Department of Economics

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

WORKING PAPER SERIES

16/014

Existence of trembling hand equilibrium in revision games with imperfect information

Sofia Moroni

June 7, 2015

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 equilibrium in revision games with imperfect information

Sofia Moroni*

Yale University

sofia.moroni@yale.edu

June 7, 2015

Latest version available here

Abstract

In revision games a group of players can move at stochastic opportunities before a deadline. Their payoffs are determined by the sequence of actions taken before the end of the game. In this paper I define trembling hand equilibrium in a large class of revision games that may feature incomplete and imperfect information, and show that trembling hand equilibria exist. Since trembling hand perfect equilibria are also Nash, existence of a Nash equilibrium follows.

*I am grateful to Johannes Hörner, Adam Kapor and Larry Samuelson for their insightful comments.

1 Introduction

A revision game consists of a group of players who can move at stochastic opportunities before a deadline. Their payoffs are determined by the sequence of actions taken before the end of the game. The opportunities to take actions arrive at Poisson rates. The space of histories in revision games is large, and as a result, existing proofs of existence do not apply. In this paper I define trembling hand equilibrium, in the spirit of the notion of perfect equilibrium for finite extensive form games by Selten (1975), and show that trembling hand equilibria exist. The games considered encompass a large class of games with stochastic timing that may feature incomplete and imperfect information. Players may have private types and may exert unobserved actions. Additionally, the identity of a player who moved may be unobserved. Since trembling hand perfect equilibria are also Nash equilibria, existence of Nash equilibrium is also established.

There are many economic situations in which players interact dynamically and must take actions before a deadline. 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. In bargaining between labor and management, 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.

Revision games are an elegant framework for modeling dynamic situations like these, in which actions may occur at any time, and in which the environment is nonstationary because of an approaching deadline. As introduced by Kamada and Kandori (2011), a revision game is built on an underlying game which is to be played at some future time. Before that time arrives, players have stochastic opportunities to choose and revise their actions. At the deadline, players receive payoffs according to the final action taken by each player.

There is a literature that applies revision games to questions in applied theory. Ambrus and Lu (2008) present a revision games model of multilateral bargaining with a deadline. Ambrus, Burns and Ishii (2014) and Moroni (2014) model online auctions as revision games. However, each of these papers derives existence and properties of equilibrium in a way that is chosen for the specific model in each paper. Because the set of histories in a revision game is uncountably infinite, it is not trivial to show that equilibria exist in these models. It has not been previously established that Nash equilibria exist in general in revision games.

The purpose of this paper is to provide a general proof of existence of equilibrium in revision games. I define a notion of trembling hand perfect equilibrium in revision games and prove that an equilibrium exists in games with incomplete information with finitely many types and finitely many actions.[1] I allow for incomplete information about agents' types, for imperfect observability of players' actions and for the action space to vary as players move. As in Selten (1975), trembling hand equilibrium will be Nash. The definition of trembling hand equilibrium is inspired by the definition by Selten for extensive form games.[2]

The key step in the proof consists of identifying strategies with functions in a L^2 space endowed with a convenient measure space. The measure space is chosen so that convergence in the weak topology—although weaker than other convergence notions—implies convergence of expected payoff of strategies for each player. All functions that represent strategies are probability measures and will be bounded in the L^2 norm. Thus, from the Banach–Alaoglu theorem, these functions must be in a set that is compact with respect to the weak-star topology L^2. Next, I define finite approximating games of the original revision game which must have an ε-constrained equilibrium (Selten [1975]). The ε-constrained equilibrium can be identified with ε-constrained strategies in the original game. As the approximating games converge to the original game these identified ε-constrained strategies must have a convergent subsequence[3], whose limit is the ε-constrained equilibrium of the original game. Once I establish that the ε-constrained equilibria exist, existence of a trembling hand perfect equilibrium follows.

2 Revision game

There are I players who participate in a game for a length of time T. A player can take an action at a time $t \in [0,T]$ if she receives an opportunity to move at that time. The opportunities to take actions arrive stochastically at Poisson rates. The timings at which each player i receives an opportunity follow a Poisson process with rate λ_i. I allow for the possibility that a player's arrival rate is private information. Each player's opportunities arrive independently from the other players' opportunities. At each opportunity a player can choose an action in her action set. I will allow for players' actions sets to depend on the history of play. The payoff that players obtain depends on the sequence of actions taken by each player.[4]

[1] All action sets are finite so the extension of extensive form trembling hand perfect equilibrium is particularly straightforward.

[2] A similar notion was proposed by Kamada and Muto (2011) in the context of revision games with complete information.

[3] In the weak-* topology which coincides with the weak topology

[4] Note that this setting is more general than revision games settings that have been studied so far. I assume that the payoffs of the agents may depend on the sequence of actions taken and not just on the last action that each player took.

Histories Let $\mathscr{A}$ denote the space of histories of arrivals from time 0 until time T. A history contains the times at which players received an opportunity and the players who received them. Let $a = ((t^1, i^1), \ldots, (t^n, i^n)) \in \mathscr{A}$, with $t^j < t^{j+1}$, denote a history with n opportunities in which player i^j has an opportunity at time t^j. Let $\mathscr{H}(a)$ denote the set of all feasible histories of play if the history of arrivals is a. A history of play in $\mathscr{H}(a)$ has the form $h = ((t^1, i^1, s^1), \ldots, (t^n, i^n, s^n))$ where s^j denotes the action player i^j takes at an arrival at time t^j.

Actions Let $m(h) = ((s^1, i^1), \ldots, (s^n, i^n))$ denote the history without the corresponding timings. I write $m_i(h,t)$ for the history of actions according to i's information up to time t given h and I denote $S_i(m_i(h,t))$ the finite set of actions available to player i at time t. The set $S_i(m_i(h,t))$ contains the last action taken by player i and it may depend on the history of moves up to time t but not on the timings at which actions were taken. The action set of each player i must be measurable with respect to i's information. The total number of actions available to each player must be finite. That is, for each i the union of all available actions after each history, $\bigcup_{h_i(t) \in \mathscr{H}_i(t), t \in [0,T]} S_i(m_i(h,t))$, is a finite set.

Information The games considered have imperfect and incomplete information. First, player i may imperfectly observe actions taken by other players. I model this informational asymmetry by letting each player have a partition of actions, such that at each history, each player only observes that an action from a partition element was taken but cannot identify the action itself. Thus, each action s_j by player j belongs to an element $P_i^s(s_j)$ of i's *action partition* and i cannot distinguish between actions in a set in his action partition.

Second, players may have private information about their "types". I model this possibility by assuming that each player may not observe the identity of another player. Thus, from player i's point of view each player j belongs to an element of player i's *type partition* $P_i^\theta(j)$. [5] This information structure allows for players that have private information about their payoffs or arrival rates or types. In this case, the partition element $P_i^\theta(j)$ represents one player and the elements $j' \in P_i^\theta(j)$ are distinct types of player $P_i^\theta(j)$. The prior probabilities over players in a type partition must sum to one: $\sum_{j' \in P_i^\theta(j)} p_i^\theta(j') = 1$.

Third, players imperfectly observe the identity of a player who took an action. For example, in an auction a player may observe that someone has placed a higher bid but cannot observe who placed the bid. Thus, I let each player have a partition over players, such that

[5] In an online auction, for instance, player i may observe the history by an opposing player. However, this player is identified only by a randomly generated code number. Hence i does not know who the opponent is and may for a prior about the player's valuation and arrival rate.

at each history, each player only observes an element of a partition to which an opposing player 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 $P_i^p(P_i^\theta(j), P_i^s(s_j))$ corresponds to the type partition element of player j after taking action s_j. Player i cannot distinguish between different players in the partition element $P_i^p(P_i^\theta(j), P_i^s(s_j))$. Since player i only observes that player j belongs to type partition element $P_i^\theta(j)$ and that action s^j belongs to action partition element $P_i^s(s_j)$, the player partition can only depend on these variables.

