The Repeated Prisoner's Dilemma in a Network

Markus Kinateder*
Universitat Autònoma de Barcelona†

16 June 2008

Abstract

An infinitely repeated discounted Prisoner's Dilemma played on a connected, undirected and fixed network is studied. Players observe their immediate neighbors' behavior only, but communicate over time the repeated game's history throughout the network. The delay in receiving this information requires the players to be more patient to sustain the same level of cooperation as in a complete network. In special cases, the network need not be connected. The reduction in the set of payoff vectors supportable by sequential equilibria, which is due to the network, is illustrated analytically and graphically for three players in a star versus a complete network. The set of sequential equilibria under strategic communication intersects with that under exogenously imposed truthtelling, though each is not a subset of the other for some range of discount factors.

JEL classification numbers: C72, C73, D85

Keywords: Repeated Game, Prisoner's Dilemma, Imperfect Monitoring, Network, Strategic Communication, Payoff Space, Centrality

1 Introduction

The Prisoner's Dilemma is a well studied game, not only in Economics, since it captures many features from reality. The players' selfish behavior leads them to play the Nash

---

[*] I am very grateful for the support received from my supervisor Jordi Massó. I benefited hugely from a conversation with Coralio Ballester. Financial support from the Spanish Ministry of Education and Science through grant SEJ2005-01481 is acknowledged.

[†] Departamento de Economía y de Historia Económica, Edificio B, Universitat Autònoma de Barcelona, 08193 Bellaterra (Barcelona), Spain; email: mkinateder@idea.uab.es

Equilibrium in strictly dominant actions. However, this is not efficient and everybody is better off if the players cooperate. Cooperation is achieved, and thus efficiency attained, by repeating the Prisoner's Dilemma forever, provided the players are patient enough. Sometimes, efficiency arises even in the finitely repeated Prisoner's Dilemma.[1]

This paper analyzes an infinitely repeated discounted Prisoner's Dilemma played on a network. Kinateder (2008) defines repeated network games based on any stage game.[2] All players in a connected and undirected network that is fixed throughout the repeated game participate in the same stage game at each point in time. A player only observes his neighbors' behavior. However, by communicating with them, he receives the entire history of the repeated game with a finite delay. A Folk Theorem obtains provided the players truthfully communicate their observations to their neighbors.

For the Prisoner's Dilemma more results obtain than in other games due to its special structure. For example, the reduction in the set of payoff vectors which is supportable by sequential equilibria as a function of the discount factor is illustrated for a three player star versus a complete network. Some sequential equilibria do not extend from exogenously imposed truthtelling to strategic communication while new ones emerge due to richer communication. The players may lie, even as part of the equilibrium, and imperfect private monitoring arises endogenously in this model. In the literature, this is frequently modelled by letting each player (in a repeated Prisoner's Dilemma) receive a distinct imperfect signal of each action profile played. Usually, a Folk Theorem obtains.[3]

In case the players follow the trigger strategy profile, which is formally defined in section 3, it is possible to relax the assumption of the network's connectedness and cooperation obtains if each group contains at least two players and they are patient enough. This resembles results obtained when pairs of players which play the repeated Prisoner's Dilemma are repeatedly and randomly rematched.[4]

The repeated network Prisoner's Dilemma is defined in the next section. It is not parameterized since this only distracts from more interesting issues, though the results can be easily generalized in this way. In section 3, the basic difference between the complete and a star network each formed by three players which follow the trigger strategy is illustrated.[5] This result is extended to any connected network, and conditions are given

---

[1] See, for example, Cooper et al. (1996) for experimental evidence, or Kreps et al. (1982) if every player assigns a small probability to his opponents being irrational.

[2] The reader is referred to this paper for additional references on repeated games not mentioned here.

[3] Examples for this are Bhaskar and Obara (2002), Ely and Välimäki (2002), Piccione (2002) and Sekiguchi (1997). Kandori (2002) surveys the imperfect private monitoring literature. Fudenberg, Levine and Maskin (1994) obtain a Folk Theorem under imperfect public monitoring given any stage game.

[4] Such a setup is studied, for example, by Kandori (1992) and by Ellison (1994).

[5] This part of the paper coincides with the Prisoner's Dilemma example given in Kinateder (2008).

under which it even holds for unconnected networks. The reduction in the payoff space supportable by sequential equilibria as a function of the discount factor which is due to the network is derived in section 4. Before concluding, results under strategic communication are given and a way to identify the most central or best informed player is proposed.

2 Preliminaries

2.1 Prisoner's Dilemma Stage Game and Network

Each player i in the finite set of players $I = \{1, ..., n\}$, where $n > 2$, has a set of pure actions $A_i = \{C, D\}$; C stands for *cooperate* and D for *defect*. The pure action space of the stage game is $A = \times_{i \in I} A_i$ with generic element a called pure action profile. To emphasize player i's role, a is written as (a_i, a_{-i}). The game can be easily extended to mixed actions as shown in Kinateder (2008). For any non-empty set of players $S \subset I$, let $A_S = \times_{i \in S} A_i$ and denote by a_S an element of this set. Player i's payoff function is a mapping $h_i : A \rightarrow \mathbb{R}$, and the payoff function $h : A \rightarrow \mathbb{R}^n$ assigns a payoff vector to each pure action profile. Given $a \in A$, player i's payoff function is

$$h_i(a) = \begin{cases} 3 & \text{if } a_j = C \text{ for all } j \in I \\ 0 & \text{if } a_i = C \text{ and } \exists\ j \in I \setminus \{i\} \text{ s.t. } a_j = D \\ 4 & \text{if } a_i = D \text{ and } a_j = C \text{ for all } j \in I \setminus \{i\} \\ 2 & \text{if } a_i = D,\ \exists\ j \in I \setminus \{i\} \text{ s.t. } a_j = D \text{ and } \exists\ l \in I \setminus \{i, j\} \text{ s.t. } a_l = C \\ 1 & \text{if } a_j = D \text{ for all } j \in I. \end{cases}$$

A player's payoff is 3 when all players choose C. It is 4 if he unilaterally chooses D. It is 2 if some other player chooses D as well while at least one player chooses C. It is 0 if he chooses C while at least one other players chooses D and it is 1 if all players choose D.

The Prisoner's Dilemma stage game in normal form is the tuple $\hat{G} \equiv (I, (A_i)_{i \in I}, (h_i)_{i \in I})$. Let the convex hull of the finite set of payoff vectors corresponding to pure action profiles in $\hat{G}$ be $co(\hat{G}) = co\{x \in \mathbb{R}^n \mid \exists\ a \in A : h(a) = x\}$.

The players in set I are the vertices of a network g, whose graph is defined as the pair (I, E), where $E \subseteq I \times I$ denotes the set of links between them. A link from player i to player j is denoted by (i, j). Graph (I, E) is undirected, that is, for all $i, j \in I$, (i, j) if, and only if, (j, i). Given network g, a path between two distinct players i and j is defined as a sequence of distinct players $i_1, ..., i_r$ with $i_1 = i$, $i_r = j$, and $(i_{l-1}, i_l) \in E$, for all $1 < l \leq r$. Its length is $r - 1$. Network g is assumed to be connected. Hence, each player is connected to at least one other player directly and to all others via paths of finite lengths. The length of the shortest path between two distinct players i and j is called *distance* between i and j and is denoted by d_{ij}. The *largest distance* between player i and any other player in g is defined by $d_i = \max_{j \in I} d_{ij}$ and the *diameter* of g is the maximal *largest distance* among all players, that is, $d = \max_{i \in I} d_i$. Finally, denote player i's set of direct neighbors by $i(1) = \{j \in I \mid d_{ij} = 1\}$, and for any $2 \leq m \leq d_i$, define his set of *m-neighbors* as $i(m) = \{j \in I \mid d_{ij} = m\}$.

