

Equilibrium Paths in Discounted Supergames

Kimmo Berg

*Systems Analysis Laboratory, Aalto University School of Science, P.O. Box 11100,
FI-00076 Aalto, Finland*

Mitri Kitti*

Department of Economics, University of Turku, FI-20014 Turku, Finland

Abstract

This paper characterizes the subgame-perfect pure-strategy equilibrium paths in discounted supergames with perfect monitoring. It is shown that all the equilibrium paths are composed of fragments called elementary subpaths. This characterization result is complemented with an algorithm for finding the elementary subpaths. By using these subpaths it is possible to generate equilibrium paths and payoffs. When there are finitely many elementary subpaths, all the equilibrium paths can be represented by a directed graph. These graphs can be used in analyzing the complexity of equilibrium outcomes. In particular, it is shown that the size and the density of the equilibrium set can be measured by the asymptotic growth rate of equilibrium paths and the Hausdorff dimension of the payoff set.

Keywords: game theory; repeated game, subgame-perfect equilibrium, equilibrium path, graph, complexity

1. Introduction

Repeated games provide the most elementary setting for analyzing dynamic interactions among self-interested agents. We consider the case where a stage game is repeated infinitely many times, players discount the future payoffs, observe perfectly each others' actions, and use pure strategies. These games have usually enormously rich sets of equilibrium strategies, which is generally thought to imply that the outcomes are hard to predict. Contrary to this intuition, we show that all the equilibrium paths are generated from a collection of subpaths. By equilibrium paths we mean infinite sequences of players' actions that are induced by subgame-perfect equilibrium strategies.

*Corresponding author. Tel.: +358 2 333 6438

Email addresses: kimmo.berg@aalto.fi (Kimmo Berg), mitri.kitti@gmail.com (Mitri Kitti)

As shown in [1, 2], an equilibrium path is such that none of the players has an incentive to deviate at any stage when the deviations lead to the paths that provide the smallest equilibrium payoff for the deviator. This idea of most severe punishments can also be utilized for characterizing the equilibrium payoffs with a set-valued fixed point equation, see [3, 4] for the case of imperfect monitoring and [10] for perfect monitoring. These results entail that in equilibrium the players take actions that are incentive compatible given the future payoffs of the strategy and the threat of receiving the smallest equilibrium payoffs after deviations.

We derive a novel characterization for equilibrium paths from the players' incentive compatibility conditions. Our main result is that all the equilibrium paths are constructed from a collection of sequences of players' actions, which will be called the elementary subpaths. We emphasize that this result characterizes all equilibrium paths simultaneously rather than gives a condition for individual paths as done in [1, 2]. We also present an algorithm for producing the elementary subpaths, and show that the equilibrium paths can be compactly represented by a directed graph. The graph offers a simple way to produce equilibrium paths and payoffs. Unlike in the literature on combinatorial games (for a survey, see [14]), where graphs are commonly used to present the players' available moves at different positions, we use graphs to describe equilibrium behavior.

The graph presentation can be used in analyzing the complexity of equilibrium paths and payoffs. In particular, we propose the asymptotic growth rate of the equilibrium paths as a measure for the complexity of equilibrium behavior. This measure tells us the rate at which the number of finitely long equilibrium paths grows when their length is increased. Since the paths of action profiles are given by strategies, the growth rate reflects the size of the equilibrium set and the increase of strategies producing the finitely long equilibrium paths.

Our approach to complexity of repeated game equilibria is new and differs from the previous literature on strategic complexity, see [17] and [8] for surveys, and on the complexity of finding winning strategies, see, e.g., [12]. We analyze the complexity of all equilibrium outcomes without relying on the complexity of individual strategies nor their computation. In particular, we can assess the complexity of different repeated games in terms of the equilibrium behavior.

The graph enables us to analyze the complexity of equilibrium payoffs in addition to paths. The payoff set can be identified as a particular fractal, whose complexity can be measured by the Hausdorff dimension determined by the graph; see [7] for more detailed analysis on the fractal properties. The phenomenon that the payoff set behaves in a rather complex manner, as fractals do, is not completely new, see [18] and [20]. We offer a more comprehensive view to the structure of equilibria: when the discount factors vary, the elementary subpaths change, which affects the graph that generates the payoffs.

The results of this paper can also be used in computing equilibrium payoffs; in [6] the methodology is applied to 2×2 games. Related computational methods have been developed in [9] and [16], who assume equal discount factors and correlated strategies; see also [5]. The characterization of equilibrium paths

enables finding the exact payoffs given by pure strategies regardless of whether the discount factors are equal or not. Hence, our methods complement the literature on the computation of supergame equilibria.

The paper is structured as follows. In Section 2 we show that the equilibrium paths consist of elementary subpaths. Section 3 deals with the computation and analysis of finite elementary sets. Illustrative examples of the main ideas and the computational methods are presented in Section 4. Conclusions are given in Section 5.

2. Equilibrium Paths and Subpaths

2.1. Notation and Definitions

We assume that there are n players, and $N = \{1, \dots, n\}$ denotes the set of players. The set of actions available for player i in the stage game is A_i . Each player is assumed to have finitely many actions. The set of action profiles is denoted by $A = \times_i A_i$. Moreover, a_{-i} denotes the action profile of players other than player i . The corresponding set of action profiles is $A_{-i} = \times_{j \neq i} A_j$. Function $u : A \mapsto \mathbb{R}^n$ gives the vector of payoffs that the players receive in the stage game when a given action profile is played, i.e., when $a \in A$ is played, player i receives payoff $u_i(a)$.

In the supergame the stage game is repeated infinitely many times, and the players discount the future payoffs with discount factors δ_i , $i \in N$. We assume perfect monitoring: all players observe the action profile played at the end of each period. A history contains the path of action profiles that have previously been played. The set of length k histories or paths is denoted by $A^k = \times_k A$. The empty path is \emptyset , i.e., $A^0 = \{\emptyset\}$. Infinitely long paths are denoted by A^∞ . When referring to the set of paths beginning with a given action profile a we use $A^k(a)$ and $A^\infty(a)$ for length k paths and infinitely long paths, respectively. Moreover, \mathcal{A} is the set of all paths, finite or infinite, and $\mathcal{A}(a)$ is the set of all paths that start with a , i.e., union of $A^k(a)$, $k = 1, 2, \dots$ and $A^\infty(a)$.

The length of path p is denoted by $|p|$. Furthermore, $i(p)$ is the initial and $f(p)$ is the final element of p . If p is infinitely long, in brief an infinite subpath, then $f(p) = \emptyset$. If p and p' are two paths then pp' is the path obtained by juxtaposing the terms of p and p' . For $p \in \mathcal{A}$, we let p_j denote the path that starts from the element $j + 1$ of p . Respectively, p^k is the path of first k elements of p . More specifically, when $p = a^0 a^1 \dots$, we have $p_1 = a^1 a^2 \dots$, $p^k = a^0 \dots a^{k-1}$, and $p_j^k = a^j \dots a^{j+k-1}$.

