Cooperation in Continuous Dilemma and Uncertain Reaction Lag

In-Uck Park*

University of Bristol and Sungkyunkwan University

May 2, 2014

Abstract. *This paper shows that cooperation can be sustained until close to the end of a finite-horizon, continuous-time prisoners' dilemma when there is informational asymmetry in how quickly players can respond. The simulated equilibrium closely replicates recent experimental results (Friedman and Oprea, 2012, AER). The core argument is extended to a class of canonical preemption games with private information on the player's payoff margin of preempting relative to being preempted, that can be applied to other well-known examples of conflict such as the centipede game.* (*JEL* Classification Codes: C72, C73, D82)

Keywords: continuous time, cooperation, prisoners' dilemma, reaction lag.

1 Introduction

The prisoners' dilemma is one of the simplest and best-known examples of a stark conflict between the self-interest of rational players and mutually beneficial cooperation. Because it is a strictly dominant strategy to "defect" (D), the prediction based on backward induction is logically compelling in that both players should always defect in any finitely repeated games, but substantial amounts of cooperation are observed both casually and in various experiments. Notwithstanding multiple ingenious theories proposed to account for observed cooperation (to be discussed later), the dilemma continues to be an active topic of investigation to date.

Recently, Friedman and Oprea (2012) report that in their experiment of a continuous-time prisoners' dilemma game, the rate of cooperation is much higher (around 90% of a one-minute experiment) than in many other previous experiments in discrete-time that also exhibited substantial levels of cooperation. They also report that coordination is typically achieved in the first few seconds by one player switching to cooperate (C) with the other quickly following suit, and then one player switches back to D near the end of the game,

*This project started while I was a Richard B. Fisher member at the Institute for Advanced Study, Princeton, and I thank their hospitality. I also thank Heski Bar-Isaac, Marco Casari, Yeon-Koo Che, Martin Cripps, Dan Friedman, Johannes Hörner, Sergey Izmalkov, Philippe Jehiel, David Levine, Bentley MacLeod, Stephen Morris, Hyun Song Shin, and Aleksey Zinger for helpful comments and suggestions; and participants at the CEPR European Summer Symposium 2012 in Gerzensee, and seminars at City University London, Korea University, Seoul National University, and University College London. All remaining errors are mine.
Email: i.park@bristol.ac.uk

again immediately followed by the other player. To my knowledge, this paper is the first to theoretically justify such a behavior as an equilibrium.

A key to sustaining cooperation in finite-horizon settings of the prisoners' dilemma is how to prevent unraveling from the end of the game. Consider two players currently cooperating. Because cooperation is a dominated option in the short term, either player may continue to cooperate only if he or she believes that doing so would induce the opponent to cooperate longer than otherwise. In discrete-time settings, this is not possible in the last period because players cannot influence each other's behavior at that point, and backward induction unravels cooperation all the way back to the first period. However, this prediction may rely on the artificial discretization of time into units for action. In fact, Bigoni, *et al.* (2013) report that behavior patterns emerge in continuous-time experiments that are qualitatively different from those observed in discrete-time experiments.

In continuous-time environments where there is no last period, players may hope to prolong their opponent's cooperation by extending their own cooperation longer. However, both players still aspire to preempt their opponent by defecting first before the game ends. How may this dilemma be reconciled to sustain cooperation in the face of eventual defection?

This paper shows that this is possible if players are unsure as to how quickly their opponent can respond when preempted. In this case, players may well be perceived as being more willing to cooperate a little longer if they can respond quicker when preempted. This creates uncertainty over the opponent's defection time, which plants the seed of doubt that fends off the aforementioned unraveling of cooperative behavior, as explained below.

When a player gets preempted, she obviously wishes to respond immediately by defecting herself as well but the earliest point at which she may do so is not pinned down. This is a well-known technical issue in continuous-time games that stems from the fact that real numbers are not well-ordered. A technique widely adopted in the literature to circumvent this problem is to introduce a small lapse of time that has to pass before the next move can take place.

It is plausible to assume that it takes a little time for a player to react to moves made by other players. In business negotiations, for example, it may reflect the time for negotiators to obtain approval on terms of a counteroffer from the headquarters, as suggested in Perry and Reny (1993). In the experiments of Friedman and Oprea (2012) it would reflect the period of time, albeit very short, for a player to perceive the opponent's changed action, determine whether and how to respond, and implement the decision. In addition, it also seems plausible that the amount of time needed to implement a reaction may vary across players because of their innate characteriztics (e.g., some people are physically more agile) and/or external factors (e.g., some firms are more efficient than others in decision making).

Therefore, this paper models *heterogeneous reaction lags as private information,* which is the key to resolving the aforementioned dilemma. The intuition is that the quicker the player can respond, the longer she is willing to continue cooperation because she can cope better when preempted. Then each player perceives the opponent's defection time as a random variable, and at any point in time toward the end, there always is a chance that the opponent, being sufficiently quick, will cooperate a little longer, which in turn justifies each player's willingness to cooperate a little longer if she is quick enough herself.

To highlight this core intuition, Section 2 illustrates a simple equilibrium in which cooperation prevails until sometime toward the end when the less quick player of the two defects first, followed by the other player. This baseline analysis is extended in two directions.

First, for broader applicability, the fundamental underlying logic is extended in Section 3 to a stylized class of canonical preemption games in which two players wish to preempt the other later than sooner because the passage of time before preemption is beneficial. Players are privately informed of their payoff margin/net gain from preempting as opposed to being preempted. Subject to a minor condition, only two classes of equilibria are shown to exist. The first class of equilibria is fully separating, in which the two players defect/preempt gradually such that the later they defect, the lower their payoff margin (type) is. These equilibria are time-invariant: An equilibrium continues to be one when the defection time is shifted by a constant. The other class is semi-pooling, in which both players defect at the beginning of the game if their types are above certain thresholds, but would wait a certain amount of time before defecting if they are of the threshold type, and gradually longer as their types are lower. Sufficient conditions for the distribution of types are identified for such equilibria to exist. In addition, cooperation in a centipede game is illustrated as another application where heterogenous payoff margins reflect varying levels of altruism a player may possess.

Second, the prisoners' dilemma game itself is further analyzed to account for the richer dynamics observed in Friedman and Oprea's (2012) experiments, particularly the transition to mutual cooperation in the early stages, as well as the eventual defection near the end of the game. This extension (presented in Appendix A for the smooth flow of exposition) requires the formalization of an extensive-form game between players with different reaction lags, which is a novel development in itself. Noteworthy is that the equilibrium thus obtained closely replicates the experimental results in Friedman and Oprea, as explained at the end of Section 2.

Section 4 concludes with a summary and a discussion on this paper's contributions to the literature.

2 An Illustration of Core Results

Suppose that two players $i = 1, 2$, play a prisoners' dilemma game (PD) represented by

1\ 2	C	D
C	(1,1)	(0,h)
D	(h,0)	(ℓ, ℓ)

where $0 < \ell < 1 < h$ over a unit time interval $[0, 1]$. At any given point in time during this interval both players may switch between C and D as they wish subject to "reaction lag" constraints explained below. Every switch in either player's action is observed by the other player instantaneously, and the players receive flow payoffs determined by the action pair prevailing at each point $t \in [0, 1]$ without discounting.

As the key result is how cooperation (C) may be sustained despite the anticipated eventual defection (D), this section illustrates an equilibrium in which both players start the game with C and defect shortly before the game ends. Once a player, say j, defects at time $t < 1$, the other player i wishes to respond as soon as possible by defecting herself as well, but the earliest point in time at which she may do so is not pinned down. As explained earlier, this issue is addressed by introducing a small reaction lag, θ_i, such that $t + \theta_i$ is the first instant of time at which player i may respond to her opponent's move at t.

The value of the reaction lag, θ_i, is private information of player i and thus is referred to as her *type*. This is a crucial difference from existing models with reaction lags[1] and is a key factor behind this paper's results, as explained in Section 4. For illustration purposes, θ_1 and θ_2 are assumed to be independent draws from a uniform distribution over a small interval $(0, \bar{\theta})$ in this section, but the main result holds for a much wider class of distributions and also for the cases in which the distribution differs between two players (see Section 3).

The aforementioned equilibrium is now illustrated. For the cleaner exposition of the core insight, it is heuristically explained here, and precise technical details are provided later in Section 3 and Appendix A. Recall that both players start with C in the considered equilibrium. In the contingency that one player, say i, switches to D at $t < 1$, it is intuitively clear that the following constitutes a continuation equilibrium: Player i keeps to D until the end of the game, and player $j \neq i$ switches to D as soon as possible, that is, at $t+\theta_j$ (provided that $t + \theta_j \leq 1$), and keeps to D until the end of the game.[2] Refer to this continuation equilibrium as the "end game" for the contingency that player i preempts player j at t. Then the end-game payoffs in this contingency are denoted by $V_i(t, \theta_i, \theta_j)$ and $v_j(t, \theta_j, \theta_i)$ for player i (the preemptor) and player j (the preempted), respectively, and can be expressed as

$$\begin{cases} V_i(t, \theta_i, \theta_j) = h\theta_j + \ell(1 - t - \theta_j) \quad \text{and} \quad v_j(t, \theta_j, \theta_i) = \ell(1 - t - \theta_j) & \text{if} \quad t \leq 1 - \theta_j \\ V_i(t, \theta_i, \theta_j) = h(1 - t) \quad \text{and} \quad v_j(t, \theta_j, \theta_i) = 0 & \text{if} \quad t > 1 - \theta_j \end{cases} \tag{1}$$

because player i (j) receives a flow payoff of h (0) until player j defects, after which both receive ℓ until the end of the game. Therefore, the payoffs from the whole game in this contingency are

$$t + V_i(t, \theta_i, \theta_j) \quad \text{and} \quad t + v_j(t, \theta_j, \theta_i) \tag{2}$$

for players i and j, respectively.

The end game described above specifies a continuation equilibrium in the contingency that either player "preempts" by defecting first at an arbitrary point of the game. Given this, to describe an equilibrium of the whole game, the point in time at which it is optimal for each player to preempt (conditional on not having been preempted already) needs to be specified as follows.

First, note that it is suboptimal for either player of any type never to defect (unless preempted) because she can increase her expected payoff by preempting the opponent slightly before her opponent is anticipated to defect (or slightly before the end of the game if the opponent is expected never to defect). Conditional on preempting the opponent, however, the players prefer defecting later than sooner because doing so would prolong the duration of a cooperation payoff of 1 prior to preemption while curtailing that of a smaller non-cooperative payoff of ℓ later. However, by delaying defection either player risks being preempted in the

[1] A notion of uncertain reaction lag is modeled by Ma and Manove (1993) in an alternating move bargaining game, but as common uncertainty to players rather than as private information.

[2] The off-equilibrium belief is that players stick to their respective strategy even when their opponents behave differently in the continuation game. See Appendix A for further details.

meantime, in which case the loss is smaller when they can respond more quickly, i.e., when their reaction lag is shorter. Therefore, quicker players are more willing to wait a little longer before defecting and thus end up defecting later in equilibrium.

Such an equilibrium can be represented by a type-contingent preemption time denoted by the function $\tau : (0, \bar{\theta}) \to [0, 1]$, at which the players will defect conditional on not having been preempted by then. As quicker players defect later, τ is a decreasing function of θ. Illustrated below is such an equilibrium that is symmetric (i.e., $\tau(\cdot)$ is common to the two players) and according to which preemption may take place until the last moment, i.e., $\tau(\theta_i) \to 1$ as $\theta_i \to 0$:

$$\tau(\theta_i) \;=\; 1 - \rho(\theta_i) \;\text{ for }\; i = 1, 2, \quad \text{where} \quad \rho(\theta_i) := \frac{h\theta_i}{1 - \ell} > \theta_i. \tag{3}$$

Here $\rho(\theta_i)$ is the remaining length of time until the end of game when a player of type θ_i preempts, which is a linear function because of the uniform distribution of θ_i assumed for this illustration. The introduction of $\rho(\cdot)$ facilitates the following exposition.

In the equilibrium specified in (3), each player perceives her opponent's defection time to be stochastic and, in particular, uniformly likely over an interval of time $(1 - \rho(\bar{\theta}), 1)$ at the end of the game. As time passes within this interval, therefore, the hazard rate at which the opponent may defect in the next instant increases. This means that when a player continues cooperating longer, the risk of being preempted increases while the potential benefit from prolonged mutual cooperation remains constant (because $\rho(\theta_i) > \theta_i$). As a result, the lower the player's type, the longer she is willing to continue cooperating because she can cope better if she gets preempted.

To verify that this indeed constitutes an equilibrium, consider the contingency that the game reaches without defection the time at which either player of a certain type $\theta_i \in (0, \bar{\theta})$ is supposed to defect, that is, $\tau(\theta_i)$. At that point, from the fact that the opponent has not defected by then, both players update their opponent's type to be uniformly distributed over $(0, \theta_i)$. Based on this updated belief, player i must find it optimal to defect at $\tau(\theta_i)$ if her type is θ_i but to wait if her type is below θ_i in the considered equilibrium.