2.2 Communication and Observations

When the Prisoner's Dilemma is played repeatedly, in each period, a player first chooses an action, in a way specified below, and then makes observations and communicates with his neighbors. He observes the actions chosen by his immediate neighbors, before they communicate him strategically the information they received one period earlier from their neighbors. Similarly, he reveals to any neighbor the action he plays, before communicating him strategically the information he received one period ago. Hence, information flows one link per period and with a $d_i - 1$ period lag, player i gets to know a "filtered" version (by strategic communication) of what every player in the network did in the first period.[6]

Since the players have perfect recall, no player $i \in I$ forgets new information he receives. This information is an element of his *set of observations*, denoted by Ob_i^t for any $t \geq 1$. It includes all possible observations that player i may make in period t; in particular, also those which are manipulated by his neighbors' lies. Denote an observation made by player i at t by $ob_i^t \in Ob_i^t$, and all players' observation profile at t by $ob^t \in Ob^t$, where $Ob^t = \times_{i \in I} Ob_i^t$. For every player i at any $t > 1$, there is a set of reports R_i^t from which he chooses a report r_i^t which he sends to all his neighbors in $i(1)$. By assumption, a player can only send one report. He reports the new information he received at $t-1$ in a strategic way, that is, possibly including a lie. A lie will be defined formally below.

Observations and reports are then defined as follows. At $t = 1$, $Ob_i^t = A_i \times A_{i(1)}$, that is, player i observes what he and all his neighbors do, while obviously $R_i^1 = \emptyset$, and hence, $r_i^1 = \emptyset$ since player i has nothing to report from a previous period. At the end of $t = 2$ player i reports to any neighbor in $i(1)$ what he observed that all his neighbors did at $t = 1$, that is, $R_i^2 = A_{i(1)}$. This can also be written as $R_i^2 = Ob_i^1 \setminus A_i$ since a player never reports what he did. Two distinct neighbors of i may not be able to observe each other. They get to know what the other did at $t = 1$ by i's report at $t = 2$. Although there may

---

[6] At the end of any $t \geq d_i$, under truthtelling, player i knows the actions played in period t by himself and all players in $i(1)$, the actions played by himself and all players in $i(1)$ and $i(2)$ at $t-1$, ..., and finally the actions played by all players at $t - d_i + 1$ and at any point in time before.

be other ways via which this pair of players obtains information about each other after a larger delay, as will become clear in a moment. At $t = 2$, player i observes what his neighbors do as well as what they report him, and thus, $Ob_i^2 = A_i \times A_{i(1)} \times_{j \in i(1)} R_j^2$, or equivalently, $Ob_i^2 = Ob_i^1 \times_{j \in i(1)} R_j^2$. In this way, a recursive structure of the observations and reports obtains, and at any $t > 1$, $R_i^t = Ob_i^{t-1} \setminus A_i$ and at any $t > 2$, $Ob_i^t = Ob_i^{t-1}$.[7]

At $t = d_i$, player i for the first time gets to know a filtered version of what the most distant players from him did at $t = 1$. It is filtered by all the players which are on the shortest paths between him and the players at distance d_i from him.

At every t, a player observes what all players that are between 1 and $t - 1$ links away from him did at some point in the past. The links are not only counted along shortest paths but also in any other way, in particular, one link might me used several times on a paths (for example, a piece of information flows forth and back between two neighbors). Hence, the information each player observes grows over time since a player receives one report from any neighbor and hands it over to all neighbors (including the one from which he received it) in the subsequent period.

Given ob_i^{t-1}, player $i \in I$ *lies* at t if his report r_i^t differs from ob_i^{t-1} as follows: he changes the action $a_j^s \in ob_i^{t-1}$, abusing notation, chosen by some player j at some $1 \leq s < t$ to any other action $b_j \in A_j \setminus \{a_j\}$. This yields a recursive dynamic process. The players organized in this way play an infinitely repeated discounted Prisoner's Dilemma.

2.3 Repeated Prisoner's Dilemma Played on a Network

In the infinitely repeated discounted Prisoner's Dilemma played on the fixed network g, thereafter called repeated network (Prisoner's Dilemma) game, at each point in discrete time, $t = 1, 2, ...$, the Prisoner's Dilemma stage game $\hat{G}$ is played.

At any t, a player's strategy prescribes him which action to choose and his communication strategy which report to send to his neighbors. Under truthtelling the players hand over the true history of the repeated game and the communication part is trivial.

Let player i's set of strategies be $F_i = \{\{f_i^t\}_{t=1}^{\infty} \mid f_i^1 \in A_i$, and for all $t > 1$, $f_i^t : \cup_{s=1}^{t-1} Ob_i^s \rightarrow A_i\}$. At any $t \geq 1$, player i's strategy $f_i = \{f_i^t\}_{t=1}^{\infty}$ prescribes him to choose an action. For $t > 1$, this is a mapping from the union of his sets of observations to his action set. The cartesian product of all players' strategy sets $F = \times_{i \in I} F_i$ constitutes the

---

[7] At $t = 3$, player i reports to every $j \in i(1)$ what he observed them doing at $t = 2$ and what they told him at $t = 2$, which includes what all players 2 links away from i did at $t = 1$. Since player i is two links away from himself, at $t = 3$, he tells every neighbor what they told him that he did at $t = 1$. In this way, the same information is reported again and again, though it may be altered by a player's lie.

strategy space of the repeated network game. A strategy profile $f = (f_1, ..., f_n)$ is an element of F. To emphasize player i's role, it is written as (f_i, f_{-i}).

Let player i's set of communication strategies be $Com_i = \{\{com_i^t\}_{t=1}^{\infty} \mid com_i^1 = \emptyset$, and for all $t > 1$, $com_i^t : \cup_{s=1}^{t-1} Ob_i^s \to R_i^t\}$. Player i's communication strategy $com_i = \{com_i^t\}_{t=1}^{\infty}$ prescribes him to send a report at any $t > 1$. This is a mapping from the union of his sets of observations to his report set. The cartesian product of all players' communication sets $Com = \times_{i \in I} Com_i$, constitutes the communication space of the repeated network game. A communication profile $com = (com_1, ..., com_n)$ is an element of Com. To emphasize player i's role, it is written as (com_i, com_{-i}).

If truthtelling is assumed, player i only selects a strategy while his communication is automatically determined: at any $t > 1$, $com_i^t : (\cup_{s=1}^{t-2} Ob_i^s) \cup Ob_i^{t-1} \to R_i^t$ is such that the mapping from Ob_i^{t-1} to R_i^t is the identity. The same holds if truthtelling arises endogenously under strategic communication. Then, at any $t \geq 1$, each pair (f, com) with $f \in F$ and $com \in Com$ recursively generates a pure action profile $a^t(f, com) = (a_1^t(f, com), ..., a_n^t(f, com))$, a report profile $r^t(f, com) = (r_1^t(f, com), ..., r_n^t(f, com))$ and a corresponding observation profile $ob^t(f, com) = (ob_1^t(f, com), ..., ob_n^t(f, com))$. These in turn determine the action profile and the report profile at $t + 1$. Finally, let $F \times Com$ be the strategy cum communication space of the repeated network game.

Given a common discount factor $\delta \in [0, 1)$, the function $H^{\delta} : F \times Com \to \mathbb{R}^n$ assigns a payoff vector to each strategy profile of the repeated network game. Given $f \in F$ and $com \in Com$, player i's payoff, $H_i^{\delta}(f, com) = (1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} h_i(a^t(f, com))$, is the $(1 - \delta)$-normalized discounted sum of stage game payoffs. Given δ and g, the repeated network Prisoner's Dilemma is defined as the normal form game $\hat{G}^{g,\delta} \equiv (I, (F_i)_{i \in I}, (Com_i)_{i \in I}, (H_i^{\delta})_{i \in I})$.

When g is complete, $i(1) = I \setminus \{i\}$ for all $i \in I$ and $\hat{G}^{g,\delta}$ is identical to the infinitely repeated discounted Prisoner's Dilemma, referred to as $\hat{G}^{\delta}$. In this case, f_i simplifies: at any $t > 1$ it maps $A^{t-1} = (\times_{i \in I} A_i)^{t-1}$ to A_i, that is, each player conditions his action choice on the history of action profiles chosen by all players between periods 1 and $t - 1$ since he observes this. Communication is then redundant.

Moreover, the players have common knowledge of the game played, the form of the network[8] and the strategy choices available to all players. Finally and importantly, each player i observes his payoff only at the end of the game.[9] If the discount factor is inter-

---

[8] For most of the results obtained, common knowledge of the network is not required.

[9] If a player observes his or all players' payoffs immediately, this is a(n imperfect) private or a public signal, respectively. Both kinds of signal are already extensively studied in repeated games (see footnote 3). The results' extension to this is obvious and the information the payoff may reveal is ignored.

preted as the probability with which the game is played again in the next period, then the probability that the repeated network game ends at T converges to 1 as T goes to infinity.

