

Asynchronicity and Coordination in Common and Opposing Interest Games

Original

Asynchronicity and Coordination in Common and Opposing Interest Games / Calcagno, Riccardo; Kamada, Yuichiro; Lovo, Stefano; Takuo, Sugaya. - In: THEORETICAL ECONOMICS. - ISSN 1933-6837. - 9:(2014), pp. 409-434. [10.3982/TE1202]

Availability:

This version is available at: 11583/2812485 since: 2020-04-16T12:22:40Z

Publisher:

Society for Economic Theory

Published

DOI:10.3982/TE1202

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

Asynchronicity and coordination in common and opposing interest games

RICCARDO CALCAGNO

Department of Economics, Finance and Control, EMLYON Business School

YUICHIRO KAMADA

Cowles Foundation for Research in Economics, Yale University

STEFANO LOVO

Finance Department, HEC, Paris, and GREGHEC

TAKUO SUGAYA

Stanford Graduate School of Business

We study games endowed with a pre-play phase in which players prepare the actions that will be implemented at a predetermined deadline. In the preparation phase, each player stochastically receives opportunities to revise her actions, and the finally revised action is taken at the deadline. In two-player “common interest” games, where there exists a best action profile for all players, this best action profile is the only equilibrium outcome of the dynamic game. In “opposing interest” games, which are 2×2 games with Pareto-unranked strict Nash equilibria, the equilibrium outcome of the dynamic game is generically unique and corresponds to one of the stage-game strict Nash equilibria. Which equilibrium prevails depends on the payoff structure and on the relative frequency of the arrivals of revision opportunities for each of the players.

Riccardo Calcagno: calcagno@em-lyon.com

Yuichiro Kamada: yuichiro.kamada@yale.edu

Stefano Lovo: lovo@hec.fr

Takuo Sugaya: tsugaya@gsb.stanford.edu

This paper is the result of a merger between two independent projects: Calcagno and Lovo’s “Preopening and Equilibrium Selection” and Kamada and Sugaya’s “Asynchronous Revision Games with Deadline: Unique Equilibrium in Coordination Games.” We thank Dilip Abreu, Gabriel Carroll, Sylvain Chassang, Drew Fudenberg, Michihiro Kandori, Fuhito Kojima, Barton Lipman, Thomas Mariotti, Sebastien Pouget, Stephen Morris, Assaf Romm, Satoru Takahashi, Tristan Tomala, Nicolas Vieille, and, particularly, Johannes Hörner and Yuhta Ishii for useful comments and suggestions on either project. We also thank seminar participant at the GDR Conference in Luminy 2009, SAET Conference in Ischia 2009, Toulouse School of Economics, Bocconi University, Research in Behavior in Games seminar at Harvard University, Student Lunch Seminar in Economic Theory at Princeton University, and the 21st Stony Brook International Conference on Game Theory for helpful comments. We are grateful to three anonymous referees and a co-editor for insightful comments and suggestions that substantially improved this paper. Stefano Lovo gratefully acknowledges financial support from the HEC Foundation and from ANR Grant ANR-10-BLAN 0112.

KEYWORDS. Revision games, pre-opening, finite horizon, equilibrium selection, asynchronous moves.

JEL CLASSIFICATION. C72, C73.

1. INTRODUCTION

This paper studies a situation in which players prepare their actions in a pre-play phase before the actions are taken at a predetermined deadline. As the deadline approaches, each player has the opportunity to prepare an action at stochastically determined (Poisson) times. At the deadline, the actions most recently prepared are taken and the players' payoffs are determined only by these actions. In this framework, called the *revision game*, Kamada and Kandori (2010) show that the addition of pre-play phase can *widen* the set of achievable payoffs. This paper uncovers another role that the pre-play phase can play. We show that it can *narrow down* the set of achievable payoffs.

We study this problem in two classes of games where coordination is an issue. The first is *common interest* games, in which there is an action profile that all players strictly prefer to all other profiles. For this class of games, we show that, in two-player games, this “best profile” is the unique outcome of the revision game. The second class of games is that of *opposing interest* games, which are 2×2 games with two Pareto-unranked strict Nash equilibria. In this class of games, we show that generically there is a unique outcome of the revision game, which corresponds to one of the strict Nash equilibria. Which equilibrium prevails in the revision game depends on the payoff structure and the relative arrival frequency of revision opportunities for each player. We prove these results by using a type of backward-induction argument in continuous time.

Three assumptions, in addition to the assumption that revision opportunities are stochastic, are crucial to our results. The first one is *observability*. If a player is unable to observe what the other player has prepared, then the revision phase has no binding force, and so the equilibrium outcomes of revision games are identical to those of static games. The second is *asynchronicity*. If all revision opportunities are synchronous, then any repetition of static Nash equilibria is subgame perfect. Hence, uniqueness does not hold if there are multiple static Nash equilibria. However, if opportunities are asynchronous, each player's preparation must be contingent on the opponent's current action (by observability). Thus a player can induce the opponent to prepare some particular action by using as a threat the possibility that she may not be able to revise her own action before the deadline. The third key ingredient is *finite actions and strict incentives*. As we will argue, uniqueness is implied by backward induction. If there are only finitely many actions and the static game best replies to pure actions are strict, then each player has a single best reply near the deadline (by asynchronicity) in the revision game, and this constitutes the starting point of our backward induction argument.

These assumptions seem natural in real-life and economic contexts where coordination is crucial. For example, such a situation arises in the daily practice of some financial markets, such as Nasdaq or Euronext, for example, where half an hour before the

opening of the market, participants are allowed to submit orders, which can be withdrawn and changed until opening time. These orders and the resulting (virtual) equilibrium trading price are publicly posted during the whole “pre-opening” period. Only orders that are still posted at the opening time are binding and, hence, executed. In this framework, it is natural to assume that orders are submitted asynchronously and that traders do not always manage to withdraw old orders or submit new orders instantaneously because it takes a certain random time to fill in the new order faultlessly. Observability holds as the posted orders are displayed on the screen, and the number of payoff-relevant orders is practically finite.^{1,2}

The rest of the paper is organized as follows. [Section 2](#) reviews the related literature. [Section 3](#) introduces the model. [Section 4](#) presents a simple but useful lemma that allows us to implement a backward-induction argument in continuous time. [Section 5](#) considers common interest games and [Section 6](#) studies two-player opposing interest games. [Section 7](#) discusses further results. Some of the proofs are relegated to the [Appendix](#).

2. LITERATURE REVIEW

Cheap talk

It is important to make a distinction between our model and cheap-talk models such as those in [Farrell \(1987\)](#), [Rabin \(1994\)](#), and [Aumann and Hart \(2003\)](#). In these models, players are allowed to be involved in pre-play nonbinding communication. In contrast, in our model, at each moment of time, the prepared action will become the final payoff-relevant action with a strictly positive probability. For this precise reason, the outcome can be affected by the addition of a revision phase in our model.

Equilibrium selection

It is instructive to compare our selected outcome with those in the equilibrium selection literature. In many works on equilibrium selection, the risk-dominant equilibria of [Harsanyi and Selten \(1988\)](#) are selected in 2×2 games. In our model, however, a different answer is obtained: a strictly Pareto-dominant Nash equilibrium is taken even when it is risk-dominated. Roughly speaking, since we assume perfect and complete information with nonanonymous players, there is only a very small “risk” of miscoordination when the deadline is far. There are three lines of the literature in which risk-dominant equilibria are selected: models of global games, stochastic learning models with myopia, and

¹Given this application, [Calcagno and Lovo \(2010\)](#) call the revision game a pre-opening game.

²[Biais et al. \(2014\)](#) present an experiment that simulates pre-opening in a financial market where the actual play is preceded by (only) one round of pre-play communication, which is either completely binding or completely nonbinding. In both specifications, players choose their actions simultaneously and there are multiple subgame perfect equilibria (SPE). Consistently, they find both Pareto efficient and Pareto inefficient outcomes are observed in the experiment.

models of perfect foresight dynamics.^{3,4} Since the model of perfect foresight dynamics seems closely related to ours, let us discuss it here.

Perfect foresight dynamics

Perfect foresight dynamics are studied by Matsui and Matsuyama (1995) in evolutionary models in which players are assumed to be patient and “foresighted.” That is, they value the future payoffs and take best responses given (correct) beliefs about the future path of play.⁵ A continuum of agents is randomly and anonymously matched over an infinite horizon according to a Poisson process. In this setup, they select the risk-dominant action profile in 2×2 games with two Pareto-ranked (static) Nash equilibria. The key difference is that they assume anonymous agents while we assume nonanonymous agents. For the “best action profile” to be selected in our model, it is important for each player to expect that if she prepares an action corresponding to the best profile, then that preparation can affect the other player’s future preparation. This strategic consideration is absent when players are anonymous.

Common interest games and asynchronous moves

Farrell and Saloner (1985) and Lagunoff and Matsui (1997) are early works on the topic of obtaining the unique outcome in common interest games.⁶ Dutta (2012) shows convergence to the unique outcome and Takahashi (2005) proves uniqueness of subgame perfect equilibria when players move asynchronously. One difference is that we assume a stochastic order of moves while they consider a fixed order. Also, we obtain a uniqueness result more generally than do Lagunoff and Matsui (1997) due to the finite horizon.

