

Games with Private Timing*

Yuichiro Kamada[†] Sofia Moroni[‡]

First version: May 8, 2014

This version: July 3, 2018

Abstract

We study a class of games in which the timing of players' moves is private information, but players have the option to disclose their moves by exerting a small cost. When the underlying game is a coordination game, we characterize the set of distributions of moving times such that the game has the following unique prediction: Players choose the best coordination equilibrium and do not disclose their action. This implies that the possibility of disclosure selects an equilibrium in which the best action profile is taken but nothing is disclosed. In games of opposing interests, we provide sufficient conditions for the first-arriving player to disclose her action. In extensions we allow for, among others, partial control over timing.

Keywords: common interest games, opposing interest games, asynchronous moves, private timing, dynamic games

*We thank David Ahn, Simon Board, Amanda Friedenberg, Drew Fudenberg, Johannes Hörner, Michihiro Kandori, Navin Kartik, Fuhito Kojima, Scott Kominers, George Mailath, Paul Milgrom, Nozomu Muto, Michael Ostrovsky, Daisuke Oyama, Larry Samuelson, Pablo Schenone, Philipp Strack, Balázs Szentes, Olivier Tercieux, Jean Tirole, Juuso Toikka, Takuro Yamashita, and seminar/conference participants at Arizona State University, Hitotsubashi University, Simon Fraser University, Stanford University, University of California, Berkeley, University of Chicago, University of Tokyo, Yale University, Microtheory and Market Design 2013 at Yonsei University for helpful comments. We thank Lingxuan Wu for his excellent research assistantship.

[†]Haas School of Business, University of California Berkeley, 545 Student Services Building, #1900 Berkeley, CA 94720-1900, e-mail: y.cam.24@gmail.com

[‡]Department of Economics, University of Pittsburgh, 4528 Wesley W. Posvar Hall, 230 South Bouquet St., Pittsburgh PA 15260, email: moroni@pitt.edu

1 Introduction

There are numerous social and economic situations in which knowledge about timing matters. A firm may want to conduct a costly investigation of the pricing strategy of its competing firm only if the rival already had an internal meeting to determine such a strategy. A salesperson at an electric appliance store may change her sales pitch depending on whether the customer has already visited a competing store. Investors may want to condition their decisions for a start-up company on whether other investors had enough time to make their investment decisions for the company. In all these situations, choice of actions depends on what one believes about the timing of the choices by other actors.

This paper considers a new class of games that we call *games with private timing* to analyze such situations. In games with private timing, there is an underlying normal-form game which we call the component game. Each of two players privately learns the time at which she is able to choose an action in a component game once and for all, but does not know the time at which her opponent takes his action. If there is no information disclosure between periods, then these games are strategically equivalent to simultaneous-move games. However, many questions arise once we introduce information revelation between periods: Do players want to disclose their own actions to the opponent? Do they want to commit to monitor the opponent's actions? What if actions are disclosed with some exogenously given probability? How do these possibilities affect the component-game actions that the players choose?

In this paper, we focus on one particular information revelation mechanism that highlights the non-triviality of these problems. We consider a setting in which each player has a costly option to disclose her own action to the opponent, who may use the information if he has not yet moved. We show that certain conditions on the timing distribution (about which we will elaborate shortly) imply uniqueness of a perfect Bayesian equilibrium (PBE) when the component game is either a common interest game (e.g., a coordination game) or an opposing interest game (e.g., a battle of the sexes).

In the common-interest setting, we characterize the set of distributions of moving times such that every dynamic game with a common-interest component game

has the following unique PBE for small enough disclosure costs: players choose the best coordination equilibrium of the component game and do not disclose their action. This implies an unexpected consequence: For those distributions in the identified set, the introduction of the disclosure option helps select a PBE that does not involve disclosure. The set of the distributions we identify is the ones that satisfy *asynchronicity* and *uncertainty*. Asynchronicity means that players do not move simultaneously. Uncertainty means that each player cannot be sure that she is the last to move. In short, asynchronicity implies less strategic uncertainty about the opponent’s choice than synchronicity when the players have the choice to disclose their actions, and uncertainty about the timing distribution implies less freedom on allowable beliefs at information sets. Thus, there are fewer issues of multiplicity caused by the freedom of cooking up beliefs at off-path information sets. In particular, we rule out equilibrium candidates in which Pareto dominated actions are taken.

When the component game is an opposing-interest game, we give sufficient conditions on the distribution of the timing of moves such that there exists a unique PBE. In this PBE, the player who happens to become the first player pays the disclosure cost. The sufficient conditions essentially state that moving times are sufficiently “dispersed” and well behaved.

To put our results into specific context, consider two firms’ decisions about a product’s concept and design. Those decisions are typically accompanied by substantial investments, such as the setup and launch of production at scale of the chosen design. These large investments make the product decisions de facto fixed once they are made. Whether the firms should make the decision secret or public is a crucial strategic choice, as it may change the opponent firm’s decision. Moreover, the timing of a firm’s decision is affected by various factors such as other commitments that the firm has and the time necessary to develop full-fledged set of prototypes to pick from. Delaying decisions may not be an option if the firm needs to secure future funding or needs to have cash in hand. Hence the firm is not in complete control of the time at which she is able to decide on her product design and this time is ex ante unknown. Furthermore, these decision times are likely not common across the two firms. First, consider a market where each firm has its own strength, and hence would like to target, through its design, one specific

segment which is different from the other firm's preferred segment. It may be profitable, for example, for different firms to target different segments in order to avoid direct competition. The game between the two firms would be expressed as a common interest game. Our result then predicts that if the firms can talk to the opponent firm at a small cost,¹ they will each target the segment they prefer without resorting to such communications.² Next, consider a market where firms benefit from compatibility even though each of them has its relative strength. As a result, they would like to target the same segment, but they each prefer a different common target. The game would be expressed as an opposing interest game. Our result implies that the first firm to have a chance to make a move chooses its preferred segment, and communicates to the opponent firm.³

The present paper is part of the literature that tries to understand the relationship between timing and economic behavior. As discussed, asynchronicity and uncertainty are the keys to our results. The role of asynchronicity in equilibrium selection is present in the literature. Lagunoff and Matsui (1997) consider asynchronous repeated games and show a uniqueness result for games in which the two players have the same payoff function.⁴ Caruana and Einav (2008) consider a finite-horizon model with switching costs and show that there is a unique equilibrium under asynchronicity. Calcagno et al. (2014) show uniqueness of equilibrium in a finite-horizon setting with asynchronicity and a stage game that is a (not

¹This cost may arise from a small probability of detection by the anti-trust agency.

²Another possibility with a similar payoff structure that may potentially involve private timing arises when two firms produce complementary products, such as CPU and RAM. Each product can be of high or low quality. Both choosing high quality yields the best outcome for both firms. If the opponent firm chooses a low quality, however, then it is costly to choose high quality because the development cost of high quality is large while the revenue is significantly capped by the fact that the complementary product's quality is low. Such a situation can be thought of as a common interest game.

³An example for this could be from telecommunication industry. In the advancement into 4G technology, Huawei favored LTE protocol while Intel favored WiMAX due to the patents they owned. If they both adopt the same protocol, they benefit from sharing satellites, basing stations easily, and tackling technical challenges collectively. In fact, a revolutionary LTE service, which Huawei promoted, went public, and then Intel followed. We note that in reality, there were more firms in the picture, notably with Ericsson and Nokia for LTE and IBM for WiMAX. Although the first version of WinMAX went public in 2006, it was not satisfactory enough so technology development continued, and then LTE went public in 2009, after which Intel followed.

⁴See also Yoon (2001), Lagunoff and Matsui (2001), and Dutta (1995).

necessarily perfect) coordination game.⁵ The basic intuition for asynchronicity helping selection in our paper is similar to that in those papers in the literature. The difference is that those papers assume perfect information, while we assume that players may not observe the actions taken by the other player. In fact, when the component game is a coordination game, no one observes any action by the opponent on the equilibrium path.

Uncertainty about timing is less present in the literature. Kreps and Ramey (1987) provide an example of an extensive-form game in which players do not have a sense of calendar time and do not know which player moves first. They argue that such situations may naturally arise in reality and show that they may give rise to a new issue in specifying players' beliefs at off-path information sets. Matsui (1989) considers a situation involving private timing in a context quite different from ours: He considers an espionage game in which, with a small probability prior to the infinite repetition of the stage game, a player can observe the opponent's supergame strategy and revise her own supergame strategy in response to it, but whether there has been such a revision opportunity is private information. The equilibrium strategies in his model have a similar flavor to the costly-disclosure option in our model, in that they use a supergame strategy in which a player signals to the opponent that he has been able to observe the supergame strategy, by taking an action that is costly in terms of instantaneous payoffs.

Our model studies how timing affects behavior, but some papers have analyzed how behavior affects timing. Ostrovsky and Schwarz (2005, 2006) consider models in which players can target their activity times but there are noises in such choices, which results in uncertainty. Park and Smith (2008) consider a timing game in which players choose their timing to be on the right "rank" in terms of moving times, and the equilibrium strategies entail mixing. Thus uncertainty about timing endogenously arises as a result of mixing by the players. The difference relative to these papers is that, in our paper, players can change their actions depending on their exogenously given moving time and observation at that point. Such conditioning, which seems to fit to the real-life examples that we mentioned, is not present in the aforementioned papers. We note that we also consider the case of

⁵Ishii and Kamada (2011) further examine the role of asynchronicity by considering a model with a mix of asynchronous and synchronous moves.

endogenous moving times in Section 5.3.

The organization of the paper is as follows. Section 2 provides the model. Sections 3 and 4 consider common interest and opposing interest component games, respectively. In Section 5, we provide discussions to deepen the understanding of our results. For example, we consider the case in which each player has some discretion regarding the choice of the moving time. Section 6 concludes. The Appendix contains proofs that are not provided in the main text and the Online Appendix provides additional discussions.

2 Model

Component Game The component game is a strategic-form game $S = (N, (A_i)_{i \in N}, (g_i)_{i \in N})$, where $N = \{1, \dots, I\}$ is the finite set of players, A_i is player i 's finite action space, and $g_i : A \rightarrow \mathbb{R}$ is player i 's payoff function, where $A := A_1 \times \dots \times A_I$.⁶

Dynamic Game In the dynamic game, time is discrete and progresses in an ascending manner. There is a countable set of times $\mathcal{T} \subset \mathbb{R}$, and each player moves once at a stochastic time $T_i \in \mathcal{T}$ which is drawn by Nature according to a commonly-known probability mass function $p(T_1, \dots, T_I)$. For any pair of events E and F such that F has positive probability, let $\text{Prob}^p(E|F)$ be the conditional probability of E given F induced by p . Let $\mathcal{T}_i = \text{supp}(T_i)$. Given a realization of times (t_1, \dots, t_I) , player i chooses an element from $A_i \times \{\text{pay}, \text{not}\}$ at time t_i , after observing her own t_i and (a_j, t_j) of every opponent j who chose (\cdot, pay) at $t_j < t_i$.⁷ If player i chooses (a_i, pay) for some $a_i \in A_i$, then she pays the cost $c > 0$.⁸ We denote by $\Gamma = (S, \mathcal{T}, p, c)$ the complete specification of the dynamic game. We will omit the reference to Γ whenever there is no room for ambiguity.

⁶In most of the paper, we focus on the two-player case except in some of the analysis, such as the existence theorem (Theorem 3), where we consider n players.

⁷We assume that disclosures succeed with probability 1. Online Appendix B.1 discusses the case in which disclosures fail with positive probability.

⁸The complete specification of the payoff structure is provided shortly.

Strategies A history h is composed of a sequence of times and action choices by all players:

$$h = (t_i, a_i, d_i)_{i \in N} \in \times_{i \in N} [\mathcal{T}_i \times A_i \times \{\text{pay, not}\}]$$

where t_i is the moving time of player i and $(a_i, d_i) \in (A_i \times \{\text{pay, not}\})$ is the profile of the action and disclosure decision of player i at that time. If $h = (t_i, a_i, d_i)_{i \in N}$ is such that $p((t_i)_{i \in N}) > 0$, we say that h is **feasible**. Let \mathcal{H} be the set of feasible histories.

Player i 's private history, h_i , is defined as

$$h_i = (N', (t_j, a_j)_{j \in N'}, t) \in \bigcup_{N' \in 2^N \setminus \{i\}} [\{N'\} \times (\times_{j \in N'} [\mathcal{T}_j \times A_j]) \times \mathcal{T}_i].$$

Here, N' is the set of players who moved before i and chose to disclose their actions. For each player $j \in N'$, (t_j, a_j) specifies j 's action and moving time. The last element of i 's private history, t , denotes player i 's moving time.

We say that a history $\tilde{h} = (\tilde{t}_j, \tilde{a}_j, \tilde{d}_j)_{j \in N}$ is **compatible** with a private history $\hat{h}_i = (N', (\hat{t}_j, \hat{a}_j)_{j \in N'}, t)$ if (i) $t = \tilde{t}_i$, (ii) $\tilde{t}_j < \tilde{t}_i$ and $\tilde{d}_i = \text{pay}$ if and only if $j \in N'$, and (iii) $\hat{t}_j = \tilde{t}_j$, and $\hat{a}_j = \tilde{a}_j$ for all $j \in N'$. The set of all possible private histories that have some feasible history compatible with them is denoted $\mathcal{H}_i := \{h_i | \exists h \in \mathcal{H} \text{ s.t. } h \text{ is compatible with } h_i\}$.

Player i 's strategy, $\sigma_i : \mathcal{H}_i \rightarrow \Delta(A_i \times \{\text{pay, not}\})$, is a map from private histories to (possibly correlated) probability distributions over A_i and disclosure decisions. Let Σ_i be the set of all strategies for player i . Define $\Sigma_{-i} = \times_{j \neq i} \Sigma_j$ and $\Sigma = \times_{i \in N} \Sigma_i$.

Payoffs If the chosen action profile is a and i chooses $d_i \in \{\text{pay, not}\}$, then her overall payoff is

$$g_i(a) - c \times \mathbb{I}_{d_i=\text{pay}}$$

with $c > 0$ which we assume to be common across players.⁹ That is, if i chooses “pay” to disclose her action, she incurs cost c . The expected payoff for player i

⁹As will become clear, the assumption that c does not vary across players is imposed only for notational simplicity and is not crucial for any of our results.

from strategy profile σ is denoted by $u_i(\sigma)$.¹⁰ A *belief* $\mu \in \Delta(\mathcal{H})$ is a probability measure over histories. A strategy profile σ induces a continuation payoff $u_i(\sigma|\mu, t)$ conditional on the belief that (i) the distribution of the past play at times strictly before t is given by μ , and (ii) the play at and after time t is given by σ .

Let the set of dynamic games defined above be denoted by \mathcal{G} .

Equilibrium Notion A strategy profile σ induces a probability distribution over the set of histories \mathcal{H} . Let $\mathcal{H}(\sigma)$ be the set of histories that have positive probability given σ . The strategy profile σ is a *weak perfect Bayesian equilibrium* (henceforth we simply call this a “*PBE*”) if, for each player i , the following two conditions hold:

1. (On-path best response) $u_i(\sigma) \geq u_i(\sigma'_i, \sigma_{-i})$ for all $\sigma'_i \in \Sigma_i$.
2. (Off-path best response) For each h_i , there exists $\mu \in \Delta(\mathcal{H})$ such that every $h \in \text{supp}(\mu)$ is feasible and compatible with h_i , and $u_i(\sigma|\mu, t) \geq u_i(\sigma'_i, \sigma_{-i}|\mu, t)$ for all $\sigma'_i \in \Sigma_i$.

