On Asynchronicity of Moves and Coordination[*]

Attila Ambrus[†]
Duke University, Department of Economics

Yuhta Ishii[‡]
Yale University and **CIE, ITAM**

March 23, 2015

Abstract

This paper shows that asynchronicity of moves can lead to a unique prediction in coordination games, in an infinite-horizon setting, under certain conditions on off-equilibrium payoffs. In two-player games we derive necessary and sufficient conditions for play ultimately being absorbed in the Pareto dominant Nash equilibrium of the stage game, for every Markov perfect equilibrium. For players patient enough, the condition is that the Pareto dominant Nash equilibrium is also risk dominant, but for lower levels of patience the condition departs from simple risk-dominance. For general n-player symmetric games with patient players, we show that a necessary and sufficient condition for the Pareto dominant Nash equilibrium to be the unique limit outcome in all symmetric Markov perfect equilibrium is a particular generalization of risk-dominance for more than two players. We provide extensions to the unique selection results to all subgame perfect Nash equilibria, and to coordination games in which different players prefer different Nash equilibria of the stage game.

Keywords: repeated games, asynchronous moves, coordination, equilibrium selection
JEL Classification: C72, C73

1 Introduction

Coordination games have been used extensively to model many economic and social applications. Examples include bank runs (Diamond and Dybvig (1983)), adoption of a new technology in the presence of network extenalities (Katz and Shapiro (1985)), macroeconomic coordination failures (Cooper and John (1988)), self-fulfilling currency crises (Obstfeld (1996)) and riots and political change (Atkeson (2000)). One of the first papers to place coordination games in a dynamic context is Farrell and Saloner (1985). They show that in coordination problems such as when firm in an industry face a decision to switch to a more efficient technology standard, sequentiality of moves "solves" the coordination problem. In contrast, in a simultaneous-move game players can get stuck in an equilibrium staying with the old standard. In particular, they study a game in which each player make one publicly observed decision sequentially, in an exogeneously determined order. In such a game, standard backward induction arguments establish the existence of a unique subgame perfect Nash equilibrium (SPNE) in which players coordinate on the Pareto superior action. While this prediction is stark, the game in which it is established is clearly stylized. In this paper we revisit the question of whether asynchronicity of moves leads to efficient coordination in an infinite-horizon model in which players can switch back and forth between actions arbitrarily many times.[1]

In particular, we study a continuous-time model, in which a stage game is played continuously, generating flow payoffs. However, players can only change their actions at idyosychratic random times, and a player is locked into the previously selected action between such opportunities. This is one of the two types of asynchronous-move games commonly investigated in the literature, the other being discrete time alternating-move games. The two types of games are closely related,[2] and we adopted the continuous-time framework because for more than two players it is analytically more tractable and we also think that it is a more natural modeling assumption in most applications than assuming a fixed deterministic order of moves. Asynchronous-move models originated in the macroeconomic literatures of staggered prices, and of costly search: see Calvo (1983) and Burdett and Judd (1983). These models have generated considerable interest in game theory since the seminal works of Maskin and Tirole (1988a, 1988b), and more recently in empirical industrial organization (Arcidiacono et al. (2010), Doraszelski and Judd (2012)). The theoretical interest stems from the fact that predictions obtained from the study of these games are often very different than those from standard repeated games with simultaneous moves.[3] In empirical applications, as Arcidiacono et al. (2013) points out, the main advantage of asynchronous-move models is that their computational complexity

---

[1]It is natural to allow for agents to switch back and forth in many of the applications mentioned above. The adopter of a new technology, in an environment with network externalities, might want to switch back to the older technology if she realizes that the new technology will not attract a sufficiently large user base. Citizens during a revolution might go back and forth between being passive versus active, depending on their updated expectations regarding the likelihood of success of the movement.

[2]For example, Maskin and Tirole (1988a) points out the equivalence of the set of Markov perfect equilibria for 2-player games in the two settings.

[3]For implications in contrast with predictions of simultaneous-move dynamic games, see for example Maskin and Tirole (1988a, 1988b), Lagunoff and Matsui (1997), Takahashi and Wen (2003), and Ambrus et al. (2014).

is greatly reduced relative to standard models. This facilitates structural estimation and counterfactual policy-simulations in complicated dynamic models with a large number of state variables. Furthermore, as Schmidt-Dengler (2006) and Einav (2010) point out, asynchronicity of moves is a more realistic modeling assumption in many situations.

Formally, we investigate when it is the case that the Pareto dominant Nash equilibrium of the stage game is the unique limit outcome in the dynamic game, in the sense that in every equilibrium of the game, independently of the initial action profile, play gets absorbed in the above outcome in finite time, with probability one.[4] For most of the paper we restrict attention to Markov perfect equilibria (MPE), in the sense of investigating when the Pareto dominant Nash equilibrium is the unique limit outcome of the game across MPE. MPE are SPNE in which players play Markovian strategies that only depend on payoff-relevant states, which in our games are the action profiles other players are locked into. Besides tractability, our motivation for focusing on MPE comes from Bhaskar and Vega-Redondo (2002) and from Bhaskar, Mailath and Morris (2013), who show that natural refinement criteria select MPE in asynchronous-move games.[5] At the end of the paper we provide a partial extension of our results to address the question of when the Pareto dominant Nash equilibrium is the unique limit outcome in all SPNE.[6]

First we provide an exact characterization of when there is a unique limit outcome in MPE in 2×2 games. We show that for a full dimensional set of payoff specifications of the stage game, there is a range of discount rates for which this is the case. In particular, when the Pareto dominant Nash equilibrium of the stage game is risk-dominant, it is the unique limit outcome of the dynamic game in MPE for discount rates close to 0.[7] Moreover, the Pareto dominant Nash equilibrium of the stage game can also be the unique limit outcome of the dynamic game when it is risk-dominated, however in this case for an intermediate range of discount rates. All in all, our necessary and sufficient condition for unique limit outcome in MPE, which is on the payoffs of the game and the discount rate, is a generalization of risk-dominance, and it applies to arbitrary levels of patience (as opposed to most existing papers on emedding games into noisy dynamic environments, which focus on equilibrium selection for high levels of patience).

Next we consider symmetric n-player coordination games, with binary action sets. We show that for sufficiently high levels of patience, generically there is a unique symmetric MPE on pure strategies, in which

---

[4]As we discuss in subsection 6.1, this does not necessarily imply uniqueness of expected payoffs across all MPE, the latter typically requires stronger conditions.

[5]While formally the results in Bhaskar and Vega-Redondo (2002) and Bhaskar, Mailath and Morris (2013) are established for deterministic alternating-move games, if one extends to notion of bounded memory to our setting along the lines of players remembering the last K action choices in the game, for some $K \geq 1$, then their proofs go through without any modification in our setting, too.

[6]As we show, the answers to these questions in our infinite-horizon setting are quite nuanced. In contrast, much stronger unique selection results hold generically in finite horizon asynchronous move games, as shown for example in Dutta (2012) and CKLS (2014).

[7]In models like ours, with random opportunities to take action corresponding arrival times of a Poisson process, only the ratio $\frac{\lambda}{r}$ matters, which is related to the neutrality of the model in what unit is used to measure time. For this reason, any comparative statics on r with a fixed λ exactly correspond to comparative statics on λ with a fixed r. For example, any result that holds for r close enough to 0 (for a fixed λ) also holds for high enough λ (for a fixed r).

in every Markovian state either the action corresponding to the Pareto dominant Nash equilibrium (α) or in every state the other action (β) is played. The condition that determines which Nash equilibrium of the stage game is selected is a generalization of risk-dominance for symmetric n-player games: α is the unique absorbing state of the unique symmetric pure strategy MPE iff the sum of payoff differences in the stage game for playing α versus β, when i other players is playing (i being summed from 0 to $n-1$, the latter corresponding to all other players playing α) is strictly positive. Conversely, when the above sum is strictly negative, β is selected to be the unique limit outcome in pure strategy symmetric MPE. Using techniques from the theory of stationary Markov chains,[8] we show that when playing α in every state is the unique symmetric pure strategy MPE then α is also the unique absorbing state in every symmetric MPE in mixed strategies, and therefore it is the unique limit outcome across all symmetric MPE. In contrast, in the region where playing β in every state is the unique symmetric pure strategy MPE, there are always symmetric mixed MPE with α being an absorbing state. Hence in this region there cannot be a unique limit outcome in symmetric MPE. Finally we conclude the analysis of symmetric n-player coordination games by showing that whenever parameters are such that α is the unique limit outcome in all symmetric MPE, then payoffs at all states approximate the Pareto dominant payoff in all equilibria when the players are sufficiently patient.

To summarize the above results, we find that asynchronicity of moves can solve coordination problems in a dynamic setting, even with infinite horizon and no restrictions on switching actions back and forth, but only if certain conditions hold for off-equilibrium payoffs. For patient players these conditions can be simply summarized through risk-dominance (a simple generalization of risk dominance, in the case of more than two players).

We also show that our results can be extended in various directions. First, there also exists a full dimensional set of parameter specifications with unique limit outcome over all SPNE (not only all MPE). As expected, this holds for a strictly smaller set of specifications than uniqueness of limit outcome for MPE. Moreover, for generic payoff specifications it can only hold for intermediate levels of patience (relative to the frequency of being able to change actions). Intuitively, when players are very patient, then in infinite-horizon games such as the ones we consider, some form of the folk theorem applies, implying a severe multiplicity of equilibria. On the other hand, if players are very impatient then either of the Nash equilibria of the stage game can become absorbing, as players are not willing to sacrifice current payoffs for future payoff improvements. Our result here complements the findings of Lagunoff and Matsui (1997), who show that in an asynchronous-move repeated game framework the Pareto-dominant equilibrium gets selected in a class of pure coordination games for high enough discount factors. However, this result only holds for a nongeneric class of coordination games, as in generic asynchronous-move games versions of the folk theorem were shown to hold by Dutta (1995) and Yoon (2001).[9] Our unique limit outcome result does not get into friction with

---

[8]For an introduction, see Chapter 4 of Gallager (2013) or Chapters 1 and 2 of Karlin and Taylor (1975).

[9]Lagunoff and Matsui (2001) point out that the genericity or non-genericity of anti-folk theorems depends on the order limits

the folk theorem, because it applies for intermediate discount rates.

We also extend some of our results to Battle of the Sexes (BoS) type games, in which players disagree which of two stage game Nash equilibria to coordinate on. In such games the conditions for unique limit outcome are stronger, and even in MPE it can only happen for intermediate levels of patience.

2 Related Literature

Besides Lagunoff and Matsui (1997), discussed above, Takahashi (2005) investigated the possibility unique equilibrium selection in a class of asynchronous-move common interest games, encompassing both finite and infinite horizons. Unique limit outcome in MPE, which is the main focus of our work, is not investigated, instead the paper provides a sufficient condition for unique limit outcome in SPNE. In the class of games that are in the intersection of the two papers, the condition we provide in our extension to SPNE is independent of the condition provided by Takahashi (2005). Our focus on MPE allows us to provide not only a sufficient, but a necessary and sufficient condition for unique limit outcome (in all MPE for 2-player games and in all symmetric MPE for symmetric n-player games).[10]

There are also several lines of literature with a similar motivation than our paper, but adopting modeling assumptions that depart from ours in important ways. Dynamic models of technology adoption (see for example Katz and Shapiro (1986), Farrell and Klemperer (1997), Young (2009) and Frick and Ishii (2014)) assume that the decision to switch from an existing technology to a new one is irreversible, as opposed to the repeated game type framework we investigate. The recent literature on asynchronous move games with a finite horizon (see Takahashi (2005), Dutta (2012) and Calcagno, Kamada, Lovo and Sugaya (2014) (CKLS)) obtains strong unique limit outcome selection results whenever the stage game has a Pareto dominant equilibrium.[11] In contrast in infinite-horizon games like the ones we investigate, we show that stronger conditions are necessary for the selection of a unique limit outcome even in MPE. Lastly, papers on games with strategic complementarities embedded into a dynamic context, such as Matsui and Matsuyama (1995), Frankel and Pauzner (2000), and Burdzy et al. (2001), typically assume a continuum of agents, implying that a player's action choice does not have any influence on future action choices of others.[12] Moreover, in Frankel and Pauzner (2000) and Burdzy et al. (2001) players observe a payoff-relevant state, which evolves according to an exogenously specified stochastic process, perfectly. With positive probability, the stochastic process eventually pushes the state to a region in which a given action becomes strictly dominant, facilitating global games type arguments for unique equilibrium selection. In contrast, as both Frankel and Pauzner

---

are taken, as for any discount factor close to 1, there is a full-dimensional set of coordination-games for which all spne payoffs have to be very close to the payoffs in the Pareto-dominant equilibrium of the stage game.

[10]Haller and Lagunoff (2000) investigate MPE of asynchronous-move games, but they do not address the question of unique limit outcome.

[11]As Dutta (2012) shows, when the order of moves is determinisitic then SPNE is generically unique with finite horizon even without assuming a Pareto dominant equilibrium.

[12]See also earlier related papers on settings with noisy evolution and learning, as in Kandori et al. (1993) or Young (1993).

(2000) and Burdzy et al. (2001) point out, an environment in which the state begins and remains forever (with probability one) in a region where no actions are strictly dominant necessarily leads to a multiplicity of equilibria in these games. Despite all of these differences, the conditions we obtain for unique limit outcome in MPE of coordination games, as the discount rate is taken to 0, are similar.[13] The reason is that players' incentives in the two types of models become similar at the respective limits, essentially boiling down to players preferring the risk-dominant action. Away from the limit typically players in our model have more complex strategic incentives than in the above models, even when we restrict attention to Markovian strategies. This relates to the feature of our model that players are non-negligible, and their current choices can influence subsequent choices of the other players (in contrast with the above papers which feature a continuum of agents).

3   The model

Let $G = (N, A, u)$ be an n-player normal form game, with the action set of each player i being $A_i = \{\alpha_i, \beta_i\}$, with the restriction that there exist exactly two strict Nash equilibria of the game: $\alpha = (\alpha_1, ..., \alpha_n)$ and $\beta = (\beta_1, ..., \beta_n)$, and $u_i(\alpha) > u_i(\beta)$ for every $i \in N$ We will refer to G as the stage game. Our restrictions imply that G is a pure coordination game, with α as the Pareto dominant Nash equilibrium.

We consider a continuous-time dynamic game with asynchronous moves, in which players play the above stage game indefinitely, and accumulate payoffs continuously, but they can only change their actions at random discrete times that correspond to realizations of random arrival processes. Formally, players are associated with independent Poisson arrival processes with arrival rate λ. A player can only change her action in the game at the realizations of her arrival process, hence in between two arrival events she is locked into the action chosen at the former one. Let a^t denote the action profile played at time t. Player i's (normalized) payoff in the dynamic game is then defined as $r \int_0^\infty e^{-r} u_i(a^t) dt$, where $u_i(\cdot)$ stands for player i's payoff in the stage game and r is the discount rate (assumed to be the same for the two players, for simplicity). We assume that the action profile at time 0, a^0, is given exogenously.

Note that the independence of the arrival processes implies that it is a 0 probability event that there is a point of time when both players can change their actions. We assume that the realizations of the arrival processes are regular, and such simultaneous arrivals indeed never happen. Hence players, when they choose actions, know what action the opponent will be locked in for a random amount of time.

We assume that arrivals are publicly observed, hence a publicly observed history of the game at time t, labeled as h^t, consists of the path of the action profile on $[0, t)$ and the set of realizations of each players' arrival processes on $[0, t)$. Let H^t be the set of such h^t. Strategies of player i are measurable mappings from

---

[13]More precisely, in our games, the Pareto dominant action profile is selected as a unique limit outcome when it is also risk dominant. Loosely speaking, the above papers find that the risk dominant equilibrium (not the Pareto dominant action profile) is selected as a unique limit outcome. Thus in both approaches, risk dominance is the key condition that leads to unique selection results.

H^t to $\{\alpha_i, \beta_i\}$, with the value at h^t interpreted as the action that the given player would choose at time t if she had an arrival at that time and the history at that time was h^t.[14]

In most of the analysis below, we focus on the set of MPE, that is SPNE in which any player's action choices only depend on the payoff relevant state at the time of the action choice, which is simply the action profile currently played by the opposing players.

4   Two-player Coordination Games

First we characterize the set of MPE for all possible payoff configurations and discount rates. This in particular yields a necessary and sufficient condition for all MPE of the game (irrespective of the initial action profile) having the property that play is eventually absorbed in α.[15] For ease of exposition, we refer to this property as having a unique limit outcome in MPE. The condition is then also a necessary condition for having a unique limit outcome in SPNE.

If strategies are Markovian, at times when there is no arrival event, each player can only have four continuation values, corresponding to the four possible action profiles in the stage game. Let $v_i(.,.)$ denote player i's continuation payoff as a function of the action profile.

We begin by establishing a lemma on the relationship between continuation payoffs at different states in any MPE, which will be useful in narrowing the possible candidate strategy profiles in MPE.

