University of Pittsburgh
Department of Economics
Dietrich School of Arts and Sciences

WORKING PAPER SERIES

20/001

Existence of Trembling hand perfect and sequential equilibrium in Stochastic Games

Sofia Moroni

February, 2020

University of Pittsburgh • Dietrich School of Arts and Sciences • Department of Economics
4700 Wesley W. Posvar Hall • 230 South Bouquet Street • Pittsburgh, PA 15260

Existence of Trembling hand perfect and sequential equilibrium in Stochastic Games

Sofia Moroni*

University of Pittsburgh

moroni@pitt.edu

February 2020

Abstract

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 stochastic games that may feature incomplete and imperfect information. These equilibria do not necessitate the use of a public correlating device. Under further regularity assumptions each stochastic game has a sequence of approximating finite games whose equilibria approximate equilibria of the limit game.

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

*I am grateful to Juan Escobar, Faruk Gul, Adam Kapor, Roger Myerson, Wolfgang Pesendorfer, Phil Reny, Luca Rigotti, Larry Samuelson, Ali Shourideh, Can Urgun and Richard Van Weelden for their insightful comments.

1 Introduction

In this paper we propose notions of trembling hand perfect and sequential equilibrium that extend Selten (1975) and Kreps and Wilson (1982)'s definitions to a large class of stochastic games with informational asymmetries. This class encompasses games with an infinite state space and finite actions in which players receive signals that may be informative about states visited, about past play or about their opponents' observations. We show that under two assumptions—one restricting the payoff function and another requiring a noisy observation of the opponents' signals—trembling hand perfect equilibria exist and, under further regularity assumptions, these equilibria are also sequential equilibria. These results are the first proofs of existence for a general class of stochastic games with infinite state spaces which do not require the existence of a public correlating device. Two particular cases that satisfy the restriction on payoffs are stochastic games with discounted payoffs and games with stochastic move opportunities, such as revision games, that feature a finite expected length.[1]

A *trembling hand perfect equilibrium* is defined as a limit of totally mixed strategies that are constrained equilibria, in which players must put positive weight on every available action. This limit is in *weak convergence of strategies*, which requires the convergence of probabilities of histories of play for each player.[2] Our assumptions guarantee that under the resulting topology the space of strategies is compact and the players' payoffs are continuous in strategies. These properties guarantee existence of a trembling hand equilibrium and imply that it is a Nash Equilibrium. Under further regularity assumptions the equilibrium is, for each player, a best response to a limit of beliefs that are consistent with totally mixed approximating strategies.[3] A strategy that satisfies the latter property of sequential rationality is what we call a *weak sequential equilibrium*. Weak limits coincide with standard limits in countable spaces, and, therefore, our notion coincides with Kreps and Wilson (1982)'s in finite and countably infinite games.

Our assumption on payoffs requires that a sequence of bounds for each period t's expected payoff be summable. It includes, in addition to discounted games, games in which payoffs decrease more slowly over time. It also applies to most of the games with stochastic

[1] This class includes repeated games and games of finite length.

[2] The function from player's signal realizations to probabilities over actions can be viewed as a function in an L^2 space. The convergence is in the weak topology of L^2, starting from each deviation from a limit strategy.

[3] The limit of beliefs is taken in the weak-* of the dual of L^∞ and it results in a finitely additive measure.

move opportunities with a deadline that have been studied in the literature. The assumption on the (un)observability of the opponents signals is more delicate. It requires that the joint distribution of the players' signals up to each period t be absolutely continuous with respect to the product measure of the marginals of each player's signal sequence up to t. This assumption—which we dub *noisy observability*—is a generalization of the absolute continuity assumption in Milgrom and Weber (1985) to the context of dynamic games. Noisy observability holds, for example, in games with countable signal spaces and in Markov stochastic games with private monitoring with uncountable signal spaces in which each player's signal has some idiosyncratic noise.[4]

Our final set of results relates to sequences of games that may approximate the original stochastic game. We define a notion of an approximating game sequence that has the feature that its constrained equilibria converge to constrained equilibria of the original game. It requires, among other conditions, that there be finer and finer countable partitions of the state and signal spaces in which the available action set and the payoffs are well behaved. We also show that any game that is Markov, has compact state and signal spaces and payoffs continuous in the state has a finite approximating game sequence.[5] These results do not require the noisy observability assumption and may be useful to prove existence in settings in which the assumption fails. They imply, in particular, that perfect information asynchronous moves games with compact state spaces and continuous payoffs (as a function of the state), with either discounted payoffs or a finite expected length have Markov perfect equilibria.

As a motivation for our interest in these games, consider a setting in which players receive stochastic opportunities to move at times between 0 and a deadline at time T. At an opportunity players choose an available action after observing a signal that is informative about a payoff relevant state, the rate at which moves will realize in the future, and/or past play of the opponents. In particular, after some actions players may not learn that an opponent received an opportunity at all.[6] This type of game can model realistic settings in which players have uncertainty about the timing of their moves and their opponents'. As the rate at which the players receive move opportunities grows large, the game approaches a continuous time game while avoiding the technical issues with the definition of strategies

[4]See section 4.1

[5]The key is the finite subcovers of the compact state and signal spaces define a suitable sequence of approximating partitions.

[6]In games in which players can choose to "keep their own action" and opportunities arrive frequently, this assumption—which requires imperfect information—is the most natural.

that arise in continuous time settings. In these games it is crucial that the timing of play be a part of the state, which must, therefore, be uncountable. Thus, due to the dimensionality of the space of histories and the informational assumptions, existing proofs of existence do not apply.

Several applied settings have been studied using games with stochastic move opportunities. 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.[7] 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.[8] Candidates must choose their stance, announce their policy proposals and disclose opposition research at strategic times before an election.[9]

The literature on stochastic games, stemming from the seminal work by Shapley (1953) focuses mostly on Markov games with almost perfect information with discounted payoffs and shows existence of Markov perfect equilibria.[10] Due to the potentially non-stationary nature of our setting the techniques developed in the stochastic games literature with almost complete information games cannot be readily applied to our setting.[11] Fink (1964), Sobel (1971) and Takahashi (1964) show existence of Markov perfect equilibria under a finite state space. Parthasarathy (1973) shows existence in a two player game with a countable set of states. Federgruen (1978) and Whitt (1980) show existence of Markov perfect

[7] Ambrus et al. (2014), Hopenhayn and Saeedi (2016), and Kapor and Moroni (2016) are examples of 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.

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

[9] Kamada and Sugaya (2020) and Kamada and Kandori (2017) consider stochastic-timing models of election campaigns.

[10] Players may move simultaneously but they observe the present state and all previous moves in each period.

[11] 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).

equilibria when the state space is countable. When the state space is uncountable, however, examples of nonexistence have been constructed. Simon (2003) provides an example of a one stage game with three players and finite actions in which there is no equilibrium. In the three stage game in Harris et al. (1995) there is no subgame perfect equilibrium in a game with a continuum of actions.[12] The standard approach for proving existence—which involves using Kakutani-Fan-Glicksberg over a compact set of actions—does not work. If one defines a topology in which strategies live in a compact space the limits of these strategies may not be strategies themselves but instead *correlated strategies*.[13] Many authors "close" the space of strategies by assuming the game has some way of correlating the players' strategies. Duggan (2012) adds potentially payoff relevant noise that is conditionally independent of the current state and actions. Harris et al. (1995) add a public signal that serves as a correlating device. Manelli (1996) adds cheap talk to a signaling game to obtain existence. Our approach, in contrast, is to assume that there is noise in the observation of the opponents' observations.

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 an ε-equilibrium in a setting with the average payoffs criterion. 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.

Our paper also relates to a literature that proposes notions of sequentially rational equilibria for dynamic games. The closest paper, Myerson and Reny (2019), considers games with infinite state and action spaces with finite horizon and defines an equilibrium notion, a perfect conditional equilibrium which is sequentially rational. A perfect conditional equilibrium is defined as a distribution over states and actions that is the limit of a net of perfect

[12]Their game can be expressed as a stochastic game in which the action is part of the state in period 3.

[13]This issue has been pointed out, for example, by Harris et al. (1995), Börgers (1991) and Simon and Stinchcombe (1989). Myerson and Reny (2019) call it "strategic entanglement".

conditional ε-equilibria (in which players optimize their payoffs up to ε conditional on every measurable set of private histories and are close to ε-equilibria of slightly perturbed games). A perfect conditional equilibrium may involve correlated strategies. The main difference with our paper is that we consider the limit of equilibria in which players are constrained to utilize perfectly mixed strategies and show that this limit—a trembling hand perfect equilibrium—is sequentially rational under further technical assumptions. Because we assume a finite action space—while theirs is infinite—and make different informational assumptions we are able to show existence of an equilibrium without correlations.

2 Stochastic Games with Incomplete and Imperfect Information

We consider a game in which players choose actions in each period $t \in \mathbb{N}$. The game can be described by the list $\Gamma = (N, (\Omega, \mathscr{F}), (S_i, \mathcal{S}_i), (X_i, A_i, g_i), (\mu^\omega, \mu^s))$ where N is a finite set of n players, Ω is a state space with Borel measurable sets in $\mathscr{F}$, S_i is a space of signals for each player i with Borel measurable sets $\mathcal{S}_i$, X_i is a finite set of actions for each player i, $X = \Pi_{i \in N} X_i$ with generic element—an *action profile*—$\mathbf{a} = (a_1, \ldots, a_n) \in X$, $A_i : \cup_{t \in \mathbb{N}} \Omega^t \times X^{t-1} \rightrightarrows X_i$ is a measurable correspondence that yields the actions available to player i as a function of the history realized states and the history of play, $g_i : \cup_{t \in \mathbb{N}} \Omega^t \times X^t \to \mathbb{R}$ is measurable in $\omega^t \in \Omega^t$ and represents the flow payoff received by player i as a function of history up to the period of play, and μ^ω and μ^s are transition probabilities which we describe in what follows.[14]

The state transitions from one period to the next according to the *state transition probability function* $\mu^\omega : \cup_{t \in \mathbb{N}} \Omega^t \times X^t \times \mathscr{F} \to [0,1]$. This function determines the probability with which each state is drawn as a function of the states that were visited previously and the actions taken by the players. That is, $\mu^\omega(Z|\omega^t, \mathbf{a}^t)$ is the probability that the state drawn at time $t+1$ is in the set $Z \in \mathscr{F}$ given that states $\omega^t = (\omega_1, \ldots, \omega_t)$ were drawn and action profiles $\mathbf{a}^t = (\mathbf{a}_1, \ldots \mathbf{a}_t)$ were played in periods 1 through t, and it is measurable as function of $(\omega^t, \mathbf{a}^t)$.

We are interested in games with imperfect and incomplete information in which each player may receive signals about previous play, about players' types, and about their payoffs. Specifically, at each time t each player i observes a signal $s_{i,t} \in S_i$ before choosing her

[14]We use the notation $Y^t := \times_{i=1}^t Y$ for any set Y.

action in $A_i(\omega^t, \mathbf{a}^{t-1})$ at time t. Player i's signals are drawn according to the *signal transition probability function* $\mu_i^s : \left(\cup_{t\in\mathbb{N}}\Omega^t \times X^{t-1}\right) \times \mathcal{S}_i \to [0,1]$, where $\mu_i^s(Z|\omega^t, \mathbf{a}^{t-1})$ is the probability that player i's period-t signal $s_{i,t}$ satisfies $s_{i,t} \in Z \in \mathcal{S}_i$ after $\omega^t \in \Omega^t$ and $\mathbf{a}^{t-1} \in X^{t-1}$, and it is measurable as a function of $(\omega^t, \mathbf{a}^{t-1})$.[15] The measurable space of joint signals of all players is $(\mathbf{S} := \times_{i\in N} S_i, \mathbf{\mathcal{S}} := \otimes_i \mathcal{S}_i)$ and $\mu^s : \left(\cup_{t\in\mathbb{N}}(\Omega^t \times X^{t-1})\right) \times \mathbf{\mathcal{S}} \to [0,1]$ denotes the joint transition probability on that space. For each player i, S_i may contain an element $\star$, which represents "no observation". When $\star$ is drawn a player is unaware that a move took place. In an application a player may be drawn to play in some periods and not others and may be unaware of the moves that took place in periods in which she was not called to play. To make our notation consistent, whenever $s_i = \star$, we assume i takes action $a_\star$ with probability 1 (and does not remember it).[16,17]

Player i's signal $s_i \in S_i$ may contain information about $s_j \in S_j$ for $j \neq i$. For example, in a game of *almost complete information* it is common knowledge that $S_i = S_j = \Omega \times X$ and at each period t, the period-t signals are given by $s_{i,t} = s_{j,t} = (\omega_t, \mathbf{a}_{t-1})$ for each player $i, j \in N$, with $\mathbf{a}_0 = \emptyset$.[18] More generally, our model encompasses many different informational settings. Players may have types that affect their likelihood of drawing states which they learn privately at time 1. Players may not observe past actions but may receive an informative signal about the action profiles realized in previous periods. Additionally, they may observe partitional or other information regarding an opponent's payoff relevant variable.

Histories A *time-t state history* consists of all the states that were visited up to time t. A generic state history is written as $\omega^t = (\omega_1, \omega_2, \ldots \omega_t)$. A *time-$t$ history of play* contains the action profiles selected by the players up to time t. A generic history of play is written as $\mathbf{a}^t = (\mathbf{a}_1, \mathbf{a}_2, \ldots, \mathbf{a}_t)$. A *time-$t$ history of signals* consists of all the signals realized up to time t. A generic time-t history of signals is written as $\mathbf{s}^t = (\mathbf{s}_1, \mathbf{s}_2, \ldots, \mathbf{s}_t)$, where $\mathbf{s}_l = \left(s_{i,l}\right)_{i\in N}$ is the profile of signals realized at time $l \leq t$. A generic *time-t history* $h = (\omega^t, \mathbf{a}^t, \mathbf{s}^t)$ is composed of a state history, a history of play, and a history of signals. We use $\hat{\omega}(h)$, $\hat{\mathbf{a}}(h), \hat{\mathbf{s}}(h)$ to denote the history of states, play and signals in history h, respectively. For

[15] Because μ_i^s depends on all previous states, each player may learn more about states that were realized before time t, at time t.

[16] Moves that a player is unaware of do not enter her private histories, as defined below.

[17] Formally, whenever a player's signal at time t is $\star$, the state is such that $A_i(\omega^t, \mathbf{a}^{t-1}) = \{a_\star\}$.

[18] These games are called "games of almost perfect information" because players observe all past moves and move simultaneously.

$\mathbf{a}^t \in X^t$, $\mathbf{a}^t_{i,\tau}$ denotes player i's action in the action profile $(\mathbf{a}^t)_\tau$, for $\tau \le t$. $\omega^t_{i,\tau}$ and $\mathbf{s}^t_{i,\tau}$ are defined analogously. $\mathbf{a}^{t,(l)}$ for $l \le t$ denotes the truncation of $\mathbf{a}^t$ up to and including period l. $\omega^{t,(l)}$ and $\mathbf{s}^{t,(l)}$ are defined analogously. The length of a history h is denoted $|h|$. The set of all time-t histories is denoted $\mathscr{H}^t$ and it is endowed with the product measure. $\mathscr{H} := \bigcup_{t\in\mathbb{N}} \mathscr{H}^t$ is the set of histories. The σ-algebra of measurable sets in $\mathscr{H}$ is denoted $\mathcal{H}$.[19]

A private history of a player i at time t, $h^t_i = \Big((s_{i,l})_{l\le t, s_{i,l}\neq\star}, (a_{i,l})_{l\le t-1, a_{i,l}\neq a_\star}\Big)$ consists of the signals $s_{i,l}$, other than $\star$, received at each period $l \le t$, and player i's actions, other than $a_\star$, up to period $t-1$. The set of player-i, time-t private histories is denoted $\mathscr{H}^t_i$ and the set of player i private histories is given by $\mathscr{H}_i := \cup_{t\in\mathbb{N}} \mathscr{H}^t_i$.[20] The set of measurable sets in $\mathscr{H}_i$ is denoted $\mathcal{H}_i$.[21] There is a unique private history of player i associated with each game history $h \in \mathscr{H}^t$ and it is denoted by $h_i(h)$.[22] Since a player's private history depends only on her signals and actions we may also write $h_i(s^t_i, \mathbf{a}^{t-1}_i)$ for i's private history after signal realization $s^t_i \in S^t_i$ and sequence of actions profiles $\mathbf{a}^{t-1}_i \in X^{t-1}_i$. For $h \in \mathscr{H}^t$, the truncation of h up to period $l \le t$ is denoted $h^{(l)}$ and player i's period-l private history is given by $h^{(l)}_i(h) := h_i(h^{(l)})$.

The correspondence A_i must be measurable with respect to each player i's private history. That is, we can write $A_i : \mathscr{H}_i \rightrightarrows X_i$.

Strategies A strategy of player i, $\sigma_i : \mathscr{H}_i \to \Delta X_i$, is a function from i's private histories to probability distributions over i's actions satisfying $\text{supp}\sigma_i(h^t_i) \subseteq A_i(h^t_i)$. The set of player i strategies is denoted Σ_i. A strategy profile is of the form $\sigma := \times_i \sigma_i \in \times_i \Sigma_i := \Sigma$.

Expected Payoffs Consider a strategy $\sigma = \times_{i\in N}\sigma_i$ and a history $(\omega^t, \mathbf{a}^t, \mathbf{s}^t)$. The strategy induces a probability, $prob(\mathbf{a}^t|\omega^t, \mathbf{s}^t)$ over histories (of actions) conditional on the realizations in $(\omega^t, \mathbf{s}^t)$. For history $h = (\omega^t, \mathbf{a}^t, \mathbf{s}^t)$,

$$prob(\mathbf{a}^t|\omega^t, \mathbf{s}^t, \sigma) := \prod_{l\le t}\prod_{i\in N} \sigma_i(\mathbf{a}^t_{i,l}|h^{(l)}_i(h)),$$

[19]The measurable sets in $\mathscr{H}$ are defined by the product $\sigma-$algebras.

[20]Notice that a player may be unaware of the number of moves that have occurred at the time that she observes her private history. Therefore a private history $h_i \in \mathscr{H}_i$ may belong to $\cap_{k\in I}\mathscr{H}^k_i$ for some set I with $|I| > 1$.

[21]The σ-algebra is defined by product σ-algebras.

[22]h_i is $\mathcal{H}/\mathcal{H}_i$ measurable.

where $\sigma_i(\mathbf{a}^t_{i,l}|h_i^{(l)}(h))$ is the probability that player i assigns to action $\mathbf{a}^t_{i,l}$ at private history $h_i^{(l)}(h)$ when she moves at time l. We call the function $prob(\mathbf{a}^t|\omega^t,\mathbf{s}^t,\sigma)$ the *probability function associated to* σ.

Abusing notation we also write $prob(h,\sigma)$ for $prob\left(\hat{\mathbf{a}}(h)|\hat{\omega}(h),\hat{\mathbf{s}}(h),\sigma\right)$. We define also, $prob_{-i}(\mathbf{a}^t|\omega^t,\mathbf{s}^t,\sigma) := \prod_{l\le t}\prod_{j\in N\setminus\{i\}}\sigma_j(a_{j,l}|h_j^{(l)}(h))$ and, $prob_i(\mathbf{a}^t|\omega^t,\mathbf{s}^t,\sigma) := \prod_{l\le t}\sigma_i(a_{i,l}|h_j^{(l)}(h))$which we also denote $prob_{-i}(h,\sigma)$ and $prob_i(h,\sigma)$, respectively.

Player i's payoff when players follow strategy σ is given by[23]

$$U_i(\sigma) := \sum_{t=1}^{\infty}\int_{h\in\mathscr{H}^t} g_i(h)\cdot prob(h,\sigma)\; d\mu^t(h),$$

where μ^t is the measure over $\mathscr{H}^t$ induced by μ^ω and μ^s and the counting measure over X.[24]

3 Equilibrium

Each function $prob(\cdot,\sigma)$ associated with a strategy $\sigma\in\Sigma$ can be viewed as a function in an L^2 space endowed with a measure $\hat{\mu}$ over $\mathscr{H}$. The measure $\hat{\mu}$ is such that the inner product between two functions $\psi,\phi:\mathscr{H}\to\mathbb{R}$ in $L^2(\mathscr{H},\hat{\mu})$ is given by

$$\langle\psi,\phi\rangle = \sum_{t=1}^{\infty}\int_{h\in\mathscr{H}^t}\delta^t\cdot\psi(h)\cdot\phi(h)\,d\mu^t(h),$$

for fixed $\delta\in(0,1)$, and the norm of $f\in L^2(\mathscr{H},\hat{\mu})$ is given by $\|f\|_2 = \langle f(h),f(h)\rangle^{1/2}$.[25]

We endow each $L^2(\mathscr{H},\hat{\mu})$ space with the topology of weak convergence.[26] Notice that

[23] $g_i(h)$ denotes $g_i\left(\hat{\omega}(h),\hat{\mathbf{a}}(h)\right)$.

[24] Formally, let $\mu^{\omega,\mathbf{s}}(\cdot|\omega^l,\mathbf{a}^l)$ be the joint measure over $\Omega\times\mathbf{S}$ induced by $\mu^\omega(\cdot|\omega^l,\mathbf{a}^l)$ and $\mu^s(\cdot|\omega^{l+1},\mathbf{a}^l)$. The measure $\mu^t(\cdot)$ on $\mathscr{H}^t$ is given by

$$\mu^t(Z) = \sum_{\mathbf{a}^t\in X^t}\int_{\Omega^t\times\mathbf{S}^t}\mathbf{1}\{(\omega^t,\mathbf{s}^t,\mathbf{a}^t)\in Z\}\,d\mu^{\omega,\mathbf{s}}(\omega_t,\mathbf{s}_t|\omega^{t-1},\mathbf{s}^{t-1},\mathbf{a}^{t-1})\cdots d\mu^{\omega,\mathbf{s}}(\omega_1,\mathbf{s}_1|\emptyset).$$

for each measurable set Z in $\mathscr{H}^t$.

[25] The measure $\hat{\mu}$ over $\mathscr{H}^t$ that yields the inner product is given by

$$\hat{\mu}(Z) = \sum_{t=1}^{\infty}\delta^t\int_{h\in\mathscr{H}^t}\mathbf{1}\{\tilde{h}\in Z,\tilde{h}^{(t)}=h\}\,d\mu^t(h),$$

for each measurable set $Z\subseteq\mathscr{H}$.

[26] The topology is characterized by its convergent nets: a net $\{f_\alpha\}$ converges to f^* in the weak topology if

for every $\sigma \in \Sigma$, $prob(\cdot,\sigma)$ is in $L^2(\mathscr{H},\hat{\mu})$ and $\left\|prob(\cdot,\sigma)\right\|_2 < \left(\frac{\delta}{1-\delta}\right)^{1/2}$. Thus, by the Banach-Alaoglu Theorem if the set

$$\Lambda(B) = \{prob(\cdot,\sigma)|\sigma \in B\},$$

for $B \subseteq \Sigma$ is closed in the weak topology of $L^2(\mathscr{H},\hat{\mu})$, then it is compact.

Trembling hand perfect equilibrium Given $\varepsilon, \nu > 0$, such that $\varepsilon > \nu$, a function $\tilde{\varepsilon}_i : \{(h_i,a_i)|h_i \in \mathscr{H}_i, a_i \in A_i(h_i)\} \to [\nu,\varepsilon]$ is called an (ε,ν)-*tremble of player* i, and $\tilde{\varepsilon} = (\tilde{\varepsilon}_i)_{i\in N}$ is called an (ε,ν)-tremble profile.

Given an ε-tremble profile $\tilde{\varepsilon}$, a strategy profile is a $\tilde{\varepsilon}$-*constrained strategy* if at each private history $h_i \in \mathscr{H}_i$ each player puts weight at least $\tilde{\varepsilon}_i(h_i,a_i)$ on action $a_i \in A_i(h_i)$. $\Sigma_i(\tilde{\varepsilon})$ denotes the set of $\tilde{\varepsilon}$-constrained strategies of player i.[27] $\Sigma(\tilde{\varepsilon})$ denotes the set of $\tilde{\varepsilon}$-constrained strategy profiles.

Definition 1. Let $\tilde{\varepsilon}$ be a (ε,ν)-tremble profile. An $\tilde{\varepsilon}$-*constrained equilibrium* $\sigma^{\tilde{\varepsilon}}$ is an $\tilde{\varepsilon}$-constrained strategy profile such that for every player $i \in N$

$$\sigma_i^{\tilde{\varepsilon}} \in \arg\max \left\{ U_i(\sigma_i', \sigma_{-i}^{\tilde{\varepsilon}}) | \sigma_i' \in \Sigma_i(\tilde{\varepsilon}) \right\}$$

We define for each t and $h \in \mathscr{H}$ and $\tilde{H} \subseteq \{1,\ldots,|h|\}$ the function

$$prob_i(h,\sigma_i,\tilde{H}) = \begin{cases} \prod_{l\in\tilde{H}} \sigma_i(\hat{\mathbf{a}}(h)_{i,l}|h_i^{(l)}(h)) & \text{if } \tilde{H} \neq \emptyset \\ 1 & \text{otherwise.} \end{cases}$$

Notice that $prob_i(h,\sigma) = prob_i(h,\sigma_i,\{1,\ldots,t\})$.

