The Effect of Correlated Inertia on Coordination *

Yuhta Ishii †   Yuichiro Kamada‡

April 15, 2011

Abstract

We study how the structure of moves influences equilibrium predictions in the context of revision games, as termed by Kamada and Kandori (2009). In our variant of revision games, two players prepare their actions at times that arrive stochastically before playing a coordination game at a predetermined deadline, at which time the finally-revised actions are implemented. The revisions are either synchronous or asynchronous. The coordination game we study is a 2×2 game with two strict Pareto-ranked Nash equilibria. We identify the condition under which the Pareto-superior payoff profile is the unique outcome of the dynamic game. Specifically, we find that uniqueness of this outcome is more easily obtained when the degree of asynchronicity is sufficiently high relative to the risk of taking the action corresponding to the Pareto-superior profile. We further show that when this degree is low the set of payoffs attainable in equilibria expands considerably.

1 Introduction

A multitude of recent papers have shown the sharp distinction in equilibrium predictions between dynamic games in which players always move asynchronously and in which moves

*We thank Drew Fudenberg and Attila Ambrus for helpful discussions.

†Department of Economics, Harvard University, email: ishii2@fas.harvard.edu

‡Department of Economics, Harvard University, email: ykamada@fas.harvard.edu

arrive simultaneously. Such assumptions seem particularly stark in many economic applications. We examine whether such results are maintained when the synchronicity of moves is neither perfectly synchronous nor perfectly asynchronous.

We study this question in the context of revision games, as termed by Kamada and Kandori (2009). More precisely, we study a model in which two players prepare their actions at times that arrive stochastically before playing a coordination game at a predetermined deadline, at which time the finally-revised actions are implemented. The coordination game we study is a 2×2 game with two strict Pareto-ranked Nash equilibria. Revision opportunities arise according to a poisson process upon whose arrival, a player revises his action alone with probability p and simultaneously with probability $1 - 2p$. When making a revision decision, players do not know whether the opponent has also received an action revision opportunity or not. Therefore our model encapsulates both purely asynchronous and perfectly synchronous move revision games as special cases. Furthermore our model allows us to vary the level of synchronicity of the dynamic game to study the effect of changes in synchronicity on equilibrium predictions.

Our paper is motivated mainly by a result due to Calcagno and Lovo (2010) and Kamada and Sugaya (2010), hereafter referred to as respectively as CL and KS, that demonstrates the impact that asynchronous moves has on equilibrium outcomes in revision games. More specifically, their main results show that in a revision game in which the stage game is a coordination game with a Pareto-dominant action profile, the payoffs corresponding to the Pareto dominant action profile is selected as the unique equilibrium payoff as the time horizon until the deadline approaches infinity. However their reasoning relies crucially on action revision opportunities arriving to each player in a purely asynchronous fashion. Our main question asks whether such selection results are an artifact of the purely asynchronous nature of moves in their models. Of course if revision moves are purely synchronous then any infinite repetition of a static Nash equilibrium constitutes a sequential equilibrium. Therefore, at least initially, it is unclear whether unique selection can still be obtained in games in which moves are not perfectly asynchronous.

With additional assumptions about the Pareto dominant action profile, we find that equilibrium selection can still be maintained even in revision games in which moves are not purely asynchronous. This additional assumption necessary for unique equilibrium selection comes in the form of a criterion termed q-dominance introduced first by Morris, Rob, and Shin (1995). More precisely, we find that Pareto dominant equilibrium selection in these games is driven by μ-dominance of the Pareto dominant Nash equilibrium where μ is an increasing function in p. This additional assumption furthermore becomes easily obtained when the structure of moves approaches that of purely asynchronous moves. For games with purely synchronous moves, it is never satisfied.

Moreover in the class of games that we analyze, μ-dominance of the Pareto dominant Nash equilibrium is not only sufficient but a necessary condition for unique selection of the Pareto dominant payoff profile. We show that whenever this condition is violated, there exists an equilibrium in which both players always play the action profile corresponding to the Pareto inferior equilibrium. This result generalizes proposition 2.1 of CL in the context of the games that we consider.[1]

On a more practical level, our model is motivated by the fact that in many repeated interactions, even if we model interactions as occurring simultaneously, players' action revisions aren't always realized with certainty. As CL suggests, this inertia of actions exists due to technological errors or simply human mistakes in implementing the action change. For example, consider two firms that must meet over time to simultaneously revise their actions on a project. Revision opportunities arrive stochastically due to constraints in the availability of attention and time for the project for the two firms. However upon meeting each other to revise actions simultaneously, action revisions of either one of the firms may not become realized due to technological failure in processing the action revision. [2] Such errors in action revision processing generate arrivals that look like asynchronous moves as well as opportunities that look like a synchronous move opportunity.

The rest of the paper is organized as follows. In section 2 we introduce the model and framework to be analyzed. Section 3 presents the main unique selection results including definitions of q-dominance that lie at the heart of our proofs. In section 4, we illustrate that the conditions for unique selection are both necessary and sufficient.

[1]However their result is more general in that it applies to component games beyond the cases that we consider here in our paper.

[2]CL describe how such technological failures actually arise in many financial markets such as the NASDAQ or Euronext.

2 The Model

Consider a two-player 2 x 2 normal-form game in Figure 2.1, referred to as a *component game*:

	α_2	β_2
α_1	$1, 1$	y, z
β_1	z, y	$0, 0$

with $y \leq 0$ and $z < 1$.

Figure 2.1: The Component Game

We let the action space for player i be $A_i = \{\alpha_i, \beta_i\}$, and let $A = A_1 \times A_2$. Player i's payoff function is denoted by $u_i : A \to \mathbb{R}$.

