Online Supplementary Appendix to: Games with Private Timing

Yuichiro Kamada[†] Sofia Moroni[‡]

July 3, 2018

B Additional Discussions

In this appendix, we provide additional discussions. The proofs of the results stated in this section are provided in Appendix C.

B.1 A Probabilistic Disclosure Model

In this subsection, we consider the possibility that disclosure of an 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, $\{a_i\} = \arg\max_{a'_i \in A_i} [q' g_i(a'_i, a_{-i}) + (1-q') g_i(a'_i, \alpha_{-i})]$ holds for any $\alpha_{-i} \in \Delta(A_{-i})$ and any $q' \in [q, 1]$, where the domain of g_i is extended in a natural manner to encompass mixed actions.

Proposition 10. *Fix a common interest component game S such that there is $q > 0$ 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 $(1+\varepsilon-q)$-dispersed for some $\varepsilon > 0$.*

[†]Haas School of Business, University of California Berkeley, 2220 Piedmont Avenue, Berkeley, CA 94720-1900, USA, 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`

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

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 riskiness of a_i^*. This last condition is captured by $(r - q)$-dominance and $(1 + \varepsilon - q)$-dispersion.

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 lack of upper hemicontinuity of the set of timing distributions inducing the unique PBE.

B.2 Multiplicity in Component Games without the Pure Stackelberg Property

In the main sections, we focused on games with the pure Stackelberg property and showed uniqueness of a PBE under certain assumptions. In order to understand the role of the pure Stackelberg property, here we consider examples of component games without the pure Stackelberg property. Although a full characterization of the set of PBE for those games is beyond the scope of this paper, analyzing those examples helps us understand the role of the pure Stackelberg property as the dynamic games in the examples involve multiple PBE. First, we provide an example of a dynamic game that has multiple PBE, where the Stackelberg action in the component game is mixed.

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

Consider the two-player component game in Figure 5 and an independent timing distribution such that $\mathcal{T}_1$ is the set of odd natural numbers and $\mathcal{T}_2$ is the set of even natural numbers. The game has a symmetric mixed Stackelberg strategy that

	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 5: Mixed Stackelberg leads to multiplicity

involves mixing with probability 1/2 on actions B and C, while the pure Stackelberg action of each player puts probability 1 on A. Corresponding to these two, 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 which avoided this type of complexity. □

The pure Stackelberg property also implies that the Nash equilibrium a^i in 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 6 with an independent timing distribution such that $\mathcal{T}_1$ is the set of odd natural numbers and $\mathcal{T}_2$ is the set of even natural numbers. 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

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

Figure 6: Multiple best responses to the Stackelberg action

without observation, plays (S, pay). If player i observes the opponent's action B_1 or B_2 before she moves, then she takes the (unique) static best response. If she observes S, then she takes B_1.

In the other PBE, each player i, upon receiving the chance to move without observation, plays (B_2, pay). If player i observes the opponent's action B_1 or B_2 before she moves, then she takes the (unique) static best response. If she observes S, then she takes B_2.

To formalize the issue with this example, let $a_i^* \in \arg\max_{a_i \in A_i} \max_{a_{-i} \in BR(a_i)} g_i(a_i, a_{-i})$. The problem with the above example is that there exist $a'_{-i} \in BR(a_i^*)$, $a_i \in A_i$ 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 fact that $-i$ has a best response to S that gives i a worse payoff than (B_1, B_1) or (B_2, B_2). This is why in our main analysis we focused on the case in which $g_i(a_i^*, a'_{-i}) > g_i(a_i, a''_{-i})$ for every $a'_{-i} \in BR(a_i^*)$, $a_i \in A_i$ and $a''_{-i} \in BR(a_i)$. That is, we considered the case in which each i has an action that guarantees herself a better payoff than any other of her actions does when the opponent best-responds. □

B.3 Sense of Calendar Time

Example 3 demonstrates why we need the dispersed potential moves condition to prove uniqueness. In that example, 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 calendar time.

