

Symposium

Bounded rationality and learning *

Introduction

Yaw Nyarko¹, Michael Woodford², and Nicholas C. Yannelis³

¹ Department of Economics, New York University, New York, NY 10003, USA

² Department of Economics, University of Chicago, Chicago, IL 60637, USA

³ Department of Economics, University of Illinois at Urbana-Champaign, IL 61801, USA

Many have objected to the use of the Nash equilibrium (or more generally, Bayesian Nash equilibrium) concept in game theory, and similarly to the use of the rational expectations concept in the theory of competitive markets, on the grounds that the theory assumes too much *sophistication* and *coordination of beliefs* on the part of decision-makers. The papers in this volume are among those which study the implications of relaxing those assumptions.

To provide a suitable context to describe the papers in this symposium, we will begin with a slight digression. Consider a repeated stage game where at each stage there are finitely many players, each with finitely many actions. First consider the case where each player seeks to maximize his expected utility at each date (i.e., he has a zero discount factor) given specified beliefs about the play of his opponents. Let us suppose that we have eliminated all dominated actions within the stage game. This means, in particular, that for each of the actions of each player there is a set of beliefs that a player could have on the actions of the other players which would rationalize that action as a best response for that player. If the stage game were to be played only once, each outcome to the game could be rationalized as the outcome of maximizing behavior under *some* beliefs of agents. Maximizing behavior by itself imposes absolutely no restriction on the outcome of the one-shot stage game! This is true even when the payoff matrices of the players are common knowledge. Introducing imperfect information over the payoffs does not diminish the set of possible utility maximizing outcomes. Indeed, by allowing type-dependent correlations in behavior, the introduction of imperfect information over payoffs may actually increase the set of possible (probability distributions of) outcomes.

Instead of a zero discount factor, let us now suppose that the players have a positive discount factor and seek to maximize the expected sum of their discounted payoffs. Then, even without eliminating strategies which are dominated in the normal form game, we may show via folk-theorem-type arguments that without placing restrictions on the beliefs of players we can rationalize just about any

* The first author thanks the C. V. Starr Center at New York University for their generosity.

outcome as being optimal under *some* beliefs of the players. Note well that in this result *we do not require* the beliefs of the players to be the same or to be “correct” in anyway.

The above mentioned class of results is well-documented in the literature. This is, for example, the central message of the rationalizability literature [see Bernheim (1984) and Pearce (1984)]. The result has been documented, in a form similar to that given above, by Feldman (1987), Jordan (1992) and Blume and Easley (1992)¹. Nyarko (1990) has also made this point in the context of an example of repeated play by firms facing imperfect information about a common demand curve they face. The conclusion in each of these papers is clear: maximizing behavior alone in most cases says nothing at all about outcomes unless we make restrictions on the beliefs of the agents in an economy or the players of a game. When we model players or agents as being in a Nash or Rational Expectations equilibrium, we are not only assuming that they are maximizing their expected utilities, but we are also saying something about their beliefs: they are correct and they are the same for each agent! This is obviously a very strong assumption.

The “learning” literature addresses this issue. In particular, the learning literature may be described as trying to impose weaker restrictions on the beliefs of players that still allow one to make predictions on the asymptotic behavior of players. The paper by Koutsougeras and Yannelis (1994), as well as those of Nyarko (1994), and of Kurz (1994 a, b) are written with this explicit objective. The papers by Koutsougeras and Yannelis (1994) and Nyarko (1994) restrict the beliefs of agents by assuming that *ex ante* beliefs of agents assign probability zero to the same events (i.e., their beliefs are mutually absolutely continuous), so that any event assigned probability zero by one agent will be assigned probability zero by all other agents. The paper by Kurz studies the question of selecting the appropriate prior beliefs of players using frequentist ideas.

Another line of criticism of equilibrium theory argues that agents are assumed to be too sophisticated in the ability to arrive at an optimal response to their environment (which may be very complex exactly because it depends upon the actions of other agents, who in turn are assumed to choose highly complex strategies). This has led to interest in analyses that assume “bounded rationality” on the part of agents, where the agents are assumed to follow simple “rules of thumb” in choosing their actions. The motivation for this is that the real world is too complicated, and players do not have the capacity to perform the difficult optimization exercises involved in maximizing their infinite-horizon payoffs. The “bounded rationality” literature accordingly seeks simple and intuitive rules that these boundedly rational players use in choosing their actions. The paper by In-Koo Cho (1994) studies a neural network algorithm while that of Metrick and Polak (1994) studies the “fictitious play” rule. Both rules are considered by the authors to be simple “rules of thumb” that one could reasonably expect players of a game to use.