We consider a *revision game* with the component game specified above, in which players prepare their actions before they actually play actions in the component game. Specifically, time is continuous, $-t \in [-T, 0]$, and the component game is played once and for all at time 0. Notice that t denotes the time left until time 0, the deadline. First, at time $-T$, players simultaneously choose actions. Between time $-T$ and 0, there is a Poisson process $\mathbf{P}$ with arrival rate $\lambda > 0$. At each arrival, there are three possible events: With probability p_1, player 1 is allowed to revise her prepared action while player 2 is not. With probability p_2, player 2 is allowed to revise her prepared action while player 1 is not. With the remaining probability $1 - p_1 - p_2$, both players are allowed to revise their actions. At $t = 0$, the action profile that has been prepared most recently is actually taken and each player receives the payoff that corresponds to the payoff specification of the component game.

We will define a strategy as a map from histories to distributions over actions. To this end, we first formalize the notion of history. Roughly, a history for player i at time $-t$ includes all the information about what has happened until time $-t$ and whether or not i has obtained an revision opportunity at $-t$, but it does not include any other information. This means that we assume whether or not a player is allowed to make a revision at time $-t$ is a private information at $-t$. That is, upon receiving an opportunity to revise, player i does not know whether player $-i$ is obtaining the opportunity at the same moment.

Formally,[3] let $t_{i,k} \leq t$ be the time when player i has received the k'th revision opportunity until time $-t$ and $x_{i,k}$ be the action prepared by player i at $t_{i,k}$.[4] A *history* for player i at t is

$$h_i(t) = \left\{ \left(\{t_{i,k}\}_{k,\ t_{i,k} \leq t} \ , \ \{a_{i,k}\}_{k,\ t_{i,k} < t} \right), \left(\{t_{j,k}\}_{k,\ t_{j,k} < t} \ , \ \{a_{j,k}\}_{k,\ t_{j,k} < t} \right),\ t \right\},$$

where $j \neq i$. Let $H_i(t)$ denote the set of all possible histories for player i at t.[5] A strategy for player i is a mapping $\sigma_i : \cup_{-t=-T}^{0} H_i(t) \to \{\emptyset\} \cup \Delta(A_i)$ where $\sigma_i(h_i(t)) \subseteq \Delta(A_i)$ if there exists k such that $t_{i,k} = t$ (i.e. at t player i receives a revision opportunity) and $\sigma_i(h_i(t)) = \emptyset$ otherwise (i.e. at $-t$ player i does not receive a revision opportunity). For any given history $h_i(t)$, let $a_i(t) := a_{i,k^*}, \in A_i$ with $k^* := \arg\max_k \{t_{i,k} < t\}$ be player i's prepared action resulting from his last revision opportunity (strictly) before t. We shall denote $a(t) := \{a_i(t)\}_{i=1,2}$ the *last prepared action profile* before time t (time $-t$ "PAP" henceforth). Note that $a_i(t)$ will be player i's payoff relevant action in $t = 0$ in the event i receives no further revision opportunities from time t included, until time 0.

A strategy profile σ^* forms a sequential equilibrium of the revision game if for all t and $h_i(t)$,[6]

$$\sigma_i^* \in \arg\max_{\sigma_i} E\left[u_i(x(0)) | h_i(t), \sigma_i, \sigma_{-i}^*\right].$$

Our main results will concern the case when T is high. We note that the model with arrival rate λ and horizon length T is essentially equivalent to the model with arrival rates λ and horizon length $\frac{T}{a}$, for any positive constant a. Hence our results for high T can be reinterpreted as those for high λ.

---

[3]The formalism here borrows largely from an ongoing research by Calcagno, Kamada, Lovo and Sugaya, with appropriate modificaitons.

[4]Notice that $-T = -t_{i,1} < \cdots < -t_{i,k} \leq -t$, that is, we count the revision opportunities from the first one $k = 1$ after the beginning of the revision game.

[5]Note that $H_i(T)$ is defined to be a singleton.

[6]Strictly speaking, $x(0)$ is the last action profile prepared before time 0, thus in this formulation players do not maximize the expected payoff prepared exactly at 0. However, since the probability that any player obtains a revision opportunity exactly at time 0 is nil, this issue does not affect the solution of the maximization problem.

3 Unique Selection of the Pareto Dominant Profile

We show in this section that even when the game is not purely asynchronous, when stronger restrictions are imposed on the action profile α beyond just pareto dominance, unique selection of the pareto dominant payoff profile can be obtained. This additional restriction is trivially obtained when the game is purely asynchronous and never satisfied when the moves are purely synchronous.

This additional restriction necessary for unique selection is the concept of q-dominance introduced first by Morris, Rob and Shin (1995) that was initially used in the literature on global games.

Definition 3.1. An action profile (a_1, a_2) is called strictly q-dominant for player i if for any action σ_{-i},

$$qu_i(a) + (1-q)u_i(a_i, \sigma_{-i}) > qu_i(a_i', \alpha_{-i}) + (1-q)u_i(a_i', \sigma_{-i})$$

for any pair of actions $a_i' \in A_i$ with $a_i \neq a_i'$.

3.1 Synchronicity Relatively Likely

Before we proceed, let us define the following parameter μ:

$$\mu = \frac{p}{1-p}.$$

In words, μ represents the probability that a player assigns to having arrived asynchronously upon an arrival. In this section we characterize sufficient conditions for unique equilibrium selection when $\mu \leq 1/2$. This corresponds to the case in which the move structure is relatively close to purely synchronous moves.