Fourth, players cannot observe the timing of the other players' moves perfectly. Thus, the timings of a player $j \neq i$, t_j, belong to a *finite* partition $P_j^t(t_j)$. If player j moves at time t_j player i only observes the corresponding element in $P_j^t(t_j)$. In an online auction this assumption would hold if the players can observe only a coarse measure of the other players timings of moves. For example, the assumption holds if players can observe the hour, minute and second at which a player has moved but cannot observe the millisecond.[6]

Finally, I also allow for the possibility that a player may not observe some or all of the other players' arrivals.

Belief over arrival histories Because players may have private information regarding their arrival rates each player i has a prior belief the probability of arrival histories. Let $\beta_{\mathscr{A}}$ denote the Borel sets of the set of arrival histories $\mathscr{A}$. Player i's prior belief is a probability measure $\mu_i : \beta_{\mathscr{A}} \to [0,1]$ that is absolutely continuous with respect to the Lebesgue measure (because opportunities arrive at Poisson rates).

Private Histories Players may not observe the actions or arrivals that have occurred up to time t. Given a history $h = ((t^1, i^1, s^1), \ldots, (t^n, i^n, s^n))$ with n arrivals, the private history up to but not including time t given i's information is denoted $h_i(t)$. Player i may not observe that an action was taken at all or may realize that an action was taken at an observable time but cannot observe exactly which action was taken or which player took it. In the first case, a given element (t^j, i^j, s^j) would not be present in i's private history. In the second case i observes elements of a partition of the actions taken by a player and the identity of the players. Thus, player i's private history at time t takes the form

$$h_i(t) = ((P_i^t(t^1), P_i^p(P_i^t(i^1), P_i^s(s^1)), P_i^s(s^1)), \ldots, (P_i^t(t^m), P_i^p(P_i^t(i^m), P_i^s(s^m)), P_i^s(s^m)))$$

[6]This technical assumption is needed to guarantee that the limit of strategies is a strategy. See proof of Lemma 1. The assumption is not needed in games with perfect observability. That is games in which the players may have private types but the identity of the players who move, the actions taken and the timing of moves is perfectly observed.

with $m \leq n$ and $t^k < t^{k+1} \leq t$. I assume perfect recall, therefore, player i's history h_i contains all opportunities and moves by player i. I denote the set of player i's private histories up to time t as $\mathscr{H}_i(t)$.

Strategies A strategy for player i, $\sigma_i : \bigcup_{t\in[0,T]} \mathscr{H}_i(t) \times t \to \Delta S_i(m_i(h,t))$, is a function from private histories to probability measures over action sets.

Payoffs The payoff of the players after history $h = ((t^1, i^1, s^1), \ldots, (t^n, i^n, s^n))$ depends on the sequence of actions that were taken but not on their timings. The payoff of each player i is given by a function $g_i(m_i(h,t))$[7]. I assume that for each i, $|g_i|$ is bounded by a constant G_i. Without loss of generality I assume $g_i > 0$ for each i. Abusing notation I denote $g(h)$ for the vector of payoffs induced by the sequence of actions taken by all players in history h.

A *revision game* is a game with the properties described above.

3 Equilibrium

A strategy $\sigma = \times_{i=1}^{I} \sigma_i$ induces a probability $prob(h|a,\sigma)$ over histories of actions taken given the history of opportunities. For a history $h = ((t^1, i^1, s^1), \ldots, (t^n, i^n, s^n))$, the probability is given by

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

where $\sigma_{i^j}(s^j|h_{i^j}(t),t)$ denotes the probability that player i^j assigns to action s^j given the private history $h_{i^j}(t)$ if she receives an opportunity at time t. The function $prob(h|a,\sigma)$ is termed the *probability function associated to* σ.

Player i's expected payoff when players follow strategy σ is given by

$$U_i(\sigma) = \int_{\mathscr{A}} \sum_{h \in \mathscr{H}(a)} g_i(h) \cdot prob(h|a,\sigma) d\mu_i(a), \tag{1}$$

where μ_i denotes the measure over arrival opportunities given each player's arrival rate and player i's prior regarding players' types.

Strategies as functions in L^2. I associate a strategy σ with the function $prob(h|a,\sigma)$ in a L^2 space endowed with a convenient measure space defined as follows. Let $\mathscr{H} = \{(h,a)|h \in \mathscr{H}(a), a \in \mathscr{A}\}$. Consider the measure ν_i over $\mathscr{H}$ defined as $d\nu_i(h,a) = g_i(h) d\mu_i(a)$. Then

[7] Note that the the player's payoff may depend on their types since different types of players are denoted with different subscripts.

$prob(h|a,\sigma)^N \in \times_{i=1}^N L^2(\mathscr{H},\nu_i)$ where the squared norm of a function $f = (f_i(h,a))_{i=1}^N \in \times_{i=1}^N L^2(\mathscr{H},\nu_i)$ is given by

$$\|f\|_2 = \sum_{i=1}^N \int_{a\in\mathscr{A}} \sum_{h\in\mathscr{H}(a)} |f_i(h,a)|^2 g_i(h) d\mu_i(a).$$

I say that f^N converges to f^* in the weak topology of L^2 if for all functions $\varphi_i \in L^2(\mathscr{H},\nu_i)$ and for all i

$$\int_{\mathscr{A}} \left(\sum_{h\in\mathscr{H}(a)} f_i^N(h,a)\cdot\varphi_i(h,a) \right) d\mu_i(a) \to \int_{\mathscr{A}} \left(\sum_{h\in\mathscr{H}(a)} f_i^*(h,a)\cdot\varphi_i(h,a) \right) d\mu_i(a). \tag{2}$$

A shorthand notation for weak convergence is

$$< f_i^N, \varphi_i > \to < f_i^*, \varphi_i > .$$

The set of probability functions associated to strategies is bounded in the L^2 norm. In fact, we have

$$\sum_{i=1}^N \int_{\mathscr{A}} \sum_{h\in\mathscr{H}(a)} |prob(h|a,\sigma)|^2 g_i(h) d\mu_i(a) \le G_i.$$

This observation will prove useful to prove existence, because it implies, by the Banach-Alaoglu theorem, that the space of probabilities associated to strategies is compact. Thus, sequences of functions associated to probabilities have convergent subsequences.

Trembling hand equilibrium An ε-constrained strategy is a strategy profile such that for each history h and time t each player i puts weight at least ε on every action in the set of available actions, $S_i(m_i(h,t))$.

I refer to strategies that put weight of at least ε on every available action after each history as *ε-constrained strategies*.

Definition 1. An *ε-constrained equilibrium* σ^ε is an ε-constrained strategy profile that is constrained optimal. That is, for every player i,

$$\int_{\mathscr{A}} \sum_{h\in\mathscr{H}(a)} g_i(h)\cdot prob(h|a,(\sigma_i',\sigma_{-i}^\varepsilon)) d\mu_i(a) \le \int_{\mathscr{A}} \sum_{h\in\mathscr{H}(a)} g_i(h)\cdot prob(h|a,\sigma^\varepsilon) d\mu_i(a), \tag{3}$$

for every ε-constrained strategy σ_i'.

