Symmetric players in repeated games: Theory and evidence

Christoph Kuzmics*, Thomas Palfrey†, and Brian Rogers‡

March 14, 2010

preliminary draft

*Managerial Economics and Decision Sciences, Kellogg School of Management, Northwestern University, c-kuzmics@kellogg.northwestern.edu.

†California Institute of Technology, trp@hss.caltech.edu

‡Managerial Economics and Decision Sciences, Kellogg School of Management, Northwestern University, b-rogers@kellogg.northwestern.edu.

1 Abstract/Introduction

In this paper we study repeated games in which the stage game is symmetric. We study symmetric, we call them attainable (from Crawford and Haller (1990) and Alos-Ferrer and Kuzmics (2008)), equilibria of such games.

We show the following. First, restricting attention to attainable strategies does not restrict the set of feasible payoff profiles. Second, there is no folk theorem for attainable equilibria: not every feasible payoff profile can be justified in an attainable equilibrium. In fact, highly asymmetric and close to efficient payoff profiles are not possible in an attainable equilibrium. Third, the set of attainable payoff profiles has positive Lebesgue measure. Fourth, we then show that ex-ante (Pareto) efficient attainable equilibria must be ex-post symmetric. Fifth, there is a unique such attainable equilibrium which dominates most (if not all) others. This is based on playing according to the so-called (well-known) Thue-Morse sequence. Sixth and finally, if any attainable equilibrium of the repeated game is focal (in the spirit of Schelling (1960) and as defined by Alos-Ferrer and Kuzmics (2008) it is probably one which is based on the meta-norm of efficiency and simplicity (based on complexity of finite automata). This gives rise to turn-taking in the 2x2 BoS and some rotation scheme more generally in nxn BoS games (to be defined in the main body of the paper). This theory is completely testable by means of lab-experiments, which we also endeavor to do.

2 Related Literature

Discuss, among others, Crawford and Haller (1990), Blume (2000), Blume and Gneezy (2000), Blume and Gneezy (2008).

3 Model

Let $\Gamma = (I, A, u)$ be a symmetric n-player stage game, which is played repeatedly at discrete points in time $t = 0, 1, 2, ...$ Players have a common discount factor δ and maximize the discounted stream of stage game payoffs, normalized by $(1 - \delta)$. The notation is taken, as much as possible, from Mailath and Samuelson (2007).

3.1 The stage game

We study a class of n-player generalizations of the well-known 2-player Battle of the Sexes game.

Definition 1. *A normal form stage game* $\Gamma = (I, A, u)$ *is an* $n \times n$ Battle of the Sexes Game *if* $A = \{1, ..., n\}$*, i.e.* $|A| = |I| = n$*, and* $u(a) = (0, ..., 0)$ *for all action profiles* $a \in A^n$ *with the property that there are two players* $i \neq j$ *such that* $a_i = a_j$*, while* $u_i(a) = x_{a_i}$ *for all* $i \in I$ *if* a *is a permutation of* $(1, 2, ..., n)$*. We furthermore require that* $x_1 < x_2 < ... < x_n$*.*

Note that all players in an $n \times n$ BoS game are symmetric, while there are no symmetric strategies. Note that for $n = 2$ this reduces to the usual 2-player BoS game.

3.2 The repeated game

The game is played repeatedly at discrete points in time $t = 0, 1, 2, ...$ In each period players only observe one of two outcomes c, n. Let $Y = \{c, n\}$. If $a \in A^n$ is played at stage t then players observe c if a is a permutation of $(1, 2, ..., n)$. Otherwise they observe n. Thus c is the "event" that players achieved coordination (a non-zero payoff vector) and n is the event that they did not achieve coordination and thus obtained 0 payoffs all. The payoff matrix is assumed to be common knowledge. Players, thus, know what payoff they got at each stage. Indeed, this information structure is induced by players observing their own payoff and nothing else.

We thus can describe public and private histories for the repeated game. The *set of public histories* is given by $\mathcal{H} = \bigcup_{t=0}^{\infty} Y^t$ with $Y^0 = \emptyset$. Player i's *set of private histories* is given by $\mathcal{H}^* = \mathcal{H}_i = \bigcup_{t=0}^{\infty} (A \times Y)^t$. Given the symmetry we have $\mathcal{H}_i = \mathcal{H}_j = \mathcal{H}^*$ for all $i, j \in I$. A pure strategy for any player i is a mapping $\sigma : \mathcal{H}^* \to A$. A behavioral strategy is a mapping $\sigma : \mathcal{H}^* \to \Delta(A)$.

It will be useful for us to represent (some) repeated game strategies by (finite) automata. Following Mailath and Samuelson (2007), an automaton (for a single player) is defined by a tuple $(\mathcal{W}, w^0, f, \tau)$, where $\mathcal{W}$ is the set of states, $w^0 \in \mathcal{W}$ is the (deterministic) initial state, $f : \mathcal{W} \to \Delta(A)$ is the output or decision function, and $\tau : \mathcal{W} \times A \times Y \to \mathcal{W}$ is the (deterministic) transition function.

Finally we need to specify payoffs. For a pure strategy profile $\sigma = (\sigma_1, ..., \sigma_n)$ payoffs are given by $u_i = (1 - \delta) \sum_{t=0}^{\infty} \delta^t u_i(a^t(\sigma))$, where $a^t(\sigma)$ is the action profile induced by strategy profile σ in period t. For mixed strategy profiles we extend u_i by simply taking expectations.

3.3 The solution concept

The key departure we make in this paper from the usual repeated games literature is that we restrict attention to attainable strategies. Attainable strategies are those that respect the symmetry structure of the game (see Crawford and Haller (1990) for symmetric strategies and Alos-Ferrer and Kuzmics (2008) for symmetry also between players). Attainability in the presence of strategy symmetry is the requirement that a player's symmetric strategies be used with the same probability by this player. We shall here focus on player symmetry. All games in this paper are such that there are no symmetric strategies. The idea of the requirement of attainability in the presence of player symmetry is perhaps best explained as follows. One (the only one?) justification for Nash equilibrium is that if a publicly observed prediction (recommendation) is made as to how to play a game, this prediction must satisfy the conditions of a Nash equilibrium in order for the prediction to (possibly) come true. Now suppose in a truly symmetric game, in which players have no commonly known names (or commonly known distinguishable characteristics)[1], a prediction (recommendation) was made that one player will (should) play one strategy and another player will (should) choose another strategy. Then, as players cannot be called by their name, they would have to figure out for themselves who plays what in this game. If they then cannot communicate it will not be possible for them to follow such a recommenda-

[1]Note that the requirement that the players have no commonly known names or characteristics is important here. If they do – for instance it is commonly known that one of the, say 2, players is a man, while the other is a woman – one can announce that the woman do one thing, while the man do another. This does not present a problem and the 2 players can play as recommended. However, if, for instance, in an anonymous auction one would like one bidder to bid one way and another another way, this will not be possible without the risk of mis-coordination.

tion without the positive probability event of both choosing the same strategy. Thus predictions (recommendations) in the presence of player symmetry must describe a single strategy profile that all symmetric players can use, in order to have any hope that play will follow the prediction. Of course, the prediction must also be a (sequential) Nash equilibrium. Otherwise at least one (in fact each) player will have an incentive to deviate. Thus, we shall only look at attainable (sequential) Nash equilibria of the repeated game.

4 Feasible Payoffs

The requirement of attainability does not to a large extent reduce the set of feasible payoffs. Of course, if both players are ex-ante symmetric and follow the same repeated game strategy, then ex-ante and in expectation they will have to get the same payoff. However, along the path of play they can get violently different payoffs.