Theorem 3.2. *Suppose $\mu \leq 1/2$ and that α is strictly μ-dominant for both players. Then in any equilibrium when at least one person is currently playing α_i both players will play α_j.*

Proof. Suppose an arrival occurs at t for player 1 and suppose that the currently prepared action profile is either (α) or (β_1, α_2). Then player 1 places probability probability $(1-\mu)$

on player 2 also arriving at time t. Thus the expected payoff of playing α_1 is equal to at least

$$\mu e^{-\lambda t} u_1(\alpha) + (1-\mu)e^{-\lambda t} u_1(\alpha_1, \sigma_2(t)) + (1 - e^{-\lambda t})\underline{u}_1.$$

However the payoff from playing β_1 is at most

$$\mu e^{-\lambda t} u_1(\beta_1, \alpha_2) + (1-\mu)e^{-\lambda t} u_1(\gamma_1, \sigma_2(t)) + (1 - e^{-\lambda t})u_1(\alpha).$$

By strong μ-dominance of α and the symmetry of the problem, there exists some $t^* > 0$ such that at all times $t \leq t^*$, both players will strictly prefer playing α_i to any other strategy when the opponent's current prepared profile is α_j no matter the history of play.

Then consider the play of player 1 at time $t \leq t^*$ when player 1's prepared action is α_1 but player 2's prepared action β_2 is not necessarily α_2. Note that player 2 if he arrives simultaneously will play α_2. Therefore playing α_1 guarantees a payoff of

$$\mu e^{-\lambda t} u_1(\alpha_1, \beta_2) + (1-\mu)e^{-\lambda t} u_1(\alpha) + (1 - e^{\lambda t})\underline{u}_1,$$

while preparing β_1 yields at most a payout of

$$\mu e^{-\lambda t} u_1(\beta) + (1-\mu)e^{-\lambda t} u_1(\beta_1, \alpha_2) + (1 - e^{\lambda t})u_1(\alpha).$$

Since $1-\mu \geq \mu$, μ-dominance implies $(1-\mu)$-dominance of α. Then by $(1-\mu)$-dominance of α for player 1 to conclude that for all $t \leq t^*$ player 1 will always choose to play α_1 whenever at least one player is playing α_i. A similar argument holds for player 2.

Now consider the following time. Let T^* be the supremum over all times $t > 0$ with the property that in any equilibrium, both players always play α whenever the currently played action profile has at least one α for all $\tilde{t} \leq t$. In other words if T^* is finite, then for any $t > T^*$, there exists some equilirium in which at least one player plays some action $\sigma_i(\tilde{t}) \neq \alpha_i$ at time $\tilde{t} < t$. However consider some time $t > T^*$. Now consider an equilibrium and let $\sigma_1'(t) \in \arg\max_{\sigma_1} u_1(\sigma_1, \sigma_2(t))$.

Let us define the following continuation value functions

$$\begin{aligned}
V_1(\alpha_1, \beta_2) &= e^{-\lambda(1-p)T^*} u_1(\alpha_1, \beta_2) + (1 - e^{-\lambda(1-p)T^*}) u_1(\alpha) \\
V_1(\beta_1, \alpha_2) &= e^{-\lambda(1-p)T^*} u_1(\beta_1, \alpha_2) + (1 - e^{-\lambda(1-p)T^*}) u_1(\alpha) \\
V_1(\beta_1, \beta_2) &\leq e^{-\lambda(1-p)T^*} u_1(\beta) + (1 - e^{-\lambda(1-p)T^*}) u_1(\alpha)
\end{aligned}$$

Note that the continuation value bounds do not have to specify the history of play since from time T^* on, the play at histories in which at least one α_i is played is given no matter the history. Furthermore note that the second continuation value assumes that player 2 continues to play α_2 whenever he has done so after time T^*. This uses the assumption that $\mu \leq 1/2$.

Consider a time $t > T^*$ and suppose that the PAP is $(\cdot, \alpha_2)$. If player 2 upon arrival at such a history plays α_2, then clearly it is a strict best response for player 1 to play α_1 for t sufficiently close to T^*. Therefore consider the situation in which player 2 plays β_2 upon simultaneous arrival at t.[7]

Playing α_1 when the opponent is currently playing α_2 gives a payout of at least

$$e^{-\lambda(t-T^*)} \left(\mu u_1(\alpha) + (1-\mu) V_1(\alpha_1, \beta_2)\right) + (1 - e^{-\lambda(t-T^*)}) \underline{u}_1,$$

whereas playing some $\sigma_1 \neq \alpha_1$ gives a payout of

$$e^{-\lambda(t-T^*)} \left(\mu V_1(\beta_1, \alpha_2) + (1-\mu) V_1(\beta_1, \beta_2)\right) + (1 - e^{-\lambda(t-T^*)}) u_1(\alpha).$$

Note that

$$\begin{aligned}
& \mu u_1(\alpha) + (1-\mu) V_1(\alpha_1, \beta_2(t)) \\
= \; & \left(\mu + (1 - e^{-\lambda(1-p)T^*})(1-\mu)\right) u_1(\alpha) + (1-\mu) e^{-\lambda(1-p)T^*} u_1(\alpha_1, \beta_2)
\end{aligned}$$

[7] Note that this is without loss of generality since if the inequalities hold strictly in this situation as well as when player 2 plays α_2 upon simultaneous arrival, the inequalities will also hold strictly for any convex combination and thus for when player 2 potentially plays a mixed strategy upon arrival at t.

