

Uncertain Reaction lag and Cooperation in Continuous Prisoners' Dilemma

IN-UCK PARK*

University of Bristol, U.K. and IAS, Princeton, U.S.A.

July 3, 2012

(Very Preliminary—do not circulate)

1 Introduction

Friedman-Oprea (2012) report experimental results of a continuous-time prisoners' dilemma game. They find that the cooperation is much higher in their experiment (above 90% of the duration of the game which is 60 seconds) than in the previous experiments in discrete time settings. They also report that, when both subjects start with Defect (D) at the beginning, typically one player switches to Cooperate (C) in the first few seconds with the other following suit within one second, and then one player switches back to D a few seconds before the end with the other quickly following suit. This pattern seems natural, yet has not been theoretically justified as an equilibrium (as opposed to π -equilibrium).

As Friedman-Oprea point out, the asynchronous nature of moves may contribute to the increased cooperation in continuous-time environments where the reaction can be quick. That is, players may be more willing to continue with C, betting on the opponent doing the same, when they can respond quickly upon finding out that their bet was wrong, relative to the situations in which wrong bets cannot be mended for a fixed duration of time as in discrete time settings.

The basic insight may be depicted by the varying efficacy of the “trigger strategy until the end” that has been examined in previous studies: it is an π -best response if the time lag to respond to “cheating” is very small (Radner, 1986; Bergin-MacLeod, 1993) and it is a best response if the time lag is nil (Simon-Stinchcombe, 1989).

A key difficulty in supporting cooperation in finite-horizon settings of the prisoners' dilemma game is to prevent unraveling from the end of the game. The aforementioned experimental data exhibit prolonged cooperation in spite of eventual defections by both players. The data also exhibit a small yet clear time lag in responding to the opponent's moves. If the length of the time lag is public information as in the papers cited above, the

*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, Dan Friedman, Bentley MacLeod, Stephen Morris, Hyun Song Shin, and Aleksey Zinger for various helpful comments and suggestions. All remaining errors are mine.

temptation to pre-empt the opponent unravels cooperation incentives all the back to the beginning. Thus, the prolonged cooperation followed by eventual defections remain as a theoretical puzzle.

In this paper I resolve this puzzle by introducing *heterogeneous reaction lags as private information*. The key intuition is that the quicker a player is in responding, the longer she is willing to continue with cooperation. 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 a sufficiently quick type, will cooperate a little longer (unless one already defected). This, in turn, justifies each player's willingness to cooperate a little longer if her reaction lag is short enough.

I characterize an equilibrium that exhibits such features, and also calculate the actual equilibrium for parameter values corresponding to the experiments of Friedman-Oprea. The calculated equilibrium matches the experimental data closely. I also discuss some modelling issues in continuous-time and the approach taken to deal with them.

As the upper bound of the reaction lag vanishes, the limit of the aforesaid equilibrium prescribes cooperation throughout the whole duration of the game, which is the solution predicted by Simon-Stinchcombe (1989) and Bergin-MacLeod (1993) for the continuous-time game of prisoners' dilemma. Therefore, the current analysis supports their predictions as the limiting equilibrium as the reaction lag vanishes, which has not been established in the literature as far as I am aware.

In my model, cooperation is sustained until toward the end by way of endogenously creating and capitalizing on the prospect at any point in time that the opponent may continue with cooperation a little longer. In a famous article, Kreps, et al. (1982) also support cooperation in a discrete-time, finitely repeated game of prisoners' dilemma based on a similar idea: if each player is perceived to be committed to a "Tit-for-Tat" strategy with a small probability, a normal player will masquerade as such a type to induce cooperation from the opponent who may entertain the idea that she may facing such a commitment type. My approach differs crucially in that it does not introduce uncertainty on the opponent's preference, and consequently, does not rely on the idea of reputation building (i.e., some type trying to imitate another type) to be perceived as committed to a particular behavior pattern conducive to cooperation. Rather, in my model each type would defect at a uniquely optimal time that maximizes the same utility function within its own physical constraint. As a result, the "end play" of the equilibrium (i.e., after the cooperation breaks down) is very simple in my model, unlike in the model of Kreps, et al. (1982) in which the authors find it "very complex" that they do not characterize it in their article.

2 Illustration

A. Game

Two players, $i = 1, 2$, play the following game over the time interval $I = (0, 1]$:

$1 \backslash 2$	C	D
C	(1,1)	(0,h)

$$D = (h, 0) \text{ (} \cdot \cdot \text{)}$$

where $0 < \cdot < 1 < h$.¹ The players start with $(D; D)$ at $t = 0$ and at any time $t \in I = (0; 1]$, they may switch to the other action and back as they wish (subject to constraints below), which is observed by the other player instantaneously. The players receive flow payouts determined by the action pairs prevailing at each point $t \in [0; 1]$ without discounting.

The players face two related constraints concerning reaction time:

- [A] If player j switches her action at $t \in (0; 1)$, then player $i \neq j$ cannot switch her action during $(t; t + \tau_i)$.
- [B] If player i switches her action at $t \in (0; 1)$, then she cannot switch her action again during $(t; t + \tau_i)$.

Part [A] means that $t + \tau_i$ is the first instance of time after t , at which player i may respond to her opponent's move at t (provided that $t + \tau_i \leq 1$). This reflects the time needed for player i to recognize the changed situation, decide how to react, and put the decision into action. The value of τ_i , which we call the "reaction lag," is private information (type) of player i , $i = 1, 2$, which is an independent draw from a common distribution represented by an atomless cdf $F(\cdot)$ on $(0; \bar{\tau})$ where $\bar{\tau} > 0$ is small, with an associated density function $f(\cdot) = F'(\cdot)$ which is assumed to be continuous: i.e.,

$$(1) \quad \tau_i \sim F \text{ on } (0; \bar{\tau}) \text{ where } \lim_{\theta \downarrow 0} F(\theta) = 0 \text{ and } f(\cdot) = F'(\cdot) \text{ is continuous.}$$

Part [B] postulates an analogous minimum time lag between any two consecutive moves by the same player, which reflects the physical constraint of the player. It seems plausible that $\tau_i \leq \tau_i$. For expositional ease (to avoid two-dimensional types), we assume $\tau_i = \tau_i$ for a small $\tau_i \in (0; 1]$.

The game is conceptually quite simple, yet players may make multiple moves in a continuous-time framework. A rigorous formulation of the players' strategies in such environments requires laying down a nontrivial analytical structure, as is done in the next section. In this section, an informal description seems to suffice: A strategy of an "agent" (a player of a given type) specifies, contingent on every possible history, a probability that she will have switched her action by every future point in time, conditional on no player has switched her action by then after the referenced history.² An agent's strategy is required to respect [A], [B], and an inter-temporal consistency condition stipulating that an agent's plan (of switching action) contingent on a possible future history does not change as the history evolves toward that particular future history and reaches it. A profile of agent strategies of the two players constitutes a perfect Bayesian equilibrium if the usual mutual best-response property is satisfied and the belief profile is consistent with Bayes rule.

B. Description of equilibrium

We present a symmetric equilibrium for the case that F is uniform, i.e., $F(\cdot) = \frac{\cdot}{\bar{\tau}}$ and $f(\cdot) = \frac{1}{\bar{\tau}}$. It is described below as a sequence of several phases.

¹We could assume $h < 2$ so that (C, C) is efficient, but not needed for our results.

²That is, it consists of an "open-loop" strategy for every history until the next move takes place.

