SELF-CONFIRMING EQUILIBRIUM

IN-KOO CHO AND THOMAS J. SARGENT

1. Introduction

A self-confirming equilibrium is the answer to the following question: what are the possible limiting outcomes of purposeful interactions among a collection of adaptive agents, each of whom averages past data to approximate moments of the conditional probability distributions of interest? If outcomes converge, a Law of Large Numbers implies that agents' beliefs about conditional moments become correct on events that are observed sufficiently often. Beliefs are not necessarily correct about events that are infrequently observed. Where beliefs are correct, a self-confirming equilibrium is like a rational expectations equilibrium. But there can be interesting gaps between self-confirming and rational expectations equilibria where beliefs of some important decision makers are incorrect.

Self-confirming equilibria interest macroeconomists because they connect to an influential 1970s argument made by Christopher Sims that advocated rational expectations as a sensible equilibrium concept. This argument defended rational expectations equilibria against the criticism that they require that agents 'know too much' by showing that we do not have to assume that agents start out 'knowing the model'. If agents simply average past data, perhaps conditioning by grouping observations, their forecasts will eventually become unimprovable.

Research on adaptive learning has shown that that the glass is 'half-full' and 'half-empty' for this clever 1970s argument. On the one hand, the argument is correct when applied to competitive or infinitesimal agents: by using naive adaptive learning schemes (various versions of recursive least squares), agents can learn every conditional distribution that they require to play best responses within an equilibrium. On the other hand, large agents (e.g., governments in macro models) who can influence the market outcome cannot expect to learn everything that they need to know to make good decisions: in a self-confirming equilibrium, large agents may base their decisions on conjectures about off-equilibrium-path behaviors which turn out to be incorrect. Thus, a rational expectations equilibrium is a self-confirming equilibrium, but not *vice versa*.

While agents' beliefs can be incorrect off the equilibrium path, the self-confirming equilibrium path still restricts them in interesting ways. For macroeconomic applications, the government's model must be such that its off-equilibrium path beliefs rationalize the decisions (its Ramsey policy or Phelps policy, in the language of [16]) that are revealed

Date: December 7, 2006.

Financial support from the National Science Foundation (ECS-0523620) and Deutsche Bank through Institute for Advanced Study is gratefully acknowledged. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation, Deutsche Bank or Institute for Advanced Study.

along the equilibrium path. The restrictions on government beliefs required to sustain self-confirming equilibria have only begun to be explored in macroeconomics, mainly in the context of some examples like those in [16]. Analogous restrictions have been more thoroughly analyzed in the context of games [10].

The freedom to specify beliefs off the equilibrium path makes the set of self-confirming equilibria generally be larger than the set of Nash equilibria, which often admit unintuitive outcomes in extensive form games [11]. A widely used idea of refining self-confirming equilibrium is to embed the decision making problem within a learning process in which decision makers estimate unknown parameters through repeated interactions, and then to identify a stable stationary point of the learning dynamics (e.g., [8, 7, 6]).

The gap between a self-confirming equilibrium and a rational expectations equilibrium can be vital for a government designing a Ramsey plan, for example, because its calculations necessarily involve projecting outcomes of counterfactual experiments. For macroeconomists, an especially interesting feature of self-confirming equilibria is that because a government can have a model that is wrong off the equilibrium path, policy that it thinks is optimal can very well be far from optimal. Even if a policy model fits the historical data correctly and is unimprovable, one cannot conclude that the policy is optimal. As a result, it requires an entirely *a priori* theoretical argument to diminish the influence of a good fitting macroeconomic model on public policy [16].

2. Formal Definitions

An agent i is endowed with strategy space A_i and state space X_i. Generic elements of A_i and X_i are called a strategy and a state, respectively. A probability distribution $\mathcal{P}_i$ over $A_i \times X_i$ describes how actions and states are related. A utility function is $u_i : A_i \times X_i \to \mathbb{R}$. Let $\mu_i(\cdot : a_i)$ be a probability distribution over X_i, which represents i's belief about the state conditioned on action a_i. Agent i's decision problem is to solve

$$\max_{a_i \in A_i} \int_{X_i} u_i(a_i, x_i) d\mu_i(x_i : a_i). \tag{2.1}$$

2.1. Single Person Decision Problems.

Here a self-confirming equilibrium is a simply a pair (a_i^*, μ_i^*) satisfying

$$a_i^* \in \arg\max_{a_i \in A_i} \int_{X_i} u_i(a_i, x_i) d\mu_i^*(x_i : a_i) \tag{2.2}$$

