
INTRODUCTION

This book gathers our joint work (through 2008) on the closely connected topics of repeated games and reputation effects, along with related papers on more general issues in game theory and dynamic games. Since this book examines, among other things, the way that long-run interactions facilitate cooperation, and our collaboration began in 1980, it is fitting that the papers appear in a jointly authored volume.

Understanding Dynamic Games: Limits, Continuity, and Robustness

Economists' interest in the applications of game theory exploded in the late 70's and early 80's. It was largely driven by the realization that game theory, in particular extensive-form games, provided the techniques for modeling such issues as commitment and credibility in entry deterrence, information and signaling in markets, and the dynamics of bargaining. The increased number and diversity of applications in turn generated many more abstract questions relating to dynamic games. These foundational and mathematical questions were the starting point of our collaboration. Our work on them in the early 1980's, examining issues such as the mathematical properties (mainly compactness and continuity) of various equilibrium concepts and their robustness to various sorts of small changes in the game, laid the groundwork for much of our later work.

Our first paper, "Subgame-Perfect Equilibrium of Finite and Infinite Horizon Games" (Chapter 1) was inspired by the contrast between the infinitely repeated prisoner's dilemma, which has a large set of subgame-perfect equilibria when players are patient, and Rubinstein's infinitely-repeated bargaining game, where the subgame-perfect equilibrium is unique. In contrast, when the horizon is finite, both games have a unique perfect equilibrium. This led us to examine if there is a difference between the finite-horizon games that explains the difference in the infinite horizon? We found that the two games are different in the sensitivity of their finite horizon equilibria to perturbations in the payoff functions in the far distant future.

After examining several motivating examples, we were led to examine the asymptotic behavior of the finite horizon epsilon-perfect equilibria. We suppose that payoffs are "continuous at infinity" — meaning that, essentially events in the distant future have only a small impact on payoffs. In this case the limits of the finite-horizon epsilon equilibria, with epsilon going to 0 as the horizon grows, correspond to the exact equilibria of the infinite horizon game. In particular, if the set of finite-horizon epsilon equilibria remains large there will be many exact equilibria in the limit, while there will be few if it shrinks. We also showed that under the same continuity assumption in finite-action games the set of strategies is compact in a topology that makes the set of equilibria closed; this implies that there is a best and a worst equilibrium for each player. Following a comment by Gerard Debreu that "it has to be more general than that", our paper "Limit Games and Limit

Equilibria” (Chapter 2) examined continuity of Nash equilibria in the limit of a family of games, for example games with a fine grid of actions converging to games with real-valued actions.

With the exception of three papers on incomplete-information bargaining, most of our joint work in the 1980’s studied other foundational issues. The papers “Open and Closed Loop Equilibria . . .”, “Finite Player Approximations to a Continuum of Players” and “When Are Nonanonymous Players Negligible” (Chapters 3, 4 and 6) studied when small players are strategically negligible in the sense of acting as if they have no effect on the future play of their opponents. These papers provide a strong contrast to the folk theorem for repeated games, as they imply that with sufficient players there is little possibility of sustaining cooperation through threats of future retaliation.

Intuitively, when players are small, we expect that the identity of a deviator cannot be easily detected, so subgame-perfection should rule out equilibria that rely on some overly broad forms of retaliation. For example, the recording industry cannot prevent anonymous music downloading by threatening to use nuclear weapons as soon as the first mp3 appears on the internet. The first paper examined the importance of feedback from future play in a two-stage simultaneous move game with many players. It showed that as the continuum limit is approached, the feedback makes little difference provided that the equilibrium of the second stage is regular in the sense of depending continuously on the first-period outcome.¹ Consequently, the subgame-perfect or “closed-loop” equilibria converge to limits that are “open loop” equilibria in which players commit to uncontingent time paths of play and thus do not try to influence one another’s future actions.²

The second paper provided a lower-hemicontinuity result, showing that under stringent technical conditions, each open loop equilibrium has an associated sequence of closed loop equilibria that converge to it as the number of players grows. Both of these papers assume that players participate anonymously, so that only the aggregate action is observed at the end of each period. The third paper asks when similar results might hold when players are identifiable, but their actions are observed with noise. We find that similar results hold. This is because intuitively, players can punish some of their opponents some of the time, but not all of them all of the time. That is, if punishments are limited to broad aggregates — for example income tax rates — it is possible to single out a few individuals to monitor — closely, and so provide them with strong incentives, but unless actions are perfectly observed, it is not possible that “most” players play this pivotal role.

“On the Robustness of Equilibrium Refinements” (Chapter 5) examined more closely the issues raised by refinements of Nash equilibrium such as subgame perfection. It showed that any refinement of Nash equilibrium that does more than

¹This result, and the need for the regularity condition, are extensions of ideas from the literature on when the limits of monopolistic competition are perfectly competitive (for example Roberts [1980]).

