Department of Economics
University of Pittsburgh
Kenneth P. Dietrich School of Arts and Sciences

WORKING PAPER SERIES

19/005

Existence of trembling hand perfect and sequential equilibrium in games with stochastic timing of moves

Sofia Moroni

July, 2018

University of Pittsburgh • Dietrich School of Arts and Sciences • Department of Economics
4901 Wesley W. Posvar Hall • 230 South Bouquet Street • Pittsburgh, PA 15260

Existence of trembling hand perfect and sequential equilibrium in games with stochastic timing of moves.

Sofia Moroni*

University of Pittsburgh

moroni@pitt.edu

July 25, 2018

Abstract

We consider continuous time games in which players have stochastic opportunities to move before a deadline. In this paper we define notions of trembling hand and sequential equilibrium and show that both types of equilibria exist in a large class of such games that may feature incomplete and imperfect information. These games model realistic non-stationary dynamic situations in which players do not know exactly when they or their opponents will be able to move. In the complete information case we establish existence of a Markov Perfect equilibrium.

KEYWORDS: existence, revision games, stochastic games, trembling hand perfect equilibrium, sequential equilibrium.

*I am grateful to Johannes Hörner, Adam Kapor, Phil Reny, Larry Samuelson, Roee Teper, Can Urgun and Richard Van Weelden for their insightful comments.

1 Introduction

In this paper we show existence of a trembling hand and sequential equilibrium in a large class of games with stochastic timing of moves. This class encompasses games that feature incomplete and imperfect information, and games in which the actions available to the players varies with the history of moves. In particular, we consider continuous time dynamic games in which players can take actions only at times in which they receive an opportunity to move. These opportunities arrive stochastically, at Poisson rates, independently across players, up to a deadline at which the game ends. Payoffs depend on the sequence of actions taken before the end of the game. These games model realistic situations in which players cannot exactly predict when they or their opponents will move, while they know with certainty that only finitely many moves can occur in any given time interval. If the rate at which opportunities realize is large then players can move at frequent opportunities and the game approaches a continuous time game.

Due to the dimensionality of the space of histories and our assumptions regarding informational asymmetries, existing proofs of existence do not apply. To prove our results, we show that an equilibrium that is the limit of "trembles", as in Selten's notion of trembling hand perfect equilibrium, exists. The equilibrium that we find is a Nash equilibrium and a Sequential Equilibrium (as we define it to extend the notion of Kreps and Wilson).[1]

Games with stochastic timing of moves can be a tractable way of representing dynamic economic situations that are non-stationary due to an approaching deadline. To give some examples, in online auctions, players must submit bids before a deadline if they want to win the good, and make inferences about opposing players' bids and valuations as the auction proceeds. Due to delays and failures of information transmission the players cannot in practice move at arbitrary times. A player who attempts to move at the very last instant may fail to do so.[2] When labor and management of a firm bargain over pay, if the two sides do not reach a deal before a deadline a strike may occur. Politicians in congress may need to reach an agreement before a deadline at which the debt ceiling binds.[3] Candidates

[1]Our notion of sequential equilibrium coincides with Kreps and Wilson (1982) in finite games. In a revision game, we ask that strategies be sequentially rational given beliefs derived from weak limits of beliefs that correspond to weakly convergent strategies. The limit beliefs are such that players may not continue to believe that an apposing player plays according to equilibrium strategies after she sees him deviate.

[2]Ambrus et al. (2014), Hopenhayn and Saeedi (2016) , and Moroni and Kapor (2015) are examples of an stochastic move models of online auctions. Relatedly, Roth and Ockenfels (2002); Ockenfels and Roth (2006) model online auctions with discrete moving times but randomness in the realization of the final move opportunity.

[3]Ambrus and Lu (2014) consider a model of bargaining with stochastic moving times.

must choose their stance, announce their policy proposals and disclose opposition research at strategic times before an election.[4]

The class of games that we study includes the class of games known as revision games with asynchronous moves.[5] These are games in which players may revise their actions in a stage game before a deadline at which the payoffs of the stage game are realized. Kamada and Kandori (2017a) were first to introduce revision games. Our results apply to more general environments than the ones considered by Kamada and Kandori (2017a), as we allow payoffs to depend on the sequence of actions taken, available actions to vary as players move and we introduce informational asymmetries. There is a growing literature that applies revision games to economic questions. Calcagno et al. (2014) study equilibrium selection in a revision game built on an opposing interest stage game. Gensbittel et al. (2017) study revision games with an underlying zero-sum stage game. Kamada and Kandori (2017b) show that cooperation can arise in several applications of revision games, such as games of price competition, exchange of goods and election campaigns.

There are several papers that consider games with stochastic timing of moves that realize at Poisson rates. Ambrus et al. (2014) and Moroni and Kapor (2015) are models of online auctions. Hopenhayn and Saeedi (2016) consider an auction model in which players learn about their valuation. The opportunities to bid realize according to a general process that includes Poisson arrivals. Kamada and Sugaya (2014) consider a model of dynamic posturing before an election. Players may change their policy announcement at an opportunity as long as it continues to be ambiguous. Kamada and Moroni (2017) analyze a game with one stochastic move opportunity. Their results apply to opportunities that arrive at Poisson rates. In the topic of monetary policy, under the widely used Calvo pricing model, firms can revise their prices only when they receive an stochastic opportunity to do so (Calvo (1983)). In industrial organization, Doraszelski and Judd (2012) and Arcidiacono et al. (2016) propose stochastic move models which have a light computational burden relative to discrete time simultaneous move counterparts. In the dynamic stochastic evolutionary games literature players receive stochastic opportunities to revise their strategies

[4]Kamada and Sugaya (2014) and Kamada and Kandori (2017b) consider stochastic-timing models of election campaigns.

[5]In our setting each player's opportunity is drawn independently from the other players'. Therefore, the event that two or more players move simultaneously, at a given time, occurs with probability zero. However, our assumptions about observability of moves and payoffs allow us to accommodate some situations with simultaneous moves. For example, players may not observe actions taken for "intervals" of time. Since during the interval players do not know what actions the players have taken, their moves in the interval are de facto simultaneous.

(see for example Sandholm (2011)).

The paper is related to the literature on stochastic games, stemming from the seminal work by Shapley (1953). In fact, the class of games that we consider can be expressed as stochastic games that transition from one moving time to another at an exponential rate over the interval of time that is left in the game. Unlike our paper, most of the literature on stochastic games considers games of perfect information. The techniques developed in the stochastic games literature with complete information cannot be readily applied to our setting.[6]

In concurrent work, Lovo and Tomala (2015) show existence of a Markov Perfect Equilibrium in games with stochastic timing of moves with complete information. Their result holds for a more general state space and transitions between states. They also allow for asynchronous and synchronous moves. Under complete information, our approach allows us to show existence of Markov Perfect Equilibria, thus providing an alternative proof for the results in Lovo and Tomala (2015) for our setting (Corollary 1).

There is a growing literature that considers stochastic games with imperfect and incomplete information. Altman et al. (2008) study stochastic games with imperfect and incomplete information with finite actions and states. They show that, under some assumptions on the nature of the transition matrix, the set of feasible strategies and non-observability of costs, a stationary equilibrium exists. Balbus et al. (2013) show existence of a stationary Markov Nash equilibrium in a setting with private signals in stochastic games with strategic complementarities. It has also been shown that in some cases equilibria may not exist: Flesch et al. (2003) provides an example of a game with unobservable actions and observable payoffs that does not have and ε-equilibrium. Ours is the first paper, to our knowledge, to show existence of trembling hand perfect and sequential equilibria in a class of stochastic games with imperfect and incomplete information and an infinite state space.

Let us now provide an outline of our argument for existence. We identify strategies with the induced probabilities of realizing a history of play. These probabilities are functions in a L^2 space endowed with the measure of moving opportunities. Under this measure, con-

[6]In rough terms, the approach of the stochastic games literature is to define a correspondence that takes expected payoffs, as a function of observed states, to expected payoffs. It is then shown that under the assumptions the correspondence has a fixed point. An equilibrium is found as a measurable selection that yields the expected payoffs of the fixed point. Under incomplete or imperfect information, however, the players' expected payoffs are not fully determined by any observable state. An unobserved action, for instance, affects payoffs. If there is incomplete information about types, players' previous actions (which one can make part of the state) may signal an opponent's type and affect expected payoffs *endogenously* in equilibrium. See, for example, Duggan (2012) and Nowak and Raghavan (1992).

vergence in the weak topology—although weaker than other convergence notions—implies convergence of expected payoff of strategies for each player. Thus, limits of equilibria, if they correspond to some strategy, must be equilibria themselves. We then define a sequence of finite games that can approximate the revision game. These finite games have a ε-constrained equilibria which can be identified with ε-constrained strategies in the revision game.[7] Probability measures must integrate to one and are, thus, bounded in the L^2 norm. We can then appeal to the Banach–Alaoglu theorem, to show that these "identified" ε-constrained strategies must have a convergent subsequence, the limit of which corresponds to an ε-constrained equilibrium of the revision game.[8] A similar argument can now be used to establish existence of a Trembling Hand Perfect Equilibrium. The type of convergence that we consider is fairly weak. For this reason most of the work in our argument is devoted to showing that the limit probability functions correspond to some strategy.[9]

The organization of the paper is as follows. In section 2 we introduce games with stochastic timing of moves. In section 3 we define a trembling hand perfect equilibrium and a weak sequential equilibrium (3.3) and prove their existence. In section 3.4 we argue that existence of a Markov perfect equilibrium in the complete information case follows.

2 Game with stochastic timing of moves

I players participate in a game in continuous time of time length T. A player can take an action from her action set at a time $t \in [0,T]$ if she receives an opportunity to move at that time. Opportunities to take actions arrive stochastically at Poisson rates and independently across players. Each player i has $n_i \in \mathbb{N}$ possible types. Different types of players may have different payoffs and different Poisson rates at which their opportunities to move realize. The set of the possible types of player i is denoted Θ_i, and $\Theta = \bigcup_{i\in\{1,\ldots,I\}} \Theta_i$ is the set of all types of players that may be present in the game. Players have distinct types, that is $\Theta_i \cap \Theta_j = \emptyset$ if $i \neq j$. The Poisson rate of a player of type $\theta_i \in \Theta_i$ for $i \in \{1,\ldots,I\}$ is

[7] Here ε is a profile of trembles that puts a lower bound on the probability on actions at each realized history.

[8] Convergence is in the weak-* topology which coincides with the weak topology.

[9] The limit strategy may depend on actions in a manner that correlates the actions of the different players and, hence, makes it non-measurable with respect to a player's information. To show that the strategies of each player can be "de-coupled" we assume that players do not observe moving times of the opposing players coarsely. In an application this assumption implies that players may observe moving times up to any fraction of a second but cannot observe finer time measures.

denoted λ_{θ_i}.[10]

Histories A *history of opportunities* of play consists of all times at which players received an opportunity to move together with the identity of the players who received them. A generic opportunity history with n opportunities is written as $a = \big((t^1,\theta^1),\ldots,(t^n,\theta^n)\big)$, with $t^j < t^{j+1}$, where $t^j \in [0,T]$ is the timing of the j'th opportunity in a and $\theta^j \in \Theta$ is the type of player who received it. We say that history a is feasible if whenever $\theta^k, \theta^{k'} \in \Theta_i$ for $k,k' \in \{1,\ldots n\}$ then $\theta^k = \theta^{k'}$. That is, a is feasible if players have persistent types. $\mathscr{A}$ denotes the set of feasible histories of opportunities.

The number of opportunities in a is denoted $|a|$. We define $t^j(a) = t^j$ and $\theta^j(a) = \theta^j$, for $j \leq |a|$ as the j'th moving time and type of player, respectively. $\theta_i(a)$ denotes the type of player i in history a. It is given by $\theta_i(a) = \theta^j$ if $\theta^j \in \Theta_i$ for some $j \in \{1,\ldots,|a|\}$ and $\theta_i(a) = \emptyset$, otherwise. $i^j(a)$ denotes the player that moves at the j'th opportunity in history h. It is given by $i^j(a) = i$ if $\theta^j \in \Theta_i$.

At each opportunity players choose an action from their set of available actions. This set may depend on the sequence of actions taken up to the move opportunity. It cannot, however, depend on the times at which these actions were taken.[11] We assume each player takes a default action at time zero. There is a unique default action for each player.

A generic history of play is written as $h = \big((t^1,\theta^1,s^1),\ldots,(t^n,\theta^n,s^n)\big)$ where s^j denotes the action player θ^j takes at an arrival at time t^j.[12] $|h| = n$ is the number of moves in h, $t^j(h) = t^j$, $\theta^j(h) = \theta^j$, and $s^j(h) = s^j$ for $j \leq |h|$ are, respectively, the j'th time, type and action in history h. $\theta_i(h)$ denotes the type of player i in history h. $m(h) = ((s^1,\theta^1),\ldots,(s^n,\theta^n))$ denotes the sequence of actions and types of players in history h.

We allow for imperfect observability of actions. Thus, a player may not observe $m(h)$ when he receives an opportunity at time t. We write $m_i(h,t)$ for player i's observation from the sequence of actions and types at an opportunity at time t, given history h.[13]

The set of actions available to i at time t, denoted $S_i(m_i(h,t))$, is a finite set that depends on $m_i(h,t)$. The total number of actions available to each player throughout the game is finite. That is, for each i, $\mathbf{S}_i \equiv \bigcup_{t\in[0,T],h\in\mathscr{H}} S_i(m_i(h,t))$, is a finite set. This assumption

[10] That is to say, the probability that a player of type θ_i has n arrivals in an interval of length $\tilde{t}$ is given by $\left(\lambda_{\theta_i}\tilde{t}\right)^n e^{-\lambda_{\theta_i}\tilde{t}}/n!$.

[11] As an example, in a proxy-bid online auction, bidders can only place bids above the second highest bid that has been placed.

[12] The tuple (t^j,θ^j,s^j) is included in h even if player θ^j does not change her action at her opportunity at time t^j.

[13] $m_i(h,t)$ depends on the information structure defined below.

allows us to represent, among other settings, games in which each player can move only finitely many times as the set of available actions can become empty after a given history of moves.

