An Extensive-Form Representation of Continuous-Time Games with Reaction Lag*

In-Uck Park
University of Bristol
Sungkyunkwan University

Siyang Xiong
University of California
Riverside, USA

September 29, 2020

Abstract. *We propose a new formulation of continuous-time games that represent them as standard game trees. It is pillared by the following two core ideas: each player (i) chooses when and how to move next each time any player makes a move, and (ii) is subject to a player-specific "reaction lag" which is the minimal amount of time that has to elapse before their next move at each such decision point. In this framework, continuous-time games can be analyzed by the same familiar tool set that we use for standard extensive-form games. As an application, we analyze the prisoners' dilemma game in this framework and shed new light on recent observations in continuous-time experiments. (JEL Codes: C72, C73, D82)*
Keywords: continuous-time, inertia, reaction lag, prisoners' dilemma

1 Introduction

Some economic interactions are ideally modelled in continuous time, such as when their timing is critical relative to other players' moves (e.g., preemption games and stock trading). A number of recent continuous-time experiments on selected games document that subjects behave in line with theoretical predictions, which is qualitatively different from their behavior in discrete-time counterparts as explained below in more detail. These observations suggest a potential role of continuous-time analysis in economic investigations, but its use in economics has been limited hitherto.

The limited use is attributed, at least partly, to technical issues stemming from the conventional way of modelling players in continuous-time games as choosing an action at ev-

*We thank Jim Bergin, Hugo Hopenhayn, Johannes Hörner, Sergey Izmalkov, Philippe Jehiel, Bentley MacLeod, Eric Maskin, Stephen Morris, and Aleksey Zinger for helpful conversations at various stages of this research project, and participants at 2012 CEPR European Summer Symposium in Economic Theory on "*Communication and Information*" in Gerzensee, Switzerland, the 2nd Lancaster Game Theory Conference in Lancaster, U.K., and seminars at European University Institute and Paris School of Economics.

ery point in time, contingently on the full history until then. Due to the well-known property that real numbers are not well-ordered,[1] this modelling practice suffers from various technical issues (e.g., in defining players' strategies and/or the outcomes from strategy profiles—see below); moreover, it impairs the analytical power of inductive reasoning which is central to economic analysis of dynamic (extensive-form) games, because the immediately preceding or succeeding moves are typically ill-defined. For example, the natural and intuitive strategy of punishing any deviator immediately, known as the "grim-trigger" strategy, is not well-defined because the first instant of time to start punishment after a deviation is not pinned down in continuous time. These are non-trivial issues that get in the way of adapting the predominant paradigm of dynamic economic analysis to continuous-time environments.

Proposed in this paper is an alternative formulation of continuous-time games that we believe models players' dynamic action choices more realistically. Although time flows continuously, it is unrealistic to think that we incessantly make choices at every point in time. Rather, we tend to review the situation as new events unfold and plan and implement the next move at an appropriate future point in time, which may well be revised if the situation evolves yet again. In addition, we face the physical limitation of being unable to carry out the next move until some minimal time has elapsed since the previous move, referred to as their "reaction lag" in this paper, the exact length of which varies across people. In our formulation, therefore, players choose the time and action for their next move as history evolves subject to their reaction lag, rather than choosing an action at every instant of time.

We show that, modelled this way, continuous-time games are free from the technical issues mentioned above and representable as standard extensive-form games where immediately preceding or succeeding moves are clearly identified. As a result, they can be analyzed by the same logic and principles for extensive-form games, which have been thoroughly established and well-understood in economics. In addition, we may consider only behavior strategies for mixed strategies as Kuhn's Theorem applies to our framework (by Aumann (1964)). Our theory also adds new light to some recent experimental observations, as will be elaborated in due course.

To illustrate the issues inherent in continuous-time games, let us consider a prisoners' dilemma game played over an interval of time. The well-known grim-trigger strategy prescribes that a player, at any point in time, cooperates if the opponent has never defected

[1] A totally ordered set S is well-ordered if every non-empty subset of S has a least element.

before, but defects otherwise. In discrete-time versions of the prisoners' dilemma game, both agents adopting this strategy defines an outcome uniquely by induction, in which both cooperate in every period. In continuous time, however, the same strategy profile is consistent with both players cooperating until and *including* a certain time point t and both defecting at every point after t. As this holds true whenever t may be, it is consistent with a continuum of outcomes that differ in the switching time.[2] Such indeterminacy stems from the lack of an immediately preceding time point which will inductively determine the player's action in the next time point. It is clear from this illustration that some structures must be imposed for a continuous-time model to be functional.

Of course, we are not the first to impose structures to overcome such pathological issues. A prominent early study is Simon and Stinchcombe (1989), who construct a continuous-time framework as the limit of discrete-time models with an infinitely fine grid, in the sense that for any strategy profile every sequence of increasingly fine grids imposed on it defines the same limit outcome. The restrictions imposed on admissible strategies for this, however, are rather complicated to formalize and verify.

Bergin (1992) and Bergin and MacLeod (1993) impose "inertia" on strategies in the sense that a player's action prevailing at any point in time must continue for a positive span of time. Thus, agents play a game by switching from an action to another at distinct time points as history unfolds (rather than choosing an action at every instant of time). This resolves the indeterminacy of outcomes because then any strategy profile divides the time into countable intervals of fixed action profiles. Yet, it falls short of reinstating inductive strategies or allowing for extensive-form representation of continuous-time games because an "earliest next time point to act" is still ill-defined as the current action may be switched to another in the arbitrarily near future.[3] For instance, the grim-trigger strategy is not representable as a legitimate strategy with inertia, either. Building upon the notion of inertia, we address these issues by introducing individual-specific reaction lag, as outlined below.

Inertia and reaction lags are not only modelling devices to resolve a technical issue,

[2]This issue was first observed by Anderson (1984) in this context and discussed subsequently by Simon and Stinchcombe (1989), Bergin (1992) and Bergin and MacLeod (1993). We refer to those papers for more detailed explanations.

[3]To elaborate, the span of time for an action to prevail can be arbitrarily short depending on the time and history. Hence, an infinite sequence of moves may occur at time points converging to a limit before the end of game (see footnote 13). Such a sequence of moves leads neither to a strategic node with a well-defined preceding node nor a terminal node at which payoffs are specified, thus fails to be a part of a game tree.

but they also depict real-world phenomena. They may arise from various sources, such as physical limitation in implementing an action, time to process information or reach a decision through a hierarchical process (e.g., in case the player is a corporation), or some combination of them. As such, inertia is real but the exact extent to which it affects a player may vary depending on their identity and circumstances.

In light of the discussions above, we propose a new formulation of continuous-time games buttressed by the following two core ideas. One is to model each player as choosing when and how to move next each time any player makes a move. The other is to model the notion of inertia explicitly as a player-specific "reaction lag" which is the minimal span of time that has to elapse before their next move at each such decision point.

Then, at the initial node and every time there is a move, each player is in a well-defined information set where they choose the time and action of their next move subject to their reaction lag. Such a choice gets implemented at the chosen time, leading to the next information set for each player, unless pre-empted by the implementation of some other player's move chosen for an earlier time point, which leads to a different next information set for every player. As such, the game is represented as a standard extensive-form game (game tree) where a player's behavior strategy specifies time and action for the next move at every information set as described above. Hence, continuous-time games can be analyzed by the same tool set as that for extensive-form games. Not only is this model simpler and familiar to study than conventional models where a choice is made at each and every time point, we believe that it enhances realism as well (see Section 2).

Reaction lags applied to the players' own moves warrant enough inertia after each move to resolve the indeterminacy issue and well-define any finite-player and finite-action game. But, to have a well-defined notion of an earliest subsequent time to react (so that inductive strategies are defined unambiguously), reaction lags need be applied to opponents' moves as well. When the length of reaction lag differs across players, however, a technical issue arises on how to apply a new reaction lag if an opponent moves while a player is already in a reaction lag triggered by a previous move. We address this issue in Section 4 based on an interpretation of reaction lag as the time needed to deliberate on the best future course of action as the game evolves with players' moves.

Starting with the influential experimental findings by Friedman and Oprea (2012), a number of experiments document subjects' behavior in continuous time which is critically different from that in discrete time in various contexts, including social dilemma (Bigoni, *et al.,* 2015), market entry games (Calford and Oprea, 2017), and network formation (Choi, *et al.,* 2020). In particular, Friedman and Oprea (2012) conduct the first experiment on a continuous-time prisoners' dilemma game and report that typically the two players cooperate until near the end of the game when one player defects first and the other follows suit immediately. In a market entry game where the firm's profit increases with their entry time up to an ideal point but starts to decline when pre-empted, Calford and Oprea (2017) specifically test the effect of inertia and report entry delays close to the ideal time when inertia is small. Such behavior is in sharp contrast to the unique equilibrium in the discrete-time version of the respective game, namely, immediate defection/entry. The authors rationalize the observed continuous-time behavior based on the notion of ϵ-equilibrium (Radner, 1980).

Our theory also rationalizes the observed behavior: when the length of each player's reaction lag is their private information, the shorter is their reaction lag the longer is the player willing to wait before defecting/entering because they can cope better in case pre-empted. To our knowledge, this is the first theoretical rationalization as equilibrium of such behavior by *fully* rational players with a firm terminal time.[4] Moreover, this equilibrium offers a testable implication: the earlier is the first defection, the longer should be the time until the follow-on defection in the sense of first-order stochastic dominance. We find this relationship in the experimental data reported by Friedman and Oprea (2012) – see Section 5.

We regard our work as complementing the recent experimental studies that demonstrate the potential usefulness of continuous-time analysis, by providing an alternative explanation of the observed behavior as equilibrium in a theoretical framework that models inertia explicitly. We hope that this will foster further work, both theoretical and experimental, facilitating more fruitful continuous-time analysis in the future when our sense of time becomes more continuous thanks to the ever-improving information technology.

We liken our approach of formulating continuous-time games to that of Harsanyi (1967) who formalized a loosely-understood notion of uncertain payoffs of players at the time, into games of incomplete information via the introduction of payoff types so that the well-

[4] Cooperation in finitely repeated prisoners' dilemma games has previously been explained due to some departure from common knowledge of the terminal date (Neyman, 1999) or from full rationality (e.g., ϵ-equilibrium (Radner, 1980), bounds on the complexity of admissible strategies (Neyman, 1985)), or the presence of a committed type to a "tit-for-tat" strategy (Kreps, *et al.,* 1982). The last explanation entails an equilibrium with a complex "end game" (the continuation game after the first defection) which is not well in line with the experimental observations.

established analytic tool for imperfect-information games could be readily applied. In parallel, we formalize the notion of inertia concretely via the introduction of reaction-lag types, transforming technically cumbersome continuous-time games into discrete-move games representable by extensive-form games that we are so competent with.

Related Literature

The idea that time is needed for players to respond to preceding moves has been introduced in previous studies, notably in continuous-time bargaining models. With a view to unifying various bargaining protocols, Perry and Reny (1993) build a continuous-time game with reaction lags similar to ours and replicate various existing bargaining outcomes from different protocols by varying the ratio of the reaction lag to own offers relative to that to opponent's offers. However, introducing reaction lags to otherwise conventional strategies does not resolve the indeterminacy problem.[5] Hence, they additionally assume a condition equivalent to the inertia condition of Bergin and MacLeod (1993). As such, their reaction lags are not an indispensable ingredient of their analytic framework, but modelled as specific features of bargaining environments that influence the equilibrium outcome.

Motivated to explain the real-world phenomena (such as initial delays, agreements near the deadline and some failures thereof) not predicted by standard bargaining models, Ma and Manove (1993) study a continuous-time, alternating-offer bargaining model with a random reaction time after each offer. By virtue of the alternating-offer rule, the game is one in extensive form even without random reaction time, for which outcome is well-defined for every strategy profile. Once again, therefore, the random reaction time is introduced as a feature of bargaining environment that leads to the desired properties of the equilibrium. In contrast, reaction lags in our approach are an inherent constituent of analytical framework that reflects the reality as well as resolving technical problems.

Yuliy Sannikov achieved the latest breakthrough in continuous-time analysis by developing a new framework based on stochastic calculus methods that overcome tractability problems: the aforementioned technical problems of continuous-time games do not arise because history is monitored imperfectly via state variables (on which strategies are defined) that drift *continuously* and stochastically. As a result, studies of *stochastic* continuous-time games proliferated relatively recently, generating fresh insights in various settings including

[5] Their counterexample involves strategies of offering 0 (1, resp.) at every rational (irrational, resp.) $t > 0$ eligible for making an offer, which is infeasible in our framework where players choose when and how to move next.

repeated games with imperfect monitoring (Sannikov, 2007), optimal contract in principal-agent problems (Sannikov, 2008) and finance (DeMarzo and Sannikov, 2006). Yet, dynamic interactions in many contexts are still primarily analyzed by *deterministic* games with discrete moves. Our approach is an attempt to facilitate such customary, deterministic analyses in continuous-time environments as well.

Kamada and Kandori (2020) study "revision games" in which players publicly revise their action choices prior to a game is actually played, at opportunities that arrive according to a Poisson process. Although they focus on a different role (how much cooperation may be sustained) of stochastic opportunities to change actions, modelling players' moves as such would warrant inertia if adapted to continuous-time games. It would be interesting to study how this alternative approach compares relative to ours, which we leave as future work.

The paper is organized as follows. Section 2 provides a conceptual illustration of our approach with a simple example. Section 3 lays down a simple benchmark model with minimal reaction lag constraints that highlights the link to the conventional approach, which is further developed in Section 4 to fully-fledged models allowing for incomplete information and mixed strategies. Section 5 applies our framework to the prisoners' dilemma game and sheds new light on recent experimental results. Section 6 concludes and Appendix contains deferred proofs.

2 Illustration of a continuous-time game as a game tree

Section 2.1 reviews the conventional formulation of a continuous-time game. Although it may be the most intuitive way to define a continuous-time game, it appears to be overly convoluted to reflect real-life decision making and technically overloaded. Section 2.2 illustrates with a simple example an alternative way of modelling real-life decision making processes, that admits representation by a standard game tree which is impossible for conventional continuous-time games.

2.1 The conventional continuous-time games

Consider a set I of agents who choose an action from their respective action set $A_i,\ i \in I$, at every instant t in a unit time interval $[0, 1]$, upon observing all players' past action choices prior to t. Formally, a history at time $t \in [0,1]$ is a function h^t from $[0,t)$ to $A \equiv \Pi_{i\in I} A_i$, that is, $h^t \in A^{[0,t)}$; and a strategy of agent i is a measurable function defined on the set of all histories:[6]

$$x_i : \cup_{t\in[0,1]} A^{[0,t)} \longrightarrow A_i. \tag{1}$$

Let $X_i^\circ \equiv (A_i)^{\cup_{t\in[0,1]} A^{[0,t)}}$ denote the set of all (pure) strategies of agent i. Then, a continuous-time game form is described by a tuple $\langle I, (A_i)_{i\in I}, (X_i^\circ)_{i\in I}\rangle$. Provided that an "outcome" (a full record of action profiles $o : [0,1] \to A$) is well-defined for every strategy profile $x \in X^\circ \equiv \Pi_{i\in I} X_i^\circ$, one describes a continuous-time game by additionally specifying a measurable utility function from the outcome space to $\mathbb{R}^I$.