Consider a repeated $n \times n$ BoS game. Consider the following strategy in automata form, $(\mathcal{W}, w^0, f, \tau)$. Let $w^0 = R \in \mathcal{W}$ be the initial state. Let $f(R)$ be the uniform distribution over all actions in A. Let $\tau(R, a, y) = R$ if $y = y_2$, and $\tau(R, a, y) = S(a)$ for some $S(a) \in \mathcal{W}$ with $S(a) \neq S(b)$ for all $a, b \in A$. Thus, this automaton represents a strategy in which all players initially randomize uniformly over all actions in every period until coordination is achieved (i.e. all players use a different action). After that, if all players use this automaton, all players' automata will now be in different states. Thus, from that point on they can, in principle, play any (asymmetric) strategy profile of the repeated game. Let $v(a)$ denote the continuation payoff for the player who played action $a \in A$ when symmetries were broken, i.e. when coordination was achieved for the first time. Let $v \in \mathbb{R}^n$ denote the vector of all these payoffs. Given the uniform distribution[2] over actions the probability of coordination in any given stage is $q = \frac{n!}{n^n} > 0$. Thus, players will eventually coordinate, symmetries are broken, and players obtain the payoff-vector v from the continuation play in the now unrestricted repeated game. Let $w(a)$ denote the ex-ante expected payoff to the player who ended up playing action a when symmetries were broken. Then

$$w(a) = qv(a) \sum_{t=0}^{\infty} \delta^t (1-q)^t = qv(a) \frac{1}{1 - \delta(1-q)}.$$

Let the set of feasible payoff-vectors of the non-symmetric repeated game be denoted by $\mathcal{F}_\Gamma \subset \mathbb{R}^n$. Note that $\mathcal{F}_\Gamma$, while in principle dependant on the discount factor δ, is actually constant for all $\delta \geq \bar{\delta}$ for some given $\bar{\delta} < 1$. This follows from a result in Sorin (1986), concisely stated also as Lemma 3.7.1 in Mailath and Samuelson (2007). In the 2-player BoS, in fact, we have that $\mathcal{F}_\Gamma(\delta) = \mathcal{F}_\Gamma$ for all $\delta \geq \frac{1}{2}$. Let $\mathcal{F}_\Gamma^a(\delta) \in \mathbb{R}^n$ denote the set of feasible attainable payoff-vectors. Then

$$\mathcal{F}_\Gamma^a = \{w \in \mathbb{R}^n | w = v \frac{q}{1 - \delta(1-q)} \text{ for some } v \in \mathcal{F}_\Gamma\}.$$

It is thus obvious that, as δ tends to 1, i.e. as players get more and more patient, the set of feasible payoffs under attainable strategies coincides with the set of feasible payoffs under all strategies.

Proposition 1. *Let Γ be an $n \times n$ Battle-of-the-Sexes game. Then*

$$\lim_{\delta \to 1} \mathcal{F}_\Gamma^a(\delta) = \mathcal{F}_\Gamma.$$

[2] Note that the uniform distribution, in fact, here maximizes the probability that symmetries are broken.

Proof: Immediate. QED

5 Attainable Equilibrium Payoffs: Not a Folk Theorem

While attainability, thus, hardly poses a restriction on the set of feasible payoffs, this section demonstrates that attainability does impose interesting restrictions on the set of equilibrium payoffs, even as δ tends to 1.

We shall here not provide a complete characterization of the set of attainable equilibrium profiles in all of our generalized n-player BoS games. While, this would be interesting in its own right, it is somewhat tangential to the objective of this paper. However, the fact that attainability does pose **some** restriction on equilibrium payoffs, and that this restriction, on the other hand, also does not eliminate every interesting payoff pair, is instrumental to this paper. Thus, we shall demonstrate exactly this by means of the simplest 2-player BoS example, which is a special case of both classes of n-player BoS games. For this example we show that highly asymmetric payoff pairs with a large total payoff are not justifiable in any attainable equilibrium of the repeated game. Also, there is a set of payoff-pairs with positive Lebesgue measure, each element of which is justifiable by some attainable equilibrium.

Consider a particular 2×2 BoS game with $x_1 = 0$ and $x_2 = 1$. I.e. the stage game is given by

	A	B
A	0,0	1,0
B	0,1	0,0

.

Note that action A weakly dominates action B within the stage game. Note that for this stage game observing y_1 or y_2, i.e. whether or not coordination has been achieved, plus one's own action is sufficient for players to know exactly what has been played. Thus, for this stage game, the repeated game is one of perfect monitoring.

To obtain an interesting upper bound on the set of attainable equilibrium payoff pairs in this repeated game, we shall appeal to a fixed point argument of an appropriate function in the set of potential attainable equilibrium payoff pairs. Consider time 0 or any time period in which symmetries have not yet been broken (i.e. so far both players always played the same strategy in each round). The potential outcomes of play in this stage then are the pure strategy combinations AA, AB, BA, and BB. Suppose AB occurs. Then the A-player obtains a payoff of 1, the B-player one of 0, and symmetries are now broken. Thus, the continuation from here on can be any (asymmetric) equilibrium of the repeated game. Thus, the continuation payoffs can be any element in $\mathcal{F}$, the set of feasible payoff-pairs. This follows from the standard and well-known folk theorem[3]. We can, in fact combine the payoff of 1 and the later and discounted continuation in $\mathcal{F}$, by giving the two players some appropriate payoff-pair in $\mathcal{F}$ right at this time 0 (after observing outcome AB). Thus, in this case after observing outcome AB at this time, continuation payoffs can be any pair $(\alpha_0, \beta_0) \in \mathcal{F}$, where $\alpha_0, \beta_0 \geq 0$ and $\alpha_0 + \beta_0 \leq 1$, where α_0 is the payoff to the A-player and β_0 is the payoff to the B-player.

Similarly, after outcome BA players can be taken to receive some payoff-pair in $\mathcal{F}$. An important restriction attainability imposes here, though, is that whatever "player 1" gets after

[3] In fact, in this game, we don't even need the folk theorem for this result. Players simple play some (asymmetric) Nash equilibrium in each stage

BA must be the same as what "player 2" gets after AB. I.e. in each case the A-player gets some α_0 and the B-player some β_0.

After outcomes AA and BB both payers receive a payoff of 0 and symmetries are not yet broken. Thus, now before players use a stage game strategy at time 1 the game is exactly as it was at time 0. Thus, the continuation payoffs given AA or BB can again be any attainably equilibrium payoff-pair. Note, however, that play does not necessarily have to continue in the same way after AA and BB. Let us denote the set of attainable equilibrium payoff-pairs by $\mathcal{E}^a(\delta)$. Thus, after AA a continuation might be any $(\alpha_1, \beta_1) \in \mathcal{E}^a(\delta)$, which, as we shall see, is a more severe restriction than just $\alpha_1, \beta_1 \geq 0$ and $\alpha_1 + \beta_1 \leq 1$. After BB continuation payoffs might also be any $(\alpha_2, \beta_2) \in \mathcal{E}^a(\delta)$, possibly different from (α_1, β_1).

Thus, at time 0 players expect the following continuation profiles: $(\alpha_0, \beta_0) \in \mathcal{F}$ to the A- and B-player, respectively, after AB or BA; discounted $(\alpha_1, \beta_1) \in \mathcal{E}^a(\delta)$ after AA; and discounted $(\alpha_2, \beta_2) \in \mathcal{E}^a(\delta)$ after BB. Any such continuation profile induces (the same) incentives for the two (symmetric) players, which governs their choice of mixed stage game strategy at time 0, which, in turn, determines their expected payoffs at stage 0. Thus, we can define a function $f : \mathcal{G} \to \mathcal{G}$, with $\mathcal{G}$ the set of all subsets of $\mathcal{F}$, which assigns to a candidate attainable equilibrium set $\mathcal{E} \in \mathcal{G}$, used for continuation after AA and BB the set of all payoff-pairs at time 0.

The set of attainable equilibrium payoff-pairs $\mathcal{E}^a$ must be a subset of the largest fixed point of f. The proof of the following proposition uses this fact in order to derive an upper bound $\bar{\mathcal{E}}^a(\delta) \subset \mathcal{F}$ such that $\mathcal{E}^a(\delta) \subset \bar{\mathcal{E}}^a(\delta)$.

