

Article

Repeated Play of Families of Games by Resource-Constrained Players

Arina Nikandrova

Department of Economics, Mathematics and Statistics, Birkbeck College, Malet Street, London WC1E 7HX, UK; E-Mail: a.nikandrova@bbk.ac.uk; Tel.: +44-(0)-20-7631-6457; Fax: +44-(0)-20-7631-6416.

Received: 4 April 2013; in revised form: 28 June 2013 / Accepted: 3 July 2013 /

Published: 11 July 2013

Abstract: This paper studies a repeated play of a family of games by resource-constrained players. To economize on reasoning resources, the family of games is partitioned into subsets of games which players do not distinguish. An example is constructed to show that when games are played a finite number of times, partitioning of the game set according to a coarse exogenously given partition might introduce new symmetric equilibrium payoffs which Pareto dominate best equilibrium outcomes with distinguished games. Moreover, these new equilibrium payoffs are also immune to evolutionary pressure at the partition selection stage.

Keywords: evolutionary stability; repeated games; bounded rationality; analogies

1. Introduction

Classical game theory assumes that economic agents are perfectly rational and have unlimited reasoning resources, which enable them to make correct decisions at all times. In reality, reasoning resources are scarce and the economic environment is complex, with agents interacting in many settings or games. If games played by agents are similar in some dimensions, distinguishing all games at all times requires undue reasoning effort. This paper examines whether an evolutionary process could weed out such a waste of resources.

Formally, it is assumed that players repeatedly play games drawn from some pre-specified family. The family of games remains fixed for the duration of play, but in every period, a new game is drawn. Due to the scarcity of reasoning resources, the family of games is partitioned into analogy classes, *i.e.*, the subsets of games that players do not distinguish. While players choose strategies in the repeated

games induced by the partitions rationally, the choice of how to partition the game set is beyond their control and is, instead, governed by some evolutionary process.

In this setting, an example is constructed demonstrating that when games are played a finite number of times, not knowing exactly which game is currently being played might make all players better off. Intuitively, coarse partitioning of the game set is a commitment device that allows players to play strategies that are not incentive-compatible when players know perfectly which game is currently being played. Such commitment can result in a Pareto improvement. Furthermore, in some instances, the new equilibrium payoffs are also immune to evolutionary pressure that operates at a partition selection level (also referred to as a partition selection meta-game).

In the context of this paper, introducing an appropriate notion of evolutionary viability poses a conceptual challenge. Ideally, the evolutionary pressure should operate on the level of partition selection only, without affecting the strategies players follow in the resulting dynamic game. Such dichotomy between decision-making and evolution of cognition is natural. On the one hand, in everyday interaction, people make conscious choices, taking their reasoning and cognition constraints as given. On the other hand, the process that shapes human reasoning operates on a different time scale and it takes countless generations for substantial changes to occur. However, in repeated games, such separation is difficult, as the partition determines the knowledge of players about the play of the game so far. Hence, an evolutionary viable strategy has to specify a partition and an equilibrium strategy in the resulting game.

The adopted definition of evolutionary viability is related to the notion of an evolutionarily stable strategy and attempts to identify partition-strategy pairs that are immune to one-off invasion by a small number of mutants with a different partition. It is assumed that players in the population are (boundedly) rational, but do not know much about the workings of the evolutionary process and, hence, are unaware of the possibility of mutation.

The idea that reasoning resources are scarce and that people use analogies to make decisions in their day-to-day interactions is not new. For example, [1] advocates that players, rather than being action rational, *i.e.*, consciously optimizing in each decision situation, are rule rational and apply “rules of thumb”. These rules of thumb emerge as a result of an evolutionary process and work well on average. For example, leaving a tip at a restaurant is rational only if one intends to come back to that restaurant and would like to incentivize staff to provide good service in future. However, customers routinely leave generous tips even when dining at restaurants where they are unlikely to come back in the near future. The rule “always tip generously” works well on average and spares its users the anguish of routinely estimating the likelihood of coming back.

Similarly, [2] suggest that the use of vague, imprecise language can be explained only if people have a vague view of the world. This paper models interactions of a player who has evolved to have an imprecise view of the world and their day-to-day interactions.