Although (1) is the most straightforward extension of extensive-form strategies as full contingent plans to continuous-time settings, it creates various technical problems including indeterminate outcomes, as discussed in the Introduction. To overcome such problems, specific restrictions are imposed on X_i° to define admissible strategies by such authors as Simon and Stinchcombe (1989) and Bergin and MacLeod (1993).

The main insight of Bergin and MacLeod (1993) is that, if an inertia condition is imposed on X_i°, a continuous-time game described above is well-defined in the sense that an outcome is uniquely determined by every strategy profile (that satisfies inertia). Throughout the paper, we impose the inertia condition which is described in detail in Section 3.3.

2.2 An illustration of real-life decision making in continuous time

The conventional continuous-time strategy defined above, albeit conceptually intuitive, does not seem to reflect our real-life decision making aptly. To see this, suppose that two players play a game over a time interval $[0,1]$ by choosing an action from $A_i = \{C, D\}$.

To facilitate illustration, let us consider a simple version of this game such that after the initial choice of actions at $t = 0$, each player may switch actions/move at most once during the game. We further focus on the case that the game starts with initial action choices (C, C) at time $t = 0$. Then, according to (1), a continuation strategy of a player, say 2, specifies an action $x_2(h^t) \in \{C, D\}$ contingently on every history h^t where $t \in (0, 1]$. The set of legitimate histories (in the sense that each player may switch actions at most once)

[6] For expositional ease, we use $A^{[0,0)}$ to denote the null history at time $t = 0$.

are divided into

$$
\begin{aligned}
H^O &= \cup_{0<t\le 1}\{h^t \mid h^t(s) = (C,C)\ \forall s < t\},\\
H^1 &= \cup_{0<t\le 1}\big(\cup_{0<t'<t}\ \{h^t \mid h^t(s) = (C,C)\ \forall s < t',\ h^t(s) = (D,C)\ \forall s \in [t',t)\}\big),\\
H^2 &= \cup_{0<t\le 1}\big(\cup_{0<t'<t}\ \{h^t \mid h^t(s) = (C,C)\ \forall s < t',\ h^t(s) = (C,D)\ \forall s \in [t',t)\}\big), \quad \text{and}\\
H^{12} &= \cup_{0<t\le 1}\big(\cup_{0<t'<t}\ \{h^t \mid h^t|_{t'} \in H^O \cup H^1 \cup H^2,\ h^t(s) = (D,D)\ \forall s \in [t',t)\}\big)
\end{aligned}
$$

where $h^t|_{t'}$ denotes h^t truncated at $t' \le t$. In words, H^O is the set of histories with no switches; H^1 those in which only player 1 switched; H^2 those in which only player 2 switched; and H^{12} those in which both switched.

Thus, for example, a strategy for player 2 to switch to D at $t=0.5$ if player 1 has not switched by then, but to not switch until the end otherwise, is described as[7]

$$
x_2(h^t) = \begin{cases} C & \forall h^t \ \text{ where } \ t < 0.5 \ \text{ and } \ h^t \in H^O \cup H^1 \\ D & \forall h^t \ \text{ where } \ t \ge 0.5, \ \ h^t|_{0.5} \in H^O, \ \text{ and } \ h^t \in H^2 \cup H^{12} \ \text{ if } t > 0.5 \\ C & \forall h^t \in H^1 \ \text{ where } \ t \in [0.5, 1) \ \text{ and } \ h^t|_{0.5} \in H^1 \\ D & \forall h^t \in H^1 \ \text{ where } \ t = 1 \ \text{ and } \ h^t|_{0.5} \in H^1. \end{cases} \tag{2}
$$

This is an overly cumbersome and convoluted way to describe such a simple strategy, due to a conventional strategy requiring that players make a choice at every point in time.

A more natural way for each player i to play the game is: (i) upon observing the initial action choices (C,C) at $t = 0$, decides on a time, say $s_i(0) \le 1$, to switch if the other player doesn't switch until then; (ii) otherwise, i.e., if the other player switches before the planned time to switch, say at $t < s_i(0)$, revises the time to switch as desired, say to $s_i(t) \in (t, 1]$. Then, each player i's (continuation) strategy is represented by a much simpler function

$$
s_i : [0,1) \to (0,1] \ \text{ such that } \ s_i(t) > t.
$$

The example strategy (2) is represented as $s_2(0) = 0.5$ and $s_2(t) = 1$ for $t \in (0, 0.5)$.

With strategies described in this way, the game can be represented by a game tree as in Figure 1.[8] Moreover, any strategy profile (s_1, s_2) determines an outcome straightforwardly

[7] We omit here the off-equilibrium contingencies arising from player 2's own deviation.

[8] After initial choices at $t = 0$, the ensuing subgame is drawn for the initial choice of (C,C) but is omitted for (C,D), (D,C) and (D,D) for visual clarity.

as follows: if $s_1(0) < s_2(0)$, the action profile at t is

$$\begin{cases} (C,C) & \forall t < s_1(0) = \min\{s_1(0), s_2(0)\} \\ (D,C) & \forall s_1(0) \leq t < s_2(s_1(0)) \\ (D,D) & \forall s_2(s_1(0)) \leq t \leq 1 \quad ; \end{cases}$$

and analogously for the cases that $s_1(0) > s_2(0)$ and $s_1(0) = s_2(0)$.

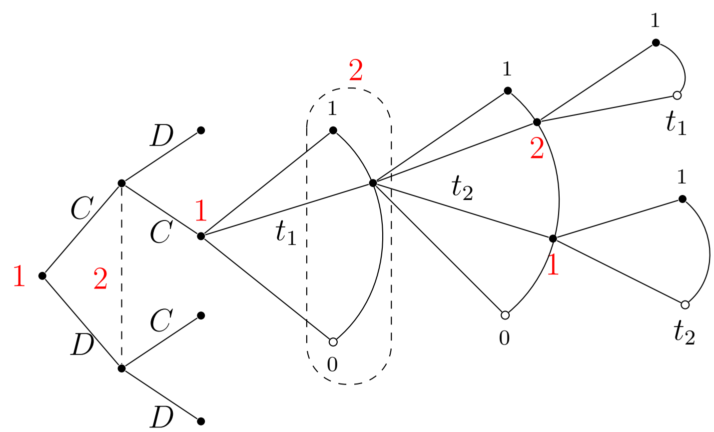

Figure 1: Game Tree

Let us consider an even simpler version of the continuation game after the initial choices (C, C) at $t = 0$, in which no further moves are permitted once there is a switch of actions by either player (possibly by both players simultaneously). Then, each player's strategy is simply a time $s_i(0) \in (0, 1]$ to switch, and the player who chooses an earlier time gets to switch at that time. This corresponds to the practice, known as the "open-loop" approach, of modelling future action choices as committed once and for all at the beginning of the game, as opposed to the "closed-loop" approach of modelling them to be determined contingently on how history evolves. The open-loop approach is simple but fails sequential rationality, yet is innocuous if only the first move matters, e.g., as in the war of attrition. Our approach illustrated above is closed-loop because the players revise their future action plans as history evolves, but retains the simplicity of open-loop approach by modelling them as picking the time of their next move every time history evolves with a move (and only the first next move matters as it updates history), which allows for representation by a game tree.

This is in a sharp contrast to conventional continuous-time games which cannot be represented by standard game trees. In a game tree, a history, or equivalently, a node has the following two defining features:

- every non-initial node has a unique immediately-preceding node, and
- every choice at a non-terminal node leads to a unique immediately-succeeding node.

Neither of these features is shared by conventional continuous-time games. Being a full record of action profiles on a half-closed-half-open interval $[0, t)$ for some t, a history in the conventional framework has no immediately-preceding history. Agents' choices at time t leads to a full record of action profiles on the closed interval $[0, t]$, which is not a legitimate history based on which to make subsequent moves in the conventional framework.

We illustrated above an alternative approach to model continuous-time games with a simple example, that admits a representation by a game tree which reflects real-life decision making sensibly for the given situation. We extend this approach to obtain a framework that represents as game trees a wide range of situations where potentially an arbitrarily large number of moves could be made by players. An advantage of our approach is obvious for this class of continuous-time games: the thoroughly established theory of extensive-form games can be readily applied to properly formulate and analyze such games. It is our hope that this approach will provide an analytic framework that can fruitfully accommodate various situations of time-sensitive dynamic interactions of economic interest.

3 A new approach – a complete-information benchmark with uniform inertia

We now formalize a new model of continuous-time games in the manner sketched in Section 2, in two stages. In this section, we introduce a simple benchmark version of the model in complete information and clarify the link to the conventional approach by showing that it is "equivalent" to the model of Bergin and MacLeod (1993) when their inertia is modified to have a uniform lower bound within each strategy (Proposition 1). In the next section, we develop it further to accommodate player-specific lower bounds on inertia as a private type and mixed strategies.

To facilitate exposition, we formulate a model for a two-player game played over a unit time interval $[0, 1]$ with a binary action set available for each player to choose from at each $t \in [0, 1]$. Thus, we let $I = \{1, 2\}$ denote the set of players and $A_i = \{a_i, b_i\}$ the set of available actions for player $i \in I$. However, it can be generalized straightforwardly to any finite-player game, with any finite action sets.[9]

Throughout the paper, for any topological space X, we consider the induced Borel sigma algebra, which we denote by $\mathcal{B}[X]$. For any metric space Y, we denote the set of probability distributions over Y by $\triangle(Y)$ and endow it with weak* topology. We endow Euclidean spaces with the Euclidean metric, and finite spaces with the discrete metric. We endow any product space with the product metric: for any two metric spaces, $\left(X, d^X\right)$ and $\left(Y, d^Y\right)$, the product space $X \times Y$ is endowed with the metric $d^{X\times Y}$ defined as

$$d^{X\times Y}\left[(x,y),(x',y')\right] = \max\left\{d^X(x,x'),\, d^Y(y,y')\right\}.$$

3.1 Histories

The two players take an action each of their choice at $t=0$, determining an *initial history* $h^0 \in A_1 \times A_2$, and may switch their actions at any time $t \in T \equiv (0,1]$ subject to some regularity conditions described below. A "move" of a player refers to either their initial choice of action at $t = 0$ or any subsequent change of actions. A "jump" refers to an instance of action switches taking place at one instant of time, described by a pair (t, I') where $t \in T$ is the time of the instance and $I' \in \mathcal{I} \equiv \{\{1\},\{2\},\{1,2\}\}$ is the set of players who switch actions at that time.

A *history,* denoted by h, is a sequence of jumps together with an initial history h^0. Thus, the set of all histories with exactly k jumps, denoted by H^k, is defined as

$$\begin{aligned}
H^0 &= A_1 \times A_2,\\
H^1 &= H^0 \times (T \times \mathcal{I}),\\
&\vdots\\
H^k &= \left\{\left[h^0, \left(t^1, I^1\right), ..., \left(t^k, I^k\right)\right] \in H^0 \times (T\times\mathcal{I})^k : 0 < t^1 < t^2 < ... < t^k \leq 1\right\},\\
&\vdots
\end{aligned}$$

where (t^j, I^j) denotes the j-th jump. Each H^k is a product metric space and we use d^k to denote the (product) metric on H^k. The space of all histories[10] is denoted by $\mathcal{H} \equiv \bigcup_{k=0}^{\infty} H^k$

[9]The possibility of changing action sets over time (contingently on history) can be accommodated straightforwardly as well, so long as they change subject to "inertia."

[10]It will be clear that only finite jumps occur due to uniform inertia.

with the metric d defined as

$$d(h,h') = \begin{cases} d^k(h,h') & \text{if } \{h,h'\} \subset H^k \text{ for some non-negative integer } k, \\ 1 & \text{otherwise.} \end{cases}$$

That is, any two histories of different "lengths" (i.e., numbers of jumps) are separated. Throughout the paper, we consider the Borel sigma algebra on $\mathcal{H}$ induced by the metric d.

Each history $h = \left[h^0, (t^1, I^1), ..., (t^k, I^k)\right]$ determines players' action profiles up to time t^k in the obvious manner. For example, under the history $\left[h^0 = (a_1, a_2), \left(\frac{1}{4}, \{1\}\right), \left(\frac{1}{3}, I\right)\right]$, the profile (a_1, a_2) prevails for $t \in [0, \frac{1}{4})$; at time $t = \frac{1}{4}$, player 1 switches to b_1 while player 2 stays put, hence the profile (b_1, a_2) prevails for $t \in [\frac{1}{4}, \frac{1}{3})$; at time $t = \frac{1}{3}$, player 1 switches to a_1, and player 2 switches to b_2. Defining a mapping $\psi : \mathcal{H} \to A_1 \times A_2$ that depicts the action profile induced by the last jump of a histiory (see Appendix A.1 for a formal definition), a complete description of players' action profiles up to time t^k for history $h = \left[h^0, (t^1, I^1), ..., (t^k, I^k)\right]$ is:

$$\begin{pmatrix} \text{the profile } h^0 \text{ prevails during time interval } [0, t^1); \\ \text{the profile } \psi\big([h^0, (t^1, I^1)]\big) \text{ prevails during } [t^1, t^2); \\ \vdots \\ \text{the profile } \psi\big([h^0, (t^1, I^1), ..., (t^{k-1}, I^{k-1})]\big) \text{ prevails during } [t^{k-1}, t^k); \\ \text{the profile at time } t^k \text{ is } \psi\big([h^0, (t^1, I^1), ..., (t^k, I^k)]\big). \end{pmatrix}$$

By $\tau_i(h)$ we denote the time of the last move of player i for history $h \in \mathcal{H}$:

$$\tau_i(h) = \begin{cases} 0 & \text{if } h = \left[h^0, (t^1, I^1), ..., (t^k, I^k)\right] \text{ and } i \notin I^1 \cup ... \cup I^k, \\ t^{k'} & \text{if } h = \left[h^0, (t^1, I^1), ..., (t^k, I^k)\right] \text{ and } i \in I^{k'}, \text{ and } i \notin I^{k'+1} \cup ... \cup I^k. \end{cases}$$

By $\tau(h)$ we denote the time of the last move/jump of history $h \in \mathcal{H}$:

$$\tau(h) \;=\; \max\{\tau_1(h), \tau_2(h)\}.$$

3.2 Strategies

A pure strategy of a player i consists of an initial action $s_i^0 \in A_i$ and a measurable function $s_i^+ : \mathcal{H} \to (0,1]$, where $s_i^+(h) \in (0,1]$ specifies the time point at which player i intends to make the next switch, conditional on reaching history h. We impose two assumptions on s_i^+.

