Supplementary Notes for "Coalition Formation With Binding Agreements"

Kyle Hyndman

Southern Methodist University

and

Debraj Ray

New York University

December 2006

In these notes we provide some supplementary results and discussion not included in the main text of the paper. In particular, we provide a proof of the existence of a Markovian equilibrium, the proof of Proposition 3, some examples for three-player games and further discussion of the efficiency result. We also provide the proof of the four player inefficiency example and construct an equilibrium with an inefficient (globally) absorbing state.

1 Existence of Markovian Equilibrium

Theorem. *Suppose that the set of states is finite. Then there exists a stationary Markov equilibrium.*

Proof. For any finite set A let $\Delta(A)$ stand for the set of all probability measures on A. For any state x, let $H(x)$ be the set of all (S, y) pairs such that $S \in \mathcal{S}(x, y)$. Restrict attention to states x such that $H(x) \neq \varnothing$.[1] For any i, let $H(x, i)$ be a subset of $H(x)$ constrained by the further requirement that $i \in S$ for every $(S, y) \in H(x, i)$. Define

$$\Sigma_i \equiv \prod_x \{\Delta(H(x)) \times [0,1]^{|H(x,i)|}\}$$

[1] That is, there is *some* state y and some approval committee S that can approve a transition from x to y. For all those x for which $H(x) = \varnothing$, there is no need to specify any strategies for the players.

Then Σ_i may be viewed as the collection of i's Markovian strategies (indeed, those for which i's response does not depend on the identity of the proposer). So $\Sigma \equiv \prod_i \Sigma_i$ is the collection of all Markovian strategy profiles σ. Clearly, Σ is identifiable with a compact convex subset of some Euclidean space.

Use the notation $V_i(x, \sigma)$ instead of V_i^σ for value functions. It is easy to see that $V_i(x, \sigma)$ is continuous in σ. Now we construct a suitable fixed-point mapping.

Fix some strategy profile σ. For each person i and each state x, define

$$B(i, x, \sigma) \equiv \underset{(T,z):T \in \mathcal{S}(x,z)}{\operatorname{argmax}} \left[\Lambda_i(x, z, T)V_i(z, \sigma) + (1 - \Lambda_i(x, z, T))V_i(x, \sigma)\right]$$

where $\Lambda_i(x, z, T) = \prod_{j \in T \setminus i} \lambda_j(x, z, T)$ is the probability of the proposal being accepted under the profile σ. Note that $\Delta(B(i, x, \sigma))$ is the set of "best-response" proposals by i at state x, given σ. Define

$$M_i(\sigma) \equiv \prod_x \Delta(B(i, x, \sigma)).$$

Now for acceptance-rejection decisions. For every state x and $(S, y) \in H(x, i)$, collect *all* values of $\lambda_i(x, y, S)$ such that (a) $\lambda_i(x, y, S) = 1$ if $V_i(y, \sigma) > V_i(x, \sigma)$, (b) $\lambda_i(x, y, S) = 0$ if $V_i(y, \sigma) < V_i(x, \sigma)$, and (c) $\lambda_i(x, y, S) \in [0, 1]$ if $V_i(y, \sigma) = V_i(x, \sigma)$. Define $L_i(\sigma)$ to be the product of these collections as (y, S) goes over all elements of $H(x, i)$ and then as x ranges over all states. These are the best-response accept-reject decisions by i, given σ.

To complete the construction, define $R_i(\sigma) \equiv M_i(\sigma) \times L_i(\sigma)$. It is easy to see that for each σ, $R_i(\sigma)$ is a nonempty, compact convex subset of Σ_i, and that it is upperhemicontinuous in σ. Then the product correspondence $R(\sigma) \equiv \prod_i R_i(\sigma)$ has all the same properties, and maps from Σ to Σ. By Kakutani's fixed point theorem, the correspondence has a fixed point which can easily be seen to be an equilibrium. ||

2 Further Discussion of the Efficiency Theorems

2.1 A Compact Set of States

Here we prove that our efficiency result extends to the case of a compact set of states, with the additional restriction that the proposer protocol is deterministic.

