

Two Problems in Applying Ljung's "Projection Algorithms" to the Analysis of Decentralized Learning*

DIEGO MORENO AND MARK WALKER

Department of Economics, University of Arizona, Tucson, Arizona 85721

1

We show that Ljung's *projection algorithms*, which have recently been used by economists to establish convergence to rational expectations equilibrium, do not seem to apply to learning or forecasting behavior that one would normally call "decentralized." If the algorithm is defined in a way that allows individuals to have differing information, then Ljung's theorem does not apply. And even if a similar theorem could be proved that would allow for differing information, there remains a Lyapunov-like condition that is central to Ljung's projection method and which requires that individual beliefs be narrowly related to the equilibrium and to one another. *Journal of Economic Literature* Classification Numbers: D83, D84, C72, D82.

In several recent studies of the way individual economic units might learn their parts of a rational expectations equilibrium, the learning process has been modeled as a recursive algorithm in which individuals' expectations depend upon their observations of variables that are influenced by others' expectations. The critical tool for establishing that the learning process converges to an equilibrium has been a theorem of Ljung [1] that gives sufficient conditions for convergence of a general class of recursive stochastic estimation algorithms. (See, in particular, Marcet and Sargent [3, 4] and Woodford [6].) It seems reasonable to hope that the same approach will also help us to understand the process by which the players in a noncooperative game might learn their parts of a Nash equilibrium, a particular kind of rational expectations equilibrium.

In the analysis of a noncooperative game, one's attention is naturally drawn to the question whether the learning process has been modeled as a "decentralized" process, in the sense that the players are allowed to have independent information and that they are allowed to act independently

* An associate editor and two referees provided helpful comments.

† E-mail: dmoreno@arizvm1.bitnet and mwalker@arizvm1.bitnet.

even if they have the same information. We show that the learning behavior described by Ljung's class of algorithms does not have this decentralized character. Thus, Ljung's approach may not be so attractive for the analysis of learning behavior in games—or indeed for the analysis of any kind of *decentralized* learning process.

Ljung's method consists essentially of two theorems, Theorems 1 and 4 in [1], which we refer to as Ljung's First and Second Theorems, respectively. Ljung's First Theorem requires one to establish that the sequence of forecasts generated by the learning algorithm contains a bounded subsequence. Because this boundedness property is generally difficult to verify, Ljung next introduces a "projection operator" according to which, whenever the algorithm yields an "obviously bad" forecast, the forecast is replaced with "something more reasonable." Ljung's Second Theorem then gives a condition that is usually easier to verify than boundedness—a kind of Lyapunov condition—which guarantees that the sequence of forecasts generated by the *projection algorithm* (but not necessarily those generated by the original algorithm) will converge to the correct forecast.

The macroeconomic applications of Ljung's method that appear in [3, 4, 6] are all applications of the Second Theorem, involving the projection operator. The idea in these applications seems to be that everyone is seeing the same data—perhaps even that a small collection of expert forecasters are generating the forecasts to be used by the participants in the economy—and that therefore the forecasts should move together. Furthermore, it is then reasonable that the forecasters will all invoke the projection operator together. But when individuals have *different* information, they will not generally arrive at "obviously bad" forecasts together, and they will therefore not invoke the projection operator together—and, therefore, as we show, Ljung's Second Theorem does not apply.

Let us try to apply Ljung's Second Theorem (to be described shortly) to a simple model in which, at each period $t = 1, 2, \dots$, each person i observes the value of a real variable $z_i(t)$ and then, based upon his observations $z_i(1), \dots, z_i(t)$, he makes a forecast $x_i(t)$ of what his next observation $z_i(t+1)$ will be. In particular, we assume that there are only two persons and that their observations and forecasts are determined by the equations

$$z(t) = Bx(t-1) + e(t), \quad (1)$$

and

$$x(t) = x(t-1) + \frac{1}{t} (z(t) - x(t-1)), \quad (2)$$

where

$$B = \begin{pmatrix} 0 & b_2 \\ b_1 & 0 \end{pmatrix},$$

and where $\{e(t)\} = \{e_1(t), e_2(t)\}$ is a sequence of independent and identically distributed random variables with expected value 0.