³The literature on global games was pioneered by Rubinstein (1989), and analyzed extensively in Carlsson and van Damme (1993), Morris and Shin (1998), and Sugaya and Takahashi (2009). They show that the lack of almost common knowledge due to incomplete information can select an equilibrium. The type of incomplete information they assume is absent in our model. Stochastic learning models with myopia are analyzed in Kandori et al. (1993) and Young (1993). They consider a situation in which players interact repeatedly, and each player’s action at each period is stochastically perturbed. The key difference between their assumptions and ours is that in their model players are myopic while we assume that players prepare actions anticipating their opponents’ future moves. In addition, the game is repeated infinitely in their models, while the game is played once and for all in our model.

⁴As an exception, Young (1998) shows that in the context of contracting, his evolutionary model does not necessarily lead to the risk-dominant equilibrium (p -dominant equilibrium in Morris et al. 1995). But he considers a large anonymous population of players and repeated interaction, so the context he focuses on is very different from the one of this paper.

⁵See also Oyama et al. (2008).

⁶Dutta’s (1995) result implies that this result in Lagunoff and Matsui (1997) is due to the lack of full dimensionality of the feasible and individually rational payoff set. See also Lagunoff and Matsui (2001), Yoon (2001), and Wen (2002). Rubinstein and Wolinsky (1995) show that even when the discount factor is arbitrarily close to 1, the set of SPE payoff vectors of the repeated games that result from the repetition of the extensive form game may not coincide with those that result from the normal form game if the individually rational payoffs are different or full dimensionality is not satisfied.

War of attrition with incomplete information

The intuition behind the result for opposing interest games is similar to that for the “war of attrition with incomplete information.”⁷ Although the structure of the equilibria in war of attrition is similar to the structure of the equilibria in our model, the reasoning is different: in our model, players use the probability of not having future revision opportunities as a “commitment power” while the literature in the war of attrition assumes the existence of “commitment types” a priori.

Switching cost

Our model assumes that no cost is associated with revision. Several papers consider a finite-horizon model with switching cost and show that a unique outcome prevails in their respective games. Typically, the existence of switching cost results in different implications on equilibrium behavior. See, for example, [Lipman and Wang \(2000\)](#) and [Caruana and Einav \(2008\)](#) for details.

Revision games

[Kamada and Kandori \(2010\)](#) introduce the model of revision games. They show that, among other things, non-Nash “cooperative” action profiles can be taken at the deadline when a certain set of regularity conditions is satisfied. Hence, their focus is on expanding the set of equilibria when the static Nash equilibrium is inefficient relative to non-Nash profiles. We ask a very different question in this paper: we consider games with multiple efficient static Nash equilibria and ask which of these equilibria is selected.⁸ What drives this difference is that the action space is finite in our paper, whereas it is not in [Kamada and Kandori \(2010\)](#). [Kamada and Sugaya \(2010b\)](#) consider a revision game model in which the players have finite action sets in the context of an election campaign. The main difference with our paper is that they assume once a player changes her action, she cannot revise it further. Thus the characterization of the equilibrium is essentially different from the analysis in the present paper because in our model, when another opportunity arrives, a player can always change her prepared action to the previously prepared action.⁹

Further results

Finally, further results beyond what we have in this paper can be found in either or both [Calcagno and Lovo \(2010\)](#) and [Kamada and Sugaya \(2010a\)](#). We refer to these papers whenever appropriate.

⁷For example, among others, see [Abreu and Gul \(2000\)](#) and [Abreu and Pearce \(2007\)](#).

⁸See also [Ambrus and Lu \(2009\)](#) for a variant of the revision games model of bargaining in which the game ends when an offer is accepted.

⁹[van Damme and Hurkens \(1996\)](#) analyze a related model of “timing games,” in which players can choose the timing of their move out of two periods and they cannot switch back.

3. THE MODEL

We consider a two-player normal-form finite game $((X_i)_{i=1,2}, (u_i)_{i=1,2})$ (the *component game*) where X_i is the finite set of player i 's actions with $|X_i| \geq 2$ and $u_i: X \rightarrow \mathbb{R}$ is player i 's utility function. Here, $X = X_1 \times X_2$ is the set of action profiles. We use a female pronoun for player 1 and a generic player i , and use a male pronoun for player 2.

Before players take actions, they need to “prepare” their actions. We model this situation as in Kamada and Kandori (2010): time is continuous, $t \in [-T, 0]$, and the component game is played once and for all at time 0. The *revision game* proceeds as follows. First, at time $-T$, the initial action profile is exogenously given.¹⁰ In the time interval $(-T, 0]$, each player independently obtains opportunities to revise her prepared action according to two random Poisson processes \mathbf{p}_1 and \mathbf{p}_2 with arrival rates λ_1 and λ_2 , respectively, where $\lambda_i > 0$, $i = 1, 2$. As the Poisson processes \mathbf{p}_1 and \mathbf{p}_2 are independent, the probability that the two players revise their actions simultaneously is zero. In other words, almost certainly only asynchronous revision opportunities arise.¹¹ 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.

So as to define the strategy space of the revision game, suppose the game has reached time t . We assume here that each player i at any time t has perfect information about all past events. In particular, she knows whether i has a revision opportunity at t but does not know whether the opponent gets an opportunity at t . Formally, for any given $t \in (-T, 0]$, let $H_i^n(t)$ and $H_i^r(t)$ denote the subset of all possible histories for player i such that she does not have a revision opportunity at t and that she does have a revision opportunity at t , respectively. Thus, the set of all possible histories for player i is $H_i := \bigcup_{t \in [-T, 0]} H_i^n(t) \cup H_i^r(t)$. A *history* for player i at $t \in (-T, 0]$ takes either of the following two forms, depending on whether i receives an opportunity at t :

$$h_i(t) = ((t_1^k, x_1^k)_{k=0}^{k_1}, (t_2^k, x_2^k)_{k=0}^{k_2}) \in H_i^n(t)$$

if player i does not have a revision opportunity at t and

$$h_i(t) = ((t_1^k, x_1^k)_{k=0}^{k_1}, (t_2^k, x_2^k)_{k=0}^{k_2}, t) \in H_i^r(t)$$

if she does, where $t_i^0 := -T$, k_i is a nonnegative integer for $i = 1, 2$, $-T < t_i^1 < t_i^2 < \dots < t_i^{k_i} < t$ for $i = 1, 2$, and $x_i^k \in X_i$ for $i = 1, 2$. The interpretation is that x_i^0 is the exogenously given action for player i at time $t_i^0 = -T$, k_i is the number of revision opportunities that i has received in the time interval $(-T, t)$, t_i^k is the time at which player i received her k th revision opportunity, and x_i^k is the action player i prepared at t_i^k .

A strategy for player i is a mapping $\sigma_i: H_i \rightarrow \{\emptyset\} \cup \Delta(X_i)$, where $\sigma_i(h_i(t)) = \emptyset$ if $h_i(t) \in H_i^n(t)$ and $\sigma_i(h_i(t)) \in \Delta(X_i)$ if $h_i(t) \in H_i^r(t)$.

¹⁰As we will see, the uniqueness results in Sections 5 and 6 become even sharper if players simultaneously choose actions at time $-T$.

¹¹We refer to Section 7 for a discussion of the role played by this assumption. See Calcagno and Lovo (2010) and Ishii and Kamada (2011) for more general processes that underlie the arrival of revision opportunities.

For any given history $h_i(t)$, let $x_i(t) := x_i^{k_i} \in X_i$ be player i 's prepared action that resulted from his last revision opportunity (strictly) before t . We denote by $x(t) := \{x_i(t)\}_{i=1,2}$ the last prepared action profile (PAP) before time t . Note that $x_i(t)$ is player i 's payoff-relevant action at $t = 0$ in the event i receives no further revision opportunities in the time interval $[t, 0]$.

Our main results concern subgame perfect equilibrium (SPE) of the revision game for the case when T is large. However, we note that the model with arrival rate (λ_1, λ_2) and horizon length T is essentially equivalent to the model with arrival rates $(a\lambda_1, a\lambda_2)$ and horizon length T/a , for any positive constant a .¹² Hence, all our results obtained for T large enough and fixed revision frequencies (λ_1, λ_2) can be obtained by keeping fixed the horizon T and having revisions frequent enough.

To avoid ambiguity, in the rest of the paper, we use the term *revision equilibrium* for a SPE of the whole revision game and *(strict) Nash equilibrium* for a (strict) Nash equilibrium of the component game.

4. BACKWARD INDUCTION IN CONTINUOUS TIME

The proofs of our main results rely on the idea of backward induction. The standard backward-induction argument starts by proving a statement for the last period and then, given the statement is true, proves the statement for the second-last period and so forth. However, this argument is not immediately applicable to our continuous-time setting, as there is no obvious definition of "second-last period." In this section, we present a lemma that allows us to implement a type of backward-induction argument in continuous time. The proof is given in [Appendix A.1](#).

LEMMA 1. *Suppose for every $t \in (-T, 0]$, there exists $\epsilon > 0$ such that statement $A_{t'}$ is true for all $t' \in (t - \epsilon, t]$ if statement $A_{t''}$ is true for every $t'' > t$. Then, for every $t \in (-T, 0]$, statement A_t is true.*

Note that the ϵ in the statement of the lemma can depend on t . Hence, in particular, the lemma goes through even though the required ϵ shrinks to zero as t approaches some finite constant and then jumps discontinuously there.¹³

5. COMMON INTEREST GAMES

In this section, we consider a component game with an action profile that strictly Pareto-dominates all other action profiles. Formally, we say that an action profile x^* is *strictly Pareto-dominant* if $u_i(x^*) > u_i(x)$ for all i and all $x \in X$ with $x \neq x^*$. We say that a game is a *common interest game* if it has a strictly Pareto-dominant action profile. Notice that if x^* is strictly Pareto-dominant, then it is a strict Nash equilibrium.