That is, we require optimality on the equilibrium path of play, while off the path we only require optimality against *some* (possibly correlated) distribution over the strategy profile of the opponents and Nature’s moves that is compatible with the observation.¹¹ Note that condition 1 implies that players best-respond to the beliefs computed by Bayes rule on the equilibrium path. In Section 5.5, we discuss

¹⁰In two-player games, for example, player i ’s expected payoff given strategy σ is given by,

$$u_i(\sigma) = \sum_{t, t' \in \mathcal{T}^2} \left(\text{Prob}^p(T_1 = t, T_2 = t') \left(\mathbb{E}_\sigma \left[g_i(a^{(t, t')}) | (d_1^{(t, t')}, d_2^{(t, t')}), T_1 = t, T_2 = t' \right] - c \times \mathbb{I}_{d_i^{(t, t')} = \text{pay}} \right) \right),$$

where $(a^{(t, t')}, (d_1^{(t, t')}, d_2^{(t, t')})) \in A \times \{\text{pay}, \text{not}\}^2$ denotes a choice in the support of σ at moving times (t, t') . The expectation is taken over probabilities over actions and disclosure choices induced by σ at (t, t') .

¹¹This allows for correlated beliefs over Nature’s moves and the opponents’ deviations. This weak notion is enough to establish uniqueness in two-player games that we consider in our main analysis. When we consider more than two players, we would need to introduce a more stringent notion of equilibrium that conforms to convex structural consistency of Kreps and Ramey (1987). This is because, if such correlations are allowed, player 1’s deviation at some time t may make player 2 believe that player 3 will play later and believe 2 has already played some inefficient action upon not disclosing, which may affect 2’s optimal decision.

what would happen if we did not impose condition 2. Existence of PBE is not trivial because the support of the times of play, \mathcal{T} , may not be finite, and we discuss this issue in Section 5.1.

Pure Stackelberg Property Stackelberg actions will play a key role in characterizing uniqueness. We say that a component game $(N, (A_i)_{i \in N}, (g_i)_{i \in N})$, with $|A_i| \geq 2$ for each $i \in N$, satisfies the *pure Stackelberg property* if for each $i \in N$, there is a strict Nash equilibrium $a^i \in A$ such that $g_i(a^i) > g_i(a)$ for all $a \neq a^i$. That is, player i 's payoff from a^i is strictly higher than the payoff from any other action profile. There are two important (exclusive and exhaustive) subclasses of component games satisfying the pure Stackelberg property. The first is *common interest games*, in which in the above definition, $a^i = a^j$ for all $i, j \in N$. Otherwise the component game is called an *opposing interest game*. The latter class includes, for example, the battle of the sexes. We analyze the former class of games in Section 3 and the latter in Section 4. In common interest games, we call a_i^i player i 's *best action* and denote it by a_i^* .

In Online Appendix B.2, we provide two examples that show that multiplicity of PBE may hold in games that do not have the pure Stackelberg property.

3 Common Interest Games

3.1 An Example

Here we consider a simple example that illustrates the intuition of the analysis that follows. There are two players, $i = 1, 2$. There is a probability distribution f over the possible moving times $\mathcal{T} = \mathbb{Z}$,¹² and it has full support and is “sparse.” Specifically, assume that there exists $\varepsilon > 0$ such that $0 < f(t) < \varepsilon$ for all $t \in \mathcal{T}$. We assume that p satisfies $p(t_1, t_2) = f(t_1) \cdot f(t_2)$ for all $t_1, t_2 \in \mathcal{T}$. For any $\varepsilon > 0$, choose an arbitrary p satisfying the stated conditions, and denote it by p^ε .

¹²We define the support to be \mathbb{Z} for simplicity. The support of the timing distribution does not have to be unbounded on the left for the argument to go through. This is because we can equivalently define a dynamic game in which there are countably many moving times in a bounded interval. For instance, the example works with a support of moving times $\{f(x) | x \in \mathbb{Z}\}$ defined by a mapping $f : \mathbb{Z} \rightarrow (-1, 1)$ with $f(x) = 1 - e^{-x}$ for $x \geq 0$ and $f(x) = e^x - 1$ for $x < 0$.

	A	B
A	2, 2	0, 0
B	0, 0	1, 1

Figure 1: A Common Interest Game

Consider the payoff matrix in Figure 1. Let S be this coordination game. We can show the following:

Proposition 1. *There exist $\bar{\varepsilon} > 0$ and $\bar{c} > 0$ such that for all $\varepsilon \in (0, \bar{\varepsilon})$ and $c \in (0, \bar{c})$, there is a unique PBE in the game $(S, \mathcal{T}, p^\varepsilon, c)$ with S , \mathcal{T} , and p^ε as specified above. On the path of the unique PBE, each player i takes (A, not) for any realization of T_i .*

That is, by introducing the option to pay the cost to disclose actions, players are able to coordinate on the best action profile (A, A) without exercising the option to disclose.

The result requires the disclosure cost $c > 0$ to be sufficiently small. In fact, the proof we will present goes through as long as $0 < c < g_i(A, A) - g_i(B, B) = 1$.

To illustrate the non-triviality of this result, below we explain three possibilities of alternative timing distributions and equilibria in these settings.¹³

Example 1. [Simultaneous-Move Game]

Suppose that $N = \{1, 2\}$, $\mathcal{T} = \{1\}$, and $p(1, 1) = 1$. That is, it is common knowledge that, at time 1, both players take actions with probability one. The component game S is the same as in Figure 1. First, no player pays the disclosure cost in any PBE because even if i pays, $-i$ does not have a chance to move after observing it. Hence, the game is strategically equivalent to the static simultaneous-move game. There are three Nash equilibria in such a game, namely (A, A) , (B, B) and a mixed equilibrium. Thus, there are three PBE in the dynamic game corresponding to these three Nash equilibria. \square

Example 2. [Deterministic Sequential-Move Game (and Forward Induction)]

¹³We note that the constructions of multiple equilibria in what follows will not be based on the wild freedom in the choice of beliefs in condition 2 of the definition of PBE. Indeed, as it will become clear, the PBE we construct are sequential equilibria as well.

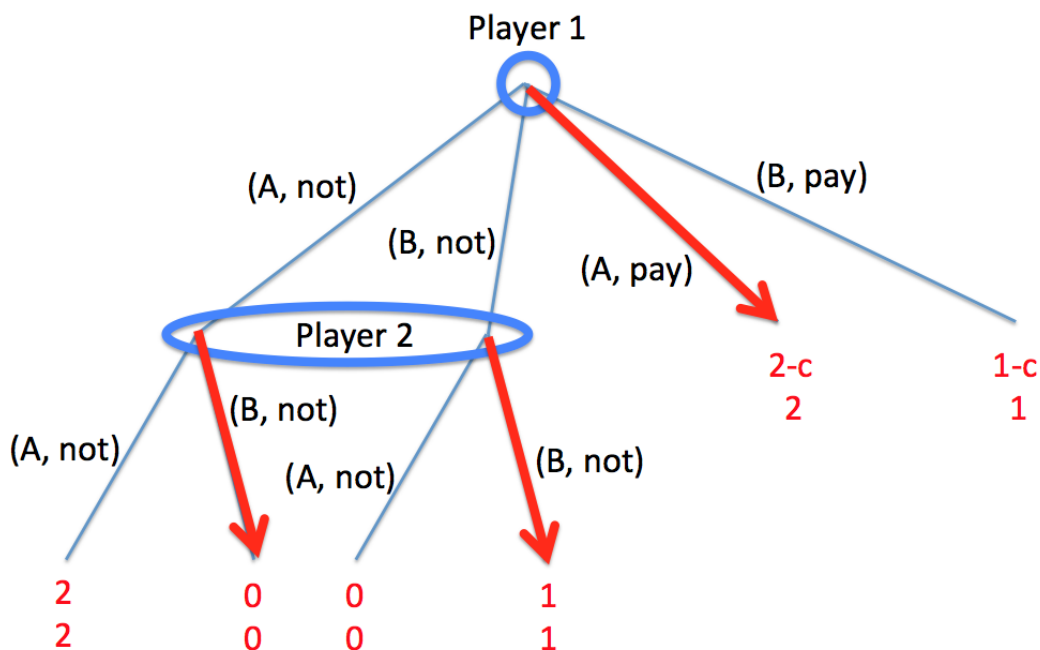


Figure 2: A Forward-Induction Argument

Suppose that $N = \{1, 2\}$, $\mathcal{T} = \{1, 2\}$, and $p(1, 2) = 1$. That is, it is common knowledge that player 1 moves at time 1 and player 2 moves at time 2 with probability one. The component game S is the same as in Figure 1. There are at least two PBE in this game. In the first PBE, each player plays (A, not) on the path of play. In the event that 2 observes 1's action, 2 takes a static best response. The second PBE is what we call the *pessimist equilibrium*. In this equilibrium, player 1 plays (A, pay) , and player 2 plays a static best response if 1 discloses her action, while he plays (B, not) if 1 does not.

Let us check that this second strategy profile constitutes a PBE. First, player 1 takes a best response given 2's strategy. Also, 2's strategy obviously specifies a best response after 1's disclosure. After no disclosure, (B, not) is a best response under the belief that 1 has played (B, not) .

Let us note that this pessimist equilibrium would be ruled out by a so-called "forward induction" argument. To see this point, consider the extensive-form representation in Figure 2 of the game in consideration. Note that we omitted actions corresponding to player 2's "pay," as they are obviously suboptimal. We

	A	B
A	3, 3	-3, 0
B	0, -3	1, 1

Figure 3: A Risky Common Interest Game

also omitted 2's actions following 1's payment, as there is obviously a unique best response (which is to play (A, not) after (A, pay) and to play (B, not) after (B, pay)). The payoffs after 1's payment are written assuming 2's best response. In this game, for player 1, (B, not) is dominated by (A, pay) . A forward induction argument would then dictate that this would imply that at player 2's information set, his belief must assign probability 0 to the right node. Given this belief, player 2's unique best response at the information set is to play (A, not) . Hence, 1 can obtain the best feasible payoff in the game by playing (A, not) , so this is a unique action that can be played in any PBE. Intuitively, if 1 does not pay then 2 should be able to infer that 1 had played A because there is no reason to play B as it is dominated. This suggests 1 should play (A, not) .

Our private-timing game rules out such an outcome *without* resorting to any "forward induction" argument. Still, we will see that the proof is based on a similar idea.¹⁴ □

Example 3. [Correlated-Move Game]

Suppose that $N = \{1, 2\}$ and $\mathcal{T} = \mathbb{Z}$. For all $t_1 \in \mathbb{Z} \setminus \{0\}$, p satisfies $\sum_{t_2 \in \mathbb{Z}} p(t_1, t_2) = \frac{1}{2} \frac{1-r}{r} r^{\frac{|t_1|+1}{2}}$ with $r \in (0, 1)$ for odd t_1 , and $\sum_{t_2 \in \mathbb{Z}} p(t_1, t_2) = 0$ for even t_1 . We also assume that $\text{Prob}^p(T_2 = t_1 - 1 | T_1 = t_1) = \text{Prob}^p(T_2 = t_1 + 1 | T_1 = t_1) = \frac{1}{2}$. That is, T_1 is positive with probability $\frac{1}{2}$, negative with probability $\frac{1}{2}$, and follows an exponential distribution with rate r on each side over the odd integers. Player 2's moving time is either right before or right after player 1's, with equal probability. These conditions imply that $\text{Prob}^p(T_1 = t - 1 | T_2 = t) : \text{Prob}^p(T_1 = t + 1 | T_2 = t) = 1 : r$ for all even $t \geq 2$ and an analogous condition holds for all even $t \leq -2$ (the ratio is 1 : 1 if $t = 0$). For the component game S , consider the game in Figure 3.

There are at least two PBE in this game when $r < 1$ is sufficiently close to 1. In the first PBE, each player plays (A, not) on the path of play. In the event that

¹⁴See Remark 2 for discussion on this similarity.

player i observes the opponent j 's action, i takes the static best response. The second PBE is one in which each player plays (B, not) on the path of play. In the event that i observes j 's action, i takes the static best response.

The second strategy profile is a PBE because a deviation by player 1 to A can only be profitable if it involves disclosure, in which case she would succeed in coordination with probability $1/2$ and miscoordinate with player 2 with probability $1/2$. Hence the expected payoff from these events is $\frac{1}{2} \cdot 3 + \frac{1}{2} \cdot (-3) - c = -c$, and it is smaller than the payoff of 1 from not deviating. A similar reasoning applies to player 2's incentive when $r < 1$ is sufficiently close to 1 so the probability that 1 moves before or after 2 is close enough to $\frac{1}{2}$ at every period he can move. The key difference from the game in Proposition 1 is that at any realization of a player's moving time, the probability that the player assigns to being the first mover is only (close to) $1/2$, while in the game in Proposition 1 it can be arbitrarily close to 1 as $t \rightarrow -\infty$.¹⁵ \square

Proof Sketch of Proposition 1

Here we provide a rough sketch of the proof of Proposition 1. The formal proof is given as a proof of Theorem 4, as Proposition 1 is a corollary of that theorem.¹⁶ The proof consists of two steps. In the first step, we prove that players must play A on the path of play of every PBE. The second step shows that no player pays the disclosure cost.

- *First Step:* Suppose that player i gets her move at time t and she has observed no action disclosed. The assumption that the timing distribution is independent across players and the probability of simultaneous moves is small implies that, if t is early enough, the probability that she assigns to the event that the opponent moves later is close to 1. At such t , the expected payoff from playing (A, pay) is close to $2 - c$, while that from playing (B, pay) or (B, not) is at most 1. Hence, at t , player i does not play B in any PBE. By

¹⁵In this example, any private information about a player's own moving time does not reveal sufficiently precise information about the order of moves. Online Appendix B.3 makes this point even clearer by considering an extreme case in which players do not have a sense of calendar time.

¹⁶In the language in Theorem 4, S is $\frac{1}{2}$ -common for each player and is q -dispersed for any $q \leq 1 - \varepsilon$, where ε is the upper bound of the probability at a given time defined in this section.

independence of the timing distribution, this is true for all times before t as well. Since $-i$ knows this, $-i$ does not play B at $t+1$ in any PBE. Applying this argument iteratively shows that, for any time, the moving player never plays B .

