Doing it Now, Later, or Never

Doing It Now, Later, or Never

Kutay Cingiz* János Flesch† P. Jean-Jacques Herings‡
Arkadi Predtetchinski§

March 14, 2016

Abstract

We study centipede games played by an infinite sequence of players. Following the literature on time–inconsistent preferences, we distinguish two types of decision makers, naive and sophisticated, and the corresponding solution concepts, naive ϵ–equilibrium and sophisticated ϵ–equilibrium. We show the existence of both naive and sophisticated ϵ-equilibria for each positive ϵ. Under the assumption that the payoff functions are upper semicontinuous, we furthermore show that there exist both naive and sophisticated 0-equilibria in pure strategies. We also compare the probability to stop of a naive versus a sophisticated decision maker and show that a sophisticated decision maker stops earlier.

Keywords: centipede games; subgame perfect ϵ-equilibria; time–inconsistent preferences; upper semi-continuous functions; sophisticated players; naive players.

JEL Classification Numbers: C70, D11, D60, D74, D91.

*Department of Economics, Maastricht University, P.O. Box 616, 6200 MD, Maastricht, The Netherlands. K.Cingiz@maastrichtuniversity.nl

†Department of Quantitative Economics, Maastricht University, PO Box 616, 6200 MD Maastricht, The Netherlands. J.Flesch@maastrichtuniversity.nl.

‡Department of Economics, Maastricht University, P.O. Box 616, 6200 MD, Maastricht, The Netherlands. P.Herings@maastrichtuniversity.nl

§Department of Economics, Maastricht University, P.O. Box 616, 6200 MD, Maastricht, The Netherlands. A.Predtetchinski@maastrichtuniversity.nl

1 Introduction

We study centipede games played by an infinite sequence of players. Each player is active only once. The active player can choose either to stop the game or to continue. As soon as the active player chooses to stop, the game ends.

One of the main application areas of our model concerns the vast literature on decision making with time–inconsistent preferences. It is customary (Strotz (1955), Pollak (1968), Peleg and Yaari (1973), Goldman (1979)) to model a decision maker with time–inconsistent preferences as consisting of a sequence of multiple selves, where day t self makes a decision on behalf of the decision maker on day t. This leads to a game played by an infinite sequence of players.

In their well-known paper *Doing it now or later*, O'Donoghue and Rabin (1999) consider a decision maker who has to decide when to execute a certain task. One important feature of their model is an exogenous deadline: once the deadline is reached, the decision maker has no choice but to execute the task. This model can be seen as a finite centipede game and is a special case of our more general model. In our model, the decision maker may have the option to never quit, whence the title of our paper.

Following the literature on decision making with time–inconsistent preferences, we distinguish two types of decision makers, naive and sophisticated, and examine the corresponding two types of solution concepts.

A naive decision maker acts under the erroneous assumption that his current self controls all future decisions in the game. Thus a naive decision maker intends to follow a strategy that maximizes his payoff over the entire continuation game, but in reality he only carries out the first action. This happens because the strategy that is optimal for the current self need not be optimal for future selves. This behavior is captured by the concept of naive equilibrium.

A sophisticated decision maker, in contrast, is fully aware that his day t self only controls the decision on day t, and that the future selves have different preferences. Thus in a sophisticated equilibrium each self of the decision maker best responds to the strategies of the future selves. Hence a sophisticated equilibrium is essentially the subgame perfect equilibrium of the game played by the selves of the decision maker "against" each other.

An example given in Flesch, Kuipers, Mashiah-Yaakovi, Schoenmakers, Solan, and Vrieze (2010), discussed in detail in the following section, shows that in general a sophisticated equilibrium need not exist. This motivates us to consider more permissive solution concepts: naive ϵ–equilibrium and sophisticated ϵ–equilibrium. A naive ϵ-equilibrium is strategy profile with the property that every player's strategy can be supported with a belief that makes this strategy and belief combination ϵ-optimal. In a sophisticated ϵ–equilibrium each player is

assumed to play an ϵ-best response to the strategies of the subsequent players.

Our results are as follows. We show that for each $\epsilon > 0$ there exists both a naive ϵ–equilibrium and a sophisticated ϵ–equilibrium. These existence results rely on mixed strategies. If we assume that each player's payoff function is upper semicontinuous, then there exist both a naive 0-equilibrium in pure strategies and a sophisticated 0-equilibrium in pure strategies.

Herings and Rohde (2006) and Luttmer and Mariotti (2006) consider how time-inconsistent decision makers interact in a market environment and give sufficient conditions for equilibrium existence. Nevertheless, Gabrieli and Ghosal (2013) point out that under standard assumptions equilibria may fail to exist in such environments. The heart of the problem is the satiation of the induced preferences of sophisticated decision makers, and the examples of non-existence are robust. On the contrary, the equilibrium existence issues taken up in this paper are at the level of the individual decision maker rather than the interaction between decision makers and existence problems can be solved by notions of ϵ-equilibrium.

One of the key results in O'Donoghue and Rabin (1999) is that a sophisticated decision maker executes a task earlier than a naive decision maker. We provide a counterpart of this result in our setup. We show that for a given sophisticated ϵ-equilibrium there exists a naive ϵ-equilibrium with the probability of stopping not higher than in the given sophisticated ϵ-equilibrium. Conversely, given a naive ϵ-equilibrium there is a sophisticated ϵ-equilibrium with the probability of stopping not smaller than in the given naive ϵ-equilibrium.

Apart from the literature on time–inconsistent decision making, our results contribute to the literature on the existence of subgame perfect ϵ–equilibrium in perfect information games, see e.g. Flesch et al. (2010), Purves and Sudderth (2011), and De Pril, Flesch, Kuipers, Schoenmakers, and Vrieze (2014). For the most part, this literature focuses on games with finitely many players. In contrast, here we consider a class of games played by infinitely many players.

Related to the infinite centipede games as considered here are so–called stopping games, see Solan (2005) and Mashiah-Yaakovi (2009). These are dynamic games where at each period of time each player can choose to stop or to continue. Our work is also related to intergenerational games, where there is a sequence of players such that each player represents an entire generation, see Phelps and Polak (1968) and Balbus, Jaśkiewicz, and Nowak (2015).

The paper is organized as follows. In Section 2, we discuss the so–called procrastination game. The game serves to illustrate some of the non–trivial aspects of our analysis and to motivate the need for the solution concepts of naive ϵ–equilibrium and sophisticated ϵ–equilibrium. In Section 3, we introduce the general model and define naive and sophisticated ϵ–equilibria. In Section 4, we focus on a special class of games in which the payoffs are

upper semi-continuous, and show the existence of naive 0-equilibrium and sophisticated 0-equilibrium in pure strategies. In Section 5, we examine the existence of naive ϵ-equilibrium and, in Section 6, the existence of sophisticated ϵ-equilibrium. In Section 7, we compare the stopping probabilities of naive and sophisticated decision makers and show that sophisticated decision makers stop earlier.

