

Voorneveld, Mark

Working Paper

The possibility of impossible stairways and greener grass

SSE/EFI Working Paper Series in Economics and Finance, No. 673

Provided in Cooperation with:

EFI - The Economic Research Institute, Stockholm School of Economics

Suggested Citation: Voorneveld, Mark (2007) : The possibility of impossible stairways and greener grass, SSE/EFI Working Paper Series in Economics and Finance, No. 673, Stockholm School of Economics, The Economic Research Institute (EFI), Stockholm

This Version is available at:

<https://hdl.handle.net/10419/56319>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

The possibility of impossible stairways and greener grass

Mark Voorneveld¹

Dept. of Econometrics and Operations Research, Tilburg University, The Netherlands

and

Dept. of Economics, Stockholm School of Economics, Sweden

Abstract. In classical game theory, players have finitely many actions and evaluate outcomes of mixed strategies using a von Neumann-Morgenstern utility function. Allowing a larger, but countable, player set introduces a host of phenomena that are impossible in finite games.

Firstly, in coordination games, all players have the same preferences: switching to a weakly dominant action makes everyone at least as well off as before. Nevertheless, there are coordination games where the best outcome occurs if everyone chooses a weakly dominated action, while the worst outcome occurs if everyone chooses the weakly dominant action.

Secondly, the location of payoff-dominant equilibria behaves capriciously: two coordination games that look so much alike that even the consequences of unilateral deviations are the same may nevertheless have disjoint sets of payoff-dominant equilibria.

Thirdly, a large class of games has no (pure or mixed) Nash equilibria. Following the proverb “the grass is always greener on the other side of the hedge”, greener-grass games model constant discontent: in one part of the strategy space, players would rather switch to its complement. Once there, they’d rather switch back.

KEYWORDS: coordination games, dominant strategies, payoff-dominance, nonexistence of equilibrium, tail events.

JEL CLASSIFICATION: C72

SSE/EFI WORKING PAPER SERIES IN ECONOMICS AND FINANCE, NO. 673

THIS VERSION: AUGUST 28, 2007

¹Corresponding author. Department of Economics, Stockholm School of Economics, Box 6501, 113 83 Stockholm, Sweden. E-mail: Mark.Voorneveld@hhs.se. Tel: +46-8-736 92 17. Fax: +46-8-31 32 07.

I thank Kaushik Basu, Stefano DeMichelis, Jörgen Weibull, and Dilan Ölcer for helpful discussions, Anders Björkelid for making Figure 1, and audiences in Tilburg, Stockholm, and Santiago de Compostela for comments. Financial support from the Netherlands Foundation for Scientific Research (NWO) and the Wallander/Hedelius Foundation is gratefully acknowledged.

1. Introduction

Game-theoretic models use infinite player sets as a useful abstraction to model interaction between large numbers of agents. The classical game-theoretic motivation for perfect competition and general-equilibrium analysis models players as a continuum, usually the unit interval with the associated Borel-Lebesgue measure: individual players have no influence on payoffs. On the other hand, it is probably easier to imagine infinitely many people as a *countable* set. For completeness, Section 5.2 recalls standard motivations for using countably infinite player sets. Let us, instead, get straight to the main message: Games with a countably infinite set of players allow for a host of practically relevant phenomena that cannot occur in finite games.

Basu (1994), for instance, considers coordination games — games where all players have the same preferences — and shows that, in contrast with the finite case, dominant-strategy equilibria need not be payoff-dominant:

CLAIM 1. *There are games where*

- (a) *there are no conflicts of interest: all players have the same payoff function,*
- (b) *each player has a strictly dominant strategy,*
- (c) *nevertheless, everybody is better off if no-one chooses the dominant strategy than in the unique Nash equilibrium, where all players choose it.*

This is out of the question in finite coordination games. Start with an arbitrary strategy profile and let players not already doing so switch to the strictly dominant strategy, one at a time. Each switch strictly increases the common payoff function, so the dominant-strategy equilibrium payoff-dominates all other strategy profiles. Example 1.1, simplifying that of Basu (1994, p. 5), establishes the validity of Claim 1.

Example 1.1 Each player $i \in \mathbb{N}$ has two pure strategies, 0 and 1, and payoff function $U : \times_{j \in \mathbb{N}} \{0, 1\} \rightarrow \mathbb{R}$ defined, for each $a \in \times_{j \in \mathbb{N}} \{0, 1\}$, by

$$U(a) = \begin{cases} \sum_{j \in \mathbb{N}} a_j \cdot 2^{-j} & \text{if } \sum_{j \in \mathbb{N}} a_j \text{ is finite,} \\ -2 + \sum_{j \in \mathbb{N}} a_j \cdot 2^{-j} & \text{if } \sum_{j \in \mathbb{N}} a_j \text{ is infinite.} \end{cases}$$

