Asynchronous Choice in Battle of the Sexes Games: Unique Equilibrium Selection for Intermediate Levels of Patience

Attila Ambrus*
Duke University, Department of Economics

Yuhta Ishii†
Harvard University, Department of Economics

September 2011

Abstract

This paper shows that in infinite-horizon asynchronous-move battle of the sexes games, for a full-dimensional set of payoff specifications, there is an intermediate range of discount rates for which every subgame perfect Nash equilibrium induces the same unique limit outcome. The latter is one of the pure Nash equilibria of the stage game, and play is absorbed in it in finite time with probability 1. We fully characterize the set of game specifications for which there is a unique limit outcome in any Markov perfect equilibrium, but show by example that these conditions are not sufficient for unique selection in the limit in all subgame perfect Nash equilibria. We also provide sufficient conditions for the existence of a unique limit outcome in any subgame perfect Nash equilibrium. Our results complement the findings of Lagunoff and Matsui (1997) and others, who show that in a class of coordination games asynchronicity of moves leads to unique equilibrium selection for high enough (as opposed to intermediate) levels of patience.

Keywords: repeated games, asynchronous moves, battle of the sexes, uniqueness of equilibrium, genericity

JEL Classification: C72, C73

*E-mail: aa231@duke.edu

†E-mail: ishii2@fas.harvard.edu

1 Introduction

Asynchronous-move repeated games have been generating considerable interest in the game theory literature since the seminal papers 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 the predictions obtained from these games are often very different than those from standard repeated games with simultaneous moves.[1] In empirical applications, as Arcidiacono et al. (2010) points out, the main advantage of asynchronous-move models is that their computational complexity 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 pointed out in Schmidt-Dengler (2006) and Einav (2010), in many situations asynchronicity of moves is a more realistic modeling assumption.

This paper analyzes battle of the sexes games (introduced by Luce and Raiffa (1957) as the simplest class of games in which players want to coordinate, but their interests differ in which outcome to coordinate on) in an infinite-horizon continuous-time framework, in which players accumulate payoffs continuously, but they can only change their actions at random discrete times, governed by independent Poisson processes.[2] Despite the stage game has two strict pure

[1]For example, while Markov perfect equilibria in standard infinitely repeated games are simply infinite repetitions of Nash equilibria of the stage game, there can be nontrivial strategic dynamics in Markov perfect equilibria of asynchronous-move games - see Maskin and Tirole (1988a, 1988b) and Lagunoff and Matsui (1997). Another feature distinguishing asynchronous games from the standard model is that expected subgame-perfect equilibrium payoffs can be strictly below players' minimax values in the stage game - see for example Takahashi and Wen (2003).

[2]This implies that the probability of two players changing their actions exactly at the same time is 0. As Maskin and Tirole (1988a) points out, these games are closely related to deterministic alternating-move games, if one restricts attention to Markovian strategies. There is a recent string of papers using similar continuous-time random-arrival models. Ambrus and Lu (2010) investigate multilateral bargaining with a deadline in such a context, while Kamada and Kandori (2011), and Calgano, Kamada, Lovo and Sugaya (2012) examine situations in which players can publicly modify their action plans before playing a normal-form game. Ambrus et al. (2012) and Kamada and Sugaya (2011) use similar frameworks to analyze on-line auctions and policy announcement games, respectively. Similar models were first considered in the macroeconomics literature, on sticky prices (Calvo (1983)) and in the search litera-

Nash equilibria (each strictly preferred by one of the players), we show that there is an intermediate range of discount rates for which for a full dimensional specification of stage game payoffs asynchronicity of moves leads play in every subgame perfect Nash equilibrium (SPNE) to the same stage game equilibrium outcome, no matter what the starting actions are.

Our result 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.[3] The main difference is that in battle of the sexes games unique selection occurs for intermediate levels of patience, not for high ones. This also implies that there is no tension between our result and the folk theorem, which was shown to hold in generic asynchronous-move games by Dutta (1995) and Yoon (1999),[4] and that unique selection is a phenomenon that is robust to any small perturbations of payoffs and discount rates.

Analyzing equilibrium possibilities for intermediate levels of patience in repeated games is a notoriously difficult task, and for this reason we only provide a set of necessary conditions, and another set of sufficient conditions, for one of the stage game equilibria be the unique limit outcome in any SPNE.

We obtain necessary conditions by providing a complete characterization of Markov perfect equilibria (MPE) in asynchronous-move battle of the sexes games. These are SPNE in which players' strategies only depend on the payoff-relevant state, meaning that at any point of time a player gets the chance to choose an action, her choice only depends on the current action choice of

ture explaining price and wage dispersion (Burdett and Judd (1983), Burdett and Mortensen (1998)).

[3]See also Farrell and Saloner (1985) and Dutta (2003) for similar results in finite-horizon asynchronous-move coordination games, and Takahashi (2005) in a class of common interest games encompassing both finite and infinite horizons. More closely related to the current paper is Calgano et al. (2012), who show that in generic finite horizon battle of the sexes games with asynchronous moves there is a unique SPNE, obtained by backward induction. The infinite horizon games we investigate in the current paper require different techniques to analyze, and indeed the results are very different.

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

the other player. As Lagunoff and Matsui (1997) show, there always exists an MPE (on mixed strategies). Hence, a necessary condition for a unique limit outcome in SPNE is that there is a unique limit outcome in MPE. Our complete characterization of MPE in these games is also of independent interest from the main point of the paper, revealing for example that one action profile (when both players play the action corresponding to their least preferred equilibrium in the stage game) cannot be part of an absorbing cycle in any SPNE, for any discount rate.[5]

We show by example that uniqueness of the limit outcome in MPE does not imply uniqueness of the limit outcome in SPNE. Hence, the above necessary conditions are not sufficient for the result. For this reason, we provide a different set of sufficient conditions for play to converge to a given player's favored Nash equilibrium outcome in the stage game, in any SPNE. We also show, by example, that these conditions still do not uniquely pin down equilibrium strategies: there can be multiple SPNE profiles, with different expected payoffs, with the feature that they all converge to the same stage game outcome. We provide an additional condition that guarantees the uniqueness of SPNE strategies as well (not just limit outcomes). Finally, we investigate that for which payoff specifications there exists a range of discount rates with unique limit outcome.