Proposition 3. *Suppose that every individual is benign, the proposer protocol is deterministic and the set of states is compact. Then in characteristic function games with permanently binding agreements, every limit payoff of every pure strategy equilibrium is efficient.*

Proof. We first provide another proof of Proposition 1. Given the restriction to a deterministic protocol and pure strategies, there is no underlying source of randomness. Therefore, we may write:

$$V_i^\sigma(h_t) = (1-\delta)u_i(x_t) + \delta V_i^\sigma(h_{t+1})$$

for all $i \in N$. It is still the case that $V_i^\sigma(h_{t+1}) \geq u_i(x_t)$, which again implies that $V_i^\sigma(h_t) \geq u_i(x_t)$. We also claim that

$$V_i^\sigma(h_{t+1}) \geq V_i^\sigma(h_t)$$

for all histories h_t, for if not, then for some history $V_i^\sigma(h_{t+1}) < V_i^\sigma(h_t)$, which, given the definition of the value function implies:

$$V_i^\sigma(h_t) < (1-\delta)u_i(x_t) + \delta V_i^\sigma(h_t)$$

which cannot be true. Therefore, we know that the equilibrium value function is monotonically increasing.

Of course, since the set of states is compact, equilibrium values must be contained in a compact set. Moreover, since values are nondecreasing over histories, we know that it must converge to some limit V_i^*, but this implies that $u_i(x_t)$ also converges to u_i^* and the equilibrium is absorbing. Indeed, it is easy to see that $V_i^\sigma(h_t)$, $V_i^\sigma(h_{t+1}) \to V_i^*$ so it must be that $V_i^* = u_i^*$.

We now show that in any equilibrium, convergence *must* occur in finite time. Notice that $V_i^\sigma(h_t) \geq \max\{u_i(x_t), V_i^\sigma(h_{t-1})\}$. Therefore, on the equilibrium path we know that

$$V_i^\sigma(h_t) \geq \max\{u_i(x_t), u_i(x_{t-1}), \ldots, u_i(x_0)\} \tag{1}$$

Given (1) and that the set of payoffs within a given coalition is Pareto efficient with respect to that coalition, we know the following:

(a) If any coalition structure is visited more than once with distinct payoffs for at least two players, it cannot be absorbing.

(b) Convergence to $\mathbf{u}^*$ must be from *below.* That is, for *all* t, $\mathbf{u}(x_t) \leq \mathbf{u}^*$.

We now show that convergence must occur in finite time. For each coalition structure π, let $U(\pi) \subset \mathbb{R}^N$ denote the set of feasible payoffs, and let $B(\mathbf{u}, \epsilon)$ denote an open ball around $\mathbf{u}$ with radius ϵ. First suppose that there exists some $\epsilon > 0$ small enough such that $\mathbf{u}^* \in B(\mathbf{u}^*, \epsilon) \cap U(\pi') \neq \phi$ for a *unique* coalition structure, π'. In this case, it is a direct result of (a) that convergence must occur in finite time.

Otherwise, for all $\epsilon > 0$, there is a set of coalition structures $\Pi(\epsilon) = \{\pi_1, \ldots, \pi_K\}$, with $K > 1$, such that $\mathbf{u}^* \in B(\mathbf{u}^*, \epsilon) \cap \left[\cap_{k=1}^K U(\pi_k)\right] \neq \phi$ *and* for all $\pi \in \Pi(\epsilon)$, there exists $t \geq T(\epsilon)$ such that $\pi(x_t) = \pi$.[2] Take the largest such set $\Pi(\epsilon)$ (which surely exists since there are only

[2]That is, along the equilibrium path, the coalition structure π is actually visited, and $\mathbf{u}(x_t) \in B(\mathbf{u}^*, \epsilon)$.

a finite number of coalition structures).