2 The Procrastination Game

Consider a decision maker who contemplates quitting smoking. On any given day the decision maker prefers quitting tomorrow to quitting today, and prefers quitting today to never quitting. This is an example of a decision maker with time–inconsistent preferences: quitting on day 2 is the best option from the perspective of day 1, but it is no longer the best option once it is considered on day 2 itself.

Following the standard approach to modeling time–inconsistent preferences, we represent the decision maker by a sequence of different selves, where day t self makes a decision on behalf of the decision maker on day t. This leads us to the following game tree, where S (stop) stands for quitting smoking and C (continue) represents the option to postpone quitting:

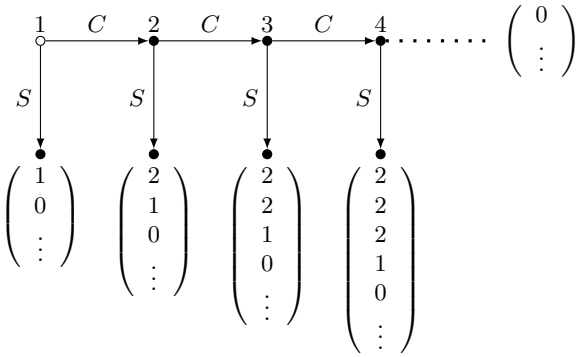

Figure 1: Procrastination game.

For the sake of concreteness we choose the following numerical values for the payoffs: the day t self of the decision maker obtains a payoff of 1 if the decision maker quits on day t, a payoff of 2 if the the decision maker quits on any day $k > t$, and 0 in all other situations, so in particular if the decision maker never quits. We refer to this situation as the Procrastination game. O'Donoghue and Rabin (1999) consider a decision maker who has to decide when to execute a task such as quitting smoking. One important feature of their model is an exogenous deadline: once the deadline is reached, the decision maker has no choice but to execute the task. In contrast, the decision maker acting in the procrastination game above has the option to never quit. Our general model as detailed in the following section captures both cases with

and without an exogenous deadline.

Whether the decision maker quits, and if so, when, depends on his type. Following much of the literature on time–inconsistent decision making we distinguish two types of decision makers: naive and sophisticated.

A naive decision maker acts under the erroneous assumption that his current self controls all future decisions in the game. Thus a naive decision maker intends to follow a strategy that maximizes his payoff over the entire continuation game, but in reality only carries out the first action. In the procrastination game, this behavior implies never quitting as the decision maker always intends to quit later.

A sophisticated decision maker, in contrast, is fully aware that his day t self only controls the decision on day t, and that future selves have different preferences. In the literature, this behavior is captured by the concept of sophisticated equilibrium, which essentially is nothing but a subgame perfect equilibrium of the game with multiple selves.

Somewhat surprisingly, the procrastination game has no sophisticated equilibrium in pure strategies, a fact already noticed (without proof) in Flesch et al. (2010), who introduced this game to show that properties of games with infinitely many players can be substantially different from those having finitely many players. For the sake of completeness, we give a short argument.

Claim 2.1. *The Procastination game has no sophisticated equilibrium in pure strategies.*

Proof. Suppose by way of contradiction that there is a sophisticated equilibrium in pure strategies. We distinguish three cases and derive a contradiction in each case. Let I be the set of players who choose action S conditional on reaching their decision nodes.

Case 1: Each player plays action C, so each player receives a payoff of 0. Player 1 would get a payoff of 1 by deviating to S.

Case 2: Exactly one player, say player t, chooses S. Since players $t+1, t+2, \ldots$ do not belong to I, player $t+1$ receives a payoff of 0 in the subgame starting in time $t+1$. A deviation to S gives player $t+1$ a payoff of 1.

Case 3: There exist two distinct players, say t_1 and t_2, who are elements of I. Without loss of generality, suppose $t_1 < t_2$. In the subgame starting in time t_1, player t_1 receives a payoff of 1 but he would get 2 by deviating to C. □

We show in Claim 6.1 that the procrastination game has no sophisticated equilibrium even when mixed strategies are considered. Non–existence of a sophisticated equilibrium in the procrastination game motivates us to consider approximate solution concepts: naive ϵ–equilibrium and sophisticated ϵ–equilibrium. Under both concepts, the decision maker is assumed to maximize his payoff up to a margin of ϵ. As we demonstrate in Section 6, the

procrastination game does have a sophisticated ϵ–equilibrium for each positive ϵ, namely the strategy profile whereby each self stops with probability ϵ.

3 The General Model

In this section, we describe infinite centipede games and define two solution concepts: naive ϵ-equilibrium and sophisticated ϵ-equilibrium.

In an infinite centipede game G, the set of players is the set $\mathbb{N}$ of natural numbers and the set of actions is $A = \{C, S\}$, where C stands for continue and S stands for stop. The game is played as follows. At time 1, player 1 chooses an action. If he chooses action S, then the game ends. If he chooses action C, then the play proceeds to time 2 where player 2 chooses an action. This is repeated as long as players choose action C. The payoff for player $i \in \mathbb{N}$ is a_t^i if the game ends at time t and a_∞^i if no one plays action S. We assume that payoffs are uniformly bounded, i.e.,

$$B = \sup_{i \in \mathbb{N}} \sup_{t \in \mathbb{N}^*} |a_t^i| < \infty, \tag{1}$$

where $\mathbb{N}^*$ stands for $\mathbb{N} \cup \{\infty\}$. By using the vector notation $a_t = (a_t^i)_{i \in \mathbb{N}}$ for every $t \in \mathbb{N}^*$, a centipede game can be represented as in Figure 2.

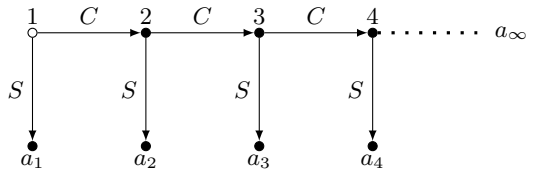

Figure 2: An infinite centipede game.

Our model can easily encompass various forms of discounted utility. For instance, when for all players the choice of C during the first t days yields instantaneous benefits $b_t \geq 0$ in day t, the instantaneous costs of stopping in day t are $c_t \geq 0$, after stopping no further instantaneous costs and benefits occur, and player i discounts benefits and costs in day t by the discount factor $\delta_t^i \in [0, 1]$, then we have

$$a_t^i = \sum_{k=1}^{t-1} \delta_k^i b_k - \delta_t^i c_t.$$

If discounting takes the standard exponential form, $\delta_t^i = (\delta)^t$ for some $\delta \in (0, 1)$, then boundedness of the sequences b_t and c_t is sufficient for $a_\infty^i = \sum_{k=1}^{\infty} \delta_k^i b_k$ to be well-defined and to obtain uniform boundedness as expressed in (1).