The intuition for why unique selection occurs for certain specifications is simple. If the off-equilibrium stage game payoff when both players play the strategies corresponding to their favored equilibria is particularly bad for one of the players, labeled the weak player, then for a not too high level of patience the latter player is better off giving in when getting the chance, and switch play to the other player's favored outcome, no matter what continuation play she expects. But then if the same out-of-equilibrium outcome is not as bad for the other player, labeled the strong player, and the level of impatience is not too low, the strong player can force play to ultimately switch to her favored equilibrium

[5] For foundations for MPE in asynchronous-move games see Bhaskar and Vega-Redondo (2002) and Bhaskar et al. (2012). For certain properties of MPE of these games see Haller and Lagunoff (2000, 2006).

outcome, by choosing the corresponding strategy. Actual equilibrium dynamics might be more complicated, and might require mixing at certain histories.

The above intuition also reveals why the level of patience has to be from an intermediate range for a unique limit outcome. If players are too impatient then either of the stage game equilibrium outcomes become absorbing states. And if players are too patient then neither of them can force the other one to play the strategy corresponding to the first player's favored equilibrium, and in the limit the folk theorem applies.

Which player becomes strong (for a range of discount rates) depends on both on- and off-equilibrium payoffs of the stage game. Increasing a player's payoff in off-equilibrium outcomes, as well as increasing her payoff in her favored Nash equilibrium makes her stronger. On the other hand, increasing her payoff in her dispreferred Nash equilibrium makes her weaker, as it makes it more attractive for her to stay in that outcome.

2 The model

Consider the following 2×2 game, referred to below as the stage game:

	α_2	β_2
α_1	a, u	b, v
β_1	c, w	d, x

Table 1: Stage Game

with the parameter restrictions $a > c$, $d > b$, $u > v$, $x > w$, $a > d$ and $x > u$. The restrictions imply that the game has two strict Nash equilibria, (α_1, α_2) and (β_1, β_2), with player 1 (P1) preferring the first one and player 2 (P2) preferring the second.

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. More precisely, assume that the action profile at time 0, a^0, is given exogenously. The 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 payoff in the dynamic game is then defined as $\int_0^\infty e^{-r} u_i(a^t)dt$, where $u_i()$ 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).

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 are measurable mappings from H^t to $\{\alpha_1, \beta_1\}$ and $\{\alpha_2, \beta_2\}$ respectively, 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.

We are interested in both all SPNE of the above game, and the set of MPE, in which a player's strategy when choosing an action only depends on the payoff relevant state, which is simply the action currently played by the other player. More precisely, there can be 8 different states (types of continuation histories from which the continuation game looks the same) at any point of time: (α_1, α_2), (α_1, β_2), (β_1, α_2), (β_1, β_2), $(., \alpha_2)$, $(., \beta_2)$, $(\alpha_1, .)$ and $(\beta_1, .)$. The first four correspond to points of time when neither player has an arrival (and one of four possible action profiles are played). The next two corresponding to points of time when P1 receives an arrival and P2 is locked into playing one of his two actions - these are the only histories at which P1 chooses an action. Finally, the last two correspond to points of time when P2 receives an arrival - the only histories at which P2 chooses an action. Hence, for each player there can only be two states (two types of continuation games) at times when she has to choose an action. From now on, whenever it does not cause confusion, instead of (α_1, α_2) and (β_1, β_2) we will simply write (α) and (β).

Lastly, we note that we can essentially map the game back to a model in discrete time, the following way. Fix a strategy profile in the dynamic game, and let $\omega_i^{t,j}$ denote the continuation payoff of player i in the continuation game starting at t in case the very first arrival in the game is at time t by player j. Then we can write player i's expected payoff at time 0 as:

$$\int_0^\infty \lambda e^{-2\lambda\tau}\left((1-e^{-r\tau})u_i(a^0)+e^{-r\tau}w_i^{\tau,1}\right)d\tau+\int_0^\infty \lambda e^{-2\lambda\tau}\left((1-e^{-r\tau})u_i(a^0)+e^{-r\tau}w_i^{\tau,2}\right)d\tau.$$

This simplifies to:

$$\frac{r}{2\lambda+r}u_i(a^0)+\frac{2\lambda}{2\lambda+r}\left(\frac{1}{2}\int_0^\infty(2\lambda+r)e^{-(2\lambda+r)\tau}(w_i^{\tau,1}+w_i^{\tau,2})d\tau\right).$$

Therefore effectively, the discount factor is $2\lambda/(2\lambda+r)$.

3 Example of unique limit outcome

Consider the following stage game:

	α_2	β_2
α_1	$4, 2$	$1, -4$
β_1	$-2, 3$	$2, 4$

with discount rate $r = 1$ and arrival rate $\lambda = 1$.

In this example, we show that α must always be the unique limit outcome regardless of the initial state in any SPNE. More strongly, the set of SPNE is unique and Markovian with a unique absorbing state at α.

The reason that the example induces a unique limit outcome is due to the following simple reason. We first demonstrate that P2 must always play α_2 whenever he obtains an opportunity to revise his action when P1 is currently playing α_1 in any SPNE. To show this, note first that the worst payoff that P1 can obtain in any SPNE with initial state α is $u = 2$. This is because P2 can guarantee himself a payoff of 2 by always playing α_2 at all histories, regardless of P1's strategy. We compare this to the best payoff that P2 could get in any SPNE beginning at the state (α_1, β_2). Let us denote this payoff by $V(\alpha_1, \beta_2)$.

Then we can show that this payoff must be bounded above by either

$$\frac{r}{\lambda + r}v + \frac{\lambda}{\lambda + r}x = \frac{1}{2}(-4) + \frac{1}{2}4 = 0 < 2$$

or

$$\frac{r}{2\lambda + r}v + \frac{2\lambda}{2\lambda + r}\left(\frac{1}{2}v_\alpha + \frac{1}{2}x\right) = \frac{1}{3}(-4) + \frac{2}{3}\left(\frac{1}{2}\frac{11}{4} + \frac{1}{2}4\right) = \frac{1}{12} < 2,$$

where v_α corresponds to the payoff of P2 with initial state α corresponding to the strategy profile in which P1 plays β_1 regardless of the history and P2 plays the action that P2 is currently playing. The first expression corresponds to the expected payoff that P2 would get if he chose to play β_2 regardless of history and P1 played β_1 whenever P2's currently played action was β_2.

Note we are not stating that such strategy profiles constitute SPNE. Rather we are only providing an upper bound on the payoff $V(\alpha_1, \alpha_2)$.[6] In the example, both of the expressions above are strictly less than u, the best payoff that P2 can obtain in any SPNE at state α. Therefore in any SPNE, P2 must play α_2 at any history in which P1's currently played action is α_1.

