

Kimmo Berg and Mitri Kitti
**Equilibrium Paths in Discounted
Supergames**

Aboa Centre for Economics

Discussion paper No. 96

Turku 2016

The Aboa Centre for Economics is a joint initiative of the economics departments of the University of Turku and Åbo Akademi University.



Copyright © Author(s)

ISSN 1796-3133

Printed in Uniprint
Turku 2016

Kimmo Berg and Mitri Kitti
Equilibrium Paths in Discounted Supergames

Aboa Centre for Economics
Discussion paper No. 96
March 2016 (first draft November 2014)

ABSTRACT

This paper examines the subgame-perfect pure-strategy equilibria in discounted supergames with perfect monitoring. It is shown that all the equilibrium paths are composed of fragments called elementary subpaths. This characterization result makes it possible to compute and analyze the equilibrium paths and payoffs by using a collection of elementary subpaths. It is also shown that all the equilibrium paths can be compactly represented by a directed graph when there are finitely many elementary subpaths. In general, there may be infinitely many elementary subpaths, but it is always possible to construct finite approximations. When the subpaths are allowed to be approximatively incentive compatible, it is possible to compute in a finite number of steps a graph that represents all the equilibrium paths. The directed 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.

JEL Classification: C72, C73

Keywords: repeated game, subgame-perfect equilibrium, equilibrium path, graph presentation of paths, complexity

Contact information

Kimmo Berg
Systems Analysis Laboratory
Aalto University School of Science
P.O. Box 11100
FI-00076 Aalto, Finland
Email: kimmo.berg (at) aalto.fi

Mitri Kitti (corresponding author)
Department of Economics
University of Turku
FI-20014, Finland
Email: mitri.kitti (at) utu.fi

Acknowledgements

We thank Ichiro Obara, Hannu Salonen, and a number of anonymous referees for their valuable comments. Mitri Kitti acknowledges funding from the Academy of Finland.

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, it is shown that all the equilibrium paths are generated from a collection of subpaths called elementary subpaths. By equilibrium paths, we mean the infinite sequences of players' actions that are induced by the subgame-perfect equilibrium strategies. The elementary subpaths offer new tools for examining the internal structure of equilibria. In particular, they can be used in analyzing the complexity of equilibria and in computing the equilibrium paths and payoffs.

Abreu [1, 2] has shown that all the equilibrium outcomes can be obtained in simple strategies. This means that it is enough to consider equilibrium paths, which are characterized by the property 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 to the deviator. This idea of most severe punishments can also be utilized in characterizing the equilibrium payoffs with a set-valued fixed-point equation, see Abreu et al. [3, 4] for the case of imperfect monitoring and Cronshaw and Luenberger [18] for perfect monitoring, and Kitti [29, 30] for generalizations to stochastic games. 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 the equilibrium paths from the players' incentive-compatibility conditions. As shown in the paper, all the equilibrium paths are constructed from a collection of sequences of players' actions, which are called the elementary subpaths. The main motivation for introducing elementary subpaths is methodological; they offer new ways to compute and analyze the equilibria. To get an idea of the concept, consider the prisoner's dilemma game of Figure 1. The action profiles are denoted by letters a to d . Playing b and then c , i.e., bc , is an elementary subpath, because it is possible to play the action profile b whenever is followed by c and then continued by any elementary subpath that starts with c . For example, it is possible to combine the elementary subpaths bc and cb to get an equilibrium path $bc bc bc \dots = (bc)^\infty$. On the other hand, combining bc , ca and aa leads to an equilibrium path $bcaaaa \dots = bca^\infty$. All paths produced in this manner are the induced outcomes of subgame-perfect equilibrium strategies.

We present an algorithm for computing the elementary subpaths and show that the equilibrium paths can be compactly represented by a directed graph. The graph in Figure 1 shows all the equilibrium paths in the game for the discount factors between $1/3$ and 0.46 . The graph offers a simple way to produce the equilibrium paths and payoffs, and analyze them. Unlike in the literature on combinatorial games [24], where graphs are commonly used to present the

	<i>L</i>	<i>R</i>
<i>T</i>	3, 3 (<i>a</i>)	0, 4 (<i>b</i>)
<i>B</i>	4, 0 (<i>c</i>)	1, 1 (<i>d</i>)

Elementary subpaths:
d, *aa*, *ba*, *bc*, *ca*, *cb*




Figure 1: Elementary subpaths and a graph of all the equilibrium paths.

players' available moves at different positions, we use graphs to describe the variety of equilibrium behavior. Moreover, we emphasize that this characterization concerns all the equilibrium paths simultaneously rather than gives a condition for individual paths as in [1, 2].

We propose two complexity measures to analyze the equilibrium paths and payoffs based on the graph presentation. The first one is the asymptotic growth rate of the equilibrium paths. This measure tells us the rate at which the number of finitely long equilibrium paths grows when their length is increased. Because 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.

The second complexity measure is the Hausdorff dimension of the payoff set. This measure reflects the density of the equilibrium payoff set, which is a fractal in general [10]. The phenomenon that the payoff set behaves in a rather complex manner, as fractals do, is not completely new, see [31] and [34]. 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 proposed complexity measures make it possible, e.g., to compare different repeated games in terms of equilibrium behavior.

Our approach to the complexity of repeated game equilibria is new, and it differs from the previous literature on strategic complexity [15, 28] and computational complexity [16, 21]. We analyze the complexity of all the equilibrium outcomes without relying on the complexity of individual strategies nor their computation. It has been shown that computing even an approximate equilibrium in a stage game is difficult [20], and the task is not any easier in repeated games [13, 33]. However, there are efficient algorithms that work for a class of repeated games when the assumption of subgame perfection is relaxed [6].

The graph presentation of elementary subpaths can be applied in producing the set of equilibrium payoffs or its approximation. The technique is applied to the systematic analysis of symmetric 2×2 games in [9]. The earlier algorithms for finding the equilibrium payoffs are based on the set-valued fixed-point iteration [5, 14, 17, 27], and they assume equal discount factors. Apart from [14],

these algorithms produce the equilibrium payoffs corresponding to correlated strategies. It should be noted that our method allows the players to have unequal discount factors, and we do not require a public correlation device. The method also produces the exact set of equilibrium payoffs when all the elementary subpaths are known.

The paper is structured as follows. In Section 2, it is shown that the equilibrium paths consist of elementary subpaths. The properties of collections of elementary subpaths are analyzed in Section 3. Section 4 deals with the computation and approximation of the elementary sets. The graph presentation and the complexity of equilibrium paths are studied in Section 5. Illustrative examples of the main ideas and the computational methods are presented in Section 6. Conclusions are given in Section 7.

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 $u : A \rightarrow \mathbb{R}^n$ 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 $u_i(a)$.

In the supergame, the stage game is repeated infinitely many times and the players discount the future payoffs with discount factors $\delta_i \in [0, 1)$, $i \in N$. 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 k -length 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 k -length 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 $A^\infty(a)$. We denote the action profiles by letters a to z , and $(bc)^\infty = bcbcbc \dots$ denotes that bc is infinitely repeated.

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 of 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 . It should be noted that a strategy usually induces different paths after different histories of past play (i.e., the strategies $\sigma|p$, $p \in A^k$, $k = 0, 1, \dots$, may all induce different paths from period $k + 1$ onwards). This is clarified in the following example.

Example 1. Consider the prisoner's dilemma game with payoffs given below.

	L	R
T	3, 3	0, 4
B	4, 0	1, 1

In this game, the action sets are $A_1 = \{T, B\}$ and $A_2 = \{L, R\}$. The four action profiles can be named as $a = (T, L)$, $b = (T, R)$, $c = (B, L)$, and $d = (B, R)$.

Consider now a variant of so called getting-even strategy. Player 1 chooses T if the number of times player 2 has chosen R is greater than the number of times player 1 has chosen T , and otherwise player 1 chooses B . Player 2 chooses R if the number of times that player 1 has chosen B is larger than the number of times player 2 has chosen R , and otherwise player 2 chooses L . In the first round and in any round in which the players have played equally many times B and R , player 1 chooses T and player 2 chooses R . If the players do not deviate in the first round, the strategy leads to the path $p = (bc)^\infty = bcbcb \dots$. For a history in which either of the players has deviated, the game will continue by the players' following either the path $(bc)^\infty$ or $(cb)^\infty$. Hence, these are the paths induced by this strategy.