Proposition 2.

$$(\alpha, \beta) \in \lim_{\delta \to 1} \mathcal{E}^a(\delta) \Rightarrow \alpha, \beta \leq \frac{3}{4}$$

Proof: Suppose $\mathcal{E}(\bar{\alpha})$ is such that any $(\alpha, \beta) \in \mathcal{E}(\bar{\alpha})$ satisfies $\alpha, \beta \leq \bar{\alpha}$. We have $\mathcal{E}(1) = \mathcal{F}$, but for $\bar{\alpha} < 1$ $\mathcal{E}(\bar{\alpha})$ is a proper subset of $\mathcal{F}$.

Suppose $(\alpha_1, \beta_1) \in \mathcal{E}(\bar{\alpha})$ is the continuation after AA and $(\alpha_2, beta_2) \in \mathcal{E}(\bar{\alpha})$ is the continuation after BB. Furthermore let $(\alpha_0, \beta_0) \in \mathcal{F}$ be the (immediately paid out) continuation after AB and BA. Let $x_A \in [0, 1]$ denote the probability players attach to pure action A in stage 0, which must be the same for both players by attainability.

At stage 0, when players contemplate their choice of x_A, expected payoffs from choosing pure action A and B are given by

$$u(A, x_A) = \frac{\alpha_1 + \beta_1}{2} \delta x_A + (1 - x_A)\alpha_0,$$

and

$$u(B, x_A) = \beta_0 x_A + (1 - x_A)\frac{\alpha_2 + \beta_2}{2}\delta.$$

This follows from the continuation payoffs and the fact that after AA and BB both players are equally likely, given attainability, to end up being the A-player or B-player at the very moment when symmetries are broken.

Players now choose x_A such that neither of them has an incentive to deviate to another (mixed) action. Thus, we are looking for a symmetric equilibrium of the following symmetric 2×2 game.

	A	B
A	$\delta\frac{\alpha_1+\beta_1}{2}, \delta\frac{\alpha_1+\beta_1}{2}$	α_0, β_0
B	β_0, α_0	$\delta\frac{\alpha_2+\beta_2}{2}, \delta\frac{\alpha_2+\beta_2}{2}$

.

Note that if this game has only pure equilibria (such as A being a dominant strategy) then the ex-ante, at time 0, expected payoff must be less than or equal to $\delta\bar{\alpha} < \bar{\alpha}$. Thus, if these were the only equilibria in this game, then $\bar{\alpha} = 0$ and $\mathcal{E}(0) = \emptyset$. Thus, in order to get something interesting we need to consider continuations such that this game has a strictly mixed equilibrium. There are two possibilities for this to happen. Either the game is a coordination game or of the Hawk-Dove variety.

In order to have a coordination game we must have $\alpha_0 \leq \frac{\alpha_2+\beta_2}{2}\delta < \frac{1}{2}$ and $\beta_0 \leq \frac{\alpha_1+\beta_1}{2}\delta < \frac{1}{2}$. In order for this game to be of the Hawk-Dove variety we must have $\alpha_0 \geq \frac{\alpha_2+\beta_2}{2}\delta$ and $\beta_0 \geq \frac{\alpha_1+\beta_1}{2}\delta$, which, by $\alpha_0 + \beta_0 \leq 1$, implies that $\alpha_0 \leq 1 - \frac{\alpha_1+\beta_1}{2}\delta > \frac{1}{2}$ (as long as δ sufficiently close to 1). Thus, for α_0 the Hawk-Dove case is less restrictive than the coordination case.

In either case the unique completely mixed symmetric equilibrium mixed action is given by

$$x_A^* = \frac{2\alpha_0 - (\alpha_2 + \beta_2)\delta}{2(\alpha_0 + \beta_0) - (\alpha_1 + \beta_1 + \alpha_2 + \beta_2)\delta}.$$

Given the continuation profile and the induced x_A^* we then have that the ex-ante, at stage 0, expected payoff to the eventual A-player (at the moment symmetries are broken) is given by

$$\alpha^* = (x_A^*)^2\alpha_1\delta + 2x_A^*(1 - x_A^*)\alpha_0 + (1 - x_A^*)^2\alpha_2\delta.$$

We are now trying to find an upper bound for the maximal α^*, given $\alpha_1, \alpha_2 \leq \bar{\alpha}$, as well as the aforementioned restrictions on the continuation profile, and subject to the incentive constraints.

In order to do so we distinguish 3 cases. Suppose first that $x_A^* < \frac{1}{2}$. Thus

$$\frac{2\alpha_0 - (\alpha_2 + \beta_2)\delta}{2(\alpha_0 + \beta_0) - (\alpha_1 + \beta_1 + \alpha_2 + \beta_2)\delta} < \frac{1}{2},$$

or, equivalently,

$$2\alpha_0 - (\alpha_2 + \beta_2)\delta < 2\beta_0 - (\alpha_1 + \beta_1)\delta.$$

Thus,

$$\alpha_0 - \beta_0 < \frac{1}{2}\left[(\alpha_2 + \beta_2)\delta - (\alpha_1 + \beta_1)\delta\right].$$

Given $\alpha_0 + \beta_0 \leq 1$, we finally have,

$$\alpha_0 < \frac{1}{2} + \frac{1}{4}\delta \leq \frac{3}{4}.$$

As long as $\alpha_0 \leq \bar{\alpha}\delta$ we must have $\alpha^* \leq \bar{\alpha}\delta$ as well. This is definitely true if $\frac{1}{2} + \frac{1}{4}\delta \leq \bar{\alpha}\delta$, i.e. if $\bar{\alpha} \geq \frac{\frac{1}{2}+\frac{1}{4}\delta}{\delta}$. Thus, for $x_A^* < \frac{1}{2}$ we have $\alpha^* < \bar{\alpha}$, and, thus, the mapping f is a contraction (for such x_A^*), if $\bar{\alpha} \leq \frac{3}{4}$ (in the limit when $\delta \to 1$).

For the second case, assume $x_A^* = \frac{1}{2}$. Then it is straightforward to show that $\alpha^* \leq \frac{1}{2} + \frac{1}{4}\delta\bar{\alpha} \leq \frac{3}{4}$.

For the third and final case assume $x_A^* > \frac{1}{2}$. Note that obviously

$$\alpha^* \leq \alpha^{**} = \max_{x \in [\frac{1}{2},1], \alpha_0, \alpha_1, \alpha_2} x^2\alpha_1\delta + 2x(1 - x)\alpha_0 + (1 - x)^2\alpha_2\delta,$$

subject to the given restrictions on the α_i's. Now α^{**} is, obviously, higher the higher all α_i's. However, there are some restrictions on these α_i's deriving from the requirement that the game at hand is either a coordination game or a Hawk-Dove game. These restrictions are more stringent than at least $\alpha_0 \leq 1 - \frac{\alpha_1+\beta_1}{2}\delta > \frac{1}{2}$.

Thus, $\alpha^{**} \leq \max_{x\in[\frac{1}{2},1],\alpha_1\in[0,\bar{\alpha}]} x^2\alpha_1\delta + 2x(1-x)\left(1 - \frac{\alpha_1}{2}\delta\right) + (1-x)^2\bar{\alpha}\delta$. But for $x > \frac{1}{2}$, this expression is maximized at $\alpha_1 = \bar{\alpha}$. Thus we have

$$\alpha^{**} \leq \max_{x\in[\frac{1}{2},1]} x^2\bar{\alpha}\delta + 2x(1-x)\left(1 - \frac{\bar{\alpha}}{2}\delta\right) + (1-x)^2\bar{\alpha}\delta.$$

The right hand side of this expression is a simple quadratic in x and is less than or equal to $\bar{\alpha}\delta$ if $\left(1 - \frac{\bar{\alpha}}{2}\delta\right) < \bar{\alpha}\delta$. This, in turn, is true if $\bar{\alpha}\delta > \frac{2}{3}$.