$\mathscr{H}(a)$ denotes the set of all feasible histories of play given a history of opportunities a. It is given by the set

$$\mathscr{H}(a) = \left\{ \left(t^k(a), \theta^k(a), s^j\right)_{j=1}^{|a|} \middle| s^j \in S_{i^j(a)}\left(m_{i^j(a)}\left(\left(\left(t^k(a), \theta^k(a), s^k\right)_{k=1}^{|a|}, t\right)\right)\right) \right\}.$$

$\bar{\mathscr{H}} = \{(h,a)|h \in \mathscr{H}(a), a \in \mathscr{A}\}$ is the set of history and opportunity history pairs and $\mathscr{H} = \{h|h \in \mathscr{H}(a), a \in \mathscr{A}\}$ is the set of histories.

Information We consider games with imperfect and incomplete information.

First, each player i may imperfectly observe actions taken by other players. This possibility is represented as follows. For each player i, there is a partition P_i^s, called the *action partition*, of the opposing players' actions. At each history, player i observes the partition element, $P_i^s(s^j)$, to which an opposing player's action, s^j, belongs, not the action itself.

Second, players may not observe players' types perfectly. Each type of player $\theta_j \in \Theta_j$ belongs to an element, $P_i^\theta(\theta_j)$, of player i's *type partition,* P_i^θ, and i cannot distinguish between players in the same partition element. For $\theta_j \in \Theta_j$, $\theta_k \in \Theta_k$ with $j \neq k$ the type partition satisfies $P_i^\theta(\theta_k) \neq P_i^\theta(\theta_j)$. Each player i has a prior belief $p_{\theta_i}(\theta_1, \ldots, \theta_I)$ over player types given her type, θ_i. These beliefs are consistent with a common prior , $p(\theta_1, \theta_2, \ldots, \theta_I)$.

Third, players may imperfectly observe the identity of a player who took an action.[14] Each player has a partition over players, such that, after an action has been taken, a player only observes an element of a partition to which the player who took an action belongs. This partition depends on the action that was taken if, for instance, some actions completely or partially reveal a player's identity while others do not. Player i's *player partition* is denoted P_i^p where, for $j \neq i$, the set $P_i^p\left(P_i^\theta(\theta_j), P_i^s(s_j)\right)$ corresponds to the partition element that i observes when player of type $\theta_j \in \Theta_j$ takes action s_j. As before, player i cannot distinguish between different players in an element of partition P_i^p.[15]

Fourth, we assume that players observe a coarse measure of the timing of the opposing

[14] As an example, in an auction a player may observe that someone has placed a higher bid but cannot observe who placed the bid.

[15] Since player i only observes that player j's type belongs to type partition element $P_i^\theta(\theta_j)$ and that action s^j belongs to action partition element $P_i^s(s_j)$, the player partition must only depend on these objects.

players' moves. Each moving time, $t_k \in [0,T]$, of a player $k \neq i$, belongs to a *finite* partition P^t. This partition bounds the preciseness with which an opponent's moving time can be observed. Each player i can at most observe the corresponding element $P^t(t_k)$ when $k \neq i$ moves at time t_k. How exact is the observation of an opponent's moving time may depend on the action that the opponent took and the time at which it is observed. That is, for each player i, there is a partition P_i^t of $[0,T]$ such that when player $j \neq i$ of type θ_j takes action s_j at time t_j, player i only observes, at a time $\tilde{t} \geq t_j$, that t_j belongs to the partition element $P_i^t(t_j, P_i^s(s_j), P_i^\theta(\theta_j), \tilde{t})$ with $P_i^t(t_j, P_i^s(s_j), P_i^\theta(\theta_j), \tilde{t}) \subseteq P_i^t(t_j, P_i^s(s_j), P_i^\theta(\theta_j), \hat{t})$ for $\hat{t} \leq \tilde{t}$. Thus, players may observe an action with a lag and/or observe the timing of some actions more precisely than others. Due to the coarse observation of the opponents' moves $t', t \in P^t(t_j)$ implies $t', t \in P_i^t(t_j, P_i^s(s_j), P_i^\theta(\theta_j), \tilde{t})$ whenever $\theta_j \notin \Theta_i$. We assume players have perfect recall and, therefore, $P_i^t(t_i, P_i^s(s_i), P_i^\theta(\theta_i), \tilde{t}) = \{t_i\}$ whenever $\theta_i \in \Theta_i$ and $\tilde{t} \geq t_i$. In an application, the finite time partition assumption would hold if players observe their moving times perfectly but observe their opponents' times to the minute, the second or any arbitrarily small unit of time. In most applications the assumption is not very restrictive.[16]

Finally, we allow for the possibility that players do not observe an opponent's opportunity to move if the latter choses to take an *unobservable action* at that opportunity. For example, in an application, opposing players may not observe that a player received an opportunity if the player chooses to keep her current action. In this case, "keep action" would be an unobservable action. The set of player i's unobservable actions is denoted $\mathring{S}_i \subseteq \bigcup_{h \in \mathscr{H}, t \in [0,T]} S_i(m_i(h,t))$, and $\mathring{S}_{-i}$ denotes the set of unobservable actions of players other than i. Stochastic payoffs are included in our setting as they can be represented as unobservable moves by nature.

Beliefs over $\mathscr{A}$ Each player i has a prior belief over histories of opportunities that realize in the game given her type. This conditional prior belief may not be common to all players. For example, a player may have private information regarding her opportunity arrival rate and the distribution of these rates may be correlated. Let $\beta_{\mathscr{A}}$ denote the Borel sets of the set of opportunity histories $\mathscr{A}$. Player i's belief over opportunity histories, if she is of type $\theta_i \in \Theta_i$, is a probability measure $\mu_{\theta_i}(\cdot) : \beta_{\mathscr{A}} \to [0,1]$ that is absolutely continuous with respect to the Lebesgue measure. These beliefs are consistent with the common prior.

[16]The finite time partition assumption does not turn a stochastic timing game into a finite game, as, even under the assumption, players continue to take their own actions in continuous time.

Private Histories Given a history $h = \left((t^1, \theta^1, s^1), \ldots, (t^n, \theta^n, s^n)\right)$ with n arrivals, the private history up to (but not including) time t is denoted $h_i(h,t)$. This is the history that player i observes before taking an action at an opportunity at time t. $h_i(h,t)$ does not coincide with h if the game has imperfect information. In particular, player i may fail to observe that an action was taken or may observe a move imperfectly. In the former case, an element (t^j, θ^j, s^j) is not present in i's private history. In the latter, i observes elements of partition of actions, times, type or identity of players that are non-singletons.

Player i's private history at time t, given history h, is

$$h_i(h,t) = \left(P_i^t\left(t^j, P_i^s(s^j), P_i^\theta(\theta^j), t\right), P_i^p\left(P_i^\theta(\theta^j), P_i^s(s^j)\right), P_i^s(s^j)\right)_{j=1,\ldots,n, s^j \notin \mathring{S}_{-i}, t^j < t}.$$

Due to perfect recall, each player i's private history, h_i, contains all opportunities and moves of player i. The set of player i's time-t private histories is denoted $\mathscr{H}_i(t)$. For any given history, h, a private history of player i does not depend on her type. Thus, for every $\theta_i \in \Theta_i$, $h_{\theta_i}(h,t)$ stands for $h_i(h,t)$.

Since $m_i(h,t)$ contains events observed by player i, m_i can also be expressed as a function of $h_i \in \cup_{t \in [0,T]} \mathscr{H}_i(t) := \mathscr{H}_i$. We will sometimes write $m_i(h_i)$ for $m_i(h,t)$ with $h_i = h_i(h,t)$.

Strategies A strategy for player i of type θ_i, $\sigma_{\theta_i} : \bigcup_{t \in [0,T]} \mathscr{H}_i(t) \times \{t\} \to \Delta S_i(m_i(h,t))$, is a function from private histories and moving times to probability measures over action sets. The set of player i strategies is denoted Σ_i. A strategy profile is of the form $\sigma = \times_{i, \theta_i \in \Theta_i} \sigma_{\theta_i} \in \times_{i, \theta_i \in \Theta_i} \Sigma_i \equiv \Sigma$.[17]

Payoffs The payoff of the players after history h depends on the sequence of actions that were taken but not on the times at which they were taken. The payoff of a player of type θ is given by $g_\theta : \mathscr{H} \to \mathbb{R}$. The function g_θ depends only on the sequence of actions taken up until the end of the game. It does not depend on the times at which these actions were taken. That is, $m(h) = m(h')$ implies $g_\theta(h) = g_\theta(h')$ for every $\theta \in \Theta$, and $h, h' \in \mathscr{H}$. We assume that $|g_\theta|$ is bounded by a constant G, for every $\theta \in \Theta$.

Consider a strategy $\sigma = \times_{\theta \in \Theta} \sigma_\theta$ and a history of move opportunities $a \in \mathscr{A}$. σ induces a probability, $prob(h|a, \sigma)$, over histories (of actions) conditional on the realization of a.

[17] Although the game is infinite there is no ambiguity in referring to mixed strategies. This is because at each opportunity there is only a finite number of actions a player can take and a player gets finitely many opportunities with probability one.

For history $h = ((t^1, \theta^1, s^1), \ldots, (t^n, \theta^n, s^n))$, this probability is

$$prob(h|a,\sigma) = \Pi_{j=1}^{n} \sigma_{\theta^j}(s^j | h_{\theta^j}(h,t^j), t^j),$$

where $\sigma_{\theta^j}(s^j|h_{\theta^j}(h,t),t)$ is the probability a player of type θ^j assigns, under σ_{θ^j}, to action s^j at private history $h_{\theta^j}(h,t)$ if she receives an opportunity at time t. We call the function $prob(h|a,\sigma)$ the *probability function associated to* σ.

Type $\theta_i \in \Theta_i$'s expected payoff when players follow strategy σ is given by

$$U_{\theta_i}(\sigma) = \int_{\mathscr{A}} \sum_{h \in \mathscr{H}(a)} g_{\theta_i}(h) \cdot prob(h|a,\sigma) d\mu_{\theta_i}(a). \tag{1}$$

A game $\mathscr{G}$ with the described characteristics will be called a *stochastic timing game*.

3 Equilibrium

Each function $prob(h|a,\sigma)$, associated to a strategy $\sigma \in \Sigma$, can be viewed as a function in L^2 endowed with the measure ν_{θ_i} over $\bar{\mathscr{H}}$ for each type $\theta_i \in \Theta_i$, defined as $d\nu_{\theta_i}(h,a) = d\mu_{\theta_i}(a)$. The norm in $L^2(\bar{\mathscr{H}}, \nu_{\theta_i})$ of a function $f : \bar{\mathscr{H}} \to \mathbb{R}$ is given by

$$\|f\|_2 = \int_{a \in \mathscr{A}} \sum_{h \in \mathscr{H}(a)} |f_i(h,a)|^2 d\mu_{\theta_i}(a).$$

If $f : \bar{\mathscr{H}} \to \mathbb{R}$ has finite L^2-norm we say it is in $L^2(\bar{\mathscr{H}}, \nu_{\theta_i})$. The set of probability functions associated to strategies is bounded in the $L^2(\bar{\mathscr{H}}, \nu_{\theta_i})$ norm, as

$$\int_{\mathscr{A}} \sum_{h \in \mathscr{H}(a)} |prob(h|a,\sigma)|^2 d\mu_{\theta_i}(a) \leq 1,$$

for every $\sigma \in \Sigma$.

For $\varphi, f \in L^2(\bar{\mathscr{H}}, \nu_{\theta_i})$, we define

$$< f, \varphi >:= \int_{\mathscr{A}} \left(\sum_{h \in \mathscr{H}(a)} f(h,a) \cdot \varphi(h,a) \right) d\mu_{\theta_i}(a).$$

We say that $\{f^m\}_m$ converges to f^* in the weak topology of $L^2(\bar{\mathscr{H}}, \nu_{\theta_i})$ if for every function $\varphi \in L^2(\bar{\mathscr{H}}, \nu_{\theta_i})$

$$< f^m, \varphi > \to < f^*, \varphi > . \tag{2}$$

By the Banach-Alaoglu theorem, the space of probabilities associated to strategies is compact in the topology of weak convergence (as it is bounded in the $L^2(\bar{\mathscr{H}}, \nu_{\theta_i})$ norm). Thus, sequences of functions associated to probabilities have subsequences that converge weakly to some limit.[18]

Trembling hand equilibrium Given $\varepsilon > 0$ and $\nu \in (0, \varepsilon)$, a function $\tilde{\varepsilon}_i : S_i(m_i(h_i)) \times \mathscr{H}_i \times [0,T] \to [\nu, \varepsilon]$ is called a ε-tremble of player i, and $\tilde{\varepsilon} = (\tilde{\varepsilon}_i)_{i \in \{1,\dots,I\}}$ is called an ε-tremble profile.

Given an ε-tremble profile $\tilde{\varepsilon} = (\tilde{\varepsilon}_i)_{i \in \{1,\dots,I\}}$, a strategy profile is an $\tilde{\varepsilon}$*-constrained strategy* if at each private history $h_i \in \mathscr{H}_i$ and time $t \in [0,T]$ each player i puts weight at least $\tilde{\varepsilon}_i(s_i, h_i, t)$ on every action s_i in her set of available actions, $S_i(m_i(h_i))$. $\Sigma_i(\tilde{\varepsilon})$ denotes the set of $\tilde{\varepsilon}$-constrained strategies of player i.[19]

Definition 1. Let $\varepsilon > 0$ and let $\tilde{\varepsilon} = (\tilde{\varepsilon}_i)_{i \in \{1,\dots,I\}}$ be an ε-tremble profile. An $\tilde{\varepsilon}$*-constrained equilibrium* σ^ε is an $\tilde{\varepsilon}$-constrained strategy profile such that for every player $i \in \{1,\dots,I\}$ and $\theta_i \in \Theta_i$

$$\sigma^\varepsilon_{\theta_i} \in \operatorname{argmax} \left\{ \int_{\mathscr{A}} \sum_{h \in \mathscr{H}(a)} g_{\theta_i}(h) \cdot prob(h|a, (\sigma'_{\theta_i}, \sigma^\varepsilon_{-\theta_i})) d\mu_{\theta_i}(a) \middle| \sigma'_{\theta_i} \in \Sigma_i(\tilde{\varepsilon}) \right\}. \tag{3}$$

Consider a history h, let $H_i(h)$ be the set of tuples, (t, θ_i, s), in history h, with $\theta_i \in \Theta_i$. These are the tuples, in history h, corresponding to player i's move opportunities. We define for $\tilde{H}(h) \subseteq H_i(h)$, the function

$$\Gamma_{\theta_i}(h, \sigma_{\theta_i}, \tilde{H}(h)) := \begin{cases} \prod_{(t,\theta_i,s) \in \tilde{H}(h)} \sigma_{\theta_i}(s|h_i(h,t),t) & \text{if } \tilde{H}(h) \neq \emptyset \\ 1 & \text{otherwise} \end{cases} . \tag{4}$$

We define, also, $\Gamma_{\theta_i}(h, \sigma_{\theta_i}) := \Gamma_{\theta_i}(h, \sigma_{\theta_i}, H_i(h))$, which is the product of the probabilities with which player i takes the actions at her opportunities in history h.