For example, the games in [Figure 1](#) are common interest games. In each case, (U, L) is strictly Pareto-dominant.

The first main result of this paper is the following.

¹²See the "arrival rate invariance" result discussed in [Kamada and Kandori \(2010\)](#).

¹³A version of the lemma that switches the order of quantifiers (so that ϵ cannot depend on t) appears in [Chao \(1919\)](#).

	L	R		L	C	R
U	2, 2	-10, 1	U	2, 2	1, 0	-4, 1
D	1, -10	1, 1	M	1, 1.5	1, 1	-3, 1
			D	1, 0	0, 1	0, 0

FIGURE 1. Common interest games.

THEOREM 1. *Consider a common interest component game and let x^* be the strictly Pareto-dominant action profile. Then for any $\epsilon > 0$, there exists $T' > 0$ such that for all $T > T'$, in all revision equilibria, $x(0) = x^*$ with probability higher than $1 - \epsilon$.*

5.1 Intuition

The proof consists of two steps. First, we show that x^* is absorbing in the revision game: since the action space is finite, the difference between $u_i(x^*)$ and i 's second best payoff is strictly positive. Therefore, when the PAP is x^* , no player wants to prepare another action, and to create a possibility that she cannot have further revision opportunities and will be forced to take a second best or even worse action profile.

Second, given the first step, each player i knows that if her opponent $-i$ has a revision opportunity while player i prepares x_i^* , then the opponent will prepare x_{-i}^* . Hence, the lower bound of the equilibrium payoff for each player is given by always preparing x_i^* whenever she receives a revision opportunity. If T is sufficiently large, then this strategy gives her a payoff very close to $u_i(x^*)$, which means x^* should be taken with high probability at the deadline in any revision equilibrium.

5.2 Proof of Theorem 1

Now we offer the formal proof. The two steps in the formal proof correspond to those in the intuitive explanation above.

Step 1. Let $m := \min_{i, x \neq x^*} (u_i(x^*) - u_i(x))$, the minimum across players of the differences between the best payoff and the second best payoff. Since X is finite and x^* is strictly Pareto-dominant, m is well defined and positive.

Fix $t \leq 0$ arbitrarily. Suppose that for all time after $t < 0$, each player i has a strict incentive to prepare x_i^* if the opponent $-i$'s prepared action is x_{-i}^* .¹⁴ Suppose also that player i obtains a revision opportunity at time $t - \epsilon$ and $-i$'s prepared action is x_{-i}^* . Then the payoff from preparing x_i^* is at least

$$u_i(x^*) - (1 - e^{-(\lambda_i + \lambda_{-i})\epsilon})M, \quad (1)$$

where $M := \max_{i, x \neq x^*} (u_i(x^*) - u_i(x)) < \infty$, because with probability at least $e^{-(\lambda_i + \lambda_{-i})\epsilon}$, no further revision opportunities arrive between $t - \epsilon$ and t , and the PAP at time t is x^* . In such a case, action x^* will be taken at the deadline by assumption. On the other hand, the payoff from preparing an action $x_i \neq x_i^*$ is at most

$$u_i(x^*) - e^{-\lambda_i(-t+\epsilon)}m, \quad (2)$$

¹⁴For $t = 0$, this is vacuously true.

because with probability $e^{-\lambda_i(-t+\epsilon)}$, player i never has a revision opportunity again and in such a case, the action profile at the deadline cannot be x^* .

Notice that expression (1) is strictly greater than expression (2) for $\epsilon = 0$. Also by the continuity of (1) and (2) with respect to ϵ , there exists $\epsilon' > 0$ such that for all $\epsilon \in (0, \epsilon')$, expression (1) is strictly greater than expression (2).¹⁵ Hence, by Lemma 1, we have that for any $t < 0$, each player i has a strict incentive to prepare x_i^* if her opponent $-i$ prepares x_{-i}^* .

Step 2. Since in any subgame perfect equilibrium, players can guarantee at least the payoff that can be obtained by always playing the action x_i^* , it suffices to show that the payoff of always preparing x_i^* converges to the strictly Pareto-dominant payoff as T goes to infinity. This implies that the probability of the action being x^* at the deadline converges to 1, as desired.

By Step 1, the action profile x^* is the absorbing state: each player has a strict incentive to prepare x_i^* if her opponent $-i$ prepares x_{-i}^* . In two-player games, since player i is the unique opponent of player $-i$, player $-i$ prepares x_{-i}^* if player i prepares x_i^* . Therefore, the payoff of always preparing x_i^* guarantees a payoff that converges to $u_i(x^*)$.

5.3 Remarks

Four remarks are in order at this stage.

First, if players choose their actions at $-T$, then we can pin down the behavior of players on the equilibrium path. In fact, in Appendix A.2, we show that in a common interest game defined as above, players prepare the strictly Pareto-dominant profile x^* at all times $t \in [-T, 0]$ on the (unique) path of play in any revision equilibrium.

Second, notice that if there exist two strict Pareto-ranked Nash equilibria in a 2×2 component game, then the game is a common interest game. Hence, in such a case, the Pareto-superior Nash equilibrium is the outcome of the revision game.¹⁶

Third, the outcome of the revision game is the strictly Pareto-dominant profile even if it is risk-dominated by another Nash equilibrium. For example, in the left payoff matrix in Figure 1, the action profile (U, L) is risk-dominated while it is the equilibrium outcome of the revision game. The key is that whenever a player prepares x_i^* (the action that corresponds to the Pareto-dominant profile), the opponent will move to the Pareto-dominant profile whenever she can revise and they stay at this profile until the deadline (Step 1 of the proof in the previous subsection). Therefore, if the remaining time is sufficiently long, then the “risk of miscoordination” by preparing x_i^* is arbitrarily small (Step 2).

Fourth, notice that we allow for the component game to be different from a pure coordination game (i.e., a game in which two players have identical payoff functions). This result is in stark contrast to the result of Lagunoff and Matsui (1997), which applies

¹⁵Note that here we again use the assumption that the action space is finite, so that the maximum payoff difference is bounded.

¹⁶Note that Kamada and Kandori (2010) prove that if each player has a strictly dominant action when the action space is finite, then it is taken in asynchronous revision games.

only to pure coordination games (see [Yoon 2001](#)). This difference comes from the different assumptions on the horizon: since their models have an infinite horizon, there can be an infinite sequence of punishments. On the other hand, in our model, there is a deadline so the incentives near the deadline can be perfectly pinned down as x^* is strictly Pareto-dominant. Hence, we can implement backward induction starting from the deadline.

5.4 n -player case

In this subsection, we examine how the result in this section can be generalized to the case of n players. The model setting and the strictly Pareto-dominant profile, denoted x^* , are analogously defined. Let I be the set of players and denote by λ_i the arrival rate of the revision opportunities for player i . Also, define $r_i = \lambda_i / \sum_{j \in I} \lambda_j$. We assume $|I| := n \geq 2$.

Step 1 of the proof is easily extended to the case of n players. However, Step 2 cannot be extended. With more than two players, if each player i currently prepares an action different from x_i^* , no single player can create a situation where it is enough for only one player to change her preparation so as to go to x^* . Hence, the proof in [Section 5.2](#) does not work. Below we show that if there are more than two players, a unique selection result is obtained when the preferences of the players are similar, where the measurement of similarity is given by the following definition.

DEFINITION 1. A common interest game is said to be a K -coordination game if for any $i, j \in I$ and $x \in X$,

$$\frac{u_i(x^*) - u_i(x)}{u_i(x^*) - \underline{u}_i} \leq K \frac{u_j(x^*) - u_j(x)}{u_j(x^*) - \underline{u}_j},$$

where $\underline{u}_i = \min_x u_i(x)$.

The minimum of K is 1 when the game is a pure coordination game, where the players have exactly the same payoff functions. As K increases, the preferences become less aligned.

Let $\alpha = \min_{i \in I} r_i$ (the smallest value of r_i) and $\beta = \min_{i \in I, i \neq j^*} r_i$ where j^* is an arbitrary member of $\arg \min_{i \in I} r_i$ (the second smallest value of r_i).

THEOREM 2. Suppose that a common interest game is a K -coordination game with the strict Pareto-dominant action profile x^* and

$$(1 - \alpha - \beta)K < 1 - \beta. \quad (3)$$

Then, for any $\varepsilon > 0$, there exists T' such that for all $T > T'$, in all revision equilibria, $x(0) = x^*$ with probability higher than $1 - \varepsilon$.

The proof is given in [Appendix A.3](#) and a detailed discussion can be found in [Kamada and Sugaya \(2010a\)](#).

	a_2	b_2		a_2	b_1
a_1	1, 1, 1	-1, -1, -1		-1, -1, -1	-1, 0, 0
b_1	-1, -1, -1	0, 0, -1		0, -1, 0	-1, -1, -1
	a_3			b_3	

FIGURE 2. A counterexample.

Several remarks regarding condition (3) are in order. The smaller K is, the more likely this condition is to be satisfied. In particular, if the game is a pure-coordination game (i.e., if $K = 1$), then it is always satisfied. In addition, for a fixed K and number of players, if the relative frequency r_i is more equally distributed, the condition is more likely to be satisfied. Note also that if we let $\lambda_i = \lambda_j$ for all $i, j \in I$ and fix $K > 1$, then it is more likely to be satisfied as the number of players becomes smaller. The condition is automatically satisfied in two-player games since $1 - \alpha - \beta = 0$, so that Theorem 2 implies Theorem 1.