A strategy for player i in the supergame is a sequence of mappings $\sigma_i^0, \sigma_i^1, \dots$ where $\sigma_i^k : A^k \mapsto A_i$. The set of strategies for player i is Σ_i . The strategy profile composed of $\sigma_1, \dots, \sigma_n$ is denoted by σ . Given a strategy profile σ and a path p , the restriction of the strategy profile after p is $\sigma|p$. The outcome path induced by σ is $(a^0(\sigma), a^1(\sigma), \dots) \in A^\infty$, where $a^k(\sigma) = \sigma^k(a^0(\sigma) \dots a^{k-1}(\sigma))$ for all k .

The average discounted payoff for player i corresponding to a strategy profile σ is

$$U_i(\sigma) = (1 - \delta_i) \sum_{k=0}^{\infty} \delta_i^k u_i(a^k(\sigma)). \quad (1)$$

Subgame perfection is defined in the usual way; σ is a subgame-perfect equilibrium (SPE) of the supergame if

$$U_i(\sigma|p) \geq U_i(\sigma'_i, \sigma_{-i}|p) \text{ for all } i \in N, p \in A^k, k \geq 0, \text{ and } \sigma'_i \in \Sigma_i.$$

This paper focuses on SPE paths and subpaths defined as below.

Definition 1. A path $p \in A^\infty$ is a subgame-perfect equilibrium path if there is an SPE strategy profile that induces it.

Definition 2. A path $p' \in \mathcal{A}(a)$ is an SPE subpath if there is an SPE path $p \in A^\infty(a)$ such that $p|^{p'} = p'$.

We shall derive a characterization for equilibrium subpaths by assuming that the set of subgame-perfect equilibrium payoffs is known. Eventually, it turns out that all we need to know from the set of equilibrium payoffs are the players' smallest payoffs. Hence, our approach is analogous with the way how the Euler equation is derived for the optimal paths in dynamic programming.

In the following, V denotes the set of equilibrium payoffs. It is assumed that V is non-empty, in which case it will also be a compact subset of \mathbb{R}^n [10]. The smallest equilibrium payoff for player i will be denoted by $v_i^-(V) = \min \{v_i : v \in V\}$. A pair (a, v) of an action profile $a \in A$ and a continuation payoff $v \in V$ is admissible with respect to V if it satisfies the incentive compatibility condition

$$(1 - \delta_i)u_i(a) + \delta_i v_i \geq \max_{a'_i \in A_i} [(1 - \delta_i)u_i(a'_i, a_{-i}) + \delta_i v_i^-(V)] \quad \forall i \in N. \quad (2)$$

This constraint says that it is better for player i to take the action a_i and get the payoffs v_i than to deviate and then obtain $v_i^-(V)$.

In the following, $C_a(V)$ denotes the set of payoffs for which the pair (a, v) is admissible. Note that the vector of the smallest payoffs that make (a, v) admissible can be found from the incentive compatibility condition. We let $\text{con}(a)$ denote this vector. It is the payoff vector in which the component $\text{con}_i(a)$, $i \in N$, is the maximum of $v_i^-(V)$ and the solution v_i of

$$(1 - \delta_i)u_i(a) + \delta_i v_i = \max_{a'_i \in A_i} [(1 - \delta_i)u_i(a'_i, a_{-i}) + \delta_i v_i^-(V)].$$

Note that $C_a(V) = \{v \in V : v \geq \text{con}(a)\}$, where the inequality means that $v_i \geq \text{con}_i(a)$ for all $i \in N$.

We define an affine mapping $B_a : \mathbb{R}^n \mapsto \mathbb{R}^n$ corresponding to an action profile $a \in A$ by setting

$$B_a(v) = (I - T)u(a) + Tv,$$

where I is $n \times n$ identity matrix and T is a $n \times n$ diagonal matrix with discount factors $\delta_1, \dots, \delta_n$ on the diagonal. These mappings are contractions because the discount factors are less than one.

2.2. Elementary Subpaths

Let us consider the set of tail payoffs that are possible after an action profile a when it begins a path $p \in \mathcal{A}(a)$. We also assume that p is an SPE subpath. If p is infinite then it is an equilibrium path itself. First, we can observe that a should be followed by a payoff that belongs to $C_a(V)$. As a is followed by p_1 , we need to consider what are the payoffs that $i(p_1)$ generates from the set of tail payoffs that are possible for $i(p_1)$ when it is followed by p_2 .

Let $W(p)$ denote the set of continuation payoffs that are possible for a after $p \in \mathcal{A}(a)$. To be more specific, when the first action profile is a , i.e., the first element of p , then $W(p)$ contains all the possible continuation payoffs that may follow a . The set $W(p)$ satisfies the recursion

$$W(p) = C_{i(p)}(V) \cap B_{i(p_1)}(W(p_1)).$$

Namely, the continuation payoff for a should belong to $C_a(V)$ and it should be generated by $i(p_1)$ from $W(p_1)$. Note that for $W(p)$ we need $W(p_1)$ for which we need $W(p_2)$, and so on. In particular, if $|p| = \infty$, the recursion is infinite and $W(p)$ becomes a singleton: the vector of average discounted payoffs corresponding to p . To complete the definition of $W(p)$ we set $W(\emptyset) = V$ and $B_\emptyset = I$. This is needed because p_1 is an empty path when $|p| = 1$. In particular, if $|p| = k$, then $p_{k-1} = f(p)$ and $|p_{k-1}| = 1$, which imply $W(p_{k-1}) = C_{f(p)}(V)$.

The following observations on $W(p)$ will form the basis for our definition of elementary subpaths. The first is that p is an SPE subpath if and only if $W(p) \neq \emptyset$, i.e., the first element of p has a non-empty set of possible continuation payoffs. The second, and more important observation is that if $W(p_1) \neq \emptyset$ and

$$B_{i(p_1)}(W(p_1)) \subseteq C_{i(p)}(V), \quad (3)$$

then at the time when the first action profile a of p is played, any SPE subpath starting with the final element of p , i.e., $f(p)$, is a possible continuation for p . For example, if abc is an SPE subpath such that $W(bc) \neq \emptyset$ and $B_b(W(bc)) \subseteq C_a(V)$, then at the time when a is played and it is known that bc will follow, any SPE subpath that begins with c is possible. However, at the time when b is played, it may matter how the path abc continues after c is played. For instance, the action profile b may require that it is followed by a path cc , which implies that abc must be followed by another c . At the time when a is played, the second c may not be required. Note also that if $V \subseteq C_a(V)$, i.e., $V = C_a(V)$, then a can be followed by any SPE subpath.

The subpaths that satisfy (3) and have no shorter subpaths satisfying this condition are elementary.

Definition 3. If $p \in \mathcal{A}(a)$ satisfies $W(p_1) \neq \emptyset$, condition (3), and there is no $k < |p|$, such that p^k satisfies these conditions, i.e., p^k is not an elementary subpath, then p is an elementary subpath, and we denote $p \in P^{|p|}(a)$.

The requirement that none of the first k elements satisfy (3) means that set P^k does not contain paths that already belong to sets P^1, \dots, P^{k-1} . For

example, if abc is an elementary subpath, then $abcd$ cannot be an elementary subpath even though it may satisfy (3). Note that infinitely long subpaths can be elementary. If $p \in P^\infty(a)$, then sets $B_{i(p_1)}(W(p_1^k))$, $k = 1, 2, \dots$, contain payoff vectors that do not belong to $C_a(V)$.

