

CIRJE-F-114

Introduction to Repeated Games with Private Monitoring

Michihiro Kandori
University of Tokyo

April 2001

Discussion Papers are a series of manuscripts in their draft form. They are not intended for circulation or distribution except as indicated by the author. For that reason Discussion Papers may not be reproduced or distributed without the written consent of the author.

Introduction to Repeated Games with Private Monitoring*

KANDORI, Michihiro

Faculty of Economics, University of Tokyo,
7-3-1 Hongo, Bunkyo-Ku, Tokyo 113-0033, Japan
kandori@e.u-tokyo.ac.jp

April 26, 2001

Abstract

We present a brief overview of recent developments on discounted repeated games with (imperfect) private monitoring. The literature explores the possibility of cooperation in a long-term relationship, where each agent receives imperfect *private* information about the opponents' actions. Although this class of games admits a wide range of applications such as collusion under secret price-cutting, exchange of goods with uncertain quality, and observation errors, it has fairly complex mathematical structure due to the lack of common information shared by players. This is in sharp contrast to the well-explored case of repeated games under public information (with the celebrated Folk Theorems), and until recently little had been known about the private monitoring case. However, rapid developments in the past few years have revealed the possibility of cooperation under private monitoring for some class of games. *Journal of Economic Literature* Classification Numbers: C72, C73, D43, D82, L13, L41

*To appear in the symposium issue of *Journal of Economic Theory* on repeated games with private monitoring. This is an outgrowth of my presentation at International Conference on Game Theory at Stony Brook 1996 and Cowles Foundation Conference on Repeated Games with Private Monitoring 2000. I thank George Mailath for helpful comments and discussion.

1 A Simple, Hard Open Question

The theory of repeated games provides a formal framework to explore the possibility of cooperation in long-term relationships, such as collusion between firms, cooperation among workers, and international policy coordination. The extensive literature has by now established that efficiency can be achieved under fairly general conditions (the Folk Theorems). However, virtually all those existing results heavily rely on one crucial assumption, which does exclude a number of important applications. The key assumption in the existing literature is that players share *common information* about each other's actions. The present article provides an overview of rapidly growing recent literature, which relaxes this restrictive assumption.

In the perfect observability case (see, for example, the Folk Theorem by Fudenberg and Maskin [19]), players commonly observe actual actions, and in the imperfect monitoring case explored by the majority of the existing literature (see Green and Porter [20], Abreu, Pearce and Stacchetti [3] and Fudenberg, Levine and Maskin [18]), players observe a common signal in each period, which is an imperfect indicator of the actions taken in the current period. For example, in the model of collusion proposed by Green and Porter, each firm secretly chooses its own output level and all the firms commonly observe the market price. The market price reflects the firms' actions (quantities supplied), but it is also subject to demand shocks. Hence the market price provides imperfect but *commonly shared* information about the actual actions taken by the firms.

In contrast, consider the situation where firms offer secret price cutting to their customers. The firms are not able to observe others' secret offers, but they can obtain some information from their own sales. When a firm's sales slump, it might be caused by the secret price cutting of its rival firms. It is not, however, a perfect indicator, as the slumped sales might also be caused by low demand. Note the similarity and difference between Green and Porter's model and the secret price cutting story. Both assume that the players' actions are imperfectly monitored, but in the former the players *publicly* observe the same signal (the market price) while in the latter each player receives *private* information (one's own sales). As we will see in detail in Section 2, this seemingly minor change makes a substantial difference in terms of the tractability of the model. In contrast to our nearly perfect understanding of the (perfect or imperfect) *public monitoring* case, our knowledge about repeated games with imperfect *private monitoring* is quite limited. However, in the past few years, this has become an active field of research, and a number of new findings have been obtained. A majority

of those were presented in the Cowles Foundation Conference on Repeated Games with Private Monitoring in April, 2000 and are put together in the current special issue of *Journal of Economic Theory*.