Assumption 1 *For each $i \in I$, we have $s_i^+(h) > \tau(h) \;\; \forall h \in \mathcal{H}$.*

Assumption 2 (uniform inertia) *For each $i \in I$ and s_i^+, there exists an $\varepsilon > 0$ such that*

$$s_i^+(h) \neq 1 \implies s_i^+(h) - \tau_i(h) \geq \varepsilon, \quad \forall h \in \mathcal{H}.$$

Assumption 1 says naturally that, upon reaching history h, players may take the next move only at a future time point. Assumption 2 says that, for each strategy of a player, there is a uniform bound $\varepsilon > 0$ on inertia in the sense that any two moves of that player must be at least ε apart. There is one exception to this requirement due to our notational convention. If the player's previous move was within ε of the end ($t\!=\!1$) so that she cannot move again before the game ends, then we mandate that she must switch at the end of the game, i.e., at $t = 1$. This is because, for notational ease, we represent a strategy as a function s_i^+ (defined above) which disallows the option of never switching in the future, and adopt the convention of identifying that option with the choice of switching actions at the end of the game (i.e., at $t = 1$). This is innocuous since the two choices are payoff-equivalent so long as the utility is the same for two outcomes that differ on a measure zero set of time, a condition we maintain throughout this paper. Note that the uniform bound ε is strategy-specific and may be arbitrarily short depending on the strategy.

Let S_i denote the set of player i's strategies satisfying Assumptions 1 and 2, and $S \equiv \Pi_{i \in I} S_i$. Given a strategy profile $(s_i^0, s_i^+)_{i \in I} \in S$, the game proceeds as follows. At the initial node they choose $h^0 = (s_1^0, s_2^0)$ and form an initial history. At the information set determined by the initial history h^0, each player $i \in \{1, 2\}$ chooses simultaneously and independently a time point $s_i^+(h^0)$ for their first intended switch. The switch intended for an earlier time point (both switches if intended for the same time) gets implemented to form a jump, leading to the next node[11] represented by

$$h' = \Big[h^0, \Big(t^1 = \min_{i \in I} \big\{s_i^+(h^0)\big\},\ I^1 = \arg\min_{i \in I} \big\{s_i^+(h^0)\big\}\Big)\Big].$$

Similarly, at history h', each player $i \in \{1, 2\}$ chooses simultaneously and independently a time point $s_i^+(h')$ for their next switch, which leads to the next node

$$h'' = \Big[h^0, (t^1, I^1), \Big(t^2 = \min_{i \in I} \big\{s_i^+(h')\big\},\ I^2 = \arg\min_{i \in I} \big\{s_i^+(h')\big\}\Big)\Big].$$

The game proceeds in an analogous manner until the end of the game is reached.

[11] To be precise, it is an information set for the player who moves at the jump if the other player's choice at h^0 is unimplemented and thus, unobserved. Such unobserved choices are strategically irrelevant as they do not affect the continuation game.

Although the uniform bound ε may be arbitrarily short as noted above, players switch actions only finitely many times during a game, with their last switches at $t = 1$ given the convention above. The set of terminal histories, denoted by $\mathcal{H}^T$, is thus defined as

$$\mathcal{H}^T \equiv \left\{ \left[h^0, \left(t^1, I^1\right), ..., \left(t^k, I^k\right) \right] \in \mathcal{H} : t^k = 1,\ I^k = \left\{1, 2\right\} \text{ and } k \in \mathbb{N} \right\}$$

and each strategy profile $(s_i^0, s_i^+)_{i \in I} \in S$ uniquely determines a terminal history in the manner described above. We denote this mapping as

$$\Lambda : S \ \rightarrow \ \mathcal{H}^T.$$

3.3 Relation to the conventional continuous-time model

The benchmark model described above, although seemingly different from the conventional continuous-time model, is "equivalent" to the model of Bergin and MacLeod (1993) provided that their notion of inertia is modified to be uniform within each strategy. We first review their model in Seciton 3.3.1, and then establish the equivalence in Section 3.3.2.

3.3.1 The Bergin-MacLeod model

It proves useful to reproduce the formulation of Bergin and MacLeod (1993). Let $o : [0, 1] \rightarrow A$ denote an outcome which records an action profile at every point $t \in [0, 1]$. Let $\mathcal{O} \equiv A^{[0,1]}$ denote the set of all outcomes. Agent i's strategy is described by a measurable function

$$x_i : \mathcal{O} \times [0, 1] \ \rightarrow \ A_i.$$

For each $(o, t) \in \mathcal{O} \times [0, 1]$, the interpretation is that at time t, agent i observes the truncated past outcome $o|_{[0,t)} \in A^{[0,t)}$ (which is a conventional history h^t) and chooses an action $x_i(o, t) \in A_i$. Hence, an agent's action depends only on the past history, in conformity with the standard, original formulation described in (1). This is captured by the first of the two assumptions Bergin and MacLeod (1993) impose on legitimate strategies:

Assumption BM1 *Given* $x_i : \mathcal{O} \times [0, 1] \ \rightarrow \ A_i$*, for any* $(o, o', t) \in \mathcal{O} \times \mathcal{O} \times [0, 1]$*, we have*

$$x_i(o, t) = x_i(o', t) \quad \text{if} \quad o(s) = o'(s) \ \text{ on a full-measure subset of } \ s \in [0, t).$$

Note that any two histories that differ on a measure 0 set are treated equivalently by every player in their strategy. Consequently, for any $o \in \mathcal{O}$, $x_i(o, \cdot)$ and $\tilde{x}_i(o, \cdot)$ are equivalent for every player if they differ on a measure 0 set of $t \in [0,1]$, because they induce the same action choices from all other players and thereby, identical total utilities for everyone. To facilitate exposition, therefore, whenever $x_i(o,\cdot)$ is a constant on a full-measure subset of any interval $[t,t') \neq \emptyset$, without loss of generality we represent it with an equivalent $x_i(o,\cdot)$ that assumes the same constant on the entire interval $[t,t')$.[12] Then, the crucial "inertia" condition introduced by Bergin and MacLeod (1993) is expressed as follows.

Assumption BM2 (inertia) *Given* $x_i : \mathcal{O} \times [0,1] \to A_i$, *for any* $(o,t) \in \mathcal{O} \times [0,1)$ *there exists an* $\varepsilon > 0$ *such that*

$$x_i(o,t) \;=\; x_i(o,t'), \quad \forall t' \in [t, t+\varepsilon).$$

This implies that if agent i takes up a certain action at time t, agent i sticks with that action for a positive span of time from t. Let X_i^{BM} denote the set of agent i's strategies satisfying Assumptions BM1 and BM2, and $X^{BM} \equiv \Pi_{i\in I} X_i^{BM}$. The main result of Bergin and MacLeod (1993) is that each $x \in X^{BM}$ uniquely determines an outcome, and we denote this mapping by

$$\Lambda^{BM} \;:\; X^{BM} \;\to\; \mathcal{O}.$$

Note that any action chosen must prevail for a positive span ε of time by inertia, but ε may vary and be arbitrarily short depending on the time t, the chosen action, and the history. Therefore, a strategy profile x may entail an infinite sequence of action switches taking place at time points converging to a limit, say $s < 1$ before the end of the game,[13] determining a conventional history $h^s \in A^{[0,s)}$. Then, the players make action choices at time s based on the history h^s, from which a second sequence of action switches ensues according to the given strategy profile x. An outcome is uniquely determined through analogous processes repeated (as needed) until the end of game is reached.

3.3.2 An equivalence

We now impose the following assumption in the Bergin-MacLeod model, which corresponds to Assumption 2 of our benchmark model.

[12]For example, $x_i(o,t)$ assuming a_i for $t \leq 0.5$ and b_i for $t > 0.5$, is represented by $x_i(o,t)$ assuming a_i for $t < 0.5$ and b_i for $t \geq 0.5$. Without this convention of representing equivalent strategies, defining inertia is more complicated in order to accommodate all strategies in the same equivalence class—see Definition 1 in page 25 of Bergin and MacLeod (1993).

[13] As a simple example, if player 1 switches actions at every time point $t = 1/2 - 1/4^n$ for all $n \in \mathbb{N}$ and player 2 never switches, an infinite sequence of switches entails converging to a limit point $s = 1/2$.

Assumption UI (uniform inertia) *Given* $x_i : \mathcal{O} \times [0,1] \;\to\; A_i$*, there exists an* $\varepsilon > 0$ *such that for any* $(o,t) \in \mathcal{O} \times [0,1)$ *there is* $\hat{t} \geq 0$ *where* $\hat{t} \leq t < \hat{t} + \varepsilon$ *and*

$$x_i(o,t) = x_i(o,t'), \quad \forall t' \in [\hat{t}, \hat{t}+\varepsilon) \cap [\hat{t}, 1].$$

This is stronger than Assumption BM2 (inertia) because there is a uniform lower bound on inertia strategy-wise: it says that within each strategy there is a minimal span ε of inertia such that any action taken up at any time point must prevail for at least ε of time (or until game ends). This implies that every strategy profile entails a finite (albeit possibly arbitrarily large) number of action switches, precluding infinite sequences of moves feasible under inertia as described above. Let X_i denote the set of i's strategies satisfiying Assumptions BM1 and UI, and $X \equiv \Pi_{i \in I} X_i$. Thus, we have $X \subset X^{BM}$.

We are now ready to state an equivalence of our benchmark model and a simplified version of the Bergin-MacLeod model. For this, we need a mapping that transforms in the obvious manner every terminal history in our model to an outcome in the Bergin-MacLeod model, denoted by

$$\eta \,:\, \mathcal{H}^T \;\to\; \mathcal{O}. \tag{3}$$

Proposition 1 *There exists a profile of surjective functions* $[\xi_i : X_i \to S_i]_{i \in I}$ *such that*

$$\Lambda^{BM}\big[(x_i)_{i \in I}\big] \;=\; \eta\big[\Lambda\big([\xi_i(x_i)]_{i \in I}\big)\big], \quad \forall (x_i)_{i \in I} \in X. \tag{4}$$

Proposition 1, proved in Appendix A.2, implies that one may represent the Bergin-MacLeod model in our framework provided that their strategies satisfy Assumption UI. Specifically, their strategy x_i is mapped to $\xi_i(x_i)$ in our framework in such a way that the outcome induced by their strategy profile $(x_i)_{i \in I}$ and the terminal history induced by the mapped strategy profile $[\xi_i(x_i)]_{i \in I}$ match via η.

Uniform inertia ensures that, whichever strategies the players adopt, the game evolves with sequential moves taken at discrete time points reaching the end within finite steps. This warrants that all immediately succeeding and preceding nodes are clearly defined where relevant and therefore, the benchmark model is representable as a game tree (cf. footnote 14).

However, although the lower bound ε of inertia is strictly positive in each strategy, a player may still choose a strategy where this lower bound is arbitrarily close to 0. Moreover, the lower bound of inertia applies from the time points of a player's own switches of actions but not from those of other player's switches of their actions. Therefore, inductive strategies that involve responding to certain moves "immediately" (i.e., as soon as possible), such as the grim-trigger strategy, remain ill-defined because the earliest instant a player may choose to act after certain moves of any player is still not pinned down.

As a matter of fact, the benchmark model captures the first of our two core ideas mentioned in Introduction but not the second, namely, interpersonal differences in their limitation in responding to changed situations without delay. We enrich the benchmark model with this feature in the next section, which also accommodates inductive strategies discussed above.

4 Extensive-form games with private reaction lag

In this section, we add two additional ingredients to the benchmark model of Section 3, namely, reaction-lag types and mixed behavior strategies.

Let $\Theta_i = (0, 1)$ denote the set of reaction-lag types of player i, and there be a common prior $\mu \in \triangle(\Theta)$ where $\Theta \equiv \Pi_{i \in I} \Theta_i$. A type $\theta_i \in \Theta_i$ encodes the minimal span of time that has to elapse after the previous move before player i may make the next move, referred to as their "reaction lag," which is private information. A key difference from the previous, benchmark model is that the reaction lag, θ_i, is a lower bound of inertia that is player-specific and thus, applies to *all* strategies that a θ_i-type agent may employ. It captures interpersonal variations in response time, which may be due to varying efficiency in information processing and decision making and/or agility in implementing the decision.

We provide a sketch of an extensive-form game. First, nature chooses a type profile $\theta \in \Theta$. At the node corresponding to the chosen θ, the two players simultaneously choose an action each from A_i to form an initial history $h^0 \in A$. At the succeeding node corresponding to h^0, the set of possibles choices for player i is the interval $[\theta_i, 1]$ of feasible time points for the next switch of actions. The two players' simultaneous choices there, say (t_1, t_2), lead to the next node corresponding to a history with a first jump at $\min\{t_1, t_2\}$, denoted by h'. At the node corresponding to h', the set of possibles choices for player i is either $[\tau(h') + \theta_i, 1]$ if her reaction lag applies or $(\tau(h'), 1]$ otherwise. The two players' simultaneous choices there determine the next jump and lead to the node of an updated history, and the game tree continues analogously.[14]

Thus, a strategy of a player $i \in \{1,2\}$ consists of a measurable initial (mixed) action strategy $\sigma_i^0 : \Theta_i \to \triangle(A_i)$ and a measurable function $\sigma_i^+ : \Theta_i \times \mathcal{H} \to \triangle(T)$ that represents a (mixed) switching strategy[15] in the sense that $\sigma_i^+(\theta_i, h) \in \triangle((\tau(h), 1])$ is a distribution of time points $t \in T = (0, 1]$ at which type θ_i intends to make the next switch, conditional on reaching history h. (Further restrictions will be specified below reflecting reaction lags). Note that we consider only behavior mixed strategies, which is innocuous since Kuhn's Theorem applies as explained later.

There are two kinds of "reaction lags", namely, in response to their own moves and to the opponent's moves. The former kind has been introduced in the benchmark model as Assumption 2 (without private types θ_i and in pure strategies), in addition to the obvious Assumption 1 that the next move takes place in the future. These two assumptions are extended to mixed strategies straightforwardly as follows: $\forall h \in \mathcal{H}$, $\forall \theta_i \in \Theta_i$, $\forall i \in I$,

Assumption 3 $\sigma_i^+(\theta_i, h) \in \triangle((\tau(h), 1])$

Assumption 4 $\sigma_i^+(\theta_i, h) \in \triangle([\tau_i(h) + \theta_i, 1] \cup \{1\})$

That is, having made the last move at time $\tau_i(h)$, player i of type θ_i may make a subsequent move at or after $\tau_i(h)+\theta_i$, subject to one exception due to the modelling convention explained earlier: if $\tau_i(h) + \theta_i > 1$ then player i switches at $t = 1$, which is equivalent to making no more moves.