Lemma 4.1. *Suppose that $v_i(\alpha_i, \beta_{-i}) \geq v_i(\beta)$. Then in any Markovian equilibrium, we must have $v_i(\alpha) > v_i(\beta_i, \alpha_{-i})$. Similarly if $v_i(\beta_i, \alpha_{-i}) \geq v_i(\alpha)$ then we must have $v_i(\beta) > v_i(\alpha_i, \beta_{-i})$.*

Proof. Assume that $v_i(\alpha_i, \beta_{-i}) \geq v_i(\beta)$. Suppose by way of contradiction that $v_i(\beta_i, \alpha_{-i}) \geq v_i(\alpha)$. Note the following:

$$
\begin{aligned}
v_i(\alpha) &\geq \frac{r}{2\lambda + r}u_i(\alpha) + \frac{2\lambda}{2\lambda + r}\left(\frac{1}{2}v_i(\alpha) + \frac{1}{2}\left(pv_i(\alpha) + (1-p)v_i(\alpha_i, \beta_{-i})\right)\right) \\
v_i(\alpha_i, \beta_{-i}) &= \frac{r}{2\lambda + r}u_i(\alpha_i, \beta_{-i}) + \frac{2\lambda}{2\lambda + r}\left(\frac{1}{2}v_i(\alpha_i, \beta_{-i}) + \frac{1}{2}\left(pv_i(\alpha) + (1-p)v_i(\alpha_i, \beta_{-i})\right)\right)
\end{aligned}
$$

for some $p \in [0, 1]$. The first expression is an inequality since if player i obtains an arrival at state α, he has an incentive to play β_i. These expressions then imply that

$$
v_i(\alpha) - v_i(\alpha_i, \beta_{-i}) \geq \frac{r}{\lambda + r}(u_i(\alpha) - u_i(\alpha_i, \beta_{-i})).
$$

Next we can write the value functions when the currently played action of player i is β_i:

$$
\begin{aligned}
v_i(\beta_i, \alpha_{-i}) &= \frac{r}{2\lambda + r}u_i(\beta_i, \alpha_{-i}) + \frac{2\lambda}{2\lambda + r}\left(\frac{1}{2}v_i(\beta_i, \alpha_{-i}) + \frac{1}{2}\left(qv_i(\beta_i, \alpha_{-i}) + (1-q)v_i(\beta)\right)\right) \\
v_i(\beta) &\geq \frac{r}{2\lambda + r}u_i(\beta) + \frac{2\lambda}{2\lambda + r}\left(\frac{1}{2}v_i(\beta) + \frac{1}{2}\left(qv_i(\beta_i, \alpha_{-i}) + (1-q)v_i(\beta)\right)\right).
\end{aligned}
$$

---

[14]This simple definition of strategies is standard in the literature, but implicitly assumes measurability of H^t. In the Supplementary Appendix we provide details on how to construct an appropriate sigma algebra on H^t and define strategies rigorously.

[15]As we show below, there is no set of parameters for which play in every MPE is eventually absorbed in β.

This then implies

$$v_i(\beta) - v_i(\beta_i, \alpha_{-i}) \geq \frac{r}{\lambda + r}(u_i(\beta) - u_i(\beta_i, \alpha_{-i})).$$

Then this means that

$$
\begin{aligned}
v_i(\alpha) - v_i(\beta_i, \alpha_{-i}) &= (v_i(\alpha) - v_i(\alpha_i, \beta_{-i})) + (v_i(\alpha_i, \beta_{-i}) - v_i(\beta)) + (v_i(\beta) - v_i(\beta_i, \alpha_{-i})) \\
&\geq \frac{r}{\lambda + r}(u_1(\alpha) - u_1(\alpha_1, \beta_2)) + \frac{r}{\lambda + r}(u_1(\beta) - u_1(\beta_1, \alpha_2)) \\
&= \frac{r}{\lambda + r}(u_i(\alpha) - u_i(\beta_i, \alpha_{-i}) + u_i(\beta) - u_i(\alpha_i, \beta_{-i})) > 0.
\end{aligned}
$$

This is a contradiction. Therefore we must have $v_i(\alpha) > v_i(\beta_i, \alpha_{-i})$. $\qquad\square$

The previous lemma is useful for narrowing the possible candidate strategy profiles in MPE.

Lemma 4.2. *Neither (β_1, α_2) nor (α_1, β_2) can be part of any absorbing cycle in any MPE. This in particular implies that in coordination games, either (α) or (β) (or both) must be absorbing states in any MPE.*

Proof. By symmetry it suffices to prove that (β_1, α_2) cannot be part of an absorbing cycle. For (β_1, α_2) to be part of an absorbing cycle, we need $v_i(\beta_1, \alpha_2) \geq v_i(\alpha)$ for both players $i = 1, 2$. Otherwise both players would prefer to play α_i whenever the opponent's currently played action is α_{-i} making (α) an absorbing state. This however would make it impossible for (β_1, α_2) to be part of an absorbing cycle.

Using Lemma 4.1, this means that $v_1(\beta) > v_1(\alpha_1, \beta_2)$ and $v_2(\beta) > v_2(\beta_1, \alpha_2)$. But then this makes β an absorbing state, making it impossible for (β_1, α_2) to be part of an absorbing cycle. $\qquad\square$

Our main interest is in the characterization of the set of parameter values for which (α) is the unique absorbing state in every MPE. We do so by characterizing all parameter values in which there is an MPE with an absorbing cycle containing (β), and then taking the complement of this set. Lemma 4.2 implies that the latter is indeed the set of parameter values in which (α) is the unique absorbing state. Lemma 4.2 leaves five possible kinds of dynamics in which (β) is an absorbing state:
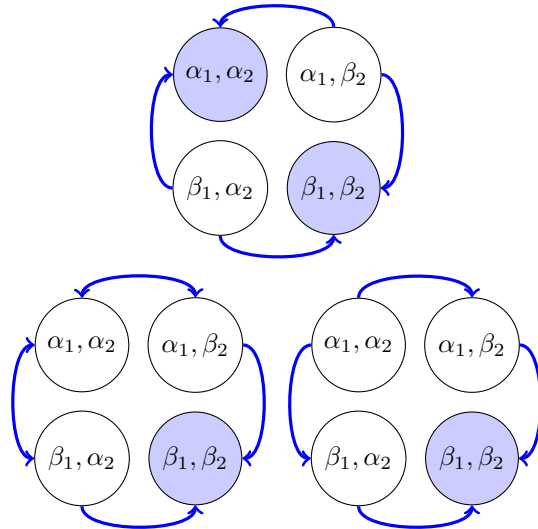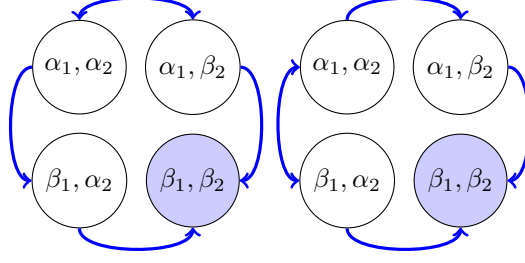
The last two dynamics are nongeneric (since only one player is mixing), and it is straightforward to show that if for some parameter values such dynamics are possible in MPE, there is also an MPE with the second or third type of dynamics. Therefore we examine necessary and sufficient conditions for the first three types of dynamics to be possible in an MPE. Since the proof of the following lemma involves a tedious case by case analysis, it is relegated to the Supplementary Appendix.

Lemma 4.3. *The necessary and sufficient conditions for an equilibrium with (β) to be an absorbing state is either*

$$\frac{r}{\lambda+r}u_i(\alpha_i,\beta_{-i}) + \frac{\lambda}{\lambda+r}u_1(\alpha) \leq u_i(\beta) \ \ \forall \ i = 1,2$$

or

$$u_i(\alpha) + \frac{\lambda}{\lambda+r}u_i(\alpha_i,\beta_{-i}) \leq u_i(\beta_i,\alpha_{-i}) + \frac{\lambda}{\lambda+r}u_i(\beta) \ \ \forall \ i = 1,2.$$

The previous results directly characterize our main object of interest, the set of parameter values for which (α) is a global absorbing state.

Proposition 4.4. *In a 2-player coordination game, α is the unique limit outcome in MPE iff:*

(i) $\frac{r}{\lambda+r}u_i(\alpha_i,\beta_{-i}) + \frac{\lambda}{\lambda+r}u_i(\alpha) > u_i(\beta)$ for some i;

and

(ii) $u_i(\beta_i,\alpha_{-i}) + \frac{\lambda}{\lambda+r}u_i(\beta) < u_i(\alpha) + \frac{\lambda}{\lambda+r}u_i(\alpha_i,\beta_{-i})$ for some i.

Proof. Lemma 4.2 implies that the set of parameter values with α being the unique limit outcome is the complement of the set of parameter values with β being an absorbing state. The result then follows from Lemma 4.3. □

The two conditions in the above proposition are rather intuitive. The first one implies that at least one of the players has a strict incentive to move out of (β) if she expects the other player to follow this move by also switching from β to α. The second condition ensures that at least one of the players has a strict incentive to stay in (α) for the time being, even if she expects the other player to switch to β forever at the first opportunity. This eliminates MPE in which play moves out of (α) because players expect each other to switch to β and both would prefer to switch first rather than second. It is straightforward that the above

8

conditions are necessary for α to be the unique limit outcome in MPE. The proof above establishes that they are also sufficient.

We next investigate the question of when it is the case that for a certain payoff specification of the game there exist discount rates such that α is the unique limit outcome in MPE. Note that the condition in Proposition 4.4 can be rewritten as:

$$\max_{i=1,2} \left\{ \frac{u_i(\alpha) - u_i(\beta_i, \alpha_{-i})}{u_i(\beta) - u_i(\alpha_i, \beta_{-i})} \right\} > \frac{\lambda}{\lambda + r} > \min_{i=1,2} \left\{ \frac{u_i(\beta) - u_i(\alpha_i, \beta_{-i})}{u_i(\alpha) - u_i(\alpha_i, \beta_{-i})} \right\} \tag{1}$$

In the limit as $r \to 0$, the above condition simplifies to α being risk-dominant. This implies that in two-player coordination games in which α is risk-dominant, there is an interval of r around 0 such that α is the unique limit outcome in MPE.

Generally though condition (1) does not require payoff parameters for which α is risk-dominant. Consider the following example, with $\varepsilon > 0$:

	α_2	β_2
α_1	$1, 1$	$-(1+\varepsilon), 0$
β_1	$0, -(1+\varepsilon)$	$0, 0$

Note that

$$\min_{i=1,2} \left\{ \frac{u_i(\alpha) - u_i(\beta_i, \alpha_{-i})}{u_i(\beta) - u_i(\alpha_i, \beta_{-i})} \right\} = \frac{1}{1+\varepsilon}$$

Furthermore

$$\min_{i=1,2} \left\{ \frac{u_i(\beta) - u_i(\alpha_i, \beta_{-i})}{u_i(\alpha) - u_i(\alpha_i, \beta_{-i})} \right\} = \frac{1+\varepsilon}{2+\varepsilon}.$$

Therefore for $\varepsilon > 0$ sufficiently small, there exists an open interval of discount rates for which α is the unique absorbing state in all MPE, even though β is risk dominant in the above example.

In order to provide an explicit characterization for the set of payoffs for which there exists some $r \in (0,1)$ with α being the unique absorbing state in all MPE, we define the following expression:

$$p_i \equiv \frac{u_i(\beta) - u_i(\alpha_i, \beta_{-i})}{u_i(\alpha) - u_i(\alpha_i, \beta_{-i})}.$$

Intuitively if p_i is very small, this means that the payoff advantage of $u_i(\alpha)$ over $u_i(\alpha_i, \beta_{-i})$ is large relative to payoff advantage of $u_i(\beta)$ over $u_i(\alpha_i, \beta_{-i})$. Thus a small p_i means that player i is more willing to move out of β with a short term cost of $u_i(\beta) - u_i(\alpha_i, \beta_{-i})$ followed by a large payoff gain upon moving to $u_i(\alpha)$.

Let $p = \min\{p_1, p_2\}$. This represents the maximum willingness to move out of the β equilibrium. The proposition below shows that a necessary and sufficient condition for all MPE to have a unique absorbing state at α is if α is $\frac{1}{1+p}$-dominant for at least one of the players. Note that small p makes the condition more easily satisfied (1-dominance is simply the definition of Nash equilibrium).

Proposition 4.5. *There exists an open parameter of discount rates in which α is the unique limit outcome in all MPE if and only if*

$$\max_{i=1,2} \left\{ \frac{u_i(\alpha) - u_i(\beta_i, \alpha_{-i})}{u_i(\beta) - u_i(\alpha_i, \beta_{-i})} \right\} > p.$$

Moreover the above inequality holds if and only if for some player i α is $1/(1+p)$-dominant.

Proof. Note that from Inequality (1), a necessary and sufficient condition for α to be the unique limit outcome in all MPE is:

$$\max_{i=1,2}\left\{\frac{u_i(\alpha)-u_i(\beta_i,\alpha_{-i})}{u_i(\beta)-u_i(\alpha_i,\beta_{-i})}\right\} > \frac{\lambda}{\lambda+r} > p. \tag{2}$$

Because $p < 1$, there exists some open parameter of discount rates in which α is the unique limit outcome in all MPE if and only if

$$\max_{i=1,2}\left\{\frac{u_i(\alpha)-u_i(\beta_i,\alpha_{-i})}{u_i(\beta)-u_i(\alpha_i,\beta_{-i})}\right\} > p.$$

Then note that the above inequality is satisfied if and only if there exists $i = 1,2$ such that

$$\frac{u_i(\alpha)-u_i(\beta_i,\alpha_{-i})}{u_i(\beta)-u_i(\alpha_i,\beta_{-i})} > p \Leftrightarrow u_i(\alpha)+pu_i(\alpha_i,\beta_{-i}) > u_i(\beta_i,\alpha_{-i})+pu_i(\beta)$$

$$\Leftrightarrow \frac{1}{1+p}u_i(\alpha)+\frac{p}{1+p}u_i(\alpha_i,\beta_{-i}) > \frac{1}{1+p}u_i(\beta_i,\alpha_{-i})+\frac{p}{1+p}u_i(\beta).$$

The last inequality corresponds exactly to α being $1/(1+p)$ dominant for player i, proving the claim. $\square$

5 Symmetric n-Player Coordination Games

In this section we consider symmetric MPE of symmetric n-player coordination games. In particular, we investigate when α is the unique limit outcome in symmetric MPE of such games, focusing on cases with patient players. If players are impatient enough, both α and β are absorbing states in all MPE, hence there cannot be a unique limit outcome. We focus on high versus intermediate levels of patience in n-player games mainly for tractability.

Note that symmetric strategies can be expressed concisely as a map : $p : K \equiv \{0,1,\ldots,n-1\} \to \Delta(\{\alpha,\beta\})$, where p_k is interpreted as the probability with which a player, who arrives facing k opponents (not including himself) playing α, plays α. Given any stationary symmetric MPE p, we can construct a Markov chain with transition matrix A on the state space $\{0,1,\ldots,n\}$ (representing the number of players currently playing α) in the obvious way:

$$A_{k,k+1} = \frac{n-k}{n}p_k$$

$$A_{k,k-1} = \frac{k}{n}(1-p_{k-1})$$

$$A_{k,k} = 1 - \frac{n-k}{n}p_k - \frac{k}{n}(1-p_{k-1}).$$

We now study properties of the Markov transition matrix A induced by an equilibrium p. Throughout this section we use terminology standard in the theory of stationary Markov chains. For readers unfamiliar with this terminology, we provide definitions in the Supplementary Appendix.

5.1 Characterization of the Set of Symmetric MPE

First we show that a symmetric MPE with α as a limit outcome exists regardless of the parameters.

Claim 5.1. *There exists a symmetric MPE (possibly mixed) in which α is an absorbing state.*

Proof. Define G^* to be a restricted game obtained from the original game by restricting each player's action set such that at any history at which $n-1$ other players are playing α, the action choice is restricted to be α. Define Σ to be the set of stationary Markovian behavioral strategies in G^*, that is the set of mappings from $\{0, ..., n-1\}$ to $\Delta(\{\alpha, \beta\})$ with the restriction that α is allocated probability 1 at $n-1$). Next, for any $\sigma \in \Sigma$ and $k \in \{0, ..., n-1\}$, define $BR_k(\sigma)$ to be the set of best response (mixed) action choices for a player i at a history at which k other players are playing α, given that all players other than i plays the stationary Markovian strategy σ, and player i herself plays according to σ at every history in which the number of other players playing α is not k. Let $BR(\sigma) = \underset{k \in \{0,...,n-1\}}{\times} BR_k(\sigma)$. It is easy to verify that all the conditions for Kakutani's fixed point theorem hold for BR, therefore it has a fixed point σ^*. Consider now the symmetric stationary Markovian strategy profile ξ in G^* in which every player plays σ^*. Because σ^* is a fixed point of BR, every player at every Markovian state best responds to other players playing σ^*, implying that ξ is a symmetric MPE of G^*. Furthermore, it also constitutes an MPE in G^*, since given ξ_{-1} playing α is a best response at Markovian state $n-1$, as it leads to the maximal feasible payoff in the game. $\square$

Our goal is to study properties of the set of equilibria when players are sufficiently patient. Towards this end, we establish the following fact which narrows the set of possible limit outcomes.

Claim 5.2. *There exists $r^* > 0$ such that for all $r < r^*$ and any symmetric stationary MPE and A its induced Markov chain, either*

1. *$\{\alpha\}$ is the unique recurrent class of A, or*

2. *$\{\beta\}$ is the unique recurrent class of A.*

Proof. See Appendix A for the proof. $\square$

The proof in the Appendix first establishes that any recurrent class in a symmetric MPE has to consist of a single state, otherwise players would have an incentive to deviate at least at one boundary state of the recurrent class. Then it is shown that such a state cannot be interior. The most involved part of the proof, in which we use techniques from the theory of stationary Markov chains, establishes that for low enough discount rates there cannot exist symmetric MPE with both α and β absorbing states. The intuition here is that very patient players should mostly care about maximizing the probability of ending up in state α versus β. Therefore in any candidate profile in which both α and β are absorbing, we show the existence of a state in which β is played with positive probability despite players having a strict incentive to

11

play α, leading to a contradiction. Note however that the possibility of equilibrium strategies with mixing probabilities potentially approaching degenerate distributions and hence slowing down transition necessitates subtle arguments here, for which we exploit certain properties of limits of matrices of transition probabilities in MPE when the discount rate is taken to 0.

Our ultimate goal is to characterize necessary and sufficient conditions of payoffs in the stage game for which, when players are sufficiently patient, $\{\alpha\}$ is the unique recurrent class in all equilibria. To understand these necessary and sufficient conditions, Claims 1 and 5.2 imply that it is sufficient to study exactly when equilibria of type 2 above exist when players are sufficiently patient. To this end, the following claim establishes a key inequality that must be satisfied in any equilibrium in which $\{\beta\}$ is the unique recurrent class of A. Before we proceed to the claim, we establish the following conventions for the remainder of this section: Define $\max \emptyset = -1$ and $v_{-1}^{\alpha} = v_{-1}^{\beta} = 0$.