- *Second Step:* Suppose that at time t , i gets her moving opportunity and has not observed any action disclosed. From the first step we know that i plays A at time t . If she plays (A, pay) , her expected payoff is at most $2 - c$. Suppose that i does not pay. If $-i$ has played before or at t , then i knows that $-i$'s action was A . So consider any time $t' \geq t$ at which $-i$ moves. Given no observation of a disclosed action, $-i$ at t' assigns positive probability that i will move after t' , so $-i$ is on the path of play. Hence, by the first step, $-i$ plays A at t' . In total, regardless of i 's belief about $-i$'s moving time, i knows that $-i$'s action is A with probability 1. Thus the expected payoff from i 's playing (A, not) is 2. Hence, it is a unique best response to play (A, not) . Since the choice of t was arbitrary, this shows that on the path of any PBE, at any t , player i does not pay the disclosure cost. \square

Remark 1. [Relation to Examples]

How does this proof relate to the three examples that we examined? In Example 1, we demonstrated that simultaneity may prevent uniqueness. In the Proof Sketch, we used the fact that the probability of simultaneous moves is small in the first step, where we concluded that $-i$ never plays B at time $t+1$. To obtain this conclusion, we could effectively ignore the possibility that i moves at $t+1$ because the probability of such an event is small. Example 2 illustrated the effect of deterministic moves on multiplicity. We used the fact that the timing distribution is uncertain in the second step of the Proof Sketch, where we concluded that $-i$ at time t' plays A because he is on the equilibrium path. Such a conclusion cannot be obtained in Example 2: If player 2 has not had any observation, he is off the equilibrium path under the pessimist equilibrium. Example 3 showed that multiplicity is possible under a highly correlated timing distribution. In particular, in the second PBE in Example 3, for any realized moving time, the moving player assigns a nontrivial probability to the event that the other player has already moved. We used the fact that the timing distribution is independent in the

first step of the Proof Sketch, where we argued that, at early enough times when i has not observed the opponent's disclosure, she assigns only a small probability to the event that the opponent has already moved. \square

Remark 2. [Similarity to the Forward Induction Argument]

In Step 1 of the Proof Sketch above, B is ruled out as it is dominated by (A, pay) . The forward induction argument that we described in Example 2 uses this fact, too. Specifically, it uses this fact to argue that player 1 should expect that player 2 (without observing disclosure) cannot expect 1 has played B , hence 2 plays A . This in turn implies that if 1 plays (A, not) she gets the best payoff. In our private-timing game, we do not need to rely on such an inference procedure to argue that player i should expect that $-i$ will play A if $-i$ has not moved yet. This is because player $-i$ will assign a positive probability of i moving later so he will be on the path of play, thus, he will take A in equilibrium due to Step 1 of the Proof Sketch. Step 1 resembles the logic of forward induction in that the action (A, pay) is used to eliminate B . However, in forward induction, the fact that 2's information set after no disclosure is reached is interpreted as containing sure information regarding 1's choice. In our model, on the other hand, the reasoning relies on the fact that 2's getting a move without observation of disclosure still leaves the possibility that 1 has not moved. \square

Remark 3. [Lack of Lower Hemicontinuity]

Suppose that $N = \{1, 2\}$ and $\mathcal{T} = \mathbb{Z}$. The timing distribution is independent across players and parameterized by $\xi \in (0, 1)$, and satisfies $\sum_{t \in \mathbb{Z}} p(1, t) = \sum_{t \in \mathbb{Z}} p(t, 2) = 1 - \xi$ and $p(t, t') > 0$ for all $t, t' \in \mathcal{T}$. For any ξ , when the disclosure cost $c > 0$ is small enough, the same logic as in the above Proof Sketch applies to show that there is a unique PBE, and in that unique PBE each player plays (A, not) .

These timing distributions converge pointwise to the distribution with $p(1, 2) = 1$ as $\xi \rightarrow 0$. However, as Example 2 shows, there are multiple equilibria under this limit distribution. Thus, there is a lack of lower hemicontinuity with respect to the timing distribution.¹⁷ One additional PBE that obtains in the limit is the

¹⁷Consider, for example, the sup norm: For two timing distributions p and p' , the distance between them is $\sup_{t, t' \in \mathcal{T}} |p(t, t') - p'(t, t')|$.

pessimist equilibrium. The reason for the lack of lower hemicontinuity is that, for any time t at which player i moves, she expects a positive probability of $-i$ moving later. This condition will be formalized as the “potential leader condition” in Definition 1 in Section 3.2. \square

3.2 General Common Interest Games

In this section, we consider general common interest games to understand the role of the timing distribution on the set of PBEs. For this purpose, we characterize the set of timing distributions such that the best action profile with no disclosure is the unique PBE outcome when the component game is a common interest game. The following two conditions are satisfied in the timing distribution of the starting example in Section 3.1.

Definition 1. The timing distribution p satisfies the **potential leader condition** if for any pair of distinct players $i, j \in N$ and all $t \in \text{supp}(T_i)$, $\text{Prob}^p(T_j > t | T_i = t) > 0$ holds.

That is to say, the potential leader condition holds if at every moving time without any past observation, each player deems it possible that the other player will have a later opportunity to move. This condition is not satisfied in the timing distribution of Example 2. This is because player 2 in that example assigns probability 0 to being the first mover (he is not a “potential leader”).

Definition 2. We say that the timing distribution p has **dispersed potential moves** if for every $q > 0$ and $t'' \in \{-\infty\} \cup \mathbb{R}$, there are $i \in N$ and $t' \in \mathbb{R}$ such that $t' > t''$ and for all $j \neq i$ and $t \in [t'', t'] \cap \mathcal{T}_i$, $\text{Prob}^p(t'' < T_j < t | T_i = t) < q$.

The dispersed potential moves condition holds if at every time t'' and small enough interval (t'', t) , there is a player who attaches small probability to the event that the other player moves before her within (t'', t') , conditional on moving at time t . Thus, every time has a vicinity and a player whose belief over the other players’ moves when moving within the vicinity is dispersed. In other words, players find it unlikely that the opponent will move in a short time interval.

The condition also requires that there is at least one player such that, if she moves early enough, she assigns only a small probability to the event that she is

the second mover (note that t'' can be equal to $-\infty$). That is, when it is early in the game, they must believe it is unlikely that the other player has not moved yet. For this reason, this condition is not satisfied in the timing distribution of Example 3. To see this formally, take $t'' = -\infty$ and $q > 0$ small enough, and observe that there is no t' satisfying the specified property. In other words, for whatever early time i moves, i assigns a nontrivial probability to her being the second mover. In fact, the “dispersed potential moves” condition implies that there is a player that on an early enough move thinks the probability of being the second mover is low. This condition is satisfied if, for example, the distributions of moves are independent across players.

The following condition is not satisfied in the starting example, but is used in the result below. See Remark 4-1 for the discussion of the connection between Proposition 1 and Theorem 1.

Definition 3. Timing distribution p is **asynchronous** if $\text{Prob}^p(T_i = T_j) = 0$ for all distinct players $i, j \in N$.

For the case of $N = \{1, 2\}$, define $T_i^< = \{t \in \mathcal{T}_i | \text{Prob}^p(T_{-i} > t | T_i = t) = 0\}$. Also, define the set $D \subseteq \Delta(\mathcal{T})$ as follows. The set D is the set of distributions such that both of the following two conditions hold:

1. p is asynchronous.
2. Either one of the following two conditions holds.
 - (a) p satisfies the potential leader condition.
 - (b) For every player i with non-empty $T_i^<$, $E \subseteq T_i^<$, $\inf_{\tilde{t} \in \mathcal{T}_{-i}} \{\text{Prob}^p(T_i \in E | T_{-i} = \tilde{t}) | \text{Prob}^p(T_i \in E | T_{-i} = \tilde{t}) > 0\} = 0$.

Condition 2b requires that, if there is an event at which a player is certain that she is the second mover, then the conditional probability that the opponent attaches to such an event can approximate zero as the conditioning on the opponent’s moving time varies.

Theorem 1. Let $N = \{1, 2\}$. Fix \mathcal{T} , and assume that p has dispersed potential moves. The timing distribution is in the set $D \subseteq \Delta(\mathcal{T})$ if and only if, for any

common interest game S , there exists $\bar{c} > 0$ such that for all $c < \bar{c}$, there is a unique PBE in the game (S, \mathcal{T}, p, c) . On the path of the unique PBE, each player i plays (a_i^*, not) at any realization of T_i .

Remark 4.

1. The proof for sufficiency of the two conditions has the same structure as that of the starting example in Section 3.1. In the example, however, we were only concerned with sufficiency. The asynchronicity condition was not satisfied in the example but the result held because the timing distribution was “close to” the one satisfying the asynchronicity condition for a fixed component game. This suggests the possibility of making a connection between the degree of commonality of the players’ interest in the component game and the level of dispersion of the two players’ moving times, where the degree of commonality is measured by the difference between each player’s best payoff and her second-best payoff. Section 5.2 explores this point by explicitly defining a commonality parameter and a dispersion parameter. Since the proof in that section implies sufficiency of conditions 1 and 2a, for sufficiency, in the proof below we refer to Theorem 4 in Section 5.2 and only prove sufficiency of conditions 1 and 2b.
2. The proof for necessity is by construction. That is, for each timing distribution violating any of the conditions characterizing the set D , we construct an example of a component game that has multiple PBE. For example, for a distribution that puts positive probability on synchronous moves at some time t , it is easy to construct a component game in which an action other than a_i^* may be a best response for each player i at time t . A tricky part of the proof is showing that there exists a PBE strategy profile (specifying all contingent plans at all times) that induces the desired play at such time t . More specifically, we show that for any distribution with synchronous moves, the component game in Figure 4 with $\alpha > 0$ has a PBE in which B is played at certain times for sufficiently large M . The necessity of condition 2 is established using an analogous argument that relies on the same component game with $\alpha = 0$.

	A	B
A	$1, 1$	$-M, -\alpha$
B	$-\alpha, -M$	$0, 0$

Figure 4: Counterexample for Necessity part

3. Condition 2 consists of two parts, and only one of them is required for a distribution to be in set D . That is, the potential leader condition (condition 2a), together with asynchronicity, is sufficient but not necessary to guarantee uniqueness. Example 4 below illustrates this point.
4. Only at times in which player i knows for sure that $-i$ has already moved (i.e times in $T_i^{<}$) can a failure to observe disclosure be off-path. In such cases it is possible, as in Example 2, that upon no disclosure i chooses an action other than a_i^* in equilibrium. Condition 2b guarantees that in the event that player i moves at a time in $T_i^{<}$ “no disclosure” is on the equilibrium path in every PBE. The reason is that 2b requires that there be times in which $-i$ thought it so unlikely that i would move in any subset of $T_i^{<}$ that the cost of disclosure was not justified. Example 4 illustrates the role of condition 2b.
5. Notice that condition 2a requires that $t < \sup_{t' \in \mathcal{T}} t'$ hold for every t , while condition 2b requires that there be infinitely many times before t . At least one of these conditions is necessary to guarantee uniqueness. In particular, if \mathcal{T} is finite then the uniqueness result does not hold. We make this point clear in Section 5.4.
6. One possible application of the theorem is a case in which the analyst only knows that the players face a common interest game, that the disclosure costs are small, and that the structure of the game is common knowledge among the players, but she does not know the cardinal utility of the players. The theorem identifies the conditions under which the analyst can be certain that the Pareto efficient outcome (i.e., a^* is played and no payment for disclosure takes place) is obtained. It is possible that the analyst’s interest is only in the actions in the component game and not in the disclosure behavior. It is relatively straightforward to identify the necessary and sufficient condition on a timing distribution p , with dispersed potential moves, such that a^* is

played with probability one in all PBE. The condition turns out to be the asynchronicity condition. \square

Example 4. [Second-Mover Game]

Suppose that the component game S is as in Figure 1. The timing distribution p over $\mathcal{T} = \mathbb{Z}$ is given by the following rule: With probability $\frac{1}{2}$, T_1 follows a geometric distribution over positive integers with parameter p , while T_2 follows a geometric distribution over nonpositive even integers with parameter p .¹⁸ With the complementary probability $\frac{1}{2}$, T_1 follows a geometric distribution over negative odd integers with parameter $p' < p$, while $T_2 = T_1 + 1$.

Note that the distribution p does not satisfy the potential leader condition, because in the first event, player 1 assigns probability 1 to the event that she is the second mover. Player 2 does not know which event he is at because his moving time is always a negative even time. However, $p' < p$ implies that the likelihood of him being in the second event becomes arbitrarily close to 1 as time goes to $-\infty$. Thus, for any common interest game, there exists a \bar{t} such that for all $t < \bar{t}$, it is not worthwhile for player 2 to pay the disclosure cost because player 2 is so sure that he is the second mover. This implies that even though player 1 knows she is the second mover at positive times, she thinks that she is still on the path of equilibrium play even if she does not observe any past disclosure.

Using arguments analogous to the ones for Proposition 1, we can conclude that the dynamic game (S, \mathcal{T}, p, c) has a unique PBE. Condition 2b allows for this type of timing distribution. \square

4 Opposing Interest Games

In the previous section, we considered common interest games. This is a class of games with the pure Stackelberg property. The rest of the games with $N = \{1, 2\}$ satisfying this property are opposing interest games. In this section, we consider two-player opposing interest games and show uniqueness of a PBE. In contrast to the case with common interest games, we show that under certain regularity conditions, players pay the disclosure cost in the unique PBE.

¹⁸If a random variable T is distributed according to a geometric distribution with parameter p over a sequence of times $\{t_k\}_{k=1}^{\infty}$, then $\text{Prob}(T = t_k) = p(1 - p)^{k-1}$ for each $k \in \mathbb{N}$.

Given an opposing interest game, let $g_i^* := \max_{a \in A} g_i(a)$ and $g_i^S := \max\{g_i(a) \mid a \in A, g_i(a) \neq g_i^*\}$ be the best and the second best payoffs, respectively, for player i . We let a_i^* be the action such that there exists a_{-i} with $g_i(a_i^*, a_{-i}) = g_i^*$.

For each tuple $(t, t_0, \underline{t}, \bar{t}) \in \mathbb{R} \times (\mathbb{R} \cup \{-\infty\})^2 \times \mathbb{R}$ such that $\text{Prob}^p(T_{-i} \geq t_0, T_i = t) > 0$, we define

$$p_i(t, t_0, \underline{t}, \bar{t}) = \text{Prob}^p(T_{-i} \in [\underline{t}, \bar{t}] \mid T_{-i} \geq t_0, T_i = t).$$

Assumption 1.

1. **(Dispersed, frequent and asynchronous potential moves)**

For each $i = 1, 2$, $\forall \varepsilon > 0$ and $t_0 \in (-\infty, \sup \mathcal{T}_i) \cup \{-\infty\}$, $\exists \bar{t}_i \in \mathcal{T}_i \cap (t_0, \infty)$ such that $\forall t \in [t_0, \bar{t}_i] \cap \mathcal{T}_i$, $p_i(t, t_0, t_0, t) < \varepsilon$,

2. **(Similar concentration)**

There is $\alpha > 0$ such that, for every $i, j \in \{1, 2\}$ and $(t_0, t_R, t_L) \in \mathcal{T}_j \times \mathcal{T}_i \times \mathbb{R}$ with $t_0 \leq t_R$, $t_L < t_R$, $\text{Prob}^p(T_{-i} \geq t_L, T_i = t_R) > 0$ and $\text{Prob}^p(T_{-j} \geq t_L, T_j = t_0) > 0$, we have $p_j(t_0, t_L, t_L, t_R) \geq \alpha p_i(t_R, t_L, t_L, t_R)$.