Thus, to summarize all 3 cases, for δ close enough to 1, $\bar{\alpha} \leq \frac{3}{4}$ ensures that the mapping f is contracting. This proves the result. QED

Proposition 2, thus, states that in an attainable equilibrium no player (after symmetries are broken) can expect to gain a payoff of more than $\frac{3}{4}$. Thus, payoffs ex-post (after symmetries are broken) cannot be too extremely asymmetric (in the ex-ante expectation) and highly efficient at the same time. I.e. payoff-pairs such as $\frac{5}{6}, \frac{1}{6}$ are not possible in an attainable equilibrium. Given that Proposition 2 provides an upper bound for the set of attainable equilibrium payoff-pairs, it is silent about whether other highly asymmetric (and inefficient) payoff-pairs are achievable in an attainable equilibrium. The next proposition states that no extremely ex-post asymmetric payoff-pairs are sustainable in an attainable equilibrium.

Proposition 3.

$$(\alpha, 0) \text{ or } (0, \beta) \in \lim_{\delta\to 1} \mathcal{E}^a(\delta) \Rightarrow \alpha = \beta = 0.$$

Proof: We do the proof for $(\alpha, 0)$. The proof for $(0, \beta)$ is completely analogous. Suppose $(\alpha, 0) \in \lim_{\delta\to 1} \mathcal{E}^a(\delta)$ and suppose $\alpha > 0$. Thus the attainable repeated game strategy which implements $(\alpha, 0)$ must be such that there are finite histories after on the induced path of play such that two things are true: First, the repeated game strategy instructs players to randomize at this history and the results of the randomization cannot lead to the eventual realized B-player getting a payoff of more than 0. Thus there must be a history with prescribed stage game mixed action $(x, 1-x)$ and $x \in (0, 1)$ and the payoff consequences from all 4 possible (positive probability) events AA, BB, AB, and BA, must result in the eventual B-person receiving zero payoff. Thus continuation payoffs after AA can be at most any $(\alpha_1, 0)\delta$, after AB and BA any $(\alpha_0, 0)$, and after BB any $(\alpha_2, 0)\delta$. This provides the following expected payoffs to players at this stage: $u(A, x) = x\delta\frac{\alpha_1}{2} + (1-x)\alpha_0$ and $u(B, x) = (1-x)\delta\frac{\alpha_2}{2}$. The ex-ante expected payoff to the eventual A-player is then given by $\alpha^* = x^2\delta\alpha_1 + 2x(1-x)\alpha_0 + (1-x)^2\delta\alpha_2$.

The proof now uses a similar fixed point argument as in the proof of Proposition 2. Suppose we know that if $(\alpha, 0) \in \mathcal{E}^a(\delta)$ then $\alpha \leq \bar{\alpha}$ for some $\bar{\alpha} \in (0, 1]$. Then $\alpha^* \geq \bar{\alpha}$ only if $\alpha_0 \geq \bar{\alpha}$. But then

$$\begin{aligned} u(A, x) &= x\delta\frac{\alpha_1}{2} + (1-x)\alpha_0 \\ &\geq x\delta\frac{\alpha_1}{2} + (1-x)\bar{\alpha} \\ &\geq 0 + (1-x)\bar{\alpha} \\ &> (1-x)\frac{\alpha_2}{2}\delta \\ &= u(B, x). \end{aligned}$$

Thus, we have for any $x \in (0,1)$ that $u(A,x) > u(B,x)$, and, thus, players have no incentive to play B, and, thus, will not randomize in an equilibrium. As this is true for any $\bar{\alpha} > 0$ we, thus, must have $\bar{\alpha} = 0$. The payoff pair $(0,0)$ is, of course, implementable by an attainable equilibrium: always play A is one. QED

Thus, having proved that certain payoff-pairs are not implementable by an attainable equilibrium, we turn to the question as to what payoff-pairs can be implemented. The next proposition states that indeed there is a set of payoff-pairs, which has positive Lebesgue measure, and is such that every payoff-pair in this set can be sustained in an attainable equilibrium. In fact the next proposition completely characterizes the set of payoff-pairs sustainable in any stationary (or "Markov") equilibrium. I.e. the attainable strategies considered are such that after any history at which symmetries are not yet broken the continuation play is exactly the same. Also after any history at which symmetries were just broken, again play is exactly the same (but possibly different from the previous case).

Proposition 4. *The set of payoff-pairs for all attainable stationary equilibria, as $\delta \to 1$, has Lebesgue-measure $\frac{1}{6}$, which is $\frac{1}{3}$ of the total Lebesgue-measure of the set of feasible payoff-pairs.*

Proof: Note that each stationary attainable strategy such strategy profile has associated with it a (normalized discounted) payoff a, the continuation payoff to the player who plays A when symmetries are broken. a must be consistent with the initial mixing probability x. Notice that, provided $\delta \geq \frac{1}{2}$, any continuation $a \in [0,1]$ is feasible.[4] Given x, at any period in which players are symmetric, the probability of the symmetry breaking is $q(x) = 2x(1-x)$. Ex ante, when mixing initially with x and playing an efficient continuation once symmetries are broken, each player has expected payoff

$$v(x,\delta) = \frac{1}{2}\sum_{t=0}^{\infty} q(1-q)^t\delta^t = \frac{q}{2(1-\delta(1-q))}.$$

In order to incentive the players to mix initially, we must have equal expected payoffs from either action, which requires

$$x\delta v + (1-x)a = (1-x)\delta v + x(1-a).$$

On the left hand side, when choosing A there are two possibilities. With probability x the other player chooses A, symmetries are not broken, and the game continues at the next date with continuation payoff v. On the other hand, with probability $(1-x)$ the other player chooses B, in which case symmetries are broken and the continuation a is realized. Similarly, when choosing B, with probability $(1-x)$ symmetries are not broken, and when they are, a continuation of $1-a$ is realized. Solving, we obtain

$$a(x,\delta) = \frac{x(1-\delta x)}{1-\delta(1-q(x))}.$$

Even though players both expect v at the beginning of the game, ex post they may obtain different payoffs, due to the fact that symmetries will eventually be broken. We want to ask which payoff profiles can be supported by Markov attainable equilibria. Thus, let us name,

[4] Follows from Sorin (1986), (Mailath and Samuelson 2007, Lemma 7.3.1).

arbitrarily, the player who eventually plays A at the time symmetries are broken as the A-player, and his opponent as the B-player, and determine whether a corresponding payoff profile of (v_A, v_B) can be obtained. Of course, the answer will be the same for the permuted payoffs (v_B, v_A), since our names are arbitrary, and continuation payoffs could be assigned in a permuted way once symmetries are broken (and players receive a payoff of 0 regardless of whether they mis-coordinate on A or on B!).

We have that

$$v_A(x, \delta) = a \sum_{t=0}^{\infty} q(1-q)^t \delta^t = a \frac{q}{1 - \delta(1-q)}$$

and

$$v_B(x, \delta) = (1-a) \sum_{t=0}^{\infty} q(1-q)^t \delta^t = (1-a) \frac{q}{1 - \delta(1-q)}.$$

For a given δ, these expressions trace out a parametric curve of equilibrium payoff profiles as x varies from zero to one, starting and ending at the origin. Notice that given a pair (v_A, v_B), any payoff along the ray pointing to the origin is also attainable, by using an inefficient continuation that gives total payoff $F < 1$. We would then have $v = \frac{F}{2} \frac{q}{1-\delta(1-q)}$ and continuation payoffs once symmetries are broken of $(a, F - a)$. Thus, the region defined by the parametric curve represents the set of Markov attainable equilibrium payoffs.