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

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

The vector of average discounted payoffs corresponding to path p is denoted by $v(p)$. Subgame perfection is defined in the usual way; σ is a subgame-perfect equilibrium (SPE) of the supgame if

$$U_i(\sigma|p) \geq U_i(\sigma'_i, \sigma_{-i}|p) \quad \forall i \in N, \quad p \in A^k, \quad 0 \leq k < \infty, \quad \text{and } \sigma'_i \in \Sigma_i.$$

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 (SPEP) if there is an SPE strategy profile that induces it.

Recall that a strategy may induce different paths after different histories of past play, including histories involving non-equilibrium behavior. A subgame-perfect equilibrium paths may be induced by a strategy after such a history.

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 $p|^{p'} = p'$.

Equilibrium strategies lead to paths of play from which no player wants to deviate given that all the other players follow their equilibrium strategies. All such paths can be implemented in simple strategies [2]. These strategies are defined by $n + 1$ paths: an initial path that the play follows and a punishment path for each player that gives the smallest equilibrium payoff to the player. The players follow the actions given by the current path unless some player makes a unilateral deviation from the path. If this happens, the play switches to the punishment path of the deviator. If more than one player deviates at the same time, the play remains on the current path and there is no punishment. Because we examine a non-cooperative game, we need not consider simultaneous deviations by multiple players.

In essence the characterization of equilibrium paths (i.e., outcomes of equilibrium behavior) in terms of simple strategies means that all we need in analyzing the equilibrium outcomes, either paths or payoffs, are the most severe punishments. What happens off the equilibrium paths has only the role of providing the incentives for the players not to deviate from the equilibrium play. These incentives can usually be provided in several ways, one of which is given by the most severe threats used in simple strategies. Hence, the strategy space of a game is more complicated than the set of equilibrium paths or payoffs: There are more equilibrium strategies than only the simple strategies that lead to the same equilibrium outcomes. However, any equilibrium path can be turned into a simple strategy. Hence, simple strategies are sufficient in characterizing the equilibrium outcomes.

Knowing all the equilibrium paths of a supergame tells everything on how strategic interaction affects the actions that can be taken in the equilibria: What can happen in the future if we have observed certain play and what are the possible histories of actions if we see certain sequences played in the future. This information contains all the restrictions for the sequences of play that come from the basic assumption that no player should be willing to deviate from the ongoing path given that the other players follow their equilibrium strategies.

We shall derive a characterization for the equilibrium subpaths by assuming that the players' smallest equilibrium payoffs are known. Finding these payoffs is discussed in Section 4.1, see also [7, 11] on the computation of the smallest equilibrium payoffs.

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 [18]. Non-emptiness follows when the stage game has a Nash equilibrium in pure strategies; [26] examines the computational complexity of pure-strategy equilibria. It should be noted that V may be non-empty even if the stage game does not have a pure-strategy Nash equilibrium. The smallest equilibrium payoff for player i is denoted by $v_i^- = v_i^-(V) = \min \{v_i : v \in V\}$ and the corresponding payoff vector is $v^-(V) = (v_1^-(V), \dots, v_n^-(V))$. Similarly, the highest equilibrium payoff for player i is $v_i^+ = \max \{v_i : v \in V\}$, and the maximum deviation payoff from action profile a is $u_i^*(a) = \max_{a'_i \in A_i} u_i(a'_i, a_{-i})$. Let Q be a compact set in \mathbb{R}^n . A pair (a, v) of an action profile $a \in A$ and a continuation payoff $v \in Q$ is

admissible with respect to Q if it satisfies the incentive-compatibility conditions

$$(1 - \delta_i)u_i(a) + \delta_i v_i \geq (1 - \delta_i)u_i^*(a) + \delta_i v_i^-(Q), \quad \forall i \in N. \quad (2)$$

According to this constraint, it is better for player $i \in N$ to take the action a_i and get the payoffs v_i than to deviate and then obtain $v_i^-(Q)$.

In the following, $C_a(Q)$ denotes the set of payoffs for which the pair (a, v) is admissible with respect to Q . Note that the vector of the smallest payoffs $\text{con}(a)$ that make (a, v) admissible can be found from the incentive-compatibility conditions:

$$(1 - \delta_i)u_i(a) + \delta_i \text{con}_i(a) = (1 - \delta_i)u_i^*(a) + \delta_i v_i^-(Q), \quad i \in N.$$

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

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

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

where I is an $n \times n$ identity matrix and T is an $n \times n$ diagonal matrix with discount factors $\delta_1, \dots, \delta_n$ on the diagonal. Note that these mappings are contractions because the discount factors are less than one. The image of a compact set of payoffs $Q \subseteq \mathbb{R}^n$ under B_a is denoted by $B_a(Q)$.

2.2. Elementary subpaths

The purpose is to define a set of subpaths that represent all SPE paths in the sense that all SPEs are obtained from these subpaths. Let us first discuss how a collection of subpaths can be used to create a set of infinitely long paths. Assume that $S \subseteq \mathcal{A}$ is a collection of subpaths some of which are possibly infinitely long. A path p can be obtained from S if for all $j = 0, 1, \dots$ there is $k_j \in \mathbb{N}$ or $k_j = \infty$ such that $p_j^{k_j} \in S$. In other words, at any stage k , there is a subpath in S that corresponds to a fragment of p beginning from the stage k . For example, if $S = \{ab, ba\}$, then $p = abab \dots = (ab)^\infty$ is a path obtained from S ; for $j = 0, 2, 4, \dots$ it holds that $p_j^2 = ab$ and for j odd it holds that $p_j^2 = ba$.

We denote the largest possible set of paths obtained from $S \subseteq \mathcal{A}$ by $\text{paths}(S)$. The largest refers to the largest set in set inclusion.

Definition 3. The set of paths corresponding to $S \subseteq \mathcal{A}$, denoted by $\text{paths}(S) \subseteq A^\infty$, is the largest set of paths $S' \subseteq A^\infty$ in set inclusion with the property that $p \in S'$ if and only if for all j there is $k_j \in \mathbb{N}$ or $k_j = \infty$ such that $p_j^{k_j} \in S$.

The main idea of this paper is to define the set S of elementary subpaths such that $\text{paths}(S)$ is exactly the set of SPE paths, i.e., S represents all SPEs. The elementary subpaths are based on the idea of considering the possible payoffs at the time when an action profile is played given that it is followed by a path of equilibrium actions. To clarify this idea, suppose that $a^0 \in A$ starts an SPE path $a^0 a^1 a^2 \dots$, $a^j \in A$ for $j = 1, 2, \dots$. Corresponding to a^0 , we can search for

an action profile a^j in $p \in S$ with the property that a^0 is incentive compatible (satisfies (2) with respect to V) given that it is followed by the sequence of action profiles $a^1 \cdots a^j$ and a^j is followed by any continuation payoff v that makes (a^j, v) admissible at stage j . The resulting subpath $p = a^0 \cdots a^j$ is an elementary subpath. The same construction can be repeated by choosing any initial action profile a^k , $k = 0, 1, 2$ on the original path $a^0 a^1 a^2 \cdots$, which produces a whole set of elementary subpaths.

In the rest of the paper, B_p denotes the composite mapping $B_{i(p)} B_{i(p_1)} \cdots B_{f(p)}$. For an infinitely long p taking $B_p(v)$ means taking an infinite sequence of affine mappings. Hence, for an infinitely long path the mapping $B_p(Q)$ becomes a singleton set containing the discounted average payoff vector of p regardless of the choice of the compact set Q . This is because

$$(I - T) \sum_{k=0}^{K-1} T^k u(a^k) + T^K v \rightarrow B_p(v) = (I - T) \sum_{k=0}^{\infty} T^k u(a^k)$$

when $K \rightarrow \infty$, $p = a^0 a^1 \cdots$, and $v \in Q$.

The property that the first element $i(p)$ of p becomes incentive compatible when the last element $f(p)$ is followed by a continuation payoff $v \in V$ that makes $(f(p), v)$ admissible can be written as

$$B_{p_1}(C_{f(p)}(V)) \subseteq C_{i(p)}(V). \quad (3)$$

The set $C_{i(p)}(V)$ contains the possible continuation payoffs of $i(p)$, while the set $B_{p_1}(C_{f(p)}(V))$ is obtained from the possible continuation payoffs of $f(p)$ when playing the action profiles in p_1 , i.e., all other action profiles of p except for the first one. Hence, the condition (3) means that the first action profile $i(p)$ of p can be played, or becomes incentive compatible, when the last action profile $f(p)$ of p is followed by any SPE payoff that is a possible continuation after $f(p)$.

We are now ready to give a precise definition for the set of elementary subpaths—briefly the elementary set.

