Private Information about the Game Horizon

Job Market Paper

Click here for the latest version

Julia Salmi*

January 8, 2024

Abstract

The starting point of this paper is the dichotomy in repeated games between finite horizon games with a commonly known ending time and infinite horizon games where the ending time is unknown. We study an environment in between where players *privately* know a deadline at which the game must end at the latest. Our main result shows that cooperation can be sustained even when there is a strong correlation between the private deadline, i.e., when the informational environment is arbitrarily close to the common knowledge of the ending time. The leading application is collaboration in a partnership before dissolution, in which we ask if cooperation can be sustained when both partners know that the relationship is going to break down.

JEL classification: C73, D82, D83

Keywords: Repeated games, backward induction, folk theorems, private information, higher-order uncertainty

*University of Copenhagen (julia.salmi@econ.ku.dk, `https://sites.google.com/view/juliasalmi`). I thank Daniel Hauser, Johannes Hörner, Jan Knoepfle, Nenad Kos, Tuomas Laiho, Peter Matthews, Pauli Murto, Anna Sanktjohanser, Peter Norman Sørensen, Juuso Toikka, Juuso Välimäki, and the seminar audience at Nordic Theory Workshop for their helpful comments.

1 Introduction

Cooperation is everywhere. There is extensive literature analyzing whether and to what extent cooperation can be sustained in economic interactions without binding contracts. It is well known that dynamic incentives can support cooperation in repeated interactions if the horizon is (possibly) infinite. However, in finite horizon games, backward induction limits the possibilities of intertemporal incentives and eliminates them altogether in games with a unique stage-game equilibrium. In this paper, we go in between the two extremes by introducing a privately observable finite ending of the game. Is cooperation possible when players are self-interested and fully rational and each player knows a finite deadline when the game must end at the latest?

Private deadlines naturally arise when individuals know how long they will personally stay in a long-term relationship. Consider a team contributing to a joint project. Each team member may have received a lucrative outside option and know already when they are leaving. Similarly, workers receive information about how structural changes are going to affect their own position in the firm. However, they may not know if their team members will be affected. Both of these situations lead to team members having private information about the dissolution of the relationship.

We exploit the idea of personal exit times to study the impact of private information about the game horizon. We model the setup as a two-player partnership game (prisoners' dilemma) that captures the conflict between self-interest and cooperation: the players choose between working and shirking where working is socially beneficial but individually costly, resulting in both players shirking in the unique stage-game equilibrium.[1]

The same stage game is played repeatedly until the first player's exit time is reached. We are interested in the case where, at the beginning of the game, each player observes her own exit time privately but not the other player's exit time. Hence, each player knows that the game is finite as the player's own exit

[1]See Benoit and Krishna (1985) for the folk theorem of finitely repeated games with multiple stage-game equilibria.

time works as a finite deadline at which the game must end at the latest. For simplicity, there is no other source of discounting.

Unavoidably, each player shirks in the last period before her personal exit time. But what happens in earlier periods? We analyze two cases separately: first, we assume that the exit times are statistically independent, and second, we introduce correlation. In both cases, we build on the standard discounted game. Therefore, we assume that the exit times are drawn from the geometric distribution with a constant hazard rate $1 - \delta$. As the game ends when either player's exit time is reached, the game would be a standard infinite horizon game with discount factor δ^2 if the players did not know their exit times.

First, we show that if the exit times are independently distributed, one can support cooperation by using standard trigger strategies, except that there is no cooperation in the last period. How is it possible that myopic optimization does not unravel backward to the penultimate period? The player who knows that the game must end in two periods thinks that it is very likely that her opponent's personal exit time is higher than hers. Therefore, she believes that her opponent believes (wrongly) that cooperation would be possible in the next period, too. In order to sustain her opponent's belief in cooperation in the last period, she works in the penultimate period. Hence, uncertainty about the other player's information prevents backward induction.

One interpretation of the result is that the case with uncorrelated exit times resembles closely the standard discounted infinite horizon game, i.e. the game where players do not observe any information about the exit times. The opposite extreme case is when the exit times are perfectly correlated, which clearly corresponds to a finite horizon game.

What happens under imperfect correlation? The most interesting case is when the correlation is strong. In that case, the trigger strategies that support cooperation without correlation are not an equilibrium. To see this, assume that the game has proceeded up to the penultimate period of at least one of the players. Now, if the correlation is strong, the player believes that it is very likely that the current period is the penultimate period of the other player as well, which means that the other player will deviate in the next period irrespective of the actions taken in the current period. Thus, a strong correlation brings backward induction back.

As it is clear that the usual trigger strategy does not work under a strong correlation, we need to check if any equilibrium strategy profile supports cooperation. The analysis of the correlated case requires handling a dynamic game of incomplete information with i) correlated private information, ii) infinite type sets, and iii) belief updating over time. More precisely, we let the exit times be the same with probability α and be independent otherwise. Our main result shows that for all $\alpha < 1$, there exists an equilibrium where the expected average payoffs approach the cooperative payoffs when the hazard rate goes to zero (i.e. the probabilistic discount rate δ goes to one). We obtain the result by constructing an equilibrium in mixed strategies where players cooperate for a long time at the beginning of the game.

In the cooperative mixed strategy equilibrium, players keep cooperating as long as the other player cooperated in the past and if their exit time is still far in the future. At some point when the game approaches the exit time, they start mixing. The mixing probabilities keep the other player's expectation about cooperation in the next period just high enough to make her indifferent in the current period. The challenge in finding such mixed strategies is that the player's expectations must depend on her exit time and her belief about the other player's exit time. This means that the mixed strategy must make different types of the opponent indifferent at the same time. The indifference conditions with fixed δ are intractable. Especially, showing that there exists a solution in valid mixing probabilities is hard. Therefore, we use an indirect approach where we first show that the set of indifference conditions has a solution in the limit as $\delta \to 1$. Then, we apply a version of the multidimensional intermediate value theorem to show that this implies that the set of indifference conditions with $\delta < 1$ also has a solution and that this solution is close to the solution of the limiting case when δ is sufficiently large, guaranteeing that the mixing probabilities are valid. The novel way we use the intermediate value theorem may be of independent interest.

What do correlated exit times mean in applications? A direct interpretation is that a common event makes both players willing to leave. Many partnerships operate on projects of various lengths: small law firms and consultancy companies may dissolve after a contract with a big client ends; there is an increased risk of divorce soon after the youngest child moves out; in research, it is uncertain if co-authorship continues after finishing a paper. We also illustrate that the analysis for the correlated case is analogous to two other cases: first, the case where players receive private signals about a common ending time, and second, the case where the personal exit times are uncorrelated but players receive imperfect signals about each others' exit times. These indirect interpretations capture situations where a common external event, such as the bankruptcy of the common employer or a major organizational change, ends the collaboration and asymmetric situations where there is almost common knowledge that one player is leaving.

Our results shed light on how dynamic incentives may sustain cooperation even in environments known to be finite. The implications include that collaborative behavior in a partnership does not need to end after a partner decides to leave. Even more strikingly, team incentives may prevail even after the other partner hears convincing rumors about her partner leaving or when a common event dissolves the partnership. Furthermore, our analysis suggests that organizations may want to use private rather than public communication with their employees in front of organizational changes. As long as employees are at least partially unsure if they share the same information with their team members, dynamic incentives may prevail even toward the end of the collaboration.

1.1 Literature

The present paper contributes to the literature on dynamic games under incomplete information by demonstrating the stark difference between games with a commonly known deadline and games with privately known deadlines. The mixed strategy equilibrium of the present paper is in the spirit of finite horizon reputation games, initiated by Kreps and Wilson (1982), Milgrom and Roberts (1982), and Kreps, Milgrom, Roberts and Wilson (1982).[2] The connection is that mixing allows the belief that the exit times are the same to go down over time. However, our environment differs from the finite horizon reputation games due to correlated private information and richer belief updating. Furthermore, no type in our setup resembles the commitment type or good type in reputation games: backward induction would be unavoidable for any realized profile of exit times if they were publicly observable.

A few papers study other kinds of incomplete information in repeated games. Neyman (1999) shows that uncertainty structures, which allow for small inconsistencies in beliefs, overcome backward induction in finite games. In Massó (1996), the chain-store paradox can be prevented when the entrants do not know their place in the sequence and the monitoring is imperfect. In Masso's information structure, it is important that the entrant does not know the number of past periods and therefore it cannot incorporate more than one long-run player. The present paper suggests a simple and intuitive information structure that breaks backward induction in games with many fully rational long-run players.

The stochastic deadline is equivalent to discounting when deadlines are unobservable before hand. Baye and Jansen (1996) provides a folk theorem for repeated games with stochastic discounting where players observe the current discount factor in each period before taking action. When the expected discount factor goes to 1, any payoffs larger than the stage Nash payoffs can be achieved in an equilibrium where grim-trigger is played when the realized discount factor is high and the stage Nash is played when it is low. The result relies on low discount factors being extremely unlikely. Nevertheless, the repeated game tends to have non-degenerate equilibria even when the discount factor goes to zero if the stage game is continuous. Bernheim and Dasgupta (1995) (see also Arribas and Urbano (2005) for a generalization) solve for the maximal rate of convergence that facilitates intertemporal incentives in continuous games when the discount factor goes to zero.

[2]See also Abreu and Pearce (2007), Atakan and Ekmekci (2013), and Fanning (2016) for reputation games with two-sided incomplete information.

