Discussion of: "Foundations of game theory," by Kenneth J. Binmore and "Refinements of Nash equilibrium" by Eric van Damme

Eddie Dekel*

> [A] unified theory for all non-cooperative games does not seem possible. The only alternative seems to be ... introducing ... information referring to personality traits, psychologies of the players, etc. [E]ven if ... examples ... throw doubt upon the universality of a concept, this does not necessarily undermine its importance. It merely establishes that care must be exerted to check whether the concept is plausible in the specific cases to which it is applied. Ideally, one should attempt to investigate the mathematical restrictions which should be placed on the domain of admissible games so that the concept is plausible.
>
> Luce and Raiffa (1957, pp. 104–5)

1 SWIMMING THROUGH THE FLOOD

Binmore believes "the foundations of game theory to be a mess." He attributes this mess in part to Bayesian foundations being "clinker-built," and in part to their failure to solve the problem of multiple equilibria. Binmore claims that foundations are "clinker-built," because they do "not address the conscious choice question"; and he observes that the problem of multiple equilibria has not been solved because refinements draw "contradictory conclusions from hypotheses said to be 'reasonable' or 'plausible'." Yet, despite this dismal view of the current state of foundations, he emphasizes that providing foundations is important since, "As regards game theory, too many of our wonderful results are wrong."

Binmore concludes that "it is necessary to model the way players think," "we need to be told something about the equilibrating process," and in fact "if the issue of conscious choice is to be faced squarely [it] would involve actively modeling the players as computing machines." Thus, he favors dynamic models of automatons that "classify solution concepts according to the environments in which it makes sense to apply them."

In contrast to Binmore, I do not think that foundations are a mess or clinker-built. Nor would I agree that our results are (in general) wrong. But current foundations are inadequate; so I do agree with Binmore insofar as he concludes that *more* foundations are needed. We need deeper foundations that examine the justifications, implications, and modifications of assumptions underlying game-theoretic tools. My view is that research using the Bayesian framework should continue:[1] assumptions (including common knowledge of rationality) should be examined further, weakened (for example, by considering that information is not correctly or costlessly processed), and added (e.g., people communicate, or strategies may have different costs of implementation).

Binmore proposes foundations that develop machine-based dynamic models of thought processes. While this is important, I will argue that furthering the Bayesian paradigm and developing other non-Bayesian approaches are no less important, despite Binmore's and van Damme's shared skepticism of psychological and behavioral assumptions. Assuming that people make decisions in accordance with an evolutionary model of perturbed automatons is no less problematic than exploring Bayesian learning, psychological assumptions, or other paradigms for decision-making; all require justification. Because I strongly agree with Binmore that we should categorize game-theoretic tools according to the environments in which they are appropriate, I view his focus on how players make decisions, and in particular on modeling players as machines, as too restrictive.

I will first review why foundations can be useful; and I will briefly argue that Binmore's goal of foundations that do not depend on intuition is unattainable: foundations inherently rely on intuition. I will then comment on the issue of conscious choice in Bayesian modeling. The main part of this chapter will raise a few questions regarding Binmore's general and specific proposals that we model how players think using explicit dynamics of game-playing machines. Because of my pessimism that any one approach will come close to solving the problems of foundations and game theory, I will conclude by supporting a more general research agenda.

To place my discussion of foundations in context I should first mention my view of its role in the broader area of economic theory. Models identify important features in economic problems, and verify the implications of those features using assumptions about human interaction. Decision theory synthesizes such assumptions to yield solution concepts. Foundations, in turn, examine these assumptions in depth. Foundations should help us comprehend which assumptions drive our results, evaluate informally (by introspection and intuition) their plausibility in various environments, analyze the robustness of results to modifications in the assumptions, and, at its best, enable us to formally analyze the assumptions using alternative and even more primitive assumptions.