With regard to the reaction lag in response to the opponent's moves, one is tempted to simply replace $\tau_i(h)$ with $\tau(h)$ in Assumption 4 and require that[16]

$$\sigma_i^+(\theta_i, h) \in \triangle([\tau(h) + \theta_i, 1] \cup \{1\}), \quad \forall h \in \mathcal{H}, \ \forall \theta_i \in \Theta_i, \ \forall i \in I, \tag{5}$$

but this suffers from a pathological problem and more structure is needed. Prior to elaborating on this issue in Section 4.2, we first establish that the model is well-defined with Assumptions 3 and 4 only; and define an equilibrium notion that continues to be valid after reactions lags are applied to opponent's moves in Sections 4.2–4.4.

[14] This is also the game tree of the benchmark model of Section 3 with one change: the choice of θ_i at the beginning is player i's strategic choice of a uniform bound ε of inertia (described in Assumption 2).

[15] Although some histories are irrelevant for each type θ_i, we keep all histories for notational ease without loss of generality.

[16] For expositional ease, we assume that agent θ_i's reaction lag to opponents' moves is of the length as that to their own moves. The analysis extends straightforwardly to the case that they differ.

4.1 Equilibrium

Let Σ_i denote the set of player i's behavior strategies $\sigma_i = (\sigma_i^0, \sigma_i^+)$ that satisfy Assumptions 3–4, and $\Sigma \equiv \Pi_{i\in I}\Sigma_i$. Given a strategy profile $\sigma = (\sigma_i)_{i\in I} \in \Sigma$ and a realized type profile $\theta = (\theta_i)_{i\in I}$, the game proceeds as follows. At the initial node the players choose an action according to $\sigma_i^0(\theta_i)$ to form an initial history $h^0 \in A_1 \times A_2$. At the information set determined by the initial history h^0, each player $i \in \{1, 2\}$ chooses simultaneously a time for their first switch according to their behavior strategy $\sigma_i^+(\theta_i, h^0) \in \triangle([\theta_i, 1])$; denoting by t_i^1 the realized choice of player i at this information set, the game proceeds to history h' defined as

$$h' = \big[h^0, \big(t^1 = \min_{i\in I}\{t_i^1\},\ I^1 = \arg\min_{i\in I}\{t_i^1\}\big)\big].$$

At the history h', players choose their next switch times t_i^2 according to their respective behavior strategy $\sigma_i^+(\theta_i, h')$, and the game proceeds to history h'' defined as

$$h'' = \big[h^0, \big(t^1, I^1\big), \big(t^2 = \min_{i\in I}\{t_i^2\},\ I^2 = \arg\min_{i\in I}\{t_i^2\}\big)\big].$$

The game proceeds in an analogous manner until a terminal history is reached, ending the game.

Generally, together with strategy sets Σ_i for $i \in I$, any measurable utility function defined on the set of terminal histories $\mathcal{H}^T$ (i.e., terminal nodes) determines a continuous-time game. In various situations, however, a player derives flow utility from the prevailing action profile at every time point and their total utility is the sum of flow utility over the duration of the game. Below, we elaborate such a model where player i's flow utility is $u_i : A \to \mathbb{R}$.

Let $\widehat{u}_i : \mathcal{H} \to \mathbb{R}$ denote the sum of realized flow utility of player i upon reaching a history $h \in \mathcal{H}$, that is, $\widehat{u}_i(h^0) = 0$ and for every $h = \big[h^0, \big(t^1, I^1\big), ..., \big(t^k, I^k\big)\big] \in \mathcal{H}$,

$$\widehat{u}_i(h) = u_i(h^0) \times t^1 + \sum_{\ell=1}^{k-1} u_i\Big(\psi\big([h^0, (t^1, I^1), \cdots, (t^\ell, I^\ell)]\big)\Big) \times \big(t^{\ell+1} - t^\ell\big).$$

Given Assumptions 3–4, it is straightforward to show that a terminal history gets reached with finite moves (hence, with probability 1) for any $(\theta, \sigma) \in \Theta \times \Sigma$. Let $\Upsilon^{(\theta,\sigma)}$ denote the distribution of terminal histories induced by (θ, σ). Define $U_i : \Theta \times \Sigma \to \mathbb{R}$ for each $i \in \{1, 2\}$ as

$$U_i\,(\theta, \sigma) \ \equiv\ \int_{\mathcal{H}^T} \widehat{u}_i(h)\, d\Upsilon^{(\theta,\sigma)},$$

that is, $U_i(\theta, \sigma)$ is the expected utility of type θ_i of player i, conditional on σ being played by the players and her opponent being of type θ_{-i}. The following result establishes the measurability of U_i.

Proposition 2 *For any $i \in I$ and any $(\theta_i, \sigma) \in \Theta_i \times \Sigma$, the function $\widehat{U}_i(\cdot \,|\, \theta_i, \sigma) : \Theta_{-i} \to \mathbb{R}$ defined as*

$$\widehat{U}_i(\theta_{-i} \,|\, \theta_i, \sigma) \;\equiv\; U_i(\theta_i, \theta_{-i}, \sigma)$$

is measurable.

The proof of Proposition 2 is relegated to Appendix A.3. In light of this result, the expected utility of type θ_i, given σ, is calculated as

$$\int_{\Theta_{-i}} \widehat{U}_i(\theta_{-i} \,|\, \theta_i, \sigma)\, \mu[d\theta_{-i}|\theta_i] \;\equiv\; \int_{\Theta_{-i}} U_i(\theta_i, \theta_{-i}, \sigma)\, \mu[d\theta_{-i}|\theta_i].$$

We are now ready to define (Bayesian Nash) equilibrium of the base model.

Definition 1 *A strategy profile $\sigma \in \Sigma$ constitutes an equilibrium if for every $i \in I$ and every $\theta_i \in \Theta_i$,*

$$\int_{\Theta_{-i}} U_i(\theta_i, \theta_{-i}, \sigma)\, \mu[d\theta_{-i}|\theta_i] \;\geq \int_{\Theta_{-i}} U_i(\theta_i, \theta_{-i}, \sigma_i', \sigma_{-i})\, \mu[d\theta_{-i}|\theta_i], \quad \forall \sigma_i' \in \Sigma_i.$$

Clearly, an equilibrium exists whenever a pure-strategy Nash equilibrium exists for the stage game. In particular, it is an equilibrium for both players to start with a stage game equilibrium $a \in A$ and never switch until the end.

4.2 Reaction lag to opponent's moves

Let us use a concrete example to illustrate the issue with the condition (5), i.e., requiring that each player i may move only after their reaction lag θ_i has elapsed since the latest move by either player. Suppose their types are $\theta_1 = 0.1$ and $\theta_2 = 0.15$ and player 1 switches at every interval of 0.1 starting from $t = 0.1$. Then, the earliest time at which player 2 may move after the initial action is $t = 0.15$ by (5), but player 1 pre-empts by switching at $t = 0.1$, pushing back the earliest time at which player 2 may move to a later time point of $t = 0.25$. But, player 1 switches again at $t = 0.2,\ 0.3$, and so on, preventing player 2 from making any move until $t = 0.35,\ 0.45$, and so on, successively, thereby blocking player 2 from moving at all throughout the entire game.

As such, the condition (5) leads to a pathological problem: a player with a shorter reaction lag may prevent the opponent from moving at all by switching actions with intervals shorter than the opponent's reaction lag. This is clearly unreasonable and artificially precludes proper implementation of perfectly sensible strategies such as the grim-trigger strategy.

To be guided on how this pathological problem may be alleviated, let us examine what the reaction lag may be capturing in reality when it is triggered by other players' moves. Other players' moves generate new information relevant for the player in question to devise a best course of play in the continuation game. Hence, the reaction lag may depict the minimum amount of time for the player to process the newly-arrived information and come up with a revised plan of how to play the continuation game optimally.

Interpreted this way, the exact length of reaction lag may differ in each instance, depending on how complex the newly-arrived information is. In principle, such variations could be incorporated by modelling the reaction lag as a function of the history (as well as the agent's type). However, it is hard to quantify precise magnitudes of such variation across histories because they depend on various factors, such as personal traits of the agent and the particular situation in question. Therefore, we abstract from such detailed variations and model the reaction lag as the minimum deliberation time for newly-available information *independently* of its (hard to quantify) complexity. In fact, insofar as complexity-dependent variations are small relative to this minimum lag, their effect should be insignificant.

With the reaction lag interpreted as such, the impact of additional moves made during a reaction lag hinges on whether the extra information generated by such moves can be subsumed in the current deliberation. If it can be subsumed fully, at the end of each reaction lag the agent comes up with a plan of how to play the continuation game optimally reflecting all the moves taken during the reaction lag. If it cannot be subsumed, upon completion of the ongoing deliberation at the end of the current reaction lag, a new reaction lag should start for the agent to deliberate on the new information generated by the moves taken during the reaction lag just ended.

We formalize these two polar cases into two alternative models below, which resolves the pathological problem mentioned above. Intermediate cases can be modelled by postulating a probability that each move during a reaction lag can be subsumed in the ongoing deliberation, for instance. One may also allow for the option of discarding the current deliberation and starting a new reaction lag each time a move takes place that cannot be subsumed in the ongoing deliberation. Although these variant models could be suitable for certain situations, the differences should be small in general.

Section 4.3 formalizes the case in which the opponent's moves during a player's reaction lag are subsumed in the current deliberation without affecting the length of the on-going reaction lag, which is referred to as the nested reaction-lag model. Section 4.4 formalize the other case in which such moves of the opponent cannot be subsumed in the current deliberation and thus, trigger another reaction lag of full length at the expiration of the current one. We call such a model the deferred reaction-lag model. Both models are well-defined in the sense described in Section 4.1 as they satisfy Assumptions 3–4, thus so is the equilibrium in Definition 1.

4.3 Nested reaction-lag model

In this model, players are assumed to deliberate on the opponent's moves taken during an ongoing reaction lag in the current reaction lag. We recursively define a function $\gamma_i : \Theta_i \times \bigcup_{k=0}^{\infty} H^k \to [0,1]$ that depicts, for every legitimate history h, the start time of the "pertinent" reaction lag during which player i of type θ_i deliberates on the history updated by the last jump of h. Note that it is the reaction lag that is either ongoing or launched at $\tau(h)$ so that, in particular, $\gamma_i(\theta_i, h) \leq \tau(h)$.

Clearly, $\gamma_i(\theta_i, h^0) = 0$ for every initial history $h^0 \in H^0$ because the initial action choices at $t = 0$ launch an initial reaction lag for every player. Next, for each history with one jump $[h^0, (t^1, I^1)] \in H^1$, define

$$\gamma_i\big(\theta_i, [h^0, (t^1, I^1)]\big) = \begin{cases} 0 & \text{if } t^1 < \theta_i, \\ t^1 & \text{if } t^1 \geq \theta_i, \end{cases} \quad \forall \theta_i \in \Theta_i \ \forall i \in I.$$

That is, the pertinent reaction lag is the initial one if the first jump takes place during the initial reaction lag for agent θ_i; but it is launched at the time of the first jump otherwise.

Given γ_i defined for all $(\theta_i, h) \in \Theta_i \times H^k$ for $k \geq 1$, define for $[h, (t^{k+1}, I^{k+1})] \in H^{k+1}$,

$$\gamma_i\big(\theta_i, \big[h, (t^{k+1}, I^{k+1})\big]\big) = \begin{cases} \gamma_i\big(\theta_i, h\big) & \text{if } t^{k+1} < \gamma_i\big(\theta_i, h\big) + \theta_i, \\ t^{k+1} & \text{if } t^{k+1} \geq \gamma_i\big(\theta_i, h\big) + \theta_i. \end{cases}$$

That is, if a new jump takes place while the pertinent reaction lag of the previous history is still going on for player i, the start time of the pertinent reaction lag remains unchanged for the updated history; but if a jump occurs after it has ended, a new reaction lag that starts at that point is the pertinent reaction lag for the updated history.

An agent θ_i deliberates on each history h during the corresponding pertinent reaction lag and chooses a behavior strategy which is a probability distribution of time points for her next move that may take place at or after the expiration of the pertinent reaction lag. Each time the opponent moves during an on-going reaction lag (so that the pertinent reaction lag stays put), the behavior strategy chosen for the updated history replaces the previously chosen behavior strategy.

We call this model the "nested reaction-lag model" and denote it by $\Sigma^{nest} \equiv \Pi_{i \in I} \Sigma_i^{nest}$, which is formally defined as

$$\Sigma_i^{nest} = \left\{ (\sigma_i^0, \sigma_i^+) \in \Sigma_i : \begin{array}{c} \sigma_i^+(\theta_i,\, h) \in \triangle([\gamma_i(\theta_i, h) + \theta_i, 1] \cup \{1\}), \\ \forall \theta_i \in \Theta_i,\, \forall h \in \mathcal{H}. \end{array} \right\}. \tag{6}$$

We regard this model as the continuous-time game formulation with a minimal structure of reaction lag that fulfils our desiderata, namely, representation in extensive form, warranting well-defined outcomes and payoffs for all strategy profiles as well as well-defined immediate responses. The grim-trigger strategy, for instance, is unambiguously defined and specifies that if an opponent defects at t, an agent i of type θ_i "immediately" follows suit by defecting at the expiration of her pertinent reaction lag, that is, at $\gamma_i(\theta_i, h) + \theta_i$ where h is the history formed by the opponent's defection at t (i.e., $\tau(h) = t$).

4.4 Deferred reaction-lag model

Next, we consider the other polar model: the opponent's moves during a player's reaction lag are not subsumed in the current deliberation and thus, trigger another reaction lag of full length at the expiration of the current one. Specifically, suppose that a history h launched agent θ_i's reaction lag at $t = \tau(h) < 1 - \theta_i$. Then, during her current reaction lag $(t, t + \theta_i)$, agent θ_i deliberates only on the information contained in the history h and comes up with an optimal response $\sigma_i^+(\theta_i, h) \in \triangle([t + \theta_i, 1])$. If there have been additional moves during the reaction lag $(t, t + \theta_i)$ and/or at time $t + \theta_i$ upon its expiration, thus forming an updated history h' as of time $t + \theta_i$, then agent θ_i goes through another, subsequent reaction lag for

$(t+\theta_i, t+2\theta_i)$ to deliberate on the history h'.