We now show that for large δ, one third of feasible payoffs represent Markov attainable equilibria. First notice that $\lim_{\delta \to 1}(v_A(\frac{1}{2}, \delta), v_B(\frac{1}{2}, \delta)) = (\frac{1}{2}, \frac{1}{2})$. Define next the area

$$A' = \lim_{\delta \to 1} \int_{x=0}^{\frac{1}{2}} v_A(x, \delta) \frac{\partial v_B(x, \delta)}{\partial x} dx,$$

which gives the limiting area under the upper lobe of equilibrium payoffs. Thus, the limiting area corresponding to Markov attainable equilibrium payoffs is

$$A = 2(A' - \frac{1}{8}) = \frac{1}{6},$$

where the area A' is computed with some straightforward but tedious calculus and algebra. QED

Proposition 2 and 4 are summarized in Figure 1.

6 Ex-ante efficiency implies ex-post symmetry

Proposition 5. *Let Γ be an $n \times n$ BoS game. Consider the repeated such game with discount factor $\delta < 1$. Let $\sigma : \mathcal{H}^* \to \Delta(A)$ be an efficient equilibrium strategy. That is the repeated game payoff $u(\sigma', \sigma', ..., \sigma')$, for σ' an equilibrium strategy, is maximized at $\sigma' = \sigma$. Then, for every $\delta < 1$ (large enough), σ must be such that*

1. *at every history at which symmetries are not yet broken σ describes the uniform distribution over the action set A*

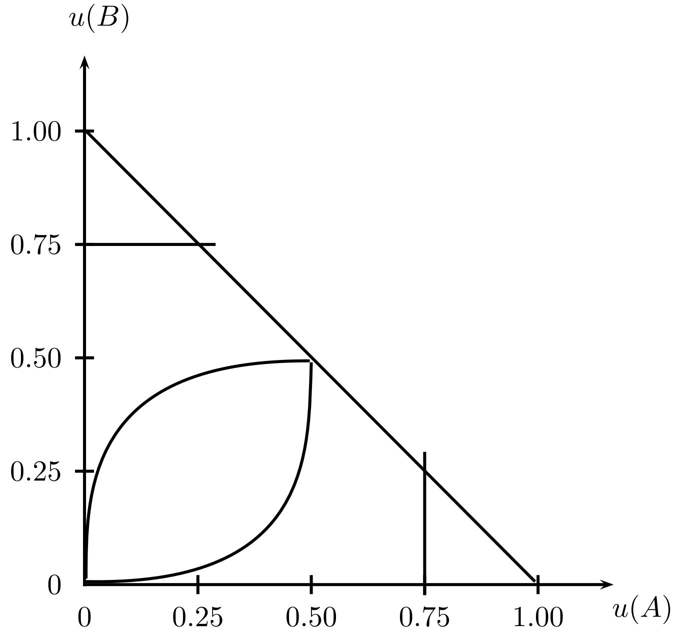

Figure 1: The oval represents the set of all payoff-pairs in attainable Markov equilibria (Proposition 4). The corners represent payoff-pairs not achievable in any attainable equilibrium (2).

2. *at every history at which symmetries are broken (for the first time) σ is an (asymmetric) repeated game strategy which yields equivalent payoffs, from the point of view of this history, regardless of what action a player played at this stage.*

Proof: Part (1) of the statement follows from two observations. First, in order to achieve a payoff-profile on the Pareto frontier the symmetry needs to be broken. I.e. the event "every player plays a different action" has to realize. No other event breaks symmetries (given our feedback, or monitoring, assumptions). Second, the uniform randomization uniquely maximizes the probability of the event "every player plays a different action". To see this let $x \in \Delta(A)$ denote a probability distribution, where x_i denotes the probability used for action $i \in A$. Then the probability that the event "every player plays a different action" is given by $P = n! \prod_{i \in A} x_i$. It is maximized when its log is maximized. Maximizing $\ln P = \ln(n!) + \sum_{i \in A} \ln(x_i)$ subject to $\sum_{i \in A} x_i = 1$ we obtain the first order conditions $\frac{1}{x_i} = c$ for some constant c for all $i \in A$. This proves part (1).

To prove part (2) recall that $\mathcal{E}^a_\Gamma(\delta)$ denotes the set of attainable equilibrium payoff-profiles. Let $\gamma = \sup_{\alpha \in \mathcal{E}^a_\Gamma(\delta)} \sum_{i \in A} \alpha_i$. This is well defined. Suppose that at a particular history action-distribution $x \in \Delta(A)$ is used. From (1) we know we must have $x_i = \frac{1}{n}$ for all $i \in A$. Now consider the expected continuation payoffs from playing action $i \in A$ at this stage. Given x has full support any action-profile can realize with positive probability. In order for σ to be efficient it has to be the case that after every such realization, in which symmetries are not broken, continuation payoff-vectors $\alpha \in \mathcal{E}^a_\Gamma(\delta)$ are such that $\sum_{i \in A} \alpha_i = \gamma$. Now consider all symmetric realization in which symmetries are broken. Here the continuation can in principle be any payoff-profile in $\mathcal{F}_\Gamma$ (provided δ is large enough). Suppose it is asymmetric, such that there are two actions $i, j \in A$ such that the continuation payoff to the i-guy is strictly greater than the continuation payoff to the j-guy. But then, in order for a player to be indifferent between all actions, this imbalance has to be equilibrated in some other events. However as all other events lead to a shared expected γ regardless of the action played, rebalancing payoffs could only be achieved by choosing some continuation to be inefficient, which provides a contradiction. QED

6.1 An example of a Pareto optimal equilibrium that fails ex post symmetry

Consider the following 3-person game, where $1 < x < 3$.

H

	H	M	L
H	0,0,0	0,0,0	0,0,0
M	0,0,0	0,0,0	0,0,0
L	0,0,0	1,x,3	0,0,0

M

	H	M	L
H	0,0,0	0,0,0	3,1,x
M	0,0,0	0,0,0	0,0,0
L	0,0,0	0,0,0	0,0,0

L

	H	M	L
H	0,0,0	0,0,0	0,0,0
M	x,3,1	0,0,0	0,0,0
L	0,0,0	0,0,0	0,0,0

This game can be thought of as one of several natural extensions of the 2-player BoS game. There are three symmetries (explain Alos-Ferrer and Kuzmics (2008) at this point?) of this game: $\{(1 \to 1, 2 \to 2, 3 \to 3), (1 \to 2, 2 \to 3, 3 \to 1), (1 \to 3, 3 \to 2, 2 \to 1)\}$. Thus any history that identifies one of the players necessarily identifies all three players, breaking all symmetries at once. The only histories, then, after which symmetries remain are HHH, MMM, and LLL.[5]

Now consider the following meta-norm, which specifies actions in any continuation game. After the first occurrence of any permutation of LMH, continue with something that is ex post symmetric (and on the Pareto frontier) by using time-averaging. After a stage in which one of the strategies is chosen by exactly one player, coordinate forever after on that player's preferred outcome. After any other profile (ie, those for which symmetries remain), continue mixing as in the previous period. It remains to be shown what mixing probabilities are implied by equilibrium given this meta-norm.

Notice that this protocol is not ex post symmetric, as it allows for the possibility of awarding per-period continuation payoffs of (any permutation of) $(1, x, 3)$.

The mixing probabilities that minimize the expected time until symmetries are broken are those that minimize the probability of realizing one of the profiles HHH, MMM, and LLL. Thus, the efficient mixing probabilities are $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$.

Our claim is that the initial equilibrium mixing implied by this protocol is $x = (x_H, x_M, x_L) = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$. For this to be true, it must be that for every action, conditional on symmetries being broken, the expected continuation values are identical. The argument for the three actions are analogous, so without loss of generality consider player 1 taking action L, and the other players choose actions according to x. Conditional on symmetries being broken, there are essentially four possibilities we must consider. First, it could be that the other players play MM or HH, in which case player one's continuation payoff is 3. This happens with probability $\frac{2}{9}$. Second, it could be the case that the other players play MH or HM, in which case player one's continuation payoff is $(1 + x + 3)/3$. This happens with probability $\frac{2}{9}$ as well. Third, it could be the case that player two plays L and player three plays M or H, in which case player one's continuation payoff is 1. This happens with probability $\frac{2}{9}$. Finally, player three could play L and player two could play M or H, in which case player one's continuation payoff is x. This also happens with probability $\frac{2}{9}$.