It may seem that there is no connection between the two approaches to learning – the “restrictions on beliefs” approach on the one hand, and the “rules of thumb” approach on the other. However, the two approaches are really not as different as one may at first think. Indeed, on a formal mathematical level the two

¹ One should consult Blume and Easley (1992) for a nice survey of the recent literature on learning in games.

approaches are identical. The reason for this is the following: It is easy to prove that under fairly generic conditions, for each rule of thumb there will always exist a set of beliefs such that the given rule maximizes that player's total discounted utility under those beliefs. In many cases that belief can be chosen so that the given rule of thumb *uniquely* maximizes that player's expected utility under those beliefs. In that case choice of a belief for a player becomes somewhat equivalent to choice of a rule of thumb. In particular, rules of thumb can almost always be rationalized in terms of some, usually "incorrect," beliefs.

For example, it is well known that in 2×2 games, fictitious play is an optimal response to the belief that your opponent chooses actions via a multinomial distribution with unknown but fixed probabilities. Of course, if contrary to your belief your opponent is, like you, engaging in fictitious play then your beliefs are incorrect. In particular you are assuming your opponent is choosing actions independently over time while in fact your opponent's actions are highly dependent upon the observed history. Not only are each agent's beliefs incorrect in this case, they even violate the mutual absolute continuity restrictions on beliefs used in the papers of Koutsougeras and Yannelis (1994) and Nyarko (1994).

Finally, a response to all of the criticisms of equilibrium theory that is attracting increasing interest is to test experimentally alternative models of economic behavior. The paper in this symposium by El-Gamal, McKelvey and Palfrey (1994) provides an example of especially careful experimental methodology.

With the above discussion as a background we will now summarize the principal results of the papers in this symposium. First we present the following example:

Example (Coin-Tossing): Consider a game with two players A and B. Player A (resp. B) has two actions to choose from at each date, TOP and BOTTOM (resp. LEFT and RIGHT). The payoffs are as below:

		Player B	
		LEFT	RIGHT
Player A	TOP	1, 1	0, 0
	BOTTOM	0, 0	1, 1

Let τ^A be realization from infinitely many independent and identical coin-tossing experiments where an outcome from {HEADS, TAILS} is chosen with equal probability. Hence τ^A is an element of $\{\text{HEADS, TAILS}\}^\infty$. Let τ^B be another realization from an i.i.d. sequence of coin-tosses, $\{\text{HEADS, TAILS}\}^\infty$, which is independent of the sequence from which τ^A was obtained. At date 0 Player A is told of τ^A (and is not informed about τ^B) and player B is told of τ^B (and is not informed about τ^A). We may consider τ^A to be player A's "type" and τ^B to be player B's type. Suppose that each agent knows how the types are drawn. Consider the following play of the game. At date n Player A looks at the n -th coordinate of his sequence of coin-tosses. If it is a HEADS he plays his first action, TOP; if it is TAILS he plays his second action, BOTTOM. Similarly, if the n -th element of τ^B is HEADS player B plays the action LEFT at date n , otherwise he plays action RIGHT. Suppose further that each agent knows that the other is choosing actions via this rule.

Clearly, at the beginning of date n each agent believes the other will choose either action with equal probability. Against this probability any action is optimal. Hence under the prescribed play each agent's strategy is a *best-response* to the others. Now fix any sample path of play. It should be clear that along any such sample path (excluding a set with probability zero) each pair of actions will occur infinitely often. In particular, the actions (TOP, RIGHT) and (BOTTOM, LEFT) will be played infinitely often. Note that these pairs of actions are *not* Nash equilibria of the stage game. ■

The papers by Koutsougeras and Yannelis (1994) and by Nyarko (1994)

Both papers consider very general, very "Bayesian" models. In these models the agents play a stage game in each and every period. The question asked in both is what can be said about the play of the game in the limit as time goes to infinity. As regards the basic structure of the game, the models of these papers are general enough to cover all the other papers in this volume as special cases. However there is a difference. These papers assume a generalization of the common prior assumption: the *ex ante* beliefs of the agents share the same probability zero sets (i.e., their beliefs are mutually absolutely continuous). This assumption is for example violated in the fictitious play model studies by Metrick and Pollack (1994), and may be violated in the model of Kurz (1994a, b). The model studied by Koutsougeras and Yannelis (1994) and by Nyarko (1994), was first addressed by Blume and Easley (1984), Feldman (1987) and Jordan (1991) under the assumption of common priors. Following Jordan (1991), Kalai and Lehrer (1993) also study the repeated "matrix" game model under assumptions which allow for differences in priors but which are stronger than the assumptions in Koutsougeras and Yannelis and by Nyarko (for example the coin-tossing example violates the Kalai and Lehrer (1993) assumptions).²