Definition 4. The set of elementary subpath $S(T)$ of the game for discount factors corresponding to T is the largest set of subpaths with the properties that for any $p \in S(T)$

1. p satisfies (3) and there is no $k < |p|$ such that p^k satisfies (3),
2. for any $j = 1, 2, \dots, |p| - 1$ there is $q \in S(T)$ such that either $p_j^k = q^k$ for some $k = 1, 2, \dots, |q|$ or $p_j = q$.

Note that the first assumption in the definition of the elementary set incorporates minimality with respect to the path length; $S(T)$ contains the shortest subpaths that satisfy (3). Hence, if $p \in S(T)$ satisfies (3) no shorter fragment of it can satisfy (3).

It is worth observing that individual elementary subpaths can be derived from SPEPs. For a given SPEP p and any $j = 0, 1, \dots$, either there is a smallest

number $k_j \in \mathbb{N}$ such that $p(j) = p_j^{k_j}$ satisfies (3) or condition (3) holds only for $p(j) = p_j$, in which case we denote $k_j = \infty$. All the subpaths $p(j)$, $j = 0, 1, \dots$, obtained this way are elementary subpaths; by construction, for each $p(j)$ and any $k = 1, \dots, |p(j)| - 1$ either $p(j)_k^l = p(j+k)^l$ for some $l = 1, 2, \dots, |p(j+k)|$, or $p(j)_k = p(j+k)$. In other words, some fragment in the beginning of $p(j)_k$ is certainly found in the beginning of $p(j+k)$. In particular, if both $p(j)$ and $p(j+k)$ are infinitely long, then this holds for $p(j)_k = p(j+k)$. Hence, we can alternatively define elementary subpaths as the shortest fragments of SPEPs that satisfy (3).

Remark 1. If p is an SPEP and k_j , $j = 0, 1, \dots$, is the smallest number such that $p_j^{k_j}$ satisfies (3), then $p_j^{k_j} \in S(T)$.

Our main result is that the elementary subpaths represent the whole set of SPE paths.

Proposition 1. *The set $\text{paths}(S(T))$ equals the set of subgame-perfect pure-strategy equilibrium paths.*

Proof. Take any path $p \in \text{paths}(S(T))$. By the definition of $\text{paths}(S(T))$, for any $j \in \mathbb{N}$ there is $k_j \in \mathbb{N}$ or $k_j = \infty$ such that $p_j^{k_j} \in S(T)$. On the other hand, having $p_j^{k_j} \in S(T)$ means that there is no profitable one-shot deviation from $i(p_j)$ for any j by the first requirement of the definition of $S(T)$. Hence, p is an SPEP.

On the other hand, take any SPEP p . As observed in Remark 1, for any $j \in \mathbb{N}$, there is an elementary subpath $p(j)$ corresponding to $i(p_j)$. Hence, any SPEP p corresponds to a collection of subpaths $S(p) = \{p(j) : j = 0, 1, \dots\}$ that satisfies (3) and is minimal in path length. Hence, the set of all SPEPs corresponds to a collection $S = \cup\{S(p) : p \text{ is an SPEP}\}$ that satisfies the two properties in the definition of $S(T)$, i.e., $S \subseteq S(T)$ because $S(T)$ is assumed to be the largest of such sets. Note that $S \subseteq S(T)$ implies $\text{paths}(S) \subseteq \text{paths}(S(T))$. Moreover, $\text{paths}(S)$ is the set of all SPEPs by construction. Because, $\text{paths}(S(T))$ are SPEPs as shown before, it follows that $\text{paths}(S) = \text{paths}(S(T))$. \square

The above result means that any subgame-perfect equilibrium path follows a 'syntax', in which for each action profile on the path there is an elementary subpath that begins with that action profile. In particular, if an outside observer has seen a certain path of past play, the observer can deduce the possible future paths, when the elementary subpaths are known. However, the players need not care about the elementary subpaths; all they have to do is to follow their equilibrium strategies, which leads to a realization that follows the syntax given by the elementary subpaths. The game may switch to another equilibrium path only if some of the players makes a unilateral deviation from the ongoing path.

It is worth noticing that the elementary set $S(T)$ is defined with respect to the set of equilibrium payoffs V , which cannot be derived through the notion of elementary paths. However, as observed in Section 4, the only information

of V that is needed in order to obtain $S(T)$ are the players' least equilibrium payoffs. Finding the least equilibrium payoffs is discussed in Section 4.1

The following example demonstrates how the equilibrium paths can be formed from the elementary subpaths. The elementary subpaths of different length and different initial action profiles are denoted by $P^k(a)$, where k is the path length and a is the initial action profile. The set of all elementary subpaths of length k is P^k . The set $S^k(T)$ denotes the set of elementary subpaths that are at most of length k and the discount factors are determined by the diagonal elements of T .

Example 2. Consider the prisoner's dilemma game of Example 1 and recall that there are four action profiles; $A = \{a, b, c, d\}$. Let the sets $P^k(a)$, $k = 1, 2$, $a \in A$, be as in Table 1. These elementary subpaths correspond to the prisoner's dilemma game for $\delta_1 = \delta_2 = \delta$ where δ is at least $1/3$ and at most $\delta \approx 0.46$. This game is further analyzed in Section 6.1. Here, the elementary subpaths are taken as given. The algorithm for finding them is presented in Section 4.

Observe in particular that for any p listed in Table 1 and any $j = 1, \dots, |p|-1$, there is an elementary subpath q such that $p_j^k = q^k$ for some $k = 1, 2, \dots, |q|$ as required in the definition of the elementary set. For example, corresponding to $p = aa$ a suitable q is aa , and for $p = bc$ the path q can be chosen to be either cb or ca .

Because $P^2(a) = \{aa\}$, i.e., aa is the only subpaths that starts with a , on any equilibrium path a is followed by another a , and the rest of the action profiles are also a 's. Moreover, aa is elementary because the second a requires a continuation payoff v such that $(I - T)u(a) + Tv \geq \text{con}(a)$, which implies that the first a can be played whenever the second is followed by any equilibrium path beginning with a .

In this example, the action profile b can be followed by the two action profiles a and c because $P^2(b) = \{ba, bc\}$. When b is followed either by a or c , it does not matter what comes after these action profiles as long as the payoffs are in $C_a(V)$ and $C_c(V)$, respectively. 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. For example, we can form SPE paths like $ddbcbcb a^\infty, d^9(cb)^\infty, d^\infty$, and a^∞ from these elementary subpaths. The relevant information on how to create all the paths from the elementary subpaths of this example is condensed in the graph of Figure 1 in the Introduction.

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

In general, a singleton action profile is an elementary subpath if and only if it is a Nash equilibrium. Thus, the one-length elementary subpaths are exactly the Nash equilibria of the stage game. In the above example, the Nash equilibrium action profile d is a singleton elementary subpath. Also, a constant sequence of action profiles such as $a \cdots a$ can be an elementary subpath only if it consists of exactly two elements, i.e., aa .

3. Properties of elementary subpaths

The main questions on elementary sets are whether they are finite, and what happens to them when the discount factors increase. These issues are considered in this section.

In general, $S(T)$ may contain infinitely many subpaths. The set $S^k(T)$, on the other hand, contains finitely many subpaths because $P^k(a)$ are finite for all k and $a \in A$. However, it can be shown that $S(T)$ is a finite set when the discount factors are small enough. Note, in particular that in such cases the set of infinitely long elementary subpaths is also finite. The finiteness of elementary subpaths is revisited in Section 4.2, where it is shown that $S(T)$ may become finite when the payoffs corresponding to equilibrium paths are bounded away from the sets of payoffs for which the players incentive-compatibility constraints are binding. Moreover, Proposition 6 shows that for any $\varepsilon > 0$, it is possible to find a finite graph representation for all the equilibrium paths if the subpaths are allowed to be ε -incentive compatible.

Proposition 2. *The set $S(T)$ is finite when δ_i , $i \in N$, are small enough.*

Proof. Let $NE \subseteq A$ stand for the pure-strategy Nash equilibria of the stage game. If $a \notin NE$, then the deviation payoffs satisfy $u_i^*(a) > u_i(a)$ for some $i \in N$. Let $N'(a)$ denote the set of players for which this inequality holds for $a \notin NE$.

Assume that the punishment payoffs $v_i^- \leq \max\{u_i(a) : a \in A\}$, $i \in N$, are exogenously given. Let us define $\delta_i(a, v^-)$ for any $a \notin NE$ and $i \in N'(a)$ as the discount factor δ_i that solves

$$(1 - \delta_i)u_i(a) + \delta_i \max_{a \in A} u_i(a) = (1 - \delta_i)u_i^*(a) + \delta_i v_i^-.$$

Note that $\max\{u_i(a) : a \in A\}$ is an upper bound for the continuation payoff of player $i \in N$. Moreover, $\delta_i(a, v^-) \in (0, 1]$ is well-defined when $u_i^*(a) > u_i(a)$, i.e., $i \in N'(a)$.