As a unilateral deviation by player i from 0 to 1 increases the payoff by $2^{-i} > 0$, regardless of the choices of the other players, action 1 strictly dominates 0. Hence, $(1, 1, \dots)$ is the unique Nash equilibrium. However, as $U(1, 1, \dots) = -1 < 0 = U(0, 0, \dots)$, the players would have been better off if they all had chosen 0: there is no payoff-dominant Nash equilibrium. \triangleleft

As a practical application of Claim 1, think of a large society of card-carrying utilitarians committed to the same social-welfare function. At the outbreak of an epidemic, an agent's ability to contribute to social welfare depends on his state of health. Penicillin provides a cure or lessens the symptoms due to malign bacteria: when facing the choice between using or abstaining from antibiotics, using them is the dominant option. But if the use of antibiotics becomes a rule, rather than an exception, bacteria quickly develop ways to survive them. This disastrously decreases the power of penicillin as a weapon against microbial foes.

Basu (1994) gives a similar example involving language. Telling the occasional white lie may well be a dominant strategy, but if lies become a routine, language ceases to have power and meaningful communication breaks down. He refers to coordination games with just two actions, satisfying (a) to (c), as “waterfall games”: the lithograph “Waterfall” by the Dutch graphic artist Maurits C. Escher, well-known for his illustrations involving impossible figures, depicts a situation with water seemingly flowing down all the way, yet ending up higher than before.

Notice that we stay as close as possible to Nash's (1951) traditional framework: players have finitely many pure strategies and evaluate outcomes of mixed strategies using a (bounded) von Neumann-Morgenstern utility function. The treatment of the other possibility results in Section 3 follows the same structure: (i) a claim is stated, (ii) it is shown to be impossible in finite games, but (iii) possible in infinite ones. Leaving precise statements to Section 3, the remainder of the introduction will provide an informal discussion of the possibility results and briefly allude to practical applications.

In Claim 1, everybody is better off if no-one chooses the dominant strategy than if everyone does: better off, but not content, since the profile where everyone plays a dominated action is not an equilibrium. The second claim makes the tension between dominant-strategy equilibria and payoff dominance in coordination games even more extreme:

CLAIM 2. There are coordination games where the best outcome occurs if everyone chooses a weakly dominated action, while the worst outcome occurs if everyone chooses the weakly dominant action.

In other words, although choosing a dominant action may seem a bright idea at first sight, it may have disastrous consequences, and everybody is best off by sticking to a dominated option. For the same reason as before, this is impossible in finite coordination games. The idea is illustrated in Figure 1: each step along the way either leaves you at the same height or one stair up the stairway, but still you started at the highest point and ended up at the lowest. Claim 2 shows

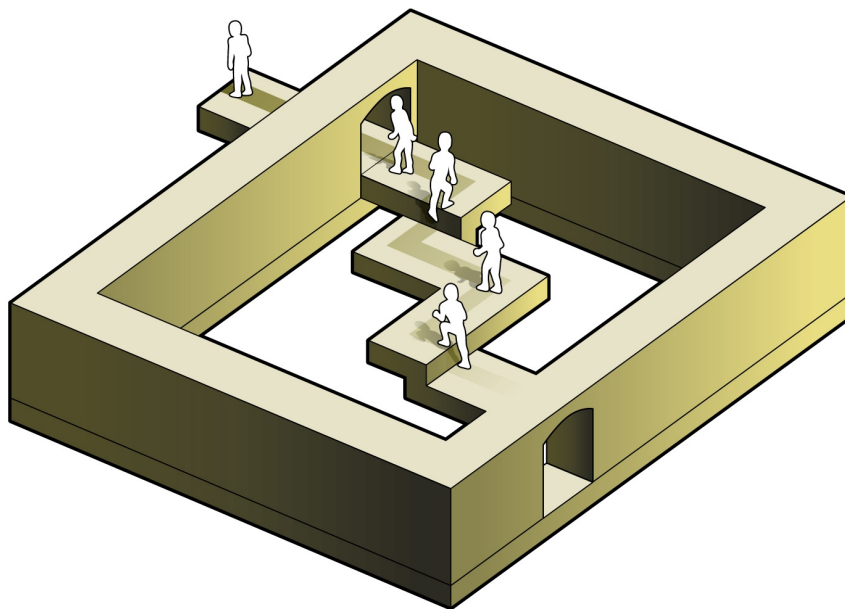


Figure 1: Good, better, . . . , worst: an impossible stairway

that insisting on admissible actions (those that are not weakly dominated) as is common in game theory (Luce and Raiffa, 1957, Kohlberg and Mertens, 1986) can be self-defeating.

Addictive substances provide a potential application. Alcoholics Anonymous (A.A., 1972, p. 7) encourages its members to forget about the past and to stay away from a drink “one day at a time”. This “one day at a time” slogan turns the life of an A.A.-member into a coordination game with independent drink/abstain decisions by a sequence of agents (“A.A.-member at time $t \in \mathbb{N}$ ”) with identical preferences, the A.A.-member’s utility. Call such a coordination game an **A.A.-game**; see Subsection 5.1 for a more detailed discussion, including a comparison with other models of addiction. Consumers of addictive substances may argue that, at any given moment, using the substance weakly dominates not doing so: unless, for instance, their senses are too numbed to notice any difference, it makes them feel slightly better. Claim 2 indicates that this temptation can be reconciled with the more politically correct view that abstinence may be the best for your welfare and heavy addiction the worst.