The setup of this paper is closely related to [3], who studies the learning process carried out by agents who are involved in many games. [3] builds a dynamic model in which players simultaneously learn how to partition the game set and which actions to choose in one-shot two-player games. Instead of being rational, players adapt their behavior through reinforcement learning. Unlike [3], this paper assumes that players are engaged in repeated games and that evolutionary pressure operates only at the partition selection stage, while subsequent decisions of players are rational.

This paper is also distantly related to the literature on commitment in games, e.g., [4]. This strand of literature assumes that in the first stage, players voluntarily and credibly commit to a subset of their pure strategies, and in the second stage, they play the game induced by their commitment. In this paper, players are limited in what they know about the game and not in what they can do in the game. More importantly, here, each player’s partition of the game set is not publicly observable and, thus, the explicit commitment mechanism is lacking.

The rest of the paper is organized as follows: Section 2 outlines an example and characterizes equilibrium payoffs under exogenously given partitions, Section 3 analyzes the evolutionary viability of coarse partitioning, and Section 4 concludes.

2. Fixed Partitions

There are two players interacting over two periods. In every period $t = 1, 2$ they play a normal form game that is randomly and independently drawn from the set $\Gamma = \{G^1, G^2\}$ according to probabilities $p^j = 1/2, \forall G^j \in \Gamma$. Each game G^j is a finite simultaneous move game and both games in Γ share the same action space. G^1 has the payoff matrix:

	A	B	C
A	1, 1	0, 0	0, -0.1
B	0, 0	2, 2	0, 0
C	-0.1, 0	0, 0	2, 2

Note that G^1 has three Nash equilibria: (A, A) , (B, B) and (C, C) . Outcomes (B, B) and (C, C) are also Pareto-efficient in this game. G^2 has the following payoff matrix:

	A	B	C
A	0, 0	10, 0	0, 0.1
B	0, 10	7, 7	0, 0
C	0.1, 0	0, 0	6, 6

The unique Nash equilibrium of this game is (C, C) . However, payoffs of the Nash action profile are Pareto-dominated by outcome (B, B) .

Players are rational, but they have limited reasoning resources, which makes distinguishing games costly. Thus, players may partition Γ into subsets of games they do not distinguish. These subsets are called analogy classes. The collection of player i ’s analogy classes is referred to as i ’s partition of Γ , and it is denoted by Π_i . With only two games in Γ , player i can either distinguish the games and have the finest partition $\mathcal{D} = \{\{G^1\}, \{G^2\}\}$ or not distinguish the games and have the coarsest partition $\mathcal{N} = \{\{G^1, G^2\}\}$.

In this section, it is assumed that each player is endowed with an exogenously given partition Π and both players partition the game set Γ in the *same* way. Thus, whenever both players do not distinguish

the games, they perceive that they are playing an “average” game $\frac{1}{2}G^1 + \frac{1}{2}G^2$ with the following payoff matrix:

	A	B	C
A	0.5, 0.5	5, 0	0, 0
B	0, 5	4.5, 4.5	0, 0
C	0, 0	0, 0	4, 4

Partition Π divides histories of dynamic game $\{G^j, p^j\}_{j=1}^2$ into equivalence classes and players, who partition Γ according to Π , can condition their continuation play only on these equivalence classes. This implies that in all respects, the fixed partition game is a standard dynamic game with the exception that the strategies available to players are restricted to some class.

Let $U_i^\Pi(s)$ denote the expected average per period payoff of player i when players partition Γ according to $\Pi = (\Pi_i, \Pi_{-i})$ and follow strategy profile $s = (s_i, s_{-i})$:

$$U_i^\Pi(s) = \mathbf{E} \left[\frac{1}{2} \sum_{t=1}^2 u_i^t \right]$$

where $\{u^t\}_{t=1}^2$ is a sequence of stage game payoffs associated with a realized outcome path and expectation is taken with respect to the measure over outcomes induced by s and the stochastic process governing the draws of the games. While playing games from Γ , player i aims to maximize $U_i^\Pi(s)$.

In the current setting, the idea of sequential rationality is captured by the notion of subgame perfect equilibrium. Let $\mathcal{E}(\Pi)$ denote the set of subgame perfect equilibrium payoffs in the fixed partition game, where all players partition Γ according to Π . The aim of this section is to compare $\mathcal{E}(\mathcal{N})$ to $\mathcal{E}(\mathcal{D})$.