Claim 5.3. *Suppose that an equilibrium q exists with $q_k < 1$ for all $k = 0, 1, \ldots, n-1$ and $q_0 = 0$. Define $\underline{k} = \max\{k < n-1 : q_k \in (0,1)\}$. Then for all $k \geq \underline{k}$,*

$$\frac{v_{k+1}^{\alpha} - v_{k+1}^{\beta}}{r} \geq \frac{1}{r + \lambda(k+2)} \left(u_{k+1}^{\alpha} - u_{k+1}^{\beta} \right) + \frac{\lambda(k+1)}{r + \lambda(k+2)} \frac{v_k^{\alpha} - v_k^{\beta}}{r}.$$

Proof. See Appendix Section A.3 for the proof. $\qquad\square$

To derive this inequality, first we observe that in a symmetric MPE with β as a unique absorbing state, expected continuation payoffs are monotonic in the state (the number of other players currently playing β). We can use this fact to provide an upper bound on v_k^{β} in terms of u_k^{β} and v_{k-1}^{β}, and a lower bound on v_k^{α} in terms of u_k^{α} and v_{k-1}^{α}. The statement in the claim is implied by these inequalities.

We are now ready to state the main proposition, which provides a clean characterization of the set of payoff parameters under which α is chosen as a unique limit outcome in all symmetric MPE for patient players.

Proposition 5.4. *If*

$$\sum_{i=0}^{n-1} u_i^{\alpha} > \sum_{i=0}^{n-1} u_i^{\beta}, \tag{3}$$

then there exists some r^* such that whenever $r < r^*$, α is the unique limit outcome in all symmetric MPE; on the other hand if

$$\sum_{i=0}^{n-1} u_i^{\alpha} < \sum_{i=0}^{n-1} u_i^{\beta},$$

then there exists r^* such that whenever $r < r^*$, a symmetric MPE with β as an absorbing state exists.[16]

Proof. We first prove the first part of Proposition 5.4: if Inequality (3) holds, there exists some $r^* > 0$ such that α is the unique limit outcome in all symmetric MPE. Suppose otherwise. By Claim 5.2, there exists

---

[16]We ignore the case in which $\sum_{i=0}^{n-1} (u_i^{\alpha} - u_i^{\beta}) = 0$ since it is non-generic.

some sequence of discount rates $r_\ell \to 0$ with corresponding equilibria p^ℓ such that p^ℓ has β as a unique absorbing state for all ℓ. Without loss of generality we can take this sequence to be such that p^ℓ converges (by Heine-Borel). Let us denote the corresponding value functions as $v_k^{\alpha,\ell}$ and $v_k^{\beta,\ell}$ for all ℓ and k.

For every ℓ, let $\underline{k}_\ell = \max\{k < n - 1 : p_k^\ell \in (0,1)\}$.[17] Because $\{0, 1, 2, \ldots, n - 2\}$ is finite, again without loss of generality, we can assume the complete sequence to be such that $\underline{k}_\ell = \underline{k}$ for some $\underline{k}$ for all ℓ (otherwise, take an appropriate subsequence). By Claim 5.3, for all $k \geq \underline{k}$ and ℓ,

$$\frac{v_{k+1}^{\alpha,\ell} - v_{k+1}^{\beta,\ell}}{r_\ell} \geq \frac{1}{r_\ell + \lambda(k+2)} \left(u_{k+1}^\alpha - u_{k+1}^\beta \right) + \frac{\lambda(k+1)}{r_\ell + \lambda(k+2)} \frac{v_k^{\alpha,\ell} - v_k^{\beta,\ell}}{r_\ell}.$$

To simplify notation, define

$$\Delta_k := \liminf_{\ell \to \infty} \frac{v_k^{\alpha,\ell} - v_k^{\beta,\ell}}{r_\ell}.$$

Using this notation, the above inequality implies

$$\Delta_{k+1} \geq \frac{1}{\lambda(k+2)} \left(u_{k+1}^\alpha - u_{k+1}^\beta \right) + \frac{k+1}{k+2} \Delta_k$$

for all $k \geq \underline{k}$. Furthermore by the definition of $\underline{k}$, $\Delta_{\underline{k}} = 0$.

Then the above set of inequalities implies that

$$\Delta_{n-1} \geq \frac{1}{\lambda n} \sum_{j=\underline{k}+1}^{n-1} (u_j^\alpha - u_j^\beta) > 0.$$

But this implies that there exists some ℓ^* such that for all $\ell > \ell^*$, $\frac{v_{n-1}^{\alpha,\ell} - v_{n-1}^{\beta,\ell}}{r_\ell} > 0$, which means that for all $\ell > \ell^*$, $p_{n-1}^\ell = 1$, a contradiction.

To prove the second part of the proposition, we show more strongly that whenever $\sum_{k=0}^{n-1}(u_k^\alpha - u_k^\beta) < 0$, the pure strategy in which all players play β with probability one at all states is indeed an equilibrium for sufficiently patient players. The proof is a straightforward computation whose details are contained in Appendix Section A.4. □

Note that Claim 5.2 implies that for low enough discount rates, the only candidate symmetric *pure* strategy MPE are the profiles in which either every player plays α at all Markovian states or every player plays β at all Markovian states. It can be seen easily that for high enough discount factors, such symmetric pure strategy equilibria indeed exist, and it is generically unique within the class of symmetric pure MPE: when $\sum_{i=0}^{n-1} u_i^\alpha > \sum_{i=0}^{n-1} u_i^\beta$ it involves everyone playing α, while in the region where $\sum_{i=0}^{n-1} u_i^\alpha < \sum_{i=0}^{n-1} u_i^\beta$ it involves everyone playing β. Furthermore in the former parameter region, Proposition 5.4 also shows that α is the unique absorbing state in all other (mixed) symmetric MPE when players are sufficiently patient. In contrast, in the latter parameter region, the above discussion and Claim 5.1 together imply the multiplicity of possible limit outcomes in symmetric (mixed) MPE.

---

[17]Recall that we let $\max \emptyset = -1$ and defined $v_{-1}^{\alpha,\ell} = v_{-1}^{\beta,\ell} = 0$ for all ℓ.

5.2 Payoffs

Note that the previous section establishes α as the unique limit outcome in all symmetric MPE when players are patient and the generalized risk dominance condition holds. This however does not necessarily imply that payoffs at all states in all symmetric MPE must approximate u_{n-1}^{α} when players are sufficiently patient. This is because equilibria may exist in which the transition to α from certain states becomes arbitrarily slow (relative to the discount rate) even as players become patient.[18] The proposition to follow rules out such possibilities.

Fix $\lambda = 1$ and define the set $E(r, u)$ as the set of payoff vectors $v = \{v_k^{\alpha}, v_k^{\beta}\}_{k=0}^{n-1}$ that are the payoffs at the corresponding states of some symmetric MPE of the game with discount rate r and stage game payoffs u.

Proposition 5.5. *Suppose that*

$$\sum_{i=0}^{n-1} u_i^{\alpha} > \sum_{i=0}^{n-1} u_i^{\beta}.$$

Then for every $\varepsilon > 0$, there exists some $r^* > 0$ such that for all $r < r^*$, $\max\left\{|v_k^{\alpha} - u_{n-1}^{\alpha}|, |v_k^{\beta} - u_{n-1}^{\alpha}|\right\} < \varepsilon$ for all $k = 0, 1, \ldots, n-1$ for all $v \in E(r, u)$.

Proof. This is implied directly by Proposition 5.4 and Proposition A.1 in the Appendix. □

6 Discussion and Extensions

6.1 Uniqueness of payoffs in MPE

The fact that α is the unique limit outcome in MPE does not necessarily imply uniqueness of expected payoffs in MPE, as we demonstrate below.

Fix some $\lambda > 0$ and consider a symmetric game with the following parameter specifications:

$$u_i(\alpha) + u_i(\alpha_i, \beta_{-i}) \geq u_i(\beta_i, \alpha_{-i}) + u_i(\beta) \ \forall i = 1, 2, \tag{4}$$

$$u_i(\beta_i, \alpha_{-i}) > u_i(\alpha_i, \beta_{-i}) \ \forall i = 1, 2. \tag{5}$$

For example these specifications are satisfied by the following stage game:

	α_2	β_2
α_1	$3, 3$	$-2, -1$
β_1	$-1, -2$	$0, 0$

We show that under these payoff specifications, there exists an open set of r's under which α is chosen as the unique limit outcome but there exist multiple equlilibria with α as an absorbing state yielding different payoffs. In fact we show the existence of two kinds of equilibria depicted in the following figure:

---

[18]Note that the results in the previous subsection do not rule out the existence of multiple mixed symmetric MPE in the region in which α is the unique limit outcome in symmetric MPE.
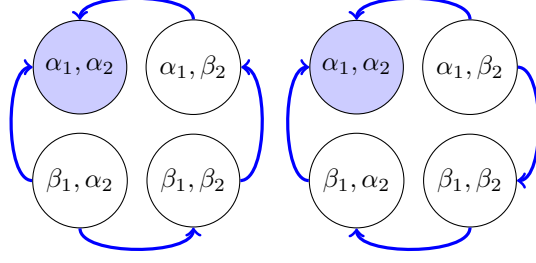
Figure 1: Multiple Equilibria

We now show the existence of the first type of equilibrium under the inqualities assumed above. It is easy to show that an equilibrium of the first type exists if

$$\frac{u_1(\alpha) - u_1(\alpha_1, \beta_2)}{u_1(\beta) - u_1(\alpha_1, \beta_2)} = \frac{u_2(\alpha) - u_2(\beta_1, \alpha_2)}{u_2(\beta) - u_2(\beta_1, \alpha_2)} > 1 + \frac{r}{\lambda} > \frac{u_1(\alpha) - u_1(\beta_1, \alpha_2)}{u_1(\beta) - u_1(\alpha_1, \beta_2)} = \frac{u_2(\alpha) - u_2(\alpha_1, \beta_2)}{u_2(\beta) - u_2(\beta_1, \alpha_2)}$$

By symmetry of payoffs, an equilibrium of the second type exists if the same inequalities hold.

Note moreover that inequality (4) and inequality (5) imply that

$$\frac{u_1(\alpha) - u_1(\alpha_1, \beta_2)}{u_1(\beta) - u_1(\alpha_1, \beta_2)} = \frac{u_2(\alpha) - u_2(\beta_1, \alpha_2)}{u_2(\beta) - u_2(\beta_1, \alpha_2)} > \frac{u_1(\alpha) - u_1(\beta_1, \alpha_2)}{u_1(\beta) - u_1(\alpha_1, \beta_2)} = \frac{u_2(\alpha) - u_2(\alpha_1, \beta_2)}{u_2(\beta) - u_2(\beta_1, \alpha_2)} \geq 1.$$

Therefore for a given λ, there exists some $\varepsilon > 0$ and some $r^* > 0$ such that for all $r \in (r^* - \varepsilon, r^* + \varepsilon)$,

$$\frac{u_1(\alpha) - u_1(\alpha_1, \beta_2)}{u_1(\beta) - u_1(\alpha_1, \beta_2)} = \frac{u_2(\alpha) - u_2(\beta_1, \alpha_2)}{u_2(\beta) - u_2(\beta_1, \alpha_2)} > 1 + \frac{r^*}{\lambda} > \frac{u_1(\alpha) - u_1(\beta_1, \alpha_2)}{u_1(\beta) - u_1(\alpha_1, \beta_2)} = \frac{u_2(\alpha) - u_2(\alpha_1, \beta_2)}{u_2(\beta) - u_2(\beta_1, \alpha_2)} \geq 1.$$

Therefore under these payoffs, λ and discount rate $r \in (r^* - \varepsilon, r^* + \varepsilon)$, both equilibria exist. But note also that inequality (1) implies that under these parameters, α must be the unique limit outcome.

The set of parameter values for which α is the unique limit outcome in MPE and expected MPE payoffs are unique can be characterized using the same strategy we used in establishing Lemma 4.3 and Proposition 4.4. We do not pursue this exercise here.

6.2   Uniqueness of limit outcome in SPNE

A Nash equilibrium of the stage game can be the unique limit outcome not only in all MPE, but in all SPNE, in both coordination games and BoS games. The Supplementary Appendix provides a concrete example. Here we provide a sufficient condition for α be the unique limit outcome in a 2x2 coordination game. The conditions are satisfied for a full-dimensional set of parameter values, implying that unique limit outcome in SPNE is not a nongeneric property.

First we provide conditions guaranteeing that whenever (α) is an absorbing state in an MPE, play eventually ends up there, from any starting state. To simplify notation, given a time t history, we let $\|h^t\|$ be the currently played action at time t.

Lemma 6.1. *Suppose that in a SPNE, both players $i = 1, 2$ play α_i whenever the opponent's currently played action is α_{-i}. Suppose further that either one of the following inequalities holds:*

$$\frac{r}{r+\lambda} u_1(\alpha_1, \beta_2) + \frac{\lambda}{r+\lambda} u_1(\alpha) > u_1(\beta)$$

$$\frac{r}{r+\lambda} u_2(\alpha_2, \beta_1) + \frac{\lambda}{r+\lambda} u_2(\alpha) > u_2(\beta).$$

Then in every subgame perfect equilibrium σ and any history h^t,

$$Pr_\sigma(\{h : \exists\ T\ s.t.\ \|h^\tau\| = \alpha\ \forall \tau \geq T\} \mid h^t) = 1$$

for any action $a \in A$. That is, the history must become absorbed at state α in finite time almost surely.

Proof. Consider all histories h^τ at which $\|h^\tau\| = (\alpha_1, \beta_2)$. Choose any arbitrary $T > 0$. Then the probability that play becomes absorbed at α within T units of time is at least:

$$\frac{1}{2}\left(1 - e^{-2\lambda T}\right) > 0.$$

Precisely the same bound can be used when starting at a history h^τ where $\|h^\tau\| = (\beta_1, \alpha_2)$.

Now consider the state (β_1, β_2). Since one of the inequalities must hold, there must exist $T/2 > 0$ and $p > 0$ such that

$$Pr\left(h : \exists\ \tau \in [t, t+T/2]\ s.t.\ \|h^\tau\| \neq \beta \mid h^t\right) > p$$

for all h^t such that $\|h^t\| = \beta$. But then with probability at least

$$p\frac{1}{2}\left(1 - e^{-\lambda T}\right)$$

play must absorb at α within T periods of t, implying the claim in the lemma. $\qquad\square$

We would like to point out that the conditions in Lemma 6.1 do not guarantee that play is surely absorbed in (α) if both players had a certain number of arrivals (even in a specific order). This is because there can be mixing in an SPNE at histories at which play has not reached yet the limit outcome.

Lemma 6.1 provides a sufficient condition under which the unique limit outcome is α in all SPNE, provided that (α) is an absorbing state in all MPE. Next we provide conditions that guarantee the latter. To simplify notation, we define v_i^β to be:

$$v_i^\beta = \frac{r}{r+\lambda} u_i(\beta) + \frac{\lambda}{r+\lambda}\left(\frac{r}{r+\lambda} u_i(\alpha_i, \beta_{-i}) + \frac{\lambda}{\lambda+r} u_i(\alpha)\right)$$

This payoff describes the payoff that player i obtains by transitioning along the cycle: $(\beta_i, \beta_{-i}) \to (\alpha_i, \beta_{-i}) \to (\alpha_i, \alpha_{-i})$ before eventually absorbing at (α_i, α_{-i}).

Lemma 6.2. *Suppose that the following inequalities both hold for either $i = 1, 2$:*

$$\frac{r}{r+\lambda} u_i(\alpha) + \frac{\lambda}{r+\lambda} u_i(\alpha_i, \beta_{-i}) > \frac{r}{r+\lambda} u_i(\beta_i, \alpha_{-i}) + \frac{\lambda}{r+\lambda} u_i(\alpha)$$

$$\frac{r}{r+\lambda} u_i(\alpha) + \frac{\lambda}{r+\lambda} u_i(\alpha_i, \beta_{-i}) > \frac{r}{r+2\lambda} u_i(\beta_i, \alpha_{-i}) + \frac{2\lambda}{r+2\lambda} \left(\frac{1}{2} u_i(\alpha) + \frac{1}{2} \max\{u_i(\beta), v_i^\beta\} \right).$$

Then in any SPNE, both players $i = 1, 2$ must play α_i whenever player $-i$'s currently played action is α_{-i}.

Proof. See the Supplementary Appendix C. □

Then the above lemma together with Lemma 6.1 immediately implies the following theorem.

Theorem 6.3. *Suppose that the following inequalities both hold for either $i = 1, 2$:*

$$\frac{r}{r+\lambda} u_i(\alpha) + \frac{\lambda}{r+\lambda} u_i(\alpha_i, \beta_{-i}) > \frac{r}{r+\lambda} u_i(\beta_i, \alpha_{-i}) + \frac{\lambda}{r+\lambda} u_i(\alpha)$$

$$\frac{r}{r+\lambda} u_i(\alpha) + \frac{\lambda}{r+\lambda} u_i(\alpha_i, \beta_{-i}) > \frac{r}{r+2\lambda} u_i(\beta_i, \alpha_{-i}) + \frac{2\lambda}{r+2\lambda} \left(\frac{1}{2} u_i(\alpha) + \frac{1}{2} \max\{u_i(\beta), v_i^\beta\} \right).$$

Suppose further that the following inequality holds for at least one $i = 1, 2$:

$$\frac{r}{r+\lambda} u_i(\alpha_i, \beta_{-i}) + \frac{\lambda}{r+\lambda} u_i(\alpha) > u_i(\beta).$$

Then in all SPNE σ,

$$Pr_\sigma(\{h : \exists \ T \ s.t. \ h^t = \alpha \ \forall t \geq T\} \mid h^0 = a^0) = 1$$

for any initial state a^0.

In the Supplementary Appendix we provide a sufficient condition for unique limit outcome in BoS games, as well as additional results related to unique limit outcome in SPNE. We show that (i) unique limit outcome in MPE does not imply unique limit outcome in SPNE; and (ii) uniqueness of limit outcome in SPNE does not imply uniqueness of SPNE or expected payoffs in SPNE. We also provide stricter sufficient conditions guaranteeing uniqueness of SPNE.