This however means that P1 should play α_1 at any history in which P2 is currently playing α_2, because by doing so, P1 can obtain his best payoff. Simple calculations reveal that the best ex ante payoff that P1 can obtain in any SPNE beginning at state (β_1, β_2) is strictly less than

$$\frac{r}{\lambda + r}b + \frac{\lambda}{\lambda + r}a.$$

[6] If we indeed examine whether such a strategy can constitute an SPNE, we can likely sharpen the sufficient conditions that we currently have.

However this is the payoff that P1 can obtain by playing α_1 at all histories, provided that P2 always plays α_2 when P1's current action is α_1. Therefore, this shows that P1 must always play α_1 regardless of the history in any SPNE. Therefore this example shows that starting at any state, the play of the game must eventually absorb at (α) almost surely, in any SPNE. The remainder of the paper examines the set of parameter values in battle of the sexes games for which this is the case.

4 Markov perfect equilibria

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 the same action profile. 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 start the analysis by establishing a lemma on the relationship between continuation payoffs at different states in any MPE, which will be useful to narrow down possible strategy profiles in MPE.

Lemma 1 *Suppose that $v_1(\alpha_1, \beta_2) \geq v_1(\beta)$. Then in any Markovian equilibrium, we must have $v_1(\alpha) > v_1(\beta_1, \alpha_2)$. Similarly if $v_2(\alpha_1, \beta_2) \geq v_2(\alpha)$ then we must have $v_2(\beta) > v_2(\beta_1, \alpha_2)$.*

Proof. Suppose that $v_1(\beta_1, \alpha_2) \geq v_1(\alpha)$. Then we can write

$$v_1(\alpha) \quad \geq \quad \frac{r}{2\lambda + r} a + \frac{2\lambda}{2\lambda + r} \left(\frac{1}{2} v_1(\alpha) + \frac{1}{2} \left(p v_1(\alpha) + (1-p) v_1(\alpha_1, \beta_2) \right) \right)$$

$$v_1(\alpha_1, \beta_2) = \frac{r}{2\lambda + r} b + \frac{2\lambda}{2\lambda + r} \left(\frac{1}{2} v_1(\alpha_1, \beta_2) + \frac{1}{2} \left(p v_1(\alpha) + (1-p) v_1(\alpha_1, \beta_2) \right) \right)$$

for some $p \in [0, 1]$. Note that the above relies on the condition that P2 conditions his play only on the opponent's currently played action. The first is an inequality since if P1 obtains an arrival at state α, he has an incentive to play β_1. These expressions then imply that

$$v_1(\alpha) - v_1(\alpha_1, \beta_2) \geq \frac{r}{\lambda + r}(a - b).$$

Now next we can write the value functions when the currently played action of P1 is β_1:

$$\begin{aligned} v_1(\beta_1, \alpha_2) &= \frac{r}{2\lambda + r} c + \frac{2\lambda}{2\lambda + r} \left(\frac{1}{2} v_1(\beta_1, \alpha_2) + \frac{1}{2} \left(q v_1(\beta_1, \alpha_2) + (1-q) v_1(\beta) \right) \right) \\ v_1(\beta) &\geq \frac{r}{2\lambda + r} d + \frac{2\lambda}{2\lambda + r} \left(\frac{1}{2} v_1(\beta) + \frac{1}{2} \left(q v_1(\beta_1, \alpha_2) + (1-q) v_1(\beta) \right) \right). \end{aligned}$$

This then implies

$$v_1(\beta) - v_1(\beta_1, \alpha_2) \geq \frac{r}{\lambda + r}(d - c).$$

Then this means that

$$\begin{aligned} v_1(\alpha) - v_1(\beta_1, \alpha_2) &= (v_1(\alpha) - v_1(\alpha_1, \beta_2)) + (v_1(\alpha_1, \beta_2) - v_1(\beta)) + (v_1(\beta) - v_1(\beta_1, \alpha_2)) \\ &\geq \frac{r}{\lambda + r}(a - b) + \frac{r}{\lambda + r}(c - d) \\ &= \frac{r}{\lambda + r}(a - c + d - b) > 0. \end{aligned}$$

This is a contradiction. Therefore we must have $v_1(\alpha) > v_1(\beta_1, \alpha_2)$. ■

Note that the lemma implies that (β_1, α_2) cannot be part of an absorbing cycle of the Markovian dynamics in any MPE. Because of the restrictions on payoffs ($u > v$, $d > c$) it is also impossible that an absorbing cycle consists of (α_1, β_2) and at most one of the stage game strict Nash equilibria. Hence, for any specification of the game, we only have the following possibilities for the Markovian dynamics implied by any MPE: (i) (α) is the unique absorbing state; (ii) (β) is the unique absorbing state; (iii) (α) and (β) are both absorbing states; (iv) there is a unique absorbing cycle involving (α_1, α_2), (α_1, β_2) and (β_1, β_2).

Our central interest is characterizing the set of parameter values for which one of the stage game Nash equilibria is the unique absorbing state in every MPE. Without loss of generality we conduct this for (α). We do it 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 1 implies that the latter is indeed the set of parameter values in which (α) is the unique absorbing state.

First we characterize the parameter region in which there is an MPE in which (β) is an absorbing state, and then the parameter region in which there is a nonsingleton absorbing cycle. Since this requires a tedious case by case analysis, the proof of the following two lemmas are in the Appendix.

Lemma 2 *The necessary and sufficient conditions for the existence of an equilibrium with* (β) *as a unique absorbing state to exist are:*

1.

$$\frac{r}{\lambda + r} v + \frac{\lambda}{\lambda + r} x \geq u.$$

2.

$$\begin{aligned} a + \frac{\lambda}{\lambda + r} b &\leq c + \frac{\lambda}{\lambda + r} d, \\ v + \frac{\lambda}{\lambda + r} x &\geq u + \frac{\lambda}{\lambda + r} w. \end{aligned}$$

Additionally if $\frac{r}{\lambda+r} v + \frac{\lambda}{\lambda+r} x < u$ *and* $d \geq \frac{r}{\lambda+r} b + \frac{\lambda}{\lambda+r} a$ *an equilibrium in which both* (α) *and* (β) *are absorbing states exist.*

Note that the above implies that for r sufficiently close to zero, there always exists a Markovian equilibrium with (β) as a unique absorbing state (symmetrically, this is also true for α).

Lemma 3 *The necessary and sufficient conditions for the existence of an MPE with a unique absorbing cycle involving* (α_1, α_2), (α_1, β_2) *and* (β_1, β_2)

are:

$$\frac{r}{\lambda+r}v+\frac{\lambda}{\lambda+r}x>u$$

and

$$\frac{r}{\lambda+r}b+\frac{\lambda}{\lambda+r}a>d.$$

We can use the previous lemmas to characterize the set of parameter values for which (α) is the unique limit outcome in MPE. Symmetric conditions characterize the parameter region in which (β) is the unique limit outcome in MPE.