To check this equilibrium condition, calculate player i's expected payoff of waiting an infinitesimal amount of time Δ from $\tau(\theta_i)$ before defecting. In the contingency that she gets preempted between $\tau(\theta_i)$ and $\tau(\theta_i) + \Delta$ (i.e., the opponent defects in between), the lower her type is, the higher her payoff is (because a lower type means that she can respond quicker). Otherwise, her payoff is the same regardless of her type because she would then preempt at $\tau(\theta_i)$ and her payoff in that case depends on the opponent's type, not her own.

Therefore, if the marginal net change in the expected payoff of player i from waiting at $\tau(\theta_i)$ is 0 when her type is θ_i, then any lower type must strictly prefer to wait. The change in the payoff for a θ_i-type agent is $\Delta(1 - \ell)$ if she does not get preempted (because the wait prolongs the duration of (C, C) by Δ, which reduces the duration of (D, D) by the same amount of time at the end of the game[3]) and it is $-h\theta_i$ if she gets preempted (because she, not her opponent, faces the so-called "sucker" payoff 0, as opposed to h, for the duration of θ_i). Because the latter happens with probability $\Delta/\rho(\theta_i)$, the net change in her payoff is first-order approximated as $\Delta(1 - \ell) - \Delta h\theta_i/\rho(\theta_i)$, which is 0 by (3).

[3]Note that this is because $\rho(\theta_i) > \theta_i$.

In sum, it is verified that, conditional on both players defecting according to (3), the marginal benefit of player i from waiting an infinitesimal amount of time from $\tau(\theta_i)$ is 0 if her type is θ_i, and therefore it is optimal for her to wait if her type is lower than θ_i. Note that the optimality still needs to be verified for a θ_i-type player to indeed defect at $\tau(\theta_i)$, by showing that delaying defection to any later time is no better, and this is addressed in the next section by straightforward calculations.

Note that $\rho(\bar{\theta}) < 1$ if $\bar{\theta}$ is sufficiently small such that in equilibrium the two players continue with C until the player with a larger reaction lag defects first, followed by the other player after the lapse of her reaction lag. Finally, the optimality of not preempting until $\tau(\bar{\theta}) = 1 - \rho(\bar{\theta})$ is straightforward because preempting before $\tau(\bar{\theta})$ curtails the duration of (C, C) while lengthening that of (D, D) later by the same amount.

To better illustrate the core insight, this section focuses on an equilibrium in which the players start the game with C as the initial action. Together with the end-game specification summarized by the aforementioned payoffs (1), this allowed for an analysis of the game in a reduced form as if the game ends with the first defection. The next section develops this reduced-form approach further and presents a more complete characterization of possible equilibria for a broader class of preemption games.

The analysis is also extended to account for the richer dynamics observed in the aforementioned experiments in Friedman and Oprea. In particular, this fuller analysis theoretically explains the transition to mutual cooperation in the early stages when they start with (D, D) as well as the eventual defection near the end of the game. This extension requires the development of an appropriate model to analyze an extensive-form game in which players with different reaction lags may make multiple moves. A detailed discussion is shown in Appendix A because it is, albeit of interest on its own, not essential for the core insight of this paper.

It is worth noting, however, that the equilibrium obtained in this extension portrays a picture very close to the experimental results reported in Friedman and Oprea (2012), who estimate the median reaction lag at 0.6 second. This paper simulates the cooperation rate of the obtained equilibrium at every $t \in (0, 1)$ for $\bar{\theta} = 1.2$ seconds. The result is presented in the first graph below in conjunction with the graph of the median cooperation rate from the experiments (Friedman and Oprea, 2012, p. 347) for comparison purposes. The parameter values used in the simulation, namely $h = 1.4$ and $\ell = .4$, correspond to the "Easy" case in the experiment.[4] The theoretical prediction is very close to the experimental results for the "Mix-a" case as well ($h = 1.8$, $\ell = .4$), although it underestimates cooperation in the "Mix-b" and "Hard" cases.

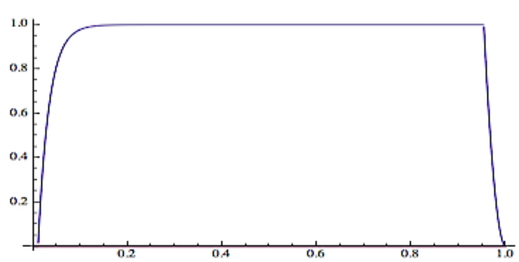

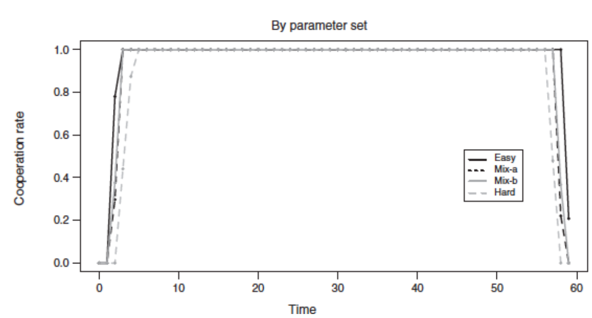

[4]The graph from the experiment is the median cooperation rate with initial states randomly assigned to all four possible action pairs, which biases the initial cooperation slightly forward.

3 A Canonical Preemption Game

The equilibrium illustrated in the previous section can be seen as a solution to a stopping game in which the players desire to preempt the opponent but prefer to do so later than sooner because the passage of time prior to stopping is beneficial. These are the core features of what may be referred to as "preemption games," examples of which include centipede games as well as the repeated prisoners' dilemma. This section provides a more complete analysis of a canonical, finite-horizon preemption game.

3.1 Model

The game is played between two players $i = 1, 2$, over a finite interval of time $[0, T]$ during which each player may "stop/defect" at most once. The game ends at the point of the first defection. Each player has a private type denoted by θ_i, $i = 1, 2$, which affects the players payoffs as detailed below. Sections 3.1 and 3.2 analyze a symmetric case in which the two players' types are drawn from the same distribution denoted by an atomless cumulative distribution function (cdf) F on $\Theta_1 = \Theta_2 = \Theta = (0, \bar{\theta})$, which is common knowledge. The associated density function f is Lipschitz-continuous and $f(\theta) > 0$ for all $\theta \in \Theta$. Note that the lower bound 0 of Θ is normalization. Section 3.3 extends the main result to asymmetric players.

The first defector is called the *winner/preemptor*, and the other player is called the *loser/preempted.* If the two players defect simultaneously, then they are said to be *evened out.* Denote player i's payoff when she is a winner, a loser, and evened out by $W(t, \theta_i, \theta_{-i})$, $L(t, \theta_i)$, and $E(t, \theta_i, \theta_{-i})$, respectively, where t is the time of defection and θ_i and θ_{-i} are the type of player i and that of the opponent. If the game ends without defection, the payoff of each player i is at most $E(T, \theta_i, \theta_{-i})$. Note that the payoff of player i as a loser, $L(t, \theta_i)$, does not depend on her opponent's type θ_{-i}. Although not essential for the result, this is the case for the two applications discussed in this paper (the prisoners' dilemma and the centipede game) and simplifies the exposition. For the case of the prisoners' dilemma, $W(t, \theta_i, \theta_{-i})$ and $L(t, \theta_i)$ are as specified in (2).

A few properties are assumed on payoffs, [u1]–[u4] below, which generalize the prisoners' dilemma game in the previous section. The first property simply states that preempting the opponent is preferred to being preempted at any given point in time, with the payoff from being evened out somewhere in between:

[u1] $W(t, \theta_i, \theta_{-i}) > E(t, \theta_i, \theta_{-i}) > L(t, \theta_i)$ for $t < T$.

The second property below just reflects the indexing convention that a lower θ value designates a desirable type or attribute of a player, which this paper refers to also as a more "composed" type, that weakly benefits her own payoff and weakly hurts the opponent's payoff:

[u2] W, E, and L are non-increasing in θ_i, non-decreasing in θ_{-i}, and Lipschitz-continuous.

Lipschitz continuity precludes a payoff from changing at infinite rates in its arguments and is assumed for technical convenience as well as realism.

The next property states that the passage of time before defection has the same positive effect on the preemptor and the preempted as long as there is enough time left for a full impact of preemption. The amount of time needed for a full impact, as represented by a non-decreasing function $\psi(\theta)$, is greater when the preempted is less composed. The second part of the property states that players prefer preempting the opponent at an appropriate point in time before the game ends, to continuing with mutual cooperation until the end:

[u3] A bounded, non-decreasing function $\psi : \Theta \to \mathbb{R}_+$ and a constant $a > 0$ exist such that

$$\begin{cases} \frac{\partial}{\partial t} W(t, \theta_i, \theta_{-i}) = a & \forall t < T - \psi(\theta_{-i}) \\ \frac{\partial}{\partial t} L(t, \theta_i) = a & \forall t < T - \psi(\theta_i) \end{cases} ; \quad \text{and} \quad W(T - \psi(\theta_{-i}), \theta_i, \theta_{-i}) > E(T, \theta_i, \theta_{-i}).$$

The particular feature that the players' payoffs increase at a *constant* rate with the passage of cooperation time, albeit valid in various applications, is not essential for the existence of the kind of equilibria analyzed in this paper, but is assumed because it simplifies exposition. In the prisoners' dilemma game in the previous section, $a = 1 - \ell$ and $\psi(\theta) = \theta$. In other applications $\psi(\theta)$ may be identically 0, as in the centipede game to be discussed later.

The last property, [u4] below, indicates that the less composed the player, the more there is for her to gain by preempting as opposed to being preempted:

[u4] $W(t, \theta_i, \theta_{-i}) - L(t, \theta_i)$ and $E(t, \theta_i, \theta_{-i}) - L(t, \theta_i)$ increase in θ_i, strictly if $t \leq T - \psi(\max\{\theta_i, \theta_{-i}\})$.

It turns out that the values of $W(t, \theta_i, \theta_{-i})$ for $t > T - \psi(\theta_{-i})$ and $L(t, \theta_i)$ for $t > T - \psi(\theta_i)$ do not matter for the class of equilibria to be characterized as long as [u1]–[u4] hold. Therefore, the underlying environment is represented by payoffs W, L, and D *defined for all* $t \in \mathbb{R}_+$, together with functions F and ψ, that satisfy [u1]–[u4] for all $T > 0$, and then equilibria are characterized for the case in which the duration of the game, T, is long enough.

Strategy A *(mixed) strategy* of player i is represented by a Lebesgue-measurable function $P_i : [0, T] \times \Theta \to [0, 1]$, where, for each $\theta_i \in \Theta$, $P_i(t|\theta_i)$ is a non-decreasing and right-continuous function that specifies the cumulative probability with which player i of type θ_i preempts her opponent at or before time t (conditional on not having been preempted first). It is a pure strategy if $P_i(\cdot|\theta_i)$ is a step function that jumps from 0 to 1 at a particular time point for every θ_i. In this case, P_i is represented by a measurable function $\tau_i : \Theta \to \mathbb{R}_+$ that specifies a type-contingent point in time at which player i preempts, that is,

$$P_i(t|\theta_i) = \begin{cases} 0 & \text{if } \; t < \tau_i(\theta_i) \\ 1 & \text{if } \; t \geq \tau_i(\theta_i) \end{cases}$$

where $\tau_i(\theta_i) > T$ is interpreted as player i of type θ_i not defecting until the game ends. Define

$$\zeta_i(t) := \int_\Theta P_i(t|\theta_i) dF$$

to be the *ex ante* cumulative probability that player i defects at or before t according to P_i provided that she does not get preempted first. Note that $\zeta_i(t)$ exists by Fubini's theorem

(Billingsley, 1995, p.234). Define the support of P_i to be that of ζ_i, i.e., $\operatorname{supp} P_i = \operatorname{supp} \zeta_i$, and $\operatorname{supp} \tau_i = \operatorname{supp} \zeta_i$ for a pure strategy.

Utility Given an opponent's strategy P_{-i}, the utility of player i of type θ_i from a non-decreasing and right-continuous function $p_i : [0,T] \to [0,1]$ that represents her own preempting strategy, is

$$\begin{aligned} u_i(p_i, P_{-i}|\theta_i) \;\; = \;\; & \int_\Theta \int_{t=0}^{T} W(t,\theta_i,\theta_{-i})\big(1 - P_{-i}(t|\theta_{-i})\big) dp_i(t) dF(\theta_{-i}) \\ & + \int_{t=0}^{T} L(t,\theta_i)\big(1 - p_i(t)\big) d\zeta_{-i}(t) \\ & + \int_\Theta \int_{t=0}^{T} E(t,\theta_i,\theta_{-i}) p_i(\{t\}) P_{-i}(\{t\}|\theta_{-i}) dt dF(\theta_{-i}) \end{aligned} \tag{4}$$

where $p_i(\{t\}) = p_i(t) - \lim_{s\uparrow t} p_i(s)$ is the point mass attached to t by p_i and similarly for $P_{-i}(\{t\}|\theta_{-i})$. Therefore, the first double integral of (4) is the expected payoff from the future contingency that player i preempts her opponent; the second (single) integral is the expected payoff from the future contingency that she gets preempted; and the last double integral is the expected payoff from the future contingency that the two players defect simultaneously. These integrals are well defined by Fubini's theorem.