Below, it is demonstrated that, in the finitely repeated setting, coarse partitioning of Γ can introduce new equilibrium payoffs, and thus, $\mathcal{E}(\mathcal{N})$ is not necessarily included in $\mathcal{E}(\mathcal{D})$. Moreover, both players may be strictly better off in the best subgame perfect equilibrium when they partition Γ according to \mathcal{N} , as compared to the best equilibrium with partition \mathcal{D} .

Suppose players are endowed with partition \mathcal{D} . Our interest lies in finding the highest subgame perfect equilibrium payoff of the induced dynamic game. By the logic of backwards induction, in the second period, a Nash equilibrium should be played in any realized stage game. However, since game G^1 has more than one Nash equilibrium, it is possible to condition the second period play on the outcome of $t = 1$. This could potentially allow the construction of inter-temporal incentives to support non-Nash outcomes in the first period. Nevertheless, here, flexibility in the second period choice of Nash equilibrium in G^1 does not suffice to sustain Pareto-efficient play in $t = 1$. A Pareto-efficient outcome in $t = 1$ involves playing either (B, B) or (C, C) if G^1 is realized and playing (B, B) if G^2 is drawn. Suppose that in $t = 1$, game G^2 is realized, and consider the interim incentive of player i to comply with Pareto-efficient play of B in $t = 1$. If player i deviates and plays A instead, he improves his immediate payoff by three. This deviation triggers a punishment in $t = 2$, which costs one if G^1 is drawn and nothing, otherwise. Hence, the expected magnitude of the future punishment is 0.5, which is less than three, the myopic incentive to deviate.

Since in $t = 1$, Pareto-efficient play is impossible, the average expected payoff in the best subgame perfect equilibrium is four. This payoff is attained by playing one of the ”good” Nash equilibria, (B, B) or (C, C) whenever game G^1 is drawn and playing the only Nash equilibrium, (C, C) whenever game G^2 is drawn.

Suppose players are endowed with partition \mathcal{N} . Then, it is as if in each t , they are playing the average game $\frac{1}{2}G^1 + \frac{1}{2}G^2$, which has two Nash equilibria: (A, A) and (C, C) . However, the Pareto-efficient outcome is (B, B) . This outcome can be sustained as a subgame perfect equilibrium outcome in $t = 1$ by the following strategy for $\forall i$:

s_i^* : Play B in $t = 1$; if (B, B) is played in the first period, play C in $t = 2$; otherwise, play A in $t = 2$.

This strategy prescribes playing a Nash equilibrium in $t = 2$ in the subgames after a deviation, as well as in the subgame where there was no deviation. In $t = 1$, the best deviation from B yields an immediate gain of 0.5, but triggers reversion to the unfavorable Nash equilibrium in $t = 2$, which costs 3.5. Hence, player i finds it optimal to follow the prescribed strategy, s_i^* , in $t = 1$.

Thus, when players do not distinguish games, the best average expected payoff is:

$$\frac{1}{2}(4.5 + 4) = 4.25$$

which exceeds four, the best expected payoff attainable when players partition the set of games according to \mathcal{D} .

3. Partition Selection

The aim of this section is to identify partitions that are viable under evolutionary pressure. To this end, the notion distantly related to an evolutionary stable strategy (ESS) is deployed.

A strategy is evolutionarily stable if there exists an $\epsilon_0 > 0$ such that for all $\epsilon < \epsilon_0$, the population playing the native strategy can resist any invasion of ϵ mutants. This definition is equivalent to requiring the native strategy to be a best response to itself, as well as to satisfy an additional stability condition. Thus, the ESS notion is a refinement of a symmetric Nash equilibrium.

In the present setting, by assumption, the evolutionary pressure operates only at the level of partitioning of Γ . Hence, the interest lies in identifying the evolutionarily stable partitions. However, the viability of a partition depends on the strategies the population and the invading mutants subsequently follow in the dynamic game, and it is difficult to separate the evolutionary selection of viable partitions of Γ from the selection of strategies in the dynamic game.

Perhaps an obvious way to define an evolutionarily stable partition of Γ is to consider the following two-stage game. In the first stage, players commit to a partition of Γ , and their choice becomes public knowledge. In the second stage, players play a subgame perfect equilibrium of the dynamic game induced by the chosen partition. An ESS could be defined as a partition profile that constitutes a symmetric (and stable) Nash equilibrium in the first stage of the game. The equilibrium partitions depend on the continuation payoffs of players, and a strategy of a player has to specify the course of actions for every, including out-of-equilibrium, selection of the partition. It could be assumed that after a deviation in the partition selection stage, the play of the dynamic game proceeds to the subgame perfect equilibrium with the lowest payoff for the deviator. This setting resembles the commitment games considered by [4].