Claim 3 is concerned with the location of payoff-dominant equilibria. Two coordination games that look so much alike that even the consequences of unilateral deviations are the same may nevertheless have different payoff-dominant equilibria:

CLAIM 3. There are pairs of virtually identical coordination games — only the two

payoff functions differ, yet the consequences of unilateral deviations are the same in both games — with different payoff-dominant equilibria.

For a potential application, consider a pair of A.A.-games, one for each of two distinct addictive goods. Claim 3 indicates that even if, at any given time, the effect of just one more consumption is the same for both commodities, there are beneficial addictions (say, an acquired taste for art) and opposite cases of harmful substances where total abstinence is well-advised.

Finally, even if one allows for mixed strategies, equilibrium existence is at stake:

CLAIM 4. *There are coordination games without (pure or mixed) Nash equilibria.*

The practical relevance is clear: much of economic analysis boils down to formulating a model and studying its equilibria. Section 4 will show that Claim 4 is not just a mathematical curiosity: there are simple assumptions on the payoffs under which games have no equilibria at all.

The remainder of the paper is structured as follows. Section 2 recalls standard game-theoretic definitions. Section 3 states Claims 2 to 4 more precisely and shows that they are impossible in finite, but possible in infinite games. The nonexistence of Nash equilibria is discussed in more detail in Section 4. Section 5 contains concluding remarks.

2. Game theoretic definitions

Our games remain as close as possible to those in Nash (1951): each player has finitely many pure strategies/actions, may use mixed strategies (probability distributions over the pure strategies) and evaluates outcomes using an integrable von Neumann-Morgenstern utility function. The only difference is that games may have a countably infinite set of players, for convenience the set \mathbb{N} of positive integers. We distinguish between these two classes of games by calling them finite (if the set of players is finite) or infinite (otherwise). Formally (Peleg, 1969), a **game** is a tuple $G = \langle N, (A_i)_{i \in N}, (u_i)_{i \in N} \rangle$. The set N is a nonempty, finite or countably infinite, set of players. Each player $i \in N$ has a nonempty, finite set A_i of pure strategies or actions. The set of mixed strategies of player $i \in N$ is

$$\Delta(A_i) := \{ \alpha_i : A_i \rightarrow \mathbb{R} \mid \alpha_i(a_i) \geq 0 \text{ for all } a_i \in A_i, \sum_{a_i \in A_i} \alpha_i(a_i) = 1 \}.$$

The mixed-strategy space is denoted by $\Delta := \times_{i \in N} \Delta(A_i)$. Each $\alpha_i \in \Delta(A_i)$ induces a probability measure on the collection 2^{A_i} of subsets of A_i in the usual way: for each $E \subseteq A_i$, $\alpha_i(E) = \sum_{a_i \in E} \alpha_i(a_i)$. Let (A, \mathcal{A}_P) be the product measurable space $\times_{i \in N} (A_i, 2^{A_i})$. That is, \mathcal{A}_P is the

product σ -algebra on the pure strategy space $A := \times_{i \in N} A_i$ generated by cylinders of the form $\times_{i \in N} E_i$ with $E_i \subseteq A_i$ for each $i \in N$ and $E_i = A_i$ for all but finitely many $i \in N$. Finally, each player $i \in N$ has a real-valued, measurable payoff/utility function u_i on (A, \mathcal{A}_P) , integrable w.r.t. each product measure $\alpha \in \Delta$ on \mathcal{A}_P . The latter assumption is made to avoid ‘trivial’ paradoxical results due to unbounded expected payoffs. As usual, we do not distinguish between a mixed strategy profile $\alpha = (\alpha_i)_{i \in N} \in \Delta$ and the product measure on \mathcal{A}_P induced by α .

In a **coordination game**, all players have the same payoff function: there is a function $U : A \rightarrow \mathbb{R}$ such that $u_i = U$ for all $i \in N$. An equilibrium of a coordination game is **payoff-dominant** if it maximizes the function U .

Let $i \in N$ and $a_i, b_i \in A_i$. Action a_i **weakly dominates** b_i if

$$u_i(a_i, a_{-i}) \geq u_i(b_i, a_{-i}) \quad \text{for all } a_{-i} \in A_{-i},$$

with strict inequality for some a_{-i} , and **strictly dominates** it if all inequalities are strict. An action is **weakly/strictly dominant** if it weakly/strictly dominates all other actions. A **dominant-strategy equilibrium** is one where each player chooses a (weakly/strictly) dominant strategy.