[18]Notice that for each θ, ν_θ is Lebesgue measurable and the density of every finite opportunity arrival history has strictly positive density. Therefore, weakly convergent sequences converge in $L^2(\bar{\mathscr{H}}, \nu_\theta)$ for every θ, and to the same limit.

[19]We allow the lower bound on the probability on each action to depend on action and the private history, as in the standard definition, so that at "off-path" histories of the THPE (defined below) beliefs may be non-uniform over actions.

Let

$$\bar{H}_i(h,\bar{\sigma}) = \left\{ \left(t^j(h), \theta^j(h), s^j(h)\right) | \theta^j(h) = \theta_i(h), \bar{\sigma}_{\theta_i(h)}(s^{\tilde{j}}(h)|h_i(h,t),t) > 0, \forall \tilde{j} \geq j \right\}.$$

$\bar{H}_i(h,\bar{\sigma})$ contains the tuples in $H_i(h)$ corresponding to moves that occur after the last deviation from $\bar{\sigma}_{\theta_i(h)}$, to a zero probability action, in history h.

We say that sequence of strategy profiles, $\{\sigma^m\}_m$ *converges to* $\bar{\sigma}$ *weakly in strategies* if for every player i and type $\theta_i \in \Theta_i$, **(a)**

$$\int_{\mathscr{A}} \left(\sum_{h \in \mathscr{H}(a)} \Gamma_{\theta_i}\left(h, \sigma^m_{\theta_i}, \bar{H}_i(h,\bar{\sigma})\right) \cdot \varphi(h,a) \right) d\mu_{\theta_i}(a) \to \int_{\mathscr{A}} \left(\sum_{h \in \mathscr{H}(a)} \Gamma_{\theta_i}\left(h, \bar{\sigma}_{\theta_i}, \bar{H}_i(h,\bar{\sigma})\right) \cdot \varphi(h,a) \right) d\mu_{\theta_i}(a), \tag{5}$$

for every $\varphi \in L^2(\bar{\mathscr{H}}, \nu_{\theta_i})$, and **(b)** $\sigma^m_{\theta_i}(s|h_{\theta_i}(h,t),t)\mathbf{1}\{(s,h,t)|\bar{\sigma}_{\theta_i}(s|h_{\theta_i}(h,t),t) = 0\}$ converges to zero almost surely.

We will say that σ^m *converges to* $\bar{\sigma}$ *weakly in strategies in the set* $B \subseteq \bar{\mathscr{H}}$ if condition (5) holds for every $\varphi_i \in L^2(\bar{\mathscr{H}}, \nu_{\theta_i})$ with support in B and **(b)** holds.

In words, a sequence of strategy profiles, σ^m, converges weakly in strategies to a limit $\bar{\sigma}$ if **(a)** the probability of a history of play after a player's last zero probability action under $\bar{\sigma}$ in the history converges to the probability under $\bar{\sigma}$ in the weak topology, and **(b)** σ^m converges to zero almost surely on the set of histories in which $\bar{\sigma}$ is equal to zero.

Definition 2. A strategy profile σ^* is a *trembling hand equilibrium (THPE)* if there exist a sequence $(\tilde{\varepsilon}^m)_{m=1,2,\ldots}$ of ε^m-tremble profiles, with $\varepsilon^m > 0$ and $\lim_{m\to 0}\varepsilon^m = 0$, and $\tilde{\varepsilon}^m$-constrained equilibria, σ^m, such that σ^m converges to σ^* weakly in strategies.

This definition is in the spirit of the Perfect Equilibrium for finite extensive form games defined by Selten (1975). Notice that we assume weak convergence of strategies instead of pointwise convergence of behavioral strategies. In finite games both types of convergence coincide. The advantage of our choice is that weak convergence of strategies is strong enough to imply convergence of expected payoffs from strategies (Lemma 2) but weak enough that all sequences have convergent subsequences. The former condition guarantees that the limits of sequences of equilibria are equilibria (Lemma 6 and Theorem 2), and the latter, the existence of a limit.

Under our weaker type of convergence it is not straightforward that the limit of a sequence should inherit desirable properties of its approximating elements. For example, while it is fairly straightforward that a THPE is a Nash Equilibrium (Lemma 6), it is not obvious that it should satisfy a notion of sequential rationality. Our notion of weak convergence in strategies is chosen precisely to obtain this property. We address this question in section 3.3 where we define a weak sequential equilibrium, as a natural extension of Kreps and Wilson's definition for finite games to our setting, and show that a THPE is a weak sequential equilibrium.[20]

We are now ready to state the main result of the paper.

Theorem 1. *A trembling hand perfect equilibrium exists.*

The proof of existence of trembling hand perfect equilibria is done in two steps. We first show that an $\tilde{\varepsilon}$-constrained equilibrium exists in every stochastic timing game. We then show that a sequence of $\tilde{\varepsilon}$-constrained equilibria must have have a limit that is a strategy.

Establishing existence of an $\tilde{\varepsilon}$-constrained equilibrium is the bulk of the proof and is done through the following steps: 1) We define a sequence of approximating games that converge to the original game, 2) we construct the probability functions associated to the approximating game strategies, and 3) we then show that from the limit of a subsequence of these probability functions we can construct an $\tilde{\varepsilon}$-constrained equilibrium of the stochastic timing game. The approximating game strategies converge weakly in strategies to the $\tilde{\varepsilon}$-constrained equilibrium.

3.1 Existence of $\tilde{\varepsilon}$-constrained equilibria

Fix an ε-tremble profile $\tilde{\varepsilon} = (\tilde{\varepsilon}_i)_{i\in\{1,\dots,I\}}$. In this section we show a $\tilde{\varepsilon}$-constrained equilibrium exists in game $\mathscr{G}$.

3.1.1 Approximating games

Fix a stochastic game $\mathscr{G}$ and consider a finite game, $\mathscr{G}^N$, in which there are 2^N periods in which the players can move. At the beginning of the game nature draws the type, $\theta_i \in \Theta_i$ of each player $i \in \{1,\dots,I\}$.

In game $\mathscr{G}^N$ players are drawn to move at stochastic opportunities. In each period either one or no player is drawn to play. Player $i \in \{1,\dots,I\}$ is drawn to play with probability

$$p_i^N(\theta_1,\dots,\theta_I) = \frac{\frac{\lambda_{\theta_i}}{2^N} e^{\sum_{j=1}^{I} -\lambda_{\theta_j}/2^N}}{C_N(\theta_1,\dots,\theta_I)},$$

[20]The notion of sequential equilibrium was defined by Kreps and Wilson (1982).

where $C_N(\theta_1,\ldots,\theta_I)$ is the probability that either one or no player is drawn in an interval of length $1/2^N$ in the game $\mathscr{G}$.[21] $p_i^N(\theta_1,\ldots,\theta_I)$ is the probability that player i receives an opportunity in an interval $I_k^N := \left[T\frac{k-1}{2^N}, T\frac{k}{2^N}\right]$ for some $k \in \{1,\ldots,2^N\}$, conditional on the event that at most one player receives an opportunity in that time interval. No player is drawn in a period with conditional probability $\left(1-\sum_{i\in\{1,\ldots,I\}} p_i^N(\theta_1,\ldots,\theta_I)\right)$.

If player i gets an opportunity to play at stage k, the actions available are given by the actions that would be available to her in the original game given the history of play in the approximating game.[22]

We refer to $\mathscr{G}^N$ as the N'th approximating game. In game $\mathscr{G}^N$, each player's prior about each type profile coincides with the prior in game $\mathscr{G}$.

Let $\mathscr{A}^N$ denote the set of histories of arrivals in $\mathscr{G}^N$ and let $\mu_{\theta_i}^N(a^N)$ denote the probability of history a^N according to player i of type θ_i's prior. Let $h^N = \left((\tilde{t}^1,\theta^1,s^1),\ldots,(\tilde{t}^n,\theta^n,s^n)\right)$ denote a history of play in the approximating game, where $\tilde{t}^k \in \left\{\frac{T}{2^N}, 2\cdot\frac{T}{2^N},\ldots,T\right\}$ denotes the time at which the k'th move took place. Each player observes a private history $h_i^N(h^N,\tilde{t})$ when they receive an opportunity to move at time $\tilde{t}$. A private history of player i at time $\tilde{t}$ is given by,

$$h_i^N(h^N,\tilde{t}) = \left(P_i^t\left(\tilde{t}^j, P_i^s(s^j), P_i^\theta(\theta^j), \tilde{t}\right), P_i^p\left(P_i^\theta(\theta^j), P_i^s(s^j)\right), P_i^s(s^j)\right)_{j=1,\ldots n, s^j\notin\mathring{S}_{-i}, \tilde{t}^j<\tilde{t}},$$

where the information partitions are defined as in the stochastic timing game. The set of player i time-t private histories is denoted $\mathscr{H}_i^N(t)$. The set of all of i's private histories is called $\mathscr{H}_i^N$.

A strategy of player type $\theta_i \in \Theta_i$, in game $\mathscr{G}^N$, is a function $\sigma_{\theta_i}^N : \bigcup_{t\in[0,T]} \mathscr{H}_i^N(t)\times t \to \Delta S_i(m_i(h^N,t))$, where S_i and m_i are defined as in the stochastic timing game.

$g(h^N)$ is the payoff vector given history of play h^N.

Given a strategy σ^N, we can define, as in $\mathscr{G}$, the probability function associated to σ^N, $prob^N(h^N|a^N,\sigma^N)$. Let $\mathscr{H}^N(a^N)$ denote the set of feasible histories given a^N in $\mathscr{G}^N$, $\bar{\mathscr{H}}^N$ denote the history and opportunity history pairs in $\mathscr{G}^N$, and $\mathscr{H}^N$ denote the set of all histories in game $\mathscr{G}^N$.

[21] This probability is given by

$$C_N(\theta_1,\ldots,\theta_I) = e^{\sum_{j=1}^I -\lambda_{\theta_j}/2^N} + \frac{\sum_{j=1}^I \lambda_{\theta_j}}{2^N} e^{\sum_{j=1}^I -\lambda_{\theta_j}/2^N}.$$

As N grows large, $C_N(\theta_1,\ldots,\theta_I)$ converges to 1.

[22] Notice that the set of available actions in the stochastic timing game only depends on the history of actions taken and not on the times at which they were taken.

The expected payoff from strategy σ^N in the N'th approximating game is given by

$$U^N_{\theta_i}(\sigma^N) = \sum_{a^N \in \mathscr{A}^N} \sum_{h^N \in \mathscr{H}(a^N)} g_{\theta_i}(h^N) \cdot prob^N(h^N | a^N, \sigma^N) \mu^N_{\theta_i}(a^N).$$

Each approximating game $\mathscr{G}^N$ is finite and therefore for any ε-tremble profile, $\tilde{\varepsilon}$, of game $\mathscr{G}^N$ there is an $\tilde{\varepsilon}$-constrained equilibrium strategy. (Selten (1975)).[23]

3.1.2 Associated strategies

Our next step is to show that as N grows large, the payoff of any strategy in $\mathscr{G}^N$ approaches the payoff of an associated strategy in $\mathscr{G}$. First note that one can map an increasingly larger set of $\mathscr{G}^N$ histories into histories of $\mathscr{G}$. In $\mathscr{G}^N$ players move at most once in each interval of length $1/2^N$. Therefore, stochastic timing game histories in which at most one player moves in each such interval can be straightforwardly mapped to histories in $\mathscr{G}^N$. The remaining histories have probabilities that converge rapidly to zero.

Define $t^N(t) : [0,T] \to \left\{T \cdot \frac{j}{2^N}\right\}_{j=1}^{2^N}$ as $t^N(t) = T \cdot \frac{j}{2^N}$ if $t \in I^N_j$. For each history $h \in \mathscr{H}$, let $h^N(h)$ be the history in $\mathscr{G}^N$ given by $h^N(h) = \left(t^N\left(t^j(h)\right), \theta^j(h), s^j(h)\right)_{j=1}^{|h|}$.

Let $\tilde{\mathscr{H}}^N(t)$ be the set of histories in $\mathscr{G}$ such that there is at most one player that moves in each interval I^N_j, $j \in \{1, \ldots, t^N(t) \cdot 2^N/T\}$, and no player that moves in $I^N_{t^N(t)\cdot 2^N/T} \cap (-\infty, t)$.

Given a $\tilde{\varepsilon}$-constrained strategy, σ^N, in game $\mathscr{G}^N$, we define *σ^N's associated strategy* in game $\mathscr{G}$, denoted $\tilde{\sigma}^N$, as follows. For each player i and $\theta_i \in \Theta_i$, if $h \in \tilde{\mathscr{H}}^N(t)$ then $\tilde{\sigma}^N_{\theta_i}(h_i(h,t),t) = \sigma^N_{\theta_i}\left(h^N_i\left(h^N(h),t\right), t^N(t)\right)$. If $h \notin \tilde{\mathscr{H}}^N(t)$ and there is no $\tilde{h} \in \tilde{\mathscr{H}}^N(t)$ with $h_i(\tilde{h},t) = h_i(h,t)$, strategy $\tilde{\sigma}^N(h_i,t)$ assigns probability ε to every available action and takes no action with the remaining probability. Thus, $\tilde{\sigma}^N$ coincides with σ^N at histories in $\mathscr{G}^N$ that can be matched to histories of $\mathscr{G}$.

Lemma 1. *If σ^N is a strategy in $\mathscr{G}^N$ and $\tilde{\sigma}^N$ is its associated strategy in $\mathscr{G}$ then, $|U_i(\tilde{\sigma}^N) - U^N_i(\sigma^N)| \to 0$ as $N \to \infty$.*

Proofs are in the appendix.

In some cases, we can also identify strategies in the original game with strategies in the approximating game.

[23] In game $\mathscr{G}^N$ a tremble profile is defined as in game $\mathscr{G}$: Given $\varepsilon > 0$ and $\nu \in (0, \varepsilon)$, a function $\tilde{\varepsilon}_i : S_i(m_i(h^N_i)) \times \mathscr{H}^N_i \times \left\{\frac{T}{2^N}, 2 \cdot \frac{T}{2^N}, \ldots, T\right\} \to [\nu, \varepsilon]$ is called a ε-tremble of player i in game $\mathscr{G}^N$.