2 FOUNDATIONS: MISPERCEPTIONS, MULTIPLICITY, AND THE PROBLEM OF INTUITION.

Foundations focus our attention on the implicit hypotheses that underlie our assumptions, and on the informal arguments that justify these assumptions. This compels us to examine whether the implicit hypotheses are consistent with other assumptions we make, and whether our informal justifications can be formalized. Several insights have been gained by carefully considering our solution concepts and assumptions. For instance, preceding Bernheim's (1984) and Pearce's (1984) work, there seemed to be little awareness that common knowledge of rationality does not imply Nash-equilibrium behavior. Also, in their talks, both Binmore and van Damme discussed the less-known fact that common knowledge of rationality (appropriately defined for extensive-form games) does not imply backwards induction (see Binmore for references).[2,3]

Another (related) reason to study foundations is to deal with the multiplicity of equilibria. If game theorists don't know which equilibrium is "right," then how do the players? If there are aspects that are not incorporated into the game and that help players know which strategies they should be playing, then explicitly incorporating these aspects may change the game. As Binmore emphasizes, how we choose an equilibrium concept is not independent of how we choose among equilibria.[4]

Most researchers agree that formal modeling is important, and are careful in their modeling assumptions: for example, what aspects of the environment are observable; whether the interaction is one-shot or repeated; whether we should assume that prices are sticky; whether it is valid to restrict attention to a subset of non-optimal mechanisms, or to linear pricing rules. Moreover, many argue that modeling assumptions should be explored (e.g., why are some aspects observable, why are prices sticky); and these researchers often reject unexamined modeling assumptions that simply yield more precise, or even "better," predictions.

By contrast, there seems to be a view that economists need not explore solution concepts because one can adopt those that are "useful." A more balanced view is that economists should be held equally responsible for their modeling assumptions and for their choice of solution concepts; the separation of assumptions about the model from those regarding the solution concept, when both ought to be appropriate for the environment, seems indefensible. I admit (and argue below) that assumptions are necessary; I want to emphasize here Binmore's point that foundations are needed – expediency does not justify leaving any fundamental assumptions unexplored.

There is, however, an important and inherent limitation of any foundations for game theory. Since all work begins with hypotheses, such as Binmore's hypothesis that automatons are – and neural networks are not – a good approximation of human decision-making, in the end we all rely on notions of intuitive appeal, introspection, and reasonableness. Binmore seems to criticize this use of intuition, while I see no way for game theory to avoid it. For example, should we reject *all* understanding gained about incredible threats because backwards induction is not the consequence of Bayesian rationality? The lack of a formal justification of an assumption cannot make it unacceptable, since assumptions at some level are always needed. Unexplored assumptions are, however, a signal to proceed with care. While a thorough comprehension of backwards induction continues to elude us, a paper can warily use backwards induction because of its intuitive appeal. It may, in fact, be an appropriate tool precisely because so many people find it appealing. (Its origin long before other equilibrium concepts supports this view.)

Of course, I agree with Binmore that an interesting and important research question is: "Why is an assumption intuitively appealing?" Without an answer, the assumption should be treated with extreme caution. Comprehension of an assumption's appeal can tell us if a result that relies on the assumption is robust and internally consistent. Most importantly, formal analysis of an assumption can explain why it appears intuitive in some contexts and not in others. In the example of backwards induction, it seems clear that a tension exists: it is appealing in some contexts, but not others; foundations should tell us why this is the case. Until foundations explain the cause of this tension, we can neither categorically reject, nor wholeheartedly accept, backwards induction.

In summary, I agree with Binmore that much can be learned by asking why an assumption is appealing. My disagreement with Binmore is more modest: I do not view dynamic models of thought either as necessary or as promising as does he. Yes, it is important to examine assumptions that underlie a model; however, the search for the most elementary assumptions is (I believe) as elusive as that for elementary particles. Just as physicists rightly do not restrict attention to the search for elementary particles, it could be extremely costly for research on foundations to await the development of a correct description of decision-making, despite the obvious importance of such a model for the discipline of decision theory.

3 WHICH FOUNDATIONS?

Binmore argues that the current state of foundations, as examplified by Aumann's (1987) "common-knowledge-of-Bayesian-rationality approach," is dismal. One major criticism he makes is that the Bayesian framework does not allow for conscious choice, since states are assumed to completely describe the world, including players' actions. This assumption in turn is used to justify the hypothesis that the players' information structure is common knowledge, which is a necessary ingredient in interpreting Bayesian foundations of solution concepts.