$$\mu_i^*(x_i : a_i^*) = \mathcal{P}_i(x_i : a_i^*). \tag{2.3}$$

(2.2) implies that the choice must be optimal given his subjective belief μ_i^*, while (2.3) says that the belief must be confirmed, conditioned on his equilibrium action a_i^*. Self-confirming equilibrium has the two key ingredients of rational expectations equilibrium: optimization and self-fulfilling property. The key difference is (2.3), which imposes a self-confirming property conditioned only on equilibrium action a_i^*. The decision maker can entertain $\mu_i(\cdot : a_i) \neq \mathcal{P}(\cdot : a_i)$, conditioned on $a_i \neq a_i^*$. In this sense, the agent can have multiple beliefs about the state conditioned on his own action.

If we strengthen things to require

$$\mu_i(\cdot : a_i) = \mathcal{P}(\cdot : a_i) \qquad \forall a_i, \tag{2.4}$$

then we attain a rational expectations equilibrium. As will be shown later, (2.4) is called the unitary belief condition [10], which is one of the three key features that distinguishes a self-confirming equilibrium from a rational expectations equilibrium or Nash equilibrium.

2.2. **Multi-person Decision Problems.** If we interpret the state space as the set of the strategies of the other players $X_i = A_{-i}$, we can naturally extend the basic definition to the situation where more than one person is making a decision. A self-confirming equilibrium is a profile of actions and beliefs, $\{(a_1^*, \mu_1^*), \ldots, (a_n^*, \mu_n^*)\}$ such that (2.2) and (2.3) hold for every $i = 1, \ldots, n$. As we move from single person to multi-person decision problems, however, (2.3) differs three ways from a Nash equilibrium, in addition to the unitary belief condition (2.4). (1) If there are more than two players, the belief of player $i \neq j, k$ about player k's strategy can be different from player j's belief about player k's strategy (failure of consistency). (2) Similarly, player i can entertain the possibility that player j and player k correlate their strategies according to an un-modeled randomization mechanism, leading to correlated beliefs. (3) If we require that a self-confirming equilibrium should admit unitary and consistent beliefs, while excluding correlated beliefs, then the self-confirming equilibrium is a Nash equilibrium [10].

2.3. **Dynamic Decision Problems.** Suppose that player i solves (2.1) repeatedly. The first step to embed self-confirming equilibria in dynamic contexts is to spell out learning rules that specify how beliefs respond to new observations. We define a learning rule as a mapping that updates belief μ_i into a new belief when new data arrive. Define $Z_i \subset X_i$ as a subspace of X_i that is observed by a decision maker. Let $\mathcal{M}_i$ be the set of probability distributions over $Z_i \times A_i$. These represent player i's belief about the state, i.e., the model entertained by player i. A learning rule is defined as

$$\mathcal{T}_i : \mathcal{M}_i \times Z_i \to \mathcal{M}_i.$$

A belief $\mu_i^* \in \mathcal{M}_i$ is a steady state of the learning dynamics if

$$a_i^* \in \arg\max_{a_i \in A_i} \int_{x_i} u_i(a_i, x_i) d\mu_i^*(x_i : a_i) \tag{2.5}$$

$$\mu_i^* = \mathcal{T}_i(\mu_i^*, z_i) \tag{2.6}$$

for every z_i in the support of $\mathcal{P}_i(x_i : a_i^*)$.

The steady state of learning dynamics is a self-confirming equilibrium for a broad class of recursive learning dynamics including Bayesian [10] and least square learning algorithms [16, 6] are self-confirming equilibria.

2.4. **Refinements.** We can study the salience of the self-confirming equilibrium by examining the stability of the associated steady states. The stability property provides a useful foundation for selecting a sensible self-confirming equilibrium [2, 12, 17, 6]. With a possible exception of the Bayesian learning algorithm, most *ad hoc* learning rules are motivated by the simplicity of some updating scheme as well as its ability to support sufficiently sophisticated behavior in the limit. By exploiting the convergence properties of learning dynamics, we can often devise a recursive algorithm to *calculate* a self-confirming equilibrium, i.e., a fixed point of the $\mathcal{T}_i$. This approach to computing equilibria has occasionally proved fruitful to compute equilibria in macroeconomics (e.g., [1]).

In principle, a player need not know the other player's payoff in order to play a self-confirming equilibrium. Self-confirming equilibria allow a player to entertain any belief conditioned on actions not used in the equilibrium. This is one of the main sources of multiplicity. In the game theoretic context, a player can delineate the set of possible actions of the other players, even if he does not have perfect foresight. If each player knows the payoff of the other players, and if it is common knowledge that every player is rational, then a player can eliminate the actions of the other players that cannot be rationalized. By exploiting the idea of sophisticated learning of [13, 14, 5] restricted the set of possible beliefs off the equilibrium path to eliminate evidently unreasonable self-confirming equilibrium.