Continuation subgame Given a strategy pair (P_1, P_2), conditional on no defection having taken place prior to $t \leq T$, define the continuation strategy $P_i^t : [t,T] \times \Theta \to [0,1]$ of player i as

$$P_i^t(s|\theta_i) := \frac{P_i(s|\theta_i) - P_i(t^-|\theta_i)}{1 - P_i(t^-|\theta_i)} \text{ if } P_i(t^-|\theta_i) < 1; \;\; \text{and } P_i^t(s|\theta_i) := 1 \text{ if } P_i(t^-|\theta_i) = 1 \tag{5}$$

where $P_i(t^-|\theta_i) = \lim_{t'\uparrow t} P_i(t'|\theta_i)$. Let

$$F_i(\theta_i|t) = \frac{\int_0^{\theta_i} 1 - P_i(t^-|\theta) dF(\theta)}{\int_\Theta 1 - P_i(t^-|\theta) dF(\theta)}, \tag{6}$$

or $F_i(\theta_i|t) = \theta_i/\bar{\theta}$ if the denominator of (6) is 0, denote the posterior belief about the type of player i conditional on no defection taking place prior to $t \leq T$. Then, given the opponent's continuation strategy P_{-i}^t, and the posterior belief about her type $F_{-i}(\theta_{-i}|t)$, the continuation utility of player i of type θ_i from a non-decreasing and right-continuous function $p_i^t : [t,T] \to [0,1]$, denoted by $u_i(p_i^t, P_{-i}^t|\theta_i)$, is defined in the same manner as above. Note that it coincides with $u_i(p_i, P_{-i}|\theta_i)$ if $t = 0$.

Definition 1 *A pair of strategies (P_1, P_2) is a <u>perfect Bayesian equilibrium</u> (PBE) if for all $t \in [0,T]$, $i = 1,2$, and all $\theta_i \in \Theta$,*

$$u_i(P_i^t(\cdot|\theta_i), P_{-i}^t \,|\, \theta_i) \;\geq\; u_i(p_i^t, P_{-i}^t \,|\, \theta_i)$$

for any non-decreasing and right-continuous function $p_i^t : [t,T] \to [0,1]$.

3.2 Equilibrium characterization

This section characterizes possible PBE's of the game. Here the following two kinds of pure-strategy PBE's turn out to be crucial:

Definition 2 *A pure strategy PBE* (τ_1, τ_2) *is* <u>*separating*</u> *if* τ_i *is continuous and strictly decreasing in* $\theta \in \Theta$ *for* $i = 1, 2$. *A pure strategy PBE* (τ_1, τ_2) *is* <u>*semi-pooling*</u> *if, for each* $i \in \{1, 2\}$, *there is a threshold type* $\widehat{\theta}_i < \bar{\theta}$ *such that* $\tau_i(\theta) = 0$ *for all* $\theta > \widehat{\theta}_i$ *and* τ_i *is continuous and strictly decreasing in* $\theta \leq \widehat{\theta}_i$.

In any separating PBE, each player waits gradually longer as her type is lower (more composed), before defecting to preempt her opponent. Any two types are separated in the sense that their defection times differ.

In any semi-pooling PBE, both players defect at the beginning of the game if their types are above a certain threshold but wait a certain amount of time if she is of the threshold type, and gradually longer if her type is lower, before defecting. Note that the trivial PBE in which both players of all types definitely defect at $t = 0$, which always exists, is a separating PBE in which the threshold type is 0 for both players. Other, nontrivial PBE's are now considered.

In any nontrivial semi-pooling PBE, as is shown later, the threshold type must be indifferent between defecting at $t = 0$ and after waiting a certain amount of time, and the threshold type may mix the two time points of defection without affecting the equilibrium condition because it constitutes a measure-zero event for the opponent. For expositional ease, therefore, an innocuous convention that, in such cases, the threshold type defects at the later point in time with certainty is adopted.

Given this convention, it is shown below that the two classes of PBE's described above are the only possible PBE's subject to a minor condition that any defection by a θ-type player takes place before $T - \psi(\theta)$, that is,

[A] $P_i(t|\theta_i) \to 1$ as $t \to T - \psi(\theta_i)$ for all $\theta_i \in \Theta$, $i = 1, 2$.

In light of [u3], the interpretation of this condition is that any preemption takes place in time for a maximum possible impact of preemption before the game ends.[5] This allows for the complete characterization of PBE's without specifying the details of partial impacts of preemptions which may vary widely depending on the circumstances. In addition, the core insight of this paper, namely that a more composed type is more willing to wait longer before preempting, must prevail even if [A] is not satisfied.[6] Consequently, PBE's that satisfy [A],[7] referred to as PBE*, are considered and characterized in the main result, Proposition 1, below. To facilitate presentation, the following definitions are introduced first:

[5]To see this more precisely, recall that all types preempt at some point before the game ends. If θ_i is supposed to defect at t and θ_{-i} at t or later, then [A] implies that $t \leq T - \psi(\theta_{-i})$ as well as $t \leq T - \psi(\theta_i)$, so that preemption by θ_i takes place at least $\psi(\theta_{-i})$ before the end of the game for any θ_{-i} that is a possible type of the opponent according to the posterior belief at time t. Note that $\psi(\theta) = \theta$ for the prisoners' dilemma game in Section 2.

[6]This paper does not formally analyze the extent to which this is true because it is beyond the purpose of this paper.

[7]This implies that both players definitely defect to preempt before the game ends, which must be the case in all PBE's because of [u1] and the second property of [u3].

Definition 3 *A pure-strategy PBE** (τ_1, τ_2) *is* ***symmetric*** *if* $\tau_1 = \tau_2$*, which is denoted by* τ*. Such a PBE** τ *is* ***time-invariant*** *if any time-shifted strategy* $\tilde{\tau} = \tau + c$ *where* $c \in \mathbb{R}$ *is also a PBE** *as long as* $\tilde{\tau}$ *satisfies [A]. Such time-shifted PBE***'s are said to be* ***equivalent****.*

Proposition 1 (a) *Every PBE** *is either a separating or semi-pooling PBE** (τ_1, τ_2) *such that* $\operatorname{supp}\tau_1 = \operatorname{supp}\tau_2$*.*

(b) *A nontrivial PBE** *may exist if*

$$\lim_{\theta \to 0} \Psi(\theta, \theta) > 0 \quad \textit{where} \quad \Psi(\theta_1, \theta_2) := \frac{a\, F(\theta_1)}{\big(W(0, \theta_2, \theta_1) - L(0, \theta_2)\big) f(\theta_1)}. \tag{7}$$

In particular, if (7) holds and T *is large enough, then a time-invariant symmetric separating PBE** *exists, and all symmetric separating PBE***'s are equivalent. In addition, for any* $\widehat{\theta} \in \Theta$*, there is a unique symmetric semi-pooling PBE** *in which* $\widehat{\theta}$ *is the common threshold type.*

(c) *If* $\Psi(\theta_1, \theta_2)$ *is Lipschitz-continuous in a neighborhood of* $(\bar{\theta}, \bar{\theta})$*, then all PBE***'s are symmetric.*

Part (a) of Proposition 1, which is proved in Appendix B, identifies two classes of pure-strategy equilibria as the only possible PBE's subject to [A]. It also suggests that the times at which the two players may defect must coincide because they both wish to preempt the other. Part (b) establishes a sufficient condition for such equilibria to exist. This is proved below by delineating the precise equilibrium conditions for the existence of a symmetric separating PBE* and showing that they are implied by (7), and then by doing the same for symmetric semi-pooling PBE*'s. In the process it is shown that all PBE*'s of the former kind are equivalent because they all solve the same autonomous differential equation, whereas there is a unique PBE* of the latter kind for every possible threshold type. In fact, these two types of PBE*'s form equilibria that change continuously in defection time. If the defection time of the former kind is brought forward so much that the upper-end type defects at the start of the game ($t = 0$), then it becomes an equilibrium of the latter kind.

Given the symmetry between the two players in the considered situation, attention is given to symmetric equilibria as they are both natural and easier to characterize. In fact, asymmetric equilibria are not viable under a mild technical condition (which the prisoners' dilemma game of Section 2 satisfies), as stated in part (c) of Proposition 1 and proved in Appendix B.

Symmetric separating PBE* Consider a symmetric separating PBE* represented by a continuous and strictly decreasing defection/preemption strategy τ, such that $\tau(\theta) > 0$ for all $\theta \in \Theta$ and its image is $Im(\tau) := (\lim_{\theta \uparrow \bar{\theta}} \tau(\theta), \lim_{\theta \downarrow 0} \tau(\theta)) \subset [0, T]$. Then its inverse function $\vartheta(t) = \tau^{-1}(t)$ is defined on $Im(\tau)$ and specifies the type of player defecting at time t. In the contingency that cooperation continues until $t \in Im(\tau)$, a $\vartheta(t)$-type player must find it optimal to defect at that point in time, and any lower type must find it optimal to wait. Identified below are the conditions under which this is indeed the case.

At $t \in Im(\tau)$, the posterior belief about the opponent's type, denoted by $F(\theta|t)$, is obtained through Bayes' rule by truncating types exceeding $\vartheta(t)$, that is, $F(\theta|t) = F(\theta)/F(\vartheta(t))$

for $\theta \leq \vartheta(t)$. Calculated as of t, the expected payoff of player i from defecting at $t+\Delta \in (t, T]$ unless preempted first, given the equilibrium strategy τ of the opponent, is

$$\int_{\theta_{-i}=0}^{\vartheta(t+\Delta)} W\big(t+\Delta, \theta_i, \theta_{-i}\big) dF(\theta_{-i}|t) + \int_{\vartheta(t+\Delta)}^{\vartheta(t)} L\big(\tau(\theta_{-i}), \theta_i\big) dF(\theta_{-i}|t). \tag{8}$$

The first integral depicts that if the type of player $-i$ is below $\vartheta(t+\Delta)$ so that she does not defect until $t+\Delta$, then player i preempts at $t+\Delta$ and receives the winner's payoff $W\big(t+\Delta, \theta_i, \theta_{-i}\big)$. The second integral captures that if the type of player $-i$ is above $\vartheta(t+\Delta)$ so that player $-i$ defects at $\tau(\theta_{-i}) < t+\Delta$, then player i receives the loser's payoff $L\big(\tau(\theta_{-i}), \theta_i\big)$.

Because $F(\theta_{-i}|t) = F(\theta_{-i})/F(\vartheta(t))$, (8) can be rewritten as

$$\frac{1}{F(\vartheta(t))}\left[\int_{\theta_{-i}=0}^{\vartheta(t+\Delta)} W\big(t+\Delta, \theta_i, \theta_{-i}\big) f(\theta_{-i})\, d\theta_{-i} + \int_{\vartheta(t+\Delta)}^{\vartheta(t)} L\big(\tau(\theta_{-i}), \theta_i\big) f(\theta_{-i})\, d\theta_{-i}\right]$$

and its first derivative with respect to Δ is

$$\begin{aligned}\frac{1}{F(\vartheta(t))}\Big[&\int_{\theta_{-i}=0}^{\vartheta(t+\Delta)} \frac{\partial W\big(t+\Delta, \theta_i, \theta_{-i}\big)}{\partial \Delta} f(\theta_{-i})\, d\theta_{-i} \\ &+\Big(W\big(t+\Delta, \theta_i, \vartheta(t+\Delta)\big) - L\big(t+\Delta, \theta_i\big)\Big) f(\vartheta(t+\Delta))\vartheta'(t+\Delta)\Big].\end{aligned} \tag{9}$$

With this evaluated at $\Delta = 0$ in light of [u3], the following is obtained:

$$\frac{1}{F(\vartheta(t))}\Big[a\, F(\vartheta(t)) + \Big(W\big(t, \theta_i, \vartheta(t)\big) - L\big(t, \theta_i\big)\Big) f(\vartheta(t))\vartheta'(t)\Big]. \tag{10}$$

Note that this is strictly decreasing in $\theta_i \leq \vartheta(t)$ by [u4]. Therefore, if (10) assumes 0 at $\theta_i = \vartheta(t)$, that is,

$$a\, F(\vartheta(t)) \;=\; -\Big(W\big(t, \vartheta(t), \vartheta(t)\big) - L\big(t, \vartheta(t)\big)\Big) f(\vartheta(t))\vartheta'(t) \tag{11}$$

$$\iff \qquad \vartheta'(t) \;=\; \frac{-\, a\, F(\vartheta(t))}{\Big(W\big(0, \vartheta(t), \vartheta(t)\big) - L\big(0, \vartheta(t)\big)\Big) f(\vartheta(t))}, \tag{12}$$

where $W(t, \vartheta(t), \vartheta(t)) - L\big(t, \vartheta(t)\big) = W(0, \vartheta(t), \vartheta(t)) - L(0, \vartheta(t))$ ensues from [u3], then either player of type below $\vartheta(t)$ strictly prefers to wait at t. Consequently, the function $\vartheta(\cdot)$ solving the differential equation (12) characterizes a PBE* as long as it is optimal for a $\vartheta(t)$-type player to defect at t. To verify this latter condition, it suffices to show that (9) is non-positive for all $\Delta \in (0, T-t]$ when $\theta_i = \vartheta(t)$. For this, first observe that (11) holds when t is replaced by $t+\Delta$ for all Δ such that $\vartheta(t+\Delta) > 0 = \inf \Theta$. The RHS of this equality increases when the second argument of W and L, $\vartheta(t+\Delta)$, is replaced with $\vartheta(t)$ by [u4]. This implies that (9) is non-positive when $\theta_i = \vartheta(t)$, as desired.

