Dynamic Games with Noisy Informational Asymmetries*

Francesco Fabbri† Sofia Moroni‡

July 30, 2024

Abstract

We define and show the existence of trembling hand perfect equilibrium and stationary Markov perfect equilibrium in infinite games with asymmetric and imperfect information. These results rely on the novel notion of sequential absolute continuity, which extends Milgrom and Weber's (1985) absolute continuity condition to dynamic games. Our approach establishes the existence of an equilibrium in a broad class of games with "noisy informational asymmetries," in which players' private information includes some idiosyncratic noise.

KEYWORDS: dynamic games, trembling hand perfect equilibrium, stationary Markov perfect equilibrium, absolutely continuous information.

JEL CODES: C72, C73.

*We are thankful to Pierpaolo Battigalli, Sylvain Chassang, Juan Escobar, Hezekiah Grayer, Faruk Gul, Adam Kapor, Alessandro Lizzeri, Benjamin Matta, Stephen Morris, Roger Myerson, Pietro Ortoleva, Wolfgang Pesendorfer, Luciano Pomatto, Giulio Principi, Phil Reny, Luca Rigotti, Larry Samuelson, Fedor Sandomirskiy, Ali Shourideh, Marciano Siniscalchi, Can Urgun, Richard Van Weelden, Mark Whitmeyer, and Leeat Yariv for their insightful comments and discussions. We are grateful to seminar participants at Princeton University, the University of Naples Federico II, the 34th Stony Brook Conference on Game Theory, the Cowles Foundation 2024 Conference on Economic Theory, and the 2024 North American Summer Meeting of the Econometric Society.

†Princeton University, Department of Economics, ffabbri@princeton.edu.

‡Princeton University, Department of Economics, smoroni@princeton.edu.

1 Introduction

We define and show the existence of trembling hand and stationary Markov perfect equilibrium in a broad class of infinite dynamic games with asymmetric and imperfect information. Our analysis reveals a key insight: in many games of economic interest, players' imprecise observations of their opponents' information—any amount of idiosyncratic noise suffices—ensures these equilibria exist.

We study a general class of dynamic games that encompasses but is not limited to multi-stage games. In our framework, a state of the world evolves stochastically based on previous states and actions, while players receive signals providing information about the history of the game, which they may not directly observe. The state and signal spaces are general measure spaces, while each player's action set is countable. Players' payoffs satisfy two properties: an upper bound on each player's expected payoff is finite, and payoffs exhibit a form of continuity in actions.

We begin by establishing the existence of constrained equilibria, where players select each available action with positive probability due to "trembles." Following Selten (1975), we define trembling hand perfect equilibria (THPE) as the limit of constrained equilibria when players' trembles vanish. This approach ensures sequential rationality, as we prove that constrained equilibria entail an optimal course of action conditional on players' beliefs at private histories. We establish that THPE exist and are subgame perfect. We then show that these existence results hold in a broad class of games where players' observations include some idiosyncratic noise. Next, we define a Markovian environment, and establish the existence of stationary Markov perfect equilibria (MPE), defined as trembling hand perfect equilibria in which each player's strategy depends solely on their current payoff-relevant information.

To illustrate the applicability of our results, consider the following class of games for which we establish the existence of THPE. The prior literature had not ascertained whether these games possess a sequentially rational equilibrium. At every period, a state of the world is drawn from some distribution over real vectors. The state is composed of individual dimensions for each player, each in turn consisting of a fundamental component and a noise term. Before moving, players observe the history of action profiles but do not directly observe the current state. Instead, each player receives a private signal consisting of their own fundamental component plus their noise term. The fundamental component may be correlated with the previous state and may depend on the previous action profile; the noise term is independent of the fundamental component, independent across periods, and is absolutely continuous with respect to the Lebesgue measure—for instance, it could follow a multivariate normal or uniform distribution. After receiving their signals, players simultaneously choose an action from a finite set, and the game proceed to the next period. Players have discounted payoffs that depend on both the current state and actions, and need not be continuous in the state.

The class of games introduced above can naturally represent various economic applications, which include, among others: (*i*) sequential oligopolistic competition where the actions are quantities or prices and signals are idiosyncratic demand shocks (e.g., Athey et al., 2004; Athey and Bagwell, 2008); (*ii*) sequential auctions with interdependent values, where the actions are bids and the signals are informative about the value that the auctioned good has for each potential buyer (e.g., Jofre-Bonet and Pesendorfer, 2003); (*iii*) global games of regime change where players attack a regime after learning about its strength (e.g., Angeletos et al., 2007).

