Dynamic Foundations for Empirical Static Games[*]

Niccolò Lomys[†]   Lorenzo Magnolfi[‡]   Camilla Roncoroni[§]

July 16, 2018

Abstract

We propose a simple estimation strategy when data on strategic interaction are interpreted as the long-run result of a history of game plays. Players interact repeatedly in an incomplete information game, possibly while learning how to play in such a game. We remain agnostic on the details of the learning process and only impose a minimal behavioral assumption describing an optimality condition for the long-term outcome of players' interaction. In particular, we assume that play satisfies a property of "asymptotic no regret" (ANR). This property requires that the time average of the counterfactual increase in past payoffs, had different actions been played, becomes approximately zero in the long run. A large class of well-known algorithms for the repeated play of the incomplete information game satisfies the ANR property. We show that, under the ANR assumption, it is possible to partially identify the structural parameters of players' payoff functions. We establish our result in two steps. First, we prove that the time average of play that satisfies ANR converges to the set of Bayes correlated equilibria of the underlying static game. To do so, we extend to incomplete information environments prior results on dynamic foundations for equilibrium play in static games of complete information. Second, we show how to use the limiting model to obtain consistent estimates of the parameters of interest.

Keywords: Empirical Games; Incomplete Information; Bayes Correlated Equilibrium; Learning in Games; No Regret; Partial Identification; Incomplete Models.

JEL Classification: C57; C70; L10.

1 Introduction

When data are generated by strategic interaction, the analyst needs to specify a solution concept to interpret players' outcomes and leverage the observables to perform identification of the game's primitives. It is common practice to assume that equilibrium play is observed. This assumption is justified whenever players are able to form correct expectations[1] on the strategic environment, and to (optimally) behave accordingly. However, for many real-world strategic environments, it is not obvious that behavior satisfies these requirements: players may need to learn how to play. An external analyst needs then to determine when the learning phase terminates, and whether it is possible to interpret subsequent behavior as the result of equilibrium play.

In this paper, we explicitly allow for the possibility that players are learning how to interact. However, we remain agnostic on the details of the learning process. Moreover, instead of assuming that players reach the ability to choose correct strategies at some given point in time, we only impose a minimal behavioral assumption describing an optimality condition for the long-term result of players' interaction. More specifically, we model players as interacting repeatedly, playing an incomplete information game, and assume that long-run outcomes satisfy a property of "asymptotic no regret" (hereafter, ANR). The ANR property requires that the time average of the counterfactual increase in past payoffs, had different actions been played, becomes approximately equal to zero in the long run. Intuitively, no matter which specific learning rule players are adopting, we assume that they eventually eliminate the regret of not having played differently in the past. After having imposed that players' behavior satisfies the ANR property, we derive the implications for identification of the game's primitives.

We show that, under the ANR assumption, it is possible to partially identify the structural parameters of players' payoff functions. We do so in two steps. First, we show that the time average of play that satisfies ANR converges to the set of Bayes correlated equilibria of the underlying static game (although, in general, it does not converge to a particular point in this set).[2] To establish this property we extend to games of incomplete information prior results on dynamic foundations for equilibrium play in static games of complete information. Second, we show how to use the limiting model to obtain consistent estimates of the parameters of interest. Our empirical approach is based on a behavioral assumption that has only implications on the predictions of a limiting model, as opposed to the full data generating process. Our approach gives rise to non-standard econometric issues, as it is not possible to fully characterize a single limit distribution of the observables, similarly to Epstein, Kaido and Seo (2016), but only the set it belongs to. Yet, we show that we can use the limiting model to obtain a consistent estimator for the parameters of interest. Since behavior is not specified, our model is incomplete in the sense of Tamer (2003).

The ANR property is weaker than the one-shot no-ex post-regret property of pure-strategy Nash equilibrium that is sometimes invoked to motivate the choice of modeling cross-sectional data as equilibrium outcomes of a static game. Indeed, this descriptive[3] interpretation of static models is often paired with the assumptions of complete information and pure-strategy Nash

---

[1]Given their information sets, which are assumed to be correctly defined by the analyst.

[2]Bayes Correlated Equilibrium is a generalization of Correlated Equilibrium to incomplete information environments developed by Bergemann and Morris (2016).

[3]In the sense of Pakes (2016)

equilibrium. The rationale for these assumptions is that the no-ex post-regret property of pure-strategy Nash equilibria reflects the stable nature of long-run outcomes.[4] Although appropriate for some environments, the static notion of no-ex post regret is a strong requirement: our work is thus complementary to standard equilibrium models of strategic interaction and provides an alternative whenever Nash equilibrium does not represent an appropriate restriction on behavior. In fact, Nash equilibrium of the static game is neither a natural long-run outcome of many simple game dynamics (for a review, see Hart and Mas-Colell (2013)), nor easy to compute in large games (Babichenko and Rubinstein (2017)).