Theorem 1: *The necessary and sufficient conditions for* (α) *to be the unique absorbing state in any MPE are:*

$$\begin{aligned}\frac{r}{\lambda+r}v+\frac{\lambda}{\lambda+r}x &< u\\ \frac{r}{\lambda+r}b+\frac{\lambda}{\lambda+r}a &> d\end{aligned}$$

and one of the following:

1.

$$a+\frac{\lambda}{\lambda+r}b>c+\frac{\lambda}{\lambda+r}d$$

2. *or*

$$v+\frac{\lambda}{\lambda+r}x<u+\frac{\lambda}{\lambda+r}w.$$

Proof. Lemmas 2 and 3 imply that the parameter region in the statement is exactly the set of parameters for which there does not exist an MPE in which (β) is an absorbing state and there does not exist an MPE in which there is an absorbing cycle involving (α_1,α_2), (α_1,β_2) and (β_1,β_2). Lemma 1 then implies that this is exactly the set of parameter values for which in all MPE the unique absorbing state is (α). ■

Intuitively, the first two conditions require that it is better for P2 if play stays in (α) forever than transitioning to (α_1,β_2) (even if P1 chooses β_1 with probability 1 at state $(., \beta_2)$), and that if P2 chooses α_2 with probability 1 at state $(\alpha_1, .)$ then it is better for P1 to transition from state (α) than staying there forever. And either of the last two conditions guarantee that there is no MPE in which P1 transitions from (α) because of the fear that otherwise P2 would transition away from this state.

Increasing a, b, u and w, and decreasing c, d, v and x make the inequalities in the statement of Theorem 1 easier to satisfy. The comparative statics in r and λ are nonmonotonic, and in fact both for low enough values of r and λ, and for high enough values of r and λ one of the conditions are violated.

5 General subgame perfect equilibria

First we demonstrate by an example that a unique limit outcome in MPE does not imply the uniqueness of limit outcome in SPNE. Then we provide a set of sufficient conditions for the uniqueness of a limit outcome in SPNE. Lastly, we provide stronger sufficient conditions for the uniqueness of SPNE strategies as well.

5.1 Example: unique limit outcome in MPE does not imply unique limit outcome in SPNE

Consider the following stage game.

	α_2	β_2
α_1	$10, 1$	$-12, -1$
β_1	$9, 1$	$1, 2$

Let λ and r be such that

$$\begin{aligned} \frac{2\lambda}{2\lambda + r} &= \frac{3}{4}, \\ \frac{\lambda}{\lambda + r} &= \frac{3}{5}. \end{aligned}$$

Note first that with these parameters, no Markovian equilibrium with β as an absorbing state exists. However we show that the following non-Markovian strategy profile with β as an absorbing state is an equilibrium. As long as no deviations by P2 have occurred,

- P1 always plays β_1.
- P2 plays β_2 whenever play immediately preceding the arrival is (β_1, α_2), or (β_1, β_2).
- P2 plays β_2 when play immediately preceding the arrival is (α_1, α_2) if either there have been no arrivals in the game or the last player to arrive was P1. Otherwise P2 plays α_2.
- P2 plays α_2 when play immediately preceding the arrival is (α_1, β_2).

If P2 has deviated, then players revert to the unique Markovian equilibrium with an absorbing state at (α).

Let us illustrate how the incentives can be checked. First suppose that P2 has deviated. Then incentives to follow the Markovian equilibrium is trivial. Therefore let us suppose that P2 has not deviated. Consider an arrival by P1 at state $(\cdot, \alpha_2)$. Then playing β_1 yields a payoff of

$$\frac{2}{5} \cdot 9 + \frac{3}{5} \cdot 1 = \frac{21}{5} = 4.2.$$

Suppose by way of contradiction that playing α_1 is indeed optimal. This then gives a payoff of:

$$V = \frac{2}{5} \cdot 10 + \frac{3}{5}\left(\frac{1}{4}(-12) + \frac{3}{4}\left(\frac{1}{2} \cdot 1 + \frac{1}{2} v_\alpha\right)\right) = \frac{973}{250} = 3.892$$

where

$$v_\alpha = \frac{2}{5} \cdot 10 + \frac{3}{5}\left(\frac{2}{5} \cdot 9 + \frac{3}{5} \cdot 1\right) = \frac{163}{25}.$$

Thus upon arrival when P2 is currently playing α_2, P1 strictly prefers to play β_1 because of the threat of P2 moving to (α_1, β_2). This threat is large enough to provide incentives for P1 to play β_1 as soon as she gets the chance.

A necessary condition for this to hold is that for P1, the path

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

is strictly preferable to the path

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

Otherwise P1 would have an incentive to play α_1 at states $(\cdot, \alpha_2)$. Straightforward calculations establish that this is indeed the case, and that P1 strictly prefers to play β_1 whenever P2's currently played action is β_2.

P2's incentives to play β_2 when P1's currently played action is β_1 are trivial. Furthermore, incentives to play α_2 when the strategy calls for it are also trivial. Thus it remains only to show that it is in her interest to play β_2 at state $(\cdot, \alpha_2)$. Note that by staying at α_2, P2 has deviated and thus receives only a continuation payoff of 1.

However if he moves to β_2, then P2 is rewarded by P1 following a strategy in the future that calls for the play of β_1 everywhere on the equilibrium path. The game was constructed in such a way to make this strictly greater than 1. The computations yield a payoff of

$$\frac{1}{4}(-1) + \frac{3}{4}\left(\frac{1}{2} \cdot 2 + \frac{1}{2} v_\alpha\right) = \frac{589}{200} = 2.945 > 1.$$

Thus we have shown that the above strategy profile is indeed an equilibrium.

The essential feature of the above example is that it is too costly for P2 to transition from (α_1, α_2) to (β_1, β_2) via (β_1, α_2), even if P1 plays β_1 with probability 1 whenever P2 plays β_2. However, P2 is willing to transition to from (α_1, α_2) to (β_1, α_2) if along the equilibrium path he can transition back to (α_1, α_2) if he gets the next arrival as well, provided that P1 is willing to transition from (α_1, α_2) to (β_1, α_2). The latter is indeed the case because transitioning via (β_1, α_2), is too costly for P1 as well, and it is in her interest to prevent it. Lastly, P2 is willing to make the temporary transition to (β_1, α_2) because otherwise he would trigger the MPE as the continuation equilibrium.