Our main result is that all the SPE paths are characterized by elementary subpaths, i.e., sets P^k and P^∞ .

Proposition 1. *A path $p \in A^\infty(a)$ is an SPE path if and only if for all $j \in \mathbb{N}$ either $p_j^k \in P^k(i(p_j^k))$ for some k or $p_j \in P^\infty(i(p_j))$.*

Proof. By the construction of P^k 's, an SPE path p satisfies one of the two conditions of the proposition.

Let us assume that for all j either $p_j^k \in P^k(i(p_j^k))$ for some k , or $p_j \in P^\infty(i(p_j))$. In that case, for any j the payoff to player i is at least $v_i^-(V)$ when the players choose action profiles such that they stay on the subpath p_j^k or p_j . We first argue that in the case when there is k such that $p_j^k \in P^k(i(p_j^k))$, the threat of reverting to the path that yields $v_i^-(V)$ to the deviator i , keeps the players on path p . In that case it does not matter for players at stage j what happens after k periods as long as the continuation payoff after these periods is in V and the penalty from deviations is $v_i^-(V)$. On the other hand, if $p_j \in P^\infty(i(p_j))$ then the players do not have any incentive to deviate either. This means that the path p is supported by the extremal penal code, i.e., there is an SPE strategy that yields p as an outcome. \square

Example 1. Let us assume that there are four action profiles; $A = \{a, b, c, d\}$. Moreover, let the sets $P^k(a)$, $k = 1, 2$, $a \in A$, be as in the table below. Now, $aa \in P^2(a)$ means that on any equilibrium path a should be followed by another a , after which it does not matter what comes next as long as it is an equilibrium path. However, since the only action profile that can follow a is a , we observe that after the first a on an equilibrium path, the rest of the action profiles are also a 's. On the other hand, b can be followed by two action profiles; ba and bc , after which any equilibrium path starting with a or c , respectively, is possible. For the action profile c the situation is symmetric to that of b . Finally, since $d \in P^1(d)$, it can be followed by any equilibrium path.

Table 1: An example of sets $P^1(a)$ and $P^2(a)$.

	a	b	c	d
P^1	\emptyset	\emptyset	\emptyset	$\{d\}$
P^2	$\{aa\}$	$\{ba, bc\}$	$\{ca, cb\}$	\emptyset

The sets P^k , $k \geq 1$, and P^∞ give us the subpaths that can follow a given initial action profile. The result of Proposition 1 says that for each action profile a on the equilibrium path there is a subpath belonging to $P^k(a)$ for some

k or $P^\infty(a)$. This means that the equilibrium paths follow a particular syntax defined by the elementary subpaths. In the rest of the paper we shall focus on the collections of elementary subpaths, which are called the elementary sets.

Definition 4. The collection of sets $P^k(a)$, $k = 1, 2, \dots$, and $P^\infty(a)$, $a \in A$, is the elementary set of the infinitely repeated game. For given discount factors T this collection is denoted by $S(T)$. Moreover, $S^k(T)$ denotes the collection of $P^j(a)$, $j = 1, \dots, k$, $a \in A$.

In the following sections particular attention will be paid to finite elementary sets, since $S(T)$ may contain infinitely many subpaths in general. However, the set $S^k(T)$ contains finitely many subpaths because P^k are finite for all k . Another observation on finiteness of elementary subpaths is an immediate consequence of contractivity of B_a , $a \in A$. Let $v(p)$ denote the vector of average discounted payoffs corresponding to p . It follows from the contractivity that if $p \in A^\infty(a)$ and $v(p_1) > \text{con}(a)$, then there is k such that p^k satisfies (3). Moreover, if there is an equilibrium path $p \in A^\infty(a)$ for which (3) fails to hold for all p^k , then $v_i(p_1) = \text{con}_i(a)$ for some i . The opposite is not true, i.e., we may have $v_i(p_1) = \text{con}_i(a)$ for some $i \in N$ at the same time when (3) holds for all p^k . More generally, we have the following result which tells that the possible infiniteness of $S(T)$ is due to subpaths close to the boundary where $v_i(p_1) = \text{con}_i(a)$ for some i .

Proposition 2. For any $\varepsilon > 0$ there is k such that $p_j^l \in P^l(i(p_j))$ for some $l \leq k$ when $p \in A^\infty(a)$, $a \in A$, and

$$v_i(p_1) \geq \text{con}_i(a) + \varepsilon \text{ for all } i \in N. \quad (4)$$

Proof. Because A is finite and B_a , $a \in A$, are contractions, for any $\rho > 0$ there is k such that the diameter of the set that is obtained by taking the image of V under a sequence $B_{a^0}, \dots, B_{a^{k-1}}$, $a^j \in A$ for all $j = 0, \dots, k-1$, has diameter less than ρ . In particular, the diameter of the set $B_{i(p_1)}(W(p_1^k))$ is less than ρ for any p . Now, ρ can be chosen such that for any $a \in A$ and $p \in A^\infty(a)$ for which (4) holds we have

$$B_{i(p_1)}(W(p_1^k)) \subseteq \{v \in V : (3) \text{ holds}\},$$

which concludes the proof. \square

Let us now consider the comparative statics of $S(T)$ with respect to T . Let T_1 and T_2 be two matrices corresponding to two different sets of discount factors. We denote $T_1 \ll T_2$ if the discount factors on the diagonal corresponding to T_2 are at least those of T_1 . With a slight abuse of notation, we denote $p \in S(T)$ when either $p \in P^k(a)$ or $p \in P^\infty(a)$ for some $a \in A$ and $k \geq 0$.

We first show a result, which is of importance itself. It tells that if the punishment payoffs remain the same for two set of discount factors T_1 and T_2 such that $T_1 \ll T_2$, then any equilibrium path for T_1 is also an equilibrium path when the discount factors are increased to T_2 . The payoffs corresponding to T_1 and T_2 are $V(T_1)$ and $V(T_2)$, respectively.

Lemma 1. *If $T_1 \ll T_2$ and $v^-(V(T_1)) = v^-(V(T_2))$, then a path $p \in A^\infty$ that is an SPEP for T_1 is an SPEP for T_2 .*

Proof. To suppress the notation let u^k denote the vector of payoffs in period $k \geq 0$, i.e., $u^k = u(a^k)$, and d^k denote the vector of deviation payoffs $\max_{a_i \in A_i} u_i(a_i, a_{-i}^k)$. Moreover, v^- stands for the vector $v^-(V(T_i))$, $i = 1, 2$.

Because u^k , $k = 0, 1, \dots$, is a payoff stream corresponding to an equilibrium path, the incentive compatibility condition (2) implies that for all $k \geq 0$ we have

$$(I - T_1)u^k + T_1 \left[(I - T_1) \sum_{j=0}^{\infty} T_1^j u^{k+j+1} \right] \geq (I - T_1)d^k + T_1 v^-.$$

By rearranging and observing that $(I - T_1)^{-1}v^- = \sum_{j=0}^{\infty} T_1^j v^-$ we get

$$S_1^k = u^k - d^k + T_1 \sum_{j=0}^{\infty} T_1^j (u^{k+j+1} - v^-) \geq \mathbf{0} \text{ for all } k = 0, 1, \dots$$

Similar expression as for S_1^k can be derived for S_2^k . For this expression we do not need the incentive compatibility condition. Indeed, the purpose is to show that $S_2^k \geq 0$, $k \geq 0$, which means that the incentive compatibility condition holds for T_2 along the SPEP path for T_1 .