The paper of Nyarko contains two results. The first says that the limit points of the empirical distributions of play are Nash equilibria of the true game. In the coin tossing example note that on each sample path (excluding a set with probability zero) the actions TOP and BOTTOM occur an average of 1/2 of the time in the limit; the same is true with the actions LEFT and RIGHT. Player A choosing TOP and BOTTOM with probability 1/2 each and B choosing actions LEFT and RIGHT with probability 1/2 each constitutes a Nash equilibrium. This therefore verifies the first result, that the limit points of the empirical distributions of play are Nash equilibria of the true game.

The second result of Nyarko states that the limit points of beliefs of agents (not conditioning on own types) are equilibria of the true underlying game. Let

² Indeed, associate with each play of an agent, i , a type, τ_i . Let μ_i denote the *ex ante* belief of player i so that $\mu_i(\cdot|\tau_i)$ denotes the beliefs of player-type τ_i . The Kalai-Lehrer assumption requires that the true play, induced by collection $\{\mu_i(\cdot|\tau_i)\}_i$, be absolutely continuous with respect to the beliefs of each agent, $\mu_i(\cdot|\tau_i)$ for each i . It is straightforward to see that this in turn implies that the set of possible types is either finite or countable, which is violated in the coin-tossing example. See Nyarko (1992) for details.

us verify this in the context of the coin-tossing example above. Fix a player, say Player A. Well, at the beginning of any date player A believes that player B will choose either LEFT or RIGHT with equal probability. However what about A's beliefs about A himself? Well, if we allow A to condition on A's own type, then A will know what he will play in that period. In particular player A will assign probability one to some pure date t strategy, TOP or BOTTOM. In either case, a pure strategy for A and probability 1/2 each on LEFT and RIGHT for B is not a Nash equilibrium. However, let us now look at A's beliefs about his own play but not conditioning on his own type. Those beliefs assign probability 1/2 to each of the actions TOP and BOTTOM (and probability 1/2 each to actions LEFT and RIGHT of B). This is indeed a Nash equilibrium. This therefore verifies the result that beliefs, not conditioning on own types, converge to a Nash equilibrium.

The results of Koutsougeras and Yannelis are related to the second result of Nyarko mentioned above, but for a more general framework. Formally, the Koutsougeras and Yannelis results make statements about the Bayesian Nash equilibria (henceforth BNE) at each date and the BNE in the limit-information game. A "limit-information game" in Koutsougeras and Yannelis is defined to be the game where the players have the information from the observation of the outcomes of the Bayesian Nash equilibria in each of the dates $n = 1, 2, 3, \dots$. For the coin-tossing model, in the limit each agent will know the entire sequence of HEADS and TAILS drawn by the other agent, so in the limit-information game there is complete information about the types of agents. (In general, however, there is not full-information of payoffs and types in the limit-information game.) The results of Koutsougeras and Yannelis are of the following kind: Let $\{x_n\}_{n=1}^{\infty}$ be the sequence of date n Bayesian-Nash equilibrium (BNE) strategies where $x_n = \{x_{i,n}\}_{i \in I}$ and where $x_{i,n}$ is the date n strategy of player i mapping agent i 's information to his/her action. Koutsougeras and Yannelis show that from $\{x_n\}_{n=1}^{\infty}$ one may extract a subsequence which converges weakly to a limiting strategy vector x^* which is a Bayesian Nash equilibrium for the limit-information game. They also show that the extraction result still holds if for all dates n , x_n is merely an ε_n -Bayesian-Nash equilibrium with $\varepsilon_n \rightarrow 0$ as $n \rightarrow \infty$. Yet another result of Koutsougeras and Yannelis states that if x_n is a pure-strategy BNE then the results just mentioned hold but where now x^* may be chosen to be a pure-strategy BNE.