3. **(Monotonicity of conditional probability)**

- (a) For each $i = 1, 2$, $p_i(t, t_0, t_0, t)$ is non-decreasing in t for $(t, t_0) \in \mathcal{T}_i \times (\mathbb{R} \cup \{-\infty\})$ with $t_0 \leq t$ such that $\text{Prob}^p(T_{-i} \geq t_0, T_i = t) > 0$,
- (b) For each $i = 1, 2$, for every $t_1 \in (\inf \mathcal{T}_{-i}, \infty)$ there is $t_0 \in \mathcal{T}_i \cap (-\infty, t_1)$ such that $p_i(t_0, -\infty, t_0, t_1) > 0$ and $p_i(t, -\infty, t, t_1)$ is non-increasing in t for $t \in \mathcal{T}_i \cap (-\infty, t_1]$.

Assumption 1 (1) says that at each time t_0 , each player i has a close potential move \bar{t}_i such that, conditional on i moving at time $t \in [t_0, \bar{t}_i]$, the opposing player is unlikely to move between t_0 and t . The assumption rules out simultaneous moves. In addition to dispersed potential moves, the condition requires sufficiently frequent potential moves. Assumption 1 (2) says that for any time interval with end points t_L and t_R in \mathcal{T} , conditional on no one having moved before the left endpoint of the interval (t_L), the probability of player j 's opponent moving in that time interval conditional on j moving weakly earlier than the right endpoint

(t_R) cannot be too low relative to the probability of player i 's opponent moving in that time interval conditional on i moving at the right endpoint, where i and j can be either the same or different. The condition is satisfied if players do not learn “too much” about the opponent’s moving time from moving either at the end or the beginning of an interval and the players’ priors are similarly concentrated across times. Assumption 1 (3a) says that, for any given t_0 , the probability that the opponent has moved in $[t_0, t]$ conditional on receiving an opportunity at t is non-decreasing in t . This condition is satisfied, for example, when the timing distribution is independent. Assumption 1 (3b) says that for each player i for any time t_1 , there is a sufficiently early moving time of player i , t_0 , at which i believes with positive probability that $-i$'s moving time is between t_0 and t_1 . Also, the probability of $-i$ moving between t_0 and t_1 is non-increasing in t_0 . Note that both Assumptions 1 (1) and 1 (3b) imply that $\inf \mathcal{T}_1 = \inf \mathcal{T}_2$. Assumption 1 is a regularity condition that essentially requires that moving times are sufficiently “dispersed” and well behaved. It is satisfied if, for example, player 1's moves have full support on the odd integers, player 2's have full support on the even integers, the moving times are independently distributed, and the following holds for each $i \in \{1, 2\}$: $\text{Prob}^p(T_i = k) \geq \text{Prob}^p(T_i = k + 2)$ for each $k = 0, 1, 2, \dots$, and there exists $\tilde{\alpha} \in (0, 1]$ such that $\text{Prob}^p(T_i = k)\tilde{\alpha} \leq \text{Prob}^p(T_{-i} = k + 1)$ for each $k = 0, 1, 2, \dots$, and analogous conditions hold for nonpositive k 's.¹⁹

Theorem 2. *Let S be a two player opposing-interest game. Suppose that (\mathcal{T}, p) satisfies Assumption 1. Then, there exists $\bar{c} > 0$ such that for every $c < \bar{c}$, the dynamic game (S, \mathcal{T}, p, c) has a unique PBE. On the path of the unique PBE, each player i plays (a_i^*, pay) without observation of a disclosure at any realization of T_i .*

As in the common interest games, sufficient uncertainty about the opponent's moves guarantees the uniqueness of a PBE. Note that if an opposing interest game is played under the timing distribution as in the deterministic-move game of Example 2, Assumption 1 (3b) is not satisfied, and there are multiple PBE.

¹⁹The following is an example of the distributions that satisfy those conditions: Player 1's moving time follows a geometric distribution with parameter $p > 0$ over positive odd numbers with probability $1/2$, it follows a geometric distribution with parameter $p > 0$ over negative odd numbers with probability $1/2$, and $\text{Prob}^p(T_2 = t) = \text{Prob}^p(T_1 = t - 1)$ for each even integer t .

5 Discussion

5.1 Existence

In the main sections, we focused on component games that satisfy the pure Stackelberg property. In other classes of games, general predictions are hard to obtain. We can establish existence, however, for any choice of component games.

Theorem 3. *Every $(S, \mathcal{T}, p, c) \in \mathcal{G}$ has a PBE.*

The proof is provided in the Appendix. A difficulty is that the support of the times of play is infinite, so the standard fixed-point argument does not apply. Moreover, since we deal with general component games, there is no obvious way to conduct a constructive proof as was possible in the main sections. The proof consists of five parts. The first part proves a lemma stating that any sequence of strategy profiles in Σ has a convergent subsequence. The second part defines finite horizon games with arbitrary length N that we call the N 'th approximating game, and the third part defines ε -constrained equilibria which exist in the N 'th approximating game. Part 4 uses the lemma in part 1 to show that an ε -constrained equilibrium exists in the original game with possibly infinite horizon by considering a subsequence of a sequence of ε -constrained equilibria in the N 'th approximating game as $N \rightarrow \infty$. Finally, part 5 again uses the lemma to show existence of a trembling-hand perfect equilibrium by considering a subsequence of a sequence of ε -constrained equilibria in the possibly infinite horizon game as $\varepsilon \rightarrow 0$ as in Selten (1975) (which is a PBE, too). The reason we use trembling-hand perfect equilibrium is that it makes the equilibrium play after off-path histories easy to handle. The approximation using finite horizon games and trembling-hand equilibria in games with stochastic opportunities and uncountable histories is analogous to the method used in Moroni (2015). The difference is that, in our setup the set of possible arrival times can have any distribution over a countable set, whereas in Moroni (2015) the distribution of arrivals is given by a Poisson process. A by-product of using this proof method is that it shows existence of a trembling-hand equilibrium. Thus, the fact that our definition of PBE is not stringent does not play a key role in proving existence.

5.2 q -Dispersed Timing Distribution and s_i -Common Interest Games

Although Theorem 1 requires asynchronicity of moves, Proposition 1 proves uniqueness of PBE allowing for a small degree of synchronicity. This is because the former considers any common interest games, while the latter applies to a fixed common interest game. This suggests that there may be a relationship between the type of common interest game we fix and the degree of synchronicity of the timing distribution when proving uniqueness of a PBE. This section provides one way to express such a relationship.

Given a common interest game S , let $g_i^* := g_i(a^*)$ be player i 's payoff from the best action profile. We also let $\underline{g}_i := \min_{a \in A} g_i(a)$ be the minimum payoff, and $g_i^S := \max\{g_i(a) | a \in A, g_i(a) \neq g_i^*\}$ be the second-highest payoff for player i . Notice that the pure Stackelberg property implies $g_i^* \neq \underline{g}_i$.

Definition 4. For any $s_i > 0$, a common interest game is s_i -**common** for i if $\frac{g_i^* - g_i^S}{g_i^* - \underline{g}_i} = s_i$.

Note that $s_i \in (0, 1]$, and it measures how good the best payoff is for player i .

Definition 5. We say that the timing structure p is q -**dispersed** if for every $t'' \in \{-\infty\} \cup \mathbb{R}$ there exist $i \in N$ and $t' \in \mathbb{R}$ such that $t' > t''$ and for all $j \neq i$ and $t \in (t'', t'] \cap \mathcal{T}_i$, $\mathbb{P}(t'' < T_j \leq t | T_i = t) < 1 - q$.

According to this definition, q -dispersion implies that the probability that the two players move at the same time at any one time t is less than $1 - q$. If a distribution is q -dispersed for every $q \in (0, 1)$ then it has dispersed potential moves and is asynchronous. In particular, q -dispersion is satisfied when T_1 and T_2 are independent and the probability of each moving time is strictly less than $1 - q$ (this includes our starting example in Section 3.1 with $\varepsilon \leq 1 - q$). Notice that in Example 3, the assumption fails when $q < 1$ is sufficiently close to 1. This is because for any $t \in \mathcal{T}$, the probability that the opponent moves earlier is at least $\frac{r}{1+r}$.

We now provide a sufficient condition on the joint distributions of T_1 and T_2 and the game S such that (a^*, not) is the only outcome of the private-timing game when the cost of disclosure is small enough.

Theorem 4. Fix a dynamic game (S, \mathcal{T}, p, c) with $N = \{1, 2\}$. Suppose that there exist $(s_i)_{i \in N} \in \mathbb{R}_{++}$ and $\varepsilon > 0$ such that S is a common interest game that is s_i -common for each $i \in N$ and p is $(1 + \varepsilon - \min_{i \in N} s_i)$ -dispersed. Then there exists $\bar{c} > 0$ such that for all $c < \bar{c}$, a^* is assigned probability one under any PBE of (S, \mathcal{T}, p, c) . Moreover, if the potential leader condition holds, then there is a unique PBE. On the path of this unique PBE, each player i takes (a_i^*, not) for any realization of $T_i \in \mathcal{T}_i$.

We note that the proof, which we provide in the Appendix, is used to establish the sufficiency part of Theorem 1.

5.3 Choice of Moving Times

In the main sections, we assumed that players do not have any control over when to move. This assumption fits many real-life situations discussed in the Introduction. In other situations, however, players may have some control over the timing of moves. Here we consider such a situation and argue that our results for common interest games go through as long as there is *some* uncertainty about timing.

To formalize the idea, we confine attention to the case where players' moving times are independent. Specifically, we consider the following two-player two-stage model. In the first stage, each player $i = 1, 2$ can simultaneously choose a timing distribution for i . Let $\mathcal{D}_i \subseteq \Delta(\mathcal{T})$ be the set of possible distributions over player i 's moving times from which i can choose her own distribution. The choice of the timing distribution is not revealed to the opponent before the dynamic game is played. In the second stage, the dynamic game (S, \mathcal{T}, p, c) is played, where p is the independent distribution comprised of the two players' choices of timing distributions. Let $(S, \mathcal{T}, (\mathcal{D}_i)_{i=1,2}, c)$ be this new two-stage dynamic game. Here are three special cases of this model:

1. If \mathcal{D}_i is a singleton for each player, then the model reduces to (S, \mathcal{T}, p, c) where p is an independent distribution where each player i 's distribution is the same as the unique element in \mathcal{D}_i .
2. For each player, with some probability q , her moving time is drawn from some distribution f over \mathcal{T} , while with the complementary probability $1 - q$,

she can freely choose her moving time from \mathcal{T} .

3. Each player can freely choose her moving time from \mathcal{T} but there is some noise, so some probability is assigned to the times close to the chosen time, while with the complementary probability the moving times can be far from the chosen time.

Corollary 2. *Let S be a two-player common interest game with the best action profile (a_1^*, a_2^*) . Suppose that, for all pairs $(f, g) \in \mathcal{D}_1 \times \mathcal{D}_2$, $\min \{ \sum_{t' > t} f(t'), \sum_{t' > t} g(t') \} > 0$ for all $t \in \mathcal{T}$ and $\text{supp}(f) \cap \text{supp}(g) = \emptyset$. Then, there exists $\bar{c} > 0$ such that for all $c < \bar{c}$, $(S, \mathcal{T}, (\mathcal{D}_i)_{i=1,2}, c)$ has a unique PBE. On the path of the unique PBE, each player i plays (a_i^*, not) at any realization of T_i .*

This proposition is a straightforward corollary of the proof of Theorem 4. It implies, in particular, that the selection result goes through even in models described in cases 2 and 3 in which the probability that the choice of a player's own moving time becomes relevant can be arbitrarily high (but is less than 1).

This corollary implies the following result that states uniqueness of a trembling-hand equilibrium in a game in which players choose their moving time. To state the result formally, consider an extensive-form game in which the following occurs.

1. First, each agent i simultaneously chooses an element of \mathcal{T}_i . Each agent observes her own choice, but does not observe other agents' choices.
2. Second, the private timing game is played. That is, if an agent chooses t in the previous stage, the agent moves at t , choosing her normal-form game action as well as whether to disclose the action. When doing so, she observes all disclosed actions at times strictly before time t .

We call this game the **game with moving-time choice**. It is characterized by (S, \mathcal{T}, c) .

In this game, let Π_i be the set of behavioral strategies of player i , $\Pi = \times_{i \in N} \Pi_i$. For $\pi_i \in \Pi_i$, let \emptyset be the null history, and h_t be a typical element of the set of histories at time t in the second stage. We extend the definition of the payoff function in a natural way by letting $u_i : \Pi \rightarrow \mathbb{R}$.

A totally mixed behavioral strategy profile $\pi^\varepsilon \in \Pi$ is an ε -**constrained equilibrium** if, for each $i \in N$ there are $\varepsilon(\tilde{a}_i) \in (0, \varepsilon)$ for each $\tilde{a}_i \in A_i \times \{\text{pay}, \text{not}\}$ and $\{w_t\}_{t \in \mathcal{T}_i}$, with $w_t \in (0, \varepsilon]$ for every $t \in \mathcal{T}_i$, such that

$$\pi_i^\varepsilon \in \arg \max_{\pi'_i \in \Pi_i} \left\{ u_i(\pi'_i, \pi_{-i}^\varepsilon) \mid \begin{array}{l} \pi'_i(\emptyset)(t) \geq w_t \text{ and } \pi'_i(h_t)(\tilde{a}_i) \geq \varepsilon(\tilde{a}_i) \text{ for every } h_t \\ \text{at each } t \in \mathcal{T}_i \text{ for all } \tilde{a}_i \in A_i \times \{\text{pay}, \text{not}\} \end{array} \right\}.$$

If $\pi^* \in \Pi$ satisfies the property that there are sequences $(\varepsilon^n)_{n=1}^\infty$ and $(\pi^n)_{n=1}^\infty$ such that π^n is an ε^n -constrained equilibrium for each $i \in N$ and $\pi^*(h_t) = \lim_{n \rightarrow \infty} \pi^n(h_t)$ for each history h_t , then we say π^* is an **extensive-form trembling-hand equilibrium**.

The result can be formally stated as follows.

Proposition 3. *Let S be a two-player common interest game with the best action profile (a_1^*, a_2^*) . Let \mathcal{T} be such that the set $(t, \sup \mathcal{T}) \cap \mathcal{T}$ is countably infinite for any $t < \sup \mathcal{T}$. Then, there exists $\bar{c} > 0$ such that for all $c < \bar{c}$, in the game with moving-time choice (S, \mathcal{T}, c) , on the path of play of every trembling-hand perfect equilibrium, each agent i plays (a_i^*, not) at any realization of T_i .*

5.4 Horizon Length

As we noted when stating Theorem 1 (Remark 4-5), it is important that the support of the moving times is infinite in at least one direction. We formalize this claim here. The proofs in this section are omitted as they are straightforward. To avoid notational complication, we restrict \mathcal{T} to be a subset of \mathbb{Z} . The results can be readily extended to more general cases.

Proposition 4. *1. For any $t^* \in \mathbb{Z}$, there exist $\mathcal{T} \subseteq \mathbb{Z}$ with $\min_{t \in \mathcal{T}} t = t^*$ and p such that for every two-player common interest game S , there is $\bar{c} > 0$ such that for every $c < \bar{c}$, the dynamic game (S, \mathcal{T}, p, c) has a unique PBE. On the path of the unique PBE, each player i plays (a_i^*, not) at any realization of T_i .*

2. For any $t^ \in \mathbb{Z}$, there exist $\mathcal{T} \subseteq \mathbb{Z}$ with $\max_{t \in \mathcal{T}} t = t^*$ and p such that for every two-player common interest game S , there is $\bar{c} > 0$ such that for every*

$c < \bar{c}$, the dynamic game (S, \mathcal{T}, p, c) has a unique PBE. On the path of the unique PBE, each player i plays (a_i^*, not) at any realization of T_i .