This criticism, which leads Binmore to conclude that Bayesian foundations are clinker-built, seems too strong. The Bayesian framework can allow (at least some aspects of) conscious choice, since it does not require that all players face the same state space. Thus each player could be uncertain about the state of the world – which is a complete description of the world *excluding* her own choice. In such a model common knowledge of rationality still leads to rationalizability, while players' choices are not determined by the state space, so that the model allows for free choice.[5] Thus, I take issue with Binmore's argument that a Bayesian Ark *must* be clinker-built.

Nevertheless, I strongly agree with Binmore that Bayesian foundations based on rationality alone have made little progress toward solving the problem of multiple equilibria. In fact, these foundations have exacerbated the problem by questioning backwards induction and showing the implausible nature of assumptions that appear necessary (within the "common-knowledge-of-Bayesian-rationality approach") to justify Nash equilibrium and its refinements. So the question is where to proceed next. Obviously it is time to examine the implications and reasonableness of additional assumptions.

Binmore proposes that we develop dynamic models of the decision-making process in order to understand how players reach an equilibrium. Aware of the various dynamics that can be proposed, he recommends using them to categorize solution concepts according to the environments in which they should be applied. His image is one of an armada of ships – each ship representing a dynamic thought process appropriate for some environments. He contrasts these dynamic models of thought processes with introspection and intuition, which are used to justify the hypotheses made in papers on learning and on refinements of Nash equilibrium. He sees introspection and intuition as providing "just cause" to those who describe game theory as pseudo-scientific.

Although I agree with the importance of dynamic models of thought, I doubt that they will save game theory from the contradictory conclusions that have provoked such criticisms. The same enlightened view of dynamic models, namely as categorizing environments, should be applied to other proposed hypotheses. Moreover, I doubt that the dynamics, let alone the models of our thought processes, will be so clearly related to "reality" that the resulting classification will be much simpler or more natural than alternative classifications based on restrictions on beliefs, physical environments, or behavioral hypotheses.[6] Thus, I fail to see the distinctive advantage of dynamic models of thought over, for example, learning models with assumptions about the way players experiment (see, for example, Fudenberg and Kreps, 1988).

I agree with Binmore that the dynamics by which equilibrium is reached provides a useful, under-explored direction for analysis.[7] Moreover, Binmore is certainly right that it would be good to open the black box of Bayesian beliefs and consider what various assumptions regarding the environment imply about the players' beliefs. However, if the Bayesian paradigm is to be challenged, it should be thoroughly overhauled, so that alternatives do not implicitly adopt a half-Bayesian model. In particular, it is equally important to explore the other sacred cow of the Bayesian paradigm: what determines preferences? In traditional short-run models it may be appropriate to take preferences as given. However, assuming fixed preferences seems questionable in precisely the evolutionary environments that Binmore advocates considering.

In any case, Binmore's more particular thesis is that we should model the dynamics of decision-making using computing machines (Turing machines or automatons for example), in order to have a model which would allow us to characterize what players would do if they were a little different. He cites interesting work in support of this thesis: Binmore and Samuelson (1990) and Fudenberg and Maskin (1990) both derive surprisingly strong results in favor of efficiency by applying an equilibrium concept based on evolutionary arguments (ESS) to an infinitely repeated game without discounting, and that is played by finite automatons.[8] Yet, I fail to see why such "metaphors for an evolutionary process," that explore beliefs using preferences based on computer science, are different at a fundamental level from psychological assumptions, or other restrictions on beliefs, strategy spaces, or learning processes.

I have argued that the Bayesian paradigm, including learning models, are no less appropriate for foundations than are models using machines. I also want to emphasize that the existing literature on automatons lies squarely within the Bayesian paradigm: recall that researchers have already made the intuitive claim that players focus attention on "simple" (e.g., stationary, Markov, or trigger) strategies. Instead of adopting Binmore's view that automatons describe human decision-making, one can take the more conventional view that they are an *assumption* within the Bayesian model: players prefer "simple" strategies. Automatons have some advantages over other *ad hoc* restrictions on strategy spaces. First, they appeared in a different literature, so are less likely to be biased in favor of results that game theorists are seeking. Second, they are easily compared to other models of machines and bounded recall in order to get a sense of their robustness. Finally, since computer scientists rank machines and problems in terms of memory and complexity, it is possible – but questionable – to assume that players' preferences rank them similarly, and hence to use traditional economic arguments regarding tradeoffs.[9]