2.4 Individual Rationality, Feasibility and Sequential Equilibrium

A player's individually rational payoff is the lowest to which he can be forced in a stage game. It obtains when he maximizes his payoff while all other players minimize it and is called *minmax* payoff. For any player $i \in I$, let his minmax payoff in pure actions be

$$\bar{\nu}_i \equiv \min_{a_{-i} \in A_{-i}} \max_{a_i \in A_i} h_i(a_i, a_{-i}). \tag{1}$$

In the Prisoner's Dilemma every player's individually rational payoff is his minmax payoff of 1. It obtains when all players choose D. Moreover, this is the unique stage game Nash Equilibrium in pure and strictly dominant actions.

The *set of feasible payoff vectors* of the repeated network game is defined as[10]

$$\mathcal{F} = \{x \in \mathbb{R}^n \mid \exists \, \{a^t\}_{t=1}^{\infty} : \, \forall \, t \geq 1, \; a^t \in A, \text{ and } \forall \, i \in I, \; x_i = (1-\delta) \sum_{t=1}^{\infty} \delta^{t-1} h_i(a^t)\}.$$

Any feasible payoff vector can be generated by a sequence of pure action profiles.

The *set of feasible and individually rational payoff vectors* is denoted by $\mathcal{F}^*$. It contains all feasible payoff vectors that are larger than or equal to $\bar{\nu} = (1, ..., 1)$, the minmax payoff vector, and is defined as

$$\mathcal{F}^* = \mathcal{F} \cap \{x \in \mathbb{R}^n \mid x \geq \bar{\nu}\}.$$

Any payoff vector in this set is a candidate to be supported by a sequential equilibrium.

A sequential equilibrium requires a strategy profile and a system of beliefs to be sequentially rational and consistent, respectively. In the repeated network game, the attention is restricted to a class of strategy profiles in which each player conditions his action choices only on his observations. In this class, each sequential equilibrium strategy profile is sequentially rational for any belief a player may have about the yet unobserved actions chosen by all other players in the most recent periods. Hence, beliefs are not

---

[10] Any payoff vector in $co(\hat{G})$ is feasible for $\delta \in (1-\frac{1}{z}, 1)$, where z is the number of vertices of $co(\hat{G})$. For any discount factor in this range, sets $\mathcal{F}$ and $co(\hat{G})$ coincide; see Fudenberg, Levine and Maskin (1994).

modelled explicitly and a sequential equilibrium is said to exist when the condition of sequential rationality is fulfilled.[11]

Definition 1. *Strategy profiles $\dot{f} \in F$ and $\dot{com} \in Com$ are a* sequential equilibrium (SE) *of $\hat{G}^{g,\delta}$, if for all $t \geq 1$ and given any $ob^t \in Ob^t$, $\{\dot{f}^{\tau}(ob^{\tau-1})\}_{\tau=t+1}^{\infty}$ and $\{\dot{com}^{\tau}(ob^{\tau-1})\}_{\tau=t+1}^{\infty}$ are such that for all $i \in I$, all $f_i \in F_i$ and all $com_i \in Com_i$,*

$$(1-\delta)\sum_{s=t+1}^{\infty}\delta^{s-1}h_i(a^s(\dot{f},\dot{com}) \quad \geq \quad (1-\delta)\sum_{s=t+1}^{\infty}\delta^{s-1}h_i(a^s(f_i,\dot{f}_{-i},com_i,\dot{com}_{-i})).$$

When g is complete this definition includes $\hat{G}^{\delta}$ and the concepts of sequential and subgame-perfect equilibrium coincide. However, equilibria of $\hat{G}^{g,\delta}$ and $\hat{G}^{\delta}$ are called sequential when Definition 1 holds, and the corresponding sets of SE strategy profiles are denoted by $SE(\hat{G}^{g,\delta})$ and $SE(\hat{G}^{\delta})$, respectively. A strategy profile is a SE if, and only if, no player's finite unilateral deviation at any point in time is profitable.[12] If truthtelling is imposed exogenously, the sets $SE^{ET}(\hat{G}^{g,\delta})$ and $SE^{ET}(\hat{G}^{\delta})$ are adored with superscript ET for *exogenous truthtelling.* Otherwise, the players use strategic communication and may or not be asked to lie. A deviation and a subsequent lie may then "cancel" each other. However, the definition of SE only says that given the observation a player has made, he does not deviate from the strategy and communication profile. This is weaker than to require a player to condition his strategy and communication on the true history of the repeated game or to let him form beliefs about it.

3 The Network makes a difference

The following example of the Prisoner's Dilemma with three players and exogenously imposed truthtelling, denoted by $\hat{com} \in Com$, illustrates how imposing a network on a set of players affects the set of SE. For this sake, consider the trigger strategy profile. It prescribes each player to cooperate as long as all players cooperate and to defect forever if any player defected. Given any network g, the trigger strategy of player i, denoted by $\hat{f}_i \in F_i$, is defined as follows: $\hat{f}_i^1 = C$, and for $t \geq 1$, given $ob_i^t \in Ob_i^t$,

$$\hat{f}_i^{t+1}(ob_i^t) = \begin{cases} D & \text{if } \exists\ 1 \leq \tau \leq t \text{ such that for } a^{\tau} \in ob_i^t,\ a_j^{\tau} = D,\ \text{while } a_{-j}^{\tau} = C \\ C & \text{otherwise.} \end{cases}$$

---

[11] A player "believes" a lie, unless this causes an inconsistency in his set of observations. Then, he may detect the liar. If the lie is part of a SE, he believes that it occurred, but not its content. A player may update his belief various times. This is no problem since a player puts a positive belief on any history.

[12] Since $\delta < 1$, a player's gain from a deviation of infinite length can be approximated by that of a finite deviation. Therefore, unilateral deviations of finite length from a strategy profile are not profitable if, and only if, it is a SE of the repeated network game; see Mailath and Samuelson (2006).

1——2——3

Figure 1: Three players form a Star

3

C

1-2	C	D
C	3, 3, 3	0, 4, 0
D	4, 0, 0	2, 2, 0

D

1-2	C	D
C	0, 0, 4	0, 2, 2
D	2, 0, 2	1, 1, 1

Figure 2: Prisoner's Dilemma for three players

Given $\hat{f} \in F$ and $c\hat{o}m \in Com$, for all $i \in I$ and all $t \geq 1$, first $a_i^t(\hat{f}, c\hat{o}m) = C$, and second, for all $a_j^\tau \in ob_i^t(\hat{f}, c\hat{o}m)$, $a_j^\tau = C$ as well for all $1 \leq \tau \leq t$ and all $j \in I$. Hence, for all $i \in I$, $H_i^\delta(\hat{f}, c\hat{o}m) = (1-\delta)\sum_{t=1}^{\infty} \delta^{t-1} h_i(a^t(\hat{f}, c\hat{o}m)) = (1-\delta)\sum_{t=1}^{\infty} \delta^{t-1} 3 = 3$.

3.1 The Players form a Star

Consider a *star* (or a line) with $n = 3$, as represented in Figure 1. The graph of g is $E = ((1,2),(2,1),(2,3),(3,2))$. Figure 2 represents $\hat{G}$ for $n = 3$, where player 1 chooses rows, player 2 columns and player 3 matrices. Suppose that the players tell the truth. The trigger strategy profile is a SE of $\hat{G}^{g,\delta}$ if, and only if, all players are patient enough, that is, δ is higher than some threshold value. Then, none of them ever deviates. Corresponding conditions on δ must be found for the truncation of the repeated network Prisoner's Dilemma at any point in time, and therefore, given any observation profile. Three classes of unilateral deviations can be identified.[13] Any deviation that may arise in the course of play can be uniquely allocated to one class. The three classes are

1) initial unilateral deviations,
2) subsequent unilateral deviations (before the initial is known by all players), and
3) unilateral deviations when the punishment takes place.

Obviously, unilateral deviations during the punishment are not profitable since all players play D. The resulting action profile is the stage game Nash Equilibrium in strictly dominant actions. Hence, every player best-replies independently of g and of δ. For the same reason, no player can deviate profitably from the trigger strategy profile in class 2. After a player's initial deviation, he and any player who knows about it are best-off to

---

[13] A SE does not impose restrictions on play after a multilateral deviation by two or more players.

play D forever (rather than to deviate and to choose C at any point in time).

