

Impact of asynchrony on the behavior of rational selfish agents

David Ilcinkas*

Département d'informatique
Université du Québec en Outaouais
Gatineau, Québec J8X 3X7, Canada

Andrzej Pelc†

Département d'informatique
Université du Québec en Outaouais
Gatineau, Québec J8X 3X7, Canada

Abstract. The behavior of rational selfish agents has been classically studied in the framework of strategic games in which each player has a set of possible actions, players choose actions simultaneously and the payoff for each player is determined by the matrix of the game. However, in many applications, players choose actions asynchronously, and simultaneity of this process is not guaranteed: it is possible that a player learns the action of another player before making its choice. Delays of choices are controlled by the adversary and each player can only secure the worst-case payoff over the adversary's decisions. In this paper we consider such asynchronous versions of arbitrary two-person strategic games and we study how the presence of the asynchronous adversary influences the behavior of the players, assumed to be selfish but rational. We concentrate on deterministic (pure) strategies, and in particular, on the existence and characteristics of pure Nash equilibria in such games. It turns out that the rational behavior of players changes significantly if the decision process is asynchronous. We show that pure Nash equilibria often exist in the asynchronous version of the game even if there were no such equilibria in the synchronous game. We also show that a mere threat of asynchrony in the game may make social optimum a rational choice while it was not rational in the synchronous game.

Keywords: Selfish agent, Nash equilibrium, asynchronous game

*This research was done during the stay of David Ilcinkas at the Université du Québec en Outaouais as a postdoctoral fellow.

†Research partially supported by NSERC grant and by the Research Chair in Distributed Computing at the Université du Québec en Outaouais.

1. Introduction

The behavior of rational selfish agents has been classically studied in the framework of strategic games. In such games each player has a set of possible actions. Each player chooses one of its actions and these choices are performed simultaneously. Each player chooses its action unaware of the choices of other players. To any set of such choices made by all players corresponds the gain of each of the players, specified by the rules of the game, known in advance to everybody. All players are selfish and rational, which means that the unique aim of each player is maximizing its own gain.

However, in many applications, players choose actions asynchronously, and simultaneity of this process is not guaranteed: it is possible that a player learns the action of another player before making its choice. Delays of choices are controlled by an adversary and it is assumed that they are arbitrarily long but finite. (We exclude the possibility of players waiting for each other forever.) The notion of an adversary models the worst-case attitude of each player: as usual in asynchronous distributed systems, an agent pessimistically expects that delays will be most detrimental for it. This guarantees that the outcome of the player's decisions will never be worse than it expected.

Consider the following example. Two stores sell the same product for the same price. Each store has two possible actions: keep the current price (K) or raise it (R). If both stores keep the price then each gets the profits from half of the customers at the current low price, which results in gain 1 for each. If one store raises the price and the other keeps it, all customers switch to the second store which results in gain 2 for this store (profit from all customers but at the current low price) and gain 0 for the first. Finally, if both raise the price then each enjoys the gain from half of the customers but at the higher price: the highest gain 3 for both (if the raise is substantial compared to the current price). This game can be represented by the following matrix in which the first entry in each cell is the gain of the player choosing rows and the second entry is the gain of the player choosing columns (Fig. 1).

	R	K
R	3;3	0;2
K	2;0	1;1

Figure 1. The price war game

This formalism corresponds to stores making their choices simultaneously. However, in reality, as prices are announced in an asynchronous way, it may happen that the price announced by one store is made public before the other store made the decision. Then the decision of the first store may well influence the decision of the second. Indeed, in our example, if the first store raised the price, the other is better off raising it as well and if the first kept the old price, the other is better off keeping it as well. Hence we need a new formalism for such asynchronous games derived from classical strategic games. The latter will be called *synchronous* games from now on.

In the asynchronous version of a two-person strategic game, the strategy of a player α is not a single action anymore but a pair (A, f) where A is one of the actions available to the player (it will be called the *basic action* of the strategy) and f is a function which maps the set of actions of the other player to the set of actions of player α . This strategy has the following meaning. If player α makes its decision first (more precisely, without seeing the action of the other player), it chooses action A . Otherwise,

if it learns that the other player chose action X , it responds by the action $f(X)$. This way, if in the underlying synchronous strategic game G player α had n possible actions and player β had k possible actions, then in the derived asynchronous game, called the *asynchronization* of G , player α has n^{k+1} possible strategies and player β has k^{n+1} possible strategies. The gain of each player for a given pair of strategies is determined in the worst-case manner, i.e., assuming that the asynchronous adversary acts in the way most detrimental for it. Consider the example of the two following strategies for the players in the asynchronous version of the game from Fig. 1. (Note that the matrix of the game is symmetric, i.e., roles of both stores are the same). Store α chooses the following strategy: perform K as the basic action (i.e., if it makes the decision without seeing the action of the other store), otherwise repeat the action of the other store (K after seeing K and R after seeing R). Store β chooses action R as the basic action and otherwise also repeats the action of the other store. Assuming these strategies, the gain of both stores in the asynchronous game is computed as follows. Store α expects the minimum of three numbers: 2 (the case when actions are simultaneous), 1 (the case when α chooses its action K first and β repeats it, and 3 (the case when β chooses its action R first and α repeats it). Hence the gain of α is 1. Similarly the gain of β is the minimum of the numbers 0,3,1, i.e., it is 0.