As an aside, I would like to discuss the distinction Binmore draws between classifying "nearby worlds" according to whether they model differences in the physical world or differences in the players' thought processes. I wonder whether, in the end, such a distinction is meaningful. As far as I know, few game theorists view the literature on reputation and incomplete information as modeling perturbations of the physical environment; rather these perturbations of the payoffs are best thought of as metaphors for perturbing the players. In fact, it seems impossible to distinguish empirically between a player whose payoffs are different and who thereby behaves "irrationally" and a player whose thought process is irrational. When Fudenberg and Levine (1989) prove that a long-run player can achieve high payoffs by "imitating" a player with "Stackelberg payoffs," isn't it reasonable to view the "Stackelberg payoffs" as a summary for a player whose reasoning leads him to play as if his payoffs were as they specify? I conjecture that results that are robust to a wide variety of "physical" perturbations will be robust to a wide variety of "mental" perturbations, and conversely.

Since I feel that we are stuck with relying on intuition at every level – both specifying the model and developing solution concepts – I would like to repeat: insights on solution concepts can be gained by other methods, in addition to formally modeling the decision-making process; and other methods are no more suspect than Binmore's program, since the next level of assumptions that underlie all the models will be justified by intuition, and the models are equally "pseudo-scientific." So I now consider possible sources for intuitive *assumptions* about decision-making other than machines.

Outside of game theory, behavioral hypotheses have proven useful: Machina's (1982) Hypothesis II clarifies the relationships among many empirical "paradoxes" (e.g., preference reversals and the Allais paradox (Safra, Segal, and Spivak, 1989)). In games, for example, one might want to model the psychological feature that the players' expectations over outcomes might affect the games' payoffs (that is, the beliefs enter the payoffs). This interesting idea was proposed and modeled by Geanakoplos, Pearce, and Stacchetti (1989). As another possibility, Kahneman and Tversky (1989) have argued that the value assigned to an object is a function of whether it is owned. A related idea is that the sequence of offers in a bargaining game effects the preferences (or behaviors) in the induced subgame. Based on this, one might argue that offers (even after being rejected and withdrawn) are *perceived* by players as relevant in that they indicate a willingness to concede in the future. This implies that in finite-horizon bargaining games, delay will occur when players are sufficiently patient (Fershtman and Seidmann, 1990).

A category of behavioral assumptions concerns communication, or cheap talk, which has been used informally to justify Nash equilibrium as a solution concept. The argument that models of communication suffer from endless debates over "intuitive" assumptions and from limited success appears to lend support for Binmore's criticisms. I disagree. Although there are cases in which it is not clear which are the right assumptions, there are certainly unambiguous cases where all assumptions yield similar conclusions. (Interestingly, some of these cases correspond to games where other refinements have led to much controversy and ambiguity.) Moreover, cheap-talk models offer interesting negative results: informal claims regarding communication yielding Nash equilibrium and Pareto optimality are not supported by formal analysis, even with fairly strong assumptions about communication.[10] Finally, since communication plays a significant role in economic environments, it seems to me that much remains to be learned by considering behavioral assumptions about how people talk and by developing new models of communication.

There are at least two tempting criticisms of psychological and behavioral assumptions: they introduce the complexity which game theory tried to abstract into the payoffs, and the flexibility allowed by introducing psychological assumptions can destroy some of the discipline that the rational-player model imposes. However, these problems also seem to arise when introducing restrictions on strategies based on computer science. Judicious use of psychological assumptions, even in the absence of a complete model of our psyche, should be no less insightful than using automatons for modeling thought processes. And I agree with Binmore that formal models should not forever avoid the complexities of reality or the care required when a restraining straight-jacket is removed.

