For Online Publication

Online Supplementary Appendix to: Games with Private Timing

Yuichiro Kamada[†] Sofia Moroni[‡]

June 30, 2023

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 game in the main text corresponds to $(S, \mathcal{T}, p, c, 1)$.

Proposition 11. *Fix a common-interest component game S, a moving-time distribution $(\mathcal{T}, p)$ and a number $r \in [0, 1]$. Suppose that there is $q \in (0, r)$ such that for each player i the game S is s_i-common with $s_i \geq 1 - r + q$, and $(\mathcal{T}, p)$ is $(1 + \varepsilon - q)$-dispersed for some $\varepsilon \in (0, q)$. Then, there is $\bar{c} > 0$ such that for all $c < \bar{c}$, the dynamic game $(S, \mathcal{T}, p, c, r)$ has a unique PBE. On the path of the unique PBE, each player i plays (a_i^*, not) at any realization of T_i.*

[†]Haas School of Business, University of California Berkeley, 2220 Piedmont Avenue, Berkeley, CA 94720-1900, USA, and Faculty of Economics, The University of Tokyo, e-mail: y.cam.24@gmail.com

[‡]School of Public and International Affairs, Princeton University, email: smoroni@princeton.edu

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 that (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 s_i-commonality and $(1 + \varepsilon - q)$-dispersion.

We note that the potential leader condition and the unlikely leader condition, which are key conditions in the main analysis, are 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 a unique PBE, as a function of r.

B.2 Sense of Calendar Time

Example 3 demonstrates why we need SAP 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 the state 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 Example 3 in Section 3.1 (see Figure 3). Figure 7 shows the extensive form of this game with payoffs at each terminal node. As in Figure 2, we omit a player's actions that are strictly suboptimal (conditional on reaching the corresponding information set) when she has observed the opponent's action.

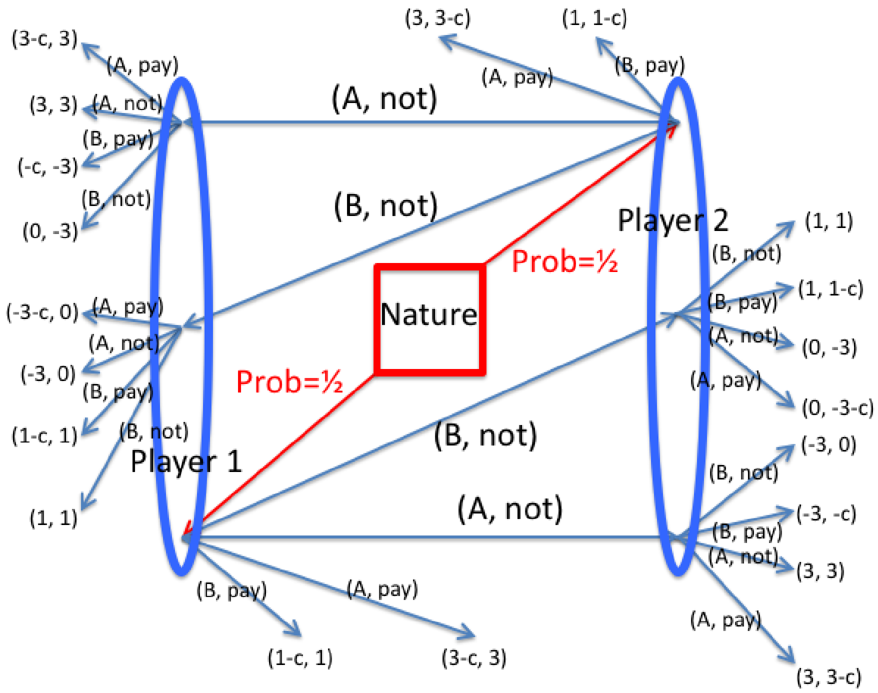

Figure 7: Extensive-form of the game with no sense of calendar time

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.3 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, which helped 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 provide a sufficient condition that guarantees

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 8: Two possible games

uniqueness of a PBE.

Specifically, there are two possible component games, $\theta = \theta_A, \theta_B$, as in Figure 8 with $\alpha > 1$. Observe that, in either game, (A, A) and (B, B) are strict Nash equilibria, but only one of them gives payoff α to each player, and which action profile gives 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 consider a finite state space Ω over which there are information partitions P_1 and P_2 of the two players and a probability distribution f. A function $\bar{\theta} : \Omega \to \{\theta_A, \theta_B\}$ specifies the realized game. That is, if the state is ω, then the realized game is $\bar{\theta}(\omega)$. Before the dynamic game starts, each player i is informed that the state belongs to cell $g \in P_i$ 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 $\omega \in \Omega$, $q_i \cdot \alpha + (1 - q_i) \cdot 1 \neq q_i \cdot 1 + (1 - q_i) \cdot \alpha$ 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 uniquely 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, assuming the opponent's static best response. 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

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

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

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

Proposition 12. *Fix $((\theta_A, \theta_B), \mathcal{T}, p, \Omega, (P_1, P_2), f)$. There exists $\bar{c} > 0$ such that for any disclosure cost $c < \bar{c}$, the dynamic game $((\theta_A, \theta_B), \mathcal{T}, p, c, \Omega, (P_1, P_2), f)$ 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, $\sigma_i^*(g)(h)(a_i(g), \text{pay}) = 1$.*
2. *For each h that contains an observation, $\sigma_i^*(g)(h)(a, \text{not}) = 1$ 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 Proofs for the Results in the Online Appendix

C.1 Proof for Proposition 11

Step 1:

Step 1-1 Fix a common-interest game such that profile a^* is $(r-q)$-dominant and a timing distribution p that is $(1+\varepsilon-q)$-dispersed. Fix a PBE of $(S, \mathcal{T}, p, c, r)$ and, for $i \in \{1,2\}$, let $N_i \subseteq \mathcal{T}_i$ be the set of times t such that there exists a private 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 is nonempty for some $i \in N$. Let $t^* := \inf(N_1 \cup N_2)$. We claim that, at time t^*, all players must choose a_i^*. To see why this holds, notice that the probability that any opponent j chooses an action other than a_j^* before time t^* is zero. Therefore, if player i chooses (a_i^*, pay) at time t^*, then $-i$ responds with a_{-i}^* with probability at least $r - q + \varepsilon \geq 1 - s_i$. Therefore, as discussed when we posed equation (9), 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 $t \in (t^*, t'] \cap \mathcal{T}_i$, $\text{Prob}^p(t^* < T_j \leq t | T_i = t, T_j \geq t^*) < q - \varepsilon$. Therefore, if player i chooses (a_i^*, pay) at time t, player $-i$ responds with a_{-i}^* with probability at least $r(1+\varepsilon-q) \geq r + \varepsilon - q$, as $q > \varepsilon$. Since the game S is s_i-common with $s_i \geq 1 - r + q$, 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, which implies $(t^*, t'] \cap N_i = (t^*, t'] \cap N_j = \emptyset$. This contradicts the definition of t^*. Therefore, N_i 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 private histories. Call this strategy σ_i'. Then, under no observation, for any realization of $T_{-i} \in \mathcal{T}_{-i}$, $-i$ is at an information set that can be reached with positive probability under σ^*, so plays $(a_{-i}^*, \cdot)$. Therefore, the payoff from $(\sigma_i', \sigma_{-i}^*)$ is g_i^*, and the deviation is profitable. This is a contradiction to the assumption that σ^* is a PBE. Hence, there is no private history of any i at which i pays the disclosure cost under σ^*.

The proposed strategy profile is indeed a PBE. At every private 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. □

C.2 Proof of Proposition 12

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 $M \subseteq (-\infty, \infty)$, conditional on receiving an opportunity at some time t and the private history up to time t, player i forms a belief about the probability that $-i$'s opportunity is in M. Let $\mu_{\sigma^c}(M, t) \in [0, 1]$ denote i's belief that $-i$'s opportunity is in set M, given that i's opportunity is at time t and i has not observed a disclosure. Let $\pi_{\sigma^c}(M, t)$ 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 M and playing according to σ^c_{-i}. If M has probability zero according to i's conditioning then $\pi_{\sigma^c}(M, t)$ 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

$$\mu_{\sigma^c}((-\infty, t], t)\, \pi_{\sigma^c}((-\infty, t], t) + (1 - \mu_{\sigma^c}((-\infty, t], t))\, (\alpha \bar{q}^i(g) + (1 - \bar{q}^i(g))) - c. \tag{10}$$

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

$$\bar{q}^i(g) \cdot 1 + (1 - \bar{q}^i(g)) \cdot \alpha. \tag{11}$$

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 $\bar{c}_1 > 0$ such that for all $c < \bar{c}_1$, (10) is strictly higher than (11) for each $i = 1, 2$ and $t < \tilde{t}$. Therefore, $\sigma_i^c(g)(h_i)$ must assign probability 1 to $(a(g), \cdot)$ for every private history h_i without observation at every time $t < \tilde{t}$.

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

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

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

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

$$\begin{aligned} &\mu_{\sigma^c}((-\infty, t], t)\, \pi_{\sigma^c}((-\infty, t], t) + \Big[\mu_{\sigma^c}\left((t, \hat{t}_F(\sigma^c)), t\right)(1-\varepsilon) \\ &+ \mu_{\sigma^c}\left([\hat{t}_F(\sigma^c), \infty), t\right)\Big]\left(\alpha \bar{q}^i(g) + (1 - \bar{q}^i(g))\right). \end{aligned} \tag{12}$$

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, $\mu_{\sigma^{c_n}}((-\infty, t], t) = \mu_{\sigma^{c_n}}\left([\hat{t}_D(\sigma^{c_n}), t], t\right)$.

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

Case 1: Suppose $\lim_{n \to \infty} s_n = 0$.

Due to independence, for every $\nu > 0$, there are $\bar{n}(\nu) < \infty$ and $t(\nu) > \hat{t}_F(\sigma^{c_n})$ such that $\mu_{\sigma^{c_n}}\left([\hat{t}_D(\sigma^{c_n}), t], t\right) = \mu_{\sigma^{c_n}}((-\infty, t], t) < \nu$ 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 (11) is strictly less than (10) 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 $\lim_{n\to\infty} s_n > 0$. 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

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

for all $t \in [\hat{t}_D(\sigma^{c_n}), \tau^n)$ and $n \geq \bar{\bar{n}}$, where the second inequality follows from the fact that it is possible that under σ^{c_n}, $-i$ would have disclosed before τ_n with positive probability.

But then, there must exist $\bar{c}_3> 0$ such that expression (10) is strictly greater than (12) 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 for each $c < \min\{\bar{c}_1, \bar{c}_2, \bar{c}_3\}$, 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 $\bar{c}_4 > 0$ such that for every $t \in \mathcal{T}_i$, (10) is strictly above (11) for $c < \bar{c}_4$ under $\sigma^c = \sigma^*$. This implies that, under σ^*, players do not have incentives to deviate to $a' \neq a_i(g)$ for $c < \bar{c}_4$, 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 $\bar{c}_5 > 0$ such that for $c < \bar{c}_5$, expression (12) is strictly less than (10) 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 $c < \min\{\bar{c}_4, \bar{c}_5\}$. This completes the proof. $\square$