²The terminology comes from the literature on optimal control and differential games.

simply eliminate weakly dominated strategies is not robust to the presence of arbitrarily small amounts of incomplete information. Intuitively, equilibrium refinements in extensive-form games gain their force from restrictions on play off of the equilibrium path that are based on the players' original understanding of the game. However, a very small ex-ante chance that a player's payoffs are different than usual can become a very large ex-post probability if the player is observed to act in a way that was ex-ante unlikely.

Reputation Effects

At the end of the 80's, we began working on the related topics of reputation effects and repeated games that have been one of our main joint interests ever since. The literature on reputation-effects explores the idea that a player in a long-run interaction has some incentive to give up current payoffs to build a reputation for playing in a certain way. This intuition is clearest when there is a single reputation-building player facing opponents who are either myopic (short run players) or strategically myopic (small players who are strategically negligible); here we might expect that a long run player will have the greatest ability to exploit reputation effects.

The first papers to model the reputation idea were Kreps and Wilson [1982], Milgrom and Roberts [1982], and the "Gang of Four" paper Kreps, Milgrom, Roberts and Wilson [1982]. The first two of these papers examined the sequential equilibria of specific finite-horizon games with specific two-point distributions of types: the long-run player was either "rational," with payoff function the sum of the per-period payoffs of the stage game, or a commitment or "crazy" type that always plays a given fixed action. The long-run player's reputation then corresponds to his opponents' subjective probability that he is the commitment type.^{3,4} The proof technique in these papers uses backwards induction to obtain a characterization of the sequential equilibrium strategies. This approach, similar to the sequential equilibrium concept, is very sensitive to the specific types of the long run player that have positive prior probability.

In such reputation-effects models, it is impossible to build a reputation for playing like a type whose prior probability is 0, so the exogenously-specified set of types corresponds to the set of possible reputations. Moreover, the prior weights on the

³The commitment can either be hard-wired (the types have no choice but to play this way) or a consequence of preferences that induce the commitment type to conform to the specified strategy along the equilibrium path; the "Robustness" paper constructs such preferences for the case of Nash equilibrium. In general preferences that can be guaranteed a priori to generate a particular play must be very carefully constructed and can seem somewhat artificial, as they must make the specified play conditionally dominant. However, as first noted by Schmidt [1992], in some cases a more restrictive equilibrium concept such as Markov-perfection implies that types with more standard preferences will act like commitment types: In Schmidt's paper this is the case for the buyer with the lowest possible valuation for a good.

⁴The "Gang of Four" paper considered reputation effects in the finitely-repeated prisoner's dilemma. Here the reputation-building player faced a long-run opponent, and the "commitment type" played the strategy "Tit for Tat."

various types are regarded as small; the point of the literature is that even a very small prior weight can have a large impact over a sufficiently long horizon. This motivated us to look for a way to characterize reputation effects for general type distributions, which in turn led to a search for a proof technique that does not require explicit characterization of the equilibrium strategies. Our paper “Reputation and Equilibrium Selection in Games with a Patient Player” (Chapter 7) provides a lower bound on equilibrium payoffs to the long run player that applies for general stage games and general type distributions, and to both finite and infinite-horizon games, and moreover holds for all of the Nash equilibria. This shows that the power of reputation effects does not rely on backwards induction or other refinements of Nash equilibrium, so that the conclusion is robust in the sense of our JET 1988 paper.

The idea of the lower bound is simple, and corresponds to the intuitive notion of reputation-building: If the long-run player’s action is observed each period, and if the short-run players assign positive probability to the long-run player being a type that always plays (say) “Up,” then if the long-run player chooses to always play “Up” the short-run players must eventually come to expect him to play “Up” in the future. More specifically, there is an integer k , independent of the equilibrium selection or the discount factor, such that there are at most k “surprises,” where “surprise” means the periods in which the short-run players do not play a best response to “Up” and yet “Up” is played. The lower bound is then the maximum payoff that the long-run player can guarantee by playing any of the actions for which there is positive prior probability of the corresponding commitment type. Thus the bound is vacuous if no commitment types have positive prior probability, and becomes more powerful as the set of positive-probability commitments grows larger, but the bound does not rely on there being positive probability of any specific commitment.

The paper also pointed out that in extensive-form games, reputation effects can be much less powerful, essentially because the long-run player may not be able to build a reputation for play at information sets that are never reached, as for example if the short-run players move first and can opt out of dealing with him. In this case we proposed a weaker lower bound, the “generalized Stackelberg payoff,” and showed by example that in some games this lower bound is attained, so that reputation effects are less powerful than if actions are observable.⁵ We returned to this idea in a recent paper with Jeff Ely that is discussed below.

