Games with Private Timing*

Yuichiro Kamada† Sofia Moroni‡

First version: May 8, 2014

This version: August 31, 2017

Preliminary

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 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 partial control over timing and imperfect disclosure.

*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, 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.

†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`

Contents

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 latter already had an internal meeting to determine such a strategy. A salesperson at an electric appliance store may change her sales talk 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 analyzes a new class of games that we call *games with private timing* to analyze such situations. In games with private timing, each of two players knows the time at which she chooses an action in a stage game once and for all, but does not know the time when her opponent does so. 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 change the outcome of the overall game?

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, in perfect Bayesian equilibrium (PBE), which we prove exists, certain conditions on the timing distribution (about which we will elaborate shortly) imply uniqueness of PBE when the stage game is either a coordination game or an opposing interests game.

In the coordination-games setting, we characterize the set of distributions of the timing of moves such that every game has the following unique PBE for small enough disclosure costs: players choose the best coordination equilibrium of the stage game and do not disclose their action. This implies an unexpected con-

	a	b
a	$0, 0$	$2, 2$
b	$1, 1$	$0, 0$

Figure 1: Targeting Game (the row player is firm A, and the column player is firm B)

sequence: for these distributions the introduction of the disclosure option helps select a PBE that does not involve disclosure.

The information revelation mechanism that we study is not only of theoretical interest but also fits the following example: Suppose that two rival firms are contemplating the prices and designs of their new competing products. They would benefit from targeting different segments of the market. Each firm has their own strength, and firm A would benefit more from targeting market segment a while having firm B target market segment b than targeting market segment b while having firm B target market segment a. Similarly, firm B would benefit more from targeting market segment b while having firm A target market segment a than targeting market segment a while having firm A target market segment b. Both firms targeting the same market segment results in harsh competition and benefits neither firm. This situation can be described as in Table 1. Since this is a coordination game, the firms would like to coordinate their actions (targeting strategies), but unfortunately communication between the firms is prohibited by law, which makes it risky to communicate because, with small probability, the communicating firm would be caught by the government (hence the cost is positive but small in expectation). If a firm nevertheless wants to communicate with the other firm, it must take into consideration whether the other firm has already decided their strategy or not, as if the other firm has already decided then it is merely costly to communicate and it would not generate any additional benefit. What we show in our main result is that, in this type of situation, the firms can target the respective "right" segment without communicating with each other.

When the stage game is an opposing-interest game (such as battle of the sexes) we give sufficient conditions on the distribution of the timing of moves such that there exists a unique equilibrium. In this equilibrium the player who happens to become the first player pays the disclosure cost.

The objective of the paper is two-fold. First, we aim to give a deeper understanding of the result described above by providing a necessary and sufficient condition for there to be a unique PBE. By doing so, we aim to demonstrate a non-triviality in the way knowledge about timing affects players' behaviors. Second, we hope to convince the reader that there are numerous open questions worth tackling in the context of private timing, both in theory and in applications. For this purpose, we complement the study of coordination games and opposing-interests games by analyzing games with a Stackelberg action and an extension of the model in which players cannot disclose their actions perfectly. In the former case we find that at the beginning of the game a player chooses the Stackelberg action. In the latter case, we show that cooperation without disclosure is the unique outcome for a coordination stage game as long as the probability that an action is observed is an increasing function of the risk dominance parameter of the stage game together with some conditions on the distribution of moving times. Of course, there are other types of stage games that are of interest, and we analyze some of them in the context with disclosure costs, while we must leave a large span of other unexplored possibilities for future research.

Let us provide a more detailed explanation for the uniqueness result for coordination games. The result hinges on two key conditions on the timing structure: *asynchronicity* and *uncertainty*. Asynchronicity means that players almost surely 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 structure implies less freedom on allowable beliefs at information sets. The implication of uncertainty deserves more explanation: With uncertainty (of the kind we will assume), players can think the opponent has not moved with some positive probability, so there are not many occurrences of off-equilibrium information sets being reached. 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.

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.[1] 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 necessarily perfect) coordination game.[2] The basic intuition for asynchronicity helping selection is similar. The difference is that these papers assume perfect information, while we assume that players may not observe the actions taken by the other player. In fact, on the equilibrium path, no one observes any action by the opponent.

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 involve 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. This equilibrium behavior has a similar flavor to the costly-disclosure option in our model.

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

[1]See also Yoon (2001), Lagunoff and Matsui (2001), and Dutta (1995).

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

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.

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 $N = \{1, \ldots, I\}$ is the finite set of players, A_i is player i's finite action space, and $g_i : A \to \mathbb{R}$ is player i's payoff function, where $A := A_1 \times \cdots \times A_I$.

Dynamic Game Time is discrete, and the dynamic game runs over time 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 $p(T_1, \ldots, T_I)$. 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 $(t_1, \ldots, t_I)$, player i chooses an element from $A_i \times \{\text{pay}, \text{not}\}$ at time t_i, observing her own t_i and (a_j, t_j) of the opponent j who chose $(\cdot, \text{pay})$ at $t_j < t_i$.[3] If player i chooses (a_i, pay) for some $a_i \in A_i$, then she pays the cost $c > 0$.[4] 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.

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

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

[3] In Section 5.10 we discuss the case in which the disclosure does not succeed with probability 1.

[4] The complete specification of the payoff structure is provided shortly.

where t^i is the moving time of player i and $(a_i, d_i) \in (A_i \times \{\text{pay}, \text{not}\})$ 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

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

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 the 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 $\hat{h}_i = \big(N', (\hat{t}_j, \hat{a}_j)_{j \in N'}, t\big)$ if (i) $t = \tilde{t}$, (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, $\sigma_i : \mathcal{H}_i \to \Delta(A_i \times \{\text{pay}, \text{not}\})$, is a map from private histories to (possibly correlated) probability distributions in A_i and disclosure decisions. Let Σ_i be the set of all strategies for player i. Define $\Sigma_{-i} = \times_{j \neq i} \Sigma_j$.

A strategy profile σ together with the a time distribution p induce a unique probability distribution over action profiles. If the distribution has a unit mass on an action profile a, then we say that a is the *outcome* of (σ, p).

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

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

with $c > 0$ which we assume to be common across players.[5] That is, if i chooses "pay" to disclose her action, she incurs cost c. The expected payoff for player i from strategy profile σ is denoted by $u_i(\sigma)$.[6] A *belief* $\mu \in \Delta(\mathcal{H})$ is a probability measure over histories. A strategy profile σ induces a continuation payoff $u_i(\sigma|\mu, t)$

[5] As 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.

[6] Formally, call $(a, (d_1, d_2)) \in A \times \{\text{pay}, \text{not}\}^2$ the choices induced by σ for (t, t') if a is the action profile played by the players and d_j is the disclosure decision by player $j = 1, 2$ when the

conditional on the belief that (i) the distribution of the past play at times strictly before t is given by μ, (ii) the play at and after time t is given by σ, and (iii) the probability distribution over Nature's move before, at, and after t is given by μ.

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 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.[7] Note that condition 1 implies that players best-respond to the beliefs computed by Bayes rule on the equilibrium path. In Section 5.7 we discuss 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.

realized times are (t, t') and players follow the strategy profile σ. Then, we define

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

where $(a^{(t,t')}, (d_1^{(t,t')}, d_2^{(t,t')}))$ denotes a choice in the support of σ at (t, t').

[7]This 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 the main section. 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 will believe 2 has already played some inefficient action upon not disclosing, and this may make it optimal for player 2 to disclose his action.

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

Figure 2: A Common-Interest Game

Pure Stackelberg Property Stackelberg actions will play a key role for characterizing uniqueness. We say that a component game $(N,(A_i)_{i\in N},(g_i)_{i\in N})$ satisfies the *pure Stackelberg property* if for each $i \in I$, 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 payoffs from all other action profiles. There are two important (exclusive and exhaustive) subcases 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 Section 5.8, we provide two examples that motivate the use of the pure Stackelberg property. These examples show that multiplicity of PBE may hold if we consider 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 $\mathcal{T} = \mathbb{Z}$, and it has full support and is sparse. Specifically, assume that there exists a small $\epsilon > 0$ such that $0 < f(t) < \epsilon$ 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 $t_1, t_2 \in \mathcal{T}$.

Consider the payoff matrix in Figure 2. Let S be this coordination game. We can show the following:

Proposition 1. *There exist $\bar{\epsilon} > 0$ and $\bar{c} > 0$ such that for all $\epsilon \in (0, \bar{\epsilon})$ and $c \in (0, \bar{c})$, there is a unique PBE in the game $(S, \mathcal{T}, p, c)$ with S, $\mathcal{T}$, and p as* *specified above. On the path of this 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) and they do not exercise the option to disclose.

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

To illustrate the non-triviality of this result, below we explain three possibilities of alternative timing structures and equilibria in these settings.[8]

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 2. 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 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. Hence there are three PBE in the dynamic game corresponding to these three Nash equilibria. □

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

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 2. 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.

[8]We 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.

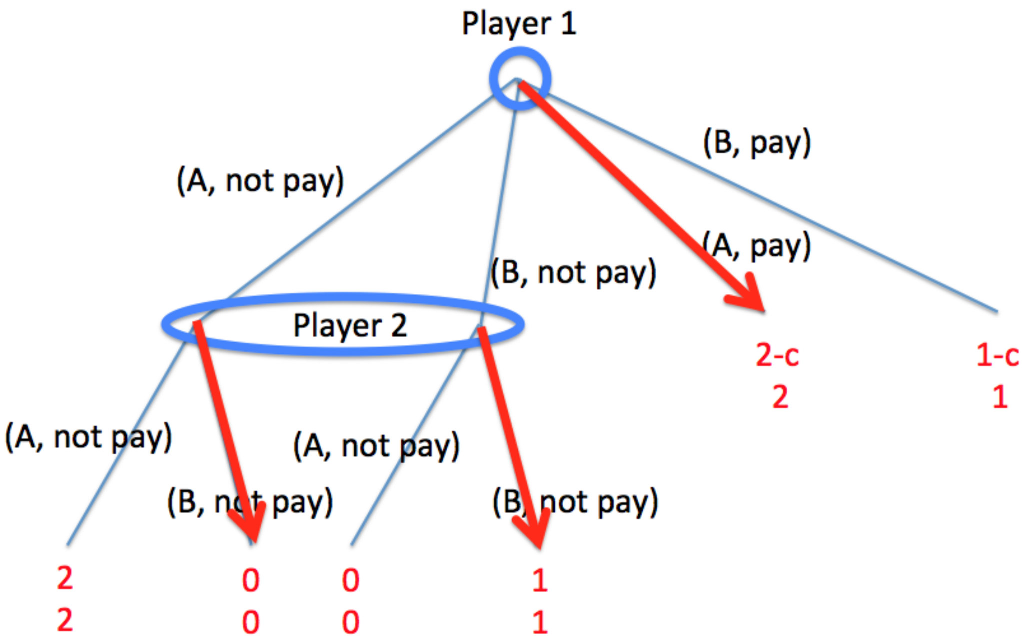

Figure 3: A Forward-Induction Argument

Let us check if 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. The only tricky part is his response after no disclosure. It is a best response given a 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 3 of the game in consideration. 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 information set h_2, player 2's belief must assign probability 0 to the right node in player 2's information set. Given this belief, player 2's unique best response at h_2 is to play (A, not). Hence, 1 can obtain the best 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

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

Figure 4: A Risky Common-Interest Game

dominated. This suggests 1 should play (A, not).