5.2 Uniqueness of limit outcome in SPNE

First we provide sufficient conditions for P2 to choose action α_2 in any history at which P1 is currently playing α_1, in any SPNE. If this is the case then the best response of P1 is also choosing action α_1 whenever P2 is currently playing α_2, establishing that (α_1, α_2) is an absorbing state in every SPNE. Subsequently, we will add a further condition that guarantees that it is the unique absorbing state.

To shorten notation define

$$v_\alpha = \frac{r}{\lambda + r}u + \frac{\lambda}{\lambda + r}\left(\frac{r}{\lambda + r}w + \frac{\lambda}{\lambda + r}x\right).$$

We will split the parameter space into two regions ($w \geq u$ and $w < u$), and first consider the easier case.

Lemma 4 *Suppose* $w \geq u$ *and that:*

$$\begin{aligned} u &> \frac{r}{\lambda + r}v + \frac{\lambda}{\lambda + r}x, \text{ and} \\ u &> \frac{r}{2\lambda + r}v + \frac{2\lambda}{2\lambda + r}\left(\frac{1}{2}x + \frac{1}{2}v_\alpha\right), \end{aligned}$$

then in any SPNE, P2 *must play* α_2 *whenever* P1*'s currently played action is* α_1.

Proof. Note first that in any equilibrium, at state (α_1, α_2), P2 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 P2 can obtain at state (α_1, β_2).

To do this, denote $V(\alpha)$, $V(\alpha_1, \beta_2)$, and $V(\beta_1, \alpha_2)$ the best possible P2 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

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

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

$$V(\alpha_1, \beta_2) \leq \frac{r}{\lambda + r}v + \frac{\lambda}{\lambda + r}x < u.$$

Therefore the best possible payoff is less than the payoff that P2 could guarantee at the state (α). Thus we have shown that P2 must play α_2 with probability one upon arrival when the currently played action of P1 is α_1 in any SPNE.

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

$$\begin{aligned} V(\alpha) &\leq \frac{r}{2\lambda + r}u + \frac{2\lambda}{2\lambda + r}\left(\frac{1}{2}\max\{V(\alpha), V(\alpha_1, \beta_2)\} + \frac{1}{2}\max\{V(\alpha), V(\beta_1, \alpha_2)\}\right) \\ &= \frac{r}{2\lambda + r}u + \frac{2\lambda}{2\lambda + r}\left(\frac{1}{2}V(\alpha) + \frac{1}{2}\max\{V(\alpha), V(\beta_1, \alpha_2)\}\right). \end{aligned}$$

Let us assume first that $V(\alpha) \geq V(\beta_1, \alpha_2)$. Then the above implies that

$$V(\alpha) \leq u.$$

However this means that

$$V(\alpha_1, \beta_2) \leq \frac{r}{2\lambda + r}v + \frac{2\lambda}{2\lambda + r}\left(\frac{1}{2}u + \frac{1}{2}x\right) < u.$$

Thus again it must be that P2 plays α_2 with probability one upon arrival in any SPNE when the currently played action of P1 is α_1.

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

$$V(\beta_1, \alpha_2) \leq \frac{r}{2\lambda + r}w + \frac{2\lambda}{2\lambda + r}\left(\frac{1}{2}V(\beta_1, \alpha_2) + \frac{1}{2}x\right).$$

This implies that

$$V(\beta_1, \alpha_2) \leq \frac{r}{\lambda + r}w + \frac{\lambda}{\lambda + r}x.$$

This immediately implies that

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

Thus in this case again, P2 must play α_2 with probability one whenever the currently played action of P1 is α_1 in any SPNE. ■

The case when $w < u$ is a bit more subtle. In the previous case, we used a punishment scheme by which all players transition to the Markovian equilibrium in which (α_1, α_2) is an absorbing state. However this only yields a punishment payoff of u for P2. This is an effective punishment against P2 when $w \geq u$ due to the fact that P1 can indeed guarantee himself a payoff of u at the state (α_1, α_2). However when $w < u$ this might not be the case anymore. Therefore it may be possible to indeed punish P2 using strategies that are non-Markovian that lower the payoff to P2 to some payoff strictly below u.

This leads to an additional problem that was non-existent in $w \geq u$. Suppose that the state is (α_1, α_2) and if P2 obtains an arrival and decides to stay at (α_1, α_2) then he gets "punished" with a continuation payoff that is strictly smaller than u. However note that this may create an incentive for P2 to play α_2 since then he may be rewarded by transitioning to an MPE with (α) as an absorbing state. Thus we must ensure against such a deviation.

Lemma 5 *Suppose $w < u$ and that*

$$\begin{aligned}\frac{r}{\lambda + r}u + \frac{\lambda}{\lambda + r}w &> \frac{r}{\lambda + r}v + \frac{\lambda}{\lambda + r}x, \text{ and} \\ \frac{r}{\lambda + r}u + \frac{\lambda}{\lambda + r}w &> \frac{r}{2\lambda + r}v + \frac{2\lambda}{2\lambda + r}\left(\frac{1}{2}x + \frac{1}{2}\max\{v_\alpha, u\}\right).\end{aligned}$$

Then in any SPNE, P1 must play α_2 in any history at which P2 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

$$V(\alpha_1, \beta_2) < \frac{r}{\lambda + r}u + \frac{\lambda}{\lambda + r}w.$$

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

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

This then implies that

$$V(\alpha_1, \beta_2) \leq \frac{r}{\lambda + r}v + \frac{\lambda}{\lambda + r}x < \frac{r}{\lambda + r}u + \frac{\lambda}{\lambda + r}w.$$

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

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

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

$$V(\alpha) \leq u.$$

This then means that

$$V(\alpha_1, \beta_2) \leq \frac{r}{2\lambda + r} v + \frac{2\lambda}{2\lambda + r}\left(\frac{1}{2}u + \frac{1}{2}x\right) < \frac{r}{\lambda + r}u + \frac{\lambda}{\lambda + r}w.$$

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

$$V(\alpha_1, \beta_2) \leq \frac{r}{2\lambda + r} v + \frac{2\lambda}{2\lambda + r}\left(\frac{1}{2}x + \frac{1}{2}v_\alpha\right) < \frac{r}{\lambda + r}u + \frac{\lambda}{\lambda + r}w.$$

Thus we have shown that $V(\alpha_1, \beta_2) < \frac{r}{\lambda+r}u + \frac{\lambda}{\lambda+r}w$. Therefore since P2 can always guarantee a payoff of

$$\frac{r}{\lambda + r}u + \frac{\lambda}{\lambda + r}w$$

at state (α) by always playing α_2 in the future, P2 must always play α_2 when P1 is currently playing α_1. ■

Next we investigate the issue of what further conditions are needed for (α) to be the unique absorbing state.