6.3 Heterogeneous Arrival Rates

The results of Section 4 generalize to environments with asymmetric arrival rates in a straightforward manner.[19] Below we provide the generalizations and relegate the proofs to the Supplmentary Appendix.

Proposition 6.4. *In a 2-player coordination game, α is the unique limit outcome in MPE iff:*

(i) $\frac{r}{\lambda_{-i}+r} u_i(\alpha_i, \beta_{-i}) + \frac{\lambda_{-i}}{\lambda_{-i}+r} u_i(\alpha) > u_i(\beta)$ for some i;

and

(ii) $u_i(\beta_i, \alpha_{-i}) + \frac{\lambda_{-i}}{\lambda_i+r} u_i(\beta) < u_i(\alpha) + \frac{\lambda_{-i}}{\lambda_i+r} u_i(\alpha_i, \beta_{-i})$ for some i.

---

[19]Our proofs for the results in Section 5 rely on the symmetry of players, hence the techniques we employ cannot be used to extend those results to environments with heterogenous arrival rates.

As in Section 4, define
$$p_i \equiv \frac{u_i(\beta) - u_i(\alpha_i, \beta_{-i})}{u_i(\alpha) - u_i(\alpha_i, \beta_{-i})}.$$
and let $p = \min\{p_1, p_2\}$.

Proposition 6.5. *There exists an open parameter of discount rates in which α is the unique limit outcome in all MPE if and only if either*

$$\frac{u_1(\alpha) - u_1(\beta_1, \alpha_2)}{u_1(\beta) - u_1(\alpha_1, \beta_2)} > \min\left\{ p_2 \frac{\lambda_2}{\lambda_1}, \frac{p_1}{p_1 \frac{\lambda_1}{\lambda_2} + (1 - p_1)} \right\}, \text{ or}$$

$$\frac{u_2(\alpha) - u_2(\beta_2, \alpha_1)}{u_2(\beta) - u_2(\alpha_2, \beta_1)} > \min\left\{ p_1 \frac{\lambda_1}{\lambda_2}, \frac{p_2}{p_2 \frac{\lambda_2}{\lambda_1} + (1 - p_2)} \right\}.$$

Note that when $\lambda_1 = \lambda_2$, the above proposition coincides with Proposition 4.5. Moreover when arrival rates are more unequal, the above inequalities are both more likely to be satisfied. Thus the set of payoff parameters under which there exists an open parameter of discount rates in which α is the unique limit outcome in all MPE expands when arrival rates are more unequal. A similar feature arises in CKLS (2014) in n-player common interest games.[20]

6.4 Battle of the Sexes games

In many situations in which participants would like to coordinate, their interests are only partially aligned, as they disagree what outcome to coordinate on. The simplest class of such games, referred to as Battle of the sexes (henceforth BoS) games, was first discussed in Luce and Raiffa (1957). Here we show that our results for pure coordination games extend to the above class of games.

Using the same notation for payoffs as we did for 2-player coordination games, the only difference in the constraints on payoff parameters is that $u_2(\beta) > u_2(\alpha)$. Thus α is player 1's preferred static Nash equilbrium and β is player 2's preferred static Nash equilibrium.

First we observe that Lemma 4.1 extends to this class of games, with the proof unchanged. This leads to the following modification of Lemma 4.2.

Lemma 6.6. *In BoS games, (β_1, α_2) cannot be part of any absorbing cycle in any MPE.*

Proof. For (β_1, α_2) to be part of an absorbing cycle, first we need $v_1(\alpha) > v_1(\beta_1, \alpha_2)$. Otherwise by Lemma 4.1, $v_1(\beta) > v_1(\alpha_1, \beta_2)$ which implies that $v_2(\beta) > v_2(\beta_1, \alpha_2)$, making β an absorbing state.

Using a symmetric argument, we must have $v_2(\beta) > v_2(\beta_1, \alpha_2)$. But if $v_1(\alpha) > v_1(\beta_1, \alpha_2)$ and $v_2(\beta) > v_2(\beta_1, \alpha_2)$, (β_1, α_2) clearly cannot be part of an absorbing cycle. $\square$

As in the coordination games analysis, we first study necessary and sufficient conditions for an equilibrium with β as a member of the absorbing cycle to exist. The possible equilibrium dynamics exhibiting an absorbing cycle with (β) as a member are the following:

---

[20]See the discussion succeeding Theorem 2 in CKLS (2014).
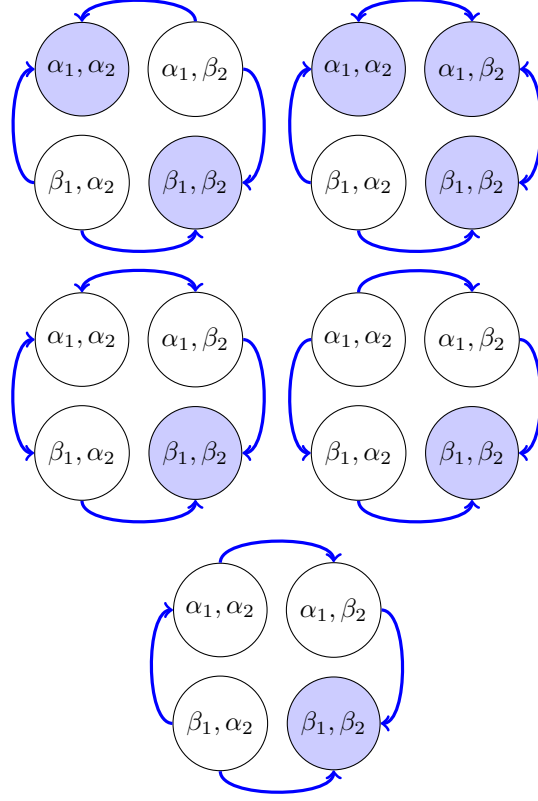
Following the same approach as in 2-player coordination games, we characterize necessary and sufficient conditions under which an equilibrium with (β) as a member of the absorbing cycle exists. Then whenever these conditions are violated, we conclude that (α) must be the unique limit outcome in all equilibria. These conditions are summarized in the following proposition whose details are postponed to the Supplementary Appendix.

Proposition 6.7. *The necessary and sufficient conditions for (α) to be the unique absorbing state in any MPE are:*

$$\frac{r}{\lambda + r} u_2(\alpha_1, \beta_2) + \frac{\lambda}{\lambda + r} u_2(\beta) \quad < \quad u_2(\alpha)$$
$$\frac{r}{\lambda + r} u_1(\alpha_1, \beta_2) + \frac{\lambda}{\lambda + r} u_1(\alpha) \quad > \quad u_1(\beta)$$

and one of the following:

1.

$$u_1(\alpha) + \frac{\lambda}{\lambda + r} u_1(\alpha_1, \beta_2) > u_1(\beta_1, \alpha_2) + \frac{\lambda}{\lambda + r} u_1(\beta)$$

2. *or*

$$u_2(\alpha_1, \beta_2) + \frac{\lambda}{\lambda + r} u_2(\beta) < u_2(\alpha) + \frac{\lambda}{\lambda + r} u_1(\beta_1, \alpha_2).$$

Proof. See Appendix D. □

19

7 References

Ambrus, A., J. Burns and Y. Ishii (2014): "Gradual bidding in Ebay-like auctions," mimeo Duke University

Arcidiacono, P., P. Bayer, J. Blevins and P. Ellickson (2013): "Estimation of dynamic discrete choice models in continuous time," mimeo Duke University.

Atkeson, A. (2000): "Discussion on Morris and Shin," IN: NBER Macroeconomics Annual, 2000.

Bhaskar, V. and F. Vega-Redondo (2002): "Asynchronous choice and Markov equilibria," Journal of Economic Theory, 103, 334-350.

Bhaskar, V., G. Mailath and S. Morris (2013): "A foundation for Markov equilibria with finite social memory," Review of Economic Studies 80, 925-948.

Burdzy, K., D. Frankel and A. Pauzner (2001): "Fast equilibrium selection by rational players living in a changing world," Econometrica, 69, 163-189.

Burdett, K. and K. Judd (1983): "Equilibrium price distortion," Econometrica, 51, 955-969.

Calcagno, R., Y. Kamada, S. Lovo and T. Sugaya (2014): "Asynchronicity and Coordination in Common and Opposing Interest Games," Theoretical Economics, forthcoming.

Calvo, G. (1983): "Staggered prices in a utility-maximizing framework," Journal of Monetary Economics, 12, 383-398.

Cooper, R. and A. John (1988): "Coordinating coordination failures in Keynesian models," Quarterly Journal of Economics 103, 441-463.

Doraszelski, U. and K. Judd (2012): "Avoiding the curse of dimensionality in dynamic stochastic games," Quantitative Economics, 3, 53-93.

Diamond, D. and P. Dybvig (1983): "Bank runs, deposit insurance, and liquidity," Journal of Political Economy 91, 401-419.

Dutta, P. (1995): "A folk theorem for stochastic games," Journal of Economic Theory, 66, 1-32.

Dutta, P. (2012): "Coordination need not be a problem," Games and Economic Behavior 76, 519-534.

Einav, L. (2010): "Not all rivals look alike: Estimating an equilibrium model of the release date timing game," Economic Inquiry 48, 369–390.

Farrell, J. and P. Klemperer (1997): "Coordination and lock-in: Competition with switching costs and network effects," IN: Handbook of Industrial Organization, 3, 1967-2072.

Farrell, J. and G. Saloner (1985): "Standardization, compatibility, and innovation," Rand Journal of Economics, 16, 70-83.

Frankel, D. and A. Pauzner (2000): "Resolving indeterminancy in dynamic settings: The role of shocks," Quarterly Journal of Economics, 115, 285-304.

Frick, M. and Y. Ishii (2014): "Innovation adoption by forward-looking social learners," mimeo Harvard University.

Gallager, R. (2013): "Stochastic processes: Theory for applications," Cambridge University Press.

Haller, H. and R. Lagunoff (2000): "Genericity and Markovian behavior in stochastic games," Econometrica, 1231-1248.

Kandori, M., G. Mailath and R. Rob (1993): "Learning, mutation, and long-run equilibria in games," Econometrica, 61, 29-56.

Karlin, S. and H. Taylor (1975): A first course in stochastic processes, 2nd edition, Academic Press.

Katz, M. and C. Shapiro (1985): "Network externalities, competition, and compatibility," American Economic Review 75, 424-440.

Katz, M. and C. Shapiro (1986): "Technology adoption in the presence of network externalities," Journal of Political Economy, 94, 822-841.

Lagunoff, R. and A. Matsui (1997): "Asynchronous choice in repeated coordination games," Econometrica, 65, 1467-1477.

Lagunoff, R. and A. Matsui (2001): "Are "anti-folk" theorems in repeated games nongeneric?", Review of Economic Design, 6, 397-412.

Luce, R. and H. Raiffa (1957): Games and Decisions: Introduction and Critical Survey. Wiley, New York.

Maskin, E. and J. Tirole (1988): "A theory of dynamic oligopoly I: Overview and quantity competition with large fixed costs," Econometrica, 56, 549-569.

Maskin, E. and J. Tirole (1988): "A theory of dynamic oligopoly II: Price competition, kinked demand curves, and education cycles," Econometrica, 56, 571-599.

Matsui, A. and K. Matsuyama (1995): "An approach to equilibrium selection," Journal of Economic Theory, 65, 415-434.

Obsfelt, M. (1986): "Rational and self-fulfilling balance-of-payments crisis," American Economic Review 76, 72-81.

Schmidt-Dengler, P. (2006): "The timing of new technology adoption: The case of MRI," mimeo London School of Economics.

Takahashi, S. (2005): "Infinite horizon common interest games with perfect information," Games and Economic Behavior, 53, 231-247.

Takahashi, S. and Q. Wen (2003): "On asynchronously repeated games," Economics Letters, 79, 239-245.

Yoon, K. (2001): "A folk theorem for asynchronously repeated games," Econometrica, 69, 191-200.

Young, P. (1993): "The evolution of conventions," Econometrica, 61, 57-84.

Young, P. (2009): "Innovation diffusion in heterogeneous populations: Contagion, social influence, and social learning," American Economic Review, 99, 1899-1924.

A n-Player Symmetric Games

We first establish the following proposition, which directly implies some of the statements in subsequent proofs.

Proposition A.1. *Let $r_\ell \to 0$ and p^ℓ be a symmetric MPE of the game with discount factor r_ℓ such that $p^\ell_{n-1} = 1$ for all ℓ. Let v^ℓ be the corresponding vector of equilibrium payoffs. Suppose that*

$$\lim_{\ell \to \infty} p^\ell = p, \lim_{\ell \to \infty} v^\ell = v.$$

Then $v(s) = u^\alpha_{n-1}$ for all s.

A.1 Proof of Proposition A.1

In this section, we prove Proposition A.1. Consider the sequence defined in the lemma. Define $\delta_\ell = \frac{n\lambda}{r_\ell + n\lambda}$. Given any matrix X, denote by $X^\top$, the transpose of X. Let

$$u = \left(u^\alpha_{n-1}, u^\beta_{n-1}, u^\alpha_{n-2}, u^\beta_{n-2}, \ldots, u^\alpha_0, u^\beta_0 \right)^\top.$$

Define for each ℓ the corresponding vector of equilibrium continuation payoffs

$$v^\ell = \left(v^{\alpha,\ell}_{n-1}, v^{\beta,\ell}_{n-1}, v^{\alpha,\ell}_{n-2}, v^{\beta,\ell}_{n-2}, \ldots, v^{\alpha,\ell}_0, v^{\beta,\ell}_0 \right)^\top$$

We can define a transition matrix A_ℓ associated with a symmetric equilibrium p^ℓ in a natural way:

$$A_\ell((\alpha,k),(\alpha,k+1)) = \frac{n-k-1}{n} p^\ell_{k+1}$$

$$A_\ell((\alpha,k),(\alpha,k-1)) = \frac{k}{n}(1-p^\ell_k)$$

$$A_\ell((\alpha,k),(\beta,k)) = \frac{1}{n}(1-p^\ell_k)$$

$$A_\ell((\alpha,k),(\alpha,k)) = 1 - \frac{n-k-1}{n} p^\ell_{k+1} - \frac{k}{n}(1-p^\ell_k) - \frac{1}{n}(1-p^\ell_k),$$

and

$$A_\ell((\beta,k),(\beta,k+1)) = \frac{n-k-1}{n} p^\ell_k$$

$$A_\ell((\beta,k),(\beta,k-1)) = \frac{k}{n}(1-p^\ell_{k-1})$$

$$A_\ell((\beta,k),(\alpha,k)) = \frac{1}{n}p^\ell_k$$

$$A_\ell((\beta,k),(\beta,k)) = 1 - \frac{n-k-1}{n} p^\ell_k - \frac{k}{n}(1-p^\ell_{k-1}) - \frac{1}{n}p^\ell_k.^{21}$$

Then the Bellman equation implies:

$$v_\ell = (1-\delta_\ell)u + \delta_\ell A_\ell v_\ell \Leftrightarrow v_\ell = (1-\delta_\ell)\left(I - \delta_\ell A_\ell\right)^{-1} u.^{22}$$

---

[21]Note that this transition matrix differs slightly from the transition matrix defined at the beginning of Section A. Here it is important to keep track of an individual's currently played action so the transition matrix must also encode that additional information. This additional information is ignored in Section A.

Furthermore we know that A_ℓ converges (because p^ℓ converges) and so let $A = \lim_{\ell \to \infty} A_\ell$.

A.1.1 Elementary Properties of $\mathcal{T}$ and Further Definitions

Consider the Markov chain induced by the transition matrix $A = \lim_{\ell \to \infty} A_\ell$. By Corollary 5.1 (p66) in Karlin and Taylor (1975), we can classify each communication class into either a transient class or a recurrent class. Let $\mathcal{T}$ be the recurrent *classes* of the Markov chain A and let T be the set of all recurrent *states*.

Lemma A.2. *The following properties hold for $\mathcal{T}$.*

1. *If (α, k) is recurrent for some $k < n - 1$, then $(\beta, k + 1)$ is also recurrent.*

2. *If (β, k) is recurrent for some $k > 0$, then $(\alpha, k - 1)$ is also recurrent.*

3. *If $(\alpha, k), (\alpha, k') \in \mathcal{R} \in \mathcal{T}$ for some $k' > k$, then $(\alpha, k'') \in \mathcal{R}$ for all $k'' = k, k + 1, \ldots, k'$.*

4. *If $(\beta, k), (\beta, k') \in \mathcal{R} \in \mathcal{T}$ for some $k' > k$, then $(\beta, k'') \in \mathcal{R}$ for all $k'' = k, k + 1, \ldots, k'$.*

Proof. Proof is obvious and thus omitted. □

Given any state s, let $\|s\|$ denote the number of players playing α at the state. For example, $\|(\alpha, k)\| = k + 1$ and $\|(\beta, k)\| = k$. Let

$$\mathcal{T}^\beta := \{\mathcal{R} \in \mathcal{T} : \exists k, (\beta, k) \in \mathcal{R}\},$$

$$\mathcal{T}^\alpha := \{\mathcal{R} \in \mathcal{T} : \exists k, (\alpha, k) \in \mathcal{R}\}.$$

Note that it may be possible that $\mathcal{T}^\beta \cap \mathcal{T}^\alpha \neq \emptyset$, but that $\mathcal{T}^\beta \cup \mathcal{T}^\alpha = \mathcal{T}$.