Consequently, if the inverse function of a solution ϑ to (12), $\tau = \vartheta^{-1}$, is defined on the entire domain Θ with the feature that $0 < \tau(\theta) < T - \psi(\theta)$ for all $\theta \in \Theta$, it indeed constitutes a symmetric separating PBE*.[8] This amounts to finding a solution ϑ to (12) such that its domain is a subset of $[0, T]$ and its image is equal to Θ.

Note that (12) is an autonomous differential equation. That is, it does not directly depend on the independent variable, t, so that its solution is time-invariant in the sense that it continues to be a solution when it is "shifted by a constant", i.e., when redefined as $\vartheta(t) = \vartheta(t+c)$ for some $c \in \mathbb{R}$. Therefore, if a solution to (12) maps onto Θ within a *bounded domain*, then it can be "shifted" to constitute a PBE* when T is large enough.

For each value of $\vartheta(t) \in (0, \bar{\theta})$, (12) specifies the instantaneous rate of change in ϑ at t, which is strictly negative. Standard results in ordinary differential equations (e.g., Theorems 3 and 4 on page 28 of Hurewicz (1958)) establish that there is a unique solution to (12) defined on a neighborhood of any $t_0 \in \mathbb{R}$ subject to an arbitrary initial condition $\vartheta(t_0) = \theta_0 \in (0, \bar{\theta})$. As the RHS of (12) is bounded away from 0 for all values of $\vartheta(t) \in (\theta_0, \bar{\theta})$, there is some finite $\underline{t} < t_0$ such that the value of a unique solution traced from t_0 converges to $\bar{\theta}$ as t tends to $\underline{t}$ from above. For $t > t_0$, there is some finite $\bar{t} > t_0$ by the same token such that the value of a unique solution traced from t_0 converges to 0 as t tends to $\bar{t}$ from below, provided that the rate at which the solution decreases is bounded away from 0 for all values of $\vartheta(t) \in (0, \theta_0)$. Because of continuity, this is the case if the RHS of (12) stays bounded away from 0 as $\vartheta(t)$ tends to 0, that is, if (7) holds. Therefore, it is established that if (7) holds, subject to any initial condition $\vartheta(t_0) = \theta_0 \in (0, \bar{\theta})$ there is a unique solution to (12) that maps a bounded time interval onto Θ. Together with the time-invariance of the solution to an autonomous differential equation, this proves the first half of Proposition 1 (b). That is, if T is large enough then the set of symmetric separating PBE*'s constitutes a continuum of equivalent PBE*'s differing in only one respect: Players wait longer before preempting in one equilibrium than in the other by a constant. In the case of the prisoners' dilemma game, because a longer wait prolongs the time for (C, C) and curtails that for (D, D) by the same amount, the equilibrium payoff is higher in the PBE* with a longer delay before preemption.

Symmetric semi-pooling PBE* Recall that a symmetric semi-pooling PBE* is represented by a common pure strategy τ such that $\tau(\theta) = 0$ for all $\theta > \widehat{\theta}$ for some (common) *threshold type* $\widehat{\theta} < \bar{\theta}$ and τ is continuous and strictly decreasing for $\theta \leq \widehat{\theta}$. Let $\widehat{t} = \tau(\widehat{\theta})$ be referred to as the *threshold time*. Here start with the following implication from the discussion above on symmetric separating PBE*:

[B] For any potential threshold type $\widehat{\theta}$ and time $\widehat{t} > 0$, provided that (7) holds and T is sufficiently large, there exists a unique symmetric separating PBE* of the game when the type space is restricted to $\widehat{\Theta} = (0, \widehat{\theta})$, denoted by $\widehat{\tau}$, with $\lim_{\theta \to \widehat{\theta}} \widehat{\tau}(\theta) = \widehat{t}$.

Therefore, a semi-pooling PBE* of the original game is obtained if for some $\widehat{\theta} \in \Theta$ and $\widehat{t} > 0$, it can be shown that $\widehat{\tau}$ identified in [B] is optimal after extending it to Θ by setting $\widehat{\tau}(\theta) = 0$ for $\theta > \widehat{\theta}$ and $\widehat{\tau}(\widehat{\theta}) = \widehat{t}$. To find the values of $\widehat{\theta}$ and $\widehat{t}$ for which this is indeed the case conditional on her opponent behaving according to the extended $\widehat{\tau}$, denoted by $\widehat{\tau}^*$, calculate

[8] To be precise, being the inverse function of a symmetric separating PBE* which is strictly monotone, $\vartheta(t)$ is guaranteed to have a well-defined derivative for almost all t and thus, (12) is valid for almost all t. Then, as $\vartheta(\cdot)$ is obtained by integrating the RHS of (12) which is a continuous function, it follows that $\vartheta'(t)$ is well-defined for all t and (12) is valid for all t.

the expected payoff of player i of type θ_i from defecting at $t = 0$ and that from defecting at $\widehat{t}$, respectively, as

$$\int_{\theta_{-i}=\widehat{\theta}}^{\bar{\theta}} E(0, \theta_i, \theta_{-i}) dF + \int_{\theta_{-i}=0}^{\widehat{\theta}} W(0, \theta_i, \theta_{-i}) dF, \tag{13}$$

and

$$\int_{\theta_{-i}=\widehat{\theta}}^{\bar{\theta}} L(0, \theta_i) dF + \int_{\theta_{-i}=0}^{\widehat{\theta}} W(\widehat{t}, \theta_i, \theta_{-i}) dF. \tag{14}$$

Because $W(\widehat{t}, \theta_i, \theta_{-i}) - W(0, \theta_i, \theta_{-i}) = a\widehat{t}$ by [u3], there is a unique value of $\widehat{t}$ for which (13) and (14) are equal when $\theta_i = \widehat{\theta}$:

$$\widehat{t} \;=\; \frac{1}{aF(\widehat{\theta})} \int_{\theta_{-i}=\widehat{\theta}}^{\bar{\theta}} E(0, \widehat{\theta}, \theta_{-i}) - L(0, \widehat{\theta})\, dF \;>\; 0 \tag{15}$$

where the inequality follows from [u1]. As $E(0, \theta_i, \theta_{-i}) - L(0, \theta_i)$ is strictly increasing in θ_i by [u4], given (15),

(i) player i strictly prefers $t = 0$ to $\widehat{t}$ for her defection if her type exceeds $\widehat{\theta}$,
(ii) the converse is true if her type is below $\widehat{\theta}$, and
(iii) she is indifferent between defecting at $t = 0$ and at $\widehat{t}$ if her type is $\widehat{\theta}$.

In addition, note that

(iv) player i of all types strictly prefers $\widehat{t}$ to any $t \in (0, \widehat{t})$ for her defection

because delaying defection until $\widehat{t}$ increases her utility as the winner by [u3], with no risk of being preempted in the meantime.

Now, consider the contingency that player i of an arbitrary type reaches $\widehat{t}$ without being preempted. At that point, if her type is $\theta_i < \widehat{\theta}$ then she finds it optimal to defect at $\widehat{\tau}(\theta_i) > \widehat{t}$ as $\widehat{\tau}$ is a symmetric separating PBE* in the continuation subgame. Recall that this is shown by proving that her expected utility always decreases as she delays defection beyond $\widehat{\tau}(\theta_i)$ (that is, (9) is non-positive for all $\Delta > 0$ for $\theta_i = \vartheta(t)$). This logic extends straightforwardly to establish that player i finds it optimal to defect at $\widehat{t}$ if her type is $\theta_i \geq \widehat{\theta}$. Together with (i)–(iv) above, this verifies that it is optimal for player i to defect as $\widehat{\tau}^*$ stipulates in the original game, conditional on her opponent doing so as well. Therefore, it is established that if (7) holds, then for any $\widehat{\theta} \in \Theta$ there is a unique symmetric semi-pooling PBE* if T is sufficiently large, in which $\widehat{\theta}$ is the threshold type for both players. This proves the second part of Proposition 1 (b). Note that the threshold time, $\tau(\widehat{\theta}) = \widehat{t}$ in (15), is strictly positive for every $\widehat{\theta} > 0$. This is because, given that a positive mass of types defect at $t = 0$, all other types must anticipate a substantial period of mutual cooperation before defection if not preempted at $t = 0$, so as for them to resist defecting at $t = 0$.

3.3 Discussion and extension

This section discusses when the key condition (7) is satisfied, presents a continuous-time version of the centipede game as another application of the theory, and extends the main result to asymmetric players.

When is (7) satisfied? In the prisoners' dilemma game illustrated in Section 2, $W(0,\theta,\theta) - L(0,\theta) = \frac{\theta}{2}h$ so that its ratio relative to θ stays bounded away from 0 as $\theta \to 0$, and so does the ratio of $F(\theta)/f(\theta) = \theta$ relative to θ, thereby satisfying (7). Generalizing this observation provides one class of environments guaranteeing (7):

$$\lim_{\theta \to 0} \frac{\theta}{W(0,\theta,\theta) - L(0,\theta)} > 0 \quad \text{and} \quad \lim_{\theta \to 0} \frac{F(\theta)}{\theta f(\theta)} > 0. \tag{16}$$

The latter inequality means that F does not vanish to an infinite order at $\theta = 0$, which is satisfied by a wide class of distribution functions including those such that $F(\theta) = \theta^n$ in a neighborhood of $\theta = 0$ for any $0 < n < \infty$.[9]

At the same time, for (7) to hold it is necessary that $W(0,\theta,\theta) - L(0,\theta) \to 0$ as $\theta \to 0$ because $F(\theta)/f(\theta) \to 0$ as $\theta \to 0$. This condition means that the benefit of preempting the opponent relative to being preempted, vanishes as both players become "maximally composed". Applied to the prisoners' dilemma game illustrated earlier, this means that the player's reaction lag may be arbitrarily short. If this condition fails (recall that the lower bound 0 of Θ in the canonical model is normalization), so that the limit in (7) is 0, then, as the defection phase nears the end, for the expected gain from immediately preempting the opponent (which stays bounded away from 0 in this case) to be balanced by the benefit of waiting in order to satisfy the incentive compatibility, the rate at which the opponent defects in the next instant needs to drop so fast that the defection phase may not terminate within a finite time.

In the case that the lower bound of the player's reaction lag is strictly greater than 0, it may still hold that $W(0,\theta,\theta) - L(0,\theta) \to 0$ as $\theta \to 0$ for some other reasons. For example, players may have varying degrees of altruism such that they do not find preempting beneficial at the most altruistic limit, in a manner akin to the players in the centipede game illustrated below.

Application to the centipede game Consider two players playing a version of the centipede game during a unit interval $[0,1]$ of time as follows: The players may choose to "stop" at any time, and the game ends at the point of the first stop by either player. The payoff for player i is $W(t,\theta_i,\theta_{-i}) = t + \theta_i$ if she solely stops at $t \in [0,1]$; $L(t,\theta_i) = t$ if her opponent solely stops at t; and $E(t,\theta_i,\theta_{-i}) = t + \theta_i/2$ if both stop at t simultaneously or the game ends without a stop and $t = T$. A possible interpretation of θ_i is the degree of altruism, with $\theta_i = 0$ designating the most altruistic type.[10] Then the analysis of the previous section applies because [A] is satisfied trivially as $\psi(\theta) \equiv 0$. In particular, there is a symmetric separating PBE in which the players stop some time during a small interval at the end of the game, gradually later as they are more altruistic.

In discrete settings, cooperation in a centipede game until close to the end can be justified by arguments based on the existence of commitment types or on boundedly complex strategy spaces, as shown for the prisoners' dilemma game by Kreps, *et al.* (1982) and Neyman (1985), respectively (cf. Section 4). In addition, Jehiel (2005) applies his *analogy-based expectation equilibrium* to the centipede game and shows that when the players identify "similar" nodes

[9] An example that fails the condition is $F(\theta) = e^{-1/\theta}$ for $\theta > 0$ and $F(0) = 0$.

[10] Note that redefining $W(t,\theta_i,\theta_{-i}) = t$ and $L(t,\theta_i) = t - \theta_i$ delivers the same result.

of the opponent, learn only the average behavior at those nodes and best respond to it, then they may cooperate until the last node at which the relevant player stops. The current analysis provides an alternative, fully rational explanation.