1.1 The Model and Applications

Let us now formally define a discounted repeated game with (imperfect) private monitoring. Players $i = 1, \dots, N$ repeatedly play the same stage game over infinite time horizon, $t = 0, 1, \dots$. In each stage, player i chooses action $a_i \in A_i$ and then observes a signal $\omega_i \in \Omega_i$. The action a_i and signal ω_i are player i 's private information. The probability of private signals $\omega = (\omega_1, \dots, \omega_N)$ depends on the current actions profile $a = (a_1, \dots, a_N)$ and is denoted by $p(\omega|a)$. Player i 's expected payoff in the stage game is given by $g_i(a) = \int_{\omega} u_i(a_i, \omega_i) p(\omega|a)$, where u_i represents player i 's realized payoff¹. Each player i maximizes discounted payoff $\sum_{t=0}^{\infty} g_i(a(t)) \delta^t$, where $a(t)$ is the action profile at t and $\delta \in (0, 1)$ is the discount factor. Note that the existing models of repeated games with public monitoring can be regarded as (degenerate) special cases of this formulation. The case with $\omega_i = a$ for all i corresponds to the *perfect monitoring* case, while the case with $\omega_1 = \dots = \omega_N$ represents the *imperfect public monitoring* case. As one can see, the model of repeated games with private monitoring is fairly simple, yet we have only limited knowledge about what the players can achieve in such a game. This is probably one of the best known long-standing open questions in economic theory.

The private monitoring case includes a number of important economic applications. We have already discussed collusion under secret price cutting. Another prominent example is exchange of goods with uncertain quality. Two players $i = 1, 2$ exchange goods, and the quality of the goods is randomly determined by the unobservable effort level of the supplier. Here, a_i and ω_i correspond respectively to player i 's effort and the quality of good she receives, both of which are her private information. Finally, in any repeated game, if players are subject to observation errors, the resulting game becomes one with private monitoring. (Here, ω_i corresponds to i 's observation of actions by other players, say, a_{-i} plus some noise².)

¹This formulation makes sure that realized payoff conveys no more information than is already contained in a_i and ω_i . In the secret price cutting story, firm i 's realized profit u_i depends on its price a_i and sales ω_i . Also note that this is without loss of generality, as we can always redefine private signal as $\omega'_i = (\omega_i, u_i)$.

²As player i 's realized payoff may well be a function of the actual actions as opposed to the observed actions, the realized payoff may convey some additional information. Hence

1.2 An Illustration of the Current State of Art - Prisoner's Dilemma

Let us now briefly summarize the current state of our knowledge about repeated games with various information structures. For the reader's convenience, I will exemplify the general results in terms of repeated prisoner's dilemma game, whose stage payoff table is given by Table 1³.

	C	D
C	$1, 1$	$-l, 1 + g$
D	$1 + g, -l$	$0, 0$

Table 1

In the *perfect monitoring* case, any outcome which dominates the minimax payoff profile⁴ $(0, 0)$ can be sustained by a subgame perfect equilibrium of the repeated game, when discount factor δ is close to 1 (the Folk Theorem by Fudenberg and Maskin [19]). This result holds for any game with a generic choice of stage game payoff functions.

Under *imperfect public monitoring*, basically the same region of payoffs can be sustained when the publicly observable signal $\omega \in \Omega$ takes on sufficiently many values. In the prisoner's dilemma game, the Folk Theorem by Fudenberg, Levine and Maskin [18] implies that for a generic signal distribution $p(\omega|a)$, the same area of payoffs as in the perfect monitoring case can be (approximately) sustained, if discount factor is sufficiently close to one, as long as Ω contains at least three elements⁵.

In contrast, we do not yet have a fully general characterization of payoffs achieved under *imperfect private monitoring*. In fact, although it deceptively

one may assume that player i 's signal consists of her observation and realized payoff (and that the latter is subject to random shocks so that it does not perfectly reveal others' actions). Alternatively, we may assume that the game is terminated with a certain probability q in each stage, and the players observe or receive the actual (total) payoff only after the game is terminated. (Note that when players do not discount, this is isomorphic to the game with discount factor $1 - q$, where the players receive no more information than their observations.)

³Here, we assume $g, l > 0$ (D is a dominant strategy) and $1 > g - l$ ((C, C) is Pareto efficient).

⁴Player i 's minimax payoff, $\min_{a_{-i}} \max_{a_i} g_i(a)$, is the payoff which she can guarantee herself in any equilibrium. It corresponds to 0 in the prisoner's dilemma game.