Let us set

$$\delta_0(v^-) = \min_{a \notin NE} \max_{i \in N'(a)} \delta_i(a, v^-).$$

If $\delta_i < \delta_0(v^-)$ for all $i \in N$, then no action profile $a \notin NE$ can be played when the punishment payoffs are v^- , because for any $a \notin NE$ there is some player $i \in N$ whose incentive-compatibility condition (2) fails to hold for any v_i such that $v_i^- \leq v_i \leq \max\{u_i(a) : a \in A\}$.

If $NE = \emptyset$, we can set v_i^- , $i \in N$, to the player's minimax payoff, i.e., v_i^- is the lower bound for the value of the least equilibrium payoff for player $i \in N$. It follows that no action profile can be played when $\delta_i < \delta_0(v^-)$ for all $i \in N$, in which case $V = \emptyset$ and $S(T) = \emptyset$. Hence, $S(T)$ is finite.

If $NE \neq \emptyset$, then v_i^- , $i \in N$, can be chosen as the player's smallest Nash equilibrium payoff in the stage game. When $\delta_i < \delta_0(v^-)$ for all $i \in N$, only the Nash equilibrium action profiles can be played. Hence, $S(T)$ is finite for discount factors small enough. \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 $k \geq 1$.

The first result, which is of importance itself, tells that if the punishment payoffs are non-increasing 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 payoff sets corresponding to T_1 and T_2 are denoted by $V(T_1)$ and $V(T_2)$, respectively. The punishment payoffs are $v^-(V(T_1))$ and $v^-(V(T_2))$. Complementing the following result, it is shown in [7] that the equilibrium paths may not be monotone in the discount factor if the punishment payoffs increase. Note also that the monotonicity of paths holds even though the equilibrium payoffs may fail to satisfy the monotone comparative statics.

Proposition 3. *If $T_1 \ll T_2$ and $v^-(V(T_1)) \geq v^-(V(T_2))$, then $paths(S(T_1)) \subseteq paths(S(T_2))$.*

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 $d^k = u^*(a^k)$ denote the vector of deviation payoffs. Moreover, v^i 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

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

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

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

Similar expression as for S_1^k can be derived for S_2^k with T_2 . The purpose of the proof is to show that $S_2^k \geq \mathbf{0}$, $k \geq 0$, 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

$$S_i^k = u^k - d^k + T_i (d^{k+1} - v^i + S_i^{k+1}), \quad \forall k \geq 0 \text{ and } i = 1, 2. \quad (4)$$

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

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 the smallest payoffs are v^2 and after that the discount factors correspond to T_1 and the smallest payoffs are v^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 (4) gives

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

Recall that the components of the vector $d^{K+1} - v^1 + S_1^{K+1}$ are non-negative. Hence, we get $S_2^K \geq S_1^K \geq \mathbf{0}$. It can now be seen from the recursion (4) and the inequalities $T^2 \gg T^1$, $d^j - v^2 \geq d^j - v^1$, $j \leq K$, that $S_2^j \geq S_1^j \geq \mathbf{0}$ 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 4. *If $T_1 \ll T_2$ and $v^-(V(T_1)) \geq v^-(V(T_2))$, then $p \in S(T_1)$ implies that there is $k \leq |p|$ such that $p^k \in S(T_2)$.*

Proof. Let us first show that any payoff $B_{p_1}(v(q))$ is an equilibrium payoff when $p \in S(T)$ and q is an SPEP such that $v(q) \in C_{f(p)}(V)$. If p is infinitely long, so is p_1 . In particular, p_1 is an SPEP, i.e., $B_{p_1}(C_{f(p)}(V))$ is a singleton SPE payoff corresponding to p_1 . Hence, the claim holds for infinitely long subpaths. Next, assume that p is finitely long.

Take an SPEP q with the corresponding payoff $v(q) \in C_{f(p)}(V)$. Any payoff in $B_{f(p)}(C_{f(p)}(V))$ is an equilibrium payoff by definition. Hence, $f(p)q$ is an SPEP. If $|p| \geq 2$, then by Proposition 1, there are $r \in S(T)$, j , and k such that $r^j = (p_{|p|-2}q)^k$. Hence, there are no profitable one-shot deviations when playing the second last action profile $p_{|p|-2}^1$ of p , i.e., $B_{p_{|p|-2}}(v(q))$ is an SPE payoff. By repeating the argument, it follows that $B_{p_1}(v(q))$ is an equilibrium payoff that corresponds to an equilibrium path p_1q . Hence, the claim holds.

In the following $P^k(a; T_i)$ denotes the set of k -length elementary subpaths corresponding to T_i , $C_a^i(V(T_i))$ 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 B_p^i is the operator B_p corresponding to T_i .

The above deduction implies that $B_{p_1}^1(v^1(q))$ is an SPE payoff for an SPE path q such that $v^1(q) \in C_{f(p)}^1(V^1)$. By Proposition 3 the same holds for $B_{p_1}^2(v^2(q))$. Moreover, if $(i(p), v^1(p_1q))$ is admissible for T_1 , so is $(i(p), v^2(p_1q))$ for T_2 , i.e., pq is an SPEP for both T_1 and T_2 . Because q can be any equilibrium

path such that $v^1(q) \in C_{f(p)}^1(V(T_1))$, this means that the condition (3) for T_1 implies that it holds also for T_2 . Hence, either p is an elementary subpath or there is $k < |p|$ such that p^k is an elementary subpath. \square

When the discount factors increase, all the subpaths that satisfy (3) still satisfy this condition if the smallest equilibrium payoffs do not increase. 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 .

4. Computation and approximation of elementary sets

An algorithm for finding the elementary sets of supergames is presented in this section. The algorithm produces all the elementary subpaths if it terminates in a finite number of steps. Otherwise, the set obtained from the algorithm can be used as an approximation. The approximations obtained this way are related to the approximate equilibria of the supergame.

The algorithm presented in this section may produce subpaths that contain non-equilibrium parts. However, these subpaths can be removed easily as is explained in Section 5.

4.1. Algorithm for finding the elementary subpaths

First, we introduce a recursive way of computing the continuation 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 $\text{con}_i(ab)$ would be below v_i^- then we simply set $\text{con}_i(ab) = v_i^-$. 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^- if it would be below that value. If $\text{con}(abc) \leq \text{con}(c)$, then any equilibrium path starting from c is an admissible continuation for abc at the time when a is played. Note that the same idea appears 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

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

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): $\text{con}(p) \leq \text{con}(a)$ is a sufficient for (3). On the other hand, if it happens that $\text{con}_i(p) > v_i^+$ for some $i \in N$, then $\text{con}(p)$ is certainly outside of $C_{i(p)}(V)$ and p cannot be an elementary subpath.

Remark 2. Condition (3) holds for $p \in A^k$ with $f(p) = a$ if $\text{con}(p) \leq \text{con}(a)$. If $\text{con}_i(p_j) > v_i^+$ for some $j = 0, \dots, |p| - 1$ and $i \in N$, then p is not an elementary subpath.

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 contain non-equilibrium parts, which can be subsequently removed. The removal of these subpaths will be explained in Section 5.1. The algorithm is demonstrated in Section 6. Let $P_*^k(a)$ denote the working set of the algorithm containing the subpaths of length k starting with a .

1. For all $a \in A$, include $a \in \hat{P}^1(a)$ if $\text{con}_i(a) \leq v_i^-$ for all $i \in N$. If $v_i^- \leq \text{con}_i(a) \leq v_i^+$ 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 $a, b \in A$, $p \in P_*^{k-1}(a)$, compute $\text{con}(q)$ for $q = pb$.
 - (a) If $\text{con}(q) \leq \text{con}(b)$ and

$$q_j \in P_*^{k-j}(i(q_j)) \text{ or } q_j^l \in \hat{P}^l(i(q_j)), \quad \text{for some } 1 \leq l \leq k - j, \quad (5)$$

$$\forall j = 1, \dots, k - 1,$$

then include q in $\hat{P}^k(a)$.

- (b) Otherwise, if $\text{con}_i(q) \leq v_i^+$ for all $i \in N$ and q satisfies Eq. (5), 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 to obtain P^k . This will be explained in Section 5.1.

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 v_i^+ , $i \in N$. Second, all parts of the subpath must satisfy the 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 because there cannot be any more elementary subpaths that have not yet been found. Moreover, in that case the elementary set is finite.