Our results apply to a broader class of games. In particular, our analysis allows for: 1) general state and signal spaces; 2) state transitions and payoffs that can depend on the entire history of states and action profiles; 3) countable action sets, as long as signals evolve continuously in a specific sense that we describe below; 4) correlation among signals, provided that the joint distribution of signals is absolutely continuous with respect to the product of their marginal distributions; 5) payoffs that may be discontinuous with respect to the history of states, but must be continuous in actions; 6) non-discounted payoffs, provided that the sum across periods of an upper bound of the players' per-period expected payoffs is finite; 7) uncertainty about whether opponents moved in the past, accommodating an unknown order of moves.

The existence of THPE, established in Theorems 1 and 2, hinges on a key assumption we term *sequential absolute continuity* (SAC). This condition extends Milgrom and Weber's (1985) absolute continuity condition for one-period Bayesian games to dynamic settings. SAC imposes two requirements on the transition probability over the history of private signal profiles conditional on the action history. First, it must be *absolutely continuous* with respect to the product of the marginal probability measures of each player's history of private signals. Second, it must be *bounded and continuous* in the action history according to a novel norm on transition probabilities, which implies continuity in the total variation norm.

To establish the existence of THPE under noisy observations, we consider settings where each player's signals consist of a fundamental component, potentially common across players, and an idiosyncratic noise term. We provide sufficient conditions for SAC in this case. Proposition 3 shows that the absolute continuity condition of SAC holds under three assumptions: (*i*) the joint distribution of the fundamental and noise components is absolutely continuous with respect to the product of their marginals; (*ii*) the joint distribution of the players' noise terms is absolutely continuous with respect to the product of their individual marginals; (*iii*) players cannot perfectly infer the fundamental component from their observed signals. Additionally, Proposition 4 shows that the boundedness and continuity requirements of SAC are met under two conditions. First, the fundamental signal must exhibit a weak form of continuity.[1] Second, the joint distribution of the private and fundamental signal histories must be absolutely continuous with respect to the product of the marginal of the fundamental signal history and a distribution over the private signal history that is independent of the action history.

A consequence of Propositions 3 and 4, Corollary 1, is that for real-valued signals, the addition of noise that is absolutely continuous with respect to the Lebesgue measure, along with certain regularity requirements, guarantees the existence of a THPE. Notably, this result implies that the addition of even an arbitrarily small amount of such noise suffices to ensure existence in games where the previous literature has otherwise demonstrated non-existence. Through two running examples, we illustrate how adding noise simultaneously resolves two issues: strategic entanglement, which generates discontinuity of the expected payoffs in players' strategies, precluding classical existence arguments; in dynamic settings, the discontinuity of signals with respect to previous actions, as seen in non-existence examples like that of Harris et al. (1995). The conditions of Corollary 1 are met by commonly used noise distributions, including independent uniform and jointly normal, making our results widely applicable across various economic models.

Our approach can be extended to the analysis of equilibria in Markov strategies. To achieve this, we adapt the notion of stationary Markov perfect equilibrium by defining them as THPE, in which each player's strategy depends only on the current payoff-relevant information. To formalize the meaning of payoff-relevant information,

[1]Specifically, Proposition 4 requires continuity with respect to the weak convergence of probability measures. See footnote 14 for a formal definition.

we introduce *Markov games*. In Markov games, the primitives, such as state and signal transitions, payoffs, available actions, and active players, are determined solely by variables specific to the current period. Moreover, the set of states of the world can be decomposed into two dimensions: a *payoff-relevant* dimension, which determines each player's payoffs and the state transition, and a *payoff-irrelevant* dimension. Private signals that provide information about the payoff-relevant state dimension are referred to as payoff-relevant signals.

For stationary Markov strategies to be optimal, we need to ensure that *bygones are bygones*, meaning that players do not find it beneficial to exploit information from previous periods. We formalize this property by requiring that when all players follow stationary Markov strategies, each player's belief about the current payoff-relevant state of the world depends only on the current payoff-relevant signal component. We call this assumption *Markov information*. We complement this condition with the *Markov payoff* assumption, which imposes that payoff-relevant signals are as informative as private histories for computing the expected payoffs.