By assumption, we have that for all $\pi \in \Pi(\epsilon)$, $\mathbf{u}^* \in U(\pi)$ *and* that there exists a state x_t along the equilibrium path such that $\pi(x_t) = \pi$ and $\mathbf{u}(x_t) \leq \mathbf{u}^*$. If $\mathbf{u}(x_t) = \mathbf{u}^*$, then convergence must have occurred (and in finite time).[3] If $\mathbf{u}(x_t) < \mathbf{u}^*$, it contradicts the Pareto efficiency of the payoff set within a given coalition structure since $\mathbf{u}(x_t)$, $\mathbf{u}^* \in U(\pi(x_t))$. Thus convergence *must* be in finite time.

To complete the proof, we must show that $\mathbf{u}^*$ is efficient. The proof of this makes use of the same benignness argument as in the main text of paper and is, therefore, omitted. ||

2.2 *Observable Behavior Strategies*

Proposition 2 fails to hold if behavior strategies are used by the players with publicly observable randomization devices. Consider the following game with four states.

$$\begin{aligned}
x_1: \quad \mathbf{u}(x_1) &= (2, 1, 1) \\
x_2: \quad \mathbf{u}(x_2) &= (2.5, 2, 2) \\
x_3: \quad \mathbf{u}(x_3) &= (2, 3, 1) \\
x_4: \quad \mathbf{u}(x_4) &= (2, 1, 3)
\end{aligned}$$

We assume that transitions from any state require the consent of all players, and that player 1 proposes whenever the state is x_1. We construct a history dependent equilibrium such that the sample path $\{x_1, x_1, \ldots\}$ has strictly positive probability.

Clearly, since all players' consent is required for any transition, x_2, x_3 and x_4 must be absorbing. Recall that behavior strategies are observable and can be conditioned upon,[4] and player 1 proposes with probability 1 from x_1. Let $p(n) = \frac{1}{(n+1)^2}$ and consider a t-period history of the form: $h_t = \{(x_1, p(i))\}_{i=1}^{t}$. That is, in each period i, player 1 proposed x_2 with probability $p(i)$ and x_1 with probability $1 - p(i)$, but the realisation was x_1. Consider the following strategies:

- Player 1 uses the randomisation device $p(t)$ in each period to make his proposal, assuming $p(i)$ has been used for all $i < t$, otherwise, he proposes x_3 or x_4 as required below.

- If player 1 uses $p(t)$ as his randomisation device and x_2 is realised, then both players 1 and 2 accept the offer. If x_1 is realised, they may either accept or reject. However, if $p(t)$ is not used, player 2 rejects the proposal. If player 2 accepts, then player 3 rejects.

[3] Once $\mathbf{u}^*$ is reached, no player would ever accept a transition giving him/her a lower one-period payoff. This follows because the limit payoff is $\mathbf{u}^*$ and so the player will never be compensated for accepting such a transition.

[4] To be sure, we need to expand the notion of a history for this to make sense.

- In the first history after $p_t \neq p(t)$ and a rejection by player i, $i = 2, 3$, player 1 offers a transition to x_{i+1} and the offer is accepted by both players. If it is rejected by $j \neq i$, player 1 continues offering x_{i+1}, while if it is rejected by player i, player 1 switches and offers x_{j+1}.

It is easy to see that these strategies constitute an equilibrium for δ large enough. Moreover, since $\sum_{t=1}^{\infty} \frac{1}{(t+1)^2} < \infty$, we have that $\prod_{t=1}^{\infty}(1 - p(t)) > 0$, which means that the path $\{x_1, x_1, \ldots\}$ has strictly positive probability and that along this path, each player's expected value converges only in the limit.

3 Other Examples For Three-Player Games

In the examples below, if a coalition structure is omitted, it means that either every player obtains an arbitrarily large negative payoff or there is some legal impediment to the formation of that coalition structure. In all of the examples of this section, we assume minimal approval committees; for example, from the singletons, players 1 and 2 can approve a transition to any state y such that $\pi(y) = \pi_1$.