It remains to show that no player has a profitable unilateral deviation from the trigger strategy profile when all players choose C. Given δ, player 2 (who is directly observed by 1 and 3) does not deviate in any period τ if, and only if,

$$\begin{aligned}
(1-\delta)\sum_{t=1}^{\infty} 3\delta^{t-1} &\geq (1-\delta)\sum_{t=1}^{\tau-1} 3\delta^{t-1} + 4(1-\delta)\delta^{\tau-1} + (1-\delta)\sum_{t=\tau+1}^{\infty} 1\delta^{t-1}, \\
(1-\delta)\sum_{t=\tau+1}^{\infty} 2\delta^{t-1} &\geq (1-\delta)\delta^{\tau-1}, \\
2\delta^{\tau+1} &\geq (1-\delta)\delta^{\tau}, \\
\delta &\geq \tfrac{1}{3}.
\end{aligned}$$

The value of $\frac{1}{3}$ is not only the threshold value for player 2 in this example but also the one for all players in a complete network. The network affects, however, the threshold value of the remaining two players in this example. Given δ, player 1 (and similarly 3) does not deviate from the trigger strategy profile in any period τ if, and only if,

$$\begin{aligned}
(1-\delta)\sum_{t=1}^{\infty} 3\delta^{t-1} &\geq (1-\delta)\sum_{t=1}^{\tau-1} 3\delta^{t-1} + 4(1-\delta)\delta^{\tau-1} + 2(1-\delta)\delta^{\tau} + (1-\delta)\sum_{t=\tau+2}^{\infty} 1\delta^{t-1}, \\
(1-\delta)\delta^{\tau} + (1-\delta)\sum_{t=\tau+2}^{\infty} 2\delta^{t-1} &\geq (1-\delta)\delta^{\tau-1},
\end{aligned}$$

which can be simplified to $2\delta+\delta^2-1 \geq 0$. The only positive solution for δ in this quadratic equation is approximately 0.414. Hence, in class 1 of the SE conditions the requirement on δ, or the players' patience, is higher in the star with three players considered here than in a complete network. This is due to the one period delay with which players 1 and 3 obtain information about each other's action choice.

This example extends to the case where $n > 3$ and the players form a star. The player at the center of the star has the same role as player 2 in this example, and for all other players the same conditions apply as for players 1 and 3 in this example.

3.2 The Repeated Prisoner's Dilemma Played in any Network

A similar result holds for any network (with $n > 3$), in which all players follow the trigger strategy. In Figure 1 it takes 2 periods until full punishment sets in, and players 1 and 3 do not deviate if $2\delta + \delta^2 - 1 \geq 0$. In any network, it takes d_i periods until full punishment sets in and no player $i \in I$ deviates if $2\delta + \delta^{d_i} - 1 \geq 0$. The time lag between the deviation and full punishment is 2 periods in the first case and d_i periods in general. In each period after i's deviation, the group of punishers increases strictly until d_i periods after it. Until then the deviator's payoff is 2 since at least one player still chooses C. From d_i periods after the deviation on he receives a payoff of 1 forever, which is expressed by δ^{d_i}.

Although this expression depends on d_i, even in very large networks the threshold value for δ is bounded above by $\frac{1}{2}$. This can be seen by taking the limit of this inequality when d_i converges to infinity. Since $\delta < 1$, the term δ^{d_i} converges to 0 and the inequality simplifies to $2\delta - 1 \geq 0$ or $\delta \geq \frac{1}{2}$. Hence, for "moderately patient" players, the trigger strategy profile is a SE in any repeated network Prisoner's Dilemma.

3.3 The Network is Unconnected

Network g is unconnected, if there are different connected and undirected components. Suppose that each component contains at least two players. Together all connected components constitute network g. All players in one component can observe each other and communicate with each other. However, they never observe any player in any other component nor do they receive information about what the player did via communication. (The distance between any pair of unconnected players is normalized to infinity.) Thus, a deviation by any player is only observed by his peer(s) in his component and not by any other player. Nevertheless, all players in the different connected components still participate in a single Prisoner's Dilemma at every $t \geq 1$.

Suppose that all players in the distinct connected components follow a modified trigger strategy profile. Any player's unilateral deviation is only punished by his companion(s), since no other player ever gets to know that it occurred though any other player's payoff whether he can observe it or not is affected by the deviation. Even then unilateral deviations are not profitable if the players are sufficiently patient. After any deviation, a player is only punished by the players in his component who observe it and choose D forever (possibly after some delay). This is a SE if, and only if,

$$(1-\delta)\sum_{t=1}^{\infty} 3\delta^{t-1} \geq (1-\delta)\sum_{t=1}^{\tau-1} 3\delta^{t-1} + 4(1-\delta)\delta^{\tau-1} + (1-\delta)\sum_{t=\tau+1}^{\infty} 2\delta^{t-1},$$

$$(1-\delta)\sum_{t=\tau+1}^{\infty} \delta^{t-1} \geq (1-\delta)\delta^{\tau-1},$$

$$\delta^{\tau+1} \geq (1-\delta)\delta^{\tau},$$

$$\delta \geq \tfrac{1}{2}.$$

All other players never observe the deviation since they only observe their payoff at the end of the game. This result shows that for moderately patient players cooperation is a SE even if a player never observes what some other participants of the game, who influence his payoff, have done. It also holds under strategic communication as follows from Theorem 1 in section 5 and when the number of players becomes arbitrarily large.

Similar results obtain in random matching models in which each pair of players plays a repeated Prisoner's Dilemma until the players are randomly rematched. Kandori (1992) obtains a Folk Theorem when the number of players is limited. Ellison (1994) extends this result to arbitrarily large populations of players. Haag and Lagunoff (2006) let each neighborhood play a repeated Prisoner's Dilemma. They analyze optimal network design when each player's discount factor is randomly drawn before the repeated game begins.[14]

4 Payoff Space as a Function of the Discount Factor

To derive the set of payoff vectors, that are supportable by SE, as a function of the discount factor in an infinitely repeated discount game with perfect monitoring is involved. A parameterized two players Prisoner's Dilemma was fully solved by Stahl (1991). Cornshaw (1997) was the first to completely solve a three player game, namely a Cournot game; he also fully solved a two country tariff war. Both papers use the dynamic programming approach introduced by Abreu, Pearce and Stacchetti (1990). A version of this result already suited to the Prisoner's Dilemma can be found in Mailath and Samuelson (2006), section 2.5. For the theoretical foundation of the construction which is applied subsequently, the reader is referred to either reference.

Each player's payoff in an infinitely repeated discounted game is decomposable into the one the player receives today and the one he receives from tomorrow on. If the continuation promise tomorrow is large enough, the player does not deviate. The payoff from tomorrow on can be decomposed in the same way. If every period's payoff is incentive compatible for all players, and none can deviate profitably, the sequence of action profiles

---

[14] Lippert and Spagnolo (2005) or Vega-Redondo, Marsili and Slanina (2005) let each linked pair of players play a bilateral repeated Prisoner's Dilemma. A link is severed by a player's deviation in the first paper, and by a player whose stochastically decaying payoff falls below some threshold in the second.

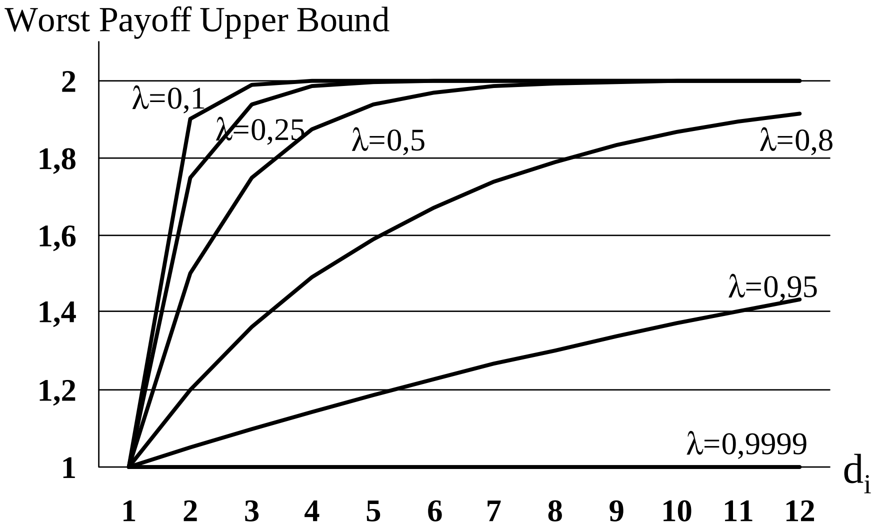

Figure 3: Upper bound of a player's worst payoff in the Prisoner's Dilemma

generates a payoff vector that is supportable by a SE and the decomposition of the payoff yields a payoff today and one tomorrow in the same set, that is, the set of payoff vectors is decomposable on itself. In this case, this set of payoff vectors is said to be self-generating. The largest self-generating set of payoff vectors is that supportable by SE.