It can be seen that S_i^k satisfies the recursion

$$S_i^k = u^k - d^k + T_i (d^{k+1} - v^- + S_i^{k+1}), \text{ for all } k \geq 0 \text{ and } i = 1, 2. \quad (5)$$

Observe that the term $u^k - d^k$ is a vector with non-positive components, which implies that the components of $d^{k+1} - v^- + S_1^{k+1}$, $k \geq 0$, are non-negative, because $S_1^k \geq \mathbf{0}$, $k \geq 0$, by the incentive compatibility.

Let us assume that for the first $K+1$ periods, i.e., for periods $k = 0, 1, \dots, K$, the discount factors are given by T_2 and after that they are given by T_1 . It holds that $T_2 = T_1 + \varepsilon$, where ε stands for the diagonal matrix $T_2 - T_1 \gg \mathbf{0}$.

The recursion (5) gives

$$\begin{aligned} S_2^K &= u^K - d^K + T_1 (d^{K+1} - v^- + S_1^{K+1}) + \varepsilon (d^{K+1} - v^- + S_1^{K+1}) \\ &= S_1^K + \varepsilon (d^{K+1} - v^- + S_1^{K+1}) \end{aligned}$$

Recall that the components of the vector $d^{K+1} - v^- + S_1^{K+1}$ are non-negative. Hence, we get $S_2^K \geq S_1^K \geq 0$. It can now be seen from the recursion (5) that $S_2^j \geq S_1^j \geq 0$ for all $j \leq K$. Letting K go to infinity we obtain the incentive compatibility condition for all $k \geq 0$, when the discount factors are given by T_2 . Hence, the result follows. \square

We can now consider the comparative statics of elementary subpaths.

Proposition 3. *If $T_1 \ll T_2$ and $v^-(V(T_1)) = v^-(V(T_2))$, then $p \in S(T_1)$ implies that there is $k \leq |p|$ such that $p^k \in S(T_2)$.*

Proof. In the following $P^k(a; T_i)$ denotes the set of length k elementary subpaths corresponding to T_i , $C_a^j(V(T_i))$ is the set of continuation payoffs for T_i , $v^i(p_k)$ is the payoff vector corresponding to p_k and T_i , and $W^i(p_k)$ is the set of possible continuations after p_k for a given T_i .

Lemma 1 shows that an equilibrium path for T_1 remains an equilibrium path for T_2 . It follows that if we take $p \in P^l(a; T_1)$, then any equilibrium path for T_1 that follows p_k , $k \leq l$, is an equilibrium path for T_2 . More formally, if we take a continuation payoff $v^1(f(p)) \in C_{f(p)}^1(V(T_1))$, then this continuation payoff corresponds to an equilibrium path. Let $v^2(f(p))$ denote the payoffs of this path for discount factors given by T_2 . Because any equilibrium path for T_1 is an equilibrium path for T_2 we have $v^2(f(p)) \in C_{f(p)}^2(V(T_2))$. By induction argument, $v^1(p_k) \in W^1(p_k)$ implies that $v^2(p_k) \in W^2(p_k)$ for all $k \leq l$.

If we have an equilibrium path for T_1 that starts with $a \in A$, i.e., there is $v^1(p) \in C_a^1(V(T_1))$ such that $(a, v^1(p))$ is admissible, then $(a, v^2(p))$ is admissible for T_2 , i.e., $v^2(p) \in C_a^2(V(T_2))$. Again this follows from Lemma 1. Hence, if $v^1(p_1) \in C_a^1(V(T_1))$, i.e., (3) holds, we also have $v^2(p_1) \in C_a^2(V(T_2))$. This means that either $p \in P^l(a; T_2)$ or there is $k \leq l$ such that $p^k \in P^k(a; T_2)$. If $p \in P^\infty(a; T_1)$ then p is an SPE path for T_1 , and therefore it is an SPE path also for T_2 . Again, either $p \in P^\infty(a; T_2)$ or there is k such that $p^k \in P^k(a; T_2)$. \square

When the discount factors increase, all the subpaths that satisfy (3) still satisfy this condition if the smallest equilibrium payoffs do not change. Note that the number of elementary subpaths and their lengths do not directly reflect the number of equilibrium paths. For example, if $abcd, abdc \in S(T_1)$ it may happen that $ab \in S(T_2)$ for $T_2 \gg T_1$, i.e., corresponding to two elementary subpaths starting with ab there is only one when the discount factors increase. Consequently, ab may be followed by other subpaths than cd or dc .

3. Computation and Analysis of Elementary Sets

We first present an algorithm for finding the subpaths that can be elementary. This process may produce subpaths that contain non-equilibrium parts. The second algorithm removes these subpaths and forms a compact graph presentation for the equilibrium paths. The graph is useful for producing the equilibrium paths and payoffs, and in analyzing the complexity of equilibrium outcomes.

3.1. Algorithm for Finding the Elementary Subpaths

First, we introduce a recursive way of computing the payoff requirements. To illustrate the main idea let us consider a subpath abc . The vector of the smallest payoffs $\text{con}(ab)$ that the players should get after ab to make the first element a incentive compatible are found by solving

$$(I - T)u(b) + T\text{con}(ab) = \text{con}(a).$$

If it happens that $\text{con}_i(a)$ would be below $v_i^-(V)$ then we simply set $\text{con}_i(a) = v_i^-(a)$. Given that $\text{con}(ab)$ is known, we can now find the smallest payoff that is required after abc to make a incentive compatible as the first action profile. This continuation payoff $\text{con}(abc)$ is found by solving

$$(I - T)u(c) + T\text{con}(abc) = \text{con}(ab).$$

Again we set the continuation to $v_i^-(V)$ if it would be below that value. If $\text{con}(abc) \leq \text{con}(c)$, then any equilibrium path starting from c is a possible continuation for abc at the time when a is played. Recall that the same idea was formulated in condition (3).

In general, we can define $\text{con}(p)$ for any $p \in A^k$, $k \geq 2$, as above. When $\text{con}(p^{k-1})$ is known and $p = p^{k-1}a^k$, we set

$$\text{con}_i(p) = \max \{ [\text{con}_i(p^{k-1}) - (1 - \delta_i)u_i(a^k)] / \delta_i, v_i^-(V) \}.$$

Now, $\text{con}(p)$ is simply the continuation payoff vector that is required after $f(p)$ to make the first action profile of p incentive compatible. The following observations are immediate. Note that the first observation relates the smallest payoffs $\text{con}(p)$ to condition (3).

Remark 1. Let $p \in A^k$ and $\bar{v}_i = \max\{v_i : v \in V\}$, $i \in N$.

- i) Condition (3) holds for $p \in A^k$ with $f(p) = a$ if and only if $W(p) \neq \emptyset$ and $\text{con}(p) \leq \text{con}(a)$.
- ii) If $\text{con}_i(p) > \bar{v}_i$ for $p \in A^k$ and some $i \in N$, then p is not an elementary subpath.