Consider the case that agent θ_i's optimal response $\sigma_i^+(\theta_i, h)$ places a full mass on $\tau(h)+\theta_i$, that is, agent θ_i chose to make an "immediate reaction" by taking an action as soon as possible in response to the move that took place and triggered the current reaction lag at $t = \tau(h)$. Then, even if the opponent has made additional moves during the current reaction lag, agent θ_i gets to carry out her immediate reaction at $\tau(h) + \theta_i$, i.e., upon expiration of the current reaction lag and before a subsequent reaction lag gets launched.[17] Since the first time to react to a previous move is pinned down as such, inductive strategies such as the grim-trigger strategy are well-defined. In addition, a quicker player cannot prevent the slower player from moving at all, resolving the aforementioned pathological problem.

We call this model the "deferred reaction-lag model" and denote it by $\Sigma^{dfr} \equiv \Pi_{i\in I}\Sigma_i^{dfr}$. To formally define Σ_i^{dfr}, we modify $\gamma_i(\theta_i, h)$ to $\widetilde{\gamma}_i(\theta_i, h)$ to depict the start time of the pertinent reaction lag for any given history h, in such a way that the last jump of h is deliberated during the corresponding pertinent reaction lag. This modified pertinent reaction lag for h is either the one launched by the last jump at $\tau(h)$, or the one to be launched by the last jump upon expiration of the ongoing reaction lag.[18] Hence, $\widetilde{\gamma}_i(\theta_i, h) \geq \tau(h)$.

Clearly, $\widetilde{\gamma}_i(\theta_i, h^0) = 0 = \gamma_i(\theta_i, h^0)$ for every $h^0 \in H^0$. For each history with one jump $[h^0, (t^1, I^1)] \in H^1$, define

$$\widetilde{\gamma}_i\big(\theta_i, [h^0, (t^1, I^1)]\big) = \begin{cases} \theta_i & \text{if } \; t^1 < \theta_i, \\ t^1 & \text{if } \; t^1 \geq \theta_i, \end{cases} \quad \forall \theta_i \in \Theta_i \; \forall i \in I.$$

Given $\widetilde{\gamma}_i$ defined for all $(\theta_i, h) \in \Theta_i \times H^k$ for $k \geq 1$, define recursively for $[h, (t^{k+1}, I^{k+1})] \in H^{k+1}$,

$$\widetilde{\gamma}_i\big(\theta_i, \big[h, (t^{k+1}, I^{k+1})\big]\big) \;=\; \begin{cases} \widetilde{\gamma}_i\big(\theta_i, h\big) & \text{if } \; t^{k+1} \leq \widetilde{\gamma}_i\big(\theta_i, h\big), \\ \widetilde{\gamma}_i\big(\theta_i, h\big) + \theta_i & \text{if } \; \widetilde{\gamma}_i\big(\theta_i, h\big) < t^{k+1} < \widetilde{\gamma}_i\big(\theta_i, h\big) + \theta_i, \\ t^{k+1} & \text{if } \; \widetilde{\gamma}_i\big(\theta_i, h\big) + \theta_i \leq t^{k+1}, \end{cases}$$

That is, if a jump occurs while the pertinent reaction lag for the previous history is still upcoming (i.e., hasn't started), then the upcoming one is the pertinent reaction lag for the

[17]Alternatively, agent θ_i may be assumed to decide whether to carry out or not the immediate reaction before the subsequent reaction lag starts. This could be a sensible variant of the model for some situations.

[18]For example, suppose a new reaction lag of agent θ_i started at t and the opponent moves at $t' \in (t, t+\theta_i)$ forming a history h. Then, $\gamma_i(h) = t$ and $\widetilde{\gamma}_i(h) = t + \theta_i$. That is, in the nested reaction-lag model the move at t' does not renew i's reaction lag and the pertinent one remains $(t, t+\theta_i)$; while in the deferred reaction-lag model, it renews i's reaction lag with a delay and the pertinent one becomes $(t+\theta_i, t+2\theta_i)$.

updated history as well. If a jump occurs after the pertinent reaction lag for the previous history has started but before it ends, then due to start upon expiration of the current reaction lag is a new one which is the pertinent reaction lag for the updated history. If a jump occurs after the pertinent reaction lag for the previous history is over, the pertinent reaction lag for the updated history is the one that starts at the time of the jump.

With $\widetilde{\gamma}_i$ defined as above, agent θ_i chooses a behavior strategy for every history h, which is a probability distribution of time points for her next move that may be taken at or after the expiration of the corresponding pertinent reaction lag, subject to the constraint that her previously chosen strategy remains valid for the time points before the pertinent reaction lag corresponding to h has started. This constraint is relevant only for histories whose corresponding pertinent reaction lags are yet to start.

To state this constraint formally, for each non-initial history h we define the "preceding" history as the truncation of h at the penultimate jump and denote it by h^-. Then, Σ_i^{dfr} is defined as

$$\Sigma_i^{dfr} = \left\{ (\sigma_i^0, \sigma_i^+) \in \Sigma_i : \begin{array}{c} \sigma_i^+(\theta_i, h) \in \triangle(\{\widetilde{\gamma}_i(\theta_i, h)\} \cup [\widetilde{\gamma}_i(\theta_i, h) + \theta_i, 1] \cup \{1\}), \text{ and} \\ \sigma_i^+(\theta_i, h)(\{\widetilde{\gamma}_i(\theta_i, h)\}) = \frac{\sigma_i^+(\theta_i, h^-)(\{\widetilde{\gamma}_i(\theta_i, h)\})}{\sigma_i^+(\theta_i, h^-)([\widetilde{\gamma}_i(\theta_i, h), 1])} \text{ if } \tau(h) < \widetilde{\gamma}_i(\theta_i, h) \\ \text{(and the denominator is non-zero),} \\ \forall \theta_i \in \Theta_i, \ \forall h \in \mathcal{H}. \end{array} \right\}. \tag{7}$$

The fraction in the second line of (7) reflects that the behavior strategy at h^- for the time point $t = \widetilde{\gamma}_i(\theta_i, h) > \tau(h)$ remains valid conditional on reaching that time point.

4.5 Kuhn's Theorem

Kuhn's Theorem (1953) facilitates analysis of finite extensive-form games that satisfy the condition known as Perfect Recall, by establishing the result that one may focus on behavior strategies without loss of generality. Aumann (1964) extended this result to extensive-form games with infinite strategy sets and countable moves under certain sequential structures of information. Since this result applies to our continuous-time models proposed above, we lose no generality by considering behavior strategies, including mixed strategies, as we have done. However, the result of Aumann (1964) does not apply to previous formulations of continuous-time games such as Simon and Stinchcombe (1989), Bergin (1992) and Bergin and MacLeod (1993), where only pure strategies are considered.

5 Application to Prisoners' Dilemma

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

$$\begin{array}{ccc} 1\backslash 2 & C & D \\ C & (1,1) & (0,h) \\ D & (h,0) & (\ell,\ell) \end{array} \tag{8}$$

where $0 < \ell < 1 < h$ over a unit time interval $[0, 1]$. We analyze this game by applying the continuous-time game framework introduced in the previous section. In particular, we demonstrate that our model accommodates an equilibrium that exhibits the typical behavior observed in the experiment by Friedman and Oprea (2012), namely, the two players cooperate until near the end of the game when one player defects first and the other follows suit immediately. To our knowledge, this is the first paper that theoretically justifies such a behavior as an equilibrium (apart from ϵ-equilibrium). Moreover, from our equilibrium emerges a clear and testable implication which is borne out by the experimental data.

We will characterize below the following equilibrium: Starting with (C, C), the two players continue cooperation until their respective time of planned defection, $\tau^*(\theta_i) \in [0, 1]$ for $i \in \{1, 2\}$, which is a decreasing function of their type θ_i. Thus, the slower player of the two (the one with a longer reaction lag) reaches her defection time first and defects at that time. The other player, being pre-empted, follows suit as soon as possible, i.e., upon expiration of her reaction lag. Thus, cooperation continues until near the end provided that the reactions lags are small (which should be the case in laboratory environments). Our analysis is valid under both the nested and deferred reaction-lag models.

Let us start with an informal insight underlying the result. 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 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.

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 pre-empt their opponent by defecting first before the game ends. How may this dilemma be reconciled to sustain cooperation in the face of eventual defection? We show that this dilemma can be resolved when player's reaction lag is private information. The intuition is that the quicker a player can respond, the longer she is willing to continue cooperation because she can cope better when pre-empted. 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.

This equilibrium provides us with a clear testable implication: the later in the game is the first defection, the shorter is the time until the opponent follows suit. We find some evidence supporting this relationship in the data reported by Friedman and Oprea (2012) on their experiments, as detailed later.

We now characterize the equilibrium for the case that the players' types, θ_1 and θ_2, are independent random variables represented by an atomless cdf F on $\Theta_i = (0, \bar{\theta})$ with full support where $\bar{\theta}$ is small. The equilibrium strategy, which is symmetric and in pure strategy, is described as follows: the initial action choice at $t = 0$ is $s_i^0 = C$, and

(a) $s_i^+(\theta_i, h) = \tau^*(\theta_i)$ if $h = (C, C) \in H^0$ where $\tau^* : \Theta_i \to (0, 1)$ is a continuous and strictly decreasing function such that $0 < \tau^*(\theta) < 1 - \theta$;

(b) $s_i^+(\theta_i, h)$ for any $h \neq (C, C)$ is for player i to switch to D at the earliest time possible, unless her current action is D in which case she stays put until $t = 1$.

Note that part (b) pertains to the history formed by the first defection along the equilibrium path, as well as all off-equilibrium contingencies. Along the equilibrium-path, therefore, one player defects at some point in the time interval $[\underline{t}, \bar{t}) \equiv [\tau^*(\bar{\theta}), \tau^*(0)) \subset (0, 1)$, soon after which the other player follows suit.

To verify that the strategy profile above constitutes an equilibrium, first consider the continuation game prompted by any history $h \neq (C, C)$, i.e., other than the equilibrium initial history. Then, by (b), both players will switch to D as soon as possible and keep to D until the end, unless their current action is already D in which case they stay put. It is clear that this strategy is optimal for both players conditional on the opponent adopting the same strategy (since D is the dominant strategy in the stage game).

Given such an off-equilibrium strategy and the switching strategy (a) after the initial history, it is straightforward to see the optimality of the initial action choice $s_i^0 = C$.

Thus, it remains to verify optimality of their strategy at the initial history $h^0 = (C, C)$, i.e., (a) above. To show it is suboptimal for either player to defect at some $t < \tau^*(\bar{\theta})$, observe that if player i did then player $-i$ would follow suit at $t + \theta_{-i} < 1$ and (D, D) would prevail until the end according to (b). If player i delayed her defection from t, say to $\tau^*(\bar{\theta})$, then player $-i$'s defection would be delayed by the same amount as well, after which (D, D) would prevail. Clearly, the latter is better for player i because it prolongs the duration of a cooperation payoff of 1 while curtailing a smaller non-cooperation payoff of $\ell < 1$ later. Hence, defecting at any $t < \tau^*(\bar{\theta})$ is suboptimal.

Finally, to verify the optimal defection time $\tau^*(\theta_i)$ for player i of type θ_i after the initial history $h^0 = (C, C)$, note that her expected payoff from choosing to defect at $t \in [\underline{t}, \bar{t})$ is

$$\int_0^{\vartheta(t)} \big[t + \theta h + (1 - t - \theta)\ell\big] dF(\theta) + \int_{\vartheta(t)}^{\bar{\theta}} \big[\tau^*(\theta) + (1 - \tau^*(\theta) - \theta_i)\ell\big] dF(\theta) \tag{9}$$

where $\vartheta(t) \equiv \tau^{*-1}(t)$, the inverse function of τ^*. Here, the first integral is her payoff when she doesn't get pre-empted because her opponent's type θ is such that $\tau^*(\theta) > t$, and the second integral is that when she gets pre-empted. Differentiating (9) with respect to t, we get

$$(1 - \ell)F(\vartheta(t)) + \Big(\vartheta(t)(h - \ell) + \theta_i \ell\Big) f(\vartheta(t))\vartheta'(t) \tag{10}$$

where $f(\theta) \equiv F'(\theta) > 0$ on $(0, \bar{\theta})$. As (10) strictly decreases in θ_i, if it assumes 0 at $\theta_i = \vartheta(t)$, i.e.,

$$\vartheta'(t) \;=\; \frac{-(1 - \ell)F(\vartheta(t))}{h\vartheta(t) f(\vartheta(t))}, \tag{11}$$

then (10) is positive for all $\theta_i < \vartheta(t)$ and negative for all $\theta_i > \vartheta(t)$. If (11) holds at all $t \in (\underline{t}, \bar{t})$, therefore, upon reaching $\underline{t}$ without a defection, it is uniquely optimal for agent θ_i to defect at $\tau^*(\theta_i)$. Consequently, the postulated equilibrium must satisfy (11) for all $t \in (\underline{t}, \bar{t})$. Conversely, if there is a solution ϑ to the differential equation (11) on some nonempty interval $(\underline{t}, \bar{t}) \subset (0, 1)$ such that $\vartheta(t) \to \bar{\theta}$ as $t \to \underline{t}$ and $\vartheta(t) \to 0$ as $t \to \bar{t}$, then there is an equilibrium described by (a) and (b) above.

Let us examine when such a solution exists. Note that (11) 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 any $c \in \mathbb{R}$. Therefore, if a solution to (11) exists defined on an open interval of length less than 1 and mapped onto $(0, \bar{\theta})$, then it can be "shifted" to constitute an equilibrium.

By standard results in ODE (Theorems 3 and 4 in Hurewicz (1958), p.28), there is a unique solution to (11) defined on a neighborhood of any $t_0 \in \mathbb{R}$ subject to an arbitrary initial condition $\vartheta(t_0) = \theta_0 \in (0, \bar{\theta})$. For the unique solution ϑ to constitute a legitimate equilibrium strategy τ^* when inverted, ϑ must map onto Θ within a bounded domain (recall the duration of the game is "normalized" to a unit interval). A sufficient condition for this to be the case is

$$\lim_{\theta \to 0} \frac{F(\theta)}{\theta f(\theta)} \; > \; 0 \tag{12}$$

because then the RHS of (11) is uniformly bounded away from 0 for $\vartheta(t) \in (0, \bar{\theta})$. The condition (12) means that F does not vanish to an infinite order at $\theta = 0$ and is satisfied by a wide class of distribution functions including those such that $F(\theta) = \theta^n$ for θ near 0 for any $n \in (0, \infty)$.[19] The equilibrium characterization above is summarized below.

Proposition 3 *Consider the game (8) played over a unit interval* $[0, 1]$ *subject to (6) or (7). If* $\bar{\theta} > 0$ *is small enough and (12) is satisfied, there is a continuum of equilibria described by (a) and (b) above, that differ only in* τ^* *by an addition of a constant. The most efficient one of these is the one with* $\lim_{\theta \to 0} \tau^*(\theta) = 1$.