4.1 Continuation Payoffs

Given g and δ, a player's continuation payoff if he deviates is called his *worst payoff*. It is determined by the largest distance between him and any other player in the network. This is the time it takes until all players punish him. It also depends on the sequence of action profiles played by the other players until they are informed about his deviation. An upper and a lower bound to a player's worst payoff can be calculated. For any $f \in F$, any player i's worst payoff in the repeated network Prisoner's Dilemma lies between the two identified bounds. The lower bound is identical to player i's minmax payoff $\bar{\nu}_i = 1$. It obtains when all players play D forever after his deviation (and is independent of g and of δ.) The upper bound depends on a player's position in g and on δ. It obtains, for example, when the players follow the trigger strategy profile. In this case, a deviator gains most since all players choose C until they become aware of his deviation. After deviating unilaterally from the trigger strategy, player i receives

$$
\begin{aligned}
(1-\delta)[2+2\delta+\dots+2\delta^{d_i-2}+1\delta^{d_i-1}+\dots] &= \\
(1-\delta)[\sum_{t=1}^{d_i-1} 2\delta^{t-1}+\sum_{t=d_i}^{\infty} 1\delta^{t-1}] &= \\
2-\delta^{d_i-1}.
\end{aligned}
$$

This upper bound of a player's worst payoff is strictly larger than 1, unless the network is complete, that is, $d_i = 1$ for all players. For different values of δ and depending on a player's position in the network it lies between 1 and 2, as depicted in Figure 3. For small values of δ, it is close to 2 even when player i's largest distance is small. Conversely, for δ close to 1, the upper bound of a player's worst payoff is close to 1 even in large networks.

4.2 Payoff Vectors supportably by Sequential Equilibria

For each action profile in the stage game, all (not necessarily credible) continuation payoffs which enforce it, such that no player deviates, are determined. A deviator is threatened with his worst payoff, that is, the players choose D forever after observing his deviation. To illustrate this, consider three players in a complete network. Player 1 cannot deviate profitably from action profile (C, C, C), if

$$(1-\delta)3+\delta x_{CCC} \quad \geq \quad (1-\delta)4+\delta x_{DCC},$$

where $x_{a_1a_2a_3}$, with subscript i suppressed, is a player's continuation payoff when each player $i = 1, 2, 3$ chooses $a_i \in A_i$. In case player 1 deviates, his continuation payoff is $x_{DCC} = 1$ (both other players observe his deviation and punish him). If he does not deviate, the continuation payoff he is promised is x_{CCC}. This inequality simplifies to $x_{CCC} \geq \frac{1}{\delta}$ and the set of continuation payoffs which enforce (C, C, C) is

$$A_{CCC} = \{(x_{CCC}, x_{CCC}, x_{CCC}) \in \mathbb{R}^3 \mid x_{CCC} \geq \tfrac{1}{\delta} \text{ for all } i \in I\}.$$

The continuation payoffs which enforce an action profile are thus determined by a simple incentive constraint which rules out any player's profitable deviation.

For each action profile, the resulting payoff is then decomposed into the one generated by the action profile today and the one from tomorrow on which only needs to be feasible and individually rational. Together with the condition on the continuation payoffs which enforce the action profile, this yields the set of decomposable payoff vectors for a given action profile. It depends on δ and gives a geometric representation of payoff vectors that

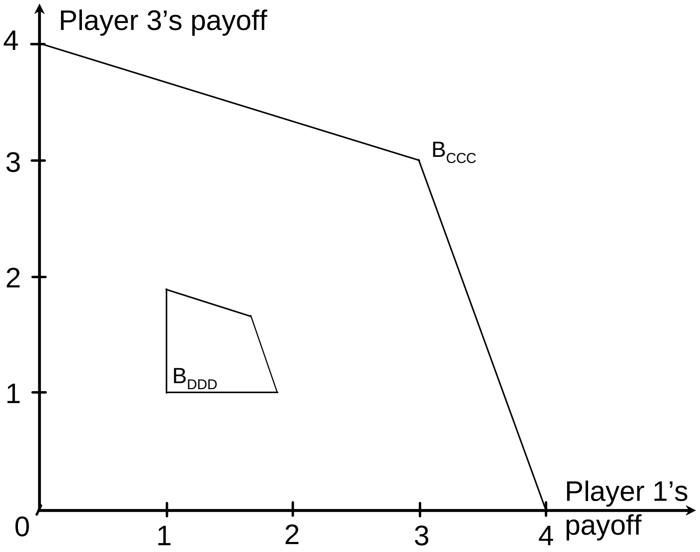

Figure 4: Set of Payoff Vectors generated by Sequential Equilibria for $\delta = \frac{1}{3}$

can be supported by SE for this action profile. In the example of the complete network formed by three players, the set of payoffs decomposable for action profile (C, C, C) is

$$B_{CCC} = \{x \in \mathcal{F}^* \mid x = (1-\delta)(3,3,3) + \delta y, \text{ for any } y \in \mathcal{F}^* \text{ and } x \in A_{CCC}\}.$$

Provided that (C, C, C) is played today any sequence of action profiles which yields an individually rational payoff can be played from tomorrow on. Then, no player has an incentive to deviate if, and only if, the payoff he receives is larger than or equal to $\frac{1}{\delta}$. For example, if $\delta = \frac{1}{4}$, then $x_{CCC} = 4$ and each player can deviate profitably from (C, C, C). For $\delta = \frac{1}{3}$, each player is just indifferent to deviate or not and chooses C.

In a similar way, the sets of continuation payoffs enforcing all other action profiles and the sets of payoffs decomposable for each are derived. The union of sets of payoffs decomposable for all action profiles in $\hat{G}$ yields the geometric representation of the set of payoff vectors that can be supported by SE for a given discount factor. In any set $B_{a_1a_2a_3}$, the payoff vector x which the players receive is decomposed into the payoff vector today, weighted by $(1-\delta)$, and the one from tomorrow on which is an element of $\mathcal{F}^*$ shrank by δ. If the action profile is sustainable, that is, $A_{a_1a_2a_3}$ is non-empty given δ, $(1-\delta)$ times the payoff vector today serves as origin to depict the shrank (or multiplied by δ) set $\mathcal{F}^*$.

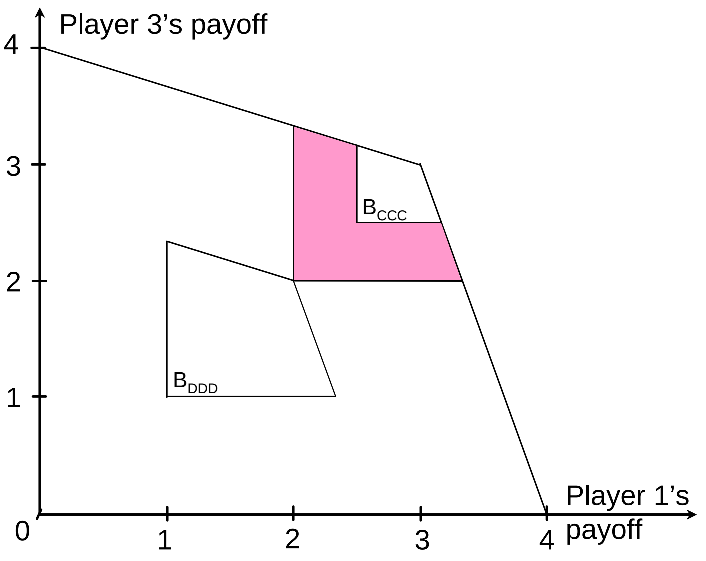

Figure 5: Set of Payoff Vectors generated by Sequential Equilibria for $\delta = \frac{1}{2}$

4.3 Reduction in the Equilibrium Payoff Space

Suppose now that the players form the network depicted in Figure 1. For player 2 nothing changes since he is immediately punished by both neighbors after deviating. However, a deviation of player 1 or 3 is more profitable since the other only observes it with a delay of one period. In this period, he is assumed to choose the same action as in the period in which the deviation took place.[15] From two periods after the deviation on he contributes to the punishment and chooses D forever.