Notice that to detect whether a subpath is elementary or not does not require knowing the whole payoff set. The above properties are efficiently utilized in the following algorithm that computes the elementary subpaths. The algorithm first generates sets \hat{P}^k that may contain subpaths that have non-equilibrium parts, and the removal of these subpaths is explained in Section 3.2. The algorithm will be demonstrated in Section 4.

1. For all $a \in A$ include $a \in \hat{P}^1(a)$ if $\text{con}_i(a) \leq v_i^-(V)$ for all $i \in N$. If $v_i^-(V) \leq \text{con}_i(a) \leq \bar{v}_i$ for all $i \in N$, and the first inequality is strict for some $i \in N$, then include a in $P_*^1(a)$. Set $k = 2$, and go to Step 2.
2. For each $a, b \in A$, $p \in P_*^{k-1}(a)$, find $\text{con}(q)$, where $q = pb$.
 - a) If $\text{con}(q) \leq \text{con}(b)$ and

$$q_j \in P_*^{k-j}(i(q_j)) \text{ or } q_j^l \in \hat{P}^l(i(q_j)), \forall j = 1, \dots, k-1, \quad (6)$$

for some $1 \leq l \leq k-j$, then include q in $\hat{P}^k(a)$.

- b) Otherwise, if $\text{con}_i(q) \leq \bar{v}_i$ for all $i \in N$ and q satisfies Eq. (6), then include q in $P_*^k(a)$.

If $P_*^k(a) = \emptyset$ for all $a \in A$ stop. Otherwise, increase k by one and repeat Step 2.

3. Remove the subpaths with non-equilibrium parts from \hat{P}^k and obtain the elementary sets P^k . This is explained in Section 3.2.

The set $P_*^k(a)$ contains the subpaths that are possibly part of elementary subpaths. The test in step 2.b) tells whether it is possible that $q = pb$ is part of an elementary subpath. First, the required continuations should not exceed the upper bounds $\bar{v}_i, i \in N$. Second, all parts of the subpath must satisfy condition (3). This means that for each $j = 1, \dots, k-1$, there is either a shorter elementary subpath starting with $i(q_j)$ or there is possibly some elementary subpath starting with $i(q_j)$, i.e., subpath in $P_*^j(i(q_j))$. If all P_*^k become empty sets the algorithm can be terminated as there cannot be any more elementary subpaths that have not yet been found. Moreover, in that case the elementary set is finite.

Remark 2. If there is k such that $P_*^k(a) = \emptyset$ for all $a \in A$, then $S(T)$ contains finitely many subpaths.

Note that the algorithm can be terminated while there still are elements in P_*^k , in which case we get an approximation for the elementary set. When the algorithm is terminated prematurely like this, some of the elementary subpaths have not been found, and we get a subset of the equilibrium paths. We can identify the missing subpaths and they give payoffs close to the boundary payoffs in some of the players' incentive compatibility conditions, as Proposition 2 suggests. We can also form an "upper bound" for the elementary set by including the remaining P_*^k into \hat{P}^k . In this case, there may be subpaths that are not incentive compatible.

Note that the algorithm uses the smallest and the highest equilibrium payoffs v^- and \bar{v} in computation, and these are typically not known in advance. The highest payoff is not a problem, and it can be replaced by the highest stage game payoff. This value affects how fast the algorithm finds the non-elementary subpaths and thus how fast the algorithm converges. The smallest payoffs are easily known for many games, but for others we have a separate algorithm for finding these payoffs. However, this issue is not covered in this paper. In general, the minmax payoffs can be used as a starting point, since they give the lower bounds to the smallest equilibrium payoffs.

3.2. Graph Presentation

The algorithm above may produce subpaths, whose incentive compatibility relies on sets P_*^k , and these sets may not form any elementary subpaths in the end when the algorithm is terminated. This means that the subpaths in \hat{P}^k may contain non-equilibrium parts. However, removing these subpaths from the sets \hat{P}^k can be done using the subpaths that we have already found, and with the same effort as forming a graph for all equilibrium paths when there are finitely many subpaths. The algorithm is presented below.

1. Form a tree of the subpaths in the sets \hat{P}^k . The root node is the empty history \emptyset .

2. Transform the tree into a graph. Each node in the tree corresponds to a node in the graph. Form the arcs between the nodes by going through them and determine the destinations for each one.
 - (a) The destinations of an inner node in the tree, i.e., node with children, are its children. Set an arc to each destination node.
 - (b) The destinations of a leaf node, i.e., node with no children, which is connected to the root node \emptyset are all the child nodes of \emptyset .
 - (c) For the other leaf nodes, i.e., for subpaths $p \in \hat{P}^k$, find the smallest $i \geq 1$ such that p_i is found in the tree. If p_i is found and it is an inner node, then we remove node p and connect $p^{|p|-1}$ to the node p_i . If p_i is not found and the longest common path with the tree is an inner node, then a part of p cannot appear on an equilibrium path and the node is removed from the graph.
3. Insert arcs and nodes for infinitely long subpaths. For each of these subpaths find largest i such that p^i is a node. Insert an arc with the label p^i from this node to a dummy node corresponding to the path.

Example 2. Let us assume that the subpaths in \hat{P}^k are c, aa, ab, bb, bab and bac ; the corresponding tree is shown in the left of Figure 1. We note that subpath bac contains a non-equilibrium part, since there are no elementary subpaths that start with ac . Thus, subpath bac cannot be part of an equilibrium path, and we see how the node is removed from the graph in the algorithm. The graph that generates all the equilibrium paths is basically formed by going through the sets \hat{P}^k . According to Step 2.(b), node c connects to nodes a, b and c . According to Step 2.(c), node aa connects to a and node ab to b , since $p_1 = a$ and $p_1 = b$ are in the tree. Thus, we loop a to itself and connect node a to b . Similarly, node bb connects to b and b loops to itself. For bab , we search $p_1 = ab$ in the tree and it is a leaf node. Thus, we search $p_2 = b$ and since it is an inner node in the tree, we connect node ba to $p_2 = b$. For bac , we search $p_1 = ac$ but it is not found. The longest common path with ac in the tree is an inner node a , and node bac is removed from the graph. The resulting graph is shown in the right of Figure 1. The last action in the node label gives the action profile that is played when the node is visited. For example, a is played in node ba .

It is straightforward to get the equilibrium paths and payoffs from the graph. The only trick is to combine the finite paths with the infinite cycles from the graph; see Section 3.3 in [6]. This is the way to generate infinite sequences from the graph. Moreover, the graph construction leads directly to the result that the finite elementary sets can be represented by a graph, and the payoff set can be identified as a particular fractal set, i.e., graph directed self-affine set [19].

Proposition 4. *When $S(T)$ contains finitely many subpaths, then all SPE paths can be represented by a graph.*

Corollary 1. *The SPE paths given by $S^k(T)$ can be presented as a graph.*

Corollary 2. *When $S(T)$ contains finitely many subpaths, the payoff set $V(T)$ is a graph directed self-affine set.*




Figure 1: An example of elementary subpaths as a tree and a graph.