Let $h \in \mathscr{H}$ and define he set

$$H^N(h) := \left\{ \tilde{h} \in \mathscr{H}^N : \tilde{h} = \left(\hat{t}^j, \theta^j(h), s^j(h)\right)_{j=1}^{|h|}, \quad t^N(\hat{t}^j) = t^N(t^j(h)) \text{ for } j = 1, \ldots, |h| \right\}.$$

$H^N(h)$ is the set of histories in $\mathscr{G}^N$ in which the times of moves are in the same length-$1/2^N$ intervals as those in h, and the player types and actions at each opportunity coincide.

We say that a type $\theta_i \in \Theta_i$ stochastic timing game strategy is *constant in* I^N if $\sigma_{\theta_i}(s|h_i(h,t),t) = \sigma_{\theta_i}(s|h_i(\tilde{h},\tilde{t}),\tilde{t})$ for every $\tilde{h} \in H^N(h)$ and $\tilde{t} \in t^N(t)$. A strategy profile $\sigma = \times_{i,\theta_i \in \Theta_i} \sigma_{\theta_i}$ is constant in I^N if σ_{θ_i} is constant in I^N for each $\theta_i \in \Theta_i$, $i \in \{1, \ldots, I\}$.

Let $\tilde{\sigma}^N$ be a game $\mathscr{G}$ strategy profile that is constant in I^N. The $\mathscr{G}^N$ *strategy associated to* $\tilde{\sigma}^N$ is defined as the strategy in $\mathscr{G}^N$ such that its associated stochastic timing game strategy is $\tilde{\sigma}^N$.

3.1.3 Candidate strategies

Let $f : \bar{\mathscr{H}} \to [0,1]$ with $f(h,a) = 1$ when (h,a) is the empty history and let $\bar{\mathscr{H}}(f) = \{(h,a) : f(h,a) > 0\}$. At histories $(h,a) \in \mathscr{H}^*(f)$ we define type $\theta_i \in \Theta_i$'s *candidate strategy associated to* f, π_{θ_i}, as

$$\pi_{\theta_i}(s|h,t,f) := \begin{cases} \frac{f\left(\left(h,(t,\theta_i,s)\right),\left(a,(t,\theta_i)\right)\right)}{f(h,a)} & \text{if } \left(\left(h,(t,\theta_i,s)\right),\left(a,(t,\theta_i)\right)\right) \in \bar{\mathscr{H}} \\ 1 & \text{otherwise} \end{cases}.$$

We say that π_{θ_i} is *measurable with respect to i's information* if, at every $(h,a) \in \mathscr{H}^*(f)$, $\pi_{\theta_i}(s|h,t,f) = \pi_{\theta_i}(s|h_i(h,t),t,f)$ for $t \in [0,T]$ and $s \in S_i(m_i(h,t))$. Notice that measurability with respect to each player i's information has to be satisfied if π_{θ_i} is to be a strategy of a player of type $\theta_i \in \Theta_i$.

Let σ^m be a sequence of strategies and let $f^m(h,a) = prob(h|a,\sigma^m)$ denote their associated probability functions.

Lemma 2. *The sequence of functions $f^m(h,a)$ converges to $f^*(h,a)$ in the weak topology of $L^2(\bar{\mathscr{H}}, \nu_{\theta_i})$ for each $i \in \{1, \ldots, I\}$ and $\theta_i \in \Theta_i$, if and only if σ^m converges weakly in strategies to π in $\mathscr{H}^*(f^*)$.*

Lemma 2 shows that weak convergence of probability functions is equivalent to weak convergence in strategies in histories in $\mathscr{H}^*(f^*)$. The Lemma would be immediate under almost-sure convergence, but it is not under weak convergence. In fact, we need to use the assumption that moves are coarsely observed to ensure the limit candidate strategy does not feature correlation between unobserved actions (which can make the candidate strategy not measurable with respect to a player's information).[24] Under the "coarse observation" assumption, players cannot condition on the exact timing of moves of the opponents, and, therefore, each player's strategy converges almost surely with respect to the opponents' timing, as shown in the proof of Lemma 2. As a consequence, we obtain the following Lemma.

Lemma 3. *Suppose $\mathscr{H}^*(f) = \bar{\mathscr{H}}$, then type $\theta_i \in \Theta_i$'s candidate strategy associated to f^*, $\pi_{\theta_i}(\cdot|\cdot,\cdot,f^*)$, is measurable with respect to i's information and satisfies $\pi_{\theta_i}(\cdot|h,t,f^*) \in \Delta(S_i(m_i(h,t)))$.*

Lemma (3) says that if $\mathscr{H}^*(f) = \bar{\mathscr{H}}$, $\pi_{\theta_i}(\cdot|\cdot,\cdot,f^*)$ is a strategy of the stochastic timing game.

Fix an ε-tremble profile, $\tilde{\varepsilon}$, in game $\mathscr{G}$ and define the tremble profile in game $\mathscr{G}^N$, ε^N, as $\varepsilon_i^N\left(s_i, h_i^N, t^N\right) = \tilde{\varepsilon}_i\left(s_i, h_i^N, t^N\right)$ for each $i \in \{1,\ldots,I\}$. Let $\{\tilde{\sigma}^{*N}\}_N$ be a sequence of strategies of $\mathscr{G}$, such that each $\tilde{\sigma}^{*N}$ is associated to an ε^N-constrained equilibrium, σ^{*N}, in game $\mathscr{G}^N$. The sequence of probability functions associated to $\tilde{\sigma}^{*N}$,

$$f^N(h,a) = prob(h|a, \tilde{\sigma}^{*N}),$$

has a subsequence that converges to a function f^* in the weak-* topology (which coincides with the weak topology of L^2 by reflexivity) by the Banach–Alaoglu theorem.

[24]For example, consider a game with two players and three actions, A, B and D, in which players do not observe each others' actions but observe the opponent's moving time. D is the default action. The set of available actions is such that players can only choose between A and B at an opportunity to move, and once a player moves, she cannot change her action again. Suppose that under σ^m each player chooses action A at the first opportunity if they observe the opponent move at a time in an interval $\left[\frac{j}{m}\cdot T, \frac{j+1}{m}\cdot T\right]$ for odd j, or, absent an observation, they themselves move at such an interval. Otherwise, they choose B. Let $\mathscr{C}$ denote the set of histories in which both players receive an opportunity to move. In every subset C of $\mathscr{C}$

$$\int_C prob\left((t^1(a), \theta^1(a), A), (t^2(a), \theta^2(a), A), \ldots |a, \sigma^m\right) d\mu_{\theta_i}(a) =$$
$$\int_C prob\left((t^1(a), \theta^1(a), B), (t^2(a), \theta^2(a), B), \ldots |a, \sigma^m\right) d\mu_{\theta_i}(a) = 1/2.$$

Therefore, the candidate strategy associated to the limit does not condition on the timing of moves and puts probability $1/2$ on both players choosing A, and probability $1/2$ on both players choosing B in all histories in $\mathscr{C}$. In this case the candidate strategy associated to the limit probability does not correspond to a strategy of the stochastic timing game, as it is not measurable with respect to the players' private histories.

Define the strategy profile $\sigma^* = \times_{i,\theta_i \in \Theta_i} \sigma_{\theta_i}$ by setting $\sigma^*_{\theta_i}(s|h_i(h,t),t) = \pi_{\theta_i}(s|h_i(h,t),t,f^*)$ for each $i \in \{1,\dots,I\}$, $\theta_i \in \Theta_i$. Notice that f^* is the limit of $\tilde{\varepsilon}$-constrained strategies, for a fixed tremble profile $\tilde{\varepsilon}$ and $\mathscr{H}^*(f^*) = \bar{\mathscr{H}}$. Therefore, by Lemma 3, σ^* is a strategy of the stochastic timing game.

Let's see that σ^* is an $\tilde{\varepsilon}$-constrained equilibrium of $\mathscr{G}$. By contradiction, suppose player type $\theta_i \in \Theta_i$ has a profitable deviation from σ^*. Then there is an $\tilde{\varepsilon}$-constrained strategy $\hat{\sigma}_{\theta_i}$ such that

$$U_{\theta_i}\left(\hat{\sigma}_{\theta_i}, \sigma^*_{-\theta_i}\right) > U_{\theta_i}(\sigma^*) + 4\delta, \tag{6}$$

for some $\delta > 0$. We will show that (6) implies that, for large enough N, there is a profitable deviation to the ε^N-constrained equilibrium of $\mathscr{G}^N$, a contradiction.

The following Lemma shows that, for each N, we can construct a strategy of player type θ_i, constant in I^N, that gives θ_i a payoff that approaches her payoff from $\hat{\sigma}_{\theta_i}$, as N tends to ∞. The key to the argument is that a strategy that coincides with $\hat{\sigma}_{\theta_i}$ in histories with at most K opportunities approximates the payoff from $\hat{\sigma}_{\theta_i}$ as K goes to ∞. However, for fixed K, a strategy's payoff over histories with at most K opportunities can be written as an integral of a function from $\mathbb{R}^K$ to $\mathbb{R}$. The Lemma then follows from the fact that measurable functions from $\mathbb{R}^K$ to $\mathbb{R}$ can be approximated almost surely by simple functions, a basis of which are functions constant in I^N, for $N \in \mathbb{N}$.

Lemma 4. *For every N there is a game-$\mathscr{G}$ $\tilde{\varepsilon}$-constrained strategy of type θ_i, $\hat{\sigma}^N_{\theta_i}$, constant in I^N, such that* $\left|U_{\theta_i}\left(\hat{\sigma}^N_{\theta_i}, \tilde{\sigma}^{*N}_{-\theta_i}\right) - U_{\theta_i}\left(\hat{\sigma}_{\theta_i}, \sigma^*_{-\theta_i}\right)\right| \to 0$ *as $N \to \infty$.*

From Lemma 4, there is $N_1 > 0$ such that for each $N \geq N_1$ there is a strategy $\hat{\sigma}^N_{\theta_i}$, constant in I^N, such that

$$U_{\theta_i}\left(\hat{\sigma}^N_{\theta_i}, \tilde{\sigma}^{*N}_{-\theta_i}\right) + \delta > U_{\theta_i}\left(\hat{\sigma}_{\theta_i}, \sigma^*_{-\theta_i}\right),$$

for $N \geq N_1$.

Let $\hat{\hat{\sigma}}^N_{\theta_i}$ be the strategy in game $\mathscr{G}^N$ associated to $\hat{\sigma}^N_{\theta_i}$. By Lemma 1 there is $N_2 \geq N_1$ such that for $N \geq N_2$,

$$U^N_{\theta_i}\left(\hat{\hat{\sigma}}^N_{\theta_i}, \sigma^{*N}_{-\theta_i}\right) + \delta > U_{\theta_i}\left(\hat{\sigma}^N_{\theta_i}, \tilde{\sigma}^{*N}_{-\theta_i}\right).$$

By Lemma 1 and the weak convergence of f^m to f^*, there is $N_3 \geq N_2$ such that for

$N \geq N_3$,

$$U_{\theta_i}(\sigma^*) = \int_{\mathscr{A}} \left(\sum_{h \in \mathscr{H}(a)} f^*(h,a) \cdot g_{\theta_i}(h) \right) d\mu_{\theta_i}(a) > U^N_{\theta_i}(\sigma^{*N}) - \delta.$$

Combining with equation (6) we obtain

$$U^N_{\theta_i}(\hat{\sigma}^N_{\theta_i}, \sigma^{*N}_{-\theta_i}) > U^N_{\theta_i}\left(\sigma^{*N}\right) + \delta,$$

for $N \geq N_4$. However, this contradicts σ^{*N} is an ε^N-constrained equilibrium in game $\mathscr{G}^N$ and shows that for any ε-tremble profile, $\tilde{\varepsilon}$, in a stochastic timing game $\mathscr{G}$, an $\tilde{\varepsilon}$-constrained equilibrium exists.

Proposition 1. *Let $\varepsilon > 0$ and let $\tilde{\varepsilon}$ be a ε-tremble profile. An $\tilde{\varepsilon}$-constrained equilibrium exists.*

3.2 Existence of trembling hand equilibrium

Once we have established that each game has a $\tilde{\varepsilon}$-constrained equilibrium, for every ε-tremble profile $\tilde{\varepsilon}$, existence of a trembling hand equilibrium follows. To establish existence we show that in the limit as $\varepsilon \to 0$, $\tilde{\varepsilon}$-constrained equilibria, where each $\tilde{\varepsilon}$ is an ε-tremble, converge weakly in strategies to a limit strategy. The main challenge is finding the limit strategy after zero probability histories. Arguing recursively on the number of deviations in each history we can recover strategies after these histories.

Lemma 5. *If a stochastic timing game has an $\tilde{\varepsilon}$-constrained equilibrium for every ε-tremble profile $\tilde{\varepsilon}$, then it has a trembling hand perfect equilibrium.*

Thus, Lemma 5 concludes the proof of our main result, Theorem 1.

3.3 Nash and weak sequential equilibrium

In this section we show that a THPE is a Nash equilibrium and that it is sequentially rational with respect to beliefs.

Definition 3. A strategy profile $\sigma^* = \times_{i,\theta_i \in \Theta_i} \sigma_{\theta_i}$ is a Nash Equilibrium if

$$U_i(\sigma^*) \geq U_i(\sigma'_{\theta_i}, \sigma_{-\theta_i}),$$

for every strategy $\sigma'_{\theta_i} \in \Sigma_i$ and player $i \in \{1,\ldots,I\}$.

Lemma 6. *A trembling hand equilibrium is a Nash equilibrium.*

The proof is straightforward. There must be a sequence of ε^m-tremble profiles, $\tilde{\varepsilon}^m$, and $\tilde{\varepsilon}^m$-constrained strategies, σ^m, that converge weakly in strategies to a THPE, σ. By Lemma 2, the players' payoffs from σ^m converge to their payoffs from σ. Thus, the fact that each σ^m does not have profitable deviations immediately implies that σ cannot have profitable deviations.

We now turn to the question of whether a THPE is sequentially rational with respect to beliefs.