From a more abstract perspective, repeated games with stochastic discounting fall into the wider category of stochastic games. In stochastic games folk theorems by Hörner, Sugaya, Takahashi and Vieille (2011) and Fudenberg and Yamamoto (2011), an important assumption is the cyclic nature of the state dynamics: eventually the game returns to a specific state. Intuitively, a cycle can be seen as one dynamic stage game which is then repeated again and again. For instance, the setup in Baye and Jansen (1996) satisfies the cyclicity condition, whereas the setup of the present paper features an absorbing state and is clearly acyclic.

The paper is also related to Rubinstein (1989), in which unbounded higher-order beliefs select a unique rationalizable action in a static game that has multiple Nash equilibria under complete information. The result demonstrates that games with almost common knowledge of the game structure may be strikingly different from complete information games. See also Carlsson and van Damme (1993) and Weinstein and Yildiz (2007).

Finally, there is large experimental literature on repeated games. The often observed pattern in finite horizon games (mostly prisoner's dilemma) is that subjects cooperate but that the probability of cooperation decreases dramatically toward the end (see e.g. Selten et al. (1997) and Bó (2005)). This phase of decreasing cooperation, known as "the end game" in the experimental literature, resembles the mixing phase of the present paper. Notice that the model of the present paper is analogous to the case where subjects are unsure whether their opponent has received the same information about the game horizon or if that information may have been left unnoticed. Selten and Stoecker (1986), Engle-Warnick and Slonim (2006), and Embrey et al. (2017) find that the end game starts earlier if the subjects have played a similar supergame before, which is consistent with the result that defection starts earlier when it is more likely that the players share the same information about the game horizon.

The rest of the paper is organized as follows. First, we present the model in Section 2, followed by the analysis of the special case without correlation in Section 3. The main result that shows that cooperation can be sustained with an arbitrary level of imperfect correlation is in Section 4. We present alternative models in Section 5 and concluding remarks in Section 6.

2 Model

Two players, $i \in \{A, B\}$, play repeated contribution game (prisoners' dilemma): in each period, the players simultaneously choose either to work, W, or shirk, S. Working is individually costly and incurs a cost of $P > 0$ to the player and a gain of $R > P$ to her opponent. Figure 1 represents the stage game.

		Player B: W	Player B: S
Player A	W	$R-P, R-P$	$-P, R$
	S	$R, -P$	$0, 0$

Figure 1: Stage game payoffs, $R > P > 0$.

Before the game starts, each player privately observes her personal exit time (or deadline), T_i. The exit times are drawn from a joint distribution $F(T_A, T_B)$. The game ends when the first player exits so that the last period is $T := \min\{T_A, T_B\}$.

A player's total payoff is the unweighted average of her per period payoffs:

$$U_i(a^T) = \frac{1}{T} \sum_{t=1}^{T} u_i(a_t),$$

where a_t is a pure action profile played in period t, $a^t := (a_1 \dots a_t)$ is a sequence of action profiles up to and including t, and $u_i(a_t)$ is the stage game payoff when profile a_t is played. There is no discounting.[3]

Monitoring is perfect: a public history $h_t = a^{t-1}$ consists of all past actions. Let H_t denote the set of all (public) histories of length t. Our solution concept is perfect Bayesian equilibrium where players' strategies are functions from histories and personal exit times to the probability of cooperation, $\sigma_i : H_t \times \mathbb{N} \to [0, 1]$.[4]

[3]Notice that a player's belief that the other player is there in the next period corresponds with the discount factor.

[4]In the analysis, we do not impose any assumptions on out-off-equilibrium beliefs. Hence, the results would remain unchanged if one used a stronger equilibrium concept, such as sequential equilibrium.

With a fixed strategy profile, $\sigma = (\sigma_A, \sigma_B)$, we let $V_i(\sigma) = \mathbb{E}[U_i(a^T)]$ denote the *ex ante* average payoff of player i where the expectation is over T_i, T_j, and γ and the possible randomization in σ.

3 Uncorrelated exit times

As a starting point, we cover uncorrelated exit times in this section: the exit times T_A and T_B are independent draws from the geometric distribution with hazard rate $1-\delta \in (0,1)$. Hence, T follows the geometric distribution with hazard rate $1-\delta^2$. The variant of the game where the players do not observe their personal exit times is the standard infinite horizon repeated game with discount factor δ^2.

We show that there is an equilibrium with cooperation whenever δ is sufficiently large. First, we propose equilibrium strategies and, then, verify that they constitute an equilibrium.

Consider a trigger strategy σ^{TR} for player i:

$$\sigma^{TR}(h_t, T_i) := \begin{cases} 1 & \text{if no } S \text{ in the previous period and } t < T_i, \\ 0 & \text{otherwise.} \end{cases}$$

Suppose that Player i follows σ^{TR} and that the current period is $T_j - 1$. By working today, Player j can make Player i work tomorrow if and only if Player i's personal exit time is larger than T_j. Hence, Player j has an incentive to work if

$$Pr(T_i > T_j | T_i \geq T_j - 1, T_j) R \geq P. \tag{1}$$

As the exit times are independent, the continuation probability equals $Pr(T_i > T_j | T_i \geq T_j - 1, T_j) = \delta^2$, implying that (1) is satisfied for all $\delta \geq \sqrt{P/R}$. We verify in Appendix A.1 that this is a necessary and sufficient condition for incentive compatibility:

Proposition 1. *Mutual play of σ^{TR} is a perfect Bayesian equilibrium strategy profile if and only if $\delta^2 \geq \frac{P}{R}$.*[5]

[5] To complete the equilibrium, one can use any off-path beliefs that are consistent with the fact that the game has not ended, e.g., that the other player's exit time is geometrically distributed over all future periods. Shirking is the unique best response after a deviation, independent of the off-path beliefs.

When both players follow σ^{TR}, there is cooperation until the last period. In the last period, the exiting player shirks and the other player works. The relative weight of the last period vanishes for long games, and hence the average *ex ante* payoff approaches the efficient payoff, $R - P$, when δ goes to 1.

By changing the information structure, the game has both a discounted infinite game and an undiscounted finite game as special cases: if players do not know their personal exit times, the game is infinite; if players observe each others' personal exit times, the game is finite. Notice that the condition for the hazard rate in Proposition 1, $\delta^2 \geq P/R$, is identical to the case when the players do not observe their own personal exit times.[6] We conclude that private information about personal exit times does not hinder intertemporal incentives, whereas public information would destroy them altogether.

4 Correlated exit times

In the previous section, we saw that one can support cooperation under uncorrelated exit times similarly to the case where players do not know their personal exit times. In contrast, it is clear that the backward induction outcome is the unique equilibrium if the exit times are perfectly correlated. In this section, we ask what happens under imperfect correlation.

We introduce correlation such that it is captured by a single parameter, α. Let the correlation state γ take value 1 with probability α and value 0 with probability $1 - \alpha$. Now, if $\gamma = 0$, T_A and T_B are independent draws from the geometric distribution with hazard rate $1 - \delta$. The game ends when the first exit time is reached, $T = \min\{T_A, T_B\}$. If $\gamma = 1$, the exit times are the same, $T_A = T_B = T$, and follow the geometric distribution with hazard rate $1 - \delta^2$. With this formulation, the ending time, T, is independent of parameter α, and hence the expected ending time is the same as in the uncorrelated game in Section 3.[7] The

[6]The exactly identical hazard rate condition relies on the shirking benefit, P, being independent of the other player's action. If shirking was more (less) profitable when the other player worked, the private information threshold for δ would be lower (higher) than the no information threshold.

[7]The main result extends to an alternative model where $T \sim Geom(1 - \delta)$ if $\gamma = 1$. The

players observe their own T_i but not the realization of γ.[8] Parameter α scales the amount of correlation in the model. If $\alpha = 0$, we are back in the uncorrelated case. For tractability, we exclude a measure-zero set of possible levels of correlation by assuming that $\alpha \notin \{a \in [0,1] : a = 1 - (\frac{P}{R})^n \text{ for some } n \in \mathbb{N}\}$.

4.1 How does correlation complicate cooperation?

First, we illustrate how the correlated case differs from the uncorrelated case. We argue that the trigger strategy profile σ^{TR} used in the previous section is not an equilibrium when the correlation is strong.

Consider the trigger strategy σ^{TR}, and assume that the game has proceeded up to the penultimate period for Player j. If the correlation is strong (α close to 1), the probability $Pr(T_i > T_j | T_i \geq T_j - 1, T_j)$ is very small for all δ. In that case, the trigger strategies do not constitute an equilibrium. Instead, if Player i follows σ^{TR}, Player j faces a temptation to shirk in the penultimate period, and the equilibrium unravels like in a finite horizon game. In this sense, backward induction applies to games with high correlation. Hence, there must be an upper bound for α such that σ^{TR} is an equilibrium for any δ. A direct calculation gives (see Appendix A.2):

Lemma 1. *The mutual play of σ^{TR} is a perfect Bayesian equilibrium strategy profile for some $\delta < 1$ if and only if $\alpha < \frac{R-P}{R+P}$.*

When knowing that cooperation is possible without correlation (Proposition 1) and impossible under perfect correlation, it goes according to the intuition that there is a threshold level of correlation that determines if σ^{TR} defines an equilibrium. Intuitively, high correlation means almost common knowledge of the finite deadline, which breaks down intertemporal incentives.