Nevertheless, exciting ideas and interesting issues will result not only by imposing assumptions from outside economics (e.g., psychology, cheap talk, and computer science) but also by using more conventional Bayesian methods to explore and relax the assumption that rationality is common knowledge. Binmore's very statement that "what [common knowledge of rationality] means is less than clear" suggests that significant questions remain unanswered. The notion of forwards induction (see Kohlberg and Mertens, 1986) is a recent and important contribution to our understanding of rationality.[11] The traditional approach to relaxing the assumption that rationality is common knowledge is based on games of incomplete information. Using this approach, Carlsson and van Damme provide an interesting justification of risk dominance. They show that in any Nash equilibrium (and in some cases, any rationalizable outcome) of a two-by-two game with incomplete information, as the uncertainty vanishes, players choose the risk-dominant equilibrium of the (almost certain) game. Their global game sets up another flotilla: depending on the environment, modeled by different forms of incomplete information, different selection criteria, such as risk dominance, may be justifiable.[12]

4 CONCLUSION: MANY ARMADAS ARE NEEDED

I disagreed with Binmore's claim that the conscious-choice problem (described in sections 2 and 3) implies that Bayesian foundations are clinker-built. However, I do agree that foundations have so far focused attention on problems, rather than providing positive results. I also agree with Binmore that it will be useful to have formal models of how players think (perhaps as machines). Yet, we will benefit by studying Bayesian learning, by incorporating into the Bayesian model assumptions drawn from many fields, and, of course, by developing new paradigms. Thus, many (contradictory) classifications of solution concepts are appropriate: armadas based on models of machines, learning, incomplete information, psychological features, and other aspects of the environment, will help us intuitively select a solution concept.

It is unlikely that criticisms about the pseudo-scientific nature of game theory will be less forceful even if we pursue dynamics and automatons wholeheartedly. Arguments over what are reasonable restrictions on beliefs will be replaced by arguments over reasonable dynamics. The same contradictory conclusions that Binmore attributes to intuitive restrictions on beliefs will arise from "plausible" dynamics. And I doubt that the dynamics – let alone our models of how we think – will be so well based in reality that we will actually be able to associate the different dynamics with clear-cut economic environments.

Game theory, especially if viewed as a branch of decision theory, would certainly benefit from an understanding of how people reach decisions. But the view that such an understanding, especially based on dynamic models of machines, can – let alone will – actually solve the problems from which decision theory suffers seems overly optimistic. I reject the view that such models are currently *necessary* for further insights, even though they are an admirable objective.

Notes

* I would like to thank Jean-Jacques Laffont for the enjoyable and exciting invitation to discuss these papers. Conversations with Elchanan Ben-Porath, Adam Brandenburger, Faruk Gul, and especially the patience and comments of Matthew Rabin and Suzanne Scotchmer have contributed to this discussion. Financial support from the Sloan Foundation the Econometric Society and NSF grant SES 88-08133 are gratefully acknowledged. My focus will be on Ken Binmore's paper, and I will relate it to two recent papers by Eric van Damme (which I received in lieu of a paper for this session).

1 The Bayesian paradigm examines the decision-theoretic implications of various hypotheses about players' beliefs and utilities, usually including variants of "rationality is common knowledge."

2 Moreover, having accepting backwards induction as a starting point, if a player sees it violated should she continue to believe in it? The traditional logic of backwards induction requires an unambiguous yes, but it is far from clear that this is reasonable. Van Damme and Noldeke show that a similar problem occurs in refinements of equilibrium that impose the "support restriction." This restriction requires that if an equilibrium leads a player to believe an event is impossible then she should never change that view, even off the equilibrium path. This ignores the fact that being off the path violates the equilibrium hypothesis used as a starting point, and hence it is again not clear that such a refinement is reasonable. The problems with the support restriction and with backwards induction both result from the difficulty in analyzing counterfactuals. The fact that very similar problems with apparently intuitive arguments arise repeatedly emphasizes the need to examine the foundations of solution concepts.

3 The literature on cheap talk (begun by Crawford and Sobel, 1982 and Farrell, forthcoming) provides another example of the insights gained from formalizing assumptions. For example, it may be appealing to assume that players choose some Pareto optimal equilibrium; and it may also be appealing to assume that we should allow *any* of the Pareto optimal equilibria. This is the approach of much of the literature on renegotiation in games (e.g. Bernheim and Ray, 1989 and Farrell and Maskin, 1989). However, Rabin (1990) has argued that explicit modeling of symmetric communication (which could be thought of as the reason that these two assumptions are jointly appealing) *cannot* justify them together (except for the trivial case of a unique Pareto dominant equilibrium). For related examples, see footnotes 5 and 11.

