Games and Economic Behavior **29,** 1–6 (1999)
Article ID game.1999.0733, available online at http://www.idealibrary.com on IDEAL®

# Introduction to the Special Issue

Rakesh Vohra

*MEDS Department, Kellogg Graduate School of Management, Northwestern University, 2001 Sheridan Road, Evanston, Illinois 60208-2009*

David K. Levine

*Department of Economics, University of California, Los Angeles, Los Angeles, California 90095*

and

Dean Foster

*Department of Statistics, The Wharton School, University of Pennsylvania, Philadelphia, Pennsylvania 19104*

Received April 19, 1999

David Blackwell was born 80 years ago in Centralia, Illinois. He graduated from high school at the age of 16 and enrolled at the University of Illinois with the intention of becoming an elementary school teacher. This as we now know was not to be. The schoolboy's loss is the scholars gain. He went on to obtain a Ph.D., a fellowship at the fabled Institute for Advanced Studies, and finally retiring, but only in the legal sense, as Professor of Statistics at the University of California.

In spite of his many mathematical accomplishments, Blackwell confesses to have no interest in research. To quote him directly:

> "Basically, I'm not interested in doing research and I never have been, I'm interested in *understanding*, which is quite a different thing."

From the acorn of such modest intentions have sprung mighty oaks. This special issue of *Games and Economic Behavior* pays tribute to them by gathering together papers that are branches of two of Blackwell's oaks: approachability and merging. Blackwell's approachability theorem can be viewed as a generalization of the minimax theorem of zero-sum games to matrix games with vector payoffs. While interesting in its own right, what is remarkable are its applications to forecasting and learning in games. One set of contributions to this special issue advances the modern theory of forecasting and learning in games, and explores the connection of these



0899-8256/99 $30.00

results to Blackwell's approachability theorem. The second of Blackwell's oaks is the merging of opinions. Merging is a precise formulation of the following basic question: when do different individuals with different priors but sharing the same data agree about their posteriors? The second set of contributions to this special issue explores the merging of opinions between different players in a game. Remarkably, it turns out also that the merging of opinions has a strong connection to calibrated learning, which in turn rests on Blackwell's first oak. So it seems an entirely fitting tribute to David Blackwell to dedicate this special issue of *Games and Economic Behavior* devoted to forecasting, learning, and merging of opinions in games to him.

To illustrate the application of the approachability theorem to forecasting and learning, consider the problem of repeatedly predicting the outcome of a coin toss. One's performance will be measured by the fraction of correct predictions. Sans knowledge of the process that governs the outcome on each round, how well can one do? An obvious solution is to predict HEADS in each round with probability a half. In the limit, at least 50% of ones predictions will be correct. Can one do better? Surprisingly, yes. The approachability theorem yields a randomized method that produces a fraction of correct predictions, $f(n)$ with the following property,

$$f(n) \geq \max\{h(n), 1 - h(n)\} - e(n),$$

where $e(n)$ goes to zero almost surely as $n$ goes to infinity. Here $h(n)$ is the fraction of HEADS realized in the first $n$ rounds.

The coin problem considered above is a special case of what computer scientists would call the *on-line decision problem*. In each period an action is chosen, a signal observed, and a utility received as a function of the action and signal. The problem is to try to forecast the signal, so as to choose the optimal action. Historical values of the actions and signals may be used to "learn" what the forthcoming signal will be.

The weakest useful criterion for successful long-run learning is that of universal consistency in learning. This considers simple empirical frequencies. That is, a learning rule is universally consistent if it is approximately optimal in the long run against the true empirical frequencies. The existence of such rules was first established in the mid 1950s by James Hannan (1957). Subsequently Blackwell (1956a) derived it as an application of his approachability theorem Blackwell (1956b). This result of Hannan's was lost and rediscovered several times, as Foster and Vohra explain in "Regret in the On-Line Decision Problem." The Hannan result admits considerable further generalization. Rustichini, in "Minimizing Regret: the General Case" uses approachability, to extend the result to the case when players receive an indirect signal about their payoffs at each round. In ‘‘Adaptive Game Playing using Multiplicative Weights’’ by Freund and Schapire the Hannan theorem is revisited using an algorithm much studied in the computer science literature. They provide a simple learning algorithm with convergence bounds that *hold for any finite number of rounds*. Further, they establish in a certain sense that no adaptive procedure for playing a game can have a better rate of convergence.