3. Applications

Self-confirming equilibria and recursive learning algorithms are powerful tools to investigate a number of important dynamic economic problems such as (1) the limiting behavior of learning systems [6, 9]; (2) the selection of plausible equilibria in games and dynamic macroeconomic models [12, 11]; (3) the incidence and distribution of rare events that occasionally arise as large deviations from self-confirming equilibria [4, 16]; and (4) formulating plausible models of how agents respond to model uncertainty [3].

Remarkably, related mathematics tie together all of these applications. The mean dynamics that propel the learning algorithms to self-confirming equilibria (in item 1) are described by ordinary differential equations (ODE) derived through an elegant stochastic approximation algorithm (e.g., [9, 17, 12]). Because the stationary point of the ODE is a self-confirming equilibrium, the stability of the ODE determines the selection criterion used to make statements about item 2 [9, 6]. Remarkably, by adding an adverse deterministic shock to that same ODE, we obtain a key object that appears in a *deterministic* control problem that identifies the large-deviation excursions in item 3 *away* from a self-confirming equilibrium [4]. Finally, that same large deviations mathematics is associated with robust control ideas that use entropy to model how agents cope with model uncertainty [15].

References

[1] Rao Aiyagari, Albert Marcet, Thomas J. Sargent, and Juha Seppälä. Optimal taxation without state-contingent debt. *Journal of Political Economy*, 110(6):1220–1254, 2002.
[2] Margaret M. Bray. Learning, Estimation and the Stability of Rational Expectations. *Journal of Economic Theory*, 26:318–339, 1982.
[3] In-Koo Cho and Ken Kasa. Learning and Model Validation. University of Illinois and Simon Fraser University, 2006.
[4] In-Koo Cho, Noah Williams, and Thomas J. Sargent. Escaping Nash Inflation. *Review of Economic Studies*, 69:1–40, 2002.
[5] Eddie Dekel, Drew Fudenberg, and David K. Levine. Payoff Information and Self-Confirming Equilibrium. *Journal of Economic Theory*, 89:165–185, 1999.
[6] George W. Evans and Seppo Honkapohja. *Learning and Expectations in Macroeconomics*. Princeton University Press, 2001.
[7] Drew Fudenberg and David M. Kreps. Learning Mixed Equilibria. *Games and Economic Behavior*, 5:320–367, 1993.
[8] Drew Fudenberg and David M. Kreps. Learning in Extensive Games, I: Self-Confirming and Nash Equilibrium. *Games and Economic Behavior*, 8:20–55, 1995.
[9] Drew Fudenberg and David Levine. *Learning in Games*. M.I.T. Press, 1998.
[10] Drew Fudenberg and David K. Levine. Self-Confirming Equilibrium. *Econometrica*, 61:523–545, 1993.
[11] David M. Kreps. Anticipated Utility and Dynamic Choice. In *Frontiers of Research in Economic Theory: The Nancy L. Schwartz Memorial Lectures, 1983-1997*. Cambridge University Press, 1998.
[12] Albert Marcet and Thomas J. Sargent. Convergence of Least Squares Learning Mechanisms in Self Referential Linear Stochastic Models. *Journal of Economic Theory*, 48:337–368, 1989.
[13] Paul Milgrom and D. John Roberts. Rationalizability, Learning and EQuilibrium in Games with Strategic Complementarities. *Econometrica*, 58(6):1255–1277, 1990.
[14] Paul Milgrom and D. John Roberts. Adaptive and Sophisticated Learning in Normal Form Games. *Games and Economic Behavior*, 3(1):82–100, 1991.
[15] C. Pandit and S. P. Meyn. Worst-Case Large-Deviations with Application to Queueing and Information Theory. *Stochastic Process and their Applications*, 116(5):724–756, May 2006.
[16] Thomas J. Sargent. *The Conquest of American Inflation*. Princeton University Press, 1999.
[17] Michael D. Woodford. Learning to Believe in Sunspots. *Econometrica*, 58:277–307, 1990.

Department of Economics, University of Illinois, 1206 S. 6th Street, Champaign, IL 61820 USA

E-mail address: `inkoocho@uiuc.edu`

URL: `http://www.business.uiuc.edu/inkoocho`

Department of Economics, New York University, 269 Mercer Street, New York, NY 10003 USA

E-mail address: `ts43@nyu.edu`

URL: `http://homepages.nyu.edu/~ts43`