3. For any $t^*, t^{**} \in \mathbb{Z}$, for all $\mathcal{T} \subseteq \mathbb{Z}$ with $\min_{t \in \mathcal{T}} t = t^*$ and $\max_{t \in \mathcal{T}} t = t^{**}$ and for every p , there exists a common interest game S such that there is $\bar{c} > 0$ such that for every $c < \bar{c}$, the dynamic game (S, \mathcal{T}, p, c) has multiple PBE.

Thus, the timing distribution having a minimum alone or a maximum alone is not a problem in equilibrium selection, but having both prevents equilibrium selection. This seeming discontinuity occurs because of the order of limits²⁰: Here we fix a timing distribution and then consider all possible common interest games. If we flip the order of limits, then we retain continuity. The next proposition makes this point clear.

Proposition 5. Consider a family of pairs (\mathcal{T}_K, p_K) defined by $\mathcal{T}_K = \{1, \dots, K\}$ and let p_K be the uniform distribution over \mathcal{T}_K , independent across players.

1. For any $K \in \mathbb{N}$, there exists a two-player common interest game S and $\bar{c} > 0$ such that, for all $c < \bar{c}$, the dynamic game $(S, \mathcal{T}_K, p_K, c)$ has multiple PBE.
2. For any two-player common interest game S , there exists K and $\bar{c} > 0$ such that, for all $c < \bar{c}$, the dynamic game $(S, \mathcal{T}_K, p_K, c)$ has a unique PBE. On the path of the unique PBE, each player i plays (a_i^*, not) at any realization of T_i .

The first part of the proposition is a corollary of the third part of Proposition 4, while the second part shows continuity of the equilibrium actions with respect to the timing distribution.

5.5 Bayes Nash Equilibrium

In the unique PBE for common interest games, even if we did not assume optimality after (observed or unobserved) deviations, deviations would not be optimal. This may suggest uniqueness might be true even under Bayes Nash equilibrium which requires only condition 1 in the definition of PBE. But this is not the case.

²⁰Consider, e.g., a metric $d(\mathcal{T}, \mathcal{T}') = \left| \frac{1}{|\mathcal{T}|+1} - \frac{1}{|\mathcal{T}'|+1} \right|$.

That is, there may exist multiple Bayes Nash equilibria. To see this, consider the component game as in Figure 1 and an independent timing distribution such that \mathcal{T}_1 is the set of odd natural numbers and \mathcal{T}_2 is the set of even natural numbers. By inspection, one can verify that the strategy profile in which each player plays (B , not) under all histories is a Bayes Nash equilibrium. This can be supported by an off-path strategy specification in which each player chooses a different action than the opponent's once the opponent deviates to disclose his action.

The point is that, without off-path optimality, which requires players to best-respond to an observed action, Step 1 of the Proof Sketch for Proposition 1 does not go through. Thus, even though off-path optimality might seem irrelevant if one only looks at the strategy profile used in the unique PBE, it implicitly plays a key role in eliminating inefficient outcomes.

6 Conclusion

This paper studied games with private timing. These games satisfy the often realistic assumption that the timing of moves is private information. We demonstrated that incentives are nontrivial in such a setting. When the component game is a coordination game and players have an option to disclose their actions with a small cost, we proved uniqueness of a perfect Bayesian equilibrium under asynchronicity, uncertainty and sufficient dispersion of moving times. For opposing games, we show that under conditions ensuring certain regularities of the timing distribution, the player who happens to be the first mover plays the action corresponding to her favorable Nash equilibrium and pays the disclosure cost. A number of discussions are provided to further understand and extend those results. Through this analysis, we hope to convey the non-triviality of the way knowledge about timing affects players' behaviors.

There are numerous open questions worth tackling in the context of private timing, both in theory and in applications. For example, we focused on games that satisfy the pure Stackelberg property, but one could consider a wider class of component games. Online Appendix B.4 considers various other component games, such as constant-sum games, games with a dominant action for each player, and a game that is solvable by iterated dominance. One prominent example that we do

not cover is the Cournot quantity-competition game. It is straightforward to show that there exists a PBE in which the first mover plays the Stackelberg action and pays the disclosure cost, but it is not clear if it is the unique equilibrium.²¹

Beyond just examining different component games, there are numerous possibilities for future research in the framework of games with private timing. First, one could investigate the effect of monitoring options. With common interest games with costly monitoring, the best action profile may not be the unique outcome. One can construct examples in which a Pareto-dominated action profile is played and no monitoring takes place. Second, one could consider a cost of secrecy.²² Third, one can consider a setting in which disclosure induces a signal about the action taken, and examine the effect of the noisiness of the signals on the set of equilibrium outcomes. In Online Appendix B.1, we consider a setting where a signal is probabilistically sent, but it is correct whenever it is generated. Another possibility that we do not study is when a signal is always generated but it is possibly incorrect. Fourth, it may be interesting to examine how private information about timing may interact with private information about the payoff functions. Online Appendix B.5 considers a simple case where two players are uncertain which of two possible coordination games is the true one and shows that, in that setting, players pay the disclosure cost. Fifth, the present paper concentrated on the case in which each player moves only once before an action profile is determined. One may want to extend this setting to the case where each player moves more than once.²³

²¹As mentioned in Section 2, Online Appendix B.2 provides two examples that show that multiplicity of PBE may hold in games that do not have the pure Stackelberg property.

²²We thank Drew Fudenberg for suggesting this possibility.

²³In Online Appendix B.6, we ask a different question about multiple moves: What happens if the private-timing game with one action is repeated many times, where after each round, players observe the entire history at that round?

A Appendix

Before we start proofs, let us define the following notation. For $\tilde{t} \in \mathcal{T}_i$, a time- \tilde{t} history of player i is a history $h_i = (N', (t_j, a_j)_{j \in N'}, t) \in \mathcal{H}_i$ with $t = \tilde{t}$. The set of time t -histories of player i is denoted $\mathcal{H}_{i,t}$.

A.1 Proof of Theorem 1

As noted in Remark 4-1, sufficiency of conditions 1 and 2(a) follows from the proof of Theorem 4. So we only prove sufficiency of conditions 1 and 2(b), and necessity here.

Sufficiency of conditions 1 and 2(b):

If each player i plays a_i^* at every opportunity, on- and off-path, incurring the disclosure cost is not a best response. In the proof of Theorem 4, we argue that only a^* is played on the equilibrium path if condition 1 holds. Fix a PBE and suppose that player i chooses an action in $A_i \setminus \{a_i^*\}$ when she receives an opportunity at histories, without observation, at times in a nonempty set $E \subseteq T_i^<$ and plays a_i^* at times without observation in $T_i^< \setminus E$. Note that i plays a_i^* at histories without an observation at times in $\mathcal{T}_i \setminus T_i^<$ because such histories are on the equilibrium path of play. From condition 2(b), for every $\bar{\varepsilon} > 0$ there is a time $\tilde{t} \in \mathcal{T}_{-i}$ such that $\text{Prob}^p(T_i \in E | T_{-i} = \tilde{t}) < \bar{\varepsilon}$. If $\bar{\varepsilon} < \frac{c}{\max_{a \in A} (g_{-i}(a^*) - g_{-i}(a))}$, then at time \tilde{t} , it is strictly suboptimal for player $-i$ to play (\cdot, pay) since the expected payoff from playing such an action is at most $g_{-i}(a^*) - c$, while the expected payoff from playing (a_{-i}^*, not) is at least $g_{-i}(a^*) - \bar{\varepsilon} \cdot (\max_{a \in A} (g_{-i}(a^*) - g_{-i}(a)))$. Thus, conditional on player i having a move at $t \in E$, it is with positive probability that player i does not observe $-i$'s action. Hence, i plays a_i^* in any PBE, and this is a contradiction to the definition of E .

Necessity of condition 1:

Now we show that condition 1 is necessary. Suppose there is t^* such that

$$\min_{i \in \{1,2\}} \text{Prob}^p(T_1 = T_2 = t^* | T_i = t^*) = \tilde{\varepsilon} > 0.$$

Define $\varepsilon_i = \text{Prob}^p(T_{-i} > t^* | T_i = t^*)$ for each $i = 1, 2$, $\tilde{\varepsilon} = \min\{\varepsilon_i | \varepsilon_i > 0, i = 1, 2\}$,²⁴ and $\varepsilon = \min\{\tilde{\varepsilon}, \tilde{\varepsilon}\}$. Consider the game in Figure 4, where M and $\alpha > 0$ are such that $(1 - \varepsilon)1 + \varepsilon(-M) < -\alpha$. At time t^* , if a player assigns probability 1 to the event that the other player chooses B conditional on having a move at t^* , then it is a best response to play B at time t^* as well. We will construct a PBE in which B is played at time t^* .

Let $c < \min\{\varepsilon\alpha, 1\}$. $A_0^i = \emptyset$ for $i \in \{1, 2\}$. We define A_j^i recursively as

$$A_j^i = \left\{ t < t^* \mid -\text{Prob}^p(T_{-i} = t^* | T_i = t, T_{-i} \notin (A_{j-1}^{-i} \cap \{T_{-i} < t\})) M \right. \\ \left. + \text{Prob}^p(T_{-i} \neq t^* | T_i = t, T_{-i} \notin (A_{j-1}^{-i} \cap \{T_{-i} < t\})) < 1 - c \right\}. \quad (1)$$

That is, A_j^i is the set of times before t^* such that it is a best response for player i to play (A, pay) , conditional on no observation, if player $-i$ plays (A, pay) at times in A_{j-1}^{-i} , plays (A, not) at times in $(A_{j-1}^{-i})^c \setminus \{t^*\}$ and (B, \cdot) at time t^* .

Now, define $A^i := \lim_{j \rightarrow \infty} A_j^i$. This limit is well defined because $A_0^i \subseteq A_1^i$ for $i \in \{1, 2\}$ by definition, and $A_{j-1}^{-i} \subseteq A_j^{-i}$ implies $A_j^i \subseteq A_{j+1}^i$ for all $j \in \mathbb{N}$. The latter follows from the fact that $\text{Prob}^p(T_{-i} = t^* | T_i = t, T_{-i} \notin A_{j-1}^{-i} \cap \{T_{-i} < t\})$ is increasing in A_{j-1}^{-i} —if player $-i$'s arrivals are drawn from a smaller set, it is weakly more likely that $-i$'s opportunity occurs at time t^* . Noting that $t^* \notin A^i$ for each $i = 1, 2$, consider the strategy profile $(\sigma_i)_{i=1,2}$ in which, for each player $i = 1, 2$, $t \in \mathcal{T}_i$ and $h_{i,t} \in \mathcal{H}_{i,t}$,

$$\sigma_i(h_{i,t}) := \begin{cases} (A, \text{pay}) & \text{if } t \in A^i \text{ and } h_{i,t} = (\emptyset, \cdot, t) \\ (A, \text{not}) & \text{if } t \in (A^i)^c \setminus \{t^*\} \text{ and } h_{i,t} = (\emptyset, \cdot, t) \\ (B, \text{not}) & \text{if } t = t^*, \varepsilon_i = 0, \text{ and } h_{i,t} = (\emptyset, \cdot, t) \\ (B, \text{pay}) & \text{if } t = t^*, \varepsilon_i > 0, \text{ and } h_{i,t} = (\emptyset, \cdot, t) \\ (A, \text{not}) & \text{if } h_{i,t} = (\{-i\}, (t', A, \text{pay})) \text{ for some } t' \in \mathcal{T}_{-i} \\ (B, \text{not}) & \text{if } h_{i,t} = (\{-i\}, (t', B, \text{pay})) \text{ for some } t' \in \mathcal{T}_{-i} \end{cases},$$

²⁴By convention, the minimum of an empty set is ∞ .

where we abuse notation to express the pure strategy by identifying the action that is assigned probability 1 by σ_i (we abuse notation in the same way in what follows).

At player i 's moving time t , under any event that happens with positive probability under σ , i 's belief is computed by Bayes rule. If the private history of player i is (\emptyset, \cdot, t) and if such a private history is assigned zero probability under σ , then i 's belief assigns probability one to the set of histories in $\{((t, a_i, d_i), (t', B, \text{not})) | a_i \in \{A, B\}, d_i \in \{\text{pay}, \text{not}\}\}$ for some $t' < t$. If the private history of player i is $(\{-i\}, (t', a_{-i}), t)$ for some $t' < t$, then i 's belief assigns probability one to the set of histories in $\{((t, a_i, d_i), (t', a_{-i}, \text{pay})) | a_i \in \{A, B\}, d_i \in \{\text{pay}, \text{not}\}\}$.

Note that player i is best-responding at all times before t^* by the definition of A_j^{-i} and continuity of expected payoffs with respect to probabilities. Also, at time t^* , both players playing B is a best response because M and $\alpha > 0$ are chosen so that $(1 - \varepsilon)1 + \varepsilon(-M) < -\alpha$. For each $i = 1, 2$, an upper bound of the payoff of (B, not) is $-(1 - \tilde{\varepsilon})\alpha$, and a lower bound of the payoff of (B, pay) is $-(1 - \tilde{\varepsilon} - \varepsilon_i)\alpha - c$. Since $c < \varepsilon\alpha \leq \varepsilon_i\alpha$, (B, pay) is a better response for i than (B, not) if the opponent chooses A at every $t > t^*$. At $t > t^*$, by Bayes rule, each player i believes that $-i$ moved at a time in $(A^{-i})^c \setminus \{t^*\}$ and, therefore, played A . Hence, (A, not) is a best response. These facts imply that each player takes a best response at each private history.

Necessity of condition 2:

Suppose, for contradiction, that there are i and a nonempty set $E \subseteq T_i^<$ such that $\inf_{\tilde{t} \in \mathcal{T}_{-i}} \{\text{Prob}^p(T_i \in E | T_{-i} = \tilde{t}) | \text{Prob}^p(T_i \in E | T_{-i} = \tilde{t}) > 0\} = \delta > 0$.