---

analysis of that case is available upon request.

[8] The players update their belief on γ after seeing their own exit time: $Pr(\gamma = 1 | T_i) = \frac{\alpha(1+\delta)\delta^{T_i - 1}}{\alpha(1+\delta)\delta^{T_i - 1} + 1 - \alpha}$. This updating is not central to our main result; see the previous footnote.

4.2 Cooperation under correlated exit times

In the absence of other sources of information, backward induction is unavoidable under high correlation, as shown by Lemma 1. In what follows, we construct another strategy, $\tilde{\sigma}$, that supports cooperation even when α is large. To facilitate cooperation, mixing is required on the equilibrium path. Mixing provides an endogenous source of information that increases the players' beliefs about the other player's exit time.

We define a *mixed trigger strategy*:

Definition 1. *A strategy $\tilde{\sigma}_i : H_t \times \mathbb{N} \to [0,1]$ is a mixed trigger strategy if it takes the following form:*

$$\tilde{\sigma}_i(W; h_t, T_i) = \begin{cases} p_\delta(T_i - t; T_i) \in [0,1] & \textit{if no } S \textit{ in the previous period,} \\ 0 & \textit{otherwise,} \end{cases} \tag{2}$$

for some function $p_\delta : \mathbb{N} \times \mathbb{N} \to [0,1]$.

A mixed trigger strategy calls for shirking whenever someone has shirked before. Otherwise, the players use a mixed action where the mixing probability depends on the remaining time before their personal exit time, $T_i - t$, and on the personal exit time, T_i. We call a mixed trigger strategy *stationary* if it does not depend on T_i but only on $T_i - t$. We use mixed trigger strategies to show our main result:

Theorem 1. *For all $\alpha < 1$ and $\epsilon > 0$, there exists $\overline{\delta} \in (0,1)$ such that there exists a perfect Bayesian equilibrium strategy profile σ with $V_i(\sigma) \geq R - P - \epsilon$ for both $i \in \{A, B\}$ if $\delta \geq \overline{\delta}$.*

Theorem 1 implies that cooperation is possible for any level of imperfect correlation. We get a folk theorem like result despite that players are almost sure that they share the same information about the finite deadline.

Theorem 1 follows once we show that there exists a function $p_\delta : \mathbb{N} \times \mathbb{N} \to [0,1]$ such that player i is indifferent between working and shirking when $p_\delta(T_i - t, T_i) \in (0,1)$, prefers working when $p_\delta(T_i - t, T_i) = 1$, and prefers shirking when $p_\delta(T_i -$

$t, T_i) = 0$. This shows that there is an equilibrium in mixed trigger strategies. Furthermore, we need to argue that the fraction of periods where players cooperate goes to one as the expected exit times go to infinity. Formally, let $\tau(T_i) := \min\{t \in \mathbb{N} : p_\delta(T_i - t, T_i) < 1\}$. We want to show that $\tau(T_i)/T_i \to 1$ as $T_i \to \infty$ when δ is large.

Instead of trying to find p_δ directly, we take a detour and consider a stationary strategy profile that makes the players indifferent for large enough realizations of the exit time. However, the constructed stationary profile is not an equilibrium because low types have an incentive to deviate. Instead, it should be viewed as a useful middle step in the argument.

4.3 Cooperative stationary strategy profile

In this subsection, we show that one can find stationary mixing probabilities that make the other player indifferent for large enough realizations of the exit time. There, we ignore the problem that such a mixed trigger strategy may not be a best-response if the realized exit time was low. Therefore, the stationary strategy we construct is not an equilibrium strategy but a useful tool in the construction. In the next subsection, we show that there exists a non-stationary mixed trigger strategy profile that is an equilibrium and that converges to the stationary strategies we construct in this subsection.

Suppose that Player i follows a stationary strategy:

$$\tilde{\sigma}_i^S(W; h_t, T_i) = \begin{cases} p_\delta(T_i - t) \in [0, 1] & \text{if no } S \text{ in the previous period,} \\ 0 & \text{otherwise.} \end{cases} \tag{3}$$

This is the stationary version of the mixed trigger strategy (2). Furthermore, suppose that $p_\delta(T_i - t) < 1$ only for $T_i - t \leq K$.

Our objective is to find $p_\delta(1), p_\delta(2), \ldots, p_\delta(K) \in (0, 1)$ that make Player j indifferent in periods $T_j - K, T_j - K + 1, \ldots, T_j - 1$ when Player i follows $\tilde{\sigma}_i^S$. Let the current period be $T_j - k > K$ and the history be such that both players have always worked. Now, Player j's updated belief that the exit times are the same, i.e. $\gamma = 1$, is (derivation in Appendix A.3):

$$\hat{\alpha}(k) := \frac{\alpha(1+\delta)\Pi_{l=k+1}^{K} p_\delta(l)}{\alpha(1+\delta)\Pi_{l=k+1}^{K} p_\delta(l) + (1-\alpha)\left(\delta^{K-k} + \delta^{-k}(1-\delta)\sum_{n=0}^{K-1}\delta^i \Pi_{l=n+1}^{K} p_\delta(l)\right)}.$$

The posterior belief $\hat{\alpha}(k)$ depends only on remaining time until T_j and not on T_j itself. This follows from stationarity and $T_j - k > K$, which guarantees that Player i has had enough time to mix before period $T_j - k$. Posterior $\hat{\alpha}(k)$ is increasing in k – and hence decreasing over time – if endogenous updating through mixing is stronger than exogenous updating through time, i.e. $\hat{\alpha}(k) - \hat{\alpha}(k-1) > 0$ if $p_\delta(k) < \delta$. Intuitively, this must hold at least for small k in order to make Player j (weakly) willing to work. Otherwise, the belief in the penultimate period would remain above α and backward induction logic would dictate that Player j shirks.

Using posterior $\hat{\alpha}(k)$, we can write the indifference condition for Player j in period $T_j - k > K$ as

$$\hat{\alpha}(k)p_\delta(k)p_\delta(k-1) + (1-\hat{\alpha}(k))\sum_{n=0}^{\infty}\beta_n p_\delta(n)p_\delta(n-1) = \frac{P}{R}, \tag{4}$$

where $\beta_n := \delta^n \Pi_{l=n+1}^{K} p_\delta(l) / (\sum_{h=0}^{\infty} \delta^h \Pi_{l=h+1}^{K} p_\delta(l))$ is the probability that $T_i = t+n$ conditional on $T_i \geq t$ and $\gamma = 0$. Overall, the left-hand side of (4) equals the probability that Player i works at least up to period $T_j - (k-1)$ conditional on working up to $T_j - (k+1)$. Here, we implicitly use that Player j is indifferent in period $T_j - (k-1)$, and hence the value in that period equals R times the probability that Player i works.

Indifference condition (4) must hold for all $k \in \{1, \ldots, K\}$. Hence, when solving for the mixing probabilities $p_\delta(1), \ldots, p_\delta(K)$, we have K equations and K unknowns. Furthermore, each $p_\delta(l)$ must be between zero and one to be a valid mixing probability. In Appendix A.4, we show that the system of equations indeed has a solution in valid mixing probabilities when we choose a suitable K:

Lemma 2. *For all $\alpha \in \left(\frac{R-P}{R+P}, 1\right)$, there exist $\bar{\delta} < 1$ and K such that for all $\delta > \bar{\delta}$, there exists numbers $p_\delta(1), p_\delta(2), \ldots, p_\delta(K) \in (0,1)$ such that (8) holds for all $k \in \{1, 2, \ldots, K\}$.*[9]

[9] Here, the parameter restriction that excludes a measure-0 set of values for α is used to get that each $p_\delta(l)$ is *strictly* between 0 and 1.

Lemma 2 implies that there exists a version of $\tilde{\sigma}_i^S$ that is incentive compatible for large realizations of personal exit times. On the implied path, the players receive payoff $R-P$ in all but the last K periods. However, $\tilde{\sigma}_i^S$ is not an equilibrium strategy because it is not optimal for either player if the exit time is small. Hence, a different strategy must be used at least for $T_i < K$, which then affects incentives of types $T_j \in \{K+1, \dots, 2K\}$ of the other player, too. Then if types $T_j \in \{K+1, \dots, 2K\}$ do not use a stationary strategy, the best response by types $T_i \in \{2K+1, \dots, 3K\}$ is non-stationary, and the argument moves on to even higher T_i and T_j. Next, we show that we can find an equilibrium in non-stationary mixed trigger strategies that converge to the same expected outcome path as the stationary mixed trigger strategy $\tilde{\sigma}_i^S$.

4.4 Proof of Theorem 1: idea

To prove Theorem 1, we construct non-stationary symmetric strategies that converge to the same limit as the stationary strategy (3) as $\delta \to 1$ for large enough exit times. Then, as small realized exit times get very unlikely and otherwise the players mix only for a finite number of periods, the expected payoffs approach the cooperative payoffs.