Theorem 3 establishes the existence of stationary Markov equilibria under a new assumption which we term Markov absolute continuity (MAC). The latter replaces SAC in Markov games, modifying it in two ways. First, it applies only to the transition of payoff-relevant signal profiles. Second, it requires that the absolute continuity condition holds with respect to the product of the marginal measures over private signals, taken not only across players but also across periods. Thus, MAC is neither weaker nor stronger than SAC.

Some applications for which our analysis yields novel existence results in Markov games are: (1) games with asynchronous moves, including asynchronous revision games (Kamada and Kandori, 2020), and dynamic cheap talk games (Renault et al., 2013); (2) stochastic games in which players receive both a public and a private shock (Balbus et al., 2013), such as in dynamic oligopolies, where the public and private shocks can be interpreted as demand and firm's costs, respectively.

Related literature. Despite their widespread applications in economics, establishing equilibrium existence in infinite games with asymmetric information has proven challenging. Even in one-period Bayesian games with finite actions, the literature has provided examples that lack equilibria. Simon (2003) presents an example of one such game that has no Bayes-Nash equilibrium and, similarly, Hellman (2014) and Simon and Tomkowicz (2018) construct examples that lack even approximate equilibria.

In the presence of uncountable states, equilibrium existence in dynamic games is established either by "closing" the strategy space with some form of correlation or imposing restrictions on the state transition. In games with almost perfect information,[2] Harris et al. (1995) constructs a two-period game featuring compact action spaces that lacks a subgame perfect equilibrium and restores existence by adding a stage-wise public signal that serves as a correlating device. He and Sun (2020) extends this existence result by requiring the state transition to be atomless. Manelli (1996) adds cheap talk to a signaling game to obtain existence.

This paper establishes existence by following a different approach. We build upon the seminal works of Milgrom and Weber (1985) and Balder (1988) by generalizing their absolute continuity condition to dynamic settings. Furthermore, we provide a sufficient condition for absolute continuity based on noisy signals, broadening their framework's applicability even in static cases.

We contribute to the literature studying sequentially rational equilibria in infinite dynamic games. The closest work is Myerson and Reny (2020), which introduces the concept of perfect conditional ε-equilibria, defined as strategy profiles that can be approximated by a net of conditional ε-equilibria.[3] These equilibria eventually assign positive probability to every possible action and almost every move of nature. Their limiting distributions, as ε tends to zero, are termed perfect conditional equilibrium distributions. However, these distributions may not always be induced by a strategy profile. In contrast, we focus on trembling hand perfect equilibria, which we define as limits of ε-constrained equilibria. These trembling hand perfect equilibria are conditional-$\bar{\varepsilon}$ equilibria, with $\bar{\varepsilon}$ vanishing as ε does.[4]

A vast literature on stochastic games stemming from the seminal work of Shapley (1953) studies the existence of stationary Markov perfect equilibria in settings where players observe the history of play and the current state, and have discounted payoffs. In this environment, existence generally requires continuity assumptions on the transition of the state across periods.[5] Levy (2013), and subsequently Levy and

[2]In games with *almost perfect information*, players may move simultaneously after perfectly observing the history of the game.

[3]In a conditional ε-equilibrium, players optimize their payoffs up to ε utils, conditional on every positive measure set of private histories.

[4]See Proposition 6 in Supplemental Appendix B.9.

[5]Duggan (2012) provides an excellent discussion of these assumptions. Nowak and Raghavan (1992), Duffie et al. (1994), and Duggan (2012) establish the existence of stationary Markov perfect equilibria under restrictions on the state transition and other conditions such as payoff-irrelevant

McLennan (2015), shows that non-existence of stationary Markov perfect equilibria may arise even in standard stochastic games with finite action sets where the state transition is absolutely continuous with respect to a fixed measure.

We also intersect the literature on stochastic games with asymmetric information, which spans from the study of folk theorems to dynamic persuasion (Aumann et al., 1995; Ely, 2017). Altman et al. (2008) studies the existence of a stationary Nash equilibrium when each player privately observes the realizations of an associated controlled Markov chain. Balbus et al. (2013) establishes the existence of a stationary Markov perfect equilibrium in a setting with strategic complementarities in the presence of public and private shocks. To the best of our knowledge, we are the first to provide conditions for the existence of trembling hand and stationary Markov perfect equilibria under general state and signal spaces.