Remark 3. If there is k such that $P_*^k(a) = \emptyset$ for all $a \in A$, then $S(T)$ contains finitely many subpaths.

If the algorithm is terminated after k steps and some of the sets $P_*^k(a)$, $a \in A$, are non-empty, it is possible to find the infinitely long elementary subpaths that are of the form $p = qb^\infty$ for some $b \in A$.

Remark 4. If the algorithm is terminated after k steps and $P_*^k(a) \neq \emptyset$ for some $a \in A$, then for each $p \in P_*^k(a)$ the path $q = pb^\infty$ belongs to $P^\infty(a)$ if $b^j \in P^j(b)$ for some $j \leq k$, $b \in A$, and $u(b) \geq \text{con}(p)$.

The above observation is useful, because it means that if all the infinitely long elementary subpaths are of the form $p = qb^\infty$, $b \in A$, where $|q| \leq k$ and $b \in A$, then they can be found when terminating the algorithm after k steps and going through the possible combinations of $q \in P_*^k(a)$ and $b^j \in P^j(b)$ for $j \leq k$.

In [9], it is found that the elementary set is finite for up to $\delta \approx 0.6$ for the repeated prisoner's dilemma, chicken, and stag hunt games, and up to $\delta \approx 0.9$ for the repeated leader game; see Table 1 in [9] and the entries without * and ** therein. For these supergames, the algorithm terminates after finitely many steps even for reasonably large discount factors.

When the discount factors are close to one, there is no guarantee that the algorithm would terminate after finitely many steps, because there may be elementary subpaths that lead to payoffs arbitrarily close to the boundary of some of the players' incentive-compatibility conditions. Consequently, there may be arbitrarily long elementary subpaths.

The algorithm can be terminated while there still are elements in P_* , in which case a subset of the elementary set is obtained, because some of the elementary subpaths may not have been found. As Proposition 5 in the following section tells, the missing subpaths give payoffs close to the boundary of some of the players' incentive-compatibility conditions. When terminating the algorithm while P_* is nonempty, it is possible to obtain approximate equilibria. This topic is further discussed in Section 4.2.

The algorithm for computing the elementary set uses the smallest and the highest equilibrium payoffs v_i^- and v_i^+ , $i \in N$. However, these payoffs are typically not known in advance. The highest payoffs v_i^+ , $i \in N$, need not be the highest equilibrium payoffs; they can be replaced by the highest stage game payoffs. These values affect how fast the algorithm finds the non-elementary subpaths and thus how fast it converges. The punishment payoffs v_i^- , $i \in N$, are easily determined for many games. However, if they are not known the following algorithm [7, 11] can be used in finding them. This method is the first that has been proposed for finding the smallest equilibrium payoffs. Hence, there is potential for further development of algorithms for computing the most severe punishments.

Initialize the lower bounds

$$l_i = \min_{a_{-i} \in A_{-i}} \max_{a_i \in A_i} u_i(a).$$

While the punishment paths are not found do

1. Find the elementary subpaths corresponding to the punishment payoffs l_i , $i \in N$.
2. Find the smallest payoffs m_i , $i \in N$, from the equilibrium paths constructed from the subpaths found in Step 1.
 If $m_i > l_i$ for some $i \in N$ then update $l_i = m_i$.
 Otherwise the payoffs v_i^- , $i \in N$, are found.

The algorithm sets the minimax values as the smallest payoffs and updates these lower bounds until the correct values are found. The elementary subpaths, or a subset of the elementary set corresponding to the current punishment payoffs l_i , $i \in N$, are computed in Step 1, as if l_i , $i \in N$, were the punishment payoffs. In Step 2, the new punishment payoffs are searched from the resulting paths. For this purpose, it is possible to utilize the graph presentation of elementary subpath, see [7, 11] for more details. Forming the graph is presented in Section 5.1.

4.2. Approximating the elementary set and approximate equilibria

As mentioned, the algorithm for finding the elementary subpaths can be terminated while there still are elements in P_*^k , in which case we get a subset of the elementary set, because some of the elementary subpaths may not have been found. As shown in the following result, the missing subpaths give payoffs close to the boundary of some of the players' incentive-compatibility conditions.

Proposition 5. *For any $\varepsilon > 0$, there is k such that $p^l \in P^l(i(p))$ for some $l \leq k$ when $p \in A^\infty(a)$ is an SPE path, $a \in A$, and*

$$v_i(p_1) \geq \text{con}_i(a) + \varepsilon, \quad \forall i \in N. \quad (6)$$

Proof. Let Q stand for the set

$$\times_{i \in N} [v_i^-, \max_{a \in A} u_i(a)].$$

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 Q under a sequence $B_{a^0}, \dots, B_{a^{k-1}}$, $a^j \in A$ for $j = 0, \dots, k-1$, has diameter less than ρ . Because $V \subseteq Q$, it follows that the diameter of the set $B_{p_1^k}(C_{f(p_1^k)}(V))$ is less than ρ for any p . This together with condition (6) implies that when choosing ρ small enough and k large enough we have

$$(1 - \delta_i) \sum_{j=0}^{k-1} \delta_i^j u_i(i(p_{j+1})) + \delta_i^k v_i \geq \text{con}_i(i(p)),$$

for all $i \in N$ and $v \in C_{f(p_1^k)}(V)$. Gathering the components of the left hand side of the above equation for $i \in N$ into a vector, it reads as $B_{p_1^k}(v)$. Hence, ρ

can be chosen such that for any $a \in A$ and $p \in A^\infty(a)$ for which (6) holds we have

$$B_{p_1^k} \left(C_{f(p_1^k)}(V) \right) \subseteq C_{i(p)}(V).$$

By Remark 1 either p^{k+1} is an elementary subpath or a shorter fragment of it is elementary. \square

The above result means that if Eq. (6) holds for all the equilibrium paths then there is a bound to the maximum length of elementary subpaths, which implies that the elementary set is finite. It also tells that an elementary subpath $p \in \mathcal{A}(a)$ that is not found when the algorithm is terminated with non-empty P_*^k and large k , is such that v_i is close to $\text{con}_i(a)$ for some $i \in N$ and any payoff $v \in B_{p_1}(V) \cap C_a(V)$. Hence, by taking k large enough we should be able to obtain a good approximation of the elementary set when the final set of subpaths is obtained either by removing or including P_*^k into the set of subpaths.

Let us first consider the outer approximation for the elementary set that is obtained by including the remaining $P_*^k \neq \emptyset$ into \hat{P}^k . In this case, the final set contains subpaths that produce all the SPE paths of the game but there may also be subpaths that do not satisfy the players' incentive-compatibility conditions at all stages.

Definition 5. The k -step outer approximation $S_O^k(T)$ of $S(T)$ is the set of subpaths obtained by including $P_*^j(a)$, $a \in A$, $j \leq k$, produced by the algorithm into $\hat{P}^j(a)$, $a \in A$, $j \leq k$, and then removing the non-equilibrium parts.

The paths obtained from $S_O^k(T)$ contain all the SPEPs. Hence, it is reasonable to call $S_O^k(T)$ an outer approximation. The outer approximation is related to approximate equilibria of the supergame. We say that a strategy profile σ is an ε -incentive-compatible equilibrium if no player has a one-shot deviation from σ that would benefit the deviating player by at most ε . A one-shot deviation from a strategy profile σ means that a player i chooses σ'_i that differs from σ only for one history of past play; there is exactly one period k in which the action prescribed by σ' is different from the action prescribed by σ_i . The usual notion of ε -equilibria would allow for more general deviations that do not benefit any player by more than ε . In other words, σ is an ε -equilibrium if

$$U_i(\sigma|p) \geq U_i(\sigma'_i, \sigma_{-i}|p) - \varepsilon, \quad \forall i \in N, \quad p \in A^k, \quad 0 \leq k < \infty, \quad \text{and} \quad \sigma'_i \in \Sigma_i.$$

Hence, ε -incentive-compatible equilibria are a subset of ε -equilibria. A path that is induced by an ε -incentive-compatible equilibrium strategy profile is an ε -incentive-compatible equilibrium path.

As stated below, $S_O^k(T)$ may give ε -incentive-compatible equilibrium paths of the supergame. However, not all of the ε -incentive-compatible equilibrium paths are necessarily obtained, because the punishment payoffs corresponding to ε -incentive-compatible equilibria may be lower than v_i^- , $i \in N$, used in the algorithm. Since $S_O^k(T)$ contains all the equilibrium paths, Proposition 6 implies that for any $\varepsilon > 0$, it is possible to compute in a finite number of steps a graph representation for the equilibrium paths if we allow the subpaths to be ε -incentive compatible.