A *strategy* for player i is a probability distribution σ^i on the set of actions $\{C, S\}$. The interpretation is that, if time i is reached, then σ^i recommends to play C with probability

$\sigma^i(C)$ and to play S with probability $\sigma^i(S)$. The set of strategies for player i is denoted by Σ^i and the set of strategy profiles is denoted by $\Sigma = \times_{i\in\mathbb{N}}\Sigma^i$. A strategy σ^i of player i is uniquely specified by the probability to stop, $\sigma^i(S)$. Hence Σ^i can be identified with $[0,1]$, and Σ can be identified with $[0,1]^{\mathbb{N}}$.

A strategy σ^i is called *pure* if either $\sigma^i(S) = 0$ or $\sigma^i(S) = 1$. Hence a pure strategy is an element of $\{0,1\}$ and a pure strategy profile is an element of $\{0,1\}^{\mathbb{N}}$.

Let σ be a strategy profile. The expected utility of player i, conditional on the game not being stopped before time t, is denoted by $u^i(\sigma|t)$ and can be calculated as:

$$u^i(\sigma|t) = \sigma^t(S)\cdot a^i_t + \sum_{k=t+1}^{\infty}\sigma^k(S)\prod_{j=t}^{k-1}\sigma^j(C)\cdot a^i_k + \prod_{j=t}^{\infty}\sigma^j(C)\cdot a^i_\infty.$$

Note that $u^i(\sigma|i+1)$ is the expected utility of player i when player i plays action C at time i.

We now define the concept of naive ϵ-equilibrium.

Definition 3.1. Let $\epsilon \geq 0$. A strategy profile $\tau^* \in \Sigma$ is called a *naive ϵ-equilibrium* if there exists a sequence $(\tau_i)_{i\in\mathbb{N}}$ of strategy profiles satisfying the following two conditions:

1. $\tau^{*,i} = \tau^i_i$ for every player $i \in \mathbb{N}$,

2. $u^i(\tau_i|i) \geq u^i(\sigma|i) - \epsilon$ for every player $i \in \mathbb{N}$ and every strategy profile $\sigma \in \Sigma$.

A naive 0-equilibrium is simply called a *naive equilibrium*.

The idea behind Definition 3.1 originates with the literature on time–inconsistent decision making. Thus suppose that, as in the procrastination game of the previous section, player i represents the day i self of a decision maker. The strategy profile τ_i can then be thought of as the complete course of actions that the day i self *intends* to carry out. Condition 2 says that τ_i is an ϵ–optimal strategy profile in the continuation game when evaluated against day i's preferences.

The decision maker is naive as he fails to realize that his day i self only controls the decision on day i, and that the strategy profile τ_i need not be ϵ–optimal for the future selves. As a result, the sequence of strategies that the naive decision maker actually carries out is $(\tau^1_1, \tau^2_2, \dots)$. This sequence is exactly τ^* by Condition 1 of Definition 3.1. Thus τ^* could be thought of as the *realized* behavior of a naive decision maker.

Naive ϵ–equilibrium could also be interpreted without a recourse to time–inconsistent decision making. It represents a situation in which player i fails to take into account the fact that he only controls a single decision node at time i, subsequent decisions being taken by other players.

For each player $i \in \mathbb{N}$, we define $A^i = \{a_i^i, a_{i+1}^i, \ldots\} \cup \{a_\infty^i\}$, which is the set of possible payoffs for player i conditional on the fact that no player has stopped the game before him. Also, we let

$$M^i = \sup A^i. \tag{2}$$

With this notation, Condition 2 of Definition 3.1 is equivalent to

$$u^i(\tau_i|i) \geq M^i - \epsilon \text{ for every player } i \in \mathbb{N}. \tag{3}$$

As an illustration, consider the Procrastination game in Figure 1. In this game, $M^i = 2$ for each player i. It can be verified that the naive ϵ-equilibria are exactly those strategy profiles τ^* for which $\tau^{*,i}(C) \geq 1 - \epsilon$ for each player i.

Definition 3.2. Let $\epsilon \geq 0$. A strategy profile $\sigma^* \in \Sigma$ is called a *sophisticated ϵ-equilibrium* if for each player $i \in \mathbb{N}$ and each strategy $\sigma^i \in \Sigma^i$:

$$u^i(\sigma^*|i) \geq u^i((\sigma^i, \sigma^{*,-i})|i) - \epsilon.$$

A sophisticated 0-equilibrium is simply called a *sophisticated equilibrium*.

A sophisticated decision maker is fully aware that his day i self only controls the decision on day i and that the future selves may have different preferences. Thus in a sophisticated ϵ–equilibrium each self of the decision maker ϵ–best responds to the strategies of the future selves. Hence a sophisticated ϵ–equilibrium is essentially a subgame perfect ϵ–equilibrium of the game.

If player i's opponents play according to $\sigma^{*,-i}$ and player i chooses action C at time i, then his payoff is $u^i(\sigma^*|i+1)$, whereas if player i chooses action S, then his payoff is a_i^i. Hence, a strategy profile σ^* is a sophisticated ϵ-equilibrium if and only if σ^* satisfies the following two inequalities for every player i:

$$u^i(\sigma^*|i) \geq u^i(\sigma^*|i+1) - \epsilon, \tag{4}$$

$$u^i(\sigma^*|i) \geq a_i^i - \epsilon. \tag{5}$$

In our illustrative example, the Procrastination game in Figure 1, there is no sophisticated ϵ-equilibrium in pure strategies for $\epsilon \in [0, 1)$ and no sophisticated equilibrium as we will show in Section 6. On the other hand, for $\epsilon > 0$, we will provide a proof that this game does admit a sophisticated ϵ-equilibrium in mixed strategies, where each player stops with probability ϵ.

4 Existence of Equilibrium in the Upper Semicontinuous Case

In this section, we establish the existence of naive and sophisticated equilibria if the payoffs in the game are *upper semi-continuous*, i.e., if for every player $i \in \mathbb{N}$

$$\limsup_{t\to\infty} a_t^i \leq a_\infty^i. \tag{6}$$

The payoffs of player i are upper semi-continuous if the payoffs of player i when stopping the game at time t with t going to infinity are less than or equal to the payoffs obtained when never stopping the game. A typical example of payoffs that are upper semi-continuous but not continuous concerns the case where stopping the game corresponds to making a costly investment, say for instance in the reduction of carbon dioxide emissions. Making the investment too late leaves too little time to recoup the costs or would not change a disastrous outcome. In such a case, not making the investment at all would be preferred to making the investment at a very late point in time. Another example of payoffs that are upper semi-continuous but not continuous results when stopping corresponds to giving in and the player derives a positive psychological benefit from never doing so. Continuity, and therefore upper semi-continuity, is satisfied in the standard model of exponential discounting when instantaneous costs and benefits are uniformly bounded.