Given a history of opportunities $a \in \mathscr{A}$ and a player $i \in \{1,\ldots,I\}$, a strategy σ induces a probability over i's opponents' play in history h given i's actions that is given by

$$prob_{-i}(h|a,\sigma) = \prod_{j\neq i} \Gamma_{\theta_j(h)}(h,\sigma_{\theta_j(h)}).$$

We define an *assessment* as a pair $(\sigma,\tilde{\mu})$ with $\sigma \in \Sigma$ and $\tilde{\mu} = \times_{i=1,\theta_i\in\Theta_i}^{I}\tilde{\mu}_{\theta_i}$ a system of beliefs over histories of play of her opponents. The belief $\tilde{\mu}_{\theta_i}(h|h_i,t)$ is the probability that i assigns to history h after observing private history $h_i \in \mathscr{H}_i(t)$ at time t, assuming she assigns probability 1 to her actions in history h.

Let h^t and a^t denote the histories h and a truncated at the last realization before t and let $h^{\geq t} = \left(t^j(h),\theta^j(h),s^j(h)\right)_{j=1,t^j(h)\geq t}^{|h|}$ be events in h that occur starting from and including time t. Define

$$\bar{\mathscr{H}}_i(h_i,t,\theta_i) = \left\{(\hat{h},\hat{a}) \in \bar{\mathscr{H}} : h_i(\hat{h},t) = h_i, t^{|\hat{a}^t|+1}(\hat{a}) = t, \theta^{|\hat{a}^t|+1}(a) = \theta_i\right\}.$$

$\bar{\mathscr{H}}_i(h_i,t,\theta_i)$ is the set of histories that player i of type θ_i deems possible (ignoring her strategy) after observing private history h_i if she moves at time t. The belief $\tilde{\mu}_{\theta_i}(h,a|h_i,t)$ has support in $\bar{\mathscr{H}}_i(h,t,\theta_i)$. Notice that it is a belief over final nodes, not time t histories consistent with h_i (as usually defined).

For a player of type $\theta_i \in \Theta_i$, the expected payoff from assessment $(\sigma,\tilde{\mu})$ at time t is given by

$$U_{\theta_i}(\sigma_{\theta_i}|h_i,t,\tilde{\mu}) = \int_{(h,a)\in\bar{\mathscr{H}}_i(h_i,t,\theta_i)} g_{\theta_i}(h)\cdot\Gamma_{\theta_i}\left(h^{\geq t},\sigma_{\theta_i}\right)\cdot d\tilde{\mu}_{\theta_i}(h,a|h_i,t).$$

for every $h_i \in \mathscr{H}_i(t)$.

Definition 4. An assessment $(\sigma, \tilde{\mu})$ is *weakly consistent* if for each player i of type θ_i

$$d\tilde{\mu}_{\theta_i}(h, a|h_i, t) = \frac{prob_{-i}(h|a, \sigma) \cdot d\mu_{\theta_i}(a)}{\int_{(\hat{h}, \hat{a}) \in \bar{\mathscr{H}}_i(h_i, t, \theta_i)} prob_{-i}(\hat{h}^t|\hat{a}^t, \sigma) \cdot d\mu_{\theta_i}(\hat{a})},$$

if $(h, a) \in \bar{\mathscr{H}}_i(h_i, t, \theta_i)$, and zero otherwise, whenever the denominator is positive.

Thus, a belief is weakly consistent if on an on-path private history h_i, the belief that i assigns to a game history h is obtained by Bayes' rule given i's information and the strategy profile σ.

Definition 5. We say that $(\sigma, \tilde{\mu})$ is *sequentially rational* if $\forall i$, $\forall \theta_i \in \Theta_i$, $\forall h_i \in \mathscr{H}_i(t)$, $\forall \sigma'_{\theta_i} \in \Sigma_i$

$$U_{\theta_i}(\sigma_{\theta_i}|h_i, t, \tilde{\mu}) \geq U_{\theta_i}\left(\sigma'_{\theta_i}|h_i, t, \tilde{\mu}\right).$$

Thus, a strategy is sequentially rational if it is a best response, given beliefs, at every private history.

Definition 6. An assessment $(\sigma, \tilde{\mu})$ is *quasi-consistent* if there is a sequence of assessments $(\sigma^n, \tilde{\mu}^n)$, weakly consistent, with σ^n totally mixed, such that σ^n converges to σ weakly in strategies, and $\tilde{\mu}^n$ converges to $\tilde{\mu}$ in the weak topology.

Definition 7. An assessment $(\sigma, \tilde{\mu})$ is a *weak sequential equilibrium* if it is quasi-consistent and sequentially rational.

In finite games, weak sequential equilibria coincide with sequential equilibria, defined by Kreps and Wilson (1982). We call our notion "weak" to emphasize that we require weak convergences instead of almost sure.

For $(h, a) \in \bar{\mathscr{H}}_i(h_i, t, \theta_i)$, $\tilde{\mu}_{\theta_i}(h^t, a^t|h_i, t)$ is type θ_i's belief density, after observing h_i, that the realized history up to t is (h^t, a^t). In a weak sequential equilibrium we will have $\tilde{\mu}_{\theta_i}(h, a|h_i, t) = \tilde{\mu}_{\theta_i}(h^t, a^t|h_i, t) \prod_{j=1, t^j(h) \geq t}^{|h|} \tilde{\sigma}^i_{\theta^j(h)}\left(s^j(h)|h_i, t^j(h)\right)$, where $\tilde{\sigma}^i$ is a strategy profile that may not coincide with σ (as it would in a sequential equilibrium in a finite game). In fact, quasi-consistency implies that $\tilde{\sigma}^i_j = \sigma_j$ if h^t does not have any deviations by j to an action not in the support of j's strategy but may differ from σ_j at other histories. As an interpretation, under a weak sequential equilibrium, once player i observes player $j \neq i$ has deviated, i may not continue to believe that j plays according to her purported strategy.

Theorem 2. *A Trembling Hand Perfect Equilibrium is a weak sequential equilibrium.*

3.4 Markov Perfect Equilibrium

We now analyze the question of the existence of a Markov Perfect Equilibrium.

Definition 8. We say that a strategy σ_{θ_i} of a player of type θ_i is Markov if it depends only on the time left in the game and the observed actions taken by the player and the opponents. A strategy profile $\sigma = \times_{i,\theta_i \in \Theta_i} \sigma_{\theta_i}$ is Markov if σ_{θ_i} is a Markov strategy for every $\theta_i \in \Theta_i$, $i \in \{1, \ldots, I\}$.

Lemma 7. *Suppose σ^n converges to σ^* weakly in strategies and σ^n is Markov, then σ^* is Markov.*

Definition 9. We say that a strategy profile is Markov Perfect Equilibrium if it is Markov and a trembling hand perfect equilibrium.

Recall that we prove existence of an $\tilde{\varepsilon}$-constrained equilibrium in a stochastic timing game, $\mathscr{G}$, by finding $\tilde{\varepsilon}$-constrained equilibria in a sequence of approximating games, $\{\mathscr{G}^N\}_{N \in \mathbb{N}}$, and finding a limit of these equilibria. Lemma 7 implies that if the approximating equilibria are Markov Perfect, so is the limit $\tilde{\varepsilon}$-constrained equilibrium of the stochastic timing game. Also from the Lemma, if the $\tilde{\varepsilon}$-constrained equilibria in $\mathscr{G}$ are Markov so is their limit. Therefore, if each finite game $\mathscr{G}^N$ has a Markov $\tilde{\varepsilon}$-constrained equilibrium, then the stochastic timing game $\mathscr{G}$ has a Markov Perfect Equilibrium.

Markov Perfect strategies condition only on the time that is left and the actions that a player has seen her opponents take. They do not condition on the exact timing in which the past moves take place. This means that the issue we encountered with "spurious" correlations in the weak limit is no longer an issue (see Lemma 2 and footnote 24). Therefore, if the equilibria in the sequence of $\tilde{\varepsilon}$-constrained equilibria of the approximating games are Markov Perfect, we do not need to assume a coarse observation of the timings of moves, (i.e that P^t is a finite partition) to establish existence. Furthermore, the argument to show existence would accommodate synchronous moves, in which more than one player may be drawn to play at a given time.

Consider a stochastic timing game $\mathscr{G}$ with perfect information–in which types, actions and timing of moves is observed perfectly. $\mathscr{G}$ can be expressed as an stochastic game in which the states represent opportunities to move. For example, a state s can be written as $s = (i, t, a^t)$ where $i \in \{1, \ldots, I\} \cup \{\emptyset\}$, $t \in [0, T]$ and a^t is the sequence of actions taken in the game so far. In state s player i is drawn to move at time t after action history a^t. The end of the game is represented by a draw of player $\emptyset$. An N-th approximating game, $\mathscr{G}^N$, can be expressed as an stochastic game in an analogous way. In game $\mathscr{G}^N$, however, the state space is finite and, therefore, an $\tilde{\varepsilon}$-constrained equilibrium in Markov strategies exists.[25] From our previous discussion, this observation entails that a stochastic timing game with perfect information has a Markov Perfect Equilibrium.

Corollary 1. *If game $\mathscr{G}$ is a stochastic timing game with perfect information, then it has a Markov Perfect Equilibrium.*

A Appendix

PROOF OF LEMMA 1:

Let $\tilde{\mathscr{H}}^N = \left\{(h,a)|a \in \mathscr{A}, h \in \mathscr{H}(a) \cap \cup_{t\in[0,T]} \tilde{\mathscr{H}}^N(t)\right\}$ and $\left(\tilde{\mathscr{H}}^N\right)^c = \mathscr{H} \setminus \tilde{\mathscr{H}}^N$. The latter set is the set of histories in $\mathscr{G}$ in which there are two or more opportunities in an interval I_k^N for some k. Each player θ_i's measure of this set, $\nu_{\theta_i}\left(\left(\tilde{\mathscr{H}}^N\right)^c\right)$, converges to zero as N converges to ∞. To see this notice that $\nu_{\theta_i}\left(\left(\tilde{\mathscr{H}}^N\right)^c\right)$ is less than

$$2^N \left(1 - \sum_{(\theta_1,\ldots,\theta_I)\in\times_j\Theta_j} p_{\theta_i}(\theta_1,\ldots,\theta_I) \left(e^{-\sum_{j=1}^I \left(\lambda_{\theta_j}\right)/2^N} + \frac{\sum_{j=1}^N \left(\lambda_{\theta_j}\right)}{2^N} e^{-\sum_{j=1}^N \left(\lambda_{\theta_j}\right)/2^N}\right)\right),$$

where $p_{\theta_i}(\theta_1,\ldots,\theta_I)$ is i's prior belief that the type of each player j is θ_j for $j \in \{1,\ldots,I\}$, given her type is θ_i. By L'Hôpital's rule the previous expression converges to zero as N converges to ∞ and, therefore, $\nu_{\theta_i}\left(\left(\tilde{\mathscr{H}}^N\right)^c\right)$ converges to zero in N._

Let $\mathscr{A}(a^N) = \{a \in \mathscr{A}|t^N(t^k(a)) = t^k(a^N)\}$ and $\tilde{\mathscr{A}}_N = \cup_{a^N\in\mathscr{A}^N}\mathscr{A}(a_N)$. $\tilde{\mathscr{A}}_N$ is the set of histories of opportunities in which the players have at most one arrival at each interval I_k^N. Since $\nu_{\theta_i}\left(\left(\tilde{\mathscr{H}}^N\right)^c\right) \to 0$ and g is bounded, we have, for each $\theta_i \in \Theta_i$

$$|\underbrace{\int_{\mathscr{A}} \sum_{h\in\mathscr{H}(a)} prob(h|a,\tilde{\sigma}^N) g_{\theta_i}(h)\, d\mu_{\theta_i}(a)}_{U_{\theta_i}(\tilde{\sigma}^N)} - \int_{\tilde{\mathscr{A}}_N} \sum_{h\in\mathscr{H}(a)} prob(h|a,\sigma^N) g_{\theta_i}(h)\, d\mu_{\theta_i}(a)| \to 0. \qquad (7)$$

Now, let $d\mu(a|\bar{\theta})$ be the density, with respect to the Lebesgue measure, of an opportunity

[25] A simple extension of Sobel (1971) shows that game $\mathscr{G}^N$ has an $\tilde{\varepsilon}$-constrained Markov equilibrium. In fact, the expected payoff (v_δ^i in his notation) is continuous in $\tilde{\varepsilon}$-constrained strategies. To see this, notice that in $\mathscr{G}^N$ a reward arrives only at the end of the game, which is an absorbing state. Thus, as in Sobel (1971), the expected payoff can be computed from the geometric series of the Markov matrix of the stochastic game and this expected payoff is continuous, even if there is no discounting, because the geometric series of a Markov matrix with an absorbent state is convergent.

history a when the player types are in the vector $\bar{\theta} = (\theta_1, \ldots, \theta_I)$. We can write,

$$\int_{\tilde{\mathscr{A}}_N} \sum_{h \in \mathscr{H}(a)} prob(h|a, \sigma^N) \cdot g_{\theta_i}(h) d\mu_{\theta_i}(a) = \sum_{\bar{\theta} \in \times_j \Theta_j} p_{\theta_i}(\bar{\theta}) \int_{\tilde{\mathscr{A}}_N} \sum_{h \in \mathscr{H}(a)} prob(h|a, \sigma^N) \cdot g_{\theta_i}(h) d\mu(a|\bar{\theta}). \tag{8}$$

Also, we can write the payoff from σ^N to θ_i as

$$U^N_{\theta_i}(\sigma^N) = \sum_{(a^N, h^N) \in \mathscr{H}^N,} g_{\theta_i}(h^N) \cdot prob(h^N|a^N, \sigma^N) \mu^N_{\theta_i}(a^N) = \sum_{\bar{\theta} \in \times_j \Theta_j} p_{\theta_i}(\bar{\theta}) \frac{\int_{\tilde{\mathscr{A}}_N} \sum_{h \in \mathscr{H}(a)} prob(h|a, \sigma^N) \cdot g_{\theta_i}(h)\, d\mu(a|\bar{\theta})}{C_N(\bar{\theta})^{2^N}},$$

where $C_N(\bar{\theta}) = e^{\sum_{j=1}^I -\lambda_{\theta_j}/2^N} + \frac{\sum_{j=1}^I \lambda_{\theta_j}}{2^N} e^{\sum_{j=1}^I -\lambda_{\theta_j}/2^N}$.

We have $C_N(\bar{\theta})^{2^N} = e^{\sum_{j=1}^I -\lambda_{\theta_j}} \left(1 + \frac{\sum_{j=1}^I \lambda_{\theta_j}}{2^N}\right)^{2^N} \to 1$, therefore, as the support of types is finite,

$$\left| \int_{\tilde{\mathscr{A}}_N} \sum_{h \in \mathscr{H}(a)} prob(h|a, \sigma^N) \cdot g_{\theta_i}(h) d\mu_{\theta_i}(a) - U^N_{\theta_i}(\sigma^N) \right| \to 0,$$

which implies $\left| U^N_{\theta_i}(\sigma^N) - U_{\theta_i}(\tilde{\sigma}^N) \right| \to 0$.