⁵In general, the folk theorem holds for a generic choice of payoff function and signal distribution, as long as $|\Omega| \geq |A_i| + |A_j| - 1$ for each pair of players $i \neq j$. One can check that the "individual and pairwise full rank conditions" assumed in Fudenberg-Levine-Maskin folk theorem is generically satisfied under this condition.

looks like a simple homework exercise, just constructing *any* equilibrium (apart from the repetition of the stage game equilibrium) is far from trivial for the reasons explained in the next section. Hence, Sekiguchi [33] came as a surprise, which was the first to construct an equilibrium that can approximately sustain the cooperative payoff $(1, 1)$ in the prisoner's dilemma game under private monitoring, assuming that the monitoring structure is nearly perfect. Sekiguchi's work initiated the rapidly growing literature, and Bhaskar and Obara [8] extended Sekiguchi's construction to support any point Pareto dominating $(0, 0)$, when monitoring is private but almost perfect. Piccione [31] introduced a completely different, useful technique to support essentially the same area under almost perfect monitoring. All those works employ some conditions on payoffs and/or information structure, but Ely and Valimaki [15], who extended Piccione's technique, managed to remove those assumptions and proved the folk theorem for the prisoner's dilemma with private monitoring, when monitoring is almost perfect. A strong result is obtained recently by Matsushima [29], who further extended Ely and Valimaki's construction to show that their folk theorem continues to hold even if monitoring is far from perfect, as long as private signals are distributed independently (given an action profile).

2 The Difficulties Associated with Private Monitoring

Why is the private monitoring case so different from the (perfect or imperfect) public monitoring case? Basically, when players do not share common information, we encounter two major difficulties. Firstly, the games under private monitoring lack the *recursive structure* in the sense of Abreu, Pearce and Stacchetti [3], so that the set of equilibria does not possess a simple characterization. Secondly, at each moment of time, players must conduct *statistical inference* (not only to detect potential deviations, as in the imperfect public monitoring case, but also) on what others are going to do, which can be quite complex. Let us explain each of them in turn.

Under public monitoring, players can condition their actions on commonly observed events. In perfect monitoring case, all strategies share this nature, and the majority of existing literature in imperfect public monitoring case focuses on such behavior, called *public strategies*. Public strategies specify actions in each stage depending only on the history of commonly observed signal (hence the history of one's own action is ignored). Equilibria in this class of strategies (*perfect public equilibrium*) turned out to be rich

enough to obtain the various Folk Theorems for the public monitoring case.

When players condition their future action plans on commonly observed events, after any history, they play a Nash equilibrium of the remaining game, which is identical to the original infinitely repeated game. This means that, after any history, the set of continuation payoffs is always equal to the equilibrium payoff set of the repeated game. This is the recursive structure explored by Abreu, Pearce and Stacchetti [3]. Note that the continuation payoffs at time t is generated by current stage game payoffs and the continuation payoffs at $t+1$. We can write down this relationship as $W_t = B(W_{t+1})$, where W_s is the set of continuation payoffs at time s . Thanks to the recursive structure, the set W^* of perfect public equilibrium payoffs in (perfect or imperfect) public monitoring case is characterized by simple fixed point (or 'self-generating') equation, $W^* = B(W^*)$.

Under private monitoring, however, such a simple characterization is no longer available. At each moment t , player i conditions her action on the history of her action and private signal, $(a_i(0), \dots, a_i(t-1), \omega_i(0), \dots, \omega_i(t-1))$, which is only known to her. We call it her *private history* and denote it by h_i^t . On the equilibrium path in a private monitoring game, the probability distribution of private histories is common knowledge, and they are taking mutual best replies. This means that the continuation play at time $t > 0$ *on the equilibrium path* is a correlated equilibrium of the repeated game, where the private histories play the role of correlation device. Note that the correlation device (h_1^t, \dots, h_N^t) becomes increasingly more complex over time, so the set of continuation payoffs (the associated correlated equilibrium payoffs) generally changes. A part of the stationary structure in the public monitoring case is lost here. Compte [11] considers a repeated prisoner's dilemma game with private monitoring where defection is irreversible, and he shows that a kind of stationarity can be recovered by introducing a correlation device at $t = 0$. He constructed an equilibrium where efficiency is achieved as the discount factor tends to 1.