3. Possibility results

In this section, Claims 2 to 4 from the introduction are stated in somewhat more detail. The claims are shown to be false in finite games, but true in infinite ones. Claim 2 shows that the recommended course of action in terms of dominant strategies or payoff-dominance may be diametrically opposed, even if everybody gets the same payoff:

CLAIM 2. *There are games where*

- (a) *there are no conflicts of interest: all players have the same payoff function,*
- (b) *each player has a weakly dominant pure strategy,*
- (c) *nevertheless, the best thing that can happen to the players (a maximizer of the utility function) occurs when everyone chooses a dominated strategy, whereas the worst thing that can happen to the players (a minimizer of the utility function) occurs when everyone chooses the dominant strategy.*

Reasoning as in Claim 1, this is out of the question in finite coordination games. Start with an arbitrary action profile and let players not already doing so switch to the dominant action,

one at a time. Each switch weakly increases the common payoff function, so the payoff in the dominant-strategy equilibrium is maximal. The next example — the mathematical equivalent of Figure 1 — proves Claim 2.

Example 3.1 Each player $i \in \mathbb{N}$ has two pure strategies, 0 and 1, and payoff function $U : \times_{j \in \mathbb{N}} \{0, 1\} \rightarrow \mathbb{R}$ defined, for each $a \in \times_{j \in \mathbb{N}} \{0, 1\}$, by

$$U(a) = \begin{cases} 2 & \text{if } \sum_{j \in \mathbb{N}} a_j \text{ is finite,} \\ 0 & \text{if } \sum_{j \in \mathbb{N}} (1 - a_j) \text{ is finite,} \\ \sum_{j \in \mathbb{N}} a_j \cdot 2^{-j} & \text{otherwise.} \end{cases}$$

As a unilateral deviation by player i from 0 to 1 leaves the payoff unchanged if $\sum_{j \in \mathbb{N}} a_j$ or $\sum_{j \in \mathbb{N}} (1 - a_j)$ is finite and increases it by $2^{-i} > 0$ otherwise, action 1 weakly dominates action 0. However, $U(0, 0, \dots) = 2 = \max_a U(a) > \min_a U(a) = 0 = U(1, 1, \dots)$. \triangleleft

According to Claim 3, coordination games that are so similar that even the consequences of unilateral deviations are the same may nevertheless have different payoff-dominant equilibria:

CLAIM 3. *There are pairs of games, differing only in their payoff functions, where*

- (a) *there are no conflicts of interest: both games are coordination games,*
- (b) *the consequences of unilateral deviations are identical in both games: if the payoff changes by some amount in one game, it changes by the same amount in the other,*
- (c) *their nonempty sets of payoff-dominant equilibria are disjoint.*

This is impossible in finite coordination games. There, the difference between the payoff functions must be a constant, so the payoff-dominant equilibria of the games coincide. To see this, suppose the two games have $n \in \mathbb{N}$ players and payoff function U and V , respectively. Consider two action profiles $a = (a_1, \dots, a_n)$ and $b = (b_1, \dots, b_n)$. We show that $U(a) - V(a) = U(b) - V(b)$. Start with a and let the first player deviate to b_1 . The consequences of unilateral deviations are required to be identical in both games:

$$U(a) - U(b_1, a_2, \dots, a_n) = V(a) - V(b_1, a_2, \dots, a_n),$$

or, after rearranging terms,

$$U(a) - V(a) = U(b_1, a_2, \dots, a_n) - V(b_1, a_2, \dots, a_n).$$

Repeating this argument n times, each time letting the next player deviate unilaterally to his action in b gives

$$\begin{aligned}
U(a) - V(a) &= U(b_1, a_2, \dots, a_n) - V(b_1, a_2, \dots, a_n) \\
&= U(b_1, b_2, a_3, \dots, a_n) - V(b_1, b_2, a_3, \dots, a_n) \\
&= \dots \\
&= U(b) - V(b).
\end{aligned}$$

The next example proves Claim 3.

Example 3.2 Let $c \in \mathbb{R}$. Each player $i \in \mathbb{N}$ has two pure strategies, 0 and 1, and payoff function $U_c : \times_{j \in \mathbb{N}} \{0, 1\} \rightarrow \mathbb{R}$ defined, for each $a \in \times_{j \in \mathbb{N}} \{0, 1\}$, by

$$U_c(a) = \begin{cases} c - \sum_{j \in \mathbb{N}} a_j \cdot 2^{-j} & \text{if } \sum_{j \in \mathbb{N}} a_j \text{ is finite,} \\ \sum_{j \in \mathbb{N}} a_j \cdot 2^{-j} & \text{if } \sum_{j \in \mathbb{N}} a_j \text{ is infinite.} \end{cases}$$