It appears that the following is major difference between the papers of Koutsougeras and Yannelis and Nyarko: In the paper of Koutsougeras and Yannelis are results on *strategies* while the results of Nyarko are on empirical distributions of play and the *beliefs* of players. However, this difference is not so great. Consider again the coin-tossing example. Fix a sample path and fix for that sample path a sub-sequence of dates where the action pair (TOP, RIGHT) is chosen at each such date. On each sample path (excluding a set with probability zero) such a sequence can clearly be chosen. One may be tempted to conclude that the Koutsougeras and Yannelis result predicts that one may be able to extract from this a sub-sequence that converges to the Bayesian Nash equilibrium of the limit-information game. Well, the limit information is complete information. The only Bayesian Nash equilibria are the Nash equilibria. Since (TOP, RIGHT) is not a Nash equilibrium one may want to conclude that this is a contradiction to the Koutsougeras and Yannelis (1994) result. This is of course not the case. The mode of convergence in Koutsougeras and Yannelis (1994) is convergence in the weak-

convergence (or convergence in distribution) of the strategies, and not convergence almost everywhere. In the coin-tossing example, the *probability distribution* of the date n strategies of agents assign probabilities $1/2$ each to TOP and BOTTOM and $1/2$ each to LEFT and RIGHT. Notice that this is the same as the *beliefs* of agents at each date. For the coin-tossing example therefore the object over which the issue of convergence is studied is the same in both the Koutsougeras and Yannelis and the Nyarko papers. There is one major difference between statements of the results of two papers. Nyarko's results are in the spirit of *upper* semi-continuity results, and are of the form "limits points of beliefs are Nash equilibria." The main results of Koutsougeras and Yannelis have the flavor of *lower* semi-continuity as well. The lower-semicontinuity results of Koutsougeras and Yannelis (1994) are most probably the first in the learning literature. In both papers the beliefs of agents over the type spaces (or the space of utility parameters) need not be product measures over the individual players' type spaces. For this reason the limits obtained in that paper are more generally correlated as opposed to Nash equilibria. [The paper of Koutsougeras and Yannelis uses the more general definition of a BNE of Aumann (1987), which includes correlated equilibria as well as Nash.]

The papers by Kurz (1994a, b)

There are two papers by Kurz in this volume. The first provides the theoretical results while the second is an application of this theory to the context of a simple Muth-type continuum-of-firms market model. These papers use frequentist ideas in explaining the formation of prior probability beliefs. [See for example Hempel (1994).] The scenario envisioned by this paper may be summarized as follows: An individual needs to form a probability belief over the probability distribution governing a stochastic process of interest, which we refer to as $\{x_t\}_{t=1}^{\infty}$. That individual is assumed to have observed a sample path of this process for a large number of times *before* forming his prior. The individual then estimates the probability distribution governing this process. This "empirical distribution" is obtained by proceeding as if the underlying process is stationary, in which case the long average number of times the process falls in a given set will determine the probability of that set under the empirical distribution. Let us denote the empirical distribution on the given sample path by m .

With the above empirical distribution m , the agent now needs to form a belief over the true probability distribution governing the underlying stochastic process. Of course there are many probability distributions for the underlying process which give the same *long run* empirical distribution, m . This is because we are not restricting attention to only stationary processes. In particular, there may be processes which have strange behavior over the first few dates but then settle down over time; or processes which behave erratically on dates farther and farther apart – i.e., at "remote" times. In either of these two cases, the long run empirical distribution may be equal to m . Kurz imposes axioms of rationality on agents requiring that their beliefs be "consistent" with the observed long-run empirical distribution: their beliefs should generate the same long run distribution and should assign positive probability to any event to which m assigns positive probability.

Kurz provides a characterization of the set of beliefs consistent with these axioms. Each of such beliefs is made up of two parts, a stationary part which is consistent with the empirical distribution and a “remainder” non-stationary part.

To further understand the Kurz result let us go back to the coin-tossing example. Let us consider the repeated play of the coordination game in the coin-tossing game discussed earlier. Suppose that player B is just as in the coin-tossing experiment, and chooses actions via the coin-toss. Suppose, however, that player A obeys the axioms of the Kurz paper. How would Player A behave in this case? Well, even to, begin to analyze this problem the Kurz method requires us to endow Player A with observations – an infinite history of them. It is not clear which data we should use and where this should come from. The natural thing to do, it seems, is to suppose that Player A has observed a previous game where the same Player B played a game. Player A then uses the data from that previous game of player B to compute an empirical distribution m . This empirical distribution will, of course, show that the actions of Player B are i.i.d. with probability of LEFT and RIGHT equal to $1/2$ in each period. The axioms of Kurz require that Player A’s beliefs can be decomposed into a stationary part and a non-stationary part. All that is required of the stationary part is that it be mutually absolutely continuous with respect to the empirical measure. In finite time (i.e., on finite-dimensional events) such a measure could be very different from m . The mutual absolute continuity assumption will, however, pin down the stationary part on tail events (i.e., events which involve distant futures and limits). But then the non-singular part will allow us degrees of freedom to vary the beliefs of agents. Hence, even in very simple problems like the coin-tossing example, the Kurz axioms allow for a rich diversity of beliefs of agents. Indeed, the beliefs of agents could violate the mutual absolute continuity assumption used by Koutsougeras and Yannelis (1994) and by Nyarko (1994).