Define for $h \in \mathscr{H}$

$$\bar{H}_i(h,\bar{\sigma}) = \left\{ l \leq |h| \middle| \bar{\sigma}_i\left(\hat{\mathbf{a}}(h)_{i,\tilde{l}}|h_i(h^{(\tilde{l})})\right) > 0, \forall \tilde{l} \in \{l\ldots,|h|-1\} \right\}.$$

$\bar{H}_i(h,\bar{\sigma})$ contains the indices of moves that occur after the "last deviation" of i from $\bar{\sigma}_i$ to a zero probability action in history h.

and only if $\langle f_\alpha,\psi\rangle \to \langle f^*,\psi\rangle$ for every $\psi \in L^2(\mathscr{H},\hat{\mu})$. This topology coincides with the weak star topology due to the reflexivity of L^2.

[27] 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.

We say that a net of strategy profiles $\{\sigma^\alpha\}$ *converges to* $\bar{\sigma}$ *weakly in strategies* if for every player i if **(a)** for every time $t \in \mathbb{N}$

$$\langle prob_i(\cdot, \sigma^\alpha, \bar{H}_i(\cdot, \bar{\sigma})), \varphi \rangle \to \langle prob_i(\cdot, \bar{\sigma}, \bar{H}_i(\cdot, \bar{\sigma})), \varphi \rangle,$$

for every $\varphi \in L^2(\mathscr{H}, \hat{\mu})$, with $\operatorname{supp}\varphi \in \mathscr{H}^t$, and **(b)** $\sigma_i^\alpha(\bar{\mathbf{a}}_{i,l}|h_i(h))\mathbf{1}\left\{h|\hat{\sigma}_i\left(\hat{\mathbf{a}}_{i,l}|h_i(h)\right)=0\right\}$ converges to zero almost surely.

In words, a sequence of strategy profiles σ^m converges weakly in strategies to a limit $\bar{\sigma}$ if **(a)** the probability of i's history of play under σ^m after i's last zero probability action (according to $\bar{\sigma}$) converges weakly to that probability under $\bar{\sigma}$, and **(b)** σ^m converges to zero almost surely on the set of histories in which $\bar{\sigma}$ is equal to zero.[28]

The motivation for using this form of convergence is two-fold. First, it generates, by definition, continuity of the probability of i's histories of play with respect to strategies.[29] In contrast, if we were to take convergence of each $\sigma_i^{\tilde{\varepsilon}}(\cdot|\cdot)$ individually, this continuity would not necessarily obtain due the players' ability to "entangle" their future actions with their previous play as function of a signal that lives in the continuum (see footnotes 13 and 43). Second, the convergence is "reset" after a zero probability action so that the limit is able to recover the strategies after histories with a deviation.[30] Indeed, without this property, the information about play after a deviation, which is present in the $\tilde{\varepsilon}$-constrained strategies due to mixing, would be lost.[31]

Definition 2. A strategy profile σ^* is a *trembling hand perfect equilibrium (THPE)* if there exist a sequence $(\tilde{\varepsilon}^m)_{m\in\mathbb{N}}$ of (ε^m, ν^m)-tremble profiles, with $\varepsilon^m > \nu^m > 0$ and $\lim_{m\to\infty}\varepsilon^m = 0$, 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 Perfect Equilibrium for finite extensive form games defined by Selten (1975). An important difference is that we assume weak convergence of

[28]Notice that the convergence "drops information" regarding actions that occurred before a given period about each player i's strategies. Since every player knows the actions she has taken, this does not preclude sequential rationality with respect to beliefs that pertain to the opposing players' past actions (see section 5.3).

[29]As discussed below, this property will be essential for existence: in conjunction with Assumption 2 it implies compactness and continuity of payoffs.

[30]The strategy can be calculated as:

$$\sigma_i(\hat{\mathbf{a}}(h)_{i,|h|}|h_i(h)) = prob_i(h, \bar{\sigma}, \bar{H}_i(h, \bar{\sigma}))/prob_i(h^{(|h|-1)}, \bar{\sigma}, \bar{H}_i(h^{(|h|-1)}, \bar{\sigma})).$$

[31]Notice that the reset does not generate issues with continuity because $\sigma_i^{\tilde{\varepsilon}}$ must converge to 0 almost surely, whenever $\bar{\sigma} = 0$ and, therefore, no "entanglement" is lost from the truncation.

strategies instead of pointwise convergence of behavioral strategies. In finite games both types of convergence coincide. The advantage of our choice of topology is that, under our assumptions,[32] weak convergence of strategies is strong enough to imply convergence of expected payoffs of strategies (Lemma 2 and Proposition 2) but weak enough that the space of strategies is closed (and therefore compact). The former condition guarantees that the best response has closed graph, which together with the latter condition implies the existence of $\tilde{\varepsilon}$-constrained equilibria via Kakutani-Fan-Glicksberg.

Under our weak notion of convergence it is also not straightforward that the limit of a sequence should inherit desirable properties from its approximating elements. For example, even though it is fairly straightforward that a THPE is Nash (Lemma 4), it is not obvious that it should satisfy a notion of sequential rationality. We address this question in section 5.3 where we define a weak sequential equilibrium, as a natural extension of Kreps and Wilson (1982)'s definition for finite games to our setting, and show that a THPE is a weak sequential equilibrium under some regularity conditions. Our notion of weak convergence in strategies is chosen purposely to obtain this property.

3.1 Existence of a trembling hand perfect equilibrium

We now introduce two assumptions—1 and 2 below—that will be allow us to show existence of an equilibrium.

Define for each player $i \in N$, $g_i^{t,max,t-1} : \Omega^t \times \mathbf{S}^t \times X^{t-1} \to \mathbb{R}$ as

$$g_i^{t,max,t-1}(\omega^t, \mathbf{s}^t, \mathbf{a}^{t-1}) = \max_{\tilde{\mathbf{a}}_t \in X} |g_i(\omega^t, \tilde{\mathbf{a}}_t, \mathbf{a}^{t-1})|,$$

and recursively, for each $j \leq t-2$, define $g_i^{t,max,j} : \Omega^{j+1} \times \mathbf{S}^{j+1} \times X^j \to \mathbb{R}$ as

$$g_i^{t,max,j}(\omega^{j+1}, \mathbf{s}^{j+1}, \mathbf{a}^j) = \max_{\mathbf{a}_{j+1} \in X} \int_{\Omega \times S} g^{t,max,j+1}(\omega^{j+2}, \mathbf{s}^{j+2}, \mathbf{a}^{j+1})\, d\mu^{\omega,s}(\omega_{j+2}, \mathbf{s}_{j+2} | \omega^{j+1}, \mathbf{s}^{j+1}, \mathbf{a}^{j+1}),$$

where $\mu^{\omega,s}(\cdot | \omega^{j+1}, \mathbf{s}^{j+1}, \mathbf{a}^{j+1})$ denotes the joint measure over $\Omega \times \mathbf{S}$ induced by μ^{ω} and μ^s. As the maximum of finitely many measurable functions, each $g_i^{t,max,j}$ is measurable.

Assumption 1. *For each $i \in N$, we have*

$$\sum_{t=1}^{\infty} \int_{\Omega \times S} g^{t,max,1}(\omega_1)\, d\mu(\omega_1, \mathbf{s}_1 | \emptyset) < \infty.$$

[32] See Assumptions 1 and 2 below.

Assumption 1 is satisfied if for each $i \in N$, the following stronger condition holds

$$\sum_{t=1}^{\infty} \sup_{h \in \mathscr{H}^t} |g_i(h)| < \infty. \tag{1}$$

Notice that (1) holds in a stochastic game with discounted payoffs of the form $g_i(h) = \delta^t v_i(h)$ for $h \in \mathscr{H}^t$ where $v_i(h)$ uniformly bounded in $\mathscr{H}$ for each $i \in N$. Condition (1) also holds for payoffs that decrease slower than geometrically in t, such as a time t flow payoff given by $g_i(h) = v_i(h) \cdot \frac{1}{t^2}$ for $h \in \mathscr{H}^t$ and $i \in N$. Assumption 1, however, is weaker than condition (1). It holds, for instance, in games with stochastic move opportunities—described in section 4.2—that have bounded payoffs and a finite expected length. Condition (1), in contrast, does not hold in such settings.

The main role of Assumption 1 is to ensure that weak convergence of strategies implies convergence of expected payoffs.[33]

The following definition relates to the (un)observability of the opponents' signals.

Definition 3 (Noisy Observability)**.** We say a game of imperfect and incomplete information has *noisy observability of the opponents' signals* if for each $\mathbf{a}^{t-1} \in X^{t-1}$,

$$\mu^t(\omega^t, \mathbf{s}^t, \mathbf{a}^{t-1}) = \hat{\mu}^{\omega}(\omega^t | \mathbf{s}^t, \mathbf{a}^{t-1}) \times \hat{\mu}^s(\mathbf{s}^t | \mathbf{a}^{t-1}),$$

where $\hat{\mu}^{\omega} : \mathbf{S}^t \times X^t \times (\otimes^t \mathcal{S}_i) \to [0,1]$ is a transition probability, $\hat{\mu}^s$ is the marginal over $\mathbf{S}^t$ of $\mu^t(\omega^t, \mathbf{s}^t, \mathbf{a}^{t-1})$ (for fixed $\mathbf{a}^{t-1}$) and

$$d\hat{\mu}^s(\mathbf{s}^t | \mathbf{a}^{t-1}) = f^s(\mathbf{s}^t, \mathbf{a}^{t-1})\, d\mu_1^t(\mathbf{s}_1^t | \mathbf{a}^{t-1}) d\mu_2^t(\mathbf{s}_2^t | \mathbf{a}^{t-1}) \cdots d\mu_n^t(\mathbf{s}_n^t | \mathbf{a}^{t-1}),$$

where $\mu_j^t(s_j^t | \mathbf{a}^{t-1})$ is the marginal of μ^t on $s_j^t \in S_j^t$, for fixed action profile $\mathbf{a}^{t-1}$, and f^s is a function in $L^1(\mathbf{S}^t, \hat{\mu}^s(\cdot | \mathbf{a}^{t-1}))$.

The noisy observability of opponents' signals condition holds if the spaces of signals are countable. When the state space is uncountable the condition precludes settings in which all players observe the state, as, crucially, players are allowed to observe their opponents' observation only with noise.

Noisy observability is related to the uniform continuity condition that guarantees existence in Milgrom and Weber (1985) and Balder (1988)'s extension in a static Bayesian

[33]This result is implied by the dominated convergence theorem (see Proposition 2 and 8).

game.[34,35] In fact, in a static Bayesian game our condition holds precisely if the distribution satisfies Milgrom and Weber (1985)'s absolute continuity.[36] Thus, noisy observation generalizes absolute continuity for the context of a dynamic stochastic games, in which players can receive signals that are informative about past play, payoff types, the probabilistic evolution of the game, etc.

Assumption 2. *The game has noisy observability of the opponents' signals.*

We are now ready to state our main result.

Theorem 1. *Let Γ be a stochastic game that satisfies Assumptions 1 and 2. Γ has a trembling hand perfect equilibrium.*

When assumptions fail: Let us discuss examples of non-existence when Assumptions 1 and 2 fail.

It is easy to construct an example of non-existence when Assumption 1 fails, even in a one player setting. Suppose there is one player who must choose an action in the set $\{C,E\}$ at each t. If she chooses C (continue) n times and E (end) at time $n+1$ her payoff is $\sum_{j=0}^{n}\frac{1}{2^j}$. If she chooses C at every t her payoff is zero. This game does not have an equilibrium because for any stopping date in which E is chosen the player increases her payoff by choosing C for one period. This stochastic game does not satisfy Assumption 1 as for each t, $g^{t,max,1} = \max_{\mathbf{a}\in X^t} g_i \geq 1$ and, therefore, the sum in Assumption 1 is infinite.

Simon (2003) provides an example of non-existence that fails the noisy observability property in a one period game, with three players and finite actions. The state space is the set of binary sequences indexed by the integers. At each information set each player

[34] Their focus on distributional strategies is also similar to our use of "prob" functions as an intermediate step to obtain existence.

[35] Milgrom and Weber (1985) consider a one-shot game in which the signal is a player's type and the state is the combination of all signals. States and signals live in a metric spaces. Balder (1988) extends Milgrom and Weber (1985) existence result to general measure (not necessarily metric) spaces, like the ones we consider in our setting. Both papers' approach is—instead of viewing strategies as functions in L^p, as we do—to associate strategies to transitional probabilities and use the compactness implied by (extensions) of Prokhorov's theorem to obtain existence. It is likely that a similar approach can be used to show existence of trembling hand perfect equilibrium in our dynamic setting. However, it would also require the definition of combined transition probabilities for sequences of actions taken over time (much like our "prob" functions). It is also not clear how to obtain our results on sequential equilibrium with their alternative approach. In our view, an advantage of using weak convergence in L^p spaces is that it is a convergence of functions, which is oftentimes easier to ascertain.

[36] In a Bayesian game the state is composed by the collection of types and $\hat{\mu}^{\omega}$ is the Dirac measure.

believes that two sequences are possible, each probability 1/2. These two sequences are defined by measure preserving involutions. Players 2 and 3 observe the same signal whereas player 1 observes a different signal. The three players play a version of matching pennies in which 1 wants to mismatch 2 and 3 while players 2 and 3 want to match player 1. Player 1's payoff depends also on the state. Their example fails Assumption 2 because players 2 and 3 observe the same signal that lives in an uncountable space. Therefore, the distributions of players' signals are not absolutely continuous with respect to the product measure of marginals.

Lifting other assumptions in our definition of a stochastic game can also lead to non-existence. Harris et al. (1995) provides an example in which the space of actions is infinite that does not have a subgame perfect equilibrium. They show existence by allowing for publicly correlated strategies.[37] In a setting in which players evaluate payoffs according to the average reward, Gillette (1957) provides an example of non-existence in a game commonly referred to as the "Big Match". The form of expected payoffs in our setting does not include average reward payoffs.

4 Applications

The bulk of the literature on stochastic games considers Markov settings with perfect almost perfect information and discounted payoffs. Ours is, to our knowledge, the first proof of existence of notions of trembling hand perfect and sequential equilibrium for general stochastic games without perfect information.

We now provide two important applications of our main theorem. The first application illustrates how a small amount of noise in the observation of the state can ensure the existence of an equilibrium in a stochastic game. The second application applies our results to games with stochastic move opportunities.

4.1 Markov stochastic games with imperfect signals

Consider a stochastic game in which the distribution of the period t state $\mu^{\omega}(\omega_t|\omega^{t-1}, \mathbf{a}^{t-1})$ depends only on the previous state visited and action (i.e. we can write $\mu^{\omega}(\omega_t|\omega_{t-1}, \mathbf{a}_{t-1})$).

[37] Our methods do not accommodate infinite action spaces. The natural extension would be to allow for distributions over actions and consider the space of "$prob(\cdot, \sigma)$" functions as we have done. This space, however—even abstracting of the issue of strategic entanglement—need not be closed (i.e. the limits of these functions may not be probability distributions).

Players observe all past play but they can only observe a noisy signal of the state before they can move in each period. More specifically, suppose that the state space Ω is a Polish vector space with distance d and player i's signal at time t is given by

$$s_{i,t} = \omega_t + \varepsilon_i,$$

where ω_t is the state at time t and $\varepsilon_{i,t}$ is a "noisy" random variable with support in a set Ξ_i, satisfying Assumption 3, below. Let the joint distribution of $(\varepsilon_i)_{i\in N}$ be denoted μ^{ε} and the marginal of ε_i be denoted μ^{ε_i}. We call a stochastic game as described a *discounted Markov stochastic game with imperfect signals*.[38]

Assumption 3. *The joint measure of $(\varepsilon_i)_{i\in N}$ is (a) independent of ω, independent across periods and absolutely continuous with respect to $\mu^{\varepsilon_1} \times \ldots \times \mu^{\varepsilon_n}$ and (b) for each $j \in N$ and each measurable $B \subseteq \Xi := \times_{i\in N\setminus\{j\}}\Xi_i$ if $\mu^{\varepsilon}(B|\varepsilon_j) > 0$ there is $\eta > 0$ such that $\mu(B - \bar{z}|\varepsilon_j - z) > 0$ where $\bar{z} = (z)_{i\in N\setminus\{j\}}$ for $z \in \Omega$ such that $d(0,z) < \eta$.*

Assumption 3 is satisfied if $\Omega \subseteq \mathbb{R}^l$ for $l \geq 1$, and $d\mu^{\varepsilon}(\varepsilon_1,\ldots,\varepsilon_n) = f^{\varepsilon}(\varepsilon_1,\varepsilon_2,\ldots,\varepsilon_n)\,d\lambda(\varepsilon_1)\ldots d\lambda(\varepsilon_n)$, where λ is the Lebesgue measure and almost-surely in ε, $f(\varepsilon) > 0$ implies $f(\tilde{\varepsilon}) > 0$ for $\tilde{\varepsilon}$ in a vicinity of ε. In applied work it is common to assume imperfect information of this form.

Notice that the restrictions imposed by Assumption 3 are different than the requirements in Duggan (2012), where all players observe the state perfectly but the state has a component which is sufficiently noisy.

Proposition 1. *A stochastic game with imperfect signals that satisfies Assumption 1 has a trembling hand perfect equilibrium.*

The result follows from Theorem 1 and the fact that Assumption 3 implies Assumption 2.[39]

4.2 Games with stochastic move opportunities

There is a growing literature that considers games with stochastic move opportunities. In these games, players are drawn to move in an interval $[0,T)$ with $T \in \mathbb{R}_+ \cup \{\infty\}$. Many of

[38] This information structure allows perfect or imperfect observation of previous actions. In fact, information about previous play can, for example, be conveyed by a coordinate of the state.

[39] The proof shows that a weaker assumption (Assumption 6 which is implied by Assumption 3) that requires absolute continuity of the joint distribution of the period t realization of the state and the signal realization profiles up to period t, with respect to the product measure of the marginals of these two objects.

these games can be modeled as a stochastic game as follows.[40] In each period j a moving time $t_j \in [0,T)$, a set of players N_j, and an underlying state γ_j in a set of states Υ are drawn. These are part of the state of the stochastic game. There is an absorbing state $\gamma^{\text{end}} \in \Upsilon$ that represents the end of the game. Once the game ends players receive zero flow payoffs. The evolution of the state is as follows. Let $\gamma_0 \in \Upsilon$ denote the initial underlying state and it is drawn from a distribution $\alpha(\cdot|\emptyset) : \Upsilon \to [0,1]$. The first moving time is drawn from a density $f_1(\cdot|\gamma_0) : [0,T] \to \mathbb{R}$, with respect to the Lebesgue measure over $[0,T]$, and the set of players who move at time t_1, N_1, are drawn according to a probability distribution $\chi_1(\cdot|\gamma_0) \in \Delta 2^N$. After t_1 is drawn a new state γ_1 is drawn according to distribution $\alpha(\cdot|\gamma_0, t_1)$. Recursively, the j'th timing of play is drawn from a density $f_j(\cdot|\gamma_{j-1}, t_{j-1}) : (t_{j-1}, T] \to \mathbb{R}$, the subset of players who move, N_j, are drawn according to $\chi_j(\cdot|\gamma_{j-1}, t_{j-1})$ and the state γ_j is drawn according to $\alpha_j(\cdot|\gamma_{j-1}, t_j)$. In our notation the period 0 state is given by $\omega_0 = \gamma_0$ and the period j state is given by $\omega_j = (\gamma_j, t_j, N_j)$. Let $\mu(|\sigma)$ denote the measure over $\mathscr{H}^t$ generated by $\{f_j\}_{j\in\mathbb{N}}$, α, and a strategy profile σ.

The set of histories in which the game ends at time j is given by $\bar{H}^j = \{h \in \mathscr{H}^j | \hat{\omega}(h)_{j,1} = \gamma^{\text{end}}, \hat{\omega}(h)_{j-1,1} \neq \gamma^{\text{end}}\}$. If each g_i is bounded, the following Lemma shows that Assumption 1 holds if the expected number of moves is finite.

Lemma 1. *Assumption 1 holds if* $\sup_{h\in\mathscr{H}} |g_i(h)| < \infty$ *and* $\sup_{\sigma\in\Sigma} \sum_{j\in\mathbb{N}} j \cdot \mu(\bar{H}^j|\sigma) < \infty$.

Assumption 2 holds if, for example, only one player is drawn to move at each moving time (the game is asynchronous), each player observes her own moving time but not her opponents', and at most one player observes the realization of the state variable γ_l for each $l \in \mathbb{N}$ if Υ is uncountable. Players may observe also a countable signal that is informative about their opponents' moving times and the state. As an example, the signal space of player i can be given by $S_i = [0,T) \times P_i$, where P_i is a countable partition of $[0,T)^{n-1}$.

These examples accommodate settings in which players move "almost" simultaneously as each player could observe their own moving time but observe others' moving times up to an arbitrarily small unit of time.

From Theorems 1 and 2 any game with stochastic moves that satisfies Assumptions 1 and 2 has a trembling hand perfect equilibrium that is a weak sequential equilibrium.

Revision Games A leading example of games with stochastic move opportunities is revision games. A revision game is a game with stochastic move opportunities in which

[40] Games in continuous time are generally not representable as stochastic games. We show that games with stochastic move opportunities are an exception.

there is an underlying normal form game that is fixed throughout. Players can revise their choices in the normal form game at their move opportunities that arrive before a fixed deadline ($T < \infty$) at constant Poisson rates. State γ^{end} is reached when a time larger than the deadline is drawn.

Revision games model settings in which, even though players do not move simultaneously at prescribed fixed times as in a repeated game, they also cannot move at every time in the continuum. Some reasons for these assumptions, in applications, include the possibility that there is randomness in the players' response time or that players' face unforeseen events that can prevent them from acting at any time at will. As the rate of arrival of opportunities tends to infinity, however, a revision game resembles a continuous time game as players are able to move often. Our results show that revision games can be a suitable option to model continuous time settings, as strategies are well-defined, and Assumption 1 holds—as discussed above[41]—and, therefore, Assumption 2 guarantees that sequentially rational equilibria exist.

Kamada and Kandori (2019) were first to introduce revision games. 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 (2019) show that cooperation can arise in several applications of revision games, such as games of price competition, exchange of goods and election campaigns. Kamada and Sugaya (2020) consider a model of dynamic posturing before an election.

In concurrent work, Lovo and Tomala (2015) show existence of a Markov Perfect Equilibrium in games with stochastic timing of moves with almost perfect information. Their result holds for a finite state space Υ and Poisson distributed opportunities.

5 Existence of Equilibrium: Analysis

In this section we show existence of a trembling hand perfect equilibrium. We then show that a THPE is a Nash Equilibrium that is sequentially rational with respect to beliefs.

[41]The expected number of moves is bounded by $\max_{i \in N} \lambda_i T$ where λ_i is the rate at which player i's opportunities to move realize.

5.1 Existence of trembling hand perfect equilibrium

The bulk of the proof of existence of trembling hand perfect equilibria consists of proving the existence of $\tilde{\varepsilon}$-constrained equilibrium. The key steps are showing that under Assumption 2 the space of strategies is closed in $L^2(\mathscr{H}, \hat{\mu})$, and hence compact—by Banach-Alaoglu as it is norm-bounded—and that players' expected payoffs are continuous with respect to the topology of weak convergence of strategies.

Fix an (ε, ν)-tremble profile $\tilde{\varepsilon} = (\tilde{\varepsilon}_i)_{i \in N}$.

The following lemma establishes an important consequence of the noisy observability of opponents' signals assumption.

Lemma 2. *Suppose Assumption 2 holds. Let $\{\sigma^\alpha\} \subseteq \Sigma(\tilde{\varepsilon})$ and let $f^\alpha := prob(\cdot, \sigma^\alpha)$ be a net that converges in the weak topology of $L^2(\mathscr{H}, \hat{\mu})$ to f^*. Then there is a strategy σ^* such that $f = prob(\cdot, \sigma^*)$ and the net $\{\sigma^\alpha\}$ has a subnet that converges weakly in strategies to σ^*. Conversely, if $\{\sigma^\alpha\}$ converges weakly in strategies to σ^* then $prob(\cdot, \sigma^\alpha)$ has a subnet that converges to $prob(\cdot, \sigma^*)$ in the weak topology of $L^2(\mathscr{H}, \hat{\mu})$.*

Lemma 2 shows that weak convergence in probabilities and weak convergence in strategies are closely related. The Lemma would be immediate under almost sure convergence, but is not under weak convergence. In fact, the assumption that players observe others' information noisily plays an important role. It ensures that the candidate limit strategy that corresponds to the limit of probability functions is indeed a strategy—i.e. it does not feature correlation between unobserved actions, which would make a player's strategy not measurable with respect to her information.[42,43] Under the noisy observation assumption,

[42]This issue, sometimes called "strategic entanglement", has been discussed by other authors. See example 2 in Milgrom and Weber (1985), and discussions in Harris et al. (1995), Börgers (1991), Simon and Stinchcombe (1989) and Myerson and Reny (2019).