This model of observation and expectation formation would arise, for example, from a two-person game with linear reaction functions (such as a linear duopoly) in which $z_i(t)$ is player i 's observation of his opponent's action at time t (say $q_{-i}(t)$, perhaps observed with noise—thus, $z_i(t) = q_{-i}(t) + e_i(t)$), and in which $x_i(t)$ is player i 's expectation about his opponent's *next* action. Let player 1's reaction function be $q_1 = a_1 + b_1 q_2$ and player 2's $q_2 = a_2 + b_2 q_1$, and assume that each player i 's action $q_i(t)$ is his best response to his *forecast* $x_i(t)$ about what his opponent will do next. Then we have

$$\begin{aligned} \begin{pmatrix} z_1(t) \\ z_2(t) \end{pmatrix} &= \begin{pmatrix} q_2(t) \\ q_1(t) \end{pmatrix} + \begin{pmatrix} e_1(t) \\ e_2(t) \end{pmatrix} \\ &= \begin{pmatrix} a_2 \\ a_1 \end{pmatrix} + \begin{pmatrix} 0 & b_2 \\ b_1 & 0 \end{pmatrix} \begin{pmatrix} x_1(t-1) \\ x_2(t-1) \end{pmatrix} + \begin{pmatrix} e_1(t) \\ e_2(t) \end{pmatrix}, \end{aligned}$$

and the Nash equilibrium and the rational expectations equilibrium coincide (we ignore boundary conditions):

$$x_i^* = q_i^* = \frac{a_i + b_i a_{-i}}{1 - b_1 b_2}, \quad i = 1, 2.$$

If each player forms his forecast by averaging his past observations, i.e., if

$$x_i(t) = \frac{1}{t} \sum_{s=1}^t z_i(s),$$

and if we work with *deviations* from the equilibrium (i.e., $a_1 = a_2 = 0$), then Eqs. (1) and (2) describe the movement of expectations and observations over time, and we can ask whether the players will eventually "learn" to have correct expectations x_i^* about one another and therefore learn to take the equilibrium actions q_i^* . In other words, if the system's law of motion is given by (1) and (2), will it converge to its equilibrium?

Stochastic difference equation systems of the form given by (1) and (2) are called by Ljung "recursive stochastic estimation algorithms." His method for analyzing the asymptotic properties of these algorithms is to *avoid* direct analysis of the difference equation system, and instead to identify an associated *differential* equation. Each of his two theorems then

gives a set of conditions on the algorithm *and* its associated differential equation which ensure that, if the differential equation has an asymptotically stable stationary point, then the asymptotic behavior of the difference equation system will be the same as that of the differential equation, i.e., the difference equation system will converge (almost surely) to the differential equation's stationary point.

Applying Ljung's method to our system (1) and (2), we first substitute $z(t)$ from (1) into (2) to obtain as the transition function the following stochastic difference equation:

$$x(t) = x(t-1) + \frac{1}{t} [(B-I)x(t-1)] + \frac{1}{t} e(t), \quad (3)$$

for which Ljung's associated differential equation (ADE) is $\dot{x} = f(x) = (B-I)x$, which has $x^* = 0$ —the equilibrium—as its unique stationary point. The stationary point's domain of attraction, which we denote by \mathcal{D} , depends on the eigenvalues of the matrix $B-I$. We assume that $b_1, b_2 < 1$ (i.e., that the reaction curves "cross the right way" for best-response behavior to be stable); hence, both eigenvalues are negative and we have $\mathcal{D} = \mathbb{R}^2$.