In the second paper of Kurz in this volume, an application of the above ideas is provided. A model with a continuum of firms is studied³. Kurz (1994a, b) shows how in that model, due to the diversity of beliefs, and in particular the non-stationary components allowed, there may be “excess volatility” in the economy over time. By weakening the assumption of rational expectations there will in general be a much *larger* set of possible outcomes.

The papers by Cho (1994) and by Metrick and Polak (1994)

The paper of Cho (1994) studies a model where the players in an infinitely repeated game use neural networks to implement their strategies. Under the classification scheme mentioned earlier, this paper therefore falls in the class of models which have as their principal motivation the use of simple “rules of thumb.” The particular issue this paper focusses on is the folk theorem. Recall that the folk theorem [see Fudenberg and Maskin (1986)] concludes that all individually rational payoffs can be obtained as the average payoffs of a subgame

³ Feldman (1988) also studies learning in the Muth model of firms, but imposes the mutual absolute continuity assumption of Koutsougeras and Yannelis (1994) and Nyarko (1994).

perfect Nash equilibrium of the infinitely repeated game with discounting. This paper takes as its starting point the observation that to obtain the folk theorem result one has to use very complex strategies, and that this may cast doubt on the appropriateness of the folk theorem as a description of the behavior of “real-life” players. Neural networks on the other hand, argues this paper, are indeed simple rules of thumb that one could expect real-life players to use. The main result of this paper is that the full menu of outcomes predicted by the folk theorem may be obtained when players are restricted to using only neural networks.

In the equilibrium of this paper note that the players are in a Nash equilibrium. In particular, each agent will know how the other is playing and then take a best-response to that. Agents are restricted to the use of simple strategies, but they optimally choose their strategies within this class. Further, the outcome studied is actually a Nash equilibrium where each player knows the strategy choice of the other. “Bounded rationality” in the sense of the Cho paper is quite distinct from an analysis of “learning,” for Cho *assumes* that players’ beliefs about the environment are correct, as in standard Nash equilibrium theory. On the other hand, restriction of the set of strategies in this way addresses some of the concerns about how agents can come to have correct information about their environment. A restriction on the set of strategies that can be chosen reduces the amount of information about a player’s environment that it is relevant to obtain, and the *knowledge* that other players choose a strategy from a restricted set provides a restriction on beliefs that can make convergence to correct beliefs easier. Thus the sort of analysis undertaken here by Cho equilibria in simple strategies provides a natural basis for an attack on the learning problem as well, even if that is not attempted here.

The paper by Metrick and Polak (1994) studies the classic fictitious play algorithm. This paper presents a very nice geometric proof of the convergence of fictitious play to Nash equilibria in 2 player–2 action games. Of course, fictitious play may be considered a very simple “learning” rule or “rule of thumb” where agents predict that the probability that any of their opponents will choose a given action is equal to the proportion of times that strategy has been used in the past. This rule of the thumb is consistent with the basic structure used in the papers of Koutsougeras and Yannelis and Nyarko but will violate their mutual absolute continuity assumption on the beliefs of agents. It should also be easy to see that the fictitious play rule can actually be implemented by a neural network of the type studied by Cho. However, unlike in the Cho paper, agents are not in a Nash equilibrium when they each use fictitious play against each other.

In passing it is interesting to note how a fictitious play agent will do against a player in the coin-tossing game. In particular, suppose that player A is a fictitious play agent and plays as in the Metrick and Polak model. Suppose that Player B plays as in the coin-tossing model. To make the coin-tossing model non-trivial let us suppose that the coin which is used in obtaining Player B’s type is not a fair coin but is a coin for which HEADS and TAILS occur with probabilities θ and $1-\theta$. Suppose that player A knows how B chooses his actions, but does not know the value of θ . Then it is easy to see that in the *long-run* fictitious play is in optimal against *any* prior belief over θ . In the short-run fictitious play is actually optimal if the prior over θ is multinomial (with the histogram of past plays slightly modified to include the prior mean over θ).