3.1 More on Inefficiency

One response to the inefficiency example of Section 4.1.1 in the main text is that the inefficient state described there will never be reached *starting from the singletons.* Setting the initial state to the singletons has special meaning: presumably this is the state from which all negotiations commence. However, this is wrong on two fronts (at least for Markov equilibria) as we now show.

3.1.1 Coordination Failures

Coordination failures, leading to inefficiency from *every* initial state, are a distinct possibility, even in three player games. Consider the following:

$$\begin{array}{lcllcl} x_0: & \pi(x_0) & = & \pi_0, & \mathbf{u}(x_0) & = & (2,2,2) \\ x_1: & \pi(x_1) & = & \pi_1, & \mathbf{u}(x_1) & = & (-1,1,1) \\ x_2: & \pi(x_2) & = & \pi_2, & \mathbf{u}(x_2) & = & (1,-1,1) \\ x_3: & \pi(x_3) & = & \pi_3, & \mathbf{u}(x_3) & = & (1,1,-1) \end{array}$$

Result 1. *Suppose that everyone proposes with equal probability at every date. Then, for $\delta \in [\frac{3}{5}, 1)$, there is an MPE in which x_i is absorbing, and from x_0, there is a transition to x_i with probability $\frac{1}{3}$ for $i = 1, 2, 3$.*

3.1.2 Convergence to Inefficiency From The Singletons

Consider the following example, which is a variation on the "failed partnership" example of Section 4.1.1.

$$\begin{aligned} x_0: \quad \pi(x_0) &= \pi_0, \quad \mathbf{u}(x_0) = (5,5,5) \\ x_1: \quad \pi(x_1) &= \pi_1, \quad \mathbf{u}(x_1) = (0,6,8) \\ x_2: \quad \pi(x_2) &= \pi_2, \quad \mathbf{u}(x_2) = (3,0,10) \\ x_3: \quad \pi(x_3) &= \pi_3, \quad \mathbf{u}(x_3) = (4,4,0) \end{aligned}$$

Result 2. *For any history-independent proposer protocol such that at x_0 each player has strictly positive probability of proposing, there exists $\bar{\delta} \in (0,1)$ such that if $\delta \geq \bar{\delta}$, all stationary Markovian equilibria involve a transition from x_0 to x_3 — and full absorption into x_3 thereafter — with strictly positive probability.*

Proof. Let $\alpha = (\alpha_1, \alpha_2, \alpha_3) \in \text{int}(\Delta)$ denote the proposers' protocol at x_0. First notice that in every equilibrium x_1 and x_2 must be absorbing. The states x_1 and x_2 give players 2 and 3, respectively, their unique maximal payoff. Moreover, at x_1 (resp. x_2) player 2 (resp. player 3) has veto power over any transition. Second, in every equilibrium, x_0 cannot be absorbing. This follows because players 2 and 3 can always initiate a transition to x_1 and earn a higher payoff.

We now proceed with the rest of the proof. First, we rule out a "cycle" by proving the following: If there is a positive probability transition from x_0 to x_3, then x_3 must be absorbing. Indeed, suppose not. Then for $i = 1, 2$, $V_i(x_0) = V_i(x_3) = 4$. But then, from x_0, player 1 will always reject a transition to x_2, which means that $V_2(x_0) \geq 5$, a contradiction.

Next suppose that the probability of reaching x_3 from the singletons is zero. Observe that $V_1(x_0) \leq 3$, for if not, x_1 is the only absorbing state reachable from x_0, implying that $V_1(x_0) \to 0$ for δ sufficiently high, a contradiction. Similarly, $V_3(x_0) \leq 8$, for if not, x_2 is the only absorbing state reachable from the singletons. But then for δ sufficiently high, $V_2(x_0) \leq 4$, implying that players 1 and 2 would initiate a transition to x_3, a contradiction. Finally, observe that since x_3 is not reached with positive probability, it must be that $V_2(x_0) \geq 4$, since otherwise, 1 would offer x_3 and it would be accepted.