Thus, for any two-person strategic (synchronous) game given by a matrix of type as in Fig. 1, the asynchronization of this game is represented by a (much larger) matrix informally described above (see Section 2 for the precise definition). The asynchronization of a synchronous game can again be considered as a strategic game with strategies replacing actions and players' gains computed as above. Hence the usual tools of game theory can be applied to the asynchronous game, giving insight concerning the behavior of rational selfish players in the (often more realistic) asynchronous environment.

The aim of this paper is to study how the presence of the asynchronous adversary influences the behavior of the players, i.e., how their decisions change when the synchronous game is replaced by its asynchronization. It should be stressed that the novelty of our approach consists in introducing asynchrony in a single (one-shot) game, as opposed to asynchronous behavior in multiple plays of a game repeated over time, that was previously studied in the literature (cf. subsection 1.2). We concentrate on deterministic (pure) strategies, and in particular, on the existence and characteristics of pure Nash equilibria in asynchronous games. (Intuitively, a pure Nash equilibrium is a pair of actions (or, in the asynchronous case, strategies) of both players, such that none of them has a reason to unilaterally change its decision). It turns out that the rational behavior of selfish players changes significantly if the decision process is asynchronous.

1.1. Our results

We first show that if a synchronous strategic game has a pure Nash equilibrium then its asynchronization also has a pure Nash equilibrium, in fact it must have many such equilibria. The converse turns out to be false: the asynchronization process may create pure Nash equilibria even when no such equilibria existed in the underlying synchronous game. Indeed, we prove that the asynchronization of every strictly competitive game¹ has a pure Nash equilibrium, and the asynchronization of every strategic game in which each player has only two decisions, also always has a pure Nash equilibrium. This should be contrasted with the well-known fact that there are many (synchronous) strictly competitive games (even

¹Intuitively, a game is strictly competitive when players have opposite interests. The example is a zerosum game in which one player always wins what the other loses. See Section 2 for precise definitions.

zerosum games), and also synchronous games with two decisions per player, that do not have any pure Nash equilibria.

The possibility of creating new pure Nash equilibria by asynchronization can be used to entice rational selfish players to socially desirable behavior. We show synchronous games in which the social optimum (a pair of actions for which the sum of the players' gains is maximized) is not a rational choice for the players but in the asynchronization of this game a pair of strategies with this pair of basic actions becomes a pure Nash equilibrium. Hence we obtain a surprising result that the mere threat of asynchrony in the decision process may make the social optimum a rational choice, while it was not a rational choice in the synchronous decision process. This is reminiscent of a recent result from [9] that another uncontrollable feature of the environment, namely the threat of the presence of Byzantine agents, may entice rational selfish players to socially desirable behavior.

We finally show a somewhat paradoxical result. Consider a strategy in the asynchronous game which involves a strictly suboptimal answer in response to some action of the other player. We call such a strategy *exotic*. An exotic strategy does not seem to be a rational choice for a player, as replacing this suboptimal answer by an optimal one, the player never loses and can sometimes increase its gain. Nevertheless we show that a pair of exotic strategies can form a pure Nash equilibrium, even in asynchronizations of some zerosum games. We also show that this type of anomaly cannot be pushed to the extreme, at least for strictly competitive games. Indeed, exotic strategies, in which answers to all actions of the other player are worst possible cannot form a pure Nash equilibrium in the asynchronization of a strictly competitive game.

1.2. Related work

The behavior of rational selfish agents has been extensively studied in game theory, in fact it is one of the main topics of this theory [11, 12]. However strategic games modeling this behavior do not take into account the possibility of asynchrony in a single (one-shot) game. M.J. Osborne [11] says explicitly:

Time is absent from the model. The idea is that each player chooses its action once and for all, and the players choose their actions “simultaneously” in the sense that no player is informed, when it chooses its action, of the action chosen by any other player. (For this reason, a strategic game is sometimes referred to as a “simultaneous-move game”.)

Games in which players take turns in making moves have also been extensively studied since the seminal paper [14], where they were introduced, under the name of extensive games with perfect information. Such games (including important examples like chess, checkers, go, or Othello) are modeled by trees rather than by matrices. However, again, asynchrony in a single play is not modeled there: the order of moves is strictly predetermined, although in some situations this order may change over time in repeated plays of the same game (cf. e.g., [8]). Asynchrony in repeated plays of a game was also investigated, e.g., in [4, 5, 6], in the context of learning algorithms, where players change their strategies asynchronously over time, based on previously acquired experience. Again, the asynchronous behavior of players does not concern a single play: they choose their current strategies simultaneously.