More specifically, consider a two-player 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

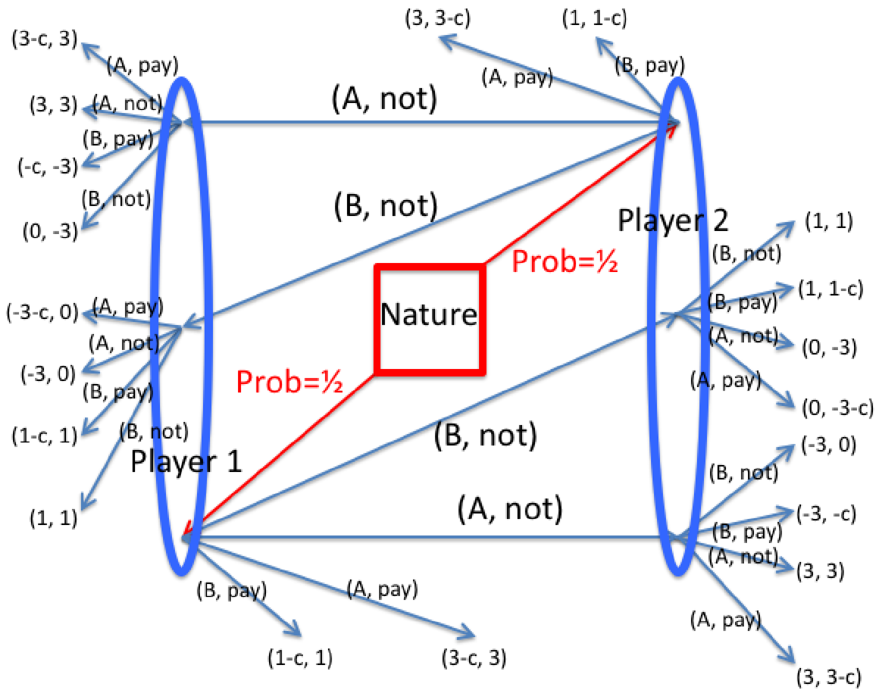

Figure 7: Extensive-Form of the Game with No Sense of Calendar Time: Nature moves first. Each player i cannot distinguish among the three possible histories (i being the first mover, $-i$ played A and did not disclose his action, and $-i$ played B and did not disclose his action). We omit the actions after the first-mover's payment.

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 in Section 3.1. Figure 7 shows the extensive form of this game with payoffs at each terminal node. As in Figure 2, 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 of a disclosure, while the other is that each player plays (B, not) under no observation of a disclosure. The second strategy profile is a PBE because if one follows it, the expected payoff is 1, while if she deviates to play (A, pay), the expected payoff reduces to $\frac{1}{2} \cdot 3 + \frac{1}{2} \cdot (-3) - c = -c$.

B.4 Various Component Games

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

B.4.1 Constant-Sum Games

When the component game is a two-player constant-sum game, we show that no player pays the disclosure cost.

Proposition 11. *Suppose that the component game S is a two-player constant-sum game. Then, for any dynamic game $(S, \mathcal{T}, p, c)$, for each i and $t \in \mathcal{T}_i$, i assigns probability zero to (a_i, pay) for any $a_i \in A_i$ at any history at any realization of T_i in any PBE. Moreover, the probability distribution over component-game action profiles under any PBE is a correlated equilibrium of S.*

The intuition is simple: For each player, one available strategy is to play a minmax strategy in the component game and never pays. This provides a lower bound of any PBE payoff, but the sum of such lower bounds is in fact also the maximum the players can achieve due to the constant-sum assumption. Hence, the lower bound is actually the unique PBE payoff, implying that no player pays in any PBE. Correlation comes from the possibility of a correlated timing distribution. Since there may be correlation in the players' strategies in equilibrium depending on the moving times, the class of private-timing constant-sum games is not equivalent to its simultaneous-move counterpart.

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

Example 7. Consider the three-player component game S with the payoff matrix given by Figure 8. Let $\mathcal{T} = \{1, 2, 3\}$ and $p(1, 2, 3) = 1$. That is, with probability one, 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,

A_3:

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

B_3:

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

Figure 8: A Three-Player Constant-Sum Game

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 distribution 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 entire component game is a constant-sum game. $\square$

B.4.2 Dominance Games

If the component game features a dominant action for each player, then each player plays their dominant action while not paying the disclosure cost.

Proposition 12. *Suppose that in the two-player component game S, each player i has a strictly dominant action a_i^D. Then, for any dynamic game $(S, \mathcal{T}, p, c)$, in any PBE, each player i plays (a_i^D, not) under any history.*

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

B.4.3 Reputation Games

Here we consider the class of "reputation games" as in Figure 9.[29] 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 aim to give a sufficient condition on the timing distribution p such that there exists a PBE in which, if player 1 receives her opportunity early

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

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

Figure 9: Reputation Game

then she "teaches" player 2 that she has taken A_1. For this purpose, we define a few conditions on p.