The key step in our first reputation effects paper — the fact that the short-run players learn to expect the commitment action after at most k surprises — relies heavily on the assumption that the commitment action is observed without

⁵The underlying notion on which the generalized Stackelberg payoff is based is that of an “epsilon-confirmed response”, where the short run players play a best response to some long-run player action that is consistent with his observed signals. This is related to our learning-theoretic notion of self-confirming equilibrium. As part of our research agenda on learning in games, we do not cover self-confirming equilibrium in this book.

noise. This rules out both moral hazard (exogenous noise in the observations) and commitment to a mixed action. This latter case is important because the best possible commitment in finite stage games is typically a mixed and not pure action.

“Maintaining a Reputation when Strategies are Imperfectly Observed” (Chapter 8) extends our earlier result and intuition, and thus the equilibrium-selection result, to the case where the observed signal is stochastic. This may be either because the long run player is using a mixed strategy or because his action is observed with some exogenous noise, as in a model of moral hazard. In this case the beliefs of the short run players are stochastic, so we provide a uniform bound on the rate that these stochastic beliefs evolve if the long-run player chooses to always play like one of the positive-probability commitment types.⁶ As Sorin [1999] points out, this amounts to a uniform version of the “merging of beliefs” theorem of Blackwell and Dubins [1962] (which is in turn an extension of the martingale convergence theorem.) Using the extension to mixed commitment actions, the paper also gives an upper bound on the long-run player’s Nash equilibrium payoff. It argues that if “commitment types have full support,” so that all possible reputations are allowed ex-ante, then in generic simultaneous move games the two bounds will coincide in the limit of discount factors going to one. In this case, reputation effects imply that the long-run player receives his Stackelberg equilibrium payoff.

When the stage game has a non-trivial extensive form, this last result does not apply, and there can be equilibria that give even a patient long run player less than his Stackelberg payoff. More strikingly, in “Bad Reputation,” Ely and Valimaki [2003] study a reputation-effects model where the unique equilibrium attains the lower bound of our RES paper, so that the long-run player does worse with this particular form of reputation building than if there were no reputation building at all.

The Ely and Valimaki example more or less turns our usual view of reputation on its head. However, their argument was for a specific two-point distribution on types and a particular game. The boundaries of this counterexample are probed in our paper “When is Reputation Bad?” (with Jeff Ely, Chapter 10). The Ely-Valimaki result does not hold for general type distributions, but we show that there is a sense in which it applies to “most” distributions that assign probability near one to the rational type. The class of games in which the results holds is expanded to a broad class of “exit games.” In these games the short-run player has the chance not to participate, in which case it is natural to assume that the play of the long-run player is not seen.

⁶Although the lower bound is computed by assuming that the long-run player chooses to always play a fixed action, it is important to note that this in general is not the equilibrium strategy when there are commitment types that can play a mixed action or if the action is imperfectly observed. Indeed, Cripps, Mailath, and Samuelson [2004] give conditions under which the long-run player’s type is eventually revealed along every path, so that reputation is asymptotically “impermanent”, even though it decays slowly enough that the long-run player does as well as if he were thought to be the commitment type.

Much of the reputation effects literature focuses on the case of a single reputation builder; the asymmetry built into this model suggests a possible outcome (Stackelberg for the reputation builder) and leads to strong results. It is less clear what the results might be without some such asymmetry, and indeed the nature of equilibrium with two equally patient reputation builders seems to depend heavily on the fine details of the model and the type distributions, as shown by Fudenberg and Maskin's [1986] incomplete-information folk theorem. There are however some results about reputation building with various sorts of asymmetries, for example several medium-to-long run players and a single reputation builder, as in "Maintaining a Reputation Against a Long-Lived Opponent" (with Marco Celentani and Wolfgang Pesendorfer, Chapter 9).⁷ Although the stage game here has simultaneous moves, the ideal reputation is typically for a strategy in the repeated game, so that off-path beliefs become relevant, as they do when the stage game has sequential moves. The paper shows that if the player trying to build a reputation has an action observed with noise, the importance of off-path beliefs is mitigated.

Repeated Games

One of the central issues in economics is identifying institutions and settings that lead to more, or less efficient outcomes. Simple static examples suggest that this is not easy; think for example of the free-rider problem or of the difficulty in providing insurance for unobserved risks. Of course in practice the incentive to free-ride and the temptation to misreport one's current state can be offset by the possibility of future considerations. Even such phenomena as the "invisible hand" of competitive markets rests ultimately on this foundation, as private property may be difficult to protect without the possibility of future retaliation for theft and violence. In game theory the simplest setting in which we can study future considerations is that of the repeated game, where the future incentives arise from equilibrium play without ex-ante commitments. Within repeated games, the simplest and most revealing case is that in which players are patient, so that future considerations matter most.⁸

There are many equilibria in repeated games when the players are patient, as shown by the folk theorem: every payoff outcome that is individually rational and socially feasible can be sustained as an equilibrium. Notice that this is a sharp characterization: outcomes that fail to be either individually rational or socially feasible cannot possibly be Nash equilibria. The folk theorem has both an upside and a downside. The upside is that a great many efficient outcomes are sustainable. The downside is that some very bad equilibria are possible as well, and that there

⁷Fudenberg and Kreps [1987] consider the asymmetric case of one large firm simultaneously facing several long-run opponents.