- 1) War-of-attrition (WoA) phase: Starting from $(D; D)$ at the beginning of the game, both players switch to C with a common “flow rate” $q > 0$ provided that the time has not passed T and the other has not switched to C already, where $T < 1 - \bar{\gamma}(h - \bar{\gamma} + \frac{1+h}{1-\ell})$ is a fixed point in time. Thus, for each player the cumulative probability that she will have switched to C at or before $t \in (0; T]$ provided that the other has not switched to C by then, is $1 - e^{-qt}$. We refer to this initial part of the game until one player switches to C at some point in time before T , as the “war-of-attrition (WoA) phase.”
- 2) Transition phase to cooperation: If one player, say i , switched to C in the WoA phase, say at $t \leq T$, then player j follows suit $\bar{\gamma}$ later, i.e., switches to C at $t + \bar{\gamma}$ (and i adheres to C until then).³ In case neither player switched to C by T , both players switch to C simultaneously at $T^\theta = T + \frac{\bar{\theta}}{1-\ell} < 1 - (\frac{h}{1-\ell} + h - \bar{\gamma})$.⁴
- 3) Cooperation phase: After the two players both switched to C as explained in 2) above, they continue with $(C; C)$ until the time reaches $\hat{T} = 1 - \frac{h}{1-\ell}$. We refer to this part of the game as the “cooperation phase.”
- 4) Defection phase: Once the time reaches \hat{T} through the cooperation phase, each player switches to D at an appropriate time point during $(\hat{T}; 1)$ depending on her type, conditional on the other not having switched to D already: Specifically, a player of type $\bar{\gamma}$ defects when the remaining time is $r(\bar{\gamma}) = \frac{h}{1-\ell}$, i.e., at $t = 1 - r(\bar{\gamma})$. Once one player, say i , defects at t , player $j \neq i$ follows suit as soon as possible, i.e., at $t + \bar{\gamma}_j$, after which both players keep to D until the end of the game. (As shown below, both will switch to D before the game ends.) We refer to this part of the game as the “defection phase.”
- 5) Off-equilibrium: If history at any t departs from the equilibrium-paths described above, then both players switch to D as soon as they can, unless they were already playing D , and keep to D until the end of the game. A caveat is in order: in the off-equilibrium contingency that one player, say i , switched to C at t in the WoA phase, yet the other player did not follow suit at $t + \bar{\gamma}$, the earliest time after this deviation that player i can switch back to C

verify that, off-the-equilibrium path, it is optimal for each player to keep or switch to D as soon as possible, given that the other will do the same (apart from the aforementioned exception case which we will deal with later). Below we verify optimality of the strategy on the equilibrium-path, backwards from the defection phase.

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 specified above.

Next, consider any point $t \in (\hat{T}; 1)$ in the defection phase with no defection having taken place from $(C; C)$ by either player until then. According to the defection strategy described above, either player may defect at any point between \hat{T} and 1 depending on her type, with a higher type defecting earlier. Letting $r = 1 - t$ denote the time remaining until the end of the game, therefore, the posterior belief of either player's type at $t = 1 - r$ is uniform on the interval $(0; \#(r))$ where $\#(r) = \frac{1-\ell}{h}r$ is the type that is supposed to defect at $t = 1 - r$, because a player of a type higher than $\#(r)$ should have defected by $t = 1 - r$. Consequently, each player anticipates that the other player will defect with a uniform distribution during the remaining period of the game, $(t; 1)$.

Given this, let's calculate the expected payoff of a player, say i , from waiting an infinitesimal amount of time, Δ , from $t = 1 - r$ before defecting. In the contingency that she gets pre-empted between t and $t + \Delta$ (i.e., player $-i$ defects in between), player i 's payoff is higher the lower her type, θ_i , is (because lower θ_i means that she can follow suit more quickly); but her payoff is the same regardless of her type if she does not get pre-empted. Therefore, if the marginal net change in player i 's expected payoff from waiting is 0 when her type is $\#(r)$, any lower type should strictly prefer to wait, as prescribed by the proposed equilibrium strategy. The change in payoff for a $\#(r)$ -type agent is $\Delta(1 - \ell)$ if she does not get pre-empted (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); and it will be $-\#(r)h$ if she gets pre-empted (because she, instead of her opponent, will suffer from the "sucker" payoff 0, as opposed to h , for a duration of the reaction lag, $\#(r)$). Since the latter happens with probability Δ/r , the net change in her payoff is first-order approximated⁵ as $\Delta(1 - \ell) - \Delta\#(r)h = r$, which is 0 because $\#(r) = \frac{1-\ell}{h}r$. Hence, any type lower than $\#(r)$ strictly prefers to wait at t as explained above. Additionally, it is a straightforward calculation to show that it is optimal for a player of type $\#(r)$ to defect at $t = 1 - r$ (outlined in the previous footnote). This verifies the optimality of the defection phase.

Cooperation phase: Consider any point in time $t < \hat{T}$ in the cooperation phase after

⁵Precisely, it is $P(r, \theta_i) = \frac{\Delta}{r} \left[\frac{\Delta}{2} + \left(r - \frac{\Delta}{2} - \theta_i \right) \ell \right] + \frac{r-\Delta}{r} \left[\frac{r-\Delta}{2} + \frac{(r-\Delta)(1-\ell)}{2h} h + \left(r - \frac{r-\Delta}{2} - \frac{(r-\Delta)(1-\ell)}{2h} \right) \ell \right]$ where θ_i is player i 's type. Its first and second derivatives wrt θ_i are calculated to be, respectively, [see cntnsPD-1.nb] $\frac{(r-\Delta)(1-\ell)-\theta_i h}{hr} \ell$ and $\frac{(\ell-1)\ell}{hr} < 0$. Note that the first derivative decreases in θ_i , and is 0 when $\theta_i = 0$ and $\theta_i = r(1-\ell)/h = \vartheta(r)$, which implies that the continuation payoff increases as θ_i increases from 0 for all $\theta_i < \vartheta(r)$, so such types should not switch to D at $t = 1 - r$. For player i of type $\vartheta(r)$, since $P(r, \vartheta(r))$ is strictly concave in $\vartheta(r)$, waiting until $\vartheta(r)$ is worse than defecting at $t = 1 - r$ so long as $r - \vartheta(r)/2 - \vartheta(r) > 0$. When $r - \vartheta(r)/2 - \vartheta(r) \leq 0$, the first term of $P(r, \vartheta(r))$ become $(\vartheta(r)/2)$ and the derivative of $P(r, \vartheta(r))$ wrt $\vartheta(r)$ is calculated to be $\ell(1-h-\ell)(r-\vartheta(r))/(hr) < 0$.

both players have switched to C . It is straightforward to verify that it is suboptimal for either player to switch to D before \hat{T} because she would be better off by doing so at \hat{T} instead, which would prolong the duration of $(C; C)$ by $\hat{T} - t^\theta$ and curtail the duration of $(D; D)$ by the same length later.

Transition phase to cooperation: Moving backward one step, consider player j after player i has switched to C at $t < T$ in the WoA phase: It is clearly optimal for j to switch to C at $t + \epsilon$ because otherwise i will switch back to D as soon as possible after $t + \epsilon$ (again, this point in time is not pinned down—see the next section) and keep to it until the game ends, which is worse. (Indeed, not switching at $t + \epsilon$ is dominated by switching at $t + \epsilon$ and then defecting at \hat{T} .) Given this, it is optimal for player i to keep to C until $t + \epsilon$ because her switching back to D before $t + \epsilon$ is worse by the same reasoning.

Consider the other case that neither player switched to C until T . Then, by switching to C at T^θ , the players get into cooperation phase until \hat{T} , followed by the defection phase. The expected payoff from this is easily calculated to be larger than that for a player who behaves differently, because then the off-equilibrium strategy prescribes that both players would switch to D as soon as possible (unless they were already playing D) and keep to D until the end of the game.