Furthermore, when a player deviates, she knows that the distribution of private histories are altered, which is not known to other players. Hence, after a deviation, the distribution of the correlation device (private histories) is no longer common knowledge, and therefore the continuation play *off the equilibrium path* is not even an equilibrium of the original game⁶. Hence, the

⁶Alternatively, one can view the continuation game at time t as a Bayesian game, where the beliefs on types h_i^t are given by *conditional* distributions $Pr(h_i^t | h_i^t)$. (This is somewhat non-standard definition, as the conditional type distributions are *not* derived by a common prior.) Then, the continuation strategy profile, which specifies the play *on and off* the equilibrium path, can be regarded as a Bayesian Nash equilibrium of this game.

recursive structure found in the public monitoring case, i.e., the property that the continuation payoff after any history is chosen from the identical set of equilibrium payoffs, is lost under private monitoring. As a result, the set of equilibria cannot be characterized by the simple self-generation condition, which played a major role in the analysis of the public monitoring case. Amarante [4] shows that certain aspects of the recursive structure survive in the private monitoring case. In particular, a version of the successive approximation method ⁷ to find equilibrium payoff set remains to be true.

The second difficulty is that checking incentives in each stage requires fairly complex *statistical inference*. To determine the best strategies at each moment of time, players must know what others are going to do. This is immediate under public monitoring when they use public strategies, as the future action plans are always common knowledge. Under private monitoring, however, each player must make a statistical inference about others' private histories to determine what they are going to do. In other words, player i should calculate conditional distribution $Pr(h_{-i}^t|h_i^t)$ by Bayes' rule in each stage t , and this can become increasingly more complex as time passes by. Hence, checking a player's incentives after all histories is typically quite demanding (even though others are using relatively simple strategies), and as a result just constructing any equilibrium (other than the repetition of the one-shot Nash equilibrium) under private monitoring turns out to be a non-trivial task.

Finally, note that even in the public monitoring case, we encounter the same difficulties as described above, once we consider strategies that are not public (i.e., the ones where current action depends on one's own past actions; we call such strategies *private*). Hence closely related techniques and results are obtained *both* for private equilibria in public monitoring case *and* for private monitoring case. Kandori [23] and Obara [30] (combined to appear in [25]) show that private equilibria can payoff-dominate any public equilibrium in repeated games with imperfect public monitoring. Mailath, Matthews and Sekiguchi [28] demonstrate various methods to construct private equilibria in finitely repeated games with public monitoring.

⁷This method considers a finite (T) repetition of stage game plus *arbitrary* terminal payoff function, whose range is a sufficiently large bounded set. It is shown that a strategy profile in the infinitely repeated game is an equilibrium if and only if it is the limit of the equilibrium of the T -stage game (as $T \rightarrow \infty$).

3 Prior Contributions

There are some prior contributions which manage to bypass the aforementioned difficulties.

No Discounting or ϵ -Rationality: Firstly, efficient equilibria under private monitoring have been obtained in the case with no discounting or approximate optimization (where ϵ loss in the average discounted payoff is tolerated), by Radner [32], Fudenberg and Levine [17] and a series of papers by Lehrer (for example, see [26]). In those settings, each player has to deviate infinitely often to get any payoff improvement, and as a result checking incentives is relatively easy. However, as this property no longer holds in discounted case with full rationality, the proposed equilibrium strategies in those works do not work, once we have any amount of discounting and full rationality. The continuity between the discounted and undiscounted cases remains to be seen.

Communication: Secondly, Compte [9] and Kandori and Matsushima [24] demonstrated that introducing communication at each stage of a repeated game solves the aforementioned difficulties, and they proved the folk theorems. At each stage, players are asked to reveal their private signals, and they can tell a lie if that is profitable. However, by constructing equilibria where one's report is used to police *other* players and does not affect one's own future payoff, players can be induced to tell the truth. Given this idea, we can construct equilibrium strategies which only depend on the *publicly* observable history of communication. This is similar to the perfect public equilibria in the public monitoring case. With analogous assumptions to Fudenberg, Levine and Maskin [18]'s pairwise full rank condition, the folk theorem is obtained when there are at least three players⁸. Recent paper by Aoyagi [1] demonstrates an alternative way to construct an efficient equilibrium with communication. He shows that a version of the secret price cutting example discussed in Section 1.1 has a special information structure to facilitate nearly efficient collusion by a simple equilibrium with communication, which works quite differently from Compte or Kandori-Matsushima.

