

The Seventh Stony Brook Summer Festival on
Game Theory
Social Learning Workshop
July 21 to July 23, 1996

Book of Abstracts

Adaptive Play in Multiplayer Bargaining Situations

Murali Agastya

University College London

Abstract

This paper studies adaptive learning in multiplayer bargaining situations. It is assumed that a multiplayer bargaining situation can be summarized by a characteristic function. Paying close attention to the delicate questions of allocation and coalition formation that arise naturally in such situations, conditions are provided under which players learn to adopt core allocations. If one now allows for small deviations from rationality, we find that only a strict subset of these allocations are (stochastically) stable. We characterize these allocations and find them to be closely related to the allocation that maximizes the Nash product over core allocations. Uniqueness is not assured if there are more than three players.

Growing Artificial Societies

Robert Axtell

Brookings Institution

Abstract

A model will be described in which a heterogeneous population of autonomous agents inhabit an artificial landscape of renewable resources and engage in production, consumption, economic exchange, sexual reproduction and cultural transmission. All agent interactions are bilateral and spatially local. For example, trade occurs between neighbors via Edgeworth barter with prices determined locally. Under a wide range of conditions the economy that emerges from such distributed interactions is essentially non-equilibrium in character. The extent to which this highly decentralized economy can approach socially-optimal outcomes is investigated. Results will be displayed as animations using software for “growing” artificial societies.

Market Experimentation and Pricing

Dirk Bergemann

Yale University and Boston University

Abstract

We develop a continuous-time model of dynamic Bayesian learning in a duopoly. The value of one product is initially unknown to the market. The market participants learn more about the true value of the product as they experiment over time. Firms are setting prices to induce experimentation with their product and the aggregate outcomes are public information.

As agents learn from the experiments of others, informational externalities arise. Surprisingly, the informational externality leads to too much learning. The firms free ride on the market as the social costs of experimentation are not appropriately reflected in the equilibrium prices. The value function of each seller displays preference for information in contrast to the buyers who are information averse.

The structure of the inefficiency of the unique Markov Perfect Equilibrium in this two-sided learning model is analyzed. The price and allocation structure of the equilibrium is obtained explicitly for a finite number of buyers as well as for a continuum of buyers. The model is extended to a locational model of product differentiation and the dynamic evolution of the market shares is analyzed.

Evolution of the Meaning of Messages in Sender-Receiver Games: An Experiment

Andreas Blume

University of Iowa

Abstract

This paper experimentally investigates the evolution of the meaning of messages in sender-receiver games of common and divergent interests. Meaning is viewed as an equilibrium phenomenon and an equilibrium as the result of a dynamic process. Hypotheses are investigated using evolutionary game theory. There is considerable support for the theory in common interest games. Results indicate that efficient communication emerges endogenously; providing population history information accelerates this process. For games with divergent interests, there is less support for the theory. Theory is not unanimous in its predictions and fails to capture certain salient features of the data. Consistent with theory, the size of the message space influences the equilibrium selected. However, the degree of divergence of interest between senders and receivers and receivers' initial posterior beliefs about sender types also influence the equilibrium selected. Finally, initial conditions play an important role by establishing coordination precedents that are decisive in both game types.

Naive Reinforcement Learning with Endogenous Aspirations

Tilman Borgers

University College London

Abstract

This paper considers a simple learning process for decision problems under risk. All behavior change derives from the reinforcing or deterring effect of instantaneous payoff experiences. Payoff experiences are reinforcing or deterring depending on whether the payoff exceeds an aspiration level or falls short of it. The aspiration level is endogenous. Over time it is adjusted into the direction of the actually experienced payoff. This paper shows that realistic aspiration level adjustments may improve the decision maker's long run performance, because they may prevent him from feeling dissatisfied with even the best of the available strategies. On the other hand, the paper also shows that in a large class of decision problems endogenous aspiration levels lead to persistent deviations from expected payoff maximization because they create "probability matching" effects.