Extension to asymmetric players The core results in the previous section extend straightforwardly to the cases in which $\bar{\theta}$, F, f, W, L, D, a, and ψ are player-specific, which is indicated by subscripts $i = 1, 2$, as stated below and proved in Appendix B. For this, the formula in (7) needs to be modified for each player $i = 1, 2$, as

$$\Psi_i(\theta_i, \theta_{-i}) \ := \ \frac{a_{-i} \, F_i(\theta_i)}{\big(W_{-i}(0, \theta_{-i}, \theta_i) - L_{-i}(0, \theta_{-i})\big) \, f_i(\theta_i)}. \tag{17}$$

Proposition 2 *Suppose that the following hold for $i = 1, 2$:*
(i) $\lim_{\theta \to 0} \Psi_i(\theta, \theta) > 0$*, and*
(ii) $\Psi_i(\theta_i, \theta_{-i})$ *is Lipschitz-continuous in a small neighborhood of each point* $(\theta_i, 0)$ *and* $(0, \theta_{-i})$ *with* $\theta_i > 0$ *and* $\theta_{-i} > 0$*.*
There exists a separating PBE* (τ_1, τ_2) if T is sufficiently large, where $\tau_i(\cdot)$ is continuous and strictly decreasing in θ_i for $i = 1, 2$, and $\operatorname{supp} \tau_1 = \operatorname{supp} \tau_2$. Moreover, it is time-invariant, that is, $(\tau_1 + c, \tau_2 + c)$ continues to be a PBE* for any $c \in \mathbb{R}$ as long as [A] is satisfied. If $\Psi_i(\theta_i, \theta_{-i})$ is Lipschitz-continuous in a small neighborhood of $(\bar{\theta}_i, \bar{\theta}_{-i})$ for $i = 1, 2$, there is a unique separating PBE* modulo time-invariance.

This result basically states that a separating PBE* exists if the value of (17) does not change at infinite rates as (θ_1, θ_2) approaches either axis. These conditions are satisfied, for instance, by the prisoners' dilemma game in Section 2 when $F_i(\theta_i) = \theta_i^n$ for $i = 1, 2$ and $n > 0$. Moreover, as they are sufficient but not necessary conditions, separating PBE*'s may exist more widely.

Proposition 2 provides a basis for extending semi-pooling PBE*'s to asymmetric players as well. For an arbitrary pair of potential threshold types $(\widehat{\theta}_1, \widehat{\theta}_2) \in \Theta_1 \times \Theta_2$ and a potential threshold time $\widehat{t}$, Proposition 2 asserts the existence of a separating PBE* $(\widehat{\tau}_1, \widehat{\tau}_2)$ for the restricted type spaces $(0, \widehat{\theta}_1)$ and $(0, \widehat{\theta}_2)$. Let $(\widehat{\tau}_1^*, \widehat{\tau}_2^*)$ denote the extension of $(\widehat{\tau}_1, \widehat{\tau}_2)$ to the entire type spaces Θ_1 and Θ_2 by stipulating that player i of type $\theta_i > \widehat{\theta}_i$ definitely defects at $t = 0$. If both players of the threshold type $\widehat{\theta}_i$, $i = 1, 2$ are indifferent between defecting at $t = 0$ and defecting at $\widehat{t}$ given their opponent's strategy $\widehat{\tau}_{-i}^*$, then $(\widehat{\tau}_1^*, \widehat{\tau}_2^*)$ constitutes a semi-pooling PBE* of the original game. Note that given the potential threshold time $\widehat{t}$, the particular type of each player who is indifferent between defecting at $t = 0$ and at $\widehat{t}$ is a continuous and increasing function of the opponent's threshold type. A fixed point of these two functions, if it exists in $\Theta_1 \times \Theta_2$, constitutes a semi-pooling PBE* with $\widehat{t}$ as the common threshold time.

4 Concluding remarks

This paper provides a new theoretical justification for the observed cooperative behavior that breaks down shortly before the end of finite-horizon prisoners' dilemma games. The core insight is that when the players are heterogenous in their payoff margin of preempting relative to being preempted, it is self-sustaining for both players to delay defection longer for the sake of prolonging the cooperation phase, when their payoff margin is lower. The differential payoff margin may stem from physical causes such as different reaction lags, or from psychological causes such as different degrees of altruism. This theory contributes to the literature on dynamic cooperation between agents of conflicting interests in various contexts, as explained below.

Several theories exist that rationalize substantial cooperation in finitely repeated prisoners' dilemma games in discrete settings. Radner (1980, 1986) points out that, for any integer k, a "trigger strategy" of switching to D in period k gets arbitrarily close to being the best response to itself when the number of repetitions is large and thus that cooperation is sustained in ϵ-equilibria. Kreps, *et al.* (1982) demonstrate that if a player is committed to a "tit-for-tat" strategy with a small probability, a normal-type player by reputation motives, cooperates to be perceived as a commitment type, thereby inducing cooperation from the opponent.[11] This result is embedded in the "incomplete information Folk theorem" of Fudenberg and Maskin (1986). Neyman (1985, 1999) shows that cooperation may arise if there are bounds on the complexity of strategies that may be used or there is some small departure from common knowledge about the number of repetitions. The present paper contributes to this literature by proposing an alternative approach based on fully rational players, which is particularly relevant in continuous-time environments.

For a continuous-time analysis, the aforementioned technical issue of the indeterminate timing of immediate responses needs to be taken care of first. Simon and Stinchcombe (1989) formalize an analytic framework allowing consecutive moves at the same point in time (therefore immediate reactions with no time lag), which supports, *inter alia*, cooperation throughout the whole duration of the prisoners' dilemma game as an equilibrium. Stipulating an "inertia" condition that actions may not be changed for a little while (which can be arbitrarily short) at any point in time, Bergin and MacLeod (1993) look for the limit of ϵ-equilibria to obtain a Folk theorem-type result supporting, in particular, the same outcome. The present paper's approach recovers this equilibrium outcome predicted by these studies as the limiting equilibrium when the reaction lag vanishes. By contrast, the application of the standard, commonly known reaction lag leads to discontinuity at the limit because it prescribes "no cooperation at all" as the unique equilibrium, however short the reaction lag is.[12]

[11] Note that their results rely on there being the "right kind" of behavior that the player may be committed to. For instance, if either player may be committed to "cooperate in every period regardless of history," then it is straightforward to see that the logic of backward induction dictates that a rational player must definitely defect in every period of any finitely repeated prisoners' dilemma game.

Ambrus and Pathak (2011) also show that early cooperation by selfish players arises if some players are motivated differently, by reciprocity in particular, in finite-horizon public good contribution games.

[12] Suppose that the reaction lag is known to be $\theta > 0$ for both players. Then, both players must defect (to preempt) at least θ before the game ends, i.e., by $T-\theta$, in any equilibrium. Fix an equilibrium and let $\hat{t} \leq T-\theta$ denote the earliest time by which both players defect (to preempt) for certain in that equilibrium. If $\hat{t} > 0$, note that neither player would defect at $\hat{t}$ with a positive probability, because if a player, say i, did, then player j would defect for certain *before* $\hat{t}$, rendering $\hat{t}$ a suboptimal time to defect for player i. Thus, both players must mix defecting continuously over a time interval leading up to $\hat{t}$. Then, as of $\hat{t}-\epsilon$ for some arbitrarily small $\epsilon > 0$, player i must assess that there is a later point in time $\hat{t}' \in (\hat{t}-\epsilon, \hat{t})$ arbitrarily close to $\hat{t}$ which is no worse time to defect than $\hat{t}-\epsilon$. But, this is impossible because player i would win for sure

In addition to abstract games like the centipede game discussed earlier, preemption games have been studied in various contexts such as market entry, patent races, and financial investment decisions. Earlier studies such as Reinganum (1981), Fudenberg and Tirole (1985) and Weeds (2002), examine rival firms' strategic timing decisions of entry in complete-information environments in which preemption becomes a dominant strategy only after the industry matures enough, and find that the optimal timing is shaped by how the payoff margin of preempting (relative to being preempted) changes over time. In the canonical preemption game analyzed here, on the other hand, the payoff margin is constant over time, and therefore both players "preempt" at the beginning of the game under complete information. However, it is shown that cooperation can be sustained if the magnitude of the payoff margin is private information.

Abreu and Brunnermeier (2003) explain financial bubbles as an equilibrium of a preemption game among traders who are privately informed about when the bubble will burst, ending the game. Hopenhayn and Squintani (2011) examine a patent race in which the competing firms' knowledge advances stochastically over time, and show that the race lasts longer when the firm's level of knowledge is private information. The present paper's approach shares the feature that players have private information, but a key difference is that the exact time at which the game ends is common knowledge and finite, which is generally perceived as making the model more susceptible to unraveling by backward induction.

Appendix A

Appendix A presents a richer analysis of the continuous-time prisoners' dilemma game introduced in Section 2, in which there are multiple changes in the players' actions, and rationalizes a PBE in which the players coordinate to cooperate through asynchronous moves in the early stages of the game, followed by defections near the end of the game in the manner explained in Section 2.

The game under consideration is one in which a player may switch between C and D repeatedly, with a reaction lag applied after each move of her own as well as that of her opponent (see below). If the reaction lag of a player, say i, is shorter than that of player j, then player i can block any move of player j by switching her own actions back and forth if player j's reaction lag is reset at each move of player i. This, however, is an artifact of applying a simple notion of a reaction lag devised essentially for alternating move games (such as bargaining games) to richer environments where players with different reaction lags may make consecutive moves.

Recall that reaction lags are devised to pin down the time of the earliest response to a previous move in continuous-time models, so that the strategy of an "immediate reaction" is well defined. Based on this fundamental idea, the question of how reaction lags apply depending on various possible histories (of how the game has been played) needs to be addressed to

by defecting at $\hat{t}-\epsilon$ but she would lose almost surely if waited until t'. Therefore, $\hat{t}=0$ must hold.

describe a precise extensive-form game under consideration. Here an extensive-form game is described in this spirit, and an equilibrium of the game showing the aforementioned features is characterized.[13]

A. Description of an extensive-form game

Each player $i \in \{1, 2\}$ is privately informed of her own reaction lag θ_i, which is an independent draw from a uniform distribution $F(\theta) = \theta/\bar{\theta}$ on $(0, \bar{\theta})$, where $\bar{\theta} > 0$ is small. At each $t > 0$, the "default action" for player i refers to the action that she has been taking in an interval immediately before t, i.e., during $(t - \epsilon, t)$ for some $\epsilon > 0$. Assume that the initial default action at $t = 0$ is given exogenously (which conforms to the experiments in Friedman and Oprea (2012)), in particular, to be D for both players. At any time during $[0, 1]$, each player may switch from the default action to the other action subject to the following rules:

(i) Upon observing the opponent's switch in actions at $t < 1$, player i determines whether to produce an "immediate response" of switching actions herself as well. If she decides to immediately respond, then her action gets switched at $t + \theta_i$ where $\theta_i > 0$ is her reaction lag/type. In this case, she is said to be "committed" from t until her action gets switched (at $t + \theta_i$).

(ii) Upon observing the opponent's switch in actions at $t < 1$, if player i decides against responding immediately, then her action must remain the same during $[t, t + \theta_i]$. That is, she may change her action only after $t + \theta_i$. In this case, she is said to be "latent" during $[t, t + \theta_i]$. Here the length of latency is set to θ_i for notational ease in the sense that the equilibrium to be presented is valid for longer lengths as well (unless too long).

(iii) Upon switching her action at t, if player i decides to switch her action again immediately, then the same issue arises such that the earliest point in time she can do so is not pinned down. As in (i), therefore, it is modeled in this case that her action gets switched after a time lag, say $\eta_i > 0$. Similarly, if player i decides against switching her actions again immediately after switching them at t, then she remains latent until $t + \eta_i$ in the manner explained in (ii). Because the value of η_i does not affect the result, it is assumed that $\eta_i = \theta_i$ in the sequel.

(iv) While player i is committed in the sense of having decided to immediately respond to the opponent's switch in actions at t (as described in (i)), if the opponent switches actions again, say at $t' \in (t, t + \theta_i)$, then the commitment of player i to switch actions gets implemented at $t + \theta_i$. That is, the decision is not affected by the opponent's switch while she is committed. In addition, when her committed switch gets implemented at $t + \theta_i$, player i determines at that point whether to switch her actions again immediately in the manner explained in (iii). Therefore, the decisions of player i described in (i) and (ii) should be understood to pertain to the case in which she is not committed at t.

[13] There are other plausible ways to treat reaction lags that support essentially the same equilibrium. For example, modelling that the immediate reaction of player i to a move at t gets implemented at some random time during $(t, t + \theta_i)$ also works with appropriate supplementary modifications.

(v) While player i is latent from t in the sense explained in (ii), if the opponent switches actions again, say at $t' \in (t, t+\theta_i]$, then player i is modeled to decide at that point whether to respond immediately (by switching actions herself) to her opponent's action switch at t' in the manner described in (i) and (ii). Alternatively, player i can be modeled to remain latent until $t+\theta_i$ in this contingency. The equilibrium to be presented is also valid in the latter model.