As one may argue that communication is readily available in a number of applications, we have to examine carefully the motivation to study the private monitoring case without communication. First note that there are some cases where communication is simply not feasible. The example of exchange of goods discussed in 1.1 may be regarded as a stylized version

⁸In the two-player case, the folk theorem can be obtained by infrequent communication. This is based on the idea of Abreu, Milgrom and Pearce [2] that delaying the release of information helps to achieve efficiency.

of the medieval long distance trade, where there was no effective means of communication between the traders living in the different areas. More importantly, even in the modern age communication to facilitate collusion between firms is often infeasible, as it is deemed illegal by the antitrust law. An important motivation to study the case without communication is to determine the effectiveness of such provision in the anti trust law, as we already know that full collusion is possible with communication under mild assumptions. Secondly, if communication is subject to some noise, the resulting game becomes again the one with private monitoring, as is in the observation error model we discussed in Section 1.1. Lastly, from the point of view of pure theory, it is important to determine what is possible if players share no common information.

Partial Observation: There are also related prior contributions to examine specific classes of private monitoring games. The case where each player's action is perfectly observed by a *subset* of players is examined by Kandori [22], Ellison [13] and Ben-Porath and Kahneman [6]. A leading example is a random matching game, where each player only observes what her opponents have done to her. In the former two papers, it is shown that efficient outcome is achieved without communication in repeated prisoner's dilemma with random matching, by means of 'contagion' strategies. Kandori also shows the folk theorem provided that what one observes in today's match is (honestly) passed on to her next match. Ben-Porath and Kahneman showed the folk theorem with communication. A general characterization of equilibrium without communication in this class is subject to the same difficulties discussed in the last section, and it is yet to be obtained.

4 Insights from Two-Period Examples

In this section we present, by means of simple two-period examples, some of the basic ideas of the papers appearing in this issue. Consider a two-stage game whose first period game is given by the prisoner's dilemma game in Table 1. At the end of the first stage, each player $i = 1, 2$ receives a private signal $\omega_i \in \{c, d\}$, whose distribution depends on the action profile in the first stage, denoted by $a \in \{C, D\} \times \{C, D\}$.

4.1 Coordinated Punishment

Firstly, let us assume that the second stage game is given by Table 2.

	X	Y
X	2, 2	0, 0
Y	0, 0	1, 1

Table 2

If errors in the signals (i) occur with small probabilities and (ii) are sufficiently *correlated*⁹, then cooperation (C, C) can be achieved by much the same way as in the perfect/public monitoring case. It is easy to check that the "coordinated punishment strategy",

- (*) playing C in the first period and then choosing X
if and only if one's own action and signal were C and c

is an equilibrium. A natural conjecture is that there should be continuity between the public monitoring case and the private monitoring case with highly correlated signals. However, Mailath and Morris [27] present an example where this fails. They go on to show that the continuity holds under some assumptions. Specifically, if an equilibrium in public a monitoring game gives strict incentives and specifies current action depending on a *finite history* of public signal, then there is a similar equilibrium in the private monitoring game with highly correlated signals

If the private signals are independent (given each action profile), in contrast, cooperation (C, C) cannot be sustained by any *pure* strategy, even though observability is nearly perfect. Suppose both players adopt strategy (*) above and assume that player 1 receives $\omega_1 = d$. By the equilibrium expectations player 1 believes that this is an "error" and opponent actually played C . As "errors" are not correlated across players and occur with small probabilities, player 1 also believes that the opponent is observing c with a high probability, as player 1 chose C in the first period. Hence 1 believes that the opponent is going to choose X with a high probability, and she is not willing to 'punish' the opponent even though her highly informative signal takes the value d .