Theorem 2 *Suppose that*

$$\begin{aligned} \frac{r}{\lambda + r}b + \frac{\lambda}{\lambda + r}a &> d, \\ \min\left\{u, \frac{r}{\lambda + r}u + \frac{\lambda}{\lambda + r}w\right\} &> \frac{r}{\lambda + r}v + \frac{\lambda}{\lambda + r}x, \text{ and} \\ \min\left\{u, \frac{r}{\lambda + r}u + \frac{\lambda}{\lambda + r}w\right\} &> \frac{r}{2\lambda + r}v + \frac{2\lambda}{2\lambda + r}\left(\frac{1}{2}x + \frac{1}{2}\max\{v_\alpha, u\}\right). \end{aligned}$$

Then in every subgame perfect equilibrium σ,

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

for any initial state a^0. That is, the history must become absorbed at state α in finite time almost surely.

Proof. By the previous lemma, we know that α must be an absorbing state. Suppose by way of contradiction that

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

for some a^0. Clearly a^0 cannot be α.

Note that a history encodes three pieces of information. First it specifies the arrival times: $t^1, t^2, \ldots$ which specify at what time some player obtains a revision opportunity. Secondly it specifies the action chosen at each of these arrival times $a^1, a^2, \ldots$. Finally it specifies the player who has arrived at each of these times. We ignore this third piece of information for the purposes of the proof. Thus we denote a history h by:

$$h = (a^0, (t^k, a^k)_{k=1}^{\infty})$$

where a^0 is the initial state. Note first that we can restrict to h such that t^k is an infinite sequence since histories with finitely many arrivals occur with probability zero under Poisson arrivals. To ease notation, let us define the following set:

$$H^k = \{h \in H : \exists t^{k'} \geq k \text{ } s.t. \text{ } a^{k'} \neq \alpha\}.$$

Define

$$H^k(a) = \{h \in H : \exists t^{k'} \geq k \text{ } s.t. \text{ } a^{k'} = a\}$$

and

$$H(a) = \bigcap_{k=0}^{\infty} H^k(a).$$

Then the above implies that

$$Pr_\sigma\left(\bigcap_{k=0}^{\infty} H^k\right) > 0.$$

But then note that

$$Pr_\sigma\left(H(\alpha_1, \beta_2) \cup H(\beta_1, \beta_2 \setminus (H(\beta_1, \alpha_2) \cup H(\alpha_1, \beta_2))) \cup H(\beta_1, \alpha_2)\right) \geq Pr_\sigma\left(\bigcap_{k=0}^{\infty} H^k\right) > 0$$

which is less than

$$Pr_\sigma(H(\alpha_1, \beta_2)) + Pr_\sigma(H(\beta_1, \beta_2) \setminus (H(\beta_1, \alpha_2) \cup H(\alpha_1, \beta_2))) + Pr_\sigma(H(\beta_1, \alpha_2)).$$

But note that $Pr_\sigma(H(\alpha_1, \beta_2)) = Pr_\sigma(H(\beta_1, \alpha_2)) = 0$ since state (β_i, α_{-i}), player i plays α_i with probability one. Thus the above implies that

$$Pr_\sigma(H(\beta_1, \beta_2) \setminus (H(\beta_1, \alpha_2) \cup H(\alpha_1, \beta_2))) > 0.$$

But if $h \in H(\beta_1, \beta_2) \setminus (H(\beta_1, \alpha_2) \cup H(\alpha_1, \beta_2))$, there exists some T such that $h^t = (\beta_1, \beta_2)$ for all $t \geq T$. Here we are using the assumption of asynchronous nature of moves. This is because play cannot transition from (β_1, β_2) to (α_1, α_2) without passing through either (α_1, β_2) or (β_1, α_2). But if this is the case, consider the subgame induced at time T. P1 is receiving a payoff of d. However he can deviate to playing α_1 receiving a payoff of

$$\frac{r}{\lambda + r} b + \frac{\lambda}{\lambda + r} a$$

which is strictly better than d. Therefore we have arrived at a contradiction. ■

We would like to point out that the conditions in Theorem 2 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.

5.3 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:

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

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}.
- P1 plays α_1 with probability $9/11$ whenever the opponent's prepared actions is β_2.
- P2 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:

$$\begin{aligned}
V(\alpha_1, \alpha_2) &= (2, 1) \\
V(\alpha_1, \beta_2) &= \left(\frac{1}{2}\varepsilon + 1, -\frac{4}{3}\right) \\
V(\beta_1, \alpha_2) &= \left(\frac{1}{3}\varepsilon + \frac{4}{3}, \frac{1}{2}\right) \\
V(\beta_1, \beta_2) &= \left(\frac{1}{2}\varepsilon + 1, \frac{1}{2}\right).
\end{aligned}$$

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 P1 at an initial state (β_1, β_2) beyond what he can achieve in the Markovian equilibrium. This is because for P1, 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 P2 is moving to α_2 with probability one upon arrival. In the above strategy profile when the state is (β_1, β_2), either P1 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 P2 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, besides uniqueness of a limit outcome in SPNE.

Theorem 3 *Suppose that*

$$\begin{aligned} \frac{r}{\lambda + r} b + \frac{\lambda}{\lambda + r} a &> d, \\ \min\left\{u, \frac{r}{\lambda + r} u + \frac{\lambda}{\lambda + r} w\right\} &> \frac{r}{\lambda + r} v + \frac{\lambda}{\lambda + r} x, \text{ and} \\ \min\left\{u, \frac{r}{\lambda + r} u + \frac{\lambda}{\lambda + r} w\right\} &> \frac{r}{2\lambda + r} v + \frac{2\lambda}{2\lambda + r}\left(\frac{1}{2} x + \frac{1}{2}\max\{v_\alpha, u\}\right). \end{aligned}$$

Suppose further that

$$b + \frac{\lambda}{\lambda + r} a > d + \frac{\lambda}{\lambda + r} c$$

or

$$w + \frac{\lambda}{\lambda + r} u < x + \frac{\lambda}{\lambda + r} v.$$

Then there exists a unique SPNE. Furthermore the unique SPNE must be the unique MPE.