The timing distribution p is such that **player 1 can be arbitrarily early** if for every $\varepsilon > 0$, there exists $t \in \mathcal{T}_1$ such that $\text{Prob}^p(T_2 < t) < \varepsilon$. Similarly, the time distribution p is such that **player 1 can be arbitrarily late** if for every $\varepsilon > 0$, there exists $t \in \mathcal{T}_1$ such that $\text{Prob}^p(t < T_2) < \varepsilon$. This last condition holds, for example, if the potential leader condition holds for player 2.[30]

Let $\Delta(p)$ be defined by $\Delta(p) = \sup_{t \in \mathbb{R}} (\text{Prob}^p(T_1 \leq t) - \text{Prob}^p(T_2 < t))$. This is a measure of the lag of player 2's moving times relative to player 1's.

Proposition 13. *Suppose that S is given by the payoff matrix as in Figure 9, and p is asynchronous, independent, and such that player 1 can be both arbitrarily early and arbitrarily late. Suppose also that*

$$\frac{1-c}{1+\beta_1} + \Delta(p) \leq \frac{\alpha_2}{\beta_2 + \alpha_2}.$$

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 maximum probability of moving at a given time 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.[31] Note that the condition implies that a low disclosure cost may prevent (A_1, pay) from being a PBE outcome. The reason is that, if the cost was too low, then player 1 would be incentivized to play (A_1, pay)

[30]The potential leader condition says that $\text{Prob}^p(T_1 > t | T_2 = t) > 0$ for every $t \in \mathcal{T}_2$. Let $t'' \in \mathcal{T}_2$ be such that $\text{Prob}^p(T_2 > t'') < \varepsilon$. Now, since $\text{Prob}^p(T_1 > t'', T_2 = t'') > 0$, there is $t' \in \mathcal{T}_1$ with $t' > t''$ such that $\text{Prob}^p(T_1 = t') > 0$. We conclude that $\text{Prob}^p(T_2 > t') \leq \text{Prob}^p(T_2 > t'') < \varepsilon$.

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

at too many realizations of her moving time, which discourages player 2 to take B_2 early in the game (which would, in turn, discourage player 1 from paying). A similar reasoning explains why it is easier to satisfy the condition when the lag is small: If player 1 is more likely to play before 2 then 1 would choose (A_1, pay) too often.

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, in contrast, 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, a 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 a 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 the 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 constructed PBE as $c \to 0$, $\sup_{t \in \mathcal{T}} \max_{i \in \{1,2\}} |\text{Prob}^p(T_i = t)| \to 0$, and $\Delta(p) \to 0$ (i.e., the timing distribution approaching the "uniform distribution.") is $\left(\frac{1}{2(1+\beta_1)}, \frac{2+2\beta_1-\beta_2}{2(1+\beta_1)^2}\right)$, 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 PBE because player 2 never plays A_2 unless he become sure that 1 has played A_1.

The implication of Proposition 13 is that the result in Online Appendix B.4.2 does not extend to iterated dominance. The reason for the difference is that, one player's static best response depends on the other player's choice (before deletion of the dominated action of the latter player), and the latter player can commit to a dominated action to incentivize the former player to play a particular action.

θ_A:

	A	B
A	α, α	$0, 0$
B	$0, 0$	$1, 1$

θ_B:

	A	B
A	$1, 1$	$0, 0$
B	$0, 0$	α, α

Figure 10: Two Possible Games

B.5 Incomplete Information

In the main text, we considered the situation in which the component game is common knowledge among players from the beginning of the dynamic game. This in particular implies that the best action profile is common knowledge in the case where the component game is a common interest game, which helped to implement the contagion argument in Step 1 of the Proof Sketch for Proposition 1. In this section, we consider a possibility of incomplete information about the component game in a simple setting, and seek for a sufficient condition to guarantee uniqueness of a PBE.

Specifically, there are two possible component games, $\theta = \theta_A, \theta_B$, as in Figure 10 with $\alpha > 1$. Observe that, in either game, (A, A) and (B, B) are strict Nash equilibria, but only one of them gives the payoff of α to each player, and which action profile gives rise to the payoff α depends on the realized game. Note that action $a \in \{A, B\}$ is $\frac{1}{1+\alpha}$-dominant in game θ_a.

To model the knowledge structure, we suppose that there is a finite state space Ω over which information partitions P_1 and P_2 of two players and a probability distribution f are defined. There exists a function $\bar{\theta} : \Omega \to \{\theta_A, \theta_B\}$ such that, if the state is ω, then the realized game is $\bar{\theta}(\omega)$. Before the dynamic game starts, each player i is informed of a cell of the partition $g \in P_i$ with probability $\sum_{\omega \in g} f(\omega)$.