[43]For example, consider a game with two players and 2 actions, A and B. Both players receive a signal that exactly coincides with a state that uniformly distributed in $\Omega = [0,1]$. Consider a sequence of strategies in which both players choose A when the state is in an interval $\left[\frac{j}{m} \cdot T, \frac{j+1}{m} \cdot T\right]$ for odd j, and choose B, otherwise. In every set $C \in \mathscr{F}$

$$\int_C prob\left((A,A)|\omega\right) d\mu(\omega) = \int_C prob\left((B,B)|\omega\right) d\mu(\omega) = 1/2.$$

This shows that 1) the 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 Ω. This strategy is not measurable with respect to the players' private histories, and the expected payoff is not continuous with respect to each player's strategy (their limit strategy is to put probability 1/2 in each action, independently of the other player's choice). This example is analogous to Example 2 in Milgrom and Weber (1985), Example 2.1 in Cotter (1991) and Example 2.1 in Stinchcombe (2011).

players cannot condition their strategies on the exact signal of their opponents because the do not observe them precisely, and, therefore, as shown in Lemma 7 in the appendix, the convergence of each player's strategy inherits the desirable properties of almost sure convergence with respect to the opponents' signals.

An important consequence of Lemma 2 is that the space of probability functions associated to strategies is closed. It also follows from the Lemma that the set

$$\Sigma_{prob}(\tilde{\varepsilon}) = \{prob(\cdot,\sigma) | \sigma \in \Sigma(\tilde{\varepsilon})\}$$

is compact.

We will say that a net $\{\sigma^\alpha \subseteq \Sigma(\tilde{\varepsilon})\}$ *converges to* $\sigma \in \Sigma(\tilde{\varepsilon})$ *weakly in probabilities* if $prob(\cdot,\sigma^\alpha)$ converges to $prob(\cdot,\sigma)$ in the weak topology of $L^2(\mathscr{H},\hat{\mu})$.[44] This property defines a topology in $\Sigma(\tilde{\varepsilon})$, under which, from our previous discussion, $\Sigma(\tilde{\varepsilon})$ is closed and compact. Also the function $\sigma \to prob(\cdot,\sigma)$ from $\Sigma(\tilde{\varepsilon})$ to $L^2(\mathscr{H},\hat{\mu})$—the latter endowed with the topology of weak convergence—is a homeomorphism. In fact, it is invertible as we can back out σ from $prob(\cdot,\sigma)$ by setting, for $h \in \mathscr{H}^t$,

$$\sigma_i\left(\hat{\mathbf{a}}(h)_{i,t}|h_i(h)\right) = \frac{1}{prob(h^{(t-1)},\sigma)} \sum_{\mathbf{a}_{-i}\in X_{-i}} prob\left(\left(\hat{\mathbf{a}}(h)_{i,t},\mathbf{a}_{-i},\hat{\omega}(h)_t,\hat{\mathbf{s}}(h)_t,h^{(t-1)}\right),\sigma\right),$$

where the denominator is non-zero $\hat{\mu}$-a.s. as σ is $\tilde{\varepsilon}$-constrained.

The following is a useful consequence of Lemma 2.

Proposition 2. *Suppose assumption 1 holds. Then, if $prob(\cdot,\sigma^\alpha)$ converges in the weak topology of $L^2(\mathscr{H},\hat{\mu})$ to $prob(\cdot,\sigma^*)$, then for every player $i \in N$, $U_i(\sigma^\alpha)$ converges to $U_i(\sigma^*)$.*

Proposition 2 shows that convergence of the probability functions associated to strategies implies convergence of each player's corresponding expected utility.[45] Thus, U_i is continuous in σ under the topology of weak convergence in probabilities.

Define *player i' s best response correspondence* $r_i : \Sigma^{\tilde{\varepsilon}}_{-i} \rightrightarrows \Sigma^{\tilde{\varepsilon}}_i$ as

$$r_i(\sigma_{-i}) \in \arg\max\{U_i(\tilde{\sigma}_i,\sigma_{-i}) | \tilde{\sigma}_i \in \Sigma_i(\tilde{\varepsilon})\}.$$

[44] Notice that weak convergence in probabilities does not coincide with the notion of convergence in probability for random variables.

[45] The argument relies on Assumption 1, which guarantees that convergence of the expected payoffs at each time t implies convergence of the infinite sum by the dominated convergence theorem.

Let $r : \Sigma(\tilde{\varepsilon}) \rightrightarrows \Sigma(\tilde{\varepsilon})$ be the Cartesian product of the r_i and let $\Sigma(\tilde{\varepsilon})$ be endowed with the topology defined by weak convergence in probabilities.

Lemma 3. *The correspondence r is non-empty, convex-valued, and has closed graph.*

The proof of Lemma 3 uses Lemma 2 and Proposition 2 to establish that $U_i(\cdot, \sigma_{-i})$ is continuous and therefore attains its maximum in σ_i, and to establish that the best response correspondence, r, is closed.

Lemma 3 and the compactness of $\Sigma(\tilde{\varepsilon})$ imply directly, by the Kakutani-Fan-Glicksberg's fixed point Theorem, that an $\tilde{\varepsilon}$-constrained equilibrium exists.[46]

Finally, by Lemma 2 the set $\Sigma_{prob} = \{prob(\cdot, \sigma) | \sigma \in \Sigma\}$ is compact. Therefore, existence of $\tilde{\varepsilon}$-constrained equilibria for every (ε, ν)-tremble profiles implies existence of a THPE. This concludes the proof of Theorem 1.

5.2 Nash Equilibrium

Definition 4. A strategy profile $\sigma^* = \times_{i \in N} \sigma_i$ is a Nash Equilibrium if

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

for every player $i \in N$.

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

The proof is straightforward. There must be a sequence of (ε^m, ν^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 and Proposition 2, the players' payoffs from σ^m converge to the payoffs from σ. Thus, the fact that each σ^m does not have a profitable deviation implies that σ cannot have profitable deviations.

5.3 Weak Sequential Equilibrium

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

Throughout this section we assume that Ω and $\{S_i\}_{i \in \mathbb{N}}$ are Polish spaces and that if , $\star \in S_i$, for a player $i \in N$, then $\star$ is an isolated point of S_i.

[46]Corollary 17.55 in Aliprantis and Border (1999).

The set of histories in $\mathscr{H}^t$ that lead to player i's private history $\tilde{h}_i$ at some time $\tilde{t} \leq t$ is given by,

$$\mathscr{H}_i^t(\tilde{h}_i) := \left\{ h \in \mathscr{H}^t | h_i(h^{(\hat{t})}) = \tilde{h}_i, \text{for some } \hat{t} \leq t \right\}.$$

For a history $h \in \mathscr{H}_i^t(\tilde{h}_i)$, $\hat{t}(h, \tilde{h}_i)$ is defined as the minimum $\hat{t}$ such that $h_i(h^{(\hat{t})}) = \tilde{h}_i$.[47] Let $\mathscr{H}_-^t := \Omega^t \times \mathbf{S}^t \times X^{t-1}$ be the set of time-t histories before players choose their time-t actions and let $\mathscr{H}_- := \cup_{t \in \mathbb{N}} \mathscr{H}_-^t$. The set of histories in $\mathscr{H}_-$ that generate private history $\tilde{h}_i$ (at the time of observation) is given by

$$\bar{\mathscr{H}_i}(\tilde{h}_i) := \left\{ h_- \in \mathscr{H}_- | h_i(h_-) = \tilde{h}_i \right\},$$

where $h_i(h_-)$ is the private history of i after $h_- \in \mathscr{H}_-$ is realized. We define also, $\bar{\mathscr{H}}_i^t(\tilde{h}_i) := \bar{\mathscr{H}_i}(\tilde{h}_i) \cap \mathscr{H}_-^t$.

We define an *assessment* as a pair $(\sigma, \tilde{\mu})$ with $\sigma \in \Sigma$, and $\tilde{\mu} = \times_{i \in N} \tilde{\mu}_{-i}$ a *system of beliefs*, with $\tilde{\mu}_{-i} = \times_{t \in \mathbb{N}} \tilde{\mu}_{-i}^t$. The latter is a probability measure over histories of states, signals and opponents' actions—present and future—given each player i's private history. That is, the belief $\tilde{\mu}_{-i}^t(Z|\tilde{h}_i)$ is the probability that i assigns to the realization of histories in a set $Z \subseteq \mathscr{H}_i^t(\tilde{h}_i)$ after observing private history $\tilde{h}_i$, assuming she assigns probability 1 to her observed actions in each history $h \in Z$.[48] Thus, $\tilde{\mu}_i^t$ is required to be the counting measure with respect to player i's actions. Notice that a system of beliefs is over histories of any length, not past histories consistent with $\tilde{h}_i$—i.e. histories in $\bar{\mathscr{H}_i}(\tilde{h}_i)$—as is usually defined.

Player i's payoff from assessment $(\sigma, \tilde{\mu})$ at private history $\tilde{h}_i$ is given by

$$U_i(\sigma_i | \tilde{h}_i, \tilde{\mu}) = \sum_{l=0}^{\infty} \int_{h \in \mathscr{H}_i^l(\tilde{h}_i)} g_i(h) \cdot prob_i(h, \sigma_i, \tilde{h}_i) \, d\tilde{\mu}_{-i}^l(h | \tilde{h}_i),$$

where $prob_i(h, \sigma_i, \tilde{h}_i) = prob_i(h, \sigma_i, \mathcal{H}_i(h, \tilde{h}_i))$ and where $\mathcal{H}_i(h, \tilde{h}_i) = \{l | l \geq \hat{t}(h, \tilde{h}_i)\}$. In words, $prob_i(h, \sigma_i, \tilde{h}_i)$ is the product of the probability of i's actions in h, given strategy σ_i, after private history $\tilde{h}_i$.

A strategy σ induces a distribution over the opponents' moves in $\mathscr{H}_-$, denoted $\alpha(\cdot|\sigma)$ and is given by $d\alpha(\omega^t, \mathbf{s}^t, \mathbf{a}^{t-1} | \sigma) = prob_{-i}(\mathbf{a}^{t-1} | \omega^{t,(t-1)}, \mathbf{s}^{t,(t-1)}) d\mu_-^t(\omega^t, \mathbf{s}^t, \mathbf{a}^{t-1})$, for $(\omega^t, \mathbf{s}^t, \mathbf{a}^{t-1}) \in \mathscr{H}_-^t$, where μ_-^t is defined analogously to μ^t, except that μ^t does not

[47] Due to perfect recall $\hat{t}$ corresponds to the timing of the last signal different from $\star$ at the time of observation.

[48] By observed actions we mean actions other than $a_\star$

"count" over the time t actions.[49,50] Let $\alpha_{h_i}(\cdot|\sigma) : \mathcal{H}_i \to [0,1]$ defined as $\alpha_{h_i}(\hat{H}_i|\sigma) = \alpha\left(\{h \in \mathscr{H} : h_i(h) \in \hat{H}_i\}\,|\sigma\right)$. That is, $\alpha_{h_i}(\cdot|\sigma)$ ascribes to each set of private histories $\hat{H}_i \in \mathcal{H}_i$, the measure of the histories in $\mathscr{H}_-$ (induced by σ) that lead to private histories in $\hat{H}_i$.

We now define a notion of Bayesian updating of beliefs requiring the analogue of an application of "Bayes' rule when possible" to our infinite setting.

Definition 5. We say that an assessment $(\sigma, \tilde{\mu})$ satisfies the *generalized Bayes' rule* if for every player i and $t \in \mathbb{N}$

$$d\alpha(h|\sigma) = d\tilde{\mu}^t_{-i}(h|\tilde{h}_i, \sigma) \times d\alpha_{h_i}(\tilde{h}_i|\sigma), \tag{BR}$$

for $h \in \mathscr{H}$, $\tilde{h}_i \in \mathscr{H}_i$. If $(\sigma, \tilde{\mu})$ satisfies the generalized Bayes' rule we also say that $\tilde{\mu}$ is the *Bayes belief* of σ.[51]

Thus, a belief system of player i satisfies the generalized Bayes' rule if it corresponds to the conditional distribution over histories induced by the strategy σ, given i's private history. In a finite setting, our generalized Bayes' rule coincides with standard Bayesian updating. In a larger game, however, Bayesian updating is not always possible due to the possibility that for a positive measure of a player i's private histories, $\hat{H}_i$, for each $\tilde{h}_i \in \hat{H}_i$, $\bar{\mathscr{H}}_i(\tilde{h}_i)$ has zero measure even when strategies are totally mixed. The natural analogy to standard Bayes' rule is that players should form a conditional belief about their opponent's past play given their observations in a manner that is consistent with the objective belief over history realizations given their opponents' strategies. This is the requirement in equation (BR). By definition, player i's belief does not contain information about i's previous play. Intuitively, whenever i is called to play she recalls her previous actions, and thus her belief can simply assign probability one to them.

For every strategy profile σ there is a belief $\tilde{\mu}$ such that $(\sigma, \tilde{\mu})$ is weakly consistent. This belief is unique almost surely in h_i.[52] If σ is $\tilde{\varepsilon}$-constrained, every $\hat{H} \in \mathcal{H}$ such that

[49] Notice that by the definition of μ^t_-, $\alpha(\cdot|\sigma)$ is the counting measure with respect to i's actions.

[50] Formally, the measure $\mu^t_-(\cdot)$ on $\mathscr{H}^t_-$ is given by

$$\mu^t(Z) = \sum_{\mathbf{a}^{t-1} \in X^{t-1}} \int_{\Omega^t \times \mathbf{S}^t} \mathbf{1}\{(\omega^t, \mathbf{s}^t, \mathbf{a}^{t-1}) \in Z\}\, d\mu^{\omega,\mathbf{s}}(\omega_t, \mathbf{s}_t|\omega^{t-1}, \mathbf{s}^{t-1}, \mathbf{a}^{t-1}) \cdots d\mu^{\omega,\mathbf{s}}(\omega_1, \mathbf{s}_1|\emptyset).$$

for each measurable set Z in $\mathscr{H}^t_-$.

[51] As we explain below this belief is unique, almost surely in h_i.

[52] See theorem 5.9 in Pollard (2002).

$\mu^t(\hat{H}) > 0$, for some t, satisfies $\alpha_{h_i}(h_i(\hat{H})|\sigma) > 0$. In analogy to the finite case, due to the almost sure uniqueness, the distribution $\tilde{\mu}$ that makes $(\sigma, \tilde{\mu})$ weakly consistent is pinned down almost surely at any such set $\hat{H}$. In fact, under Assumption A5.b, $\tilde{\mu}$ can be calculated explicitly.[53]

Definition 6. We say that $(\sigma^\varepsilon, \tilde{\mu}^\varepsilon)$, where σ^ε is an $\tilde{\varepsilon}$-constrained equilibrium is a $\tilde{\varepsilon}$-*constrained sequentially rational* if $\forall i, \forall \tilde{h}_i \in \mathscr{H}_i$ and $\forall \tilde{\sigma}_i^\varepsilon \in \Sigma_i(\tilde{\varepsilon})$

$$U_i(\sigma_i^\varepsilon|\tilde{h}_i, \tilde{\mu}^\varepsilon) \geq U_i(\tilde{\sigma}_i^\varepsilon|\tilde{h}_i, \tilde{\mu}^\varepsilon).$$

We say that $(\sigma, \tilde{\mu})$, with $\sigma \in \Sigma$, is *sequentially rational* if $\forall i, \forall \tilde{h}_i \in \mathscr{H}_i$ and $\forall \sigma_i' \in \Sigma_i$

$$U_i(\sigma_i|\tilde{h}_i, \tilde{\mu}) \geq U_i(\sigma_i'|\tilde{h}_i, \tilde{\mu}).$$

The following Proposition shows that an $\tilde{\varepsilon}$-constrained equilibrium is part of a constrained sequentially rational assessment.

Proposition 3. *Let $\tilde{\varepsilon}$ be an (ε, ν)-tremble, for $\varepsilon, \nu > 0$, let σ^ε be a $\tilde{\varepsilon}$-constrained equilibrium, and let $\tilde{\mu}^\varepsilon$ be the belief such that $(\sigma^\varepsilon, \tilde{\mu}^\varepsilon)$ satisfies the generalized Bayes' rule. The assessment $(\sigma^\varepsilon, \tilde{\mu}^\varepsilon)$ is constrained sequentially rational.*

From Proposition 3 a THPE is arbitrarily close to a strategy that is sequentially rational over a set strategies that is arbitrarily close to Σ_i. In many cases—including in every countable game—this property will imply that a THPE is sequentially rational with respect to the limit belief. However, in the general case, a difficulty arises due to the discontinuity of players' payoffs with respect to beliefs over the opponents' actions. In fact, each player i's strategy may become strategically entangled with her belief in the limit, thus, yielding the discontinuity. Since the limit of beliefs may yield a distribution that cannot be represented by an L^1 function, assumption 2 is no longer sufficient to ensure the "detangling" of the limits. In what follows we define convergence of beliefs and introduce a continuity and regularity "by parts" requirement on the densities of players' signals that ensures the desired continuity of payoffs.

We say that a system of beliefs $\tilde{\mu}^n$ converges to μ^* in the weak$-*$ topology of $(L^\infty)^*$ if for every $t \in \mathbb{N}$, each player i, $\tilde{\mu}_{-i}^{n,t}(\cdot|\tilde{h}_i)$ converges to $\mu_{-i}^{*,t}(\cdot|h_i)$ in the weak-* topology

[53] See Lemma 12.

of $(L^\infty(\mathscr{H}_i^t(\tilde{h}_i),\mu^t))^*$ for every $\tilde{h}_i$ in a full measure subset of $\mathscr{H}_i$.[54] $\mu^*(\cdot|\tilde{h}_i)$ is a finitely additive measure for each $\tilde{h}_i \in \mathscr{H}_i$.[55]

Definition 7. 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 of $(L^\infty)^*$.

Definition 8. An assessment $(\sigma,\tilde{\mu})$ is a *weak sequential equilibrium* if it is quasi-consistent and sequentially rational.

Our definition of sequential equilibrium is closely related to Kreps and Wilson (1982)'s. One important difference is that we define beliefs to be over *past and future play of the opponents* at the time of observation. Kreps and Wilson (1982) in contrast, define the belief over histories of play up to the present. In a finite game as in Kreps and Wilson (1982), this distinction is irrelevant, as defining the belief in either way yields the same set of equilibria. In fact, the two notions coincide in countable games. Because each player's own strategy can be "entangled" with their future play in the limit, however, our definition is needed in order to find existence outside of countable games. Finally, because the topology that is used for convergence matters for issues of compactness and continuity, we had to require "weaker" types of convergence than Kreps and Wilson (1982), who use almost-sure convergence.

Let $\tilde{g}_i^t := \sup_{h\in\mathscr{H}^t}|g_i(h)|$ and define

$$\bar{g}_i^t(h) = \begin{cases} \tilde{g}_i^t & \text{if } |g_i(h)|>0 \\ 0 & \text{otherwise.} \end{cases}$$

Assumption 4. *For each $i \in N$, $\sum_{l=1}^{\infty}\max_{a^l\in X^l}\int_{\Omega^l\times\mathbf{S}^l}\bar{g}_i^l(\omega^l,\mathbf{a}^l,\mathbf{s}^l)\,d\mu^l(\omega^l,\mathbf{a}^l,\mathbf{s}^l)<\infty$.*

Assumption 4 implies Assumption 1.

We now introduce a technical assumption that requires continuity and boundedness of the probability density introduced in Assumption 2. Define $T_i(s_i^t|\mathbf{a}^{t-1}) := h_i(s_i^t,\mathbf{a}_i^{t-1})$, where $h_i(s_i^t,\mathbf{a}_i^{t-1})$ is the private history of i after signal history s_i^t and i's action history ,$\mathbf{a}_i^{t-1}$. Let $\gamma_{h_i}(\cdot|\mathbf{a}^{t-1}):\mathscr{H}_i\to[0,1]$ be the measure defined by $\gamma_{h_i}(\hat{H}_i|\mathbf{a}^{t-1}) =$

[54] $(L^\infty(\mathscr{H}_i^t(\tilde{h}_i),\mu^t))^*$ denotes the dual space of $L^\infty(\mathscr{H}_i^t(\tilde{h}_i),\mu^t))$. The unit ball in $(L^\infty)^*$, characterized by the norm $\|T\|_{(L^\infty)^*}=\sup\{T(x)|x\in L^\infty,\|x\|_\infty\le 1\}$, is compact in the weak-$*$ topology.

[55] See, for example, Dunford and Schwartz (1957) for the characterization of $(L^\infty)^*$. It is implicit in this definition of convergence that $\tilde{\mu}^n(\cdot|\tilde{h}_i)$ belongs to $(L^\infty(\mathscr{H}_i(\tilde{h}_i),\hat{\mu}))^*$ for each $\tilde{h}_i$.

$\int_{S_i^t} \mathbf{1}\{T_i^{-1}(\hat{H}_i|\mathbf{a}^{t-1})\} d\mu_i^t(s_i^t|\mathbf{a}^{t-1})$. That is to say, γ_{h_i} measures the set of i's signals that induce a set of i's private histories for a fixed sequence of action profiles. For $t, \tilde{t} \in \mathbb{N}$ and $\mathbf{a}^{t-1} \in X^t, s_i^t \in S_i^t$ define $\mathscr{A}_{\tilde{t}}(\mathbf{a}^{t-1}) = \{\tilde{\mathbf{a}}^{\tilde{t}-1} \in X^{\tilde{t}-1} | T_i^{-1}(h_i(s_i^t, \mathbf{a}^{t-1})|\mathbf{a}^{\tilde{t}-1}) \neq \emptyset\}$.[56] $\mathscr{A}_{\tilde{t}}(\mathbf{a}^{t-1})$ is the set of action profiles that are possible to have realized given an observation of player i induced by a sequence of signals and actions. $\mathscr{A}_{\tilde{t}}(\mathbf{a}^{t-1})$ may not be empty for $\tilde{t} \neq t$ in a game with past unobserved moves (that is a game in which signal $\star$ may be drawn). Suppose that Assumption 2 holds, we define for $\tilde{t}, t, \hat{t} \in \mathbb{N}$ with $\hat{t} \leq t$, $s_i^t \in S_i^t$, $\mathbf{s}_{-i}^{\tilde{t}} \in \mathbf{S}_i^{\tilde{t}}$, and $\mathbf{a}^{t-1} \in X^{t-1}$ the function

$$\bar{\ell}^s_{\tilde{t},\hat{t}}(s_i^t, \mathbf{a}^{t-1}, \mathbf{s}_{-i}^{\tilde{t}}, \tilde{\mathbf{a}}^{\tilde{t}-1}) = \log f^s(\mathbf{s}_{-i}^{\tilde{t}}, T_i^{-1}(h_i(s_i^{t,(\hat{t})}, \mathbf{a}^{t,(\hat{t}-1)})|\tilde{\mathbf{a}}^{\tilde{t}-1}), \tilde{\mathbf{a}}^{\tilde{t}-1}). \tag{2}$$

Assumption 5. *Suppose Assumption 2 holds and for each i and for each t there is a countable collection of sets $\mathcal{P}_{t,i} = \{P_{t,i}^k\}_{k \in \mathbb{N}}$ with $P_{t,i}^k \subseteq S_i^t$ for each $k \in \mathbb{N}$ and such that $\mu_i^t(S_i^t|\mathbf{a}^{t-1}) = \mu_i^t(\cup_{k \in \mathbb{N}} P_{t,i}^k|\mathbf{a}^{t-1})$ and,*

A5.a *$\forall \nu > 0, k \in \mathbb{N}, s_i^t \in P_{t,i}^k$ there is $\delta > 0$ such that if $\hat{s}_i^t \in B(s_i^t, \delta) \cap P_{t,i}^k$ then*

$$|\ell^s_{\tilde{t},\hat{t}}(s_i^t, \mathbf{a}^t, \mathbf{s}_{-i}^{\tilde{t}}, \tilde{\mathbf{a}}^{\tilde{t}-1}) - \ell^s_{\tilde{t},\hat{t}}(\hat{s}_i^t, \mathbf{a}^t, \mathbf{s}_{-i}^{\tilde{t}}, \tilde{\mathbf{a}}^{\tilde{t}-1})| < \nu$$

$\forall \hat{t} \leq t, \tilde{t} \in \mathbb{N}, \tilde{\mathbf{a}}^{\tilde{t}-1} \in \mathscr{A}_{\tilde{t}}(\mathbf{a}^t), \mathbf{s}_{-i}^{\tilde{t}} \in \mathbf{S}_{-i}^{\tilde{t}}$.[57]

A5.b *There is a measure $\beta_i : \mathcal{H}_i \to [0,1]$ such that $\gamma_{h_i}(\tilde{h}_i|\mathbf{a}^{t-1})$ has Radon-Nikodym derivative $b_i(\tilde{h}_i, \mathbf{a}^{t-1})$ with respect to β_i, and for every $k \in \mathbb{N}$, $b_i(h_i(s_i^t, \mathbf{a}_i^{t-1}), \mathbf{a}^{t-1})$ is continuous in $s_i^t \in P_{t,i}^k$, for all $\mathbf{a}^{t-1} \in X^{t-1}$.*[58]

A5.c *For every $k \in \mathbb{N}$, $\hat{t} \in \mathbb{N}$ and $\mathbf{a}^{t-1} \in X^{t-1}$ there is $R^k(\cdot, \mathbf{a}^{t-1}, \hat{t}) \in L^1(P_{t,i}^k, \mu_i^t(\cdot|\mathbf{a}^{t-1}))$ such that*

$$\left| \frac{f^s(\hat{s}_i^t, \mathbf{s}_{-i}^t, \mathbf{a}^{t-1})}{f^s\left(\hat{s}_i^{t,(\hat{t})}, \mathbf{s}_{-i}^{t,(\hat{t})}, \mathbf{a}^{t-1,(\hat{t}-1)}\right) \hat{b}_{i,\hat{t}}(\hat{s}_i^t, \mathbf{a}^{t-1})} \right| \leq R^k(\hat{s}_i^t, \mathbf{a}^{t-1}, \hat{t}), \tag{3}$$

for every $\hat{s}_i^t \in P_{t,i}^k \cap \check{S}_i^t(\mathbf{a}^{t-1})$ and $\mathbf{s}_{-i}^t \in \mathbf{S}_{-i}^t$, where $\hat{b}_{i,\hat{t}}(s_i^t, \mathbf{a}^{t-1}) :=$

---

[56] Notice that $\mathscr{A}_{\tilde{t}}(\mathbf{a}^{t-1})$ does not depend on the choice of $s_i^t \in S_i^t$.