A criterion stronger than universal consistency is that of calibration with respect to own forecasts. In the coin flipping example, suppose in each round one predicts the *probability* of observing heads. Calibration means that on those occasions when a probability $p$ of HEADS was announced, HEADS should have been realized in a proportion $p$ of those periods. Rules satisfying this stronger property were first described by Foster and Vohra (1998). Moreover, Foster and Vohra (1997) showed that when all players follow universally calibrated rules, play converges in the long-run time average sense to the set of correlated equilibria. Because of the importance of this result, several variations are considered in this volume. Two simple universally calibrated algorithms together with simple proofs, are in Foster ‘‘A Proof of Calibration via Blackwell’s Approachability Theorem,’’ and in Fudenberg and Levine ‘‘An Easier Way to Calibrate.’’ The former is particularly interesting, because it shows that Blackwell’s approachability method can be used to establish calibration results as well as universal consistency results. Fudenberg and Levine, by contrast, give a direct algebraic proof. A broad overview of the literature on calibration and the related literature on the on-line decision problem in computer science is in Foster and Vohra’s ‘‘Regret and the Online Decision Problem.’’ Finally, Fudenberg and Levine ‘‘Conditional Universal Consistency’’ consider a relatively simple variation on fictitious play, and show how it may be universally calibrated with respect to a much broader class of ‘‘conditional’’ checking rules. Basically, these checking rules require calibration conditional on prespecified events, such as the previous period realization of the random variable. These conditional checking rules also play an important role in connecting calibrated learning with merging as Kalai *et al*. show in ‘‘Calibrated Forecasting and Merging.’’

In addition to the papers appearing in this volume, there is other work on calibration. Hart and Mas-Colell (1997) use Blackwell’s approachability criterion to show that a simple regret based procedure is universally consistent. Fudenberg and Levine (1998) give a general method of converting any universally consistent procedure to a universally calibrated procedure. Hart and Mas-Colell add two innovations to this: first, they observe that the eigenvector problem Fudenberg and Levine propose to solve at every step can be replaced with a simple multiplication of the matrix times the previously chosen probabilities of actions. Second, they simplify this to playing actions in proportion to past regret. This loses the property of universal calibration, but at least if everyone plays this way, all are calibrated, and play in the long-run time average sense converges to the set of correlated equilibria.

A second paper on calibration is Lehrer's (1997) ''Any Inspection is Manipulable.'' This paper makes two contributions. The first is a generalization of the approachability theorem to infinite-dimensional spaces. The second is to apply the theorem to derive forecasting procedures that pass simultaneously any countable collection of history dependent checking rules. History based checking rules are restricted to using the previous sequence of outcomes to specify what future events will be checked for calibration. This is a generalization of the Fudenberg and Levine result that considers only particular countable classes of history-dependent checking rules. A third paper, by Sandroni *et al*. (1999), ''Calibration with Many Checking Rules,'' extends this result to include forecast based checking rules. Forecast based checking rules use the previous history of forecasts to specify what future events will be checked.

The second of Blackwell's oaks is merging. As we indicated above, merging is a precise formulation of the following basic question: when do different individuals with different priors but sharing the same data agree about their posteriors? Blackwell and Dubins (1962) showed that merging occurs when the priors are absolutely continuous with respect to each other; roughly speaking, they assume that the support of the individuals priors must be the same. In ''Calibrated Forecasting and Merging'' by Kalai *et al*. this criterion for merging is shown to be equivalent to the notion of calibration with respect to checking rules. The idea of calibration is to compare empirical frequencies with theoretical probabilities. A checking rule specifies the events for which the frequencies should be checked. This paper closes the circle between the two of Blackwell's oaks.

The fact that in an infinite parameter space there is no single measure that all other measures are absolutely continuous with respect to has important consequences for learning. These are made precise in Miller and Sanchirico in ''The Role of Absolute Continuity in Merging of Opinions.''

In his paper ''Merging, Reputation, and Repeated Games with Incomplete Information,'' Sylvain Sorin, uses merging to pour new wine into old bottles. He uses the result of Blackwell and Dubins and their refinements due to Kalai and Lehrer to characterize the equilibria of two-person repeated games with incomplete information and known own payoffs. This paper shows how the modern literature on learning is closely connected to the earlier literature on reputation, in which an impatient player is manipulated by a more patient player who exploits the equilibrium learning procedure of the less patient player.