First, we provide necessary and sufficient conditions for the existence of naive equilibrium without continuity assumptions on the payoffs.

Theorem 4.1. *The following statements are equivalent:*
(i) For every player $i \in \mathbb{N}$, the set $A^i = \{a_i^i, a_{i+1}^i, \ldots\} \cup \{a_\infty^i\}$ has a maximum.
(ii) There exists a naive equilibrium in pure strategies.
(iii) There exists a naive equilibrium.

Proof. $(i \to ii)$ Consider a player i. If the maximum of A^i is a_∞^i, then let τ_i be the pure strategy profile that always chooses action C. If the maximum of A^i is not a_∞^i, but some a_t^i with $t \in \mathbb{N}$, then let τ_i be the pure strategy profile that always chooses action C, except at time t, where it chooses action S. It is clear that the pure strategy profile τ^* defined by $\tau^{*,i} = \tau_i^i$ for every $i \in \mathbb{N}$ is a naive equilibrium.
$(ii \to iii)$ Obvious.
$(iii \to i)$ Suppose that there exists a naive equilibrium. Take a player i. By (3), there exists a strategy profile τ_i such that $u^i(\tau_i|i) \geq M^i$. It follows from (2) that $u^i(\tau_i|i) = M^i$ and that there exists a $t \in \{i, i+1, \ldots\} \cup \{\infty\}$ such that $a_t^i = M^i$. Thus, the set A^i has a maximum. □

In view of the above theorem, a naive equilibrium does not always exist. A concrete example is the game in Figure 3, which we will consider later. However, we have the following existence result for games with upper semicontinuous payoffs.

Theorem 4.2. *There exists a naive equilibrium in pure strategies if the payoffs are upper semicontinuous.*

Proof. Due to (6), the set $A^i = \{a_i^i, a_{i+1}^i, \ldots\} \cup \{a_\infty^i\}$ has a maximum for every player $i \in \mathbb{N}$. Therefore, by Theorem 4.1, there exists a naive equilibrium in pure strategies. □

Now we turn to the existence of a sophisticated equilibrium. The proof of the following result employs a truncation approach similar to that in Fudenberg and Levine (1983). One crucial difference however is that we do not assume the payoffs to be continuous, but only upper semicontinuous.

Theorem 4.3. *There exists a sophisticated equilibrium in pure strategies if the payoffs are upper semicontinuous.*

Proof. Consider a centipede game G. For every $T \in \mathbb{N}$, we define the T-period truncated game G_T which is identical to G except for one modification: if all players $1, \ldots, T$ choose to continue, then, regardless of future play, the payoff of each player $i \in \mathbb{N}$ is equal to a_{T+1}^i. Since the payoffs cannot change after time T, this game is essentially a T-period game.

For every $T \in \mathbb{N}$, the truncated game G_T admits a pure sophisticated equilibrium σ_T in which $\sigma_T^i(S) = 1$ for every player $i \geq T + 1$. Indeed, due to the payoffs in G_T, we can set $\sigma_T^i(S) = 1$ for every player $i \geq T + 1$ and then determine $\sigma_T^T(S), \ldots, \sigma_T^1(S)$ by backward induction. If a player is indifferent between playing action C and action S then either action can be taken. The set of pure strategy profiles, as mentioned earlier, can be identified with the infinite Cartesian product $\{0,1\}^{\mathbb{N}}$ and is thus a compact metrizable topological space. Hence the sequence $(\sigma_T)_{T=1}^\infty$ has an accumulation point $\bar{\sigma} \in \{0,1\}^{\mathbb{N}}$. By taking a subsequence if necessary, we can assume that $(\sigma_T)_{T=1}^\infty$ converges to the strategy profile $\bar{\sigma}$. We distinguish two cases.

Case 1: Suppose that there are only finitely many players who play action S in the strategy profile $\bar{\sigma}$. So, there exists a time t such that for all players $i \geq t$, $\bar{\sigma}^i(C) = 1$. We prove that $\bar{\sigma}$ induces a sophisticated equilibrium for the subgame of G starting at time t. So we need to show that $u^i(\bar{\sigma}|i) \geq a_i^i$ for all $i \geq t$. For every $i, T \in \mathbb{N}$, let

$$m_T^i = \min\{k \geq i | \sigma_T^k(S) = 1\},$$

so m_T^i is the first player at time i or later who stops in the strategy profile σ_T. We have for

every $i \geq t$ that

$$\begin{aligned} u^i(\bar{\sigma}|i) &= a^i_\infty && (7)\\ &\geq \limsup_{T\to\infty} a^i_T && (8)\\ &\geq \limsup_{T\to\infty} a^i_{m^i_T} && (9)\\ &= \limsup_{T\to\infty} u^i(\sigma_T|i) && (10)\\ &\geq a^i_i. && (11) \end{aligned}$$

Equality (7) holds, since no one stops from time t onwards in the strategy profile $\bar{\sigma}$; (8) follows from the assumption of the theorem; (9) follows from the definition of limit superior, because $m^i_T \to \infty$ as $T \to \infty$, for every $i \geq t$; (10) holds as $u^i(\sigma_T|i) = a^i_{m^i_T}$ for every i and T; and finally (11) is true since σ_T is a sophisticated equilibrium in G_T. Hence, we have $u^i(\bar{\sigma}|i) \geq a^i_i$ for all $i \geq t$, as desired. This means that $\bar{\sigma}$ induces a sophisticated equilibrium for the subgame of G starting at time t. Now we can use backward induction from time t to obtain a sophisticated equilibrium in G.

Case 2 : Suppose that there are infinitely many players who play action S in the strategy profile $\bar{\sigma}$. Take an arbitrary player i. Let

$$n^i = \min\{k > i | \bar{\sigma}^k(S) = 1\},$$

so n^i is the first player at time $i+1$ or later who stops in the strategy profile $\bar{\sigma}$. Since $(\sigma_T)_{T=1}^\infty$ converges to $\bar{\sigma}$, there exists $T \geq n^i$ such that for all $j \leq n^i$ we have $\sigma^j_T = \bar{\sigma}^j$. Because σ_T is a sophisticated equilibrium in the game G_T, player i does not have a profitable deviation from σ_T in G_T. It follows that player i does not have a profitable deviation from $\bar{\sigma}$ in the game G. We conclude that $\bar{\sigma}$ is a sophisticated equilibrium of G. □

5 Existence of Naive ϵ-Equilibrium

We know from the previous section that a naive equilibrium does not always exist. The following theorem deals with the existence of naive ϵ-equilibrium.