Social Learning and Costly Information Acquisition

Roberto Burguet

Instituto de Análisis Económico (CSIC), Barcelona, Spain

Abstract

Short-lived agents want to predict a random variable θ and have to decide how much effort to devote to collect private information and consequently how much to rely on public information. The latter is just a noisy average of past predictions. It is shown that costly information acquisition prevents an unbounded accumulation of public information if (and only if) the marginal cost to acquire information is positive at zero ($C'(0) > 0$). When $C'(0) = 0$ public precision at period n , τ_n , tends to infinity with n but the rate of convergence of public information to θ is slowed down with respect to the exogenous information case. At the market outcome agents acquire too little private information and there is underinvestment in public information. This happens either with respect to a first best benchmark or, for n large, with respect to a second best benchmark. For high discount factors the limit point of market public precision always falls short of the welfare benchmarks whenever $C'(0) > 0$. Otherwise, if $C'(0) = 0$ public precision accumulates in an unbounded way both at the first and second best solutions. More public information may hurt at either the market or second best solutions.

**Structure and Strategy in Collective Action:
Communication and Coordination in Social Networks**

Michael Suk-Young Chwe

University of Chicago

Abstract

Which social structures are conducive to collective action? In this paper, each individual in a group wants to participate only if joined by others; exactly how many compatriots are necessary is given by her “threshold.” Social structure here is thought of as a communication network or graph. As time progresses, people explicitly communicate their thresholds to each other using the network. At a given time, each person, based on his knowledge of other people’s thresholds, and his knowledge of what other people know, decides whether or not to participate. Knowledge of others is modelled as a partition over a state space, and hence the approach is similar to correlated equilibrium, as opposed to adaptation or evolution. Some basic results are: participation is “clumpy” and typically does not grow smoothly; common knowledge is equivalent to a complete subgraph; organizations are good at forming common knowledge; “strong” links are better when thresholds are low and “weak” links are better when thresholds are high; and wide dispersion of insurgents is good for collective action, but too much dispersion can be bad.

This paper is available at <http://www.spc.uchicago.edu/wwwchwe/>.

Reinforcement-Based Learning vs. Bayesian Learning: A Comparison

David J. Cooper

University of Pittsburgh

Abstract

This paper compares a reinforcement-based learning model (Roth and Erev, 1995) with a Bayesian learning model (Cooper, Garvin, and Kagel 1995a) for their abilities to characterize data from limit pricing experiments. We find that the Bayesian model outperforms the reinforcement-based model. This is due to the Bayesian model's ability to incorporate greater cognitive sophistication. We discuss the process of selecting learning models to describe experimental data.

Learning with Hazy Beliefs

Dean Foster* (joint w/P. Young[†])

Abstract

Players are rational if they always choose best replies given their beliefs. They are good predictors if the difference between their beliefs and the distribution of the others' actual strategies goes to zero over time. Learning is deterministic if beliefs are fully determined by the initial conditions and the observed data. (Bayesian updating is a particular example.) If players are rational, good predictors, and learn deterministically, there are many games for which neither beliefs nor actions converge to a Nash equilibrium. In fact, with high probability they are not even close to a Nash equilibrium. We introduce an alternative approach to learning called prospecting in which players are rational and good predictors, but beliefs have a small random component. In any finite game, and from any initial conditions, prospecting players learn to play arbitrarily close to Nash equilibrium with probability one.

*University of Pennsylvania

[†]The Johns Hopkins University

Word-of-Mouth Learning

Drew Fudenberg

Harvard University

Abstract

This paper analyzes a simple model of rational word-of-mouth learning, in which successive generations of agents make once-and-for-all choices between two alternatives. Before making a decision, each new agent samples N old ones and asks them which choice they used and how satisfied they were with it. If the sampling is “unbiased” or “proportional” in the sense that each person in the population is equally likely to be sampled, each player samples two or more others, and there is some information (however little) in payoff observations, the long run outcome is one in which everyone is doing the same thing. Moreover, if the payoff observation is sufficiently informative, the long run outcome is efficient, where this “sufficient informativeness” condition is quite weak if the only source of aggregate uncertainty relates to which choice is better. We also investigate a range of other sampling rules, such as those that over-represent popular or successful choices, and determine which ones favor global convergence towards efficiency.

Social Learning and Rigid Behavior

Joseph E. Harrington, Jr.

The Johns Hopkins University

Abstract

A source of heterogeneity among people is their adaptiveness to changes in the environment. This paper explores under what circumstances we might expect social learning to result in agents engaging in relatively rigid practices.