On the other hand, in the parallel and distributed computing community, asynchrony is one of the main topics of study [2], as most of the real distributed systems are at least partly asynchronous. An important source of asynchrony are delays of communication between processors (agents), similarly as in our example of competing stores. With the advent of internet, a lot of economic activity (among often

selfish agents) is conducted via this intrinsically asynchronous medium, whence the natural combination of rational selfish agents' behavior with the asynchronous environment. However, while various algorithmic issues coming from game theory, also those concerning the behavior of selfish agents linked by communication networks (cf. e.g., [1, 3, 7, 9]), have been investigated in the algorithmic community, particularly after the seminal paper of Papadimitriou [13], the impact of asynchrony on the behavior of such agents, in the sense of *uncertainty* concerning the time of choosing strategies by players in a single game, has apparently not been studied. This paper is meant as a first step in this direction.

2. Terminology and notation

We consider two-person strategic games. In such a game each of two players α and β has a set of possible actions $\mathcal{A}_\alpha = \{A_1, \dots, A_n\}$ and $\mathcal{A}_\beta = \{B_1, \dots, B_k\}$, respectively. The gains of the players are given by $n \times k$ matrices $[a_{ij}]$ and $[b_{ij}]$, where a_{ij} (resp. b_{ij}) is the gain of player α (resp. β) when player α chooses action A_i and player β chooses action B_j . For convenience, both matrices are represented together in one table as shown in the following example, in which each player has two possible actions:

	A_1	A_2
B_1	$a_{11}; b_{11}$	$a_{12}; b_{12}$
B_2	$a_{21}; b_{21}$	$a_{22}; b_{22}$

Figure 2. A 2×2 strategic game

An action A_i of player α *strictly dominates* its action A_r if $a_{ij} > a_{rj}$, for all $j \leq k$. Strict domination for player β is defined similarly. A pair of actions A_i of player α and B_j of player β is a *pure Nash equilibrium* (PNE), if $a_{ij} \geq a_{rj}$ for all $r \leq n$ and $b_{ij} \geq b_{is}$ for all $s \leq k$.

A game is *strictly competitive* if, for any indices i, j, i', j' , the inequality $a_{ij} \geq a_{i'j'}$ implies $b_{ij} \leq b_{i'j'}$. An important subclass of strictly competitive games are *zerosum games* that satisfy $a_{ij} = -b_{ij}$, for all i and j . Such games can be interpreted as requiring, for any pair of actions, that any gain of one player is equal to the loss of the other player.

Consider a two-person strategic game as above, where actions of players are assumed to be chosen simultaneously. We call such a game a *synchronous game*. Given such a game G we define its *asynchronization*. This is another strategic game G^* defined as follows. Actions of player α in this game (called *strategies* to avoid confusion with actions of game G) are pairs (A_i, f) , where $A_i \in \mathcal{A}_\alpha$ and $f : \mathcal{A}_\beta \rightarrow \mathcal{A}_\alpha$. Action A_i in strategy (A_i, f) is called the *basic action* of this strategy. This is the action player α chooses if it didn't previously learn the action of player β . On the other hand, the function f defines the answer of player α to every action of player β . Similarly, strategies of player β are pairs (B_j, g) , where $B_j \in \mathcal{A}_\beta$ and $g : \mathcal{A}_\alpha \rightarrow \mathcal{A}_\beta$. For convenience, sets of actions $\{A_1, \dots, A_n\}$ and $\{B_1, \dots, B_k\}$ of both players in game G are considered to be ordered and thus strategies (A_i, f) and (B_j, g) in G^* can be noted as the following sequences of actions in game G : $(A_i, f(B_1), f(B_2), \dots, f(B_k))$ and $(B_j, g(A_1), g(A_2), \dots, g(A_n))$, respectively. Also, by a slight abuse of notation, we will treat functions f and g as having indices of actions, rather than actions themselves, as arguments and values, thus writing, e.g., $f(j) = i$ instead of $f(B_j) = A_i$. For fixed orders of actions of both players this abuse does not

yield any ambiguities. The gain of player α corresponding to strategies (A_i, f) and (B_j, g) is defined as the number $\Gamma_\alpha((A_i, f), (B_j, g)) = \min\{a_{ij}, a_{f(j)j}, a_{ig(i)}\}$. Similarly, the gain of player β corresponding to strategies (A_i, f) and (B_j, g) is defined as the number $\Gamma_\beta((A_i, f), (B_j, g)) = \min\{b_{ij}, b_{f(j)j}, b_{ig(i)}\}$. The reader is invited to check that these definitions correspond to those intuitively explained in the example given in the introduction. The table in Fig. 3 gives the asynchronization of the synchronous price war game described in Fig. 1.