If F is uniform, for example, the most efficient equilibrium is characterized by $\tau^*(\theta) = 1 - \frac{h}{1-\ell}\theta$.

According to the equilibria described above, if the first defection takes place at time $t \in (\tau^*(\bar{\theta}), \tau^*(0))$ by one player, the other player's reaction lag is distributed according to the truncated cdf $F(\theta)/F(\vartheta(t))$ on $(0, \vartheta(t))$. Since a slower player defects earlier according to $\vartheta(t)$, the earlier is the first defection, the longer should be the time until the follow-on defection in the sense of first-order stochastic dominance. We find that this relationship is indeed observed in the experimental data reported by Friedman and Oprea (2012) on subjects who played the prisoners' dilemma game precisely described above.

Specifically, using their data, we plot in Figure 2 the cdf's of the time gaps between the first and follow-on defections for two groups of experiment rounds that differ in the times of the first defection: the first group (plotted in blue) comprise rounds where the first defection

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

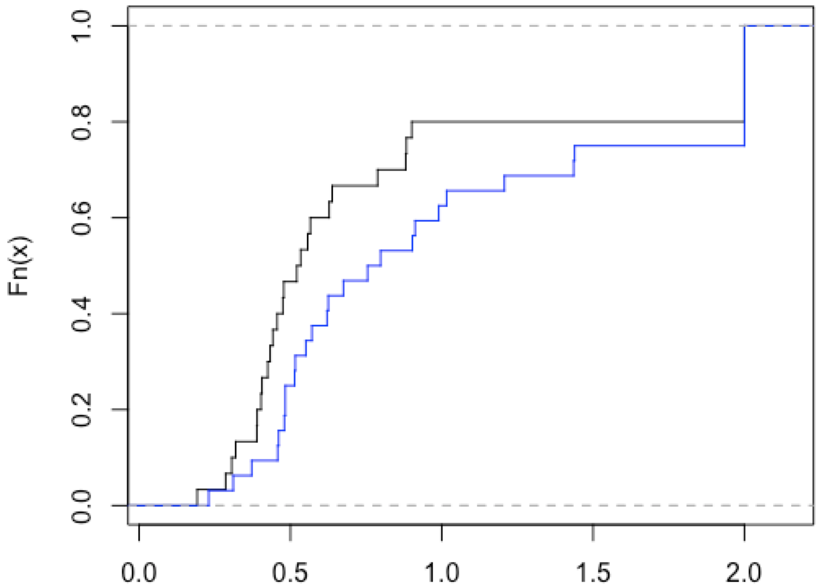

Figure 2

took place between 1.25 and 2 seconds before the game ends, while for the second group (in black) it took place between 0.5 and 1.25 second before the game ends. It appears visibly clear that the distribution of time gaps for the first group (blue, earlier first defections) first-order stochastically dominates that for the second group (black).[20]

To confirm the statistical significance, we ran the Kolmogorov-Smirnov test and reject the null hypothesis with 85% confidence level that the two groups follow the same distribution (Kolmogorov-Smirnov $D = 0.3104,\ n1 = 30,\ n2 = 32,\ P < 0.15$). We also ran the Mann-Whitney U test, which revealed that the second group is significantly larger than the first group: median latencies in the two groups are 0.5275 and 0.7765, and the distributions in the two groups differ significantly (Mann-Whitney $U = 343,\ n1 = 30,\ n2 = 32,\ P < 0.1$ two-tailed).

6 Conclusion

We proposed a new framework for modelling continuous-time games, which represents them as standard extensive-form games so that the familiar and predominant analytic apparatus can be readily applied. It differs from the conventional framework in two key respects: players choose when and how to move next each time there is a move by a player (rather

[20]We plot the cdf's for rounds 15–32 to capture the behavior of subjects experienced in playing the game.

than at every instant of time); and they are subject to a player-specific reaction lag which is the minimal time that has to elapse before their next move from each such decision point. We believe this approach enhances realism of the model, as well as resolving inherent technical problems of the conventional framework (e.g., indeterminacy of outcomes and precluding inductive strategies such as the grim-trigger strategy).

While our approach is different critically as outlined above, it is clearly rooted in and builds upon previous studies. In particular, Proposition 1 shows that our benchmark model in Section 3 is equivalent to a simplified version of Bergin and MacLeod (1993), i.e., with the additional assumption of uniform inertia. In a similar vein, our fully-fledged model in Section 4 may be viewed as an elaboration of their model with additional ingredients of reaction-lag types and mixed behavior strategies. That is, we could have incorporated the same new ideas of our fully-fledged model by imposing additional structures and ingredients on the model of Bergin and MacLeod.[21] In this sense, we follow the tradition of Simon and Stinchcombe (1989) and Bergin and MacLeod (1993), which impose structures on the conventional continuous-time model in order to make it better represent strategic scenarios in reality. Compared with these two classical papers, what we achieve is that we identify the additional structures needed in a continuous-time model for it to be representable as a standard extensive-form game which we are much better equipped to study.

We show that our framework sheds new light on some recent experimental observations. In particular, it explains the gradual unraveling of cooperation near the end of continuous-time prisoners' dilemma games reported by Friedman and Oprea (2012). It is our hope that this approach will be fruitfully extended to accommodate various other situations of intertemporal economic interactions of interest.

[21] However, such an approach would have required tremendous notational complications, while the current models in Sections 3 and 4 are much simpler notationally.

Appendix

A.1 Definition of ψ

To keep track of the action being played, we define an auxiliary function $\iota_i : A_i \times \{0, 1\}$ as

$$\iota_i(a_i, 0) \;=\; \iota_i(b_i, 1) \;=\; a_i \quad \text{and} \quad \iota_i(b_i, 0) \;=\; \iota_i(a_i, 1) \;=\; b_i$$

for each $i \in I$, where the first argument of ι_i is the action prior to a jump and the second argument depicts whether player i jumps (1) or not (0). Then, we define

$$[\psi : \mathcal{H} \longrightarrow A_1 \times A_2] \;\equiv\; \begin{pmatrix} \psi_1 : \mathcal{H} \longrightarrow A_1, \\ \psi_2 : \mathcal{H} \longrightarrow A_2 \end{pmatrix}$$

inductively on $k = 0, 1, 2, \cdots$, as

$$\psi(h^0) = h^0, \;\; \forall h^0 \in H^0,$$

$$\text{and for each} \;\; h = \big[h^-, (t^{k+1}, I^{k+1})\big] \;\; \text{where} \;\; h^- \in H^k,$$

$$\psi_1(h) = \begin{cases} \iota_1\big(\psi_1(h^-), 1\big) & \text{if } 1 \in I^{k+1} \\ \iota_1\big(\psi_1(h^-), 0\big) & \text{if } 1 \notin I^{k+1} \end{cases} \quad \text{and} \quad \psi_2(h) = \begin{cases} \iota_2\big(\psi_2(h^-), 1\big) & \text{if } 2 \in I^{k+1} \\ \iota_2\big(\psi_2(h^-), 0\big) & \text{if } 2 \notin I^{k+1}. \end{cases}$$

A.2 Proof of Proposition 1

Every terminal history in our model can be translated naturally to an outcome in the Bergin-MacLeod model, and we record it by the following function.

$$\eta : \mathcal{H}^T \longrightarrow \mathcal{O}.$$

We first describe η rigorously. For each $h = \big[h^0, (t^1, I^1), ..., (t^{k-1}, I^{k-1}), (t^k = 1, I^k = \{1, 2\})\big] \in \mathcal{H}^T$,

$$\eta(h)[t] = \psi\big([h^0, (t^1, I^1), ..., (t^{k'}, I^{k'})]\big), \;\; \forall k' \in \{0, ..., k-1\}, \;\; \forall t \in [t^{k'}, t^{k'+1}).$$

That is, starting at time $t^{k'}$, agents play the action profile $\psi\big(\big[h^0, (t^1, I^1), ..., (t^{k'}, I^{k'})\big]\big)$ until time $t^{k'+1}$ at which point they switch to the action profile $\psi\big(\big[h^0, (t^1, I^1), ..., (t^{k'}, I^{k'}), (t^{k'+1}, I^{k'+1})\big]\big)$.

Furthermore, for each non-terminal history $h \in \mathcal{H} \diagdown \mathcal{H}^T$, we let $\widehat{h} \in \mathcal{H}^T$ denote the terminal history which is a simple extension from h with no further jumps except the convention of both players jumping at $t = 1$: specifically,

$$\widehat{h^0} = \left[h^0, \left(t^1 = 1, I^1 = \{1,2\}\right)\right],$$
$$\text{and } \forall h = \left[h^0, \left(t^1, I^1\right), ..., \left(t^k, I^k\right)\right] \in \mathcal{H} \diagdown \left[\mathcal{H}^T \cup \left\{h^0\right\}\right],$$
$$\widehat{h} = \left[h^0, \left(t^1, I^1\right), ..., \left(t^k, I^k\right), \left(t^{k+1} = 1, I^{k+1} = \{1,2\}\right)\right].$$

We are now ready to define the profile of surjective functions $[\xi_i : X_i \to S_i]_{i \in I}$ which establish Proposition 1. For each $i \in I$, define $\xi_i : X_i \to S_i$ as follows. Fixing any $x_i \in X_i$, define

$$\xi_i(x_i) \equiv \left(\overline{s}_i^0,\ \ \overline{s}_i^+ : \mathcal{H} \to (0,1]\right),$$

where $\overline{s}_i^0$ and $\overline{s}_i^+$ are described in details as follows. First,

$$\overline{s}_i^0 = x_i(o, t = 0),\ \forall o \in \mathcal{O},$$

and $\overline{s}_i^0$ is well-defined due to Assumption BM1.[22] That is, following the strategy, x_i dictates that agent i play $x_i(o, t = 0)$ at time 0, and hence, for our targeted strategy $\xi_i(x_i) \equiv \left(\overline{s}_i^0,\ \ \overline{s}_i^+\right) \in S_i$, we require that agent i play $x_i(o, t = 0)$ under $\xi_i(x_i)$. Second,

$$\forall h \in \mathcal{H} \diagdown \mathcal{H}^T,$$
$$\text{define } E = \left\{t > \tau(h) : x_i\big(\eta(\widehat{h}), \tau(h)\big) \neq x_i\big(\eta(\widehat{h}), t\big)\right\},$$
$$\overline{s}_i^+(h) = \begin{cases} \min E, & \text{if } E \neq \varnothing; \\ 1, & \text{if } E = \varnothing; \end{cases}$$

That is, upon reaching history h at time $\tau(h)$, agent i plays $x_i\big(\eta(\widehat{h}), \tau(h)\big) = \psi_i(h)$ at time $\tau(h)$. Due to inertia, players must stick to $\psi(h)$ for, at least, a short while after $\tau(h)$, and given player $-i$ not taking another move, $\overline{s}_i^+(h)$ is the earliest time point at which agent i may take the next move (as described by $x_i\big(\eta(\widehat{h}), \tau(h)\big) \neq x_i\big(\eta(\widehat{h}), t\big)$). This completes the definition of $\xi_i(x_i) \equiv \left(\overline{s}_i^0,\ \ \overline{s}_i^+\right)$. It is straightforward to show that $\xi_i(x_i)$ is surjective, and we omit the details.

Finally, we prove

$$\Lambda^{BM}\left[(x_i)_{i \in I}\right] = \eta\left[\Lambda\left([\xi_i(x_i)]_{i \in I}\right)\right],\ \forall (x_i)_{i \in I} \in X. \tag{13}$$

Fix any $(x_i)_{i \in I} \in X$. Due to Assumption UI (uniform inertia), the lengths of inertia are bounded below by $\varepsilon_1 > 0$ and $\varepsilon_2 > 0$ for x_1 and x_2, and as a result, both players can take

[22] Assumption BM1 requires $x_i(o, t = 0) = x_i(o', t = 0)$ for any $(o, o') \in \mathcal{O} \times \mathcal{O}$.

finite moves for the time interval $[0,1]$. Then, we prove (13) by induction on the (sequentially numbered) jumps. Recall $\xi_i(x_i) \equiv \left(\overline{s}_i^0,\ \overline{s}_i^+ : \mathcal{H} \to (0,1]\right)$ for each $i \in I$.

First, under (x_1, x_2), players choose action profile $[x_1(o, t=0),\ x_2(o, t=0)]$ (for any $o \in \mathcal{O}$) at time 0, and analogously, under (x_1, x_2), players choose action profile $\left[\overline{s}_1^0 = x_1(o, t=0),\ \overline{s}_2^0 = x_2(o, t=0)\right]$. Furthermore, this generates the initial history $h^0 = [\overline{s}_1^0,\ \overline{s}_2^0]$ for $[\xi_1(x_1),\ \xi_2(x_2)]$ and outcome $\widehat{h^0}$ for (x_1, x_2) up to time $t = 0$.

Second, since both players will stay put with h^0 for a short while due to inertia, $\overline{s}_i^+(h^0) = \min\left\{t > \tau(h^0) : x_i\big(\eta(\widehat{h}^0), \tau(h^0)\big) \neq x_i\big(\eta(\widehat{h}^0), t\big)\right\}$ is the first time point agent i will take the next move. There are three cases to consider: (i) $\overline{s}_1^+(h^0) < \overline{s}_2^+(h^0)$; (ii) $\overline{s}_1^+(h^0) > \overline{s}_2^+(h^0)$; (iii) $\overline{s}_1(h^0) = \overline{s}_2^+(h^0)$.

In case (i), the first jump occurs at $\overline{s}_1^+(h^0)$, which generates the history $h^1 = [h^0, (t^1 = \overline{s}_1^+(h^0), I^1 = \{1\})]$ for $[\xi_1(x_1),\ \xi_2(x_2)]$ and outcome $\widehat{h^1}$ for (x_1, x_2) up to time $t = \overline{s}_1^+(h^0)$. Or quivalently, both players take the action profile h^0 at time 0, and stay put until time $\overline{s}_1^+(h^0)$, and at time $\overline{s}_1^+(h^0)$, player 1 switches her action, while player 2 does not.

A similar description as above applies to case (ii), which we omit.

In case (iii), the first jump occurs at $\overline{s}_1^+(h^0)$, which generates the history $h^1 = [h^0, (t^1 = \overline{s}_1^+(h^0), I^1 = \{1,2\})]$ for $[\xi_1(x_1),\ \xi_2(x_2)]$ and outcome $\widehat{h^1}$ for (x_1, x_2) up to time $t = \overline{s}_1^+(h^0)$. Or quivalently, both players take the action profile h^0 at time 0, and stay put until time $\overline{s}_1^+(h^0)$, and at time $\overline{s}_1^+(h^0)$, both player 1 and player 2 switch their actions.