Our 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.[9]

$\square$

Example 3. [Correlated-Move Game]

Suppose that $N = \{1, 2\}$, $\mathcal{T} = \mathbb{Z}$ and for all $t_1 \in \mathbb{Z}\backslash\{0\}$, 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. Also we 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 $\frac{1}{2}$ and negative with probability $\frac{1}{2}$, and follows an exponential distribution with a rate r on each side. Player 2's moving time is 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 a component game S, consider the game in Figure 4.

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

The reason for this second strategy profile to be a PBE is that player 1 deviating to play A could improve her payoff only when she paid the disclosure cost, and in such a case she would miscoordinate with player 2 with probability $1/2$ and succeed in coordination 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 $\frac{1}{2}$. The key

[9] See Remark 2 for a further explanation about this similarity.

difference from the starting example is that at any realization of a player's moving time, the probability that the player can assign to her being the first mover is only (close to) 1/2, while in the starting example it can be arbitrary close to 1 as $t \to -\infty$. $\square$

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, because Proposition 1 is a corollary of that theorem. The proof consists of two steps. In the first step, we prove that players play only A on the path of play in any PBE. The second step shows that no player pays the disclosure cost.

- *First Step:* Suppose that player i gets her move at time t that is early enough in the game and she has observed no action disclosed. The assumption that the timing distribution is independent across players implies that the probability that she assigns to the event that the opponent will move later is close to 1. Hence 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 never plays B. This is true for all times before t. Since $-i$ knows this, $-i$ never plays B at $t + 1$. We can go forward to show 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. If she plays (A, pay), the expected payoff she can get is at most $2 - c$. Suppose that i does not pay. 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'. Since $-i$ was on the path of play at and before t, 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 in any PBE, player i never pays 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 Sketch of Proof, 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 ignored the probability that i moves at $t+1$ because it is small. Example 2 illustrated the effect of deterministic move on multiplicity. We used the fact that the timing structure is uncertain in the second step of the Sketch of Proof, 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 under a highly correlated timing structure multiplicity is possible. In particular, in the second PBE in Example 3, for any realized moving time, the moving player assigns a high probability to the event that the other player has already moved. We used the fact that the timing structure is independent in the first step of the Sketch of Proof, 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 the Sketch of Proof above, Step 1 rules out B as it is dominated by (A, pay). The forward induction argument 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 she plays (A, not) she gets the best payoff. In our private timing game, we do not need to use 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 $-i$ assigns a positive probability of i moving later so she is on the path of play, thus, Step 1 shows that the player must take A in equilibrium. Step 1 of our proof 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 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 players without any observation are always on the path of play. Thus, 2's getting a move without observation of disclosure still leaves the possibility that 1 has not moved. □

Remark 3. [**Lack of Lower Hemicontinuity**]

Suppose that $N = \{1, 2\}$ and $\mathcal{T} = \mathbb{Z}$. The time distribution is independent across players and is parameterized by ξ, 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 $\xi > 0$, the same logic as the above Sketch of Proof applies to show that there is a unique PBE, and in that unique PBE each player plays (A, not).

These time distributions converge pointwise to the distribution with $p(1, 2) = 1$ as $\xi \to 0$. 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 structure.[10] One additional PBE that we obtain in the limit is the pessimist equilibrium. The reason for the lack of lower hemicontinuity is that the game is not finite. In particular, for any time t that 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. □

3.2 General Common Interest Games

In this section we generalize the findings we have discussed so far in order to understand the role of the timing distribution on the set of PBEs for common interest games. For this purpose, we characterize the set of timing distributions such that the best action profile with no disclosure is the PBE outcome when the component game is a common interest game. The following two conditions are satisfied in the timing structure of the starting example.

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.

This condition is not satisfied in the timing structure 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").

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

Definition 2. We say that the time structure p has **frequent 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 $t \in [t'', t'] \cap \mathcal{T}_i$, $\text{Prob}^p(t'' < T_j \leq t | T_i = t) < q$.

This condition is satisfied if, for example, moves take place only at integer times (or some other sparse set of times), or if the distributions of moves are independent across players.

This condition is not satisfied in the timing distribution of Example 3. The reason is that for $t'' = -\infty$ and $q > 0$ small enough, 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 "frequent potential moves" condition implies that if a player moves early enough then she thinks the probability of being the second mover is low.

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. Time distribution p is **asynchronous** if $\text{Prob}^p(T_i = T_{-i}) = 0$.

Define $T_i^< = \{t \in \mathcal{T}_i | \text{Prob}^p(T_{-i} > t) = 0\}$. 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, for every $t \in \mathcal{T}$:

1. p is asynchronous.

2. Either one of the following two conditions holds.

   (a) p satisfies the potential leader condition.

   (b) Otherwise, for every $E \subseteq T_i^<$, $\inf_{\tilde{t} \in \mathcal{T}_{-i}} \{\text{Prob}^p(T_i \in E | T_{-i} = \tilde{t}) | \text{Prob}^p(T_i \in E | T_{-i} = \tilde{t}) > 0\} = 0$.

Theorem 1. *Fix $\mathcal{T}$, and assume that p has frequent potential moves. The time 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 $c < \bar{c}$ the only PBE of $(S, \mathcal{T}, p, c)$ is such that each player plays (a_i^*, not).*

Remark 4.

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

Figure 5: Counterexample for Necessity part

1. The proof for sufficiency of the two conditions has the same structure as that of the starting example. In the example, however, we only considered this sufficiency part. There, the asynchronicity condition was not satisfied 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 for making the connection between the degree of commonality of the component game and that of concentration of 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 concentration parameter. Since the proof in that section implies the sufficiency of conditions 1 and 2(a), for sufficiency, in the proof below we refer to Theorem 4 in Section 5.2.

2. The proof for necessity is by construction. That is, for each of the conditions characterizing set D, we construct an example of a component game that has multiple PBE, for each timing distribution violating that condition. For example, for a distribution that has a positive probability of synchronous moves at some time t, it is easy to construct an example of a component game in which an action other than a_i^* may be a best response 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. Specifically, we consider the component game as in Figure 5 and show that a PBE strategy profile in which B is played at certain period exists if M is sufficiently large. The proof of necessity of condition 2 is conducted using an analogous manner using the same component game. The details are relegated to the Appendix.

3. Condition 2 consists of two parts, and only one of them needs to be satis-

fied. That is, the potential leader condition (condition 2a) is sufficient but not necessary to guarantee uniqueness. Example 4, after the proof of the theorem, to illustrate this point.

4. Notice that condition 2a requires that $t < \sup_{t' \in \mathcal{T}} t'$ holds 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 5.6.

5. One possible application of the theorem is the case in which the analyst only knows that the players face a common interest games and the structure of the game is common knowledge among 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 of the disclosure cost takes place) is obtained. On the other hand, it is possible that the analyst's interest is only in the actions in the component game and not in the disclosure behavior. This may happen because the costs of disclosure are low, or because the analyst does not observe the size of the cost. It is relatively straightforward to identify the sufficient and necessary condition on the time distribution p such that a^* is the outcome of (σ, p) for all PBE σ. 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 2. The timing distribution p over $\mathcal{T} = \mathbb{Z}$ is given by the following rule: With probability $\frac{1}{2}$, T_1 follows the discrete exponential distribution over positive integers with rate λ, while T_2 follows the discrete exponential distribution over nonpositive even integers with rate λ. With the complementary probability $\frac{1}{2}$, T_1 follows the discrete exponential distribution over negative odd integers with rate $\lambda' < \lambda$, 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, $\lambda' < \lambda$ 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 the argument as above, 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 satisfying this property are opposing interest games. In this section, we consider two-player opposing interest games, show uniqueness of PBE, and discuss the difference of the structure of this unique PBE from that of common interest games. In particular, we show that under regularity conditions, players pay the disclosure cost in the unique PBE.

Given an opposing interest game, let g_i^* and g_i^S be the best and the second best payoffs, respectively, for player i. By the definition of pure Stackelberg property, $g_i^S \neq g_i^*$ holds for each 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

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

Assumption 1.

1. For each $i = 1, 2$, $p_i(t, t_0, t_0, t)$ is non-decreasing in t for $(t, t_0) \in \mathbb{R} \times (\mathbb{R} \cup \{-\infty\})$ with $t_0 \leq t$ such that $\text{Prob}^p(T_{-i} \geq t_0, T_i = t) > 0$,

2. 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 $\forall t \in [t_0, \bar{t}_i] \cap \mathcal{T}_i$, $p_i(t, t_0, t_0, t) < \varepsilon$,

3. There is $\alpha > 0$ such that, for every $i, j \in \{1, 2\}$ and $(t_1, t_2, t_3) \in \mathcal{T}_j \times \mathcal{T}_i \times \mathbb{R}$ with $t_1 \leq t_2$, $t_3 \leq t_2$, $\text{Prob}^p(T_{-i} \geq t_3, T_i = t_2) > 0$ and $\text{Prob}^p(T_{-j} \geq t_3, T_j = t_1) > 0$, we have $p_j(t_1, t_3, t_3, t_2) \geq \alpha p_i(t_2, t_3, t_3, t_2)$.

4. 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, for any given t_0, the probability that the opponent has moved in $[t_0, t]$ conditional on receiving an opportunity at t is nondecreasing in t. This condition is satisfied, for example, when the time distribution is independent. Assumption 1 (2) says that for any t_0 before which no one has moved, if a player's moving time is after t_0 but close enough to it, then the probability that the opponent has moved in $[t_0, t]$ can be made arbitrarily small. The assumption rules out simultaneous moves. Assumption 1 (3) says that for any time interval with end points in $\mathcal{T}$, conditional on no one having moved before the left endpoint of the interval, the probability of player i's opponent moving in that time interval conditional on i moving at the left endpoint cannot be too low relative to the probability of player j's opponent moving in that time interval conditional on j moving at the right endpoint, where i and j can be either the same or different. Assumption 1 (4) says that for any time t_i, each player i believes that there is a positive probability of the opponent moving before t_i if i moves at sufficiently early. Note that both assumptions (2) and (4) imply that $\inf \mathcal{T}_1 = \inf \mathcal{T}_2$.

An example of a time distribution satisfying all the above assumption is the one in which player 1's moving time is distributed according to the exponential distribution over the negative odd numbers, player 2's moving time is distributed according to the exponential distribution over the negative even numbers, and these two distributions are independent.

Theorem 2. *Let S be an 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, and in that PBE, each player* $i \in I$ *chooses* (a_i^*, pay) *under the history* $(\emptyset, \cdot)$.

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. The complication 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 section. 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, and the third part defines ε-constrained equilibria which exist in the N'th approximation 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 the sequence as $N \to \infty$), and part 5 again uses the lemma to show existence of trembling-hand perfect equilibrium. (by considering a subsequence of the sequence as $\varepsilon \to 0$) as in Selten (1975) (which is PBE too). The reason we use extensive-form 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 therefore uncountable histories) is due to 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 this proof method is that it shows existence of trembling-hand equilibrium. Thus, the fact that our definition of PBE is not stringent does not play a key role in proving the existence theorem.