To understand the need for condition (3), consider the game in Figure 2, where $\lambda_1 = \lambda_2 = \lambda_3 = \lambda$.

Consider the Markov perfect strategy profile in which each player's preparation depends only on the PAP (not on time t), and she prepares a different action than her current action if and only if the PAP includes exactly two a 's. It is straightforward to verify that this strategy profile is a SPE, and induces different outcomes depending on the exogenously given action profile at time $-T$.¹⁷ The key is that when the PAP is (a_i, a_j, b_k) , i prefers preparing b_i , which guarantees the payoff of 0, to sticking to a_i to bet on the lottery among the payoffs 1, 0, and -1 . This is because in this lottery, -1 is assigned a higher probability than the other two payoffs due to the possibility of no one getting a future revision opportunity. What is important is that the moves are stochastic, so there is no event at which, at the PAP (a_i, a_j, b_k) , player i is sure that the next mover is k and, hence, sticking to a_i induces (a_i, a_j, a_k) . That is, player i is exposed to the risk of player j switching to b_j in the future and i getting the payoff -1 .¹⁸ Note that player j would not switch to b_j if such a strategy would hurt her as well. But in this example, moving from (a_i, a_j, b_k) to (a_i, b_j, b_k) gives opposite consequences to players i and j (bad for player i but good for player j). That is, the fact that preferences are diverse is also important, which is the reason why we need to restrict to a class of games like K -coordination games.

Finally, we note that the above example implies that condition (3) is tight in the sense that for any nonpositive number, we can find a pair that consists of a component game and a profile of arrival rates such that the difference of left and right hand sides of condition (3), $(1 - \beta) - (1 - \alpha - \beta)K$, is equal to that number. In this game, for example, we have $K = 2$ and $r_i = \frac{1}{3}$ for all i . Hence, the difference is zero.¹⁹

¹⁷Specifically, the outcome is the same as the given action profile x at $-T$ unless x has exactly two a 's, in which case it is either an action profile with three a 's or exactly one a , depending on who moves the first revision opportunity.

¹⁸Such a player j does not exist in two-player games. This is why the unique selection result always holds when $n = 2$.

¹⁹Notice that we can replace -1 's in the above example by $-1 - l$ for any $l > 0$. This makes the left hand side of the above inequality strictly negative ($-l/3$).

	<i>L</i>	<i>R</i>		<i>L</i>	<i>R</i>
<i>U</i>	3, 3	0, 1	<i>U</i>	2, 1	0, 0
<i>D</i>	0, 5	1, 6	<i>D</i>	0, 0	1, 2

FIGURE 3. Opposing interest games.

	<i>L</i>	<i>R</i>
<i>U</i>	$u_1(U, L), u_2(U, L)$	$u_1(U, R), u_2(U, R)$
<i>D</i>	$u_1(U, D), u_2(U, D)$	$u_1(D, R), u_2(D, R)$

FIGURE 4. General opposing interest game.

6. OPPOSING INTEREST GAMES

In the previous section, we analyzed games in which a single action profile is best for both players. Now we turn to the class of games in which different players have different “best” action profiles. Examples of games that we consider in this section are given in Figure 3.

Generally, we consider two-player component games as in Figure 4 with two strict Nash equilibria (U, L) and (D, R) such that

$$u_1(U, L) > u_1(D, R) \quad \text{and} \quad u_2(U, L) < u_2(D, R). \quad (4)$$

The first inequality implies that player 1 strictly prefers (U, L) to (D, R) among pure Nash equilibria while the second implies that player 2’s preference is opposite. Note that since (U, L) and (D, R) are strict Nash equilibria of this component game, these conditions imply that (U, L) gives player 1 a strictly better payoff than any other action profile and that (D, R) gives player 2 a strictly better payoff than any other action profile.

Let

$$t_1^* = -\frac{1}{\lambda_1 + \lambda_2} \ln \left(\frac{\lambda_1 u_1(D, R) - u_1(U, R)}{\lambda_2 u_1(U, L) - u_1(D, R)} + \frac{u_1(U, L) - u_1(U, R)}{u_1(U, L) - u_1(D, R)} \right) \quad (5)$$

and

$$t_2^* = -\frac{1}{\lambda_1 + \lambda_2} \ln \left(\frac{\lambda_2 u_2(U, L) - u_2(U, R)}{\lambda_1 u_2(D, R) - u_2(U, L)} + \frac{u_2(D, R) - u_2(U, R)}{u_2(D, R) - u_2(U, L)} \right). \quad (6)$$

THEOREM 3. *Suppose that the component game of the revision game satisfies condition (4). If $t_1^* \neq t_2^*$, then there exists a unique revision equilibrium for all T . As $T \rightarrow \infty$, there exist two scenarios:*

(i) *If $t_1^* > t_2^*$, then the equilibrium payoffs converge to $u_i(U, L)$.*

(ii) *If $t_1^* < t_2^*$, then the equilibrium payoffs converge to $u_i(D, R)$.*

Notice that $t_1^* = t_2^*$ happens only for a knife-edge set of parameters. In this non-generic case, the revision game has multiple equilibria.²⁰

Theorem 3 states that for almost all parameter values, there is a unique revision equilibrium payoff and the outcome at the deadline will form one of the strict Nash equilibria in the component game with probability that converges to 1 as T increases. Which Nash equilibrium is prepared depends on a joint condition on the payoff function (u) and the ratio of arrival rates (λ_1/λ_2), as t_1^* and t_2^* depend on these parameters. In **Figure 3**, if $\lambda_1 = \lambda_2$, then $t_1^* > t_2^*$ in the left game and $t_1^* = t_2^*$ in the right game. Hence if $\lambda_1 = \lambda_2$, then (U, L) is the (limit) outcome in the left game, while the theorem does not cover the case in the right game. However, if $\lambda_1 < \lambda_2$, then the theorem implies that in the right game, the (limit) outcome is (U, L) . Similarly, if $\lambda_1 > \lambda_2$, then the (limit) outcome is (D, R) .

In the proof, we completely pin down the behavior at any time t in the unique revision equilibrium. In particular, players prepare the action that corresponds to the limit payoff profile for a sufficiently long time on the path of play, and this action profile is absorbing. This implies that if they were to choose actions simultaneously at $-T$ and if T were large enough, then they would choose these actions and never revise them on the path of play.

In **Section 6.1**, we provide an interpretation of this result. **Section 6.2** provides the proof, and **Section 6.3** fully describes the equilibrium dynamics, including off-path plays.

6.1 Interpretation of *Theorem 3*

The first step of the proof of **Theorem 3** shows that when t is close to zero, each player strictly prefers to prepare a best response in the component game to the last prepared action of her opponent. Hence, in the games of **Figure 3**, when the time is close to zero, players will move away from PAP (U, R) , to reach either (U, L) or (D, R) and then stay there until the deadline.²¹ If t is further from 0 and we assume that after t each player prepares a best response in the component game to her opponent's last prepared action, then player i 's expected continuation payoff from PAP (U, R) gets closer to a convex combination of $u_i(U, L)$ and $u_i(D, R)$ since the probability that no players revise their actions between t and 0 gets smaller. Hence, there is a finite time t^* such that, when the PAP is (U, R) , one player, whom we call the *strong player*, becomes indifferent at time t^* between (a) preparing a best response in the component game to her opponent's prepared action and (b) preparing the action necessary to form her preferred Nash equilibrium in the component game. Strictly before t^* , the strong player strictly prefers choice (b) in all PAPs. As the proof in the next subsection clarifies, $t^* = \min\{t_1^*, t_2^*\}$ is the time such that the strong player is indifferent between these two actions. In other words, $t_1^* > t_2^*$ means that player 1 can stick to non-Nash profiles longer than player 2 to induce player 2 to coordinate on her own preferred Nash equilibrium. This is why we call player 1 the strong player.

²⁰See **Kamada and Sugaya (2010a)** for a characterization of the set of revision equilibrium payoffs for the case $t_1^* = t_2^*$.

²¹Note that the incentive is strict at the deadline $t = 0$.

To see how this “strength” is affected by the parameters of the model, we consider two special cases. First, suppose that $\lambda_1 = \lambda_2$. In this case, $t_1^* > t_2^*$ is equivalent to

$$\frac{u_2(D, R) - u_2(U, R)}{u_2(D, R) - u_2(U, L)} > \frac{u_1(U, L) - u_1(U, R)}{u_1(U, L) - u_1(D, R)}.$$

The formula compares how strongly each player likes (U, R) relative to the two Nash equilibria. If player 1 likes it more, then she suffers less from miscoordination at (U, R) , so as a consequence, the inequality is more likely to be satisfied. If player 1 likes (D, R) less, then she expects less from moving away from (U, R) to (D, R) , and so the inequality is more likely to be satisfied if we decrease $u_1(D, R)$.²²

Second, consider the case with symmetric payoff functions $u_1(U, L) = u_2(D, R)$, $u_1(D, R) = u_2(U, L)$, and $u_1(U, R) = u_2(U, R)$. In this case, $t_1^* > t_2^*$ is equivalent to $\lambda_1 < \lambda_2$. This means that $\lambda_1 < \lambda_2$ implies (U, L) is the outcome of the revision game and that $\lambda_1 > \lambda_2$ implies (D, R) is the outcome of the revision game. More generally, $|t_1^*/t_2^*|$ is increasing in λ_1/λ_2 : if player 1’s relative frequency of the arrival of revision opportunities compared to player 2’s frequency decreases, then player 1’s commitment power becomes stronger, so (U, L) is more likely to be selected.