However

$$
\begin{aligned}
& \mu V_1(\beta_1,\alpha_2)+(1-\mu)V_1(\beta_1,\beta_2) \\
\leq\ & \left(\mu(1-e^{-\lambda(1-p)T^*})+(1-e^{-\lambda(1-p)T^*})(1-\mu)\right)u_1(\alpha) \\
+\ & \mu e^{-\lambda(1-p)T^*}u_1(\beta_1,\alpha_2)+(1-\mu)e^{-\lambda(1-p)T^*}u_1(\beta_1,\beta_2)
\end{aligned}
$$

Note that

$$
\begin{aligned}
& \mu e^{-\lambda(1-p)T^*}u_1(\alpha)+(1-\mu)e^{-\lambda(1-p)T^*}u_1(\alpha_1,\beta_2) \\
>\ & \mu e^{-\lambda(1-p)T^*}u_1(\beta_1,\alpha_2)+(1-\mu)e^{-\lambda(1-p)T^*}u_1(\beta_1,\beta_2)
\end{aligned}
$$

by strong μ-dominance. We then have

$$\mu u_1(\alpha)+(1-\mu)V_1(\alpha_1,\beta_2)>\mu V_1(\beta_1,\alpha_2)+(1-\mu)V_1(\beta_1,\beta_2).$$

Therefore there exists some $\tilde{T}>T^*$ such that at all times $t<\tilde{T}$, player 1 strictly prefers to play α_1 when the currently prepared action of player 2 is α_2. Similarly, we can choose $\tilde{T}$ such that both players strictly prefer to play α_i when the currently prepared action of the opponent is α_{-i}. Using exactly the same argument as before we can also show that given these strategies, both players would continue to play α_i whenever they are currently playing α_i. However this contradicts the definition of T^* and hence we are done. □

An easy consequence of the above theorem is the following corollary.

Corollary 3.3. *Suppose $\mu\leq 1/2$. Suppose further that α is strictly μ-dominant for both players. Then the set of payoffs shrinks to $\{u(\alpha)\}$ as the time horizon approaches infinity.*

The proof is a simple consequence of the above theorem. Theorem 3.2 implies that by always playing α_i from time T on, player i can guarantee himself a payoff of

$$(1-e^{-\lambda(1-p)T})u_i(\alpha)+e^{-\lambda(1-p)T}\underline{u}_i.$$

Note that this converges to $u_i(\alpha)$ as $T\to\infty$.

3.2 Synchronicity Relatively Unlikely

We investigate the question of whether the above selection result can be extended to the case of $\mu > 1/2$. Note that the above discussion shows that when α is strictly $(1-\mu)$-dominant, unique selection of the Pareto dominant payoff profile is obtained. However, $(1-\mu)$-dominance is a very strong assumption in the case of μ close to 1. We would like selection criteria that approximates those that are necessary for unique selection in the purely asynchronous case which is simply that α is Pareto dominant for μ close to 1.

Since we already have a selection theorem for when α is strictly $(1-\mu)$-dominant as well as μ-dominant, it remains to analyze the case in which α and β are both strictly μ-dominant. In this case, near the deadline both players want to match the opponent's currently played action no matter what the opponent does.

Theorem 3.4. *Suppose $\mu > 1/2$ and that α and β are both strictly μ-dominant. Then there exists a time T^* such that at all times $t < T^*$, players always match the play of the opponent's PAP in any equilibrium. Furthermore in any equilibrium at any time $t > T^*$, players play α when at least one of the players is currently playing α in the PAP.*

Proof. As in the previous proofs, note that there exists some $t^* > 0$ such that for all $t < t^*$, players strictly prefer to match the currently played action of the opponent in any equilibrium. Let T^* be the sup of all such times. Then we can calculate the value functions at such a time. Note that the T^* above must be finite. It is easy to see this from looking at the payoffs after a history in which the PAP is (α, β). Let the following continuation values be defined at time T^*.

$$
\begin{aligned}
V(\alpha) &= 1 \\
V(\beta, \alpha) &= e^{-2\lambda p T^*}(q(T^*)u(\beta, \alpha) + (1 - q(T^*))u(\alpha, \beta)) + \frac{1 - e^{-2\lambda p T^*}}{2} \\
V(\alpha, \beta) &= e^{-2\lambda p T^*}(q(T^*)u(\alpha, \beta) + (1 - q(T^*))u(\beta, \alpha)) + \frac{1 - e^{-2\lambda p T^*}}{2} \\
V(\beta) &= 0
\end{aligned}
$$

where $q(t)$ is the following probability:

$$q(t) = \frac{1 + e^{-2\lambda(1-2p)t}}{2}$$

representing the probability of an even number of simultaneous arrivals. This calculation can easily be seen in the following way. The probability of an even number of simultaneous arrivals conditional on the event that all arrivals are simultaneous over a time interval t is the following:

$$\begin{aligned}
\sum_{k=0}^{\infty} \frac{e^{-\lambda(1-2p)t}(\lambda(1-2p)t)^{2k}}{(2k)!} &= e^{-\lambda(1-2p)t}\frac{e^{\lambda(1-2p)t} + e^{-\lambda(1-2p)t}}{2} \\
&= \frac{1 + e^{-2\lambda(1-2p)t}}{2}.
\end{aligned}$$

The particular functional form of this probability does not matter. The important thing here for the proof is that $q(0) = 1$, $q(t) > 1/2$ for all t, and that $q(t) \to 1/2$ monotonically as $t \to \infty$.

First we prove the following:

Lemma 3.5.

$$\mu u(\alpha) + (1-\mu)V(\alpha, \beta) > \mu V(\beta, \alpha) + (1-\mu)V(\beta).$$

This means that both players will strictly prefer to play α for some time a little bit for T^* when the currently played action of the opponent is α.