[5] For instance, after HML, even though payoffs are $(0, 0, 0)$, symmetries are broken because player one can be declared "the player for whom, if the other two players exchanged their actions, would achieve a payoff of three."

In summary, for any action played, there is a $\frac{1}{9}$ chance that symmetries are not broken. If, however, they are broken, there are four equally likely possibilities, that involve continuation payoffs of $1, x, 3$, or $(1 + x + 3)/3$. Thus, uniform mixing is played in equilibrium, as claimed.

The ability to construct such a protocol requires that we be able to treat the three actions in a symmetric way in the continuation games after symmetries are broken. If we are to construct a protocol that is not ex post symmetric, then to achieve Pareto-optimal mixing, it must be that these asymmetries are constructed in a balanced way in order to equalize incentives for the initial mixing. In some settings this may not be possible, in which case the only Pareto-optimal equilibria involve ex post symmetric outcomes.

Consider, for example, a requirement that the meta-norm treat realizations of H at least as well as realizations of M, which are treated at least as well as realizations of L, in any period in which symmetries are initially broken. This precludes, for instance, awarding a continuation payoff of three to player one following LMM, which the construction above utilizes. Under this condition, any meta-norm that fails ex post symmetry will imply that $x_H \geq x_M \geq x_L$, with at least one inequality strict. The only Pareto efficient equilibria will be these that are ex post symmetric.

7 Implementing symmetric continuations

In this section we address the question as to how one would implement the ex-post symmetric play for specific games with discounting.

We are at the stage in the repeated game where symmetries are just completely broken and players are ready to continue with an ex-post payoff symmetric (but strategy asymmetric) continuation strategy. First, consider any discount factor $\delta < 1$. It follows from Lemma 3.7.1 in Mailath and Samuelson (2007) originally due to Sorin (1986) that indeed an exact payoff symmetric continuation can be constructed from an appropriate sequence of pure stage game action profiles, provided δ is large enough[6]. Indeed it is easy to see that there are many such possible constructions.

We shall here be interested in the limiting case when $\delta \to 1$. We shall do two things in this section. First, for the special case, the 2×2 BoS, we will identify a unique special continuation protocol that has the property that for any other protocol there is a $\bar{\delta} < 1$ such that for all $\delta \geq \bar{\delta}$ the special protocol provides more symmetric payoffs than the other protocol. This special protocol involves playing actions in a particular sequence which turns out to be a known sequence, the Thue-Morse sequence. This sequence is the obvious way to play if we pair the meta-norm of Pareto-efficiency in the game at hand with the meta-norm of Pareto-efficiency even in nearby (slightly discounted) games. Another obvious way to play based on the meta-norm of Pareto-efficiency can be found when pairing it with the meta-norm of simplicity. Among all Pareto-efficient repeated game strategy profile there is an arguably unique simplest one, when simplicity is defined as state-simplicity of the automaton inducing the repeated game strategy.

Consider again the particular 2×2 BoS game with $x_1 = 0$ and $x_2 = 1$ from before. Consider symmetries broken in some period, redefined as period 0, with a "winner" who played A, and a "loser" who played B. There is now a multitude of possible continuations (always playing pure stage game Nash equilibria). As before let $\{y_t\}_{t=0}^{\infty}$ (with $y_0 = 1$) denote one possible continuation (from the point of view of the winner, A). We could, as before, work out the (per

[6]Exact bounds on δ can be given for the result to hold.

period) payoff to A, denoted $u^y(A) = (1-\delta)\sum_{t=0}^{\infty}\delta^t y_t$. Then the payoff to the loser, B, is given by $u^y(B) = 1-u^y(A)$. It will be useful to work with the difference between the two as a function of δ. Define $\Delta^y(\delta) = u^y(A)(\delta) - u^y(B)(\delta)$.

We already know that for any $\frac{1}{2} \leq \delta < 1$ there is a sequence $y \in \{0,1\}^\infty$ such that $\Delta^y(\delta) = 0$.

We are now interested in the following question. Is there a special sequence, denoted $y^s \in \{0,1\}^\infty$, such that for any periodic sequence $y \in \{0,1\}^\infty$ there is a $\bar{\delta} < 1$ such that for all $\delta \in [\bar{\delta}, 1)$ we have that $\Delta^{y^s}(\delta) < \Delta^y(\delta)$?

The answer is yes. In fact, it is given by the Thue-Morse sequence, denoted z.[7] Set $z_0 = 1$, and define the sequence recursively by $z_{2k} = z_k$ and $z_{2k+1} = 1 - z_k$.

It is well known that this sequence can alternatively be constructed as follows. Set $z_0 = 1$. Proceed iteratively, at each step replacing every instance of 1 with 10, and replacing every instance of 0 with 01. Another alternative construction is the following. Let us generally call $\{z_t\}_{t=0}^{2^k-1}$ the block of size 2^k. Let $\{1-z_t\}_{t=0}^{2^k-1}$ denote the inverse of this block (i.e. 1's are replaced by 0's and vice versa). We have the block of size 1. The block of size 2 is then given by the block of size 1 and then the inverse of the block of size 1. In general, the block of size 2^k is given by the block of size 2^{k-1} (which accounts for its first half) plus the inverse of this very same block (which accounts for the second half. Yet another well-known equivalent characterization of the Thue-Morse sequence is the following. Let $z_t = 1$ whenever the binary expansion of t has an even number of 1's and let $z_t = 0$ otherwise.

A simple lemma will prove useful.

Lemma 1. *Let $z \in \{0,1\}^\infty$ (with $z_0 = 1$ as the first element) be the Thue-Morse sequence. Let $\mathbf{1}(t)$ be the number of 1's in the sequence z up to and including time t. Define $\mathbf{0}(t)$, analogously, as the number of 0's in z up to time t. If t is odd then*

1. *$z_t = 1 - z_{t-1}$, and*
2. *$\mathbf{1}(t) = \mathbf{0}(t)$.*

Proof: The first statement follows from directly from the definition, $1 - z_{2k+1} = z_{2k} = z_k$. The second statement follows from the first (by induction). QED

The Thue-Morse sequence z has the property that, for sufficiently large δ, it awards the 1 to the player who has the lower present discounted payoff.

Proposition 6. *For every t there exists a $\bar{\delta} < 1$ such that the following is true. If $z_t = 1$ then $\sum_{k=0}^{t-1}\delta^k z_k < \sum_{k=0}^{t-1}\delta^k(1-z_k)$ for all $\delta > \bar{\delta}$, and if $z_t = 0$ then $\sum_{k=0}^{t-1}\delta^k z_k > \sum_{k=0}^{t-1}\delta^k(1-z_k)$ for all $\delta > \bar{\delta}$.*

Proof: The statement is definitely true for $t = 1$ and $t = 2$. Suppose now that $t \geq 3$ is odd. Then $\mathbf{0}(t-2) = \mathbf{1}(t-2)$ by the second part of Lemma 1. By the first part we then have $z_t = 1 - z_{t-1}$. Suppose $z_{t-1} = 1$. Then $\mathbf{1}(t-1) = \mathbf{0}(t-1) + 1$ and, thus, there is a $\bar{\delta}$ such that for all $\delta > \bar{\delta}$ we have $\sum_{k=0}^{t-1}\delta^k z_k > \sum_{k=0}^{t-1}\delta^k(1-z_k)$. Given $z_t = 0$ the statement is true in this case. Now suppose that $z_{t-1} = 0$. Then $\mathbf{1}(t-1) = \mathbf{0}(t-1) - 1$ and, thus, there is a $\bar{\delta}$ such that for all $\delta > \bar{\delta}$ we have $\sum_{k=0}^{t-1}\delta^k z_k < \sum_{k=0}^{t-1}\delta^k(1-z_k)$. Given $z_t = 1$ the statement is true also in this case. This proves the statement for t odd.