5.2 q-Concentrated Time Distribution and s_i-Common Interest Games

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 $g_i^S = \max_{a \in A \setminus \{a^*\}} g_i(a)$ be the second-highest payoff for player i. We assume $g_i^* \neq \underline{g}_i$.

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 measures how good the best payoff is for player i.

We will say that a time structure p is q-concentrated if at every time t'' there is a sufficiently close future time t' such that for all time t between t'' and t', i moving at t implies that j moves between t' and t with probability less than q. Formally,

Definition 5. We say that the time structure p is q**-concentrated** 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 $t \in [t'', t'] \cap \mathcal{T}_i$, $\mathbb{P}(t'' \leq T_j \leq t | T_i = t) < q$.

In particular, q-concentration is satisfied when T_1 and T_2 are independent and the probability of each moving time is strictly less than q (this includes our starting example with $\epsilon \leq q$). Notice that in Example 3, the assumption fails when $q > 0$ is sufficiently small. This is because for any $t \in \mathcal{T}$, the probability that the opponent moves earlier is at least $\frac{r^2}{1+r^2}$. The definition of q-concentration is more stringent in that it pertains to conditional probabilities at any time not just early in the game.

We now characterize the joint distributions of T_1 and T_2 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* $(S, \mathcal{T}, p, c)$. *Suppose that there exist* $(s_i)_{i \in N} \in \mathbb{R}_{++}$ *and* $\epsilon > 0$ *such that* S *is a common interest game that is* s_i*-common for each* $i \in N$ *and that* $N = \{1, 2\}$ *and* p *is* $(\min_{i \in N} s_i - \epsilon)$*-concentrated. Then there exists* $\bar{c} > 0$ *such that for all* $c < \bar{c}$, a^* *is the outcome 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 we use the proof provided below to prove the sufficiency part of Theorem 1.

Proof.

Step 1:

Step 1-1: Fix a common interest game that is $(s_i)_{i \in N}$-common, $\epsilon \in (0, \min_{i \in N} s_i)$, and a time structure p that is $(s_i - \epsilon)_{i \in N}$-concentrated. Fix a PBE and take $c \in (0, \min_{i \in N} \left[(1 - (s_i - \epsilon))g_i^* + (s_i - \epsilon)\underline{g}_i - g_i^S\right])$. Notice that, by the definition of s_i, $(1 - (s_i - \epsilon))g_i^* + (s_i - \epsilon)\underline{g}_i - g_i^S = \epsilon(g_i^* - \underline{g}_i) > 0$. 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 $t^* := \inf_{t \in \bigcup_{i \in N} N_i(a^*)} t$.

Step 1-2: By the definition of q-concentration, there must exist $i \in N$ and $t' > t^*$ such that for $j \neq i$ and $t \in [t^*, t'] \cap \mathcal{T}_i$, $\mathbb{P}(t^* \leq T_j \leq t | T_i = t) < s_i - \epsilon$. Our choice of c implies that (a_i^*, pay) would give such i a strictly higher payoff, 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.

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. As we have shown above, the outcome under σ^* must be a^*, so i's payoff from σ^* is $g_i^* - c$. 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 $(a_i^*, \cdot)$. Hence the outcome under $(\sigma_i', \sigma_{-i}^*)$ must be 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. Hence there is no time at which any player pays the disclosure cost. □

5.3 Sense of Time

Example 3 demonstrates why we need the frequent potential moves condition to prove uniqueness. The idea of the example was that any private information about a player's own moving time does not reveal sufficiently precise information about the order of moves. To make this point even clearer, here we consider an extreme case in which players do not have a sense of time.

More specifically, consider an extensive-form game in which the Nature chooses one of the two states with probability 1/2 each. Player 1 moves first in the first state, and player 2 moves first in the second state. Players do not know the order of moves unless the opponent reveals the action, so if the strategy profile assigns probability one to no one revealing any action, then at each information set, each player assigns probability 1/2 to being the first mover. The set of available actions and the payoff functions are exactly the same as in the starting example. Figure 6 shows the extensive form of this game with payoffs at each terminal node.

As in Figure 3, we omit the actions that are strictly suboptimal (conditional on reaching the corresponding information set) when a player knows she is the second mover.

As in Example 3, there are at least two PBE. One is that each player plays (A, not) under no observation, while the other is that each player plays (B, not) under no observation. The second strategy profile is a PBE because if one follows it, the expected payoff is 1, while if she deviates to play (A, not), the expected

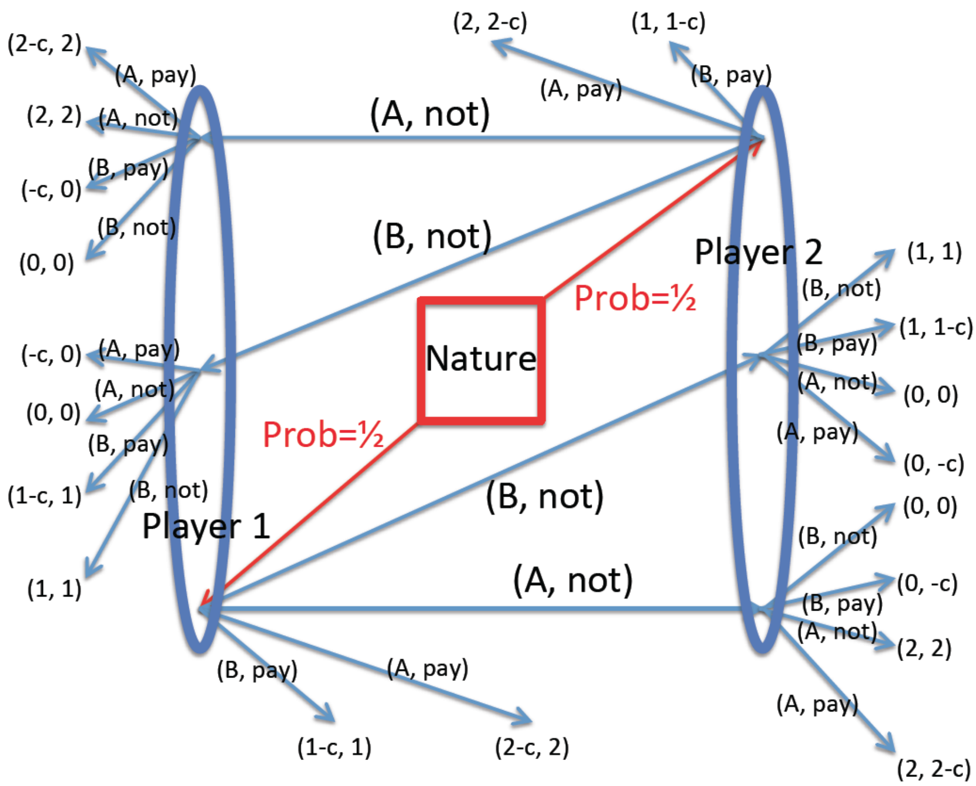

Figure 6: Extensive-Form of the Game with No Sense of Time: Nature moves first. Each player i cannot distinguish among the three possible histories (i being the first mover, $-i$ played A and did not disclose, and $-i$ played B and did not disclose). We omit some actions after the first-mover's payment.

payoff reduces to $\frac{1}{2}(-c) + \frac{1}{2}(2 - c) = 1 - c$.

5.4 Choice of Moving Times

In the main section we assumed that players do not have any control over when to move. This assumption fits many real-life situations discussed in the Introduction. However, in other situations players may have some control over the timing of moves. Here we consider such a situation and argue that our basic results 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, consider a model in which before the dynamic game starts, each player i 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. Let $(S, \mathcal{T}, (\mathcal{D}_i)_{i=1,2}, c)$ be the new 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 high probability is assigned to the times close to the chosen time, but with some probability the moving times can be far from the chosen time.

Corollary 2. *Let S be a common interest game with the best action profile (a_1^*, a_2^*). Suppose that $supp(f) \cap supp(g) = \emptyset$ and $\min\{\sum_{t'>t} f(t'), \sum_{t'>t} g(t')\} > 0$ for all t for all pairs $(f, g) \in \mathcal{D}_1 \times \mathcal{D}_2$. Then, $(S, \mathcal{T}, (\mathcal{D}_i)_{i=1,2}, c)$ has a unique PBE, and on the path of this PBE, each player plays (a^*, not) at any t without observation of a disclosure.*

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 is arbitrarily high (but is less than 1).

5.5 Repeated Games

Here we apply our uniqueness result to the setting with repeated interactions. Consider a sequence of countable sets of moving times, $(\mathcal{T}^{(1)}, \mathcal{T}^{(2)}, \dots)$ where $\mathcal{T}^k \subseteq [k-1, k)$ for each $k = 1, 2, \dots$, and a sequence of probability distributions $(p_1, p_2, \dots)$ where $p_k \in \Delta((\mathcal{T}^{(k)})^2)$ for each $k = 1, 2, \dots$. We consider the situation in which for each interval $[k-1, k)$, moving times of two players are

	C_2	D_2
C_1	$1, 1$	$-s_1, d_2$
D_1	$d_1, -s_2$	$0, 0$

Figure 7: Prisoner's Dilemma

drawn according to p_k, and they move observing the outcomes at the moving opportunities at times in $[0, k-1)$ (including the actions that are not disclosed in that time interval). We suppose that, for each k, $p_k \in D$ where D is as defined in Section 3.2.

The component game is a prisoner's dilemma game as in Figure 7 with $d_i > 1$, $s_i > 0$, and $d_i - s_i < 2$. As in the base model, we assume that at each opportunity each player i chooses from $\{C_i, D_i\} \times \{\text{pay}, \text{not}\}$. The discount rate is $\rho > 0$. The payoff from the action profile in time interval $[k-1, k)$ materializes at time k. We call this model of repeated interactions as the **repeated private-timing prisoner's dilemma**.

We consider the following class of supergame strategy profile: a strategy profile is a T**-grim trigger** if each player i plays C_i forever if there has been no D_i in the first T periods, and plays D_i otherwise. This does not pin down the strategies at the first T periods. The following proposition shows that with asynchronicity and private timing, there is a unique PBE in the class of such strategies. The proof is given in the Appendix.

Proposition 3. *In the repeated private-timing prisoner's dilemma, there exists $\bar{\rho} > 0$ such that for all $\rho < \bar{\rho}$ and all $T < \infty$, the only PBE in the class of T-grim trigger strategies has the outcome such that (C_1, C_2) is played at all the realized moving opportunities.*

Remark 5. Four remarks are in order.

1. In contrast to the result in Proposition 3, if the moves at each period are simultaneous, then a folk theorem holds even with the restriction to the set of equilibria that play grim trigger far in the future.

2. The choice of T can be independent of the discount rate ρ. In particular, there is no restriction on the size of $e^{-\rho T}$. This means that, under the private

timing environment, only a slight punishment at the end of the T periods (from the viewpoint of period 0) can be useful in sustaining cooperation as a unique outcome.

3. The idea of using the threat at the end of a finite horizon is present in the literature (e.g., Benoit and Krishna (1985)), where these threats are used to enlarge the set of payoffs. In contrast, in our setting we use these to obtain cooperation as a *unique* outcome.