Proof. First assume that $0 \geq u(\alpha, \beta) \geq u(\beta, \alpha)$. Then

$$\begin{aligned}
q(T^*)u(\beta, \alpha) + (1 - q(T^*))u(\alpha, \beta) &\leq \frac{1}{2}(u(\beta, \alpha) + u(\alpha, \beta)) \equiv \bar{u} \\
q(T^*)u(\alpha, \beta) + (1 - q(T^*))u(\beta, \alpha) &\geq \frac{1}{2}(u(\beta, \alpha) + u(\alpha, \beta)) \equiv \bar{u}.
\end{aligned}$$

Therefore

$$\mu V(\alpha) + (1-\mu)V(\alpha, \beta) \geq \mu + (1-\mu)e^{-2\lambda T^*}\bar{u} + (1-\mu)\frac{1 - e^{-2\lambda pt}}{2}$$

whereas

$$\mu(V(\beta,\alpha)+(1-\mu)V(\beta)\leq \mu e^{-2\lambda pT^*}\bar{u}+\mu\frac{1-e^{-2\lambda pt}}{2}.$$

However note that

$$\mu>(\mu-1/2)(1-e^{-2\lambda pT^*})$$

and therefore since $\bar{u}\leq 0$,

$$\mu>(2\mu-1)(\frac{1-e^{-2\lambda pt}}{2}+e^{-2\lambda pT^*}\bar{u})$$

which implies that

$$\mu+(1-\mu)e^{-2\lambda T^*}\bar{u}+(1-\mu)\frac{1-e^{-2\lambda pt}}{2}>\mu e^{-2\lambda pT^*}\bar{u}+\mu\frac{1-e^{-2\lambda pt}}{2},$$

which then implies that

$$\mu V(\alpha)+(1-\mu)V(\alpha,\beta)>\mu V(\beta,\alpha)+(1-\mu)V(\beta).$$

Now case 2, in which we assume that $u(\beta,\alpha)>u(\alpha,\beta)$. Then notice that

$$\begin{aligned} q(T^*)u(\beta,\alpha)+(1-q(T^*))u(\alpha,\beta) &\leq u(\beta,\alpha)\\ q(T^*)u(\alpha,\beta)+(1-q(T^*))u(\beta,\alpha) &\geq u(\alpha,\beta). \end{aligned}$$

Therefore

$$\mu V(\alpha)+(1-\mu)V(\alpha,\beta)\geq \mu+(1-\mu)\left(e^{-2\lambda pT^*}u(\alpha,\beta)+\frac{1-e^{-2\lambda pt}}{2}\right),$$

and

$$\mu V(\beta,\alpha)+(1-\mu)V(\beta)\leq \mu\left(e^{-2\lambda pT^*}u(\beta,\alpha)+\frac{1-e^{-2\lambda pt}}{2}\right).$$

But note that

$$\mu+(1-\mu)e^{-2\lambda pT^*}u(\alpha,\beta)-\mu e^{-2\lambda pT^*}u(\beta,\alpha)>(2\mu-1)\frac{1-e^{-2\lambda pT^*}}{2}$$

by μ-dominance of α and this concludes our proof. □

So there exists some $T_1^* > T^*$ such that for all $t \leq T_1^*$, both players play α when the opponent's PAP is α in any equilibrium. Now consider a history at which the PAP is (α, β) at some $t \leq T_1^*$. Then the player will know that the opponent will play α. Suppose by way of contradiction that

$$\mu V(\beta) + (1-\mu)V(\beta, \alpha) > \mu V(\alpha, \beta) + (1-\mu)V(\alpha).$$

Then at t sufficiently close to T^*, both players prefer to play β after such histories.

Furthermore the above inequality implies that

$$V(\beta) > V(\alpha, \beta)$$

This implies that no matter what the opponent does upon simultaneous arrival, both players prefer to strictly play β in any equilibrium at times sufficiently close to T^* when the opponent's PAP is β.

This however contradicts the definition of T^*. Therefore we must have

$$\mu V(\beta) + (1-\mu)V(\beta, \alpha) = \mu V(\alpha, \beta) + (1-\mu)V(\alpha).$$

Note the above pins down T^* exactly and consequently $V(\beta, \alpha)$ and $V(\alpha, \beta)$. Furthermore the above indifference condition implies that $V(\beta) > V(\alpha, \beta)$.

Now consider time $t \leq T^1$. At such a time, we know behavior of both players when the PAP is such that the opponent is playing α. We need to determine behavior when the PAP is (α, β). By playing α we can calculate the lower bound:

A.1. if no other arrivals occur

$$e^{-\lambda(t-T^*)}(\mu V(\alpha, \beta) + (1-\mu)V(\alpha))$$

A.2. if exactly one arrival occurs and the first arrival was asynchronous and the player decides to play α in the second arrival

$$\lambda\mu e^{-\lambda(t-T^*)}(t-T^*)(pV(\alpha,\beta)+pV(\alpha)+(1-2p)V(\alpha))$$

A.3. if exactly one arrival occurs and the first arrival was asynchronous and the player decides to play β in the second arrival

$$\lambda\mu e^{-\lambda(t-T^*)}(t-T^*)(pV(\beta)+pV(\alpha)+(1-2p)V(\beta,\alpha))$$

A.4. if exactly one arrival occurs and the first arrival was simultaneous then

$$\lambda(1-\mu)e^{-\lambda(t-T^*)}(t-T^*)V(\alpha)$$

A.5. if two or more arrivals arrives

$$(1-e^{-\lambda(t-T^*)}(1+t-T^*))\underline{u}$$

Note that (3.2) and (3.2) are equal because of the indifference condition.