WoA phase: The game starts with both players playing D at $t = 0$. Consider either player, say i , at $t < T$ in the contingency that neither player has switched to C by then. Relative to the option of switching to C right now (i.e., at t), let's compare that of waiting for an infinitesimal amount of time Δ and switching to C at $t + \Delta \leq T$ (provided that the other player has not switched in between). If the other player, j , switches in between, which happens with probability $1 - e^{-q\Delta}$, player i would enjoy a payoff of h until she follows suit herself later, whereas she would get a payoff of 0 for the same duration if she switched to C at t . If player j did not switch in between, which happens with probability $e^{-q\Delta}$, the wait simply delay the transition to the cooperation phase: this prolongs the initial state of $(D; D)$ by Δ and curtails the duration of $(C; C)$ by the same amount of time in the cooperation phase, costing $1 - \ell$ both players for a duration of Δ . Therefore, the expected net gain from waiting is first-order approximated⁶ by $(1 - e^{-q\Delta})h - e^{-q\Delta}\Delta(1 - \ell)$ and its derivative wrt Δ is $qh - (1 - \ell)$ when evaluated at $\Delta = 0$. By setting $q = \frac{1-\ell}{\theta h}$, therefore, we ensure that both players are indifferent between switching to C at any point $t < T$ and waiting until a later time $t^\theta \leq T$, establishing the optimality of the strategy in the WoA phase.


It remains to show that a player, i , is indifferent between switching to C at T and not switching at any $t \leq T$. Since the two options are identical if the other player, j , switches to C before T , consider the case that player j does not switch until T . (The contingency that player j switches exactly at T is a measure 0 event, so is left out in calculation.) By switching at T , player i receives a payoff of $\hat{T} - T - \epsilon$ (calculated from a payout of 0 from T to $T + \epsilon$, then 1 until \hat{T}) before the defection stage starts. By not switching at T , she

⁶A full expression is $\int_0^\Delta qe^{-q\tau}(\tau\ell + \theta h + (\hat{T} - t - \tau - \theta) + V_{\hat{T}})d\tau + e^{-q\Delta}(\ell + (\hat{T} - t - \epsilon - \theta) + V_{\hat{T}}) = \frac{e^{-q\Delta}}{q}(1 - \ell - \theta qh - e^{q\Delta}(1 - \ell - \theta q(h - 1) - q(\hat{T} - t)))$ where $V_{\hat{T}}$ denotes the continuation payoff from \hat{T} . Its derivative with respect to Δ is $-e^{-q\Delta}(1 - \ell - \theta qh)$ which is identically 0 when $q = \frac{1-\ell}{\theta h}$.

receives the same payoff of $\hat{T} - T - \bar{\ell}$ (calculated from a payout of $\bar{\ell}$ from T to $T^\theta = T + \frac{\bar{\theta}}{1-\ell}$, then 1 until \hat{T}) before the defection stage starts, establishing the desired indifference.

D. Discussion of equilibrium

This equilibrium portrays a picture very close to the experimental results reported in Friedman and Oprea (2012, AER) who estimate the median reaction lag at 0.6 second. We simulated the cooperation rate of our equilibrium at every $t \in (0;1)$ for $\bar{\ell} = .015$ min or 0.9 second, which is illustrated in the first graph below in parallel with the corresponding graph of the experimental results. The parameter values used for our equilibrium, $h = 1.4$ and $\bar{\ell} = .4$, correspond to the “Easy” case of the experiment.⁷ The theoretical prediction is very close to the experimental results for the “Mix-a” case as well ($h = 1.8$; $\bar{\ell} = .4$), although it under-predicts cooperation in the “Mix-b” and “Hard” cases.



Random $\bar{\ell}_i$: We note that the equilibrium strategy in the defection phase is straightforwardly extended to the cases that the reaction time of a $\bar{\ell}_i$ -type player is a random variable

⁷The graph from experiment is the median cooperation rate with initial states randomly assigned to all four possible action pairs, which biases the initial cooperation forward a little. But, our graph is for the first switch to C , which also biases forward a little.

with a mean τ_i , as long as its support is narrow enough to ensure that a reaction of a τ_i -type player, upon observing a defection by the opponent at $t < 1 - r(\tau_i)$, is realized before the end of the game. This is because the analysis only relies on the expected payoffs which remain the same in these variants of the model. Consequently, the equilibrium described in the previous section extends straightforwardly to this case as well, with the flow rate of switching in the WoA phase modified accordingly (slower because the time lag between the first and second switches to C will be longer to ensure that the slowest type can make it for sure).

Flickering/blitzing: It is conceivable that once in the cooperation phase, a player with a very short reaction lag (i.e., very low τ_i and thus, low τ_i) would switch back and forth as quickly as she can, so as to benefit from her opponent not being able to respond in case her reaction lag is longer than τ_i . This is a byproduct of extending the standard practice of modelling a reaction lag in continuous-time settings, to allow for a seemingly realistic feature that players may differ in the length of their reaction lag. Consequently, the possibility that a player may block her opponent from responding by “flickering” appears to be an artificial, rather than realistic, concern. One can resolve this issue, for instance, by introducing “exhaustion” to the effect that if a player switches repeatedly within an interval less than $\bar{\tau}$, then she bears a cost for each switch that increases sufficiently rapidly with the number of repetition; or conversely, a player can reduce her reaction time below τ_i at a reasonable cost. Both costs do not seem unrealistic. However, we take a different approach described below.⁸

Punishing the opponent when she does not follow suit to cooperation: Consider the off-equilibrium contingency that one player, say i , switched to C at t in the WoA phase, yet the other player did not follow suit at $t + \bar{\tau}$. Then, player i wishes to switch back to D at the earliest time after t , but such a time does not exist because the real numbers are not well-ordered and consequently, player i 's equilibrium strategy in this continuation game is not determined in this manner. One can go around this technical issue by postulating a more involved continuation equilibrium. Specifically, player i switches back to D at $t + \bar{\tau} + \epsilon$ for a small $\epsilon > 0$, then both players switch to C at some later time before \hat{T} , say T^0 , followed by the cooperation and defection phases of the equilibrium-path. If player i switches back to D either before or after $t + \bar{\tau}$, then the continuation would be for both players to keep to D until the end. Provided that T^0 is sufficiently late, player $j \neq i$ prefers following the equilibrium-path to invoking this continuation game by failing to follow suit at $t + \bar{\tau}$. However, we take an alternative approach as explained below.

Reaction lag as the time to absorb updated situations: Although it serves the purpose, the off-equilibrium strategy above is devised to circumvent a purely technical issue inherent in continuous-time analysis. As a matter of fact, when the anticipated act of following suit

⁸Another possibility is to introduce “rm moves” which each player may commit to a certain action with certainty at a predetermined action time at least θ_i prior to that time, regardless of what happens in between. (This does not include probabilistic moves such as in the WoA phase, so does not include closed-loop strategies.) However, this seems to create another issue in the defection stage because players would want to commit to defect at a certain time to reduce reaction time when pre-empted. It is not clear that our equilibrium can be modified to accommodate this.

does not take place, it seems plausible that the player who was anticipating it would find herself in an unexpected situation and, therefore, she would need a little time to absorb the updated situation before deciding how react. This line of reasoning may be extended a little further: if a player anticipates a move from her opponent with a positive probability at a certain future point in time and intends to react differently depending on whether the anticipated move indeed takes place or not, then upon reaching that future point in time, she will need a little time to absorb which of the two possible contingencies has actually realized before reacting, regardless of whether the anticipated move indeed took place at time t or not. Interpreting reaction lag in this way as a small lapse of time needed to absorb “updates” of the relevant situation, one is led to apply the reaction lag defined earlier in [A] when the opponent fails to take an anticipated move, as well as when the opponent takes a move (expectedly or unexpectedly).

This is the approach adopted in the next section to develop a formal model and formulate players’ strategies. Note, however, that whether a player anticipates a move by her opponent at a certain time or not is determined by the other player’s strategy. To accommodate this interdependence, first we formulate “new plans” as building blocks of an agent’s strategy, where a new plan (i) starts either when the opponent makes a move or when the previous plan expires, and (ii) prescribes probabilities that the agent will have switched to the other action by every future points in time until its expiration time (which can be the end of the game). Then, we require that an equilibrium strategy profile be “internally consistent” in the sense that whenever a plan has its expiration before the end of the game, it is the first instance of time at which a move is anticipated by the opponent with a positive probability.