Given an arbitrary mixed-strategy profile, Kolmogorov's Zero-One Law (cf. Billingsley, 1995, p. 287, or Williams, 1991, p. 46) implies that the sum $\sum_{j \in \mathbb{N}} a_j$ converges almost surely or diverges almost surely. In the former case, each player's unique best reply is to choose action 0 with probability one. In the latter case, each player's unique best reply is to choose action 1 with probability one. Therefore, the game has two Nash equilibria, both in pure strategies: $(0, 0, \dots)$ and $(1, 1, \dots)$. For each $a \in \times_{j \in \mathbb{N}} \{0, 1\}$,

$$U_c(a) \leq \max\{U(0, 0, \dots), U(1, 1, \dots)\} = \max\{c, 1\},$$

so $(0, 0, \dots)$ is payoff dominant if $c > 1$ and $(1, 1, \dots)$ is payoff dominant if $c < 1$: the location of the payoff-dominant equilibrium changes depending on c . The consequences of unilateral deviations do not: for each player $i \in \mathbb{N}$ and each action profile a_{-i} of his fellow players,

$$U_c(0, a_{-i}) - U_c(1, a_{-i}) = \begin{cases} 2^{-i} & \text{if } \sum_{j \in \mathbb{N}} a_j \text{ is finite,} \\ -2^{-i} & \text{if } \sum_{j \in \mathbb{N}} a_j \text{ is infinite,} \end{cases}$$

is independent of c . In other words, coordination games with payoff function U_c for some $c \in \mathbb{R}$ are identical in terms of unilateral deviations and consequently have identical best-response correspondences, but may have different payoff-dominant equilibria. \triangleleft

But life gets worse: there may be no equilibria at all.

CLAIM 4. *There are coordination games without (pure or mixed) Nash equilibria.*

Of course, Nash’s (1951) celebrated existence result rules this out in the finite setting. The next example proves Claim 4.

Example 3.3 Each player $i \in \mathbb{N}$ has two pure strategies, 0 and 1, and payoff function $U : \times_{j \in \mathbb{N}} \{0, 1\} \rightarrow \mathbb{R}$ defined, for each $a \in \times_{j \in \mathbb{N}} \{0, 1\}$, by

$$U(a) = \begin{cases} \sum_{j \in \mathbb{N}} a_j \cdot 2^{-j} & \text{if } \sum_{j \in \mathbb{N}} a_j \text{ is finite,} \\ -\sum_{j \in \mathbb{N}} a_j \cdot 2^{-j} & \text{if } \sum_{j \in \mathbb{N}} a_j \text{ is infinite.} \end{cases}$$

Consider an arbitrary mixed-strategy profile. By Kolmogorov’s Zero-One Law, $\sum_{j \in \mathbb{N}} a_j$ converges almost surely or diverges almost surely. In the first case, each player strictly prefers action 1. As the sum converges almost surely, there are players not choosing this unique best reply. In the second case, each player strictly prefers action 0. As the sum diverges almost surely, there are players not choosing this unique best reply. Hence, there is no Nash equilibrium. \triangleleft

Section 4 shows that Claim 4 is not just a mathematical curiosity: there are simple assumptions on the payoffs under which games have no equilibria at all.

4. Tail events and greener-grass games

What drives the previous results is that payoffs are defined in terms of whether or not infinitely many people choose a certain action. Such events are common in everyday issues: “If we don’t eventually stop using oil/fishing for cod in the Baltic Sea/..., a catastrophe will occur.” Individual players and, by induction, an arbitrary finite number of players cannot affect the occurrence of such events: if, for instance, only finitely many people choose a certain action, nothing you — or for that matter, any finite number of players — can do, will make this number infinite. In the language of probability theory, events of this type are called tail events. Subsection 4.1 provides preliminaries on tail events; see Billingsley (1995) or Williams (1991) for further details. Subsection 4.2 shows that tail events easily give rise to games without (pure or mixed) Nash equilibria. This is done by constructing, for each non-trivial tail event, a so-called **greener-grass game**, following the English proverb that “the grass is always greener on the other side of the hedge”: if the tail event, say E , occurs (with probability one), players would rather jump to a strategy profile outside of E , and vice versa.

4.1. Tail events

Let \mathbb{N} be the player set. As each player $i \in \mathbb{N}$ has a finite set A_i of actions, we can assume w.l.o.g. that A_i is a finite subset of \mathbb{R} ; the examples use $A_i = \{0, 1\}$. For each $i \in \mathbb{N}$, let $X_i : A \rightarrow \mathbb{R}$ be the random variable specifying player i 's action: $X_i(a) = a_i$ for each $a \in A$. The tail σ -algebra is the σ -algebra

$$\mathcal{A}_T := \bigcap_{i=1}^{\infty} \sigma(\{X_j : j \geq i\}),$$

where $\sigma(\{X_j : j \geq i\})$ is the σ -algebra generated by the random variables X_i, X_{i+1}, \dots . An event $E \in \mathcal{A}_T$ is a **tail event**. By definition, whether a tail event occurs is independent of the realization of finitely many of the X_i 's: only tails matter. As all X_i are measurable in the product σ -algebra \mathcal{A}_P , it follows that $\mathcal{A}_T \subset \mathcal{A}_P$.

A convenient tool for tail events is Kolmogorov's Zero-One Law, already used in Examples 3.2 and 3.3: if a tail event is defined in terms of independent stochastic variables (and the players choose their strategies independently), it occurs with probability zero or with probability one.