However, Bhaskar and van Damme [7] show that cooperation (C, C) can be sustained with a large probability by (i) mixed strategies in stage one and (ii) public randomization in stage two. If the players mix in the first period, the players receive *correlated* information, (a_1, ω_1) and (a_2, ω_2) , even though the signals are independent. Note that a_i and ω_j ($i \neq j$) are highly correlated, as errors are rare. With this correlation device, the players can

⁹We say that we have an error if we have $\omega_1 = d$ when $a_2 = C$, and so on. The errors are positively correlated if my opponent is more likely to get an error when I get one.

utilize the coordinated punishment, where player i plays X if and only if $(a_i, \omega_i) = (C, c)$. The public randomization in the second stage, in contrast, is necessary for a somewhat subtle reason. In the mixed strategy equilibrium, the players are indifferent between C and D , and the equilibrium payoff should be, in particular, equal to the payoff associated with D . When the signals are accurate, this is detected with a large probability and punishment is triggered in the second period. Hence, if the punishment is severe, the overall payoff becomes low. If we mitigate the punishment by public randomization, however, we can increase the equilibrium payoff, and the efficient payoffs can be approximately achieved¹⁰.

The coordinated punishment idea, originally proposed in the earlier version of Bhaskar and van Damme [7] in a two period example, was substantially extended by Sekiguchi [33] to infinitely repeated games. Specifically, he showed that the symmetric efficient payoff can be sustained in the prisoner's dilemma game with private monitoring, provided that the signals are sufficiently accurate and the discount factor is close to 1. The equilibrium is basically a mixture of the trigger strategy and permanent defection, where the original game is divided into K independent repeated games, each of which is played in every K period. This has the same effects as restarting the game anew in each period with a certain probability. We can see in this construction the crucial features of the above example; the use of mixed strategy and public randomization (i.e., restarting the game anew). Bhaskar and Obara [8] managed to provide a very much simplified version of Sekiguchi's equilibrium and showed that asymmetric payoffs can also be sustained. They also consider the prisoner's dilemma with N players. Sekiguchi [34] relaxes some assumptions on the information structure in his original paper, by introducing a new method of identifying equilibrium paths without fully constructing the equilibrium strategies.

Compte [10] shows that it is vital to restart the game in their equilibria: the grim trigger strategies, where the game is *never* started anew, achieve no cooperation when discount factor is close to unity under private monitoring. As this is true even if the observability is nearly perfect, Compte's result suggests that a certain discontinuity exists between perfect and almost perfect private monitoring case. A recent paper by Ely [14], in contrast, shows that the discontinuity is resolved if we view the grim trigger strategies as degenerate correlated equilibria. Namely, he shows that there exists a se-

¹⁰Essentially the same issue arises in the literature on the Stackelberg game under observation errors (Bagwell [5] and van Damme and Hurkens [12]). See Bhaskar and van Damme [7] for details.

quence of correlated equilibria under private monitoring converging to the grim trigger strategies, as the signaling noise goes to zero.

4.2 Uncoordinated Punishment

Let us go back to the two stage example and suppose that the second stage game is now given by Table 3.

	X_2	Y_2
X_1	5, 6	1, 5
Y_1	6, 0	0, 1

Table 3

This is essentially the matching pennies game, with unique equilibrium being the equal mix of each action. Note that X_i is a rewarding action that gives high payoffs to the opponent, while Y_i offers low payoffs and can potentially be used as punishment. Assume that signals are independent and have small probabilities of errors, and consider the following action plan in the second stage. If player i 's signal was d , she plays the punishing action Y_i for sure. If the signal was c , on the other hand, she mixes X_i and Y_i in such a way that the overall probability of taking X_i or Y_i is just equal to $1/2$ (the mixed strategy equilibrium of the second stage game). Given this action plan, the opponent has an incentive to play C in the first stage (as long as the gain from deviation g is not too large), because defection increases the probability of punishment Y_i . As a result, cooperation (C, C) in the first stage followed by the unique (mixed strategy) equilibrium in the second stage can be sustained as an equilibrium. This is the basic idea in Kandori [21]¹¹. The crux of the matter is that each player is *indifferent* between rewarding and punishing actions, so that she does not object to play the latter when she receives a 'bad' signal. The basic idea may be phrased as "*uncoordinated punishment*", to be contrasted to the coordinated punishment discussed above.

¹¹Kandori [21] considers the same stage game repeated in each stage; a strictly dominated action C is introduced to the game in Table 3, where (C, C) achieves the symmetric efficient payoffs. By the same argument, it is shown that (C, C) can be sustained in the first stage. This shows that cooperation can be sustained in a finitely repeated game *even though the stage game has a unique equilibrium*, when the monitoring structure is private. Also note that repeating the two-stage equilibrium is an equilibrium of the infinitely repeated version of this game, and to the best of my knowledge this is the first example of an equilibrium in private monitoring game which is not a repetition of one-shot Nash equilibrium.