The sets of decomposable payoffs for each action profile are derived in Appendix A. Given them, for each discount factor the set of payoff vectors generated by SE can be drawn. Player 2's payoff is not depicted since for him the outcome is the same as for any player in a complete network. The network in Figure 1 only matters for players 1 and 3. The payoff space contains the point $(0, 0)$ since it obtains when players 1 and 3 choose C while player 2 chooses D. In a two player Prisoner's Dilemma this point is not attainable. Finally, not all equilibrium payoff vectors which are identified are necessarily feasible (see footnote 10). This is not taken into account in Figures 4, 5 and 6.

In Figure 4, the set of equilibrium payoff vectors is depicted for $\delta = \frac{1}{3}$. Only sets B_{DDD} and B_{CCC} are non-empty in the complete network. In the star network, the payoff $(3, 3, 3)$

---

[15] Obviously, a deviating player's worst payoff is largest if the player which observes the deviation with a delay of one period chooses C in the period after the deviation. However, whether this player chooses C or D yields almost the same result, and hence, it is assumed that he repeats the action.

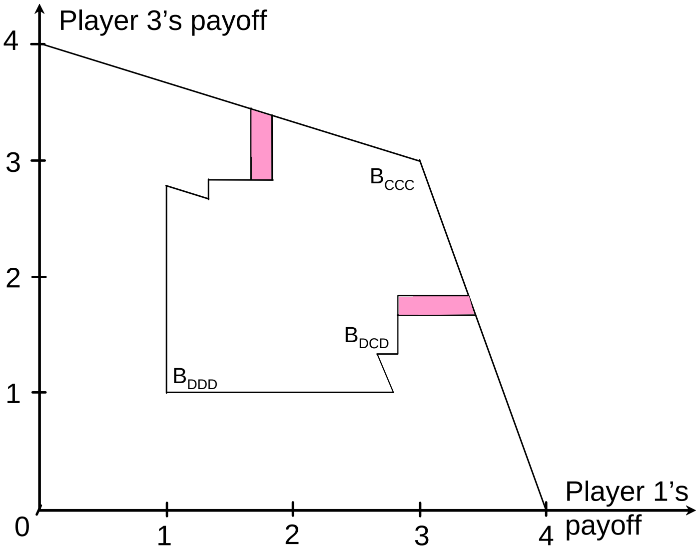

Figure 6: Set of Payoff Vectors generated by Sequential Equilibria for $\delta = \frac{2}{3}$

cannot be supported by a SE. This is the reduction in the set of equilibrium payoff vectors for $\delta = \frac{1}{3}$ due to the network. The set B_{DDD} in the star network is identical to that in the complete one. All other sets are empty for both networks.

In Figure 5, the set of equilibrium payoff vectors is depicted for $\delta = \frac{1}{2}$. Set B_{DDD} is again identical for both networks. All other sets, except of B_{CCC}, are empty. The shaded area obtains only for the complete network and the remaining one also for the star.

In Figure 6, the set of equilibrium payoff vectors is depicted for $\delta = \frac{2}{3}$. The sets B_{DDD} and B_{DCD} coincide for both networks and overlap. All other sets differ for the two networks, but this does not affect the set of equilibrium payoff vectors since all of them, except of B_{CCC}, are contained in B_{DDD} and B_{DCD}. Again set B_{CCC} is larger in the complete network and the shaded area cannot be supported by SE in the star network.

To extend this result to other games is very involved since the punishment structure in the Prisoner's Dilemma is quite simple. In most other games, a credible punishment cannot contain a conversion to the minmax profile forever and a compensation phase must follow once the deviator's entire gain is eliminated. This complicates the geometric representation even in the case of complete networks as illustrated by Stahl (1991) and Cornshaw (1997). However, for the Prisoner's Dilemma played by more than three players the sets of decomposable payoffs can be derived analytically given any network.

5 Strategic Communication

In this section, the Prisoner's Dilemma played on any network is extended to strategic communication. This is challenging, due to the bilateral communication structure.

At any t, a unilateral deviation is any player's non-compliance with the strategy, with the communication or with both while all other players choose the action prescribed by the strategy profile and send a report according to the communication profile. A unilateral deviation might be of finite but not of infinite length (see footnote 12). As assumed before, a player always observes correctly what his neighbors do. The extension of the trigger strategy profile to endogenous truthtelling under strategic communication is shown in Theorem 1. In this case, any deviation from the communication profile is a lie.

Theorem 1. *Let $\hat{G}$, g, $\hat{f} \in F$, $\hat{com} \in Com$ and $\delta \in [0,1)$ be given. Then, $(\hat{f}, \hat{com}) \in SE^{ET}(\hat{G}^{g,\delta})$ if, and only, if $(\hat{f}, \hat{com}) \in SE(\hat{G}^{g,\delta})$.*

Proof. Suppose that $(\hat{f}, \hat{com}) \in SE^{ET}(\hat{G}^{g,\delta})$. Then, unilaterally choosing a different action than prescribed by $\hat{f}$ after any history given that every player tells the truth is not profitable for any player. Hence, unilaterally choosing a different action than prescribed by $\hat{f}$ is neither profitable for any player given any observation even if it includes lies.

Next, lies and simultaneous deviations are show not to be profitable given any observation profile. Three cases might occur. First, no player has yet chosen D at any time. If a player claims that another chose D, punishment starts as if the player who lied has deviated himself. Since he does not deviate, he neither lies. He lies and deviates simultaneously since given his observations, the trigger strategy asks him to choose C.

Suppose next that at least one player chose D already. Any other player who observes this should choose D himself to punish the deviator. If he lies and claims that there was no deviation and simultaneously deviates by choosing C he is worse off since at least the deviator chooses D by $\hat{f}$. Finally, suppose that all players choose D. Then, all of them are indifferent to tell the truth or not since to claim that any player chose C instead of D even if this is true does not change the sequence of action profiles played. Again to simultaneously lie and to choose C is not profitable.

Finally, lies and deviations should not be profitable for any observation a player may have made. Again a player is indifferent to lie or not even if after some initial deviation some other player who knows about it and should play D chooses C instead. All players report this out of indifference and are better off to choose D instead of C. Hence, lies (without a simultaneous deviation) do neither occur after any history which already includes a sequence of deviations and/or lies.

Suppose that $(\hat{f}, \hat{com}) \in SE(\hat{G}^{g,\delta})$. Then, no player ever deviates from $\{a^t(\hat{f}, \hat{com})\}_{t=1}^{\infty}$, and obviously he would neither deviate if truthtelling was imposed exogenously. $\square$

If the players in any network follow the trigger strategy profile and $\delta \geq \frac{1}{2}$, as shown in section 3.2, no player ever deviates from the strategy or communication profile or from both. Thus, truthtelling arises endogenously (since in this particular case a deviation from the communication profile must be a lie). For many networks, $(\hat{f}, \hat{com})$ is a SE and truthtelling obtains under strategic communication for some values of $\delta \in (\frac{1}{3}, \frac{1}{2})$.

However, truthtelling does not arise endogenously in general. Two kinds of lies may occur. A player might claim that there was a deviation when there was none. This starts the punishment phase and affects him as if he had deviated. As long as he cannot deviate profitably by choosing a different action, he neither can lie profitably in this way. Another kind of lie is not to reveal a deviation and to deviate simultaneously, that is, not to punish it. Suppose that any deviation (whether profitable or not) is punished by the players. If the deviator chooses C instead of D, the player who should report this has an incentive to lie and to deviate if he is the only one who observes this. Then, the players continue to follow the initially prescribed sequence of action profiles. To report it would start punishment and the monitor's payoff would drop to 1 forever after some time (provided punishment involves playing D forever). By not reporting the deviation from D to C, his payoff is at least as large as the one he would receive if the deviation had never taken place and this, by individual rationality, is strictly larger than 1.

For any sequence of action profiles it depends on the strategy and communication profile which generates it whether it is a SE. A player who should report a deviation from C to D and punish it, is better off not to do this and to deviate himself, if he is the only monitor and this remains undetected. However, if the strategy and communication profile prescribes him this, then it is a SE. This behavior is also efficient since each player's payoff is above 1, the payoff each receives if punishment starts forever.