Now consider an upper bound on the payoff from playing β:

B.1. if no other arrivals occur

$$e^{-\lambda(t-T^*)}(\mu V(\beta)+(1-\mu)V(\beta,\alpha))$$

B.2. if exactly one arrival occurs and the first arrival was asynchronous:

$$\lambda\mu e^{-\lambda(t-T^*)}(t-T^*)U$$

where U is the maximum of the following four terms:

$$V(\beta) \tag{3.1}$$

$$pV(\beta) + pV(\alpha,\beta) + (1-2p)V(\alpha,\beta) \tag{3.2}$$

$$pV(\beta,\alpha) + pV(\beta) + (1-2p)V(\beta,\alpha) \tag{3.3}$$

$$pV(\beta,\alpha) + pV(\alpha,\beta) + (1-2p)V(\alpha) \tag{3.4}$$

B.3. if exactly one arrival occurs and the first arrival was simultaneous then

$$\lambda(1-\mu)e^{-\lambda(t-T^*)}(t-T^*)\max\{pV(\beta)+pV(\alpha)+(1-2p)V(\alpha,\beta), pV(\beta,\alpha)+(1-p)V(\alpha)\}$$

B.4. if two or more arrivals arrives

$$(1-e^{-\lambda(t-T^*)}(1+t-T^*))V(\alpha)$$

The case in which two or more arrivals comes is negligible since it is a second order term for t sufficiently close to T^*. So we ignore that term for now. Also note that terms (A.1) and (B.1) are exactly the same because of the indifference condition. Moreover (A.2) and (A.3) are again exactly the same because of the indifference condition. So the relevant terms for the calculation are (A.2),(A.3), (A.4) versus (B.2) and (B.3). Therefore if we can show that

$$(A.2) + (A.4) > (B.2) + (B.3)$$

then we have shown that there exists some $T^{**} > T^*$ such that whenever $t \in (T^*, T^{**}]$, if an arrival comes at t and the PAP is such that at least one person is playing α, then the player will choose to play α. Let us prove the above statement.

First consider the expressions (3.2) - (3.4). Note that each of these expressions is dominated by either $pV(\alpha,\beta) + pV(\alpha) + (1-2p)V(\alpha))$ or $(pV(\beta) + pV(\alpha) + (1-2p)V(\beta,\alpha)$. Since (A.4) $>$ (B.3), we have the above statement for when U is equal to one the expressions (3.2) - (3.4). So to prove the above statement all we have left to check is the case in which

$U = V(\beta)$. Consider first the case in which

$$(B.3) = \lambda(1-\mu)e^{-\lambda(t-T^*)}(t-T^*)(pV(\beta,\alpha)+(1-p)V(\alpha)).$$

In this case we have

$$\begin{aligned}
& (A.2)+(A.4) > (B.2)+(B.3) \\
\Leftrightarrow\ & \mu(p+(1-2p)V(\beta,\alpha))+(1-\mu) > (1-\mu)(pV(\beta,\alpha)+(1-p)) \\
\Leftrightarrow\ & \mu p+\mu(1-2p)V(\beta,\alpha)+(1-\mu)p > (1-\mu)pV(\beta,\alpha) \\
\Leftrightarrow\ & \mu^2+\mu(1-\mu)V(\beta,\alpha)+(1-\mu)\mu > (1-\mu)\mu V(\beta,\alpha) \\
\Leftrightarrow\ & \mu > 0.
\end{aligned}$$

Therefore for this case the above is satisfied. Finally let us check the case in which

$$(B.3) = \lambda(1-\mu)e^{-\lambda(t-T^*)}(t-T^*)(pV(\beta)+pV(\alpha)+(1-2p)V(\alpha,\beta)).$$

In this case,

$$\begin{aligned}
& (A.2)+(A.4) > (B.2)+(B.3) \\
\Leftrightarrow\ & 1 > (1-2\mu)V(\alpha,\beta)
\end{aligned}$$

Using the indifference condition, we can rewrite the above inequality as

$$\begin{aligned}
& 1 > (1-\mu)V(\alpha,\beta)-\mu V(\alpha,\beta) \\
\Leftrightarrow\ & 1 > (1-\mu)V(\alpha,\beta)+(1-\mu)-(1-\mu)V(\beta,\alpha) \\
\Leftrightarrow\ & \mu > (1-\mu)(V(\alpha,\beta)-V(\beta,\alpha)) \\
\Leftrightarrow\ & \mu > (1-\mu)e^{-2\lambda pT^*}(2q(T^*)-1)(u(\alpha,\beta)-u(\beta,\alpha)).
\end{aligned}$$

Note that one of the two inequalities above is satisfied trivially when μ is sufficiently close to 1/2 or when μ is sufficiently close to 1. More precisely, when

$$1 > (u(\alpha,\beta) + u(\beta,\alpha))\left(\frac{1}{2} - \mu\right) \tag{3.5}$$

$$\Leftrightarrow \quad \frac{1}{2} - \frac{1}{u(\alpha,\beta) + u(\beta,\alpha)} > \mu \tag{3.6}$$

and $u(\alpha,\beta) > u(\beta,\alpha)$ then the first inequality is satisfied, and when

$$\frac{\mu}{1-\mu} > u(\alpha,\beta) - u(\beta,\alpha) \tag{3.7}$$

the second inequality is satisfied.

The result above then allows to us to continue the backward induction in the same way that we did in the previous section. Note that the above discussion establishes the existence of some $T^{**} > T^*$ such that at all times $t \leq T^**$, both players play α when the opponent's PAP is α. Furthermore at all times $t \in (T^*, T^{**}]$, players continue to play α when their own PAP is α. As in the arguments previously, take T^{**} to be the supremum over all such times.

Then the continuation values at $t \leq T^{**}$ are given by the following:

$$\begin{aligned} W(\alpha,t) &= 1 \\ W(\beta,\alpha,t) &= e^{-\lambda(1-p)(t-T^*)}V(\beta,\alpha) + (1-e^{-\lambda(1-p)(t-T^*)})V(\alpha) \\ W(\alpha,\beta) &= e^{-\lambda(1-p)(t-T^*)}V(\alpha,\beta) + (1-e^{-\lambda(1-p)(t-T^*)})V(\alpha) \end{aligned}$$

Let $t > T^{**}$ and consider a history at which the PAP of player 2 is α_2. It is easy to show after such histories that player 1 will strictly prefer to play α_1. Because of symmetry, so will player 2.