Additionally, we make the modelling choice that requires every new plan to prevail at least for a duration equal to the player’s reaction lag, so as to ensure that she can react to the updated situation uninterrupted by the opponent’s moves. This modelling choice is for the sake of resolving the aforementioned issue of flickering, and is not needed if either cost discussed there (on quick/frequent switches) is introduced. We make this choice so that our formulation of strategies resolves both of the technical issues explained above, namely, flickering and the nonexistence of the optimal time to switch back to D when the opponent fails to follow suit.

3 A Formal Model

The game starts with $(D; D)$ at $t = 0$. For $t \in (0; 1)$, a history $h^t = (h_1^t; h_2^t)$ consists of two sequences of switches/moves, one for each player: $h_i^t = (\tau_1^i; \tau_2^i; \dots; \tau_n^i)$ is a list of sequential times at which player i switched her action, where $0 < \tau_1^i < \dots < \tau_n^i \leq t$. Hence, the odd-number-subscripted times are when the player switched from D to C and the even-number-subscripted times are when the player switched from C to D . Let $h_i^0 = \emptyset$; $i = 1; 2$, for completeness. A history $h^{t'}$ is an *extension* of h^t if $t < t'$ and the two histories are identical up to and including time t ; and $h^{t'}$ is a *simple extension* of h^t if, in addition, neither player switched after t according to $h^{t'}$. Given a history, a “jump” refers to an instance of either one player switching or both players switching at the same time. Histories need to satisfy

certain properties to be consistent with admissible strategies to be formalized below.

A *plan at h^t* of player i of type τ_i , “agent τ_i ” for short, is a right-continuous non-decreasing function $p_i(\cdot|h^t; \tau_i) : (t; T_i(h^t; \tau_i)] \rightarrow [0; 1]$ where $p_i(t^+|h^t; \tau_i) := \lim_{s \nearrow t} p_i(s|h^t; \tau_i) = 0$ and $T_i(h^t; \tau_i) \in (t; 1]$, called the “expiration-time”, is the point in time until which the agent τ_i considers this plan to be valid. Abusing notation, we use $p_i(\cdot|h^t; \tau_i)$ to mean a plan, i.e., as a shorthand for $(p_i(\cdot|h^t; \tau_i); T_i(h^t; \tau_i))$. The value $p_i(s|h^t; \tau_i)$ denotes the probability that, conditional on the history h^t has been reached, the agent τ_i will have switched by $s \in (t; T_i(h^t; \tau_i)]$ without previous switches by either player after h^t .

An *agent-strategy* of player i of type τ_i , denoted by $\sigma_i(\tau_i)$, is a collection of plans $(p_i(\cdot|h^t; \tau_i); T_i(h^t; \tau_i))$, one for every possible history h^t where $t \in [0; 1]$, together with a subset $\overset{N}{\sigma}_i(\tau_i) \subset \sigma_i(\tau_i)$ of “new plans”, with the following properties:

- [C1] The initial plan, $p_i(\cdot|h^0; \tau_i)$, is in this subset, i.e., $p_i(\cdot|h^0; \tau_i) \in \overset{N}{\sigma}_i(\tau_i)$.
- [C2] If $p_i(\cdot|h^t; \tau_i) \in \overset{N}{\sigma}_i(\tau_i)$ and $T_i(h^t; \tau_i) < 1$, then $p_i(\cdot|h^{T_i(h^t, \theta_i)}; \tau_i) \in \overset{N}{\sigma}_i(\tau_i)$ when $h^{T_i(h^t, \theta_i)}$ is an extension of h^t with no jumps during $(t; T_i(h^t; \tau_i))$, although there may be a jump at $T_i(h^t; \tau_i)$. In this case, we say $p_i(\cdot|h^{T_i(h^t, \theta_i)}; \tau_i)$ is a “new plan due to expiration.”
- [C3] If $p_i(\cdot|h^t; \tau_i) \in \overset{N}{\sigma}_i(\tau_i)$ and either (i) it is a new plan due to expiration, or (ii) h^t contains a move by player $-i$ during $(t - \tau_i; t]$,⁹ then

$$(2) \quad \min\{1; t + \tau_i\} \leq T_i(h^t; \tau_i) \text{ and } p_i(s|h^t; \tau_i) = 0; \forall s \in (t; t + \tau_i) \cap (0; 1):$$

Furthermore, $p_i(\cdot|h^{t'}; \tau_i) \notin \overset{N}{\sigma}_i(\tau_i)$ if $h^{t'}$ is an extension of h^t and $t' < t + \tau_i$. In this case, we say $p_i(\cdot|h^{t'}; \tau_i)$ is a “new plan due to situation update.”

- [C4] If $p_i(\cdot|h^t; \tau_i) \in \overset{N}{\sigma}_i(\tau_i)$ and neither (i) nor (ii) above holds where $t > 0$, then h^t contains a move by player i and $p_i(s|h^t; \tau_i) = 0 \forall s \in (t; t + \tau_i) \cap (t; T_i(h^t; \tau_i))$. In this case, we say $p_i(\cdot|h^t; \tau_i)$ is a “new plan solely due to own move.”
- [C5] If $p_i(\cdot|h^t; \tau_i) \in \overset{N}{\sigma}_i(\tau_i)$ and $h^{t'}$ is an extension of h^t with a first jump at $t' < T_i(h^t; \tau_i)$ since t , then $p_i(\cdot|h^{t'}; \tau_i) \in \overset{N}{\sigma}_i(\tau_i)$ unless $p_i(\cdot|h^t; \tau_i)$ is a new plan due to situation update and $t' < t + \tau_i$; in the latter case, the jump at t' is by player $-i$ and $p_i(\cdot|h^{t+\theta_i}; \tau_i) \in \overset{N}{\sigma}_i(\tau_i)$ where $h^{t+\theta_i}$ is any extension of $h^{t'}$.

The condition [C3] stipulates physical constraints of player i in responding to an updated situation namely, that *i*) it takes at least τ_i for her to absorb the updated situation and react to it, and that *ii*) each new plan lasts at least for a duration of τ_i if prompted by an updated situation, so that she has a chance to respond to that update of situation. The latter case of [C5] stipulates that if the other player takes another move before player i has a chance to respond to the updated situation, then yet another new plan starts at $t + \tau_i$ for player i (this respects the principle that any updated situation warrants a separate reaction lag for info processing and consequently, a new plan). These place some constraints on the

⁹This means that all “situation updates” that take place during $(t, t + \theta_i]$, i.e., during the reaction lag prompted by the previous situation update(s), are processed together during the same reaction lag $(t + \theta_i, t + 2\theta_i)$ and responded collectively at $t + 2\theta_i$.

histories to define $\pi_i(\cdot)$ on (these constraints are rather complex to fully describe), e.g., agent i 's move cannot be taken within τ_i of a move by player $-i$, unless it was preceded by a sequence of new plans that lasted only for τ_i each due to player $-i$'s move within the reaction lag of each new plan. These are common knowledge.

In addition, we require that an agent-strategy $\pi_i(\cdot)$ with $\pi_i^N(\cdot)$ satisfy “inter-temporal consistency” conditions:

- [C6] If $\pi_i(\cdot|h^t; \cdot) \in \pi_i(\cdot)$ and $h^{t'}$ is a simple extension of h^t such that $t^\theta < T_i(h^t; \cdot)$, then $\pi_i(\cdot|h^{t'}; \cdot) \in \pi_i(\cdot)$ has the same expiration-time as $\pi_i(\cdot|h^t; \cdot)$, i.e., $T_i(h^{t'}; \cdot) = T_i(h^t; \cdot)$, and $\pi_i(\cdot|h^{t'}; \cdot)$ is obtained from $\pi_i(\cdot|h^t; \cdot)$ via Bayesian updating:

$$(3) \quad \pi_i(s|h^{t'}; \cdot) = \frac{\pi_i(s|h^t; \cdot) - \pi_i(t^\theta|h^t; \cdot)}{1 - \pi_i(t^\theta|h^t; \cdot)} \quad \forall s \in [t^\theta; T_i(h^t; \cdot)]$$

provided that the denominator is non-zero. In this case, we say that $\pi_i(\cdot|h^{t'}; \cdot)$ is a “truncated(?) plan” of $\pi_i(\cdot|h^t; \cdot)$.