[7] See Thue (1906, 1912) and Morse (1921). Allouche and Shallit (1999) provide a useful discussion.

We now turn to t even. We know the statement is true for $t = 1$ and $t = 2$. Suppose the statement is true for all $k \leq t-1$. As t is even it can be written as $2k$ for the appropriate k. We have

$$\begin{aligned}\sum_{i=0}^{2k-1} \delta^i z_i &= \sum_{i=0}^{k-1} \delta^{2i} z_{2i} + \sum_{i=0}^{k-1} \delta^{2i+1} z_{2i+1} \\ &= \sum_{i=0}^{k-1} (\delta^2)^i z_i - \delta \sum_{i=0}^{k-1} (\delta^2)^i z_i \\ &= (1-\delta) \sum_{i=0}^{k-1} (\delta^2)^i z_i\end{aligned}$$

Since $k \leq t-1$, we know that if $z_k = 1$, then there is a $\bar{\delta}$ such that $\sum_{i=0}^{k-1} (\delta^2)^i z_i < \sum_{i=0}^{k-1} (\delta^2)^i (1 - z_i)$ for all $\delta^2 < \bar{\delta}$ and so under the same condition $\sum_{i=0}^{2k-1} \delta^i z_i < \sum_{i=0}^{2k-1} \delta^i (1 - z_i)$. By definition, $z_{2k} = z_k = 1$, and the case is satisfied for all $\delta > \sqrt{\bar{\delta}}$. The argument is parallel for $z_k = 0$. QED

The difference in payoffs for this sequence turns out to be relatively simple to describe. It is given by $\Delta^z(\delta) = (1-\delta) \prod_{k=0}^{\infty} \left(1 - \delta^{2^k}\right)$. This can be verified by observing that the difference in payoffs after the first block is $\Delta^0(\delta) = 1$. After the second block it is $\Delta^1(\delta) = \Delta^0(1-\delta)$, after the third $\Delta^2(\delta) = \Delta^1(1-\delta^2)$, and generally after the $k+1$-st it is $\Delta^k(\delta) = \Delta^{k-1}(1-\delta^{2^{k-1}})$. Finally the desired difference is $\Delta^z(\delta) = (1-\delta) \lim_{k\to\infty} \Delta^k(\delta)$, which can be expressed as the infinite product given above.

We now show that if y is given by the infinite repetition of a given block (of z) of size 2^k for some k, then $\Delta^y(\delta) > \Delta^z(\delta)$.

Suppose we take the block of size 2 and repeat it infinitely. This is tit-for-tat. Then for the difference we have a 1 every other period beginning with period 0 and a -1 every other period beginning in period 1. Thus $\Delta^T = (1-\delta)\frac{1-\delta}{1-\delta^2}$. This is simply $\frac{(1-\delta)}{1-\delta^2}\Delta^1(\delta)$. In general it is true that if y is the repetition of the block of size 2^k we get $\Delta^y(\delta) = \frac{1-\delta}{1-\delta^{2^k}} \Delta^{k-1}(\delta) = \frac{1-\delta}{1-\delta^{2^k}} \prod_{l=0}^{k-1} \left(1 - \delta^{2^{k-1}}\right)$.

Note that the inequality $\Delta^y(\delta) > \Delta^z(\delta)$ is, thus, given by

$$\frac{1-\delta}{1-\delta^{2^k}} \prod_{l=0}^{k-1} \left(1 - \delta^{2^{k-1}}\right) > (1-\delta) \prod_{l=0}^{\infty} \left(1 - \delta^{2^l}\right),$$

which reduces to

$$\frac{1}{1-\delta^{2^k}} > \prod_{l=k}^{\infty} \left(1 - \delta^{2^l}\right),$$

which after multiplying by $(1-\delta^{2^k})$, is clearly true.

Indeed we can show that for any sequence y with a given periodicity of, say, k, there is a $\bar{\delta} < 1$ (can be made to depend only on k) such that the sequence y is less symmetric than the special sequence for all $1 > \delta > \bar{\delta}$.

A sequence y has periodicity k if $y_{t+k} = y_t$ for all $t \geq 0$. Let $\Delta^y(\delta)$ denote the difference in payoff for the winner versus the loser if sequence y is played after symmetries are broken. Note that for all balanced[8] periodic sequences y we have that $\lim_{\delta\to 1} \Delta^y(\delta) = 0$.

[8]A sequence y is balanced if $\lim_{T\to\infty} \frac{1}{T} \sum_{t=0}^{T} y_t = \frac{1}{2}$.

Proposition 7. *For every k there exists a $\bar{\delta} < 1$ such that, for every sequence y with periodicity k, $\Delta^y(\delta) > \Delta^z(\delta)$ whenever $\delta > \bar{\delta}$.*

Proof: A sequence of periodicity k is characterized by its first k entries. Given its structure it is straightforward to express $\Delta^y(\delta)$ as a function of these first k entries. In fact,

$$\Delta^y(\delta) = \frac{1-\delta}{1-\delta^k} \sum_{t=0}^{k-1} \delta^t w_t,$$

where $w_t = 2y_t - 1$, i.e. it is $+1$ when $y_t = +1$ and is -1 when $y_t = 0$.

Note that $\sum_{t=0}^{k-1} \delta^t w_t$ is thus some polynomial in δ of degree $k-1$. What we are interested in here is whether, and if so how, it converges to 0 as δ tends to 1. Let $A^0(\delta) = \sum_{t=0}^{k-1} \delta^t w_t$. Suppose first that $A^0(1) \neq 0$. Then this sum, obviously, does not tend to 0. Suppose now that $A^0(1) = 0$. Then $A^0(\delta)$ can be factored (divided) by $(1-\delta)$. Let, thus, $A^0(\delta) = (1-\delta)A^1(\delta)$, where $A^1(\delta)$ is another polynomial in δ, but of degree $k-2$. We can, again, either have $A^1(1) = 0$ or not. In the latter case, the sum of interest then tends to zero in the same fashion as $1-\delta$ does. In the former case, we can factor out another $1-\delta$. Repeating this argument $k-1$ times we obtain that

$$\Delta^y(\delta) \geq \frac{(1-\delta)^k}{1-\delta^k}.$$

It remains to be shown that $\Delta^z(\delta)$ (for our special sequence) tends to zero faster than any such $\Delta^y(\delta)$ for some periodic y. Recall that $\Delta^z(\delta) = (1-\delta)\lim_{k\to\infty} \Delta^k(\delta)$, where $\Delta^k(\delta) = \Delta^{k-1}(1-\delta^{2^{k-1}})$. Alternatively we can write

$$\Delta^k(\delta) = \prod_{j=0}^{k-1} \left(1-\delta^{2^j}\right).$$

Note that, for any $j \geq 1$ the expression $\left(1-\delta^{2^j}\right)$ can be written as the product of $\left(1-\delta^{2^{j-1}}\right)$ and $\left(1+\delta^{2^{j-1}}\right)$. The former term can then, by the same argument, factorized into another two such terms. Repeating this, and doing this for all terms in the above product we obtain an alternative representation of $\Delta^k(\delta)$, given by

$$\Delta^k(\delta) = (1-\delta)^k \prod_{j=0}^{k-1} \left(1+\delta^{2^j}\right)^{k-j-1}.$$

We still have $\Delta^z(\delta) = (1-\delta)\lim_{k\to\infty} \Delta^k(\delta)$. Note that we also have $\Delta^k(\delta) > \Delta^{k+1}(\delta)$ for all k. Thus we also have $\Delta^z(\delta) < (1-\delta)\Delta^k(\delta)$.