[57] We use the convention that $\ell^s_{\tilde{t},\hat{t}}(s_i^t, \mathbf{a}^t, s_{-i}^{\tilde{t}}, \tilde{\mathbf{a}}^{\tilde{t}-1}) - \ell^s_{\tilde{t},\hat{t}}(\hat{s}_i^t, \mathbf{a}^t, s_{-i}^{\tilde{t}}, \tilde{\mathbf{a}}^{\tilde{t}-1}) = 0$ if $\ell^s_{\tilde{t},\hat{t}}(s_i^t, \mathbf{a}^t, s_{-i}^{\tilde{t}}, \tilde{\mathbf{a}}^{\tilde{t}-1}), \ell^s_{\tilde{t},\hat{t}}(\hat{s}_i^t, \mathbf{a}^t, s_{-i}^{\tilde{t}}, \tilde{\mathbf{a}}^{\tilde{t}-1}) = -\infty$.

[58] Due to the Radon-Nikodym Theorem such β_i exists. The restrictiveness of the assumption lies in the continuity over s_i^t in each cell element $P_{t,i}^k$.

$b_i\left(h_i\left(s_i^{t,(\hat{t})}, \mathbf{a}^{t-1,(\hat{t}-1)}\right), \mathbf{a}^{t-1,(\hat{t}-1)}\right)$ *and* $\check{S}_i^t(\mathbf{a}^t) := \left\{s_i^t \in S_i^t | \hat{b}_{i,\hat{t}}(s_i^t, \mathbf{a}^t) \neq 0\right\}$.

A stochastic move opportunity game with partitional observation of the opponents' moving times and Poisson arrivals, such as a revision game, satisfies Assumption 5. The assumption requires that in each cell of an almost-sure countable partition of each player's signal space, a) the log of the density of the players' signals satisfies a continuity requirement, b) the density over private histories of each player's individual signal is continuous with respect to the signals that generate the private histories, and c) there is "sufficient boundedness" (above and below) of the previously mentioned densities.

Theorem 2. *Let Γ be a stochastic game that satisfies Assumptions 4 and 5. Then every trembling hand perfect equilibrium of Γ is a weak sequential equilibrium.*

6 Approximating Games

In this section we study conditions under which stochastic games admit approximating games whose equilibria approach equilibria of the original game.

Let us assume that Ω and S_i for $i \in N$ are Polish spaces.

Define the correspondence $\mathbf{A} : \cup_{t\in\mathbb{N}}\Omega^t \rightarrow \cup_{t\in\mathbb{N}}X^t$ as $\mathbf{A}(\omega^t) = \left\{\mathbf{a}^t \in X^t | (\mathbf{a}^t)_{i,\tau} \in A_i\left(\omega^{t,(\tau)}, \mathbf{a}^{t,(\tau-1)}\right), \text{ for } \tau \leq t\right\}$, if $\omega^t \in \Omega^t$. $\mathbf{A}$ is a correspondence that yields for every length-t vector of realized states, the set of length-t action profiles that are feasible given each player i's feasible actions correspondence, A_i.

Definition 9. We say that a stochastic game Γ *has an approximating game sequence* if there are sequences of countable partitions of the state space Ω, $\{\mathscr{P}_n^{\Omega}\}_{n\in\mathbb{N}}$, and of each space of signals S_i, $\{\mathscr{P}_n^{S_i}\}_{n\in\mathbb{N}}$, such that

(a) $\mathbf{A}(\omega^t) = \mathbf{A}(\tilde{\omega}^t)$ for each $\tilde{\omega}^t \in P_n^{\Omega,t}(\omega^t)$, where $P_n^{\mathscr{C},t}(\omega^t)$, for $\mathscr{C} \in \{\Omega, \{S_i\}_{i\in N}\}$, denotes the element of the partition $\mathscr{P}_n^{\mathscr{C},t} := \times_{l=1}^t \mathscr{P}_n^{\mathscr{C}}$ to which ω^t belongs to.

(b) For every $\varepsilon > 0$, and $\omega^t \in \Omega^t$, $\mathbf{a}^t \in X^t$ there is n such that

$$|g_i(\omega^t, \mathbf{a}^t) - g_i(\tilde{\omega}^t, \mathbf{a}^t)| < \varepsilon,$$

for every $\tilde{\omega}^t \in P_n^{\Omega,t}(\omega^t)$, $i \in N$.

(c) For every $\mathbf{a}^t \in X^t$, $B_1 \subseteq \Omega^t$, $B_2 \subseteq \mathbf{S}^t$, open sets, the sets $B_{1,n} = \{\omega^t | \omega^t \in P_1 \in \mathscr{P}_n^{\Omega,t}, P_1 \subseteq B_1\}$ and $B_{2,n} = \{\mathbf{s}^t | \mathbf{s}^t \in P_2 \in \mathscr{P}_n^{\mathbf{S},t}, P_2 \subseteq B_2\}$, with $\mathscr{P}_n^{\mathbf{S},t} := \times_{l=1}^t \times_{i\in N} \mathscr{P}_n^{S_i}$, satisfy[59]

$$\lim_{n\to\infty} \mu^{\omega,s}(B_{1,n}, B_{2,n}|\mathbf{a}^t) = \mu^{\omega,s}(B_1, B_2|\mathbf{a}^t)$$

In words, a game has an approximating game sequence if there are sequences of uncountable partitions of the state and signal spaces such that (a) along the sequence the available actions are the same at every element of the partition, (b) each player's flow payoff gets closer and closer within each partition cell as the sequence evolves and (c) for open sets $B_1 \subseteq \Omega^t$ and $B_2 \subseteq \mathbf{S}^t$, the measure of the partition cells contained in B_1 and B_2 converges to the measure of B_1 and B_2, respectively.

We say that the stochastic game *has a finite approximating game sequence if* $\mathscr{P}_n^{\Omega}$ and $\mathscr{P}_n^{S_i}$ are finite for each $i \in N$ and $n \in \mathbb{N}$.

Consider a game Γ that has an approximating game sequence. We define a *n'th approximating game of* Γ, as a stochastic game $\Gamma_n = (N, (\Omega_n, \mathscr{F}_n), (S_{i,n}, \mathbb{S}_i), (X_i, A_i, g_i), (\mu_n^{\omega}, \mu_n^{s}))$, such that the state space Ω_n and signal spaces, $S_{i,n}$, are of the form $\Omega_n = \{\omega_p\}_{p\in\mathscr{P}_n^{\Omega}}$ and $S_{i,n} = \{s_{i,p_i}\}_{p_i\in\mathscr{P}_n^{S_i}}$, with $\omega_p \in p$ and $s_{i,p_i} \in p_i$ for each $p \in \mathscr{P}_n^{\Omega}$ and $p_i \in \mathscr{P}_n^{S_i}$. The measure over Ω_n^t and $\mathbf{S}_n^t$, induced by μ_n^{ω} and μ_n^{s}, denoted $\mu_n^{\omega,s}$, satisfies $\mu_n^{\omega,s}(\tilde{\omega}^t, \tilde{\mathbf{s}}^t|\mathbf{a}^t) = \mu^{\omega,s}(P_n^{\Omega,t}(\tilde{\omega}^t), P_n^{\mathbf{S},t}(\tilde{\mathbf{s}}^t)|\mathbf{a}^t)$ for each $t \in \mathbb{N}$.

Given a strategy in the approximating game Γ_n, σ_n, we can define *an associated strategy in* Γ, $\hat{\sigma}_{i,n}(a_i|h_i(\omega^t, \mathbf{s}^t, \mathbf{a}^t)) = \sigma_{i,n}\left(a_i|h_i(\omega_n(\omega^t), \mathbf{s}_n(\mathbf{s}^t), \mathbf{a}^t)\right)$, where $\omega_n(\omega^t) \in P_n^{\Omega,t}(\omega^t) \cap \Omega_n^t$ and $\mathbf{s}_n(\mathbf{s}^t) \in P_n^{s,t}(\mathbf{s}^t) \cap \mathbf{S}_n^t$.

In the following Proposition convergence is in the weak topology of $L^2(\mathscr{H}, \hat{\mu})$

Proposition 4. *Let Γ be a stochastic game that satisfies Assumption 1 and let Γ_n be a sequence of approximating games. For each n, let σ_n be a $\tilde{\varepsilon}$-constrained equilibrium of Γ_n for some (ε, ν)-tremble, $\tilde{\varepsilon}$, and let $\hat{\sigma}_n$ be its associated strategy in Γ. If there is a strategy σ^* in Γ such that $prob(\cdot, \hat{\sigma}^n) \to prob(\cdot, \sigma^*)$, and for each $i \in N$, $prob_i(\cdot, \hat{\sigma}_i^n) \to prob_i(\cdot, \sigma_i^*)$ and $prob_{-i}(\cdot, \hat{\sigma}_{-i}^n) \to prob_{-i}(\cdot, \sigma_{-i}^*)$ then σ^* is a $\tilde{\varepsilon}$-constrained equilibrium of Γ.*

Proposition 4 can be useful in settings where it is easier to derive properties of equilibria of an approximating game that—due to the convergence in the Proposition—translate to properties of the equilibria of the original stochastic game. It may also be a step to show existence. For example, if a game has an approximating sequence our results imply that

[59] $\mu^{\omega,\mathbf{s}}(\cdot|\omega^l, \mathbf{a}^l)$ is the joint measure over $\Omega \times \mathbf{S}$ induced by $\mu^{\omega}(\cdot|\omega^l, \mathbf{a}^l)$ and $\mu^{s}(\cdot|\omega^{l+1}, \mathbf{a}^l)$.

equilibria exist along the sequence under Assumption 1. Thus, convergence to a strategy as required by the Proposition, implies existence of an $\tilde{\varepsilon}$-equilibrium in the orginal stochastic game. By Lemma 7 in the appendix the noisy observability condition implies that the equilibria of an approximating sequence have a convergent subsequence that converges to a strategy as required by Proposition 4.

We say that a game is *Markovian in payoffs and actions* if for each player i the payoff function g_i depends only on the last state visited and action profile, and the available action correspondence A_i depends only on the present state.[60] That is, for $\omega^t \in \Omega^t$ and $\mathbf{a}^t \in X^t$, we can write $g_i(\omega^t, \mathbf{a}^t) \equiv \tilde{g}_i((\omega^t)_t, (\mathbf{a}^t)_t)$ and $A_i(\omega^t, \mathbf{a}^{t-1}) \equiv \tilde{A}_i((\omega^t)_t, (\mathbf{a}^{t-1})_t)$ for some $\tilde{g}_i : \Omega \times X \to \mathbb{R}$ and $\tilde{A}_i : \Omega \to X_i$.[61]

The following result establishes that if a game is Markovian in payoffs and actions, and the payoffs and the measure of states and signals are well behaved, then the stochastic game has a finite approximating game sequence.

Proposition 5. *Let* $\Gamma = (N, (\Omega, \mathscr{F}), (S_i, \mathcal{S}_i), (X_i, A_i, g_i), (\mu^\omega, \mu^s))$ *be a stochastic game Markovian in payoffs and actions. If* Γ *satisfies Assumption 1,* Ω *is compact and, for each* $i \in N$*,* S_i *is compact,* A_i *is upper-hemicontinuous and* g_i *is continuous, then* Γ *has a finite approximating game sequence.*

For the proof we exploit the fact that by compactness and continuity, the state and signal spaces can be covered by finitely many sets chosen so that within each set the players' flow payoffs are at most at any given (small) distance apart. These sets define the approximating partitions as we take the distance between payoffs within each of these sets to zero.

Propositions 4 and 5 suggest an alternative approach to show existence of a trembling hand perfect equilibrium in stochastic games under Assumption 2.[62] The alternative approach, however, requires stronger continuity and metrizability conditions than we need to obtain our main result.

A stochastic game is said to be *Markov* if it is Markovian in payoffs and actions and the measure over states in each period depend only on the previous period's state and actions. Strategies are Markov if they only condition on the period's state. A *Markov game has*

[60] Traditionally, in a Markovian game the distribution of states only depends on the last state that is drawn (and there are no signals due to an almost perfect information assumption). We do not need these conditions for Proposition 5.

[61] Notice that one can always define a non-Markovian game to be Markovian by enlarging the state space to comprise all sequences of states visited. The Markovian assumption does have bite, however, in our next result because it requires a compact state space.

[62] Assumption 2 implies the convergence requirements of Proposition 4.

discounting if for each player i, $g_i(\omega^t, \mathbf{a}^t) = \delta^t v_i((\omega^t)_t, (\mathbf{a}^t)_t)$ for some bounded function $v_i : \Omega \times X \to \mathbb{R}$.

Corollary 1. *Let Γ be a asynchronous Markov stochastic game of almost perfect information that satisfies the assumptions in Proposition 5. If either Γ has discounting or has a finite expected length then it has a trembling hand perfect equilibrium in Markov strategies.*

The result follows from the fact that finite Markov stochastic games have Markov Perfect equilibria.[63] By Proposition 5 the game has a finite approximating game sequence. By Proposition 4 the $\tilde{\varepsilon}$-constrained Markov equilibria of the approximating game converge to equilibria of the original game. The convergence required in Proposition 4 follows from the asynchronicity of the players' strategies and the independence of player's strategies on previous realized states along the approximating sequence.

Corollary 1 implies, in particular, that asynchronous revision games have a Markov Perfect Equilibrium.

A Appendix: Proofs

PROOF OF PROPOSITION 1

We will show that the joint measure of the players' signals up to period t is absolutely continuous with respect to the product measure of player's marginals over their signals, thus satisfying assumption 2.

We will show that the following weaker assumption, is sufficient for Assumption 2 and we will then show that Assumption 3 implies Assumption 6,

Assumption 6. *(a) Assumption 3 (a) holds and*

(b) for each $j \in N$ and $\mathbf{a}^{t-1} \in X^{t-1}$, the joint measure of the period t realization of the state ω_t, and the sequence of profiles of signal realizations up to time t, $\mathbf{s}^t$, conditional of j's signal realizations $\mathbf{s}^t_j$, $\mu^{\Omega,\mathbf{S}^t}(\cdot|\mathbf{s}^t_j, \mathbf{a}^{t-1})$, is absolutely continuous with

[63] By Sobel (1971) a finite Markov game of almost perfect information with discounting has a Markov Perfect equilibrium. A simple extension of Sobel (1971) shows that a finite approximating game has an $\tilde{\varepsilon}$-constrained Markov equilibrium in a game that ends in finite time with probability 1. In fact, the expected payoff (v^i_δ in his notation) is continuous in $\tilde{\varepsilon}$-constrained strategies. To see this, notice that the end of the game can be modelled as 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.

respect to the product measure of the corresponding marginals, $\mu^{\Omega}(\cdot|\mathbf{s}^t_j, \mathbf{a}^{t-1})$ *and* $\mu^{\mathbf{S}^t}(\cdot|\mathbf{s}^t_j, \mathbf{a}^{t-1})$.

In the rest of the proof of Proposition 1, we omit the dependence on $\mathbf{a}^{t-1}$ for ease of notation.

Lemma 5. *Assumption 6 implies Assumption 2 in a Markov stochastic game with imperfect signals.*

Proof. Let μ^P denote the product measure of the marginals and let μ^J be the joint measure. Suppose set B is such that $\mu^J(B) > 0$. Let us show that under Assumption 6, $\mu^P(B) > 0$.

As B is a measurable set $B \subseteq \mathbf{S}^t$ it can be written as

$$B = \{(\mathbf{s}_1, \mathbf{s}_2, \ldots, \mathbf{s}_t) \in \mathbf{S}^t | \mathbf{s}_1 \in \hat{\mathbf{S}}^1, \mathbf{s}_2 \in \hat{\mathbf{S}}^2(\mathbf{s}_1), \ldots, \mathbf{s}_t \in \hat{\mathbf{S}}^t(\mathbf{s}_1, \ldots, \mathbf{s}_{t-1})\},$$

where for each $\tau \leq t$, $(\mathbf{s}_1, \ldots, \mathbf{s}_{\tau-1}) \in \mathbf{S}^\tau$, $\hat{\mathbf{S}}^\tau(\mathbf{s}_1, \ldots, \mathbf{s}_{\tau-1})$ is a measurable set in $\mathbf{S}$.[64] Similarly, $\hat{\mathbf{S}}^\tau$ ca be written as

$$\hat{\mathbf{S}}^\tau(\mathbf{s}_1, \ldots, \mathbf{s}_{\tau-1}) = \{(s_1, \ldots, s_n) \in \mathbf{S} | s_i \in \tilde{S}^\tau_i(\mathbf{s}_1, \ldots, \mathbf{s}_{\tau-1})(s_1, \ldots, s_{i-1}), i \in 1, \ldots, n\}.$$

Slightly abusing language, we will say that a set $\hat{B}$ is contained in $\times_{j=1}^k S_j$ if $\hat{B} = \hat{\hat{B}} \times (\times_{j=k+1}^n S_j)$ for some $\hat{\hat{B}} \subseteq \times_{j=1}^k S_j$. We argue by induction on $t \in \mathbb{N}$ and the largest k such that $\hat{\mathbf{S}}^t(\mathbf{s}_1, \cdots, \mathbf{s}_{t-1})$ is contained in $\times_{j=1}^k S_j$ for every $(\mathbf{s}_1, \cdots, \mathbf{s}_{t-1}) \in \mathbf{S}^{t-1}$.

Suppose $t = 1$ and $k = 1$ and let $\widetilde{B}$ be such that $B = \widetilde{B} \times (\times_{j=1}^n S_j)$. By the definition of the marginal distribution $\mu^J(B) = \int_{\widetilde{B}} d\mu_1^1(\mathbf{s}_1) = \mu^P(B)$.

Suppose that the result holds true for time t and $k-1$ or time $t-1$ and $k = n$. We will show that $\mu^P(B) > 0$ for t and k, and t and $k = 1$, respectively.

Let $B^{t-1} = \{(\mathbf{s}_1, \mathbf{s}_2, \ldots, \mathbf{s}_{t-1}) \in \mathbf{S}^{t-1} | \mathbf{s}_1 \in \hat{\mathbf{S}}^1, \mathbf{s}_2 \in \hat{\mathbf{S}}^2(\mathbf{s}_1), \ldots, \mathbf{s}_t \in \hat{\mathbf{S}}^{t-1}(\mathbf{s}_1, \ldots, \mathbf{s}_{t-2})\}$ and let $B^{t-1,k-1} = \{(s_1, s_2, \ldots, s_{k-1}) \in S_1 \times S_2 \times \ldots S_{k-1} | s_i \in \tilde{S}^t_i(\mathbf{s}_1, \ldots, \mathbf{s}_{t-1})(s_1, \ldots, s_{i-1}), i \in 1, \ldots, k-1\}$ and $B^{t-1}_k(\mathbf{s}_1, \ldots, \mathbf{s}_{t-1})(s_1, \ldots, s_{k-1}) = \tilde{S}^t_k(\mathbf{s}_1, \ldots, \mathbf{s}_{t-1})(s_1, \ldots, s_{k-1})$. In what follows $\mu(\cdot)$ denotes the joint measure, according to the primitives, of the evaluated variables and $\mu(\cdot|\cdot)$ denotes a conditional measure.

[64] See Lemma 4.46 in Aliprantis and Border (1999), as a Polish Space is second countable.

$\mu^J(B) > 0$ implies

$$0 < \int_{\mathbf{s}^{t-1}\in B^{t-1}} \int_{s^{k-1}\in B^{t-1,k-1}(\mathbf{s}^{t-1})} \int_{s_k\in B_k(s^{k-1})} d\mu(s_k|s^{k-1},\mathbf{s}^{t-1})\, d\mu(s^{k-1}|\mathbf{s}^{t-1}) d\mu(\mathbf{s}^{t-1}) =$$
$$\int_{\mathbf{s}^{t-1}\in B^{t-1}} \int_{s^{k-1}\in B^{t-1,k-1}(\mathbf{s}^{t-1})} \int_{s_k\in B_k(s^{k-1})} d\mu(s_k|s^{k-1},\mathbf{s}^{t-1})\, \hat{f}(s^{k-1},\mathbf{s}^{t-1}) d\mu_1^1(s_1^{k-1}|\mathbf{s}_1^{t-1})\ldots\mu_{k-1}^1(s_{k-1}^{k-1}|\mathbf{s}_{k-1}^{t-1}),$$

for some measurable function $\hat{f}:\mathbf{S}^{t-1}\times(\times_{j\leq k-1}S_j)\to\mathbb{R}$, by the Radon Nikodym Theorem and the induction hypothesis. Now,

$$\int_{s^{k-1}\in B_k(s^{k-1})} d\mu(s_k|s^{k-1},\mathbf{s}^{t-1}) =$$
$$\int_{\omega_t\in\Omega}\int_{\varepsilon_k\in B_k(s^{k-1})-\omega_t} f^{\varepsilon}(s_1^{k-1}-\omega_t,\ldots,s_{k-1}^{k-1}-\omega_t,\varepsilon_k)\, d\mu^{\varepsilon_k}(\varepsilon_k)\, d\mu(\omega_t|s^{k-1},\mathbf{s}^{t-1}) =$$
$$\int_{\omega_t\in\Omega}\int_{\varepsilon_k\in B_k(s^{k-1})-\omega_t} f^{\varepsilon}(s_1^{k-1}-\omega_t,\ldots,s_{k-1}^{k-1}-\omega_t,\varepsilon_k)\check{f}(\omega_t,s^{k-1},\mathbf{s}^{t-1})\, d\mu^{\varepsilon_k}(\varepsilon_k)\, d\mu^{\omega}(\omega_t|\mathbf{s}_k^{t-1}),$$

where the second equality follows by Assumption 6, which implies $d\mu(\omega_t|s^{k-1},\mathbf{s}^{t-1}) = \check{f}(\omega_t,s^{k-1},\mathbf{s}^{t-1})d\mu^{\omega}(\omega_t|\mathbf{s}_k^{t-1})$, for some measurable function $\check{f}:\Omega\times S_1\times\ldots S_{k-1}\times\mathbf{S}^{t-1}\to\mathbb{R}$. Now, since

$$\int_{\omega_t\in\Omega}\int_{\varepsilon_k\in B_k(s^{k-1})-\omega_t} d\mu^{\varepsilon_k}(\varepsilon_k)\, d\mu^{\omega}(\omega_t|\mathbf{s}_k^{t-1}) = \mu_k^1(B_k(s^{k-1})|\mathbf{s}_k^{t-1}),$$

it follows that

$$\mu^P(B) = \int_{\mathbf{s}^{t-1}\in B^{t-1}} \int_{s^{k-1}\in B^{t-1,k-1}(\mathbf{s}^{t-1})} \int_{s_k\in B_k(s^{k-1})} d\mu_k^1(s_k|\mathbf{s}_k^{t-1}) d\mu_{k-1}^1(s_{k-1}^{k-1}|\mathbf{s}_{k-1}^{t-1})\ldots\mu_1^1(s_1^{k-1}|\mathbf{s}_1^{t-1})\cdot$$
$$d\mu_k^{t-1}(\mathbf{s}_k^{t-1})\ldots d\mu_1^{t-1}(\mathbf{s}_1^{t-1}) > 0.$$

□

Lemma 6. *Assumption 3 implies Assumption 6.*

Proof. Let $B = \{(\omega,\mathbf{s}_{-j}^t)|\omega\in\hat{\Omega},\mathbf{s}_{-j}^t\in\hat{\mathbf{S}}_{-j}^t(\omega)\}$ where $\hat{\Omega}\subseteq\Omega$ and $\hat{\mathbf{S}}^t(\omega)\subseteq\mathbf{S}_{-j}^t$ for each $\omega\in\hat{\Omega}$. The joint measure of B, conditional on $\mathbf{s}_j^t$, is given by,

$$\int_{\omega^{t-1}\in\Omega^{t-1}}\int_{\omega_t\in\hat{\Omega}}\int_{\mathbf{s}_{-j}^t\in\hat{\mathbf{S}}_{-j}^t(\omega_t)} d\mu\left(\mathbf{s}_{-j}^t|\omega_t,\omega^{t-1},\mathbf{s}_j^t\right) d\mu\left(\omega_t|\omega^{t-1},\mathbf{s}_j^t\right) d\mu(\omega^{t-1}|\mathbf{s}_j^t) =$$
$$\int_{\omega^{t-1}\in\Omega^{t-1}}\int_{\omega_t\in\hat{\Omega}}\int_{\varepsilon_{-j}^t\in\hat{\mathbf{S}}_{-j}^t(\omega_t)-\bar{\omega}(\omega^{t-1},\omega_t)} \prod_{\tau\leq t} f^{\varepsilon}(\varepsilon_{-j,\tau}^t,\mathbf{s}_{j,\tau}^t-\omega_\tau^t)d\mu^{\varepsilon}\left(\varepsilon_{-j,\tau}^t\right) d\mu\left(\omega_t|\omega^{t-1},\mathbf{s}_j^t\right) d\mu(\omega^{t-1}|\mathbf{s}_j^t),$$

where $\bar{\omega}(\omega^{t-1},\omega_t) = \left((\omega_1^{t-1})_{i\in N\backslash\{j\}},\ldots,(\omega_{t-1}^{t-1})_{i\in N\backslash\{j\}},(\omega_t)_{i\in N\backslash\{j\}}\right)$ and $\omega^t = (\omega^{t-1},\omega_t)$. This implies

$$\int_{\omega^{t-1}\in\Omega^{t-1}}\int_{\omega_t\in\hat{\Omega}} d\mu(\omega_t|\omega^{t-1},\mathbf{s}_j^t)d\mu(\omega^{t-1}|\mathbf{s}_j^t) > 0.$$

It also implies that it is without loss to assume that for each $\omega_t \in \hat{\Omega}$, there is $\omega^{t-1}(\omega_t) \in \Omega^{t-1}$ such that

$$\int_{\varepsilon_{-j}^t\in\hat{\mathbf{S}}_{-j}^t(\omega_t)-\bar{\omega}(\omega^{t-1}(\omega_t),\omega_t)}\prod_{\tau\le t} f^{\varepsilon}(\varepsilon_{-j,\tau}^t,\mathbf{s}_{j,\tau}^t-\omega_\tau^t(\omega_t))d\mu^{\varepsilon}\left(\varepsilon_{-j,\tau}^t\right) > 0,$$

where $\omega^t(\omega_t) = (\omega^{t-1}(\omega_t),\omega_t)$.