The population is comprised of two generic behavioral types: rigid agents always choose the same action while adaptive agents always choose the action most appropriate for the current environment. The advantage to being adaptive is clear while the advantage to being rigid is that we assume “practice makes perfect” in that more frequent usage of an action improves an agent’s performance with that action. Agents compete for advancement in a hierarchical society. Those with high social status comprise the set of role models for the next generation of agents. Each new agent adopts a behavioral rule based upon a confluence of two forces: an innate bias and a desire to imitate the behavior of agents with high social status. However, incoming agents do not get to observe successful agents’ behavioral rules but only their behavior and the context within which they acted. This can create systematic biases in the imitation dynamic. The population to which this process converges is characterized.

Depending on the parameter values, we find that society can be either dominated by rigid behavior or adaptive behavior or there can be significant coexistence of both behavioral patterns. However, there is a clear bias in favor of rigid behavior. Rigid behavior is more likely to dominate, the more hierarchical is society. While it is often the case that a more volatile environment makes adaptive behavior more widespread, there is a wide set of circumstances under which an increase in the volatility of the environment results in society moving from a predominance of adaptive behavior to a predominance of rigid behavior.

Correlated Learning

Sergiu Hart* (joint w/A. Mas-Colell†)

Abstract

We propose simple adaptive procedures that, in any finite n -person game, converge to the set of correlated equilibria of the game. Each player needs to know only his own payoff function and the history of past actual moves of all players.

*The Hebrew University, Israel

†Universitat Pompeu Fabra, Spain

Money as a Social Convention: A Dynamic Bilateral Trade Game with a Single Intermediating Commodity

Phillip M. Johnson

University of California at Los Angeles

Abstract

A long-standing question in monetary theory asks whether a medium of exchange could possibly or certainly emerge spontaneously from the behavior of economic agents without the compulsion of a monetary authority. Previous research on money as a medium of exchange (eg., Oh (1989) and Kiyotaki and Wright (1989)) has depended on either initial conditions or steady state conditions to achieve results. These approaches are not satisfactory for addressing the question of emergence, in the former case because the determining element is exogenous, the latter case there are multiple equilibria without criteria for selecting between them. We construct an N -person non-cooperative anonymous game from an economy with many commodities and bilateral exchange. The environment is restricted so that in addition to direct barter opportunities there is a commodity - a token commodity with no value in consumption - which can be used for intermediate exchanges. We make use of the recent advances in evolutionary learning in games pioneered in Kandori, Mailath, and Rob (1993) to characterize conditions under which monetary trade is typically observed in the long-run despite any initial conditions and how informational assumptions affect the transition time.

Endogenous Relationships

George Mailath

University of Pennsylvania

Abstract

Models of evolutionary dynamics and social learning have become important in recent years. The majority of these models assume a “uniform” environment in which all players interact symmetrically. Some more recent work has focused on local interactions. In this setting, players only observe and interact with a (potentially) small set of neighboring players (for example, a firm only observes its own clientele). Models that place players on a lattice treat the neighbors as exogenous - they are best thought of as models of exogenous relationships. On the other hand, most relationships in economics are endogenous. Taking explicit account of the endogeneity of the relationships can lead to new insights into the operation of social learning. In particular, there may be feedback effects between the actions an agent takes to influence its relationships with other agents (such as, their identity) and other aspects of that agent’s behavior.

Strategic Behavior with General Local Interaction

Stephen Morris

University of Pennsylvania

Abstract

This paper introduces techniques for analyzing strategic behavior in general local interaction structures without fixed geometric form. The analysis exploits an equivalence between incomplete information games and local interaction games.

Adaptive Behavior, Idiosyncratic Risk, and Aggregate Uncertainty

Arthur J. Robson

University of Western Ontario, Canada

Abstract

This paper considers a class of adaptive stochastic learning rules which have been applied in psychology. They generalize biological selection and have become relevant to economic theory as a consequence of their use in evolutionary game models. The environment considered here is non-strategic but includes both idiosyncratic risk and aggregate uncertainty. In the long run, individuals turn out to be more averse to aggregate uncertainty than to precisely comparable idiosyncratic risk and they may violate first-order stochastic dominance. The implications of this are discussed.