2 Model

We study dynamic games played in countably many periods $t \in \mathbb{N} := \{1, 2, \dots\}$. Formally, any such game is represented by the following list of objects:

$$\Gamma = (N, \Omega, \mathcal{M}, (X_i, S_i, A_i, g_i)_{i \in N}, \mu, \gamma),$$

where

- - $N := \{1, \dots, n\}$ is a finite set of n players.
  - Ω is a measurable set of states of the world, and S_i is a measurable set of private signals for each player $i \in N$. $S := \prod_{i \in N} S_i$ denotes the set of signal profiles.
  - X_i is a countable, compact metric space, endowed with its Borel σ-algebra, representing the action set of each player $i \in N$. The set of action profiles is $X := \prod_{i \in N} X_i$, with generic element $a = (a_1, \dots, a_n)$. The set of histories of action profiles up to period $t \in \mathbb{N}$ is $X^t := \prod_{\ell \leq t} X$, with generic element $a^t = (a_1, \dots, a_t)$. For $\ell \leq t$, the element $a^{t,(\ell)} \in X^\ell$ denotes the truncation of history a^t up to and including period ℓ. The action $a^t_{i,\ell}$ corresponds to player i's move in period-ℓ action profile a^t_ℓ. The sets Ω^t, S^t, and their corresponding elements, such as $\omega^{t,(\ell)}$ or $s^t_{i,\ell}$, are defined analogously.[6]

---

and payoff-relevant noise. Parthasarathy and Sinha (1989), and Nowak (2003) assume stronger restrictions than Levy (2013) to prove existence. He and Sun (2017) unifies these results by assuming the "decomposable coarser transition kernel" condition on the state transition.

[6] For $t \geq 1$, we use the notation $Y^t := \prod_{\ell=1}^{t} Y$ for any set Y; $Y^0 := \{\emptyset\}$.

- $\mathcal{M}(\cdot)$ maps each measurable set to its σ-algebra. For instance, $\mathcal{M}(\Omega)$ denotes the σ-algebra of measurable subsets of Ω. We endow product spaces with their product σ-algebra, subsets of measurable spaces with their relative σ-algebra,[7] and we assume all singleton sets are measurable.

- The correspondence $A_i : \cup_{t\in\mathbb{N}} S_i^t \times X_i^{t-1} \rightrightarrows X_i$ is non-empty closed-valued, weakly measurable,[8] and specifies the *actions available* to player $i \in N$ as a function of i's private signal and action history.

- The function $g_i : \cup_{t\in\mathbb{N}} \Omega^t \times X^t \to \mathbb{R}$ represents the *flow payoff* received by player $i \in N$ as a function of the history of the states of the world and action profiles. It is measurable and bounded.

- The map $\mu : \cup_{t\in\mathbb{N}\cup\{0\}} \Omega^t \times X^t \to \Delta(\Omega)$ is the *state transition probability*[9] which determines the probability of a new state as a function of the history of the states of the world and action profiles. That is, $\mu(Z|\omega^t, a^t)$ is the probability that the period $t+1$-state belongs to the set $Z \in \mathcal{M}(\Omega)$ given $(\omega^t, a^t) \in \Omega^t \times X^t$.

- The *signal transition function* $\gamma : \cup_{t\in\mathbb{N}} \Omega^t \times X^{t-1} \to S$ is measurable and determines the private signal profiles as a function of the history of the states of the world and action profiles. Denote by $\gamma_i : \cup_{t\in\mathbb{N}} \Omega^t \times X^{t-1} \to S_i$ the projection of γ onto i's private signal in S_i.

Following the representation of Myerson and Reny (2020), we assume that the state of the world evolves stochastically, while signals are deterministic functions of the history of the game. As the state space is general, this representation is equivalent to an alternative framework in which the signal realization is also stochastic.

For ease of exposition, in the main body of the paper, we focus on games where players are informed about whether their opponents have moved in the past, i.e., they observe the period $t \in \mathbb{N}$ in which the game is being played. In an extension, we model games where the order of moves may be unknown. This is achieved by restricting the information available in private histories, which do not record any information from inactive periods. See Section 6 for more details.

[7] For every $Z \in \mathcal{M}(Y)$, *Z's relative σ-algebra* is $\mathcal{M}(Z) = \{B \cap Z | B \in \mathcal{M}(Y)\}$.