Proposition 6. *For any $\varepsilon > 0$, there is k such that paths $(S_O^k(T))$ are ε -incentive-compatible equilibrium paths.*

Proof. The purpose is to show that choosing k large enough, the violation of the incentive compatibility condition can be made arbitrarily small for any $p \in \text{paths}(S_O^k(T))$, which implies that p is ε -incentive compatible equilibrium path for any $\varepsilon > 0$ when k is large enough.

Take a path $p = a^0 a^1 \cdots \in \text{paths}(S_O^k(T))$ and let $u^*(a^0)$ denote the vector of deviation payoffs as before. Moreover, w_i^k denotes player i 's smallest payoff corresponding to paths $(S_O^k(T))$, and the player's smallest stage game payoff is u_i^- . The corresponding payoff vectors are w^k and u^- , respectively. Note that $u_i^- \leq w_i^k$ for all $i \in N$.

By the definition of $\text{con}(p_k)$, we have

$$(I - T) \sum_{j=0}^k T^j u(a^j) + T^{k+1} \text{con}(p_k) = (I - T)u^*(a^0) + Tv^-. \quad (7)$$

Moreover, the step 2 (b) of the algorithm assures that $\text{con}(p_k) \leq v^+$.

The largest possible violation of the incentive-compatibility condition at the time when a^0 is played and is followed by p_1 , is obtained when the players' payoffs after p are w_i^k , $i \in N$. It follows from (7) that

$$\begin{aligned} (I - T) \sum_{j=0}^k T^j u(a^j) + T^{k+1} w_i^k &= (I - T)u^*(a^0) + Tv^- - T^{k+1} [\text{con}(p_k) - w^k] \\ &\geq (I - T)u^*(a^0) + Tv^- - T^{k+1} (v^+ - w^k) \\ &\geq (I - T)u^*(a^0) + Tv^- - T^{k+1} (v^+ - u^-). \end{aligned}$$

The components of $T^{k+1}(v^+ - u^-) \geq \mathbf{0}$ can be made arbitrarily small by choosing k large enough. Hence, (a, v) becomes ε -admissible for $v = (I - T) \sum_{j=0}^k T^j u(a^{j+1}) + T^{k+1} w^k$, when k is large enough. To be specific, it holds that

$$(1 - \delta_i)u_i(a^0) + \delta_i v_i \geq (1 - \delta_i)u_i^*(a^0) + \delta_i \hat{v}_i^- - \varepsilon, \quad \forall i \in N,$$

where $\hat{v}_i^- \leq w_i^k \leq v_i^-$ is the smallest ε -incentive-compatible equilibrium payoff for player $i \in N$. The same deduction holds for any action profile a belonging to any subpath of $S_O^k(T)$. Hence, for any $\varepsilon > 0$, it is possible to find k large enough such that any $p \in \text{paths}(S_O^k(T))$ is an ε -incentive-compatible equilibrium path. \square

In addition to forming an outer approximation of $S(T)$ it is possible to form an inner approximation by removing the paths belonging to P_* from \hat{P} . The set of subpaths obtained this way can be used in obtaining a subset of SPEPs.

Definition 6. The k -step inner approximation is the set $S_I^k(T)$ obtained when terminating the iteration after k steps and excluding $P_*^j(a)$, $a \in A$, $j \leq k$, from $\hat{P}^j(a)$, $a \in A$, $j \leq k$.

The inner approximation is related to ε -strict incentive-compatible equilibria in an analogous manner as the outer approximation is related to ε -incentive-compatible equilibria. A strategy profile σ is called an ε -strict incentive-compatible equilibrium if all one-shot deviations from σ for any player lead to payoffs that are worse than the deviating player's original payoff by at least ε . Note that the more common notion of ε -strict equilibrium means that a strategy profile σ satisfies

$$U_i(\sigma|p) \geq U_i(\sigma'_i, \sigma_{-i}|p) + \varepsilon, \quad \forall i \in N, \quad p \in A^k, \quad 0 \leq k < \infty, \quad \text{and } \sigma'_i \in \Sigma_i.$$

These strategies would allow for more general deviations than only one-shot deviations provided that they lead to payoffs that are worse than the original payoff by at least ε . Any strategy that is not ε -strict incentive-compatible equilibrium cannot be ε -strict equilibrium either. Hence, ε -strict equilibria are a subset of ε -strict incentive-compatible equilibria.

An ε -strict incentive-compatible equilibrium path is a path of action profiles induced by an ε -strict incentive-compatible equilibrium strategy. The inner approximation can be used in producing all these paths. However, the set of paths obtained from $S_I^k(T)$ may also contain other SPEPs.

Proposition 7. *For any $\varepsilon > 0$ there is k such that all ε -strict incentive-compatible equilibrium paths are included in the set $\text{paths}(S_I^k(T))$.*

Proof. For a path $p = a^0 a^1 \dots \in \text{paths}(S_I^k(T))$ and for any $j = 0, 1, \dots$ there is $l^1 \leq k$ such that

$$(I - T) \sum_{j=0}^{l_1} T^j u(a^j) + T^{l_1+1} \text{con}(f(p_j^{l_1})) \geq (I - T)u^*(a^0) + Tv^-.$$

On the other hand, the same holds for a^{j+l_1+1} and some $l_2 \geq k$. Hence, for any a^j , we can choose $k' = l_1 + \dots + l_r \geq k$, where r is the smallest number greater than k for which $l_1 + \dots + l_r \geq k$, and it holds that

$$\begin{aligned} (I - T) \sum_{j=0}^{k'} T^j u(a^0) + T^{k'+1} v(p_{k'+1}) &\geq (I - T)u^*(a^0) + Tv^- \\ &\quad + T^{k'+1} [v(p_{k'+1}) - \text{con}(f(p_{k'}))]. \end{aligned}$$

Because

$$0 \leq T^{k'+1} [v(p_{k'+1}) - \text{con}(f(p_{k'}))] \leq T^{k+1} [v(p_{k+1}) - \text{con}(f(p_k))],$$

and k can be chosen large enough such that

$$T^{k+1} [v(p_{k+1}) - \text{con}(f(p_k))] \leq (\varepsilon, \dots, \varepsilon),$$

it is possible to create all paths $p = a^0 a^1 \dots$ for which it holds that

$$(I - T)u(a^0) + Tv(p_1) \geq (I - T)u^*(a^0) + Tv^- + (\varepsilon, \dots, \varepsilon) \quad (8)$$

from the paths of $S_I^k(T)$.

A path $p = a^0 a^1 \dots$ that is an ε -strict incentive compatible satisfies

$$\begin{aligned} (I - T)u(a^j) + Tv(p_{j+1}) &\geq (I - T)u^*(a^j) + T\hat{v}^- + (\varepsilon, \dots, \varepsilon) \\ &\geq (I - T)u^*(a^j) + Tv^- + (\varepsilon, \dots, \varepsilon), \quad \forall j = 0, 1, \dots, \end{aligned}$$

where $\hat{v}^- \geq v^-$ is the vector of least ε -strict incentive-compatible equilibrium payoffs. Hence, the paths that satisfy the condition (8) contain all ε -strict incentive-compatible equilibrium paths. \square

5. Graph presentation and the complexity of equilibria

The algorithm presented in the previous section may produce subpaths, whose incentive compatibility relies on sets P_*^k . These sets do not necessarily form any elementary subpaths 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 have already been found and with the same effort as forming a graph for all the equilibrium paths when there are finitely many subpaths. This section describes how to form the graph presentation. The graph is useful for producing the equilibrium paths and payoffs and in analyzing the complexity of equilibrium outcomes.

5.1. Forming the graph presentation

The algorithm for forming the graph 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 remove node p and connect node $p^{|p|-1}$ 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 3. 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 2. We note that subpath bac contains a non-equilibrium part, because 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 during 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 a . According to Step 2.(c), node aa connects to a and node ab to b , because $p_1 = a$ 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 2. 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 .




Figure 2: An example of elementary subpaths as a tree and a graph.

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 [9]. 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 [10, 32]. Proposition 2 guarantees that the graph presentation is possible when the discount factors are small enough and the stage game has pure-strategy Nash equilibria.

Proposition 8. *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 represented by a graph.*

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

Corollary 3. *When the discount factors δ_i , $i \in N$, are small enough, $S(T)$ can be represented by a graph, and the payoff set $V(T)$ is a graph-directed self-affine set.*

5.2. Complexity of equilibrium outcomes

The complexity of equilibrium paths and payoffs can be analyzed with the graph presentation of the elementary subpaths. We emphasize that this approach to the complexity of equilibria differs fundamentally from previous literature on the topic, where the focus has been on computational complexity [22] or the complexity of individual strategies [28, 15]. Here, complexity refers to the complexity of the set of all possible equilibrium paths.

