on an abstract process of thinking and need to collect and process information. In this case we find that adaptive learning makes them discard any strategic consideration and their actions converge to the strategic equilibrium that would emerge with $\alpha = 0$. Though the equilibrium itself is independent on the value agents attach to coordination, convergence does still depend on it.

This result is important, as it shows that if agents learn from their forecast errors, they are led to neglect strategic considerations in their actions. In all cases where a value of α different from zero is socially inefficient, this means a gain in efficiency. For example, Angeletos and Pavan (2007) show that in beauty contest economies private motives for coordination are not warranted from a social perspective, and the equilibrium that emerges under incomplete private information is inefficient. We have shown that adaptive learning can solve this problem and lead agents to an equilibrium that is efficient from a social perspective.

We have then considered the possibility of agents coordinating using a sunspot component, and we have shown that learning dynamics rule out such possibility in this contest: even if agents use an extraneous variable to try improve their forecasts, over time they learn to discard such component as irrelevant for the economy, provided the condition for learnability of the fundamental equilibrium holds.

7 Conclusions

In this paper we have considered the problem of learning and coordination for agents when their actions are strategic complements or substitutes. Under complete information about the exogenous fundamental, but uncertainty about other players' actions, both under adaptive expectations, adaptive learning and educative learning, agents can learn the fundamental, symmetric equilibrium, but specific conditions for learnability differ. In case of adaptive expectations and educative learning, the required condition is that agents do not value coordination too much or too little, because in both cases they would generate instability. In case of adaptive learning, instead, the requirement is only that agents do not value coordination too much. Adaptive learning therefore converges for a larger set of economies. In a setting with heterogeneous agents, moreover, we find that what matters for convergence is only the average characteristic of the population, in all cases.

Under incomplete and private information about the fundamental, we find that if agents update their beliefs through adaptive learning the general condition for convergence is modified, and depends on an additional component. For the special case in which the exogenous fundamental has zero mean, the condition for learnability reduces to the one we found under complete information. Interestingly, though, agents' beliefs do not converge towards the optimal values implied by the game theoretical, strategic equilibrium, but they converge instead towards values that neglect the strategic component of interactions: in situations where the private value of coordination is unwarranted from a social perspective, adaptive learning therefore leads the economy towards a more efficient equilibrium.

Finally, we have shown that sunspot components are not learnable by agents in this setting, and can not therefore enter in the solution under learning dynamics.

References

- [1] Angeletos, G.-M., Pavan, A., 2007. Efficient use of information and social value of information. *Econometrica* 75, 1103-1142.
- [2] Angeletos, G.-M., Pavan, A., 2007. Supplement to "Efficient use of information and social value of information" (*Econometrica* 75, 1103-1142). *Econometrica Supplementary Material*.
- [3] Angeletos, G.-M., Hellwig, C., Pavan, A., 2007. Dynamic global games of regime change: learning, multiplicity, and the time of attack. *Econometrica* 75, 711-756.
- [4] Azariadis, C., 1981. Self-fulfilling prophecies. *Journal of Economic Theory* 25, 380-396.
- [5] Beggs, A., 2009. Learning in Bayesian games with binary actions. *The BE Journal of Theoretical Economics - Advances*, 33.
- [6] Branch, W. A., McGough, B., 2009. A New Keynesian model with heterogeneous expectations. *Journal of Economic Dynamics and Control* 33, 1036-1051.
- [7] Carton, J., Guse, E., 2010. Replicator Dynamic Learning in Muth's Model of Price Movements. Working Papers 10-18, Department of Economics, West Virginia University.
- [8] Cass, D., Shell, K., 1983. Do sunspot matter? *Journal of Political Economy* 91, 193-227.
- [9] Crawford, V. P., 1995. Adaptive dynamics in coordination games. *Econometrica* 63, 103-143.
- [10] Evans, G. W., Guesnerie, R., 1993. Rationalizability, strong rationality, and expectational stability. *Games and Economic Behavior* 5, 632-646.
- [11] Evans, G.W., Honkapohja, S., 1994. On the local stability of sunspot equilibria under adaptive learning rules. *Journal of Economic Theory* 64, 142-161.
- [12] Evans, G.W., Honkapohja, S., 2001. *Learning and Expectations in Macroeconomics*. Princeton University Press.
- [13] Evans, G.W., Honkapohja, S., 2003. Expectational stability of stationary sunspot equilibria in a forward-looking model. *Journal of Economic Dynamics and Control* 28, 171-181.
- [14] Evans, G.W., McGough, B., 2005. Stable sunspot solutions in models with predetermined variables. *Journal of Economic Dynamics and Control* 29, 601-625.
- [15] Gaballo, G., 2010. Eductive learning matches adaptive learning. Mimeo.
- [16] Guesnerie, R., 1992. An exploration of eductive justifications of the rational expectations hypothesis. *American Economic Review* 82, 1254-78.
- [17] Guesnerie, R., 1986. Stationary sunspot equilibria in an N-commodity world, *Journal of Economic Theory* 40, 103-128.
- [18] Kreps, D., 1998. Anticipated Utility and Dynamic Choice. In D.P. Jacobs, E. Kalai, and M. Kamien, eds., *Frontiers of Research in Economic Theory*, Cambridge: Cambridge University Press, 242-74.
- [19] Marcet, A., Sargent, T.J., 1989. Convergence of least-squares learning in environments with hidden state variables and private information. *Journal of Political Economy*, 1306-1322.
- [20] Marimon, R., McGrattan, E., 1992. On adaptive learning in strategic games. *Economics Working Paper 24*, Universitat Pompeu Fabra.
- [21] Morris, S., Shin, H. S., 2001. Global games: theory and applications. *Cowles Foundation Discussion Papers 1275R*.
- [22] Morris, S., Shin, H. S., 2002. The social value of public information. *American Economic Review* 92, 1521-1534.
- [23] Townsend, R.M., 1983. Forecasting the forecasts of others. *Journal of Political Economy* 91, 546-588.
- [24] Woodford, M., 1990. Learning to believe in sunspots. *Econometrica* 58, 277-307.