We now form a partial order on sets of states. Let $U, U' \subseteq S$. Then we define:

$$U > U' \iff \|s\| > \|s'\| \text{ for all } s \in U, s' \in U'.$$

By Lemma A.2, and because we assumed $p_0 = 0$ and $p_{n-1} = 1$ all along the sequence, we can then enumerate the sets in $\mathcal{T}^\alpha$ and $\mathcal{T}^\beta$ in the following manner. If $\{(\beta, 0)\} \in \mathcal{T}$, then we enumerate according to the following:

$$\{(\alpha, n - 1)\} = \mathcal{R}_{m+1}^\alpha > \mathcal{R}_m^\alpha > \cdots > \mathcal{R}_1^\alpha,$$

$$\mathcal{R}_m^\beta > \cdots > \mathcal{R}_1^\beta > \mathcal{R}_0^\beta = \{(\beta, 0)\}.$$

Otherwise, we enumerate in the following manner:

$$\{(\alpha, n - 1)\} = \mathcal{R}_{m+1}^\alpha > \mathcal{R}_m^\alpha > \cdots > \mathcal{R}_1^\alpha > \mathcal{R}_0^\alpha,$$

$$\mathcal{R}_m^\beta > \cdots > \mathcal{R}_1^\beta > \mathcal{R}_0^\beta.$$

---

[22] Note that $(I - \delta_\ell A_\ell)$ is necessarily invertible since $\sum_{t=0}^\infty \delta_\ell^t A_\ell^t$ exists and constitutes an inverse.

Definition A.3. Given a state s, we say that $\mathcal{R} \in \mathcal{T}$ neighbors s if $s \in \mathcal{R}$ or if $s \notin T$ and either

1. $\mathcal{R} > \{s\}$ and there exists no $\mathcal{R}' \in \mathcal{T}$ such that $\mathcal{R} > \mathcal{R}' > \{s\}$, or

2. or $\mathcal{R} < \{s\}$ and there exists no $\mathcal{R}' \in \mathcal{T}$ such that $\mathcal{R} < \mathcal{R}' < \{s\}$.

Given a state s, let $\mathcal{N}(s)$ be the set of recurrent classes that neighbor s. Note that the cardinality of $\mathcal{N}(s)$ equals 1 if and only if s is recurrent. Otherwise, it can be either 2, 3, or 4.[23] Given any recurrent class $\mathcal{R} \in \mathcal{T}$ and any state s, let

$$A^*_{s\mathcal{R}} := \sum_{s' \in \mathcal{R}} A^*_{ss'}.$$

A.1.2 Properties of A

We first establish some elementary properties of A.

Lemma A.4. Let $\mathcal{R} \in \mathcal{T}$. Then $\mathcal{R}$ is aperiodic. Furthermore $A^* := \lim_{t \to \infty} A^t$ exists.

Proof. Suppose otherwise. Then there exists some state $s \in \mathcal{R}$ such that

$$d(s) = \gcd\{t : A^t_{ss} > 0\} > 1,$$

which in particular implies that $A_{ss} = 0$. This means however that according to the transition matrix A, all players must change their actions with probability one from their corresponding actions in the state s. This is then a contradiction since s can no longer be recurrent. Thus $\mathcal{R}$ is aperiodic.

For a proof of the last statement, see the characterization of A^* in Karlin and Taylor (1975), in Theorems 1.2 and 3.1 and Problem 6 (p.77). $\qquad\square$

Lemma A.5. $Av = v$ and $A^*v = v$.

Proof. Note the following identity:

$$\delta_\ell A_\ell v^\ell + (1 - \delta_\ell)u = (\delta_\ell A_\ell P_\ell + (1 - \delta_\ell)I)\, u = P_\ell u = v^\ell.$$

This then implies that $Av = v$. This therefore implies that $A^t v = v$ for all $t > 0$ which means that $A^*v = v$. $\qquad\square$

The above lemma allows for the use of important information regarding A^* from standard Markov chain theory to obtain information about the structure of P and thus the structure of v.

Lemma A.6. *Suppose that $\mathcal{R} \in \mathcal{T}$ and that $s, s' \in \mathcal{R}$. Then $v(s) = v(s')$.*

---

[23]Recall that we assumed that A has a recurrent class at $\{\alpha, n-1\}$ and $\{\beta, 0\}$, which implies that $\mathcal{N}(s)$ is at least 2 for transitory states. Note also that in both directions from s there can be two different recurring classes, depending on player i playing α or β.

Proof. Let the unique invariant distribution of the recurrent class $\mathcal{R}$ be π (defined as a row vector) whose support is obviously in $\mathcal{R}$. Denote the submatrix of A^* corresponding to the rows of states in $\mathcal{R}$ by $A^* \mid_{\mathcal{R}}$.

Theorem 1.2 of Karlin and Taylor (1975) implies that

$$A^* \mid_{\mathcal{R}} = \left(\begin{array}{ccccccc} 0 & \cdots & 0 & G & 0 & \cdots & 0 \end{array} \right)$$

where

$$G = \begin{pmatrix} \pi \\ \pi \\ \vdots \\ \pi \end{pmatrix}.$$

By Lemma A.5, since $A^*v = v$, $v(s) = v(s')$. $\qquad\square$

Due to the above lemma, given a recurrent class $\mathcal{R}$, we write $v(\mathcal{R})$ for the payoff associated with any state in recurrent class $\mathcal{R}$.

Lemma A.7. *Let s be any state. Then*

$$v(s) = \sum_{\mathcal{R} \in \mathcal{N}(s)} A^*_{s\mathcal{R}} v(\mathcal{R}).$$

In other words, $v(s)$ is a convex combination of the elements of the set $\{v(\mathcal{R})\}_{\mathcal{R} \in \mathcal{N}(s)}$.

Proof. This is trivial if s is recurrent. So assume that s is transient. By Lemma A.5, $v = A^*v$. Furthermore note that if s' is transient, $A^*_{ss'} = 0$ for all s. Thus

$$v(s) = A^*_s v = \sum_{s' \in T} A^*_{ss'} v(s') = \sum_{\mathcal{R} \in \mathcal{T}} A^*_{s\mathcal{R}} v(\mathcal{R}).$$

Then it is easy to see that $A^*_{s\mathcal{R}} = 0$ for any $\mathcal{R} \notin \mathcal{N}(s)$. Therefore

$$v(s) = \sum_{\mathcal{R} \in \mathcal{N}(s)} A^*_{s\mathcal{R}} v(\mathcal{R}).$$

This concludes the proof. $\qquad\square$

A.1.3   Concluding the Proof

To simplify notation, let us define the following: $v\left(\mathcal{R}^{\beta}_{m+1}\right) := u^{\alpha}_{n-1}$ and if $\mathcal{R}^{\beta}_0 = \{(\beta, 0)\}$, $v(\mathcal{R}^{\alpha}_0) := v\left(\mathcal{R}^{\beta}_0\right)$.[24]

Lemma A.8. $v\left(\mathcal{R}^{\alpha}_j\right) \geq v\left(\mathcal{R}^{\beta}_j\right)$ *for all j.*

---

[24]Note that we define these accordingly even though $\mathcal{R}^{\beta}_{m+1}$ and $\mathcal{R}^{\alpha}_0$ are undefined.

Proof. We prove this by induction. Let us first prove that $v\left(\mathcal{R}_0^\alpha\right) \geq v\left(\mathcal{R}_0^\beta\right)$. If $\mathcal{R}_0^\beta = \{(\beta,0)\}$, this follows by definition. Now suppose that $(\beta,0) \in \mathcal{R}_0^\beta$ but that $\{(\beta,0)\} \neq \mathcal{R}_0^\beta$. Then note that $\mathcal{R}_0^\alpha = \mathcal{R}_0^\beta$ and so trivially $v\left(\mathcal{R}_0^\alpha\right) \geq v\left(\mathcal{R}_0^\beta\right)$. Finally suppose that $(\beta,0) < \mathcal{R}_0^\beta$ and that $v\left(\mathcal{R}_0^\alpha\right) < v\left(\mathcal{R}_0^\beta\right)$. Let $k = \min\{k' : (\alpha,k') \in \mathcal{R}_0^\alpha\}$. Then $p_k = 1$ and so using the equation $Av = v$, we get:

$$
v\left(\mathcal{R}_0^\alpha\right) = v_k^\alpha \geq v_k^\beta \geq \frac{1}{n}v_k^\alpha + \frac{k}{n}v\left(\mathcal{R}_0^\alpha\right) + \frac{n-k-1}{n}v_{k+1}^\beta
$$
$$
= \frac{k+1}{n}v\left(\mathcal{R}_0^\alpha\right) + \frac{n-k-1}{n}v\left(\mathcal{R}_0^\beta\right) > v\left(\mathcal{R}_0^\alpha\right),
$$

which is a contradiction. This shows that $v\left(\mathcal{R}_0^\alpha\right) \geq v\left(\mathcal{R}_0^\beta\right)$.

Suppose now that $v\left(\mathcal{R}_j^\alpha\right) \geq v\left(\mathcal{R}_j^\beta\right)$. We want to prove that $v\left(\mathcal{R}_{j+1}^\alpha\right) \geq v\left(\mathcal{R}_{j+1}^\beta\right)$. Suppose otherwise so that

$$
v\left(\mathcal{R}_{j+1}^\alpha\right) < v\left(\mathcal{R}_{j+1}^\beta\right).
$$

We examine two cases:

$$
v\left(\mathcal{R}_j^\beta\right) \geq v\left(\mathcal{R}_{j+1}^\alpha\right), \tag{6}
$$
$$
v\left(\mathcal{R}_j^\beta\right) < v\left(\mathcal{R}_{j+1}^\alpha\right). \tag{7}
$$

Consider first case (6). In this case note that

$$
v\left(\mathcal{R}_{j+1}^\alpha\right) = \min\left\{v\left(\mathcal{R}_j^\beta\right), v\left(\mathcal{R}_{j+1}^\beta\right), v\left(\mathcal{R}_j^\alpha\right), v\left(\mathcal{R}_{j+1}^\alpha\right)\right\}.
$$

Let $k = \min\left\{k' : (\alpha,k') \in \mathcal{R}_1^\alpha\right\}$. By the definition of a recurrent class, $p_k = 1$. Note that because $Av = v$,

$$
v_k^\beta = \frac{1}{n}v_k^\alpha + \frac{n-k-1}{n}v_{k+1}^\beta + \frac{k}{n}\left(p_{k-1}v_k^\beta + (1-p_{k-1})v_{k-1}^\beta\right)
$$
$$
\geq \frac{1}{n}v\left(\mathcal{R}_{j+1}^\alpha\right) + \frac{n-k-1}{n}v\left(\mathcal{R}_{j+1}^\beta\right) + \frac{k}{n}v\left(\mathcal{R}_{j+1}^\alpha\right)
$$

where the last inequality follows from Lemma A.7. Then $v_k^\beta > v_k^\alpha$ which implies that for ℓ sufficiently large, $p_k^\ell = 0$. But this contradicts the assumption that $p_k = 1$.

Now consider case (7). In this case,

$$
v\left(\mathcal{R}_j^\beta\right) = \min\left\{v\left(\mathcal{R}_j^\beta\right), v\left(\mathcal{R}_{j+1}^\beta\right), v\left(\mathcal{R}_j^\alpha\right), v\left(\mathcal{R}_{j+1}^\alpha\right)\right\}.
$$

Let $\overline{k} = \max\left\{k' : (\beta,k') \in \mathcal{R}_j^\beta\right\}$. Then let $k = \min\left\{k' > \overline{k} : v_{k'}^\beta > v\left(\mathcal{R}_j^\beta\right)\right\}$. Note that $v_{k-1}^\beta = v\left(\mathcal{R}_j^\beta\right)$. Because $Av = v$, $p_{k-1} = 0$ since otherwise $v_{k-1}^\beta > v\left(\mathcal{R}_j^\beta\right)$. Furthermore $p_k > 0$ since otherwise, $v_k^\beta = v\left(\mathcal{R}_j^\beta\right)$. Therefore $v_k^\alpha \geq v_k^\beta > v\left(\mathcal{R}_j^\beta\right)$. Moreover because $Av = v$,

$$
v_{k-1}^\alpha = \frac{1}{n}v_{k-1}^\beta + \frac{k-1}{n}v_{k-2}^\alpha + \frac{n-k}{n}\left(p_k v_k^\alpha + (1-p_k)v_{k-1}^\alpha\right) > v\left(\mathcal{R}_j^\beta\right) = v_{k-1}^\beta.
$$

This however contradicts the earlier observation that $p_{k-1} = 0$. $\qquad\square$

To complete the proof we now show that $v\left(\mathcal{R}_j^\alpha\right) \geq v\left(\mathcal{R}_j^\beta\right) = u_{n-1}^\alpha$ for all j, which implies, by Lemma A.7, that $v(s) = u_{n-1}^\alpha$ for all s. Suppose otherwise. Let

$$j^* = \min\left\{ j : v\left(\mathcal{R}_{j'}^\beta\right) = u_{n-1}^\alpha \ \forall j' > j \right\}.$$

Also let

$$k = \max\{k' : (\beta, k') \in \mathcal{R}_{j^*}^\beta\},$$

$$\underline{k} = \max\left\{ k' \geq k : v_{k''}^\beta = u\left(\mathcal{R}_{j^*}^\beta\right) \ \forall k'' = k, \ldots, k' \right\}.$$

By Lemma A.8, we know that $v\left(\mathcal{R}_{j^*}^\alpha\right) \geq v\left(\mathcal{R}_{j^*}^\beta\right)$ and that $v\left(\mathcal{R}_j^\alpha\right) = v\left(\mathcal{R}_j^\beta\right) = u_{n-1}^\alpha$ for all $j > j^*$. These observations along with Lemma A.7 imply that $v_{\underline{k}+1}^\beta > v\left(\mathcal{R}_{j^*}^\beta\right)$. For this to be the case, it must be that $p_{\underline{k}+1} > 0$, otherwise $v_{\underline{k}+1}^\beta = \left(\mathcal{R}_{j^*}^\beta\right)$ by $Av = v$. But this implies that $v_{\underline{k}+1}^\alpha \geq v_{\underline{k}+1}^\beta > v\left(\mathcal{R}_{j^*}^\beta\right)$. Furthermore $p_{\underline{k}} = 0$ since otherwise, $v_{\underline{k}}^\beta > v\left(\mathcal{R}_{j^*}^\beta\right)$. But since $Av = v$,

$$v_{\underline{k}}^\alpha = \frac{1}{n} v_{\underline{k}}^\beta + \frac{\underline{k}}{n} v_{\underline{k}-1}^\alpha + \frac{n - \underline{k} - 1}{n}\left(p_{\underline{k}+1} v_{\underline{k}+1}^\alpha + (1 - p_{\underline{k}+1}) v_{\underline{k}}^\alpha \right) > v\left(\mathcal{R}_{j^*}^\beta\right).$$

This however implies that $p_{\underline{k}} = 1$ which is a contradiction. This completes the proof of Proposition A.1.

A.2 Proof of Claim 5.2

We can now prove Claim 5.2 which proceeds via two steps:

1. Any recurrent class of any equilibrium must be either $\{\alpha\}$ or $\{\beta\}$.

2. There exists some r^* such that for all $r < r^*$, there exists no equilibrium with both $\{\alpha\}$ and $\{\beta\}$ as recurrent classes.

Here we show the first step of the proof above.

Lemma A.9. *Suppose that a stationary symmetric MPE has a recurrent class $\mathcal{C}$ consisting of multiple states. Then $\mathcal{C}$ must be connected (i.e. $k < k'$ and $k, k' \in \mathcal{C}$ implies $\ell \in \mathcal{C}$ for all $\ell \in (k, k')$). Furthermore if $\mathcal{C} = \{\underline{k}, \underline{k} + 1, \ldots, \overline{k}\}$, then $v_k^\alpha = v_k^\beta$ for every $\underline{k} \leq k \leq \overline{k} - 1$.*

Proof. The proof is immediate and thus omitted. □

Lemma A.10. *Suppose that $\mathcal{C}$ is a recurrent class of some symmetric stationary MPE. Then $\mathcal{C}$ consists of exactly one state.*

Proof. Suppose otherwise and let $\mathcal{C} = \{\underline{k}, \underline{k} + 1, \ldots, \overline{k}\}$. By the definition of a recurrent class, $p_{\underline{k}-1} = 1$ and $p_{\overline{k}} = 0$. Furthermore we must have $v_k^\alpha = v_k^\beta$ for all $k = \underline{k}, \ldots, \overline{k} - 1$.

First suppose that $\underline{k} = 0$ and $\overline{k} = n - 1$. Then we must have $v_{n-1}^\alpha = v_{n-1}^\beta$, $v_{n-2}^\alpha = v_{n-2}^\beta$. Note that because it is optimal to always play α in the future, v_{n-1}^α must be some convex combination of u_{n-1}^α and

v_{n-2}^α, i.e. there exists some $\lambda \in (0,1)$ such that $v_{n-1}^\alpha = \lambda u_{n-1}^\alpha + (1-\lambda)v_{n-2}^\alpha$. Similarly, there exists some $\mu \in (0,1)$ such that $v_{n-1}^\beta = \mu u_{n-1}^\beta + (1-\mu)v_{n-2}^\beta$. Clearly, $v_{n-1}^\beta > u_{n-1}^\beta$ and $v_{n-1}^\alpha < u_{n-1}^\alpha$. Therefore the above expressions imply that $v_{n-1}^\alpha > v_{n-2}^\alpha$ and $v_{n-1}^\beta < v_{n-2}^\beta$. This then contradicts indifference at all states.

Next consider the case in which either $0 < \underline{k}$ or $\overline{k} < n$. The proofs are exactly symmetric so we consider only the case in which $\overline{k} < n$. Consider the state at which a player is currently playing α and $\overline{k}$ others are playing α. By the definition of a recurrent class, $\sigma(\overline{k}) = 0$ and therefore $v_{\overline{k}}^\alpha \leq v_{\overline{k}}^\beta$.

Note however that there exists some $\lambda \in (0,1]$ such that $v_{\overline{k}}^\alpha \geq \lambda u(\alpha, \overline{k}) + (1-\lambda)v_{\overline{k}-1}^\alpha$ since by optimality, the player cannot have a profitable deviation to playing α forever until play absorbs at the recurrent class. Note however that because always playing α is optimal once play has entered the recurrent class, $v_{\overline{k}-1}^\alpha < u(\alpha, \overline{k})$. Therefore $v_{\overline{k}}^\alpha > v_{\overline{k}-1}^\alpha$.