The paper by El-Gamal, McKelvey and Palfrey (1994)

This paper is an experimental study of the centipede game of Rosenthal (1992). The “rational” or backward induction solution is to “take” immediately resulting in a low payoff – an action, typically not observed in the experiments. This paper supposes that some fraction, q , of the population of possible opponents are “irrational.” Each pair of experimental subjects plays the game twice. The authors look at two hypotheses about the behavior of the rational participants of this game: the myopic and the sequential hypotheses. After playing the first game the agents will receive information which sheds light on whether their opponent is irrational. In the myopic case the agents, after playing the first game, do not update their prior probabilities in light of this new information. In the sequential model, on the other hand, the students update their probabilities. The paper then proceeds to first compute the sub-game perfect Nash equilibrium under these two hypotheses and then to test which of the two hypotheses best fits the data from the experiments, and indeed whether their model fits the experimental data.

Their first main result is that given the observed experimental data, the model of irrationality used is actually not good enough in discriminating between the two hypotheses, myopic versus sequential. The authors then change the model of irrationality. In the new model there are two types of irrational individuals. Those who randomize over the available actions at each node as in the previous model; and, in addition, there are the altruists who pass at each node in an attempt to implement the cooperative solution. With this new model the authors conclude that the sequential model was accepted. In particular, in solving this game the students actually do the correct Bayesian updating!

On such a positive note we end this summary of the papers. Enjoy.

References

- Aumann, R.: Correlated equilibrium as an expression of Bayesian rationality. *Econometrica* **55**, 1–18 (1987)
- Bernheim, B. D.: Rationalizable strategic behavior. *Econometrica* **52**, 1007–1028 (1984)
- Blume, L., Easley D.: What has the rational learning literature taught us? Manuscript, Cornell University 1992
- Cho, In-Koo: Bounded rationality, neutral network and Folk Theorem in repeated games with discounting. *Econ. Theory* **4**, 935–957 (1994)
- El-Gamal, A. A., McKelvey, R. D., Palfrey, T. R.: Learning in experimental games. *Econ. Theory* **4**, 901–922 (1994)
- Feldman, M.: An example of convergence to rational expectations with heterogeneous beliefs. *Int. Econ. Rev.* **28**, 635–650 (1987)
- Fudenberg, D., Maskin, E.: The Folk theorem with discounting or with incomplete information. *Econometrica* **54**, 522–554 (1986)
- Hempel, F.: Some thoughts about the foundations of statistics. In: Morgenthaler, S., Ronchetti, E., Stahel, W. A. (eds) *Directions in data analysis and robustness*. Basel: Birkhäuser 1994
- Jordan, J. S.: Bayesian learning in normal form games. *Games Econ. Behav.* **3**, 60–81 (1991)
- Jordan, J. S.: Bayesian learning in games: A nonbayesian perspective. Manuscript, University of Minnesota 1992
- Kalai, E., Lehrer, E.: Rational learning leads to Nash equilibrium. *Econometrica* **61**, 5, 1019–1046 (1993)
- Koutsougeras, L., Yannelis, N. C.: Convergence and approximation results for non-cooperative Bayesian games: Learning theorems. *Econ. Theory* **4**, 843–857 (1994)

- Kurz, M.: On the structure and diversity of rational beliefs. *Econ. Theory* **4**, 877–900 (1994a)
- Kurz, M.: On rational belief equilibria. *Econ. Theory* **4**, 859–876 (1994b)
- Metrick, A., Polak, B.: Fictitious play in 2×2 games: A geometric proof of convergence. *Econ. Theory* **4**, 923–933 (1994)
- Nyarko, Y.: Bayesian learning without common priors and convergence to Nash equilibrium. C. V. Starr Working Paper No. 90-45, New York University (1990)
- Nyarko, Y.: Bayesian learning without common priors and convergence to Nash equilibrium. C. V. Starr Working Paper No. 92-25, New York University, (revised Summer version of Fall 1993)
- Nyarko, Y.: Bayesian learning leads to correlated equilibria in normal form games. *Econ. Theory* **4**, 821–841 (1994)
- Pearce, D.: Rationalizable strategic behavior and the problem of perfection. *Econometrica* **52**, 1029–1050 (1984)
- Rosenthal, R.: Games of perfect information, predatory pricing, and the chain store paradox. *J. Econ. Theory* **25**, 92–100 (1982)