Proof. We have already shown that the first two inequalities imply that P2 must play α_2 at any history when P1 is currently playing α_1 and P1 must play α_1 at states when P2 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

$$V_1(\beta) < \frac{r}{\lambda + r}b + \frac{\lambda}{\lambda + r}a.$$

Suppose first that $V_1(\beta) \leq V_1(\beta_1, \alpha_2)$. Then

$$V_1(\beta_1, \alpha_2) \leq \frac{r}{\lambda + r}c + \frac{\lambda}{\lambda + r}d.$$

If $V_1(\beta) \geq V_1(\alpha_1, \beta_2)$ then this implies that

$$V_1(\beta) \leq \frac{r}{\lambda + r}d + \frac{\lambda}{\lambda + r}\left(\frac{r}{\lambda + r}c + \frac{\lambda}{\lambda + r}a\right).$$

But clearly the above is strictly less than

$$\frac{r}{\lambda + r}b + \frac{\lambda}{\lambda + r}a.$$

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

$$V_1(\beta_1, \alpha_2) \leq \frac{r}{\lambda + r}c + \frac{\lambda}{\lambda + r}a.$$

Additionally

$$V_1(\alpha_1, \beta_2) \leq \frac{r}{\lambda + r}b + \frac{\lambda}{\lambda + r}a.$$

This then implies that

$$\begin{aligned} V_1(\beta) &\leq \frac{r}{2\lambda + r}d + \frac{2\lambda}{2\lambda + r}\left(\frac{1}{2}V_1(\alpha_1, \beta_2) + \frac{1}{2}V_1(\beta_1, \alpha_2)\right) \\ &< \frac{r}{\lambda + r}b + \frac{\lambda}{\lambda + r}a. \end{aligned}$$

Finally let us assume that $V_1(\beta) > V_1(\beta_1, \alpha_2)$. Therefore

$$V_1(\beta) \leq \frac{r}{2\lambda + r}d + \frac{2\lambda}{2\lambda + r}\left(\frac{1}{2}V_1(\beta) + \frac{1}{2}\max\{V_1(\beta), V_1(\alpha_1, \beta_2)\}\right).$$

One can show that this is less than $\frac{r}{\lambda+r}b + \frac{\lambda}{\lambda+r}a$.

Thus we have shown that in all of these cases,

$$V_1(\beta) < \frac{r}{\lambda+r}b + \frac{\lambda}{\lambda+r}a.$$

However the right hand side of the above inequality denotes the payoff that P1 can guarantee when she transitions play to (α_1, β_2). Therefore it must be that in any SPNE, P1 must play α_1 with probability one at any history in which P2 is currently playing β_2. This establishes the claim. ■

6 Discussion

Payoff specifications compatible with unique limit outcome

The conditions we provided for the uniqueness of limit outcome imply restrictions jointly on the payoffs and the discount rate. A natural question to ask is whether it is true for generic payoff specifications of the battle of the sexes game that there is an intermediate range of discount rates such that there is a unique limit outcome in SPNE.

The answer to this question is negative. Consider the following battle of sexes games:

	α_2	β_2
α_1	$1, 0$	$-\frac{3}{2}, -1$
β_1	$-2, \frac{3}{4}$	$0, 1$

Our characterization of MPE establishes that there is no discount rate at which there is a unique limit outcome in MPE. Moreover, this remains true for all small perturbations of stage game payoffs, too.

More generally, a simple manipulation of our necessary and sufficient condition for unique limit outcome reveals that the limit outcome cannot be unique in MPE for any discount rate whenever:

$$\begin{aligned} u_2(\beta_1, \beta_2) + u_2(\alpha_1, \beta_2) &< u_2(\beta_1, \alpha_2) + u_2(\alpha_1, \alpha_2) \\ u_1(\alpha_1, \beta_2) + u_1(\alpha_1, \alpha_2) &> u_1(\beta_1, \beta_2) + u_1(\beta_1, \alpha_2). \end{aligned}$$

That is, if for both players it is true that the strategy corresponding to the player's least preferred stage game Nash equilibrium is risk-dominant, then there is always multiplicity of limit outcomes in equilibrium.

Expected equilibrium payoffs when there is a unique limit outcome

Because uniqueness of the limit outcome occurs at an intermediate range of discount rates, the strong player (whose favored stage game Nash equilibrium is selected as the limit outcome) is not necessarily better off. To see this, consider again the example in Section 3:

	α_2	β_2
α_1	$4, 2$	$1, -4$
β_1	$-2, 3$	$2, 4$

Following the arguments of Section 3, we can show that the unique SPNE is Markovian and follows the following set of dynamics with α as the unique absorbing state.

However, the ex ante payoffs of P1 at state (β) is:

$$\begin{aligned} & \frac{r}{2\lambda + r}d + \frac{2\lambda}{2\lambda + r}\left(\frac{1}{2}\left(\frac{r}{\lambda + r}c + \frac{\lambda}{\lambda + r}a\right) + \frac{1}{2}\left(\frac{r}{\lambda + r}b + \frac{\lambda}{\lambda + r}a\right)\right) \\ = \; & \frac{1}{3} \cdot 2 + \frac{2}{3}\left(\frac{1}{2}\left(\frac{1}{2}(-2) + \frac{1}{2}(4)\right) + \frac{1}{2}\left(\frac{1}{2}(1) + \frac{1}{2}4\right)\right) \\ = \; & \frac{11}{6} < 2. \end{aligned}$$

That is, not only P2 but also P1 is worse off at state (β) than if play stayed forever at state (β).

Possible extensions

Our analysis could be easily extended to allowing for different discount rates and/or arrival rates for the players. Intuitively, increasing one of the player's discount rate makes the player weaker, making it more likely that the other player's favorite stage game equilibrium becomes the unique limit outcome in SPNE. Increasing the arrival rate of a player also makes that player weaker, as it makes the other player more like a Stackelberg leader in the game, since that player is credibly locked into playing her current action for a relatively longer time.

Characterizing MPE and providing a sufficient condition for unique limit outcome in SPNE could be done similarly as in our basic setting, at the cost of more complicated notation and more complicated equilibrium conditions. As the additional qualitative insights from this exercise are limited, we do not pursue this direction formally here.

A much less straightforward direction would be extending the investigation beyond battle of the sexes games, and characterizing more generally the set of games for which there is an intermediate range of discount rates for which the limit outcome is unique in SPNE. As we showed above, such a range might not exist even for games within the battle of the sexes class, suggesting that this extension might be difficult.

A Appendix

A.1 Proof of Lemma 2

Below we first provide necessary conditions for existence of MPE in which (β_1, β_2) is the unique absorbing state. Then it is trivial to show that these necessary conditions are also sufficient. We find it convenient to describe the type of equilibrium graphically. The arrow connecting two states between which one of the players can transition play indicates the player's action conditional on the other player playing the action consistent with these states. If the arrow points to one state, it implies choosing the corresponding action with probability 1, otherwise mixing between the two actions. The states colored in blue represent the elements of the unique absorbing cycle.