Suppose that K is a number of mixing periods that satisfies Lemma 2. We use a version of the non-stationary mixed trigger strategy (2) where $p_\delta(T_i - t; T_i) = 0$ if $T_i < K$ and $p_\delta(T_i - t; T_i) = 1$ if $T_i - t > K$. Now, if Player i follows the suggested strategy, Player j with $T_j > K$ is indifferent in period $T_j - k$ if the following holds:

$$\begin{aligned}&\hat{\alpha}(k|T_j)p_\delta(k|T_j)p_\delta(k-1|T_j)\\&+(1-\hat{\alpha}(k|T_j))\sum_{n=0}^{\infty}\beta_n(T_j)p_\delta(n|T_j-k+n)p_\delta(n-1|T_j-k+n)=\frac{P}{R},\end{aligned}\tag{5}$$

where $\beta_n(T_j) := \delta^n \pi_{l=n+1}^K p_\delta(l|T_j - k + n)/(\sum_{h=0}^{\infty} \delta^h \pi_{l=h+1}^K p_\delta(l|T_j - k + h))$ is the probability that $T_i = T_j - k + n$ conditional on $T_i \geq T_j - k$ and $\gamma = 0$, and $\hat{\alpha}(k|T_j) := \frac{\alpha(1+\delta)\Pi_{l=k+1}^K p_\delta(l|T_j)}{\alpha(1+\delta)\Pi_{l=k+1}^K p_\delta(l)+(1-\alpha)\left(\delta^{K-k}+\delta^{-k}(1-\delta)\sum_{n=0}^{K-1}\delta^i \Pi_{l=n+1}^K p_\delta(l|T_j-k+n)\right)}$ is the posterior probability for $\gamma = 1$. Notice that (5) is the non-stationary version of (4).

The system of equations (5) is hard to solve. Especially showing directly that the solution consists of valid mixing probabilities is intractable. Therefore, we use an indirect approach: we notice that the limit of (5) coincides with the limit of the stationary indifference condition (4) as $\delta \to 1$. Therefore, we guess that we can approximate (5) with the limit of the stationary problem, which we know to have a solution in $(0,1)^K$. Then, the following lemma turns out to be useful:

Lemma 3. *Let $C \subset \mathbb{R}^n$ be compact and $F : C \to \mathbb{R}^n$ be a continuous function. Suppose that there exists $x^* \in int(C)$ such that $F(x^*) = 0$ and that satisfies the following "sign condition": there exists $\bar{\epsilon} > 0$ and $M > 0$ such that for all $\epsilon \in (0, \bar{\epsilon})$ and for all $y \in \mathbb{R}^n$ such that $||y|| < \epsilon$, there exists $x \in C$ such that $F(x) = y$ and $||x - x^*|| < M\epsilon$.*

Let $(F^m)_{m \in \mathbb{N}}$ be a sequence of continuous functions from C to $\mathbb{R}^n$ that converges to F uniformly. Then, for all $\epsilon > 0$, there exists $\bar{m}$ such that for all $m > \bar{m}$, there exists $z^m \in C$ such that $F^m(z^m) = 0$ and $||z^m - x^*|| < \epsilon$.

We prove the lemma in Appendix A.5. The key step in the proof is that the sign condition allows us to use the multidimensional intermediate value theorem.

Consider a system of K equations and unknows that results from (5) after letting all $p(l|T_i \neq T_j)$ be arbitrary fixed numbers between 0 and 1. Lemma 3 implies that this system has a solution close to the solution of the limit of the stationary strategy profile as $\delta \to 1$, guaranteeing that the required mixing probabilities are between 0 and 1. As the argument works for all $T_j > K$, this verifies that there exists an equilibrium in non-stationary mixed trigger strategies such that players cooperate in all except the last K periods for large enough δ. Theorem 1 follows as the *ex ante* payoff from such a strategy profile converges to $R - P$ as $\delta \to 1$. The detailed proof is in Appendix A.6.

5 Related environments

Here, we present two alternative models and show how the analysis of the main model extends to them.

5.1 Privately observable common shock that ends the game

Consider a team working in a firm that is about to go into bankruptcy. However, with a small probability, each worker is unaware of the coming end. Here, we show this kind of small uncertainty may be enough for building intertemporal incentives for cooperation.

Let the game be otherwise the same as in the main model but assume that instead of private ending times, there is a common event that ends game. The common ending time T is a random draw from the geometric distribution with hazard rate $1-\delta$. Before the game starts, each player observes T with probability α (independently). Otherwise, the player learns nothing and holds the prior $T \sim Geom(1-\delta)$. We let $T_i \in \mathbb{N} \cup 0$ denote the player's type where $T_i = 0$ means Player i has not learned T, and otherwise $T_i = T$.

If Player i has learned T and if α is close to 1, she believes that it is very likely that Player j has also learned T. Therefore, similar to the main model with high correlation, one cannot support cooperation using a pure-action trigger strategy. Instead, mixing is required.

We argue that version of the mixed trigger strategy (2) works in this environment, too. Specifically, we set $p_\delta(T_i - t; T_i) = 1$ if $T_i = 0$ (uninformed player), $p_\delta(T_i - t; T_i) = 0$ if $T_i \in \{1, \ldots, K\}$, and $p_\delta(T_i - t; T_i) = p(T_i - t)$ if $T_i > K$ with some natural number K. Furthermore, we want the strategy to be such that $p(T_i - t) = 1$ if $T_i - t > K$, and incentive compatibility requires that $p(0) = 0$. We denote this strategy by $\tilde{\sigma}^C$.

If Player i follows the suggested mixed trigger strategy, informed Player j is indifferent in period $T_j - k$ if

$$\frac{(1-\alpha) + \alpha \Pi_{l=k-1}^{K} p(l)}{(1-\alpha) + \alpha \Pi_{l=k+1}^{K} p(l)} = \frac{P}{R}. \tag{6}$$

This indifference condition is almost identical to the stationary indifference condition in the main model (4) evaluated at $\delta = 1$. It follows that we can follow similar steps as in the proof of Lemma 2 to show that there exist valid mixing probabilities $p(1), \ldots, p(K)$ such that the indifference condition holds for all $k \in \{1, \ldots, K\}$. Furthermore, the uninformed player wants to follow the trigger strategy if δ is sufficiently high. Therefore we get the following (proof in Appendix B.1):

Proposition 2. *Suppose each player observes the common deadline with probability $\alpha < 1$. Then, there exists $\overline{\delta} \in (0,1)$ and mixing probabilities $p(1), \ldots, p(K)$ such that the mutual play of $\tilde{\sigma}^C$ is a perfect Bayesian equilibrium if $\delta \geq \overline{\delta}$.*[10]

The intuition behind Proposition 2 is largely the same as for Theorem 1: mixing allows the belief that the exit times are not the same to go up toward the end of the, maintaining intertemporal incentives for cooperation.

A natural question is what happens if the events of being informed are correlated or if they depend on T. The argument behind Proposition 2 remains practically unchanged as long as $Pr(T_i = 0|T_j)$ is bounded away from zero for all T_j. The only difference is that the mixing probabilities (and potentially the number of mixing periods) would depend on T if the condition $Pr(T_i = 0|T_j) = Pr(T_i = 0|T_j')$ did not hold for all $T_j, T_j' > K$.

5.2 Uncorrelated exit times with correlated signals

Now, suppose that there is only a one-sided probability of leaving. Let the game end at $T \sim Geom(1-\delta)$ when Player A leaves. Before the game starts, Player A observes T for sure and Player B observes it with probability α. This model extends the scope of applications to include asymmetric situations where there is almost common knowledge that one player is leaving.

Again, we need mixing when α is high. Let the players follow a mixed trigger strategy (2) with $p_\delta(T-t,T) = p^A(T-t)$ if $T > K$ and $p_\delta(T-t,T) = 0$ otherwise for Player A, $p_\delta(T-t,T) = p^B(T-t)$ if $T > K$ and $p_\delta(T-t,T) = 0$ otherwise for informed Player B, $p_\delta(T-t,T) = t$ for uninformed Player B. Let $(\tilde{\sigma}^A, \tilde{\sigma}^B)$ denote this strategy profile. To make Player A indifferent in period $T-k$, the informed Player B's mixing probabilities, $(p^B(1), \ldots, p^B(K))$, must satisfy:

$$\frac{(1-\alpha) + \alpha \Pi_{l=k-1}^{K} p^B(l)}{(1-\alpha) + \alpha \Pi_{l=k+1}^{K} p^B(l)} = \frac{P}{R}.$$

[10]Notice that the mixing probabilities $p(1), \ldots, p(K)$ are independent of δ.

This is the same condition as in the previous subsection because the information environment is the same for Player A.

Next, the informed Player B is indifferent in period $T-k$ if Player A's mixing probabilities satisfy $p^A(k)p^A(k-1) = P/R$. This condition can be satisfied only for $l \leq 2$, which implies that $p^B(1) = 0$. We clearly have $p^A(k) \neq p^B(k)$. Despite this asymmetry, we can follow similar steps as in the main model to show that the constructed strategy profile is an equilibrium (details in Appendix B.2):

Proposition 3. *Suppose that Player B observes the exit time of Player A with probability $\alpha < 1$. Then, there exists $\overline{\delta} \in (0,1)$ and mixing probabilities $p(1), \ldots, p(K)$ such that $(\tilde{\sigma}^A, \tilde{\sigma}^B)$ is a perfect Bayesian equilibrium if $\delta \geq \overline{\delta}$.*

As a final remark, notice it is not the player who knows the true ending time, Player A, who builds a reputation in the equilibrium. Instead, by mixing, the informed Player B increases Player A's belief that Player B is uninformed.

5.3 Bounded beliefs

Here, we show through an example that the main results do not hinge on the unbounded support of the players' beliefs about the other player's personal exit time.