In contrast, the ANR property is satisfied by a large class of well-known algorithms for the repeated play of the underlying one-shot game, once they are extended to games of incomplete information. This class includes simple adaptive heuristics, fictitious-play-like dynamics, more sophisticated learning rules involving active experimentation, calibrated learning, and several equilibrium dynamics. Since we do not fully specify what the behavior of players is or what they do to play according to this minimal long-run requirement, we depart from the current literature on empirical dynamic games that typically imposes the Markov perfect (or related) solution concept (for a review, see Ackerberg, Benkard, Berry and Pakes (2007)).

Related Literature. Our work is related to the literature on learning in games, especially to Hart and Mas-Colell (2000, 2013), whose convergence results we extend to incomplete information environments. In contrast to these authors, the emphasis of our work is on connecting learning dynamics to the inference problem of an external observer. Recent contributions in computer science offer both related theoretical results (Hartline, Syrgkanis and Tardos (2015)), and connection to empirical work (Nekipelov, Syrgkanis and Tardos (2015)), although the latter is specialized to online auction environments. Instead, we consider a general model where the primitive to be recovered is not the payoff type of players, but rather the structural features of payoffs, in line with the econometric literature on empirical games.

We are not the first to leverage on results in the literature on learning in games to perform empirical analysis. Lee and Pakes (2009) develop a learning-based procedure to compute counterfactuals in dynamic games. Several recent advances in the estimation of dynamic games investigate tractable and less restrictive empirical models (e.g., Doraszelski, Lewis and Pakes (2018)). Our paper proposes a valid descriptive approach that is complementary to these structural methods.

Magnolfi and Roncoroni (2017) also consider estimation of discrete games under the assumption of Bayes correlated equilibrium behavior. Although this paper proposes the use of a similar estimation technique, the motivation is very different. In fact, Magnolfi and Roncoroni (2017) exploit the link between equilibrium behavior and information to establish that Bayes correlated equilibrium allows to estimate static discrete games under minimal assumptions on information. In this paper, instead, we motivate the use of Bayes correlated equilibrium as a behavioral restriction when the data are generated by repeated interaction.

---

[4]For instance, Ciliberto and Tamer (2009) argue as follows: "The idea behind cross-section studies is that in each market, firms are in a long-run equilibrium.

References

Ackerberg, Daniel, C. Lanier Benkard, Steven Berry, and Ariel Pakes (2007), "Econometric Tools for Analyzing Market Outcomes." James J. Heckman and Edward E. Leamer eds., *Handbook of Econometrics*. Elsevier.

Babichenko, Yakov and Aviad Rubinstein (2017), "Communication Complexity of Approximate Nash Equilibria." *STOC 2017*, 878–889.

Bergemann, Dirk and Stephen Morris (2016), "Bayes Correlated Equilibrium and the Comparison of Information Structures in Games." *Theoretical Economics*, 11, 487—-522.

Ciliberto, Federico and Elie Tamer (2009), "Market Structure and Multiple Equilibria in Airline Markets." *Econometrica*, 77, 1791–1828.

Doraszelski, Ulrich, Gregory Lewis, and Ariel Pakes (2018), "Just Starting Out: Learning and Equilibrium in a New Market." *The American Economic Review*, 108, 565–615.

Epstein, Larry, Hiroaki Kaido, and Kyoungwon Seo (2016), "Robust Confidence Regions for Incomplete Models." *Econometrica*, 84, 1799–1838.

Hart, Sergiu and Andreu Mas-Colell (2000), "A Simple Adaptive Procedure Leading to Correlated Equilibrium." *Econometrica*, 68, 1127–1150.

——— (2013), "Simple Adaptive Strategies: From Regret-Matching to Uncoupled Dynamics." *World Scientific Series in Economic Theory*.

Hartline, Jason, Vasilis Syrgkanis, and Eva Tardos (2015), "No-Regret Learning in Bayesian Games." *NIPS 2015*.

Lee, Robin S. and Ariel Pakes (2009), "Multiple Equilibria and Selection by Learning in an Applied Setting." *Economics Letters*, 104, 13–16.

Magnolfi, Lorenzo and Camilla Roncoroni (2017), "Estimation of Discrete Games with Weak Assumptions on Information." *Working Paper*.

Nekipelov, Denis, Vasilis Syrgkanis, and Eva Tardos (2015), "Econometrics for Learning Agents." *EC 2015*.

Pakes, Ariel (2016), "Methodological Issues in Analyzing Market Dynamics." Frank Thuijsman and Florian Wagener eds., *Advances in Dynamic and Evolutionary Games: Theory, Applications, and Numerical Methods*. Springer International Publishing.

Tamer, Elie (2003), "Incomplete Simultaneous Discrete Response Model with Multiple Equilibria." *The Review of Economic Studies*, 70, 147–165.