⁸While the various theorems we discuss are stated only for the limits of discount factors arbitrarily close to one, their qualitative conclusions apply more generally. That is, for discount factors on the order of 0.99 or even 0.95, the best equilibria in games where the folk theorem holds are typically more efficient than the best equilibria in games where the equilibrium payoffs are bounded away from efficiency.

is little predictive power in the theory, but there is a widespread feeling that the more efficient equilibria are relatively more common.⁹

The simplest version of the folk theorem is immediate: Suppose that our equilibrium notion is that of Nash equilibrium and players are infinitely patient, meaning that they care only about the time average of payoffs. Then any payoff that is individually rational (in the sense of giving each player at least their minmax payoff) and socially feasible can be sustained by the threat to minmax any player who deviates. This theorem is problematic in three respects: it assumes infinite patience rather than discounting, and so simply ignores any short term benefits of deviating; it relies on a threat that may be not be credible, namely to minmax a deviating player even if that is very costly, and it assumes that players can perfectly observe that any deviations from equilibrium play.

An important challenge in subsequent work on game theory has been to relax these assumptions to make the theorem more relevant. Aumann and Shapley [1976] extended the folk theorem to subgame-perfect equilibrium, while maintaining the assumption that player care only about the time average of their payoffs, and Rubinstein [1979] extended it to the case where players rank payoff streams by their time average but break ties using the “overtaking criterion.”¹⁰ Notice though, that the overtaking criterion is substantially less appealing than discounting, as punishments in the far distant future have the same impact on payoffs as payoffs that are immediate.

Friedman [1971] proved the first general result about the subgame perfect equilibria of repeated games with discounting. Friedman showed that, given any outcome that strictly Pareto dominates some static Nash equilibrium, there is a critical discount factor such that perpetual play of that outcome can be sustained by a subgame-perfect equilibrium when the discount factor exceeds the critical level. This, however, falls short of the full folk theorem for many games. When individually rational payoffs are less than any static Nash equilibrium, not only does this open the door to equilibria worse than any considered by Friedman, but the range of payoffs on the Pareto frontier that can be sustained may be significantly broadened.

Friedman’s result left open the question of whether the full folk theorem applies in the limit of discounting case. “The Folk Theorem in Repeated Games with Discounting or with Incomplete Information” (Fudenberg and Maskin [1986]) shows

⁹This feeling is supported by both informal and formal experiments, for example, Axelrod [1984] and DalBo [2005]. It has generated many theoretical explanations. Of these the most prominent line is the evolutionary approach, starting with Axelrod and Hamilton [1981], followed by Fudenberg and Maskin [1990], Binmore and Samuelson [1992], and Nowak and Sigmund [1992]. There have also been attempts to explain cooperation using models of reputation, notably Kreps et al [1992], but in contrast to the case of one-long run player the equilibrium selection here depends delicately on the support of the possible commitment types (Fudenberg and Maskin [1986].)

¹⁰The overtaking criteria ranks two infinite sequences of payoffs by their time averages when those averages exist; it is more sensitive in that $(1, 0, 0, \dots)$ is preferred to $(0, 0, 0, \dots)$. More formally, sequence $\{x_t\}$ is preferred to $\{y_t\}$ by the overtaking criterion if there is a T such that for all $T' > T$ we have $\sum_{t=1}^{T'} x_t > \sum_{t=1}^{T'} y_t$.

that it does with the addition of a modest (“generic”) full-dimension assumption on the set of feasible payoffs of the game.¹¹ With Eric Maskin’s kind permission we have reprinted that paper (Chapter 11). The main question left open by this result is understanding what happens when monitoring is not perfect, so that players receive only noisy signals of one another’s actions. This topic has been a focus of our collaboration and of the repeated games literature as a whole.

Obviously, if players receive no information at all about the play of their opponents, then repeating the game has no effect, so the key question with imperfectly observed actions is how the nature of the information structure influences the equilibrium set, a question that does not arise with observed actions. The early results on the role of the discount factor in these games showed that the folk theorem need not hold. For example, it was known from the work of Green and Porter (1984) and Abreu, Pearce and Stachetti (1986) that the folk theorem need not apply when players are restricted to using strongly symmetric strategies.¹² In particular the best equilibrium payoff is bounded away from efficiency uniformly in the discount factor. Subsequent work obtained similar inefficiency results in specific games without imposing a symmetry condition, for example, the partnership game by Radner, Myerson and Maskin [1986].