Learning to do Backward and Forward Induction

Larry Samuelson

University of Wisconsin

Abstract

This paper examines processes by which players adapt their behavior in, or learn how to play, extensive form games. Some modeling issues concerning the effect of the extensive form on the learning process are first considered. Do players observe strategies, or only outcomes? What do players have to learn? What do players know about behavior at information sets they cannot observe, and how do they learn about such behavior? Attention then turns to the implications of such learning processes for equilibrium refinement and selection issues. Does the evolutionary process yield outcomes that satisfy backward or forward induction criteria? Results are mixed in both cases, but suggest that backward induction plays a less important role than is currently common, and forward induction a somewhat more important role, in evaluating equilibria.

On the Evolution of Imitative Behavior

Karl H. Schlag

University of Bonn

Abstract

We analyze the evolution of behavior for choosing among actions with uncertain payoffs where rare mutations occur. Players belonging to a large population use in each of a sequence of rounds an action and then receive a payoff. Occasionally a player has the opportunity to observe the action of another (randomly selected) player and the payoff he received. Each player has a revision behavior, that determines the future action of the player based on the observation. Evolution is modelled using the replicator dynamic with the revision behaviors as replicators. Analysis is restricted to the choice between two actions and to populations with at most two different behavioral rules present.

Non-stationary improving rules characterized by Schlag [12] can prevent mutations from spreading. These are the only such rules that have this property for any distribution over payoffs. Mutations that survive against these rules follow an imitative behavior.

Evolving Aspirations and Cooperation

Fernando Vega-Redondo

**Universidad de Alicante and
Instituto Valenciano de Investigaciones Económicas**

Abstract

A model of “satisficing” behavior in the repeated Prisoners Dilemma is studied. Each player has an *aspiration* at each date, and takes an action. [S]he switches from the action played in the previous period only if the achieved payoff fell below the aspiration level (with a probability that depends on the shortfall). Aspirations are updated in each period, according to payoff experience in the previous period. In addition, aspirations are subjected to random perturbations around the going level, with a small “tremble” probability. For sufficiently slow updating of aspirations, and small tremble probability, it is shown that in the long run both players cooperate most of the time.

SOCIAL LEARNING WORKSHOP

Organized by Peyton Young

July 21 to July 23, 1996

Table of Contents

Murali Agastya - Monday 4:00 PM

Adaptive Play in Multiplayer Bargaining Situations

Robert Axtell - Tuesday 1:00 PM

Growing Artificial Societies

Dirk Bergemann - Monday 11:15 AM

Market Experimentation and Pricing

Andreas Blume - Monday 1:40 PM

Evolution of the Meaning of Messages in Sender-Receiver Games: An Experiment

Larry Blume - Tuesday 8:30 AM

Local Interactions

Tilman Borgers - Sunday 1:40 PM

Naive Reinforcement Learning with Endogenous Aspirations

Roberto Burguet - Monday 11:55 AM

Social Learning and Costly Information Acquisition

Michael Chwe - Tuesday 1:40 PM

Structure and Strategy in Collective Action: Communication and Coordination in Social Networks

David J. Cooper - Monday 2:20 PM

Reinforcement-Based Learning vs. Bayesian Learning: A Comparison

- Dean Foster (P. Young) - Sunday 4:40 PM**
Learning with Hazy Beliefs
- Drew Fudenberg - Monday 10:00 AM**
Word-of-Mouth Learning
- Edward Glaeser - Tuesday 11:15 AM**
Crime and Social Interaction
- Joseph E. Harrington, Jr. - Monday 4:40 PM**
Social Learning and Rigid Behavior
- Sergiu Hart (A. Mas-Colell) - Sunday 4:00 PM**
Correlated Learning
- Phillip M. Johnson - Monday 3:20 PM**
Money as a Social Convention: A Dynamic Bilateral Trade Game with a Single Intermediating Commodity
- Ehud Kalai - Monday 8:30 AM**
Rational Interactive Learning
- George Mailath - Tuesday 10:00 AM**
Endogenous Relationships
- Stephen Morris - Tuesday 2:20 PM**
Strategic Behavior with General Local Interaction
- Arthur J. Robson - Sunday 3:20 PM**
Adaptive Behavior, Idiosyncratic Risk, and Aggregate Uncertainty
- Larry Samuelson - Sunday 10:30 AM**
Learning to do Backward and Forward Induction
- Karl H. Schlag - Sunday 2:20 PM**
On the Evolution of Imitative Behavior
- Fernando Vega-Redondo - Sunday 1:00 PM**
Evolving Aspirations and Cooperation
- Peyton Young - Sunday 9:00 AM**
Learning, Rationality, and Game Dynamics