4. With perfect monitoring, backwards induction implies cooperation for all T periods because of asynchronicity. Our result is similar, but we show that such results are true even in the presence of uncertainty about timing. $\square$

5.6 Horizon Length

As we noted when stating Theorem 1, it is important that the support of the moving times is infinite in at least one direction. We make this claim formal here. 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 exists $\mathcal{T} \subseteq \mathbb{Z}$ with $\min_{t\in\mathcal{T}} t = t^*$ and p such that every common interest game S, the dynamic game $(S, \mathcal{T}, p, c)$ has a unique PBE, and in that PBE each player i plays (a_i^*, not) at any T_i.*

2. For any $t^* \in \mathbb{Z}$, there exists $\mathcal{T} \subseteq \mathbb{Z}$ with $\max_{t\in\mathcal{T}} t = t^*$ and p such that every common interest game S, the dynamic game $(S, \mathcal{T}, p, c)$ has a unique PBE, and in that PBE each player i plays (a_i^*, not) at any T_i.

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

Thus, the time 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 limits[11]: here we fix a

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

time 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 $(\mathcal{T}_k, p_k)$ defied by $\mathcal{T}_k = \{1, \ldots, K\}$ 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 common interest game S such that the dynamic game $(S, \mathcal{T}_K, p_K, c)$ has multiple PBE.*
2. *For any common interest game S, there exists K such that the dynamic game $(S, \mathcal{T}_K, p_K, c)$ has a unique PBE and in that PBE each player i plays (a_i^*, not) at any 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.

5.7 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. That is, there may exist multiple Bayes Nash equilibria. To see this, consider the component game as in Figure 2 and an independent timing structure 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.

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

	A	B	C
A	$1, 1$	$0, 0$	$0, 0$
B	$0, 0$	$3, 0$	$0, 3$
C	$0, 0$	$0, 3$	$3, 0$

Figure 8: Mixed Stackelberg leads to multiplicity

5.8 Multiplicity in Component Games without the Pure Stackelberg Property

In the main sections we focused on games with the pure Stackelberg property. In order to understand the role of that assumption, here we consider component games without the pure Stackelberg property, and show there may be multiple PBE under those component games under certain timing structures.

First, we provide an example of the dynamic game in which the Stackelberg action in the component game is mixed and there are multiple PBE.

Example 5 (Mixed Stackelberg Leads to Multiplicity)**.**

Consider the two-player component game in Figure 8 and an independent timing structure such that $\mathcal{T}_1$ is the set of odd natural numbers and $\mathcal{T}_2$ is the set of even natural numbers. There are at least two PBE in the dynamic game. In one PBE, each player, upon receiving the chance to move without observation, plays (A, not). If the player observes the opponent's action before she moves, then she takes the (unique) static best response.

In the other PBE, each player, upon receiving the chance to move without observation, plays $(\frac{1}{2}B + \frac{1}{2}C, \text{not})$, where $\frac{1}{2}B + \frac{1}{2}C$ denotes the half-half mixing of actions B and C. Again, if the player observes the opponent's action before she moves, then she takes the (unique) static best response.

The reason for multiplicity in this example is that the mixed Stackelberg action profile Pareto-dominates the pure one, but there is no way to disclose the deviation to such a mixed action because only the *realized* action can be disclosed. In the main sections we restricted attention to the case in which the Stackelberg action is pure to avoid this type of complexity. □

The pure Stackelberg property also implies that the Nash equilibrium a^i in

	S	B_1	B_2
S	$-1,-1$	$5,0$	$-1,0$
B_1	$0,5$	$2,1$	$0,0$
B_2	$0,-1$	$0,0$	$1,2$

Figure 9: Multiple best responses to the Stackelberg action

the component game is strict. If i's opponent has multiple best replies in the component game, the dynamic game can have multiple PBE. The next example shows this point.

Example 6 (Multiple Best Responses to the Stackelberg Action)**.**

Consider the two-player component game in Figure 9 and an independent timing structure such that $\mathcal{T}_1$ is the set of odd natural numbers and $\mathcal{T}_2$ is the set of even natural numbers. Notice that, in the component game, S is the Stackelberg action for each player, and B_1 and B_2 are both best replies to S.

There are at least two PBE in the dynamic game when the disclosure cost satisfies $c \leq 1$. In one PBE, each player, upon receiving the chance to move without observation, plays (S, pay). If player i observes the opponent's action before she moves, then she takes a static best response, where if she observes S then she takes B_1.

In the other PBE, each player, upon receiving the chance to move without observation, plays (B_i, pay). If player i observes the opponent's action before she moves, then she takes a static best response, where if she observes S then she takes B_2.

To formalize what the problem was, let $a_i^* \in \arg\max_{a_i \in A_i, a_{-i} \in BR(a_i)} g_i(a_i, a_{-i})$. The problem with the above example is that there exist $a'_{-i} \in BR(a_i^*)$, a_i and $a''_{-i} \in BR(a_i)$ such that $g_i(a_i^*, a'_{-i}) < g_i(a_i, a''_{-i})$. In the component game in this example, this corresponds to the situation where $-i$ has a best response to S that gives i a worse payoff than (B_1, B_1) or (B_2, B_2), in which $-i$ takes a best response. This is why in our main analysis we focus on the case in which $g_i(a_i^*, a'_{-i}) \geq g_i(a_i, a''_{-i})$ for any $a'_{-i} \in BR(a_i^*)$, a_i and $a''_{-i} \in BR(a_i)$. That is, we consider the case in which each i has an action that guarantees herself a better payoff than any other action of hers does when the opponent best-responds. $\square$

5.9 Various Component Games

Here we consider various classes of component games outside the ones that satisfy the pure Stackelberg property.

5.9.1 Constant-Sum Games

Proposition 6. *Suppose that the component game S is a 2-player constant-sum game. Then, for any dynamic game $(S, \mathcal{T}, p, c)$ with any cost $c > 0$, for each i and $t \in \mathcal{T}_i$, i assigns probability zero to (a_i, pay) for any $a_i \in A_i$ after any history under any PBE. The probability distribution over component-game action profiles under any PBE is a correlated equilibrium of S.*

Proof. Fix a constant-sum component game. Let $U = g_1(a) + g_2(a)$ for some $a \in A$. For each player i, consider a minmax strategy α_i. Consider player i's strategy that plays (a_i, not) with probability $\alpha_i(a_i)$ for each $a_i \in A_i$ conditional on any history. This strategy gives a lower bound of player i's expected payoff under any PBE. This lower bound is her minmax value of the component game. This is true for both players, so the sum of the payoffs from the dynamic game under any equilibrium path is U.

This implies that no player assigns positive probability to (a_i, pay) for any a_i under any equilibrium path, as otherwise the sum of the payoffs must be strictly less than U.

As a result, the game $(S, \mathcal{T}, p, c)$ is equivalent to the simultaneous-move game with a correlated randomization device, so in particular the probability distribution over equilibrium action profiles is equal to that of correlated equilibria of the stage game. □

The intuition is simple: It is costly for a player to reveal the choice of an action, and it can only decrease the player's own payoff because the opponent, if he gets the move afterwards, will best-respond and the preferences are perfectly misaligned in constant-sum games.

The result does not generalize to the cases with more than two players. The next example illustrates this point.

Example 7. Consider the component game S with the payoff matrix given by Figure 10. Let $\mathcal{T} = \{1, 2, 3\}$ and $p(1, 2, 3) = 1$. That is, with probability one,

A_3:

	A_2	B_2
A_1	$2, 2, -4$	$2, 2, -4$
B_1	$0, 0, 0$	$1, 1, -2$

B_3:

	A_2	B_2
A_1	$2, 2, -4$	$2, 2, -4$
B_1	$0, 0, 0$	$1, 1, -2$

Figure 10: A Three-Player Constant-Sum Game

player i moves at time i. Notice that the payoffs for players 1 and 2 are those of common interest games and their payoffs do not depend on 3's actions. Then, the dynamic game $(S, \mathcal{T}, p, c)$ has a PBE in which player 1 pays the disclosure cost, as we have seen in Example 2. In general, for any two-player component game and timing structure such that there exists a PBE in which some player pays the disclosure cost under some history on the equilibrium path, we can "add in" a third player such that (i) such a player's actions do not affect the first two players' payoffs and (ii) the whole component game is a constant-sum game.

5.9.2 Dominance Games

Proposition 7. *Suppose that in the component game S, each player i has a strictly dominant action a_i^D. Then, for any timing structure p, each player i plays (a_i^D, not) under any history.*

Proof. Fix a PBE. Suppose that after a given history at time t, player i plays (a_i, pay) for some action $a_i \in A_i$. Then, if $-i$ moves at $t' \leq t$, then (a_i, pay) with $a_i = a_i^D$ is i's strict best response at t because $-i$'s action is independent of i's. If $-i$ moves at $t' > t$, then again (a_i, pay) with $a_i = a_i^D$ is i's strict best response at t because $-i$'s best response at t' is (a_{-i}, not). Thus, conditional on i playing (a_i, pay), we must have $a_i = a_i^D$.

Next, suppose that after a given history at time t, player i plays (a_i, not) for some action $a_i \in A_i$. Then, irrespective of $-i$'s action, (a_i, not) with $a_i = a_i^D$ is i's strict best response at t because $-i$'s action is independent of i's. Thus, conditional on i playing (a_i, not), we must have $a_i = a_i^D$.

The above argument implies that each player i assigns probability 1 to $\{(a_i^D, \text{pay}), (a_i^D, \text{not})\}$ under any history on the equilibrium path. Given such a strategy of player i, $-i$'s payoff from playing (a_{-i}^D, pay) is $g_{-i}(a_{-i}^D, a_i^D) - c$, while the one from plying (a_{-i}^D, not) is $g_{-i}(a_{-i}^D, a_i^D)$. Since the latter is strictly greater than the former, each player plays (a_{-i}^D, not) under any history. $\square$

	A_2	B_2
A_1	$1, 1$	$-\beta_1, 1-\beta_2$
B_1	$1+\alpha_1, -\alpha_2$	$0, 0$

Figure 11: Reputation Game

The intuition is again simple. Since each player moves only once, there is no way to incentivize the opponent to play non-dominant action. Given such a consideration, there is no incentive to play a non-dominant action, or to pay a cost to disclose actions.

5.9.3 Reputation Games

Here we consider the class of "reputation games" as in Figure 11.[12] We assume that $\alpha_1, \alpha_2, \beta_1, \beta_2 \geq 0$ and $\beta_2 \leq 1$. Note that B_1 dominates A_1, and after eliminating A_1, B_2 dominates A_2. So (B_1, B_2) is a unique Nash equilibrium of this component game. Player 1 prefers (A_1, A_2) to (B_1, B_2), and A_2 is a best response to A_1. In what follows, we consider a PBE in which if player 1 receives her opportunity early then she "teaches" player 2 that she has taken A_1. Specifically, we aim to give a sufficient condition on the time distribution p such that, for an open set of payoff structures as the one shown in Figure 11, such a PBE exists. For this purpose, we define a few conditions on p.