However, it is somewhat unnatural to assume that the choice of partition is public. If the chosen partition remains private, the only proper subgame of the two-stage game is the game itself and subgame perfection has no bite. In this case, the appropriate equilibrium concept needs to specify the beliefs players hold in every period.

To abstract from the issues related to learning of opponents' partitions, as well as to avoid the evolutionary selection at the level of the dynamic game strategies, the following simplifying assumptions are made. The evolutionarily viable strategy is defined as a tuple $\sigma = (\Pi, s)$ consisting of a partition, Π , and a subgame perfect equilibrium strategy, s , such that, when agents in the population partition Γ according to Π and in the induced game, play s , no mutant with a different partition can do better than incumbents in a paring with incumbents. While playing the induced game, agents from the population believe that with probability 1 all other agents partition Γ according to Π and play strategy s .

The standard definition of ESS applies to symmetric strategies. Here, the symmetry requirement is imposed on the way agents partition Γ , but not on the strategy s . Hence, for a strategy $\sigma = (\Pi, s)$ to be evolutionarily stable, it is necessary that a mutant in the role of any player cannot do better than an agent from the population in the same role.¹

In the partition selection meta-game, the overall utility of a player $\hat{U}_i(\sigma)$, where $\sigma = (\Pi, s)$ with $\Pi = (\Pi_i, \Pi_{-i})$ and $s = (s_i, s_{-i})$, takes into account the payoffs this player derives from playing the dynamic game induced by partition profile Π , $U_i^\Pi(s)$, as well as the cost of sustaining Π_i . It is assumed that $\hat{U}_i(\sigma)$ is lexicographic, firstly increasing in the payoffs derived from playing the games from Γ and, secondly, decreasing in the cardinality of the partition.²

Definition. In the partition selection game, strategy profile $\sigma = (\Pi, s)$ is **partition deviation stable** if and only if $s = (s_i, s_{-i})$ is a subgame perfect equilibrium in the dynamic game induced by $\Pi = (\Pi_i, \Pi_{-i})$, where Π is such that both players partition Γ in the same way, i.e., $\Pi_1 = \Pi_2$, and for each player $i, i = 1, 2$:

1. $U_i^{(\Pi_i, \Pi_{-i})}(s_i, s_{-i}) > U_i^{(\tilde{\Pi}_i, \Pi_{-i})}(\tilde{s}_i, s_{-i})$, or:
2. $U_i^{(\Pi_i, \Pi_{-i})}(s_i, s_{-i}) = U_i^{(\tilde{\Pi}_i, \Pi_{-i})}(\tilde{s}_i, s_{-i})$ and: $|\Pi_i| \leq |\tilde{\Pi}_i|$,

for all alternative strategies of $i, \tilde{\sigma}_i = (\tilde{\Pi}_i, \tilde{s}_i)$.

Section 2 demonstrates that players endowed with coarse partition \mathcal{N} could be better off than players who partition Γ more finely. It turns out that a population of agents endowed with a coarse partition might be immune to one-off invasion by mutants with finer partitions.

Consider the partition selection meta-game and suppose that in this game, strategy σ_i^* is partition deviation stable. Strategy σ_i^* prescribes choosing \mathcal{N} and, in the induced dynamic game, playing according to s_i^* defined in the previous section. Suppose player i is a mutant who can distinguish the games. By assumption, this player should behave optimally at every stage of the dynamic game. If in $t = 1$, this player does not exploit his superior information and mimics the play prescribed by s_i^* , in $t = 2$, he has no profitable deviation either as (C, C) is a Nash equilibrium strategy profile in both games. Thus, for the deviation to be profitable,³ it has to involve a deviation from the play prescribed by s_i^* in $t = 1$.

¹Alternatively, it is possible to symmetrize the setting and define agent's payoff from the repeated play of games from Γ as the average of payoffs across all possible roles. The requirement that a mutant in the role of any player cannot do better than the population is more stringent.

²Similar modification was used by [5] in their study of evolutionary stability in a repeated Prisoner's Dilemma Game played by finite automata.