Suppose that $\tilde{T}$ is drawn from the geometric distribution with hazard rate $1-\delta$. Again, we let $\gamma = 1$ with probability α and $\gamma = 0$ otherwise. If $\gamma = 1$, $T_A = T_B = \tilde{T}$. If $\gamma = 0$, T_A and T_B are independent draws from the uniform distribution on $\{n \in \mathbb{N} : |n - \tilde{T}| \leq M\}$ with some $M \in \mathbb{N}$.

Notice that each player knows the other player's exit time is at most $2M$ periods apart from the player's own exit time. We still get a result analogous to Proposition 1:

Proposition 4. *There exists a triplet $(\bar{M}, \bar{\alpha}, \bar{\delta})$ with $\bar{M} \in \mathbb{N}$ $\bar{\alpha} > 0$, and $\bar{\delta} < 1$ such that the mutual play of σ^{TR} is a perfect Bayesian equilibrium if $M > \bar{M}$, $\alpha < \bar{\alpha}$, and $\delta > \bar{\delta}$.*

Proposition 4 follows by noticing that $Pr(T_i > T_j | T_i \geq T_j - 1, T_j)$ approaches 1 as $M \to \infty$, $\alpha \to 0$, and $\delta \to 1$. Hence, the optimality condition (1) is satisfied for large M and δ and small α.

Similar to the main model, the trigger strategy σ^{TR} is not an equilibrium for large α even when $M \to \infty$ and $\delta \to 1$. One can support cooperation by using mixed trigger strategies for all $\alpha < 1$ if M and δ are large enough. This can be easily seen by noticing that the limit of this alternative model as $M \to \infty$ and $\delta \to 1$ is the same as the limit of the main model as $\delta \to 1$. Therefore, one can use Lemma 3 to show that the indifference conditions of the alternative model can be simultaneously satisfied for any $\alpha < 1$ when M and δ are large.

6 Concluding remarks

Here, we conclude with a few important remarks, starting with the applied implications and ending with a discussion of the key assumptions behind the results.

Despite the game theoretic focus of the paper, the implications derived from our research shed light on the practical considerations that organizations must address in maintaining effective cooperation, particularly in the face of impending changes. Consider an organization where teamwork is essential, but it is difficult to reward helping others. In such situations, organizations must rely on the organizational culture where workers trust that if they help a team member today, they will get help tomorrow. The expectations about the durability of the team are the key to whether such intertemporal incentives may be effective: if there is an organizational change ahead or if you or your colleague are about to retire or change jobs, your colleague may not have time – or willingness – to pay back your good deed. The results of the present paper suggest that private, rather than public, communication in organizations helps maintain cooperative norms in front of major changes.

In practice, while it may be necessary to inform individuals about how organizational changes will impact their roles, there typically is much more freedom in designing if they should be informed about the fates of their peers. An implication of our main result (Theorem 1) is that even if each worker can infer fairly accurately what the changes are going to mean for the team as a whole, dynamic considerations may facilitate cooperation. Maintaining an element of dynamic uncertainty can foster greater cooperation among team members, as the continuation probability for future interactions remains a key driver of cooperative equilibrium. In contrast, if it is common knowledge among the workers that the collaboration is about to end, dynamic incentives for unselfish behavior unravel. Therefore, telling each worker privately about an upcoming termination may be less disruptive to cooperation than a public announcement.

More broadly, this paper suggests that the set of environments where cooperation between self-interested agents can be self-sustaining is more extensive than expected based on the existing studies. Our study demonstrates that even when individuals know of a finite deadline and possess only a modest level of uncertainty regarding their peers' beliefs, the cooperative equilibrium endures.

The critical feature in the analysis is that the continuation probability for the other player staying in the game *in the near future* must be bounded away from zero. We hypothesize that the findings of the present paper generalize to many other distributions for the exit times that satisfy this feature. The individual optimality condition in the cooperative equilibrium asks if by working today, one can induce the other player to work tomorrow. Implications for the more distant future do not enter into the optimality condition. For instance, whether beliefs (of any fixed order) are bounded or unbounded is unlikely to affect the result.

We conclude by suggesting an environment where the result will potentially fail. Suppose the exit times become increasingly synchronized over time (or that the players learn each other's exit times by waiting). In such an environment, the number of mixing periods must increase as the probability of holding the same exit time increases. Therefore, even if a mixed trigger strategy exists, it does not guarantee expected payoffs close to the cooperative payoffs because the cooperative face ends early. This counter-example suggests that a necessary condition for the positive result of this paper is that the continuation probability, $Pr(T_i > T_j | T_i \geq T_j - 1, T_j)$, is uniformly bounded away from zero. The exact conditions for joint distributions that enable long-run cooperative payoffs warrant further investigation.

A Proofs

A.1 Proof of Proposition 1

Proof. Suppose that Player i follows σ^{TR}. First notice that the one-shot deviation principle applies here. Player j does not have a profitable one-shot deviation at $t < T_j$ if the following holds

$$\delta(R-P) + (1-\delta)(-P) + \delta w^*_{T_j}(t+1) \geq \delta R + (1-\delta)0,$$

where $w^{TR}_{T_j}(t+1)$ is player j's value in period $t+1$ after mutual play of W and when both players follow σ^{TR} and Player j's personal exit time is T_j. The value $w^{TR}_{T_j}(t)$ is

$$w^{TR}_{T_j}(t) = \frac{\delta R(1-\delta^{-(T_j-t+1)}) - P(1-\delta^{-(T_j-t)})}{1-\delta}.$$

It is decreasing in t when $\delta \geq \frac{P}{R}$. If $\delta < \frac{P}{R}$, it is clear that cooperation cannot be sustained. Therefore, it is enough to check for incentive compatibility in the penultimate period when the continuation value is $w^{TR}_{T_j}(T_j) = \delta R$

In period $T_j - 1$, Player j is willing to cooperate if and only if

$$\delta(R-P) + (1-\delta)(-P) + \delta^2 R \geq \delta R \iff \delta^2 \geq \frac{P}{R}.$$

□

A.2 Proof of Lemma 1

Proof. Suppose that Player i follows σ^{TR}. As in the proof of Proposition 1 without correlation, the IC constraint for Player j is satisfied for all $t < T_j$ if it is satisfied at $t = T_j - 1$ The best response by Player j is to play W at $T_j - 1$ if and only if $Pr(T_i > T_j | T_i \geq T_j - 1, T_j)R \geq P$. This is equivalent to

$$Pr(\gamma = 0 | T_i \geq T_j - 1, T_j) \geq \frac{P}{\delta^2 R}, \tag{7}$$

where the conditional probability equals

$$Pr(\gamma = 0 | T_i \geq T_j - 1, T_j) = \frac{(1-\alpha)(1-(1-\delta)\delta)}{\alpha(1+\delta)\delta + (1-\alpha)}.$$

To get this, notice that after seeing T_j but before the game starts, Player j poses probability $\frac{\alpha(1-\delta^2)\delta^{2(T_j-1)}}{\alpha(1-\delta^2)\delta^{2(T_j-1)}+(1-\alpha)(1-\delta)\delta^{T_j-1}}$ on $\gamma = 1$.

We get the upper bound for α by plugging in the conditional probability and evaluating (7) at the limit $\delta = 1$, which yields the condition in the statement. □

A.3 Derivation of the posterior belief

In Appendix, we drop the sub-index δ from mixing probabilities $p_\delta(l)$ and $p_\delta(l|T_i)$.

When Player i follows $\tilde{\sigma}_i^S$, the posterior belief $Pr(\gamma = 1|T_j, a_t = (W, W)\forall t < T_j - k)$ is:

$$\frac{\alpha Pr(T_j|\gamma = 1)\Pi_{l=k+1}^K p(l)}{\alpha Pr(T_j|\gamma = 1)\Pi_{l=k+1}^K p(l) + (1-\alpha)Pr(T_j|\gamma = 0)\sum_{T_i=1}^\infty Pr(T_i|\gamma = 0)\Pi_{l=T_i-T_j+k+1}^K p(l)}$$

$$= \frac{\alpha(1-\delta^2)\delta^{2(T_j-1)}\Pi_{l=k+1}^K p(l)}{\alpha(1-\delta^2)\delta^{2(T_j-1)}\Pi_{l=k+1}^K p(l) + (1-\alpha)(1-\delta)^2\delta^{T_j-1}\sum_{T_i=1}^\infty \delta^{T_i-1}\Pi_{l=T_i-T_j+k+1}^K p(l)}.$$

Notice that $\sum_{T_i=1}^\infty \delta^{T_i-1}\Pi_{l=T_i-T_j+k+1}^K p(l) = \sum_{T_i=T_j-k}^\infty \delta^{T_i-1}\Pi_{l=T_i-T_j+k+1}^K p(l)$ because $\Pi_{l=T_i-T_j+k+1}^K p(l) = 0$ if $T_i < T_j + k$. Then, by using $n = T_i - T_j + k$ as the running index, the infinite sum can be rewritten as $\sum_{n=0}^\infty \delta^{n+T_j-k}\Pi_{l=n+1}^K p(l)$. The formula for $\hat{\alpha}(k)$ follows after a straightforward calculation once one plugs in the rewritten infinite sum and uses that $\Pi_{l=n+1}^K p(l) = 1$ if $T_i \geq K$.