4 A simple and natural example concerns the literature on communication: as Myerson (1986) and Forges (1986) have shown, introducing communication changes the game. So, using communication to select among outcomes (in the

modified game) can lead to an outcome that is a correlated, rather than Nash, equilibrium of the original game.

5 A problem is that in such a model there seems to be no appealing analog to the common prior assumption (CPA), let alone a justification for it. But, at the least, as shown by Brandenburger and Dekel (1986), conscious choice seems to be allowed if one is willing to forego the CPA.

A similar point can be made from another, more speculative, perspective. Following his "conscious choice" criticism of Bayesian foundations Binmore discusses the Bayesian model of "knowledge." He claims that foundations must address how knowledge is attained. Binmore therefore proposes modeling knowledge as a consequence of proving theorems, and for this purpose an appropriate logical foundation is the modal logic of provability, (G). Brandenburger, Dekel, and Geanakoplos (forthcoming) show that the set of equilibria attained by allowing for information structures which, in a sense, correspond to (G) (in that they drop the assumption that knowledge implies truth) is equal to the set of equilibria attained using the standard partition model (that corresponds to S5) so long as the CPA is not imposed. Thus, if one is willing to forgo the CPA, Binmore's recommendation that information and knowledge be modeled using (G) need not lead to different sets of equilibria than those which result from the standard partition model corresponding to S5.

6 Dynamic models seem very sensitive to their specification. For example, assuming discrete versus continuous time, or including mutations and other forms of noise can have significant (and non-robust) implications. Can we really say what economic environments correspond to different models of mutation? What forms do mutations in economic environments take? Moreover, the properties of "limit points" may depend on whether one looks at attractors, stable points, asymptotically stable points, accumulation points, cycles, etc. Which characterizes the equilibrium we are at (if, in fact, we are in a convergent environment)?

7 This view is substantiated by the fact that the very originators of evolutionary dynamics, after turning to game theory, have become concerned about what qualifies as reasonable equilibria: "theories of almost limitless craziness can no longer be ruled out on commonsense grounds. If we observe an animal ... standing on its head instead of running away from a lion, it may be doing so to show off to a female. It may even be showing off to the lion." But natural "selection alone is entitled to judge" (Dawkins, 1989, p. 313). Nevertheless, since nature offers biologists many more *real* life experiments than those offered economists, I hope economists do not abandon intuitively informed arguments. A model that, *in the end*, does not convince one of its intuitive appeal should and will be ignored.

8 These papers are interesting and thought provoking. However, it is not clear that the use of infinitely repeated games without discounting is reasonable. Such games may be appropiate as an approximation for how players perceive long finite games. However, since the players receive payoffs and reproduce according to the actual game, not their perceptions, this approximation does not justify using such games in evolutionary models. Moreover, in terms of the

dynamics, it is disconcerting to (implicitly) use an evolutionary model that involves repetition of a game that never ends. An evolutionary model of overlapping generations seems much more "intuitive," and a promising avenue of research that has yet to be investigated. Finally, ESS is no more a "libration" than are other refinements: it is informally motivated using evolutionary arguments, but formal dynamics yield quite different predictions. My point is that this type of research can offer insights, despite hypothesizing a dynamic process that cannot be met in any environment, and despite failing Binmore's "libration" test.

9 Perhaps I should address the criticism that my argument simply lumps everything into the Bayesian paradigm, thereby achieving nothing. My point is that many aspects of how we think can and should be explored by taking beliefs and utilities as a starting point, and imposing assumptions. I conclude that automatons are only one way of formalizing specific types of assumptions. Regarding the more speculative search for new paradigms, I think that focusing on machine-based models of human thinking is no more promising than other, e.g., behavior or psychology-based, lines of research.

10 A commonly known language appears insufficient to guarantee equilibrium outcomes (Farrell, 1988; Rabin, 1990). Rabin, however, has shown that in combination with natural behavioral hypotheses, the set of payoffs a player can expect under rationalizability coincides with the set of expected Nash-equilibrium payoffs to that player. So, in one sense, the distinction between rationalizability and Nash equilibrium disappears due to communication (see also footnote 4).