A trembling hand perfect equilibrium (THPE) as the limit of a sequence of ε-constrained equilibria in the weak topology. Consider a history $h = ((t^1, i^1, s^1), \dots, (t^n, i^n, s^n))$ and let $H_i(h)$ denote the set of tuples (t, i, s) in history h at which player i has had the opportunity to move in history h. Define

$$\Gamma_i(h, \sigma_i) = \prod_{(t_k, i, s_k) \in H_i(h)} \sigma_i(s_k | h_i(t_k), t_k). \tag{4}$$

$\Gamma_i(h, \sigma_i)$ is the product of the probabilities with which player i takes the actions in history h at the relevant opportunities.

A sequence of strategies, σ^m *converge to $\bar{\sigma}$ in the weak topology* if and only if for every player i,

$$\int_{\mathscr{A}} \left(\sum_{h \in \mathscr{H}(a)} \Gamma_i(h, \sigma_i^m) \cdot \varphi_i(h, a) \right) d\mu_i(a) \to \int_{\mathscr{A}} \left(\sum_{h \in \mathscr{H}(a)} \Gamma_i(h, \bar{\sigma}_i) \cdot \varphi_i(h, a) \right) d\mu_i(a) \tag{5}$$

for every $\varphi_i \in L^2(\mathscr{H}, \mu_i)$.

I will say that σ^m *converges to $\bar{\sigma}$ in the weak topology in the set $B \subseteq \mathscr{A}$* if condition (5) holds for every $\varphi_i \in L^2(\mathscr{H}, \mu_i)$ with support in arrival histories in B.

Definition 2. A strategy profile σ^* is a *trembling hand equilibrium* if there exists a sequence $(\varepsilon^m)_{m=1,2,\dots}$, with $\varepsilon^m > 0$ and $\lim_{m \to 0} \varepsilon^m = 0$, and ε^m-constrained equilibria σ^{ε^m} such that σ^{ε^m} converges to σ^* in the weak topology.

This definition is in the spirit of perfect equilibrium in extensive form games defined by Selten (1975). The difference is that convergence is of the product of probabilities instead of the behavioral strategies themselves. The type of convergence that I propose is weaker and will allow me to establish existence of a trembling hand equilibrium.

3.1 Existence of ε-constrained equilibria

The proof of existence of trembling hand perfect equilibria is done in two steps. I first show that an ε-constrained equilibrium exists in every revision game. I then show that a sequence of ε-constrained equilibria must have have a limit in the weak topology.

Establishing existence of an ε-constrained equilibrium is most of the work and is done through the following steps. I define a sequence of approximating games that converge to the original game. I prove that the probability functions associated to the approximating game strategies must have a convergent subsequence. I then show that the limit probability functions of the subsequence is a probability function that corresponds to a an equilibrium in the original game.

3.1.1 Preliminaries

In this section I provide some useful results that establish a connection between convergence of probability functions associated to strategies and weak convergence of strategies.

Given a function $f^*:\mathscr{H}\to[0,1]$, the *candidate strategy associated to* f^* is defined as

$$\pi_i(s|h,t,f^*)=\frac{f^*\left(\left(h,(t,i,s)\right),\left(a,(t,i)\right)\right)}{f^*(h,a)}$$

where h is a history in which the players' arrivals is given by a and the latest arrival in a occurs before time t. I say that π_i is well defined at history $(h,(t,i,s))$ whenever $f^*(h,a)>0$. If $f^*(h,a)=0$ π_i is not well defined and its value is undetermined. If h is the empty history $f^*(h,a)$ is taken to be 1.

Note that in order for π_i to *correspond to a strategy* at histories in which it is well defined it must be measurable with respect to i's information and the probabilities over actions must sum to one. For π_i to be measurable with respect to i's information we must have $\pi_i(s|h,t,f^*)=\pi_i(s|h_i(t),t,f^*)$ where $h_i(t)$ is player i's private history at time t given history h.

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

The set of arrival histories in which the last opportunity belongs to player i is denoted $\mathscr{A}_i$.

Lemma 1. *The sequence of functions $f^m(h,a)$ converges to $f^*(h,a)$ in the weak topology, if and only if for each player i the candidate strategy $\pi_i(\cdot|\cdot,f^*)$ associated to f^*corresponds to a strategy for histories in which it is well defined and we have*

$$\int_{\mathscr{A}}\sum_{h\in\mathscr{H}(a)}\left(\Gamma_i(h,\sigma_i^m)-\Gamma_i(h,\pi_i)\right)\varphi_i(h,a)\,d\mu_i(a)\to 0, \tag{6}$$

$$\int_{\mathscr{A}}\sum_{h\in\mathscr{H}(a)}\left(\left(\prod_{j\neq i}\Gamma_j(h,\sigma_j^m)\right)-\left(\prod_{j\neq i}\Gamma_j(h,\pi_j)\right)\right)\varphi_i(h,a)\,d\mu_i(a)\to 0, \tag{7}$$

for every $\varphi_i(h,a)\in L^2(\mathscr{H},\mu_i)$ with support in the set of histories in which π is well defined.

Proof. I will argue recursively based on the number of arrivals in a. If φ_i has support in the subset of histories in which only player i receives an arrival it is immediate that (2) holds if and only if (6) holds. Now, suppose the result holds for every φ_i with support in histories with at most k arrivals and let's see that the result must also hold for histories with $k+1$ arrivals. Let $\mathscr{A}_i^{k+1}$ denote the set of arrival histories with $k+1$ arrivals in which player i receives the last opportunity. Let $\mathscr{H}^\pi$ denote the set of histories in which the profile π is well defined. From the definition of weak convergence in equation (2), f^m converges to f^* in $\mathscr{H}^\pi$, if and only if for every $\varphi_i \in L^2(\mathscr{H}, \nu_i)$ with support in $\mathscr{H}^\pi$,

$$\int_{\mathscr{A}_i^{k+1}} \sum_{h\in\mathscr{H}(a)} \left(\left(\prod_{j\neq i} \Gamma_j(h,\sigma_j^m) \right) \cdot \Gamma_i(h,\sigma_i^m) - \left(\prod_{j\neq i} \Gamma_j(h,\pi_j) \right) \cdot \Gamma_i(h,\pi_i) \right) \cdot \varphi_i(h,a)\, d\mu_i(a) \to 0. \tag{8}$$

From the definition of $\Gamma_i(h,\sigma_i)$ in equation (4), $\left(\prod_{j\neq i}\Gamma_j(h,\sigma_j^m)\right)$ is the multiplication of strategies that depend on private histories with at most k arrivals (since player i moves last in arrival histories in $\mathscr{A}_i^{k+1}$. Thus, from the induction hypothesis we have,

$$\int_{\mathscr{A}_i^{k+1}} \sum_{h\in\mathscr{H}(a)} \left(\left(\prod_{j\neq i} \Gamma_j(h,\sigma_j^m) \right) \cdot \Gamma_i(h,\pi_i) - \left(\prod_{j\neq i} \Gamma_j(h,\pi_j) \right) \cdot \Gamma_i(h,\pi_i) \right) \cdot \varphi_i(h,a)\, d\mu_i(a) \to 0. \tag{9}$$

Combining equations (8) and (9) we obtain f^m converges to f^* in $\mathscr{A}_i^{k+1}$ if and only if

$$\int_{\mathscr{A}_i^{k+1}} \left(\left(\prod_{j\neq i} \Gamma_j(h,\sigma_j^m) \right) \left(\Gamma_i(h,\sigma_i^m) - \Gamma_i(h,\pi_i) \right) \right) \cdot \varphi_i(h,a)\, d\mu_i(a) \to 0. \tag{10}$$