Theorem 5.1. *For every $\epsilon > 0$, there exists a naive ϵ-equilibrium in pure strategies.*

Proof. Let $\epsilon > 0$ be given. Consider a player i. We distinguish two cases in order to define a pure strategy profile τ_i.

Case 1: $a^i_\infty \geq M^i - \epsilon$, where M^i is given in (2). In this case, let τ_i be the pure strategy profile that always chooses action C.

Case 2: $a^i_\infty < M^i - \epsilon$. In this case, there exists $t \in \{i, i+1, \ldots\}$ such that $a^i_t \geq M^i - \epsilon$. Let τ_i be the pure strategy profile that always chooses action C, except at time t, where it chooses action S.

Now define the pure strategy profile τ^* by setting $\tau^{*,i} = \tau^i_i$ for every player $i \in \mathbb{N}$. Then, the strategy profiles τ^* and τ_i, for every $i \in \mathbb{N}$, satisfy Condition 1 of Definition 3.1 and inequality (3), so τ^* is a naive ϵ-equilibrium. □

As an illustration, consider a game where the payoff for player $i \in \mathbb{N}$ is $1 - 1/(t - i + 1)$ if the game ends at time $t > i$, and 0 if the game ends at time $t \leq i$ or if no one stops. The game is given in Figure 3.

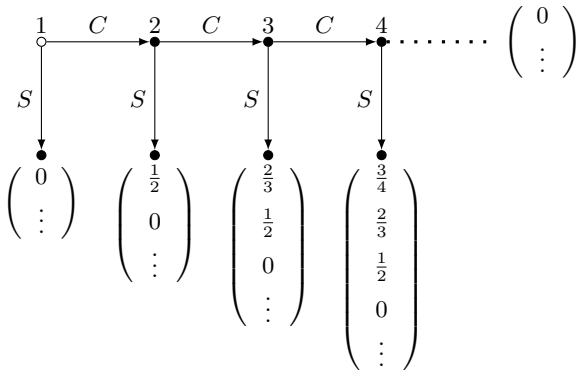

Figure 3: A game without naive equilibrium.

By Theorem 4.1, there is no naive equilibrium in this game. On the other hand, the pure strategy profile τ^* that always chooses action C is a naive ϵ-equilibrium for every $\epsilon > 0$. Indeed, let $\epsilon > 0$. For every player i, take a time t_i such that $t_i > i$ and $a^i_{t_i} \geq 1 - \epsilon$. Define τ_i to be the pure strategy profile that always chooses action C, except at time t_i, where it chooses action S. Then, for every player $i \in \mathbb{N}$ it holds that $\tau^{*,i} = \tau^i_i$ and inequality (3) is satisfied, so τ^* is a naive ϵ-equilibrium as claimed. Notice also that the strategy profile τ^* is a sophisticated equilibrium of the game.

6 Existence of Sophisticated ϵ-Equilibrium

In this section, we examine the existence of a sophisticated ϵ-equilibrium. The following claim establishes that the Procrastination game has neither a sophisticated equilibrium nor a pure sophisticated ϵ-equilibrium.

Claim 6.1. *The Procrastination game in Figure 1 has the following properties:*

1. *It admits no sophisticated ϵ-equilibrium in pure strategies for any $\epsilon \in [0, 1)$.*

2. It admits no sophisticated equilibrium.

Proof. First we prove part 1. Take an $\epsilon \in [0, 1)$ and suppose by way of contradiction that σ is a sophisticated ϵ-equilibrium in pure strategies. Consider some player $i \in \mathbb{N}$. For any pure strategy τ^i it holds that $u^i(\sigma|i) \geq u^i((\tau^i, \sigma^{-i})|i) - \epsilon$. Because pure strategy profiles can only induce payoffs 0, 1 or 2 to any player in the game and because $\epsilon < 1$, we must have $u^i(\sigma|i) \geq u^i((\tau^i, \sigma^{-i})|i)$. Therefore, pure strategy profile σ is a sophisticated equilibrium. This contradicts Claim 2.1.

Now we prove part 2. Assume to the contrary that σ is a sophisticated equilibrium. For every player k, let

$$p(\sigma|k) = \prod_{i=k}^{\infty} \sigma^i(C)$$

be the probability that the game never stops, provided that it has not been stopped before time k and that the players play according to σ.

Assume first that there are two players i and j, with $i < j$, such that σ^i and σ^j are not pure. Since σ^i is not pure, we have $u^i((C, \sigma^{-i})|i) = u^i((S, \sigma^{-i})|i)$. Note that $u^i((S, \sigma^{-i})|i) = 1$ and

$$u^i((C, \sigma^{-i})|i) = (1 - p(\sigma|i + 1)) \cdot 2,$$

so $p(\sigma|i + 1) = 0.5$. By a similar argument, we obtain for player j that $p(\sigma|j + 1) = 0.5$. But then

$$p(\sigma|i + 1) = \sigma^{i+1}(C)\sigma^{i+2}(C) \cdots \sigma^j(C)p(\sigma|j + 1)$$

yields $\sigma^j(C) = 1$, which is a contradiction to the fact that σ^j is not pure.

Therefore, there is a time t such that in the subgame that starts at t, the strategy profile σ is a sophisticated equilibrium in pure strategies. By backward induction, we can then construct a sophisticated equilibrium in pure strategies for the whole game. This is however in contradiction with Claim 2.1. □

The main result of this section is the following theorem on the existence of a sophisticated ϵ-equilibrium.

Theorem 6.2. *For every $\epsilon > 0$, there exists a sophisticated ϵ-equilibrium.*

Proof. Take a centipede game G. Recall the definition of B from (1). Choose a number δ so that $0 < \delta \leq \min\{\frac{\epsilon}{2B}, 1\}$. Define the normal-form game G^* as the game with set of players $\mathbb{N}$, set of strategies for player $i \in \mathbb{N}$ equal to

$$\Sigma^{*i} = \{\sigma^i \in \Sigma^i \mid \sigma^i(S) \in [\delta, 1]\},$$

and utility function v^i of player $i \in \mathbb{N}$ defined by

$$v^i(\sigma) = u^i(\sigma \mid i)$$

for each σ in $\Sigma^* = \times_{i\in\mathbb{N}}\Sigma^{*i}$. The strategy sets of all players are non-empty, convex, and compact. Each player's utility function is affine, so quasi-concave, in that player's own strategy, since

$$v^i(\sigma) = \sigma^i(S)a_i^i + (1 - \sigma^i(S))u^i(\sigma|i+1), \tag{12}$$

where $u^i(\sigma|i+1)$ is unaffected by the strategy σ^i of player i. Crucially, each player's utility function is continuous on Σ^*. The key to the continuity of the payoff functions is that under any strategy profile in Σ^*, the probability that no one ever stops is zero.

