

Dynamic Foundations for Empirical Static Games*

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

February 15, 2020

Abstract

We develop an approach to identification and estimation of a game’s primitives when players interact repeatedly in a one-shot incomplete information game and the only restriction on behavior is an asymptotic no regret (hereafter, ANR) condition. 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. Well-known algorithms for the repeated play of a one-shot incomplete information game satisfy the ANR property. Under the ANR assumption, we (partially) identify the structural parameters of the one-shot game. We establish our result in two steps. First, we prove that the empirical distribution of play that satisfies ANR converges to the set of Bayes (coarse) correlated equilibrium predictions of the underlying one-shot game. To do so, we generalize 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. We apply our method to data on pricing behavior in an online platform.

Keywords: Empirical Games; Incomplete Information; Bayes (Coarse) Correlated Equilibrium; Learning in Games; No Regret; Partial Identification; Incomplete Models; Robust Predictions.

JEL Classification: C57; C70; L10.

*We thank Dirk Bergemann for excellent feedback that greatly improved the paper. We had valuable conversations with Tilman Börgers, Giacomo Calzolari, Mira Frick, Ryota Iijima, Rida Laraki, Alessandro Lizzeri, Juan Ortner, Larry Samuelson, and Eran Shmaya. We received helpful comments from seminar audiences at the University of Mannheim, EEA–ESEM 2019, the 2019 North American Winter Meeting of the Econometric Society, and the 2018 Lisbon Meetings in Game Theory and Applications. The first author acknowledges funding from the European Research Council Grant N^o 714147, and from the ANR under grant ANR-17-EURE-0010 (Investissements d’Avenir program).

[†]Toulouse School of Economics, University of Toulouse Capitole; niccolo.lomys@tse-fr.eu.

[‡]Department of Economics, University of Wisconsin-Madison; magnolfi@wisc.edu.

[§]Department of Economics, University of Warwick; c.roncoroni@warwick.ac.uk.

1 Introduction

When data are generated by strategic interaction, an external 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 expectations¹ 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.

In this paper, we explicitly allow for the possibility that players may need time to adapt to the strategic environment. Instead of assuming that players choose equilibrium strategies at each point in time, we only impose a behavioral assumption describing a minimal optimality condition for the long-run outcome of players' interaction. More specifically, we model players as playing repeatedly a one-shot incomplete information game and assume that the long-run outcome of their interaction satisfies 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 arbitrarily small in the long run. Intuitively, no matter how players are adapting, we assume that they eventually eliminate the regret for not having played differently in the past. After having imposed that players' behavior satisfies the ANR property, we derive the implications for the identification of the game's primitives.

We show that, under the ANR assumption, it is possible to partially identify the structural parameters of the one-shot incomplete information game. We do so in two steps. First, we show that the empirical distribution (i.e. the time average) of actions, signals, and states that satisfies ANR converges to the set of Bayes correlated equilibria (Bergemann and Morris (2016)), or its coarse analog—depending on the specific regret notion we consider—of the underlying one-shot game (although, in general, converge is not to a particular point in this set). To establish this property, we generalize to games of incomplete information prior results on dynamic foundations for equilibrium play in static games of complete information (Hart and Mas-Colell (2013) and references therein). 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 for 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, 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 satisfied by a large class of well-known algorithms for the repeated play of the underlying one-shot game, once they are appropriately extended to

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

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 [Akerberg, Benkard, Berry and Pakes \(2007\)](#)).

The ANR property is weaker than the static 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 (in the sense of [Pakes \(2016\)](#)) interpretation of static models is often paired with the assumptions of complete information and pure-strategy Nash 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.² 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.

We apply our method to data on pricing behavior by sellers in a decentralized online platform. In this environment, a large population of sellers interacts strategically over time. These sellers may (but do not need to) display a limited degree of sophistication. Hence, relying on the ANR assumption for sellers’ price setting behavior is an appealing approach to recover the distribution of sellers’ marginal costs under weak assumptions on behavior. The large number of sellers in this environment generates a severe curse of dimensionality, which makes it impossible to apply standard approaches in the estimation of games literature. We instead reformulate the sellers’ problem as an aggregative game ([Nocke and Schutz \(2018\)](#)), whereby a seller’s profits only depend on rivals’ prices through an aggregate statistic. With this approach, we significantly reduce the dimensionality, and can thus proceed to recover the distribution of sellers’ marginal costs. We use this primitive to investigate a range of market-design counterfactuals.

Related Literature. Our work is related to a large and growing literature at the intersection of economics and computer science studying regret learning and regret minimization. The regret learning framework was originally developed for single-agent decision problems. However, a series of papers have pointed out significant implications of regret-based learning procedures in strategic settings. Seminal game theory papers in which players try to minimize regret include [Hannan \(1957\)](#), [Foster and Vohra \(1997\)](#),

²For 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.”