Let p_i denote the probability of a transition from x_0 to x_i for $i = 0, 1, 2$. By assumption, $p_3 = 0$ and we have just shown that $p_1, p_2 > 0$. Given p_i, write the equilibrium value functions and take the limit as $\delta \to 1$ to obtain:

$$\begin{aligned} \bar{v}_1(x_0) &= \tfrac{3p_2}{1-p_0} && \leq 3 \\ \bar{v}_2(x_0) &= \tfrac{6p_1}{1-p_0} && \geq 4 \\ \bar{v}_3(x_0) &= \tfrac{8p_1+10p_2}{1-p_0} && \leq 8 \end{aligned} \tag{2}$$

From the third equation in (2), we see that $p_2 = 0$, which then implies that the first equation is satisfied with strict inequality. Therefore, player 1 strictly prefers to propose x_2, and the offer will be accepted by player 3. Hence, $p_2 > \alpha_1 > 0$, a contradiction. It then follows that for δ sufficiently high the same conclusion may be drawn. ||

3.2 Cyclical Equilibria

Next, equilibrium *cycles* become a distinct possibility:

$$\begin{array}{rcccl} x_0 : \pi(x_0) & = & \pi_0, \; \mathbf{u}(x_0) & = & (1,1,1) \\ x_1 : \pi(x_1) & = & \pi_1, \; \mathbf{u}(x_1) & = & (0,2,1) \\ x_2 : \pi(x_2) & = & \pi_2, \; \mathbf{u}(x_2) & = & (\frac{1}{2},4,1) \end{array}$$

Result 3. *Suppose that everyone proposes with equal probability at every date. Then there is an equilibrium with the following transitions:*

$$x_0 \underset{2/3}{\longrightarrow} x_1 \underset{1}{\longrightarrow} x_2 \underset{2/3}{\longrightarrow} x_0$$

3.3 Dynamic Inefficiency In Every Equilibrium

Though we did not formally prove this for characteristic functions, every Markovian equilibrium must exhibit full *dynamic* efficiency from *some* initial state. This is no longer true for games with externalities:

$$\begin{array}{rcccl} x_0 : \pi(x_0) & = & \pi_0, \; \mathbf{u}(x_0) & = & (1,1,1) \\ x_1 : \pi(x_1) & = & \pi_1, \; \mathbf{u}(x_1) & = & (10,0,0) \\ x_2 : \pi(x_2) & = & \pi_2, \; \mathbf{u}(x_2) & = & (0,10,0) \\ x_3 : \pi(x_3) & = & \pi_3, \; \mathbf{u}(x_3) & = & (0,0,10) \end{array}$$

If x_i, $i = 1,2,3$ were absorbing, then for $j \neq i$, $V_j(x_i) = 0$. However, notice that in every Markovian equilibrium, for all $i = 1,2,3$, $V_i(x_0) \geq 1$. Therefore, j *must* accept a proposal from x_i to x_0, hence a profitable deviation exists. Finally, it can be shown that any cyclical Markovian equilibrium must necessarily spend time at x_0. We have therefore proved:

Result 4. *Suppose that everyone proposes with equal probability at every date. Then every Markovian equilibrium exhibits dynamic inefficiency from every initial state.*

4 Strong Inefficiency in Four Player Games

Consider the following four player game:

$$\begin{array}{llllll} x_1: & \pi(x_1) & = & \{\{1,2\},\{3\},\{4\}\}, & \mathbf{u}(x_1) & = & (4,4,4,4) \\ x_2: & \pi(x_2) & = & \{\{1\},\{2\},\{3\},\{4\}\}, & \mathbf{u}(x_2) & = & (5,5,5,5) \\ x_3: & \pi(x_3) & = & \{\{1\{,\{2\},\{3,4\}\}, & \mathbf{u}(x_3) & = & (0,0,10,10) \\ x_4: & \pi(x_4) & = & \{\{1,2\},\{3,4\}\}, & \mathbf{u}(x_4) & = & (2,2,2,2) \end{array}$$

and recall from the main text

Observation 2. *For δ sufficiently high, every stationary Markov equilibrium is inefficient starting from any initial state.*