We assume the following genericity condition: For each player $i \in \{1, 2\}$, for any $g \in P_i$, $\sum_{\omega \in g, \bar{\theta}(\omega)=\theta_A} f(\omega) \neq \frac{1}{2} \cdot \sum_{\omega \in g} f(\omega)$. This assumption implies that, at each state $\omega \in \Omega$, $q_i \cdot \alpha + 1 \cdot (1 - q_i) \neq 1 \cdot q_i + \alpha \cdot (1 - q_i)$ holds where q_i is the probability i assigns to game θ_A before the dynamic game starts but after she observes the cell of her information partition. Hence, each player i strictly prefers to take some action over the other conditional on her signal, assuming that the opponent best-responds to her action. Let $q^i_a(g)$ denote the probability that player i believes that the game is θ_a at cell $g \in P_i$. Let $a_i(g) \in \{A, B\}$ denote the action a satisfying $q^i_a(g) > \frac{1}{2}$, which exists by assumption. That is, $(a_i(g), a_i(g))$ is the action profile that i strictly prefers conditional on observing the cell of her information partition g. Let $\bar{q}_i(g) = \max\{q^i_A(g), q^i_B(g)\}$.

We assume that player i always has some uncertainty about what player $-i$ believes to be the best action profile. Formally, we assume that there exists $\varepsilon > 0$ such that

$$\max_{i\in\{1,2\}, g_i\in P_i, a\in\{A,B\}, g_{-i}\in P_{-i}} \mathrm{Prob}^p\left(a_{-i}(g_{-i}) = a \middle| a_i(g_i) = a\right) < 1 - \varepsilon.$$

Consider an asynchronous timing distribution with probability distribution $(\mathcal{T}, p)$ that is independent across players. Specifically, suppose that $\mathrm{supp}(T_1) \cap \mathrm{supp}(T_2) = \varnothing$ and $\mathrm{supp}(T_1) \cup \mathrm{supp}(T_2) = \mathbb{Q}$. Also we suppose that for any $t, t' \in \mathbb{R}$ with $t < t'$, $\mathrm{Prob}^p(T_i \in (t, t')) > 0$ holds for each $i = 1, 2$. Note that, because probabilities have finite measures, for all $t' \in \mathbb{R}$ and $i = 1, 2$, $\lim_{t\to t'} \mathrm{Prob}^p(T_i \in (t, t')) = 0$.

We denote the incomplete-information dynamic game specified above by $((\theta_A, \theta_B), \mathcal{T}, p, c, \Omega, (P_1, P_2))$. PBE is defined in an analogous manner as in Section 2. We let $\sigma_i(g)(h)$ be the distribution over actions when the observed information cell is g and the history is h.

Proposition 14. *Fix $((\theta_A, \theta_B), \mathcal{T}, p, \Omega, (P_1, P_2))$. There exists $\bar{c} > 0$ such that for any disclosure cost $c < \bar{c}$, the dynamic game $((\theta_A, \theta_B), \mathcal{T}, p, c, \Omega, (P_1, P_2))$ has a unique PBE σ^*. This σ^* satisfies the following for each player i and each $g \in P_i$.*

1. *For each h that contains no observation, $\sigma^*_i(g)(h)(a_i(g), pay) = 1$.*
2. *For each h that contains an observation, $\sigma^*_i(g)(h)(a, not) = 1$ where a is the static best response to the first-mover's action.*

The proposition implies that, given no observation, players always pay to disclose their actions. The reason is that there is a lack of common knowledge about which action is optimal, and thus there is always a risk of miscoordination when there is no disclosure, even on the equilibrium path. This risk can be avoided by paying a small cost, and thus players prefer to pay.

	C_2	D_2
C_1	$1, 1$	$-s_1, m_2$
D_1	$m_1, -s_2$	$0, 0$

Figure 11: Prisoner's Dilemma

B.6 Repeated Games

Here we apply our uniqueness result for common interest games 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 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 S is a prisoner's dilemma game as in Figure 11 with $m_i > 1$, $s_i > 0$, and $m_i - s_i < 2$ for each $i = 1, 2$. As in the base model, we assume that each player i chooses from $\{C_i, D_i\} \times \{\text{pay}, \text{not}\}$ at each opportunity. 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**. It is characterized by $\big(S, (\mathcal{T}^{(k)})_{k\in\mathbb{N}}, (p_k)_{k\in\mathbb{N}}, c, \rho\big)$.