First, say that p is such that **player 1 has no more than an ϵ-mass** if, for every $t \in \mathcal{T}_1$, $\sum_{t' \in \mathcal{T}_2} p(t, t') \leq \epsilon$ holds.

Next, the time distribution p is such that **player 1 can be arbitrarily early** if for every $\epsilon > 0$, there exists $t \in \mathcal{T}_1$ such that $\text{Prob}^p(T_2 < t) < \epsilon$. Similarly, the time distribution p is such that **player 1 can be arbitrarily late** if for every $\epsilon > 0$, there exists $t \in \mathcal{T}_1$ such that $\text{Prob}^p(t < T_2) < \epsilon$. This condition holds, for example, if the potential leader condition holds.

Finally, we say that the time distribution p is such that **player 2 is at most Δ-lagged** if for every $t \in \mathbb{R}$, $\text{Prob}^p(T_1 \leq t) - \text{Prob}^p(T_2 \leq t) \leq \Delta$.

Proposition 8. *Suppose that S is given by the payoff matrix as in Figure 11, and p is asynchronous, independent, and such that there exists $\epsilon > 0$ such that player 1*

[12]This game is called a reputation game as it is used in the literature on repeated games and reputation (see Mailath and Samuelson (2006)).

has no more than an ϵ-mass, player 1 can be both arbitrarily early and arbitrarily late, and player 2 is at most Δ-lagged. Suppose also that

$$\frac{1-c}{1+\beta_1} + \Delta + \epsilon \leq \frac{\alpha_2}{\beta_2 + \alpha_2}.$$

Then, the dynamic game $(S, \mathcal{T}, p, c)$ has a PBE in which (A_1, pay) is played with positive probability.

Note that it is easier for the condition to hold if the cost is high, the lag is small, the mass is small, the relative desirability of (A_1, B_2) for player 1 (measured by $\frac{1}{1+\beta_1}$) is small, or B_2 is safe (measured by $\frac{\alpha_2}{\beta_2+\alpha_2}$) for player 2.[13]

The proof is by construction. Specifically, we construct a PBE in which player 1 commits to A_1 and discloses the action early in the game, while she plays (B_1, not) later in the game. Player 2, on the other hand, always plays (B_2, not) unless (A_1, pay) has been observed. Player 1 does not have an incentive to play (A_1, pay) later in the game because she assigns a sufficiently high probability to being the second mover, in which case 2 has played B_2. Then cutoff time point of the switch from (A_1, pay) to (B_1, not) is pinned down by player 1's incentives, and the inequality in the statement of the proposition is used to guarantee that player 2 always takes a best response given such time cutoff. We impose that player 1 can be arbitrarily early to ensure there is a time at which (A_1, pay) is played, and that player 1 can be arbitrarily late to ensure that player 2 cannot be sure that he is a second mover which may happen if player 1 does not have a potential moving time after the cutoff time. Asynchronicity and independence simplify our computations.

The limit expected payoff profile under the unique PBE as $c \to 0$, $\sup_{t\in\mathcal{T}} \max_{i\in\{1,2\}} |\text{Prob}^p(T_i = t) \to 0$, and $\sup_{t\in R} |\text{Prob}^p(T_1 > t) - \text{Prob}^p(T_2 > t)| \to 0$ (i.e., the time distribution approaching the "uniform distribution.") is $(5/4, 15/8)$, which is a convex combination of three pure action profiles except (B_1, A_2). The action profile (B_1, A_2) cannot be played in this equilibrium because player 2 never plays A_2 unless he become sure that 1 has played A_1.

The implication of Proposition 8 is that the result in Section 5.9.2 does not

[13]Note that B_2 is a best response if player 1 assigns probability no more than $\frac{\alpha_2}{\beta_2+\alpha_2}$ to B_1.

extend to iterated dominance. The reason for the difference is that, one player's static best response depends on the second player's choice (before deletion of the dominant action of the second player), and the second player can commit to a dominated action to incentivize the second player to play a particular action.

5.10 A Probabilistic Disclosure Model

In this subsection we consider the possibility that disclosure of the action is successful with probability less than one. Specifically, if player j moves after i and i plays (a_i, pay), then j's private information contains information about a_i and i's moving time with probability r, while with the complementary probability his private information does not contain a_i or i's moving time; so in particular j does not observe whether i has moved or not. Formally, define the dynamic game $(S, \mathcal{T}, p, c, r)$ which is the extension of the standard game $(S, \mathcal{T}, p, c)$ such that disclosure is successful with probability r. The standard game corresponds to $(S, \mathcal{T}, p, c, 1)$.

An action profile $a \in A$ is said to be q-dominant if for each player i, $\{a_i\} = \arg\max_{a'_i \in A_i} [q' g_i(a'_i, \alpha_{-i}) + (1 - q') g_i(a'_i, \alpha_{-i})]$ holds for any $\alpha_{-i} \in \Delta(A_{-i})$ and any $q' \in [q, 1]$.

Proposition 9. *Fix a common-interest game S such that the best action profile a^* is strictly $(r - q)$-dominant and the dynamic game $(S, \mathcal{T}, p, c, r)$ is such that p is $(q - \varepsilon)_{i \in N}$-concentrated for some $\varepsilon > 0$. Then there exists a unique PBE, and in this PBE, each player i plays (a_i^*, not) at each $t \in \mathcal{T}_i$.*

The proof is given in the Appendix. The argument is similar to the one for the case with $r = 1$, in that we first show that playing $(a_i, \cdot)$ with $a_i \neq a_i^*$ is worse than playing (a_i^*, pay), and then show that (a_i^*, not) gives a higher payoff than (a_i^*, pay). We need an extra condition to ensure (a_i^*, pay) generates a high payoff when $r < 1$ because paying is less likely to affect the opponent's action when r is small. For (a_i^*, pay) to give rise to a higher payoff than $(a_i, \cdot)$ for $a_i \neq a_i^*$, it suffices that the probability of the opponent observing the player's action is high relative to the safety of a_i^*. This last condition is captured by $(r - q)$-dominance and $(q - \epsilon)$-concentration.

We note that the potential leader condition- one of the key conditions for the main analysis- is not relevant for the analysis in this section. The reason is that not observing the opponent's action is always on the path of any PBE if $r < 1$. This implies that at $r = 1$ there is a discontinuity of the set of the timing distributions inducing the unique PBE.

6 Conclusion

This paper studied games with private timing. These are games in which 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 perfect Bayesian equilibrium under asynchronicity and uncertainty of timing. We applied the model to a repeated prisoner's dilemma and showed that cooperation is a unique outcome when timing is private, with only a mild assumption on the conditional behavior at far distant future.

We focused on games that satisfy the pure Stackelberg property, but one could consider a wider class of component games. One prominent example 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.

There are numerous possibilities for future research in the framework of games with private monitoring. 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.[14] 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. For example, one could analyze the difference between possibly incorrect signals that are sent with probability one, and correct signals that are sent with probability less than one.

[14]We thank Drew Fudenberg for suggesting this possibility.

Fourth, 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.

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. and L. Samuelson (2006): *Repeated Games and Reputations: Long-Run Relationships*, 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.

A Appendix

A.1 Proof of Theorem 1

As noted in Remark 4-1, the sufficiency of the conditions 1 and 2(a) follow 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 $\inf_{\tilde{t}\in\mathcal{T}_{-i}}\{\text{Prob}^p(T_i \in E|T_{-i} = \tilde{t})|\text{Prob}^p(T_i \in E|T_{-i} = \tilde{t}) > 0\} = 0$ for every $E \subseteq T_i^<$, only (a_i^*, not) is played by each player i. The reason is that condition 2(b) implies that there exists $\bar{\epsilon} > 0$ such that for all $\epsilon < \bar{\epsilon}$, there exists a time $\tilde{t}$ for which $\text{Prob}^p(T_i \in E|T_{-i} = \tilde{t}) < \epsilon$. By taking $\bar{\epsilon} < \frac{c}{\max_{a\in A}(g_{-i}(a^*)-g_{-i}(a))}$, for all $\epsilon < \bar{\epsilon}$, it is strictly suboptimal for player $-i$ to play $(\cdot, \text{pay})$ since the expected payoff from playing such an action is at most $g_{-i}(a^*) - c$, while the expected payoff from playing a_{-i}^* is at least $g_{-i}(a^*) - \epsilon \cdot (\max_{a\in A}(g_{-i}(a^*) - g_{-i}(a)))$. Thus, conditional on having a move at $t \in T_i^<$ in any PBE, it is with positive probability that player i does not observe anything.

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 $i = 1, 2$, $\tilde{\tilde{\varepsilon}} = \min\{\varepsilon_i|\varepsilon_i > 0, i = 1, 2\}$[15], and $\varepsilon = \min\{\tilde{\tilde{\varepsilon}}, \tilde{\varepsilon}\}$. Consider the game in Figure 5, where M is such that $(1-\varepsilon)1 + \varepsilon(-M) < 0$. 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 show that we can construct a PBE in which B is played at time t^*.

[15]By convention, the minimum of an empty set is ∞.

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

$$A_j^i = \{t < t^* | -\text{Prob}^p(T_{-i} = t^* | T_i = t, T_{-i} \notin A_{j-1}^{-i})M$$

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

That is, A_j^i is the set of times before t^* such that it is a best response for player i to play $(A,\ \text{pay})$ if player $-i$ plays $(A,\ \text{pay})$ at times in A_{j-1}^{-i}, plays $(A,\ \text{not})$ at times in $(A_{j-1}^{-i})^c$ and $(B,\ \cdot)$ at time t^*.

Now, define $A^i := \lim_{j\to\infty} A_j^i$. 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 $j \in \mathbf{N}$. The latter follows from $\text{Prob}^p(T_{-i} = t^* | T_i = t, T_{-i} \notin A_{j-1}^{-i})$ 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^*. Consider the strategy profile in which, for each player $i = 1, 2$, $t \in \mathcal{T}_i$ and $h_{i,t} \in \mathcal{H}_{i,t}$,

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

We now check that each player takes a best response at each private history. 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. At time t^* both players play $(B,\ \text{pay})$ because M is chosen so that $(1 - \varepsilon)1 + \varepsilon(-M) < 0$.

Necessity of condition 2:

Suppose, for contradiction, that there is 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 given by Figure 5. Define:

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

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

We will show that for any $M \in [0, \infty)$, we can find $\bar{c} > 0$ such that when the component game S is the common interest game in Figure 5 and $c < \bar{c}$, there is a PBE σ of $(S, \mathcal{T}, p, c)$ such that the following hold. For each $t \in \mathcal{T}_{-i}$ and $h_{-i,t} \in \mathcal{H}_{-i,t}$,

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

where we abuse notation to express the pure strategy by identifying the action that is assigned probability 1 by σ_{-i} (we use the same abuse of notation in what follows) Also, for each $t \in \mathcal{T}_i$ and $h_{i,t} \in \mathcal{H}_{i,t}$,

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

Now we specify beliefs. For each private history of player $-i$, the strategy profile σ_i and Bayes rule completely pins down the probability distribution over histories, and we define player $-i$'s belief at each private history to be such a probability distribution. Also, except at times in $E \cup \tilde{E}_i$, for each private history of player i, the strategy profile σ_{-i} and Bayes rule completely pins down the probability distribution over histories, and we define player i's belief at each of those private histories to be such a probability distribution. For each time $t \in E \cup \tilde{E}_i$, if the private history is not $(\varnothing, \cdot)$, then again the strategy profile σ_{-i} and

Bayes rule completely pins down the probability distribution over histories, and we define player i's belief at each of those private histories to be such a probability distribution. If the private history is $(\varnothing, \cdot)$, 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 the probability distribution over the history that assigns probability 1 to $((t^*(t), B, \text{not})_{-i}, (t, B, \text{not})_i)$. Thus, in the off-path histories in which at times in $E \cup \tilde{E}_i$ player i does not observe (A, pay), she believes that $-i$ played (B, not). This procedure fixes a probability distribution over histories for any given private history of any player.