Piccione’s paper [31] introduced a general technique to construct uncoordinated punishment in infinitely repeated games. Note that, if we had more than two stages, constructing a similar example to the one presented above would be fairly complex, as computing one’s beliefs about the opponent’s private history is a tedious task even after a few stages (the second difficulty discussed in Section 2). Piccione introduced an ingenious idea, which makes this problem irrelevant. He considers the repeated prisoner’s dilemma with private monitoring, and constructed an equilibrium strategy represented by an automaton with countably many states. Piccione showed that it is possible to construct an automaton in such a way that each player is *always indifferent* to play C and D *no matter which state the opponent is in*. Piccione’s construction is similar to the two period game where the second stage game is given in Table 4.

	X	Y
X	1, 1	0, 1
Y	1, 0	0, 0

Table 4

Note that in Table 4 each player is always indifferent between X and Y no matter what the opponent does. Hence she can reward (by playing X) or punish (Y) the opponent according to her private signal, irrespective of the opponent’s behavior in the second stage game. Likewise, in Piccione’s equilibrium each player does not have to compute her beliefs about what her opponent has been observing, which provides a drastic simplification of the analysis (the second difficulty mentioned in Section 2 is resolved). Piccione endogenously derives continuation games similar to Table 4 by showing that the system of dynamic programming equations for value functions has a relevant solution. With some restrictions, Piccione showed the folk theorem for the prisoner’s dilemma, when monitoring error tends to zero.

Piccione’s construction is substantially simplified independently by Ely and Valimaki [15] and Obara [30] (to appear in Kandori and Obara [25]). They showed that construction similar to Piccione’s can be obtained by just two states, and this sweeping simplification broke new ground and provided the possibility to extend the analysis in various directions. Ely and Valimaki managed to remove information or payoff restrictions for the previous folk theorems for the prisoner’s dilemma with almost perfect monitoring, and they also examine more general stage games. Obara emphasizes that the same construction can be used to construct private equilibria in public

monitoring case¹², and showed that sometimes private strategy equilibria dominates public equilibria. Obara [30] and a recent paper by Ely and Valimaki [16] characterize the maximum payoffs associated with those class of equilibria.

A recent paper by Matsushima [29] further extends the above ideas to prove the folk theorem for the prisoner's dilemma even though the monitoring is far from perfect. This paper combines the idea of Ely, Valimaki and Obara and Abreu, Milgrom and Pearce [2]. The latter showed that delaying the release of information can improve efficiency. Matsushima redefines the stage game as the T times repetition of the original stage game, where information is pooled for a statistical testing to determine future payoffs. Matsushima showed that it is possible to modify Abreu et al's statistical testing to make 'always cooperate for T periods' and 'always defect for T periods' indifferent and all other strategies strictly worse. This effectively reduces the stage game strategies into those two. Given this, and when T is large, the resulting stage game is the one with almost perfect monitoring (note that deviating T times can easily be detected, when T is large), and the two-state automata construction of Ely-Valimaki-Obara, which works when monitoring is nearly perfect, can be applied to prove the folk theorem.

Bhaskar (see [7] and [8]) questions the robustness of those works built on the uncoordinated punishment idea. He shows that those mixed strategy equilibria do not admit Harsanyi's purification, if the payoff perturbations are additively separable, as the repeated game payoffs are.

¹²An advantage of this class of equilibria is that the beliefs about the opponent's state is irrelevant, and this means that the strategies work irrespective of the degree of correlation of the signals. In particular, they also work when the signals are perfectly correlated, which is nothing but the public monitoring case.