Now consider a player currently playing β against $\overline{k}$ others playing α. An optimal strategy for such a player is to always play β and therefore $v_{\overline{k}}^\beta > u(\beta, \overline{k})$. Furthermore there exists some $\mu \in (0,1]$ such that $v_{\overline{k}}^\beta = \mu u(\beta, \overline{k}) + (1-\mu)v_{\overline{k}-1}^\beta$. Therefore $v_{\overline{k}-1}^\beta > v_{\overline{k}}^\beta$. But the above inequalities imply: $v_{\overline{k}}^\alpha > v_{\overline{k}-1}^\alpha = v_{\overline{k}-1}^\beta > v_{\overline{k}}^\beta$. This means that $\mathcal{C}$ cannot be a recurrent class, which is a contradiction. $\qquad\square$

Lemma A.11. *Let k be a recurrent state of some stationary symmetric MPE. Then $k = 0$ or $k = n$.*

Proof. Suppose there exists a symmetric MPE such that there exists an interior recurrent state $0 < k < n$. Then $v_k^\beta = u_k^\beta$ and $v_{k-1}^\alpha = u_{k-1}^\alpha$. Suppose first that $u_k^\beta \geq u_{k-1}^\alpha$. Consider a player currently playing α when facing $k-1$ opponents currently playing α.

By deviating and playing β forever the player obtains a payoff at least $\lambda u_{k-1}^\beta + (1-\lambda)v_k^\beta = \lambda u_{k-1}^\beta + (1-\lambda)u_k^\beta$ for some $\lambda \in (0,1]$. But this strict convex combination is strictly greater than $u_k^\beta \geq u_{k-1}^\alpha$ which means that this player has a strictly positive profitable deviation, which is a contradiction.

Similarly consider the case in which $u_k^\beta < u_{k-1}^\alpha$. Then a player currently playing β against k others playing α can deviate to playing α forever in the future giving him a payoff at least: $\lambda u_k^\alpha + (1-\lambda)u_{k-1}^\alpha$ for some $\lambda \in (0,1]$. This is again strictly greater than $u_{k-1}^\alpha > u_k^\beta$, which yields a strictly positive profitable deviation. This proves the claim. $\qquad\square$

Thus far we have shown that in any equilibrium under any parameters, the only recurrent classes are either $\{\alpha\}$ or $\{\beta\}$. To complete the proof of Claim 5.2, we need to show that there exists some $r^* > 0$ such that for all $r < r^*$, there does not exist any equilibria p with $p_0 = 0$ *and* $p_{n-1} = 1$. Suppose otherwise. Then there exists a sequence $r_\ell \to 0$ with corresponding equilibria p^ℓ (with continuation payoff vector v^ℓ) such that $p_0^\ell = 0$ and $p_{n-1}^\ell = 1$ for all ℓ. Now take a subsequence such that both p^ℓ and v^ℓ converge. But note that $\lim_{\ell \to \infty} v_0^{\beta,\ell} = u_0^\beta$ since $p_0^\ell = 0$ for all ℓ. However Proposition A.1 implies that $\lim_{\ell \to \infty} v_0^{\beta,\ell} = u_{n-1}^\alpha$, a contradiction. This completes the proof of Claim 5.2.

A.3 Proof of Claim 5.3

We prove the following two lemmata which together immediately imply Claim 5.3.

Lemma A.12. *Suppose that an equilibrium q exists with $q_k < 1$ for all $k = 0, 1, \ldots, n-1$ and $q_0 = 0$. Then for every $k \leq n-1$,*

$$v_k^\beta \leq \frac{r}{r + \lambda(k+1)} u_k^\beta + \frac{\lambda(k+1)}{r + \lambda(k+1)} \left(\frac{1}{k+1} v_k^\beta + \frac{k}{k+1} v_{k-1}^\beta \right).$$

Proof. We first show that v_k^β is strictly decreasing in k. Note first that $v_0^\beta = u_0^\beta$ and since always playing β is an optimal strategy, $v_0^\beta > v_1^\beta$ trivially. Now suppose that there exists some k such that $v_{k+1}^\beta \geq v_k^\beta$. Let k^* be the minimal such k. Note from the above discussion that $k^* \geq 1$. Again because an optimal strategy is to play β forever, there must be some $\lambda \in (0, 1)$ such that

$$(1 - \lambda)u_{k^*+1}^\beta + \lambda v_{k^*}^\beta \geq v_{k^*+1}^\beta \geq v_{k^*}^\beta,$$

which implies that $u_{k^*}^\beta > u_{k^*+1}^\beta \geq v_{k^*+1}^\beta \geq v_{k^*}^\beta$. But now consider the payoff $v_{k^*}^\beta$. There must exist some $a, b > 0$ with $a + b \leq 1$ such that $v_{k^*}^\beta = au_{k^*}^\beta + bv_{k^*-1}^\beta + (1 - a - b)v_{k^*+1}^\beta$. But since $u_{k^*}^\beta, v_{k^*-1}^\beta > v_{k^*}^\beta$ and $v_{k^*+1}^\beta \geq v_{k^*}^\beta$,

$$v_{k^*}^\beta = au_{k^*}^\beta + bv_{k^*-1}^\beta + (1 - a - b)v_{k^*+1}^\beta > v_{k^*}^\beta,$$

which is a contradiction.

With the above observation, we can conclude the proof. Clearly the statement holds for v_0^β, so consider any $k \geq 1$. Then

$$v_k^\beta = \frac{r}{r + \lambda n} u_k^\beta + \frac{\lambda n}{r + \lambda n} \left(\frac{1}{n} v_k^\beta + \frac{k}{n} \left(p_{k-1} v_k^\beta + (1 - p_{k-1})v_{k-1}^\beta \right) \right.$$

$$\left. + \frac{n - k - 1}{n} \left(p_k v_{k+1}^\beta + (1 - p_k)v_k^\beta \right) \right)$$

$$\leq \frac{r}{r + \lambda n} u_k^\beta + \frac{\lambda n}{r + \lambda n} \left(\frac{1}{n} v_k^\beta + \frac{k}{n} v_{k-1}^\beta + \frac{n - k - 1}{n} v_k^\beta \right),$$

where the last inequality follows from the observation that v_k^β is strictly decreasing in k. Then the above implies:

$$v_k^\beta \leq \frac{r}{r + \lambda(k+1)} u_k^\beta + \frac{\lambda(k+1)}{r + \lambda(k+1)} \left(\frac{1}{k+1} v_k^\beta + \frac{k}{k+1} v_{k-1}^\beta \right).$$

$\square$

Lemma A.13. *Suppose that an equilibrium q exists with $q_k < 1$ for all $k = 0, 1, \ldots, n-1$ and $q_0 = 0$. Define $\underline{k} = \max\{k < n-1 : q_k \in (0,1)\}$.[25] Then for all $k > \underline{k}$,*

$$v_k^\alpha \geq \frac{r}{r + \lambda(k+1)} u_k^\alpha + \frac{\lambda(k+1)}{r + \lambda(k+1)} \left(\frac{1}{k+1} v_k^\beta + \frac{k}{k+1} v_{k-1}^\alpha \right).$$

---

[25]Throughout we use the convention that $\max \emptyset = -1$.

Proof. We first show that $p_{k+1}v_{k+1}^\alpha + (1-p_{k+1})v_k^\alpha \geq v_k^\alpha$ for all $k > \underline{k}$. Note that the above is trivial if $\underline{k} < k < n-2$ or if $p_{n-1} = 0$. So it remains to prove the claim when $p_{n-1} \in (0,1)$ and $k = n-2$. In this case, note that

$$v_{n-1}^\alpha = \frac{r}{r+(n-1)\lambda}u_{n-1}^\alpha + \frac{(n-1)\lambda}{r+(n-1)\lambda}\left(p_{n-1}v_{n-1}^\alpha + (1-p_{n-1})v_{n-2}^\alpha\right).$$

Because $u_{n-1}^\alpha > \left(p_{n-1}v_{n-1}^\alpha + (1-p_{n-1})v_{n-2}^\alpha\right)$, $v_{n-1}^\alpha \geq v_{n-2}^\alpha$. Therefore, $p_{n-1}v_{n-1}^\alpha + (1-p_{n-1})v_{n-2}^\alpha \geq v_{n-2}^\alpha$.

Now consider any $k > \underline{k}$.

$$
\begin{aligned}
v_k^\alpha &= \frac{r}{r+\lambda n}u_k^\alpha + \frac{\lambda n}{r+\lambda n}\left(\frac{1}{n}v_k^\beta + \frac{k}{n}\left(p_k v_k^\alpha + (1-p_k)v_{k-1}^\alpha\right) + \frac{n-k-1}{n}\left(p_{k+1}v_{k+1}^\alpha + (1-p_{k+1})v_k^\alpha\right)\right) \\
&\geq \frac{r}{r+\lambda n}u_k^\alpha + \frac{\lambda n}{r+\lambda n}\left(\frac{1}{n}v_k^\beta + \frac{k}{n}v_{k-1}^\alpha + \frac{n-k-1}{n}\left(p_{k+1}v_{k+1}^\alpha + (1-p_{k+1})v_k^\alpha\right)\right) \\
&\geq \frac{r}{r+\lambda n}u_k^\alpha + \frac{\lambda n}{r+\lambda n}\left(\frac{1}{n}v_k^\beta + \frac{k}{n}v_{k-1}^\alpha + \frac{n-k-1}{n}v_k^\alpha\right),
\end{aligned}
$$

where the last inequality obtains from the above observation. Rearranging, we obtain the desired inequality:

$$v_k^\alpha \geq \frac{r}{r+\lambda(k+1)}u_k^\alpha + \frac{\lambda(k+1)}{r+\lambda(k+1)}\left(\frac{1}{k+1}v_k^\beta + \frac{k}{k+1}v_{k-1}^\alpha\right).$$

$\square$

A.4 Proof of the Second Part of Proposition 5.4

Proof. Consider the strategy profile in which all players play β with probability one at every state. Denote by v_k^β the continuation value for a player currently playing β and k others currently playing α, conditional on all players playing β in the future. Similarly define v_k^α as the continuation value for a player currently playing α and k others are playing α conditional on all players playing β in the future. We show that $v_k^\beta > v_k^\alpha$ for all k when r is sufficiently close to 0.

Let us define the following expression:

$$\Gamma_k = \prod_{i=1}^{k}(i\lambda + r).$$

Now note that for every $k = 0, 1, \ldots, n-1$,

$$v_k^\beta = \frac{r}{(k+1)\lambda+r}u_k^\beta + \frac{(k+1)\lambda}{(k+1)\lambda+r}\left(\frac{1}{k+1}v_k^\beta + \frac{k}{k+1}v_{k-1}^\beta\right),$$

$$v_k^\alpha = \frac{r}{(k+1)\lambda+r}u_k^\alpha + \frac{(k+1)\lambda}{(k+1)\lambda+r}\left(\frac{1}{k+1}v_k^\beta + \frac{k}{k+1}v_{k-1}^\alpha\right).$$

Using the above, we can write:

$$v_{k+1}^\beta - v_{k+1}^\alpha = \frac{r}{(k+2)\lambda+r}(u_{k+1}^\beta - u_{k+1}^\alpha) + \frac{k\lambda}{(k+1)\lambda+r}(v_k^\beta - v_k^\alpha)$$

Then it is easy to show that

$$\frac{v_k^\beta - v_k^\alpha}{r} = \sum_{i=0}^{k}\frac{k!}{i!}\frac{\Gamma_i\lambda^{k-i}}{\Gamma_{k+1}}(u_i^\beta - u_i^\alpha).$$

30

Now note that the right hand side of the above inequality evaluated at exactly $r = 0$ is equal to

$$\sum_{i=0}^{k}(u_i^\beta - u_i^\alpha) > 0.$$

Thus we can choose $r^* > 0$ such that for all $r < r^*$ and for all k,

$$\frac{v_k^\beta - v_k^\alpha}{r} > 0.$$

But this means that for all k and all $r < r^*$,

$$v_k^\beta > v_k^\alpha.$$

Thus for all $r < r^*$, the pure strategy profile in which all players play β at all histories is indeed an equilibrium of the game. $\qquad\square$

B Supplementary Appendix: Not for Publication

B.1 A Formal Definition of Strategies

To define strategies, first we define a probability space $(\Omega, P, \mathcal{F})$, with a sigma-algebra $\mathcal{F}$ and a probability measure P. Define for every i, a pair of stochastic processes $X_i : [0, \infty) \times \Omega \to \mathbb{N}$ and $Y_i : [0, \infty) \times \Omega \to \Delta(A_i)$ defined on the probability space $(\Omega, P, \mathcal{F})$. X_i is defined to be a Poisson process with arrival rate λ while Y_i is defined with respect to X_i in the following manner: Given jump times $0 = \tau_0 < \tau_1 < \tau_2 < \dots$ determined by X_i, we define for all $t \in [\tau_k, \tau_{k+1})$, $Y_t := Z_k$ where $\{Z_k\}_{k=0}^{\infty}$ is a sequence of independently identically distributed random variables which are uniform distributions over A_i. Then define $Q_i := (X_i, Y_i)$.

Finally define $Q := (Q_1, \dots, Q_n)$ where the collection $\{Q_i\}_{i=1}^{n}$ are independent and let $\mathcal{F}^* := \{\mathcal{F}_t^*\}_{t \geq 0}$ be the natural filtration generated by the stochastic process Q. Then we define a strategy for player i formally to be a stochastic process $\sigma_i : [0, \infty) \times \Omega \to \Delta(A_i)$ that is adapted to the filtration $\mathcal{F}^*$.

Note that the above formulation of strategies assumes that arrivals are publicly observed, with the interpretation that $\sigma_i(t, \omega)$ is the mixed action that player i would choose given information $\mathcal{F}_t$ at time t if she indeed obtained an arrival at time t.

B.2 2-Player Coordination Games

First note that Lemma 4.1 and Lemma 4.2 generalize to environments with asymmetric arrival rates with minor modifications. With this observation, we prove the following lemma which immediately implies Lemma 4.3 and Proposition 6.4.

Lemma B.1. *The necessary and sufficient conditions for an equilibrium with (β) to be an absorbing state is either*
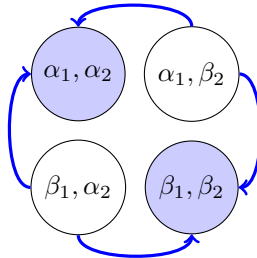
$$\frac{r}{\lambda_{-i} + r} u_i(\alpha_i, \beta_{-i}) + \frac{\lambda_{-i}}{\lambda_{-i} + r} u_i(\alpha) \leq u_i(\beta) \ \ \forall \ i = 1, 2$$

or

$$u_i(\beta_i, \alpha_{-i}) + \frac{\lambda_{-i}}{\lambda_i + r} u_i(\beta) \geq u_i(\alpha) + \frac{\lambda_{-i}}{\lambda_i + r} u_i(\alpha_i, \beta_{-i}) \ \ \forall \ i = 1, 2.$$

Proof of Lemma B.1. Before we begin, let us define $\Lambda = \lambda_1 + \lambda_2$ to simplify notation.

Case 1: Myopic Best Response

In order for this strategy profile to constitute an MPE, the indicated actions have to be best responses for when the opponent is currently playing α, and when the opponent is currently playing β. Note that the first requirement holds for any parameter specification. The condition for the second one is:

$$\frac{r}{\lambda_{-i} + r} u_i(\alpha_i, \beta_{-i}) + \frac{\lambda_{-i}}{\lambda_{-i} + r} u_i(\alpha) \quad \leq \quad u_i(\beta) \ \ \forall \ i = 1, 2.$$

Case 2: Pure Strategy, β Unique Absorbing State
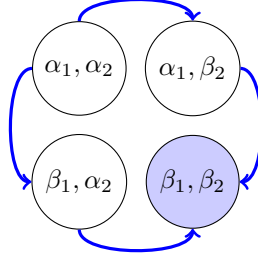
In this strategy profile, it is trivially a best response for each player i to play β_i whenever the opponent's currently played action is β_{-i}. The displayed strategies are optimal if and only if, when the other player is currently playing α, players prefer to move out of state (α) first (otherwise at least one player would like to deviate at this state):

$$u_i(\beta_i, \alpha_{-i}) + \frac{\lambda_{-i}}{\lambda_i + r} u_i(\beta) \quad \geq \quad u_i(\alpha) + \frac{\lambda_{-i}}{\lambda_i + r} u_i(\alpha_i, \beta_{-i}) \ \ \forall \ i = 1, 2.$$

Case 3: Mixed Strategy, β Unique Absorbing State
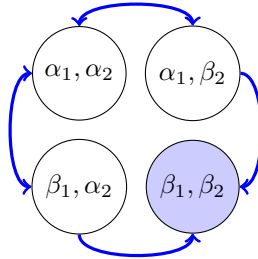
The condition for this strategy profile to constitute an MPE is:

$$u_i(\beta_i, \alpha_{-i}) + \frac{\lambda_{-i}}{\lambda_i + r} u_i(\beta) \quad \geq \quad u_i(\alpha) + \frac{\lambda_{-i}}{\lambda_i + r} u_i(\alpha_i, \beta_{-i}) \ \ \forall \ i = 1, 2.$$

Note that given that players mix when the opponent is currently playing α_{-i}, by Lemma 4.1, β must be an absorbing state. So the only necessary and sufficient conditions for the above dynamics to be an equilibrium again concern the incentives of players whenever the opponent's currently played action is α_{-i}.