PROOF OF LEMMA 2:

We say that $\varphi \in B \subseteq \bar{\mathscr{H}}$ if $\varphi \in L^2(\bar{\mathscr{H}}, \nu_{\theta_i})$, for every i, and φ has support on B. We say that f^m converges to f^* in B if $< f^m, \varphi > \to < f^*, \varphi >$ for every $\varphi \in B$.

We will show that **(i)** $f^m(h,a)$ converges to $f^*(h,a)$ in the weak topology of $L^2(\bar{\mathscr{H}}, \nu_{\theta_i})$ for each $i \in \{1, \ldots, I\}$ if and only if

$$\int_{\mathscr{A}} \sum_{h \in \mathscr{H}(a)} \left(\Gamma_{\theta_i(a)}(h, \sigma^m_{\theta_i(a)}) - \Gamma_{\theta_i(a)}(h, \pi_{\theta_i(a)}) \right) \varphi(h,a)\, d\mu_{\theta_i}(a) \to 0, \tag{9}$$

and,

$$\int_{\mathscr{A}} \sum_{h \in \mathscr{H}(a)} \left(\left(\prod_{k \neq i} \Gamma_{\theta_k(a)} \left(h, \sigma^m_{\theta_k(a)} \right) \right) - \left(\prod_{k \neq i} \Gamma_{\theta_k(a)} \left(h, \pi_{\theta_k(a)} \right) \right) \right) \varphi(h,a)\, d\mu_{\theta_i}(a) \to 0, \tag{10}$$

for every $\varphi \in \mathscr{H}^*(f^*)$, $i \in \{1, \ldots, I\}$ and $\theta_i \in \Theta_i$. We then show **(ii)**: conditions (9) and (10) are equivalent to σ^m converges weakly in strategies to π in $\mathscr{H}^*(f^*)$.

Proof of (i)

We argue recursively on the number of opportunities. Let $\mathscr{A}^k$ denote the set of opportunity histories in which players obtain at most k opportunities. Let $\bar{\mathscr{H}}^k = \{(h,a) : a \in \mathscr{A}^k, h \in \mathscr{H}(a)\}$.

It is immediate that the result holds for $\varphi \in \bar{\mathscr{H}}^1 \cap \mathscr{H}^*(f^*)$. That is, if $< f^m, \varphi > \to < f^*, \varphi >$ holds for $\varphi \in \bar{\mathscr{H}}^1 \cap \mathscr{H}^*(f^*)$ if and only if (9) and (10) hold for $\varphi \in \bar{\mathscr{H}}^1 \cap \mathscr{H}^*(f^*)$.

Suppose that $< f^m, \varphi > \to < f^*, \varphi >$ for $\varphi \in \bar{\mathscr{H}}^k \cap \mathscr{H}^*(f^*)$ if and only if (9) and (10) hold for $\varphi \in \bar{\mathscr{H}}^k \cap \mathscr{H}^*(f^*)$. We now show the equivalence must hold also for $\varphi \in \bar{\mathscr{H}}^{k+1} \cap \mathscr{H}^*(f^*)$.

Let $\bar{\mathscr{H}}_i^{k+1}$ ($\bar{\mathscr{H}}_{-i}^{k+1,}$) denote the subset of $\bar{\mathscr{H}}^{k+1}$ in which player i receives (does not receive) the last opportunity. From the induction hypothesis, for each $i \in \{1,\dots,I\}$, (9) holds if $\varphi \in \bar{\mathscr{H}}_{-i}^{k+1} \cap \mathscr{H}^*(f^*)$ and (10) holds for $\varphi \in \bar{\mathscr{H}}_i^{k+1} \cap \mathscr{H}^*(f^*)$.

We now show that (9) holds for $\varphi \in \bar{\mathscr{H}}_i^{k+1} \cap \mathscr{H}^*(f^*)$ for each $i \in \{1,\dots,I\}$. From the definition of weak convergence in equation (2), f^m converges to f^* in $\bar{\mathscr{H}}_i^{k+1} \cap \mathscr{H}^*(f^*)$, if and only if for every $\varphi \in \bar{\mathscr{H}}_i^{k+1} \cap \mathscr{H}^*(f^*)$,

$$\int \sum_{h \in \mathscr{H}(a)} \left(\left(\prod_{j \neq i} \Gamma_{\theta_j(a)}(h, \sigma^m_{\theta_j(a)}) \right) \cdot \Gamma_{\theta_i(a)}(h, \sigma^m_{\theta_i(a)}) - \left(\prod_{j \neq i} \Gamma_{\theta_j(a)}(h, \pi_{\theta_j(a)}) \right) \cdot \Gamma_{\theta_i(a)}(h, \pi_{\theta_i(a)}) \right) \cdot \varphi(h,a)\, d\mu_{\theta_i}(a) \to 0. \tag{11}$$

To simplify notation, in what follows, we omit the dependence of $\theta_j(a)$ on a.

$\left(\prod_{j \neq i} \Gamma_{\theta_j(a)}(h, \sigma^m_{\theta_j(a)}) \right)$ is the product of strategies that depend on private histories with at most k arrivals (since player i moves last in histories in $\bar{\mathscr{H}}_i^{k+1}$). By the induction hypothesis we have,

$$\int \sum_{h \in \mathscr{H}(a)} \left(\left(\prod_{j \neq i} \Gamma_{\theta_j}(h, \sigma^m_{\theta_j}) - \prod_{j \neq i} \Gamma_{\theta_j}(h, \pi_{\theta_j}) \right) \cdot \Gamma_{\theta_i}(h, \pi_{\theta_i}) \right) \cdot \varphi(h,a)\, d\mu_{\theta_i}(a) \to 0. \tag{12}$$

Combining equations (11) and (12) we obtain f^m converges to f^* in $\bar{\mathscr{H}}_i^{k+1} \cap \mathscr{H}^*(f^*)$ if and only if for every $\varphi \in \bar{\mathscr{H}}_i^{k+1} \cap \mathscr{H}^*(f^*)$,

$$\int \sum_{h \in \mathscr{H}(a)} \left(\left(\prod_{j \neq i} \Gamma_{\theta_j}(h, \sigma^m_{\theta_j}) \right) \left(\Gamma_{\theta_i}(h, \sigma^m_{\theta_i}) - \Gamma_{\theta_i}(h, \pi_{\theta_i}) \right) \right) \cdot \varphi(h,a)\, d\mu_{\theta_i}(a) \to 0. \tag{13}$$

Since the game has a finite time partition, players do not observe the exact timing of moves of opposing players. If player i moves at time $\tilde{t}$ a player $j \in -i := \{1,\dots,I\} \setminus \{i\}$ who moves after $\tilde{t}$ can at most observe that $\tilde{t}$ is in a set $P \in P^t(\tilde{t})$. This means that the strategies of players in $-i$ must be constant in $\tilde{t} \in P \in P^t(\tilde{t})$. Consider a history $(h,a) \in$

$\bar{\mathscr{H}}_i^{k+1} \cap \mathscr{H}^*(f^*)$. Let

$$H^{k+1}(h,a) = \left\{ (\tilde{h},\tilde{a}) \in \bar{\mathscr{H}}_i^{k+1} \cap \mathscr{H}(f^*) : \tilde{h} = \left(\tilde{t}, \theta^l(h), s^l(h)\right)_{l=1}^{k+1}, \tilde{t} \in P^t\left(t^l(h)\right), l = 1,\dots,k+1 \right\}.$$

$H^{k+1}(h,a)$ is the set of all histories with $k+1$ opportunities in which the l'th player to receive an opportunity, is player type $\theta^l(h)$, moves at time $t \in P^t(t^l(h))$ and takes action $s^l(h)$ for $l \in \{1,\dots,k+1\}$. Each player j is unable to distinguish between the histories in $H^{k+1}(h,a)$ in which the opposing players' moving times vary while her times remain fixed. Thus, by measurability with respect to private histories, player j's strategies must be constant in each of these subsets of histories of $H^{k+1}(h,a)$.

The time partition is finite and there are finitely many actions, therefore, there is a finite subset of $\bar{\mathscr{H}}_i^{k+1} \cap \mathscr{H}^*(f^*)$, denoted $\hat{H}$, such that $\bar{\mathscr{H}}_i^{k+1} \cap \mathscr{H}^*(f^*) = \bigcup_{(h,a)\in\hat{H}} H^{k+1}(h,a)$. Let $A^{k+1}(h,a) = \{\tilde{a} : (\tilde{h},\tilde{a}) \in H^{k+1}(h,a)\}$ and let $A_i^{k+1}(h,a)$ ($A_{-i}^{k+1}(h,a)$) be i's (players' in $-i$) move opportunities in opportunity histories in $A^{k+1}(h,a)$.

We can restate condition (13) as f^m converges to f^* in $\bar{\mathscr{H}}_i^{k+1} \cap \mathscr{H}^*(f^*)$ if and only if for every $\varphi \in \bar{\mathscr{H}}_i^{k+1} \cap \mathscr{H}^*(f^*)$, $(\hat{h},\hat{a}) \in \hat{H}$, and $\eta > 0$ there is $\bar{m}$ such that for all $m \geq \bar{m}$,

$$\left| \int_{A^{k+1}(\hat{h},\hat{a})} \left(\Gamma_{\theta_i}(h(a),\sigma_{\theta_i}^m) - \Gamma_{\theta_i}(h(a),\pi_{\theta_i})\right) \left(\prod_{j\neq i} \Gamma_{\theta_j}(h(a),\sigma_{\theta_j}^m) \right) \cdot \varphi(h(a),a)\, d\mu_{\theta_i}(a) \right| < \frac{\eta}{2}, \tag{14}$$

where $h(a)$ is the unique h such that $(h,a) \in H^{k+1}(\hat{h},\hat{a})$. $h(a)$ is unique because all histories in $H^{k+1}(\hat{h},\hat{a})$ feature the same actions.

We have

$$\int_{A^{k+1}(\hat{h},\hat{a})} \left(\Gamma_{\theta_i}(h(a),\sigma_{\theta_i}^m) - \Gamma_{\theta_i}(h(a),\pi_{\theta_i})\right) \left(\prod_{j\neq i} \Gamma_{\theta_j}(h(a),\sigma_{\theta_j}^m) - \prod_{j\neq i} \Gamma_{\theta_j}(h(a),\pi_{\theta_j}) \right) \cdot \varphi(h(a),a)\, d\mu_{\theta_i}(a) = \tag{15}$$

$$\int_{A_i^{k+1}(\hat{h},\hat{a})} \left(\Gamma_{\theta_i}(h(a),\sigma_{\theta_i}^m) - \Gamma_{\theta_i}(h(a),\pi_{\theta_i})\right) \left(\int_{A_{-i}^{k+1}(\hat{h},\hat{a})} \left(\prod_{j\neq i} \Gamma_{\theta_j}(h(a),\sigma_{\theta_j}^m) - \prod_{j\neq i} \Gamma_{\theta_j}(h(a),\pi_{\theta_j}) \right) \cdot \varphi(h(a),a)\, d\mu_{\theta_i}(a_{-i}) \right) d\mu_{\theta_i}(a_i),$$

where $d\mu_{\theta_i}(a_i)$ integrates over i's opportunity realizations, $a_i \in A^{k+1}(\hat{h},\hat{a})$, and $d\mu_{-i}(a_{-i})$ over the realizations of players other than i.[26] The last equality follows from $(\Gamma_i(h(a),\sigma_i^m) - \Gamma_i(h(a),\pi_i))$ constant in a_{-i}, for each m at histories in $H^{k+1}(\hat{h},\hat{a})$.

[26] In the incomplete information case we cannot separate de differentials in the integrals as in equation (15). The correct expression, however, would be a sum of weighted integrals of the same form (see, for example, equation (8)).

From the induction hypothesis there is $\bar{\bar{m}}$ such that for $m \geq \bar{\bar{m}}$

$$\left|\left(\int_{A_{-i}^{k+1}(\hat{h},\hat{a})}\left(\prod_{j\neq i}\Gamma_{\theta_j}(h(a),\sigma_{\theta_j}^m)-\prod_{j\neq i}\Gamma_{\theta_j}(h(a),\pi_{\theta_j})\right)\cdot\varphi(h(a),a)\,d\mu_{\theta_i}(a_{-i})\right)\right|<\frac{\eta}{2}, \tag{16}$$

almost surely in $a_i \in A_i^{k+1}(\hat{h},\hat{a})$, because $\left(\prod_{j\neq i}\Gamma_{\theta_j}(h,\sigma_{\theta_j}^m)\right)$ and $\left(\prod_{j\neq i}\Gamma_{\theta_j}(h,\pi_{\theta_j})\right)$ are constant in $a_i \in A_i^{k+1}(\hat{h},\hat{a})$ (the former because σ^m is a strategy profile and the latter because of the induction hypothesis). Notice also that equation (16) holds almost surely in $a_i \in A_{-i}^{k+1}(\hat{h},\hat{a})$ for every $(\hat{h},\hat{a})$ and, therefore, holds almost surely in $a_i \in \cup_{(\hat{h},\hat{a})}A_{-i}^{k+1}(\hat{h},\hat{a})$. From $\left|\Gamma_{\theta_i}(h(a),\sigma_{\theta_i}^m)-\Gamma_{\theta_i}(h(a),\pi_{\theta_i})\right|\leq 1$, this implies that the absolute value of equation (15) is less than $\eta/2$ for $m \geq \bar{\bar{m}}$, which together with equation (14), yields

$$\left|\int_{A^{k+1}(\hat{h},\hat{a})}\left(\Gamma_{\theta_i}(h(a),\sigma_{\theta_i}^m)-\Gamma_{\theta_i}(h(a),\pi_{\theta_i})\right)\left(\prod_{j\neq i}\Gamma_{\theta_j}(h(a),\pi_{\theta_j})\cdot\varphi(h(a),a)\,d\mu_{\theta_i}(a_{-i})\right)d\mu_{\theta_i}(a_i)\right|<\eta, \tag{17}$$

for $m \geq \max\{\bar{m},\bar{\bar{m}}\}$. Notice that $\prod_{j\neq i}\Gamma_{\theta_j}(h,\pi_{\theta_j})>0$ in $A^{k+1}(\hat{h},\hat{a})$ since, by definition $H^{k+1}(\hat{h},\hat{a})\subseteq\mathscr{H}^*(f^*)$ and $((h,(\cdot)),(a,(\cdot)))\in\mathscr{H}^*(f^*) \implies (h,a)\in\mathscr{H}^*(f^*)$. Therefore, the last expression implies condition (9).