[8] A correspondence $\phi : Z \rightrightarrows Y$ that maps a measurable space Z to a topological space Y is *weakly measurable* if, for every closed subset $B \subseteq Y$, the set $\{z \in Z : \phi(z) \subseteq B\}$ is measurable. See Definition 18.1 in Aliprantis and Border (2006) and the ensuing discussion.

[9] For measurable spaces Z and Y, the function $\xi : Z \to \Delta(Y)$ is a *transition probability* if $\xi(B|z)$ is measurable in z for every $B \in \mathcal{M}(Y)$; $\Delta(Y)$ is the set of probability measures over Y. A *transition measure* is defined analogously when $\xi : Z \to \mathcal{M}(Y)$ and $\xi(\cdot|z)$ is a measure for every $z \in Z$.

2.1 Histories, Strategies, and Expected Payoffs

Histories. For $t \in \mathbb{N}$, a *period-t history* $h = (\omega^t, a^{t-1})$ is composed of a *history of the states of the world* $\omega^t \in \Omega^t$, and a *history of action profiles* $a^{t-1} \in X^{t-1}$. The set $\mathcal{H}^t$ contains every period-t history, and $\mathcal{H} := \cup_{t\in\mathbb{N}\cup\{0\}}\mathcal{H}^t$ contains every history.

Private Histories. For $i \in N$, $t \in \mathbb{N}$, $(s_i^t, a_i^{t-1}) \in S_i^t \times X_i^{t-1}$ is a *period-t private history* of player i. Denote by $\mathcal{H}_i^t := S_i^t \times X_i^{t-1}$ the set of period-t private histories of player i, and by $\mathcal{H}_i := \cup_{t\in\mathbb{N}}\mathcal{H}_i^t$ the set of all private histories of player i. The set $\mathcal{H}_{A_i} := \{(s_i^t, a_i^{t-1}) \in \mathcal{H}_i |\, a_{i,\ell}^{t-1} \in A_i(s_i^{t,(\ell)}, a_i^{t-1,(\ell-1)}),\, \ell \le t-1\}$ denotes player i's *available private histories* according to the action correspondence A_i.

Strategies. A *strategy* of player $i \in N$ is a transition probability $\sigma_i : \mathcal{H}_{A_i} \to \Delta(X_i)$ that maps i's available private histories to probability distributions over i's actions satisfying $\operatorname{supp}\sigma_i(h_i) \subseteq A_i(h_i)$ for every $h_i \in \mathcal{H}_{A_i}$. Denote by Σ_i the set of player i strategies,[10] and by $\sigma = (\sigma_j)_{j\in N} \in \prod_j \Sigma_j := \Sigma$ a strategy profile.

Conditional measures. For $i \in N$, conditional on player i's actions $a_i^\tau \in X_i^\tau$ and signals $s_i^t \in S_i^t$, for $t, \tau \in \mathbb{N}\cup\{0\}$ with $\tau \le t$, the strategy $\sigma_i \in \Sigma_i$ *induces a transition probability* over player i's action history $a_i^t \in X_i^t$ as follows

$$p_i(a_i^t|s_i^t, a_i^\tau, \sigma_i) := \prod_{\ell=\tau+1,\dots,t} \sigma_i\big(a_{i,\ell}^t|s_i^{t,(\ell)}, a_i^{t,(\ell-1)}\big) \tag{1}$$

if $(s_i^t, a_i^{t,(t-1)}) \in \mathcal{H}_{A_i}$, $a_i^{t,(\tau)} = a_i^\tau$; zero otherwise.[11] That is, $p_i(a_i^t|s_i^t, a_i^\tau, \sigma_i)$ is the probability of i's actions in history a_i^t that are a continuation of a_i^τ if s_i^t is realized and i plays according to σ_i. The strategy profile $\sigma \in \Sigma$ induces $p(a^t|s^t, a^\tau, \sigma) := \prod_{j\in N} p_j(a_j^t|s_j^t, a_j^\tau, \sigma_j)$ for $(s^t, a^t) \in S^t \times X^t$, $a^\tau \in X^\tau$. When $\tau = 0$, hence $a^\tau = \emptyset$, we write $p_i(a_i^t|s_i^t, \sigma_i) := p_i(a_i^t|s_i^t, \emptyset, \sigma_i)$ and $p(a^t|s^t, \sigma) := p(a^t|s^t, \emptyset, \sigma)$.