We consider the following class of supergame strategy profile: A strategy profile is a **T-grim trigger** if each player i plays a grim-trigger strategy after the first T-periods. That is, at each period $t > T$, each player i plays C_i if there has been no D_j, $j = 1, 2$ in the past, 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.

Proposition 15. *There exists $\bar{\rho} > 0$ such that for all $\rho < \bar{\rho}$ and all $T < \infty$, there is $\bar{c} > 0$ such that for all $c < \bar{c}$, for any repeated private-timing prisoner's dilemma $\big(S, (\mathcal{T}^{(k)})_{k\in\mathbb{N}}, (p_k)_{k\in\mathbb{N}}, c, \rho\big)$, if a T-grim trigger strategy profile is a PBE, then it 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 15, if the moves at each period are simultaneous, then a folk-theorem type result holds even with the restriction to the set of equilibria that play grim trigger far in the future. Formally, for any nonnegative payoff and feasible payoff vector, for any $\varepsilon > 0$, there is T such that there is a strategy profile in the class of T-grim trigger such that the average discounted payoff profile is in the ε-neighborhood of the given payoff vector.

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. In a game with perfect monitoring and asynchronicity of moves, backwards induction implies cooperation for all T. What we show here is that such results are true even in the presence of uncertainty about timing if the players have the option to disclose their move (and choose not to) in the current period. $\square$

C Proofs for the Results in the Online Appendix

C.1 Proof for Proposition 10

Step 1:

Step 1-1 Fix a common interest game such that profile a^* is strictly $(r-q)$-dominant and a timing distribution p that is $(1+\varepsilon-q)$-dispersed. 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 $t^* := \inf_{t \in \bigcup_{i \in N} N_i(a^*)} t$. At time t^* all players must choose a_i^*. In fact, the probability that any opponent j chooses an action other than a_j^* before time t^* is zero. Therefore, if at time t^* player i chooses (a_i^*, pay) the opponents respond with a_{-i}^* with probability at least $r - q + \varepsilon$. Therefore, because a^* is $(r-q)$-dominant, player i must choose a_i^* at time t^*.

Step 1-2: By the definition of $(1+\varepsilon-q)$-dispersion, there must exist $i \in N$ and $t' > t^*$ such that for $j \neq i$ and $t \in (t^*, t'] \cap \mathcal{T}_i$, $\mathbb{P}(t^* < T_j \leq t | T_i = t) < q - \varepsilon$. If player i chooses (a_i^*, pay) at time t, player $-i$ responds with a_{-i}^* with probability at least $r(1+\varepsilon-q) \geq r+\varepsilon-q$ because $q > \varepsilon$ holds as p is $(1+\varepsilon-q)$-dispersed. In the component game the action profile a^* is strictly $(r-q)$-dominant, therefore, 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, there exists $\bar{c} > 0$ such that for all $c < \bar{c}$, player i prefers to play a_i^* at time t. Thus, $(t^*, t'] \cap N_i(a^*) = (t^*, t'] \cap N_j(a^*) = \emptyset$. This contradicts the definition of t^*. Thus, $N_i(a^*)$ 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, σ^* must assign probability one to 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 $(a_j^*, \cdot)$. Hence $(\sigma_i', \sigma_{-i}^*)$ must assign probability one to a_j^*. 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 history at which any player pays the disclosure cost under σ^*.

Step 3:

The proposed strategy profile is indeed a PBE. At every history at time t under which there has not been an observation about player j, each player i's belief assigns probability one to the the event in which j plays (a_j^*, not) at j's moving time. Hence, (a_i^*, not) is i's best response at t. $\square$

C.2 Proof of Proposition 11

Fix a constant-sum component game. Let $U = g_1(a) + g_2(a)$ for $a \in A$. 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 (correlation of strategies arises because actions depend on the timing of moves, which can be, itself, correlated across players). Thus, in particular, fixing a PBE, the probability distribution over the action profiles is equal to that of some correlated equilibrium of the component game. $\square$

C.3 Proof of Proposition 12

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 strictly better than (a_i, pay) with $a_i \neq a_i^D$ 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 strictly better than (a_i, pay) with $a_i \neq a_i^D$ at t because $-i$'s best response at t' is (a_{-i}^D, 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 strictly better than (a_i, not) with $a_i \neq a_i^D$ because $-i$'s action is independent of i's unless $-i$ assigns probability one to the event that he is the second mover even without observation. Even if he attaches probability one to such an event, then he plays (a_{-i}^D, not), which he would play even if i pays. Thus, overall, $-i$'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 arguments imply 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 playing (a_{-i}^D, not) is $g_{-i}(a_{-i}^D, a_i^D)$. Since the latter is strictly greater than the former, player $-i$ plays (a_{-i}^D, not) under any history in any PBE. $\square$