Player i does not observe the timings of moves of the opposing players perfectly. She observes only that the timings belong to an element of a finite partition of the time interval. Let A_i^{k+1} denote a set of arrival histories of player i contained in the set of histories with $k+1$ arrivals. Let $A_{-i}^{k+1}(a_i)$ the set of possible arrival histories of players other than i contained in the set of histories with $k+1$ arrivals given that player i arrivals are given by $a_i \in A_i^{k+1}$. For each $a_i \in A_i^{k+1}$ there is a finite partition of $A_{-i}^{k+1}(a_i)$, denoted $\tilde{P}_{-i}^t(a_i)$, that represents the partition elements of the opposing players' times of moves that are observable by player i. The timings of i's moves are also not perfectly observable by the other players. Let $\tilde{P}_i^t$ be the partition of times of player i's arrivals contained in $\mathscr{A}_i^{k+1}$ as observed by the opposing players. Let $\tilde{A} \in \tilde{P}_i^t$. We can restate condition (10) as f^m converges to f^* in $\mathscr{A}_i^{k+1}$ if and only if for every $\tilde{A} \in \tilde{P}_i^t$ and $\nu > 0$ there is $\bar{m}$ such that for all $m \geq \bar{m}$,

$$\left| \int_{\tilde{A}} \sum_{B \in P_{-i}^t(a_i)} \left(\Gamma_i(h, \sigma_i^m) - \Gamma_i(h, \pi_i)\right) \left(\int_B \prod_{j \neq i} \Gamma_j(h, \sigma_j^m) \cdot \varphi_i(h, a) \, d\mu_i(a_{-i}) \right) d\mu_i(a_i) \right| < \nu.$$

The parenthesis inside the integral is constant in $a_i \in \tilde{A}$, since the strategies of players other than i must be measurable with respect to $\tilde{P}_i^t$. Now, from the induction hypothesis, there is $\bar{\bar{m}}$ such that for $m \geq \bar{\bar{m}}$

$$\left| \left(\int_B \prod_{j \neq i} \Gamma_j(h, \sigma_j^m) \cdot \varphi_i(h, a) \, d\mu_i(a_{-i}) \right) - \left(\int_B \prod_{j \neq i} \Gamma_j(h, \pi_j) \cdot \varphi_i(h, a) \, d\mu_i(a_{-i}) \right) \right| < \frac{\nu}{2},$$

for every $a_i \in \tilde{A}$. Thus, we have, for $m \geq \max\{\bar{m}, \bar{\bar{m}}\}$

$$\left| \int_{\tilde{A}} \sum_{B \in P_{-i}^t(a_i)} \left(\Gamma_i(h, \sigma_i^m) - \Gamma_i(h, \pi_i)\right) \left(\int_B \prod_{j \neq i} \Gamma_j(h, \pi_j) \cdot \varphi_i(h, a) \, d\mu_i(a_{-i}) \right) d\mu_i(a_i) \right| < \frac{\nu}{2}. \tag{11}$$

For histories h in which π is well defined a condition analogous to (9), in which the integral is taken over the set of arrival histories with k arrivals implies

$$\left(\int_B \prod_{j \neq i} \Gamma_j(h, \pi_j) \cdot \varphi_i(h, a) \, d\mu_i(a_{-i}) \right) > 0,$$

for each $B \in P_{-i}^t(a_i)$ and $a_i \in \mathscr{A}_i^{k+1}$ and equation (11) holds if and only if condition (6) holds.

By an analogous argument we can now show that condition (7) must hold for histories in which i is not the last player to move. That is, we now show that

$$\left| \int_{\mathscr{A}_{-i}^{k+1}} \sum_{h \in \mathscr{H}(a)} \left(\left(\prod_{j \neq i} \Gamma_j(h, \sigma_j^m) \right) - \left(\prod_{j \neq i} \Gamma_j(h, \pi_j) \right) \right) \varphi_i(h, a) \, d\mu_i(a) \right| \to 0, \tag{12}$$

where $\mathscr{A}_{-i}^{k+1}$ denotes the set of histories with $k+1$ arrivals in which i is not the last player to move. To see this note that from the induction hypothesis,

$$\left| \int_{\mathscr{A}_{-i}^{k+1}} \sum_{h \in \mathscr{H}(a)} \left(\left(\prod_{j \neq i} \Gamma_j(h, \pi_j) \right) \Gamma_i(h, \sigma_i^m) - \left(\prod_{j \neq i} \Gamma_j(h, \pi_j) \right) \Gamma_i(h, \pi_i) \right) \varphi_i(h, a) \, d\mu_i(a) \right| \to 0.$$

Combining the previous expression with condition (8) we obtain that f^m converges weakly to f^* in $\mathscr{A}_{-i}^{k+1}$ if and only if,

$$\left| \int_{\mathscr{A}_{-i}^{k+1}} \sum_{h \in \mathscr{H}(a)} \Gamma_i(h, \sigma_i^m) \left(\left(\prod_{j \neq i} \Gamma_j(h, \sigma_j^m) \right) - \left(\prod_{j \neq i} \Gamma_j(h, \pi_j) \right) \right) \varphi_i(h,a)\, d\mu_i(a) \right| \to 0.$$

Arguing as before because players' moves are only partially observed condition (12) is implied from the previous expression. Thus, the result holds for arrival histories with $k+1$ events.

Suppose now that φ_i has support in the set of histories in which π is not well defined then (8) follows immediately. In fact, from the definition of f^m, $f^*(h,a) = 0$ at every $h \in (\mathscr{H}^\pi)^c$ and f^m must converge to zero in the weak topology in $(\mathscr{H}^\pi)^c$.

Finally, since $\Gamma_i(h, \sigma_i^m)$ is measurable with respect to i's information so is $\Gamma_i(h, \pi_i)$. Because $\Gamma_i(h, \sigma_i^m)$ sums to one over the available actions of i π_i must also sum up to one over actions. □

In the proof of Lemma 1 I use the assumption that the timing of moves are not perfectly observed to guarantee that the strategy associated to the limit of $\{f^m\}_m$, f^*, is measurable with respect to each player's information. If the game has perfect observability of moves–that is the players observe the timings of moves, the actions taken and the players who took them–measurability with respect to a player's information is not an issue. In the perfect observability case each player's strategy depends on all of the arguments of the limit function f^* and thus the candidate strategy associated to f^* is always a strategy. Games with perfect observability of moves include games with incomplete information about players' arrival rates or their payoffs. By arguments analogous to the ones presented below I can establish existence of a trembling hand perfect equilibrium in games with perfect observability of moves.