A sufficient condition (Peleg, 1969) for the existence of mixed-strategy Nash equilibria in games as defined above is that payoffs are continuous in the product topology \mathcal{T}_P on the pure strategy space. Endowing each A_i with the discrete topology 2^{A_i} , the product topology \mathcal{T}_P is the topology whose basis consists of the cylinders of the form $\times_{i \in N} E_i$ with $E_i \subseteq A_i$ for each $i \in N$ and $E_i = A_i$ for all but finitely many $i \in N$.

The product topology and the tail σ -algebra, however, have only trivial sets in common:

Proposition 4.1 *A tail event $E \in \mathcal{A}_T$, other than \emptyset or the entire pure-strategy space A , is neither open nor closed in the product topology.*

Proof. Let $E \in \mathcal{A}_T \setminus \{\emptyset, A\}$. Suppose E is closed, i.e., its complement E^c is open: $E^c \in \mathcal{T}_P$. Let $x \in E, y \in E^c$. This is possible, as $E \notin \{\emptyset, A\}$. Construct a sequence $(s_n)_{n \in \mathbb{N}}$ in A as follows. For each $n \in \mathbb{N}$, the strategy profile $s_n = (s_{n1}, s_{n2}, \dots)$ has coordinates $s_{nk} = y_k$ if $k \leq n$ and $s_{nk} = x_k$ if $k > n$. That is, s_n is obtained from x by exchanging its first n coordinates by those of y . Then $\lim_{n \rightarrow \infty} s_n = y \in E^c$. However, as x lies in the tail event E , and each s_n differs from x only in the first n coordinates, also $s_n \in E$. As $(s_n)_{n \in \mathbb{N}}$ is a convergent sequence in the closed set E , its limit y must lie in E as well. Contradiction.

Next, suppose E is open, i.e., E^c is closed. By definition of a σ -algebra, also $E^c \in \mathcal{A}_T$. By the previous step, no such closed set exists. \square

Consequently, payoff functions defined in terms of tail events are typically not continuous in the product topology. This violation of Peleg's (1969) sufficient condition for equilibrium existence is exploited in Proposition 4.2 to provide a simple construction of a large class of games without Nash equilibria.

4.2. Greener-grass games

As the saying goes, the grass is always greener on the other side of the hedge. This dissatisfaction, no matter what happens, is the intuition behind our **greener-grass games**. Given some tail event E , other than \emptyset or the entire pure-strategy space A , we construct a game G_E without Nash equilibria by making sure that if the tail event E occurs (with probability one), the players would rather jump to a strategy profile in its complement E^c , and vice versa.

Proposition 4.2 *For each $i \in \mathbb{N}$, let $A_i \subset \mathbb{R}$ be a finite set of (at least two) pure strategies. Let $E \in \mathcal{A}_T$ be a tail event, $E \notin \{\emptyset, A\}$. Then there is a greener-grass game $G_E = \langle \mathbb{N}, (A_i)_{i \in \mathbb{N}}, (u_i)_{i \in \mathbb{N}} \rangle$ without (pure or mixed) Nash equilibria.*

Proof. Let $x \in E, y \in E^c$. This is possible, as $E \notin \{\emptyset, A\}$. For each $i \in \mathbb{N}$, define $u_i : A \rightarrow \mathbb{R}$ by

$$u_i(a) = \begin{cases} 1 & \text{if } (a \in E \text{ and } a_i = y_i) \text{ or } (a \in E^c \text{ and } a_i = x_i), \\ 0 & \text{otherwise.} \end{cases} \quad (1)$$

The game $G_E = \langle \mathbb{N}, (A_i)_{i \in \mathbb{N}}, (u_i)_{i \in \mathbb{N}} \rangle$ has no Nash equilibria. Indeed, given an arbitrary mixed-strategy profile, Kolmogorov's Zero-One Law implies that the tail event E either occurs with probability one or with probability zero. Hence, each player $i \in \mathbb{N}$ has a unique best reply: pure strategy y_i in the first case and pure strategy x_i in the second. So a Nash equilibrium must be in pure strategies. But pure Nash equilibria do not exist. Suppose $a \in A$ is a Nash equilibrium. If $a \in E$, the unique best reply of each $i \in \mathbb{N}$ is y_i : $a_i = y_i$ for all $i \in \mathbb{N}$, i.e., $a = y \notin E$, a contradiction. Similarly, if $a \in E^c$, then $a_i = x_i$ for all $i \in \mathbb{N}$, i.e., $a = x \in E$, a contradiction. \square

A possible reaction to the nonexistence of Nash equilibria is that one may not be able to make everybody content, but that there may be approximate equilibria (Radner, 1980), where players get a payoff pretty close to that of a best response. Consider, for instance, Example 3.3. Let $\varepsilon > 0$. An **approximate** or **ε -equilibrium** is a strategy profile in which players cannot gain more than ε from unilateral deviation. As the effect, in absolute value, of a unilateral deviation by an arbitrary player j is at most 2^{-j} , players $j \in \mathbb{N}$ with $2^{-j} < \varepsilon$ automatically receive a

payoff that is at most ε away from their best response. Assign to these players an arbitrary mixed strategy. The reduced game between the players j with $2^{-j} \geq \varepsilon$ has only finitely many players and therefore a Nash equilibrium. Combining this with the previously fixed strategies of players with a large index gives an ε -equilibrium of the game.