C.4 Proof of Proposition 13

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 $h_{1,t} \in \mathcal{H}_{1,t}$, let

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

Also, for each $t \in \mathcal{T}_2$ and $h_{2,t} \in \mathcal{H}_{2,t}$, let

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

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 profile, (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 of the form $(\emptyset, \cdot, t)$. 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

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

so the history $(\emptyset, \cdot, t)$ 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,

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

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,

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

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

$$\mathrm{Prob}^p(t < T_1 \leq t^* | t < T_1) \cdot \beta_2 \leq \mathrm{Prob}^p(t^* < T_1 | t < T_1) \cdot \alpha_2. \tag{7}$$

Since $\mathrm{Prob}^p(t < T_1 \leq t^* | t < T_1)$ is nonincreasing in t and $\mathrm{Prob}^p(t^* < T_1 | t < T_1)$ is nondecreasing in t, (7) holds for all $t \leq t^*$ if

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

which is equivalent to

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

or

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

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

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

By the definition of $\Delta(p)$, this implies

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

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. □

C.5 Proof of Proposition 14

For each disclosure cost $c > 0$, fix an arbitrary PBE σ^c. We first prove that there exists $\bar{c} > 0$ such that for all $c < \bar{c}$, $\sigma^c = \sigma^*$ must hold. Then we show that σ^* is indeed a PBE.

Step 1: Showing that $\sigma^c = \sigma^*$ if σ^c is a PBE.

Given strategy profile σ^c, for each set $A \subseteq (-\infty, \infty)$, conditional on receiving an opportunity at some time t and the history up to time t, player i forms a belief about the probability that $-i$'s opportunity is in A. Let $\mu_{\sigma^c}(A, t) \in [0, 1]$ denote i's belief that $-i$'s opportunity is in set A, given that i's opportunity is at time t and i has not observed a disclosure. Let $\pi_{\sigma^c}(A, t)$ denote c plus the expected payoff of player i conditional on her moving at time t, not observing a disclosure, playing $(a_i(g), \text{pay})$ at time t where g is the observed cell, and $-i$ moving at a time in set A and playing according to σ^c_{-i}. If A has probability zero according to i's conditioning then $\pi_{\sigma^c}(A, t)$ is taken to be zero.

Conditional on the realized cell $g \in P_i$ and receiving an opportunity at time $t \in \mathcal{T}_i$ without an observation, a lower bound on player i's payoff from playing $(a_i(g), \text{pay})$ at time t is given by

$$\mu_{\sigma^c}((-\infty, t], t)\, \pi_{\sigma^c}((-\infty, t], t) + (1 - \mu_{\sigma^c}((-\infty, t], t))\,(\alpha \bar{q}^i(g) + (1 - \bar{q}^i(g))) - c. \quad (8)$$

Under the same conditioning, an upper bound of the payoff from $(a', \cdot)$ where $a' \neq a_i(g)$ is

$$\bar{q}^i(g) \cdot 1 + (1 - \bar{q}^i(g)) \cdot \alpha. \quad (9)$$

This expression is an upper bound because it assumes that, conditional on each action by i, player $-i$ takes an action that maximizes i's payoff.

Now, noting that the state space is finite, there exists $\delta > 0$ such that for all $g \in P_i$ and $i = 1, 2$, $\bar{q}^i(g) > \frac{1}{2} + \delta$ holds. Thus, by asynchronicity and independence, there exist $\tilde{t} > -\infty$ and $\bar{c} > 0$ such that for all $c < \bar{c}$, (8) is strictly higher than (9) for each $i = 1, 2$ and $t < \tilde{t}$. Therefore, $\sigma^c_i(g)(h)$ must assign probability 1 to $(a(g), \cdot)$ for every history h without observation at every time $t < \tilde{t}$.

Let $\mathcal{H}_i^{\emptyset, t}$ be the set of i's time-t private histories that have no observation. Define the following two pieces of notation:

$$\hat{t}_F(\sigma^c) = \inf_{i \in \{1,2\},\, g \in P_i} \{t \in \mathbb{R} \cup \{-\infty, \infty\} | \sigma^c_i(g)(h)(a_i(g), \cdot) < 1 \text{ for some } h \in H_i^{\emptyset, t}\};$$

$$\hat{t}_D(\sigma^c) = \inf_{i \in \{1,2\},\, g \in P_i} \{t \in \mathbb{R} \cup \{-\infty, \infty\} | \sigma^c_i(g)(h)(a_i(g), \text{pay}) < 1 \text{ for some } h \in H_i^{\emptyset, t}\}.$$

By definition, $\hat{t}_F(\sigma^c) \geq \hat{t}_D(\sigma^c)$ holds. Also, as we argued, $\hat{t}_F(\sigma^c) > -\infty$ holds.