1998, 1999), and [Hart and Mas-Colell \(2000\)](#). In particular, our work is related to [Hart and Mas-Colell \(2000, 2013\)](#) and [Hart \(2005\)](#), whose convergence results we generalize to incomplete information environments.

Recent contributions in computer science ([Hartline, Syrgkanis and Tardos \(2015\)](#) and [Caragiannis, Kaklamanis, Kanellopoulos, Kyropoulou, Lucier, Paes Leme and Tardos \(2015\)](#)) offer related theoretical results on the convergence of no-regret learning dynamics in environments with incomplete information. In contrast to these authors, who focus on price of anarchy and efficiency results, the emphasis of our work is on connecting learning dynamics to the inference problem of an external observer. Moreover, because of motivation, modeling, and technical differences, our setting is different from that studied by the computer science literature. Therefore, our analyses of regret-based learning procedures in settings with incomplete information are complementary, but not equivalent. We refer to Section 3 for further discussion.

Another strand of the computer science literature pursues the connection to empirical work ([Nekipelov, Syrgkanis and Tardos \(2015\)](#) and [Nisan and Noti \(2017a,b\)](#)). [Nekipelov et al. \(2015\)](#) are the first to suggest the regret-based approach to econometrics. In an online auctions environment, they characterize (and perform inference on) the set of valuations consistent with a given level of regret, but do not rely on an equilibrium concept. [Nisan and Noti \(2017a,b\)](#) evaluate a similar approach in experimental data and propose adjustments to the no regret estimation procedure. We also perform inference on the distribution of payoff types when the path of play has an asymptotic no regret property. In contrast to these papers, however, we do so by leveraging convergence results to interpret the data through the lens of the static equilibrium notion of Bayes correlated equilibrium or its coarse analog.

More broadly, regret learning and regret minimization is now the leading approach in online learning (see [Shalev-Shwartz \(2012\)](#)) and in multi-armed bandit problems (see [Bubeck and Cesa-Bianchi \(2012\)](#), [Lattimore and Szepesvári \(2019\)](#), and [Slivkins \(2019\)](#)). These concepts are also a cornerstone in the algorithmic game theory literature (see [Nisan, Roughgarden, Tardos and Vazirani \(2007\)](#), [Shoam and Leyton-Brown \(2008\)](#), and [Roughgarden \(2016\)](#)). An excellent discussion of regret minimization from the online machine learning perspective can be found in [Cesa-Bianchi and Lugosi \(2006\)](#). Regret minimization has also been used in designing treatment rules (e.g., [Manski \(2004\)](#) and [Stoye \(2009\)](#)) and in forecast aggregation (e.g., [Arieli, Babichenko and Smorodinsky \(2018\)](#) and [Babichenko and Garber \(2018\)](#)).

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.

We also develop a novel approach to deal with the curse of dimensionality inherent in the estimation of large games. In our application, we reformulate the game in an aggregative form as in [Nocke and Schutz \(2018\)](#). Hence, players’ payoffs only depend on their competitors’ actions via an aggregate statistic. We thus achieve a dramatic reduction in the dimension of the (Bayes correlated) equilibrium distribution, without redefining the equilibrium concept as in [Weintraub, Benkard and Van Roy \(2008\)](#).

[Magnolfi and Roncoroni \(2019\)](#), [Syrngkanis, Tamer and Ziani \(2018\)](#), and [Gualdani and Sinha \(2020\)](#) also consider estimation under the assumption of Bayes correlated equilibrium behavior for models of discrete games, auctions, and single-agent decision problems respectively. Although this paper proposes the use of a similar estimation technique, the motivation is very different. In fact, [Magnolfi and Roncoroni \(2019\)](#), [Syrngkanis et al. \(2018\)](#) and [Gualdani and Sinha \(2020\)](#) exploit the link between equilibrium behavior and information to establish that Bayes correlated equilibrium allows to estimate games under weak assumptions on information. In this paper, instead, we motivate estimation under Bayes correlated equilibrium when the data are generated by repeated interaction with weak assumptions on behavior.

Road Map. In Section 2, we present the theoretical framework and formalize the notions of regret and asymptotic ε -regret for the repeated play of a one-shot incomplete information game. In Section 3, we study the convergence properties of ε -regret dynamics to the set of Bayes (coarse) correlated ε -equilibria of the one-shot game. In Section 4, we specialize the theoretical model to study what features of the underlying economic environment we can empirically recover, and how, under weak assumptions on behavior when the one-shot game is played repeatedly over time. In Section 5, we present our empirical application and outline our aggregative approach to sidestep the curse of dimensionality at the estimation stage. In Section 6, we further discuss our main theoretical and econometric results and present extensions and robustness checks. In Section 7, we conclude. In Appendix A, we discuss regret-minimizing algorithms for the repeated play of the one-shot incomplete information game. Omitted proofs are in Appendix B.