These results are reminiscent of findings in the bargaining literature. Player i ’s bargaining power increases in the disagreement payoff $u_i(U, R)$, decreases with the steepness of preference over the two “agreement outcomes” ($|u_1(U, L) - u_1(D, R)|$ for player 1 and $|u_2(D, R) - u_2(U, L)|$ for player 2), and increases in the ability to commit $1/\lambda_i$ to a proposal.

6.2 Proof of Theorem 3

In this subsection, we provide a proof of the convergence of the equilibrium payoff in Theorem 3. The proof consists of the following three steps. In the first step, we show that there is t^* finite such that after t^* , each player strictly prefers preparing her best response in the component game to her opponent’s action. Note that, starting from t^* , players do not change their actions as soon as the PAP forms a strict Nash equilibrium. In the second step, we consider the strategies before t^* . First we show that, before t^* , the strong player prefers preparing the action consistent with her preferred strict Nash equilibrium, irrespective of the prepared action of the weak player. Second, we show that when the strong player’s prepared action is the one consistent with her preferred Nash equilibrium, the weak player prefers to accommodate and prepares the action that will form such a strict Nash equilibrium. In the third step, we show that when T is sufficiently larger than $|t^*|$, there is enough time for the strong player to prepare the action

²²Note that a risk-dominated Nash equilibrium in the component game may be the (limit) outcome of the revision equilibrium. Consider the payoff matrix

	L	R
U	$2 + \epsilon, 1$	$0, 0$
D	$2\epsilon, 0$	$1, 2$

with $\epsilon > 0$. The action pair (U, L) is risk-dominated by (D, R) , while it is the (limit) outcome of the revision equilibrium when $\lambda_1 = \lambda_2$.

consistent with her preferred Nash equilibrium and for the weak player to accommodate so that the probability with which at t^* the PAP forms such a Nash equilibrium can be made arbitrarily close to 1 by increasing T . If this PAP is reached by t^* , players will stick to it until the end of the game.²³

Step 1. First, for each player i , we define t_i^* to be the infimum of times t such that given that each player prepares her best response in the component game to her opponent's action at any $t' > t$, i strictly prefers to prepare a best response in the component game to any other action. Since the incentive to take a static best response in the component game is strict at the deadline, this is true for t close enough to 0. By this definition and continuity of the expected payoffs (with respect to probabilities and so to time), player i must be indifferent between the two actions at t_i^* given that (i) the PAP at t_i^* is (U, R) and that (ii) each player prepares a best response in the component game to her opponent's action at time $t > t_i^*$. Then, from a straightforward calculation contained in [Appendix A.3](#), we show that for each $i = 1, 2$, t_i^* defined in this way coincides with t_i^* defined in (5) and (6).

Step 2. Suppose without loss of generality (w.l.o.g.) that $t_1^* > t_2^*$ and fix $t \in (-T, 0]$. Suppose that the following statements are true for any $t' > t$:

- (i) We have $t_1^* \leq t'$ or player 1 strictly prefers preparing U at t' whatever her opponent's current prepared action is.
- (ii) We have $t_1^* \leq t'$ or player 2 strictly prefers preparing L at t' when player 1's current prepared action is U .

These two statements are trivially true for t' close enough to 0. We show that there exists $\epsilon > 0$ such that these two statements are true also for all $t' \in (t - \epsilon, t]$, which proves that the statements are true for any t , by [Lemma 1](#).

Step 2.1. First, consider player 1's incentive when she obtains an opportunity at time $t < t_1^*$ (if $t > t_1^*$, the conclusion trivially holds; see [Appendix A.4](#) for the case of $t = t_1^*$). Suppose first that player 2 is currently preparing L or has a chance to revise strictly after time t but strictly before time t_1^* . If player 1 prepares action U , then statements (i) and (ii) and Step 1 imply that the action profile at the deadline is (U, L) , which gives player 1 the largest possible payoff that she can obtain in this revision game. On the other hand, if she prepares D , then there is a positive probability that she will obtain no other chances to revise. In such a case, the action profile at the deadline is not (U, L) . Hence, player 1 receives a payoff strictly less than her best possible payoff $u_1(U, L)$.

Suppose next that the current action of player 2 is R and he will not have any chance to revise strictly after time t but strictly before time t_1^* . In this case, player 1's expected payoff is the same as the continuation payoff when player 2's prepared action is R at time t_1^* .²⁴ Hence, player 1 must be indifferent between U and D at t_1^* by Step 1.

Overall, player 1 is strictly better off by preparing U at time t . Hence statement (i) is true at time t .

²³The intuition behind this proof idea is analogous to the one provided in [Kamada and Sugaya's \(2010a\)](#) "three-state example." We thank an anonymous referee for suggesting the way to extend it.

²⁴Note that the probability of player 2 getting a revision opportunity at t_1^* is zero.

Step 2.2. Now consider player 2's incentive when he obtains an opportunity at time $t < t_1^*$ (again, the case of $t > t_1^*$ is trivial; see [Appendix A.4](#) for the case of $t = t_1^*$). Suppose that player 1's current action is U (note that statement (ii) concerns only such a case). If player 2 prepares L , then statements (i) and (ii) and Step 1 imply that neither player changes her action in the future. Hence, the action profile at the deadline is (U, L) , which leads to the payoff $u_2(U, L)$. On the other hand, suppose that he prepares R . Player 2 prepares L if he obtains a revision opportunity strictly after time t but strictly before time t_1^* , which results in the payoff of $u_2(U, L)$. If he does not obtain any revision opportunity within that interval, then his expected payoff is the same as his continuation payoff given action profile (U, R) at time t_1^* . The latter is strictly less than $u_2(U, L)$, since, by the assumption that $t_2^* < t_1^*$, player 2 has a strict incentive to prepare L given that player 1 is preparing U at all $t > t_1^*$.

Overall, player 2 is strictly better off by preparing L when player 1 prepares U at time t . Hence, statement (ii) is true at time t .

Step 2.3. By continuity (of expected payoffs with respect to time), Steps 2.1 and 2.2 imply that there exists $\epsilon > 0$ such that for all $t' \in (t - \epsilon, t]$, both statements (i) and (ii) hold. Thus, by [Lemma 1](#), we have the desired result.

Step 3. Statement (i) in Step 2 shows that at any $t < t_1^*$, player 1 prepares U . Hence, for any $t' < t_1^*$, the probability that player 1's prepared action is U at t' converges to 1 as T increases. If player 1's prepared action is U at t' , then between t' and t_1^* , by statement (ii), player 2 must prepare L , and by statement (i), player 1 keeps preparing U . Hence, the probability that the PAP at t_1^* is (U, L) can be made arbitrarily close to 1 by setting T large enough. Since the probability of revision at time t_1^* is zero, Step 1 implies that if the PAP at t_1^* is (U, L) , then the players keep preparing (U, L) until the deadline.

6.3 Equilibrium dynamics

The proof in the previous subsection characterizes the strong player's equilibrium strategy fully but the weak player's equilibrium strategy only after the strong player prepares the action that corresponds to the strong player's preferred Nash equilibrium.²⁵ Here we provide a full characterization of the equilibrium dynamics, which implies that the equilibrium strategy is unique. The proof of the result stated in this subsection is provided in [Calcagno and Lovo \(2010\)](#) and [Kamada and Sugaya \(2010a\)](#).

The equilibrium dynamics are summarized in [Figure 5](#) for the case $t_1^* > t_2^*$. The dynamics consist of three phases. In each phase, the arrow that originates from an action pair x represents what players will prepare if they are given an opportunity to revise during that phase when the PAP is x . More specifically, an arrow from (x_i, x_{-i}) to (x'_i, x_{-i}) means that if player i is given an opportunity to revise when the PAP is $x \in \{(x_i, x_{-i}), (x'_i, x_{-i})\}$, then player i would prepare x'_i . If a player does not switch her action, then there is no arrow that corresponds to that strategy. Hence, in particular, if

²⁵If the players choose their actions simultaneously at $-T$, then it is common knowledge that the strong player prepares the action that corresponds to the strong player's preferred Nash equilibrium at $-T$. Hence, the proof is sufficient to fully characterize the path of play in the revision equilibrium.

is in the expression for t^* , the derivation of which is left as an exercise for the interested reader.

Second, consider the effect of different degrees of asynchronicity. For this purpose, in addition to the two independent processes specified in [Section 3](#), consider another independent Poisson process \mathbf{p}_{12} with arrival rate $\lambda_{12} > 0$, at which both players revise simultaneously. For simplicity, we assume the Poisson process is homogeneous. At the time of decision that corresponds to each revision opportunity, player i does not know whether such an opportunity is driven by the process \mathbf{p}_i or by \mathbf{p}_{12} . If $\lambda_1 = \lambda_2 = 0$ and $\lambda_{12} > 0$, then all revision opportunities are synchronous and it is straightforward to show that any repetition of a Nash equilibrium is an equilibrium of the revision game. The following result shows that a small degree of asynchronicity is not enough to rule out multiple equilibria when the component game has multiple strict Nash equilibria.