3.1.2 Approximating games

Consider a game, denoted $\mathscr{G}^N$, in which there are 2^N stages in which the players can move. In each period either one or no player is drawn to play in that period. Let $C(N)$ be the probability that only one or no player has an arrival at an interval of length $1/2^N$ according to the Poisson process in the original game. $C(N)$ can be computed explicitly as:

$$C(N) = e^{\sum_{j=1}^N -\lambda_j/2^N} + \frac{\sum_{j=1}^N \lambda_j}{2^N} e^{\sum_{j=1}^N -\lambda_j/2^N}. \tag{13}$$

In each period player $i \in \{1,\dots,I\}$ is the one drawn to play with probability

$$\frac{\frac{\lambda_i}{2^N} e^{\sum_{j=1}^{N} -\lambda_j/2^N}}{C(N)}.$$

This probability is the probability that player i gets an arrival in an interval $I_k^N \equiv \left[T\frac{k}{2^N}, T\frac{k+1}{2^N}\right]$ for a given $k \in \{0,\dots,2^N-1\}$, conditional on the event that only one player receives an opportunity in that time interval. No player is drawn in a period with probability

$$\frac{e^{\sum_{j=1}^{N} -\lambda_j/2^N}}{C(N)}.$$

If player i gets an opportunity to play at stage k the actions available are given by the actions that would be available in the original game given the history of play in the approximating game.[8] Let $\mathscr{A}^N$ denote the set of histories of arrivals in this simplified game and let $\mu_i^N(a^N)$ denote the probability of history a^N according to player i's prior. Let $h^N = \left((\tilde{t}^1, i^1, s^1), \dots, (\tilde{t}^n, i^n, s^n)\right)$ denote a history of play in the approximating game. In this case $\tilde{t}^k \in \{1,\dots,2^N\}$ denotes the period in which each player moved. The function $g(h^N)$ is the payoff vector given the history of play h^N which as before depends only on the last actions of each player. Given a strategy σ^N let $prob^N(h^N|a^N,\sigma^N)$ denote the associated probability vector. I refer to $\mathscr{G}^N$ as the *N'th approximating game* and define $\mathscr{H}^N(a^N)$ as the set of feasible histories given a^N in $\mathscr{G}^N$.

Each approximating game $\mathscr{G}^N$ is finite and therefore it has an ε-constrained equilibrium strategy, σ^N (Selten (1975)). Let $U_i^N(\sigma^N)$ denote the payoff to player i of strategy σ^N .

Let σ^N be an ε-constrained equilibrium of $\mathscr{G}^N$. For each player i, there are no profitable deviations from equilibrium strategies which means

$$\sum_{a^N} \sum_{h^N \in \mathscr{H}(a^N)} g_i(h^N) \cdot prob^N(h^N|a^N, \sigma^N)\mu_i^N(a^N) \geq \tag{14}$$
$$\sum_{a^N} \sum_{h^N \in \mathscr{H}(a^N)} g_i(h^N) \cdot prob^N(h^N|a^N, \left((\sigma_i^N)', \sigma_{-i}^N\right))\mu_i^N(a^N),$$

for every ε-constrained strategy of player i $(\sigma_i^N)_i$.

[8] Note that this is well defined because the actions available at each time only depend on the history of actions and not on the timing at which the actions were taken.

3.1.3 Sequence of strategies

I next describe how strategies in the simplified game approximate strategies of the original game. Histories of the original game in which only one player moves in each interval of length $1/2^N$ have a natural counterpart history in the simplified game. For histories that cannot be mapped from the simplified game to the original game, that is histories in which more than one player moves in one of the $1/2^N$ length intervals, the strategy of the original game is chosen arbitrarily. The key observation is that as N grows large, the probability of histories that cannot be mapped to histories of the simplified game converges rapidly to zero. Thus, the strategies in the simplified game approximate strategies in the original game.

Given strategy σ^N in the N'th approximating game, I define σ^N*'s associated strategy* in the original game, denoted $\tilde{\sigma}^N$ as follows. Consider a history of the original game at time t given by $h_i(t) = \big((P_i^t(t^1), P_i^p(P_i^t(i^1), P_i^s(s^1)), P_i^s(s^1)), \ldots, (P_i^t(t^m), P_i^p(P_i^t(i^m), P_i^s(s^m)), P_i^s(s^m))\big)$ for $i \in \{1, \ldots, I\}$ and let $k \in \{0, \ldots, 2^N - 1\}$ be such that $t \in I_k^N = \big(T\frac{k}{2^N}, T\frac{k+1}{2^N}\big]$. Assume that under history $h_i(t)$ at most one player has moved in each interval I_j^N with $j \in \{0, \ldots, k-1\}$ and no player has moved so far during I_k^N. Define $\tilde{\sigma}_i^N$ at history $h_i(t)$ as $\sigma_i^N(h_i^N)$ for history h_i^N of the simplified game given by $h_i^N(t) = ((t^N(t^1), P_i^p(P_i^t(i^1), P_i^s(s^1)), P_i^s(s^1)), \ldots, (t^N(t^m), P_i^p(P_i^t(i^m), P_i^s(s^m)), P_i^s$ with $t^N(t^k) = j$ if $t^k \in I_j^N$. That is $\tilde{\sigma}^N$ coincides with σ^N after histories in the simplified game that can be matched to histories to the original game. For private histories $h_i(t)$ in which two players have moved in an interval I_j^N with $j < k$ or some player has moved in interval I_k^N, strategy $\tilde{\sigma}^N(h_i(t))$ assigns probability ε to every available action and takes no action with the remaining probability.

Lemma 2. *Let σ^N be a strategy of $\mathscr{G}^N$ for each N and let $\tilde{\sigma}_N$ denote its associated strategy in the original game. As $N \to \infty$ we have $|U_i(\tilde{\sigma}^N) - U_i^N(\sigma^N)C(N)^{2^N}| \to 0$.*

Proof. The measure of the set of histories in which there are more than one arrival in an interval $\big[T\frac{k}{2^N}, T\frac{k+1}{2^N}\big]$ for $k \in \{0, 2^N - 1\}$ converges to zero as N converges to ∞. In fact, the probability that there is more than one arrival by any two players at any interval I_k^N, $1 - C(N)^{2^N}$, is less than or equal than

$$2^N \left(1 - e^{-\sum_{j=1}^N \lambda_j / 2^N} - \frac{\sum_{j=1}^N \lambda_j}{2^N} e^{-\sum_{j=1}^N \lambda_j / 2^N}\right),$$

which by L'Hôpital's rule approaches zero as N goes to ∞. Thus let $\tilde{\mathscr{A}}_N$ denote the set of histories in which the players have at most one arrival at each interval I_k^N. Then, we have

$$|\underbrace{\int_{\mathscr{A}} \sum_{h^T \in \mathscr{H}^T(a^T)} prob_i(h^T|a^T, \tilde{\sigma}^N) g(h^T) d\mu_i(a^T)}_{U_i(\tilde{\sigma}^N)} - \int_{\tilde{\mathscr{A}}_N} \sum_{h^T \in \mathscr{H}^T(a^T)} prob_i(h^T|a^T, \sigma^N) g(h^T) d\mu_i(a^T)| \to 0,$$

as $N \to \infty$. We conclude noting that

$$\begin{aligned} \int_{\tilde{\mathscr{A}}_N} \sum_{h^T \in \mathscr{H}^T(a^T)} prob_i(h^T|a^T, \sigma^N) \cdot g(h^T) d\mu_i(a^T) &= \sum_{a^N} \sum_{h^N \in \mathscr{H}(a^N)} g_i(h^N) \cdot prob_i(h^N|a^N, \sigma^N) \mu_i^N(a^N) \cdot C(N)^{2^N} \\ &= U_i^N(\sigma^N) C(N)^{2^N}, \end{aligned}$$

where $C(N)$ is defined by equation (13). □

Conversely, in some cases we can identify strategies in the original game with strategies in the approximating game. Consider a history $h = \big((t^1, i^1, s^1), \ldots, (t^m, i^m, s^m)\big)$ and define $H^N(h) = \{\tilde{h} : \tilde{h} = \big((\hat{t}^1, i^1, s^1), \ldots, (\hat{t}^m, i^m, s^m)\big), \quad t^N(\hat{t}^j) = t^N(t^j) \text{ for } j \in \{1, \ldots, m\}\}$. That is, $H^N(h)$ is the set histories with the same number of arrivals as h, in which the same players have arrivals in the same order and take the same actions as in h, and such that opportunity times are in the same I_k^N intervals as the opportunity times in history h. We say that a strategy of the original game is *constant in the N'th intervals* if $\sigma_i(s|h_i(t), t) = \sigma_i(s|\tilde{h}_i(t), t)$ for every $\tilde{h} \in H^N(h)$. When a strategy in the original game is constant in the N'th intervals it can be associated to a strategy in N'th game by a construction that is analogous to the one that associates σ^N to $\tilde{\sigma}^N$.