References

- [1] M. Aoyagi, Collusion in dynamic Bertrand oligopoly with correlated private signals and communication, *J. Econ. Theory*, this issue.
- [2] D. Abreu, P. Milgrom and D. Pearce, Information and timing in repeated partnerships, *Econometrica* **59** (1991), 1713-1733.
- [3] D. Abreu, D. Pearce and E. Stacchetti, Toward a theory of discounted repeated games with imperfect monitoring, *Econometrica* **58** (1990), 1041-1064.
- [4] M. Amarante, Recursive structure and equilibria in games with private monitoring, mimeo., 1997.
- [5] K. Bagwell, Commitment and observability in games, *Games Econ. Behav.* **8** (1995), 271-280.
- [6] E. Ben-Porath and M. Kahneman, Communication in repeated games with private monitoring, *J. Econ. Theory* **70** (1996), 281-297.
- [7] V. Bhaskar and E. van Damme, Moral hazard and private monitoring, *J. Econ. Theory*, this issue.
- [8] V. Bhaskar and I. Obara, Belief-based equilibria in the repeated prisoners' dilemma with private monitoring, *J. Econ. Theory*, this issue.
- [9] O. Compte, Communication in repeated games with imperfect private monitoring, *Econometrica* **66** (1998), 597-626.
- [10] O. Compte, On failing to cooperate when monitoring is private, *J. Econ. Theory*, this issue.
- [11] O. Compte, Sustaining cooperation without public observations, *J. Econ. Theory*, this issue.
- [12] E. van Damme and S. Hurkens, Games with imperfectly observable commitment, *Games Econ. Behav.* **21** (1997), 282-308.
- [13] G. Ellison, Cooperation in the prisoner's dilemma with anonymous random matching, *Rev. Econ. Stud.* **61** (1994), 567-588.
- [14] J. C. Ely, Correlated equilibrium and private monitoring, mimeo., 2000.

- [15] J. C. Ely and J. Välimäki, A robust folk theorem for the prisoner's dilemma, *J. Econ. Theory*, this issue.
- [16] J. C. Ely and J. Välimäki, Notes on private monitoring with non-vanishing noise, mimeo., 2000.
- [17] D. Fudenberg and D. Levine, An approximate folk theorem with imperfect private information, *J. Econ. Theory* **54** (1991), 26-47.
- [18] D. Fudenberg, D. Levine, and E. Maskin, The folk theorem with imperfect public information, *Econometrica* **62** (1994), 997-1040.
- [19] D. Fudenberg and E. Maskin, The folk theorem in repeated games with discounting or with incomplete information, *Econometrica* **54** (1986), 533-554.
- [20] E. Green and R. Porter, Noncooperative collusion under imperfect price information, *Econometrica* **52** (1984), 87-100.
- [21] M. Kandori, Cooperation in finitely repeated games with imperfect private information", mimeo., 1991.
- [22] M. Kandori, Social norms and community enforcement, *Rev. Econ. Stud.* **59** (1991), 63-80.
- [23] M. Kandori, Check your partner's behavior by randomization: new efficiency results on repeated games with imperfect monitoring, CIRJE Discussion Paper F-49, University of Tokyo, 1999.
- [24] M. Kandori and H. Matsushima, Private observation, communication and collusion, *Econometrica* **66** (1998), 627-652.
- [25] M. Kandori and I. Obara, Efficiency in repeated games revisited: The role of private strategies, mimeo., 2000.
- [26] E. Lehrer, Nash equilibria of n-player repeated games with semi-standard information, *Int. J. Game Theory* **19** (1990), 191-217.
- [27] G. J. Mailath and S. Morris, Repeated games with almost-public monitoring, *J. Econ. Theory*, this issue.
- [28] G. J. Mailath, S. A. Matthews, and T. Sekiguchi, Private strategies in repeated games with imperfect public monitoring, CARESS Working Paper #01-10, University of Pennsylvania, 2001.

- [29] H. Matsushima, The folk theorem with private monitoring and uniform sustainability, CIRJE Discussion Paper F-84, University of Tokyo, 2000.
- [30] I. Obara, Private strategy and efficiency: Repeated partnership game revisited, mimeo., 1999.
- [31] M. Piccione, The repeated prisoner's dilemma with imperfect private monitoring, *J. Econ. Theory*, this issue.
- [32] R. Radner, Repeated partnership games with imperfect monitoring and no discounting, *Rev. Econ. Stud.* **53** (1986), 43-58.
- [33] T. Sekiguchi, Efficiency in the prisoners' dilemma with private monitoring, *J. Econ. Theory* **76** (1997), 345-361.
- [34] T. Sekiguchi, Robustness of efficient equilibria in repeated games with imperfect private monitoring, mimeo., 1999.