Ljung's First Theorem requires that the forecasts lie in a bounded area infinitely often, i.e., that we can find a bounded subsequence of $\{x(t)\}$. This is usually a difficult condition to check, even for systems as simple as ours. To avoid this difficulty, Ljung provides his Second Theorem, for which he modifies the forecasting algorithms (2) into the following *projection algorithm*, which we denote by (2'):

Let D be a bounded subset of \mathbb{R}^2 with a nonempty interior;

Let C be a closed subset of the interior of D , called the algorithm's *target set*;

Let $\{p(t)\}$ be a sequence of points¹ in C ;

And instead of (2), the forecast $x(t)$ is given by

$$x(t) = \begin{cases} \bar{x}(t) & \text{if } \bar{x}(t) \in D \\ p(t) & \text{if } \bar{x}(t) \notin D, \end{cases} \quad \begin{array}{l} (2'a) \\ (2'b) \end{array}$$

where $\bar{x}(t) = x(t-1) + (1/t)(z(t) - x(t-1))$.

Ljung and Söderstrom [2] suggest that $p(t)$ can be taken to be, for instance, $x(t-1)$, and they cite Ljung [1] as providing a proof of convergence. It is clear, however, that Ljung's Second Theorem does not apply to such a projection algorithm: The rule $p(t) = x(t-1)$ does not yield an algorithm of the form (2') because it does not guarantee that $p(t)$ will lie in a closed subset of the interior of D . (The distance from any such closed subset C to the exterior of D must be positive, and therefore we might have $x(t-1) \in D \setminus C$.)

Algorithm (2') differs from (2) in that the forecasts in (2') are made to lie in a prespecified bounded set D by substituting, for any forecast $\tilde{x}(t)$ that lies outside D , a point $p(t)$ that does lie in D . (There is an important technical reason, which we describe momentarily, why the sequence $\{p(t)\}$ must be bounded away from the exterior of D .) It is natural to interpret the set D as the *support of the forecaster's beliefs*: When the forecast obtained lies outside this set, the forecaster ignores it (as an "outlier"—an atypical, perhaps mistaken, value), and uses an alternative prediction $p(t)$ instead. It is important to note, in particular, that the use of a projection algorithm in a model of learning behavior is a *behavioral assumption* and not simply an analytical device for studying convergence properties.

Ljung's Second Theorem requires, roughly, that trajectories of the ADE that start in C never leave the set D . This condition is formalized in terms of a "Lyapunov-like" function U that appears in the following condition.

Condition L. The system (1) and (2') is said to satisfy Condition L if there is a twice continuously differentiable nonnegative real valued function $U: \mathbb{R}^2 \rightarrow \mathbb{R}_+$ that satisfies

$$(L.1) \quad \forall x \in D \setminus C : U'(x)f(x) < 0, \text{ and}$$

$$(L.2) \quad \exists c_1, c_2, \text{ with } 0 \leq c_1 < c_2, \text{ such that } \forall x \in C : U(x) \leq c_1 \text{ and } \forall x \notin D : U(x) \geq c_2.$$

LJUNG'S SECOND THEOREM. *Let the ADE of a projection algorithm (such as the algorithm defined by (1) and (2')) be $\dot{x} = f(x)$, and let x^* be an asymptotically stable stationary point of f (i.e., the domain of attraction \mathcal{D} of x^* under f is an open set). If the algorithm's target set is a subset of \mathcal{D} , and if the algorithm satisfies Condition L, then for any $x(0)$ the sequence $\{x(t)\}$ converges almost surely to x^* .*

If the forecasts are described by a projection algorithm such as (2'), then at each time t each individual must choose $x_i(t) = \tilde{x}_i(t)$, or else each must choose $x_i(t) = p_i(t)$. Each person must therefore know whether or not $\tilde{x}(t) \in D$ in order to know whether he should invoke the projection operator and thereby revert to $p_i(t)$. In particular, he must know the other person's unprojected forecast $\tilde{x}_{-i}(t)$. This is clearly unsatisfactory. If our intent is to explain how individuals might attain equilibrium even when they do not start out knowing what one another will do, then we do not want to assume that they know one another's forecasts.