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 in which at times in $E \cup \tilde{E}_{-i}$ 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 $-i$, the payoff of playing (A, pay) at times in E_{-i} if there has been no observation 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 $c \in (0, \delta(M + 1))$ $-i$'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

$$c < \bar{c} := \min_i g_i^* - g_i^S. \tag{2}$$

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

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

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 $(\emptyset, \cdot)$ under $\sigma^*(c)$. Note that assumption 4 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 before which player i chooses to play (a_i^*, pay) at any private history $(\emptyset, \cdot)$ 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 $\tau \leq \bar{t}(c)$ 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 $\bar{t}(c) = \max_i \sup \mathcal{T}_i$. Thus, for each player i,

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

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 (2) and equation (2), $\exists \bar{t}_i^1 \in \mathcal{T}_i \cap (\bar{\bar{t}}(c), \infty)$ such that $\forall t \in (\bar{\bar{t}}(c), \bar{t}_i^1] \cap \mathcal{T}_i$,

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

Define $\bar{p}_i(\bar{t}) = p_i(\bar{t}, \bar{\bar{t}}(c), \bar{\bar{t}}(c), \bar{t})$. Also from assumption 1 (2) and equation (3), there exists $\bar{t}_i^2 \in \mathcal{T}_i$ such that $\forall t \leq \bar{t}_i^2$, $t \in \mathcal{T}_i$,

$$1 - p_i(t, \bar{\bar{t}}(c), t, \bar{t}(c)) = p_i(t, \bar{\bar{t}}(c), \bar{\bar{t}}(c), \bar{t}(c)) - p_i(t, \bar{\bar{t}}(c), t, \bar{t}(c)) \leq p_i(t, \bar{\bar{t}}(c), \bar{\bar{t}}(c), t) < \delta$$

with $\delta := \frac{\left(g_i^* - c\right)(1 - \bar{p}_i(\bar{t}_i^1)) + \underline{g}_i \bar{p}_i(\bar{t}_i^1) - g_i^S}{g_i^* - g_i^S}$, which is strictly positive by (4). This implies that, $\forall t \leq \bar{t}_i^2$, $t \in \mathcal{T}_i$,

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

Since $p_i(t, \bar{\bar{t}}(c), \bar{\bar{t}}(c), t) \leq \bar{p}_i(\bar{t}_i^1)$ $\forall t \in (\bar{\bar{t}}(c), \bar{t}_i^1] \cap \mathcal{T}_i$ (by assumption 1 (1)), this implies that for all $t \in (\bar{\bar{t}}(c), \min\{\bar{t}_i^1, \bar{t}_i^2\}] \cap \mathcal{T}_i$,

$$p_i(t, \bar{\bar{t}}(c), t, \bar{t}(c))g_i^S + (1 - p_i(t, \bar{\bar{t}}(c), t, \bar{t}(c)))g_i^* < (g_i^* - c)\left(1 - p_i(t, \bar{\bar{t}}(c), \bar{\bar{t}}(c), t)\right) + \underline{g}_i p_i(t, \bar{\bar{t}}(c), \bar{\bar{t}}(c), t). \tag{5}$$

Note that, if $\bar{\bar{t}}_i(c) < \bar{\bar{t}}_{-i}(c)$, an upper bound on i's payoff from (a_i^*, not) at $t \in (\bar{\bar{t}}(c), \min\{\bar{t}_i^1, \bar{t}_i^2, \bar{\bar{t}}_{-i}(c)\}] \cap \mathcal{T}_i$ is given by

$$p_i(t, t, t, \bar{t}(c))g_i^S + (1 - p_i(t, t, t, \bar{t}(c)))g_i^*. \tag{6}$$

If $\bar{\bar{t}}_i(c) = \bar{\bar{t}}_{-i}(c)$, then an upper bound on i's payoff from (a_i^*, not) at $t \in (\bar{\bar{t}}(c), \min\{\bar{t}_i^1, \bar{t}_i^2\}] \cap \mathcal{T}_i$ is

$$p_i(t, \bar{\bar{t}}(c), \bar{\bar{t}}(c), \bar{t}(c))g_i^S + (1 - p_i(t, \bar{\bar{t}}(c), \bar{\bar{t}}(c), \bar{t}(c)))g_i^*. \tag{7}$$

Both (6) and (7) are no more than the left hand side of (5). Thus, there exists $\tau > 0$ such that the left hand side of (5) is an upper bound on the payoff of playing (a_i^*, not) at $t \in \mathcal{T}_i \cap (\bar{\bar{t}}(c), \bar{\bar{t}}(c) + \tau]$.

Similarly, if $\bar{\bar{t}}_i(c) < \bar{\bar{t}}_{-i}(c)$, then a lower bound on i's payoff from (a_i^*, pay) is $(g_i^* - c)$ at $t \in (\bar{\bar{t}}(c), \min\{\bar{t}_i^1, \bar{t}_i^2, \bar{\bar{t}}_{-i}(c)\}] \cap \mathcal{T}_i$ which is greater than the right hand side of (5). If $\bar{\bar{t}}_i(c) = \bar{\bar{t}}_{-i}(c)$ a lower bound on i's payoff from (a_i^*, pay) is the right hand side of (5). Thus, again, there exists $\tau > 0$ such that the right hand side of (5) 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 (4), there exist $\check{t}_i^1 \in \mathbb{R} \cap (-\infty, \bar{\tau})$ such that for all $t \in \mathcal{T}_i \cap (-\infty, \check{t}_i^1]$ we have $p_i(t, -\infty, t, \bar{\tau}) > 0$. Therefore, since $p_i(t, -\infty, t, \bar{\tau}) \leq p_i(t, -\infty, t, \bar{t}(c))$ for each t, 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 (2), for any $\beta > 0$, there is $\check{\check{t}}_i^2 \in \mathcal{T}_i$ such that $\forall t \in (-\infty, \check{\check{t}}_i^2] \cap \mathcal{T}_i$,

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

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

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

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 $\bar{\bar{t}}(c) = -\infty$.

Step 2: Showing $\bar{\bar{t}}(c) < \bar{t}(c)$. Here we show that $\bar{\bar{t}}(c) < \bar{t}(c)$. To see this, suppose on the contrary that $\bar{\bar{t}}(c) = \bar{t}(c)$. If player i obtains an opportunity before $\bar{\bar{t}}(c)$, then she plays (a_i^*, pay) under $\sigma^*(c)$. From inequality (2), there exists $\varepsilon > 0$ such that $(1 - \varepsilon)(g_i^* - c) + \varepsilon \underline{g}_i > g_i^S$. For such ε, assumption 1 (2) implies that, for each $i \in \{1, 2\}$, there is $\bar{t}_i \in \mathcal{T}_i \cap (\bar{\bar{t}}(c), \infty)$ such that $\forall t \in (\bar{\bar{t}}(c), \bar{t}_i] \cap \mathcal{T}_i \neq \emptyset$, $p_i(t, \bar{\bar{t}}(c), \bar{\bar{t}}(c), t) < \varepsilon$ and therefore,

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

for $i \in \{1, 2\}$. Hence, there exists $\tau > 0$ such that each player i would choose $(a_i^*, \cdot)$ at $t \in (\bar{t}(c), \bar{t}(c) + \tau)$. This contradicts the assumption that $\bar{t}(c)$ is the supremum time before which every player i chooses $(a_i^*, \cdot)$ at any private history $(\emptyset, \cdot)$. Therefore, we must have $\bar{\bar{t}}(c) < \bar{t}(c)$.

Step 3: Defining k, $\hat{t}(c)$, and $a(c)$. 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 two 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

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

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 $\sigma^*(c)$. Since assumption 1 (1) implies $p_k(t, \bar{\bar{t}}(c), \bar{\bar{t}}(c), t)$ is non-decreasing in t, for all $t \in \mathcal{T}_k \cap (\hat{t}(c), \infty)$ we have $p_k(t, \bar{\bar{t}}(c), \bar{\bar{t}}(c), t) \geq \frac{g_k^* - c - g_k^S}{g_k^* - \underline{g}_k} \equiv a(c)$. The function $a(c)$ is decreasing in c.

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

Step 4. Case a). <u>Suppose $\bar{\bar{t}}_{-j}(c) > \bar{\bar{t}}_j(c)$.</u>

Note that

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

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

$$p_j(\tilde{t}_j, \tilde{t}_j, \tilde{t}_j, \hat{t}(c)) g_j^S + (1 - p_j(\tilde{t}_j, \tilde{t}_j, \tilde{t}_j, \hat{t}(c))) g_j^* \geq g_j^* - c. \tag{8}$$

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 $\tilde{t}_j \in \mathcal{T}_j$ and for every $\tilde{\varepsilon} > 0$ there is $t_k \in \mathcal{T}_k \cap (\hat{t}(c), \infty)$ such that $p_j(\tilde{t}_j, \tilde{t}_j, \tilde{t}_j, \hat{t}(c)) \geq p_j(\tilde{t}_j, \tilde{t}_j, \tilde{t}_j, t_k) - \tilde{\varepsilon}$. By assumption 1 (3), there exists $\alpha > 0$ such that $p_j(\tilde{t}_j, \tilde{t}_j, \tilde{t}_j, t_k) \geq \alpha p_k(t_k, \tilde{t}_j, \tilde{t}_j, t_k)$ which yields $p_j(\tilde{t}_j, \tilde{t}_j, \tilde{t}_j, \hat{t}(c)) \geq \alpha p_k(t_k, \tilde{t}_j, \tilde{t}_j, t_k) - \tilde{\varepsilon}$. By the continuity of the probability with respect to increasing sets, $\forall t_k \in \mathcal{T}_k \cap (\hat{t}(c), \infty)$, $\forall \varepsilon > 0$,

$\exists \tau > \bar{\bar{t}}(c)$ such that $\forall \tilde{t}_j \in \mathcal{T}_j \cap [\bar{\bar{t}}(c), \tau]$,

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

These observations imply that for any $\varepsilon, \tilde{\varepsilon} > 0$ there are $\tau > \bar{\bar{t}}(c)$ such that $\forall t_k \in \mathcal{T}_k \cap (\hat{t}(c), \infty)$ and $\forall \tilde{t}_j \in \mathcal{T}_j \cap [\bar{\bar{t}}(c), \tau]$

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

and (8) hold. However, there exists $\bar{\epsilon} > 0$ such that for all $\tilde{\varepsilon} < \bar{\epsilon}$ and $\varepsilon < \bar{\epsilon}$ there exists $\tau > \bar{\bar{t}}(c)$ and $\tilde{c} > 0$ such that for all $c < \tilde{c}$ and $\forall \tilde{t}_j \in \mathcal{T}_j \cap [\bar{\bar{t}}(c), \tau]$, we have

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

as $a(c) > a(g_k^* - g_k^S) = 0$ and $a(c)$ is decreasing in c. This contradicts (8).