	<i>RRR</i>	<i>RRK</i>	<i>RKK</i>	<i>RKR</i>	<i>KKR</i>	<i>KKK</i>	<i>KRK</i>	<i>KRR</i>
<i>RRR</i>	3;3	3;3	0;2	0;2	0;2	0;2	0;2	0;2
<i>RRK</i>	3;3	3;3	0;2	0;2	0;1	0;1	0;1	0;1
<i>RKK</i>	2;0	2;0	0;0	0;0	0;1	0;1	0;1	0;1
<i>RKR</i>	2;0	2;0	0;0	0;0	0;2	0;2	0;2	0;2
<i>KKR</i>	2;0	1;0	1;0	2;0	0;0	0;1	0;1	0;0
<i>KKK</i>	2;0	1;0	1;0	2;0	1;0	1;1	1;1	1;0
<i>KRK</i>	2;0	1;0	1;0	2;0	1;0	1;1	1;1	1;0
<i>KRR</i>	2;0	1;0	1;0	2;0	0;0	0;1	0;1	0;0

Figure 3. The asynchronization of the price war game

Strategic games that are asynchronizations of some synchronous games are called asynchronous games. If a game G^* is the asynchronization of a synchronous game G , the game G is called the underlying synchronous game for G^* .

3. Existence of pure Nash equilibria in asynchronous games

We first prove the following basic result.

Theorem 3.1. If a synchronous game G has a PNE then its asynchronization G^* has a PNE as well.

Proof:

Consider a synchronous game G with action sets $\mathcal{A}_\alpha = \{A_1, \dots, A_n\}$ and $\mathcal{A}_\beta = \{B_1, \dots, B_k\}$ and player gains given by matrices $[a_{ij}]$ and $[b_{ij}]$. Let (A_i, B_j) be a PNE in G . Consider any of the following strategies in the asynchronization G^* of G : a strategy (A_i, f) of player α , such that $f(B_j) = A_i$ and a strategy (B_j, g) of player β , such that $g(A_i) = B_j$. (Notice that there are $n^{k-1} \cdot k^{n-1}$ such pairs of strategies.) We prove that any such pair of strategies is a PNE in the asynchronous game G^* . By definition of the gain Γ_α in G^* we have $\Gamma_\alpha((A_i, f), (B_j, g)) = a_{ij}$ and $\Gamma_\alpha((A_r, f'), (B_j, g)) \leq a_{rj}$, for any strategy (A_r, f') of player α in game G^* . Since (A_i, B_j) is a PNE in G , we have $a_{rj} \leq a_{ij}$. Hence $\Gamma_\alpha((A_r, f'), (B_j, g)) \leq \Gamma_\alpha((A_i, f), (B_j, g))$. A similar argument shows that $\Gamma_\beta((A_i, f), (B_s, g')) \leq \Gamma_\beta((A_i, f), (B_j, g))$, for any strategy (B_s, g') of player β in game G^* . Hence the pair of strategies $((A_i, f), (B_j, g))$ is a PNE in the asynchronous game G^* . \square

The above result shows that the existence of a PNE is preserved in the process of asynchronization, in fact a PNE in the synchronous game yields many PNE in its asynchronization. How about the converse of the above proposition? It turns out that it is not true: asynchronizations of many games without PNE may have PNE.

First consider strictly competitive games. It is well known that many such games do not have PNE. Examples can be easily found even among zerosum games. Any game

	B_1	B_2
A_1	$a; -a$	$b; -b$
A_2	$c; -c$	$d; -d$

Figure 4. A zerosum game

with $a < b, c$ and $d < b, c$, does not have a PNE. Nevertheless we have the following result.

Theorem 3.2. The asynchronization of every strictly competitive synchronous game has a PNE.

Proof:

Consider a strictly competitive synchronous game G with action sets $\mathcal{A}_\alpha = \{A_1, \dots, A_n\}$ and $\mathcal{A}_\beta = \{B_1, \dots, B_k\}$ and player gains given by matrices $[a_{ij}]$ and $[b_{ij}]$. Take an action A_i of player α with the property

$$\min_{j \leq k} a_{ij} = \max_{r \leq n} \min_{j \leq k} a_{rj}.$$

Similarly, take an action B_j of player β with the property

$$\min_{i \leq n} b_{ij} = \max_{s \leq k} \min_{i \leq n} b_{is}.$$

Consider the following strategies in the asynchronization G^* of G :

- the strategy (A_i, f) of player α , where f optimizes the response of player α to all actions of player β : formally, for any $s \leq k$ we require $a_{f(s)s} = \max_{r \leq n} a_{rs}$;
- the strategy (B_j, g) of player β , where g optimizes the response of player β to all actions of player α : formally, for any $r \leq n$ we require $b_{rg(r)} = \max_{s \leq k} b_{rs}$.

We prove that the pair of these strategies is a PNE in the game G^* . By definition of the gain Γ_α in G^* we have $\Gamma_\alpha((A_i, f), (B_j, g)) = \min\{a_{ij}, a_{f(j)j}, a_{ig(i)}\}$. Since f optimizes the response of player α to action B_j , we have $a_{ij} \leq a_{f(j)j}$. Since g optimizes the response of player β to action A_i , we have $b_{ij} \leq b_{ig(i)}$. By the property of strict competitiveness of game G we have $a_{ij} \geq a_{ig(i)}$. Hence $\Gamma_\alpha((A_i, f), (B_j, g)) = a_{ig(i)}$.