3.3. Complexity of Equilibrium Outcomes

The complexity of equilibrium paths and payoffs can be analyzed with the graph. The graph can be represented by its $m \times m$ adjacency matrix D , where m is the number of nodes in the graph, and $D_{ij} = 1$ if there is an arc from node i to j and otherwise $D_{ij} = 0$. The eigenvalues and eigenvectors can be used in counting the number of walks in a graph [11], where a walk means any sequence of nodes using the arcs of the graph. The element $d_{ij}^{(k)}$ of the matrix D^k is equal to the number of walks of length k from node i to j . Here, we are interested in the walks originating from the root node, which is given index 1 and the rest of the nodes are indexed with $j = 2, \dots, m$. The number of k length equilibrium paths is

$$y(k) = \sum_{j=2}^m D_{1j}^k.$$

Asymptotically, the number of equilibrium paths satisfies

$$y(k) \approx y_0 \rho^k(D), \tag{7}$$

where y_0 is a constant and $\rho(D)$ is the largest eigenvalue of D ; see, e.g., Theorem 2.2.2 in [11]. Hence, $\rho(D)$ is the asymptotic growth rate. This measure tells how large the set of equilibrium paths is, and it can be used for comparing different games.

It is also possible to measure the complexity of the payoff set using the graph. One of the fractal measures is the Hausdorff dimension, which tells intuitively how the equilibrium payoffs fill the space. The Hausdorff dimension s can be estimated from the graph by solving¹[19]

$$\rho(\delta^s D) = 1,$$

¹The exact dimension can be defined when so called open set condition holds, which means

assuming that the players have a common discount factor δ . The Hausdorff dimension corresponds to the value s for which the largest eigenvalue of matrix $\delta^s D$ is one. Thus, we can analyze the payoff set with the eigenvalues of the weighted adjacency matrix. In [7] it is shown how to estimate the Hausdorff dimension when the players have different discount factors.

Example 3. Let us examine the graph of Figure 1. The adjacency matrix is

$$D = \begin{pmatrix} 0 & 1 & 1 & 1 & 0 \\ 0 & 1 & 1 & 1 & 0 \\ 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 1 & 0 \end{pmatrix} \quad (8)$$

with the largest eigenvalue $\varphi = (1 + \sqrt{5})/2 \approx 1.618$, which is the golden ratio. The asymptotic growth rate of equilibrium paths is $\rho(D) = \varphi$. When the discount factor is $\delta = 1/2$, the Hausdorff dimension is $s = \log_2 \varphi \approx 0.694$.

4. Numerical Examples

4.1. Prisoner's Dilemma Game

In this section we demonstrate how the algorithm finds the elementary subpaths, how the graph is constructed, and how to analyze the equilibria with the graph. We examine a prisoner's dilemma game with a common discount factor $\delta = 1/2$ and the stage game payoffs as below. The payoff sets of repeated prisoner's dilemma have previously been studied in [21], [22], and [18]. Here, we show the exact paths that can be played in the game and analyze what happens when the discount factor increases.

	<i>L</i>	<i>R</i>
<i>T</i>	3, 3	0, 4
<i>B</i>	4, 0	1, 1

We denote the action profiles from left to right and top to bottom as a, b, c , and d ; for example (T, R) is b . The punishment path is the infinite repetition of d , which is denoted by d^∞ . The corresponding payoffs are $v_i^-(V) = 1$, $i = 1, 2$.

Let us find the elementary set for this game. For this purpose finite paths are classified into elementary and non-elementary sets, and those which belong to P_*^k . We neglect the sets \hat{P}^k and use P^k instead, since in this example there are no subpaths with non-equilibrium parts. In Step 1 of algorithm in Section 3.1, we calculate $\text{con}(p)$ for one length paths p . For example, $\text{con}(d) = (1, 1)$, and d is an elementary subpath since $\text{con}_i(d) \leq v_i^-(V)$, $i = 1, 2$. For paths a , b and c , we have $v^-(V) \leq \text{con}(p) \leq \bar{v}$ and $P_*^1 = \{a, b, c\}$. Table 2 gives the

that the payoffs that are mapped in the graph do not overlap. This condition holds when the discount factor is less than $1/2$. In general, there are techniques for estimating lower and upper bounds for the dimension [13].

payoff requirements for one and two length paths. The elementary subpaths are denoted by +, the non-elementary by -, and those that belong to $P_*^k(a)$, $k = 1, 2$, $a \in A$, by *.

Table 2: Finding elementary subpaths with $|p| \leq 2$.

path	con(path)	path	con(path)	path	con(path)
a	$(2, 2)^*$	b	$(2, 1)^*$	c	$(1, 2)^*$
aa	$(1, 1)^+$	ba	$(1, 1)^+$	ca	$(1, 1)^+$
ab	$(4, 1)^-$	bb	$(4, 1)^-$	cb	$(2, 1)^+$
ac	$(1, 4)^-$	bc	$(1, 2)^+$	cc	$(1, 4)^-$
ad	$(3, 3)^*$	bd	$(3, 1)^*$	cd	$(1, 3)^*$

In the first run of Step 2, we examine $\text{con}(pb)$ for all $b \in A$ and $p \in P_*^1$, and these are the two length paths in Table 2. For instance, for $\text{con}(ab)$ we need $\text{con}(a) = (2, 2)$ and $u(b) = (0, 4)$, and we get

$$\begin{aligned} \text{con}_1(ab) &= \max \left\{ \left[2 - \left(1 - \frac{1}{2} \right) \cdot 0 \right] / \frac{1}{2}, v_1^-(V) \right\} = \max \{4, 1\} = 4, \\ \text{con}_2(ab) &= \max \left\{ \left[2 - \left(1 - \frac{1}{2} \right) \cdot 4 \right] / \frac{1}{2}, v_2^-(V) \right\} = \max \{0, 1\} = 1. \end{aligned}$$

We find that aa , ba , bc , ca and cb are elementary, but since ad , bd , and cd belong to P_*^2 , we have not yet found all the elementary subpaths. We can immediately observe that ad is incentive compatible only when it is followed by an infinite repetition of a , i.e., $P^\infty(a) = \{ada^\infty\}$, since no other action profile gives the required payoff $(3, 3)$. Thus, we do not consider other subpaths starting with ad , and it is removed from the set P_*^2 , which is a minor deviation from the algorithm.² Moreover, we only need to examine paths beginning with b or c , since the game is symmetric. Let us consider the three and four length paths beginning with cd . For example, cdb belongs to $P_*^3(c)$ because $\text{con}_i(cdb) \leq 3$, for all $i \in N$, $d \in P^1(d)$ and $b \in P_*^1(b)$.

Now, we can see that the only possible paths starting with cd are cda and cdb . The only continuation to cda is aa , since the only elementary subpaths starting with a are aa and ad , and ad gives lower payoff than the required $(1, 3)$. Thus, $P^\infty(c) = \{cda^\infty\}$ and we do not need to consider other subpaths starting with cda . From length four paths, we can observe that $P^4(c) = \{cdba\}$ and $P^\infty(c) = \{cldbda^\infty, cda^\infty\}$. As earlier, $cdbd$ can only be followed by a^∞ and no other subpaths starting with $cdbd$ need to be considered. Hence, there are no longer paths to be searched for and we have found the elementary set.