Conditional on a history of action profiles $a^{t-1} \in X^{t-1}$ and states of the world $\omega^\tau \in \Omega^\tau$, for $t, \tau \in \mathbb{N}\cup\{0\}$ with $\tau \le t$, the probability measure over Ω^t induced by the state transition probability μ is

$$d\mu_\omega^t(\omega^t|\omega^\tau, a^{t-1}) := \prod_{\ell=\tau}^{t-1} d\mu(\omega_{\ell+1}^t|\omega^{t,(\ell)}, a^{t-1,(\ell)})$$

[10] Lemma 4 shows that the set of strategies is non-empty. That is, there exists a measurable selector, with support on the available actions set, as a function of $h_i \in \mathcal{H}_i$. This result follows from the weak measurability of A_i.

[11] We use the convention that $\prod_{t=\tau}^{\tau-1} y_t = 1$ for $(y_t)_{t\in\mathbb{N}\cup\{0\}}$ sequence in $\mathbb{R}$, and $\tau \in \mathbb{N}$.

if $\omega^{t,(\tau)} = \omega^\tau$; zero otherwise. When $\tau = 0$, hence $\omega^\tau = \emptyset$, we write $d\mu^t_\omega(\omega^t|a^{t-1}) := d\mu^t_\omega(\omega^t|\emptyset, a^{t-1})$.

For every $t \in \mathbb{N}$, $(\omega^t, a^{t-1}) \in \mathscr{H}^t$, denote the history of realized signals up to period t by $\gamma^t(\omega^t, a^{t-1}) := (\gamma(\omega^{t,(1)}), \gamma(\omega^{t,(2)}, a^{t-1,(1)}), \cdots, \gamma(\omega^t, a^{t-1}))$. To ease notation, we write $p(a^t|\omega^t, \sigma)$ and $p(a^t|\omega^t, a^\tau, \sigma)$, instead of $p(a^t|\gamma^t(\omega^t, a^{t,(t-1)}), \sigma)$ and $p(a^t|\gamma^t(\omega^t, a^{t,(t-1)}), a^\tau, \sigma)$, respectively.

For $a^{t-1} \in X^{t-1}$, the measure $\mu^t_\omega(\cdot|a^{t-1})$ and the function γ induce a probability measure over private signal profiles, $\mu^t_s(B|a^{t-1}) := \mu^t_\omega(\{\omega^t : \gamma^t(\omega^t, a^{t-1}) \in B\}|a^{t-1})$, for $B \in \mathcal{M}(S^t)$. Let $\mu^t_{s_i}(\cdot|a^{t-1})$ be the marginal of μ^t_s on player i's signals in S^t_i.

Finally, for every $t \in \mathbb{N}$, we assume throughout the existence of a transition probability $\mu^t_{s_i|s_i} : S^{t-1}_i \times X^{t-1} \to \Delta(S_i)$ satisfying $d\mu^t_{s_i}(s^t_i|a^{t-1}) = d\mu^t_{s_i|s_i}(s^t_{i,t}|s^{t,(t-1)}_i, a^{t-1}) \times d\mu^{t-1}_{s_i}(s^{t,(t-1)}_i|a^{t-1,(t-2)})$.

Expected payoffs. For every $t \in \mathbb{N}$, a strategy profile $\sigma \in \Sigma$, a signal history function γ^t, and the state transition probability μ induce a probability $P^t(\cdot|\sigma)$ over $\Omega^t \times X^t$. Player i's *expected payoff* is $U_i(\sigma) := \sum_{t\in\mathbb{N}} U_{i,t}(\sigma)$, where $U_{i,t}(\sigma) := \int_{\Omega^t \times X^t} g_i(\omega^t, a^t)\, dP^t(\omega^t, a^t|\sigma)$.

Similarly, for $\tau \in \mathbb{N}$, we define player i's *continuation expected payoff* from $\sigma \in \Sigma$ after history $(\omega^\tau, a^{\tau-1}) \in \mathscr{H}^\tau$ as $U_i(\sigma|\omega^\tau, a^{\tau-1}) := \sum_{t\geq\tau} U_{i,t}(\sigma|\omega^\tau, a^{\tau-1})$, where $U_{i,t}(\sigma|\omega^\tau, a^{\tau-1}) := \sum_{a^t\in X^t} \int_{\Omega^t} g_i(\omega^t, a^t)\, dP^t(\omega^t, a^t|\sigma, \omega^\tau, a^{\tau-1})$ and $P^t(\cdot|\sigma, \omega^\tau, a^{\tau-1})$ is a probability over $\Omega^t \times X^t$ induced by σ and $(\omega^\tau, a^{\tau-1})$.