The game G^* therefore has a Nash equilibrium, say σ^*, by Theorem 6.2 of Peleg and Yaari (1973). In view of equation (12), if $u^i(\sigma^*|i+1) < a_i^i$ then $\sigma^{*,i}(S) = 1$, and if $u^i(\sigma^*|i+1) > a_i^i$ then $\sigma^{*,i}(S) = \delta$.

We show that σ^* is a sophisticated ϵ-equilibrium of G. Take any player $i \in \mathbb{N}$. If $u^i(\sigma^*|i+1) < a_i^i$ then $\sigma^{*,i}(S) = 1$, so $u^i(\sigma^*|i) = v^i(\sigma^*) = a_i^i$ and (4) and (5) are satisfied. If $u^i(\sigma^*|i+1) = a_i^i$ then $u^i(\sigma^*|i) = v^i(\sigma^*) = a_i^i$, so (4) and (5) are satisfied. So assume that $u^i(\sigma^*|i+1) > a_i^i$. We have that $\sigma^{*,i}(S) = \delta$, so

$$\begin{aligned}
u^i(\sigma^*|i) &= v^i(\sigma^*) \\
&= \sigma^{*,i}(S)a_i^i + (1 - \sigma^{*,i}(S))u^i(\sigma^*|i+1) \\
&= \delta\, a_i^i + (1-\delta)\, u^i(\sigma^*|i+1) \\
&= \delta(a_i^i - u^i(\sigma^*|i+1)) + u^i(\sigma^*|i+1) \\
&\geq -\delta 2B + u^i(\sigma^*|i+1) \\
&\geq u^i(\sigma^*|i+1) - \epsilon.
\end{aligned}$$

Hence, inequality (4) is satisfied. Furthermore, it holds that

$$\begin{aligned}
u^i(\sigma^*|i) - a_i^i &= v^i(\sigma^*) - a_i^i \\
&= \sigma^{*,i}(S)a_i^i + (1 - \sigma^{*,i}(S))u^i(\sigma^*|i+1) - a_i^i \\
&= (1 - \sigma^{*,i}(S))(u^i(\sigma^*|i+1) - a_i^i) \\
&\geq 0,
\end{aligned}$$

where the inequality follows by $u^i(\sigma^*|i+1) > a_i^i$. Hence, inequality (5) is also satisfied. □

We remark that one could give a direct proof of the theorem using truncations of the game tree. To do so, let G_T be the game as in the proof of Theorem 4.3. One could show that G_T

has a sophisticated ϵ–equilibrium σ_T that is an element of Σ^*. One can further show that any accumulation point of the sequence $\{\sigma_T\}_{T\in\mathbb{N}}$ is a sophisticated ϵ–equilibrium of the original game G.

According to the definition of B in (1), the payoffs are uniformly bounded. One might wonder whether it would be enough to assume only that the payoffs are bounded for each player separately. The following example shows that this weaker assumption would not suffice for the existence of a sophisticated ϵ-equilibrium.

Consider the game with the following payoffs for every player i: If the game ends before time i then player i's payoff is 0. If the game ends at time i then player i's payoff is 2^i. If the game ends after time i then player i's payoff is 2^{i+1}. Finally, if the game never ends then player i's payoff is 0. The game tree is given in Figure 4.

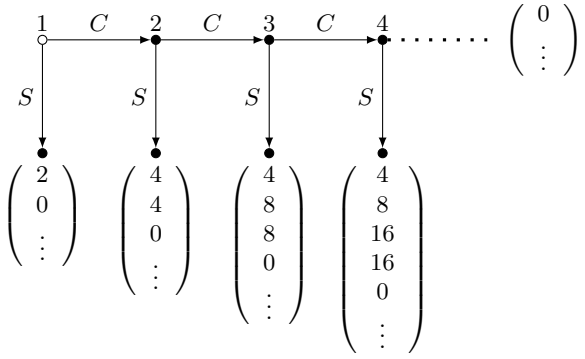

Figure 4: A centipede game without a sophisticated ϵ-equilibrium.

It is clear that, for each $i \in \mathbb{N}$, $\sup_{t\in\mathbb{N}^*} |a_t^i| = 2^{i+1}$. At the same time it holds that $\sup_{i\in\mathbb{N}} \sup_{t\in\mathbb{N}^*} |a_t^i| = \infty$.

Claim 6.3. *For each $\epsilon > 0$, the game in Figure 4 admits no sophisticated ϵ-equilibrium.*

Proof. Take $\epsilon > 0$ and suppose by way of contradiction that there is a sophisticated ϵ-equilibrium σ for the game in Figure 4. As in the proof of Claim 6.1, let $p(\sigma|k) = \prod_{i=k}^{\infty} \sigma^i(C)$ denote the probability that the game never stops, given that it has not stopped before time k and the players follow the strategy profile σ.

Take a player $i \in \mathbb{N}$. Since σ is a sophisticated ϵ–equilibrium, inequality (5) implies that $u^i(\sigma|i) \geq 2^i - \epsilon$. On the other hand, since 2^{i+1} is the highest payoff player i can get and since he gets 0 if the game never stops, we have the following upper bound on the payoff: $u^i(\sigma|i) \leq (1 - p(\sigma|i)) \cdot 2^{i+1} + p(\sigma|i) \cdot 0$. Combining these facts and rearranging terms, we find that

$$p(\sigma|i) \leq \frac{1}{2} + \frac{\epsilon}{2^{i+1}}.$$

Notice that the sequence $\{p(\sigma|i)\}_{i\in\mathbb{N}}$ is non–decreasing and bounded and hence has a limit. Furthermore, the preceding inequality implies that

$$\lim_{i\to\infty} p(\sigma|i) \leq \frac{1}{2}.$$

Consider any player $i \in \mathbb{N}$. For each $j > i$ it holds that

$$p(\sigma|i) = \prod_{t=i}^{\infty} \sigma^t(C) = \prod_{t=i}^{j-1} \sigma^t(C) \cdot p(\sigma|j).$$

Taking the limit as j approaches infinity, we obtain

$$p(\sigma|i) = \lim_{j\to\infty} \prod_{t=i}^{j-1} \sigma^t(C) \cdot \lim_{j\to\infty} p(\sigma|j) = p(\sigma|i) \cdot \lim_{j\to\infty} p(\sigma|j) \leq p(\sigma|i) \cdot \frac{1}{2},$$

which implies that $p(\sigma|i) = 0$.

Thus for each $i \in \mathbb{N}$ it holds that $u^i(\sigma|i) = \sigma^i(S)\cdot 2^i + (1-\sigma^i(S))\cdot 2^{i+1}$ and $u^i(\sigma|i+1) = 2^{i+1}$. Since σ is a sophisticated ϵ–equilibrium, it holds by inequality (4) that $u^i(\sigma|i) \geq u^i(\sigma|i+1) - \epsilon$, and therefore $\sigma^i(S) \leq 2^{-i}\epsilon$.