Proposition 1. *Let $\hat{G}$, g and $\delta \geq \frac{2}{3}$ be given such that one player in g has only one monitor. Then, there is $(f, \hat{com}) \in SE^{ET}(\hat{G}^{g,\delta})$ such that $(f, \hat{com}) \notin SE(\hat{G}^{g,\delta})$ and there is $(f, \tilde{com}) \in SE(\hat{G}^{g,\delta})$ such that $(f, \tilde{com}) \notin SE^{ET}(\hat{G}^{g,\delta})$.*

Proof. Let g be such that one player has only one monitor and let f be a modified trigger strategy: each player chooses C and only the monitored player is prescribed to choose D every 1000 periods. All players punish any unilateral deviation from f by choosing D forever apart from the monitor and the monitored player who only punish profitable unilateral deviations. Then, $(f, \hat{com}) \in SE^{ET}(\hat{G}^{g,\delta})$ for $\delta \geq \frac{2}{3}$, as can be easily verified.

However, there is a history and a corresponding observation profile for which the monitor has a profitable lie. Suppose that the player deviates and chooses C instead of D, as prescribed by f. Then, his monitor is better off to tolerate the deviation and would claim that it did not occur. His payoff in the remainder of the game is strictly individually rational, and thus above 1, the payoff he would receive with some delay by starting the punishment. Thus, $(f, \hat{com}) \in SE^{ET}(\hat{G}^{g,\delta})$ but $(f, \hat{com}) \notin SE(\hat{G}^{g,\delta})$ since the monitor has a profitable lie if he is not forced to tell the truth.

Obviously, there is a communication profile $\tilde{com} \in Com$ which validates the monitor's lie. Hence, $(f, \tilde{com}) \in SE(\hat{G}^{g,\delta})$, but $(f, \tilde{com}) \notin SE^{ET}(\hat{G}^{g,\delta})$ since under $\tilde{com}$ the monitor does not have to reveal unprofitable unilateral deviations of the monitored player. □

Thus, there are SE under exogenously imposed truthtelling which are not sustainable under strategic communication and there are SE under strategic communication which are not compatible with truthtelling.

5.1 Related Literature

Strategic communication in networks can be modelled in various ways.[16] The approach taken in this paper also relates to the literature on communication in repeated games.[17]

Few papers combine both ideas. Ben-Porath and Kahneman (1996) study SE of infinitely repeated discounted games in which the players form a (not necessarily connected) network. The players publicly announce their own action choices and observations made about their neighbors in a strategic way, that is, including lies. When each group contains three or more players unilateral deviations are detectable, and hence, do not occur in equilibrium. In Ben-Porath and Kahneman (2003) this idea is extended. Since monitoring is costly, only one monitor is assigned to every player. After an incompatible announcement, which in equilibrium does not occur, both players are punished and the monitor is substituted. In comparison, the network in this paper is connected, though as seen in section 3.3, the trigger strategy is a SE also in an unconnected network, as long as each group contains at least two players. For the Prisoner's Dilemma, this result is stronger than Ben-Porath and Kahneman's, though theirs holds for any stage-game.

Renault and Tomala (1998) show how to sustain uniform Nash Equilibria—which is a weaker concept than SE—in finitely and infinitely repeated undiscounted games when the players form a 2-connected and directed graph. Since this implies that there are two

---

[16] See Hagenbach and Koessler (2008) for one possibility and further references.

[17] See for example Compte (1998), Kandori and Matsushima (1998) or Kandori (2003), which all resolve imperfect monitoring by communication in form of public announcements.

distinct paths between any pair of players, lies are prevented in equilibrium. In their model, however, the payoff accumulation stops during a communication phase.

6 The most central or best informed Player

There are two ways to identify the most central or best informed player. Depending on the communication profile $com \in Com$, one or both determine this player's location.

Under exogenously imposed truthtelling, the most central player is the one whose largest distance is smallest. He receives the information what all other players did in the past first and thus observes unilateral deviations first.[18] The second concept is Bonacich centrality, as defined by Ballester, Calvó-Armengol and Zenou (2006). Roughly, it counts the number of paths of different length which start in any player $i \in I$, weighted by δ. The player from which more paths stem is most central. Under strategic communication, he receives the most information which includes what his neighbors tell him that he told them that they told him and so on. Moreover, the other players in the network accumulate more information about him than about any other player.

If the players tell the truth, only largest distances matter while if they may lie, both concepts are important. Bonacich centrality identifies the player with more information and largest distances the one who first receives (possibly wrong) information about what all other players did in the past. Ballester, Calvó-Armengol and Zenou (2006) show that the two concepts do not coincide, and usually, identify different players as being most central. To formalize the results introduced in this section is left for future research.

7 Conclusion

Although the Prisoner's Dilemma is a well-studied game, there are still new results to explore. This paper studies the imposition of a network on a set of impatient players which play the repeated Prisoner's Dilemma. The players' level of patience required to sustain the trigger strategy profile as a SE is larger even in a simple star network compared with a complete one, each formed by three players. This result extends to any network. For sufficiently patient players, the trigger strategy profile is a SE in an unconnected network and under strategic communication. Most other strategy profiles which are SE under exogenously imposed truthtelling are not robust to strategic communication and some player has a profitable lie. However, new SE arise due to richer communication.

---

[18] To identify a unilateral deviation, a player needs to know a period's entire action profile.

The equilibrium payoff space as a function of the discount factor is derived for the three player star network and it is shown that there is a reduction compared to the same payoff space depicted for a complete network (for some range of discount factors). Hence, the delay caused by the network reduces the scope for cooperation. This result is important since in reality economic agents rarely obtain all relevant information immediately after it is generated, and usually, suffer from a delay. Henceforth, a model which assumes perfect monitoring overestimates the amount of cooperation sustainable in reality.

References

Abreu, D., D. Pearce and E. Stacchetti (1990), “Toward a Theory of Discounted Repeated Games with Imperfect Monitoring,” *Econometrica* 58, 1041-1063.

Ballester, C., A. Calvó-Armengol and Y. Zenou (2006), “Who’s who in Networks. Wanted: The Key Player,” *Econometrica* 74, 1403-1417.

Ben-Porath, E. and M. Kahneman (1996), “Communication in Repeated Games with Private Monitoring,” *Journal of Economic Theory* 70, 281-297.

Ben-Porath, E. and M. Kahneman (2003), “Communication in Repeated Games with Costly Monitoring,” *Games and Economic Behavior* 44, 227-250.

Bhaskar, V. and I. Obara (2002), “Belief-Based Equilibria in the Repeated Prisoners’ Dilemma with Private Monitoring,” *Journal of Economic Theory* 102, 40-69.

Compte, O. (1998), “Communication in Repeated Games with Imperfect Private Monitoring,” *Econometrica* 66, 597-626.

Cooper, R., D. DeJong, R. Forsythe and T. Ross (1996), “Cooperation without Reputation: Experimental Evidence from Prisoner’s Dilemma Games,” *Games and Economic Behavior* 12, 187-218.

Cornshaw, M. (1997), “Algorithms for Finding Repeated Game Equilibria,” *Computational Economics* 10, 139-168.

Ellison, G. (1994), “Cooperation in the Prisoner’s Dilemma with Anonymous Random Matching,” *Review of Economic Studies* 61, 567-588.

Ely, J. and J. Välimäki (2002), “A Robust Folk Theorem for the Prisoner’s Dilemma,” *Journal of Economic Theory* 102, 84-105.

Fudenberg, D., D. Levine and E. Maskin (1994), "The Folk Theorem with Imperfect Public Information," *Econometrica* 62, 997-1039.

Haag, M. and R. Lagunoff (2006), "Social Norms, Local Interaction, and Neihborhood Planning," *International Economic Review* 47, 265-296.

Hagenbach, J. and F. Koessler (2008), "Strategic Communication Networks," *mimeo.*

Kandori, M. (1992), "Social Norms and Community Enforcement," *Review of Economic Studies* 59, 63-80.

Kandori, M. (2002), "Introduction to Repeated Games with Private Monitoring," *Journal of Economic Theory* 102, 1-15.

Kandori, M. (2003), "Randomization, Communication, and Efficiency in Repeated Games with Imperfect Public Monitoring," *Econometrica* 71, 345-353.

Kandori, M. and H. Matsushima (1998), "Private Observation, Communication and Collusion," *Econometrica* 66, 627-652.