It turns out, however, that these counter-examples are all very special: In some cases, such as the Green and Porter oligopoly model, the failure of the folk theorem is due to the external assumption of strong symmetry, which requires that all players use the same actions in every period, regardless of the history, and the folk theorem applies when the symmetry assumption is relaxed. In other cases, such as Radner, Myerson, and Maskin, the folk theorem fails because of the assumption that there are only two possible observations in each period, “success” or “failure,” and the folk theorem applies when the set of possible observations is sufficiently large. The common feature of these two sorts of counterexamples to the folk theorem is that it is not possible to construct equilibrium strategies that reward one player while punishing another. Our paper with Eric Maskin, “The Folk Theorem in Repeated Games with Imperfect Public Information” (Chapter 12) shows that the ability to make such transfers is the key to the folk theorem.

This paper has a long and somewhat unusual history, as it grew out of the merger of two earlier papers. The first, “Discounted Repeated Games with One-Sided Moral Hazard,” by Fudenberg and Maskin, (unpublished 1986) observed that the key to the dynamic-programming methodology of Abreu, Pearce and Stachetti [1986, 1990] is not the restriction to pure strategies but rather the restriction to “perfect public equilibria” or “PPE.” It used mixed strategies to construct “equilibrium polygons” that approximate the set of feasible individually rational payoffs when the discount factors are close to one. The point of the construction was that the edges of the polygon can be made smaller as the players become more patient, which allows the

¹¹It also requires strengthening “individually rational” to “strictly individual rational.”

¹²Mailath and Samuelson [2007] extend this observation to general games.

polygon to more closely approximate the set of feasible individually rational payoffs. This was the first general positive result about the folk theorem with discounting and with imperfect information.¹³ However, the polygonal construction proved difficult to extend to general games.

The second paper, “Folk Theorem with Unobservable Actions,” [unpublished 1988], was written while the first paper was under review at *Econometrica*; it arose from conversations trying to understand how the intuition for the polygon construction could be made more general. This paper replaced the complicated polygon constructions by smooth sets; intuitively, these smooth sets can be viewed as the limits of equilibrium polygons as the edges become vanishingly small. The paper then introduced the concept of local generation. The earlier concept of “self-generation” is due to Abreu, Pearce and Stachetti [1986, 1990]: roughly it means that the payoffs in the set are equilibrium payoffs in which the continuation payoffs following each possible current observation are in the same set. A set is locally generated if every element of the set has an open neighborhood that is self-generating for some discount factor. Because the discount factor can vary with the point in question, local generation can be verified pointwise, yet it also gives a sufficient condition for payoffs to be supported by PPE as the discount factor goes to one. Finally, we used the fact that smooth sets are locally linear to show that sets that are “enforceable on tangent hyperplanes” are locally generated.¹⁴

When the reports came back on the first paper, we decided to merge the two of them, and then spent several years working with Maskin to generalize the information conditions that imply payoffs can be enforced on tangent hyperplanes. One simple sufficient condition is “pairwise full rank,” which states briefly that (a) the various actions of a given player can be statistically distinguished from one another (this is the “individual full rank” condition) and (b) the actions of any one player can be statistically distinguished from those of any other (“pairwise identifiability.”) Note that these information conditions are purely qualitative, and do not require bounds on the informativeness of the signals. This is because the theorem applies to the limit as the discount factor goes to 1. The informativeness of the signals does matter for any fixed discount factor, as shown by Kandori [1992b].¹⁵

We then noticed that the same “smooth set/tangent hyperplane” methods that we had used to prove the folk theorem could be extended to characterize the limit of

¹³Radner [1981,85] had previously proved a folk theorem for repeated principal-agent games with time-average and then discounted payoffs, respectively; the Fudenberg and Maskin paper differed in considering general payoff functions and allowing more than two possible public signals.

¹⁴A set is enforceable on tangent hyperplanes if at every point v in the set, there is an action profile a that (1) yields payoffs that are separated from the set by the tangent hyperplane, and (2) can be enforced a with continuation payoffs that lie in the tangent hyperplane. Working independently, Matsushima [1989] also hit upon the method of using a smooth boundary and calculus to do computations.

¹⁵Kandori shows that if the set of PPE payoffs for a given discount factor is convex, it is not made larger (that is, it is weakly decreased) when the information is “garbled” in the sense of Blackwood.

the PPE payoffs in games with long-run and short-run players, where the standard folk theorem does not apply. This led to “Efficiency and Observability in Games with Long-Run and Short-Run Players” (Chapter 13). One motivation for the paper was to study games with long-run and short-run players; we discuss this in more detail below. The paper also made a methodological contribution by extending enforcement on hyperplanes to enforcement on half-spaces. This extension applies whether or not there are short-run players, and does not require any conditions on the nature of the public monitoring technology, so it allows a characterization of the limit set of PPE payoffs in cases when the folk theorem does not apply due to a failure of the identification conditions, and it makes it easier to look for alternative sufficient conditions for the folk theorem to apply.

The half-space idea is that the limit of the set of PPE payoffs is the solution to a family of linear programming problems, where each problem corresponds to finding the highest self-generating half-space in a given direction.¹⁶ This technique has since been extended to a number of related contexts, first by Kandori and Matsushima [1988] in their study of repeated games with private monitoring and communication, and most recently by Sannikov and Skrypacz [2007] in their study of discrete-time repeated games with short time periods.