Now take $t \in \mathbb{N}$ such that $2^{1-t}\epsilon < 1$. Since $1 - p(\sigma|t)$ is the probability that the game eventually stops conditional on time t being reached, we have

$$1 - p(\sigma|t) = \sum_{j=t}^{\infty} \sigma^j(S) \prod_{i=t}^{j-1} \sigma^i(C) \leq \sum_{j=t}^{\infty} \sigma^j(S) \leq \sum_{j=t}^{\infty} 2^{-j}\epsilon \leq 2^{1-t}\epsilon,$$

so $p(\sigma|t) \geq 1 - 2^{1-t}\epsilon > 0$, contradicting $p(\sigma|t) = 0$. □

7 Sophisticates Stop Earlier

One of the key results in O'Donoghue and Rabin (1999) is that a sophisticated decision maker executes a task earlier than a naive decision maker. In this section we derive the counterpart of this result in our setup. The comparison of sophisticated and naive decision makers in our setup is somewhat complicated by the fact that in general there might exist multiple sophisticated ϵ–equilibria and multiple naive ϵ–equilibria. We thus have to compare two sets of equilibria. We achieve this by showing that [1] given a sophisticated ϵ–equilibrium there exists a naive ϵ–equilibrium with the probability to stop not greater than in the sophisticated ϵ–equilibrium at any given time, and [2] given a naive ϵ–equilibrium there exists a sophisticated ϵ–equilibrium with the probability to stop not smaller than in the naive ϵ–equilibrium at any given time.

We restrict attention to the case $\epsilon > 0$ to guarantee the existence of naive and sophisticated ϵ–equilibria. Indeed, recall that the Procrastination game of section 2 has no sophisticated

equilibrium (see Claim 6.1), while it does admit a naive equilibrium, namely playing C at every period. The game depicted in Figure 3 admits a sophisticated equilibrium but not a naive equilibrium.

Theorem 7.1. *Let some $\epsilon > 0$ be given. For each sophisticated ϵ-equilibrium σ^*, there exists a naive ϵ-equilibrium τ^* such that for all $i \in \mathbb{N}$, $\tau^{*,i}(S) \leq \sigma^{*,i}(S)$.*

Proof. Let σ^* be a sophisticated ϵ-equilibrium. For every $i \in \mathbb{N}$, we define the strategy profile $\tau_i \in \Sigma$ as follows. If $u^i(\sigma^*|i) \geq M^i - \epsilon$, then let $\tau_i = \sigma^*$. Otherwise, it holds that $u^i(\sigma^*|i) < M^i - \epsilon$. As $a_i^i - \epsilon \leq u^i(\sigma^*|i)$ by inequality (5), it then holds that $a_i^i < M^i$. Consequently, there exists $t \in \{i+1, i+2, \dots\} \cup \{\infty\}$ such that $a_t^i \geq M^i - \epsilon$. If $t = \infty$ then we define τ_i by letting $\tau_i^j(S) = 0$ for all $j \in \mathbb{N}$, while if $t \in \{i+1, i+2, \dots\}$ we define $\tau_i^t(S) = 1$ and $\tau_i^j(S) = 0$ for all $j \neq t$. It is easy to see that the strategy profile τ_i satisfies inequality (3). Now define $\tau^* \in \Sigma$ by letting $\tau^{*,i} = \tau_i^i$ for each $i \in \mathbb{N}$. Then τ^* is a naive ϵ–equilibrium. Since $\tau^{*,i}(S)$ is either equal to $\sigma^{*,i}(S)$ or 0, it holds for every $i \in \mathbb{N}$ that $\tau^{*,i}(S) \leq \sigma^{*,i}(S)$. □

Theorem 7.2. *Let some $\epsilon > 0$ be given. For each naive ϵ-equilibrium τ^*, there exists a sophisticated ϵ-equilibrium σ^* such that for all $i \in \mathbb{N}$, $\tau^{*,i}(S) \leq \sigma^{*,i}(S)$.*

Proof. Fix a naive ϵ-equilibrium τ^* and let $(\tau_i)_{i\in\mathbb{N}}$ be as in Definition 3.1. Choose a number δ so that $0 < \delta \leq \min\{\frac{\epsilon}{2B}, 1\}$. For $i \in \mathbb{N}$ let $\pi^i = \max\{\delta, \tau^{*,i}(S)\}$. Define the normal–form game G^{**} as the game with set of players $\mathbb{N}$, set of strategies for player $i \in \mathbb{N}$ equal to

$$\Sigma^{**i} = \{\sigma^i \in \Sigma^i \mid \sigma^i(S) \in [\pi^i, 1]\},$$

and utility function v^i of player $i \in \mathbb{N}$ defined by

$$v^i(\sigma) = u^i(\sigma \mid i)$$

for each σ in $\Sigma^{**} = \times_{i\in\mathbb{N}}\Sigma^{**i}$. The strategy sets of all players are non-empty, convex, and compact. The utility functions of all players are continuous on Σ^{**} and affine, so quasi-concave, in their own strategy. The game G^{**} therefore has a Nash equilibrium, say σ^*, by Theorem 6.2 of Peleg and Yaari (1973). The payoff functions can be written as

$$v^i(\sigma) = \sigma^i(S)a_i^i + (1 - \sigma^i(S))u^i(\sigma|i+1),$$

where $u^i(\sigma|i+1)$ is unaffected by the strategy σ^i of player i. It follows that if $u^i(\sigma^*|i+1) < a_i^i$ then $\sigma^{*,i}(S) = 1$, and if $u^i(\sigma^*|i+1) > a_i^i$ then $\sigma^{*,i}(S) = \pi^i$.

We show that σ^* is a sophisticated ϵ-equilibrium of G.

Take any player $i \in \mathbb{N}$. If $u^i(\sigma^*|i+1) < a_i^i$ then $\sigma^{*,i}(S) = 1$, so $u^i(\sigma^*|i) = v^i(\sigma^*) = a_i^i$ and (4) and (5) are satisfied. If $u^i(\sigma^*|i+1) = a_i^i$ then $u^i(\sigma^*|i) = v^i(\sigma^*) = a_i^i$, so (4) and (5) are satisfied. So assume that $u^i(\sigma^*|i+1) > a_i^i$. We have that $\sigma^{*,i}(S) = \pi^i$.