Consider the game shown in Figure 4 with $\alpha = 0$. Define:

$$E_{-i} = \{\tilde{t} \in \mathcal{T}_{-i} | \text{Prob}^p(T_i \in E | T_{-i} = \tilde{t}) > 0\}, \text{ and}$$

$$\tilde{E}_i = \{t \in T_i^< | \text{Prob}^p(T_{-i} \in E_{-i} | T_i = t) = 1\}.$$

Thus, $t \in E_{-i}$ if (i) $t \in \mathcal{T}_{-i}$ and (ii) conditional on $-i$ moving at time t , there is a positive probability that i moves at a time in set E . Notice that E_{-i} is non-empty. Also, $t \in \tilde{E}_i$ if (i) $t \in T_i^<$ and (ii) the probability that $-i$ moves at E_{-i} given that i moves at time t is 1. We will show that for any $M \in [0, \infty)$, we can

find $\bar{c} > 0$ such that if the component game S is the common interest game in Figure 4 and $c < \bar{c}$, then σ , defined as follows, is a PBE of (S, \mathcal{T}, p, c) . For each $t \in \mathcal{T}_{-i}$ and $h_{-i,t} \in \mathcal{H}_{-i,t}$,

$$\sigma_{-i}(h_{-i,t}) := \begin{cases} (A, \text{pay}) & \text{if } t \in E_{-i} \text{ and } h_{-i,t} = (\emptyset, \cdot, t) \\ (A, \text{not}) & \text{if } t \notin E_{-i} \text{ and } h_{-i,t} = (\emptyset, \cdot, t) \\ (A, \text{not}) & \text{if } h_{-i,t} = (\{i\}, (t', A, \text{pay})) \text{ for some } t' \in \mathcal{T}_i \\ (B, \text{not}) & \text{if } h_{-i,t} = (\{i\}, (t', B, \text{pay})) \text{ for some } t' \in \mathcal{T}_i \end{cases},$$

Also, for each $t \in \mathcal{T}_i$ and $h_{i,t} \in \mathcal{H}_{i,t}$,

$$\sigma_i(h_{i,t}) := \begin{cases} (B, \text{not}) & \text{if } t \in E \cup \tilde{E}_i \text{ and } h_{i,t} = (\emptyset, \cdot, t) \\ (A, \text{not}) & \text{if } t \notin E \cup \tilde{E}_i \text{ and } h_{i,t} = (\emptyset, \cdot, t) \\ (A, \text{not}) & \text{if } h_{i,t} = (\{-i\}, (t', A, \text{pay})) \text{ for some } t' \in \mathcal{T}_{-i} \\ (B, \text{not}) & \text{if } h_{i,t} = (\{-i\}, (t', B, \text{pay})) \text{ for some } t' \in \mathcal{T}_{-i} \end{cases}.$$

Now we specify beliefs. First, for each player $j = 1, 2$, if a private history is $(\{-j\}, (t', a_{-j}), t)$ for some time t' , then j assigns probability one to the set of histories $\{((t, a_j, d_j), (t', a_{-j}, \text{pay})) | a_j \in \{A, B\}, d_j \in \{\text{pay}, \text{not}\}\}$. For each private history (\emptyset, \cdot, t) of player $-i$, $-i$'s belief is computed by Bayes rule. Also, except at times in $E \cup \tilde{E}_i$, for each private history of player i , i 's belief is computed by Bayes rule. If the private history is (\emptyset, \cdot, t) and $t \in E \cup \tilde{E}_i$, take an arbitrary element $t^*(t)$ of $\{t' \in \mathcal{T}_{-i} \mid t' < t, p(t', t) > 0\}$. We define player i 's belief at private histories at time $t \in E \cup \tilde{E}_i$ to be a probability distribution over the history that assigns probability 1 to the set of histories $\{((t^*(t), B, \text{not}), (t, a_i, d_i)) \mid a_i \in \{A, B\}, d_i \in \{\text{pay}, \text{not}\}\}$. Thus, in the off-path histories at times in $E \cup \tilde{E}_i$ at which player i does not observe (A, pay) , she believes that $-i$ played (B, not) at time $t^*(t)$.

We now check that each player takes a best response at each private history. First, it is straightforward to check that σ_i and σ_{-i} specify best responses after private histories in which there has been an observation of an action taken by the opponent. In what follows, we consider each player's action after a private history in which there has not been any observation. In the off-path history at times in

$E \cup \tilde{E}_{-i}$ at which player i has not observed A , player i 's belief is that $-i$ played B and, therefore, she best-responds with (B, not) . At all other private histories, player i believes that player $-i$ will play or has played (A, \cdot) and best-responds with (A, not) . For player $-i$, if there has been no observation, the payoff of playing (A, pay) at times in E_{-i} is $1 - c$. At time $t' \in E_{-i}$, the payoff of playing (A, not) is at most $\delta \cdot (-M) + (1 - \delta) \cdot 1$. Thus, for $c \in (0, \min\{1, \delta(M + 1)\})$, $-i$'s best response is to play (A, pay) at all times $\tilde{t} \in E_{-i}$. At every time $\tilde{t} \notin E_{-i}$, choosing (A, not) is a best response for player $-i$ as player i 's strategy and Bayes rule indicate that player $-i$'s belief at such a time must assign probability 1 to the event that player i plays (A, not) if $-i$ chooses (A, not) . \square

A.2 Proof of Theorem 2

For each $c > 0$, fix an arbitrary PBE $\sigma^*(c)$ of the game (S, \mathcal{T}, p, c) , which we know exists from Theorem 3. In what follows, we only consider $c > 0$ such that

$$c < \bar{c} := \min_i (g_i^* - g_i^S) / 2. \quad (2)$$

By Assumption 1 (1), for each $i = 1, 2$, there exists $\tau \in (\inf \mathcal{T}_i, \infty)$ such that for all $c < \bar{c}$ and for all $t \in \mathcal{T}_i \cap (-\infty, \tau) \neq \emptyset$,

$$(1 - p_i(t, -\infty, -\infty, t))g_i^* + p_i(t, -\infty, -\infty, t)g_i - c > g_i^S$$

holds and thus player i plays (a_i^*, \cdot) at any $t \in \mathcal{T}_i \cap (-\infty, \tau) \neq \emptyset$ at any private history (\emptyset, \cdot, t) under $\sigma^*(c)$. Note that Assumption 1 (3b) implies that $\inf \mathcal{T}_1 = \inf \mathcal{T}_2$ so each player $i = 1, 2$ must play (a_i^*, \cdot) early enough in the game.

Let $\bar{t}(c)$ denote the supremum time t over all histories with no observation before which each player i chooses (a_i^*, \cdot) under $\sigma^*(c)$. Let $\bar{\bar{t}}_i(c)$ be the supremum time t before which player i chooses to play (a_i^*, pay) at any private history (\emptyset, \cdot, t) under $\sigma^*(c)$ and let $\bar{\bar{t}}(c) = \min_i \bar{\bar{t}}_i(c)$. Since $\inf \mathcal{T}_1 = \inf \mathcal{T}_2$, the discussion above implies that there is $\tau \in (\inf \mathcal{T}, \infty)$ such that $\tau < \bar{t}(c)$ for every $c < \bar{c}$. Fix an arbitrary choice of such τ and denote it by $\bar{\tau}$.

We consider two cases.

Case 1: Suppose $\bar{t}(c) = \max_i \sup \mathcal{T}_i$. This implies that, for each player i ,

$$p_i(\tilde{t}, t, t, \bar{t}(c)) = 1 \quad (3)$$

for every $t \in \mathbb{R} \cup \{-\infty\}$ and $\tilde{t} \in \mathcal{T}_i$ such that $\text{Prob}^p(T_{-i} \geq t, T_i = \tilde{t}) > 0$. Suppose for contradiction that $\bar{\bar{t}}(c) < \bar{t}(c)$ holds. Assume without loss of generality that $\bar{\bar{t}}_i(c) \leq \bar{\bar{t}}_{-i}(c)$. From Assumption 1 (1) and equation (2), $\exists \bar{t}_i^1 \in \mathcal{T}_i \cap (\bar{\bar{t}}(c), \infty)$ such that $\forall t \in (\bar{\bar{t}}(c), \bar{t}_i^1] \cap \mathcal{T}_i$,

$$g_i^S < g_i^* (1 - p_i(t, \bar{t}(c), \bar{\bar{t}}(c), t)) + \underline{g}_i p_i(t, \bar{t}(c), \bar{\bar{t}}(c), t) - c. \quad (4)$$

Now, an upper bound on the payoff of (a_i^*, not) at every time $t \in \mathcal{T}_i$ is g_i^S because $\bar{t}(c) = \max_i \sup \mathcal{T}_i$ implies the opposing player chooses a_{-i}^* at all moving times.

If $\bar{\bar{t}}_i(c) < \bar{\bar{t}}_{-i}(c)$, then a lower bound on i 's payoff from (a_i^*, pay) at $t \in (\bar{\bar{t}}(c), \min\{\bar{t}_i^1, \bar{\bar{t}}_{-i}(c)\}] \cap \mathcal{T}_i$ is $(g_i^* - c)$, which is greater than the right-hand side of (4). If $\bar{\bar{t}}_i(c) = \bar{\bar{t}}_{-i}(c)$, a lower bound on i 's payoff from (a_i^*, pay) at $t \in (\bar{\bar{t}}(c), \bar{t}_i^1] \cap \mathcal{T}_i$ is the right-hand side of (4). Thus, there exists $\tau > 0$ such that the right-hand side of (4) is a lower bound on the payoff of playing (a_i^*, pay) at $t \in \mathcal{T}_i \cap (\bar{\bar{t}}(c), \bar{\bar{t}}(c) + \tau]$.

Overall, there is $\tau > 0$ such that for all $t \in (\bar{\bar{t}}(c), \bar{\bar{t}}(c) + \tau] \cap \mathcal{T}_i$, each player i would play (a_i^*, pay) , which contradicts the definition of $\bar{\bar{t}}(c)$.

Case 2: Suppose $\bar{t}(c) < \sup \mathcal{T}_i$ for some $i \in \{1, 2\}$.

Step 1: Early enough, (a_i^*, pay) is played. Here we show that $\bar{\bar{t}}(c) \neq -\infty$. Suppose on the contrary that $\bar{\bar{t}}(c) = -\infty$. By Assumption 1 (3b), there exists $\check{t}_i^1 \in \mathbb{R} \cap (-\infty, \bar{\tau})$ such that for all $t \in \mathcal{T}_i \cap (-\infty, \check{t}_i^1]$, we have $p_i(t, -\infty, t, \bar{\tau}) > 0$. Since $p_i(t, -\infty, t, \bar{\tau}) \leq p_i(t, -\infty, t, \bar{t}(c))$ for each t , due to the fact $\bar{\tau} < \bar{t}(c)$, there are $\check{c} > 0$ and $\beta > 0$ such that for all $t \in \mathcal{T}_i \cap (-\infty, \check{t}_i^1]$ and $c \leq \check{c}$,

$$g_i^S p_i(t, -\infty, t, \bar{t}(c)) + (1 - p_i(t, -\infty, t, \bar{t}(c)))g_i^* < g_i^* - c - \beta.$$

By Assumption 1 (1), for any $\beta > 0$, there is $\check{t}_i^2 \in \mathcal{T}_i$ such that $\forall t \in (-\infty, \check{t}_i^2] \cap \mathcal{T}_i$,

$$(g_i^* - \underline{g}_i) p_i(t, -\infty, -\infty, t) < \beta.$$

The above two inequalities imply that for all $t \in (-\infty, \min\{\check{t}_i^1, \check{t}_i^2\}] \cap \mathcal{T}_i$,

$$g_i^S p_i(t, -\infty, t, \bar{t}(c)) + (1 - p_i(t, -\infty, t, \bar{t}(c))) g_i^* < g_i^* (1 - p_i(t, -\infty, -\infty, t)) + \underline{g}_i p_i(t, -\infty, -\infty, t) - c.$$

As in Case 1, the left-hand side of the previous expression is an upper bound on the payoff from (a_i^*, not) while the right-hand side is a lower bound on the payoff from (a_i^*, pay) . Thus, there is τ such that for all $t \in (-\infty, \tau]$, each player i chooses (a_i^*, pay) , which contradicts $\bar{t}(c) = -\infty$.

Step 2: Defining k , $\hat{t}(c)$, and $a(c)$ and showing $\bar{\bar{t}}(c) < \bar{t}(c)$. Let player k be such that for each $t_{-k} \in \mathcal{T}_{-k}$ such that $\sigma_{-k}(h_{-k})(a_{-k}^*, \cdot) < 1$ with $h_{-k} = \emptyset$ where h_{-k} is the private history at time t_{-k} , we can find $t_k \leq t_{-k}$ such that $\sigma_k(h_k)(a_k^*, \cdot) < 1$ with $h_k = \emptyset$ where h_k is the private history at time t_k (If both players satisfy such a condition, let $k = 1$). In other words, k is the first player to play an action other than a_k^* under some history without observation. Define

$$\hat{t}(c) := \inf\{t \in \mathcal{T}_k \cap [\bar{t}(c), \infty) \mid (1 - p_k(t, \bar{\bar{t}}(c), \bar{\bar{t}}(c), t))(g_k^* - c) + p_k(t, \bar{\bar{t}}(c), \bar{\bar{t}}(c), t)(\underline{g}_k - c) \leq g_k^S\}.$$

The left-hand side of the inequality of the previous expression is a lower bound on the payoff from (a_k^*, \cdot) . Note that $\bar{t}(c) \geq \hat{t}(c)$, since $\hat{t}(c)$ is the earliest possible time at which player k could choose an action other than a_k^* under $\sigma^*(c)$. Notice that by Assumption 1 (1), $\lim_{t \rightarrow \bar{\bar{t}}(c), t \in \mathcal{T}_k} p_k(t, \bar{\bar{t}}(c), \bar{\bar{t}}(c), t) = 0$ which implies $\hat{t}(c) > \bar{\bar{t}}(c)$. This shows that $\bar{\bar{t}}(c) < \bar{t}(c)$. Since Assumption 1 (3a) implies that $p_k(t, \bar{\bar{t}}(c), \bar{\bar{t}}(c), t)$ is non-decreasing in t , for all $t \in \mathcal{T}_k \cap [\hat{t}(c), \infty)$, we have $p_k(t, \bar{\bar{t}}(c), \bar{\bar{t}}(c), t) \geq \frac{g_k^* - c - g_k^S}{g_k^* - g_k} \equiv a(c)$. The function $a(c)$ is decreasing in c . Therefore, there are $\tilde{c} > 0$ and $\gamma > 0$ such that for $c < \tilde{c}$, $a(c) > \gamma$. Fix an arbitrary choice of such a pair (\tilde{c}, γ) .

Let $j \in \{k, -k\}$ be such that $\bar{\bar{t}}_{-j}(c) \geq \bar{\bar{t}}_j(c)$.

Step 3. Case a). Suppose $\bar{t}_{-j}(c) > \bar{t}_j(c)$.

We start by fixing $\varepsilon, \tilde{\varepsilon} \in (0, \alpha \cdot \gamma/2)$, where α is an arbitrary choice of “ α ” that satisfies the condition in Assumption 1 (2).