References

- Akerberg, 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.
- Arieli, Itai, Yakov Babichenko, and Rann Smorodinsky (2018), “Robust Forecast Aggregation.” *Proceedings of the National Academy of Sciences*, 115, 67–96.
- Babichenko, Yakov and Dan Garber (2018), “Learning of Optimal Forecast Aggregation in Partial Evidence Environments.” *Working Paper*.
- Bergemann, Dirk and Stephen Morris (2016), “Bayes Correlated Equilibrium and the Comparison of Information Structures in Games.” *Theoretical Economics*, 11, 487–522.

- Bubeck, Sebastien and Nicolò Cesa-Bianchi (2012), “Regret Analysis of Stochastic and Nonstochastic Multi-Armed Bandit Problems.” *Foundations and Trends in Machine Learning*, 5, 1–122.
- Caragiannis, Ioannis, Christos Kaklamanis, Panagiotis Kanellopoulos, Maria Kyropoulou, Brendan Lucier, Renato Paes Leme, and Éva Tardos (2015), “Bounding the Inefficiency of Outcomes in Generalized Second Price Auctions.” *Journal of Economic Theory*, 157, 343—388.
- Cesa-Bianchi, Nicolò and Gábor Lugosi (2006), “Prediction, Learning, and Games.” *Cambridge University Press*.
- 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.
- Foster, Dean P. and Rakesh V. Vohra (1997), “Calibrated Learning and Correlated Equilibrium.” *Games and Economic Behavior*, 21, 40–55.
- (1998), “Asymptotic Calibration.” *Biometrika*, 85, 379–390.
- (1999), “Regret in the On-Line Decision Problem.” *Games and Economic Behavior*, 29, 7–35.
- Gualdani, Cristina and Shruti Sinha (2020), “Identification and Inference in Discrete Choice Models with Imperfect Information.” *Working Paper*.
- Hannan, James (1957), “Approximation to Bayes Risk in Repeated Play.” M. Dresher, A.W. Tucker, and P. Wolfe eds., *Contributions to the Theory of Games, Vol. 3*. Princeton University Press.
- Hart, Sergiu (2005), “Adaptive Heuristics.” *Econometrica*, 73, 1401–1430.
- 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.” Eric Maskin ed. *World Scientific Series in Economic Theory*. World Scientific.
- Hartline, Jason, Vasilis Syrgkanis, and Eva Tardos (2015), “No-Regret Learning in Bayesian Games.” *NIPS 2015: Advances in Neural Information Processing Systems*, 3061–3069.
- Lattimore, Tor and Csaba Szepesvári (2019), “Bandit Algorithms.” *Preprint available at <https://tor-lattimore.com/downloads/book/book.pdf>*.
- 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 (2019), “Estimation of Discrete Games with Weak Assumptions on Information.” *Working Paper*.
- Manski, Charles F. (2004), “Statistical Treatment Rules for Heterogeneous Populations.” *Econometrica*, 72, 1221–1246.

- Nekipelov, Denis, Vasilis Syrgkanis, and Eva Tardos (2015), “Econometrics for Learning Agents.” *EC '15: Proceedings of the 2015 ACM Conference on Economics and Computation*, 1–18.
- Nisan, Noam, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani (2007), “Algorithmic Game Theory.” *Cambridge University Press*.
- Nisan, Noam and Gali Noti (2017a), “An Experimental Evaluation of Regret-Based Econometrics.” *WWW '17 Companion: Proceedings of the 26th International Conference on World Wide Web*, 73–81, International World Wide Web Conferences Steering Committee.
- (2017b), “A Quantal Regret Method for Structural Econometrics in Repeated Games.” *EC '17: Proceedings of the 2017 ACM Conference on Economics and Computation*, 123–123, ACM.
- Nocke, Volker and Nicolas Schutz (2018), “Multiproduct-Firm Oligopoly: An Aggregative Games Approach.” *Econometrica*, 86, 523–557.
- 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.
- Roughgarden, Tim (2016), “Twenty Lectures on Algorithmic Game Theory.” *Cambridge University Press*.
- Shalev-Shwartz, Shai (2012), “Online Learning and Online Convex Optimization.” *Foundations and Trends in Machine Learning*, 4, 107–194.
- Shoam, Yoav and Kevin Leyton-Brown (2008), “Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations.” *Cambridge University Press*.
- Slivkins, Aleksandrs (2019), “Introduction to Multi-Armed Bandits.” *Foundations and Trends in Machine Learning*, 1, 1–286.
- Stoye, Jörg (2009), “Minimax Regret Treatment Choice with Finite Samples.” *Journal of Econometrics*, 151, 70–81.
- Syrgkanis, Vasilis, Elie Tamer, and Juba Ziani (2018), “Inference on Auctions with Weak Assumptions on Information.” *Working Paper*.
- Tamer, Elie (2003), “Incomplete Simultaneous Discrete Response Model with Multiple Equilibria.” *The Review of Economic Studies*, 70, 147–165.
- Weintraub, Gabriel Y, C Lanier Benkard, and Benjamin Van Roy (2008), “Markov Perfect Industry Dynamics with Many Firms.” *Econometrica*, 76, 1375–1411.