The tree of elementary subpaths is presented in Figure 2. The destinations of leaf nodes are indicated next to them. Using the graph algorithm without

²This kind of removal of subpaths from P_*^k could be included in the method. To simplify the exposition, we have excluded it from the algorithm presented in Section 3.1.

Table 3: Finding elementary subpaths with $3 \leq |p| \leq 4$.

path	con(path)	path	con(path)
<i>cda</i>	$(1, 3)^*$	<i>cdba</i>	$(1, 1)^+$
<i>cdb</i>	$(2, 2)^*$	<i>cdbb</i>	$(4, 1)^-$
<i>cdc</i>	$(1, 6)^-$	<i>cdbc</i>	$(1, 4)^-$
<i>cdd</i>	$(1, 5)^-$	<i>cdbd</i>	$(3, 3)^*$

Step 3, we get the directed graph composed of solid arcs in Figure 2. Each node denotes what is played when the node is visited.

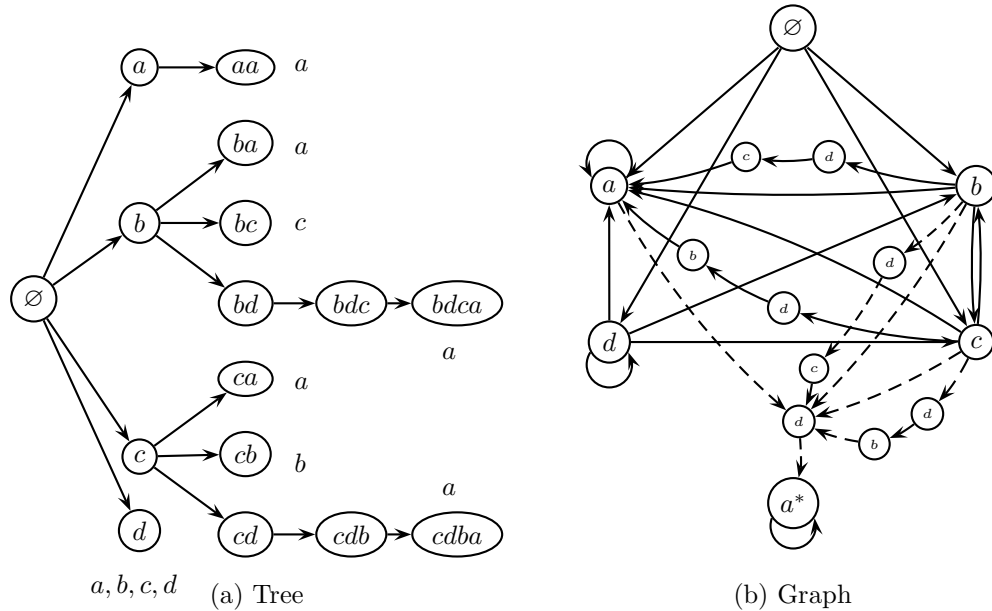


Figure 2: Tree of finite elementary subpaths and a graph of all equilibrium paths.

To get all the SPE paths of the game, we add nodes and arcs corresponding to the infinitely long elementary subpaths to the graph:

$$P^\infty = \{ada^\infty, bda^\infty, cda^\infty, bdcda^\infty, cdbda^\infty\}.$$

We need another node to distinguish whether *d* is played after *a*, *b*, or *c* or not. For example, if *ad* is played then a^∞ must follow and *ad* cannot be played any more. We denote this extra node as a^* , and by adding the new nodes and arcs we get the graph of Figure 2 in which the dashed arcs are also included.

An approximation of the payoff set is shown in the left of Figure 3. The payoff set consists of similar patterns in different scales, which shows the fractal

nature. The set is constructed by combining finite paths from the graph to the infinite cycles starting from the final nodes of the paths. The dashed and solid lines represent the payoff requirements of the right-hand side of the incentive compatibility condition (2) for the first and second columns of the game, respectively. We can see that there are payoff points on these lines, and these correspond to the paths in P^∞ , such as ada^∞ , bda^∞ and cda^∞ . This is the role of the infinitely long elementary subpaths; some part of the path gives exactly the payoff requirement.




Figure 3: The payoff sets for $\delta = 0.5$ and $\delta = 0.58$.

The payoff set is sparse and the Hausdorff dimension is zero. The largest eigenvalue of the adjacency matrix is one, and the number of k length paths increases subexponentially in k . In fact, the value of $\delta = 1/2$ is exactly the limit when the Hausdorff dimension changes from zero and the growth rate becomes exponential. For example, when $\delta = 0.51$, a subpath $adaaaa$ becomes elementary, and it is possible to play d repeatedly as long as at least four a 's are played after it. In this case, the dimension is $s \approx 0.42$ and the growth rate is $\rho \approx 1.32$.

When the discount factor increases, there will be more and more equilibrium and elementary subpaths, and the graph grows larger. When $\delta = 0.58$ the graph has over one hundred nodes and the payoff set is shown in Figure 3. The payoff set is much more complex, the estimate of the Hausdorff dimension is $s \approx 1.37$ and the paths increase at rate $\rho \approx 2.09$. With higher discount factor values, the sets P_*^k do not become empty for reasonable k since there are always elementary subpaths in the proximity of the payoff requirement values as the payoff set becomes dense.

4.2. Sierpinski Game

In this example we demonstrate an interesting feature of equilibrium payoffs, which is captured by the Hausdorff dimension. The payoff set becomes more complex when the discount factor is increased, even though the elementary set remains the same. We call the following Sierpinski game because the payoff set will be the celebrated Sierpinski triangle. The payoffs are given below and $\delta = 1/2$. We also denote $a = (T, L)$, $b = (C, M)$, and $c = (B, R)$.

	L	M	R
T	$2 - \sqrt{3}, 1$	$-1, -1$	$-1, -1$
C	$-1, -1$	$1, 2 - \sqrt{3}$	$-1, -1$
B	$-1, -1$	$-1, -1$	$0, 0$

In this game there are three pure-strategy Nash equilibria that are the corner points of the payoff set, which is illustrated in Figure 4. The equilibrium paths are all combinations of these three points, and the graph consists of all transitions between the three nodes. Here, the dummy node \emptyset is omitted as redundant. The payoff set is the Sierpinski triangle, whose Hausdorff dimension is $s = \log 3 / \log 2 \approx 1.585$. The dimension tells that the set does not quite fill the two dimensional space but it is more complex than one dimensional set.