To recap, if there is an action switch by either player, then both players determine whether to switch their actions as an "immediate response." If either player decides on such a response, then that player's action gets switched after her reaction lag. If either player decides against it, then that player's action is locked until her reaction lag expires, unless her opponent switches actions in between, at which point she needs to decide whether to respond immediately or not.

Explained up to now is how a player's decision to switch actions as an immediate reaction to a preceding switch in actions can be modeled. Now the question of how other switches of actions can be modeled and how a player's strategy can be represented are addressed.

Informally, at the beginning of the game, each player has a "plan" concerning when to switch her action (from the initial default action, D, to C) conditional on her opponent having not switched by then. Such a plan is represented by a non-decreasing and right-continuous function defined on $[0, 1]$ specifying the cumulative probability that she switches (not preceded by the opponent) by each future point in time.[14] Note that a plan is contingent on the type of the player. When a switch takes place by either player, say at t, both players discard their previous plans that become obsolete from the switch and adopt new plans drawn out for an updated history including the latest switch(es) at t. Such a new plan reflects the player's decision on whether to immediately respond to the latest switch, as described in (i)–(iii) above.

A *strategy* of a player consists of a family of (type-contingent) plans, one for every possible history of the game. This family of plans should abide by constraints implied by (i)–(v) and satisfy an inter-temporal consistency condition *á la* Bayes' rule.

Formally, a "history" at $t \in [0, 1]$, denoted by h^t, is a full record of all moves made prior to t. Let $|h^t| < t$ denote the time of the latest move in h^t. Consider player i of type θ_i, called an "agent θ_i". If $|h^t| + \theta_i \geq t$, then she is either committed to an immediate reaction to a previous move or latent "at h^t," i.e., as of t with history h^t. She is said to be *active* at h^t otherwise. A history $h^{t'}$ which extends h^t is a *simple extension* of h^t if $|h^{t'}| = |h^t|$. A history h^t is θ_i-*compatible* if the switches in h^t are consistent with the reaction lag θ_i.

An agent θ_i's plan at a θ_i-compatible history h^t, denoted by $P_i(\cdot|h^t, \theta_i)$, is a non-decreasing and right-continuous function from $[t, 1]$ to $[0, 1]$ such that:

- if she is committed at h^t, then $P_i(s|h^t, \theta_i)$ jumps from 0 to 1 at some $s \geq t$ where $s - \theta_i$ is the time of a previous switch, after which player i's action did not change.

[14] This is a standard way of representing strategies in continuous-time that has been used to analyze preemption games and war of attrition games. A potential drawback of this approach is possible loss of information in passing to the continuous-time limit of discrete-time strategies as exemplified in the so-called "grab the dollar" game (see Fudenberg and Tirole, 1985). This issue does not arise in the current analysis.

- if she is latent at h^t, then $P_i(s|h^t,\theta_i)$ assumes 0 and is continuous at $s = |h^t| + \theta_i \geq t$.

In addition, inter-temporal consistency requires that

- if she is committed at h^t, so that $P_i(\cdot|h^t,\theta_i)$ jumps to 1 at some $s \geq t$, then $P_i(\cdot|h^{t'},\theta_i)$ also jumps to 1 at s for any θ_i-compatible extension $h^{t'}$ of h^t with $t' \leq s$.[15]

- if she is not committed at h^t, for any simple extension $h^{t'}$ of h^t, $P_i(\cdot|h^{t'},\theta_i)$ is obtained from $P_i(\cdot|h^t,\theta_i)$ by Bayes' rule:

$$P_i(s|h^{t'},\theta_i) = \frac{P_i(s|h^t,\theta_i) - \lim_{s'\uparrow t'} P_i(s'|h^t,\theta_i)}{1 - \lim_{s'\uparrow t'} P_i(s'|h^t,\theta_i)}$$

  for all $s \in [t', 1]$ if the denominator is non-zero, and $P_i(s|h^{t'},\theta_i) = 1$ otherwise.

A pair of strategies of the two players is a *perfect Bayesian equilibrium* if, after every possible history, *i*) the belief profile of the types of these two players is consistent with Bayes' rule in the manner extending (6), and *ii*) the continuation part of the strategy pair satisfies the usual mutual best-response property given the belief profile.

B. The description of an equilibrium

A PBE is described below as a sequence of several phases, and involves three key points in time denoted by $0 < t^A < t^C < t^D < 1$.

1) *War-of-attrition (WoA) phase:* The initial plan of each player, regardless of her type, is $P_i(t|h^0,\theta_i) = 1 - e^{-qt}$ for $t \leq t^A$, $i = 1,2$, where $q = \frac{1-\ell}{\bar{\theta}h}$ and $t^A < 1 - \bar{\theta}(h - \ell + \frac{1+h}{1-\ell})$ is a fixed point in time. That is, starting the game with (D,D), both players switch to C with a common "flow rate" q so long as the time has not passed t^A and the other has not switched to C already. Here this initial part of the game until one player switches to C at some point before t^A is referred to as the "war-of-attrition" (WoA) phase.

2) *Transition phase to cooperation:* If one player, say i, switches to C in the WoA phase, say at $t \leq t^A$, then player j follows suit $\bar{\theta}$ later, that is, switches to C at $t + \bar{\theta}$ (and i adheres to C until then).[16] If neither player switches to C by t^A, then both players switch to C simultaneously[17] at $t^C = t^A + \frac{\bar{\theta}}{1-\ell} < 1 - \left(\frac{h}{2(1-\ell)} + h - \ell\right)\bar{\theta}$.

3) *Cooperation phase:* After both players switch to C as explained in 2), they continue with (C,C) until $t^D = 1 - \frac{h}{1-\ell}\bar{\theta}$. This part of the game is referred to as the "cooperation" phase.

[15] Inter-temporal consistency does not apply to θ_i-compatible extension $h^{t'}$ of h^t with $t' > s$, because the implementation of the committed switch at s must have voided the plan $P_i(\cdot|h^t,\theta_i)$.

[16] That player j follows suit at $t + \bar{\theta}$ regardless of her type is consistent with the idea that she benefits by following suit as late as possible without provoking punishment from player i. In addition, this prevents any signalling of player j's type, so that the subsequent defection phase is based on the initial information asymmetry and is independent of the exact equilibrium path realized.

[17] This relies on perfect coordination, unlike other parts of the equilibrium. Alternatively, an asymmetric equilibrium is also possible. At a certain point in time between t^A and t^D one player i switches to C, and player j follows suit $\bar{\theta}$ later. Then, another $\bar{\theta}$ later i switches back to D and returns to C yet another $\bar{\theta}$ later, ensuring the timing such that players are indifferent between switching and not at t^A.

4) *Defection phase:* Once t^D is reached through the cooperation phase, each player of type θ defects at $\tau(\theta) = 1 - \frac{h}{1-\ell}\theta \in (t^D, 1)$, conditional on not being preempted. Once one player, say i, defects at $t \in (t^D, 1)$, player $j \neq i$ responds immediately such that her action switches to D at $t + \theta_j$, after which both players keep to D until the end of the game. (As shown later both will switch to D before the game ends.) This part of the game is referred to as the "defection" phase.

5) *Off-equilibrium:* If the history at any t departs from the equilibrium-path described above, then both players switch to D as soon as they can, unless they are already playing D, and keep to D until the end of the game. Here a caveat is in order: In the off-equilibrium contingency that one player, say i, switches to C at t in the WoA phase, but the other player does not follow suit at $t + \bar{\theta}$, the earliest point in time after this deviation at which player i can switch back to D is not pinned down. To address this problem, the following is stipulated in this contingency:

   5-i) Player i switches back to D at $t + \bar{\theta} + \epsilon$ for a small $\epsilon > 0$, and then both players switch to C at some later point before t^D, say t'', followed by the cooperation and defection phases of the equilibrium-path. Provided that t'' is sufficiently late, player $j \neq i$ prefers following the equilibrium-path to invoking this continuation game by failing to follow suit at $t + \bar{\theta}$.

   5-ii) If player i switches back to D either before or after $t + \bar{\theta} + \epsilon$, or player $j \neq i$ switches after $t + \bar{\theta}$, then both players switch to D as soon as they can, unless they are already playing D, and keep to D until the end of the game.

C. The verification of the equilibrium

Now the equilibrium conditions of the aforementioned symmetric strategy profile are verified to determine that it constitutes a PBE. Given that D is a strictly dominant strategy of the static game, it is straightforward to verify that in the off-the-equilibrium path, it is optimal for each player to keep or switch to D as soon as possible if the other does the same (apart from the aforementioned exception case of 5-i), which is addressed separately). Now the optimality of the strategy on the equilibrium-path is verified backward from the defection phase. Recall that $F(\theta) = \theta/\bar{\theta}$.

Defection phase: In the defection phase, once one player defects, again from the fact that D is a strictly dominant strategy of the static game, it is optimal for the other player to follow suit as soon as possible and then for both to keep to D until the end, as stipulated by the equilibrium above.

Now consider the time point $t = \tau(\theta)$ for some $\theta \in (0, \bar{\theta})$ with no defection having taken place from (C, C) by either player until then. As τ is an inverse function of the solution to (12) for the current case, the arguments in Section 3 are applied to verify it to be optimal for either player to defect at $t = \tau(\theta)$ if her type is θ or higher and wait otherwise.

Cooperation phase: Consider any point in time $t < t^D$ in the cooperation phase after both players switch to C. It is straightforward to verify that it is suboptimal for either player to switch to D before t^D, say at $t' \in [t, t^D)$, because she is better off by switching to D at t^D

instead, which would prolong the duration of (C,C) by $t^D - t'$ and curtail the duration of (D,D) by the same length later.

Transition phase to cooperation: Moving backward one stage, consider player j after player i has switches to C at $t < t^A$ in the WoA phase. It is clearly suboptimal for j to switch to C before $t+\bar{\theta}$ because then i switches back to D as an immediate response and keeps to it until the game ends as the off-equilibrium strategy 5) prescribes. Here not switching to C at $t+\bar{\theta}$ would lead to an outcome described in 5-i), which is a continuation equilibrium based on 5-ii) but is worse for player j than following the equilibrium by switching to C at $t+\bar{\theta}$. Therefore, it is optimal for j to switch to C at $t+\bar{\theta}$. Given this, it is optimal for player i to keep to C until $t+\bar{\theta}$ because her switching back to D before $t+\bar{\theta}$ is worse due to the off-equilibrium strategy described in 5).

Consider the other case in which neither player switches to C until t^A. Then, by switching to C simultaneously at t^C, the players enter the cooperation phase until t^D, followed by the defection phase. The expected payoff from this is easily calculated to be greater than that for a player who behaves differently because then the off-equilibrium strategy prescribes that both players switch to D as soon as possible (unless they are already playing D) and keep to D until the end of the game.

WoA phase: As the game starts with (D,D), the *ex ante* payoff of player i of type θ_i for switching to C at $t \le t^A$, unless j switches first, is

$$\begin{aligned}
&\int_0^t qe^{-qx}\big(x\ell + \bar{\theta}h + (t^D - x - \bar{\theta}) + V_{t^D}\big)dx \;+\; e^{-qt}\Big(t\ell + (t^D - t - \bar{\theta}) + V_{t^D}\Big)\\
&= \frac{e^{-qt}}{q}\Big(1-\ell-\bar{\theta}qh - e^{qt}(1-\ell-\bar{\theta}q(h-1)-qt^D)\Big) + V_{t^D}
\end{aligned}$$

where V_{t^D} denotes the equilibrium continuation payoff (of player i of type θ_i) when the game reaches t^D with (C,C). The derivative of this with respect to t is $-e^{-qt}(1-\ell-\bar{\theta}qh) = 0$ for all $t \le t^A$ because $q = (1-\ell)/(\bar{\theta}h)$. Therefore, both players of all types are indifferent between switching to C at any $t \le t^A$.

As the probability that each player switches to C by t^A (unpreceded by the opponent) is less than 1 in the considered equilibrium, it remains to be shown that both players are indifferent between switching to C at t^A and not switching at any $t \le t^A$. Because these two options are identical if the other player, j, switches to C before t^A, consider the case in which player j does not switch until t^A. (The contingency that player j switches exactly at t^A is a null event and thus is left out of the calculation.) By switching at t^A, player $i \neq j$ receives an *ex ante* payoff of $t^A\ell + (t^D - t^A - \bar{\theta}) + V_{t^D}$. By not switching at any $t \le t^A$, she receives the same payoff from a payout of ℓ until $t^C = t^A + \frac{\bar{\theta}}{1-\ell}$ followed by 1 until t^D, establishing the desired indifference.

Appendix B

Proof of Proposition 1 Consider an arbitrary PBE* (P_1, P_2). Recall that $\zeta_j(t) = \int_\Theta P_j(t|\theta)dF$ is the ex-ante probability that player $j \in \{1,2\}$ defects at or before t according to P_j, presuming that player $i \neq j$ does not preempt player j. Define

$$\text{supp}^\circ P_i(\cdot|\theta) := \{t \,|\, P_i(\{t\}|\theta) > 0 \text{ or } t \in int(\text{supp}\, P_i(\cdot|\theta))\}$$

where $P_i(\{t\}|\theta)$ is the point mass that $P_i(\cdot|\theta)$ places at t and *int* designates the interior of a set. Here a series of properties are derived that ζ_j and $\text{supp}^\circ P_i(\cdot|\theta)$ should satisfy for $i \neq j$, which will eventually prove part (a) of Proposition 1.