A.4 Proof of Lemma 2

Proof. We can rewrite the indifference condition (4) as:

$$\frac{(1-\alpha)\delta^{-k+1}(1-\delta)\sum_{j=0}^\infty \delta^j \Pi_{l=j}^K p_\delta(l) + \alpha(1+\delta)\Pi_{l=k-1}^K p_\delta(l)}{(1-\alpha)\delta^{-k-1}(1-\delta)\sum_{j=0}^\infty \delta^j \Pi_{l=j}^K p_\delta(l) + \alpha(1+\delta)\Pi_{l=k+1}^K p_\delta(l)} = \frac{P}{R}, \tag{8}$$

which we simplify further by using $A(\delta) := (1-\delta)\sum_{j=0}^\infty \delta^j \Pi_{l=j}^K p(l)$. Notice that $A(\delta) \in (0, 1)$ and $A(\delta) \to 1$ as $\delta \to 1$.

We want Player j to be indifferent in periods $T_j - K, T_j - K + 1, \ldots, T_j - 1$, which means that (8) must hold for all $k \in \{1, 2, \ldots, K\}$. Notice that (8) has a recursive structure so that the denominator when $k = k'$ is the numerator of $k = k' - 1$. Using the recursive structure, we can derive the following equations when K is even:

$$(1-\alpha)\delta^{-k+1}A(\delta)+\alpha(1+\delta)\Pi_{l=k-1}^{K}p(l)=\left(\frac{P}{R}\right)^{\frac{K-k}{2}+1}\Big((1-\alpha)\delta^{-K-1}A(\delta)+\alpha(1+\delta)\Big)$$

for $k\in\{2,4,\dots,K\}$, (9)

$$(1-\alpha)\delta^{-k+1}A(\delta)+\alpha(1+\delta)\Pi_{l=k-1}^{K}p(l)=\left(\frac{P}{R}\right)^{\frac{K-k+1}{2}}\Big((1-\alpha)\delta^{-K}A(\delta)+\alpha(1+\delta)p(K)\Big)$$

for $k\in\{3,\dots,K-1\}$, (10)

$$(1-\alpha)\delta^{-2}A(\delta)+\alpha(1+\delta)\Pi_{l=2}^{K}p(l)=(1-\alpha)A(\delta)\frac{R}{P}, \tag{11}$$

where the last equation is the indifference condition for the penultimate period, $k=1$. This gives a system of K equation and unknowns.

One can solve $p(K)$ by observing that the LHS of the equations for $k=1$ and $k=3$ are the same. This gives:

$$p(K)=\frac{1-\alpha}{\alpha}\left(\left(\frac{R}{P}\right)^{\frac{K}{2}}-\delta^{-K}\right)\frac{A(\delta)}{1+\delta}. \tag{12}$$

Next, the equation for $k=K$, gives $p(K-1)$ when we plug in $p(K)$:

$$p(K-1)=\frac{\frac{P}{R}\Big((1-\alpha)\delta^{-K+1}A(\delta)+\alpha(1+\delta)\Big)-(1-\alpha)A(\delta)\delta^{-K+1}}{(1-\alpha)A(\delta)\left(\left(\frac{R}{P}\right)^{\frac{K}{2}}-\delta^{-K}\right)}. \tag{13}$$

Next, we calculate the necessary conditions for K that come from $p(K),p(K-1)\in(0,1)$. In (12), $p(K)>0$ always. A direct calculation gives that there exists $\tilde{\epsilon}(\delta)$ such that $\tilde{\epsilon}(\delta)\to 0$ as $\delta\to 1$ and $p(K)<1$ if $\alpha>\frac{(R/P)^{K/2}-1}{(R/P)^{K/2}+1}+\tilde{\epsilon}(\delta)$.

Similarly, for $p(K-1)>0$ in (13) a sufficient condition when δ is close to one is that $\alpha>\frac{R-P}{R+P}$, which holds by assumption. Then, $p(K-1)>1$ is implied by $\alpha<\frac{(R/P)^{K/2+1}-1}{(R/P)^{K/2+1}+1}-\hat{\epsilon}(\delta)$ with some function $\hat{\epsilon}$ such that $\hat{\epsilon}(\delta)\to 0$ as $\delta\to 1$. Combining the previous restrictions gives that $p(K),p(K-1)\in(0,1)$ if and only if $\alpha\in(\frac{(R/P)^{K/2}-1}{(R/P)^{K/2}+1}+\epsilon(\delta),\frac{(R/P)^{K/2+1}-1}{(R/P)^{K/2+1}+1}-\hat{\epsilon}(\delta)))=:(\underline{u}(K)+\tilde{\epsilon}(\delta),\overline{u}(K)-\hat{\epsilon}(\delta))$.

Based on the above reasoning, we make a guess that there exists $\epsilon(\delta)$ such that $\epsilon(\delta)\to 0$ as $\delta\to 0$ and that there exists $\bar{\delta}<1$ such that $p(k)\in(0,1)$ for all $k\in\{1,\dots,K\}$ if $\alpha\in(\underline{u}(K)+\epsilon(\delta),\overline{u}(K)-\epsilon(\delta))$ and if $\delta>\bar{\delta}$.

We have already argued that the above holds for $k=K,K-1$. We need to check that $p(k)\in(0,1)$ also for all lower k. We start by checking that $p(k)>0$

when δ is large. This can be seen by using an induction argument: we show that $p(k-1) > 0$ under the assumption that $p(l) > 0$ for all $l \geq k$. This is enough as we know that $p(K) > 0$.

Suppose first that k is even. Recall that $\Pi_{l=k+1}^{K} p_\delta(l) > 0$ by assumption. Hence, (8) gives that $p(k)p(k-1) > 0$ is equivalent to

$$\alpha(1+\delta)\left(\frac{P}{R}\right)^{\frac{K-k}{2}} - (1-\alpha)\delta^{-k-1}A(\delta)\left(\delta^2 - \left(\frac{P}{R}\right)^{\frac{K-k}{2}}\right) > 0, \tag{14}$$

where we have used (9) to evaluate the denominator of (8). By using the lower bound $\underline{u}(K)$ for α and taking the limit as $\delta \to 1$, one sees that (14) always holds for large enough δ.

Suppose next that k is odd. Again, use (9), now together with (12), to rewrite (8), which then can be used to solve for $p(k)p(k-1)$. For this case, the resulting formula is always strictly positive when $\delta^{-1} \leq \left(\frac{R}{P}\right)^{1/2}$. Hence, $p(k-1)$ is guaranteed to be positive also for off k when δ is sufficiently close to 1.

Now, we turn to checking that $p(k) < 1$. To do this, take the difference of the recursive indifference conditions (9) and (10) for k and $k+1$ (it is useful to plug in $p(K)$). Using that $p(l) > 0$ for all l, this gives a condition for $p(k-1) < 1$. For k even, the condition is hardest to achieve when α is large and the converse is true when k is odd. Hence we need to check that the condition is satisfied (weakly) when $\alpha = \overline{u}(K)$ and $\alpha = \underline{u}(K)$ respectively for $\delta \to 1$. A direct calculation verifies that the conditions are indeed satisfied in both cases.

Therefore, we conclude that there exists $\epsilon(\delta)$ that goes to zero as $\delta \to 1$ such that all $p(k) \in (0,1)$ if $\alpha \in (\underline{u}(K) + \epsilon(\delta), \overline{u}(K) - \epsilon(\delta))$.

Finally, notice that for all $\alpha \in (\frac{R-P}{R+P}, 1) \setminus \{a \in [0,1] : a = 1 - (\frac{P}{R})^n \text{ for some } n \in \mathbb{N}\}$, there exists $\overline{\delta} < 1$ such that $\alpha \in (\underline{u}(K) + \epsilon(\delta), \overline{u}(K) - \epsilon(\delta))$ for some even K. This follows as the union of all $(\underline{u}(K), \overline{u}(K))$ over all even natural numbers K contains all real numbers in $a \in (\frac{R-P}{R+P}, 1)$, excluding $\{a \in [0,1] : a = 1 - (\frac{P}{R})^n \text{ for some } n \in \mathbb{N}\}$. Therefore, we conclude that for all $\alpha \in (\frac{R-P}{R+P}, 1) \setminus \{a \in [0,1] : a = 1 - (\frac{P}{R})^n \text{ for some } n \in \mathbb{N}\}$, there exists an even number $K(\alpha)$ such that there exists valid mixing probabilities $(p(1), p(2), \ldots, p(K))$ that make the players indifferent in periods $T_i - K, T_i - K + 1, \ldots, T_i - 1$. Furthermore, $K(\alpha)$ is finite for all $\alpha < 1$. $\square$

A.5 Proof of Lemma 3

Proof. For each $\alpha \in \mathbb{R}^n$, define isohypersurface $X^i(\alpha)$ as:

$$X^i(\alpha) := \{x \in C : F_i(x) = \alpha \text{ and } (x - x^*)'b \leq (z - x^*)'b \\ \text{for all } b \in \mathbb{R}^n, z \in C \text{ s.t. } F_i(z) = \alpha.\}$$

The condition $(x - x^*)'b \leq (z - x^*)'b$ for all $z \in C : F_i(z) = \alpha$ means that if $F_i(x) = \alpha$ has multiple solutions, we choose the one that is closest to x^*.