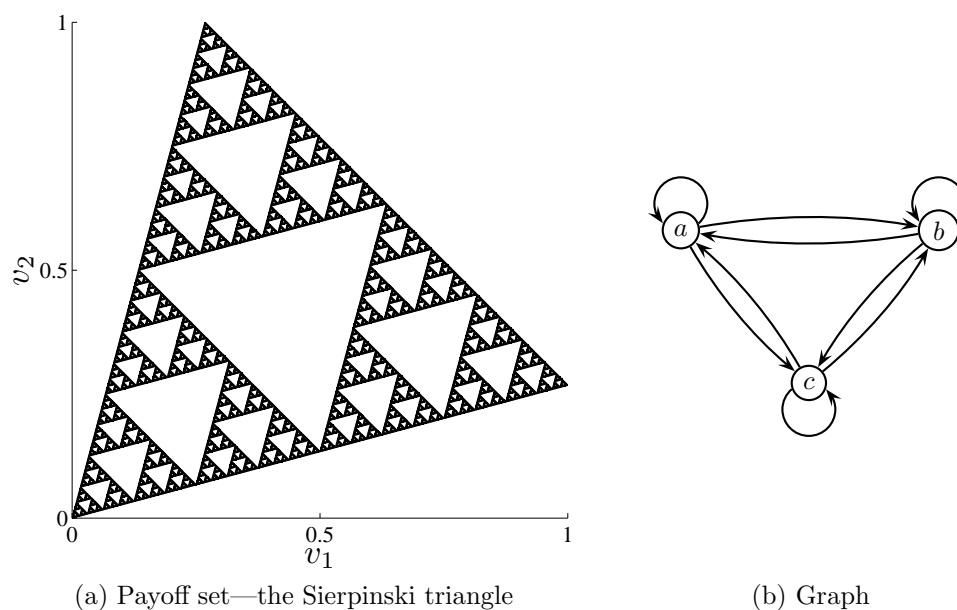


Figure 4: Sierpinski triangle as the payoff set and the graph presentation of SPE paths.

When the discount factor is increased a little from $\delta = 1/2$, the elementary set does not change. However, the payoff set becomes more complex. Eventually, the payoff set fills the triangle defined by the three Nash equilibria. This happens

when $\delta > 2/3$, and then the Hausdorff dimension becomes two. This happens even if the set of elementary subpaths remains the same when the discount factor increases. For example, we can replace minus ones by a small enough number to guarantee that there will be no more equilibrium paths when δ increases. This observation gives an important insight into the folk theorem [15]. One reason for the fact that any feasible payoff above min-max levels can be achieved as an equilibrium outcome is that the payoffs are less contracted under the mappings B_a , $a \in A$, when the discount factor increases. Moreover, the payoff set may enlarge even when the set of equilibrium paths and strategies remains the same.

5. Conclusion

Our main result is that the equilibrium paths are composed of sequences of players' action profiles which we call the elementary subpaths. The result holds for all equilibrium paths corresponding to pure strategies in discounted repeated games with infinite time horizon. The elementary subpaths are of particular interest because they can be used in analyzing the complexity of equilibrium outcomes for different games, and in constructing all equilibrium paths and the corresponding payoffs [6].

The characterization result for equilibrium paths is complemented with an algorithm for finding the elementary subpaths, i.e., the elementary set. When there are finitely many elementary subpaths the algorithm produces all of them after finitely many steps. In general, the elementary subpaths which are not found when the algorithm is terminated prematurely, are the ones which give payoffs close to the boundary payoffs in some of the players' incentive compatibility conditions. Even when the algorithm is terminated prematurely, we can form an approximation for the elementary set and identify the missing subpaths to a certain degree.

The final step of the algorithm transforms the elementary subpaths into a directed graph, which is a compact representation of the equilibrium paths. The graph can be used in generating equilibrium outcomes and analyzing their complexity. We provide two complexity measures: the asymptotic growth rate and the Hausdorff dimension. The asymptotic growth rate measures how fast the number of paths increases as they become longer. The larger the rate, the faster the number of possible finitely long equilibrium paths grows as the stage game is repeated. The Hausdorff dimension, on the other hand, measures how the payoff set fills the space, and hence serves as a measure for the complexity of equilibrium payoff set.

Acknowledgements

We thank Ichiro Obara and Hannu Salonen for their valuable comments. Kitti acknowledges funding from the Academy of Finland.

References

- [1] D. Abreu, Extremal equilibria of oligopolistic supergames, *Journal of Economic Theory* 39 (1986) 191–225.
- [2] D. Abreu, On the theory of infinitely repeated games with discounting, *Econometrica* 56 (1988) 383–396.
- [3] D. Abreu, D. Pearce, E. Stacchetti, Optimal cartel equilibria with imperfect monitoring, *Journal of Economic Theory* 39 (1986) 251–269.
- [4] D. Abreu, D. Pearce, E. Stacchetti, Toward a theory of discounted repeated games with imperfect monitoring, *Econometrica* 58 (1990) 1041–1063.
- [5] D. Abreu, Y. Sannikov, An algorithm for two player repeated games with perfect monitoring, Working paper (2011).
- [6] K. Berg, M. Kitti, Computing equilibria in discounted 2×2 supergames, *Computational Economics* (to appear) (2011).
- [7] K. Berg, M. Kitti, Fractal geometry of equilibrium payoffs in discounted supergames, Working paper (2012).
- [8] K. Chatterjee, H. Sabourian, Game theory and strategic complexity, in: R.A. Meyers (Ed.), *Encyclopedia of Complexity and System Science*, Springer, 2009.
- [9] M.B. Cronshaw, Algorithms for finding repeated game equilibria, *Computational Economics* 10 (1997) 139–168.
- [10] M.B. Cronshaw, D.G. Luenberger, Strongly symmetric subgame perfect equilibria in infinitely repeated games with perfect monitoring, *Games and Economic Behavior* 6 (1994) 220–237.
- [11] D. Cvetković, P. Rowlinson, Simić, *Eigenspaces of graphs*, Cambridge University Press, 1997.
- [12] E.D. Demaine, *Playing Games with Algorithms: Algorithmic Combinatorial Game Theory*, *Mathematical Foundations of Computer Science (MFCS 2001)*, Springer, 2001.
- [13] G.A. Edgar, J. Golds, A fractal dimension estimate for a graph-directed iterated function system of non-similarities, *Indiana University Mathematics Journal* 48 (1999) 429–447.
- [14] A. Fraenkel, Combinatorial games: Selected bibliography with a succinct gourmet introduction, *Electronic Journal of Combinatorics. Dynamic Survey* 2 (2012).
- [15] D. Fudenberg, E. Maskin, The folk theorem in repeated games with discounting and incomplete information, *Econometrica* 54 (1986) 533–554.

- [16] K. Judd, Ş. Yeltekin, J. Conklin, Computing supergame equilibria, *Econometrica* 71 (2003) 1239–1254.
- [17] E. Kalai, Bounded rationality and strategic complexity in repeated games, in: T. Ichiishi, A. Neyman, Y. Tauman (Eds.), *Game Theory and Applications*, Academic Press, 1990, pp. 131–157.
- [18] G.J. Mailath, I. Obara, T. Sekiguchi, The maximum efficient equilibrium payoff in the repeated prisoners' dilemma, *Games and Economic Behavior* 40 (2002) 99–122.
- [19] R.D. Mauldin, S.C. Williams, Hausdorff dimension in graph directed constructions, *Transactions of the American Mathematical Society* 309 (1988) 811–829.
- [20] H. Salonen, H. Vartiainen, Valuating payoff streams under unequal discount factors, *Economics Letters* 99 (2008) 595–598.
- [21] S. Sorin, On repeated games with complete information, *Mathematics of Operations Research* 40 (1986) 147–160.
- [22] D.O. Stahl, The graph of prisoners' dilemma supergame payoffs as a function of the discount factor, *Games and Economic Behavior* 3 (1991) 368–384.