3.1.4 Limits of ε-constrained strategies

Consider the sequence of strategies of the original game $\tilde{\sigma}^N$, associated to ε-constrained equilibria in the approximating game. As $N \to \infty$ and their associated L^2 functions $f^N(\cdot)$ are given by,

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

The sequence of functions is bounded in L^2 norm and therefore has a sub-sequence which I denote f^N which converges to a function $f^* \in \big(L^2(\mathscr{H}, \nu)\big)^N$ in the weak-* topology, which by reflexivity of L^2 corresponds to the weak topology in L^2.[9]

In summary the functions $prob_i(h^T|a^T, \tilde{\sigma}^N)$ have a subsequence that converges in the weak topology to a limit f^*. Now, it follows as a Corollary of Lemma 1 that the candidate

[9] This statement follows from the Banach–Alaoglu theorem.

strategy associated to f^* is indeed a strategy, because for ε-constrained strategies $\pi_i(s|h,t,f^*)$ is well defined at every history, h.The weak limit of ε-constrained strategies represents a strategy of the game in a full-measure set of arrival histories.

Thus, we have established that the limit of ε-constrained strategies also constitutes a strategy. In the following section I show that the candidate strategy associated to the limit is also an ε-constrained equilibrium.

3.1.5 Limit strategy is an ε-constrained equilibrium

Let's now see that the candidate strategy associated to f^* is an ε-constrained equilibrium of the game. Let σ^* be the candidate strategy associated to f^*. By contradiction, suppose player i has a profitable deviation from σ^*. Then there is an ε-constrained strategy σ_i' for player i with associated probability function $prob(h|a,(\sigma_i',\sigma_{-i}^*))$, such that

$$\int_{\mathscr{A}} \left(\sum_{h\in\mathscr{H}(a)} g_i(h)\cdot prob_i(h|a,(\sigma_i',\sigma_{-i}^*)) \right) d\mu(a) \;>\; \int_{\mathscr{A}} \left(\sum_{h\in\mathscr{H}(a)} f_i^*(h,a)\cdot g_i(h) \right) d\mu(a) + 6\delta, \tag{15}$$

for $\delta > 0$. I will show that this implies that there is a profitable deviation to a an ε-constrained equilibrium of an N-th approximating game, which is a contradiction.

In the following Lemma I show that one can construct a strategy of player i that is constant in the intervals $[T\frac{k}{2^N}, T\frac{k+1}{2^N}]$ and the payoff that player i obtains from using this strategy is close to the payoff from σ_i'.

Lemma 3. *Let $\delta > 0$. There is a sequence of ε-constrained strategies of player i $\left(\sigma_i^N\right)'$, which are constant in the N'th intervals, and there is $\bar{N}$ such that for all $N \geq \bar{N}$ the difference in i's payoff between the strategy profiles $\left(\left(\sigma_i^N\right)', \tilde{\sigma}_{-i}^N\right)$ and $(\sigma_i', \sigma_{-i}^*)$ is at most δ.*

Proof. Because players' opportunities arrive at Poisson rates the probability that a private history of player i has more than K moves becomes vanishingly small as K becomes large.

Denote $\mathscr{H}_i^K$ for the set of private histories such that for every history in $\mathscr{H}_i^K$ player i is the last one to move and at player i's last move her private history has at most K elements. Consider a strategy of player i $\left(\sigma_i^K\right)'$ that coincides with σ_i' at histories in $\mathscr{H}_i^K$ and corresponds strategy $\tilde{\sigma}_i^N$ at all other histories. Because the measure of histories outside $\mathscr{H}_i^K$ converges to zero as K grows and the strategies are bounded, there is K big enough such that the difference in payoff from $\left(\left(\sigma_i^K\right)', \sigma_{-i}^*\right)$ and $\left(\sigma_i', \sigma_{-i}^*\right)$ is at most $\delta/3$. Let's now fix K so that the difference in payoff from $\left(\left(\sigma_i^K\right)', \sigma_{-i}^*\right)$ and $\left(\sigma_i', \sigma_{-i}^*\right)$ is at most $\delta/3$.

Since $\tilde{\sigma}_k^N$ converges to σ_k^* for each $k \neq i$, from condition (7) in Lemma 1 there is N_1 such that for all $N \geq N_1$ the difference in i's payoff from $\left(\left(\sigma_i^K\right)', \sigma_{-i}^*\right)$ and $\left(\left(\sigma_i^K\right)', \tilde{\sigma}_{-i}^N\right)$ is at most $\delta/3$.

The private histories in $\mathscr{H}_i^K$ can be represented as elements of $[0,T]^{K \cdot I}$, by keeping track of the timings of opportunities of each player. The distribution of these timings is Lebesgue measurable. The strategy $\left(\sigma_i^K\right)'$restricted to histories in $\mathscr{H}_i^K$, represented as a measurable function in $[0,T]^{K \cdot I}$ can be approximated by simple functions that are constant in rectangles with each side corresponding to an interval $[T\frac{k}{2^N}, T\frac{k+1}{2^N}]$ for some $k \in \{0,\ldots,2^N-1\}$. More precisely, there is a sequence of candidate strategies $\left(\sigma_i^{K,N}\right)'$, constant in rectangles with sides of length $1/2^N$, that converges almost surely to $\left(\sigma_i^K\right)'$ in histories with at most K events. In histories with more than K events the two strategies coincide. The sequence of candidate strategies can be chosen so that they correspond to ε-constrained strategies. If for instance we have $\sum_{s \in S_i(m_i(h,t))} \left(\sigma_i^{K,N}\right)'(s|h_i(t),t) > 1$ since $\sum_{s \in S_i(m_i(h,t))} \left(\sigma_i^{K,N}\right)'(s|h_i(t),t) \to_k 1$ almost surely we can sum and subtract a function that converges to zero almost surely so that $\sum_{s \in S_i(m_i(h,t))} \left(\sigma_i^{K,N}\right)'(s|h_i(t),t) = 1$. An analogous construction allows to ensure that the strategies put weight at least ε on all available actions. Thus, there is N_2 such that if $N \geq N_2$ the payoff from strategy profile $\left(\left(\sigma_i^K\right)', \tilde{\sigma}_{-i}^N\right)$ is at most $\delta/3$ apart from the payoff from strategy profile $\left(\left(\sigma_i^{K,N}\right)', \tilde{\sigma}_{-i}^N\right)$.

Combining the observations above we obtain that for $N \geq \max\{N_1,N_2\}$ the difference in i's payoff from strategy profile $\left(\sigma_i', \sigma_{-i}^*\right)$ and strategy profile $\left(\left(\sigma_i^{K,N}\right)', \tilde{\sigma}_{-i}^N\right)$ is at most δ. □

From the previous Lemma there are strategies $\left(\left(\sigma_i^{K,N}\right)', \tilde{\sigma}_{-i}^N\right)$ such that

$$\int_{\mathscr{A}} \left(\sum_{h \in \mathscr{H}(a)} g_i(h) \cdot prob\left(h|a, \left(\left(\sigma_i^{K,N}\right)', \tilde{\sigma}_{-i}^N\right)\right)\right) d\mu_i(a) + \delta \quad > \quad \int_{\mathscr{A}} \left(\sum_{h \in \mathscr{H}(a)} g_i(h) \cdot prob(h|a, (\sigma_i', \sigma_{-i}^*))\right) d\mu_i(a),$$

for strategy profile $\left(\left(\sigma_i^{K,N}\right)', \tilde{\sigma}_{-i}^N\right)$ that is constant in the intervals $[T\frac{k}{2^N}, T\frac{k+1}{2^N}]$ for $k \in \{0,\ldots,2^N-1\}$.