It is natural, then, to define a *decentralized* projection algorithm as one in which each individual has a projection operator that he invokes *independently*, whenever his *own* forecast lies outside a given set (say, the support of his own beliefs about the others' equilibrium actions). In our example, then, we would have, for $i = 1, 2$, the following, denoted by (2'')

Let D_i be a bounded subset of \mathbb{R} with a nonempty interior.

Let C_i be a nonempty closed subset of the interior of D_i .

Let $\{p_i(t)\}$ be a sequence of points in C_i .

The forecast $x_i(t)$ is given by

$$x_i(t) = \begin{cases} \tilde{x}_i(t) & \text{if } \tilde{x}_i(t) \in D_i \\ p_i(t) & \text{if } \tilde{x}_i(t) \notin D_i, \end{cases} \quad \begin{matrix} (2''a) \\ (2''b) \end{matrix}$$

where $\tilde{x}_i(t) = x_i(t-1) + (1/t)(z_i(t) - x_i(t-1))$.

Is the algorithm (2'') of the form (2'), so that Ljung's Second Theorem can be applied to it? This would require both that $D = D_1 \times D_2$ and $C = C_1 \times C_2$. But now suppose that at some time t we have, say, $\tilde{x}_1(t) \notin D_1$ and $\tilde{x}_2(t) \in D_2$. Then $\tilde{x}(t) \notin D$ and the algorithm (2') requires that $x(t) = p(t) \in C$. But according to (2''), only person 1 invokes the projection operator: $x_1(t) = p_1(t) \in C_1$, but $x_2(t)$ need not lie in C_2 . Thus, $x(t)$ will lie in the set $C_1 \times D_2$, but not necessarily in C , as algorithm (2') requires. In other words, a *decentralized* projection algorithm is not a projection algorithm of the kind to which Ljung's Second Theorem applies.³

Could we remove the set C from the definition of a projection algorithm, i.e., could we allow $C = D$, so that the definition would include decentralized projection algorithms? It seems not. The requirement in algorithm (2') that each term $p(t)$ lie in a given closed set C which is in turn contained in the interior of D (i.e., that the sequence $\{p(t)\}$ be bounded away from the exterior of D) is of fundamental importance in the proof of Ljung's Second Theorem (see [5, Remark 1]): It guarantees that if the projection algorithm were invoked infinitely many times, $\{x(t)\}$ would have a subsequence in the (compact) set C and would therefore have a limit point in the interior of D . This interior limit point plays a crucial role in proving that, if Condition L is satisfied, the projection operator could in fact only be invoked finitely many times, so that an algorithm of the form (2') will eventually be of the form (2'a), and will therefore behave asymptotically like its ADE. Existence of such an interior limit point is not guaranteed for a decentralized algorithm of the form (2''), and it is not clear how the proof of Ljung's Second Theorem could be altered to include decentralized algorithms.

There is a second difficulty in applying Ljung's Second Theorem to decentralized forecasting problems. Suppose that, after all, we could prove a version of the Second Theorem that would cover decentralized projection algorithms. Now consider the implications of Condition L. This condition

³ Marcet and Sargent [4] mistakenly apply Ljung's Second Theorem to a projection operator of the form (2''). See Remark 1 in Moreno and Walker [5] for an analysis of the obstacles to adapting the theorem's proof to operators of the form (2'').