By Assumption 3 (b), for each $\omega_t \in \hat{\Omega}$ there is $\delta(\omega_t)$ such that

$$\int_{\varepsilon_{-j}^t\in\hat{\mathbf{S}}_{-j}^t(\omega_t)-\bar{\omega}(\widetilde{\omega}^{t-1},\widetilde{\omega}_t)}\prod_{\tau\le t} f^{\varepsilon}(\varepsilon_{-j,\tau}^t,\mathbf{s}_{j,\tau}^t-\omega_\tau^t)d\mu^{\varepsilon}\left(\varepsilon_{-j,\tau}^t\right) > 0,$$

if $d\left((\widetilde{\omega}^{t-1},\widetilde{\omega}_t),(\omega^{t-1}(\omega_t),\omega_t)\right) < \delta(\omega_t)$ which implies, for $\omega_t \in \hat{\Omega}$,

$$\mu(\hat{\mathbf{S}}_{-j}^t(\omega_t)|\mathbf{s}_j^t) = \int_{\widetilde{\omega}^t\in\Omega^t}\int_{\varepsilon_{-j}^t\in\hat{\mathbf{S}}_{-j}^t(\omega_t)-\bar{\omega}(\widetilde{\omega}^t)}\prod_{\tau\le t} f^{\varepsilon}(\varepsilon_{-j,\tau}^t,\mathbf{s}_{j,\tau}^t-\widetilde{\omega}_\tau^t)d\mu^{\varepsilon}\left(\varepsilon_{-j,\tau}^t\right) > 0.$$

This completes our proof. □

PROOF OF LEMMA 1

Let M be such that $\sup_{h\in\mathscr{H}}|g_i(h)| < M$. Recall that the state at time j is given by $\omega_j = (\gamma_j, t_j, N_j)$. Define

$$\bar{g}(\mathbf{a}^{t-1},\omega^t,\mathbf{s}^t) = \begin{cases} M & \text{if } \gamma_t \neq \gamma^{end} \\ 0 & \text{otherwise.}\end{cases}$$

Then, for each player i, $|g_i^{t,max,t-1}(\mathbf{a}^{t-1},\omega^t,\mathbf{s}^t)| \le \bar{g}(\mathbf{a}^{t-1},\omega^t,\mathbf{s}^t)$.

Let L_σ be a random variable that represents the length of the game given $\mu(|\sigma)$. We have

$$\sum_{t=1}^{\infty}\int_{\Omega\times S} g^{t,max,1}(\omega_1)\,d\mu(\omega_1,\mathbf{s}_1|\emptyset) \le M\cdot\sup_{\sigma\in\Sigma}\sum_{t=1}^{\infty}\sum_{\tilde{t}>t}\mathbb{P}(L_\sigma=\tilde{t}) = M\cdot\sup_{\sigma\in\Sigma}\sum_{\tilde{t}=1}^{\infty}(\tilde{t}-1)\mathbb{P}(L_\sigma=\tilde{t}),$$

where the inequality obtains from the definition of $g^{t,max,j}$, and the equality is obtained by changing the order of the summations. The right hand side is finite under the second hypothesis of Lemma 1.

PROOF OF LEMMA 2

Let us first establish a useful Lemma. Let $Y = \times_{i\in N} Y_i$ be a measurable space with Borel σ-algebra $\mathscr{B}_Y = \otimes_{i\in N}\mathscr{B}_{Y_i}$, where $\mathscr{B}_{Y_i}$ is the Borel σ-algebra in Y_i, and measure $\tilde{\mu}$ such that

$$d\tilde{\mu}(y_1,y_2,\dots,y_n) = g(y_1,y_2,\dots,y_n)d\kappa_1(y_1)d\kappa_2(y_2)\cdots d\kappa_n(y_n),$$

with each $\kappa_i : Y_i \to [0,1]$ a finite measure and $g \in L^1(Y,\tilde{\mu})$. Let $\kappa_{-i} = \otimes_{j\in N\setminus\{i\}}\kappa_j$.

In what follows convergence is in the weak topology of $L^2(Y,\tilde{\mu})$.

Lemma 7. *Let $\{f_i^\alpha\}_\alpha$, with $f_i^\alpha : Y_i \to \mathbb{R}$, be a uniformly bounded net of functions (indexed by α) that converges to f_i^* for each i. Then $\left\{\prod_{i\in N} f_i^\alpha\right\}_\alpha$ has a subnet that converges to $\prod_{i\in N} f_i^*$.*

Proof of Lemma 7. We need to show that for every $\psi \in L^2(Y,\tilde{\mu})$

$$\int_{y\in Y}\prod_{i\in N} f_i^\alpha(y_i)\psi(y)\,d\tilde{\mu}(y) \to \int_{y\in Y}\prod_{i\in N} f_i^*(y_i)\psi(y)\,d\tilde{\mu}(y).$$

Since f_i^α is uniformly bounded and simple functions are dense in L^2 it is sufficient to show that the equality holds for test functions of the form $\psi(y) = \prod_{i=1}^n \mathbf{1}\{B_i\}$ for $B_i \in \mathscr{B}_{Y_i}$ for each $i \in \{1,\dots,n\}$.[65]

We argue by induction. If f_1^α converges weakly to f_1^* then, by the definition of weak convergence we have

$$\int_{y\in Y} f_1^\alpha(y_1)\psi(y)\,d\tilde{\mu}(y) \to \int_{y\in Y} f_1^*(y_1)\psi(y)\,d\tilde{\mu}(y).$$

Suppose $\prod_{i=1}^{j-1} f_i^\alpha(y_i)$ converges to $\prod_{i=1}^{j-1} f_i^*(y_i)$ in the weak topology of $L^2(Y,\tilde{\mu})$. Let's see that there is a subnet of $\prod_{i=1}^{j} f_i^\alpha(y_i)$ that converges to $\prod_{i=1}^{j} f_i^*(y_i)$.

[65]Suppose $\int(f^\alpha - f^*)\varphi_s d\mu$ for every simple function φ_s. Let $\varphi \in L^2$, there is a sequence of simple functions $\{\varphi_n\}$ such that $\varphi_n \to \varphi$ in L^2. Thus,

$$\int(f^\alpha - f)\varphi d\mu = \int(f^\alpha - f)(\varphi_n - \varphi)d\mu + \int(f^\alpha - f)\varphi_n d\mu. \qquad (*)$$

The first integral converges to zero uniformly in α since $f^\alpha - f$ is uniformly bounded and $\varphi_n \to \varphi$ in L^2. Therefore, $(*)$ converges to zero.

Since f_j^α is uniformly bounded and $\tilde{\mu}$ is finite, it is in $L^\infty(Y_j,\kappa_j) = (L^1(Y_j,\kappa_j))^*$ and—by Banach-Alaoglu—it is contained in a compact set. Thus, there is a subnet of f_j^α that converges weakly to some function f_j^{**} in $L^\infty(Y_j,\kappa_j)$. Passing to the subnet yields,

$$\int_{y_j\in Y_j} f_j^\alpha(y_j)\phi(y_j)\,d\kappa_j(y_j) \to \int_{y_j\in Y_j} f_j^{**}(y_j)\phi(y_j)\,d\kappa_j(y_j), \tag{4}$$

Since $g(\cdot,y_{-j})$ is in $L^1(Y_j,\kappa_j)$, evaluating (4) on $\phi = g(\cdot,y_{-j})\mathbf{1}\{B\}$ for $B\in\mathscr{B}_Y$ and integrating over Y_{-j} under measure κ_{-j} shows that we must have $f_j^* = f_j^{**}$, $\tilde{\mu}$-almost surely.[66]

By the induction hypothesis it is sufficient to show that

$$\int_{y\in Y}\left(\prod_{i=1}^{j-1} f_i^\alpha(y_i)\left[f_j^\alpha(y_j)-f_j^*(y_j)\right]\psi(y)\right)d\tilde{\mu}(y)\to 0.$$

Rewriting the left hand side of the previous expression using the Radon-Nikodym derivatives yields

$$\Lambda^\alpha(g) = \int_{y_{-j}\in Y_{-j}}\prod_{i=1}^{j-1} f_i^\alpha(y_i)\left(\int_{y_j\in Y_j}\left[f_j^\alpha(y_j)-f_j^*(y_j)\right]\psi(y_j,y_{-j})g(y_j,y_{-j})\,d\kappa_j(y_j)\right)d\kappa_{-j}(y_{-j}). \tag{5}$$

Let $\varepsilon>0$, and let $\bar{M}>0$ be a bound on $\left|\prod_{i=1}^{j-1} f_i^\alpha(y_i)\left[f_j^\alpha(y_j)-f_j^*(y_j)\right]\right|$. The set of simple functions with support on $\{B_1\times\cdots\times B_n | B_i\in\mathscr{B}_{Y_i}\}$ is dense in $L^1(Y,\tilde{\mu})$. There is $\varphi(y)=\sum_{k=1}^K a_k\mathbf{1}\{B_j^k\}1\{B_{-j}^k\}$, with $a_k\in\mathbb{R}$, $B_j^k\in\mathscr{B}_{Y_j}$ and $B_{-j}^k\in\mathscr{B}_{Y_{-j}}$ for each $k\in\{1,\ldots,K\}$, such that $\|g-\varphi\|_1<\frac{\varepsilon}{2\bar{M}}$. The expression in (5) can be rewritten as

$$\begin{aligned}&\Lambda^\alpha(g)-\Lambda^\alpha(\varphi)+\Lambda^\alpha(\varphi)\le\\ &\sum_{k=1}^K a_k\left(\int_{y_{-j}\in Y_{-j}}\mathbf{1}\{\tilde{B}_{-j}^k\}\prod_{i=1}^{j-1} f_i^\alpha(y_i)\,d\kappa_{-j}(y_{-j})\right)\left(\int_{y_j\in Y_j}\left[f_j^\alpha(y_j)-f_j^*(y_j)\right]\mathbf{1}\{\tilde{B}_j^k\}\,d\kappa_j(y_j)\right)\\ &\qquad+\bar{M}\|g-\varphi\|_1,\end{aligned} \tag{6}$$

where $\tilde{B}_{-j}^k = B_{-j}^k\cap B_1\times\ldots B_{j-1}$ and $\tilde{B}_j^k = B_j^k\cap B_j$. The first parenthesis in each term of the summation is bounded. Therefore, by (4) there is $\bar{\alpha}$ such that for every $\alpha\ge\bar{\alpha}$ the summation in (6) is less than $\varepsilon/2$. This shows that for every $\varepsilon>0$ there is $\bar{\alpha}$ such that for every $\alpha\ge\bar{\alpha}$, $\Lambda^\alpha(g)\le\varepsilon$.

[66] By Theorem 4.48 in Aliprantis and Border (1999), $\phi(\cdot,y_{-j})$ is measurable in Y_j.

□

Next is the proof of Lemma 2.

Claim 1. *Let $\{\sigma^\alpha\} \subseteq \Sigma$ be a net. There is σ^* and a subnet $\{\sigma^{\hat{\alpha}}\}$ such that*

$$\langle prob_i(\cdot, \sigma_i^{\hat{\alpha}}, \tilde{H}_i(\cdot, \sigma^*)), \psi\rangle \to \langle prob_i(\cdot, \sigma_i^*, \tilde{H}_i(\cdot, \sigma^*)), \psi\rangle \tag{7}$$

for every $\psi \in \mathscr{H}^t$, $t \in \mathbb{N}$ and $i \in N$ (i.e. $\sigma^{\hat{\alpha}}$ converges weakly in strategies to σ^*).

Proof of Claim 1. We will argue by induction on the length of the history of play.

Consider first the set of histories of length one, denoted $\mathscr{H}^1$, and let $\hat{\mu}^1$ denote the restriction of $\hat{\mu}$ to $\mathscr{H}^1$. Since $\hat{\mu}^1$ is a finite measure, the restriction of $prob_i$ to $\mathscr{H}^1$, denoted $prob_i^1$, is in $L^2(\mathscr{H}^1, \hat{\mu}^1)$. Therefore, there is a subnet of σ^α and a function $prob_i^{*,1} \in L^2(\mathscr{H}^1, \hat{\mu}^1)$, such that $\langle prob_i^1(\cdot, \sigma_i^\alpha), \psi\rangle = \langle prob_i(\cdot, \sigma_i^\alpha), \psi\rangle \to \langle prob_i^{*,1}(\cdot), \psi\rangle$ for ψ with support in $\mathscr{H}^1$. Each player i' s strategy after the empty history is defined as $\sigma_i^*((\hat{\mathbf{a}}(h))_{1,i}|\emptyset) = prob_i^{*,1}(h)$. Thus, condition (7) is satisfied for ψ with support in $\mathscr{H}^1$.

Now suppose we have defined $\sigma_i^*(\cdot|h_i(h))$ at each $h \in \mathscr{H}^{t-1}$ and that there is a subnet of σ_i^α (which slightly abusing notation we also denote σ_i^α) such that for every $j \leq t-1$, $\langle prob_i(\cdot, \sigma_i^\alpha, \tilde{H}_i(h, \sigma^*)), \psi\rangle \to \langle prob_i(\cdot, \sigma_i^*, \tilde{H}_i(h, \sigma^*)), \psi\rangle$ for ψ with support in $\mathscr{H}^j$.

As before, we can find a subnet of σ^α and a function $prob_i^{*,t}$ such that, passing to the subnet, $\langle prob_i(\cdot, \sigma_i^\alpha, \tilde{H}_i(h, \sigma^*)), \psi\rangle \to \langle prob_i^{*t}(\cdot), \psi\rangle$ for every ψ with support in $\mathscr{H}^t$ (notice that $\tilde{H}_i(h, \sigma^*)$ depends on h via $h^{(t-1)}$ and, therefore, it is well defined by the induction hypothesis).

We now define $\sigma^*(\cdot|h_i(h))$ for $h \in \mathscr{H}^t$ by setting

$$\sigma_i^*((\hat{\mathbf{a}}(h))_{t,i}|h_i(h)) = \begin{cases} prob_i^{*,t}(h) & \text{if } prob_i^{*,t-1}(h^{(t-1)}) = 0 \\ prob_i^{*,t}(h)/prob_i^{*,t-1}(h^{(t-1)}) & \text{if } prob_i^{*,t-1}(h^{(t-1)}) > 0 \end{cases}.$$

By definition, $\langle prob_i(\cdot, \sigma_i^\alpha, \tilde{H}_i(h, \sigma^*)), \psi\rangle \to \langle prob_i(\cdot, \sigma_i^*, \tilde{H}_i(h, \sigma^*)), \psi\rangle$ for ψ with support in $\mathscr{H}^t$. Furthermore, the resulting σ_i^* is a strategy. In fact, σ_i^α measurable with respect to i's information implies that σ_i^* is measurable with respect to i's information almost surely. Also, since σ_i^α is a probability over available actions, so is σ_i^*. Thus, we conclude that σ^α converges weakly in strategies to σ^*. □

Claim 2. *If σ^α converges weakly in strategies to σ^* then $prob(\cdot, \sigma^\alpha)$ has a subnet that converges to $prob(\cdot, \sigma^*)$ in the weak topology of $L^2(\mathscr{H}, \hat{\mu})$.*

Proof. Proof of Claim 2

We first show that

$$\langle prob(\cdot,\sigma^{\alpha}),\psi\rangle \to \langle prob(\cdot,\sigma^{*}),\psi\rangle \tag{8}$$

holds for ψ with support in $\mathscr{H}^t$ for each t.

Since $prob(h,\sigma^{\alpha})$ is uniformly bounded, it is enough to show (8) holds for $\psi = \mathbf{1}\{B\}$ for B measurable under $\hat{\mu}$.[67]

To see that condition (8) holds for ψ with support in $\mathscr{H}^t$, notice that by Assumption 2 we can write for every $\hat{\sigma} \in \Sigma$

$$\int_{\mathscr{H}_t} prob(h,\hat{\sigma})\psi(h)\,d\mu^t(h) = \int_{X^t\times \mathbf{S}^t} prob(h,\hat{\sigma})\int_{\Omega^t}\psi(h)d\hat{\mu}^{\omega}(\omega^t|\mathbf{s}^t,\mathbf{a}^t)d\hat{\mu}^{s}(\mathbf{s}^t|\mathbf{a}^t).$$

Since $\int_{\Omega^t}\psi(h)d\hat{\mu}^{\omega}(\omega^t|\mathbf{s}^t,\mathbf{a}^t)$ is bounded and measurable with respect to $\mathbf{s}^t$ for each $\mathbf{a}^t$, Lemma 7 implies that $prob(h,\sigma^{\alpha})$ has a subnet such that,

$$\int_{\mathscr{H}_t} prob(h,\sigma^{\alpha})\psi(h)\,d\mu^t(h) \to \int_{\mathscr{H}_t} prob(h,\sigma^{*})\psi(h)\,d\mu^t(h).$$

Now, to see that $prob(\cdot,\sigma^{\alpha})$ converges in the weak topology of $L^2(\mathscr{H},\hat{\mu})$ notice that

$$\delta^t\int_{\mathscr{H}^t} prob(h,\sigma^{\alpha})\psi(h)\,d\mu^t(h) \le \delta^t.$$

Therefore, the dominated convergence theorem yields the desired result. □

Lemma 7 and Claims 1 and 2 establish the second statement of Lemma 2. The first statement follows from Claim 1 that establishes the existence of the limit σ^*. Claim 2 implies that that a subnet of f^{α} converges to $prob(\cdot,\sigma^*)$. Thus, we must have $f^* = prob(\cdot,\sigma^*)$.

PROOF OF PROPOSITION 2

Define for each player $i \in N$ and $t \in \mathbb{N}$, $g_i|_{\mathscr{H}^t}(h) := g_i(h)\cdot\mathbf{1}\{h\in\mathscr{H}^t\}$ for $h \in \mathscr{H}$. Since $g_i|_{\mathscr{H}^t} \in L^2(\mathscr{H},\hat{\mu})$ we have

$$\langle prob(\cdot,\sigma^{\alpha}),g_i|_{\mathscr{H}^t}\rangle \to \langle prob(\cdot,\sigma^{*}),g_i|_{\mathscr{H}^t}\rangle.$$

Notice that we can write $U_i(\sigma) = \sum_{t\in\mathbb{N}}\frac{1}{\delta^t}\langle prob(\cdot,\sigma),g_i|_{\mathscr{H}^t}\rangle$. Lemma 8 (below) shows

[67] See footnote 65

that $|\frac{1}{\delta^t}\langle prob(\cdot,\sigma), g_i|_{\mathscr{H}^t}\rangle| \leq \int_{\Omega\times\mathbf{S}} g^{t,max,1}(\omega_1,\mathbf{s}_1)\,d\mu(\omega_1,\mathbf{s}_1|\emptyset)$ for every $\sigma\in\Sigma$ and $t\in\mathbb{N}$. Therefore, since by Assumption 1, $\sum_{t=1}^{\infty}\int_{\Omega\times\mathbf{S}} g^{t,max,1}(\omega_1,\mathbf{s}_1)\,d\mu(\omega_1,\mathbf{s}_1|\emptyset) < \infty$. By the dominated convergence theorem $U_i(\sigma^\alpha)\to U_i(\sigma^*)$.

Lemma 8. *For every $\sigma\in\Sigma$, we have*

$$\left|\frac{1}{\delta^t}\langle prob(\cdot,\sigma), g_i|_{\mathscr{H}^t}\rangle\right| \leq \int_{\Omega\times\mathbf{S}} g^{t,max,1}(\omega_1,\mathbf{s}_1)\,d\mu^{\omega,s}(\omega_1,\mathbf{s}_1|\emptyset)$$

Proof of Lemma 8. First note that

$$\begin{aligned}\left|\frac{1}{\delta^t}\langle prob(\cdot,\sigma), g_i|_{\mathscr{H}^t}\rangle\right| &= \left|\sum_{\mathbf{a}^{t-1}\in X^{t-1}}\int_{\Omega^t\times\mathbf{S}^t}\sum_{\mathbf{a}_t\in X} g_i(\omega^t,\mathbf{a}^t)prob(\mathbf{a}^t|\mathbf{s}^t,\sigma)d\mu^{\omega,s}(\omega^t,\mathbf{s}^t|\mathbf{a}^{t-1})\right| \\ &\leq \sum_{\mathbf{a}^{t-1}\in X^{t-1}}\int_{\Omega^t\times\mathbf{S}^t} g_i^{t,max,t-1}(\omega^t,\mathbf{s}^t,\mathbf{a}^{t-1})prob(\mathbf{a}^{t-1}|\mathbf{s}^{t-1},\sigma)d\mu^{\omega,s}(\omega^t,\mathbf{s}^t|\mathbf{a}^{t-1}),\end{aligned}$$

where in the previous equation $\mathbf{a}^t = (\mathbf{a}^{t-1},\mathbf{a}_t)$.

We argue by induction, assume that $\left|\frac{1}{\delta^t}\langle prob(\cdot,\sigma), g_i|_{\mathscr{H}^t}\rangle\right| \leq G^l$ where

$$G^l = \sum_{\mathbf{a}^l\in X^l}\int_{\Omega^{l+1}\times\mathbf{S}^{l+1}} g_i^{t,max,l}(\omega^{l+1},\mathbf{s}^{l+1},\mathbf{a}^l)prob(\mathbf{a}^l|\mathbf{s}^l,\sigma)\,d\mu^{\omega,s}(\omega^{l+1},s^{l+1}|\mathbf{a}^l).$$

for some $l\in\{2,\dots,t\}$. We now show that

$$\left|\frac{1}{\delta^t}\langle prob(\cdot,\sigma), g_i|_{\mathscr{H}^t}\rangle\right| \leq G^{l-1}.$$

In fact, we will show that $G^l \leq G^{l-1}$. We have

$$\begin{aligned}G^l &= \sum_{\mathbf{a}^{l-1}\in X^{l-1}}\int_{\Omega^l\times\mathbf{S}^l}\sum_{\mathbf{a}_l\in X} prob(\mathbf{a}^l|\mathbf{s}^l,\sigma)\int_{\Omega\times\mathbf{S}} g_i^{t,max,l}(\omega^{l+1},\mathbf{s}^{l+1},\mathbf{a}^l)\,d\mu^{\omega,s}(\omega_{l+1},s_{l+1}|\omega^l,\mathbf{s}^l,\mathbf{a}^l)\,d\mu^{\omega,s}(\omega^l,\mathbf{s}^l|\mathbf{a}^{l-1}) \\ &\leq \sum_{\mathbf{a}^{l-1}\in X^{l-1}}\int_{\Omega^l\times\mathbf{S}^l}\sum_{\mathbf{a}_l\in X} prob(\mathbf{a}^l|\mathbf{s}^l,\sigma)g^{t,max,l-1}(\omega^l,\mathbf{s}^l,\mathbf{a}^{l-1})\,d\mu^{\omega,s}(\omega^l,\mathbf{s}^l|\mathbf{a}^{l-1}) = G^{l-1}.\end{aligned}$$

Thus, we obtain,

$$\left|\frac{1}{\delta^t}\langle prob(\cdot,\sigma), g_i|_{\mathscr{H}^t}\rangle\right| \leq G^1 = \int_{\Omega\times\mathbf{S}} g^{t,max,1}(\omega,\mathbf{s})\,d\mu^{\omega,s}(\omega,\mathbf{s}|\emptyset).$$

$\square$

PROOF OF LEMMA 3

It is immediate that $r_i(\sigma_{-i})$ is convex for each $i \in N$ and $\sigma_{-i} \in \Sigma_{-i}(\tilde{\varepsilon})$. Let us now show that r is non-empty and that it has closed graph.