THEOREM 4. *Fix a component game $((X_i)_{i=1,2}, (u_i)_{i=1,2})$. There exists $\epsilon > 0$ such that if the arrival rates satisfy $\lambda_1, \lambda_2 \in (0, \epsilon)$ and $\lambda_{12} > 1/\epsilon$, then for every strict Nash equilibrium of the component game $x^N \in X$ and every horizon length T , there exists a revision equilibrium in which each player i prepares action x_i^N at all revision opportunities at $t \leq 0$.*

The proof is given in [Appendix A.6](#) and a detailed discussion can be found in [Calcagno and Lovo \(2010\)](#). Note that ϵ in the theorem is required to be strictly positive. This means that a small degree of asynchronicity is not enough to eliminate multiple equilibria. This raises the question of how much asynchronicity is needed to obtain equilibrium uniqueness in a revision equilibrium. [Ishii and Kamada \(2011\)](#) characterize the parameter regions such that multiplicity persists in common interest games.

APPENDIX

A.1 Proof of [Lemma 1](#)

PROOF. Suppose that the premise of the lemma holds. Let t^* be the supremum of t such that A_t is false (if the supremum does not exist (i.e., it is negative infinity), we are done). Then it must be the case that for any $\epsilon > 0$, there exists $t' \in (t^* - \epsilon, t^*]$ such that $A_{t'}$ is false. But by the definition of t^* , there exists $\epsilon' > 0$ such that statement $A_{t'}$ is true for all $t' \in (t^* - \epsilon', t^*]$ because the premise of the lemma is true. This contradiction proves the result. \square

A.2 A sharper result for the case when the players move at $-T$

PROPOSITION 1. *Suppose that the players choose their actions at $-T$ and consider a component game of a revision game with a strictly Pareto-dominant action profile x^* . Then there exists T' such that for all $T > T'$, in all SPE, $x(0) = x^*$ with probability 1.*

PROOF. Suppose without loss of generality that $\lambda_1 \leq \lambda_2$. Consider first the case of $\lambda_1 < \lambda_2$. Fix an SPE strategy profile where player 1 prepares $x_1 \neq x_1^*$ at $-T < 0$. In this

case, player 1's expected payoff at time $-T$ is at most

$$u_1(x^*) - e^{-\lambda_1 T} m,$$

as with probability $e^{-\lambda_1 T}$, player 1 has no further revision opportunities.²⁶ On the other hand, one possible deviation is to prepare x_1^* for all $[-T, 0]$, and in that case her expected payoff is

$$u_1(x^*) - e^{-\lambda_2 T} M,$$

since by Step 1 in the proof of [Theorem 1](#), it follows that player 2 will switch to x_2^* as soon as he has a chance to revise, and afterward the PAP never changes.^{27,28} However, the assumption that $\lambda_1 < \lambda_2$ implies that for sufficiently large T , the latter value becomes strictly greater than the former, implying that in any SPE, player 1 must prepare x_1^* when T is sufficiently large. Given this, player 2 has a strict incentive to prepare x_2^* at $-T$ as that would give him the highest possible expected payoff in equilibrium, while preparing some other action results in a strictly lower payoff because there is a strictly positive probability that he has no chance to revise the action in the future.

Suppose $\lambda_1 = \lambda_2 \equiv \lambda$. Fix a revision equilibrium and let $V_i^t(x)$ be player i 's value from the revision equilibrium when the PAP is x at time t . Let $v_1^t(x_2) = \max_{x_1 \neq x_1^*} V_1^t(x_1, x_2)$ be player 1's maximum value at t when player 2 prepares x_2 conditional on player 1 not preparing x_1^* . It suffices to show that, for any $x_2 \neq x_2^*$, there exists \bar{t} such that $V_1^{\bar{t}}(x_1^*, x_2) > v_1^{\bar{t}}(x_2)$, that is, player 1 prefers preparing x_1^* given player 2 preparing x_2 at time \bar{t} .

Let us explain why finding one \bar{t} for each x_2 is sufficient (and \bar{t} can be different for different x_2 , but must be independent of a particular equilibrium we fix). Suppose player 2 prepares x_2 and player 1 receives a revision opportunity at time $t < \bar{t}$. If player 2 has a revision opportunity by \bar{t} , then preparing x_1^* at time t (and continuing to prepare x_1^* until \bar{t}) gives player 1 the highest payoff $u_1(x^*)$ by Step 1 of [Theorem 1](#). On the other hand, if player 2 does not have a revision opportunity by \bar{t} , then player 2 will prepare x_2 at time \bar{t} . Thus, in this case, preparing x_1^* at time t (and continuing to prepare x_1^* until \bar{t}) is strictly better than any other strategies since we have $V_1^{\bar{t}}(x_1^*, x_2) > v_1^{\bar{t}}(x_2)$. In total, since the probability that player 1 cannot move between t and \bar{t} is strictly positive, preparing x_1^* at time t is strictly better than any other strategies.

Given the above discussion, it suffices to derive a contradiction by assuming that there exists $x_2 \neq x_2^*$ with

$$V_1^t(x_1^*, x_2) \leq v_1^t(x_2) \quad \text{for all } t \leq 0. \quad (7)$$

Arbitrarily fix $x_2 \neq x_2^*$ such that (7) holds. Take any $S < 0$. By induction, we now show that for all integer $n \geq 0$,

$$v_1^{nS}(x_2) \leq u_1(x^*) - (n + 1 - ne^{\lambda S})me^{n\lambda S}. \quad (8)$$

²⁶Recall from the proof of [Theorem 1](#) that $m := \min_{i, x \neq x^*} (u_i(x^*) - u_i(x))$.

²⁷Recall again from the proof of [Theorem 1](#) that $M := \max_{i, x \neq x^*} (u_i(x^*) - u_i(x)) < \infty$.

²⁸Here we use the fact that there are only two players. If there are two or more opponents, Step 1 cannot be used to conclude that *all* the opponents will switch to actions prescribed by x^* .

First, with $n = 0$, (8) is equivalent to

$$v_1^0(x_2) \leq u_1(x^*) - m.$$

This inequality holds by the definitions of v_1^0 and m , so (8) is true for $n = 0$.

Second, suppose that (8) holds for $n = k$. We show that (8) holds also for $n = k + 1$. To see this, note that player 1's value at time kS when player 2 prepares x_2 is bounded from above by $v_1^{kS}(x_2)$ from (7). Therefore, $v_1^{(k+1)S}(x_2)$ can be bounded by

$$\begin{aligned} u_1(x^*) - \underbrace{e^{\lambda S}}_{\text{Pr of 2 not moving by } kS} &\times \underbrace{(u_1^*(x_2) - v_1^{kS}(x_2))}_{\substack{\text{from (7),} \\ \text{this is the least loss compared to } u_1(x^*)}} \\ &- \underbrace{(1 - e^{\lambda S})}_{\text{Pr of 2 moving by } kS} \times \underbrace{e^{\lambda S}}_{\text{Pr of 1 not moving by } kS} \times \underbrace{me^{k\lambda S}}_{\text{the least loss for } x \neq x^*} \\ &= u_1(x^*) - e^{\lambda S} \times \underbrace{(k + 1 - ke^{\lambda S})me^{k\lambda S}}_{\text{by (8) with } n=k} - (1 - e^{\lambda S})me^{(k+1)\lambda S} \\ &= u_1(x^*) - (k + 2 - (k + 1)e^{\lambda S})me^{(k+1)\lambda S}, \end{aligned}$$

which is (8) with $n = k + 1$.

Therefore, (8) holds for all integers $n \geq 0$. On the other hand, the lower bound of player 1's continuation payoff at t is

$$u_1(x^*) - e^{\lambda t} M,$$

since this is the expected payoff she gets when she sticks to x_1^* after time t until the deadline. This and (8) (and $e^{\lambda S} < 1$) imply that when n is sufficiently large, we have that

$$v_1^{nS}(x_2) \leq u_1(x^*) - (n + 1 - ne^{\lambda S})me^{n\lambda S} < u_1(x^*) - e^{\lambda nS} M \leq V_i^{nS}(x_1^*, x_2).$$

Since this is the desired contradiction, the proof is complete. \square

A.3 Proof of Theorem 2

Let $\underline{v}_i^t(k)$ be the infimum of player i 's payoff at t in subgame perfect equilibrium strategies and histories such that there are at least k players who prepare the action corresponding to x^* and no player receives a revision opportunity at t . By mathematical induction with respect to $k = n, \dots, 0$, we show that $\lim_{t \rightarrow -\infty} \underline{v}_i^t(k) = u_i(x^*)$ for all $i \in I$.

The proof for $k = n$. Step 1 of the proof of Theorem 1 is valid with an arbitrary number of players if we replace player $-i$ with the set J of players other than player i and replace λ_{-i} with $\sum_{j \in J} \lambda_j$. Hence, x^* is the absorbing state with n players. Since x^* is absorbing, $\underline{v}_i^t(n) = u_i(x^*)$ for all i and t , as desired.

Inductive argument. Suppose $\lim_{t \rightarrow -\infty} \underline{v}_i^t(k + 1) = u_i(x^*)$ for all $i \in I$ with $k + 1 \leq n$. Given this inductive hypothesis, we show that $\lim_{t \rightarrow -\infty} \underline{v}_i^t(k) = u_i(x^*)$ for all $i \in I$. For simple notation, let $\Lambda := \sum_{i \in I} \lambda_i$ be the summation of the arrival rates, let $\alpha_1 := \min_{i \in I} r_i$

be the smallest r_i , and let $\beta := \min_{j \in I, j \neq j^*} r_j$ with $j^* \in \arg \min_{j \in I} r_j$ (take one arbitrarily if there are multiples) be the second smallest r_i .²⁹ In addition, let

$$\bar{K} = \max_{i, j \in I, x \in X, x \neq x^*} \frac{u_i(x^*) - u_i(x)}{u_j(x^*) - u_j(x)} < \infty$$

be the maximum ratio of the range of utilities between the players. The denominator of the maximand is always strictly positive since we assume $|X_j| \geq 2$ for each j and x^* is strictly Pareto-dominant.