This gives us sufficient conditions for a mixed equilibrium to exist. These inequalities are also necessary for the following reason. If such an equilibrium exists, there exists p_1 and p_2 such that for every i,

$$\frac{r}{\lambda_{-i}+r}u_i(\beta_i,\alpha_{-i}) + \frac{\lambda_{-i}}{\lambda_{-i}+r}u_i(\beta) = v_i(\beta_i,\alpha_{-i}) = v_i(\alpha)$$

$$= \frac{r}{\lambda_{-i}+r}u_i(\alpha) + \frac{\lambda_{-i}}{\lambda_{-i}+r}\left(p_{-i}v_i(\alpha) + (1-p_{-i})v_i(\alpha_i,\beta_{-i})\right).$$

This means that for each i, $v_i(\alpha) > v_i(\alpha_i,\beta_{-i})$. Therefore,

$$v_i(\alpha) \geq \frac{r}{\lambda_{-i}+r}u_i(\alpha) + \frac{\lambda_{-i}}{\lambda_{-i}+r}v_i(\alpha_i,\beta_{-i})$$

$$v_i(\alpha_i,\beta_{-i}) = \frac{r}{\Lambda+r}u_i(\alpha_i,\beta_{-i}) + \frac{\lambda_i}{\Lambda+r}v_i(\beta) + \frac{\lambda_{-i}}{\Lambda+r}\left(p_{-i}v_i(\alpha) + (1-p_{-i})v_i(\alpha_i,\beta_{-i})\right)$$

$$\geq \frac{r}{\Lambda+r}u_i(\alpha_i,\beta_{-i}) + \frac{\lambda_i}{\Lambda+r}v_i(\beta) + \frac{\lambda_{-i}}{\Lambda+r}v_i(\alpha_i,\beta_{-i})$$

Thus

$$\frac{r}{\lambda_{-i}+r}u_i(\beta_i,\alpha_{-i}) + \frac{\lambda_{-i}}{\lambda_{-i}+r}u_i(\beta) \geq \frac{r}{\lambda_{-i}+r}u_i(\alpha) + \frac{\lambda_{-i}}{\lambda_{-i}+r}\left(\frac{r}{\lambda_i+r}u_i(\alpha_i,\beta_{-i}) + \frac{\lambda_i}{\lambda_i+r}v_i(\beta)\right).$$

Rearranging we obtain the desired inequalities.

Therefore the necessary and sufficient conditions for the existence an equilibrium with (β) as an absorbing state is either

$$\frac{r}{\lambda_{-i}+r}u_i(\alpha_i,\beta_{-i}) + \frac{\lambda_{-i}}{\lambda_{-i}+r}u_i(\alpha) \leq u_i(\beta) \ \ \forall\, i=1,2;$$

or

$$u_i(\beta_i,\alpha_{-i}) + \frac{\lambda_{-i}}{\lambda_i+r}u_i(\beta) \geq u_i(\alpha) + \frac{\lambda_{-i}}{\lambda_i+r}u_i(\alpha_i,\beta_{-i}) \ \ \forall\, i=1,2.$$

□

B.2.1   Proof of Proposition 6.5

Proof. We can rewrite the conditions required for the selection of α as the unique absorbing state in all MPE as follows:

$$\frac{\lambda_{-i}}{\lambda_{-i}+r} > p_i \text{ for some } i,$$

$$\frac{u_i(\alpha) - u_i(\beta_i,\alpha_{-i})}{u_i(\beta) - u_i(\alpha_1,\beta_{-i})} > \frac{\lambda_{-i}}{\lambda_i+r} \text{ for some } i.$$

Now given $i,j \in \{1,2\}$, for both of the following inequalities to hold for some open set of discount rates r,

$$\frac{\lambda_{-j}}{\lambda_{-j}+r} > p_j$$

$$\frac{u_i(\alpha) - u_i(\beta_i,\alpha_{-i})}{u_i(\beta) - u_i(\alpha_1,\beta_{-i})} > \frac{\lambda_{-i}}{\lambda_i+r},$$

34

we need

$$\frac{u_i(\alpha) - u_i(\beta_i, \alpha_{-i})}{u_i(\beta) - u_i(\alpha_1, \beta_{-i})} > \frac{\lambda_{-i}}{\lambda_i + r^*}$$

for $r^* = \frac{\lambda_{-j}(1-p_j)}{p_j}$. Substituting for r^*, we obtain:

$$\frac{u_i(\alpha) - u_i(\beta_i, \alpha_{-i})}{u_i(\beta) - u_i(\alpha_1, \beta_{-i})} > \frac{p_j \lambda_{-i}}{p_j \lambda_i + \lambda_{-j}(1 - p_j)}.$$

Thus there exists an open set of discount rates in which α is selected as the unique absorbing state in all MPE if and only if:

$$\frac{u_i(\alpha) - u_i(\beta_i, \alpha_{-i})}{u_i(\beta) - u_i(\alpha_1, \beta_{-i})} > \min_{j=1,2} \left\{ \frac{p_j \lambda_{-i}}{p_j \lambda_i + \lambda_{-j}(1 - p_j)} \right\}$$

$$= \min \left\{ \frac{p_i}{p_i + (1 - p_i)\frac{\lambda_{-i}}{\lambda_i}}, p_{-i}\frac{\lambda_{-i}}{\lambda_i} \right\}$$

for some i. This completes the proof. $\qquad\square$

B.3 Standard Definitions and Propositions from Markov Chain Theory

Recall the following definitions and propositions from standard Markov chain theory. Let Ω be an arbitrary finite state space and $L : \Omega \to \Omega$ a Markov transition matrix.

Definition B.2. State j is accessible from state j (abbreviated $i \to j$) if there exists some t such that $L^t_{ij} > 0$.

Definition B.3. Two distinct states i and j communicate if i is accessible from j and j is accessible from i.

Because the "communicates" relation is an equivalence relation, we can now partition the state space into communication classes:

Definition B.4. A communication class $\mathcal{C}$ of states is a non-empty set of states such that

1. for each $i, j \in \mathcal{C}$ with $i \neq j$, i communicates with j

2. and for each $i \in \mathcal{C}$, there exists no $j \notin \mathcal{C}$ such that i communicates with j.

We finally need to define *recurrence* and *transience* of a state. To do this, let f^t_i denote the probability with which the Markov chain returns to state i (when starting initially from i) for the first time at time t.

Definition B.5. A state is recurrent if and only if

$$\sum_{t=0}^{\infty} f^t_i = 1.$$

A state that is non-recurrent is called transient.

The following proposition is trivial from the definition of transience.

Proposition B.6. *Let i be a state. If there exists j such that $i \to j$ and i is not accessible from j, then i is transient.*

In other words, if i is recurrent and j is accessible from i, then i and j must belong to the same communication class. The following proposition allows us to classify each communication class into either a transient class or a recurrent class.

Proposition B.7. *For a finite state Markov chain with a communication class $\mathcal{C}$, either all states in $\mathcal{C}$ are recurrent or all states in $\mathcal{C}$ are transient.*

Finally we introduce the notion of periodicity of a state.

Definition B.8. Let i be a state. Then the *period* of i is

$$d(i) = \gcd\{t > 0 : L_{ii}^t > 0\}.$$

A state i is called *aperiodic* if $d(i) = 1$. If $d(i) > 1$, then i is called *periodic*.

As with recurrence and transience, all states in the same communication class have the same period. Therefore we can classify a communication class to be either aperiodic or periodic.

Proposition B.9. *Let $\mathcal{C}$ be a communication class of a finite state Markov chain. Then if $i, j \in \mathcal{C}$, $d(i) = d(j)$.*

C Unique Limit Outcome in SPNE

C.1 Coordination Games

Proof of Lemma 6.2. Note first that in any equilibrium, at state α, player i can guarantee himself a payoff of at least $\frac{r}{r+\lambda} u_i(\alpha) + \frac{\lambda}{r+\lambda} u_i(\alpha_i, \beta_{-i})$ by playing α_i upon any arrival. We now compute an upper bound on the payoff that player i can obtain when player i is at a history h^τ such that $\|h^\tau\| = (\beta_i, \alpha_{-i})$, where $\|h^\tau\|$ is the currently played action profile at time τ.

To do this, let $v_i(a^0)$ be the best possible player i SPNE payoff given initial state a^0. We now show that if the inequalities in the premises of the lemma hold for player i, then

$$v_i(\beta_i, \alpha_{-i}) < \frac{r}{r+\lambda} u_i(\alpha) + \frac{\lambda}{r+\lambda} u_i(\alpha_i, \beta_{-i}).$$

Note that this then immediately implies that player i must play α_i at all histories and all SPNE whenever player $-i$'s currently played action is α_{-i}.

Suppose first that $v_i(\beta_i, \alpha_{-i}) \geq v_i(\beta)$. Then

$$v_i(\beta_i, \alpha_{-i}) \leq \frac{r}{2\lambda + r} u_i(\beta_i, \alpha_{-i}) + \frac{2\lambda}{2\lambda + r} \left(\frac{1}{2} v_i(\beta_i, \alpha_{-i}) + \frac{1}{2} u_i(\alpha) \right)$$

which then implies that

$$v_i(\beta_i, \alpha_{-i}) \le \frac{r}{\lambda + r} u_i(\beta_i, \alpha_{-i}) + \frac{\lambda}{\lambda + r} u_i(\alpha) < \frac{r}{r + \lambda} u_i(\alpha) + \frac{\lambda}{r + \lambda} u_i(\alpha_i, \beta_{-i}).$$

Next let $v_i(\beta) > v_i(\beta_i, \alpha_{-i})$. Then

$$v_i(\beta_i, \alpha_{-i}) \le \frac{r}{2\lambda + r} u_i(\beta_i, \alpha_{-i}) + \frac{2\lambda}{2\lambda + r} \left(\frac{1}{2} v_i(\beta) + \frac{1}{2} u_i(\alpha) \right).$$

Furthermore

$$v_i(\beta) \le \frac{r}{2\lambda + r} u_i(\beta) + \frac{2\lambda}{2\lambda + r} \left(\frac{1}{2} v_i(\beta) + \frac{1}{2} \max\{v_i(\beta), v_i(\alpha_i, \beta_{-i})\} \right)$$

Then it is easy to see that if $v_i(\beta) \ge v_i(\alpha_i, \beta_{-i})$ then $v_i(\beta) \le u_i(\beta)$ and $v_i(\beta) \le v_i^\beta$ if $v_i(\beta) < v_i(\alpha_i, \beta_{-i})$. Therefore

$$v_i(\beta) \le \max\{u_i(\beta), v_i^\beta\}.$$

This then implies that

$$v_i(\beta_i, \alpha_{-i}) \le \frac{r}{2\lambda + r} u_i(\beta) + \frac{2\lambda}{2\lambda + r} \left(\frac{1}{2} u_i(\alpha) + \frac{1}{2} \max\{u_i(\beta), v_i^\beta\} \right)$$

$$< \frac{r}{2\lambda + r} u_i(\alpha) + \frac{\lambda}{\lambda + r} u_i(\alpha_i, \beta_{-i}).$$

This then proves that player i plays α_i whenever the opponent's currently played action is α_{-i} in all SPNE. But then because α is Pareto dominant, player $-i$ must also play α_{-i} whenever player i's currently played action is α_i. □

C.2   BoS Games

As in coordination games, it remains to provide sufficient conditions for players to play α_i at all histories and all SPNE given that the currently played action of player $-i$ is α_{-i}. Because α is the best payoff for player 1, we provide sufficient conditions under which player 2 chooses action α_2 in all SPNE whenever the currently played action of player 1 is α_1. Having established this property, then it is easy to see that player 1 must choose α_1 whenever player 2's currently played acton is α_2 in all SPNE. Then with the help of Lemma 6.1, we can establish that the unique limit outcome in all SPNE must be α.[26]

To shorten notation define

$$v_\alpha = \frac{r}{\lambda + r} u_2(\alpha) + \frac{\lambda}{\lambda + r} \left(\frac{r}{\lambda + r} u_2(\alpha_2, \beta_1) + \frac{\lambda}{\lambda + r} u_2(\beta) \right).$$

Intuitively this is the payoff that player 2 would receive if play began at state (α_1, α_2) and player 1 always played β_1 while player 2 played a myopic best response strategy. The payoff to this strategy will play a key role in determining whether or not (α_1, α_2) will indeed be selected as the unique limit outcome in all SPNE.

---

[26]Exactly symmetric conditions can be provided for β to be the unique limit outcome in all SPNE.

We will split the parameter space into two regions ($u_2(\alpha_2, \beta_1) \geq u_2(\alpha)$ and $u_2(\alpha_2, \beta_1 < u_2(\alpha))$), and first consider the easier case in which $u_2(\alpha_2, \beta_1) \geq u_2(\alpha)$.

Lemma C.1. *Suppose $u_2(\alpha_2, \beta_1) \geq u_2(\alpha)$ and that:*

$$u_2(\alpha) > \frac{r}{\lambda + r} u_2(\beta_2, \alpha_1) + \frac{\lambda}{\lambda + r} u_2(\beta), \text{ and}$$

$$u_2(\alpha) > \frac{r}{2\lambda + r} u_2(\beta_2, \alpha_1) + \frac{2\lambda}{2\lambda + r} \left(\frac{1}{2} u_2(\beta) + \frac{1}{2} v_\alpha \right),$$

then in any SPNE, player 2 must play α_2 whenever player 1's currently played action is α_1.

Proof. Note first that in any equilibrium, at state (α_1, α_2), player 2 can guarantee himself a payoff of at least u by playing α_1 upon any arrival. We now compute an upper bound on the payoff that player 2 can obtain at state (α_1, β_2).

To do this, let $v(\alpha)$, $v(\alpha_1, \beta_2)$, and $v(\beta_1, \alpha_2)$ the best possible player 2 SPNE payoffs at the respective states. Note that we do not require that $v(\alpha)$, $v(\alpha_1, \beta_2)$, and $v(\beta_1, \alpha_2)$ be the payoffs at their respective states in the same SPNE. We use these values to bound the continuation payoffs at the state (α_1, β_2) in any SPNE.

Note that

$$v(\alpha_1, \beta_2) \leq \frac{r}{2\lambda + r} u_2(\beta_2, \alpha_1) + \frac{2\lambda}{2\lambda + r} \left(\frac{1}{2} \max\{v(\alpha_1, \beta_2), v(\alpha)\} + \frac{1}{2} u_2(\beta) \right).$$

Suppose first that $v(\alpha_1, \beta_2) \geq v(\alpha)$. Then we have

$$v(\alpha_1, \beta_2) \leq \frac{r}{\lambda + r} u_2(\beta_2, \alpha_1) + \frac{\lambda}{\lambda + r} u_2(\beta) < u_2(\alpha).$$

Therefore the best possible payoff is less than the payoff that player 2 could guarantee at the state (α). Thus we have shown that player 2 must play α_2 with probability one upon arrival when the currently played action of player 1 is α_1 in any SPNE.

Suppose now that $v(\alpha) > v(\alpha_1, \beta_2)$. Let us bound the payoff $v(\alpha)$:

$$v(\alpha) \leq \frac{r}{2\lambda + r} u_2(\beta_2, \alpha_1) + \frac{2\lambda}{2\lambda + r} \left(\frac{1}{2} \max\{v(\alpha), v(\alpha_1, \beta_2)\} + \frac{1}{2} \max\{v(\alpha), v(\beta_1, \alpha_2)\} \right)$$

$$= \frac{r}{2\lambda + r} u_2(\alpha) + \frac{2\lambda}{2\lambda + r} \left(\frac{1}{2} v(\alpha) + \frac{1}{2} \max\{v(\alpha), v(\beta_1, \alpha_2)\} \right).$$

Let us assume first that $v(\alpha) \geq v(\beta_1, \alpha_2)$. Then the above implies that $v(\alpha) \leq u_2(\alpha)$. However this means that

$$v(\alpha_1, \beta_2) \leq \frac{r}{2\lambda + r} u_2(\beta_2, \alpha_1) + \frac{2\lambda}{2\lambda + r} \left(\frac{1}{2} u_2(\alpha) + \frac{1}{2} u_2(\beta) \right) < u_2(\alpha).$$

Thus again it must be that player 2 plays α_2 with probability one upon arrival in any SPNE when the currently played action of player 1 is α_1.

Finally assume that $v(\alpha) > v(\alpha_1, \beta_2)$ and $v(\beta_1, \alpha_2) > v(\alpha)$. Then

$$v(\beta_1, \alpha_2) \leq \frac{r}{2\lambda + r} u_2(\alpha_2, \beta_1) + \frac{2\lambda}{2\lambda + r} \left(\frac{1}{2} v(\beta_1, \alpha_2) + \frac{1}{2} u_2(\beta) \right).$$

This implies that

$$v(\beta_1, \alpha_2) \leq \frac{r}{\lambda + r} u_2(\alpha_2, \beta_1) + \frac{\lambda}{\lambda + r} u_2(\beta).$$

This immediately implies that

$$v(\alpha_1, \beta_2) \leq \frac{r}{2\lambda + r} u_2(\beta_2, \alpha_1) + \frac{2\lambda}{2\lambda + r} \left(\frac{1}{2} v_\alpha + \frac{1}{2} u_2(\beta) \right) < u_2(\alpha).$$

Thus in this case again, player 2 must play α_2 with probability one whenever the currently played action of player 1 is α_1 in any SPNE. $\qquad \square$

We now proceed to the case when $u_2(\alpha_2, \beta_1) < u_2(\alpha)$. Note that in this case, player 2 can no longer guarantee a payoff of $u_2(\alpha)$ at states (α) and (α_2, β_1). This complicates the analysis.

Lemma C.2. *Suppose $u_2(\alpha_2, \beta_1) < u_2(\alpha)$ and that*

$$\frac{r}{\lambda + r} u_2(\alpha) + \frac{\lambda}{\lambda + r} u_2(\alpha_2, \beta_1) > \frac{r}{\lambda + r} u_2(\beta_2, \alpha_1) + \frac{\lambda}{\lambda + r} u_2(\beta), \text{ and}$$

$$\frac{r}{\lambda + r} u_2(\alpha) + \frac{\lambda}{\lambda + r} u_2(\alpha_2, \beta_1) > \frac{r}{2\lambda + r} u_2(\beta_2, \alpha_1) + \frac{2\lambda}{2\lambda + r} \left(\frac{1}{2} u_2(\beta) + \frac{1}{2} \max\{v_\alpha, u_2(\alpha)\} \right).$$

Then in any SPNE, player 1 must play α_2 in any history at which player 2 is currently playing α_1.