Before proceeding with the proof of Observation 2, we first allay any concerns about existence (or lack thereof) and provide an example displaying inefficiency. Figure 1 graphically depicts one such set of equilibrium transitions. In words, x_1 is absorbing; from state x_2,

Figure 1: An Inefficient Stationary Markov Equilibrium

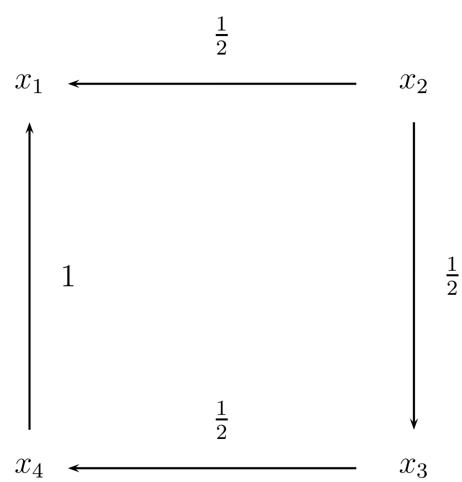

players 1 and 2 initiate a transition back to x_1, while players 3 and 4 initiate a transition to x_3; from x_3 players 1 and 2 implement a transition to x_4 when they propose, while players 3 and 4 leave the state unchanged when they propose; finally, from x_4, all players propose a return to x_1.

To verify that this description constitutes an equilibrium, begin with state x_1. Obviously players 3 and 4 do not benefit from changing the state to x_4, which is all they can unilaterally do. Players 1 and 2 can (bilaterally) change the state to x_2, by the presumed minimality of approval committees. If they do so, the subsequent trajectory will involves a stochastic path back to x_1.[5] Some fairly obvious but tedious algebra reveals that the Markov value function

[5] We are arguing in the spirit of the one-shot deviation principle, in which the putative equilibrium

$V_i(x, \delta)$ satisfies

$$V_i(x_2) = 5 - 3\delta + \frac{\delta^2(1+\delta)}{2\left(1 - \frac{\delta}{2}\right)}$$

for $i = 1, 2$. This value converges (as it must) to that of the absorbing state — 4 — as delta goes to 1, but the important point is that the convergence occurs "from below", which means that $V_i(x_2, \delta)$ is strictly smaller than $V_i(x_1, \delta) = 4$ for all delta close enough to 1.[6] Perhaps more intuitively but certainly less precisely, the move to state x_2 starts off a stochastic cycle through the payoffs 5, 0 and 2 before returning to absorption at 4, which is inferior to being at 4 throughout. This verifies that players 1 and 2 will relinquish the opportunity at x_1 to switch the state to x_2. It also proves that once at state x_2, players 1 and 2 will want to return to the safety of x_1 if they get a chance to move.

On the other hand, players 3 and 4 will want to move the state from x_2 to x_3. Proving this requires more value-function calculation. A second round of tedious algebra reveals that

$$V_i(x_3, \delta) - V_i(x_2, \delta) = 5 - 6\delta + \delta^2$$

for $i = 3, 4$. This difference vanishes (as it must) as δ approaches 1, but once again the important point is that the difference is strictly positive for all δ close to 1 (indeed, for all δ), which justifies the move of 3 and 4.

That 1 and 2 must want to move away as quickly as possible from state x_3, and 3 and 4 not at all, is self-evident. That leaves x_4. At this state players 3 and 4 receive their worst payoffs, and will surely want to move to x_1, and indeed, players 1 and 2 will want that as well.[7] Our verification is complete.

With the issue of existence of an inefficient equilibrium now dispelled, we now turn our attention to Observation 2, the proof of which proceeds in a number of steps.

Step 1: x_3 and x_4 are not absorbing.