Take $\varepsilon > 0$ arbitrarily. Since $\lim_{t \rightarrow -\infty} v_i^t(k+1) = u_i(x^*)$, there exists T_0 such that for all $t \leq T_0$ and $i \in I$, $v_i^t(k+1) \geq u_i(x^*) - \varepsilon$. Consider the situation where k players prepare actions corresponding to x^* at $t = T_0 + \tau_1$ with $\tau_1 \leq 0$, that is, $t \leq T_0$. Then, if player j who is not preparing x_j^* at time t can move first by T_0 , then she yields at least $u_j(x^*) - \varepsilon$ by preparing x_j^* . This implies each player i will at least yield

$$u_i(x^*) - \bar{K}\varepsilon. \quad (9)$$

Therefore,

$$v_i^t(k) \geq \alpha_1(1 - e^{\tau_1 \Lambda})(u_i(x^*) - \bar{K}\varepsilon) + (1 - \alpha_1(1 - e^{\tau_1 \Lambda}))\underline{u}_i \quad (10)$$

for all $i \in I$. Note that $\alpha_1(1 - e^{\tau_1 \Lambda})$ is the minimum probability that player j , who is not preparing x_j^* at t , can move first by T_0 and that we assume that if such player j does not move, the worst payoff \underline{u}_j realizes.

Taking τ_1 sufficiently large (in absolute value) in (10), there exists T_1 such that for all $\tau_2 \leq 0$,

$$v_i^t(k) \geq \alpha_1 u_i(x^*) + (1 - \alpha_1)\underline{u}_i - \bar{K}\varepsilon \quad (11)$$

for all $i \in I$, where $t = T_0 + T_1 + \tau_2 \leq T_0 + T_1$.

Consider $v_i^t(k)$ with $t = T_0 + T_1 + \tau_2$. Then we can compute lower bounds of player i 's payoff in different cases as follows.

- If player j , who is not preparing x_j^* at t , will move first by $T_0 + T_1$, then a lower bound is $u_i(x^*) - \bar{K}\varepsilon$ by the same argument as in (9).
- If player i herself will move first by $T_0 + T_1$, then a lower bound is $\alpha_1 u_i(x^*) + (1 - \alpha_1)\underline{u}_i - \bar{K}\varepsilon$ since there are two possibilities:
 - If player i is preparing x_i^* at t , then by staying at x_i^* , player i keeps the situation that there are k players preparing the actions corresponding to x^* . In this case, by (11), player i 's payoff is bounded by $\alpha_1 u_i(x^*) + (1 - \alpha_1)\underline{u}_i - \bar{K}\varepsilon$ for τ_2 sufficiently large in absolute value.
 - If player i is not preparing x_i^* at t , then by preparing at x_i^* , player i creates the situation that there are $k+1$ players preparing the actions corresponding to x^* . In such a case, the inductive hypothesis guarantees that $v_i^t(k)$ is at least $u_i(x^*) - \varepsilon \geq \alpha_1 u_i(x^*) + (1 - \alpha_1)\underline{u}_i - \bar{K}\varepsilon$.

²⁹Note that $\alpha_1 = \alpha$, defined right before the statement of Theorem 2.

- If player j , who is preparing x_j^* at t , will move first by $T_0 + T_1$, then as in the first subcase of the second case, player j can guarantee herself $\alpha_1 u_j(x^*) + (1 - \alpha_1) \underline{u}_j - \bar{K} \varepsilon$. By the definition of K , when player j gets at least $\alpha_1 u_j(x^*) + (1 - \alpha_1) \underline{u}_j - \bar{K} \varepsilon$, player i 's payoff u_i should satisfy

$$\begin{aligned} \frac{u_i(x^*) - \underline{u}_i}{u_i(x^*) - \underline{u}_i} &\leq K \frac{u_j(x^*) - (\alpha_1 u_j(x^*) + (1 - \alpha_1) \underline{u}_j - \bar{K} \varepsilon)}{u_j(x^*) - \underline{u}_j} \\ &= K(1 - \alpha_1) + \frac{K \bar{K} \varepsilon}{u_j(x^*) - \underline{u}_j}, \end{aligned}$$

that is,

$$\begin{aligned} u_i &\geq u_i(x^*) - K(1 - \alpha_1)(u_i(x^*) - \underline{u}_i) - K \bar{K} \frac{u_i(x^*) - \underline{u}_i}{u_j(x^*) - \underline{u}_j} \varepsilon \\ &\geq (1 - K(1 - \alpha_1))u_i(x^*) + K(1 - \alpha_1)\underline{u}_i - \bar{K}^3 \varepsilon. \end{aligned}$$

In total, player i 's value satisfies

$$\begin{aligned} \underline{v}_i^t(k) &\geq \alpha_1(1 - e^{\tau_2 \Lambda})(u_i(x^*) - \bar{K} \varepsilon) \\ &\quad + (\beta(1 - e^{\tau_2 \Lambda}) + e^{\tau_2 \Lambda})(\alpha_1 u_i(x^*) + (1 - \alpha_1) \underline{u}_i - \bar{K} \varepsilon) \\ &\quad + (1 - \alpha_1 - \beta)(1 - e^{\tau_2 \Lambda})((1 - K(1 - \alpha_1))u_i(x^*) + K(1 - \alpha_1)\underline{u}_i - \bar{K}^3 \varepsilon). \end{aligned}$$

Taking τ_2 sufficiently large, there exists T_2 such that at $t = T_0 + T_1 + T_2 + \tau_3$ with $\tau_3 \leq 0$,

$$\begin{aligned} \underline{v}_i^t(k) &\geq (\alpha_1 + \beta \alpha_1 + (1 - (\alpha_1 + \beta))(1 - K(1 - \alpha_1)))u_i(x^*) \\ &\quad + (1 - (\alpha_1 + \beta \alpha_1 + (1 - (\alpha_1 + \beta))(1 - K(1 - \alpha_1))))\underline{u}_i - \bar{K}^3 \varepsilon. \end{aligned}$$

Defining

$$\alpha_2 := \alpha_1 + \beta \alpha_1 + (1 - (\alpha_1 + \beta))(1 - K(1 - \alpha_1)),$$

we have

$$\underline{v}_i^t(k) \geq \alpha_2 u_i(x^*) + (1 - \alpha_2) \underline{u}_i - \bar{K}^3 \varepsilon.$$

Recursively, for each $M = 1, 2, \dots$, there exist T_0, T_1, \dots, T_M such that at $t \leq T_0 + T_1 + \dots + T_M$,

$$\underline{v}_i^t(k) \geq \alpha_M u_i(x^*) + (1 - \alpha_M) \underline{u}_i - \bar{K}^{2M-1} \varepsilon$$

with

$$\alpha_M = \alpha_1 + \beta \alpha_{M-1} + (1 - (\alpha_1 + \beta))(1 - K(1 - \alpha_{M-1}))$$

or

$$(\alpha_M - 1) = (\beta + (1 - (\alpha_1 + \beta))K)(\alpha_{M-1} - 1).$$

By condition (3), α_M is monotonically increasing and converges to 1. Taking M sufficiently large and $\varepsilon > 0$ sufficiently small yields the result.

A.4 Proof of *Theorem 3*: Derivation of t_i^*

We provide a derivation of t_1^* . The value of t_2^* can be found in a symmetric manner. By the definition of t_1^* , assuming that both players prepare best responses to the PAP at any time strictly after t_1^* , the payoff from playing a best response against R at t_1^* and playing otherwise must be equal. Thus, it must be the case that

$$u_1(D, R) = \underbrace{e^{(\lambda_1 + \lambda_2)t_1^*} u_1(U, R)}_{\text{nobody moves until 0}} + \underbrace{\frac{\lambda_1}{\lambda_1 + \lambda_2} (1 - e^{(\lambda_1 + \lambda_2)t_1^*}) u_1(D, R)}_{\text{player 1 moves first}} + \underbrace{\frac{\lambda_2}{\lambda_1 + \lambda_2} (1 - e^{(\lambda_1 + \lambda_2)t_1^*}) u_1(U, L)}_{\text{player 2 moves first}}.$$

Solving this equation for t_1^* , we obtain the desired expression.

A.5 Checking the induction argument for the case $t = t_1^*$

In Step 2 of the proof of *Theorem 3*, we referred to an appendix. Here we prove that when $t = t_1^*$, there exists $\epsilon > 0$ such that the statements (i) and (ii) hold for all $t' \in (t - \epsilon, t]$, assuming that the statements hold for all time strictly after t .

First, the existence of such ϵ for statement (ii) is given by continuity: Given that both players prepare static best responses strictly after t_1^* , we know that player 2 has a strict incentive to prepare a static best response at t_1^* . By continuity, there exists some $\epsilon > 0$ such that for all time in $(t_1^* - \epsilon, t_1^*]$, player 2 has a strict incentive to prepare a static best response against player 1's currently prepared action, no matter what we assume about player 1's strategy in the time interval $(t_1^* - \epsilon, t_1^*]$.