***r* is non-empty:**

Define the set $\Lambda_i(\sigma_{-i}, \tilde{\varepsilon}) := \{prob(\cdot, \sigma_i, \sigma_{-i}) | \sigma_i \in \Sigma_i(\tilde{\varepsilon}),\}$ for $\sigma_{-i} \in \Sigma_{-i}(\tilde{\varepsilon})$. $\Lambda_i(\sigma_{-i}, \tilde{\varepsilon})$ is closed. In fact, let $\{prob(\cdot, \sigma^{\alpha}, \sigma_{-i})\} \subseteq \Lambda_i(\sigma_{-i}, \tilde{\varepsilon})$ be a net that converges to $prob(\cdot, \tilde{\sigma}_i, \tilde{\sigma}_{-i})$ (By Lemma 2 such a strategy , $\tilde{\sigma}_i, \tilde{\sigma}_{-i}$, exists). Also by Lemma 2, $(\sigma_i^{\alpha}, \sigma_{-i})$ has a subnet that converges weakly in strategies to $(\tilde{\sigma}_i, \tilde{\sigma}_{-i})$. This implies $\sigma_{-i} = \tilde{\sigma}_{-i}$ almost surely.

Since $\Lambda_i(\sigma_{-i}, \tilde{\varepsilon})$ is closed, it is compact. Now, by Proposition 2, $U_i(\cdot, \sigma_{-i})$ is continuous in $prob(\cdot, \sigma_i, \sigma_{-i}) \in \Lambda_i(\sigma_{-i}, \tilde{\varepsilon})$ (with the topology of weak convergence of $L^2(\mathscr{H}, \hat{\mu})$). Thus, $U_i(\cdot, \sigma_{-i})$ attains its maximum over $\Lambda_i(\sigma_{-i}, \tilde{\varepsilon})$ (which is homeomorphic to $\Sigma_i(\tilde{\varepsilon})$).

***r* has closed graph:**

Let $\{\sigma^{\alpha}\}$ and $\{\hat{\sigma}^{\alpha}\}$ be nets such that $\hat{\sigma}^{\alpha} \in r(\sigma^{\alpha})$, $\hat{\sigma}^{\alpha} \to \hat{\sigma}^*$ and $\sigma^{\alpha} \to \sigma^*$ weakly in probabilities. Let us show that $\hat{\sigma}^* \in r(\sigma^*)$ (i.e. r has closed graph).

$\hat{\sigma}^{\alpha} \in r(\sigma^{\alpha})$ implies that for every $\sigma_i \in \Sigma_i(\tilde{\varepsilon})$

$$U_i(\hat{\sigma}_i^{\alpha}, \sigma_{-i}^{\alpha}) \geq U_i(\sigma_i, \sigma_{-i}^{\alpha}) \tag{9}$$

Also, by Lemma 7 there are subnets of $\hat{\sigma}^{\alpha}$ and σ^{α} that converge weakly in strategies to $\hat{\sigma}^*$ and σ^*, respectively.

Now, by the definition of weak convergence in strategies, for each player $i \in N$ and $\sigma_i \in \Sigma_i(\tilde{\varepsilon})$, passing to the subnet, $(\hat{\sigma}_i^{\alpha}, \sigma_{-i}^{\alpha})$ and $(\sigma_i, \sigma_{-i}^{\alpha})$ converge weakly in strategies to $(\hat{\sigma}_i^*, \sigma_{-i}^*)$ and $(\sigma_i, \sigma_{-i}^*)$ respectively. Then, by Lemma 7, there are subnets of $\hat{\sigma}^{\alpha}$ and σ^{α} such that passing to the subnet, $prob(\cdot, \hat{\sigma}_i^{\alpha}, \sigma_{-i}^{\alpha})$ and $prob(\cdot, \sigma_i, \sigma_{-i}^{\alpha})$ converge to $prob(\cdot, \hat{\sigma}_i^*, \sigma_{-i}^*)$ and $prob(\cdot, \sigma_i, \sigma_{-i}^*)$, respectively, in the weak topology of $L^2(\mathscr{H}, \hat{\mu})$. Finally, by Proposition 2, equation (9) implies $\hat{\sigma}_i^* \in r_i(\sigma_i^*)$ for each $i \in N$.

PROOF OF LEMMA 4

Let σ^* be a THPE and let σ^m be a sequence of $\tilde{\varepsilon}^m$-constrained equilibria with each $\tilde{\varepsilon}^m$ an (ε^m, ν^m)-tremble with $\varepsilon^m \to 0$ converging weakly in strategies to σ^*.

Then for every player i and every $\sigma_i^\varepsilon \in \Sigma_i(\tilde{\varepsilon}^m)$

$$U_i(\sigma^m) \geq U_i(\sigma_i^\varepsilon, \sigma_{-i}^m). \tag{10}$$

Let $\tilde{\sigma}_i \in \Sigma_i$. There is a sequence $\{\tilde{\sigma}_i^{\varepsilon^m}\}_m$, with $\tilde{\sigma}_i^{\varepsilon^m} \in \Sigma_i(\tilde{\varepsilon}^m)$, that converges a.s. to $\tilde{\sigma}_i$. By Proposition 2 σ^m and $(\tilde{\sigma}^m, \sigma_{-i}^m)$ converge weakly in probabilities to σ^* and $(\tilde{\sigma}_i, \sigma_{-i}^*)$, respectively. By Proposition 2 $U_i(\sigma^m) \to U_i(\sigma^*)$ and $U_i(\tilde{\sigma}_i^m, \sigma_{-i}^m) \to U_i(\tilde{\sigma}_i, \sigma_{-i})$, which, by (10), yields

$$U_i(\sigma^*) \geq U_i(\tilde{\sigma}_i, \sigma_{-i}^*).$$

PROOF OF PROPOSITION 3

The following Lemma establishes Proposition 3

Lemma 9. *Let $\tilde{\varepsilon}$ be a (ε, ν)-tremble, let σ^ε be an $\tilde{\varepsilon}$-constrained equilibrium, and let $\tilde{\mu}^\varepsilon$ its corresponding weakly consistent belief, then*

$$U_i(\sigma^\varepsilon | h_i(h), \tilde{\mu}^\varepsilon) \geq U_i(\sigma_i', \sigma_{-i}^\varepsilon | h_i(h), \tilde{\mu}^\varepsilon). \tag{11}$$

for every $\sigma_i' \in \Sigma_i(\tilde{\varepsilon})$, $\hat{\mu}$-almost surely in $h \in \mathscr{H}$.

Proof of Lemma 9. Suppose there is a set $\tilde{H} \subseteq \mathscr{H}_i$, with $\hat{\mu}\left(\{h \in \mathscr{H} | h_i(h) \in \tilde{H}\}\right) > 0$ and a strategy σ_i' for each $\tilde{h}_i \in \tilde{H}$, such that

$$U_i\left(\sigma^\varepsilon | \tilde{h}_i, \tilde{\mu}^\varepsilon\right) < U_i\left(\sigma_i', \sigma_{-i}^\varepsilon | \tilde{h}_i, \tilde{\mu}^\varepsilon\right), \tag{12}$$

for $\tilde{h}_i \in \tilde{H}$.

Let $d\mu_i^\varepsilon(\tilde{h}_i) = \left(\sum_{l=0}^\infty \int_{h \in \mathscr{H}_i^l(\tilde{h}_i)} prob_{-i}(h, \sigma^\varepsilon, \tilde{h}_i)\, d\mu_{-i}^l(h|\tilde{h}_i)\right) \cdot prob_i(\tilde{h}_i, \sigma_i^\varepsilon)\, d\alpha_{h_i}(\tilde{h}_i)$ where $prob_i(h_i, \sigma_i^\varepsilon) := prob_i(h, \sigma_i^\varepsilon)$ for h such that $h_i(h) = \tilde{h}_i$.

Define the strategy σ_i' to be equal to σ_i^ε for every $h \notin \tilde{H}$ and $\sigma_i'(\tilde{h}_i)$ after $\tilde{h}_i$ for each $\tilde{h}_i \in \tilde{H}$. Since $prob_{-i}(h, \sigma^\varepsilon, \tilde{h}_i) > 0$ and $prob_i(h, \sigma_i^\varepsilon) > 0$ for every strategy $\sigma^\varepsilon \in \Sigma_i(\tilde{\varepsilon})$, equation (12) implies

$$\int_{\tilde{h}_i \in \tilde{H}} U_i(\sigma_i | \tilde{h}_i, \tilde{\mu}^\varepsilon) d\mu_i^\varepsilon(\tilde{h}_i) < \int_{\tilde{h}_i \in \tilde{H}} U_i(\sigma_i'(h_i), \sigma_{-i}^\varepsilon | \tilde{h}_i, \tilde{\mu}^\varepsilon)\, d\mu_i^\varepsilon(\tilde{h}_i).$$

However, from from the definition of α_{h_i} and $\mu_{-i}^l(h|h_i)$ the previous expression can be written as

$$\sum_{t=0}^{\infty}\int_{h\in\mathscr{H}_i(\tilde{H})\cap\mathscr{H}^t} g_i(h)\cdot prob(h,\sigma^{\varepsilon})\,d\mu^t(h) < \sum_{t=0}^{\infty}\int_{h\in\mathscr{H}_i(\tilde{H})\cap\mathscr{H}^t} g_i(h)\cdot prob(h,\sigma_i',\sigma_{-i}^{\varepsilon})\,d\mu^t(h),$$

which contradicts σ^{ε} is $\tilde{\varepsilon}$-constrained equilibrium: strategy σ_i' is a profitable deviation for player i.

□

B Online Appendix

Proof of Theorem 2

Let us first establish two useful Lemmas. In what follows Y_1 and Y_2 are Polish spaces with Borel sets $\mathscr{Y}_1$ and $\mathscr{Y}_2$, and η_1 and η_2 are finite Borel measures over Y_1 and Y_2, respectively.

Lemma 10. *Suppose that*

L.10.a $\{\phi^m\}_{m\in\mathbb{N}}$ *is a sequence of functions from $Y_1\times Y_2$ to $\mathbb{R}$, where $\phi^m(y_1,\cdot)\in L^1(Y_2,\eta_2)$ for each $y_1\subseteq\bar{Y}_1\subseteq Y_1$, and such that the family of functions $\mathscr{F}=\{p^m:\bar{Y}_1\to L^1(Y_2,\eta_2)\,|\,p^m(y_1)=\phi^m(y_1,\cdot),m\in\mathbb{N}\}$ is equicontinuous in $y_1\in\bar{Y}_1$.*[68]

L.10.b *There is $M:\bar{Y}_1\to\mathbb{R}$ such that , $\left\|\phi^m(y_1,\cdot)\right\|_{L^1(Y_2,\eta_2)}\le M(y_1)$ for each $y_1\in\bar{Y}_1$.*

Then there exists a subsequence $\{m_r\}_{r\in\mathbb{N}}$ and a family of finitely additive finite measures $\{\eta^*(\cdot,y_1)\}_{y_1\in Y_1}$ with $\eta^*(\cdot,y_1):\mathscr{Y}_2\to[0,1]$, for each $y_1\in Y_1$, such that

L.10.1 *$d\eta^{m_r}(\cdot,y_1):=\phi^{m_r}(y_1,\cdot)d\eta_2(\cdot)\to d\eta^*(\cdot,y_1)$ for every $y_1\in\bar{Y}_1$ in the weak-* topology of $(L^\infty(Y_2,\eta_2))^*$,*

L.10.2 *$\eta^*(B_2,y_1)$ is measurable in y_1 for each $B_2\in\mathscr{Y}_2$*

Proof. First notice that by condition L.10.b, due to the compactness of the unit ball in the weak-* topology of $\left(L^\infty(Y_2,\eta_2)\right)^*$, for each $y_1\in Y_1$ there is a subsequence $\{\tilde{m}_r\}_{r\in\mathbb{N}}$ (that may depend on y_1) and a finitely additive measure $\hat{\eta}^*(\cdot,y_1)$ such that

$$\int_{Y_2}\mathbf{1}\{B_2\}\,d\eta^{\tilde{m}_r}(y_2,y_1)\to\int_{Y_2}\mathbf{1}\{B_2\}\,d\hat{\eta}^*(y_2,y_1). \tag{13}$$

[68]For the purposes of this equicontinuity, $L^1(Y_2,\eta_2)$ is endowed with its norm topology.

for every $B_2 \in \mathscr{Y}_2$.

Since Y_2 is a Polish space, by L.10.a, for each $k \in \mathbb{N}$ there is a countable open cover of Y_1, $\mathcal{V} = \{\mathscr{V}_n^k\}_{n\in\mathbb{N}}$, such that[69]

$$\left\|\phi^m(y_1,\cdot) - \phi^m(y_1',\cdot)\right\|_{L^1(Y_2,\eta_2)} < \frac{1}{k}$$

for every $y_1, y_1' \in V_n^k \cap \bar{Y}_1$ and $m \in \mathbb{N}$.

Let $\hat{Y} = \{y_{k,n}\}_{k,n\in\mathbb{N}}$ be such that $y_n^k \in V_n^k \cap \bar{Y}_1$ for $k,n \in \mathbb{N}$. We can find a subsequence $\{m_r\}_{r\in\mathbb{N}}$ such that (13) holds for a limit measure $\hat{\eta}(\cdot,y_1)$ for every $y_1 \in \hat{Y}$.[70] Define $\eta^*(\cdot,y_1) = \hat{\eta}^*(\cdot,y_1)$ for $y_1 \in \hat{Y}$. Now, since (13) also holds for any subsequence of $\{m_r\}$, we define $\eta^*(\cdot,y_1)$, for $y_1 \in \bar{Y}_1 \setminus \hat{Y}$, as the limit, $\hat{\eta}(\cdot,y_1)$, in (13) of a subsequence $\{\tilde{m}_r(y_1)\}_{r\in\mathbb{N}}$ of $\{m_r\}_{r\in\mathbb{N}}$. Define $\eta^*(\cdot,y_1)$ as the zero measure for $y_1 \in Y_1 \setminus \bar{Y}_1$.

The following claim establishes L.10.2 as continuous functions are Borel measurable.

Claim 3. *For every $\varepsilon > 0$ and $y_1 \in \bar{Y}_1$, if $y_1', y_1 \in V_n^k \cap \bar{Y}_1$ for $k \geq \lceil 6/\varepsilon \rceil$ then $|\eta^*(B_2,y_1) - \eta^*(B_2,y_1')| < \varepsilon$. Therefore, $\eta^*(B_2,y_1)$, restricted to $\bar{Y}_1$, is continuous in y_1 for every $B_2 \in \mathscr{Y}_2$.*

Proof of Claim 3.
Let $\varepsilon > 0$, let $y_1 \in \bar{Y}_1$ and $k \geq \lceil 6/\varepsilon \rceil$. Let us see that $|\eta^*(B_2,y_1) - \eta^*(B_2,y_1')| < \varepsilon$ for $y_1, y_1' \in V_n^k \cap \bar{Y}_1$.

We have

$$|\eta^*(B_2,y_1) - \eta^*(B_2,y_1')| \leq |\eta^*(B_2,y_1) - \eta^*(B_2,y_n^k)| + |\eta^*(B_2,y_1') - \eta^*(B_2,y_n^k)|.$$

Let us show that $|\eta^*(B_2,y_1) - \eta^*(B_2,y_n^k)| \leq \frac{\varepsilon}{2}$. The proof that $|\eta^*(B_2,y_1') - \eta^*(B_2,y_n^k)| \leq \frac{\varepsilon}{2}$ is analogous.

$$\begin{aligned}|\eta^*(B_2,y_1) - \eta^*(B_2,y_n^k)| \leq |\eta^*(B_2,y_1) - \eta^{\tilde{m}_r(y_1)}(B_2,y_1)| + \\ |\eta^{\tilde{m}_r(y_1)}(B_2,y_1) - \eta^{\tilde{m}_r(y_1)}(B_2,y_n^k)| + |\eta^{\tilde{m}_r(y_1)}(B_2,y_n^k) - \eta^*(B_2,y_n^k)|.\end{aligned}$$

By the definition of $\{\tilde{m}_r(y_1)\}$, there is $\bar{r}_1 \in \mathbb{N}$ such that for every $r \geq \bar{r}_1$, $|\eta^*(B_2,y_1) - \eta^{\tilde{m}_r(y_1)}(B_2,y_1)| \leq \frac{\varepsilon}{6}$. Since $\{\tilde{m}_r(y_1)\}$ is a subsequence of $\{m_r\}$ there is $\bar{r}_2$ such that

[69] A Polish space is Lindelöf.

[70] This follows straightforwardly by means of a Cantor diagonal argument.

$|\eta^{\tilde{m}_r(y_1)}(B_2,y_n^k)-\eta^*(B_2,y_n^k)|\leq\frac{\varepsilon}{6}$, $\forall r\geq\bar{r}_2$. Finally, since $y_1,y_n^k\in V_n^k\cap\bar{Y}_1$, by L.10.a,

$$|\eta^{\tilde{m}_r(y_1)}(B_2,y_1)-\eta^{\tilde{m}_r(y_1)}(B_2,y_n^k)|\leq\left\|\phi^{\tilde{m}_r(y_1)}(y_1,\cdot)-\phi^{\tilde{m}_r(y_1)}(y_1',\cdot)\right\|_{L^1(Y_2,\eta_2)}\leq\frac{\varepsilon}{6},$$

$\forall r\geq\max\{\bar{r}_1,\bar{r}_2\}$. This concludes the proof of Claim 3. $\square$

The following Claim establishes L.10.1 by L.10.b and footnote 65.

Claim 4. *For every $B_2\in\mathscr{Y}_2$ and $y_1\in\bar{Y}_1$,*

$$\int_{Y_2}\mathbf{1}\{B_2\}\,d\eta^{m_r}(y_2,y_1)\to\int\mathbf{1}\{B_2\}\,d\eta^*(y_2,y_1). \tag{14}$$

Proof of Claim 4.
Let $\varepsilon>0$, $y_1\in\bar{Y}_1$ and $k\geq\lceil\frac{12}{\varepsilon}\rceil$. Let n be such that $y_1\in V_n^k$. We will show that there is $\bar{r}\in\mathbb{N}$ such that for $r\geq\bar{r}$, $|\eta^{m_r}(B_2,y_1)-\eta^*(B_2,y_1)|<\varepsilon$ (which corresponds to (14)).

$$\begin{aligned}|\eta^{m_r}(B_2,y_1)-\eta^*(B_2,y_1)|\leq|\eta^{m_r}(B_2,y_1)-\eta^{m_r}(B_2,y_n^k)|+|\eta^{m_r}(B_2,y_n^k)-\eta^*(B_2,y_n^k)|\\+|\eta^*(B_2,y_n^k)-\eta^*(B_2,y_1)|\end{aligned}$$

First, $|\eta^{m_r}(B_2,y_1)-\eta^{m_r}(B_2,y_n^k)|\leq\left\|\phi^{m_r}(y_1,\cdot)-\phi^{m_r}(y_n^k,\cdot)\right\|_{L^1(Y_2,\eta_2)}\leq\varepsilon/12$ by the definition of V_n^k and y_n^k. Second, by the definition of $\{m_r\}_r$, there is $\bar{r}\in\mathbb{N}$ such that $|\eta^{m_r}(B_2,y_n^k)-\eta^*(B_2,y_n^k)|\leq\varepsilon/4$ for $r\geq\bar{r}$ and, finally, by Claim 3, since $k\geq\lceil\frac{12}{\varepsilon}\rceil$, $|\eta^*(B_2,y_n^k)-\eta^*(B_2,y_1)|\leq\varepsilon/2$. This concludes the proof of Claim 4. $\square$

$\square$

Lemma 11. *Suppose that*

L.11.a *$\{f^m\}_{m\in\mathbb{N}}$ is a bounded sequence of functions in $L^\infty(Y_1,\eta_1)$ such that $f^m\to f^*$ in the weak-* topology of $L^\infty(Y_1,\eta_1)$,*

L.11.b *Conditions L.10.1 and L.10.2 hold with $\bar{Y}_1$ a full η_1-measure set and L.10.b holds for $M\in L^1(Y_1,\eta_1)$.*

Then for every $h\in L^\infty(Y_1\times Y_2,\eta_1\otimes\eta_2)$

$$\int_{Y1\times Y_2}h(y_1,y_2)f^{m_r}(y_1)d\eta^{m_r}(y_2,y_1)d\eta_1(y_1)\to\int_{Y1\times Y_2}h(y_1,y_2)f^*(y_1)d\eta^*(y_2,y_1)d\eta_1(y_1), \tag{15}$$

as $r\to\infty$.

Proof. We start by proving a simple consequence of L.11.a and L.11.b.

Claim 5. *For every $B_1 \in \mathscr{Y}_1$ and $B_2 \in \mathscr{Y}_2$*

$$\int_{Y1\times Y_2} \mathbf{1}\{B_1 \times B_2\} f^{m_r}(y_1) d\eta^*(y_2,y_1) d\eta_1(y_1) \to \int_{Y1\times Y_2} \mathbf{1}\{B_1 \times B_2\} f^*(y_1) d\eta^*(y_2,y_1) d\eta_1(y_1), \tag{16}$$

Proof of Claim 16. $\int_{Y_2} \mathbf{1}\{B_2\} d\eta^*(y_2,y_1) = \eta^*(B_2,y_1)$ is a measurable function in $L^1(Y_1,\eta_1)$ by L.11.b. Therefore, by L.11.a, (16) follows. □

Finally, notice that is it enough to prove to that (15) holds for $h(y_1,y_2) = \mathbf{1}\{y_1 \in B_1, y_2 \in B_2\}$ for $B_1 \in \mathscr{Y}_1$ and $B_2 \in \mathscr{Y}_2$ (by footnote 65 and L.11.b). From this observation and Claim 5, (15) is equivalent to

$$\int_{Y1} \mathbf{1}\{B_1\} f^{m_r}(y_1) \left(\eta^{m_r}(B_2,y_1) - \eta^*(B_1,y_1)\right) d\eta_1(y_1) \to 0. \tag{17}$$

Let $\ell^{m_r}(y_1) = \eta^{m_r}(B_2,y_1) - \eta^*(B_1,y_1)$. From L.11.b, $\ell^{m_r}(y_1) \to 0$ and $|\ell^{m_r}(y_1)| \leq 2M(y_1)$ for every $y_1 \in \bar{Y}_1$. By the boundedness condition in L.11.a, equation (17) then follows by the the dominated convergence theorem, as $\bar{Y}_1$ is of full measure.

□

Define $T_i(s_i^t|\mathbf{a}^{t-1}) := h_i(s_i^t, \mathbf{a}_i^{t-1})$, where $h_i(s_i^t, \mathbf{a}_i^{t-1})$ is the private history of i after signal history s_i^t and i's action history ,$\mathbf{a}_i^{t-1}$.

Let $b_i(\tilde{h}_i, \mathbf{a}^{t-1})$ be the Radon-Nikodym derivative of $\gamma_{h_i}(\tilde{h}_i|\mathbf{a}^{t-1})$ with respect to a measure β_i which exist by Assupmtion 5.