On the other hand, for any strategy (A_r, f') of player α , we have $\Gamma_\alpha((A_r, f'), (B_j, g)) \leq a_{rg(r)}$, by the definition of Γ_α . By the definition of action A_i we have $a_{rg(r)} \leq a_{ig(i)}$. Consequently, we get $\Gamma_\alpha((A_r, f'), (B_j, g)) \leq \Gamma_\alpha((A_i, f), (B_j, g))$. A similar argument shows that, for any strategy (B_s, g') of player β we have $\Gamma_\beta((A_i, f), (B_s, g')) \leq \Gamma_\beta((A_i, f), (B_j, g))$. This proves that the pair of strategies (A_i, f) and (B_j, g) is a PNE in the game G^* . \square

We do not know if the above result generalizes to arbitrary synchronous games. However, it remains true for 2×2 synchronous games, i.e., games in which each of the two players has two possible actions. As we have seen, there are 2×2 synchronous games without a PNE.

Theorem 3.3. The asynchronization of every 2×2 synchronous game has a PNE.

Proof:

Consider a 2×2 synchronous game G , as given in Fig. 2. If G has a PNE, we are done by Theorem 3.1. Suppose that G does not have a PNE. Consider the gains of the player α . The maxima of the first and the second column cannot be in the same row, as the strategy corresponding to this row would then be part of a PNE. Without loss of generality, the maximum of the first column is in the second row and the maximum of the second column is in the first row. Now consider the gains of player β . The maximum of the first row must be in the first column, for otherwise the pair of actions (A_1, B_2) would be a PNE. The maximum of the second row must be in the second column, for otherwise the pair of actions (A_2, B_1) would be a PNE. Thus the general form of the game G is as shown in Fig. 5, where \bar{z} denotes a number greater or equal to z .

	B_1	B_2
A_1	$x; \bar{a}$	$\bar{y}; a$
A_2	$\bar{x}; b$	$y; \bar{b}$

Figure 5. A general form of a 2×2 game without PNE

Now consider the matrix of the asynchronization G^* of G , see Fig. 6. We only show a part of its entries, those that are necessary for the proof. (The reader is invited to verify that they correspond to the definition). Stars and blanks denote entries not calculated, and $m\{p, q\}$ stands for $\min\{p, q\}$. We claim that

	$B_1B_1B_1$	$B_1B_1B_2$	$B_1B_2B_2$	$B_1B_2B_1$	$B_2B_2B_1$	$B_2B_2B_2$	$B_2B_1B_2$	$B_2B_1B_1$
$A_1A_1A_1$		$x;*$					$m\{x, \bar{y}\};*$	
$A_1A_1A_2$		$x;*$					$m\{x, y\};*$	
$A_1A_2A_2$		$x;*$					$m\{x, y\};*$	
$A_1A_2A_1$	$*;m\{\bar{a}, b\}$	$x;m\{\bar{a}, b\}$	$*;m\{a, b\}$	$*;m\{a, b\}$	$*;a$	$*;a$	$m\{x, \bar{y}\};a$	$*;a$
$A_2A_2A_1$	$*;b$	$m\{\bar{x}, y\};b$	$*;b$	$*;b$	$*;m\{a, b\}$	$*;m\{a, \bar{b}\}$	$y;m\{a, \bar{b}\}$	$*;m\{a, b\}$
$A_2A_2A_2$		$m\{\bar{x}, y\};*$					$y;*$	
$A_2A_1A_2$		$m\{x, y\};*$					$y;*$	
$A_2A_1A_1$		$m\{x, y\};*$					$y;*$	

Figure 6. The game G^*

in every case, one of the four strategy pairs: $((A_1, A_2, A_1), (B_1, B_1, B_2))$, $((A_1, A_2, A_1), (B_2, B_1, B_2))$, $((A_2, A_2, A_1), (B_1, B_1, B_2))$, $((A_2, A_2, A_1), (B_2, B_1, B_2))$, is a PNE. Indeed, by the definition, the gains of the players corresponding to these pairs of strategies are as shown in Fig. 6. Hence, if $a \leq b$ then $((A_1, A_2, A_1), (B_1, B_1, B_2))$ or $((A_2, A_2, A_1), (B_1, B_1, B_2))$ must be a PNE, and if $b < a$ then $((A_1, A_2, A_1), (B_2, B_1, B_2))$ or $((A_2, A_2, A_1), (B_2, B_1, B_2))$ must be a PNE. \square

Theorems 3.2 and 3.3 show that a PNE may appear in the asynchronization of a synchronous game that does not have a PNE. On the other hand, there are examples of synchronous games G that have a PNE but the asynchronization of G has a PNE whose basic actions were not a PNE in G . This gives the

possibility to use asynchrony (or even only the threat of it) to entice some socially desirable behavior of players. Consider the following game

	B_1	B_2
A_1	2; 3	0; 0
A_2	3; 1	1; 2

Figure 7. Game H whose social optimum is not a PNE

It has one pair of actions (A_2, B_2) that form a PNE, however the social optimum of this game is the pair (A_1, B_1) which is not a PNE. Player α has an incentive to unilaterally deviate from this pair of actions. (In fact, the action A_1 is strictly dominated by A_2 .) On the other hand, the pair (A_1, B_1) not only maximizes the sum of gains of the players but in fact *each of the players* has larger gain when choosing this pair than if they choose the pair (A_2, B_2) which is the unique PNE. (In this respect the situation is similar to the game in the Prisoner's Dilemma). So it is particularly "regretful" that players cannot "agree" on the pair (A_1, B_1) .