Assume that the elementary subpaths (or their approximation) is turned into a graph. A graph is 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 of the graph can be used in counting the number of walks in the graph [19], 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

$$y(k) \approx y_0 \rho^k(D), \tag{9}$$

where y_0 is a constant and $\rho(D)$ is the largest eigenvalue of D ; see, e.g., Theorem 2.2.2 in [19]. 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 [32]

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

¹The exact dimension can be defined when so called open set condition holds, which means 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 [23].

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 [10], it is shown how to estimate the Hausdorff dimension when the players have different discount factors.

Example 4. Consider the graph of Figure 2. The adjacency matrix is

$$D = \begin{pmatrix} 0 & 1 & 1 & 1 & 0 \\ 0 & 1 & 1 & 1 & 0 \\ 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 1 & 0 \end{pmatrix} \quad (10)$$

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 $s = \log_2 \varphi \approx 0.694$.

6. Examples

6.1. Prisoner's dilemma game

In this section, it is demonstrated how the algorithm finds the elementary subpaths, how the graph is constructed, and how to analyze the equilibria with the graph. Consider the prisoner's dilemma game with a common discount factor $\delta = 1/2$ and the stage game payoffs are as in Example 1. The payoff sets of the repeated prisoner's dilemma have previously been studied in [35], [36] and [31]. Here, we show the exact paths that can be played in the game and analyze what happens when the discount factor increases.

In this game, the punishment path is the infinite repetition of d , which is denoted by d^∞ . The corresponding payoffs are $v_i^- = 1$, $i = 1, 2$. Let us now 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, because in this example there are no subpaths with non-equilibrium parts. In Step 1 of algorithm in Section 4.1, we calculate $\text{con}(p)$ for one-length paths p . For example, $\text{con}(d) = (1, 1)$ and d is an elementary subpath because $\text{con}_i(d) \leq v_i^-$, $i = 1, 2$. For paths a , b and c , we have $v^- \leq \text{con}(p) \leq v^+$ and $P_*^1 = \{a, b, c\}$. Table 2 gives the 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 $*$.

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 get

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

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	$(1, 3)^*$

At this point aa , ba , bc , ca and cb are found to be elementary, but because ad , bd , and cd belong to P_*^2 , all the elementary subpaths are not yet found. It can be observed that ad is incentive compatible only when it is followed by an infinite repetition of a , i.e., $P^\infty(a) = \{ada^\infty\}$, because no other action profile gives the required payoff $(3, 3)$. Thus, we need not consider other subpaths starting with ad and it is removed from the set P_*^2 , which is a minor deviation from the algorithm. Note that this way to deduce infinitely long elementary subpaths corresponds to the observation made in Remark 4. On the other hand, it could be tested whether a subpath belonging to P_*^k , $k \geq 2$, could be part of an equilibrium path only when it is followed by a continuation payoff produced by paths belonging to P^j (or \hat{P}^j), $j < k$. This would provide another way to remove ad from P_*^2 in this example. To simplify the exposition, this step is excluded from the algorithm presented in Section 4.1.

Due to the symmetry of the game, only the paths beginning with either b or c need to be analyzed. Consider the three and four-length paths beginning with cd in Table 3. 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 $b \in P_*^1(b)$.

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

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

Now, it can be seen that the only possible paths starting with cd are cda and cdb . The only continuation to cda is aa , because the only elementary subpaths starting with a are aa and ad , and ad gives lower payoff than the required $(1, 3)$. Thus, $cda^\infty \in P^\infty(c)$ and we need not consider other subpaths starting with cda . From four-length 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 the elementary set has been found:

$\{d, aa, ba, bc, ca, cb, cdba, bdca, ada^\infty, bda^\infty, cda^\infty, bdcda^\infty, cdbda^\infty\}$. It is possible to create all the SPE paths with these elementary subpaths. For example, the SPE paths $d^k ada^\infty$, $d^k bda^\infty$, $d^k cda^\infty$, $d^k bdcda^\infty$, $d^k cdba^\infty$, and d^∞ , are all obtained by combining the elementary subpaths. In particular, these paths cannot themselves be elementary.

The tree of finite elementary subpaths is presented in Figure 3. The destinations of leaf nodes are indicated next to them. Using the graph algorithm without Step 3, we get the directed graph composed of solid arcs in Figure 3. Each node denotes what is played when the node is visited.

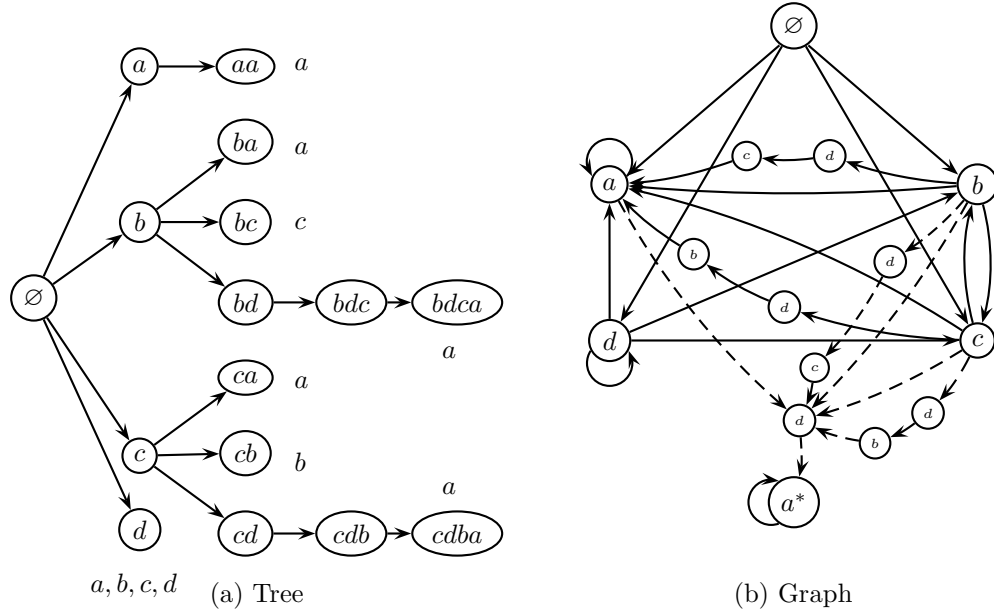


Figure 3: Tree of finite elementary subpaths and a graph of all the equilibrium paths.

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. This extra node is denoted by a^* and after adding the new nodes and arcs we get the graph of Figure 3 in which the dashed arcs are also included.

An approximation of the payoff set is shown in the left of Figure 4. The payoff set consists of similar patterns in different scales, which shows the fractal nature of equilibrium payoffs. 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 4: 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 4. 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 , because there are always elementary subpaths in the proximity of the payoff requirement values as the payoff set becomes dense.

6.2. 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. The following game is called the Sierpinski game because the payoff set is 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 5. 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 action profiles that correspond to the payoff vector $(-1, -1)$ cannot be played because there is no payoff vector v for which the incentive-compatibility conditions (2) would hold. For instance, if we take $a = (B, L)$, the incentive-compatibility condition becomes $(1 - \delta)(-1, -1) + \delta v \geq (1 - \delta)(2 - \sqrt{3}, 0) + \delta(0, 0)$. Because $v_i \leq 1$, $i = 1, 2$, this condition cannot hold unless δ is at least $(3 - \sqrt{3})/(4 - \sqrt{3}) > 1/2$.

In this example, the payoff set is the Sierpinski triangle. Its 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 5: 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 [25]. One reason for the fact that any feasible payoff above minimax 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.

6.3. Three-player game

In this example, there are three players: the row player, the column player, and the matrix player. The third player chooses from the three alternatives α , β and γ , while the other two players have two actions from which to choose. Altogether, there are twelve action profiles and the payoffs are listed below.

	L	R		L	R
T	3, 3, 0	$\varepsilon, 3 + \varepsilon, -10$	T	3, 3, 0	$0, 3 + \varepsilon, -10$
B	3, 3, 3	0, 3, 0	B	$3, \varepsilon, -10$	3, 0, 0
	α			β	
	L	R		L	R
T	0, 0, 0	3, -10, 0			
B	-10, 0, 3	-10, -10, 0			
	γ				

The parameter ε is assumed to be non-negative. Its choice affects the incentive-compatibility conditions: the larger the parameter, the larger the continuation payoff requirement for the action profiles involving ε in the payoff vector.