- [C7] If $\pi_i(\cdot|h^t; \cdot) \in \pi_i^N(\cdot)$ is a new plan due to situation update and $h^{t'}$ is a non-simple extension of h^t such that $t^\theta < t + \tau_i$, then $\pi_i(s|h^{t'}; \cdot) = \pi_i(s|h^t; \cdot) \quad \forall s \in (t^\theta; t + \tau_i]$ and $T_i(h^{t'}; \cdot) = t + \tau_i$. ($h^{t'}$ is illegitimate if any move after t was by player i .)

A *strategy* of player i , denoted by π_i , is a collection of agent-strategies, one for each type $\theta_i \in (0; \bar{\theta})$, i.e., $\pi_i \in \Sigma := \prod_{\theta \in (0, \bar{\theta})} \Sigma(\theta)$ where $\Sigma(\theta)$ is the set of all agent-strategies of a player i of type θ that satisfy [C1]–[C7]. Although we do not restrict the strategies to “legitimate” histories (relative to the given strategy), only such histories will be relevant in our analysis.

To determine the outcome from a strategy profile, we now formalize the set of all possible histories. Given a type profile $(\theta_1; \theta_2) \in (0; \bar{\theta})^2$, possible first jumps are represented by points in $J_1 := I^3$ where $I = (0; 1]$. For expositional ease, denote the first, second and third copy of I in J_1 by I_x , I_y and I_z , respectively. Then, the interpretation of points in J_1 is: a point $x \in I_x = (0; 1]$ is the contingency that only agent θ_1 switches (from D to C) at $t = x$ (without previous moves); a point $y \in I_y = (0; 1]$ is the contingency that only agent θ_2 switches (from D to C) at $t = y$ (without previous moves); and a point $z \in I_z = (0; 1]$ is the contingency that both agents θ_1 and θ_2 switch (from D to C) at $t = z$ (without previous moves). For $j_1 \in J_1$, let $|j_1| \in (0; 1]$ denote the time that the jump j_1 takes place, i.e., without reference to the player(s) who make it.

Inductively, for any admissible history $h^{j_n} = (j_1; j_2; \dots; j_n)$, the set of possible next jumps can be represented by $J_{n+1}(h^{j_n}) := I_x(h^{j_n}) \times I_y(h^{j_n}) \times I_z(h^{j_n}) \subset \mathcal{I} := I^3$ where each $I_{(\cdot)}(h^{j_n})$ is an interval of the form $[\tau; 1]$ where $\tau > |j_n|$. This process stops for some n bounded above because any two moves by the same player is at least τ_i apart. Thus, allowing illegitimate histories, we may define the set of “outcomes” as $\mathcal{O}(\theta_1; \theta_2) = (\bigcup_{n=1}^{M(\theta_1, \theta_2)} \mathcal{I}^n) \cup \{\emptyset\}$ for a type profile $(\theta_1; \theta_2)$, where $M(\theta_1; \theta_2)$ is the maximum number of jumps that may take place in any history (including off-equilibrium) according to any $(\theta_1(\theta_1); \theta_2(\theta_2))$, and \emptyset is the contingency that no player ever switched until the end of game. As $M(\theta_1; \theta_2) \rightarrow \infty$ as $\min\{\theta_1; \theta_2\} \rightarrow 0$, the set of outcomes across all type profiles

is $\mathcal{O} = (\bigcup_{n=1}^T \mathcal{I}^n) \cup \{\emptyset\}$. Note that a history and all simple extensions of it correspond to the same outcome in \mathcal{O} .

Consider a strategy profile (π_1, π_2) . We restrict attention to profiles that satisfy the following measurability conditions. For each player $i = 1, 2$, the initial plans $p_i(h^0; \pi_i)$, $\pi_i \in (0, 1)$, and other new plans that start at simple extensions of h^0 , define a non-decreasing right-continuous function $P_i(h^0)$, called player i 's plan at h^0 , that represents the cumulative probability that player i (of any type) will have switched without a previous switch by either player, i.e., $P_i(s|h^0)$ is the integral of $p_i(s|h^0; \pi_i)$ or relevant new plans that succeed it, relative to the prior distribution F on π_i . For any simple extension h^t of h^0 , a posterior belief, denoted by $F_i(h^t)$, on π_i is determined by Bayes rule if t is in the support of $P_i(h^0)$, or by arbitrary assignment otherwise. For any history $h^{j_1 j} = (j_1)$, a posterior belief, denoted by $F_i(h^{j_1 j})$, is determined in the same manner. Similarly, for any admissible history $h^{j_1 j_2 \dots j_n}$, player i 's plan at $h^{j_1 j_2 \dots j_n}$, $P_i(h^{j_1 j_2 \dots j_n})$, is determined analogously, and so is a posterior belief at any simple extension of it or at any extension of the form $h^{j_1 j_2 \dots j_n j_{n+1}}$.

We now evaluate the expected payoff the agent π_i from a strategy profile (π_1, π_2) . In this discussion, measurability issues arise at places. We proceed by presuming that necessary measurability conditions are satisfied, because they are indeed satisfied in the equilibrium I characterize.

Consider a strategy profile (π_1, π_2) and an agent π_i . Then, $P_i(h^0)$ and player i 's initial plan $p_i(h^0; \pi_i)$ and other new plans that start at simple extensions of h^0 , define for each player (more precisely, for player $-i$ and for agent π_i) a cumulative distribution function for her to have made a first jump by any point in time $t \in (0, 1]$ during the game. These two cdf's determine a cdf for each of I_x, I_y and I_z in J_1 (it is a discrete cdf for I_z), and thereby one Stiltjes measure each on I_x, I_y and I_z in J_1 (Kolmogorov-Fomin p.362). For each jump in the support of this measure, a posterior belief is determined for each player as mentioned above.

Consider a time point $t \in (0, 1)$ with no previous jumps, denoted by h^t as a history in H^t , at which a posterior belief $F_i(h^t)$ is determined as explained above. Then, from the plans starting from the history h^t and its simple extensions, Stiltjes measures can be obtained on $I_x \cap (t, 1], I_y \cap (t, 1]$ and $I_z \cap (t, 1]$ in J_1 analogously to above.

Inductively, consider any admissible history $h^{j_1 j_2 \dots j_n} = (j_1, j_2, \dots, j_n) \in \mathcal{I}^n$, together with associated posteriors $F_i(h^{j_1 j_2 \dots j_n})$ obtained above. Then, $P_i(h^{j_1 j_2 \dots j_n})$ and agent π_i 's plan at $h^{j_1 j_2 \dots j_n}$, together with other plans that start at a simple extension of $h^{j_1 j_2 \dots j_n}$, determine a cdf for each of I_x, I_y and I_z in the $(n+1)$ -th copy of \mathcal{I} (with its support contained in $[|j_n|, 1]$), and thereby one Stiltjes measure each on I_x, I_y and I_z the $(n+1)$ -th copy of \mathcal{I} . For any $h^{j_1 j_2 \dots j_n \tau}$ which is a simple extension of $h^{j_1 j_2 \dots j_n}$, Stiltjes measures can be obtained on $I_x \cap (|j_n| + \tau, 1], I_y \cap (|j_n| + \tau, 1]$ and $I_z \cap (|j_n| + \tau, 1]$ in the $(n+1)$ -th copy of \mathcal{I} in an analogous manner.