³A player, who chooses to distinguish the games, but subsequently plays in the same manner as s^* prescribes, obtains $U_i^{(D, \mathcal{N})}(s^*) = U_i^{\mathcal{N}}(s^*)$ from the play of the induced dynamic game. However, he sustains higher cardinality of the partition, and hence, the overall utility of this player is lower than the utility from playing σ_i^* .

In $t = 1$, the best response to B is to play B in G^1 and to play A in G^2 . In $t = 2$, the best response to C is C in both games, while the best response to A is to play A in G^1 and to play C in G^2 . Thus, mutant i , who responds optimally to s_i^* in every stage of the game and does not mimic the play prescribed by s_i^* , obtains the expected payoff of:

$$\frac{1}{2} \left[\frac{1}{2} \left(2 + \frac{2+6}{2} \right) + \frac{1}{2} \left(10 + \frac{1+0.1}{2} \right) \right] = 4.1375$$

from the repeated play, which is less than 4.25, the payoff from following s_i^* . Thus, a mutant with a finer partition obtains lower payoff than the population, and the strategy that prescribes not distinguishing the games and then playing according to s_i^* is partition deviation stable in the partition selection meta-game.

The example indicates that a necessary condition for a strategy profile $\sigma^* = (\mathcal{N}, s^*)$ to be partition deviation stable is that, in the last period on the equilibrium path, s^* prescribes playing a Nash equilibrium in every game, even if games are not distinguished. Otherwise, mutant i with a finer partition definitely has a profitable deviation in $t = 2$.

4. Final Remarks

This paper studies the interaction of substantially rational agents who are involved in a repeated play of normal form games drawn from some fixed family. Reasoning resources are assumed to be costly, and hence, players do not necessarily distinguish all games. The primary interest lies in identifying equilibrium payoffs that are consistent with evolutionary pressure that shapes the constraints on reasoning of agents in the long-run.

In this setting, the most striking result is that, when games are played a finite number of times, coarse partitioning of the game set might shift outwards the Pareto frontier of the dynamic game, thus introducing new symmetric equilibrium payoffs that Pareto dominate the best equilibrium outcomes with distinguished games. This implies that coarse partitioning of the game set eliminates an unnecessary waste of reasoning resources as well as commits players to act in their common interest. Given this, it does not come as a surprise that the new equilibrium payoffs could be immune to evolutionary pressure at the partition selection stage.

Deriving the general conditions under which coarse partitioning of the game set is not only beneficial for society as a whole, but also stable in an evolutionary sense is a challenging task. The necessary conditions can be broadly summarized as follows. In order to generate a Pareto improvement through coarse partitioning, it is necessary that in some games, the Pareto-efficient action profile is not a Nash equilibrium. Furthermore, in these games, it should be impossible to provide inter-temporal incentives for achieving the Pareto efficiency in early periods, when games are distinguished. In contrast, coarse partitioning should provide the flexibility in the choice of future Nash equilibria that is sufficient for sustaining Pareto-efficient play in early periods. Finally, a necessary condition for coarse partitioning to be evolutionary stable is that, in the last period on the equilibrium path, players play a Nash equilibrium in every game, even if games are not distinguished. If this last condition does not hold, a mutant with a finer partition can improve his last period payoff relative to the payoff of the population.

Acknowledgments

I am grateful to Robert Evans, Jonathan Thomas and Flavio Toxvaerd for helpful comments at various stages of this work. All errors remain my own.

Conflict of Interest

The author declares no conflict of interest.

References

1. Aumann, R.J. Rationality and bounded rationality. *Games Econ. Behav.* **1997**, *21*, 2–14.
2. Lipman, B.L. Why is language vague? Working Paper, 2009. Available online: <http://sws.bu.edu/blipman/Papers/vague5.pdf> (accessed on 8 July 2013).
3. Mengel, F. Learning across games. *Games Econ. Behav.* **2012**, *74*, 601–619.
4. Renou, L. Commitment games. *Games Econ. Behav.* **2009**, *66*, 488–505.
5. Binmore, K.G.; Samuelson, L. Evolutionary stability in repeated games played by finite automata. *J. Econ. Theory* **1992**, *57*, 278–305.

© 2013 by the author; licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution license (<http://creativecommons.org/licenses/by/3.0/>).