Step 4. Case b). Suppose $\bar{\bar{t}}_j(c) = \bar{\bar{t}}_{-j}(c)$. By the definition of $\bar{\bar{t}}(c)$, for every $\tau \in (\bar{\bar{t}}, \bar{\bar{t}}_{-j}(c)]$ there is $\tilde{t}_j \in \mathcal{T}_j \cap [\bar{\bar{t}}(c), \tau]$ such that

$$\begin{aligned} p_j(\tilde{t}_j, \bar{\bar{t}}(c), \bar{\bar{t}}(c), \hat{t}(c)) g_j^S + \left(1 - p_j(\tilde{t}_j, \bar{\bar{t}}(c), \bar{\bar{t}}(c), \hat{t}(c))\right) g_j^* \geq \\ \left(g_j^* - c\right) \left(1 - p_j(\tilde{t}_j, \bar{\bar{t}}(c), \bar{\bar{t}}(c), \tilde{t}_j)\right) + \underline{g}_i p_j(\tilde{t}_j, \bar{\bar{t}}(c), \bar{\bar{t}}(c), \tilde{t}_j). \end{aligned} \tag{9}$$

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 $\tilde{t}_j \in \mathcal{T}_j$ and for every $\tilde{\varepsilon} > 0$ there is $t_k \in \mathcal{T}_k$ such that $p_j(\tilde{t}_j, \bar{\bar{t}}(c), \bar{\bar{t}}(c), \hat{t}(c)) \geq p_j(\tilde{t}_j, \bar{\bar{t}}(c), \bar{\bar{t}}(c), t_k) - \tilde{\varepsilon}$. By assumption 1 (3), there exists $\alpha > 0$ such that $p_j(\tilde{t}_j, \bar{\bar{t}}(c), \bar{\bar{t}}(c), t_k) \geq \alpha p_k(t_k, \bar{\bar{t}}(c), \bar{\bar{t}}(c), t_k)$ which yields $p_j(\tilde{t}_j, \bar{\bar{t}}(c), \bar{\bar{t}}(c), \hat{t}(c)) \geq \alpha p_k(t_k, \bar{\bar{t}}(c), \bar{\bar{t}}(c), t_k) - \tilde{\varepsilon} \geq \alpha a(c) - \tilde{\varepsilon}$, for every $\tilde{t}_j \in \mathcal{T}_j$ and for some $t_k \in \mathcal{T}_k$.

By Assumption 1 (2), for every $\varepsilon > 0$ $\exists \bar{t}_k \in \mathcal{T}_k \cap (\bar{\bar{t}}(c), \infty)$ such that $\forall t \in [\bar{\bar{t}}(c), \bar{t}_k] \cap \mathcal{T}_k$, $p_j(t, \bar{\bar{t}}(c), \bar{\bar{t}}(c), t) < \varepsilon$.

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

and (9) hold.

However, there exists $\bar{\epsilon} > 0$ such that for all $\tilde{\varepsilon} < \bar{\epsilon}$ and $\varepsilon < \bar{\epsilon}$, there exists $\tilde{c} > 0$ such that for all $c < \tilde{c}$, we have

$$p_j(\tilde{t}_j, \bar{\bar{t}}(c), \bar{\bar{t}}(c), \hat{t}(c))g_j^S + \left(1 - p_j(\tilde{t}_j, \bar{\bar{t}}(c), \bar{\bar{t}}(c), \hat{t}(c))\right) g_j^* \leq (\alpha a(c) - \tilde{\varepsilon})\, g_j^S + \\ + (1 - (\alpha a(c) - \tilde{\varepsilon}))\, g_j^* < \left(g_j^* - c\right)(1 - \varepsilon) + \underline{g}_k \varepsilon \leq \\ \left(g_j^* - c\right)\left(1 - p_j(\tilde{t}_j, \bar{\bar{t}}(c), \bar{\bar{t}}(c), \tilde{t}_j)\right) + \underline{g}_k p_j(\tilde{t}_j, \bar{\bar{t}}(c), \bar{\bar{t}}(c), \tilde{t}_j),$$

as $a(c) > a(g_k^* - g_k^S) = 0$ and $a(c)$ is decreasing in c. This contradicts (9).

A.3 Proof of Theorem 3

As explained in the main text of the paper, the proof consists of five parts.

Part 1: Lemma

Lemma 10. *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 $\sigma^{n_k} \to \sigma$ pointwise as $k \to \infty$.*

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

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

as $k \to \infty$.

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 $x_k = \left((t_j^{x_k}, a_j^{x_k}, d_j^{x_k})_{j\in\{1,\dots,I\}}\right)$ 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 $N_i^{x_k} = \{j \in \{1, \dots, I\} | d_j = \text{pay and } t_j < t_i\}$ and let $h_i(x_k) =$

$\left(N_i^{x_k}, (t_j^{x_k}, a_j^{x_k}, d_j^{x_k})_{j \in N_i^{x_k}}\right) \in \mathcal{H}_{i,t_i^{x_k}}$. Define f_i^n for each $i \in \{1, \ldots, I\}$ $n \in \mathbb{N}$ 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 $\{\sigma^{n_k^1}\}_{k\in\mathbb{N}}$ of $\{\sigma^n\}_{n\in\mathbb{N}}$ 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 $i \in \{1, \ldots, I\}$. Now, recursively, given a sequence $\{\sigma^{n_k^m}\}_{k\in\mathbb{N}}$, there exists its subsequence $\{\sigma^{n_k^{m+1}}\}_{k\in\mathbb{N}}$ such that

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

for some value of $f_i(x_{m+1})$.

Now, consider the sequence $\{\sigma^{n_k^k}\}_{k=1}^{\infty}$. 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 $\{f_i^{n_k^{\tilde{k}}}(x_{\tilde{k}})\}_{k=\tilde{k}}^{\infty}$.

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 complete the proof.

□

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 $\mathcal{T} = \{t_k\}_{k\in\mathbb{Z}}$.[16] Let $\mathcal{T}^N := \{t_k\}_{k=-N}^{N}$.

Define

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

as the probability that the play times for all players are in $\mathcal{T}^N$.

Let $\underline{N}$ be the smallest integer N such that $\tilde{p}^N > 0$. By the definition of $\tilde{p}^N$, every player gets a move in $\mathcal{T}^N$. Also, by the definition of $\{t_k\}_{k\in\mathbb{Z}}$, $\underline{N} < \infty$ holds, and $\tilde{p}^N \to 1$ as $N \to \infty$.

For each $N \geq \underline{N}$, define the distribution of $\mathbf{t}$ 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 $\mathbf{t} \in \times_{i\in\{1,\ldots,I\}}\mathcal{T}^N$.

For each $N \geq \underline{N}$, the N'th approximating game of a private timing game $\Gamma = (S, \mathcal{T}, p, c)$ is defined by a 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 $h = (t_i, a_i, d_i)_{i\in\{1,\ldots,I\}}$ that is compatible with $h_{i,t}$ such that (i) $t_j \in \mathcal{T}^N$ for $j \in \{1, \ldots, I\}$, and (ii) $\text{Prob}^p(T_i = t_i \ \forall i \in \{1, \ldots, I\} | T_j = t_j \ \forall j \in \{1, \ldots, I\} \text{ s.t. } d_j = \text{pay and } t_j < t_i) > 0$.

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 for 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 utility in the N'th approximating game of strategy profile σ^N is denoted $u_i^N(\sigma^N)$.

Part 3: Defining 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, for each pair $(a_i, d_i) \in A_i \times \{\text{pay}, \text{not}\}$, $\sigma_i(h_{i,t})(a_i, d_i) \geq \varepsilon$ 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

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

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 $\varepsilon_n \to 0$ as $n \to \infty$ such that $\sigma^{\varepsilon_n} \to \sigma^*$ 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 11. *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 $\{\sigma^N\}_{\underline{N}}^{\infty}$, where $\sigma^N \in \Sigma^N$ for each N, of ε-constrained equilibria in the N'th approximation game. Such a sequence exists as each N'th approximating game is finite. Fix $\varepsilon \in (0, \bar{\varepsilon})$ and let $\{\sigma^N\}_{\underline{N}}^{\infty}$ be a sequence of ε-constrained equilibria.

We define the strategy profile $\tilde{\sigma}^N \in \Sigma$ of the original game *corresponding* to a strategy profile $\sigma^N \in \Sigma^N$ of the N'th approximation game 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 10, 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., $\hat{\sigma}_i^N(h_{i,t}) := \sigma_i(h_{i,t})$ for each history $h_{i,t} \in \mathcal{H}_{i,t}^N$. Note that $\sigma_i'^N$ is an ε-constrained strategy because σ_i' is an ε-constrained strategy.

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

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

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 Trembling Perfect Equilibria

Proposition 12. *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 $\varepsilon_n \to 0$ as $n \to \infty$. Proposition 11 implies that, for each $n \geq \underline{N}$, there exists an ε_n-constrained equilibrium $\sigma^{\varepsilon_n} \in \Sigma$. Lemma 10 then implies that there must be a convergent sub-sequence 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 13. *In every game $\Gamma = (S, \mathcal{T}, p, c)$, a PBE exists.*

$\square$

A.4 Proof of Proposition 3

Proof. Fix $T < \infty$. We use a mathematical induction to prove the result. Consider the time interval $[k-1, k)$ with $1 \leq k \leq T$ and suppose that for all $l > k$, it is true that, given the history that only (C_1, C_2) has been observed for all the realized moving opportunities in $[0, l-1)$, the only action profile that can be played in the

	C_2	D_2
C_1	$\frac{e^{-\rho(k-t_1)}}{1-e^{-\rho}}, \frac{e^{-\rho(k-t_2)}}{1-e^{-\rho}}$	$-e^{-\rho(k-t_1)}s_1, e^{-\rho(k-t_2)}d_2$
D_1	$e^{-\rho(k-t_1)}d_1, -e^{-\rho(k-t_2)}s_2$	$0, 0$

Figure 12: Continuation payoffs at times $t_1, t_2 \in [k-1, k)$ after only (C_1, C_2) has been observed

	C_2	D_2
C_1	$e^{-\rho(k-t_1)}, e^{-\rho(k-t_2)}$	$-e^{-\rho(k-t_1)}s_1, e^{-\rho(k-t_2)}d_2$
D_1	$e^{-\rho(k-t_1)}d_1, -e^{-\rho(k-t_2)}s_2$	$0, 0$

Figure 13: Continuation payoffs at times $t_1, t_2 \in [k-1, k)$ after at least one D_i is observed

time interval $[l-1, l)$ in any PBE in the class of T-grim trigger is (C_1, C_2). Also suppose that, under all other histories, the only action profile that can be played in the time interval $[l-1, l)$ in any PBE in the class of T-grim trigger is (D_1, D_2).