Note that

$$p_j(t, t, t, \hat{t}(c))g_j^S + (1 - p_j(t, t, t, \hat{t}(c)))g_j^*$$

is an upper bound on the payoff from playing (a_j^*, not) at $t \in [\bar{t}_j(c), \min\{\bar{t}_{-j}(c), \hat{t}(c)\}] \cap \mathcal{T}_j$ and $g_j^* - c$ is the payoff from playing (a_j^*, pay) instead. By the definition of $\bar{t}(c)$, for every $\tau \in (\bar{t}_j(c), \bar{t}_{-j}(c)]$, there is $\tilde{t}_j^\tau \in \mathcal{T}_j \cap [\bar{t}(c), \min\{\tau, \hat{t}(c)\}]$ such that the former is no less than the latter, i.e.,

$$p_j(\tilde{t}_j^\tau, \tilde{t}_j^\tau, \tilde{t}_j^\tau, \hat{t}(c))g_j^S + (1 - p_j(\tilde{t}_j^\tau, \tilde{t}_j^\tau, \tilde{t}_j^\tau, \hat{t}(c)))g_j^* \geq g_j^* - c. \quad (5)$$

This is because, otherwise, player j would play (a_j^*, pay) instead of (a_j^*, not) at every $t \in \mathcal{T}_j \cap [\bar{t}(c), \min\{\tau, \hat{t}(c)\}]$, contradicting the definition of $\bar{t}(c)$. In what follows, we will draw a contradiction to this inequality. Notice that for every $t_j \in \mathcal{T}_j \cap [\bar{t}(c), \infty)$, we have

$$p_j(t_j, t_j, t_j, \hat{t}(c)) \geq p_j(t_j, \bar{t}(c), t_j, \hat{t}(c)) \geq p_j(t_j, \bar{t}(c), \bar{t}(c), \hat{t}(c)) - p_j(t_j, \bar{t}(c), \bar{t}(c), t_j).$$

Now, from Assumption 1 (1), $\exists \tau^\varepsilon \in \mathcal{T}_j \cap (\bar{t}(c), \infty)$ such that

$$p_j(t_j, \bar{t}(c), \bar{t}(c), t_j) < \tilde{\varepsilon} \quad \forall t_j \in [\bar{t}(c), \tau^\varepsilon] \cap \mathcal{T}_j,$$

which yields

$$p_j(t_j, t_j, t_j, \hat{t}(c)) \geq p_j(t_j, \bar{t}(c), \bar{t}(c), \hat{t}(c)) - \tilde{\varepsilon} \quad \forall t_j \in [\bar{t}(c), \tau^\varepsilon] \cap \mathcal{T}_j.$$

Fix $\tilde{t}_j^{\tau^\varepsilon} \in \mathcal{T}_j \cap [\bar{t}(c), \min\{\tau^\varepsilon, \hat{t}(c)\}]$ satisfying equation (5).

By the definition of $\hat{t}(c)$ and the continuity of the probability with respect to decreasing sets, there is $t_k^\varepsilon \in \mathcal{T}_k \cap [\hat{t}(c), \infty)$ such that $p_j(\tilde{t}_j^{\tau^\varepsilon}, \bar{t}(c), \bar{t}(c), \hat{t}(c)) \geq p_j(\tilde{t}_j^{\tau^\varepsilon}, \bar{t}(c), \bar{t}(c), t_k^\varepsilon) - \varepsilon$.

By Assumption 1 (2), there exists $\alpha > 0$ such that $p_j(\tilde{t}_j^{\tau^\varepsilon}, \bar{t}(c), \bar{t}(c), t_k^\varepsilon) \geq$

$\alpha p_k(t_k^\varepsilon, \bar{t}(c), \bar{t}(c), t_k^\varepsilon)$, which yields

$$p_j(\tilde{t}_j^{\varepsilon}, \tilde{t}_j^{\varepsilon}, \tilde{t}_j^{\varepsilon}, \hat{t}(c)) \geq p_j(\tilde{t}_j^{\varepsilon}, \bar{t}(c), \bar{t}(c), t_k^\varepsilon) - \varepsilon - \tilde{\varepsilon} \geq \alpha p_k(t_k^\varepsilon, \bar{t}(c), \bar{t}(c), t_k^\varepsilon) - \varepsilon - \tilde{\varepsilon} \geq \alpha a(c) - \varepsilon - \tilde{\varepsilon}.$$

However, these inequalities imply that for all $c < \tilde{c}$,

$$\begin{aligned} & p_j(\tilde{t}_j^{\varepsilon}, \tilde{t}_j^{\varepsilon}, \tilde{t}_j^{\varepsilon}, \hat{t}(c)) g_j^S + (1 - p_j(\tilde{t}_j^{\varepsilon}, \tilde{t}_j^{\varepsilon}, \tilde{t}_j^{\varepsilon}, \hat{t}(c))) g_j^* \leq (\alpha a(c) - \varepsilon - \tilde{\varepsilon}) g_j^S + \\ & + (1 - (\alpha a(c) - \varepsilon - \tilde{\varepsilon})) g_j^* < (\alpha \gamma - \varepsilon - \tilde{\varepsilon}) g_j^S + (1 - (\alpha \gamma - \varepsilon - \tilde{\varepsilon})) g_j^* < g_j^* - c. \end{aligned}$$

This contradicts (5).

Step 3. Case b). Suppose $\bar{t}_j(c) = \bar{t}_{-j}(c)$. By the definition of $\bar{t}(c)$, for every $\tau \in (\bar{t}, \hat{t}(c)]$, there is $\tilde{t}_j^\tau \in \mathcal{T}_j \cap [\bar{t}(c), \tau]$ such that²⁵

$$\begin{aligned} & p_j(\tilde{t}_j^\tau, \bar{t}(c), \bar{t}(c), \hat{t}(c)) g_j^S + (1 - p_j(\tilde{t}_j^\tau, \bar{t}(c), \bar{t}(c), \hat{t}(c))) g_j^* \geq \quad (6) \\ & g_j^* (1 - p_j(\tilde{t}_j^\tau, \bar{t}(c), \bar{t}(c), \tilde{t}_j^\tau)) + \underline{g}_j p_j(\tilde{t}_j^\tau, \bar{t}(c), \bar{t}(c), \tilde{t}_j^\tau) - c. \end{aligned}$$

To see this, note that the left-hand side is an upper bound on the payoff of (a_j^*, not) at time \tilde{t}_j^τ . The right-hand side is a lower bound on the payoff from (a_j^*, pay) . In what follows, we will draw a contradiction to this inequality.

By the definition of $\hat{t}(c)$ and the continuity of the probability with respect to decreasing sets, for every $\tilde{\varepsilon} > 0$, there is $t_k^{\tau, \tilde{\varepsilon}} \in \mathcal{T}_k \cap [\hat{t}(c), \infty)$ such that $p_j(\tilde{t}_j^\tau, \bar{t}(c), \bar{t}(c), \hat{t}(c)) \geq p_j(\tilde{t}_j^\tau, \bar{t}(c), \bar{t}(c), t_k^{\tau, \tilde{\varepsilon}}) - \tilde{\varepsilon}$. By Assumption 1 (2), there exists $\alpha > 0$ such that $p_j(\tilde{t}_j^\tau, \bar{t}(c), \bar{t}(c), t_k^{\tau, \tilde{\varepsilon}}) \geq \alpha p_k(t_k^{\tau, \tilde{\varepsilon}}, \bar{t}(c), \bar{t}(c), t_k^{\tau, \tilde{\varepsilon}})$ which yields $p_j(\tilde{t}_j^\tau, \bar{t}(c), \bar{t}(c), \hat{t}(c)) \geq \alpha p_k(t_k^{\tau, \tilde{\varepsilon}}, \bar{t}(c), \bar{t}(c), t_k^{\tau, \tilde{\varepsilon}}) - \tilde{\varepsilon} \geq \alpha a(c) - \tilde{\varepsilon}$. Since the previous inequality holds for every $\tilde{\varepsilon} > 0$, we obtain that for every $\tau < \tau^1$, $p_j(\tilde{t}_j^\tau, \bar{t}(c), \bar{t}(c), \hat{t}(c)) \geq \alpha a(c)$.

By Assumption 1 (1), for every $\varepsilon > 0$, $\exists \tau^\varepsilon \in \mathcal{T}_j \cap (\bar{t}(c), \min\{\tau, \hat{t}(c)\})$ such that $\forall t \in [\bar{t}(c), \tau^\varepsilon] \cap \mathcal{T}_j$, $p_j(t, \bar{t}(c), \bar{t}(c), t) < \varepsilon$.

These observations imply that for any $\varepsilon > 0$, there are $\tau^\varepsilon \in (\bar{t}(c), \hat{t}(c)]$ and $\tilde{t}_j^{\tau^\varepsilon} \in \mathcal{T}_j \cap [\bar{t}(c), \tau^\varepsilon]$ such that the latter satisfies (6), $p_j(\tilde{t}_j^{\tau^\varepsilon}, \bar{t}(c), \bar{t}(c), \hat{t}(c)) \geq \alpha a(c)$, and $p_j(\tilde{t}_j^{\tau^\varepsilon}, \bar{t}(c), \bar{t}(c), \tilde{t}_j^{\tau^\varepsilon}) < \varepsilon$.

²⁵Clearly, \tilde{t}_j^τ and many of the variables we define below depend on c as well. We omit the dependence on c for ease of notation.

However, for every $c < \tilde{c}$ and $\varepsilon < \alpha\gamma$, we obtain

$$\begin{aligned} p_j(\tilde{t}_j^\varepsilon, \bar{t}(c), \bar{t}(c), \hat{t}(c))g_j^S + (1 - p_j(\tilde{t}_j^\varepsilon, \bar{t}(c), \bar{t}(c), \hat{t}(c)))g_j^* &\leq \alpha a(c)g_j^S + \\ &+ (1 - \alpha a(c))g_j^* < (g_j^* - c)(1 - \varepsilon) + \underline{g}_j\varepsilon \leq \\ g_j^*(1 - p_j(\tilde{t}_j^\varepsilon, \bar{t}(c), \bar{t}(c), \tilde{t}_j^\varepsilon)) + \underline{g}_j p_j(\tilde{t}_j^\varepsilon, \bar{t}(c), \bar{t}(c), \tilde{t}_j^\varepsilon) - c. \end{aligned}$$

This contradicts (6).

A.3 Proof of Theorem 3

As explained in the main text of the paper, the proof consists of five parts. In this proof, to avoid confusion with the index $N \in \mathbb{N}$ for the N 'th approximating game, we denote by $\{1, \dots, I\}$ the set of players.

Part 1: Convergent sequence

Lemma 6. *For every sequence of strategy profiles $\{\sigma^n\}_{n \in \mathbb{N}}$ with $\sigma^n \in \Sigma$ for each $n \in \mathbb{N}$, there exist $\sigma \in \Sigma$ and a convergent subsequence $\{\sigma^{n_k}\}_{k \in \mathbb{N}}$ of $\{\sigma^n\}_{n \in \mathbb{N}}$ such that $\sigma^{n_k} \rightarrow \sigma$ pointwise as $k \rightarrow \infty$.*

Proof. Let $\tilde{A} = \times_{i=1}^n \tilde{A}_i$, where $\tilde{A}_i = A_i \times \{\text{pay, not}\}$. Define $\tilde{\mathcal{T}} = \{\mathbf{t} \in \mathcal{T}^I \mid p(\mathbf{t}) > 0\}$. Since \mathcal{T} is countable and A is finite, $\tilde{\mathcal{T}} \times \tilde{A}$ is also countable.

We will show that there exist σ and a subsequence $\{\sigma^{n_k}\}_{n_k \in \mathbb{N}}$ of $\{\sigma^n\}_{n \in \mathbb{N}}$ such that, for every player i , private history $h_{i,t} \in \mathcal{H}_{i,t}$ and $\tilde{a}_i \in \tilde{A}_i$, we have

$$\sigma_i^{n_k}(\tilde{a}_i \mid h_{i,t}) \rightarrow \sigma_i(\tilde{a}_i \mid h_{i,t})$$

as $k \rightarrow \infty$.

Since $\tilde{\mathcal{T}} \times \tilde{A}$ is countable, there is an ordering of elements in it, denoted $\tilde{\mathcal{T}} \times \tilde{A} = \{x_k\}_{k \in \mathbb{N}}$. For each k , let $x_k = ((t_j^{x_k}, a_j^{x_k}, d_j^{x_k})_{j \in \{1, \dots, I\}})$ where $t_j^{x_k}$, $a_j^{x_k}$, and $d_j^{x_k}$ denote player j 's moving time, action, and payment decision, respectively, under x_k . Define $N_i^{x_k} = \{j \in \{1, \dots, I\} \mid d_j^{x_k} = \text{pay and } t_j^{x_k} < t_i^{x_k}\}$ and let $h_i(x_k) = (N_i^{x_k}, (t_j^{x_k}, a_j^{x_k}, d_j^{x_k})_{j \in N_i^{x_k}}) \in \mathcal{H}_{i,t_i^{x_k}}$. Define f_i^n for each $i \in \{1, \dots, I\}$ and $n \in \mathbb{N}$

as

$$f_i^n(x_k) := \sigma_i^n(\tilde{a}_i^{x_k} | h_i(x_k))$$

where $\tilde{a}_i^{x_k} = (a_i^{x_k}, d_i^{x_k})$ for each $x_k \in \tilde{\mathcal{T}} \times \tilde{A}$.

For x_1 , there exists a subsequence $\{\sigma_i^{n_k^1}\}_{k \in \mathbb{N}}$ of $\{\sigma_i^n\}_{n \in \mathbb{N}}$ such that

$$f_i^{n_k^1}(x_1) \rightarrow_{k \rightarrow \infty} f_i(x_1)$$

for some value of the limit, $f_i(x_1)$, for each $i \in \{1, \dots, I\}$. Now, recursively, given a sequence $\{\sigma_i^{n_k^m}\}_{k \in \mathbb{N}}$, there exists its subsequence $\{\sigma_i^{n_k^{m+1}}\}_{k \in \mathbb{N}}$ such that

$$f_i^{n_k^{m+1}}(x_{m+1}) \rightarrow_{k \rightarrow \infty} f_i(x_{m+1})$$

for some value of the limit, $f_i(x_{m+1}) \in [0, 1]$.

Now, consider the sequence $\{\sigma_i^{n_k^k}\}_{k=1}^\infty$. To see that this sequence has a limit, note that for each $\tilde{k} < \infty$, we must have

$$f_i^{n_k^k}(x_{\tilde{k}}) \rightarrow_{k \rightarrow \infty} f_i(x_{\tilde{k}})$$

because for each $k \geq \tilde{k}$, $\{f_i^{n_k^k}(x_{\tilde{k}})\}_{k=\tilde{k}}^\infty$ is a subsequence of $\{f_i^{n_k^k}(x_{\tilde{k}})\}_{k=\tilde{k}}^\infty$.

By the definition of the f_i^n function, this implies that

$$\sigma_i^{n_k^k}(\tilde{a}_i^{x_{\tilde{k}}} | h_i(x_{\tilde{k}})) \rightarrow_{k \rightarrow \infty} f_i(x_{\tilde{k}}).$$

Now, define σ_i by $\sigma_i(\tilde{a}_i^{x_{\tilde{k}}} | h_i(x_{\tilde{k}})) = f_i(x_{\tilde{k}})$ for each $x_{\tilde{k}}$. Since $\bigcup_{\tilde{k}=1}^\infty \{(\tilde{a}_i^{x_{\tilde{k}}}, h_i(x_{\tilde{k}}))\} = \tilde{A}_i \times \mathcal{H}_i$ and $\sum_{\tilde{a}_i^{x_{\tilde{k}}} \in \tilde{A}_i} \sigma_i^{n_k^k}(\tilde{a}_i^{x_{\tilde{k}}} | h_i(x_{\tilde{k}})) = 1$ for each $\tilde{k} \in \mathbb{N}$ and $k \geq \tilde{k}$, we must have $\sigma_i \in \Sigma_i$. This completes the proof. \square

Part 2: Finite approximating games

We now define a sequence of finite games that approximates the original game. Let $N \in \mathbb{N}$. If \mathcal{T} is finite, then the standard fixed-point argument shows that there exists a PBE. Hence we consider the case in which \mathcal{T} is infinite. Since \mathcal{T} is countable, we can write $\mathcal{T} = \{t_k\}_{k \in \mathbb{Z}}$.²⁶ Let $\mathcal{T}^N := \{t_k\}_{k=-N}^N$.