By an analogous argument we can show that condition (10) holds for $\varphi \in \bar{\mathscr{H}}_{-i}^{k+1}\cap\mathscr{H}^*(f^*)$. We need to show that

$$\left|\int\sum_{h\in\mathscr{H}(a)}\left(\left(\prod_{j\neq i}\Gamma_{\theta_j}(h,\sigma_{\theta_j}^m)\right)-\left(\prod_{j\neq i}\Gamma_{\theta_j}(h,\pi_{\theta_j})\right)\right)\varphi(h,a)\,d\mu_{\theta_i}(a)\right|\to 0, \tag{18}$$

From the induction hypothesis,

$$\left|\int\sum_{h\in\mathscr{H}(a)}\left(\left(\prod_{j\neq i}\Gamma_{\theta_j}(h,\pi_{\theta_j})\right)\Gamma_{\theta_i}(h,\sigma_{\theta_i}^m)-\left(\prod_{j\neq i}\Gamma_{\theta_j}(h,\pi_{\theta_j})\right)\Gamma_{\theta_i}(h,\pi_{\theta_i})\right)\varphi(h,a)\,d\mu_{\theta_i}(a)\right|\to 0.$$

Combining the previous expression with (11) we obtain that f^m converges weakly to f^* in $\bar{\mathscr{H}}_{-i}^{k+1}\cap\mathscr{H}^*(f^*)$ if and only if,

$$\left|\int\sum_{h\in\mathscr{H}(a)}\Gamma_{\theta_i}(h,\sigma_{\theta_i}^m)\left(\left(\prod_{j\neq i}\Gamma_{\theta_j}(h,\sigma_{\theta_j}^m)\right)-\left(\prod_{j\neq i}\Gamma_{\theta_j}(h,\pi_{\theta_j})\right)\right)\varphi(h,a)\,d\mu_{\theta_i}(a)\right|\to 0.$$

Arguing as before, because players' moves are only coarsely observed, by the induction hypothesis,

$$\left| \left(\int \left(\Gamma_{\theta_i}(h, \sigma^m_{\theta_i}) - \Gamma_{\theta_i}(h, \pi_{\theta_i}) \right) \cdot \varphi(h,a)\, d\mu_{\theta_i}(a_i) \right) \right| \to 0,$$

almost surely in a_{-i} for a_{-i} such that $(h,(a_i,a_{-i})) \in \bar{\mathscr{H}}^{k+1}_{-i} \cap \mathscr{H}^*(f^*)$. Thus, (18) follows and the result holds for arrival histories with $k+1$ events.

Finally, it is clear from equation (15) that if conditions (9) and (10) hold (as (10) implies (16)), equation (15) converges to zero in m and, therefore, f^m converges to f^* in $\bar{\mathscr{H}}_i \cap \mathscr{H}^*(f^*)$. An analogous argument shows that f^m converges to f^* in $\bar{\mathscr{H}}_{-i} \cap \mathscr{H}^*(f^*)$.

Proof of (ii)

Let us now see why (ii) holds. Weak convergence in strategies in $\mathscr{H}^*(f^*)$ follows immediately from equation (9) for each $i \in \{1,\dots,I\}$. The converse is true by an argument analogous the proof of **(i)**, using induction and the coarse observation of the opponents' moves, which implies that each Γ_{θ_i} converges weakly in i's moving times but almost surely in the opponents' moving times.

PROOF OF LEMMA 3:

Since f^N converges to f^*in weak strategies then from Lemma 2, $\tilde{\sigma}^{*N}$ converges weakly in strategies to π. The former implies that the probability over actions under π_{θ_i} sums to 1 at each private history. The latter implies that π_{θ_i} is measurable with respect to i's information.

PROOF OF LEMMA 4:

Fix $\delta > 0$ a player $i \in \{1,\dots,I\}$ and $\theta_i \in \Theta_i$.

Because players' opportunities arrive at Poisson rates, the probability that a player observes more than K moves in her private history becomes vanishingly small as K tends to ∞. Let $\mathscr{H}_i^K$ denote the set of private histories of player i that have at most K observations of moves, and $\bar{\mathscr{H}}^K$ the set of opportunity history and history pairs with at most K moves. Let $\hat{\sigma}^K_{\theta_i}$ be a player i strategy that coincides with $\hat{\sigma}_{\theta_i}$ at histories in $\mathscr{H}_i^K$ and puts equal weight on all available actions at all other histories. $\hat{\sigma}^K_{\theta_i}$ is constant in I^N at histories in $\mathscr{H}_i \setminus \mathscr{H}_i^K$, for all N. The measure $\nu_{\theta_i}\left(\left(\bar{\mathscr{H}}^K\right)^c\right)$ converges to zero as $K \to \infty$. Therefore, there is $\bar{K}$ large enough that $\left| U_{\theta_i}\left(\hat{\sigma}^K_{\theta_i}, \sigma^*_{-\theta_i}\right) - U_{\theta_i}\left(\hat{\sigma}_{\theta_i}, \sigma^*_{-\theta_i}\right) \right| \leq \delta/3$ for $K \geq \bar{K}$. We now fix K such that the previous inequality holds.

Since $\tilde{\sigma}^N_{\theta_k}$ converges to $\sigma^*_{\theta_k}$ weakly in strategies for $\theta_k \in \Theta_k$, from condition (10) in the proof of Lemma 2, there is N_1 such that for all $N \geq N_1$, $\left| U_{\theta_i}\left(\hat{\sigma}^K_{\theta_i}, \sigma^*_{-\theta_i}\right) - U_{\theta_i}\left(\hat{\sigma}^K_{\theta_i}, \tilde{\sigma}^N_{-\theta_i}\right) \right| \leq \delta/3$.

The private histories in $\mathscr{H}_i^K$ can be represented as elements of $[0,T]^{K\cdot I\cdot|\mathbf{S}|}$, where $|\mathbf{S}|$ is the total number of available actions to all players, by keeping track of the timings of opportunities of each player and their actions. The probability function associated to $\hat{\sigma}_{\theta_i}^K$, restricted to histories in $\mathscr{H}_i^K$, is a measurable function from $[0,T]^{K\cdot I\cdot|\mathbf{S}|}$ to $\mathbb{R}$. Thus, it can be approximated almost surely by simple functions that are constant in rectangles in which each side corresponds to an interval $I_k^N = [T\frac{k}{2^N}, T\frac{k+1}{2^N}]$ for $k \in \{1,\ldots,2^N\}$. More precisely, there is a sequence of candidate strategies $\hat{\sigma}_{\theta_i}^{K,N}$, constant in I^N, that converges almost surely to $\hat{\sigma}_{\theta_i}^K$ in histories with at most K events. In histories with more than K events $\hat{\sigma}_{\theta_i}^{K,N}$ is defined to coincide with $\hat{\sigma}_{\theta_i}^K$. The sequence $\hat{\sigma}_{\theta_i}^{K,N}$ can be chosen to be of $\tilde{\varepsilon}$-constrained strategies: if, for instance, it turns out $\sum_{s\in S_i(m_i(h,t))} \hat{\sigma}_{\theta_i}^{K,N}(s|h_i(h,t),t) > 1$, we can add a function (constant in I^N) that converges to zero almost surely to it and obtain $\sum_{s\in S_i(m_i(h,t))} \hat{\sigma}_{\theta_i}^{K,N}(s|h_i(h,t),t) = 1$.[27] An analogous construction ensures that the strategy is $\tilde{\varepsilon}$-constrained. Thus, there is N_2 such that for $N \geq N_2$ such that $\left|U_{\theta_i}\left(\hat{\sigma}_{\theta_i}^K, \tilde{\sigma}_{-\theta_i}^N\right) - U_{\theta_i}\left(\hat{\sigma}_{\theta_i}^{K,N}, \tilde{\sigma}_{-\theta_i}^N\right)\right| \leq \delta/3$.

Combining the inequalities we obtain that for $N \geq \max\{N_1, N_2\}$, $\left|U_{\theta_i}\left(\hat{\sigma}_{\theta_i}, \sigma_{-\theta_i}^*\right) - U_{\theta_i}\left(\hat{\sigma}_{\theta_i}^{K,N}, \tilde{\sigma}_{-\theta_i}^N\right)\right| \leq \delta$.

PROOF OF LEMMA 5:

Let $\tilde{\varepsilon}^m$ be a sequence of ε^m-tremble profiles with $\varepsilon^m \to 0$ and let σ^m be a sequence of $\tilde{\varepsilon}^m$-tremble profiles. For each $\theta_i \in \Theta_i$, the sequence $\Gamma_{\theta_i}(h, \sigma_{\theta_i}^m)$ must have a weakly convergent subsequence. Passing to the subsequence, let $\Gamma_{\theta_i}^*(h)$ denote its limit. We can define the candidate strategy associated to Γ^* as

$$\pi_{\theta_i}(s|h,t,\Gamma^*) = \frac{\Gamma^*\Big(\big(h,(t,\theta_i,s)\big)\Big)}{\Gamma^*(h)},$$

if $\big(h,(t,\theta_i,s)\big) \in \mathscr{H}$, and $\pi_{\theta_i}(s|h,t,\Gamma^*) = 1$ if $\big(h,(t,\theta_i,s)\big) \notin \mathscr{H}$, whenever $\Gamma^*(h) > 0$.

If $\Gamma^*(h) > 0$, then $\Gamma_{\theta_i}(h, \pi_{\theta_i}(\cdot|\cdot,\cdot,\Gamma^*)) = \Gamma_{\theta_i}^*(h)$. In the set of histories such that $\Gamma^*(h) > 0$ and $\Gamma^*\Big(\big(h,(t,\theta_i,s)\big)\Big) = 0$, $\sigma_{\theta_i}^m(s|h_i(h,t),t) \to 0$ almost surely.

For an arbitrary history $h \in \mathscr{H}$, let h^j denote the history h truncated at and including the j'th move. Let $H^0 = \mathscr{H}$, $H^1 = \{h \in \mathscr{H} | \exists j < |h|, \Gamma^*(h^j) = 0\}$ and $\bar{H}^1 = \{h \in \mathscr{H} | \Gamma^*(h) =$

[27] We can add probability $S(s,h,t) = 1 - \sum_{s\in S_i(m_i(h,t))} \hat{\sigma}_{\theta_i}^{K,N}(s|h_i(h,t),t)$ to an action that is assigned weight less than $1 - S(s,h,t)$ if $S(s,h,t) > 0$ or subtract $|S(s,h,t)|$ from an action with weight greater than $-S(s,h,t)$ if $S(s,h,t) < 0$. These operations can be done in a manner that $\hat{\sigma}_{\theta_i}^{K,N}$ remains measurable.

$0, \Gamma^*(h^j) > 0$ for $j < |h|\}$. For $h \in H^1$ define

$$j^1(h) = \min\left\{ j \in \{1,\ldots,|h|-1\} | \Gamma^*_{\theta_i}\left(h^j\right) = 0 \right\}.$$

$j^1(h)$ can be interpreted as the first move of θ_i in which she deviates from $\pi_{\theta_i}(\cdot|\cdot,\cdot,\Gamma^*)$ to a zero probability action in history $h \in H^1$.

Now define for $h \in \bar{H}^1 \cup H^1$,

$$\Gamma^1_{\theta_i}(h,\sigma^m_{\theta_i}) = \begin{cases} \Pi^{|h|}_{j>j^1(h),\theta^j(h)=\theta_i} \sigma^m_{\theta^j(h)}\left(s^j(h)|h_{\theta^j(h)}\left(h,t^j(h)\right),t^j(h)\right) & \text{if } h \in H^1 \\ 1 & \text{if } h \in \bar{H}^1 \end{cases}.$$

$\Gamma^1_{\theta_i}(h,\sigma^m_{\theta_i})$ has a limit point in the weak topology, $\Gamma^{1,*}_{\theta_i}(h)$. From $\Gamma^{1,*}_{\theta_i}(h)$ we can define the candidate strategy $\pi_{\theta_i}(s|h,t,\Gamma^{1,*})$, at every $\left(h,(t,\theta_i,s)\right) \in H^1$ such that $\Gamma^{1,*}_{\theta_i}(h) > 0$.

Recursively, having defined $\Gamma^{n-1,*}_{\theta_i}$, let $H^n = \{h \in \mathscr{H} | \exists j < |h|, \Gamma^{n-1,*}(h^j) = 0\}$, $\bar{H}^n = \{h \in \mathscr{H} | \Gamma^{n-1,*}(h) = 0, \Gamma^{n-1,*}(h^j) > 0$ for $j < |h|\}$

$$j^n(h) = \min\left\{ j \in \{1,\ldots,|h|-1\} | \Gamma^{n-1,*}_{\theta_i}\left(h^j\right) = 0 \right\},$$

and

$$\Gamma^n_{\theta_i}(h,\sigma^m_{\theta_i}) = \begin{cases} \Pi^{|h|}_{j>j^n(h),\theta^j(h)=\theta_i} \sigma^m_{\theta^j(h)}\left(s^j(h)|h_{\theta^j(h)}\left(h,t^j(h)\right),t^j(h)\right) & \text{if } h \in H^n \\ 1 & \text{if } h \in \bar{H}^n \end{cases}.$$

As before, let $\Gamma^{n,*}_{\theta_i}$ denote a limit point of $\Gamma^n_{\theta_i}$. From $\Gamma^{n,*}_{\theta_i}(h)$ we can define the candidate strategy $\pi_{\theta_i}(s|h,t,\Gamma^{n,*})$, at every $\left(h,(t,\theta_i,s)\right) \in H^n$, whenever $\Gamma^{n,*}_{\theta_i}(h) > 0$.

Notice that $\cup_{n\in\mathbb{N}} H^n = \mathscr{H}$ and $H^{n+1} \subseteq H^n$ for each n.

Therefore, from construction σ^m converges weakly in strategies to a strategy σ^*, defined as $\sigma^*_{\theta_i}(s|h,t) = \pi_{\theta_i}(s|h,t,\Gamma^{n,*})$ if $\left(h,(t,\theta_i,s)\right) \in H^n \setminus H^{n+1}$. If $\sigma^*_{\theta_i}(s|h,t) = 0$ and $\left(h,(t,\theta_i,s)\right) \in H^n \setminus H^{n+1}$ then we have $\Gamma^{n,*}(h) > 0$ and $\Gamma^{n,*}\Big(\left(h,(t,\theta_i,s)\right)\Big) = 0$, which implies $\sigma^m_{\theta_i}(s|h_i(h,t),t) \to 0$ almost surely.