To sum, we showed above that, under (x_1, x_2) and $[\xi_1(x_1),\ \xi_2(x_2)]$, both players take the same action profile at time 0, and for the first jump, they take the same jump at the same time point. Or equivalently, (x_1, x_2) and $[\xi_1(x_1),\ \xi_2(x_2)]$ generate the same outcome and history, respectively, up to the same time of the first jump. Then, inductively, by repeating the same argument as above, we can show that (x_1, x_2) and $[\xi_1(x_1),\ \xi_2(x_2)]$ generate the same outcome and history, respectively, up to the same time of the second jump, the third jump, and so on. Since only finite jumps occur under both under (x_1, x_2), we conclude that (13) holds.

A.3 Proof of Proposition 2

We will prove a stronger result than Proposition 2. Recall that $\Upsilon^{[\theta_i,\sigma_i]_{i\in I}}$ denote the distribution of terminal histories induced by $[\theta_i, \sigma_i]_{i\in I}$. For each $h \in \mathcal{H}$, let $\Upsilon_h^{[\theta_i,\sigma_i]_{i\in I}}$ denote the distribution of terminal histories induced by $[\theta_i, \sigma_i]_{i\in I}$ conditional on reaching history h.

Given $(\theta_i, \sigma_1, \sigma_2)$, define a value function $V_i(\theta_{-i}, h \mid \theta_i, \sigma_1, \sigma_2)$ as follows.

$$V_i(\cdot \mid \theta_i, \sigma_1, \sigma_2) : \Theta_{-i} \times \mathcal{H} \longrightarrow \mathbb{R}, \tag{14}$$
$$V_i(\theta_{-i}, h \mid \theta_i, \sigma_1, \sigma_2) \equiv \int_{\mathcal{H}^T} \widehat{u}(h') \Upsilon_h^{(\theta_1,\theta_2,\sigma_1,\sigma_2)}[dh'] - \widehat{u}(h),$$

i.e., given players taking (σ_1, σ_2), $V_i(\theta_{-i}, h \mid \theta_i, \sigma_1, \sigma_2)$ is the expected (future) utility (after h) of type θ_i conditional on reaching h, which is the (future) value function of θ_i upon reaching h.

It is straightforward to see that

$$\widehat{U}_i(\theta_{-i} \mid \theta_i, \sigma_1, \sigma_2) \equiv V_i\left(\theta_{-i}, h^0 \mid \theta_i, \sigma_1, \sigma_2\right),$$

and hence, Proposition 2 follows from the following result which is proved below.

Proposition 4 *Consider any model $\langle \Sigma \equiv \Pi_{i\in I}\Sigma_i \rangle$ which satisfies Assumptions 3 and 4. For any $i \in I$ and any $(\theta_i, \sigma_1, \sigma_2) \in \Theta_i \times \Sigma_1 \times \Sigma_2$, the function $V_i(\theta_{-i}, h \mid \theta_i, \sigma_1, \sigma_2)$ as defined in (14) is measurable.*

Define a function $\Xi : \triangle(\mathbb{R}) \times \triangle(\mathbb{R}) \to \triangle(\mathbb{R} \times \mathbb{R})$ such that for any $\mu_1 \in \triangle(\mathbb{R})$ and $\mu_2 \in \triangle(\mathbb{R})$, we have

$$\Xi(\mu_1, \mu_2)[\{(x_1, x_2) : x_1 \leq c, x_2 \leq d\}] = \mu_1[\{x_1 : x_1 \leq c\}] \times \mu_2[\{x_2 : x_2 \leq d\}], \forall \ (c, d) \subset \mathbb{R}\times\mathbb{R},$$

i.e., for two independent real random variables, x_1 and x_2 with distributions μ_1 and μ_2, respectively, $\Xi(\mu_1, \mu_2)$ is the joint distribution of (x_1, x_2).

We need the following technical lemma to prove Proposition 4, and the proof of Lemma 1 is relegated to Section A.4.

Lemma 1 *Consider a measurable space X and two metric spaces Y_1 and Y_2, which induces the probability spaces $(\triangle(Y_1), \mathcal{B}[\triangle(Y_1)])$ and $(\triangle(Y_2), \mathcal{B}[\triangle(Y_2)])$. Then, for any bounded measurable function $f : X \times Y_1 \times Y_2 \to \mathbb{R}$, and any measurable functions listed as follows,*

$$\rho_1 : X \to \triangle(Y_1),$$
$$\rho_2 : X \to \triangle(Y_2),$$

the function $\Pi : X \to \mathbb{R}$ such that

$$\Pi(x) = \int f(x, y_1, y_2) d\Xi(\rho_1[x](dy_1), \rho_2[x](dy_2))$$

is measurable.

Proof of Proposition 4: Fix any $i \in I$ and any $(\theta_i, \sigma_1, \sigma_2)$. For any $a, b \in \Theta_{-i}$ such that $0 < a < b < 1$, define

$$\begin{aligned} V_{i,(a,\,b]}\left(\cdot \mid \theta_i, \sigma_1, \sigma_2\right) &: (a,b] \times \mathcal{H} \longrightarrow \mathbb{R}, \\ V_{i,(a,\,b]}\left(x, h \mid \theta_i, \sigma_1, \sigma_2\right) &\equiv V_i\left(x, h \mid \theta_i, \sigma_1, \sigma_2\right), \end{aligned}$$

i.e., $V_{i,(a,\,b]}\left(\cdot \mid \theta_i, \sigma_1, \sigma_2\right)$ denote the function of $V_i\left(\cdot \mid \theta_i, \sigma_1, \sigma_2\right)$ with the restricted domain $(a,b] \times \mathcal{H}$. It suffices to show that $V_{i,(a,\,b]}\left(\cdot \mid \theta_i, \sigma_1, \sigma_2\right)$ is measurable for any $a, b \in \Theta_{-i}$ such that $0 < a < b < 1$.[23]

From now on, we fix any $a, b \in \Theta_{-i}$ such that $0 < a < b < 1$ and show $V_{i,(a,\,b]}\left(\cdot \mid \theta_i, \sigma_1, \sigma_2\right)$ is measurable. Pick any $\delta > 0$ such that $\delta < \min\{\theta_i, a\}$. For each positive integer n, define

$$\mathcal{H}^n = \left\{h = \left[h^0, \left(t^1, I^1\right), ..., \left(t^k, I^k\right)\right] \in \mathcal{H} \diagdown H^0 : 1 - n \times \delta < t^k \leq 1 - (n-1) \times \delta\right\}.$$

Clearly, $\mathcal{H} \diagdown H^0 = \bigcup_{n=1}^{N} \mathcal{H}^n$ for some positive integer N, and

$$\begin{aligned} (a,b] \times \mathcal{H} &= \left((a,b] \times H^0\right) \cup \left((a,b] \times \left(\mathcal{H} \diagdown H^0\right)\right) \\ &= \left((a,b] \times H^0\right) \cup \left(\bigcup_{n=1}^{N} (a,b] \times \mathcal{H}^n\right). \end{aligned}$$

We first show $V_{i,(a,\,b]}\left(\cdot \mid \theta_i, \sigma_1, \sigma_2\right)$ is measurable on $\left(\bigcup_{n=1}^{N} (a,b] \times \mathcal{H}^n\right)$, and then show it is measurable on $\left((a,b] \times H^0\right)$. We prove the former inductively.

For each $\mathcal{H}^n$, we further paritition it into two parts:

$$\begin{aligned} \mathcal{H}^{(n,2)} &= \left\{ h = \left[h^0, \left(t^1, I^1\right), ..., \left(t^k, I^k\right)\right] \in \mathcal{H}^n : \begin{array}{c} \forall i \in I,\ \exists k' \leq k, \text{ such that } i \in I^{k'} \text{ and} \\ 1 - n \times \delta < t^{k'} \leq 1 - (n-1) \times \delta \end{array} \right\}, \\ \mathcal{H}^{(n,1)} &= \mathcal{H}^n \diagdown \mathcal{H}^{(n,2)}, \\ \mathcal{H}^n &= \mathcal{H}^{(n,1)} \cup \mathcal{H}^{(n,2)}, \end{aligned}$$

[23] To see this, consider a strictly decreasing sequence $\{a_1, a_{2,...}\}$ and a strictly increasing sequence $\{b_1, b_{2,...}\}$ such that $a_1 = b_1 = \frac{1}{2}$ and $\lim_{n\to\infty} a_n = 0$ and $\lim_{n\to\infty} b_n = 1$. Note that

$$\Theta_{-i} = \left(\cup_{n=1}^{\infty} (a_{n+1}, a_n]\right) \cup \left(\cup_{n=1}^{\infty} (b_n, b_{n+1}]\right).$$

Consider any measurable set $E \subset \mathbb{R}$, and the measurability of $V_i\left(\cdot \mid \theta_i, \sigma_1, \sigma_2\right)$ on any $(a,b] \times \mathcal{H}$ implies that $V^{-1}_{i,(a_{n+1},a_n]}\left(E \mid \theta_i, \sigma_1, \sigma_2\right)$ and $V^{-1}_{i,(b_n,b_{n+1}]}\left(E \mid \theta_i, \sigma_1, \sigma_2\right)$ are measurable, and as a result,

$$V_i^{-1}\left(E \mid \theta_i, \sigma_1, \sigma_2\right) = \left(\cup_{n=1}^{\infty} V^{-1}_{i,(a_{n+1},a_n]}\left(E \mid \theta_i, \sigma_1, \sigma_2\right)\right) \cup \left(\cup_{n=1}^{\infty} V^{-1}_{i,(b_n,b_{n+1}]}\left(E \mid \theta_i, \sigma_1, \sigma_2\right)\right)$$

is measurable, i.e., $V_i\left(\cdot \mid \theta_i, \sigma_1, \sigma_2\right)$ is measurable on $\Theta_{-i} \times \mathcal{H}$.

i.e., $\mathcal{H}^n$ is the set of histories in which the last jump occurs in the interval $(1 - n \times \delta, 1 - (n-1) \times \delta]$; $\mathcal{H}^{(n,2)}$ is the subset of histories in $\mathcal{H}^n$ such that both players jump in the interval $(1 - n \times \delta, 1 - (n-1) \times \delta]$; $\mathcal{H}^{(n,1)}$ is set of histories in $\mathcal{H}^n$ such that only one player jumps in this interval.

Inductively, we prove that $V_{i,(a,\, b]}(\cdot \mid \theta_i, \sigma_1, \sigma_2)$ is measurable on each $(a, b] \times \mathcal{H}^n$. We use the first two steps to show $V_{i,(a,\, b]}(\cdot \mid \theta_i, \sigma_1, \sigma_2)$ is measurable on each $(a, b] \times \mathcal{H}^1$.

Step 1: consider $n = 1$, we show $V_{i,(a,\, b]}(\cdot \mid \theta_i, \sigma_1, \sigma_2)$ is measurable on $(a, b] \times \mathcal{H}^{(1,2)}$. For any $h \in \mathcal{H}^{(1,2)}$, both players jump in the interval $(1 - \delta, 1]$. Since $\delta < \min\{\theta_i, \theta_{-i}\}$ for any $\theta_{-i} \in (a, b]$, players cannot jump anymore, and as a result, we have

$$V_{i,(a,\, b]}\left(\theta_{-i},\ h = \left[h^0, \left(t^1, I^1\right), ..., \left(t^k, I^k\right)\right] \mid \theta_i, \sigma_1, \sigma_2\right) = u_i\left(\psi\left(h\right)\right) \times \left[1 - t^k\right],$$
$$\forall \left(\theta_{-i},\ h\right) \in (a, b] \times \mathcal{H}^{(1,2)},$$

and clearly, $V_{i,(a,\, b]}(\cdot \mid \theta_i, \sigma_1, \sigma_2)$ is measurable on $(a, b] \times \mathcal{H}^{(1,2)}$.

Step 2: we show $V_{i,(a,\, b]}(\cdot \mid \theta_i, \sigma_1, \sigma_2)$ is measurable on $(a, b] \times \mathcal{H}^{(1,1)}$. For any $h \in \mathcal{H}^{(1,1)}$, only one player jump in the interval $(1 - \delta, 1]$. If one more jump occurs, we proceed to the histories in $\mathcal{H}^{(1,2)}$. We now apply Lemma 1 to show measurability of $V_{i,(a,\, b]}(\cdot \mid \theta_i, \sigma_1, \sigma_2)$ on $(a, b] \times \mathcal{H}^{(1,1)}$. Consider

$$X = (a, b] \times \mathcal{H}^{(1,1)},$$
$$Y_1 = Y_2 = [0, 1],$$
$$\rho_1\left[\left(\theta_{-i},\ h\right)\right](y_1) = \sigma_1\left(\theta_1,\ h\right),$$
$$\rho_2\left[\left(\theta_{-i},\ h\right)\right](y_2) = \sigma_2\left(\theta_2,\ h\right),$$
$$f(\left(\theta_{-i},\ h\right), y_1, y_2) = V_{i,(a,\, b]}\left(\theta_{-i}, \eta\left(h, y_1, y_2\right) \mid \theta_i, \sigma_1, \sigma_2\right) + u_i\left[\psi(h)\right] \times \left[\min\{y_1, y_2\} - \tau(h)\right],$$
$$\text{where } \eta : \mathcal{H}^{(1,1)} \times Y_1 \times Y_2 \to \mathcal{H}^{(1,2)} \text{ such that}$$
$$\eta\left(h = \left[h^0, \left(t^1, I^1\right), ..., \left(t^k, I^k\right)\right], y_1, y_2\right)$$
$$= \begin{cases} \left[h^0, \left(t^1, I^1\right), ..., \left(t^k, I^k\right), \left(t^{k+1} = y_1, I^{k+1} = \{1\}\right)\right] & \text{if } t^k < y_1 < y_2, \\ \left[h^0, \left(t^1, I^1\right), ..., \left(t^k, I^k\right), \left(t^{k+1} = y_1, I^{k+1} = \{1, 2\}\right)\right] & \text{if } t^k < y_1 = y_2, \\ \left[h^0, \left(t^1, I^1\right), ..., \left(t^k, I^k\right), \left(t^{k+1} = y_2, I^{k+1} = \{2\}\right)\right] & \text{if } t^k < y_2 < y_1, \\ \left[h^0, \left(t^1, I^1\right), ..., \left(t^k, I^k\right), \left(t^{k+1} = 1, I^{k+1} = \{1, 2\}\right)\right] & \text{if } t^k \geq \min\{y_1, y_2\}. \end{cases}$$

It is straightforward to show that ρ_1, ρ_2 and η are measurable, which further implies f is measurable, i.e., all of the assumptions in Lemma 1 hold. As a result, for any $x = (\theta_{-i}, h) \in$

$(a,b] \times \mathcal{H}^{(1,1)} = X$, we have

$$V_{i,(a,\, b]}\left(\theta_{-i}, h \mid \theta_i, \sigma_1, \sigma_2\right) = \Pi\left(x\right) = \int f(x, y_1, y_2) d\Xi\left(\rho_1\left[x\right]\left(dy_1\right), \rho_2\left[x\right]\left(dy_2\right)\right),$$

and $V_{i,(a,\, b]}\left(\cdot \mid \theta_i, \sigma_1, \sigma_2\right)$ is measurable on $(a,b] \times \mathcal{H}^{(1,1)}$.