It is easy to see that $V_i(x_4) \geq 2$ for $i = 1, 2$. Moreover, since players 1 and 2 can initiate a transition from x_3 to x_4, x_3 is easily seen to be not absorbing. Similarly, $V_j(x_1) \geq 4$ for $j = 3, 4$; therefore, since players 3 and 4 can achieve x_1 from x_4, x_1 is not absorbing.

Step 2: x_2 absorbing implies x_2 is globally absorbing.

---

strategies are subsequently followed. Even though the one-shot deviation principle needs to be applied with care when coalitions are involved, there are no such dangers here as all coalitional members have common payoffs.

[6] We verify this by differentiating $V_i(x_2, \delta)$ with respect to δ and evaluating the derivative at $\delta = 1$.

[7] Because we've developed the state space model at some degree of abstraction, we've allowed any player to make a proposal to any coalition, whether or not she is a member of that coalition. This is why players 1 and 2 ask 3 and 4 to move along. Nothing of qualitative import hinges on allowing or disallowing this feature. The transition from x_4 back to x_1 would still happen, but more slowly.

Suppose that x_2 is absorbing. Then clearly from x_1, players 1 and 2 would induce x_2. Moreover, since x_3 and x_4 are not absorbing, if x_2 is not reached, then x_1 must be reached infinitely often. But then 1 or 2 would get a chance to propose with probability 1 and would then take the state to x_2, a contradiction.

Step 3: x_2 cannot be globally absorbing.

If x_2 is globally absorbing then, from x_2, players 3 and 4 can get a payoff of 10 for some period of time, by initiating a transition to x_3, followed by, *at worst*, 2 for one period and 4 for another period, before returning to x_2, where it will get 5 forever thereafter.[8] This sequence of events is clearly better for players 3 and 4 than remaining at x_2.

Step 4: x_1 absorbing implies x_1 globally absorbing.

Steps 2 and 3 imply that x_2 cannot be absorbing. Moreover, Step 1 tells us that neither x_3 nor x_4 can be absorbing. In particular, from x_2 players 3 and 4 initiate a transition to x_3, from x_3 players 1 and 2 initiate a transition to x_4 and (at least) players 3 and 4 initiate a transition back to x_1. Therefore, if x_1 is absorbing, it is globally absorbing.

Step 5: Every equilibrium is inefficient.

First suppose that we had an equilibrium in which x_1 is not absorbing. Then from the above analysis, nothing is absorbing. Now consider x_2. If players 1 and 2 always accept an offer of a transition from x_1 to x_2, then 3 and 4 will strictly prefer to initiate a transition from x_2 to x_3: in so doing, they can achieve an average payoff of at least $\frac{10+2+4}{3} = \frac{16}{3} > 5$. However, it is easily seen that players 1 and 2 earn an average payoff strictly less than 4 in this case. Therefore, players 1 and 2 would rather keep the state at x_1, contradicting the presumption that x_1 was not absorbing.

The only remaining possibility is one in which players 1 and 2 are indifferent between a x_1 and x_2 and players 3 and 4 are indifferent between x_2 and x_3. If such an equilibrium were to exist, it must be that $V_i(x_1) = V_i(x_2) = 4$ for $i = 1, 2$, and $V_j(x_2) = V_j(x_3) = 5$ for $j = 3, 4$. Therefore, if such an equilibrium were to exist, it would also be inefficient since players spend a non-negligible amount of time at the inefficient states x_1 and x_4.

Thus either x_1 is the unique absorbing state or there is a sequence of inefficient cyclical equilibria depending on $\delta_n \nearrow 1$ such that players 1 and 2 are indifferent between x_1 and x_2 and players 3 and 4 are indifferent between x_2 and x_3. ||

[8] Surely, players 1 and 2 must initiate a transition to x_4 with some positive probability; otherwise, x_3 would be absorbing (which Step 1 shows to be impossible). However, once at x_4, under the assumption that any player can propose to move to any state, and the fact that (by Step 2) from x_1 there would be an immediate transition to x_2, there is no need to even pass through the intermediate state x_1.