Note that, before $\hat{t}_F(\sigma^c)$, both players play according to $(a_i(g), \cdot)$ when g is the observed cell. Moreover, with probability at least ε, conditional on i choosing $(a_i(g), \text{not})$, player $-i$ would choose $a' \neq a_i(g)$ under such a strategy profile (in which case i's payoff is zero). These two facts imply that an upper bound of the expected payoff from playing $(a_i(g), \text{not})$ at $t \in \mathcal{T}_i \cap (-\infty, \hat{t}_F(\sigma^c))$ is

$$\begin{aligned}
&\mu_{\sigma^c}\left((-\infty, t], t\right) \pi_{\sigma^c}\left((-\infty, t], t\right) + \Big[\mu_{\sigma^c}\left((t, \hat{t}_F(\sigma^c)), t\right)(1-\varepsilon) \\
&+ \mu_{\sigma^c}\left([\hat{t}_F(\sigma^c), \infty), t\right)\Big]\left(\alpha \bar{q}^i(g) + (1 - \bar{q}^i(g))\right). \qquad (10)
\end{aligned}$$

Suppose that there is a sequence $\{c_n\}_{n\in\mathbb{N}}$ with $c_n \to 0$ such that for each disclosure cost c_n, σ^{c_n} is different from σ^*. Because σ^{c_n} is not σ^*, $\hat{t}_D(\sigma^{c_n}) < \infty$ for each $n \in \mathbb{N}$. From the definition of $\hat{t}_D(\sigma^{c_n})$ and the players' Bayesian belief updates, $\mu_{\sigma^{c_n}}\left((-\infty, t], t\right) = \mu_{\sigma^{c_n}}\left([\hat{t}_D(\sigma^{c_n}), t], t\right)$.

Because probabilities are bounded in $\mathbb{R}$, $\text{Prob}^p\left(T_{-i} \in (\hat{t}_D(\sigma^{c_n}), \hat{t}_F(\sigma^{c_n}))\right)$ has a convergent subsequence. Passing to the subsequence, we have the following cases.

Case 1: Suppose $\lim_{n\to\infty} \text{Prob}^p\left(T_{-i} \in (\hat{t}_D(\sigma^{c_n}), \hat{t}_F(\sigma^{c_n}))\right) = 0$.

Due to independence, for every $\nu > 0$, there are $\bar{n}(\nu) < \infty$ and $t(\nu) > \hat{t}_F(\sigma^{c_n})$ such that $\mu_{\sigma^{c_n}}\left([\hat{t}_D(\sigma^{c_n}), t], t\right) = \mu_{\sigma^{c_n}}\left((-\infty, t], t\right) < \nu$ for every $t \in (\hat{t}_F(\sigma^{c_n}), t(\nu))$ and $n \geq \bar{n}(\nu)$.

This implies that there exist $\bar{\nu} > 0$ and $\bar{c}_2 > 0$ such that, if $\nu < \bar{\nu}$ and $c < \bar{c}_2$, then expression (9) is strictly less than (8) for $t \in (\hat{t}_F(\sigma^{c_n}), t(\nu))$. That is, i chooses $a_i(g)$ for $t \in (\hat{t}_F(\sigma^{c_n}), t(\nu))$ under σ^{c_n} where g is the observed cell, which contradicts the definition of $\hat{t}_F(\sigma^{c_n})$.

Case 2: Suppose $\lim_{n\to\infty} \text{Prob}^p\left(T_{-i} \in (\hat{t}_D(\sigma^{c_n}), \hat{t}_F(\sigma^{c_n}))\right) > 0$. There is $\bar{\bar{n}} < \infty$ and $\lambda > 0$ such that for all $n \geq \bar{\bar{n}}$, $\text{Prob}^p\left(T_{-i} \in (\hat{t}_D(\sigma^{c_n}), \hat{t}_F(\sigma^{c_n}))\right) > \lambda$. Let $\tau^n \in (\hat{t}_D(\sigma^{c_n}), \hat{t}_F(\sigma^{c_n}))$ be such that $\text{Prob}^p\left(T_{-i} \in (\tau^n, \hat{t}_F(\sigma^{c_n}))\right) > \lambda/2$ for each $n \geq \bar{\bar{n}}$. Such τ^n exists because we have $\lim_{t\to t'} \text{Prob}^p(T_i \in (t, t')) = 0$ for all $t' \in \mathbb{R}$ and $i = 1, 2$ and p is an independent distribution. We have