We impose two requirements on players' expected payoffs throughout our analysis. First, we assume that a bound on the sum of per-period expected payoffs is finite. Formally, for each $i \in N$,

$$\sum_{t\in\mathbb{N}} \sup_{\sigma\in\Sigma} \left|U_{i,t}(\sigma)\right| < \infty. \qquad \text{(boundedness)}$$

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

$$\sum_{t\in\mathbb{N}} \sup_{(\omega^t,a^t)} |g_i(\omega^t, a^t)| < \infty. \qquad (2)$$

Notice that (2) holds in games with discounted payoffs, i.e., $g_i(\omega^t, a^t) = \delta^t \cdot u_i(\omega^t, a^t)$, where $\delta \in (0, 1)$ and $u_i(\omega^t, a^t)$ is bounded for $i \in N$, $(\omega^t, a^t) \in \Omega^t \times X^t$. Moreover, (2) is also satisfied in games where payoffs decrease in t slower than geometrically, e.g., $g_i(\omega^t, a^t) = \frac{1}{t^2} \cdot u_i(\omega^t, a^t)$. In general, our boundedness condition is weaker than (2) as it holds in a class of games with stochastic move opportunities that end in finite time with probability 1 and have bounded expected length, while (2) does not (see Lemma 3).[12]

We also require expected payoffs conditional on past signals to be continuous in past actions, which holds immediately when the action set, X, is finite. If X is infinite, assume there is a transition probability $\mu^t_{\omega|s} : S^t \times X^{t-1} \to \Delta(\Omega^t)$ such that $d\mu^t_\omega(\omega^t|a^{t-1}) = d\mu^t_{\omega|s}(\omega^t|s^t, a^{t-1}) \times d\mu^t_s(s^t|a^{t-1})$, for $t \in \mathbb{N}$. For every $i \in N$, $t \in \mathbb{N}$, the function $\hat{g}_{i,t} : S^t \times X^t \to \mathbb{R}$ defined as

$$\hat{g}_{i,t}(s^t, a^t) := \int_{\Omega^t} g_i(\omega^t, a^t)\, d\mu^t_{\omega|s}(\omega^t|s^t, a^{t,(t-1)}) \qquad \text{(continuity)}$$

is continuous in a^t for every s^t in a μ^t_s-full measure set.[13] Notice that we require payoff continuity with respect to the history of action profiles, not signals or states. This assumption is weaker than continuity in the action history when the state transition is modeled as Nature's moves, which implicitly assumes continuity in the state as well. This latter form of continuity is also present in the literature; it is assumed by Harris et al. (1995), He and Sun (2020), and Myerson and Reny (2020), among others.

2.2 An Application

We preview some of our findings by formalizing an interesting class of games for which we establish the existence of a sequentially rational equilibrium. Existence in this class follows from Corollary 1 in Section 3.3.

Application 1 (Dynamic games with Lebesgue signals)**.** Each state of the world $\omega \in \Omega$ can be written as $\omega = (\hat{s}, \epsilon)$:

- $\hat{s} = (\hat{s}_i)_{i \in N}$, where $\hat{s}_i \in \mathbb{R}^{\ell_i}$, $\ell_i \in \mathbb{N}$, represents the *fundamental signal component* of player $i \in N$;
- $\epsilon = (\epsilon_i)_{i \in N}$, where $\epsilon_i \in \mathbb{R}^{\ell_i}$ represents the *idiosyncratic noise term* of player $i \in N$.

Each player $i \in N$ observes a private signal $\gamma_i(\hat{s}, \epsilon) = \hat{s}_i + \epsilon_i$, which combines their fundamental component and noise term. Furthermore, each $\hat{s}_i$ and ϵ_i may or may not affect payoffs.

Assume that: (a) for every $t \in \mathbb{N}$, there exists a continuous and bounded density $f^t(\hat{s}^t, \epsilon^t_1, \ldots, \epsilon^t_n, a^{t-1})$ such that the joint distribution of $\hat{s}^t$ and ϵ^t, conditional on the

[12] Inspired by Fudenberg and Levine (1983), one can define *continuity at infinity* in our setting by requiring $\sup_{\sigma,\sigma'} |\sum_{t=\tau}^{\infty} U_{i,t}(\sigma) - U_{i,t}(\sigma')| \to 0$ as $\tau \to \infty$. Our boundedness condition is stronger as it implies $\sum_{t=\tau}^{\infty} \sup_{\sigma,\sigma'} |U_{i,t}(\sigma) - U_{i,t}(\sigma')| \to 0$ as $\tau \to \infty$.