In the previous 5 paragraphs we determined, for each admissible history $h^t \in H^t$, a posterior belief $F_i(h^t)$ on π_i and a Stiltjes measure, denoted by $\mu_i(h^t)$ on the set of next jumps to be generated by player $-i$ and agent π_i according to (π_1, π_2) . Note that this can be done in the same manner when the prior on π_i was restricted to $[\varepsilon, \bar{\pi}]$ for any small $\varepsilon > 0$.

Suppose we did this with the prior on π_i restricted to $[\varepsilon, \bar{\pi}]$. Then, there is an upper

bound of the number of jumps that may occur in any admissible history. Let M_ϵ denote the tight upper bound. For any admissible history $h^{j_{M_\epsilon-1}j} = (j_1, j_2, \dots, j_{M_\epsilon-1})$, a continuation value can be calculated for each agent i by integrating the flow payouts over the time period $(j_{M_\epsilon-1}|; 1]$ relative the aforementioned Stiltjes measure (with the residual measure applied to the contingency that there is no further move until the end of the game), which is well-defined because the payout is a simple function over the period $(j_{M_\epsilon-1}|; 1]$. Similarly, for any simple extension $h^{j_{M_\epsilon-1}j+\tau}$ of $h^{j_{M_\epsilon-1}j}$, a continuation value of agent i can be calculated analogously. In this manner, one can calculate a continuation value of agent i for all history with $M_\epsilon - 1$ jumps.

Then, consider a history h^t with $M_\epsilon - 2$ jumps. Recall that a Stiljes measure is determined on the set of next jumps from this history to be generated by player $-i$ and agent i according to (π_1, π_2) , as explained above. By integrating the value of agent i at these possible next jumps obtained above (with the residual measure applied to the contingency that there is no further move until the end of the game), one can calculate the value of agent i at h^t . In this manner, one can calculate a continuation value of agent i at every history with $M_\epsilon - 2$ jumps (or more jumps). Continuing recursively, one can calculate the value of agent i at all histories.

Recall that these values of agent i have been calculated with the prior on π_i restricted to $[\pi_i^-]$. For each h^t , take the limit of agent i 's value at h^t this obtained as $\pi_i^- \rightarrow 0$. This limit must exist as the prior F on π_i is atomless. We define this limit as the value of agent i at h^t from the strategy profile (π_1, π_2) .

Insofar as continuation values can be calculated as above for all admissible strategy profiles, we can define perfect Bayesian equilibrium. However, this is not guaranteed because measurability needs to be checked at various stages for the integrations to be well-defined that determine continuation values, although failure of measurability would be very rare if at all.¹⁰ Since it is not crucial for the purpose of this article to fully characterize the environments in which the measurability holds for all relevant cases, we proceed by focusing on strategy profiles for which the measurability holds whenever needed to define and establish an equilibrium. This is sufficient for our purpose.

Consider a strategy profile (π_1, π_2) , together with a profile of posterior beliefs defined for every (admissible) history as explained above. We introduce an “internal-consistency condition”: for any plan $(p_i(\cdot|h^t; \pi_i); T_i(h^t; \pi_i))$ at history h^t of π_i , the expiration-time $T_i(h^t; \pi_i)$ is the first time point in $(t; 1]$ at which player $-i$ is expected to make a jump with a strictly positive probability as calculated from the plans in π_i starting from h^t and its simple extensions relative to the posterior $F_i(h^t)$, unless it is prior to the lapse of reaction lag from the latest start of a new plan due to situation update, in which case $T_i(h^t; \pi_i)$ is the end of the relevant reaction lag; if player $-i$ is not expected to make a jump with a strictly positive probability during $(t; 1]$, then $T_i(h^t; \pi_i) = 1$. A strategy profile (π_1, π_2) is said to be “compatible” if every plan satisfies this condition. In this case we also say that π_i is compatible with π_{-i} .

A compatible strategy profile (π_1, π_2) , together with a belief profile, is a *perfect Bayesian*

¹⁰A bounded real function is integrable if measurable (Kolmogorov and Fomin, p.298), and existence of non-measurable subsets of real numbers rely on the Axiom of Choice (Wikipedia).

equilibrium (PBE) if for any history h^t , (i) the posterior belief on the players' types are obtained from (π_1, π_2) via Bayes rule whenever possible, and (ii) $\sigma_i(\cdot)$ is optimal for player i of type θ_i in the continuation game among all agent-strategies that are compatible with σ_{-i} .

4 Equilibrium

Consider the following symmetric strategy for both players where $q = (1 - \pi) = (\pi h)$ and $\hat{T} < 1 - \pi$ will be determined later:

- a) Until some $T \in (0; \hat{T} - 2\pi]$, player i switches to C with a flow rate $q > 0$ provided there has been no previous move. That is, the initial plan is $p_i(t|h^0; \theta_i) = 1 - e^{-qt}$ and $T_i(h^0; \theta_i) = T$ for all $\theta_i \in (0; \pi)$ and $i = 1, 2$; and if h^t is a simple extension of h^0 and $s < T$, then $T_i(h^s; \theta_i) = T$ and $p_i(t|h^s; \theta_i) = 1 - e^{-q(t-s)}$ for $t \in (s; T)$.
- b) If player i switched to C at $s \leq T$, then player i 's new plan is to keep to C until the expiration-time $s + \pi$; and player j 's new plan is to keep D until the same expiration-time $s + \pi$ at which point to switch to C with certainty (provided that i did not move between s and $s + \pi$).
- c1) If one player switched to C at $s \leq T$ and the other at $s + \pi$, then both players adopt identical new plans (that start at $s + \pi$) according to which they keep C (so long as the other does) until the expiration-time \hat{T} .
- c2) If no switch has been made at every $t \leq T$, both players adopt identical new plans (that start at T) according to which they keep D (so long as the other does) until the expiration-time T^θ at which point both switch to C with certainty, where $T^\theta = T + \frac{\theta}{1-\theta} < \hat{T}$.
- c2') If both switch to C at T^θ without any previous switches, then both players adopt identical new plans (that starts at T^θ) according to which they keep to C (so long as the other does) until the expiration-time \hat{T} .
- d) If \hat{T} is reached through a history described in c1) or in c2'), then both players adopt identical, type-contingent new plans (that start at \hat{T}) according to which either player of type $\theta \in (0; \pi)$ keeps to C until the expiration-time $1 - \pi(\theta) > \hat{T}$ at which point she switches to D with certainty (provided that the other has not switched since \hat{T}), where $\pi(\theta)$ is a continuously increasing function to be determined below with $\pi(0) = 0$ and $\pi(\pi) = 1 - \hat{T}$.
- e) If player i switched to D at $s > \hat{T}$ as per d), then player i 's new plan (that starts at s) is to keep to D until the expiration-time 1, and player j 's new plan (that starts at s) is to switch to D at $s + \pi_j$.

- f) If player i switched to D at $s > \hat{T}$ as per d) and player j switched to D at a later time $t < 1$, then both players adopt identical new plans (that start at t) according to which both players keep to D until the expiration-time 1.

We now describe off-equilibrium strategy:

- g1) If player i switched to C at $s < T$ as per a) and $s + \bar{\tau}$ has been reached without j having followed suit, then player i 's new plan (that starts at $s + \bar{\tau}$) is to switch back to D at her expiration-time $s + \bar{\tau} + \tau_i$ with certainty, and player j 's new plan (that starts at $s + \bar{\tau}$) is to keep to D until her expiration-time 1.
- g2) If player i switched to C at $s < T$ as per a) and $s + \bar{\tau}$ has been reached without j having followed suit, and player i switched back to D at $s + \bar{\tau} + \tau_i$, then both players adopt identical new plans (that start at $s + \bar{\tau} + \tau_i$) according to which both keep to D until the expiration-time 1.
- g3) If history at any t is inconsistent with any of the contingencies described above, then both players' continuation plans dictate that they switch to D as soon as they can, unless they were already playing D , and then keep to D until the end.

Note that a)-g3) fully describe a strategy profile. We now verify that it constitutes a PBE. Given that D is a strictly dominant strategy of the stage game, it is straightforward that g1)–g3) describe, respectively, a PBE of the continuation game off-the-equilibrium path.