Now consider the asynchronization of this game (see Fig. 8). (Again we show only entries necessary for the argument). The pair of strategies $((A_1, A_2, A_2), (B_1, B_1, B_2))$ is a PNE in the game H^* .

	$B_1 B_1 B_1$	$B_1 B_1 B_2$	$B_1 B_2 B_2$	$B_1 B_2 B_1$	$B_2 B_2 B_1$	$B_2 B_2 B_2$	$B_2 B_1 B_2$	$B_2 B_1 B_1$
$A_1 A_1 A_1$		2;*						
$A_1 A_1 A_2$		2;*						
$A_1 A_2 A_2$	*;1	2;1	*;0	*;0	*;0	*;0	*;0	*;0
$A_1 A_2 A_1$		2;*						
$A_2 A_2 A_1$		1;*						
$A_2 A_2 A_2$		1;*						
$A_2 A_1 A_2$		1;*						
$A_2 A_1 A_1$		1;*						

Figure 8. Game H^*

This pair is of the form $((A_1, f), (B_1, g))$, where functions f and g give optimal responses to the other player's actions. Now suppose that players *think* that they are playing the asynchronous game H^* , e.g., because they communicate via a potentially asynchronous channel. Then none of them has an incentive to unilaterally deviate from the pair $((A_1, A_2, A_2), (B_1, B_1, B_2))$, and playing these strategies becomes a rational choice. Fix this choice of strategies for both players in the game H^* . Now suppose that the adversary decides to deliver the actions of both players in a synchronous way, i.e., forbidding each of them to see the other's action before deciding its own. For this adversary, players in fact (unknowingly) play the synchronous game H , and in this game they play the pair of basic actions of their strategies from H^* , i.e., the pair (A_1, B_1) . As we know, this pair is a social optimum in the game H (and, moreover increases the gain of each player with respect to the PNE (A_2, B_2)). So, in this case, the mere threat of asynchrony made rational choosing the social optimum, while this social optimum was not rational without this threat. In [9] a somewhat similar phenomenon was discovered: the threat of the presence of Byzantine players enticed rational selfish players to make socially desirable choices by reducing the

price of anarchy. In both cases, the threat of some feature of the environment (Byzantine players in [9], asynchrony in our case) may entice rational selfish players to better collaboration.

4. Exotic strategies in asynchronous games

Strategies of players in asynchronous games discussed so far were of the form (A_i, f) or (B_j, g) , where functions f and g give (one of the possibly many) optimal responses to the other player's actions. How about other strategies? Consider a synchronous game G with action sets $\mathcal{A}_\alpha = \{A_1, \dots, A_n\}$ and $\mathcal{A}_\beta = \{B_1, \dots, B_k\}$ and player gains given by matrices $[a_{ij}]$ and $[b_{ij}]$. A strategy (A_i, f) of player α is called *exotic* if, for some $s \leq k$, we have $a_{f(s)s} < \max_{r \leq n} a_{rs}$; similarly for player β . An exotic strategy does not seem to be a rational plan of playing, as a player may clearly improve it by changing the function f to the function f' such that, for all $s \leq k$, the equality $a_{f'(s)s} = \max_{r \leq n} a_{rs}$ holds. So it is somewhat surprising that, as we will show, a pair of exotic strategies may be a PNE in an asynchronous game. Not only that. In fact we present an example of a synchronous game G whose asynchronization G^* has a PNE formed by exotic strategies (A_i, f) and (B_j, g) , where $f(B_j)$ is *not* an optimal response to action B_j and $g(A_i)$ is *not* an optimal response to action A_i . So the exotic character of these strategies is not due to bad responses to some actions that “will not happen”: indeed they involve bad responses to the basic action of the other player, and still these strategies form a PNE in G^* .

Consider the very simple zerosum game G in which each player can choose heads or tails. Each of them puts a coin with the chosen side up on the table. If sides are different, the player having chosen heads pays \$1 to the other player. If sides are the same, nobody pays. The matrix of the game is shown in Fig. 9, where A_1 and B_1 are actions of choosing heads and A_2 and B_2 are actions of choosing tails.

	B_1	B_2
A_1	0; 0	-1; 1
A_2	1; -1	0; 0

Figure 9. The heads-tails game

In this game actions A_2 and B_2 are strictly dominating and they form the unique PNE. Now consider the game G^* (see Fig. 10).