Proof. The proof follows along the same lines as the previous lemma. Define as in the previous lemma the payoffs $v(\alpha)$, $v(\alpha_1, \beta_2)$, and $v(\beta_1, \alpha_2)$. We show that it must always be the case that

$$v(\alpha_1, \beta_2) < \frac{r}{\lambda + r} u_2(\alpha) + \frac{\lambda}{\lambda + r} u_2(\alpha_2, \beta_1).$$

Suppose first that $v(\alpha_1, \beta_2) \geq v(\alpha)$. Then

$$v(\alpha_1, \beta_2) \leq \frac{r}{2\lambda + r} u_2(\beta_2, \alpha_1) + \frac{2\lambda}{2\lambda + r} \left(\frac{1}{2} v(\alpha_1, \beta_2) + \frac{1}{2} u_2(\beta) \right).$$

This then implies that

$$v(\alpha_1, \beta_2) \leq \frac{r}{\lambda + r} u_2(\beta_2, \alpha_1) + \frac{\lambda}{\lambda + r} u_2(\beta) < \frac{r}{\lambda + r} u_2(\alpha) + \frac{\lambda}{\lambda + r} u_2(\alpha_2, \beta_1).$$

Next let $v(\alpha) > v(\alpha_1, \beta_2)$. Thus

$$v(\alpha) \leq \frac{r}{2\lambda + r} u_2(\alpha) + \frac{2\lambda}{2\lambda + r} \left(\frac{1}{2} v(\alpha) + \frac{1}{2} \max\{v(\alpha), v(\beta_1, \alpha_2)\} \right).$$

Suppose that $v(\alpha) \geq v(\beta_1, \alpha_2)$. Then the above implies that

$$v(\alpha) \leq u_2(\alpha).$$

39

	α_2	β_2
α_1	$2, 1$	$\varepsilon, -4$
β_1	$1, 0$	$1, 2$

Figure 2: Stage Game

This then means that

$$v(\alpha_1, \beta_2) \leq \frac{r}{2\lambda + r} u_2(\beta_2, \alpha_1) + \frac{2\lambda}{2\lambda + r} \left(\frac{1}{2} u_2(\alpha) + \frac{1}{2} u_2(\beta) \right) < \frac{r}{\lambda + r} u_2(\alpha) + \frac{\lambda}{\lambda + r} u_2(\alpha_2, \beta_1).$$

Finally suppose that $v(\beta_1, \alpha_2) > v(\alpha) > v(\alpha_1, \beta_2)$. Then it is easy to see that

$$v(\alpha_1, \beta_2) \leq \frac{r}{2\lambda + r} u_2(\beta_2, \alpha_1) + \frac{2\lambda}{2\lambda + r} \left(\frac{1}{2} u_2(\beta) + \frac{1}{2} v_\alpha \right) < \frac{r}{\lambda + r} u_2(\alpha) + \frac{\lambda}{\lambda + r} u_2(\alpha_2, \beta_1).$$

Thus we have shown that $v(\alpha_1, \beta_2) < \frac{r}{\lambda+r} u_2(\alpha) + \frac{\lambda}{\lambda+r} u_2(\alpha_2, \beta_1)$. Therefore since player 2 can always guarantee a payoff of

$$\frac{r}{\lambda + r} u_2(\alpha) + \frac{\lambda}{\lambda + r} u_2(\alpha_2, \beta_1)$$

at state (α) by always playing α_2 in the future, player 2 must always play α_2 when player 1 is currently playing α_1. $\qquad\square$

C.3   Uniqueness of Strategies in SPNE

First we show that uniqueness of a limit outcome in SPNE does not imply uniqueness of SPNE. Consider the following example. Suppose that the stage game is the following:

with $\varepsilon > 0$ and let $\lambda = 1$ and $r = \frac{1}{2}$. We choose $\varepsilon > 0$ so that

$$\frac{1}{2} \cdot 1 + \frac{1}{4} \cdot 1 + \frac{1}{4} \cdot 2 > \frac{1}{2} \cdot 2 + \frac{1}{2} \varepsilon > 1.$$

Due to the lemma above, we know that in any SPNE both players play α_i whenever the opponent is currently playing α_{-i}.

In the above equilibrium, there exists a Markovian equilibrium with the following dynamics.

- Both players play α_i with probability one whenever the opponent's prepared action is α_{-i}.

- player 1 plays α_1 with probability 9/11 whenever the opponent's prepared actions is β_2.

- player 2 plays α_2 with probability $3\varepsilon/(2 - \varepsilon)$ whenever the opponent's prepared action is β_1.

One can easily check that these strategies constitute a SPNE. Furthermore it is easy to see that the payoffs at each state must be the following:

$$v(\alpha_1, \alpha_2) = (2, 1)$$
$$v(\alpha_1, \beta_2) = \left(\frac{1}{2}\varepsilon + 1, -\frac{4}{3}\right)$$
$$v(\beta_1, \alpha_2) = \left(\frac{1}{3}\varepsilon + \frac{4}{3}, \frac{1}{2}\right)$$
$$v(\beta_1, \beta_2) = \left(\frac{1}{2}\varepsilon + 1, \frac{1}{2}\right).$$

In fact the above is the unique MPE for this set of parameters.

However there also exists the following non-Markovian equilibrium that makes use of the above Markovian strategy upon deviations.

- Both players play α_i upon arrival as long as everyone has thus far played α_j upon arrival.

- If at some point, a deviation occurred, then players play according to the Markovian equilibrium above.

Note that this non-Markovian equilibrium raises the payoff of player 1 at an initial state (β_1, β_2) beyond what he can achieve in the Markovian equilibrium. This is because for player 1, going through the path

$$(\beta_1, \beta_2) \to (\beta_1, \alpha_2) \to (\alpha_1, \alpha_2)$$

on the way from (β_1, β_2) to (α_1, α_2) is better than moving to α_1 and using the path

$$(\alpha_1, \beta_2) \to (\alpha_1, \alpha_2)$$

as long as player 2 is moving to α_2 with probability one upon arrival. In the above strategy profile when the state is (β_1, β_2), either player 1 arrives first giving him a payoff equal to the Markovian payoff at state (β_1, β_2) of $V_1(\beta_1, \beta_2)$ or with strictly positive probability player 2 arrives first giving a payoff strictly better than $V_1(\beta_1, \beta_2)$. Note that the "punishment" via reversion to the Markovian equilibrium is necessary since otherwise, players would find it beneficial to simply wait at (β_1, β_2) for the other player to arrive and move out for him instead of moving out himself.

We conclude this section by providing sufficient conditions for uniqueness of SPNE strategies in addition to uniqueness of a limit outcome in SPNE.

Proposition C.3. *Suppose that*

$$\frac{r}{\lambda + r}u_1(\alpha_1, \beta_2) + \frac{\lambda}{\lambda + r}u_1(\alpha) > u_1(\beta),$$
$$\min\left\{u_2(\alpha), \frac{r}{\lambda + r}u_2(\alpha) + \frac{\lambda}{\lambda + r}u_2(\alpha_2, \beta_1)\right\} > \frac{r}{\lambda + r}u_2(\beta_2, \alpha_1) + \frac{\lambda}{\lambda + r}u_2(\beta), \text{ and}$$
$$\min\left\{u_2(\alpha), \frac{r}{\lambda + r}u_2(\alpha) + \frac{\lambda}{\lambda + r}u_2(\alpha_2, \beta_1)\right\} > \frac{r}{2\lambda + r}u_2(\beta_2, \alpha_1) + \frac{2\lambda}{2\lambda + r}\left(\frac{1}{2}u_2(\beta) + \frac{1}{2}\max\{v_\alpha, u_2(\alpha)\}\right).$$

41

Suppose further that

$$u_1(\alpha_1, \beta_2) + \frac{\lambda}{\lambda + r} u_1(\alpha) > u_1(\beta) + \frac{\lambda}{\lambda + r} u_1(\beta_1, \alpha_2)$$

or

$$u_2(\alpha_2, \beta_1) + \frac{\lambda}{\lambda + r} u_2(\alpha) < u_2(\beta) + \frac{\lambda}{\lambda + r} u_2(\beta_2, \alpha_1).$$

Then there exists a unique SPNE.

Proof. We have already shown that the first two inequalities imply that player 2 must play α_2 at any history when player 1 is currently playing α_1 and player 1 must play α_1 at states when player 2 is currently playing α_2. We proceed again as in the previous proofs.

Suppose that the first inequality holds. The case for which the second inequality holds proceeds along the same lines and so we omit the proof. We will show that

$$v_1(\beta) < \frac{r}{\lambda + r} u_1(\alpha_1, \beta_2) + \frac{\lambda}{\lambda + r} u_1(\alpha).$$

Suppose first that $v_1(\beta) \leq v_1(\beta_1, \alpha_2)$. Then

$$v_1(\beta_1, \alpha_2) \leq \frac{r}{\lambda + r} u_1(\beta_1, \alpha_2) + \frac{\lambda}{\lambda + r} u_1(\beta).$$

If $v_1(\beta) \geq v_1(\alpha_1, \beta_2)$ then this implies that

$$v_1(\beta) \leq \frac{r}{\lambda + r} u_1(\beta) + \frac{\lambda}{\lambda + r} \left(\frac{r}{\lambda + r} u_1(\beta_1, \alpha_2) + \frac{\lambda}{\lambda + r} u_1(\alpha) \right).$$

But clearly the above is strictly less than

$$\frac{r}{\lambda + r} u_1(\alpha_1, \beta_2) + \frac{\lambda}{\lambda + r} u_1(\alpha).$$

Suppose next that $v_1(\beta) \leq v_1(\beta_1, \alpha_2)$ and $v_1(\beta) < v_1(\alpha_1, \beta_2)$. Again we have

$$v_1(\beta_1, \alpha_2) \leq \frac{r}{\lambda + r} u_1(\beta_1, \alpha_2) + \frac{\lambda}{\lambda + r} u_1(\alpha).$$

Additionally

$$v_1(\alpha_1, \beta_2) \leq \frac{r}{\lambda + r} u_1(\alpha_1, \beta_2) + \frac{\lambda}{\lambda + r} u_1(\alpha).$$

This then implies that

$$
\begin{aligned}
v_1(\beta) &\leq \frac{r}{2\lambda + r} u_1(\beta) + \frac{2\lambda}{2\lambda + r} \left(\frac{1}{2} v_1(\alpha_1, \beta_2) + \frac{1}{2} v_1(\beta_1, \alpha_2) \right) \\
&< \frac{r}{\lambda + r} u_1(\alpha_1, \beta_2) + \frac{\lambda}{\lambda + r} u_1(\alpha).
\end{aligned}
$$

Finally let us assume that $v_1(\beta) > v_1(\beta_1, \alpha_2)$. Therefore

$$v_1(\beta) \leq \frac{r}{2\lambda + r} u_1(\beta) + \frac{2\lambda}{2\lambda + r} \left(\frac{1}{2} v_1(\beta) + \frac{1}{2} \max\{v_1(\beta), v_1(\alpha_1, \beta_2)\} \right).$$

42

One can show that this is less than $\frac{r}{\lambda+r}u_1(\alpha_1\beta_2) + \frac{\lambda}{\lambda+r}u_1(\alpha)$.

Thus we have shown that in all of these cases,

$$v_1(\beta) < \frac{r}{\lambda+r}u_1(\alpha_1,\beta_2) + \frac{\lambda}{\lambda+r}u_1(\alpha).$$

However the right hand side of the above inequality denotes the payoff that player 1 can guarantee when she transitions play to (α_1, β_2). Therefore it must be that in any SPNE, player 1 must play α_1 with probability one at any history in which player 2 is currently playing β_2. This establishes the claim. $\qquad\square$

D Battle of Sexes Games: MPE

Case 1: Myopic Best Response
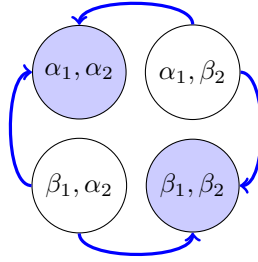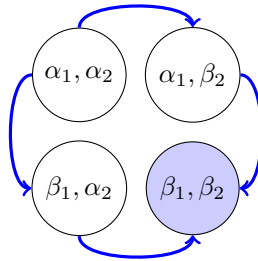

In this strategy profile, it is trivially a best response for player 1 to play α_1 whenever player 2's current action is α_2. Furthermore it is a best response for player 1 to play β_1 whenever player 2's current action β_2 if and only if

$$u_1(\beta) \geq \frac{r}{r+\lambda}u_1(\alpha_1,\beta_2) + \frac{\lambda}{r+\lambda}u_1(\alpha).$$

Symmetric arguments show that incentive compatibility for player 2 is satisfied if and only if

$$u_2(\alpha) \geq \frac{r}{r+\lambda}u_2(\beta_2,\alpha_1) + \frac{\lambda}{r+\lambda}u_2(\beta).$$

*Case 2: Pure Strategy, β Unique Absorbing State*

In this strategy profile, it is trivially a best response for each player i to play β_i whenever the opponent's currently played action is β_{-i}. The displayed strategies are optimal if and only if, when the other player is

currently playing α, players prefer to move out of state (α) first (otherwise at least one player would like to deviate at this state):

$$u_i(\beta_i, \alpha_{-i}) + \frac{\lambda}{r+\lambda} u_i(\beta) \geq u_i(\alpha) + \frac{\lambda}{r+\lambda} u_i(\alpha_i, \beta_{-i}) \ \forall i = 1, 2.$$

Case 3: Asymmetric Pure Strategy, β Unique Absorbing State
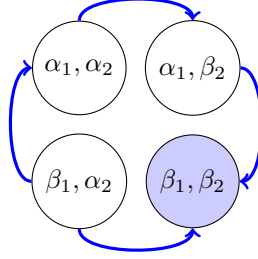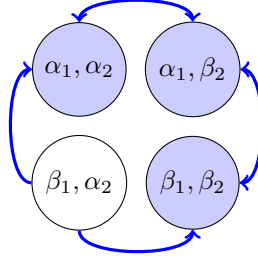

It is clear that both players prefer to play β_i whenever the opponent is current playing β_{-i}. Thus it remains to check incentives when the opponent is currently playing α_{-i}. These incentives correspond to the following inequalities:

$$u_1(\beta_1, \alpha_2) + \frac{\lambda}{r+\lambda} u_1(\beta) \leq u_1(\alpha) + \frac{\lambda}{r+\lambda} u_1(\alpha_1, \beta_2),$$
$$u_2(\alpha) \leq \frac{r}{r+\lambda} u_2(\beta_2, \alpha_1) + \frac{\lambda}{r+\lambda} u_2(\beta).$$

Case 4: Mixed Strategy, Absorbing Cycle

The above equilibrium exists if and only if

$$\frac{r}{r+\lambda} u_1(\alpha_1, \beta_2) + \frac{\lambda}{r+\lambda} u_1(\alpha) \geq u_1(\beta)$$
$$\frac{r}{r+\lambda} u_2(\beta_2, \alpha_1) + \frac{\lambda}{r+\lambda} u_2(\beta) \geq u_2(\alpha).$$

By Lemma 4.1, given that player 1 mixes when player 2 is currently playing β_2, player 1 strictly prefers to play α_1 whenever player 2 is currently playing α_2. Similarly, player 2 strictly prefers to play β_2 whenever player 1 is currently playing β_1. Thus the only necessary and sufficient conditions for the above dynamics to be an equilibrium concern the incentives of players when the opponent is playing the less-preferred action.

That the inequalities are sufficient for incentive compatibility is quite clear. They are also necessary for the following reason. Suppose that an equilibrium of the above form exists. Then

$$v_1(\alpha_1, \beta_2) = v_1(\beta) = u_1(\beta),$$

$$v_2(\beta_2, \alpha_1) = v_2(\alpha) = u_2(\alpha).$$

Therefore

$$
\begin{aligned}
u_1(\beta) = v_1(\alpha_1, \beta_2) &= \frac{r}{r+\lambda} u_1(\alpha_1, \beta_2) + \frac{\lambda}{r+\lambda} \left(p_2 v_1(\alpha) + (1-p_2) v_1(\alpha_1, \beta_2) \right) \\
&\leq \frac{r}{r+\lambda} u_1(\alpha_1, \beta_2) + \frac{\lambda}{r+\lambda} u_1(\alpha), \\
u_2(\alpha) = v_2(\beta_2, \alpha_1) &= \frac{r}{r+\lambda} u_2(\beta_2, \alpha_1) + \frac{\lambda}{r+\lambda} \left(p_1 v_2(\beta) + (1-p_2) v_2(\beta_2, \alpha_1) \right) \\
&\leq \frac{r}{r+\lambda} u_2(\beta_2, \alpha_1) + \frac{\lambda}{r+\lambda} u_2(\beta).
\end{aligned}
$$

Case 5: Mixed Strategy, β Unique Absorbing State
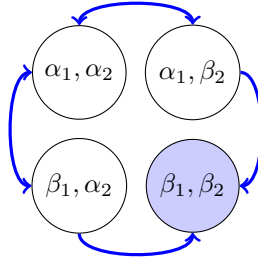
Finally the necessary and sufficient conditions for the above equilibrium to exist is given by the following set of inequalities:

$$u_1(\beta_1, \alpha_2) + \frac{\lambda}{r+\lambda} u_1(\beta) \geq u_1(\alpha) + \frac{\lambda}{r+\lambda} u_1(\alpha_1, \beta_2)$$

and either

$$u_2(\alpha) \geq \frac{r}{r+\lambda} u_2(\beta_2, \alpha_1) + \frac{\lambda}{r+\lambda} u_2(\beta)$$

$$u_2(\beta_2, \alpha_1) + \frac{\lambda}{r+\lambda} u_2(\beta) \geq u_2(\alpha) + \frac{\lambda}{r+\lambda} u_2(\alpha_2, \beta_1),$$

or

$$u_2(\alpha) \leq \frac{r}{r+\lambda} u_2(\beta_2, \alpha_1) + \frac{\lambda}{r+\lambda} u_2(\beta)$$

$$u_2(\beta_2, \alpha_1) + \frac{\lambda}{r+\lambda} u_2(\beta) \leq u_2(\alpha) + \frac{\lambda}{r+\lambda} u_2(\alpha_2, \beta_1)$$

This analysis together concludes the proof.