²⁶Note that the times are not necessarily ordered, i.e., t_k may not be monotonic in k .

Define

$$\tilde{p}^N = \text{Prob}^p(t^i \in \mathcal{T}^N \text{ for every } i \in \{1, \dots, I\}),$$

the probability that the moving times of all players are in \mathcal{T}^N .

Let \underline{N} be the smallest integer N' such that $\tilde{p}^N > 0$ for all $N \geq N'$. By the definition of \tilde{p}^N , every player gets a move in \mathcal{T}^N with positive probability. Also, by the definition of $\{t_k\}_{k \in \mathbb{Z}}$, $\underline{N} < \infty$ holds, and $\tilde{p}^N \rightarrow 1$ as $N \rightarrow \infty$.

For each $N \geq \underline{N}$, define the distribution of $\mathbf{t} \in \times_{i \in \{1, \dots, I\}} \mathcal{T}^N$ in the N 'th approximating game which we denote p^N as

$$p^N(\mathbf{t}) = \frac{p(\mathbf{t})}{\tilde{p}^N},$$

for all $\mathbf{t} \in \times_{i \in \{1, \dots, I\}} \mathcal{T}^N$.

For each $N \geq \underline{N}$, we define the N 'th approximating game of the private-timing game $\Gamma = (S, \mathcal{T}, p, c)$ as the triple $\Gamma^N = (S, \mathcal{T}^N, p^N, c)$.

A private history of player i at time $t \in \mathcal{T}^N$, $h_{i,t} \in \mathcal{H}_{i,t}$, is said to be *feasible* in the N 'th approximating game if there exists a history of play $h = (t_i, a_i, d_i)_{i \in \{1, \dots, I\}}$ that is compatible with $h_{i,t}$ such that (i) $t_j \in \mathcal{T}^N$ for $j \in \{1, \dots, I\}$, and (ii) $\text{Prob}^p(T_i = t_i \forall i \in \{1, \dots, I\} | T_j = t_j \forall j \in \{1, \dots, I\} \text{ s.t. } d_j = \text{pay and } t_j < t_i) > 0$.

Let $\mathcal{H}_{i,t}^N$ denote the set of i 's private histories that are feasible in the N 'th approximating game at time $t \in \mathcal{T}^N$. Each player i 's strategy is defined as a function from $\bigcup_{t \in \mathcal{T}^N} \mathcal{H}_{i,t}^N$ to $\Delta(\tilde{A}_i)$. The set of strategies of player i in the N 'th approximating game is denoted Σ_i^N . Let $\Sigma^N = \times_{i \in I} \Sigma_i^N$. Player i 's expected payoff in the N 'th approximating game from strategy profile σ^N is denoted $u_i^N(\sigma^N)$.

Part 3: Defining an ε -constrained strategies and equilibrium

Definition 6. For a given game $\Gamma = (S, \mathcal{T}, p, c)$, an ε -constrained strategy $\sigma_i \in \Sigma_i$ is a strategy such that, there are $\nu \in (0, \varepsilon)$ and $\hat{\varepsilon} : \mathcal{H}_i \rightarrow [\nu, \varepsilon]$ such that for each pair $(a_i, d_i) \in A_i \times \{\text{pay, not}\}$, $\sigma_i(h_{i,t})(a_i, d_i) \geq \hat{\varepsilon}(a_i, d_i, h_{i,t})$ for every $h_{i,t} \in \mathcal{H}_{i,t}$ and $t \in \mathcal{T}_i$.

Definition 7. For a given game $\Gamma = (S, \mathcal{T}, p, c)$, an ε -constrained equilibrium is a strategy profile $\sigma \in \Sigma$ such that for each i , σ_i is a best response to σ_{-i} among

ε -constrained strategies.

Definition 8 (Trembling-Hand Perfect Equilibrium). For a given game $\Gamma = (S, \mathcal{T}, p, c)$, a strategy profile $\sigma^* \in \Sigma$ is a **trembling-hand perfect equilibrium** if, for each i , there is a sequence of ε_n -constrained equilibria σ^{ε_n} with $\varepsilon_n \rightarrow 0$ as $n \rightarrow \infty$ such that $\sigma^{\varepsilon_n} \rightarrow \sigma^*$ pointwise.

Part 4: Existence of ε -constrained equilibrium

For every $N \geq \underline{N}$, the N 'th approximating game, as it is finite, has an ε -constrained equilibrium.

Proposition 7. *In every game $\Gamma = (S, \mathcal{T}, p, c)$ there exists $\bar{\varepsilon} > 0$ such that for all $\varepsilon \in (0, \bar{\varepsilon})$ there exists an ε -constrained equilibrium in Γ .*

Proof. Let $\frac{1}{2 \cdot \max_{i \in \{1, \dots, I\}} |A_i|} > \bar{\varepsilon} > 0$. Such $\bar{\varepsilon}$ is small enough that for every $\varepsilon \in (0, \bar{\varepsilon})$, there exists a sequence $\{\sigma^N\}_{N=\underline{N}}^\infty$ with $\sigma^N \in \Sigma^N$ for each N of ε -constrained equilibria in Γ^N . Such a sequence exists as each Γ^N is finite. Fix $\varepsilon \in (0, \bar{\varepsilon})$ and let $\{\sigma^N\}_{N=\underline{N}}^\infty$ be a sequence of ε -constrained equilibria, such that $\sigma_i(h_{i,t})(a_i, d_i) \geq \varepsilon(a_i, d_i, h_{i,t}) \geq \nu$ for fixed $\nu \in (0, \varepsilon)$.

We define the strategy profile $\tilde{\sigma}^N \in \Sigma$ of Γ *corresponding* to a strategy profile $\sigma^N \in \Sigma^N$ of Γ^N in the following manner. The following conditions hold for each $i \in \{1, \dots, I\}$ and $t \in \mathcal{T}$. (i) At a history $h_{i,t} \in \mathcal{H}_{i,t}^N$, $\tilde{\sigma}_i^N(h_{i,t}) = \sigma_i^N(h_{i,t})$. (ii) At a history $h_{i,t} \in \mathcal{H}_{i,t} \setminus \mathcal{H}_{i,t}^N$ we set $\tilde{\sigma}_i^N(h_{i,t})(a_i, d_i) = 1/(2|A_i|)$ for each $(a_i, d_i) \in \tilde{A}_i$.

For each $N \geq \underline{N}$, define $\tilde{\sigma}^N$ to be the strategy profile in the original game corresponding to σ^N . By Lemma 6, there exists $\tilde{\sigma} \in \Sigma$ such that the sequence $\{\tilde{\sigma}^N\}_{N=\underline{N}}^\infty$ has a convergent sequence converging to $\tilde{\sigma}$.

We next show that $\tilde{\sigma}$ is an ε -constrained equilibrium of the original game. By contradiction, assume it is not an ε -constrained equilibrium. Then, there exist $\delta > 0$, a player i and a ε -constrained strategy σ'_i such that

$$u_i(\sigma'_i, \tilde{\sigma}_{-i}) \geq u_i(\tilde{\sigma}_i, \tilde{\sigma}_{-i}) + \delta.$$

For each $N \geq \underline{N}$, define $\sigma_i'^N \in \Sigma^N$ of the N 'th approximating game by restricting attention to the relevant set of private histories, i.e., $\sigma_i'^N(h_{i,t}) := \sigma_i'(h_{i,t})$ for

each history $h_{i,t} \in \mathcal{H}_{i,t}^N$. Note that $\sigma_i^{\prime N}$ is an ε -constrained strategy because σ_i^{\prime} is an ε -constrained strategy.

Since $\tilde{p}^N \rightarrow 1$ as $N \rightarrow \infty$, there exists $\bar{N} \in [\underline{N}, \infty)$ such that for all $N \geq \bar{N}$,

$$u_i^N(\sigma_i^{\prime N}, \sigma_{-i}^N) \geq u_i^N(\sigma_i^N, \sigma_{-i}^N) + \frac{\delta}{2}$$

holds. However, this contradicts the assumption that σ^N is an ε -constrained equilibrium of the N 'th approximating game. \square

Part 5: Existence of a Trembling-hand perfect equilibrium

Proposition 8. *In every game $\Gamma = (S, \mathcal{T}, p, c)$, a trembling-hand perfect equilibrium exists.*

Proof. Fix a sequence $\{\varepsilon_n\}_{n \in \mathbb{N}}$ such that $\varepsilon_n > 0$ for each $n \in \mathbb{N}$ and $\varepsilon_n \rightarrow 0$ as $n \rightarrow \infty$. Proposition 7 implies that, for each $n \geq \underline{N}$, there exists an ε_n -constrained equilibrium $\sigma^{\varepsilon_n} \in \Sigma$. Lemma 6 then implies that there must be a convergent subsequence of the sequence of ε_n -constrained equilibria, $\{\sigma^{\varepsilon_n}\}_{n=\underline{N}}^\infty$. By the definition of trembling-hand perfect equilibrium, the limit of the subsequence will be a trembling-hand perfect equilibrium of the original game. \square

Since a trembling-hand perfect equilibrium is a PBE, we have the following result.

Corollary 9. *In every game $\Gamma = (S, \mathcal{T}, p, c)$, a PBE exists.*

\square

A.4 Proof of Theorem 4

Step 1:

Step 1-1: Fix a common interest game that is $(s_i)_{i \in N}$ -common, $\varepsilon \in (0, \min_{i \in N} s_i)$, and a timing structure p that is $(1 + \varepsilon - \min_{i \in N} s_i)$ -dispersed. Fix a PBE and take $c \in (0, \min_{i \in N} \left[\varepsilon(g_i^* - \underline{g}_i) \right])$. Notice that, by the definition of s_i , $(1 - (s_i - \varepsilon))g_i^* + (s_i - \varepsilon)\underline{g}_i - g_i^S = \varepsilon(g_i^* - \underline{g}_i) > 0$. Let $N_i(a^*) \subseteq \mathcal{T}_i$ be the set of times t such that there exists a history under which the fixed PBE designates a probability distribution

over player i 's actions at t that assigns strictly positive probability to an action that is not a_i^* . For contradiction, we suppose that $N_i(a^*)$ is nonempty for some $i \in N$. Let $t^* := \inf_{t \in \bigcup_{i \in N} N_i(a^*)} t$. Any player i who moves at time t^* must choose a_i^* . In fact, the probability that any opponent j chooses an action other than a_j^* before time t^* is zero. Therefore, a lower bound on i 's payoff from (a_i^*, pay) at time t^* is $g_i^S + \varepsilon(g_i^* - \underline{g}_i) - c$ while an upper bound on the payoff from $a'_i \neq a_i^*$ is g_i^S , which is strictly smaller because $\varepsilon(g_i^* - \underline{g}_i) > c$.

Step 1-2: By the definition of q -dispersion, there must exist $i \in N$ and $t' > t^*$ such that for $j \neq i$ and $t \in (t^*, t'] \cap \mathcal{T}_i$, $\mathbb{P}(t^* < T_j \leq t | T_i = t) < s_i - \varepsilon$. Our choice of c implies that (a_i^*, pay) would give such i a strictly higher payoff than playing any action other than a_i^* . Thus, i would not take an action different from a_i^* at any time in $(t^*, t']$. But then, for any $t \in (t^*, t'] \cap \mathcal{T}_j$, j 's payoff from (a_j^*, pay) is $g_j^* - c$, which is strictly greater than the best feasible payoff from any other action, which is g_j^S . Thus, $(t^*, t'] \cap N_i(a^*) = (t^*, t'] \cap N_j(a^*) = \emptyset$. This contradicts the definition of t^* . Hence, $N_i(a^*)$ is empty for each i . Notice that if p has dispersed potential moves, it is q -dispersed for every q and therefore, the previous arguments go through for any profile of $(s_i)_{i \in N}$. Hence, this establishes the sufficiency of conditions 1 and 2(a) of Theorem 1.

Step 2:

Assume now the potential leader condition. Suppose for contradiction that under the fixed PBE that we denote here by σ^* , there exist t and i such that there is a positive ex-ante probability with which i pays the disclosure cost at t . Player i 's payoff from such σ^* is strictly less than g_i^* . But consider i 's deviation to playing (a_i^*, not) with probability 1 at all the information sets at time t that can be reached with positive probability under σ^* , while no change is made to the distribution of actions conditional on other histories. Call this strategy σ'_i . Then, for any $j \neq i$ and any realization of $T_j \in \mathcal{T}_j$, j is at an information set that can be reached with positive probability under σ^* , so plays (a_j^*, \cdot) . Hence $(\sigma'_i, \sigma_{-i}^*)$ must assign probability one to a^* . Hence the payoff from $(\sigma'_i, \sigma_{-i}^*)$ is g_i^* , so the deviation is profitable. This is a contradiction to the assumption that σ^* is a PBE. Therefore, there is no time at which any player pays the disclosure cost.

References

- BENOIT, J.-P. AND V. KRISHNA (1985): “Finitely Repeated Games,” *Econometrica*, 53, 905–22.
- CALCAGNO, R., Y. KAMADA, S. LOVO, AND T. SUGAYA (2014): “Asynchronicity and Coordination in Common and Opposing Interest Games,” *Theoretical Economics*, 9, 409–434.
- CARUANA, G. AND L. EINAV (2008): “A Theory of Endogenous Commitment,” *Review of Economic Studies*, 75, 99–116.
- DUTTA, P. K. (1995): “A Folk Theorem for Stochastic Games,” *Journal of Economic Theory*, 66, 1–32.
- ISHII, Y. AND Y. KAMADA (2011): “The Effect of Correlated Inertia on Coordination,” Mimeo.
- KREPS, D. M. AND G. RAMEY (1987): “Structural Consistency, Consistency, and Sequential Rationality,” *Econometrica*, 55, 1331–1348.
- LAGUNOFF, R. AND A. MATSUI (1997): “Asynchronous Choice in Repeated Coordination Games,” *Econometrica*, 65, 1467–1477.
- (2001): “Are “Anti-Folk Theorems” in repeated games nongeneric?” *Review of Economic Design*, 6, 297–412.
- MAILATH, G. J. AND L. SAMUELSON (2006): *Repeated Games and Reputations: Long-Run Relationships*, New York : Oxford University Press.
- MATSUI, A. (1989): “Information Leakage Forces Cooperation,” *Games and Economic Behavior*, 1, 94–115.
- MORONI, S. (2015): “Existence of trembling hand equilibrium in revision games with imperfect information,” Mimeo.
- OSTROVSKY, M. AND M. SCHWARZ (2005): “Adoption of Standards under Uncertainty,” *The RAND Journal of Economics*, 36, 816–832.

——— (2006): “Synchronization under uncertainty,” *International Journal of Economic Theory*, 2, 1–16.

PARK, A. AND L. SMITH (2008): “Caller Number Five and related timing games,” *Theoretical Economics*, 3, 231–256.

SELTEN, R. (1975): “Reexamination of the perfectness concept for equilibrium points in extensive games,” *International Journal of Game Theory*, 4, 22–55.

YOON, K. (2001): “A Folk Theorem for Asynchronously Repeated Games,” *Econometrica*, 69, 191–200.