Generically, note first that there are only three possible types of dynamics for any MPE with (β_1, β_2) as the unique absorbing state:

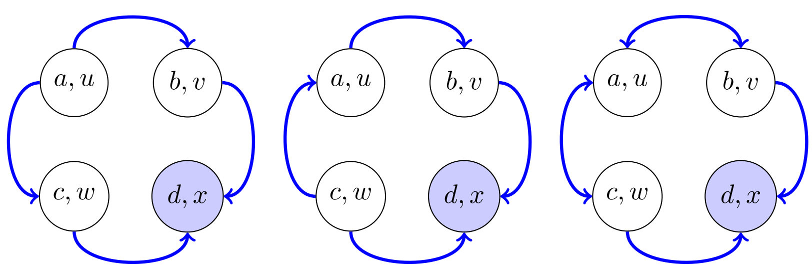

It is easy to see that the first set of dynamics exist if and only if

$$\begin{aligned} c + \frac{\lambda}{\lambda + r} d &\geq a + \frac{\lambda}{\lambda + r} b \\ v + \frac{\lambda}{\lambda + r} x &\geq u + \frac{\lambda}{\lambda + r} w. \end{aligned}$$

Similarly the second set of dynamics exist if and only if

$$\begin{aligned} \frac{r}{\lambda + r} v + \frac{\lambda}{\lambda + r} x &\geq u \\ a + \frac{\lambda}{\lambda + r} b &\geq c + \frac{\lambda}{\lambda + r} d. \end{aligned}$$

The necessary and sufficient conditions for existence of the third set of dynamics are only slightly more subtle. First for P1 to mix, we need

$$c + \frac{\lambda}{\lambda + r} d \geq a + \frac{\lambda}{\lambda + r} b.$$

This is because if P2 plays α_2 with probability one whenever he is playing α_1, then the unique best response for P1 is to play α_1 whenever P2 is playing α_2. For him to have an incentive to play β_1 when P2's current action is α_2, P1 needs to prefer playing β_1 when P2 plays β_2 with probability one regardless of the history. This is exactly the condition above.

Note that this is the first condition that is necessary and sufficient for the first set of dynamics. Since if we also have

$$v + \frac{\lambda}{\lambda + r} x \geq u + \frac{\lambda}{\lambda + r} w,$$

we showed necessary and sufficient conditions for existence of equilibria of the form of the first kind, let us assume the contrary. We have thus far,

$$\begin{aligned} c + \frac{\lambda}{\lambda + r} d &\geq a + \frac{\lambda}{\lambda + r} b \\ v + \frac{\lambda}{\lambda + r} x &< u + \frac{\lambda}{\lambda + r} w. \end{aligned}$$

With these conditions, similar to the argument we use for P1, in order for P2 to have incentives to mix, he must be willing to play β_2 when P1 plays α_1 with probability one whenever his prepared action is α_2. This means that

$$u < \frac{r}{\lambda + r} v + \frac{\lambda}{\lambda + r} x.$$

So the set of parameters for which the third kind of dynamics can occur as an MPE that are not covered by the first case is exactly

$$\begin{aligned} c + \frac{\lambda}{\lambda + r} d &\geq a + \frac{\lambda}{\lambda + r} b \\ v + \frac{\lambda}{\lambda + r} x &< u + \frac{\lambda}{\lambda + r} w \\ u &< \frac{r}{\lambda + r} v + \frac{\lambda}{\lambda + r} x. \end{aligned}$$

Taking the union of these sets of parameters gives us the set of parameters in the statement of Lemma 2.

We can perform a similar analysis to determine the set of parameters for which there exists an MPE in which both (α) and (β) are the absorbing states.

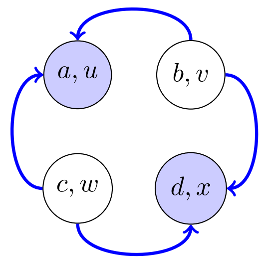

The necessary and sufficient conditions for this to occur are:

$$\begin{aligned} u &\geq \frac{r}{\lambda + r} v + \frac{\lambda}{\lambda + r} x \\ d &\geq \frac{r}{\lambda + r} b + \frac{\lambda}{\lambda + r} a. \end{aligned}$$

A.2 Proof of Lemma 3

By lemma 1, the only possible set of dynamics supporting an MPE with (α), (α_1, β_2), and (β) as elements of the absorbing cycle is the following:

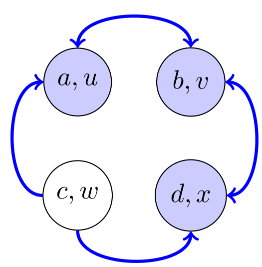

For P1 to mix, it must be that

$$\frac{r}{\lambda + r} b + \frac{\lambda}{\lambda + r} a > d.$$

Otherwise P1, would play β_1 with probability one whenever P2's currently played action is β_2 regardless of what P1 plays when P2's currently played action is α_1.

Similarly for P2 to have an incentive to mix, it must be that

$$\frac{r}{\lambda + r} v + \frac{\lambda}{\lambda + r} x > u.$$

Thus the necessary conditions are

$$\begin{aligned} \frac{r}{\lambda + r} b + \frac{\lambda}{\lambda + r} a &> d \\ \frac{r}{\lambda + r} v + \frac{\lambda}{\lambda + r} x &> u. \end{aligned}$$

It is easy to see that these conditions are indeed also sufficient.

B References

Ambrus, A. and S. Lu (2010): "A continuous-time model of multilateral bargaining," mimeo Harvard University.

Ambrus, A., J. Burns and Y. Ishii (2012): "Gradual bidding in Ebay-like auctions," mimeo Duke University

Arcidiacono, P., P. Bayer, J. Blevins and P. Ellickson (2010): "Estimation of dynamic discrete choice models in continuous time," mimeo Duke University.

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 (2012): "A foundation for Markov equilibria with finite social memory," mimeo University College London.

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

Burdett, K. and D. Mortensen (1998): "Wage differentials, employer size, and unemployment," International Economic Review, 39, 257-273.

Calgano, R., Y. Kamada, S. Lovo and T. Sugaya (2012): "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.

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

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

Dutta, P. (2003): "Coordination need not be a problem," mimeo Columbia University.

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

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

Haller, H. and R. Lagunoff (2006): "Markov perfect equilibria in repeated asynchronous choice games," mimeo Georgetown University.

Kamada, Y. and M. Kandori (2011): "Revision games," mimeo Harvard University.

Kamada, Y. and T. Sugaya (2011): "Valence Candidates and Ambiguous Platforms in Policy Announcement Games," mimeo Harvard University.

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.

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.