This topic continues to be of interest to us as well. Recently, working with Satoru Takahashi in “Perfect Public Equilibrium when Players are Patient,” (Chapter 16) we expanded the results to give an iterative algorithm for computing the limit equilibrium payoffs when the set determined by the original algorithm does not satisfy the full-dimension condition, either because the set of feasible payoffs is itself lower-dimensional or because of restrictions such as symmetry; we also show how to use this approach to give sufficient conditions for the existence of exactly efficient PPE for discount factors close to but less than one.

It is interesting to note that there is a methodological difference between these papers and the previous literature on the folk theorem with imperfect monitoring that parallels the difference between our work on reputation effects and the earlier work on that topic. In both cases, the earlier work is constructive, while our papers have sacrificed that level of detail to make it easier to characterize the equilibrium payoffs.¹⁷

Private Information

While the case of publicly observed signals is important, many informational imperfections involve private information. For example, as Fudenberg and Tirole [1991] point out, in the Stigler model of oligopoly with price search by consumers, duopolists observe only private information about their own sales, and there are no

¹⁶This turns out to be independent of the discount factor, provided it is positive.

¹⁷That said, one could readily construct equilibrium strategies using the FLM apparatus. The earlier work on repeated games with imperfect monitoring concentrated on the construction particular equilibria, whereas FLM constructed a large set of equilibria at one time.

informative public signals. In addition, situations where players are matched against different opponents and observe only the outcomes of their own matches, as in the literature on “community enforcement,” can also be viewed as a class of repeated games with private monitoring.¹⁸

The study of private monitoring with discounting has proved difficult, even in the case where a fixed set of players faces one another every period, and it is still not known whether a full folk theorem applies. In the undiscounted case, there are quite satisfactory results, especially those of Lehrer [1988, 1992]. A very general set of results for the undiscounted case is our own paper “An Approximate Folk Theorem with Imperfect Private Information” (Chapter 14). In addition to proving a full folk theorem in the undiscounted case with private information and relatively mild informational assumptions, the paper shows that the result holds for the discounted subgame-perfect case, provided we relax the equilibrium concept to approximate equilibrium. The underlying ideas are standard ones in the undiscounted literature: to put periods into blocks in order to aggregate information, followed by communication rounds in which information is shared among players, followed by either continuation or punishment blocks. The problem with using this in the discounted case is that as the end of the block is approached, depending on events that occurred early, players may be able to make small gains by deviating. More recent constructions by Horner and Olszewski [2008] have been able to extend the block approach to the case of exact discounted equilibria, albeit proving less general results.

While the general private information discounted case is difficult and proved elusive for many years, it has recently become a major topic of research as shown by Kandori’s [2002] survey and the large number of significant papers since then. Inspired by his talk at the associated Cowles conference, we realized that we could extend our results on public monitoring to games with private monitoring and communication, provided that the departure from public monitoring is small, as recorded in the small note “The Nash-Threats Folk Theorem with Communication and Approximate Common Knowledge in Two Player Games” (Chapter 15).

Long Run and Short Run Players

Although a great deal of research on repeated games has focused on the case where players are equally patient, the asymmetric case where there is one patient player and one or more impatient players or strategically negligible players has a great deal of economic significance. This is true in games such as the classical chain-store paradox, where an incumbent faces a series of different rivals, and has applications to firms who face a series of different customers. It is also important in political economy, where the long-run player is the government and the short-run players are not impatient per se, but behave as if they are impatient because they are too small to have an impact on future policy. For example, nobody would believe that

¹⁸See Kandori [1992] and Ellison [1993].

their personal decision whether or not to pay their taxes would have an impact on future income tax rates.

In long-run versus short-run games, the short-run players cannot be coerced through threats of future action, so the relevant benchmark is the Stackelberg equilibrium attained through commitment. Does repeated play allow the long-run player to do as well as this? This question also arises in the reputation-effects literature, which is concerned primarily with whether reputation implies that the long-run player does get the Stackelberg payoff, that is, whether this payoff is uniquely selected. However, characterizing the set of possibilities in the non-reputational case is also of substantial importance.

Fudenberg, Kreps and Maskin [1990] provided the first general analysis of games with long-run and short-run players. They characterized the best (and worst) subgame-perfect equilibrium payoff to the long-run player in the repeated game. In general, this best subgame-perfect equilibrium payoff is somewhere between the pure strategy and mixed strategy Stackelberg payoffs. They then showed that this payoff is obtainable if the long-run player is patient enough. The original motivation for “Efficiency and Observability with Long-Run and Short-Run Players” (Chapter 13) was to generalize this result to allow for imperfect public information and many long-run players. We also introduced the notion of moral hazard mixing games, with one long-run player facing one or more short-run opponents, and showed that in the presence of short-run players, imperfectly observed actions lead to a smaller set of equilibrium payoffs, even in the limit of discount factors tending to 1, and that more accurate signals allowing a higher limit equilibrium payoffs for the long-run player.¹⁹ This is a sharp contrast to games where all players are long run, where only the qualitative properties of the information matter for the limit equilibrium payoffs.