Steps 1 and 2 prove that $V_{i,(a,\, b]}\left(\cdot \mid \theta_i, \sigma_1, \sigma_2\right)$ is measurable on each $(a,b] \times \mathcal{H}^1$.

Step 3: Assume $V_{i,(a,\, b]}\left(\cdot \mid \theta_i, \sigma_1, \sigma_2\right)$ is measurable on $(a,b] \times \left[\cup_{n=1}^{l} \mathcal{H}^n\right]$ for some positive integer l, and we will use the next two steps to show that $V_{i,(a,\, b]}\left(\cdot \mid \theta_i, \sigma_1, \sigma_2\right)$ is measurable on $(a,b] \times \mathcal{H}^{l+1}$.

Step 4: We show $V_{i,(a,\, b]}\left(\cdot \mid \theta_i, \sigma_1, \sigma_2\right)$ is measurable on $(a,b] \times \mathcal{H}^{(l+1,2)}$. For any $h \in \mathcal{H}^{(l+1,2)}$, any additional jump would lead to histories in $\left[\cup_{n=1}^{l} \mathcal{H}^n\right]$. We now apply Lemma 1 to show measurability of $V_{i,(a,\, b]}\left(\cdot \mid \theta_i, \sigma_1, \sigma_2\right)$ on $(a,b] \times \mathcal{H}^{(l+1,2)}$. Consider

$$X = (a,b] \times \mathcal{H}^{(l+1,2)},$$

$$Y_1 = Y_2 = \left[0,1\right],$$

$$\rho_1\left[\left(\theta_{-i},\ h\right)\right]\left(y_1\right) = \sigma_1\left(\theta_1,\ h\right),$$

$$\rho_2\left[\left(\theta_{-i},\ h\right)\right]\left(y_2\right) = \sigma_2\left(\theta_2,\ h\right),$$

$$f(\left(\theta_{-i},\ h\right), y_1, y_2) = V_{i,(a,\, b]}\left(\theta_{-i}, \eta\left(h, y_1, y_2\right)\ \mid \theta_i, \sigma_1, \sigma_2\right) + u_i\left[\psi(h)\right] \times \left[\min\left\{y_1, y_2\right\} - \tau(h)\right],$$

where $\eta : \mathcal{H}^{(l+1,2)} \times Y_1 \times Y_2 \to \cup_{n=1}^{l} \mathcal{H}^n$ such that

$$\eta\left(h = \left[h^0, \left(t^1, I^1\right), ..., \left(t^k, I^k\right)\right], y_1, y_2\right)$$

$$= \begin{cases} \left[h^0, \left(t^1, I^1\right), ..., \left(t^k, I^k\right), \left(t^{k+1} = y_1, I^{k+1} = \{1\}\right)\right] & \text{if } t^k < y_1 < y_2, \\ \left[h^0, \left(t^1, I^1\right), ..., \left(t^k, I^k\right), \left(t^{k+1} = y_1, I^{k+1} = \{1,2\}\right)\right] & \text{if } t^k < y_1 = y_2, \\ \left[h^0, \left(t^1, I^1\right), ..., \left(t^k, I^k\right), \left(t^{k+1} = y_2, I^{k+1} = \{2\}\right)\right] & \text{if } t^k < y_2 < y_1, \\ \left[h^0, \left(t^1, I^1\right), ..., \left(t^k, I^k\right), \left(t^{k+1} = 1, I^{k+1} = \{1,2\}\right)\right] & \text{if } t^k \geq \min\left\{y_1, y_2\right\}. \end{cases}$$

It is straightforward to show that ρ_1, ρ_2 and η are measurable, which further implies f is measurable, i.e., all of the assumptions in Lemma 1 hold. As a result, for any $x = (\theta_{-i}, h) \in (a,b] \times \mathcal{H}^{(l+1,2)} = X$, we have

$$V_{i,(a,\, b]}\left(\theta_{-i}, h \mid \theta_i, \sigma_1, \sigma_2\right) = \Pi\left(x\right) = \int f(x, y_1, y_2) d\Xi\left(\rho_1\left[x\right]\left(dy_1\right), \rho_2\left[x\right]\left(dy_2\right)\right),$$

and $V_{i,(a,\, b]}\left(\cdot \mid \theta_i, \sigma_1, \sigma_2\right)$ is measurable on $(a,b] \times \mathcal{H}^{(l+1,2)}$.

Step 5: We now show $V_{i,(a,\, b]}\left(\cdot \mid \theta_i, \sigma_1, \sigma_2\right)$ is measurable on $(a,b] \times \mathcal{H}^{(l+1,1)}$. For any $h \in \mathcal{H}^{(l+1,1)}$, any additional jump would lead to histories in $\left[\cup_{n=1}^{l} \mathcal{H}^n\right] \cup \mathcal{H}^{(l+1,2)}$. Previously steps shows that $V_{i,(a,\, b]}\left(\cdot \mid \theta_i, \sigma_1, \sigma_2\right)$ is measurable on $(a, b] \times \left(\left[\cup_{n=1}^{l} \mathcal{H}^n\right] \cup \mathcal{H}^{(l+1,2)}\right)$, and a similar argument as Step 4 impies that $V_{i,(a,\, b]}\left(\cdot \mid \theta_i, \sigma_1, \sigma_2\right)$ is measurable on $(a, b] \times \mathcal{H}^{(l+1,1)}$. Furthermore, Steps 4 and 5 prove that $V_{i,(a,\, b]}\left(\cdot \mid \theta_i, \sigma_1, \sigma_2\right)$ is measurable on $(a, b] \times \mathcal{H}^{l+1}$.

To sum, Steps 1-5 prove that $V_{i,(a,\, b]}\left(\cdot \mid \theta_i, \sigma_1, \sigma_2\right)$ is measurable on $(a, b] \times (\mathcal{H} \diagdown H^0)$. Finally, we show $V_{i,(a,\, b]}\left(\cdot \mid \theta_i, \sigma_1, \sigma_2\right)$ is measurable on $(a, b] \times H^0$. Starting from the initial history $h = \varnothing$, any additional jump would lead to a history in $(\mathcal{H} \diagdown H^0)$. Previously steps shows that $V_{i,(a,\, b]}\left(\cdot \mid \theta_i, \sigma_1, \sigma_2\right)$ is measurable on $(a, b] \times (\mathcal{H} \diagdown H^0)$, and a similar argument as in Step 4 impies that $V_{i,(a,\, b]}\left(\cdot \mid \theta_i, \sigma_1, \sigma_2\right)$ is measurable on $(a, b] \times H^0$. ■

A.4 Proof of Lemma 1

We need a series of lemmas to prove Lemma 1, and the proofs of these lemmas can be found in Appendix A.4.1.

Lemma 2 *Consider a measurable space X and a separable metric space Y, which induces the probability space $(\triangle(Y), \mathcal{B}[\triangle(Y)])$. Then, for any bounded measurable function $f : X \times Y \to \mathbb{R}$ and any measurable function $k : X \to \triangle(Y)$, the function $\Sigma : X \to \mathbb{R}$ such that*

$$\Sigma(x) = \int_Y f(x, y)\, d\left[k[x](y)\right]$$

is measurable.

Lemma 3 *Consider two independent real random variables, x and y with distributions $\mu_x \in \triangle(\mathbb{R})$ and $\mu_y \in \triangle(\mathbb{R})$, respectively. Then, the funcion $\Xi : \triangle(\mathbb{R}) \times \triangle(\mathbb{R}) \to \triangle(\mathbb{R} \times \mathbb{R})$ such that*

$$\Xi\left(\mu_x, \mu_y\right)\left[\{(x, y) : x \leq c, y \leq d\}\right] = \mu_x\left[\{x : x \leq c\}\right] \times \mu_y\left[\{y : y \leq d\}\right], \forall\ (c, d) \subset \mathbb{R} \times \mathbb{R}, \tag{15}$$

(i.e., the joint distribution) is continuous, and hence also measurable.

Lemma 4 *Consider two random variables, x and y with distributions $\mu_x \in \triangle(X)$ and $\mu_y \in \triangle(Y)$, respectively. For any two measurable functions, $g' : Z \to \triangle(X)$ and $g'' : Z \to \triangle(Y)$, define $\gamma : Z \to \triangle(X \times Y)$, such that $\gamma(z) = \Xi\left[g'(z), g''(z)\right]$ is measurable.*

Proof of Lemma 1: Consider $y = (y_1, y_2)$. Since $k_1 : X \to \triangle(Y_1)$ and $k_2 : X \to \triangle(Y_2)$ are measurable, by Lemma 4, the function $k : X \to \triangle(Y)$ such that

$$k(x) = \Xi\left(k_1[x], k_2[x]\right)$$

is measurable. Then,

$$\Pi(x) = \int f(x, y_1, y_2) d\Xi(k_1[x](y_1), k_2[x](y_2)) = \int f(x,y) d[k[x](y)],$$

is measurable, as implied by Lemma 2. ■

A.4.1 proofs of Lemmas 2-4

We need the following theorem to prove Lemma 2.

Theorem 1 (17.25 Theorem, Kechris (1995)) *Consider a measurable space X and a separable metric space Y, which induces the probability space $(\triangle(Y), \mathcal{B}[\triangle(Y)])$. Then, for any bounded measurable function $f : X \times Y \to \mathbb{R}$, the function $\digamma : X \times \triangle(Y) \to \mathbb{R}$ such that*

$$\digamma(x, \mu) = \int_Y f(x, y) d\mu(y)$$

is measurable.

Proof of Lemma 2: Since k is measurable, the function $\tau : X \to X \times \triangle(Y)$ such that $\tau(x) = (x, k[x])$ is measurable. Furthermore, by Theorem 1, the function $\digamma$ is measurable, so $\Sigma = \digamma \circ \tau$ is measurable. ■

Proof of Lemma 3: For real random variables, convergence in weak* topology is equivalent to convergence in distribution. Consider any $(\mu_x, \mu_y) \in \triangle(\mathbb{R}) \times \triangle(\mathbb{R})$, and any sequence $\{(\mu_x^n, \mu_y^n)\}$ such that $\mu_x^n \to \mu_x$ and $\mu_y^n \to \mu_y$ in distribution, i.e.,

$$\begin{aligned} \forall\ (c,d) &\subset \mathbb{R} \times \mathbb{R}, \\ \lim_{n\to\infty} \mu_x^n [\{x : x \le c\}] &= \mu_x [\{x : x \le c\}], \qquad (16) \\ \lim_{n\to\infty} \mu_y^n [\{y : y \le d\}] &= \mu_y [\{y : y \le d\}], \qquad (17) \end{aligned}$$

which immediately implies

$$\begin{aligned} \forall\ (c,d) &\subset \mathbb{R} \times \mathbb{R}, \\ \lim_{n\to\infty} \Xi(\mu_x^n, \mu_y^n) [\{(x,y) : x \le c, y \le d\}] &= \lim_{n\to\infty} (\mu_x^n [\{x : x \le c\}] \times \mu_y^n [\{y : y \le d\}]), \qquad (18) \\ &= \lim_{n\to\infty} \mu_x^n [\{x : x \le c\}] \times \lim_{n\to\infty} \mu_y^n [\{y : y \le d\}], \\ &= \mu_x [\{x : x \le c\}] \times \mu_y [\{y : y \le d\}], \qquad (19) \\ &= \Xi(\mu_x, \mu_y) [\{(x,y) : x \le c, y \le d\}], \qquad (20) \end{aligned}$$

where (18) and (20) follow from (15), and (19) follows from (16) and (17). Therefore, the sequence $\left\{\Xi\left(\mu_x^n, \mu_y^n\right)\right\}$ converges to $\Xi\left(\mu_x, \mu_y\right)$ in distribution, or equivalently, in weak* topology. ■

Proof of Lemma 4: Since g' and g'' are measurable, the function $\eta : Z \to \triangle(X) \times \triangle(Y)$ such that $\eta(z) = [g'(z), g''(z)]$ is measurable. Furthermore, by Lemma 3, the function Ξ is measurable, so $\gamma = \Xi \circ \eta$ is measurable. ■

References

Anderson, Robert. 1984. "Quick Response Equilibria," University of California, Berkeley, C.R.M. I.P. #323.

Aumann, Robert. 1964. "Mixed and Behavior Strategies in Infinite Extensive Games" In *Advances in Game Theory,* edited by M. Dresher, L. S. Shapley, and A. W. Tucker, pp. 627–650, Princeton University Press, Princeton.

Bergin, James. 1992. "A Model of Strategic Behaviour in Repeated Games." *Journal of Mathematical Economics,* **21**: 113–153.

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

Bigoni, Maria, Marco Casari, Andrzej Skrzypacz, and Giancarlo Spagnolo. 2015. "Time Horizon and Cooperation in Continuous Time" *Econometrica,* **83**: 587–616.

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

Calford, Evan, and Ryan Oprea. 2017. "Continuity, Inertia, and Strategic Uncertainty: A Test of the Theory of Continuous Time Games" *Econometrica,* **85**: 915–935.

Choi, Syngjoo, Sanjeev Goyal, and Frederic Moisan. 2020. "Network Formation in Large Groups," Cambridge Working Papers in Economics: 1935.

DeMarzo, Peter, and Yuliy Sannikov. 2006. "Optimal Security Design and Dynamic Capital Structure in a Continuous-Time Agency Model," *Journal of Finance,* **61**: 2681–2724.

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

Harsanyi, John. 1967. "Games with Incomplete Information Played by "Bayesian" Player, Part I. The Basic Model" *Management Science,* **14**: 159–182.

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

Kamada, Yuichiro, and Michihiro Kandori. 2020. "Revision Games," *Econometrica,* **88**: 1599–1630.

Kechris, Alexander. 1995. *Classic Descriptive Set Theory,* Graduate Texts in Mathematics 156, Springer, New York, USA.

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.

Kuhn, H. W., 1953. “Extensive games and the problem of information,” *Annals of Mathematics Study,* **28**: 193–216.

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.

Sannikov, Yuliy. 2007. “Games with Imperfectly Observable Actions in Continuous Time,” *Econometrica.,* **75**: 1285–1329.

Sannikov, Yuliy. 2008. “A Continuous- Time Version of the Principal: Agent Problem,” *Review of Economic Studies,* **75**: 957–984.

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