Many of the procedures inspired by considerations of calibration and merging can be motivated by appeals to stimulus response models of learning and behavior. In Rustichini ‘‘Optimal Properties of Stimulus Response Learning Models,’’ two very broad classes of stimulus response type procedures are considered. The paper gives a precise characterization of when procedures of each type converge to optimality in terms of the information available to the players.

The calibration and learning papers study the theoretical properties of calibration and learning procedures. Of equal importance are how the procedures behave dynamically. One set of results we have already mentioned is the convergence of calibrated procedures to correlated equilibrium. The remaining contributions in this issue address other dynamic issues.

The paper by Michel Benaim and Morris Hirsch entitled ‘‘Mixed Equilibria and Dynamical Systems Arising from Fictitious Play in Perturbed Games,’’ studies the behavior of fictitious play for infinitely repeated games of incomplete information with randomly perturbed payoffs. The study of such games was motivated by a desire to provide a justification for mixed strategy equilibria. This paper uses developments in dynamical systems and stochastic approximation theory to study the limit set of the dynamic process induced by fictitious play in several classes of such games.

One of the first and most well-known rules for repeated game playing is best response. In ‘‘A Note on Best Response Dynamics,’’ Edward Hopkins studies the relationship between the continuous time best response dynamic, its perturbed version, and evolutionary dynamics in relation to mixed strategy equilibria. He shows that a variety of learning dynamics when aggregated over a large population of players, can be approximated in continuous time by adjustments in the direction of a positive definite matrix times the payoff vector. He calls these dynamics ‘‘positive definite adjustment’’ or PDA dynamics. Under these dynamics he shows that if the payoff matrix is negative definite a fully mixed equilibrium is stable; if it is positive definite, it is unstable.

The paper by Abraham Neyman and Daijiro Okada entitled ‘‘Strategic Entropy and Complexity in Repeated Games’’ brings to the study of bounded rationality ideas from yet another subject that has been close to David Blackwell’s heart: Information Theory. Different communities of scholars have proposed different restrictions on rationality as models of bounded rationality. Computer scientists use the language of finite state automata and complexity theory. Statisticians limit the class of probabilistic models the players might believe. Experimental game theorist’s appeal to findings on the cognitive missteps made by humans. Alas, these restrictions live in different types of spaces. How can we compare them, or work with more than one restriction at a time? Okada and Neymann suggest that one consider the entropy of the strategies implied by the restrictions. They provide theorems that use this entropy of the strategies to predict the outcome in terms of payoff of the repeated game.

## REFERENCES


Blackwell, D. (1956a). ''Controlled Random Walks,'' in *Proceedings International Congress of Mathematicians*, Vol. 3, pp. 336–338, Amsterdam: North Holland, 1954.

Blackwell, D. (1956b). ''An Analog of the Minmax Theorem for Vector Payoffs,'' *Pacific J. Math* **6**, 1–8.

Blackwell, D., and Dubins, L. (1962). ''Merging of Opinions with Increasing Information,'' *Ann. Math. Statist.* **38**, 882–886.

Foster, D., and Vohra, R. (1997). ''Calibrated Learning and Correlated Equilibrium,'' *Games Econ. Behavior* **21**, 40–55.

Foster, D., and Vohra, R. (1998). ''Asymptotic Calibration,'' *Biometrika* **85**, No. 2, 379–390.

Fudenberg, D., and Levine, D. K. (1998). *The Theory of Learning in Games*. Cambridge, MA: MIT Press.

Hannan, J. (1957). ''Approximation to Bayes Risk in Repeated Plays,'' in *Contributions to the Theory of Games* (M. Dresher, A. W. Tucker, and P. Wolfe, Eds.), Vol. 3, pp. 97–139. Princeton: Princeton Univ. Press.

Hart, S., and Mas-Colell, A. (1997). ''A Simple Adaptive Procedure leading to Correlated Equilibrium,'' manuscript.

Lehrer, E. (1997): ''Any Inspection is Manipulable,'' manuscript.

Sandroni, A., Smorodinsky, R., and Vohra, R. V. (1999). ''Calibration with Many Checking Rules,'' manuscript.