For repeated games in general, continuous time limits have become of great interest following Sannikov’s [2007] pioneering work on using continuous time methods to compute the set of PPE payoffs for games with imperfect public monitoring and all long-run players. The continuous time setup together with information that follows a diffusion process makes it possible to compute the PPE payoffs for given interest rates, rather than simply in the limit as the interest rate goes to zero. Applying these methods to games with a long-run versus short-run player, Faingold and Sannikov [2005] shows that only trivial equilibria are possible.²⁰ This led us to wonder what it was about diffusion processes that gave rise to such a sharp and seemingly anomalous result. We investigated this in “Continuous Time Limits of Repeated Games with Imperfect Public Monitoring” (Chapter 17), where we showed that the conclusion depends quite crucially on the assumption that the information

¹⁹Mailath and Samuelson [2007] extend this observation to moral hazard mixing with multiple long run players.

²⁰By way of contrast Faingold [2005] shows that with reputation effects continuous time makes little difference.

corresponds to a diffusion whose volatility is independent of the actions played. For example with Poisson information,²¹ non-degenerate limits can be obtained, and even in the diffusion case non-degenerate limits are possible if the volatility of the signal depends on the actions.

References

- Abreu, D., P. Milgrom, and D. Pearce (1991). Information and timing in repeated partnerships. *Econometrica*, 59, 1713–33.
- Abreu, D., D. Pearce, and E. Stachetti (1986). Optimal cartel equilibria with imperfect monitoring. *Journal of Economic Theory*, 39, 251–269.
- Abreu, D., D. Pearce, and E. Stachetti (1990). Towards a theory of discounted repeated games with imperfect monitoring. *Econometrica*, 58, 1041–1064.
- Aumann, R., and L. Shapley (1976). Long term competition — a game theoretic analysis,” mimeo.
- Axelrod, R. (1984). The evolution of cooperation, Basic Books, New York.
- Axelrod, R., and W. Hamilton (1981). The evolution of cooperation. *Science*, 211, 1390–1396.
- Binmore, K., and L. Samuelson (1992). Evolutionary stability in repeated games played by finite automata. *Journal of Economic Theory*, 57, 278–305.
- Blackwell, D. (1951). Comparison of experiments. In *Proceedings of the Second Berkeley Symposium on Mathematical Sciences and Probability*, 93–102. University of California Press.
- Celentani, M., D. Fudenberg, D. Levine, and W. Pesendorfer (1996). Maintaining a reputation against a long-lived opponent. *Econometrica*, 64, 691–704.
- Compte, O. (1998). Communication in repeated games with imperfect private monitoring. *Econometrica*, 66, 597–626.
- Dal Bo, P. (2005). Cooperation in the shadow of the future: Experimental evidence from infinitely repeated games. *American Economic Review*, 1591–1604.
- Ellison, G. (1993). Cooperation in the prisoner’s dilemma with anonymous random matching. *Review of Economic Studies*, 61, 567–588.
- Ely, J., D. Fudenberg, and D. Levine (2008). When is reputation bad. *Games and Economic Behavior*, 63, 498–526.
- Ely, J., J. Horner, and W. Olszewski (2005). Belief-free equilibria in repeated games. *Econometrica*, 73, 377–415.
- Ely, J., and J. Välimäki (2003). Bad reputation. *Quarterly Journal of Economics*, 118, 785–814.
- Faingold, E. (2005). Building a reputation under frequent decisions, mimeo.
- Faingold, E., and Y. Sannikov (2005). Equilibrium degeneracy and reputation effects, mimeo.
- Friedman, H. (1971). A no-cooperative equilibrium for supergames. *Review of Economic Studies*, 38, 1–12.
- Fudenberg, D., and D. Kreps (1987). Reputation and simultaneous opponents. *Review of Economic Studies*, 54, 541–568.
- Fudenberg, D., D. Kreps, and D. K. Levine (1988). On the robustness of equilibrium refinements. *Journal of Economic Theory*, 44, 354–380.
- Fudenberg, D., D. Kreps, and E. Maskin (1990). Repeated games with long-run and short-run players. *Review of Economic Studies*, 57, 555–573.

²¹Poisson information in repeated games was first studied in Abreu, Milgrom and Pearce [1991].