However, the payoffs in the examples were chosen for mathematical convenience only. It is not difficult to construct games where even approximate equilibria do not exist:

Example 4.3 Each player $i \in \mathbb{N}$ has two pure strategies, 0 and 1, and payoff function $u_i : \times_{j \in \mathbb{N}} \{0, 1\} \rightarrow \mathbb{R}$ defined, for each $a \in \times_{j \in \mathbb{N}} \{0, 1\}$, by

$$u_i(a) = \begin{cases} 1 & \text{if } (\limsup_{k \rightarrow \infty} \frac{a_1 + \dots + a_k}{k} > 1/2 \text{ and } a_i = 0) \\ & \text{or } (\limsup_{k \rightarrow \infty} \frac{a_1 + \dots + a_k}{k} \leq 1/2 \text{ and } a_i = 1), \\ 0 & \text{otherwise.} \end{cases}$$

Let $\varepsilon \in (0, 1/2)$ and let $\alpha \in \Delta$ be a mixed-strategy profile. By Kolmogorov's Zero-One Law, the tail event $\{a \in A : \limsup_{k \rightarrow \infty} \frac{a_1 + \dots + a_k}{k} > 1/2\}$ occurs with probability zero or one. In the first case, choosing action 0 gives payoff 0 and action 1 gives payoff 1, so if α is an ε -equilibrium, each player chooses action 1 with probability at least $1 - \varepsilon > 1/2$. This contradicts that the tail event occurs with probability zero. Similarly, in the second case, choosing action 0 gives payoff 1 and action 1 gives payoff 0, so if α is an ε -equilibrium, each player chooses action 1 with probability at most $\varepsilon < 1/2$. This contradicts that the tail event occurs with probability one. Hence, there is no ε -equilibrium, let alone a Nash equilibrium. \triangleleft

As an aside, Hart and Schmeidler (1989, Example 2) show that correlated equilibria need not exist either.

5. Concluding remarks

5.1. A.A.-games and other models of addiction

Suppose a decision-maker has to decide at each moment in discrete time whether to take a drink (action 1) or not (action 0). Given the person's uncertain life-length, common modeling practice is to treat such problems as having an infinite horizon, properly discounting consequences of future actions. This makes the decision-maker's utility function $U : \times_{j \in \mathbb{N}} \{0, 1\} \rightarrow \mathbb{R}$ of the additively separable type used in the examples in Sections 1 and 3; discounting by powers of $1/2$ was for mathematical convenience only.

The philosophy of Alcoholics Anonymous is to deal with the temptations of alcohol by treating each decision independently, “one day at a time”. Thus, one can model these independent decisions as taken by a sequence of agents representing the A.A.-member at different times $t \in \mathbb{N}$, all with his best interests, measured by his utility function U , in mind. This results in the A.A.-games from the introduction: coordination games with a countably infinite player set as a model of a single decision-maker with an uncertain life-span making a series of independent choices in discrete time.

A.A.-games differ from existing models of addiction in a number of ways. Firstly, those following the work of Becker and Murphy (1988) model consumption of addictive goods as a single decision, maximizing life-time utility. Secondly, those that do use a multiple-selves approach often rely on the work of Strotz (1955) with preferences changing over time due to hyperbolic discounting. Our games illustrate that there is room for counterintuitive results like Claim 2 already in the context of constant preferences.

The model of Basu (2000, Chapter 3) is a symbiosis: he sometimes speaks of act-optimizing behavior by independent agents (our A.A.-game), sometimes of rule-optimizing behavior by a single agent (life-time utility maximization).

5.2. *Other motivations for countable player sets*

For completeness, we summarize some other standard motivations for using countably infinite player sets. Firstly, if we consider all past, current, and future generations, the number of people may well be infinite (Basu, 1994). Secondly, the framework with a countably infinite number of players facilitates modeling situations where the population is large, but the role of individual players cannot be ignored. In contrast, in the standard nonatomic model, the player set is the unit interval $[0, 1]$ endowed with its usual Lebesgue measure, and individuals (having measure zero) are insignificant. Thirdly, authors like Myerson (1998) and Milchtaich (2004) model interaction under population uncertainty. Examples include Internet auctions, large elections, and traffic congestion. They argue that even if you are convinced that the number of fellow players is finite, you may not know an upper bound, i.e., you may never be quite sure that there isn’t just one more player in the game. In such cases, the underlying set of players is appropriately modeled as infinite. This is analogous to the motivation that some repeated games, even if they objectively speaking end after a finite number of rounds, are most appropriately modeled as infinitely repeated games. Such an infinite horizon is appropriate if players, after each round, believe — possibly erroneously — that there may be an additional

round. Fourthly, one commonly defines equilibria of games with incomplete information as Nash equilibria of a strategic game where players are identified by their ‘name’ and their type. Hence, even if the number of players of a Bayesian game is finite, allowing for a countably infinite state/type space introduces an infinite number of players in the associated strategic game.