Now consider a history at which the PAP of player 1 is α_1. By playing α_1, player 1 can guarantee himself a payoff of:

$$e^{-\lambda(t-T^**)}(\mu W(\alpha_1,\beta_2,T^{**}) + (1-\mu)W(\alpha,T^{**})) + (1-e^{-\lambda(t-T^*)})\underline{u}_1$$

whereas playing β_1 gives him a payoff at most:

$$e^{-\lambda(t-T^{*}*)}(\mu W(\beta, T^{**}) + (1-\mu)W(\beta_1, \alpha_2, T^{**})) + (1 - e^{-\lambda(t-T^{*})})u_1(\alpha).$$

We obtain a contradiction of the definition of T^{**} if

$$\mu W(\alpha_1, \beta_2, T^{**}) + (1-\mu)W(\alpha, T^{**}) > \mu W(\beta, T^{**}) + (1-\mu)W(\beta_1, \alpha_2, T^{**}).$$

Therefore we must have

$$\mu W(\alpha_1, \beta_2, T^{**}) + (1-\mu)W(\alpha, T^{**}) = \mu W(\beta, T^{**}) + (1-\mu)W(\beta_1, \alpha_2, T^{**}).$$

But in this case we can use the exact same argument that started the induction at T^* again contradicting the definition of T^{**}. Therefore such a T^{**} must be infinite. This means that in any equilibrium, both players must play α when either player's PAP is α at all times $t > T^*$. At times $t \leq T^*$, both players match the opponent's PAP. □

We conclude this discussion with the statement of the main theorem.

Theorem 3.6. *Suppose that α is strictly μ-dominant and that either $\mu \leq 1/2$ or that either condition 3.5 or 3.7 is satisfied. Then the set of payoffs shrinks to $\{u(\alpha)\}$ as the time horizon approaches infinity. If α is $(1-\mu)$-dominant, we obtain the same conclusion without any additional assumptions.*

Note that the theorem above leaves open cases in which both α and β are strictly μ-dominant but neither conditions 3.5 nor 3.7 are satisfied. Our conjecture thus far is that even in this region of parameters, a unique selection theorem obtains. However it currently remains an open question.

4 Non-selection

In the previous section, we developed sufficient conditions for unique selection of the Pareto dominant payoff profile. We show in this section that those conditions are necessary and sufficient. Before we begin the discussion, note first that the construction of an equilibrium in which the Pareto dominant payoff profile of $u(\alpha)$ is an easy exercise. The question that we answer in this section is whether other payoff profiles can be obtained in equilibrium for arbitrarily large time horizons when α is no longer strictly μ-dominant.

Theorem 4.1. *If β is $(1-\mu)$-dominant, then there exists an equilibrium in which both players always choose β_i regardless of the history and time of arrival.*

Note that proposition 2.1 in CL proposes a very similar statement. The proposition in the context of our model and notation states that for any generic game, there exists some $\bar{\mu}$ such that for all $\mu \leq \bar{\mu}$, the infinite repetition of any static Nash equilibrium is an equilibrium outcome. Note that in the context of our model, this statement is a consequence of the theorem above. Therefore our theorem is more general when attention is restricted to the class of games that we consider in our paper. However their result applies to to any revision game with a component game that is not of the form considered in this paper.

Proof. We use the one stage deviation principle adapted to these games to prove the above proposition. Consider a history at which the currently prepared action profile is such that the opponent has prepared α_{-i} and suppose that at such a history player i obtains an arrival at time t. Now suppose that the opponent is playing the strategy of always playing β_i regardless of the history and time of arrival. Then a "one-stage" deviation of playing α_i at time t and playing β_i upon further arrivals in the future gives a payout of

$$V_\mu(t) \equiv \big(e^{-\lambda t}u_1(\alpha) + e^{-\lambda(1-p)t}(1-e^{-\lambda pt})(u_1(\beta_i,\alpha_{-i}) + u_1(\alpha_i,\beta_{-i})) + (1-2e^{-\lambda(1-p)t} + e^{-\lambda t})u_1(\beta)\big)$$

with probability μ and

$$V_{1-\mu}(t) \equiv \big(e^{-\lambda(1-p)t}u_1(\alpha_i,\beta_{-i}) + (1-e^{-\lambda(1-p)t})u_1(\beta)\big)$$

with probability $(1-\mu)$. By playing according to the specified strategy, player i obtains an expected payout of

$$V(t) \equiv \mu(e^{-\lambda(1-p)t}u_1(\beta_i,\alpha_{-i}) + (1-e^{-\lambda(1-p)t})u_1(\beta)) + (1-\mu)u_1(\beta).$$

Note that

$$V(0) > \mu V_\mu(0) + (1-\mu)V_{1-\mu}(0)$$

by $(1-\mu)$-dominance of β. Note further that

$$\begin{aligned} & V(t) - \mu V_\mu(0) - (1-\mu)V_{1-\mu}(0) \\ = \; & \mu e^{-\lambda t}(u_i(\beta_i, \alpha_{-i}) - u_i(\alpha)) + e^{-\lambda(1-p)t}\left(1 - \mu e^{-\lambda p t}\right)(u_i(\beta) - u_i(\alpha_i, \beta_{-i})) \end{aligned}$$

which is greater than zero if and only if

$$(u_i(\beta_i, \alpha_{-i}) - u_i(\alpha)) + \left(e^{\lambda p t} - \mu\right)(u_i(\beta) - u_i(\alpha_i, \beta_{-i})) > 0.$$

Since β is a Nash equilibrium, the above expression is increasing in t and because it is greater than zero at $t = 0$, we can conclude that

$$V(t) - \mu V_\mu(0) - (1-\mu)V_{1-\mu}(0) > 0$$

for all t. Therefore we have shown that playing β_i whenever the opponent has prepared α_{-i} is incentive compatible against always playing β_{-i}. Showing that playing β_i is a best response whenever the opponent has prepared β_{-i} is trivial. □

References

[1] **Abreu, Dilip and Faruk Gul** "Bargaining and Reputation" *Econometrica*, 2000, *68*, 85-117.