By construction σ^m converges weakly in strategies to σ^*.

PROOF OF LEMMA 6:

Let σ^* be a trembling hand equilibrium and let σ^m be the sequence of $\tilde{\varepsilon}^m$-constrained equilibria that converge to σ^*. Let $\hat{\sigma}_\theta$ be a strategy of player of type $\theta \in \Theta_i$. For each ε^m, we can construct a strategy, $\hat{\sigma}^m_\theta$, that is $\tilde{\varepsilon}^m$-constrained and converges almost surely to $\hat{\sigma}_\theta$, as $m \to \infty$. We construct this strategy, by setting, at each history h, the probability of every action $s_i \in S_i(m_i(h_i))$, taken with probability less than $\tilde{\varepsilon}^m_i(s_i, h_i, t)$ under $\hat{\sigma}_\theta$, to $\tilde{\varepsilon}^m_i(s_i, h_i, t)$ and subtracting the appropriate probability from the action that is taken with the highest probability. This procedure results in a strategy for ε^m small enough.

From Lemma 2 it follows that $prob(h|a, \sigma^m)$ and $prob\left(h|a, \left(\hat{\sigma}^m_i, \sigma^m_{-i}\right)\right)$ converge weakly to $prob(h|a, \sigma^*)$ and $prob(h|a, (\hat{\sigma}_i, \sigma^*_{-i}))$, respectively. From the definition of weak convergence, by taking limits on both sides of (3), we obtain

$$U_\theta(\hat{\sigma}_\theta, \sigma^*_{-\theta}) \leq U_\theta(\sigma^*).$$

As strategy $\hat{\sigma}_\theta$ is arbitrary, σ^* must be a Nash equilibrium of the stochastic timing game.

PROOF OF THEOREM 2:

First, let's see that sequential rationality holds for $\tilde{\varepsilon}$-constrained equilibria. Let $\tilde{\varepsilon}$ be a profile of (ε, ν)-trembles, let σ^ε be a $\tilde{\varepsilon}$-constrained equilibrium and let $(\sigma^\varepsilon, \tilde{\mu}^\varepsilon)$ be a weakly consistent assessment. We want to show that for every player $i \in \{1, \ldots, I\}$ and every type $\theta_i \in \Theta_i$

$$U_{\theta_i}\left(\sigma^\varepsilon_{\theta_i}|h_i, t, \tilde{\mu}^\varepsilon\right) \geq U_{\theta_i}\left(\hat{\sigma}^\varepsilon_{\theta_i}|h_i, t, \tilde{\mu}^\varepsilon\right), \tag{19}$$

for all $\hat{\sigma}^\varepsilon_{\theta_i} \in \Sigma_i(\tilde{\varepsilon})$ and $h_i \in \mathscr{H}_i(t)$. Suppose the previous inequality does not hold in a set $\tilde{H}_i$ of i's private history and time pairs, and define $\mathscr{H}_i(\tilde{H}_i) = \cup_{(h_i,t)\in\tilde{H}_i}\mathscr{H}_i(h_i, t, \theta_i)$. If $\mathscr{H}_i(\tilde{H}_i)$ is a μ_{θ_i}-zero measure set, we can re-define $\sigma^\varepsilon_{\theta_i}$ in $\tilde{H}_i$ so that the inequality holds for all private histories. The resulting strategy profile σ^ε would remain a $\tilde{\varepsilon}$-constrained equilibrium.

Suppose then that $\mathscr{H}_i(\tilde{H}_i)$ has strictly positive μ_{θ_i}-measure and there exists a strategy $\hat{\sigma}^\varepsilon_{\theta_i}$ such that

$$U_{\theta_i}\left(\sigma^\varepsilon_{\theta_i}|h_i, t, \tilde{\mu}^\varepsilon\right) < U_{\theta_i}\left(\hat{\sigma}^\varepsilon_{\theta_i}|h_i, t, \tilde{\mu}^\varepsilon\right), \tag{20}$$

for $(h_i, t) \in \tilde{H}_i$.

Define $d\mu_{\theta_i}(h_i, t) = \left(\int_{(\hat{h},\hat{a})\in\mathscr{H}_i(h_i,t,\theta_i)} prob_{-i}(\hat{h}^t|\hat{a}^t, \sigma^\varepsilon)\cdot d\mu_{\theta_i}(\hat{a})\right)\cdot\Gamma_{\theta_i}(h_i, \sigma^\varepsilon_{\theta_i})$, where $\Gamma_{\theta_i}(h_i, \sigma^\varepsilon_{\theta_i}) = \Gamma_{\theta_i}(h^t, \sigma^\varepsilon_{\theta_i})$ for h such that $h_i(h, t) = h_i$. Since $prob_{-i}(\hat{h}^t|\hat{a}^t, \sigma^\varepsilon) > 0$ and $\Gamma_{\theta_i}(h^t, \sigma^\varepsilon_{\theta_i}) > 0$ for every strategy $\sigma^\varepsilon \in \Sigma(\tilde{\varepsilon})$, equation (20) implies

$$\int_{(h_i,t)\in\tilde{H}_i} U_{\theta_i}\left(\sigma^\varepsilon_{\theta_i}|h_i, t, \tilde{\mu}^\varepsilon\right) d\mu_{\theta_i}(h_i, t) < \int_{(h_i,t)\in\tilde{H}_i} U_{\theta_i}\left(\hat{\sigma}^\varepsilon_{\theta_i}|h_i, t, \tilde{\mu}^\varepsilon\right) d\mu_{\theta_i}(h_i, t).$$

Now, because $\tilde{\mu}^\varepsilon$ satisfies Bayes' rule at all histories the previous expression can be written as

$$\int_{(h,a)\in\mathscr{H}_i(\tilde{H}_i)} prob\left(h|\left(\sigma^{\varepsilon}_{\theta_i},\sigma^{\varepsilon}_{-\theta_i}\right),a\right) g_{\theta_i}(h)\, d\mu_{\theta_i}(a) < \int_{(h,a)\in\mathscr{H}_i(\tilde{H}_i)} prob\left(h|\left(\hat{\sigma}^{\varepsilon}_{\theta_i},\sigma^{\varepsilon}_{-\theta_i}\right),a\right) g_{\theta_i}(h)\, d\mu_{\theta_i}(a),$$

which contradicts σ^{ε} is $\tilde{\varepsilon}$-constrained equilibrium: the strategy $\hat{\hat{\sigma}}^{\varepsilon}_{\theta_i}$ that plays according to $\hat{\sigma}^{\varepsilon}_{\theta_i}$ at $(h_i,t)\in\tilde{H}_i$ and according to $\sigma^{\varepsilon}_{\theta_i}$ elsewhere is a profitable deviation for type θ_i.

Now, let σ be a THPE and let σ^m be a sequence that converges to σ weakly in strategies, where σ^m is $\tilde{\varepsilon}^m$-constrained, with $\tilde{\varepsilon}^m$, an ε^m-tremble profile such that $\varepsilon^m\to 0$. Let $\tilde{\mu}^m$ be the belief profile weakly consistent with σ^m and let $\tilde{\mu}$ be the weak limit of $\tilde{\mu}^m$. Suppose that

$$U_{\theta_i}\left(\sigma_{\theta_i}\middle|h_i,t,\tilde{\mu}\right) < U_{\theta_i}\left(\hat{\sigma}_{\theta_i}\middle|h_i,t,\tilde{\mu}\right) \tag{21}$$

for $(h_i,t)\in\hat{H}_i$, with $\mathscr{H}_i(\hat{H}_i)$ a non-zero μ_{θ_i}-measure set.

Now, define $d\hat{\mu}_{\theta_i}(h_i,t,\tilde{\sigma}_{\theta_i}) = \left(\int_{(\hat{h},\hat{a})\in\mathscr{H}_i(h_i,t)} 1\cdot d\mu_{\theta_i}(\hat{a})\right)\cdot\Gamma_{\theta_i}(h_i,\tilde{\sigma}_{\theta_i},\bar{H}(h_i,\sigma))$. Equation (21) implies,

$$\int_{(h_i,t)\in\hat{H}_i} U_{\theta_i}(\sigma_{\theta_i}|h_i,t,\tilde{\mu})\, d\hat{\mu}_{\theta_i}(h_i,t,\sigma_{\theta_i}) < \int_{(h_i,t)\in\hat{H}_i} U_{\theta_i}(\hat{\sigma}_{\theta_i}|h_i,t,\tilde{\mu})\, d\hat{\mu}_{\theta_i}(h_i,t,\sigma_{\theta_i}) \tag{22}$$

As in the proof of Lemma 1 we can find a sequence of strategies $\hat{\sigma}^m_{\theta_i}$ that converges almost surely to $\hat{\sigma}_{\theta_i}$ and is $\tilde{\varepsilon}^m$-constrained. As in the proof of Lemma 2 because opposing players do not observe i's moving times $\tilde{\mu}^m$ is "constant" in small intervals of i's moving times. Therefore, by an argument similar to the one in Lemma 2 we obtain that

$$\int_{(h_i,t)\in\hat{H}_i} U_{\theta_i}(\sigma^{\varepsilon}_{\theta_i}|h_i,t,\tilde{\mu}^{\varepsilon})\, d\hat{\mu}_{\theta_i}(h_i,t,\sigma^{\varepsilon}_{\theta_i})$$

converges to the left hand side of (22), while

$$\int_{(h_i,t)\in\hat{H}_i} U_{\theta_i}(\hat{\sigma}^{\varepsilon}_{\theta_i}|h_i,t,\tilde{\mu}^{\varepsilon})\, d\hat{\mu}_{\theta_i}(h_i,t,\sigma^{\varepsilon}_{\theta_i})$$

converges to the right hand side. This contradicts equation (19).

References

ALIPRANTIS, C. D. AND K. C. BORDER (1999): *Infinite dimensional analysis*, vol. 32006, Springer.

ALTMAN, E., K. AVRACHENKOV, N. BONNEAU, M. DEBBAH, R. EL-AZOUZI, AND D. S. MENASCHE (2008): "Constrained cost-coupled stochastic games with independent state processes," *Operations Research Letters*, 36, 160–164.

AMBRUS, A., J. BURNS, AND Y. ISHII (2014): "Gradual bidding in ebay-like auctions," *Working Paper*, 1–50.

AMBRUS, A. AND S. E. LU (2014): "A continuous model of multilateral bargaining with random arrival times," *American Economics Journal: Microeconomics*.

ARCIDIACONO, P., P. BAYER, J. R. BLEVINS, AND P. B. ELLICKSON (2016): "Estimation of Dynamic Discrete Choice Models in Continuous Time with an Application to Retail Competition," *The Review of Economic Studies*, 83, 889–931.

BALBUS, Ł., K. REFFETT, AND Ł. WOŹNY (2013): "Markov Stationary Equilibria in Stochastic Supermodular Games with Imperfect Private and Public Information," *Dynamic Games and Applications*, 3, 187–206.

CALCAGNO, R., Y. KAMADA, S. LOVO, AND T. SUGAYA (2014): "Asynchronicity and coordination in common and opposing interest games," *Theoretical Economics*, 9, 409–434.

CALVO, G. A. (1983): "Staggered prices in a utility-maximizing framework," *Journal of Monetary Economics*, 12, 383–398.

DORASZELSKI, U. AND K. L. JUDD (2012): "Avoiding the curse of dimensionality in dynamic stochastic games," *Quantitative Economics*, 3, 53–93.

DUGGAN, J. (2012): "Noisy Stochastic Games," *Econometrica*, 80, 2017–2045.

FINK, A. M. (1964): "Equilibrium in a stochastic n-person game," *J. Sci. Hiroshima Univ. Ser. A-I Math.*, 28, 89–93.

FLESCH, J., F. THUIJSMAN, AND O. J. VRIEZE (2003): "Stochastic games with non-observable actions," *Mathematical Methods of Operations Research*, 58, 459–475.

GENSBITTEL, F., S. LOVO, J. RENAULT, AND T. TOMALA (2017): "Zero-sum revision games," *Games and Economic Behavior*.

HOPENHAYN, H. AND M. SAEEDI (2016): "Bidding Dynamics in Auctions," .

KAMADA, Y. AND M. KANDORI (2017a): "Revision Games, Part I: Theory," *Working Paper*.

——— (2017b): "Revision Games, Part II: Applications and Robustness," *Working Paper*.

KAMADA, Y. AND S. MORONI (2017): "Games with Private Timing," Tech. rep.

KAMADA, Y. AND N. MUTO (2014): "Multi-Agent Search with Deadline," *Working Paper*, 1–77.

KAMADA, Y. AND T. SUGAYA (2014): "Valence Candidates and Ambiguous Platforms in Policy Announcement Games," *University of California Berkeley, mimeo*.

KREPS, D. M. AND R. WILSON (1982): "Sequential equilibria," *Econometrica: Journal of the Econometric Society*, 863–894.

LOVO, S. AND T. TOMALA (2015): "Markov Perfect Equilibria in Stochastic Revision Games," .

MORONI, S. AND A. KAPOR (2015): "Sniping in Proxy Auctions with Deadlines," *Working Paper*.

NOWAK, A. S. AND T. E. S. RAGHAVAN (1992): "Existence of stationary correlated equilibria with symmetric information for discounted stochastic games," *Mathematics of Operations Research*, 17, 519–526.

OCKENFELS, A. AND A. E. ROTH (2006): "Late and multiple bidding in second price Internet auctions: Theory and evidence concerning different rules for ending an auction," *Games and Economic Behavior*, 55, 297–320.

ROTH, A. AND A. OCKENFELS (2002): "Last minute bidding and the rules for ending second price auctions: evidence from eBay and Amazon auctions on the Internet," *American Economic Review*.

SANDHOLM, W. H. (2011): "Stochastic Evolutionary Game Dynamics: Foundations, Deterministic Approximation, and Equilibrium Selection," *Evolutionary Game Dynamics: American Mathematical Society Short Course, January 4-5, 2011, New Orleans, Louisiana*, 69, 111.

SELTEN, R. (1975): "Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games," *International journal of game theory*.

SHAPLEY, L. S. (1953): "Stochastic games," *Proceedings of the national academy of sciences*, 39, 1095–1100.

SOBEL, M. J. (1971): "Noncooperative Stochastic Games," *The Annals of Mathematical Statistics*, 42, 1930–1935.