In particular, start with two properties that ζ_j imposes on $\text{supp}^\circ P_i(\cdot|\theta)$:

[C] If ζ_j places no point mass at t' and $\zeta_j(t') = \zeta_j(t'') < 1$ where $t' < t'' \leq T$, then $\text{supp}^\circ P_i(\cdot|\theta_i) \cap [t', t'') = \emptyset$ for all $\theta_i \in \Theta$.

[D] If ζ_j is discontinuous at $t > 0$, $\text{supp}^\circ P_i(\cdot|\theta_i) \cap [t, t+\epsilon] = \emptyset$ for all θ_i for some $\epsilon > 0$.

Property [C] ensues because player i of any type would get a higher payoff defecting at t'' than at any $t \in [t', t'')$ by prolonging cooperation with no risk of being preempted in the meantime. To see [D], suppose that ζ_j places a positive point mass at t, say $\mu > 0$, and let $dH(\cdot|t)$ denote the type distribution of player j who would defect at t. Then, as the probability that ζ_j places on the interval $[t-\epsilon, t+\epsilon]$ converges to μ as $\epsilon \to 0$, the payoff gain of player i of type θ_i from defecting at $t - \epsilon$ rather than at any point in $[t, t+\epsilon]$ is first-order approximated by an amount no less than $\mu \int_{\theta_j} W(t, \theta_i, \theta_j) - E(t, \theta_i, \theta_j) dH(\cdot|t) > 0$ as $\epsilon \to 0$. Hence, player i would not defect at any point in a small interval $[t, t+\epsilon]$, proving [D].

Next, we show that

[E] If $t' \in \text{supp}^\circ P_i(\cdot|\theta')$ then $P_i(t'|\theta'') = 1$ for all $\theta'' > \theta'$.

To prove by contradiction, suppose to the contrary that

$$t' \in \text{supp}^\circ P_i(\cdot|\theta') \quad \text{and} \quad t'' \in \text{supp}^\circ P_i(\cdot|\theta''), \quad \text{where} \quad t' < t'' \quad \text{and} \quad \theta' < \theta''. \tag{18}$$

Then, ζ_j is continuous at t'' by [D]. Let μ' denote the point mass placed by ζ_j at t', which is 0 if $t' > 0$ by [D], and let $\mu'' = \zeta_j(t'') - \zeta_j(t')$ and $\mu''' = 1 - \zeta_j(t'')$.

The net change in player i's expected payoff from defecting at t'' rather than at t', can be decomposed into the following three parts. First, relative to the types of player j who would defect at t', the distribution of which is represented by $H(\theta_j|t')$, the aforesaid net change of player i of type θ is

$$\mu' \cdot \int_{\theta_j} \big(L(t', \theta) - E(t', \theta, \theta_j)\big) dH(\cdot|t')$$

which is strictly larger for $\theta = \theta'$ than for $\theta = \theta''$ if $\mu' > 0$ by [u4].

Second, relative to any type θ_j of player j who would defect at $t \in (t', t'')$, it is

$$L(t, \theta) - W(t', \theta, \theta_j)$$

which is strictly larger for θ' than for θ'' because (i) $W(t, \theta', \theta_j) - L(t, \theta') < W(t, \theta'', \theta_j) - L(t, \theta'')$ by [u4] and (ii) $W(t', \theta', \theta_j) - W(t, \theta', \theta_j) = W(t', \theta'', \theta_j) - W(t, \theta'', \theta_j)$ by [u3].

Third, relative to the types of player j who would defect after t'', the distribution of which is denoted by $F(\theta_j|t'')$, the aforesaid net change of player i of type θ is

$$\mu''' \cdot \Big[\int_{\theta_j} W(t'', \theta, \theta_j) dF(\cdot|t'') - \int_{\theta_j} W(t', \theta, \theta_j) dF(\cdot|t'') \Big] = \mu'''(t'' - t')a \geq 0$$

due to [u3], hence it is the same and positive for θ' and for θ''.

Therefore, unless $\mu' = \mu'' = 0$, the total net change in payoff is strictly larger for θ' and consequently, the fact that θ''-type weakly prefers to defect at t'' than at t' (i.e., $t'' \in \operatorname{supp}^\circ P_i(\cdot|\theta'')$) would imply that θ'-type strictly prefers to do so, contradicting the supposition $t' \in \operatorname{supp}^\circ P_i(\cdot|\theta')$ of (18). If $\mu' = \mu'' = 0$, on the other hand, both types should strictly prefer to defect at t'' if $\mu''' > 0$, again contradicting (18). Finally, if $\mu' = \mu'' = \mu''' = 0$ then as player j will have defected for sure by t', it would be suboptimal for player i of any type to defect any time after t', contradicting $t'' \in \operatorname{supp}^\circ P_i(\cdot|\theta'')$ of (18). Thus, (18) above cannot hold in any PBE*, which proves [E].

Now, suppose that $t', t'' \in \operatorname{supp}^\circ P_i(\cdot|\theta')$ where $t' < t''$ for some $\theta' \in (0, \bar{\theta})$. Then, [E] implies that $P_i(t'|\theta'') = 1$ for all $\theta'' > \theta'$ and that $P_i(t|\theta) = 0$ for all $t < t''$ and $\theta < \theta'$, so that ζ_i is constant on $t \in (t', t'')$, which in turn implies that ζ_j is constant on (t', t'') by [C] and thus, on $[t', t'')$ by right-continuity. If $t' > 0$ then ζ_j must be continuous at t' by [D] because $t' \in \operatorname{supp}^\circ P_i(\cdot|\theta')$, but this would contradict $t' \in \operatorname{supp}^\circ P_i(\cdot|\theta')$ by [C]. Hence, $t' = 0$ must hold. Therefore, together with [C], this establishes that

[F] $P_i(\cdot|\theta)$ places the full mass at a single time point, except possibly for one type, say $\hat{\theta}$, who may split the full mass between two time points, one of which is 0.

At this point, the aforementioned convention is imposed that this type ($\hat{\theta}$) places the full mass to the latter time point without affecting the equilibrium condition, so that every PBE* is a pure strategy equilibrium.

Thus, now any PBE* may be denoted by non-inceasing functions τ_1, $\tau_2 : \Theta \to [0, T]$. Let $\tau_i(0) = \lim_{\theta \downarrow 0} \tau_i(\theta)$ and $\tau_i(\bar{\theta}) = \lim_{\theta \uparrow \bar{\theta}} \tau_i(\theta)$ for $i = 1, 2$. Then, $\tau_1(0) = \tau_2(0)$ because it is suboptimal for either player of any type to not defect until the opponent defects (to preempt) definitely. As any open interval outside of $\operatorname{supp} \tau_i$ is outside of $\operatorname{supp} \tau_{-i}$ as well by [C],

[G] $\operatorname{supp} \tau_1 = \operatorname{supp} \tau_2$ and hence, $\operatorname{supp} \zeta_1 = \operatorname{supp} \zeta_2$.

If ζ_i placed a positive point mass at $t > 0$, player j would not defect during $[t, t + \epsilon]$ for a small $\epsilon > 0$ by [D] and thus, $\zeta_j(t) = 1$ because $\zeta_j(t) < 1$ would contract ζ_i placing a point mass at t by [C]. This would imply $t = \tau_j(0) = \tau_i(0)$, but then it would be suboptimal for player i to defect at $t = \tau_i(0)$ because the opponent will have defected for sure by then. Hence, we deduce that ζ_i may place a positive point mass only at $t = 0$. Moreover, if $\zeta_i(0) > 0$ while $\zeta_j(0) = 0$, then $\zeta_j(t) > 0$ for all $t > 0$ by [G], but player j would prefer to defect at $t = 0$ than at $t = \epsilon$ for sufficiently small ϵ because her gain from doing so is first-order approximated by $\zeta_i(0) \cdot \int E(0, \theta_j, \theta_i) - L(0, \theta_j) dH(\theta_i|0) > 0$ where $H(\theta_i|0)$ is the distribution of player i types who would defect at $t = 0$. Therefore,

[H] Either both ζ_1 and ζ_2 are atomless, or both place a positive mass at and only at $t = 0$.

As the final step for proving part (a), suppose that τ_i is discontinuous at some $\widehat{\theta}_i \in \Theta$. As τ_i is non-increasing, this means that $\tau_i(\widehat{\theta}_i^+) = \lim_{\theta \downarrow \widehat{\theta}_i} \tau_i(\theta) < \tau_i(\widehat{\theta}_i^-) = \lim_{\theta \uparrow \widehat{\theta}_i} \tau_i(\theta)$. Then, as ζ_1 and ζ_2 is constant on $(\tau_i(\widehat{\theta}_i^+), \tau_i(\widehat{\theta}_i^-))$ by [C], if $\tau_i(\widehat{\theta}_i^+) > 0$ so that ζ_1 and ζ_2 places no point mass at $\tau_i(\widehat{\theta}_i^+)$ by [H], then player i of types $\theta_i > \widehat{\theta}_i$ arbitrarily close to $\widehat{\theta}_i$ would benefit by waiting until $\tau_i(\widehat{\theta}_i^-)$ due to [A] and [u3] because the gain from extending cooperation with the types of player j who would not defect until $\tau_i(\widehat{\theta}_i^-)$ is linear in the waiting time while the loss from being preempted by the types of player j who would defect just before $\tau_i(\widehat{\theta}_i^+)$ is negligible. As this contradicts optimality of such types of player i, it follows that if τ_i is discontinuous at $\widehat{\theta}_i \in \Theta$, then $\tau_i(\theta_i) = 0$ for all $\theta_i > \widehat{\theta}_i$ and furthermore, τ_i is continuous and strictly decreasing for $[\widehat{\theta}_i, \bar{\theta})$ with $\tau_i(\widehat{\theta}_i) = \tau_i(\widehat{\theta}_i^-)$ by [H] and the convention adopted. This establishes that a PBE* is semi-pooling if both players defect at $t = 0$ with a positive probability, and is separating otherwise. Together with [G], this proves part (a) of Proposition 1.

Recall that part (b) of the proposition has been proved in Section 3.

To prove part (c), consider an arbitrary PBE* (τ_1, τ_2) such that $\operatorname{supp}\tau_1 = \operatorname{supp}\tau_2$ by part (a). Let $\bar{t} = \lim_{\theta \to 0} \tau_i(\theta)$ and $\underline{t} = \lim_{\theta \uparrow \widehat{\theta}_i} \tau_i(\theta_i)$ where $\widehat{\theta}_i$ is player i's threshold type for a semi-pooling PBE* and $\widehat{\theta}_i = \bar{\theta}$ for a separating PBE*. Apart from the trivial PBE* in which both players of all types defect at $t = 0$ for sure (which is symmetric), the inverse function $\vartheta_i = \tau_i^{-1}$ is well-defined on $(0, \widehat{\theta}_i)$. Moreover, the incentive compatibility conditions for types in $(0, \widehat{\theta}_i)$ are calculated analogously to those for symmetric separating PBE* that led to (12) in Section 3, producing an autonomous differential equation system as follows:

$$\vartheta_i'(t) = -\Psi(\vartheta_i(t), \vartheta_{-i}(t)) = \frac{-aF(\vartheta_i(t))}{\Big(W\big(0, \vartheta_{-i}(t), \vartheta_i(t)\big) - L\big(0, \vartheta_{-i}(t)\big)\Big) f(\vartheta_i(t))} \quad \text{for} \quad i = 1, 2. \tag{19}$$

If the considered PBE* (τ_1, τ_2) is separating but not symmetric, let $t^* = \sup\{t | \vartheta_1(t) \neq \vartheta_2(t)\} \in (\underline{t}, \bar{t}]$ and $\vartheta_1(t^*) = \vartheta_2(t^*) = \theta^* \in (0, \bar{\theta}]$. Since $\Psi(\theta_1, \theta_2)$ is locally Lipschitz-continuous in $\Theta \times \Theta$ by the properties assumed for W, L and f, as well as in a neighbourhood of $(\bar{\theta}, \bar{\theta})$ as assumed for part (c) of the proposition, by Theorems 3 and 4 of page 28 of Hurewicz (1958), there is a unique solution to the ODE system (19) subject to the initial condition $\vartheta_1(t^*) = \vartheta_2(t^*) = \theta^*$. However, symmetry of the ODE system (19) between players 1 and 2 would imply that $(\widetilde{\vartheta}_1, \widetilde{\vartheta}_2) = (\vartheta_2, \vartheta_1)$ should also constitute a solution to (19) subject to the same initial condition, contradicting the uniqueness of solution. This proves that any separating PBE* is symmetric.