Next, we define $C_i(\epsilon') \subset C$ such that all of its elements are in between of $X^i(\epsilon')$ and $X^i(-\epsilon')$. The formal definition is: $C_i(\epsilon') := \{x \in C : x_j \in [\min\{\overline{x}^i_j(x,j), \underline{x}^i_j(x,j)\}, \max\{\overline{x}^i_j(x,j), \underline{x}^i_j(x,j)\}]$ for all j where $\overline{x}(x,j)$ and $\underline{x}(x,j)$ are such that $\overline{x}^i(x,j) \in X^i(\epsilon')$ and $\overline{x}^i_k(x,j) = x_k$ for all $k \neq j$ and $\underline{x}^i(x,j) \in X^i(-\epsilon')$ and $\underline{x}^i_k(x,j) = x_k$ for all $k \neq j$. If $\overline{x}(x,j)$ or $\underline{x}(x,j)$ does not exist, use $\min\{C_j\}$ and $\max\{C_j\}$ for the lower and upper bars. Let $C(\epsilon') := \cap_{i=1}^n C_i(\epsilon)$. Figure 2 illustrates the constructed set $C(\epsilon')$ when $n = 2$.

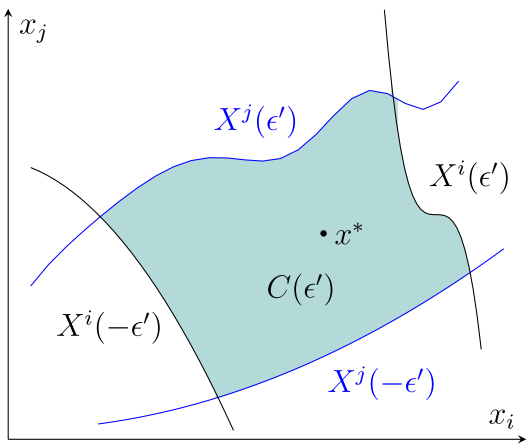

Figure 2: $C_i(\epsilon')$ is defined as the region between $X^i(\epsilon')$ and $X^i(-\epsilon')$, and similarly $C_j(\epsilon')$ is the region between $X^j(\epsilon')$ and $X^j(-\epsilon')$. The intersection of $C_i(\epsilon')$ and $C_j(\epsilon')$ gives $C(\epsilon')$.

By construction, $|F_i(x)| \leq \epsilon'$ for all $x \in C(\epsilon')$. Then, we have

$$\max_{x \in C} ||x, x^*|| = \max_{x \in C \setminus int(C)} ||x, x^*|| \leq M||F(x)|| \leq M\sqrt{n}\epsilon,$$

where the second to last inequality comes from the sign condition by the following argument. $x \in C \setminus int(C)$ implies that $x \in X^i(\epsilon') \cup X^i(-\epsilon')$ for some i. By the construction of $X^i(\epsilon')$ and $X^i(-\epsilon')$, we know that each $x \in X^i(\epsilon') \cup X^i(-\epsilon')$ is the solution closest to x^* for the problem $F(x) = y'$ with some y' such that $|y'_i| = \epsilon'$ and $|y'_j| \leq \epsilon'$. Now, the sign condition gives that for small enough ϵ', $C(\epsilon')$ is compact and satisfies: $x \in C(\epsilon') \Rightarrow ||x - x^*|| \leq M\sqrt{n}\epsilon'$.

Next, choose large enough m such that $F_i^m(x) > 0$ for all $x \in X^i(\epsilon')$ and $F_i^m(x) < 0$ for all $x \in X^i(-\epsilon')$ for all i (such m exists by uniform convergence). Now the multidimensional intermediate value theorem (Poincare-Miranda theorem) implies that $F_i^m(x) = 0$ has a solution in $C(\epsilon)$.[11] Finally, the claim follows by setting $\epsilon' = \epsilon/(2M\sqrt{n})$, which implies that $||x - x^*|| \leq \epsilon/2 < \epsilon$. □

A.6 Proof of Theorem 1

Proof. The claim follows directly from Lemma 1 if $\alpha < \frac{R-P}{R+P}$. Hence, we focus on the case where $\alpha > \frac{R-P}{R+P}$.[12]

Suppose that $p(l|T_i) = 1$ for all $l \in (K, T_i]$ with some K. We verify this guess as part of the proof. We prove the existence of a non-stationary trigger strategy by showing that for all $p(l|T_i \neq T_j) \in [0,1]$ when $l \leq K$, we can find $p(1|T_j), \ldots, p(K|T_j) \in (0,1)$ such that the indifference condition holds for type T_j for all $k \in \{1, \ldots, K\}$.

First, take the limit of the stationary indifference condition (8) as $\delta \to 0$:

$$\frac{(1-\alpha) + 2\alpha\Pi_{l=k-1}^{K} p(l)}{(1-\alpha) + 2\alpha\Pi_{l=k+1}^{K} p(l)} - \frac{P}{R} = 0, \tag{15}$$

which defines a set of continuous equations with a unique solution p^* where $p^*(l) \in (0,1)$ for all $l \in \{1, \ldots, K\}$.

The indifference condition for type T_j in period $T_j - k$ in the non-stationary

[11] The usual formulation of Poincare-Miranda theorem assumes that $F_i^m(x) < 0$ when $x_i = -1$ and $F_i^m(x) > 0$ when $x_i = 1$. The formulation used in the proof is equivalent to the standard formulation after space change.

[12] Notice that $\alpha = \frac{R-P}{R+P}$ belongs to the excluded set of measure-0.

problem is

$$\frac{(1-\alpha)\delta^{-k+1}(1-\delta)\sum_{j=0}^{\infty}\delta^j\Pi_{l=j}^{K}p(l|T_j-k+j)+\alpha(1+\delta)\Pi_{l=k-1}^{K}p(l|T_j)}{(1-\alpha)\delta^{-k-1}(1-\delta)\sum_{j=0}^{\infty}\delta^j\Pi_{l=j}^{K}p(l|T_j-k+j)+\alpha(1+\delta)\Pi_{l=k+1}^{K}p(l|T_j)}-\frac{P}{R}=0, \tag{16}$$

which is continuous in $p(l|T_j)$.

When one sets $p(l|T_j)$ equal to $p(l)$, (16) converges to (15) as $\delta \to 1$ for any $p(l|T_i \neq T_j) \in [0,1]$. Furthermore, the convergence is uniform when $(1-\alpha)\delta^{-k-1}(1-\delta)\sum_{j=0}^{\infty}\delta^j\Pi_{l=j}^{K}p(l|T_j-k+j)+\alpha(1+\delta)\Pi_{l=k+1}^{K}p(l|T_j)$ and $(1-\alpha)+2\alpha\Pi_{l=k+1}^{K}p(l)$ are bounded away from zero, which is guaranteed when all $p(l)$ are non-negative.

Then, we show that (15) satisfies the sign condition in Lemma 3. First notice that the following set of equations has a solution in $\mathbb{R}^K$ for all $r \in \mathbb{R}^K$ such that $r_k > -P/R$ for all $k \in \{1,\ldots,K\}$:[13]

$$\frac{(1-\alpha)+2\alpha\Pi_{l=k-1}^{K}p(l)}{(1-\alpha)+2\alpha\Pi_{l=k+1}^{K}p(l)}-\frac{P}{R}=r_k. \tag{17}$$

Let p^r denote that solution. Define a decreasing sequence (r_i) such that $r_i \to 0^K$ as $i \to \infty$. The corresponding sequence (p^{r_i}) is bounded, and hence it has a convergent subsequence. Let $\tilde{p}$ be the limit of such a subsequence. By the continuity of (15), this then implies that $\tilde{p}$ is a solution to (15), and hence we must have $\tilde{p} = p^*$. Hence (15) satisfies the sign condition.

We have shown that all conditions in Lemma 3 are satisfied. Hence, we conclude that for all $\epsilon > 0$, there exists $\bar{\delta} < 1$ such that the set of K equations defined by (16) and fixing $p(l|T_i \neq T_j)$ to be any numbers between 0 and 1 has a solution that is at most ϵ away from p^* when $\delta > \bar{\delta}$. Then, we choose ϵ small enough that all $p(l|T_j)$ in the solution must be in $(0,1)$ as $p^*(l) \in (0,1)$. Finally, since the argument holds for all $T_j > K$ and for any $p(l|T_i \neq T_j) \in [0,1]$ for $l \leq K$, this implies that there exists a solution such that (16) holds for all $T_j > K$ in non-stationary mixing probabilities that converge to p^* as $\delta \to 1$. Therefore, we conclude that

[13] One easy way of seeing this is to observe that the first term can be assigned to take any values in $\mathbb{R}$ just by changing $p(k-1)$ as long as its denominator and all other $p(l)$ are non-zero. Then a recursive argument shows that these conditions are satisfied.

for large enough δ the resulting non-stationary mixed trigger strategy profile is an equilibrium.

This proves Theorem 1 as the *ex ante* payoff from such a strategy profile converges to $R - P$ as $\delta \to 1$. $\square$

B Derivations for the related models

B.1 Proof of Proposition 2

The proof follows largely the same steps as the proof of Theorem 1.