$$\mu_{\sigma^{c_n}}\left((t, \hat{t}_F(\sigma^{c_n})), t\right) \geq \mu_{\sigma^{c_n}}\left((\tau^n, \hat{t}_F(\sigma^{c_n})), \tau^n\right) > \text{Prob}^p(T_{-i} \in (\tau^n, \hat{t}_F(\sigma^{c_n}))) > \lambda/2$$

for all $t \in [\hat{t}_D(\sigma^{c_n}), \tau^n)$ and $n \geq \bar{\bar{n}}$.

But then, there must exist $\bar{c}_3 > 0$ such that expression (8) is strictly greater than (10) for all $t \in (\hat{t}_D(\sigma^{c_n}), \tau^n)$ and $c_n \leq \bar{c}_3$. Fix n such that $c_n \leq \bar{c}_3$. The previous statement contradicts the definition of $\hat{t}_D(\sigma^{c_n})$ because it implies that i strictly prefers $(a_i(g), \text{pay})$ to $(a_i(g), \text{not})$ at times in $(\hat{t}_D(\sigma^{c_n}), \tau^n)$ where g is the observed cell.

Thus, we have shown that if there is a PBE, then it must be σ^*.

Step 2: Showing that σ^* is a PBE.

We now show that σ^* is indeed a PBE. We have $\mu_{\sigma^*}((-\infty, t], t) = 0$, where the right-end point of the interval in the first argument can be taken to be closed due to asynchronicity. As before, there is $\bar{c}_0 > 0$ such that for every $t \in \mathcal{T}_i$, (8) is strictly above (9) for $c < \bar{c}_0$ under $\sigma^c = \sigma^*$. This implies that, under σ^*, players do not have incentives to deviate to $a' \neq a_i(g)$ for $c < \bar{c}_0$, where g is the observed cell. Also, given a disclosure, it is immediate that taking a static best reply in the component game is optimal.

Furthermore, under σ^*, we have $\hat{t}_F(\sigma^*) = \hat{t}_D(\sigma^*) = \infty$, which implies $\mu_{\sigma^*}((-\infty, t], t) = 0$ and $\mu_{\sigma^*}\big((t, \hat{t}_F(\sigma^*)), t\big) = 1$ for all $t \in \mathcal{T}_i \cap (-\infty, \hat{t}_F(\sigma^*))$. Thus, under σ^*, there exists $\bar{c}_1 > 0$ such that for $c < \bar{c}_1$, expression (10) is strictly less than (8) for every $t \in \mathcal{T}_i$ when $\sigma^c = \sigma^*$ in these equations. That is, players do not have incentives to deviate to non-disclosure under σ^* and, therefore, σ^* is a PBE whenever $c < \min\{\bar{c}_1, \bar{c}_0\}$. This completes the proof. $\square$

C.6 Proof of Proposition 15

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 in which 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 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). Fix any $k = 1, 2, \ldots, T$. We first prove that, given the history in which only (C_1, C_2)

	C_2	D_2
C_1	$\frac{e^{-\rho(k-t_1)}}{1-e^{-\rho}}, \frac{e^{-\rho(k-t_2)}}{1-e^{-\rho}}$	$-e^{-\rho(k-t_1)}s_1, e^{-\rho(k-t_2)}m_2$
D_1	$e^{-\rho(k-t_1)}m_1, -e^{-\rho(k-t_2)}s_2$	$0, 0$

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	$e^{-\rho(k-t_1)}, e^{-\rho(k-t_2)}$	$-e^{-\rho(k-t_1)}s_1, e^{-\rho(k-t_2)}m_2$
D_1	$e^{-\rho(k-t_1)}m_1, -e^{-\rho(k-t_2)}s_2$	$0, 0$

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

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 12. Note that, if $\frac{e^{-\rho(k-t_i)}}{1-e^{-\rho}} > e^{-\rho(k-t_i)}m_i$ for each player i, this game is a common interest game. This holds if and only if $e^{-\rho} > 1 - \frac{1}{m_i}$. Since $m_i > 1$, 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, at a history in the time interval $[k-1, k)$ in which 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 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 $t_1, t_2 \in [k-1, k)$ is given by the one in Figure 13. Note that action D_i is a strictly dominant action for each i. Thus, Proposition 12 implies that the only action that can be played in any PBE is (D_1, D_2). □