Now, because f^N converges to f^* in the weak topology, for N big enough,

$$\int_{\mathscr{A}} \sum_{h^T \in \mathscr{H}^T(a^T)} f^*(h,a) \cdot g_i(h,a)\, d\mu_i(a^T) > \int_{\mathscr{A}} \sum_{h^T \in \mathscr{H}^T(a^T)} f^N(h,a) \cdot g_i(h,a)\, d\mu_i(a) - \delta \quad (16)$$

By Lemma 2 for N big enough,

$$\int_{\mathscr{A}} \sum_{h \in \mathscr{H}(a)} prob\left(h|a, \left(\left(\sigma_i^{K,N}\right)', \tilde{\sigma}_{-i}^N\right)\right) \cdot g_i(h,a) d\mu(a) < U_i^N\left(\left(\sigma_i^{K,N}\right)'', \sigma_{-i}^N\right) C(N)^{2^N} + \delta$$

where $(\sigma_i^{K,N})''$ is the strategy in the N'th simplified game that is associated to $\left(\sigma_i^{K,N}\right)'$. Also by Lemma2 for big enough N,

$$\int_{\mathscr{A}} \left(\sum_{h^T \in \mathscr{H}^T(a^T)} f_i^*(h^T, a^T) \cdot g_i(h^T) \right) d\mu(a^T) > U_i^N(\sigma^N) C(N)^{2^N} - \delta.$$

Combining with equation (15) we obtain

$$U_i^N(\sigma^N) C(N)^{2^N} + \delta < U_i^N((\sigma^{K,N})_i'', \sigma_{-i}^N) C(N)^{2^N}.$$

However, this last equation contradicts that σ^N is an equilibrium of the simplified game $\mathscr{G}^N$. Thus, the original game must have an ε-constrained equilibrium that corresponds to the strategies associated to f^*.

3.2 Existence of trembling hand equilibrium

I next present the main result of the paper. Once we have established that each game has a ε-constrained equilibrium existence of a trembling hand equilibrium follows. In the following theorem I show that I can construct a trembling hand perfect equilibrium from the limit of the L^2 functions associated to ε-constrained equilibria of the original game.

Theorem 1. *A trembling hand equilibrium exists.*

Proof. Let σ^m be sequence of ε_m-constrained equilibria with ε_m converging to zero. The sequence of functions $f^m(h,a) = prob(h|a, \sigma^m)$ in L^2 induced by σ^m must have a convergent subsequence in the weak topology since the sequence of functions is bounded in L^2 and hence weak-* compact. Passing to the subsequence, let f^* denote its limit. By Lemma 1, in the set $H^+ = \{(h,a) : \pi_i(s|h,t,f^*) \text{ is well defined}\}$ the strategies σ^m converge weakly to the candidate strategy associated to f^*, $\pi_i(s|h,t,f^*)$ in all histories in H^+. The complication is that it is not immediately obvious how to recover the limit strategies when π_i is not well defined. These are the strategies of the players in the off-path histories. Let $\mathscr{A}^j$ denote the set of arrival histories with exactly j arrivals and define $H_j^1 = \{(h,a) : a \in \mathscr{A}^j, f^*(h,a) = 0\}$. For a history (h,a) let $(h,a)^j$ denote the history composed by the tuples in (h,a) up to and including the j'th one. Define

$$H^1 \quad = \quad \left\{(h,a) : \exists j \text{ such that } a \in \mathscr{A}^k \text{ for } k \geq j+1 \text{ and } (h,a)^j \in H_j^1\right\}.$$

$\bar{H}^1$ is the set of histories such that there is j such that $f^*((h,a)^j) = 0$ but $f^*\left((h,a)^{j-1}\right) > 0$. These are the histories after an off-path move. Let

$$\bar{H}^1 = \left\{(h,a) : \exists j \text{ such that } a \in \mathscr{A}^j,\, (h,a)^j \in H_j^1 \text{ and } (h,a)^{j-1} \notin \left(H_{j-1}^1\right)\right\}.$$

$\bar{H}^1$ is the set of histories that end at the first off-path move. Let $(h,a)^{-1}$ denote the history that contains all the elements in (h,a) except for the last element in (h,a).

In order to recover the strategies after histories after an off-path action let $h = ((t^1,i^1,s^1),\ldots,(t^n,i^n,s^n))$ and define

$$f^{1,m}(h,a) = \begin{cases} f^m(h,a) & \text{if } h \notin H^1 \\ \Pi_{j=1,(h,a)^j \notin \bar{H}^1}^n \sigma_{i^j}^m(s^j|h_{i^j}(t^j),t^j) & \text{if } h \in H^1 \end{cases}$$

$f^{1,m}$ is constructed so that after histories in which a player took an off-path action the probability function associated to the histories after the deviation "omits" the probability associated to the zero probability action. In histories without off-path moves the sequence corresponds to the probability function associated to the sequence of ε-constrained equilibria. The sequence has a subseqence of f^m, f^{1,M_m^1}, with $M_m^1 \subseteq \mathbb{N}$, that converges to a limit $f^{1,*}(h,a)$ in the weak topology. From $f^{1,*}$ we can define a candidate strategy for each player i as

$$\sigma_i^{1,*}(s|h_i(t),t,f^{1,*}) = \begin{cases} \frac{f^{1,*}((h,(t,i,s)),(a,(t,i)))}{f^{1,*}(h,a)} & \text{if } (h,a) \notin \bar{H}^1 \\ \frac{f^{1,*}((h,(t,i,s)),(a,(t,i)))}{f^{1,*}\left((h,a)^{-1}\right)} & \text{otherwise.} \end{cases}$$

By Lemma 1, the sequence of strategies $\{\sigma^{M_m^1}\}_{m\in\mathbb{N}}$ converges in the weak topology to $\sigma^{1,*}$on histories in which $\sigma^{1,*}$ is well defined. That is, it converges in histories with no or only one deviation from $\sigma^{1,*}$. Also, σ_1^* must correspond to a strategy.

Recursively, in order to recover the strategies of the players after histories with N deviations I define as before

$$H_j^N = \{(h,a) : a \in \mathscr{A}^j, f^{N-1,*}(h,a) = 0\},$$

$$H^N = \left\{(h,a) \in H^{N-1} : \exists j \text{ such that } a \in \mathscr{A}^N \text{ for } k \geq j+1 \text{ and } (h,a)^j \in H_j^N\right\},$$

$$\bar{H}^N = \left\{(h,a) \in H^{N-1} : \exists j \text{ such that } a \in \mathscr{A}^j,\ (h,a)^j \in H_j^{N-1} \text{ and } (h,a)^{j-1} \notin \left(H_{j-1}^{N-1}\right)\right\},$$

$$f^{N,M_m^N}(h,a) = \begin{cases} f^{N-1,M_m^{N-1}}(h,a) & \text{if } h \notin H^N \\ \Pi_{j=1,(h,a)^j \notin \bar{H}^N}^n \sigma_{ij}^{M_m^{N-1}}(s^j|h_{ij}(t^j),t^j) & \text{if } h \in H^N \end{cases}.$$