If $\pi^i = \delta$, then

$$\begin{aligned} u^i(\sigma^*|i) &= v^i(\sigma^*) \\ &= \sigma^{*,i}(S)a_i^i + (1-\sigma^{*,i}(S))u^i(\sigma^*|i+1) \\ &= \delta\, a_i^i + (1-\delta)\, u^i(\sigma^*|i+1) \\ &\geq -\delta 2B + u^i(\sigma^*|i+1) \\ &\geq u^i(\sigma^*|i+1) - \epsilon. \end{aligned}$$

If $\pi^i = \tau^{*,i}(S)$, we have the following chain of inequalities:

$$\begin{aligned} u^i(\sigma^*|i) - u^i(\sigma^*|i+1) &= \sigma^{*,i}(S)a_i^i + (1-\sigma^{*,i}(S))u^i(\sigma^*|i+1) - u^i(\sigma^*|i+1) \\ &= \sigma^{*,i}(S)(a_i^i - u^i(\sigma^*|i+1)) \\ &= \tau^{*,i}(S)(a_i^i - u^i(\sigma^*|i+1)) && (13) \\ &\geq \tau^{*,i}(S)(a_i^i - M^i) && (14) \\ &= \tau^{*,i}(S)a_i^i + (1-\tau^{*,i}(S))M^i - M^i \\ &\geq \tau^{*,i}(S)a_i^i + (1-\tau^{*,i}(S))u^i(\tau_i|i+1) - M^i && (15) \\ &= u^i(\tau_i|i) - M^i \\ &\geq -\epsilon, && (16) \end{aligned}$$

where equality (13) holds since we assume $\sigma^{*,i}(S) = \pi^i = \tau^{*,i}(S)$. Inequality (14) holds since $u^i(\sigma^*|i+1)$ is an expectation of the payoffs under a probability distribution over the set $\{a_{i+1}^i, a_{i+2}^i, \cdots\} \cup \{a_\infty^i\}$, a subset of A^i, and hence is bounded above by $M^i = \sup A^i$. Inequality (15) follows since $u^i(\tau_i|i+1)$ is likewise bounded above by M^i, and inequality (16) follows from inequality (3).

Hence, σ^* satisfies inequality (4).

Moreover, we have

$$\begin{aligned} u^i(\sigma^*|i) - a_i^i &= \sigma^{*,i}(S)a_i^i + (1-\sigma^{*,i}(S))u^i(\sigma^*|i+1) - a_i^i \\ &= (1-\sigma^{*,i}(S))(u^i(\sigma^*|i+1) - a_i^i) \\ &\geq 0, \end{aligned}$$

so inequality (5) is also satisfied. □

Notice that the Theorems 7.1 and 7.2 do not preclude the possibility that there exist a naive ϵ–equilibrium τ^* and a sophisticated ϵ–equilibrium σ^* such that $\tau^{*,i}(S) > \sigma^{*,i}(S)$ for all $i \in \mathbb{N}$. The trivial game where all payoffs are 0 would yield an example.

8 Conclusion

In this paper, we have examined infinite centipede games with uniformly bounded payoffs. We have looked at two solution concepts, naive and sophisticated ϵ-equilibria, depending on the type of the decision maker.

Regarding a naive decision maker, we show that there does not always exist a naive 0-equilibrium. We provide the necessary and sufficient conditions for the existence of a naive 0-equilibrium and show that upper semicontinuity of payoffs is sufficient for existence. Finally, we show that a naive ϵ-equilibrium always exists, for any $\epsilon > 0$.

For a sophisticated decision maker, we also show that sophisticated equilibria do not always exist. We show the existence of a sophisticated 0-equilibrium in pure strategies when payoffs are upper semicontinuous. Moreover, we show that for every $\epsilon > 0$, there exists a sophisticated ϵ-equilibrium.

We also examine the connection between naive and sophisticated decision makers. Let some $\epsilon > 0$ be given. We show that for every sophisticated ϵ-equilibrium there exists a naive ϵ-equilibrium such that the stopping probability of every player in the sophisticated ϵ-equilibrium strategy is higher than in the naive ϵ-equilibrium strategy. Additionally, we show that for every naive ϵ-equilibrium there exists a sophisticated ϵ-equilibrium such that the stopping probability of every player in the sophisticated ϵ-equilibrium strategy is higher than in the naive ϵ-equilibrium.

References

[1] Balbus Ł., Jaśkiewicz A., Nowak A. S. (2015). Existence of Stationary Markov Perfect Equilibria in Stochastic Altruistic Growth Economies. Journal of Optimization Theory and Applications, 165(1): 295-315.

[2] De Pril J., Flesch J., Kuipers J., Schoenmakers G., Vrieze K. (2014). Existence of Secure Equilibrium in Multi-Player Games with Perfect Information. Mathematical Foundations of Computer Science, Proceedings Part II, 213-225.

[3] Flesch J., Kuipers J., Mashiah-Yaakovi A., Schoenmakers G., Solan E., Vrieze K. (2010). Perfect-Information Games with Lower-Semicontinuous Payoffs. Mathematics of Operations Research, 35(4): 742-755.

[4] Fudenberg D., Levine D. (1983). Subgame-Perfect Equilibria of Finite- and Infinite-Horizon Games. Journal of Economic Theory, 31(2): 251-268.

[5] Gabrieli T., Ghosal S. (2013). Non-Existence of Competitive Equilibria with Dynamically Inconsistent Preferences. Economic Theory, 52(1): 299-313.

[6] Goldman S. M. (1979). Intertemporally Inconsistent Preferences and the Rate of Consumption. Econometrica, 47(3): 621-626.

[7] Herings P. J. J., Rohde K. I. M. (2006). Time-Inconsistent Preferences in a General Equilibrium Model. Economic Theory, 29(3): 591-619.

[8] Luttmer E. G. J., Mariotti T. (2006). Competitive Equilibrium when Preferences Change over Time. Economic Theory, 27(3): 679-690.

[9] Mashiah-Yaakovi A. (2009). Periodic Stopping Games. International Journal of Game Theory 38(2): 169-181.

[10] O'Donoghue T., Rabin M. (1999). Doing It Now or Later. The American Economic Review 89(1): 103-124.

[11] Peleg B., Yaari M. E. (1973). On the Existence of a Consistent Course of Action When Tastes Are Changing. The Review of Economic Studies 40(3): 391-401.

[12] Phelps E. S., Pollak R. A. (1968). On Second-Best National Saving and Game-Equilibrium Growth. The Review of Economic Studies 35(2): 185-199.

[13] Pollak R. A. (1968). Consistent planning. The Review of Economic Studies, 35(2): 201-208.

[14] Purves R. A., Sudderth W. D. (2011). Perfect Information Games with Upper Semicontinuous Payoffs. Mathematics of Operations Research, 36(3): 468-473.

[15] Solan E. (2005). Subgame-Perfection in Quitting Games with Perfect Information and Differential Equations. Mathematics of Operations Research, 30(1): 51-72.

[16] Strotz R. H. (1955). Myopia and Inconsistency in Dynamic Utility Maximization. The Review of Economic Studies, 23(3): 165-180.