	$B_1 B_1 B_1$	$B_1 B_1 B_2$	$B_1 B_2 B_2$	$B_1 B_2 B_1$	$B_2 B_2 B_1$	$B_2 B_2 B_2$	$B_2 B_1 B_2$	$B_2 B_1 B_1$
$A_1 A_1 A_1$	0;0	0;0	-1;0	-1;0	-1;1	-1;1	-1;0	-1;0
$A_1 A_1 A_2$	0;0	●0;0	-1;0	-1;0	-1;0	-1;0	-1;0	-1;0
$A_1 A_2 A_2$	0;-1	0;-1	-1;-1	-1;-1	-1;0	-1;0	-1;0	-1;0
$A_1 A_2 A_1$	0;-1	0;-1	-1;-1	-1;-1	-1;1	-1;1	-1;0	-1;0
$A_2 A_2 A_1$	1;-1	0;-1	0;-1	1;-1	-1;-1	-1;0	-1;0	-1;-1
$A_2 A_2 A_2$	1;-1	0;-1	0;-1	1;-1	0;-1	●0;0	●0;0	0;-1
$A_2 A_1 A_2$	0;-1	0;-1	0;-1	0;-1	0;-1	●0;0	●0;0	0;-1
$A_2 A_1 A_1$	0;-1	0;-1	0;-1	0;-1	-1;-1	-1;0	-1;0	-1;-1

Figure 10. The asynchronization of the heads-tails game

The five PNE in this game are marked with a bullet. The four PNE with basic actions A_2, B_2 are those yielded by the unique PNE in game G , see Theorem 3.1. Consider the PNE $((A_1, A_1, A_2), (B_1, B_1, B_2))$. The strategies that form this PNE prescribe the following behavior of both players: choose “heads” as the basic action, respond “heads” if you see “heads” and respond “tails” if you see “tails”. Thus they have basic actions that are strictly dominated in the synchronous game, and they involve suboptimal responses to the other player’s basic action. Curiously, “improving” the response part by changing the suboptimal response to the other player’s basic action to the optimal response “tails” (by one or both players) does not lead to a PNE.

Hence we are confronted with a situation in which some strategies are intuitively irrational (changing an exotic strategy by switching all suboptimal responses to optimal ones can never hurt a player and sometimes can increase its gain) and yet they form a PNE, so they seem to be a rational choice. This paradox does not occur in synchronous games. What is its solution in the asynchronous setting? This is not so much a formal problem as an issue concerning intuitions of the notion of rationality, similar to the on-going discussion in game theory on how well the notion of Nash equilibrium captures the rational behavior of players. As for the above mentioned paradox in the asynchronous setting, our feeling is that rejection of exotic strategies as irrational should prevail and the possibility that they form a PNE should be considered as a peculiarity of formalism, similarly, e.g., to the existence of Lebesgue non-measurable sets. Thus we would like to advise the players: use only non-exotic strategies, and among those use ones that can give a PNE. Applied to our example, this advice gives the pair of strategies $((A_2, A_2, A_2), (B_2, B_2, B_2))$, i.e., the natural suggestion: play “tails”, no matter what happened.

Nevertheless it is natural to ask how far can this paradox go. We show that, at least for strictly competitive games, it cannot be pushed to the extreme. Consider a synchronous game G . We say that the behavior of player α is *oblivious* with respect to action B_j of player β , if $a_{ij} = a_{i'j}$, for all $i, i' \leq n$. Similarly, the behavior of player β is oblivious with respect to action A_i of player α if $b_{ij} = b_{ij'}$, for all $j, j' \leq k$. A pair of actions (A_i, B_j) is non-oblivious, if either the behavior of player α is non-oblivious with respect to B_j or the behavior of player β is non-oblivious with respect to A_i . A strategy (A_i, f) of player α in an asynchronous game G^* is a *worst-reply strategy* based on A_i , if $a_{f(j)j} = \min_{i \leq n} a_{ij}$, for all $j \leq k$. A worst-reply strategy for player β is defined similarly.

Theorem 4.1. Let G be a strictly competitive game. Let (A_i, B_j) be a non-oblivious pair of actions in G . Then a pair of worst-reply strategies based on these actions cannot be a PNE in the asynchronous game G^* .

Proof:

Consider a strictly competitive synchronous game G with action sets $\mathcal{A}_\alpha = \{A_1, \dots, A_n\}$ and $\mathcal{A}_\beta = \{B_1, \dots, B_k\}$ and player gains given by matrices $[a_{ij}]$ and $[b_{ij}]$. The theorem follows from the following claim.

Claim. Let (A_i, B_j) be any pair of actions in G . Let A_p be one of the worst responses to B_j for player α , i.e., $a_{pj} = \min_{r \leq n} a_{rj}$. Let $B_{p'}$ be one of the worst responses to A_i for player β . Let A_o be one of the best responses to B_j for player α , i.e., $a_{oj} = \max_{r \leq n} a_{rj}$. Let B_q be an answer to A_o worse or equal to B_j for player β , i.e., $b_{oq} \leq b_{oj}$. Consider the strategies $(A_i, f), (B_j, g)$ in game G^* , for which $f(j) = p, g(i) = p'$ and $g(o) = q$. If the pair $(A_i, f), (B_j, g)$ is a PNE in the game G^* then $a_{rj} = a_{r'j}$, for all $r, r' \leq n$.