We first prove that, given the history that only (C_1, C_2) has been observed for all the realized moving opportunities in $[0, k-1)$, the only action profile that can be played in the time interval $[k-1, k)$ in any PBE is (C_1, C_2). For that purpose, note that, given the induction hypotheses, the payoff matrix that players face at their respective moving times $t_1, t_2 \in [k-1, k)$ are given by the one in Figure 13. Note that, if $\frac{e^{-\rho(k-t_i)}}{1-e^{-\rho}} > e^{-\rho(k-t_1)}d_i$ for each player i, this game is a common-interest game. This holds if and only if $e^{-\rho} > 1 - \frac{1}{d_i}$. Since $d_i > 1$, there exists $\bar{\rho} > 0$ such that for all $\rho < \bar{\rho}$, the component game is a common-interest game. Thus, since $p_k \in D$, Theorem 1 implies that the only action that can be played in any PBE is (C_1, C_2).

We next prove that, given the history that it is not true that only (C_1, C_2) has been observed for all the realized moving opportunities in $[0, k-1)$, the only action profile that can be played in the time interval $[k-1, k)$ in any PBE is (D_1, D_2). For that purpose, note that, given the induction hypotheses, the payoff matrix that players face at their respective moving times $t_1, t_2 \in [k-1, k)$ are given by the one in Figure 13. Note that action D_i is a strictly dominant action for each i. Thus, Proposition 7 implies that the only action that can be played in any PBE is (D_1, D_2). □

A.5 Proof of Proposition 8

Proof. Define $t^* \in \mathbb{R}$ to be the supremum of $t \in \mathcal{T}_1$ that satisfy

$$\text{Prob}^p(T_2 < t) \cdot (-\beta_1) + (1 - \text{Prob}^p(T_2 < t)) \cdot 1 \geq c.$$

Notice that t^* is finite because p is such that player 1 can be arbitrarily early.

For each $t \in \mathcal{T}_1$ and $h_{1,t} \in \mathcal{H}_{1,t}$,

$$\sigma_1(h_{1,t}) := \begin{cases} (A_1, \text{pay}) & \text{if } t \leq t^* \text{ and } h_{1,t} = (\emptyset, \cdot) \\ (B_1, \text{not}) & \text{if } t > t^* \text{ and } h_{1,t} = (\emptyset, \cdot) \\ (B_1, \text{not}) & \text{if } h_{1,t} = (\{2\}, (t', a_2, \text{pay})) \text{ for some } a_2 \in \{A_2, B_2\} \text{ and } t' \in \mathcal{T}_2 \end{cases},$$

where we abuse notation to express the pure strategy by identifying the action that is assigned probability 1 by σ_1 (we use the same abuse of notation in what follows). Also, for each $t \in \mathcal{T}_2$ and $h_{2,t} \in \mathcal{H}_{2,t}$,

$$\sigma_2(h_{2,t}) := \begin{cases} (B_2, \text{not}) & \text{if } h_{2,t} = (\emptyset, \cdot) \\ (A_2, \text{not}) & \text{if } h_{2,t} = (\{1\}, (t', A_1, \text{pay})) \text{ for some } t' \in \mathcal{T}_1 \\ (B_2, \text{not}) & \text{if } h_{2,t} = (\{1\}, (t', B_1, \text{pay})) \text{ for some } t' \in \mathcal{T}_1 \end{cases}.$$

Note that since player 1 can be arbitrarily late, the two conditions in the definition of PBE completely specify all the off-path beliefs. Since t^* is finite, under this strategy (A_1, pay) is played with ex ante strictly positive probability. Thus, once we show incentive compatibility of this strategy profile, the proof is complete.

Now we check that each player i takes a best response at each $t \in \mathcal{T}_i$. First, under the history of the form $h_{i,t} = (\{-i\}, (t', a_{-i}, \text{pay}))$ for some $a_{-i} \in \{A_{-i}, B_{-i}\}$ and $t' \in \mathcal{T}_{-i}$, it is straightforward that the players choose a best response. Thus, in what follows, we only consider private history $(\emptyset, \cdot)$. Specifically, for each player, we consider (i) the case in which the player receives an opportunity at $t \leq t^*$ and (ii) the case in which the player receives an opportunity at $t > t^*$.

Player 1's incentive:

Case (i)

First, suppose that player 1 receives an opportunity at $t \leq t^*$. If 1 plays

(A_1, pay), then her expected payoff is, by asynchronicity and independence,

$$[\text{Prob}^p(T_2 < t) \cdot (-\beta_1) + (1 - \text{Prob}^p(T_2 < t)) \cdot 1] - c$$

$$\geq [\text{Prob}^p(T_2 < t^*) \cdot (-\beta_1) + (1 - \text{Prob}^p(T_2 < t^*)) \cdot 1] - c \geq 0,$$

by the definition of t^*. If 1 plays (A_1, not), then her payoff is $-\beta_1$. If 1 plays (B_1, pay), then her payoff is $0 - c = c$. Finally, if 1 plays (B_1, not), then her payoff is 0. Overall, it is a best response to choose (A_1, pay).

Case (ii)

Second, suppose that player 1 receives an opportunity at $t > t^*$. If 1 plays (A_1, pay), then her expected payoff is, again by asynchronicity and independence,

$$[\text{Prob}^p(T_2 < t) \cdot (-\beta_1) + (1 - \text{Prob}^p(T_2 < t)) \cdot 1] - c < 0,$$

by the definition of t^*. The payoffs to other actions are the same as in the case of $t \leq t^*$. Hence, it is a best response to choose (B_1, not).

Player 2's incentive:

Case (i)

First, suppose that player 2 receives an opportunity at $t \leq t^*$. In this case, player 2's belief assigns probability 0 to player 1 having moved because of Bayes rule and the assumption that player 1 can be arbitrarily late implies that there exists $t \in \mathcal{T}_1$ such that

$$-\beta_1 \text{Prob}^p(T_2 < t) + 1(1 - \text{Prob}^p(T_2 < t)) < c,$$

so the history $(\emptyset, \cdot)$ is on the path of play. Thus, if 2 plays (A_2, pay), then his payoff is $-\alpha_2 - c$. If 2 plays (A_2, not), then his payoff is, by asynchronicity and independence,

$$\text{Prob}^p(t < T_1 \leq t^* | t < T_1) \cdot 1 + \text{Prob}^p(t^* < T_1 | t < T_1) \cdot (-\alpha_2).$$

If 2 plays (B_2, pay), then his payoff is $0 - c = -c$. Finally, if 2 plays (B_2, not), then his payoff is, by independence,

$$\text{Prob}^p(t < T_1 \leq t^*|t < T_1) \cdot (1 - \beta_2) + \text{Prob}^p(t^* < T_1|t < T_1) \cdot 0.$$

Since this last expression is nonnegative because $1 - \beta_2 \geq 0$, it suffices to show that the payoff from (A_2, not) is no more than the one from (B_2, not). To see this, note that this condition is equivalent to

$$\text{Prob}^p(t < T_1 \leq t^*|t < T_1) \cdot \beta_2 \leq \text{Prob}^p(t^* < T_1|t < T_1) \cdot \alpha_2. \tag{10}$$

Since $\text{Prob}^p(t < T_1 \leq t^*|t < T_1)$ is nondecreasing in t and $\text{Prob}^p(t^* < T_1|t < T_1)$ is nonincreasing in t, (10) holds for all $t \leq t^*$ if

$$\lim_{t \to -\infty} \text{Prob}^p(t < T_1 \leq t^*|t < T_1) \cdot \beta_2 \leq \lim_{t \to -\infty} \text{Prob}^p(t^* < T_1|t < T_1) \cdot \alpha_2,$$

which is equivalent to

$$\text{Prob}^p(T_1 \leq t^*) \cdot \beta_2 \leq \text{Prob}^p(t^* < T_1) \cdot \alpha_2,$$

or

$$\text{Prob}^p(T_1 \leq t^*) \leq \frac{\alpha_2}{\beta_1 + \alpha_2}.$$

Now, note that by the definition of t^*, we have

$$-\beta_1 \text{Prob}^p(T_2 < t^*) + 1(1 - \text{Prob}^p(T_2 < t^*)) \geq c,$$

or

$$\text{Prob}^p(T_2 < t^*) \leq \frac{1 - c}{1 + \beta_1}.$$

Since player 2 is at most Δ-lagged, this implies

$$\text{Prob}^p(T_1 < t^*) \leq \frac{1 - c}{1 + \beta_1} + \Delta.$$

Moreover, since player 1 has no more than an ϵ-mass, we have that

$$\text{Prob}^p(T_1 \leq t^*) \leq \frac{1-c}{1+\beta_1} + \Delta + \epsilon.$$

By the assumption in the statement of the proposition, we then have that

$$\text{Prob}^p(T_1 \leq t^*) \leq \frac{\alpha_2}{\beta_1 + \alpha_2},$$

showing that the payoff from (A_2, not) is no more than the one from (B_2, not). This completes the proof for this case.

Case (ii)

Second, suppose that player 2 receives an opportunity at $t > t^*$. Bayes rule and the assumption that 1 can be arbitrarily late imply that player 2's belief assigns probability 1 to player 1 playing (B_1, not). Thus, if 2 plays (A_2, pay), then his payoff is $-\alpha_2 - c$. If 2 plays (A_2, not), then his payoff $-\alpha_2$. If 2 plays (B_2, pay), then his payoff $0 - c = -c$. Finally, if 2 plays (B_2, not), then his payoff 0. Overall, playing (B_2, not) is a best response.

Since we have examined the incentives at all the private histories, the proof is now complete. □

A.6 Proof for Proposition 9

Proof. **Step 1:**

Step 1-1 Fix a common interest game such that profile a^* is strictly $(r-q)$-dominant and a time structure p that is $(q-\epsilon)_{i \in N}$-concentrated. Fix a PBE of $(S, \mathcal{T}, p, c, r)$ and 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 $t^* := \inf_{t \in \bigcup_{i \in N} N_i(a^*)} t$.

Step 1-2: By the definition of q-concentration, there must exist $i \in N$ and $t' > t^*$ such that for $j \neq i$ and $t \in [t^*, t'] \cap \mathcal{T}_i$, $\mathbb{P}(t^* \leq T_j \leq t | T_i = t) < q - \epsilon$. If player i chooses (a_i^*, pay) at time t player $-i$ responds with a_{-i}^* with probability at least $r - q + \varepsilon$. Because the action profile a^* is strictly $(r-q)$-dominant the payoff of playing a_i^* is strictly above the payoff from any other action if $-i$ plays a_{-i}^* with probability at least $(r-q)$. Thus, for small enough c player i prefers to play a_i^* at time 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 u_j^{SS}. 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.

Step 2:

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. As we have shown above, the outcome under σ^* must be a^*, so i's payoff from σ^* is $u_i^* - c$. 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 $(a_i^*, \cdot)$. Hence the outcome under $(\sigma_i', \sigma_{-i}^*)$ must be a^*. Hence the payoff from $(\sigma_i', \sigma_{-i}^*)$ is u_i^*, so the deviation is profitable. This is a contradiction to the assumption that σ^* is a PBE. Hence there is no time at which any player pays the disclosure cost.

Step 3:

The proposed equilibrium is indeed a PBE. At every history each player i believes that the opposing player has played according to a^* and best responds by choosing (a_i^*, not). $\square$