Kinateder, M. (2008), "Repeated Games Played in a Network," *mimeo.*

Kreps, D., P. Milgrom, J. Roberts and R. Wilson (1982), "Rational Cooperation in the Finitely Repeated Prisoners' Dilemma," *Journal of Economic Theory* 27, 245-252.

Lippert, S. and G. Spagnolo (2005), "Networks of Relations and Social Capital," *mimeo.*

Mailath, G. and L. Samuelson (2006), *Long-Run Relationships*, Oxford University Press.

Piccione, M. (2002), "The Repeated Prisoner's Dilemma with Imperfect Private Monitoring," *Journal of Economic Theory* 102, 70-83.

Renault, J. and T. Tomala (1998), "Repeated Proximity Games," *International Journal of Game Theory* 27, 539-559.

Sekiguchi, T. (1997), "Efficiency in Repeated Prisoner's Dilemma with Private Monitoring," *Journal of Economic Theory* 76, 345-361.

Stahl, D. (1991), "The Graph of the Prisoner's Dilemma Supergame Payoffs as a Function of the Discount Factor," *Games and Economic Behavior* 3, 368-384.

Vega-Redondo, F., M. Marsili and R. Slanina (2005), "Clustering, Cooperation, and Search in Social Networks," *Journal of the European Economic Association* 3, 628-638.

Appendix A Enforceability and Decomposability

The reader is referred to Mailath and Samuelson (2006), section 2.5, for a detailed derivation of the concepts used in this appendix. For each action profile, the set of continuation payoffs that enforce it and the set of payoffs that are decomposable is obtained, respectively. It is assumed that a player who observed a deviation chooses D forever. In case in the network depicted in Figure 1, player 1 or 3 deviates, the other of the two players is assumed to choose the same action again in the subsequent period since he did not yet observe the deviation (see footnote 15). This is the only difference with respect to Mailath and Samuelson in which punishment obviously starts immediately while in the star network full punishment sets in only after a delay of one period.

a) Action profile $a = (C, C, C)$. For player 2, as derived above, $x_{CCC} \geq \frac{1}{\delta}$ and for players 1 and 3, the incentive constraint not to deviate is

$$(1-\delta)3 + \delta x_{CCC} \quad \geq \quad (1-\delta)4 + \delta[(1-\delta)2 + \delta 1],$$

where $[(1-\delta)2 + \delta 1] = x_{DCC} = x_{CCD}$. This simplifies to $x_{CCC} \geq \frac{1}{\delta} + (1-\delta)$. Hence,

$$A_{CCC} = \{(x_{CCC}, x_{CCC}, x_{CCC}) \in \mathbb{R}^3 \mid x_{CCC} \geq \tfrac{1}{\delta} \text{ for 2, and } x_{CCC} \geq \tfrac{1}{\delta} + (1-\delta) \text{ for 1 and 3}\}.$$

The set of payoffs decomposable using (C, C, C) and $\mathcal{F}^*$ is then

$$B_{CCC} = \{x \in \mathcal{F}^* \mid x = (1-\delta)(3,3,3) + \delta y, \text{ for any } y \in \mathcal{F}^* \text{ and } x \in A_{CCC}\}.$$

b) Action profile $a = (D, D, D)$. Any feasible and individually rational payoff is incentive compatible. Thus, $A_{DDD} = \mathcal{F}^*$ and

$$B_{DDD} = \{x \in \mathcal{F}^* \mid x = (1-\delta)(1,1,1) + \delta y, \text{ for any } y \in \mathcal{F}^*\}.$$

c) Action profile $a = (C, D, D)$. Only a deviation by player 1 might be profitable. The incentive constraint for him not to deviate is

$$(1-\delta)0 + \delta x_{CDD} \quad \geq \quad (1-\delta)1 + \delta x_{DDD},$$

where $(1-\delta)1+\delta x_{DDD}=1$. This simplifies to $x_{CDD}\geq\frac{1}{\delta}$. Hence,

$$A_{CDD}=\{(x_{CDD},x_{CDD},x_{CDD})\in\mathbb{R}^3\mid x_{CDD}\geq\tfrac{1}{\delta}\text{ for }1\}.$$

The set of payoffs decomposable using (C,D,D) and $\mathcal{F}^*$ is

$$B_{CDD}=\{x\in\mathcal{F}^*\mid x=(1-\delta)(0,2,2)+\delta y,\text{ for any }y\in\mathcal{F}^*\text{ and }x\in A_{CDD}\}.$$

d) Action profile $a=(D,D,C)$. Only a deviation by player 3 might be profitable and in analogy with the previous case,

$$A_{DDC}=\{(x_{DDC},x_{DDC},x_{DDC})\in\mathbb{R}^3\mid x_{DDC}\geq\tfrac{1}{\delta}\text{ for }3\}.$$

The set of payoffs decomposable using (D,D,C) and $\mathcal{F}^*$ is

$$B_{DDC}=\{x\in\mathcal{F}^*\mid x=(1-\delta)(2,2,0)+\delta y,\text{ for any }y\in\mathcal{F}^*\text{ and }x\in A_{DDC}\}.$$

e) Action profile $a=(D,C,D)$. Only a deviation by player 2 might be profitable and in analogy with cases c) and d),

$$A_{DCD}=\{(x_{DCD},x_{DCD},x_{DCD})\in\mathbb{R}^3\mid x_{DCD}\geq\tfrac{1}{\delta}\text{ for }2\}.$$

The set of payoffs decomposable using (D,C,D) and $\mathcal{F}^*$ is

$$B_{DCD}=\{x\in\mathcal{F}^*\mid x=(1-\delta)(2,0,2)+\delta y,\text{ for any }y\in\mathcal{F}^*\text{ and }x\in A_{DCD}\}.$$

f) Action profile $a=(C,C,D)$. The incentive constraint for player 2 not to deviate is

$$(1-\delta)0+\delta x_{CCD}\quad\geq\quad(1-\delta)1+\delta x_{CDD},$$

where $(1-\delta)1+\delta x_{CDD}=1$. This simplifies to $x_{CCD}\geq\frac{2}{\delta}-1$. The same condition holds for player 1 since player 3 chooses D anyway. Hence,

$$A_{CCD}=\{(x_{CCD},x_{CCD},x_{CCD})\in\mathbb{R}^3\mid x_{CCD}\geq\tfrac{2}{\delta}-1\text{ for 1 and 2}\}.$$

The set of payoffs decomposable using (C,C,D) and $\mathcal{F}^*$ is

$$B_{CCD}=\{x\in\mathcal{F}^*\mid x=(1-\delta)(0,0,4)+\delta y,\text{ for any }y\in\mathcal{F}^*\text{ and }x\in A_{CCD}\}.$$

g) Action profile $a=(D,C,C)$. The incentive constraints for player 2 and 3 are derived analogously as in the previous case. Hence,

$$A_{DCC} = \{(x_{DCC}, x_{DCC}, x_{DCC}) \in \mathbb{R}^3 \mid x_{DCC} \geq \tfrac{2}{\delta} - 1 \text{ for 2 and 3}\}.$$

The set of payoffs decomposable using (D, C, C) and $\mathcal{F}^*$ is

$$B_{DCC} = \{x \in \mathcal{F}^* \mid x = (1-\delta)(4,0,0) + \delta y, \text{ for any } y \in \mathcal{F}^* \text{ and } x \in A_{DCC}\}.$$

h) Action profile $a = (C, D, C)$. The incentive constraints for players 1 and 3 not to deviate is

$$(1-\delta)0 + \delta x_{CDC} \quad \geq \quad (1-\delta)1 + \delta[(1-\delta)2 + \delta 1],$$

where $(1-\delta)2 + \delta 1 = x_{DDC} = x_{CDD}$. This simplifies to $x_{CCD} \geq \frac{2}{\delta} - \delta$. Hence,

$$A_{CDC} = \{(x_{CDC}, x_{CDC}, x_{CDC}) \in \mathbb{R}^3 \mid x_{CDC} \geq \tfrac{2}{\delta} - \delta \text{ for 1 and 3}\}.$$

The set of payoffs decomposable using (C, D, C) and $\mathcal{F}^*$ is

$$B_{CDC} = \{x \in \mathcal{F}^* \mid x = (1-\delta)(0,4,0) + \delta y, \text{ for any } y \in \mathcal{F}^* \text{ and } x \in A_{CDC}\}.$$