Focus on the equilibrium path. By the same token, once one player switches to D at $t > \hat{T}$, it is a continuation PBE for the other to switch to D as soon as possible, and for both to keep to D until the end as e) and f) prescribe. Next we verify optimality of the remaining part of the strategy profile in reverse time order.

Optimality of d): For notational ease, let $\#(r) = \tau^1(r)$ denote the type of either player that would switch back to D at $t = 1 - r$, i.e., when the remaining time is r , unless the other player switched back first. We proceed by postulating that $\#(r) < r \Leftrightarrow (\tau_i) > (\tau_j)$ for all $r \in (0; 1 - \hat{T})$, which we will confirm later for an open set of environments. Consider $t = 1 - r \in (\hat{T}; 1)$, conditional on the game has been played according to c1) and $(C; C)$ was played up to t (excluding t). By d), the posterior belief of either player's type at $t = 1 - r$ is represented by the cdf $F|_r(\cdot) = F(\cdot) = F(\#(r))$ on $(0; \#(r))$. Therefore, calculated at $t = 1 - r$, the expected continuation payoff of player i 's strategy of switching to D at $t + \Delta \in (1 - r; 1)$ unless j switches first, is

$$(4) \quad \int_{\theta_j=0}^{\theta(r-\Delta)} \left[\Delta + \tau_j h + (r - \Delta - \tau_j) \right] dF|_r(\tau_j) + \int_{\theta(r-\Delta)}^{\theta(r)} \left[r - (\tau_j) + \max\{(\tau_j) - \tau_i; 0\} \right] dF|_r(\tau_j):$$

The first term depicts that if player j 's type is below $\#(r - \Delta)$ so that she would not switch to D until $t + \Delta$, then player i receives a payout of 1 until she switches to D at $t + \Delta$, then player j follows suit τ_j later (during which time player i receives h), and then i receives τ_i from $(D; D)$ for the remaining period. Note that j follows suit before the game ends

because we postulated that $\theta_j > 0$. The second term captures that if player j 's type is above $\#(r - \Delta)$ so that player j switches to D at $1 - \theta_j < t + \Delta$, then player i receives 1 from $(C; C)$ during $(1 - r; 1 - \theta_j)$, then 0 for a duration of θ_i until i also switches to D in response to j 's switch, and then receives θ from $(D; D)$ for the remaining period.

From $F|_{r(\theta_j)} = F(\theta_j) = F(\#(r))$, the first term of (4) is

$$\frac{1}{F(\#(r))} \int_{\theta_j=0}^{\theta_j(r-\Delta)} \left[\Delta(1 - \theta) + \theta_j(h - \theta) + r\theta \right] f(\theta_j) d\theta_j$$

and thus, its first derivative wrt Δ is

$$\frac{1}{F(\#(r))} \left[(1 - \theta) \left(F(\#(r - \Delta)) - \Delta f(\#(r - \Delta)) \#'(r - \Delta) \right) - \left(\#(r - \Delta)(h - \theta) + r\theta \right) f(\#(r - \Delta)) \#'(r - \Delta) \right].$$

Similarly, for Δ small so that $\theta_j - \theta_i > 0$, the first derivative of the second term of (4) wrt Δ is

$$\frac{1}{F(\#(r))} \left[\Delta + (r - \Delta - \theta_i) \theta \right] f(\#(r - \Delta)) \#'(r - \Delta).$$

Summing the two up and rearranging, the first derivative of (4) wrt Δ is

$$(5) \quad \frac{1}{F(\#(r))} \left[(1 - \theta) F(\#(r - \Delta)) - (\#(r - \Delta)(h - \theta) + \theta_i \theta) f(\#(r - \Delta)) \#'(r - \Delta) \right].$$

Evaluating this at $\Delta = 0$, we get

$$(6) \quad \frac{1}{F(\#(r))} \left[(1 - \theta) F(\#(r)) - (\#(r)(h - \theta) + \theta_i \theta) f(\#(r)) \#'(r) \right].$$

Note that this strictly decreases in θ_i . Therefore, if (6) assumes 0 at $\theta_i = \#(r)$, i.e.,

$$(7) \quad (1 - \theta) F(\#(r)) = \#(r) h f(\#(r)) \#'(r);$$

then either player of type below $\#(r)$ would strictly prefer to wait at $t = 1 - r$. Consequently, the function $\#(\cdot)$ that solves the differential equation (7) is a candidate for equilibrium $\#$. The additional condition for it to satisfy is that (5) decreases in Δ when $\theta_i = \#(r)$. For example, this is satisfied under uniform prior, i.e., when $F(\cdot) = \cdot$: the solution to (7) is $\#(r) = r(1 - \theta) = h$ and therefore, the expression in the bracket of (5) when $\theta_i = \#(r)$ is calculated to be $-\Delta(1 - \theta)^2 \theta = -(h^2 \theta)$ as desired.

Actually, we also need to take care of large Δ such that $r - \Delta - \#(r) < 0$ so that $\max\{\theta_j - \theta_i; 0\} = 0$ for some θ_j in (4): since the argument above implies i) that a player of type $\#(r)$ finds it suboptimal to wait until $t + \Delta^\theta$ such that $r - \Delta^\theta - \#(r) = 0$, and ii) at $t + \Delta^\theta = 1 - (r - \Delta^\theta)$ a player of type $\#(r - \Delta^\theta)$ would strictly prefer not to wait further and, therefore, neither does a player of type $\#(r)$ (because the payoff from switching to D is the same for both types while that from waiting is lower for higher types who is slow to respond when the other player switches to D), it follows that a player of type $\#(r)$ finds it suboptimal to wait past $t + \Delta^\theta$, either.

In conclusion, if $\#(\cdot)$ solves the differential equation (7) and (5) decreases in Δ when $\#_i = \#(r)$, then d) is optimal when $\# = \#^{-1}$. We have seen that when F is uniform then $\#(\cdot) = \frac{h}{1-\ell} > \#$ as postulated. Thus, this postulation is verified at minimum for an open set of F including the uniform distribution.

Optimality of c1): If either player switches to D at some point $t \in (s + \bar{\cdot}; \hat{T})$, then compared with the strategy of switching at \hat{T} , it reduces the duration of $(C; C)$ by $\hat{T} - t$ and increases the duration of $(D; D)$ by the same length. Therefore, it is optimal for either player to keep C at least until \hat{T} so long the other player does the same.

Optimality of b): It is optimal for j to switch to C at $s + \bar{\cdot}$ because otherwise $(D; D)$ will prevail during $(s + \bar{\cdot}; 1)$ as per g1), which is worse. Given this, it is optimal for i to keep C until $s + \bar{\cdot}$ because if she switches back to D at $t < s + \bar{\cdot}$ then $(D; D)$ will prevail during $(t + \bar{\cdot}; 1)$ as per g3), which is worse.

Optimality of a): At any $t < T$, player i 's continuation payoff of switching to C at $t + \Delta < T$, unless j switched first, is (see cntnsPD-1.nb)

$$(8) \quad \int_0^\Delta q e^{-q\tau} (\bar{\cdot} + \bar{\cdot} h + (\hat{T} - t - \bar{\cdot} - \bar{\cdot}) + V_{\hat{T}}) d\bar{\cdot} + e^{-q\Delta} (\Delta \bar{\cdot} + (\hat{T} - t - \Delta - \bar{\cdot}) + V_{\hat{T}}) \\ = \frac{e^{-q\Delta}}{q} \left(1 - \bar{\cdot} - \bar{\cdot} q h - e^{q\Delta} (1 - \bar{\cdot} - \bar{\cdot} q (h - 1) - q(\hat{T} - t)) \right) + V_{\hat{T}}.$$

The first term of (8) is for the contingency that player j switches at $t + \bar{\cdot} < t + \Delta$, in which case i receives $\bar{\cdot}$ during $(t; t + \bar{\cdot})$, then h until $\bar{\cdot}$ later when she also switches to D , and then 1 from $(C; C)$ until \hat{T} , followed by the continuation payoff from \hat{T} , denoted by $V_{\hat{T}}$. Similarly, the second term is for the alternative contingency that player i switches first to D at $t + \Delta$. The derivative of (8) wrt Δ is

$$-e^{-q\Delta} (1 - \bar{\cdot} - \bar{\cdot} q h) = 0$$

for all Δ because $q = (1 - \bar{\cdot}) = (\bar{\cdot} h)$. Therefore, either player is indifferent between switching to D at $t < T$ and waiting until $t + \Delta < T$, justifying the strategy described in a).