11 Another recent example is Gul (1990), who focuses on results that do not rely on equilibrium hypotheses.

12 I doubt, however, that *these* forms of incomplete information correspond to real features of a modeling environment, though such a correspondence is necessary for a classification of solution concepts to be useful.

References

Aumann, Robert J. (1987), "Correlated Equilibrium as an Expression of Bayesian Rationality," *Econometrica*, 55(1): 1–18.

Bernheim, Douglas B. (1984), "Rationalizable Strategic Behavior," *Econometrica*, 52(4): 1007–28.

Bernheim, Douglas B. and Debraj Ray (1989), "Collective Dynamic Consistency in Repeated Games," *Games and Economic Behavior*, (4): 295–326.

Binmore, Kenneth G. and Larry Samuelson (1990), "Evolutionary Stability in Repeated Games Played by Finite Automata," University of Michigan Discussion Paper.

Brandenburger, Adam and Eddie Dekel (1986), "Bayesian Rationality in Games," Working Paper, Churchill College, Cambridge University.

Brandenburger, Adam, Eddie Dekel, and John Geanakoplos (forthcoming), "Correlated Equilibrium with Generalized Information Structures," *Games and Economic Behavior*.

Crawford, Vincent and Joel Sobel (1982), "Strategic Information Transmission, *Econometrica*, 50: 1431–51.
Dawkins, Richard (1989), *The Selfish Gene*, 2nd edition, Oxford University Press.
Farrell, Joseph (1988), "Communication, Coordination, and Nash Equilibrium," *Economic Letters*, 27: 209–14.
(forthcoming), "Meaning and Credibility in Cheap Talk Games," in M. Demster (ed.), *Mathematical Models in Economics*, Oxford University Press.
Farrell, Joseph and Eric Maskin (1989): "Renegotiation in Repeated Games," *Games and Economic Behavior*, 1(4): 327–60.
Fershtman, Chaim and Daniel J. Seidmann (1990), "Deadline Effects and Inefficient Delay in Bargaining with Endogenous Commitment," MEDS, KGSM Discussion Paper, Northwestern University.
Forges, Françoise (1986), "An Approach to Communication Equilibria," *Econometrica*, 54(6): 1375–86.
Fudenberg, Drew and David M. Kreps (1988), "Learning, Experimentation and Equilibrium in Games," Working Paper, Massachusetts Institute of Technology.
Fudenberg, Drew and David Levine (1989), "Reputation and Equilibrium Selection in Games with a Patient Player," *Econometrica*, 57(4): 759–78.
Fudenberg, Drew and Eric Maskin (1990), "Evolution and Cooperation in Noisy Repeated Games," *American Economic Review*, 80(2): 274–9.
Geanakoplos, John, David Pearce, and Ennio Stacchetti (1989), "Psychological Games and Sequential Rationality," *Games and Economic Behavior*, 1(1): 60–79.
Gul, Faruk (1990), "Rational Strategic Behavior and the Notion of Equilibrium" Working Paper, GSB, Stanford University.
Kahneman, Daniel and Amos Tversky (1989), "Reference Theory of Choice and Exchange," Working Paper, Department of Psychology, University of California, Berkeley.
Kohlberg, Elon, and Jean-Francois Mertens (1986): "On the Strategic Stability of Equilibria," *Econometrica*, 54(5): 1003–38.
Luce, Duncan R. and Howard Raiffa (1957), *Games and Decisions*, New York: John Wiley.
Machina, Mark J. (1982), "'Expected Utility' Analysis without the Independence Axiom," *Econometrica*, 50: 277–323.
Myerson, Roger (1986), "Multistage Games with Communication," *Econometrica*, 54(2): 323–58.
Pearce, David (1984), "Rationalizable Strategic Behavior and the Problem of Perfection," *Econometrica*, 52(4): 1029–50.
(1987), "Renegotiation-Proof Equilibria: Collective Rationality and Intertemporal Cooperation," Cowles Foundation Discussion Paper No. 855, Yale University.
Rabin, Matthew (1990), "A Model of Pregame Communication," mimeo, Department of Economics, University of California, Berkeley.
Safra, Zvi, Uzi Segal, and Avia Spivak (1989), "Preference Reversal and Nonexpected Utility Behavior," *American Economic Review*, 80(4): 922–30.