Lemma 12. *Let σ be an $\tilde{\varepsilon}$-constrained strategy for (ε,ν) tremble $\tilde{\varepsilon}$. Under Assumption 5 the following holds*

L.12.1 *the measure $\alpha_{h_i}(\cdot|\sigma)$ in Definition 5 is given by*

$$d\alpha_{h_i}(\tilde{h}_i|\sigma) = \upsilon(\tilde{h}_i,\sigma) d\beta_i(\tilde{h}_i),$$

where

$$\upsilon(\tilde{h}_i,\sigma) = \sum_{\substack{t\in\mathbb{N} \\ \tilde{\mathbf{a}}^{t-1}\in X^{t-1}}} \int_{\mathbf{S}_{-i}^t} prob_{-i}(\tilde{\mathbf{a}}_{-i}^{t-1}|\mathbf{s}_{-i}^{t-1},\sigma) \tilde{f}^s(\mathbf{s}_{-i}^t, \tilde{h}_i, \tilde{\mathbf{a}}^{t-1})\, d\mu_{-i}(\mathbf{s}_{-i}^t|\tilde{\mathbf{a}}^{t-1}),$$

and $\tilde{f}^s(\mathbf{s}_{-i}^t, \tilde{h}_i, \tilde{\mathbf{a}}^{t-1}) = f^s(\mathbf{s}_{-i}^t, T_i^{-1}(\tilde{h}_i|\tilde{\mathbf{a}}^{t-1}), \tilde{\mathbf{a}}^{t-1}) b_i(\tilde{h}_i, \tilde{\mathbf{a}}^{t-1}) \cdot \mathbf{1}\{T_i^{-1}(\tilde{h}_i|\tilde{\mathbf{a}}^{t-1}) \neq \emptyset\}.$

L.12.2 *The measure $\tilde{\mu}^t_{-i}(\cdot|\tilde{h}_i,\sigma)$ is given by,*

$$\tilde{\mu}^t_{-i}(h|\tilde{h}_i,\sigma) = \varphi(\mathbf{a}^t,\mathbf{s}^t,\tilde{h}_i,\sigma)\, d\hat{\mu}^{\omega}(\omega^t|\mathbf{s}^t,\tilde{h}_i,\mathbf{a}^{t,(t-1)})d\mu_{-i}(\mathbf{s}^t_{-i}|\mathbf{a}^{t,(t-1)})\, d\mu^t_i(s^t_i|\tilde{h}_i,\mathbf{a}^{t,(t-1)}), \tag{18}$$

for $h = (\omega^t,\mathbf{a}^t,\mathbf{s}^t)$, where

$$\varphi(\mathbf{a}^t,\mathbf{s}^t,\tilde{h}_i,\sigma) = \frac{1}{\upsilon(\tilde{h}_i,\sigma)} prob_{-i}(h,\sigma) f^s(\mathbf{s}^t,\mathbf{a}^{t,(t-1)}) b_i(\tilde{h}_i,\mathbf{a}^{t,(\hat{t}(h,\tilde{h}_i)-1)})\mathbf{1}\{h \in \mathscr{H}^t_i(\tilde{h}_i)\}, \tag{19}$$

and $\mu^t_i(\cdot|\tilde{h}_i,\mathbf{a}^{t,(t-1)}) : S^t_i \to [0,1]$ is the measure such that $\mu^t(s^t_i|\mathbf{a}^{t,(t-1)}) = \mu^t_i(\cdot|\tilde{h}_i,\mathbf{a}^{t,(t-1)}) \times \gamma_{h_i}(\tilde{h}_i|\mathbf{a}^{t,(\hat{t}-1)})$, where $\hat{t}(h,\tilde{h}_i)$ is the minimum $\check{t} \in \mathbb{N}$ such that $h_i(s^{t,(\check{t})}_i,\mathbf{a}^{t,(\check{t})}) = \tilde{h}_i$.

Proof. To establish L.12.1 notice that from the definition of $\alpha_{h_i}(\cdot|\sigma)$ we have

$$\alpha_{h_i}(\hat{H}_i|\sigma) = \sum_{t\in\mathbb{N}} \int_{(\omega^t,\mathbf{s}^t,\mathbf{a}^{t-1})\in\mathscr{H}^t_i(\hat{H}_i)} prob_{-i}(\mathbf{a}^{t-1}|\mathbf{s}^{t-1}_{-i},\sigma) f^s(\mathbf{s}^t_{-i},s^t_i,\mathbf{a}^{t-1}) d\hat{\mu}^{\omega}(\omega^t|\mathbf{s}^t,\tilde{h}_i,\mathbf{a}^{t-1}) d\mu_{-i}(\mathbf{s}^t_{-i}|\mathbf{a}^{t-1})\, d\mu^t_i(s^t_i|\mathbf{a}^{t-1}) \tag{20}$$

Notice first that the integral with respect to ω^t integrates to 1. Notice also that we can write

$$d\mu^t_i(s^t_i|\mathbf{a}^{t-1}) = d\gamma_i(s^t_i|\mathbf{a}^{t-1},\tilde{h}_i) \times d\gamma_{h_i}(\tilde{h}_i|\mathbf{a}^{t-1})$$

with $\gamma_i(s^t_i|\mathbf{a}^{t-1},\tilde{h}_i) = \mathbf{1}\{T^{-1}_i(\tilde{h}_i|\mathbf{a}^{t-1}) = s^t_i\}$. Replacing into (20), we obtain $\alpha_{h_i}(\hat{H}_i|\sigma) = \int_{\hat{H}_i} \upsilon(\tilde{h}_i,\sigma)\, d\beta_i(\tilde{h}_i)$.

An analogous argument shows L.12.2 by the definition of γ_{h_i} and $\mu^t_i(\cdot|\tilde{h}_i,\mathbf{a}^{t,(t-1)})$. □

Let σ be THPE and let σ^m be a sequence that converges to σ weakly in strategies, where σ^m is an $\tilde{\varepsilon}^m$-constrained equilibrium, with $\tilde{\varepsilon}^m$ an (ε^m,ν^m)-tremble profile such that $\varepsilon^m \to 0$. Let $\tilde{\mu}^m$ be the belief profile, weakly consistent with σ^m. By Lemma 12, $\tilde{\mu}^{m,t}_{-i}$ is given by equation (18).

Fix t and $\hat{t} \le t$. Define the function

$$\hat{\phi}^m_{t,\hat{t}}(\mathbf{a}^t,\mathbf{s}^t) = \frac{1}{\upsilon\left(h_i\left(s^{t,(\hat{t})}_i,\mathbf{a}^{t,(\hat{t})}\right),\sigma^m\right)} prob_{-i}(\mathbf{a}^t_{-i}|\mathbf{s}^t_{-i},\sigma^m)\cdot f^s(\mathbf{s}^t,\mathbf{a}^t), \tag{21}$$

with υ defined in L.12.1.

Lemma 13. *Under Assumption 5, for each t, player i, $P_{t,i} \in \mathcal{P}_{t,i}$, $\hat{t} \le t$ and $\mathbf{a}^t \in X^t$,*

L.13.1 *the family of functions*

$$\mathscr{F}_{\hat{\phi}} = \left\{ p^m : P_t \cap \check{S}_i^t(\mathbf{a}^t) \to L^1(\mathbf{S}_{-i}^t, \mu_{-i}(\cdot|\mathbf{a}^{t,(t-1)})) | p^m(s_i^t) = \hat{\phi}_{t,\hat{t}}^m(\mathbf{a}^t, s_i^t, \mathbf{s}_{-i}^t), m \in \mathbb{N} \right\}$$

is equicontinuous at every $s_i^t \in P_{t,i} \cap \check{S}_i^t(\mathbf{a}^t)$*, with* $\check{S}_i^t(\mathbf{a}^t) := \left\{ s_i^t \in S_i^t | \hat{b}_{i,\hat{t}}(s_i^t, \mathbf{a}^t) \neq 0 \right\}$*, and* $\hat{b}_{i,\hat{t}}(s_i^t, \mathbf{a}^t) := b_i\left(h_i\left(s_i^{t,(\hat{t})}, \mathbf{a}^{t,(\hat{t}-1)}\right), \mathbf{a}^{t,(\hat{t}-1)}\right)$.

L.13.2 *There is* $\bar{M} \in L^1(\mathbf{S}_i^t, \mu_i^t(\cdot|\mathbf{a}^{t,(t-1)}))$ *such that* $\left\| \hat{\phi}_{t,\hat{t}}^m(s_i^t, \cdot) \right\|_{L^1(\mathbf{S}_{-i}^t, \mu_{-i}(\cdot|\mathbf{a}^{t,(t-1)}))} < \bar{M}(s_i^t)$ *for every* $s_i^t \in P_t \cap \check{S}_i^t(\mathbf{a}^t)$.

Proof. In what follows $\|\cdot\|_1$ denotes $\|\cdot\|_{L^1(\mathbf{S}_{-i}^t, \mu_{-i}(\cdot|\mathbf{a}^{t-1}))}$. For $\tilde{t} \in \mathbb{N}$ and $\mathbf{a}^{\tilde{t}-1} \in X^{\tilde{t}-1}$, $\|\cdot\|_{\infty,\tilde{t},\mathbf{a}^{\tilde{t}-1}}$ denotes the norm in $L^\infty(\mathbf{S}_{-i}^{\tilde{t}}, \mu_{-i}^{\tilde{t}}(\cdot|\mathbf{a}^{\tilde{t}-1}))$.

The following Claim establishes L.13.1.

Claim 6. *Let* $\varepsilon > 0$*. For every* $\hat{s}_i^t \in P_t \cap \check{S}_i^t(\mathbf{a}^t)$ *there is* $\delta > 0$ *such that for* $\hat{s}_i'^t \in B(\hat{s}_i^t, \delta)$ $\left\| \hat{\phi}_{t,\hat{t}}^m(\hat{s}_i^t, \cdot) - \hat{\phi}_{t,\hat{t}}^m(\hat{s}_i'^t, \cdot) \right\|_1 \leq \varepsilon$, $\forall m \in \mathbb{N}$.

Proof of Claim 6.

For $\hat{s}_i^t \in P_t \cap \check{S}_i^t(\mathbf{a}^t)$, $\hat{t} \leq t$ and $\tilde{t} \in \mathbb{N}$, let

$$\bar{f}_{\hat{t}}^s(\hat{s}_i^t, \mathbf{s}_{-i}^{\tilde{t}}, \tilde{\mathbf{a}}^{\tilde{t}-1}) := \tilde{f}^s\left(\mathbf{s}_{-i}^{\tilde{t}}, h_i\left(\hat{s}_i^{t,(\hat{t})}, \tilde{\mathbf{a}}_i^{\tilde{t},(\hat{t}-1)}\right), \tilde{\mathbf{a}}^{\tilde{t}-1}\right).$$

Notice we can write

$$\hat{\phi}_{t,\hat{t}}^m(\mathbf{a}^t, \hat{s}_i^t, \mathbf{s}_{-i}^t) = \frac{prob_{-i}(\mathbf{a}_{-i}^t | \mathbf{s}_{-i}^t, \sigma^m) f^s(\hat{s}_i^t, \mathbf{s}_{-i}^t, \mathbf{a}^{t,(t-1)})}{\sum_{\tilde{t} \in \mathbb{N}, \tilde{\mathbf{a}}^{\tilde{t}-1} \in X^{\tilde{t}-1}} \int_{\tilde{\mathbf{s}}_{-i}^{\tilde{t}} \in \tilde{\mathbf{S}}_{-i}^{t}} prob_{-i}(\mathbf{a}^{\tilde{t}-1} | \tilde{\mathbf{s}}_{-i}^{t-1}, \sigma^m) \bar{f}_{\hat{t}}^s(\hat{s}_i^t, \tilde{\mathbf{s}}_{-i}^{\tilde{t}}, \tilde{\mathbf{a}}^{\tilde{t}-1}) \, d\mu_{-i}^{\tilde{t}}(\tilde{\mathbf{s}}_{-i}^{\tilde{t}} | \tilde{\mathbf{a}}^{\tilde{t}-1})}. \quad (22)$$

Let $\hat{\upsilon}(\hat{s}_i^t, \sigma^m)$ denote the denominator on the right hand side of (22) as a function of $\hat{s}_i^t$ and σ^m.

From Assumption 5, there is $R(\hat{s}_i^t, \mathbf{a}^{t,(t-1)}, \hat{t}) \in L^1(P_{t,i}, \mu_i^t(\cdot|\mathbf{a}^{t,(t-1)}))$ such that

$$\left| \frac{f^s(\hat{s}_i^t, \mathbf{s}_{-i}^t, \mathbf{a}^{t,(t-1)})}{f^s\left(\hat{s}_i^{t,(\hat{t})}, \mathbf{s}_{-i}^{t,(\hat{t})}, \mathbf{a}^{t,(\hat{t}-1)}\right) \hat{b}_{i,\hat{t}}(\hat{s}_i^t, \mathbf{a}^t)} \right| \leq R(\hat{s}_i^t, \mathbf{a}^{t,(t-1)}, \hat{t}), \quad (23)$$

for every $\hat{s}_i^t \in P_t \cap \check{S}_i^t(\mathbf{a}^t)$, $\mathbf{s}_{-i}^t \in \mathbf{S}_{-i}^t$ and $\hat{t} \leq t$. From (22) we have

$$\left\| \hat{\phi}_{t,\hat{t}}^m(\mathbf{a}^t, \hat{s}_i^t, \mathbf{s}_{-i}^t) \right\|_1 \leq \left\| \frac{prob_{-i}(a_{-i}^{\hat{t}} | s_{-i}^{\hat{t}}, \sigma^m) f^s(\hat{s}_i^{t,(\hat{t})}, \mathbf{s}_{-i}^{t,(\hat{t})}, \mathbf{a}^{t,(\hat{t}-1)})}{\hat{\upsilon}(\hat{s}_i^t, \sigma^m)} \cdot \frac{f^s(\hat{s}_i^t, \mathbf{s}_{-i}^t, \mathbf{a}^{t,(t-1)})}{f^s\left(\hat{s}_i^{t,(\hat{t})}, \mathbf{s}_{-i}^{t,(\hat{t})}, (\mathbf{a}^{t,(\hat{t}-1)}\right)} \right\|_1$$
$$\leq \left\| \frac{prob_{-i}(a_{-i}^{\hat{t}} | s_{-i}^{\hat{t}}, \sigma^m) f^s(\hat{s}_i^{t,(\hat{t})}, \mathbf{s}_{-i}^{t,(\hat{t})}, \mathbf{a}^{t,(\hat{t}-1)}) \hat{b}_{i,\hat{t}}(\hat{s}_i^t, \mathbf{a}^t)}{\hat{\upsilon}(\hat{s}_i^t, \sigma^m)} \right\|_1 \cdot R(\hat{s}_i^t, \mathbf{a}^{t,(t-1)}, \hat{t}) \leq R(\hat{s}_i^t, \mathbf{a}^{t,(t-1)}, \hat{t}), \tag{24}$$

where the third inequality follows the the fact that the numerator inside the norm is a term in the sum in $\hat{\upsilon}(\hat{s}_i^t, \sigma^m)$.

Let $t \in \mathbb{N}$ and $\hat{s}_i^t \in P_t \cap \check{S}_i^t(\mathbf{a}^t)$. From Assumption 5, for every $\nu > 0$ there is $\delta > 0$ such that for $\hat{s}_i'^t \in B(\hat{s}_i^t, \delta) \cap P_t \cap \check{S}_i^t(\mathbf{a}^t)$, we can write

$$\begin{aligned} f^s(\hat{s}_i'^t, \mathbf{s}_{-i}^t, \mathbf{a}^{t,(t-1)}) &= (1 + \varphi(\mathbf{s}_{-i}^t, \hat{s}_i'^t)) \cdot f^s(\hat{s}_i^t, \mathbf{s}_{-i}^t, \mathbf{a}^{t,(t-1)}) \\ \bar{f}_{\tilde{t}}^s(\hat{s}_i'^t, \mathbf{s}_{-i}^{\tilde{t}}, \mathbf{a}^{\tilde{t}-1}) &= (1 + \varphi_{\tilde{t}}'(\mathbf{s}_{-i}^{\tilde{t}}, \hat{s}_i'^t)) \cdot \bar{f}_{\tilde{t}}^s(\hat{s}_i^t, \mathbf{s}_{-i}^{\tilde{t}}, \mathbf{a}^{\tilde{t}-1}), \end{aligned} \tag{25}$$

where $\left\| \varphi(\mathbf{s}_{-i}^t, \hat{s}_i'^t) \right\|_{\infty, t, \mathbf{a}^{t-1}} \leq \nu$ and $\left\| \varphi_{\tilde{t}}'(\mathbf{s}_{-i}^{\tilde{t}}, \hat{s}_i'^t) \right\|_{\infty, \tilde{t}, \mathbf{a}^{\tilde{t}-1}} \leq \nu$ for every $\tilde{t} \in \mathbb{N}$. Let ν is small enough that $\frac{2\nu}{1-\nu} \cdot R(\hat{s}_i^t, \mathbf{a}^{t,(t-1)}, \hat{t}) \leq \varepsilon$.

$$\left\| \hat{\phi}_{t,\hat{t}}^m(\hat{s}_i'^t, \cdot) - \hat{\phi}_{t,\hat{t}}^m(\hat{s}_i^t, \cdot) \right\|_1 \leq \left\| \hat{\phi}_{t,\hat{t}}^m(\hat{s}_i^t, \cdot) \frac{1+\nu}{1-\nu} - \hat{\phi}_{t,\hat{t}}^m(\hat{s}_i^t, \cdot) \right\|_1 = \frac{2\nu}{1-\nu} \cdot \left\| \hat{\phi}_{t,\hat{t}}^m(\hat{s}_i^t, \cdot) \right\|_1 \leq \varepsilon,$$

where the first inequality follows from equations (22) and (25) and the last inequality follows from the choice of ν and equation (24). □

L.13.2 follows from equation (24) and $R(\hat{s}_i^t, \mathbf{a}^{t,(t-1)}, \hat{t}) \in L^1(P_{t,i}, \mu_i^t(\cdot | \mathbf{a}^{t,(t-1)}))$. □

For each $t \in \mathbb{N}$, $\hat{t} \leq t$ and $\mathbf{a}^t \in X^t$, we define the measure over $\mathbf{S}_{-i}^t$, $\hat{\eta}_{t,\hat{t}}^m(\cdot, \hat{s}_i^t, \mathbf{a}^t)$ as $d\hat{\eta}_{t,\hat{t}}^m(\mathbf{s}_{-i}^t, \hat{s}_i^t, \mathbf{a}^t) = \hat{\phi}_{t,\hat{t}}^m(\mathbf{a}^t, \hat{s}_i^t, \mathbf{s}_{-i}^t)\, d\mu_{-i}(\mathbf{s}_{-i}^t | \mathbf{a}^{t,(t-1)})$.

Lemma 13 implies that for $Y_2 = \mathbf{S}_{-i}^t$ and $Y_1 = S_i^t$, $\phi^m(\cdot) = \hat{\phi}_{t,\hat{t}}^m(\cdot, \mathbf{a}^t)$, conditions L.10.a and L.10.b hold. Therefore, by Lemma 10, there is a subsequence $\{m_r\}_r$ and a finitely additive measure $\hat{\eta}_{t,\hat{t}}^*$ such that $\hat{\eta}_{t,\hat{t}}^{m_r}(\cdot, \hat{s}_i^t, \mathbf{a}^t)$ converges to $\hat{\eta}_{t,\hat{t}}^*$ in the weak-$*$ topology of $L^\infty(\mathbf{S}_{-i}^t, \mu^t(\cdot | \mathbf{a}^t))$ for every $\hat{s}_i^t \in P_{t,i} \cap \check{S}_i^t(\mathbf{a}^t), P_{t,i} \in \mathcal{P}_{t,i}$.

Notice that we can write,

$$\tilde{\mu}_{-i}^t(\omega^t, \mathbf{a}^t, \mathbf{s}^t | \tilde{h}_i, \sigma^{m_r}) = \mu^\omega(\omega^t | \mathbf{s}^t, \mathbf{a}^{t,(t-1)}) \times \eta_{t,\hat{\hat{t}}}^{m_r}(\mathbf{s}^t, \mathbf{a}^t) \times b(\tilde{h}_i, \mathbf{a}^{t,(\hat{\hat{t}}-1)}) \cdot \mu_i^t(s_i^t | \tilde{h}_i, \mathbf{a}^{t,(t-1)}), \tag{26}$$

where $\hat{\hat{t}} := \hat{\hat{t}}(\mathbf{a}^t, \mathbf{s}^t, \tilde{h}_i)$ is the smallest $\tilde{t}$ such that $h_i(\mathbf{a}^t, \mathbf{s}^t) = \tilde{h}_i$.

We define, analogously,

$$\mu^{*,t}_{-i}(\omega^t, \mathbf{a}^t, \mathbf{s}^t|\tilde{h}_i) = \mu^{\omega}(\omega^t|\mathbf{s}^t, \mathbf{a}^{t,(t-1)}) \times \eta^*_{t,\hat{\hat{t}}}(\mathbf{s}^t, \mathbf{a}^t) \times b(\tilde{h}_i, \mathbf{a}^{t,(\hat{\hat{t}}-1)}) \cdot \mu^t_i(s^t_i|\tilde{h}_i, \mathbf{a}^{t,(t-1)}).$$

Lemma 14. *Under Assumption 5, $\tilde{\mu}(\cdot|\cdot, \sigma^{m_r})$ converges to μ^* in the weak-$*$ topology of $(L^{\infty})^*$.*

Proof. Let $f \in L^{\infty}(\mathscr{H}^t_i(\tilde{h}_i), \mu^t)$. Let us show that

$$\int f(h) d\tilde{\mu}^t_{-i}(h|\tilde{h}_i, \sigma^{m_r}) \to \int f(h) d\mu^{*,t}_{-i}(h|\tilde{h}_i), \tag{27}$$

almost surely in h_i. The integral over Ω^t via $\mu^{\omega}(|\mathbf{s}^t, \mathbf{a}^{t,(t-1)})$ in the definitions of $\tilde{\mu}$ and μ^* yields a function in $L^{\infty}(\mathbf{S}^t_i, \mu^t_{-i}(\cdot|\mathbf{a}^{t,(t-1)}))$ $\mu^t_i(\cdot|\mathbf{a}^{t,(t-1)})$-almost surely in s^t_i. Therefore, by Lemma 10, for every $\mathbf{a}^t \in X^t$, $P_{t,i} \in \mathcal{P}_{t,i}$[71] (omitting the dependence inside $\hat{\hat{t}}(\mathbf{a}^t, \mathbf{s}^t, \tilde{h}_i)$)

$$\int_{\mathbf{S}^t_{-i} \times \Omega^t} f(\omega^t, \mathbf{a}^t, \mathbf{s}^t_{-i}, s^t_i) d\mu^{\omega}(\omega^t|\mathbf{s}^t_{-i}, s^t_i, \mathbf{a}^{t,(t-1)})\, d\eta^{m_r}_{t,\hat{\hat{t}}}(\mathbf{s}^t_{-i}, s^t_i, \mathbf{a}^t) \to$$
$$\int_{\mathbf{S}^t_{-i} \times \Omega^t} f(\omega^t, \mathbf{a}^t, \mathbf{s}^t_{-i}, s^t_i) d\mu^{\omega}(\omega^t|\mathbf{s}^t_{-i}, s^t_i, \mathbf{a}^{t,(t-1)})\, d\eta^*_{t,\hat{\hat{t}}}(\mathbf{s}^t_{-i}, s^t_i, \mathbf{a}^t),$$

$\forall s^t_i \in P_{t,i} \cap \check{S}^t_i(\mathbf{a}^t)$, and hence, also $b(\tilde{h}_i, \mathbf{a}^{t,(\hat{t}(\mathbf{a}^t,\tilde{h}_i)-1)}) \cdot \mu^t_i(\cdot|\tilde{h}_i, \mathbf{a}^{t,(t-1)})$-almost surely, β_i-almost surely in $\tilde{h}_i$.

Thus, by L.13.2 and Lemma 11 (where $\eta_2(\cdot) = b(\tilde{h}_i, \mathbf{a}^{t,(\hat{t}(\mathbf{a}^t,\tilde{h}_i)-1)}) \cdot \mu^t_i(\cdot|\tilde{h}_i, \mathbf{a}^{t,(t-1)})$), equation (27) follows. □

Recall that σ is a THPE and σ^{m_r} is a sequence that converges to σ weakly in strategies, such that σ^{m_r} is an $\tilde{\varepsilon}^{m_r}$-constrained equilibrium, with $\tilde{\varepsilon}^{m_r}$ an $(\varepsilon^{m_r}, \nu^{m_r})$-tremble profile such that $\varepsilon^{m_r} \to 0$. $\tilde{\mu}^{m_r}$ is the system of beliefs that is weakly consistent with σ^{m_r}. We have shown that there is a system of beliefs μ^* such that $\tilde{\mu}^{m_r}$ converges to μ^* in $(L^{\infty})^*$. Then, by definition, (σ, μ^*) is quasi-consistent. Let us now show that σ is sequentially rational, and hence, that (σ, μ^*) is a weak sequential equilibrium.

By introduction, suppose that there is t, $\mathbf{a}^t \in X^t$ and $\hat{S}_i \subseteq S^t_i$, with $\mu^t_i(\hat{S}_i|\mathbf{a}^t) > 0$, and a

[71] Notice that the equicontinuity required by L.10.a holds, by L.13.1 after replacing $\hat{t}$ by $\hat{\hat{t}}(\mathbf{a}^t, \mathbf{s}^t, \tilde{h}_i)$ because $\star$ is an isolated point of the signal space, and therefore, $\hat{\hat{t}}$ is constant in a small enough vicinity of s^t_i.

player i strategy $\hat{\sigma}_i$ such that

$$U_i(\sigma_i|h_i(s_i^t,\mathbf{a}^t),\mu^*) < U_i(\hat{\sigma}_i|h_i(s_i^t,\mathbf{a}^t),\mu^*), \tag{28}$$

for $s_i^t \in \hat{S}_i$.

By weak consistency we have for every $\tilde{\sigma}_i^{m_r} \in \Sigma_i(\tilde{\varepsilon}^{m_r})$ and $\tilde{h}_i \in \mathscr{H}_i$

$$U_i(\tilde{\sigma}_i^{m_r}|\tilde{h}_i,\tilde{\mu}^{m_r}) = \sum_{l=0}^{\infty}\int_{h\in\mathscr{H}_i^l(\tilde{h}_i)} g_i(h)\cdot prob_i(h,\tilde{\sigma}_i^{m_r},\tilde{h}_i)\,d\tilde{\mu}_{-i}^{m_r,l}(h|\tilde{h}_i). \tag{29}$$

Let $\hat{\sigma}_i^m$ be a sequence of $\tilde{\varepsilon}^m$-constrained strategies that approaches $\hat{\sigma}_i$, $\hat{\mu}$-almost surely.