Second, consider statement (i). Take the ϵ that we took in the previous paragraph. Suppose first that player 2 does not receive any revision opportunity in the time interval (t_1^*) . In this case, player 1's expected payoff is $u_1(D, R)$ irrespective of her preparation at t' and at any time in (t_1^*) by the definition of t_1^* (i.e., player 1 is indifferent between two actions at t_1^*). The remaining case is when player 2 receives at least one revision opportunity in the time interval (t_1^*) . We divide this case in two subcases: (a) If player 1 does not receive any opportunity in that time interval, then her payoff is $u_1(U, L)$ if 1 prepares U at t' (as we have shown that player 2 prepares a static best reply at t') and at most $\max_{a_2} u_1(D, a_2) e^{(\lambda_1 + \lambda_2)t_1^*} + u_1(U, L)(1 - e^{(\lambda_1 + \lambda_2)t_1^*})$ if she prepares D , which is strictly smaller than $u_1(U, L)$. (b) If player 1 receives an opportunity, then preparing D is better than U by at most $u_1(U, L) - \min_a u_1(a) < \infty$. Since the ratio of the probability of subcase (b) to that of subcase (a) approaches zero as t_1^* , this shows the existence of $\epsilon > 0$ such that statement (i) holds for all time $(t_1^* - \epsilon, t_1^*]$ assuming that statements (i) and (ii) hold for all time strictly after t_1^* .

A.6 Proof of *Theorem 4*

Let $P_J(t)$ denote the probability that from t on, players in set $J \subseteq \{1, 2\}$ have some revision opportunities in $(t, 0]$ while players in $\{1, 2\} \setminus J$ have none. Note that given that

after t , players' continuation strategies consist in preparing x^N , player i 's expected continuation payoff from PAP $x(t) = x = (x_i, x_{-i})$ is

$$\pi(x, t) := P_{\emptyset}(t)u_i(x) + P_{\{1,2\}}(t)u_i(x^N) + P_{\{i\}}(t)u_i(x_i^N, x_{-i}) + P_{\{-i\}}(t)u_i(x_i, x_{-i}^N). \quad (12)$$

Now suppose player i has a revision opportunity at t . With probability $\lambda_{12}/(\lambda_{12} + \lambda_i)$, player $-i$ is simultaneously revising (and will prepare x_{-i}^N), and with probability $\lambda_i/(\lambda_{12} + \lambda_i)$, she is not. Thus, if at t , the PAP is $x(t) = (x_i', x_{-i})$, player i 's expected continuation payoff from preparing x_i^N at t is

$$\frac{\lambda_{12}}{\lambda_{12} + \lambda_i} \pi(x^N, t) + \frac{\lambda_i}{\lambda_{12} + \lambda_i} \pi((x_i^N, x_{-i}), t). \quad (13)$$

If instead she prepares $x_i \neq x_i^N$, her continuation payoff is

$$\frac{\lambda_{12}}{\lambda_{12} + \lambda_i} \pi((x_i, x_{-i}^N), t) + \frac{\lambda_i}{\lambda_{12} + \lambda_i} \pi((x_i, x_{-i}), t). \quad (14)$$

Using expression (12), we have that (13) is strictly larger than (14) if and only if

$$\begin{aligned} & P_{\emptyset}(t) \left(\frac{\lambda_{12}}{\lambda_{12} + \lambda_i} (u_i(x^N) - u_i(x_i, x_{-i}^N)) + \frac{\lambda_i}{\lambda_{12} + \lambda_i} (u_i(x_i^N, x_{-i}) - u_i(x_i, x_{-i})) \right) \\ & + P_{\{-i\}}(t) \left(\frac{\lambda_{12}}{\lambda_{12} + \lambda_i} (u_i(x^N) - u_i(x_i, x_{-i}^N)) + \frac{\lambda_i}{\lambda_{12} + \lambda_i} (u_i(x^N) - u_i(x_i, x_{-i}^N)) \right) > 0. \end{aligned}$$

Because x^N is a strict Nash equilibrium, $(u_i(x^N) - u_i(x_i, x_{-i}^N)) > 0$. Hence, because $P_{\emptyset}(t) > 0$ and $P_{\{-i\}}(t) \geq 0$, if $\lambda_{12}/(\lambda_{12} + \lambda_i) < 1$ is close enough to 1, then the previous inequality is satisfied for all t .

REFERENCES

- Abreu, Dilip and Faruk Gul (2000), "Bargaining and reputation." *Econometrica*, 68, 85–117. [413]
- Abreu, Dilip and David Pearce (2007), "Bargaining, reputation and equilibrium selection in repeated games with contracts." *Econometrica*, 75, 653–710. [413]
- Ambrus, Attila and Shih-En Lu (2009), "A continuous-time model of multilateral bargaining." Unpublished paper. [413]
- Aumann, Robert J. and Sergiu Hart (2003), "Long cheap talk." *Econometrica*, 71, 1619–1660. [411]
- Biais, Bruno, Christophe Bisière, and Sebastien Pouget (2014), "Equilibrium discovery and preopening mechanisms in an experimental market." *Management Science*, 60, 753–769. [411]
- Calcagno, Riccardo and Stefano Lovo (2010), "Preopening and equilibrium selection." Unpublished paper. [411, 413, 414, 424, 425, 426]

Carlsson, Hans and Eric van Damme (1993), "Global games and equilibrium selection." *Econometrica*, 61, 989–1018. [412]

Caruana, Guillermo and Liran Einav (2008), "A theory of endogenous commitment." *Review of Economic Studies*, 75, 99–116. [413]

Chao, Yuen R. (1919), "A note on 'Continuous mathematical induction.'" *Bulletin of the American Mathematical Society*, 26, 17–18. [415]

Dutta, Prajit K. (1995), "A folk theorem for stochastic games." *Journal of Economic Theory*, 66, 1–32. [412]

Dutta, Prajit K. (2012), "Coordination need not be a problem." *Games and Economic Behavior*, 76, 519–534. [412]

Farrell, Joseph (1987), "Cheap talk, coordination, and entry." *RAND Journal of Economics*, 18, 34–39. [411]

Farrell, Joseph and Garth Saloner (1985), "Standardization, compatibility, and innovation." *RAND Journal of Economics*, 16, 70–83. [412]

Harsanyi, John C. and Reinhard Selten (1988), *A General Theory of Equilibrium Selection in Games*. MIT Press, Cambridge, Massachusetts. [411]

Ishii, Yuhta and Yuichiro Kamada (2011), "Correlated inertia and coordination." Unpublished paper. [414, 426]

Kamada, Yuichiro and Michihiro Kandori (2010), "Revision games." Unpublished paper. [410, 413, 414, 415, 417]

Kamada, Yuichiro and Takuo Sugaya (2010a), "Asynchronous revision games with deadline: Unique equilibrium in coordination games." Unpublished paper. [413, 418, 421, 423, 424, 425]

Kamada, Yuichiro and Takuo Sugaya (2010b), "Policy announcement game: Valence candidates and ambiguous policies." Unpublished paper. [413]

Kandori, Michihiro, George J. Mailath, and Rafael Rob (1993), "Learning, mutation, and long run equilibria in games." *Econometrica*, 61, 29–56. [412]

Lagunoff, Roger and Akihiko Matsui (1997), "Asynchronous choice in repeated coordination games." *Econometrica*, 65, 1467–1477. [412, 417]

Lagunoff, Roger and Akihiko Matsui (2001), "Are 'anti-folk theorems' in repeated games nongeneric?" *Review of Economic Design*, 6, 397–412. [412]

Lipman, Barton L. and Ruqu Wang (2000), "Switching costs in frequently repeated games." *Journal of Economic Theory*, 93, 149–190. [413]

Matsui, Akihiko and Kiminori Matsuyama (1995), "An approach to equilibrium selection." *Journal of Economic Theory*, 65, 415–434. [412]

Morris, Stephen, Rafael Rob, and Hyun Song Shin (1995), "p-dominance and belief potential." *Econometrica*, 63, 145–167. [412]

Morris, Stephen and Hyun Song Shin (1998), "Unique equilibrium in a model of self-fulfilling currency attacks." *American Economic Review*, 88, 587–597. [412]

Oyama, Daisuke, Satoru Takahashi, and Josef Hofbauer (2008), "Monotone methods for equilibrium selection under perfect foresight dynamics." *Theoretical Economics*, 3, 155–192. [412]

Rabin, Matthew (1994), "A model of pre-game communication." *Journal of Economic Theory*, 63, 370–391. [411]

Rubinstein, Ariel (1989), "The electronic mail game: A game with almost common knowledge." *American Economic Review*, 79, 385–391. [412]

Rubinstein, Ariel and Asher Wolinsky (1995), "Remarks on infinitely repeated extensive-form games." *Games and Economic Behavior*, 9, 110–115. [412]

Sugaya, Takuo and Satoru Takahashi (2009), "Coordination failure in repeated games with private monitoring." Unpublished paper. [412]

Takahashi, Satoru (2005), "Infinite horizon common interest games with perfect information." *Games and Economic Behavior*, 53, 231–247. [412]

van Damme, Eric and Sjaak Hurkens (1996), "Commitment robust equilibria and endogenous timing." *Games and Economic Behavior*, 15, 290–311. [413]

Wen, Quan (2002), "A folk theorem for repeated sequential games." *Review of Economic Studies*, 69, 493–512. [412]

Yoon, Kiho (2001), "A folk theorem for asynchronous repeated games." *Econometrica*, 69, 191–200. [412, 418]

Young, H. Peyton (1993), "The evolution of conventions." *Econometrica*, 61, 57–84. [412]

Young, H. Peyton (1998), "Conventional contracts." *Review of Economic Studies*, 65, 773–792. [412]