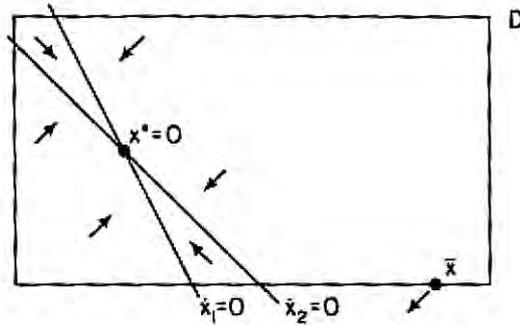


FIGURE 1

effectively requires that the trajectories of the ADE, when starting in C , never leave the set D ; if we choose C_i very close to D_i , the trajectories of the ADE at the boundary of D must point inward. To see the implications of this assumption for our example, suppose that the sets D_1 and D_2 are bounded intervals in \mathbb{R} . Figure 1 gives the phase diagram of the ADE.

The lines $\dot{x}_1 = 0$ and $\dot{x}_2 = 0$ are the reaction functions of players 1 and 2, respectively, in the two-person game if b_1 and b_2 are both negative and the stability condition $b_1 b_2 < 1$ is satisfied. At point \bar{x} , on the boundary of D , the ADE will follow a trajectory that will drive the system out of the set D ; therefore, Condition L is not satisfied here. It is only when the upper left and lower right corners of D lie between the graphs of the functions $\dot{x}_1 = 0$ and $\dot{x}_2 = 0$, as in Fig. 2, that Condition L is satisfied. If D_1 and D_2 are the real intervals (α_1, β_1) and (α_2, β_2) , then Condition L requires that the following inequalities be satisfied (when b_1 and b_2 are both negative):

$$b_2 \beta_1 < \alpha_2 < \frac{1}{b_1} \beta_1 \quad \text{and} \quad \frac{1}{b_1} < \beta_2 < b_2 \alpha_1.$$

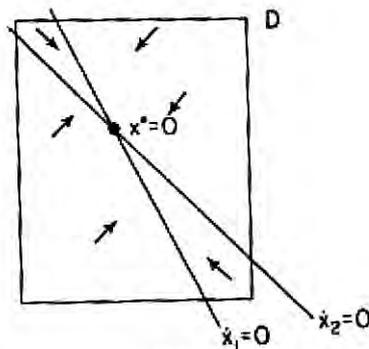


FIGURE 2

By imposing Condition L, then, we are effectively requiring that the individual beliefs have supports that are not independent of one another. If each firm arrives at its support D_i independently, Condition L will not generally be satisfied.

An alternative interpretation is that Condition L is a kind of *consistent beliefs condition*, similar to the consistency of prior beliefs that plays such an important role in justifying Nash equilibrium via common-knowledge arguments. But the primary reason for studying models of learning is to explain the attainment of equilibrium when individuals do *not* have full information about one another. Therefore Condition L, and Ljung's projection operators, do not seem to be appropriate for modeling *decentralized* processes of decision making and expectation formation.

It may be, however, that decentralized projection operators such as (2^o) are in fact good descriptions of the way individuals form expectations. If so, it is an important question for further research whether there is a decentralized analogue of Ljung's Second Theorem, a theorem that will give us conditions under which *decentralized* recursive stochastic estimation algorithms converge.

REFERENCES

1. L. LJUNG, Analysis of recursive stochastic algorithms, *IEEE Trans. Automatic Control* AC-22 (1977), 551-575.
2. L. LJUNG AND T. SÖDERSTRÖM, "Theory and Practice of Recursive Identification," MIT Press, Cambridge, MA, 1983.
3. A. MARCET AND T. SARGENT, Convergence of least-squares learning mechanisms in self-referential linear stochastic models, *J. Econ. Theory* 48 (1989), 337-368.
4. A. MARCET AND T. SARGENT, Convergence of least-squares learning in environments with hidden state variables and private information, *J. Polit. Econ.* 97 (1989), 1306-1322.
5. D. MORENO AND M. WALKER, Convergence theorems for a class of recursive stochastic algorithms, in "The Economics of Informational Decentralization: Complexity, Efficiency, and Stability" (J. Ledyard, Ed.), Kluwer Academic Publishers, Amsterdam, in press.
6. M. WOODFORD, Learning to believe in sunspots, *Econometrica* 58 (1990), 277-307.