Denote $f^{N,*}$ the limit in the weak topology of a subsequence $\{f^{N,M_m^N}\}_{m\in\mathbb{N}}$ of $\{f^{N,M_m^{N-1}}\}_{m\in\mathbb{N}}$. The candidate strategy associated to $f^{N,*}$ is defined as

$$\sigma_i^{N,*}(s|h_i(t),t,f^{N,*}) = \begin{cases} \frac{f^{N,*}((h,(t,i,s)),(a,(t,i)))}{f^{N,*}(h,a)} & \text{if } (h,a) \notin \bar{H}^N \\ \frac{f^{N,*}((h,(t,i,s)),(a,(t,i)))}{f^{N,*}\left((h,a)^{-1}\right)} & \text{otherwise.} \end{cases}$$

By Lemma 1 the subsequence $\left\{\sigma^{M_m^N}\right\}_{m\in\mathbb{N}}$ converges in the weak topology to $\sigma^{N,*}$ on the set in which $\sigma^{N,*}$ is well defined. By construction, the strategy $\sigma^{N,*}$ is well defined in histories with at most N player arrivals. Furthermore, when $\tilde{N} > N$, $f^{N,*}$ coincides with $f^{\tilde{N},*}$, in all histories with at most N arrivals. This observation follows from the fact that $f^{\tilde{N},M_m^{\tilde{N}}}(h,a)$ is a subsequence of $f^{N,M_m^N}(h,a)$ whenever a has at most N arrivals and the corresponding subsequences are convergent. Therefore, the subsequences must have the same limit for those histories. As a result, $\sigma^{N,*}$ coincides with $\sigma^{\tilde{N},*}$ in histories with at most N player opportunities.

Define the strategy σ^{**} as follows. For each history h and each player i $\sigma_i^{**}(s|h_i(t),t) = \sigma_i^{N,*}(s|h_i(t),t,f^{N,*})$ where N is taken so that it is greater than the total number of arrivals in h (note that by Lemma 1 σ_i^{**} is measurable with respect to i's information so there is no problem with choosing h arbitrarily).

Now consider the sequence of strategies $\{\tilde{\sigma}^m\}_{m\in\mathbb{N}}$ defined as

$$\tilde{\sigma}^m = \sigma^{M_m^m}.$$

Let's see that the limit $\tilde{\sigma}_m$ is a trembling hand equilibrium. Note that $\tilde{\sigma}_m$ is an $\varepsilon^{M_m^m}$-constrained equilibrium with $\varepsilon^{M_m^m} \to 0$.

I now show that $\tilde{\sigma}_m$ converges to σ^{**} in the weak topology. Let $\delta > 0$ and let $\varphi_i \in L^2(\mathscr{H}, \mu_i)$. There is K sufficiently large such that the set of arrival histories with more than K arrivals, defined as $\mathscr{A}^{>K} = \{a : a \in \mathscr{A}^k, k > K\}$, has sufficiently small measure, such that for each m,

$$\int_{A^{>K}} \left(\sum_{h \in \mathscr{H}(a)} (\Gamma_i(h, \tilde{\sigma}_i^m) - \Gamma_i(h, \sigma_i^{**})) \cdot \varphi_i(h,a) \right) d\mu_i(a) < \frac{\delta}{2}.$$

Now, for $\tilde{\sigma}^k$ with $k \geq K$ we know that $\tilde{\sigma}^{k,*}$ is well defined in histories arising from arrival histories in $\left(\mathscr{A}^{>K}\right)^c$, that $\sigma^{k,*} = \sigma^{**}$ on those histories and that, therefore, there is $\bar{M}$ such that for all $m \geq \bar{M}$

$$\int_{\left(\mathscr{A}^{>K}\right)^c} \left(\sum_{h \in \mathscr{H}(a)} (\Gamma_i(h, \tilde{\sigma}_i^m) - \Gamma_i(h, \sigma_i^{**})) \cdot \varphi_i(h,a) \right) d\mu_i(a) < \frac{\delta}{2}.$$

Thus, we conclude that

$$\int_{\mathscr{A}} \left(\sum_{h \in \mathscr{H}(a)} (\Gamma_i(h, \tilde{\sigma}_i^m) - \Gamma_i(h, \sigma_i^{**})) \cdot \varphi_i(h,a) \right) d\mu_i(a) < \delta.$$

In conclusion, $\tilde{\sigma}^m$ converges to σ^{**} in the weak topology and, therefore, σ^{**} is a trembling hand perfect equilibrium. □

3.3 A Trembling Hand Perfect Equilibrium is a Nash Equilibrium.

In this section I show that a trembling hand perfect equilibrium is also a Nash equilibrium.

Definition 3. A strategy profile $\sigma = (\sigma_1, \ldots, \sigma_I)$ is a Nash Equilibrium if

$$U_i(\sigma_i, \sigma_{-i}) \geq U_i(\sigma_i', \sigma_{-i})$$

for every σ_i' and every player i.

Let me now show that a trembling hand perfect equilibrium is also a Nash equilibrium of this game. In fact, let σ^* be a trembling hand equilibrium and let σ^m be the sequence of ε^m-constrained equilibria that converge to σ^*. Let σ_i' be a strategy of player i. We can construct a strategy $(\sigma_i')^m$ that is ε^m-constrained and converges almost surely to σ_i'.[10] From Lemma

[10] We can construct the approximating strategy by setting the probability of all actions that are taking with probability zero to ε^m and subtracting the appropriate probability from the action that is taken with the greatest probability. For small enough ε^m this construction can be done.

1 it follows that $prob(h|a,\sigma^m)$ and $prob\left(h|a,\left((\sigma_i')^m,\sigma_{-i}^m\right)\right)$ converge to $prob(h|a,\sigma^*)$ and $prob(h|a,(\sigma_i',\sigma_{-i}^*))$ in the weak topology, respectively. From the definition of weak convergence, by taking limits on both sides of (3), we can see that σ^* must be a Nash equilibrium of the original game, that is,

$$\int_{\mathscr{A}} \sum_{h^T\in\mathscr{H}^T(a^T)} g_i(h^T)\cdot prob(h^T|a^T,(\sigma_i',\sigma_{-i}^*))d\mu_i(a^T) \le \int_{\mathscr{A}} \sum_{h^T\in\mathscr{H}^T(a^T)} prob(h^T|a^T,\sigma^*)\cdot g_i(h^T)d\mu_i(a^T)$$

for every strategy of player i, σ_i'.

And, thus, we establish that trembling hand perfect equilibria must indeed be Nash as stated in the following Lemma.

Lemma 4. *A trembling hand equilibrium is a Nash equilibrium.*

References

[1] Charalambos D Aliprantis and Kim C Border. *Infinite dimensional analysis*, volume 32006. Springer, 1999.

[2] Attila Ambrus, James Burns, and Yuhta Ishii. Gradual bidding in ebay-like auctions. *Working Paper*, pages 1–50, 2014.

[3] Attila Ambrus and Shih En Lu. A continuous model of multilateral bargaining with random arrival times. *American Economics Journal: Microeconomics*, 2014.

[4] Riccardo Calcagno, Stefano Lovo, Yuichiro Kamada, and Takuo Sugaya. Asynchronous Revision Games with Deadline : Unique Equilibrium in Coordination Games. *Forthcoming Theoretical Economics*, pages 1–59, 2013.

[5] Yuichiro Kamada and Michihiro Kandori. Revision games. *Working Paper*, pages 1–23, 2011.

[6] Yuichiro Kamada and Nozomu Muto. Multi-Agent Search with Deadline. *Working Paper*, pages 1–77, 2014.

[7] Sofia Moroni. Proxy auctions with a deadline. *Working Paper*, 2014.

[8] Reinhard Selten. Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games. *International journal of game theory*, 1975.