5.3. Greener-grass games are potential games

The construction of greener-grass games in the proof of Proposition 4.2 gives players different payoff functions. We can restrict attention to coordination games: the next result shows that players act *as if* they all have the same payoff function. In terms of Monderer and Shapley (1996), greener-grass games are ordinal potential games.

Proposition 5.1 *Each greener-grass game $G_E = \langle \mathbb{N}, (A_i)_{i \in \mathbb{N}}, (u_i)_{i \in \mathbb{N}} \rangle$ constructed in Proposition 4.2 is an ordinal potential game: the function $U : A \rightarrow \mathbb{R}$ with $U = \sum_{j \in \mathbb{N}} 2^{-j} u_j$ is such that for all $i \in \mathbb{N}$, all $a_{-i} \in A_{-i}$, and all $a_i, b_i \in A_i$:*

$$u_i(a_i, a_{-i}) - u_i(b_i, a_{-i}) > 0 \quad \Leftrightarrow \quad U(a_i, a_{-i}) - U(b_i, a_{-i}) > 0. \quad (2)$$

Also the coordination game $\langle \mathbb{N}, (A_i)_{i \in \mathbb{N}}, (U)_{i \in \mathbb{N}} \rangle$ has an empty set of Nash equilibria.

Proof. Let $i \in \mathbb{N}$, $a_{-i} \in A_{-i}$, and $a_i, b_i \in A_i$. By definition, tail events do not depend on single coordinates: $(a_i, a_{-i}) \in E$ if and only if $(b_i, a_{-i}) \in E$. Substitution in (1) gives that $u_j(a_i, a_{-i}) = u_j(b_i, a_{-i})$ for all $j \in \mathbb{N} \setminus \{i\}$. By definition of U :

$$\begin{aligned} U(a_i, a_{-i}) - U(b_i, a_{-i}) &= \sum_{j \in \mathbb{N}} 2^{-j} [u_j(a_i, a_{-i}) - u_j(b_i, a_{-i})] \\ &= 2^{-i} [u_i(a_i, a_{-i}) - u_i(b_i, a_{-i})] > 0 \\ &\Leftrightarrow u_i(a_i, a_{-i}) - u_i(b_i, a_{-i}) > 0. \end{aligned}$$

The game $\langle \mathbb{N}, (A_i)_{i \in \mathbb{N}}, (U)_{i \in \mathbb{N}} \rangle$ has no Nash equilibria: reasoning as in the proof of Proposition 4.2, Kolmogorov’s Zero-One Law allows us to restrict attention to pure Nash equilibria. By (2), these must be pure Nash equilibria of the greener-grass game G_E , which has none. \square

References

Alcoholics Anonymous World Services, 1972. A brief guide to Alcoholics Anonymous. New York.

- Basu, K., 1994. Group rationality, utilitarianism, and Escher's Waterfall. *Games and Economic Behavior* 7, 1–9.
- Basu, K., 2000. *Prelude to political economy*. Oxford, UK: Oxford University Press.
- Becker, G.S., Murphy, K.M., 1988. A theory of rational addiction. *Journal of Political Economy* 96, 675–700.
- Billingsley, P., 1995. *Probability and Measure*, 3rd ed., New York: John Wiley & Sons.
- Hart, S., Schmeidler, D., 1989. Existence of correlated equilibria. *Mathematics of Operations Research* 14, 18–25.
- Kohlberg, E., Mertens, J.-F., 1986. On the strategic stability of equilibria. *Econometrica* 54, 1003–1037.
- Luce, R.D., Raiffa, H., 1957. *Games and decisions*. New York: John Wiley and Sons.
- Milchtaich, I., 2004. Random-player games. *Games and Economic Behavior* 47, 353–388.
- Monderer, D., Shapley, L.S., 1996. Potential games. *Games and Economic Behavior* 14, 124–143.
- Myerson, R.B., 1998. Population uncertainty and Poisson games. *International Journal of Game Theory* 27, 375–392.
- Nash, J.F., 1951. Non-cooperative games. *Annals of Mathematics* 54, 286–295.
- Peleg, B., 1969. Equilibrium points for games with infinitely many players. *Journal of the London Mathematical Society* 44, 292–294.
- Radner, R., 1980. Collusive behaviour in noncooperative epsilon-equilibria of oligopolies with long but finite lives. *Journal of Economic Theory* 22, 136–154.
- Strotz, R. H., 1955. Myopia and inconsistency in dynamic utility maximization. *Review of Economic Studies* 23, 165–180.
- Williams, D., 1991. *Probability with martingales*. Cambridge, UK: Cambridge University Press.