Finally, we thus have

$$\begin{aligned}
\Delta^z(\delta) &< (1-\delta)\Delta^k(\delta) \\
&< (1-\delta)^{k+1}\prod_{j=0}^{k-1}\left(1+\delta^{2^j}\right)^{k-j-1} \\
&< (1-\delta)^{k+1}\prod_{j=0}^{k-1}2^{k-j-1} \\
&< (1-\delta)^{k+1}2^{\sum_{j=0}^{k-1}k-j-1} \\
&< (1-\delta)^{k+1}2^{\frac{k(k-1)}{2}}.
\end{aligned}$$

Thus implies that $\Delta^z(\delta)$ tends to zero, when δ tends to 1, at least an order faster than $\Delta^y(\delta)$ when y has a given, but arbitrary, periodicity k. This completes the proof. QED

In fact we can provide a very accurate approximation of $\Delta^z(\delta)$ that makes it very clear that it tends to zero faster than any power of $(1-\delta)$.

First, note that $2\Delta^z(\frac{1}{2})$ is a given, positive number. It is roughly 0.35 (are there better ways to express this number?).

Fix $\frac{1}{2}<\delta<1$. There will be a $J(\delta)>0$ such that $1-\delta^{2^j}<\frac{1}{2}$ for all $j<J(\delta)$ and $1-\delta^{2^j}\geq\frac{1}{2}$ for all $j\geq J(\delta)$.

We can then write

$$\begin{aligned}
\Delta^z(\delta) &= (1-\delta)\prod_{k=0}^{\infty}\left(1-\delta^{2^k}\right) \\
&= (1-\delta)\prod_{k=0}^{\lceil J(\delta)\rceil-1}\left(1-\delta^{2^k}\right)\prod_{k=\lceil J(\delta)\rceil}^{\infty}\left(1-\delta^{2^k}\right) \\
&\approx (1-\delta)\prod_{k=0}^{\lceil J(\delta)\rceil-1}\left(1-\delta^{2^k}\right)2\Delta^z(\frac{1}{2}) \\
&\approx 0.35(1-\delta)\prod_{k=0}^{\lceil J(\delta)\rceil-1}\left(1-\delta^{2^k}\right).
\end{aligned}$$

To find $J(\delta)$ we need to solve $1-\delta^{2^{J(\delta)}}=\frac{1}{2}$ or, equivalently $\delta^{2^{J(\delta)}}=\frac{1}{2}$. Taking (natural) logs twice and rearranging we obtain

$$J(\delta)=\frac{1}{\ln 2}\ln\left(\frac{\ln 2}{-\ln\delta}\right).$$

Note that $J(\delta)$ tends to infinity as δ tends to 1. It does, so, however very slowly (quantify?).

Thus we also obtain that $\Delta^z(\delta)$ tends to zero faster than any power of $(1-\delta)$.

8 Complexity

Proposition 8. *Let Γ be an $n\times n$ BoS game. Consider the repeated such stage game and the limiting case where the discount factor $\delta\to 1$. Consider all efficient (attainable) strategy profiles $\{\sigma\}^n$. Let $(\mathcal{W},w^0,f,\tau)$ be its automaton representation. Then $|\mathcal{W}|\geq n+1$.*

Proof: Note first that there must be a state $w \in W$ such that $f(w)$ is uniform over A by Proposition 5. Note also that in order for this automaton to eventually lead to an ex-post symmetric continuation we must have that all actions must be played purely after some history. I.e. for every action i we must have a state w_i such that $f(w_i)$ attaches probability 1 on a single action. QED

Define the complexity of a strategy σ the smallest number of states in any automaton that implements this strategy.

Note that there are less complex automata that implement (inefficient) strategies. For instance consider the highly inefficient automaton with just a single state, describing to play action 1. This automaton could be employed by everyone (and is even an equilibrium in an appropriate $n \times n$ BoS game) yielding a discounted payoff of 0.

Note that for an $n \times n$ BoS there are, thus, $n! * (n-1)!$ least complex automata, which implement an efficient strategy. They all share an initial randomization state and have some protocol as to how to switch from one state (and action) to another, such that after every history in which symmetries are broken, they play a perfectly coordinated action profile $1, 2, ..., n$ or some permutation thereof.

Note that for the 2×2 BoS this leaves a unique least complex automaton, which one could call tit-for-tat. Well, actually two. One in which after the first AB play continues with BA and one in which play continues with AB (and then alternates).

Now consider the 3×3 BoS game. Note that there are essentially 2 (each 6-times, explain!) different least complex automata that implement efficient strategies. All share the same set of states, say $\{R, 1, 2, 3\}$, where R is the initial state with $f(R)$ uniform over $A = \{1, 2, 3\}$. Suppose, w.l.o.g. that state i describes to play action i. The automata now can differ in the sequence of states. We thus could have $1 \to 2 \to 3 \to 1$ ("clockwise") or $3 \to 2 \to 1 \to 3$ ("counterclockwise").

In special (knife-edge) cases there is yet another 4-state automaton that implements an efficient strategy. Suppose $x_1 = 0, x_2 = 1$, and $x_3 = 2$. Then the automaton with the same 4 states but with the following transition protocol also implements an efficient strategy: $1 \to 3 \to 1$ and $2 \to 2$. Note that one could argue that this automaton is even simpler (less complex) than the previous ones, as it has the following property: there will be a time (finite in expectation) at which we can downsize the automaton to either 2 or even only 1 state (depending on circumstances). Note that the previous automata could be downsized to 3 states, but no less. Thus, in the long run this last automaton is less complex (uses less resources) than the former ones.

Note, finally, that another inefficient automaton, which may, under some circumstance, not be too inefficient, can be eventually downsized to 1 state under all contingencies. This is the 4-state automaton in which $1 \to 1$, $2 \to 2$, and $3 \to 3$. If x_1, x_2, x_3 are all close to each other relative to 0, such an automaton is not too inefficient (when δ not too close to 1) and is even more simple that tit-for-tat and the other automata mentioned above.

9 Experimental Evidence

TBD

References

Allouche, J.-P., and J. Shallit (1999): "The Ubiquitous Prouhet-Thue-Morse Sequence," in *Sequences and Their applications, Proc. SETA'98*, ed. by T. Ding, C. Helleseth, and H. Niederreiter, p. 116. Springer-Verlag, New York.

Alos-Ferrer, C., and C. Kuzmics (2008): "Hidden symmetries and focal points," Mimeo.

Blume, A. (2000): "Coordination and Learning With a Partial Language," *Journal of Economic Theory*, 95, 1–36.

Blume, A., and U. Gneezy (2000): "An Experimental Investigation of Optimal Learning in Coordination Games," *Journal of Economic Theory*, 90, 161–172.

——— (2008): "Cognitive Forward Induction and Coordination Without Common Knowledge: An Experimental Study," mimeo.

Crawford, V., and H. Haller (1990): "Learning How to Cooperate: Optimal Play in Repeated Coordination Games," *Econometrica*, 58, 571–596.

Mailath, G., and L. Samuelson (2007): *Repeated Games and Reputations: Long-run Relationships*. Oxford University Press, Oxford, UK.

Morse, M. (1921): "Recurrent Geodesics on a Surface of Negative Curvature," *Trans. Amer. Math. Soc.*, 22, 84–100.

Schelling, T. C. (1960): *The Strategy of Conflict*. Harvard University Press, Cambridge.

Sorin, S. (1986): "On repeated games with complete information," *Mathematics of Operations Research*, 11, 147–160.

Thue, A. (1906): "Über unendliche Zeichenreihen," *Norske vid. Selsk. Skr. Mat. Nat. Kl.*, 7, 1–22, Reprinted in Selected Mathematical Papers of Axel Thue (Ed. T. Nagell). Oslo: Universitetsforlaget, pp. 139-158, 1977.

——— (1912): "Über die gegenseitige Lage gleicher Teile gewisser Zeichenreihen," *Norske vid. Selsk. Skr. Mat. Nat. Kl.*, 1, 1–67, Reprinted in Selected Mathematical Papers of Axel Thue (Ed. T. Nagell). Oslo: Universitetsforlaget, pp. 413-478, 1977.