Optimality of c2): The expected payoff of either player from switching to C at $t = T$ without previous moves, is $\hat{T} - T - \bar{\cdot}$ from evaluating (8) at $\Delta = 0$ and $t = T$. That from not switching at T but switching at T^θ along with the other player is $(T^\theta - T) \bar{\cdot} + (\hat{T} - T^\theta) + V_{\hat{T}}$. The two options are equivalent if $T^\theta = T + \frac{\bar{\theta}}{1-\ell} < \hat{T}$ (ensure the inequality holds). Switching at some $t = (T; T^\theta)$, or not switching at T^θ is worse because then $(D; D)$ would prevail for at least $(t + \bar{\cdot}; 1)$ and $(T^\theta + \bar{\cdot}; 1)$, respectively.

Optimality of c2'): The same argument applies as that given for optimality of c1) above.

This equilibrium can be extended to asymmetric players.

Suppose the two players are different: h_i and $\bar{\cdot}_i$ are the utility levels of player $i = 1; 2$ for relevant state. In addition, let F and f represent the distribution of player 1's type on

$(0; \bar{c}_1)$ and let G and g represent the distribution of player 2's type on $(0; \bar{c}_2)$. Let $\#_i(r)$ denote player i 's type who defects at $t = 1 - r$, for $i = 1, 2$.

Then, the marginal net gain of player 2 from infinitesimal wait before defecting, i.e., formula (6) in the symmetric case, becomes

$$(9) \quad \frac{1}{F(\#_1(r))} \left[(1 - \bar{c}_2)F(\#_1(r)) - (\#_1(r)(h_2 - \bar{c}_2) + \bar{c}_2 \bar{c}_2) f(\#_1(r)) \#_1'(r) \right]:$$

Thus, the equilibrium condition that this value is 0 at $\bar{c}_2 = \#_2(r)$ becomes

$$(10) \quad (1 - \bar{c}_2)F(\#_1(r)) = (\#_1(r)(h_2 - \bar{c}_2) + \#_2(r) \bar{c}_2) f(\#_1(r)) \#_1'(r):$$

The corresponding condition for player 1 is

$$(11) \quad (1 - \bar{c}_1)G(\#_2(r)) = (\#_2(r)(h_1 - \bar{c}_1) + \#_1(r) \bar{c}_1) g(\#_2(r)) \#_2'(r):$$

We need to find the solution to the system of differential equations (10) and (11).¹¹ Setting $a_i = (1 - \bar{c}_i) = \bar{c}_i$ and $b_i = (h_i - \bar{c}_i) = \bar{c}_i$ for $i = 1, 2$, the system of differential equations becomes:

$$(12) \quad a_2 F(\#_1(r)) = (b_2 \#_1(r) + \#_2(r)) f(\#_1(r)) \#_1'(r); \quad 0 < a_2 < b_2$$

$$(13) \quad a_1 G(\#_2(r)) = (b_1 \#_2(r) + \#_1(r)) g(\#_2(r)) \#_2'(r); \quad 0 < a_1 < b_1$$

The solution to this system, if exists, can be illustrated in a *phase plane*, treating $\#_1$ in the x -axis and $\#_2$ in the y -axis. For notational ease, set $\#_1(r) = x(t)$ and $\#_2 = y(t)$. Then, (12) and (13) become

$$(14) \quad \dot{x} = \frac{a_2}{(b_2 x + y)} \frac{F(x)}{f(x)} \quad \text{and} \quad \dot{y} = \frac{a_1}{(b_1 y + x)} \frac{G(y)}{g(y)}.$$

We may set

$$F(x) = (x)x \quad \text{and} \quad G(y) = (y)y:$$

Then, the system (14) becomes

$$(15) \quad \dot{x} = \frac{a_2 x}{(b_2 x + y)} \frac{(x)}{f(x)} \quad \text{and} \quad \dot{y} = \frac{a_1 y}{(b_1 y + x)} \frac{(y)}{g(y)}.$$

For any $y \geq 0$, the first equation determines a solution $x(t|y)$ and for any $x \geq 0$, the second equation determines a solution $y(t|x)$. These determine the gradient vectors at every point $(x; y) > 0$ in the first quadrant of the phase plane. We assume that F and G satisfy

$$(16) \quad \lim_{x \downarrow 0} \frac{(x)}{f(x)} > 0 \quad \text{and} \quad \lim_{y \downarrow 0} \frac{(y)}{g(y)} > 0:$$

¹¹These are *autonomous* DE's in the sense that the functions to solve for are in one variable only, i.e., time. I gather that these are relatively well-behaved.

This condition is widely satisfied so long as F and G do not vanish to an infinite order at 0,¹² including $F(x) = x^n$ for all $0 \leq n < \infty$.

Since $\dot{x} > 0$ and $\dot{y} > 0$ for all $(x, y) \gg 0$, starting from any initial point $(x_0, y_0) \gg 0$, the solution to (15) is unique and, as time goes back,

(i) will reach $(0, 0)$ within finite time, and

(ii) does not hit either axis until before reaching $(0, 0)$.

The conclusion (i) follows because \dot{x} and \dot{y} are bounded away from 0 by (16),¹³ and (ii) follows because if it hit either axis, then the solution starting from that point would be unique and move along that axis in either direction of time, contradicting the supposition that the axis were reached from outside of the axis.

Let $\bar{x}(t)$ and $\bar{y}(t)$ be the solution to (15) from the starting point (\bar{x}_1, \bar{x}_2) . Let $t_0 < 0$ be such that $\bar{x}(t_0) = \bar{y}(t_0) = 0$. Then, $\#_1(r) = \bar{x}(t_0 + r)$ and $\#_2(r) = \bar{y}(t_0 + r)$ constitute a unique solution to the system (12) and (13).

References

- Bergin, James, and W. Bentley MacLeod. 1993. "Continuous Time Repeated Games." *International Economic Review* **34**: 21–37.
- 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.
- Friedman, Daniel, and Ryan Oprea. 2012. "A Continuous Dilemma." *American Economic Review* **102**: 337–363.
- 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.
- Simon, Leo K. and Maxwell B. Stinchcombe. 1989. "Extensive Form Games in Continuous Time: Pure Strategies." *Econometrica* **57**: 1171–1214.

¹²An example that vanishes to an infinite order at 0 is $F(x) = e^{-1/x}$ if $x > 0$ and $F(0) = 0$. Then, $\xi(x)/f(x) = x \rightarrow 0$ as $x \rightarrow 0$.

¹³To be complete: If $y(t) \leq x(t)$ for some t and thus $y(t) = \gamma(t)x(t)$ with $\gamma(t) \leq 1$, then $ax/(bx + y) = a/(b + \gamma(t)) \geq a/(b + 1)$. So, this coefficient is bounded away from 0 and thus \dot{x} is bounded away from 0. Thus, as long as $y \leq x$, $x(t)$ keeps going down as t decreases at a rate bounded away from 0. It must get to 0 (or become smaller than $y(t)$ in a finite time). Similarly, if $x(t) \leq y(t)$, $y(t)$ keeps going down as t decreases. One could say that perhaps it crosses the line $y = x$ infinitely many times, thus never getting to $(0, 0)$, but by above it can spend only a finite amount of time on each side of the line $y = x$.