[2] — **and David Pearce** "Bargaining, Reputation and Equilibrium Selection in Repeated Games with Contracts" *Econometrica*, 2007, *75*, 653-710.

[3] **Ambrus, Attila and Shih-En Lu**, "A continuous-time model of multilateral bargaining," 2009, *mimeo.*

[4] **Aumann, Robert and Seigu Hart** "Long Cheap Talk" *Econometrica*, 2003, *71*, 1619-1660.

[5] **Benoit, Jean-Pierre and Vijay Krishna** "Finitely Repeated Games," *Econometrica*, 1985, *53*, 905-922.

[6] **Calcagno, Riccardo and Stefano Lovo** "Preopening and Equilibrium Selection", 2010, mimeo.

[7] **Carlsson, Hans and Eric van Damme** "Global Games and Equilibrium Selection," *Econometrica*, 1993, *61*, 989-1018.

[8] **Caruana, Guillermo and Liran Einav** "A Theory of Endogenous Commitment," *Review of Economic Studies*, 2008, *75*, 99-116.

[9] **Chao, Yuen R.** "A Note on "Continuous Mathematical Induction"" *Bulletin of the American Mathematical Society*, 1919, *26*, 17-18.

[10] **van Damme, Eric and Sjaak Hurkens** "Commitment Robust Equilibria and Endogenous Timing", *Games and Economic Behavior*, 1996, *15*, 290-311.

[11] **Dutta, Prajit K.** "A Folk Theorem for Stochastic Games," *Journal of Economic Theory*, 1995, 66, 1-32.

[12] **Dutta, Prajit K.** "Coordination Need Not Be A Problem," 2003, mimeo.

[13] **Farrell Joseph** "Cheap talk coordination and entry", *RAND Journal of Economics*, 1987, *18*, 34-39.

[14] **Farrell, Joseph, and Garth Saloner**, "Standardization, Compatibility, and Innovation," *RAND Journal of Economics*, 1985, *16*, 70-83.

[15] **Fudenberg, Drew and Wojciech Olszewski** "Repeated Games with Asynchronous Monitoring" *Games and Economic Behavior*, 2011, *72*, 86-99.

[16] **Harsanyi, John, C. and Reinhard Selten** "A general Theory of Equilibrium Selection in Games," MIT press, Cambridge, MA. 1988.

[17] **Kamada, Yuichiro and Michihiro Kandori**, "Revision Games," 2010, mimeo.

[18] — **and Takuo Sugaya** "Asynchronous Revision Games with Deadline: Unique Equilibrium in Coordination Games," 2010a, mimeo.

[19] — **and** —, "Policy Announcement Game: Valence Candidates and Ambiguous Policies," 2010b, mimeo.

[20] **Kandori, Michihiro, George J. Mailath, and Rafael Rob** "Learning, Mutation, and Long Run Equilibria in Games," *Econometrica*, 1993, *61*, 29-56.

[21] **Lagunoff, Roger and Akihiko Matsui** "Asynchronous Choice in Repeated Coordination Games " *Econometrica*, 1997, *65*, 1467 - 1477.

[22] **Lipman, Barton, L. and Ruqu Wang** "Switching Costs in Frequently Repeated Games," *Journal of Economic Theory*, 2000, *93*, 149-190.

[23] **Matsui, Akihiko and Kiminori Matsuyama** "An Approach to Equilibrium Selection," *Journal of Economic Theory*, 1995, *65*, 415-434.

[24] **Morris, Stephen, Rafael Rob, and Hyun Song Shin** "p-Dominance and Belief Potential," *Econometrica*, 1995, *63*, 145-167.

[25] — **and Hyun Song Shin** "Unique Equilibrium in a Model of Self-Fulfilling Currency Attacks" *American Economic Review*, 1998, *88*, 587-597.

[26] **Oyama, Daisuke, Satoru Takahashi, and Josef Hofbauer** "Monotone Methods for Equilibrium Selection under Perfect Foresight Dynamics," *Theoretical Economics*, 2008, *3*, 155-192.

[27] **Rabin, Matthew** "A model of pre-game communication", *Journal of Economic Theory*, 1994, *63*, 370-391.

[28] **Rubinstein, Ariel** "The Electronic Mail Game: A Game with Almost Common Knowledge," *American Economic Review*, 1989, *79*, 389-391.

[29] **Sugaya, Takuo and Satoru Takahashi** "Coordination Failure in Repeated Games with Private Monitoring," 2009, mimeo

[30] **Takahashi, Satoru** “Infinite Horizon Common Interest Games with Perfect Information," *Games and Economic Behavior*, 2005, *53*, 231-247.

[31] **Wen, Quan** “A Folk Theorem for Repeated Sequential Games," *Review of Economic Studies*, 2002, *69*, 493-512.

[32] **Yoon, Kiho** “A Folk Theorem for Asynchronous Repeated Games," *Econometrica*, 2001, *69*, 191-200.

[33] **Young, Peyton H.** “The Evolution of Conventions,” *Econometrica*, 1993, *61*, 57-84.

[34] — “Conventional Contracts,” *Review of Economic Studies*, 1998, *65*, 773-792.