Similar to the main model, (6) has a recursive structure so that the denominator when $k = k'$ is the numerator of $k = k' - 2$. This gives the following equations when K is even:

$$(1-\alpha) + \alpha \Pi_{l=k-1}^{K} p(l) = \left(\frac{P}{R}\right)^{\frac{K-k}{2}+1} \quad \text{for } k \in \{2, 4, \ldots, K\}$$
$$(1-\alpha) + \alpha \Pi_{l=k-1}^{\hat{t}} p(l) = \left(\frac{P}{R}\right)^{\frac{K-k+1}{2}} (1 - \alpha + \alpha p(K)) \quad \text{for } k \in \{3, \ldots, K-1\}$$
$$(1-\alpha) + \alpha \Pi_{l=2}^{K} p(l) = (1-\alpha)\frac{R}{P}.$$

One can solve $p(K)$ by observing that the LHS of the equations for $k = 1$ and $k = 3$ are the same. This gives:

$$(1-\alpha)\frac{R}{P} = \left(\frac{P}{R}\right)^{\frac{K-2}{2}} (1 - \alpha + \alpha p(K)) \iff p(K) = \frac{1-\alpha}{\alpha}\left(\left(\frac{R}{P}\right)^{\frac{K}{2}} - 1\right). \tag{18}$$

Now, the equation for $k = K$, gives $p(K-1)$:

$$(1-\alpha) + \alpha p(K) p(K-1) = \frac{P}{R} \iff p(K-1) = \frac{\frac{P}{R} - (1-\alpha)}{(1-\alpha)\left(\left(\frac{R}{P}\right)^{\frac{K}{2}} - 1\right)}. \tag{19}$$

Next, we calculate the necessary and sufficient conditions for $p(K), p(K-1) \in [0, 1]$. In (18), $p(K) > 0$ always. A direct calculation gives that $p(K) \leq 1$ if and only if $1 - \alpha \leq (P/R)^{K/2}$.

Similarly, $p(K-1) \geq 0$ in (19) is equivalent to $1-\alpha \leq P/R$ and $p(K-1) \leq 1$ is equivalent to $1-\alpha \geq (P/R)^{K/2+1}$. Combining the previous restrictions gives that $p(K), p(K-1) \in [0,1]$ if and only if $1-\alpha \in [(P/R)^{K/2+1}, (P/R)^{P/2}] =: [\underline{v}(K), \overline{v}(K)]$.

Next, we check that $p(k) \in [0,1]$ for all $k \in \{1, 2, \ldots, K\}$ if $1-\alpha \in [\underline{v}(K), \overline{v}(K)]$. It follows directly from the recursive indifference condition (6) that $p(k) > 0$. This can be seen by first observing that $p(k)p(k-1) \leq 0$ would imply $(1-\alpha)R/P \geq (1-\alpha) + \alpha\Pi_{l=k+1}^{K} p(l)$, but we know that $(1-\alpha) + \alpha\Pi_{l=2}^{K} p(l) = (1-\alpha)R/P$ and that $\Pi_{l=k+1}^{K} p(l) > \Pi_{l=2}^{K} p(l)$ for all $k > 1$. Hence, $p(k) > 0$ if $p(l) \in (0,1)$ for $l > k$. Then $p(k)p(k-1) > 0$ for all k implies $p(k) > 0$ as we know that $p(K) > 0$, which rules out negative $p(k)$.

We also get from (6) that $p(k) \leq 1$. To see this, take the difference of the recursive indifference conditions for k and $k+1$. If k is odd, that results in the right-hand side being negative for all $1-\alpha \leq \overline{v}(K)$ (it is useful to plug in $p(K)$). The left-hand side is negative if and only if $p(k-1) < 1$. The same expression for k even has a negative right-hand side when $1-\alpha > \underline{v}(K)$ and the left-hand side is negative if and only if $p(k-1) < 1$.

The union of $[\underline{v}(K), \overline{v}(K)]$ over all even natural numbers $K(\alpha)$ (and zero) contains all positive real numbers. Therefore, we conclude that we can always find an even number $K(\alpha)$ such that there exist valid mixing probabilities $(p(1), p(2), \ldots, p(K(\alpha)))$ that make the informed players indifferent in periods $T_i - K(\alpha), T_i - K(\alpha) + 1, \ldots, T_i - 1$. This verifies the guess we made at the beginning of the analysis. Furthermore, $K(\alpha)$ is finite for all $\alpha < 1$. Finally, notice that the incentive compatibility holds for the uninformed type when δ is sufficiently large. Hence, we have found an equilibrium where players receive payoff $R-P$ in all but the last $K(\alpha)$ periods whenever the personal exit times are above $K(\alpha)$, which happens almost surely when $\delta \to 1$. □

B.2 Proof of Proposition 3

The proof for this case is largely identical to that of Proposition 2. Therefore, we only comment on the key differences. First, Player A's indifference condition is otherwise identical to that of the informed players in the previous model, except that Player A knows that an informed Player B will shirk at $T-1$. Therefore, one needs to 'postpone' the mixing probabilities from the previous case by one for this case: $p^B(k) = p(k-1)$ where $p(l)$ is as in the proof of Proposition 2. Hence, informed Player B starts mixing in period $T-K-1$.

Second, informed Player B's indifference condition is different: $(p^A(1), \ldots, p^A(K))$ solve $p^A(k)p^A(k-1) = P/R$ for all $k \in 2, \ldots, K+1$. As $p^(K+1) = 1$ under our construction, we get an explicit solution: $P^A(k) = P/R$ for all k even and $P^A(k) = P/R$ for all k odd. As the incentive compatibility condition is satisfied for the uninformed Player B when δ is large, we conclude that the proposed strategy profile is a perfect Bayesian equilibrium for all large enough δ.

□

References

Abreu, Dilip and David Pearce, "Bargaining, Reputation, and Equilibrium Selection in Repeated Games with Contracts," *Econometrica*, 2007, *75* (3), 653–710.

Arribas, Iván and Amparo Urbano, "Repeated games with probabilistic horizon," *Mathematical Social Sciences*, 2005, *50* (1), 39–60.

Atakan, Alp E. and Mehmet Ekmekci, "A two-sided reputation result with long-run players," *Journal of Economic Theory*, 2013, *148* (1), 376–392.

Baye, Michael R. and Dennis W. Jansen, "Repeated Games with Stochastic Discounting," *Economica*, 1996, *63* (252), 531–541.

Benoit, Jean-Pierre and Vijay Krishna, "Finitely Repeated Games," *Econometrica*, 1985, *53* (4), 905–922.

Bernheim, B Douglas and Aniruddha Dasgupta, "Repeated games with asymptotically finite horizons," *Journal of Economic Theory*, 1995, *67* (1), 129–152.

Bó, Pedro Dal, "Cooperation under the Shadow of the Future: Experimental Evidence from Infinitely Repeated Games," *American Economic Review*, December 2005, *95* (5), 1591–1604.

Carlsson, Hans and Eric van Damme, "Global Games and Equilibrium Selection," *Econometrica*, 1993, *61* (5), 989–1018.

Embrey, Matthew, Guillaume R Fréchette, and Sevgi Yuksel, "Cooperation in the Finitely Repeated Prisoner's Dilemma*," *The Quarterly Journal of Economics*, 08 2017, *133* (1), 509–551.

Engle-Warnick, J. and R.L. Slonim, "Inferring repeated-game strategies from actions: evidence from trust game experiments.," *Economic Theory*, 2006, *28*, 603–632.

Fanning, Jack, "Reputational Bargaining and Deadlines," *Econometrica*, 2016, *84* (3), 1131–1179.

Fudenberg, Drew and Yuichi Yamamoto, "The folk theorem for irreducible stochastic games with imperfect public monitoring," *Journal of Economic Theory*, 2011, *146* (4), 1664–1683.

Hörner, Johannes, Takuo Sugaya, Satoru Takahashi, and Nicolas Vieille, "Recursive methods in discounted stochastic games: An algorithm for $\delta \rightarrow 1$ and a folk theorem," *Econometrica*, 2011, *79* (4), 1277–1318.

Kreps, David M and Robert Wilson, "Reputation and imperfect information," *Journal of Economic Theory*, 1982, *27* (2), 253 – 279.

_ **, Paul Milgrom, John Roberts, and Robert Wilson**, "Rational cooperation in the finitely repeated prisoners' dilemma," *Journal of Economic Theory*, 1982, *27* (2), 245 – 252.

Massó, Jordi, "A Note on Reputation: More on the Chain-Store Paradox," *Games and Economic Behavior*, 1996, *15* (1), 55–81.

Milgrom, Paul and John Roberts, "Predation, reputation, and entry deterrence," *Journal of Economic Theory*, 1982, *27* (2), 280 – 312.

Neyman, Abraham, "Cooperation in Repeated Games When the Number of Stages is not Commonly Known," *Econometrica*, 1999, *67* (1), 45–64.

Rubinstein, Ariel, "The Electronic Mail Game: Strategic Behavior under "Almost Common Knowledge."," *American Economic Review*, 1989, *79* (3), 385–91.

Selten, Reinhard and Rolf Stoecker, "End behavior in sequences of finite Prisoner's Dilemma supergames A learning theory approach," *Journal of Economic Behavior & Organization*, 1986, *7* (1), 47–70.

_ **, Michael Mitzkewitz, and Gerald Uhlich**, "Duopoly Strategies Programmed by Experienced Players," *Econometrica*, 1997, *65* (3), 517–556.

Weinstein, Jonathan and Muhamet Yildiz, "A Structure Theorem for Rationalizability with Application to Robust Predictions of Refinements," *Econometrica*, 2007, *75* (2), 365–400.