Next, suppose that (τ_1, τ_2) is semi-pooling but not symmetric with the threshold types $\widehat{\theta}_i \in \Theta$ for $i = 1, 2$, and the common threshold time $\widehat{t} > 0$. If $\widehat{\theta}_1 = \widehat{\theta}_2$, then the same argument as above leads to a contradiction to the uniqueness of the solution. Suppose otherwise, say $\widehat{\theta}_2 > \widehat{\theta}_1$. Then, the incentive compatibility that thresholds types are indifferent between defecting at $t = 0$ and at $t = \widehat{t}$ produce conditions analogous to (15):

$$\frac{1}{aF(\widehat{\theta}_2)} \int_{\theta_2 = \widehat{\theta}_2}^{\bar{\theta}} E(0, \widehat{\theta}_1, \theta_2) - L(0, \widehat{\theta}_1)\, dF \ = \ \widehat{t} \ = \ \frac{1}{aF(\widehat{\theta}_1)} \int_{\theta_1 = \widehat{\theta}_1}^{\bar{\theta}} E(0, \widehat{\theta}_2, \theta_1) - L(0, \widehat{\theta}_2)\, dF \tag{20}$$

where the first equality is for player 1 and the latter for player 2. But, this is impossible because $F(\widehat{\theta}_2) > F(\widehat{\theta}_1)$ and $E(0, \widehat{\theta}_1, \theta) - L(0, \widehat{\theta}_1) < E(0, \widehat{\theta}_2, \theta) - L(0, \widehat{\theta}_2)$ by [u4], so that the RHS of (20) is greater than the LHS. This verifies that any semi-poiling PBE* is symmetric, proving part (c). $\square$

Proof of Proposition 2 Suppose that $\bar{\theta}$, F, f, W, L, D, a and ψ are player-specific, which is indicated by subscripts $i = 1, 2$. Then, the incentive compatibility condition, that generalises (12) obtained for symmetric players, is written separately for the two players as:

$$\vartheta_1'(t) = -\Psi_1(\vartheta_1(t), \vartheta_2(t)) = \frac{-a_2 F_1(\vartheta_1(t))}{\Big(W_2\big(0, \vartheta_2(t), \vartheta_1(t)\big) - L_2\big(0, \vartheta_2(t)\big)\Big) f_1(\vartheta_1(t))}, \quad (21)$$

$$\vartheta_2'(t) = -\Psi_2(\vartheta_2(t), \vartheta_1(t)) = \frac{-a_1 F_2(\vartheta_2(t))}{\Big(W_1\big(0, \vartheta_1(t), \vartheta_2(t)\big) - L_1\big(0, \vartheta_1(t)\big)\Big) f_2(\vartheta_2(t))} \quad (22)$$

for players 2 and 1, respectively. Due to condition (i) of Proposition 2, [u2], [u4], and Lipschitz-continuity of f_i, the RHS of (21) and (22) are negative and bounded away from 0 when $\vartheta_1(t) = \vartheta_2(t) > 0$. Thus, by [u2] and [u4],

[X] there is $x < 0$ such that $\vartheta_i'(t) < x$ if $\vartheta_i(t) \geq \vartheta_{-i}(t) > 0$ for $i = 1, 2$, and

[Y] for arbitrarily small $\epsilon > 0$, there is $y < 0$ such that $y < \vartheta_i'(t) < 0$ if $\vartheta_i(t) < \vartheta_{-i}(t)$ and $\vartheta_i(t) < \bar{\theta}_i - \epsilon$ for $i = 1, 2$.

Here, [X] follows because the denominator of $\Psi_i(\theta_i, \theta_{-i})$ is bounded above (while $F(\theta_i)$ increases) as θ_i increases from θ_{-i}, and [Y] ensues because $\Psi_i(\theta_i, \theta_{-i})$ is bounded when $\theta_i = \theta_{-i} < \bar{\theta}_i - \epsilon$ and the denominator increases as θ_{-i} increases from θ_i.

Fix the value of ϵ in [Y] at a level less than $\min\{\bar{\theta}_1/4, \bar{\theta}_2/4\}$. Consider the following initial condition: $(\vartheta_1(t_0), \vartheta_2(t_0)) = \beta(\epsilon, \bar{\theta}_2) + (1-\beta)(\bar{\theta}_1, \epsilon)$ for some t_0 where $\beta \in (0, 1)$, which we refer to as the "initial condition β." Possible values of this initial condition for various $\beta \in (0, 1)$ comprises the set of all convex combinations of the two points $(\epsilon, \bar{\theta}_2)$ and $(\bar{\theta}_1, \epsilon)$, which is denoted by $J \subset \Theta_1 \times \Theta_2$. Let $\nabla = \{(\theta_1, \theta_2) \in \Theta_1 \times \Theta_2 \,|\, (\theta_1, \theta_2) \geq (\theta_1', \theta_2')$ for some $(\theta_1', \theta_2') \in J\}$, and $K = bdry(\nabla) \setminus J$ be the points in the boundary of ∇ that are not in J.

As the RHS of (21) and (22) are locally Lipschitz-continuous in $(\vartheta_1(t), \vartheta_2(t)) \in \Theta_1 \times \Theta_2$, by Theorems 3 and 4 on page 28 of Hurewicz (1958), there is a unique solution $(\vartheta_1, \vartheta_2)$ to (21)-(22) subject to any given initial condition $\beta \in (0, 1)$. Let $(\underline{t}, \bar{t})$ denote the domain of the solution $\vartheta_i, = 1, 2$ (the dependence of domain on β is suppressed for notational simplicity). Let $\vartheta_i(\underline{t}) = \lim_{t \to \underline{t}} \vartheta_i(t)$ and $\vartheta_i(\bar{t}) = \lim_{t \to \bar{t}} \vartheta_i(t)$

As the RHS of (21) and (22) are negative and bounded away from 0 for $(\vartheta_1(t), \vartheta_2(t)) \in \nabla$, it follows that $\underline{t} > -\infty$ and $(\vartheta_1(\underline{t}), \vartheta_2(\underline{t})) \in K$. By [X] and [Y] above, $\vartheta_1(\underline{t}) < \bar{\theta}_1$ for β sufficiently close to 1 and $\vartheta_2(\underline{t}) < \bar{\theta}_2$ for β sufficiently close to 0. As the solution to (21)-(22) is continuous in the initial condition by Theorem 8 on page 29 of Hurewicz (1958), $(\vartheta_1(\underline{t}), \vartheta_2(\underline{t})) = (\bar{\theta}_1, \bar{\theta}_2)$ must hold for some value of $\beta \in (0, 1)$, denoted by β^*.

Consider the unique solution $(\vartheta_1, \vartheta_2)$ to (21)-(22) subject to the initial condition β^*. As the RHS of (21) and (22) are negative, $(\vartheta_1(\bar{t}), \vartheta_2(\bar{t})) \in \mathbb{R}^2_+ \setminus \mathbb{R}^2_{++}$. It now remains to show that $\bar{t} < \infty$ and $(\vartheta_1(\bar{t}), \vartheta_2(\bar{t})) = 0$.

As time passes from t_0, by property [X] above, the total duration of time for which $(\vartheta_1(t), \vartheta_2(t))$ may stay on or above the 45^o line of the first quadrant before hitting either axis is bounded above, and so is the time for which $(\vartheta_1(t), \vartheta_2(t))$ may stay below the 45^o line of the first quadrant. This establishes that $\bar{t} < \infty$.

At last, to reach a contradiction, suppose that $\vartheta_2(\bar{t}) = \theta_2^* > 0$ while $\vartheta_1(\bar{t}) = 0$. (The case that $\vartheta_1(\bar{t}) > \vartheta_2(\bar{t}) = 0$ is symmetric.) Note that $(\vartheta_1, \vartheta_2)$ extended to the domain $(\underline{t}, \bar{t}]$ continues to be a solution to the system (21)-(22) when this system is extended to the points on the axes by $\Psi_i(\theta_i, 0) = \lim_{\theta_{-i}\to 0} \Psi_i(\theta_i, \theta_{-i})$ for $\theta_i > 0$ and $\Psi_i(0, \theta_{-i}) = \lim_{\theta_i \to 0} \Psi_i(\theta_i, \theta_{-i})$ for $\theta_{-i} > 0$, subject to the same initial condition β^*. As the RHS of (21) and (22) are Lipschitz-continuous in a neighbourhood of $(0, \theta_2^*)$ by assumption (ii) of Proposition 2, $(\vartheta_1, \vartheta_2)$ must be the unique solution subject to the initial condition $(\vartheta_1(\bar{t}), \vartheta_2(\bar{t})) = (0, \theta_2^*)$. However, because the RHS of (21) is 0 when $\vartheta_1(t) = 0$ and $\vartheta_2(t) > 0$, so that $\vartheta_i'(\bar{t}) = 0$, it follows that the solution to (21)-(22) for the initial condition $(\vartheta_1(\bar{t}), \vartheta_2(\bar{t})) = (0, \theta_2^*)$ must satisfy $\vartheta_i(t) = 0$ in a neighborhood of $\bar{t}$, contradicting the uniqueness of solution. Therefore, it is deduced that $(\vartheta_1(\bar{t}), \vartheta_2(\bar{t})) = (0, 0)$ must hold, establishing the existence of a separating PBE* when T is sufficiently large with the features described in Proposition 2. If $\Psi_i(\theta_i, \theta_{-i})$ is Lipschitz-continuous in a neighbourhood of $(\bar{\theta}_i, \bar{\theta}_{-i})$ for $i = 1, 2$, then there is a unique value of $\beta^* \in (0, 1)$ identified above, by the uniqueness of the solution to the system (21)-(22) with the initial condition $(\vartheta_1(\underline{t}), \vartheta_2(\underline{t})) = (\bar{\theta}_1, \bar{\theta}_2)$ and thus, a unique separating PBE* subject to the initial condition $(\vartheta_1(\underline{t}), \vartheta_2(\underline{t})) = (\bar{\theta}_1, \bar{\theta}_2)$.

Finally, time-invariance is a trivial consequence of (21)-(22) being an autonomous differential equation system. □

References

Abreu, Dilip, and Markus Brunnermeier. 2003. "Bubbles and Crashes." *Econometrica* **71**: 173–204.

Ambrus, Attila, and Parag A. Pathak. 2011. "Cooperation over finite horizons: A theory and experiments." *Journal of Public Economics* **95**: 500–512.

Bergin, James, and W. Bentley MacLeod. 1993. "Continuous Time Repeated Games." *International Economic Review* **34**: 21–37.

Billingsley, Patrick. 1995. *Probability and Measure,* 3rd ed., John Wiley & Sons, NJ., USA.

Friedman, Daniel, and Ryan Oprea. 2012. "A Continuous Dilemma." *American Economic Review* **102**: 337–363.

Fudenberg, Drew, and Eric Maskin. 1986. "The Folk Theorem in Repeated Games with Discounting and Incomplete Information." *Econometrica* **54**: 533–554.

Fudenberg, Drew, and Jean Tirole. 1985. "Preemption and Rent Equalization in the Adoption of New Technology." *Review of Economic Studies* **52**: 383–401.

Hopenhayn, Hugo, and Francesco Squintani. 2011. "Preemption Games with Private Information." *Review of Economic Studies* **78**: 667–692.

Hurewicz, Witold. 1958. *Lectures on Ordinary Differential Equations,* MIT Press, Cambridge, MA., USA.

Jehiel, Philippe. 2005. "Analogy-based Expectation Equilibrium." *Journal of Economic Theory* **123**: 81–104.

Kreps, David M., Paul Milgrom, John Roberts, and Robert Wilson. 1982. "Rational Cooperation in the Finitely Repeated Prisoners' Dilemma." *Journal of Economic Theory* **27**: 245–252.

Ma, C-T. Albert, and Michael Manove. 1993. "Bargaining with Deadlines and Imperfect Player Control." *Econometrica* **61**: 1313–1339.

Neyman, Abraham. 1985. "Bounded Complexity Justifies Cooperation in the Finitely Repeated Prisoners' Dilemma." *Economics Letters* **19**: 227–229.

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

Perry, Motty, and Philip Reny. 1993. "A Non-cooperative Bargaining Model with Strategically Timed Offers." *Journal of Economic Theory* **59**: 50–77.

Radner, Roy. 1980. "Collusive Behaviour in Non-Cooperative Epsilon-Equilibria in Oligopolies with Long but Finite Lives." *Journal of Economic Theory* **22**: 136–154.

Radner, Roy. 1986. "Can Bounded Rationality Resolve the Prisoners Dilemma?" In *Contributions to Mathematical Economics,* edited by Andreu Mas-Colell and Werner Hildenbrand, 387–399. Amsterdam: North-Holland.

Reinganum, Jennifer. 1981. "On the Diffusion of New Technology: A Game-theoretic approach." *Review of Economic Studies* **48**: 395–405.

Simon, Leo K. and Maxwell B. Stinchcombe. 1989. "Extensive Form Games in Continuous Time: Pure Strategies." *Econometrica* **57**: 1171-1214.

Weeds, Helen. 2002. "Strategic Delay in a Real Options Model of R&D Competition." *Review of Economic Studies* **69**: 729–747.