Since each σ^m is a $\tilde{\varepsilon}$-constrained equilibrium, by Proposition 3,

$$U_i(\sigma_i^{m_r}|h_i(s_i^t,\mathbf{a}^t),\tilde{\mu}^{m_r}) \geq U_i(\hat{\sigma}_i^{m_r}|h_i(s_i^t,\mathbf{a}^t),\tilde{\mu}^{m_r}),$$

for every $s_i^t \in \hat{S}_i$.

Define $\hat{H}_i(a_i^t,s_i^t,\sigma) = \{l \leq t|\sigma_i(a_{i,\tilde{l}}^t|h_i(a_i^{t,(\tilde{l})},s_i^{t,(\tilde{l})}) > 0, \forall \tilde{l} \in \{l,\ldots,t\}\}$. We define $:=$ $prob_i((a_i^t,s_i^t),\sigma_i^{m_r},\hat{H}_i(a_i^t,s_i^t,\sigma^{m_r}))$. Notice that $prob_i^{<,t}(a_i^t,s_i^t,\sigma^{m_r}) > 0$ almost surely in s_i^t.[72] Therefore,

$$\int_{s_i^t\in\hat{S}_i} prob_i^{<,t}(a_i^t,s_i^t,\sigma^{m_r})\cdot U_i(\sigma_i^{m_r}|h_i(s_i^t,\mathbf{a}^t),\tilde{\mu}^m)\,d\mu_i^t(s_i^t|\mathbf{a}^t) \geq \int_{s_i^t\in\hat{S}_i} prob_i^{<,t}(a_i^t,s_i^t,\sigma^{m_r})\cdot U_i(\hat{\sigma}_i^m|h_i(s_i^t,\mathbf{a}^t),\tilde{\mu}^m)\,d\mu_i^t(s_i^t|\mathbf{a}^t),$$

Replacing the definition of $U_i(\sigma_i^{m_r}|h_i(s_i^t,\mathbf{a}^t),\tilde{\mu}^{m_r})$ in the previous expression, we can re-write it as

$$\sum_{l\in\mathbb{N}}\int_{s_i^t\in\hat{S}_i}\int_{h\in\mathscr{H}_i^l(h_i(s_i^t,a_i^t))} g_i(h)\cdot prob_i^t(h,\sigma_i^{m_r})\,d\tilde{\mu}_{-i}^{m_r,l}(h|h_i(s_i^t,a_i^t))\,d\mu_i^t(s_i^t|\mathbf{a}^t) \geq \sum_{l\in\mathbb{N}}\int_{s_i^t\in\hat{S}_i}\int_{h\in\mathscr{H}_i^l(h_i(s_i^t,a_i^t))} g_i(h)\cdot prob_i^t(h,\hat{\hat{\sigma}}_i^{m_r})\,d\tilde{\mu}_{-i}^{m_r,l}(h|h_i(s_i^t,a_i^t))\,d\mu_i^t(s_i^t|\mathbf{a}^t), \tag{30}$$

where for strategy $\tilde{\sigma} \in \Sigma_i$, $prob_i^t(h,\tilde{\sigma}_i) := prob_i\left(h,\tilde{\sigma}_i,\hat{H}_i((a_i^t,s_i^t),\sigma_i)\cup\{t+1,\ldots,|h|\}\right)$ and $\hat{\hat{\sigma}}_i^{m_r}$ is the strategy that coincides with $\sigma_i^{m_r}$ at all private histories in $\mathscr{H}_i\setminus\{h_i(s_i^t,\mathbf{a}^t)|s_i^t \in$

[72] Notice that we are abusing notation as $prob_i()$ was previously defined over $h\in\mathscr{H}$. However, the new notation is also valid as its value depends only on i's actions and signals.

$\hat{S}_i\}$, and coincides with $\hat{\sigma}^{m_r}$ after private histories in $\{h_i(s_i^t, \mathbf{a}^t)|s_i^t \in \hat{S}_i\}$.

Passing to a subsequence if necessary, $prob_i^t(h, \sigma_i^{m_r})$ converges to $prob_i^t(h, \sigma_i)$ in the weak-$*$ topology of $L^\infty(\mathscr{H}, \hat{\mu})$ and $prob_i^t(h, \hat{\hat{\sigma}}^{m_r})$ converges to $prob_i^t(h, \hat{\hat{\sigma}}_i)$ in the same topology, where $\hat{\hat{\sigma}}_i$ is the strategy that coincides with σ_i at all private histories in $\mathscr{H}_i \setminus \{h_i(s_i^t, \mathbf{a}^t)|s_i^t \in \hat{S}_i\}$, and coincides with $\hat{\sigma}_i$ after private histories in $\{h_i(s_i^t, \mathbf{a}^t)|s_i^t \in \hat{S}_i\}$. These statements follow from the weak convergence in strategies of $\sigma_i^{m_r}$ to σ.

Now, the following Claim yields a contradiction of (28)

Claim 7. *Equation* (30) *implies*

$$\sum_{l\in\mathbb{N}} \int_{s_i^t \in \hat{S}_i} \int_{h \in \mathscr{H}_i^l(h_i(s_i^t, a_i^t))} g_i(h) \cdot prob_i^t(h, \sigma_i)\, d{\mu^*}_{-i}^l(h|h_i(s_i^t, a_i^t))\, d\mu_i^t(s_i^t|\mathbf{a}^t) \geq$$
$$\sum_{l\in\mathbb{N}} \int_{s_i^t \in \hat{S}_i} \int_{h \in \mathscr{H}_i^l(h_i(s_i^t, a_i^t))} g_i(h) \cdot prob_i^t(h, \hat{\hat{\sigma}}_i)\, d\tilde{\mu}_{-i}^{*,l}(h|h_i(s_i^t, a_i^t))\, d\mu_i^t(s_i^t|\mathbf{a}^t),$$

Proof of Claim 7.
Notice that by equation (26), the right hand side of (30) is equal to

$$\sum_{l\in\mathbb{N}} \int_{s_i^t \in \hat{S}_i} \int_{h \in \mathscr{H}_i^l(h_i(s_i^t, a_i^t))} g_i(h) \cdot prob_i^t(h, \sigma_i^{m_r})\, d\mu^\omega(\omega^l|\mathbf{s}^l, \mathbf{a}^{l,(l-1)})\, d\eta^{m_r}(\mathbf{s}^l, \mathbf{a}^l)\, d\mu_i^t(s_i^l|\mathbf{a}^{l,(l-1)}).$$

Now, for each l

$$G(\sigma^{m_r}) := \int_{s_i^t \in \hat{S}_i} \int_{h \in \mathscr{H}_i^l(h_i(s_i^t, a_i^t))} g_i(h) \cdot prob_i^t(h, \sigma_i^{m_r})\, d\mu^\omega(\omega^l|\mathbf{s}^l, \mathbf{a}^{l,(l-1)})\, d\eta^{m_r}(\mathbf{s}^l, \mathbf{a}^l)\, d\mu_i^t(s_i^l|\mathbf{a}^{l,(l-1)})$$
$$\to \int_{s_i^t \in \hat{S}_i} \int_{h \in \mathscr{H}_i^l(h_i(s_i^t, a_i^t))} g_i(h) \cdot prob_i^t(h, \sigma_i)\, d\mu^\omega(\omega^l|\mathbf{s}^l, \mathbf{a}^{l,(l-1)})\, d\eta^*(\mathbf{s}^l, \mathbf{a}^l)\, d\mu_i^t(s_i^l|\mathbf{a}^{l,(l-1)}).$$

The convergence follows from Lemmas 11 and 13 since the integral of g with respect to ω^t is bounded almost surely in $\mathbf{s}^l$ and that $\mu^l(\cup_{p\in\mathcal{P}_{t,i}} p) = 1$

Now, to see that the sum over l converges to the corresponding limit notice that for each l, $G(\sigma^{m_r}) \leq \max_{a^l \in X^l} \int_{\Omega^l \times \mathbf{S}^l} \bar{g}_i^l(\omega^l, \mathbf{a}^l, \mathbf{s}^l)\, d\mu^l(\omega^l, \mathbf{a}^l, \mathbf{s}^l)$ for every $r \in \mathbb{N}$, and therefore, the dominated convergence theorem applies. The convergence of the right hand side of equation (30) follows analogously. $\square$

PROOF OF PROPOSITION 4

Let $\Sigma_{i,n}(\tilde{\varepsilon})$ denote the set of player i's $\tilde{\varepsilon}$ constrained strategies in Γ_n and let $\Sigma_n = \times_{i\in N}\Sigma_{i,n}(\tilde{\varepsilon})$.

Let $\tilde{\varepsilon}$ be a (ε, ν)-tremble and let $\{\sigma_n\} \subseteq \Sigma_n(\tilde{\varepsilon})$ be a sequence of $\tilde{\varepsilon}$-constrained strategies. Let $\hat{\sigma}_n$ be the associated game-Γ strategy and suppose that $prob(\cdot|\hat{\sigma}_n) \to prob(\cdot|\sigma^*)$, $prob_i(\cdot|\hat{\sigma}_n) \to prob_i(\cdot|\sigma^*)$ and $prob_{-i}(\cdot|\hat{\sigma}_n) \to prob_{-i}(\cdot|\sigma^*)$ weakly in strategies for some $\sigma^* \in \Sigma$.

Define $g_{i,n}(\omega^t, \mathbf{a}^t) \coloneqq g_i(\tilde{\omega}^t, \mathbf{a}^t)$ for each $\omega^t \in P_n^{\Omega,t}(\tilde{\omega}^t)$ with $\tilde{\omega}^t \in \Omega_n^t$. Abusing notation we also write, for $h \in \mathscr{H}$, $g_{i,n}(h)$, for $g_{i,n}(\hat{\omega}^t(h), \hat{\mathbf{a}}^t(h))$.

Player i's payoff in game Γ_n from strategy profile σ_n can be written as a function of $\hat{\sigma}$ as,

$$U_{i,n}(\hat{\sigma}_n) \coloneqq \sum_{t \in \mathbb{N}} \int_{\mathscr{H}} g_{i,n}(h) prob(h|\hat{\sigma}_n) \, d\mu^t(h)$$

We now show that σ^* is a $\tilde{\varepsilon}$-constrained equilibrium. That is, let us show that

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

for every $\sigma_i' \in \Sigma_i(\tilde{\varepsilon})$. Fix $\sigma_i' \in \Sigma_i(\tilde{\varepsilon})$.

Lemma 15. *There is a sequence $\{\sigma_{i,n_k}'\}_{k \in \mathbb{N}} \subseteq \Sigma_{i,n_k}(\tilde{\varepsilon})$ and associated sequence of strategies in Γ, $\{\hat{\sigma}_{i,n_k}'\}_{k \in \mathbb{N}}$, such that $\left\| \hat{\sigma}_{i,n_k}' - \sigma_i' \right\|_2 \to 0$*

Proof of Lemma 15. We will show that we can construct a sequence of strategies, $\sigma_{i,n_k}' \in \Sigma_{i,n_k}(\tilde{\varepsilon})$, and associated sequence of strategies $\hat{\sigma}_{i,n_k}' \in \Sigma_i(\tilde{\varepsilon})$ such that for each n, $\left\| \hat{\sigma}_{i,n_k}' - \sigma_i' \right\|_2 < \frac{1}{k}$, which yields the desired result.

First, $\sigma_i' : \mathscr{H}_i \to \mathbb{R}^{X_i}$ is a measurable function in $L^2(\mathscr{H}_i, \hat{\mu}_i)$, where $\hat{\mu}_i = \hat{\mu} \circ h_i(\cdot)$. Therefore, σ_i' can be approximated by a sequence of simple functions. Let $\{\alpha_{i,k}\}_{k \in \mathbb{N}}$ be a sequence of simple functions that converges to σ_i' in the norm of L^2. We can assume that each $\alpha_{i,k}$ is a strategy. Each $\alpha_{i,k}$ is of the form

$$\alpha_{i,k} = \sum_{l \in I_k} c_k^l \cdot \mathbf{1}\{C_k^l\},$$

for some finite set of indices I_k, real constants $\{c_k^l\}_{l \in I_k}$ and open sets $\{C_k^l\}_{l \in I_k}$ with $C_k^l \subseteq \mathscr{H}_i$ for each $k \in \mathbb{N}$ and $l \in I_k$.[73]

[73] σ_i' can be approximated by sequences of simple functions with indicators over measurable sets. Due to the outer regularity of a measure over Borel sets of a metric space, there is also an approximating sequence of simple functions with indicators over open sets.

Let $h = (\omega^t, \mathbf{s}^t, \mathbf{a}^t) \in \mathscr{H}$ and define

$$P_n^{\mathscr{H}}(h) = \left\{\tilde{h} \in \mathscr{H} | \tilde{h} \in P_n^{\Omega,t}(\omega^t) \times P_n^{\mathbf{S},t}(\mathbf{s}^t) \times \{\mathbf{a}^t\}\right\}$$

and $P_n^{\mathscr{H}_i}(h_i) = h_i\left(P_n^{\mathscr{H}}(h)\right)$ for any h satisfying $h_i(h)$. $\mathscr{P}_n^{\mathscr{H}_i}$ denotes the partition of $\mathscr{H}_i$ comprised by the collection of sets $P_n^{\mathscr{H}_i}(h_i)$ for $h_i \in \mathscr{H}_i$.

Define

$$\alpha_{i,k,n}(h_i) = \begin{cases} \alpha_{i,k}(h_i) & P_n^{\mathscr{H}_i}(h_i) \subseteq C_l \text{ for some } l \in I_k \\ \frac{1}{|A_i(h_i)|} & \text{otherwise.} \end{cases}$$

Notice that $\alpha_{i,k,n}$ is constant in every cell of $\mathscr{P}_n^{\mathscr{H}_i}$. Therefore, it has a "natural" associated strategy in Γ_n.

Due to the L^2 convergence of $\alpha_{i,k}$ to σ_i', for each $k \in \mathbb{N}$, there is $k_1(k)$ such that $\left\|\alpha_{i,k_1(k)} - \sigma_i'\right\|_2 < \frac{1}{2k}$. Now, due to condition (c) in Definition 9, there is also $k_2(k)$ such that $\left\|\alpha_{i,k_1(k),k_2(k)} - \alpha_{i,k_1(k)}\right\|_2 < \frac{1}{2k}$. Defining the sequence of strategies $\hat{\sigma}_{i,n_k} := \alpha_{i,k_1(k),k_2(k)}$ concludes the proof of Lemma 15. □

Let $\{\hat{\sigma}'_{i,n_k}\}_{k\in\mathbb{N}} \subseteq \Sigma_i(\tilde{\varepsilon})$ be the sequence that exists from Lemma 15.

Since each σ_n is a $\Sigma_n(\tilde{\varepsilon})$-constrained equilibrium

$$U_{i,n_k}\left(\hat{\sigma}_{n_k}\right) \geq U_{i,n_k}\left(\hat{\sigma}'_{i,n_k}, \hat{\sigma}_{-i,n_k}\right) \tag{31}$$

Point (b) in Definition 9 implies $g_{i,n}$ converges to g_i almost surely. Thus, $U_{i,n_k}\left(\hat{\sigma}_{n_k}\right) \to U_i(\sigma^*)$ and $U_{i,n_k}\left(\hat{\sigma}'_{i,n_k}, \hat{\sigma}_{-i,n_k}\right) \to U_i(\sigma_i^*, \sigma_{-i}^*)$ which by equation (31) yields the desired conclusion.

PROOF OF PROPOSITION 5

Because the game is Markovian in payoffs and actions the correspondence $\mathbf{A}$ is only a function of the state. Also, $\mathbf{A}$ is upper hemicontinuous so for each $\tilde{C} \subseteq X$, the set $\Lambda(\tilde{C}) = \{\omega \in \Omega | \tilde{C} = \mathbf{A}(\omega)\}$ is closed and hence compact.[74]

Let $\varepsilon > 0$. Define the set $B^{\Omega}(\omega, \varepsilon) = \left\{\tilde{\omega} | \left|g_i(\omega, \mathbf{a}) - g_i(\tilde{\omega}, \mathbf{a})\right| < \varepsilon \forall i \in N, d(\omega, \tilde{\omega}) < \varepsilon\right\}$ where d is the distance in Ω. Define also, for each $i \in N$, $\varepsilon > 0$, $B^{S_i}(s_i, \varepsilon) = \left\{\tilde{s}_i \in S_i | d(s_i, \tilde{s}_i) < \varepsilon\right\}$.

[74]This follows directly from $\mathbf{A}$'s closed graph.

For every $\tilde{C} \in X$ and $\varepsilon > 0$, $\left\{B^{\Omega}(\omega,\varepsilon)\right\}_{\omega\in\Omega}$ is a an open cover of $\Lambda(\tilde{C})$. By compactness, there is an open subcover $\left\{B^{\Omega}(\omega_i^{\tilde{C}},\varepsilon)\right\}_{i\in\{1,\ldots,I(\tilde{C})}$ for a finite set of states $\{\omega_i^{\tilde{C}}\}_{i\in\{1,\ldots,I(\tilde{C})\}}$. We can define a measurable partition, $\{P_i(\tilde{C},\varepsilon)\}_{i\in\{1,\ldots,I(\tilde{C})\}}$, of $\Lambda(\tilde{C})$ recursively as follows. Define

$$P_1(\tilde{C},\varepsilon) = B^{\Omega}(\omega_1^{\tilde{C}},\varepsilon)\cap\Lambda(\tilde{C})$$

and for each $j \in \{2,\ldots,I(\tilde{C})\}$,

$$P_j(\tilde{C},\varepsilon) = B^{\Omega}(\omega_j^{\tilde{C}},\varepsilon)\cap\Lambda(\tilde{C})\setminus\bigcup_{l<j}P_l(\tilde{C},\varepsilon).$$

The collection of sets $\mathscr{P}(\varepsilon) = \{P_l(\tilde{C},\varepsilon)\}_{l\in\{1,\ldots,I(\tilde{C})\},\tilde{C}\in 2^X}$ is a partition of Ω. A measurable partition of S_i, $\mathscr{P}_i(\varepsilon)$, in which elements in a cell of partition are at a distance of at most ε is defined analogously.

Define the sequences of partitions $\{\mathscr{P}_n^{\Omega}\}_{n\in\mathbb{N}}$ where $\mathscr{P}_n^{\Omega} = \mathscr{P}(1/n)$ and $\{\mathscr{P}_n^{S_i}\}_{n\in\mathbb{N}}$ where $\mathscr{P}_n^{S_i} = \mathscr{P}(1/n)$ for each $i \in N$. The sets Ω_n and $S_{i,n}$ are any selection from the corresponding partitions.

Let us see that Γ has an approximating game sequence. Conditions (a) and (b) in Definition 9 are satisfied by construction. To see that (c) holds, let $B_1 \subseteq \Omega^t$ and $B_2 \subseteq \mathbf{S}^t$ be open sets, and $\mathbf{a}^t \in X^t$. Let $B_{1,n} = \{\omega^t|\omega^t \in P_1 \in \mathscr{P}_n^{\Omega,t}, P_1 \subseteq B_1\}$ and $B_{2,n} = \{\mathbf{s}^t|\mathbf{s}^t \in P_2 \in \mathscr{P}_n^{\mathbf{S},t}, P_2 \subseteq B_2\}$. We can write

$$\mu(B_1,B_2|\mathbf{a}^t) = \int \mathbf{1}\{B_1\}\mathbf{1}\{B_2\}\,d\mu(\omega^t,\mathbf{s}^t|\mathbf{a}^t), \tag{32}$$

where $\mathbf{1}\{B_1\}$ is an indicator function that yields 1 if $\omega^t \in B_1$ and zero otherwise, and $\mathbf{1}\{B_2\}$ is defined analogously. We also can write

$$\mu(B_{1,n},B_{2,n}|\mathbf{a}^t) = \int \mathbf{1}\{B_{1,n}\}\mathbf{1}\{B_{2,n}\}\,d\mu(\omega^t,\mathbf{s}^t|\mathbf{a}^t). \tag{33}$$

Let us show that

$$\mathbf{1}\{B_{1,n}\}\mathbf{1}\{B_{2,n}\} \to \mathbf{1}\{B_1\}\mathbf{1}\{B_2\} \tag{34}$$

almost surely. In fact, if $\omega^t \notin B_1$ or $\mathbf{s}^t \notin B_2$ then $\mathbf{1}\{B_1\}\mathbf{1}\{B_2\} = \mathbf{1}\{B_{1,n}\}\mathbf{1}\{B_{2,n}\} = 0$ for every $n \in \mathbb{N}$. Let $\omega^t \in B_1$ and $\mathbf{s}^t \in B_2$. For $\bar{\nu} > 0$ and $j \le t$, let $B\left((\omega^t)_j,\bar{\nu}\right)$ and let $B\left((\mathbf{s}^t)_j,\bar{\nu}\right)$

be the balls centered at $(\omega^t)_j$ and $(\mathbf{s}^t)_j$, respectively, of radius $\bar{\nu}$. Because B_1 and B_2 are open sets there is $\bar{\nu}$ small enough that $\times_{j\leq t} B\left((\omega^t)_j, \bar{\nu}\right) \subseteq B_1$ and $\times_{j\leq t} B\left((\mathbf{s}^t)_j, \bar{\nu}\right) \subseteq B_2$.

Let $\bar{n}$ be such that $\frac{1}{\bar{n}} \leq \frac{\bar{\nu}}{4}$. By the triangle inequality $P_n^{\Omega,t}\left(\omega^t\right) \subseteq \times_{j\leq t} B\left((\omega^t)_j, \bar{\nu}\right) \subseteq B_1$ and $P_n^{\mathbf{S},t}\left(\mathbf{s}^t\right) \subseteq \times_{j\leq t} B\left((\mathbf{s}^t)_j, \bar{\nu}\right) \subseteq B_2$ for every $n \geq \bar{n}$. Therefore, $\mathbf{1}\{B_1\}\mathbf{1}\{B_2\} = \mathbf{1}\{B_{1,n}\}\mathbf{1}\{B_{2,n}\} = 1$ for every $n \geq \bar{n}$.

By the Dominated Convergence Theorem, and equations (32) and (33), condition (c) in Definition 9 follows. $\square$

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*.

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.

BALDER, E. J. (1988): "Generalized equilibrium results for games with incomplete information," *Mathematics of Operations Research*, 13, 265–276.

BÖRGERS, T. (1991): "Upper hemicontinuity of the correspondence of subgame-perfect equilibrium outcomes," *Journal of Mathematical Economics*, 20, 89–106.

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

COTTER, K. D. (1991): "Correlated equilibrium in games with type-dependent strategies," *Journal of Economic Theory*, 54, 48–68.

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

DUNFORD, N. AND J. T. SCHWARTZ (1957): *Linear Operators. Part 1: General Theory*, New York Interscience.

FEDERGRUEN, A. (1978): "On n-person stochastic games by denumerable state space," *Advances in Applied Probability*, 10, 452–471.

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*.

GILLETTE, D. (1957): "Stochastic games with zero stop probabilities," *Contributions to the Theory of Games*, 3, 179–187.

HARRIS, C., P. RENY, AND A. ROBSON (1995): "The existence of subgame-perfect equilibrium in continuous games with almost perfect information: A case for public randomization," *Econometrica: Journal of the Econometric Society*, 507–544.

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

KAMADA, Y. AND M. KANDORI (2017): "Revision Games, Part II: Applications and Robustness," *Working Paper*.

——— (2019): "Revision Games," *Econometrica: Journal of the Econometric Society. Forthcoming.*

KAMADA, Y. AND T. SUGAYA (2020): "Optimal Timing of Policy Announcements in Dynamic Election Campaigns," *Quarterly Journal of Economics*, Forthcoming.

KAPOR, A. AND S. MORONI (2016): "Sniping in Proxy Auctions with Deadlines," .

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," .

MANELLI, A. M. (1996): "Cheap talk and sequential equilibria in signaling games," *Econometrica: Journal of the Econometric Society*, 917–942.

MILGROM, P. R. AND R. J. WEBER (1985): "Distributional strategies for games with incomplete information," *Mathematics of operations research*, 10, 619–632.

MYERSON, R. B. AND P. J. RENY (2019): "Perfect Conditional ε-Equilibria of Multi-Stage Games with Infinite Sets of Signals and Actions," *Econometrica: Journal of the Econometric Society*, Forthcoming.

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.

PARTHASARATHY, T. (1973): "Discounted, positive, and noncooperative stochastic games," *International Journal of Game Theory*, 2, 25–37.

POLLARD, D. (2002): *A user's guide to measure theoretic probability*, vol. 8, Cambridge University Press.

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*.

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.

SIMON, L. K. AND M. B. STINCHCOMBE (1989): "Extensive form games in continuous time: Pure strategies," *Econometrica: Journal of the Econometric Society*, 1171–1214.

SIMON, R. S. (2003): "Games of incomplete information, ergodic theory, and the measurability of equilibria," *Israel Journal of Mathematics*, 138, 73–92.

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

STINCHCOMBE, M. B. (2011): "Correlated equilibrium existence for infinite games with type-dependent strategies," *Journal of Economic Theory*, 146, 638–655.

TAKAHASHI, M. (1964): "Equilibrium points of stochastic non-cooperative \$ n \$-person games," *Journal of Science of the Hiroshima University, Series AI (Mathematics)*, 28, 95–99.

WHITT, W. (1980): "Representation and approximation of noncooperative sequential games," *SIAM Journal on Control and Optimization*, 18, 33–48.