It can be seen that each player's minimax payoff is zero. This is because (T, L, γ) is a Nash equilibrium and each player can guarantee at least zero payoff by playing T , L , or γ . When $\varepsilon > 0$, the game has two pure-strategy Nash equilibria; (B, L, α) and (T, L, γ) . For $\varepsilon = 0$, all action profiles in which the least payoff for the players is larger than -10 are Nash equilibria.

When the discount factors are small enough, the action profiles in which at least one of the players gets the payoff -10 cannot be played no matter what the continuation payoffs are. This is simply because the largest continuation payoff is only 3: a player gets a negative total payoff by first receiving -10 and then 3 given that the discount factor is small enough. In particular, this is the case when we choose $\delta_1 = \delta_2 = \delta_3 = 1/3$. For these discount factors, there are six action profiles that can be played (T, L, α) , (B, L, α) , (B, R, α) , (T, L, β) , (B, R, β) , and (T, L, γ) . Hence, the upper bound for the asymptotic growth rate is 6 (this corresponds to $\varepsilon = 0$).

In this example, the equilibrium payoff set lies in the cube $[0, 3] \times [0, 3] \times [0, 3]$. See Figure 6 for the payoff set when $\varepsilon = 0.1$ and $\delta_i = 1/3$ for all $i \in N$. The Hausdorff dimension of the payoff set is approximately 1.61, which is a little below 1.63 corresponding to $\varepsilon = 0$. Indeed, one way to obtain an upper bound for the asymptotic growth rate and the Hausdorff dimension is to assume that all the six action profiles can be played in any order, i.e., assume $\rho(D) = 6$ and calculate $s = -\log \rho(D) / \log \delta \approx 1.63$.




Figure 6: Payoff set for $\varepsilon = 0.1$ and $\delta_1 = \delta_2 = \delta_3 = 1/3$.

7. Conclusions

The main result of this work is that the pure-strategy equilibrium paths of repeated games are composed of sequences of action profiles, which are called the elementary subpaths. This means that there is a set of fragments of equilibrium paths which can be used in producing all the equilibrium paths. 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 the equilibrium paths and the corresponding payoffs [9].

The elementary subpaths are defined relative to the set of equilibrium payoffs: a subpath is elementary if its first action profile can be played when the last action profile of the subpath is followed by any possible continuation payoff at that stage, and all the action profiles in the subpath correspond to some elementary subpaths. However, we present an algorithm for finding the elementary subpaths that requires knowing only the smallest equilibrium payoffs. When there are finitely many elementary subpaths, the algorithm produces the whole set of elementary subpaths, the elementary set, in a finite number of steps. It is shown that the elementary set of a repeated game is finite at least when the discount factors are small enough. Moreover, the set of equilibrium paths increases monotonically, when the players' discount factors increase, and the smallest equilibrium payoffs do not increase.

The algorithm can also be used in computing approximate equilibrium paths. When the algorithm is terminated such that not all the elementary subpaths have been found, the missing elementary subpaths lead to equilibrium payoffs for which some of the players' incentive-compatibility conditions are close to be binding. In such cases, it is possible to form finite inner and outer approximations of the elementary set: approximations that give either a subset of all the equilibrium paths or both equilibrium paths and some approximate equilibria. In particular, for any $\varepsilon > 0$, it is possible to compute in a finite number of steps a graph that represents all the equilibrium paths of the repeated game when the subpaths are allowed to be ε -incentive compatible.

The final step of the algorithm for finding the elementary set transforms the found subpaths into a directed graph, which is a compact representation of all the equilibrium paths. The graph can be used in generating the 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 higher 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 the equilibrium payoff set. This paper lays foundations for further research, e.g., on extending the methodology to stochastic games [8] or mixed strategies [12], on designing algorithms for computing equilibria and minimum payoffs [9, 11], and on the analyses of the fractal properties of payoff sets [10].

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, *Theoretical Economics* 9 (2014) 313–338.
- [6] G. Andersen, V. Conitzer, Fast Equilibrium Computation for Infinitely Repeated Games, *Proceedings of the Twenty-Seventh AAAI Conference on Artificial Intelligence* (2013) 53–59.
- [7] K. Berg, Extremal pure strategies and monotonicity in repeated games, *Computational Economics* (to appear) (2016).
- [8] K. Berg, Elementary subpaths in discounted stochastic games, *Dynamic Games and Applications* (to appear) (2015).
- [9] K. Berg, M. Kitti, Computing equilibria in discounted 2×2 supergames, *Computational Economics* 41 (2013) 71–78.
- [10] K. Berg, M. Kitti, Fractal geometry of equilibrium payoffs in discounted supergames, *Fractals* 22 (2014).
- [11] K. Berg, M. Kärki, An algorithm for computing the minimal pure-strategy subgame-perfect equilibrium payoffs in repeated games, *Working paper* (2014).
- [12] K. Berg, G. Schoenmakers, Construction of mixed subgame-perfect equilibria in repeated games, *Working paper* (2014).
- [13] C. Borgs, J. Chayes, N. Immorlica, A.T. Kalai, V. Mirrokni, C. Papadimitriou, The myth of the Folk theorem, *Games and Economic Behavior* 70 (2010), 34–43.
- [14] A. Burkov, B. Chaib-draa, An approximate subgame-perfect equilibrium computation technique for repeated games, *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence* (2010) 729–736.
- [15] K. Chatterjee, H. Sabourian, Game theory and strategic complexity, in: R.A. Meyers (Ed.), *Encyclopedia of Complexity and System Science*, Springer, 2009.
- [16] X. Chen, X. Deng, S.-H. Teng, Settling the complexity of computing two-player Nash equilibria, *Journal of the ACM* 56 (3) (2009).

- [17] M.B. Cronshaw, Algorithms for finding repeated game equilibria, *Computational Economics* 10 (1997) 139–168.
- [18] 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.
- [19] D. Cvetković, P. Rowlinson, Simić, *Eigenspaces of graphs*, Cambridge University Press, 1997.
- [20] C. Daskalakis, On the complexity of approximating a Nash equilibrium, *ACM Transactions on Algorithms* 9 (3) (2013).
- [21] C. Daskalakis, P. Goldberg, C.H. Papadimitriou, The complexity of computing a Nash equilibrium, *SIAM Journal on Computing* 39 (1) (2009) 195–259.
- [22] E.D. Demaine, *Playing Games with Algorithms: Algorithmic Combinatorial Game Theory*, *Mathematical Foundations of Computer Science (MFCS 2001)*, Springer, 2001.
- [23] 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.
- [24] A. Fraenkel, Combinatorial games: Selected bibliography with a succinct gourmet introduction, *Electronic Journal of Combinatorics. Dynamic Survey* 2 (2012).
- [25] D. Fudenberg, E. Maskin, The folk theorem in repeated games with discounting and incomplete information, *Econometrica* 54 (1986) 533–554.
- [26] G. Gottlob, G. Greco, F. Scarcello, Pure Nash equilibria: hard and easy games, *Journal of Artificial Intelligence Research* 24 (2005) 357–406.
- [27] K. Judd, Ş. Yeltekin, J. Conklin, Computing supergame equilibria, *Econometrica* 71 (2003) 1239–1254.
- [28] 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.
- [29] M. Kitti, Conditional Markov equilibria in discounted dynamic games, *Mathematical Methods of Operations Research* 78 (2013) 77–100.
- [30] M. Kitti, Subgame perfect equilibria in discounted stochastic games, *Journal of Mathematical Analysis and Applications* 435 (2016) 253–266.
- [31] 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.

- [32] R.D. Mauldin, S.C. Williams, Hausdorff dimension in graph directed constructions, *Transactions of the American Mathematical Society* 309 (1988) 811–829.
- [33] J.H. Nachbar, W.R. Zame, Non-computable strategies and discounted repeated games, *Economic Theory* 8 (1996) 103–122.
- [34] H. Salonen, H. Vartiainen, Valuating payoff streams under unequal discount factors, *Economics Letters* 99 (2008) 595–598.
- [35] S. Sorin, On repeated games with complete information, *Mathematics of Operations Research* 11 (1986) 147–160.
- [36] 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.

The **Aboa Centre for Economics (ACE)** is a joint initiative of the economics departments of the Turku School of Economics at the University of Turku and the School of Business and Economics at Åbo Akademi University. ACE was founded in 1998. The aim of the Centre is to coordinate research and education related to economics.

Contact information: Aboa Centre for Economics,
Department of Economics, Rehtorinpellonkatu 3,
FI-20500 Turku, Finland.

www.ace-economics.fi

ISSN 1796-3133