We now prove the claim. Let $f' : \mathcal{A}_\beta \rightarrow \mathcal{A}_\alpha$ be such that $f'(j) = o$. Let b denote the gain of player α in the game G^* for the pair of strategies $((A_i, f), (B_j, g))$. Let c denote the gain of player

α in the game G^* for the pair of strategies $((A_o, f'), (B_j, g))$. By definition of the game G^* we have $b = \min\{a_{ij}, a_{ip'}, a_{pj}\}$ and $c = \min\{a_{oj}, a_{oq}\}$. By definition of A_p we have $a_{pj} \leq a_{ij}$. By definition of $B_{p'}$ we have $b_{ip'} \leq b_{ij}$, hence by strict competitiveness of game G we have $a_{ip'} \geq a_{ij}$. It follows that $b = a_{pj}$. By definition of B_q we have $b_{oq} \leq b_{oj}$. Hence by strict competitiveness of game G we have $a_{oq} \geq a_{oj}$. Hence $c = a_{oj}$. Suppose that the pair $(A_i, f), (B_j, g)$ is a PNE in the game G^* . Then $c \leq b$. It follows that $a_{rj} = a_{r'j}$, for all $r, r' \leq n$. This proves the claim. Similarly we prove that $b_{is} = b_{is'}$, for all $s, s' \leq k$, and the theorem follows. \square

5. Conclusion

We introduced the concept of asynchronization of two-person strategic games, a tool to deal with asynchrony in the behavior of rational selfish agents. Our study was focused on the existence of pure Nash equilibria in asynchronous strategic games. We showed that pure Nash equilibria exist in the asynchronization of a game if they exist in the game itself, and that they may also exist in the asynchronization of a game that does not have such equilibria itself. Moreover we established the existence of such equilibria for important classes of asynchronous games. Our results lead to the following conjecture.

Conjecture. Every two-person asynchronous game has a pure Nash equilibrium.

Moreover, we conjecture that such an equilibrium can always be found among pairs of non-exotic strategies. This was the case for all games that we considered.

Our notion of asynchronization of a strategic game can be naturally extended to n -person games. In this case the possibilities that a player should consider is learning actions of some arbitrary (possibly empty) subset of the other players and deciding its own action on this basis. This significantly increases the number of possible strategies, however it does not necessarily make finding pure Nash equilibria a non-tractable task: we have seen in Theorems 3.1 and 3.2 that PNE can be sometimes found quite fast in spite of the fact that the size of the asynchronization of a game is exponential in the size of the underlying synchronous game, already for two players. This more general notion of asynchronization of n -person strategic games could be used to model the behavior of rational selfish agents acting in such intrinsically asynchronous economic processes as, e.g., internet auctions. Another interesting generalization in the context of asynchrony is to consider mixed actions and strategies both in the underlying synchronous game and in its asynchronization, i.e., to consider probabilistic instead of deterministic decisions.

References

- [1] Albers S., Eilits S., Even-Dar E., Mansour Y., Roditty L.: On Nash equilibria for a network creation game, *Proc. 17th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2006.
- [2] Lynch N.: *Distributed Algorithms*, Morgan Kaufmann Publ., Inc., 1996.
- [3] Czumaj A., Vöcking B.: Tight bounds for worst-case equilibria. *Proc. 13th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2002, 413-420.
- [4] Friedman E.: Asynchronous Learning in Decentralized Environments: A Game Theoretic Approach (2004), in *Collectives and the Design of Complex Systems*, edited by K. Tumer and D. Wolpert, Springer-Verlag
- [5] Friedman E., Shenker S.: *Synchronous and asynchronous learning by responsive learning automata*, mimeo, 1996.

- [6] Friedman E., Shenker S.: *Learning and Implementation on the Internet*, manuscript 1997.
- [7] Koutsoupias E., Papadimitriou C.: Worst-case equilibria. *Proc. 16th Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, LNCS 1563, 1999, 404-413.
- [8] Kun A., Boza G., Scheuring I.: Asynchronous snowdrift game with synergistic effect as a model of cooperation, *Behavioral Ecology*, **17**, 2006, 633-641.
- [9] Moscibroda T., Schmid S., Wattenhofer R.: When selfish meets evil: Byzantine players in a virus inoculation game, *Proc. 25th Annual ACM Symposium on Principles of Distributed Computing (PODC)*, 2006, 35-44.
- [10] Nash J.F., Jr.: Non-cooperative games, *Annals of Mathematics*, **54**, 1951, 286-295.
- [11] Osborne M.J.: *An introduction to game theory*, Oxford University Press, 2004.
- [12] Osborne M.J., Rubinstein A.: *A Course in Game Theory*, MIT Press, 2000.
- [13] Papadimitriou C.: Algorithms, games and the Internet, *Proc. 33rd ACM Symposium on Theory of Computing (STOC)*, 2001, 749-753.
- [14] von Neumann J., Morgenstern O.: *Theory of games and economic behavior*, Princeton University Press, 1944.