- Fudenberg, D., and D. K. Levine (1983). Subgame-perfect equilibria of finite and infinite horizon games. *Journal of Economic Theory*, 31, 251–268.
- Fudenberg, D., and D. K. Levine (1986). Limit games and limit equilibria. *Journal of Economic Theory*, 38, 261–279.
- Fudenberg, D., and D. K. Levine (1988). Open-loop and Closed-loop equilibria of dynamic games with many players. *Journal of Economic Theory*, 44, 1–18.
- Fudenberg, D., and D. K. Levine (1988). Finite player approximations to a continuum of players, mimeo.
- Fudenberg, D., and D. K. Levine (1989). Reputation and equilibrium selection in games with a single long-run player. *Econometrica*, 57, 759–778.
- Fudenberg, D., and D. K. Levine (1991). An approximate folk theorem for games with imperfect private information. *Journal of Economic Theory*, 54, 26–47.
- Fudenberg, D., and D. K. Levine (1991). Maintaining a reputation when strategies are imperfectly observed. *Review of Economic Studies*, 59, 561–581.
- Fudenberg, D., and D. K. Levine (2007). The Nash-Threats folk theorem with communication and approximate common knowledge in two-player games. *Journal of Economic Theory*, 132, 461–473.
- Fudenberg, D., and D. K. Levine (2007). Continuous time limits of repeated games with imperfect public monitoring. *Review of Economic Dynamics*, 10, 173–192.
- Fudenberg, D., D. K. Levine, and E. Maskin (1994). The folk theorem in repeated games with imperfect public information. *Econometrica*, 62, 997–1039.
- Fudenberg, D., D. K. Levine, and W. Pesendorfer (1998). When are non-anonymous players negligible? *Journal of Economic Theory*, 79, 46–71.
- Fudenberg, D., D. K. Levine, and S. Takahashi (2007). Perfect public equilibrium when players are patient. *Games and Economic Behavior*, 61, 27–49.
- Green, E., and R. Porter (1984). Non-cooperative collusion under imperfect price information. *Econometrica*, 52, 87–100.
- Hörner, J., and W. Olszewski (2005). The folk theorem for games with private almost-perfect monitoring, mimeo.
- Hörner, J., and W. Olszewski. How robust is the folk theorem with imperfect public monitoring? mimeo.
- Kandori, M. (1992a). Social norms and community enforcement. *Review of Economic Studies*, 59, 63–80.
- Kandori, M. (1992b). The use of information in repeated games with imperfect monitoring. *Review of Economic Studies*, 59, 581–594.
- Kandori, M. (2002). Introduction to repeated games with private monitoring. *Journal of Economic Theory*, 102[1], 1–15.
- Kandori, M., and H. Matsushima (1998). Private observation, Communication, and Collusion. *Econometrica*, 66, 627–652.
- Kreps, D., and R. Wilson (1982). Reputation and imperfect information. *Journal of Economic Theory*, 27, 253–279.
- Kreps, D., P. Milgrom, J. Roberts, and R. Wilson (1982). Rational cooperation in the finitely repeated prisoner's dilemma. *Journal of Economic Theory*, 27, 245–252, 486–502.
- Lehrer, E. (1988). Repeated games with stationary bounded recall strategies. *Journal of Economic Theory*.
- Lehrer, E. (1992). Two-player repeated games with nonobservable actions and observable payoffs. *Mathematics of Operations Research*.
- Mailath, G., M. Cripps, and L. Samuelson (2004). Imperfect monitoring and impermanent reputations. *Econometrica*, 72, 407–432.

-
- Matsushima, H. (1989). Efficiency in repeated games with imperfect monitoring. *Journal of Economic Theory*, 48, 428–442.
- Milgrom, P., and J. Roberts (1982). Predation, reputation, and entry deterrence. *Journal of Economic Theory*, 27, 280–312.
- Nowak, M., and K. Sigmund (1992). Tit for tat in heterogeneous populations. *Nature*, 355, 250–253.
- Radner, R. (1981). Monitoring cooperative agreements in a repeated principal-agent relationship. *Econometrica*, 49, 1127–1148.
- Radner, R. (1985). Repeated principal-agent games with discounting. *Econometrica*, 53, 1173–1198.
- Radner, R., R. Myerson, and E. Maskin (1986). An example of a repeated partnership games with discounting and with uniformly inefficient equilibria. *Review of Economic Studies*, 53, 59–70.
- Roberts, K. (1980). The limit points of monopolistic competition. *Journal of Economic Theory*, 22, 246–278.
- Rubinstein, A. (1979). Equilibrium in supergames with the overtaking criterion. *Journal of Economic Theory*, 21, 1–9.
- Sannikov, Y. (2007). Games with imperfectly observed actions in continuous time. *Econometrica*, 75, 1285–1329.
- Schmidt, K. (1993). Reputation and equilibrium selection in repeated games of conflicting interests. *Econometrica*, 61, 325–351.
- Sekiguchi, T. (1997). Efficiency in repeated prisoner's dilemma with private monitoring. *Journal of Economic Theory*, 76(2), 345–361.
- Sorin, S. (1999). Merging, reputation, and repeated games with incomplete information. *Games and Economic Behavior*, 29, 274–308.
- Stigler, G. (1964). A theory of oligopoly. *Journal of Political Economy*, 72, 44–61.