[13] For every $t \in \mathbb{N}$, $B \in \mathcal{M}(S^t)$ is a μ^t_s*-full measure set* if $\mu^t_s(B|a^{t-1}) = 1$ for every $a^{t-1} \in X^{t-1}$.

action profile history $a^{t-1} \in X^{t-1}$, can be written as

$$d\mu^t_{\hat{s},\epsilon}(\hat{s}^t, \epsilon^t_1, \ldots, \epsilon^t_n | a^{t-1}) = f^t(\hat{s}^t, \epsilon^t_1, \ldots, \epsilon^t_n, a^{t-1})\, d\mu^t_{\hat{s}}(\hat{s}^t | a^{t-1}) \times d\lambda^t(\epsilon^t)$$

where $\mu^t_{\hat{s}}$ denotes the marginals with respect to $\hat{s}^t$ and λ^t denotes the Lebesgue measure over real vectors, respectively; (b) for every $i \in N$, and private signal $s^t_i \in S^t_i$, the set of fundamental components which, combined with some noise term, can yield s^t_i has positive measure; (c) for every $t \in \mathbb{N}$, $d\mu^t_{\hat{s}}(\cdot | a^{t-1})$ is continuous in the action profile history $a^{t-1} \in X^{t-1}$ in the topology of *weak convergence of probability measures*.[14] ◂

The class of games included in Application 1 encompasses a wide range of economic applications. It can model dynamic oligopolistic competition (Athey et al., 2004; Athey and Bagwell, 2008), where firms repeatedly interact in a market, setting quantities or prices based on idiosyncratic demand signals. The framework also applies to sequential auctions with interdependent values (Jofre-Bonet and Pesendorfer, 2003; Aoyagi, 2003; Skrzypacz and Hopenhayn, 2004), where buyers submit bids informed by signals about their value for the auctioned good. Moreover, it extends to contexts such as currency attacks (Morris and Shin, 1998), where strategic decisions are informed by private signals about fundamental values, and global games of regime change (Angeletos et al., 2007), where players decide to attack a regime based on signals about its evolving strength. Notice that the noise structures in these applications conform to our assumptions, typically featuring additive independent normal or uniform noise terms.

2.3 Examples: Payoff Discontinuity and Non-existence

Without additional requirements on signal transitions, equilibrium existence is not guaranteed. The following two examples illustrate potential payoff discontinuities in infinite games with asymmetric information, emphasizing how these discontinuities can lead to non-existence. The first example demonstrates that players' payoffs may be discontinuous in their strategies, a phenomenon known as *strategic entanglement*.[15]

[14] A sequence of probability measures over Y, $(\rho_n)_{n\in\mathbb{N}}$, $\rho_n \in \Delta(Y)$, *converges weakly* to $\rho \in \Delta(Y)$ if $\lim_{n\to+\infty} \int f d\rho_n = \int f d\rho$ for every bounded, continuous function $f : Y \to \mathbb{R}$.

[15] This issue was identified by Simon and Stinchcombe (1989), Börgers (1991), and Harris et al. (1995), among others. The term "strategic entanglement" was coined by Myerson and Reny (2020). Example 1 is analogous to Example 2 in Milgrom and Weber (1985), Example 2.1 in Cotter (1991), and Example 2.1 in Stinchcombe (2011).

This discontinuity prevents the application of standard fixed-point arguments to prove existence. The second example adapts the non-existence case by Harris et al. (1995) to our framework with countable actions. We revisit these examples later to provide intuition for our results.

Example 1 (Strategic entanglement)**.** Consider a two-player game where both players observe a public signal s, drawn from a uniform in $[0,1]$, before they choose an action A or B. The following sequence of strategies generates a payoff discontinuity. For each player $i \in \{1,2\}$, and $n \in \mathbb{N}$, define $\sigma_i^n(A|s) = 1$ if $s \in [(k-1)/2^n, k/2^n]$ for odd k, and $\sigma_i^n(A|s) = 0$ otherwise. Players choose the same action as a function of the signal, but as n increases, players switch actions over progressively finer intervals. Figure 1 illustrates the first three strategies in this sequence.

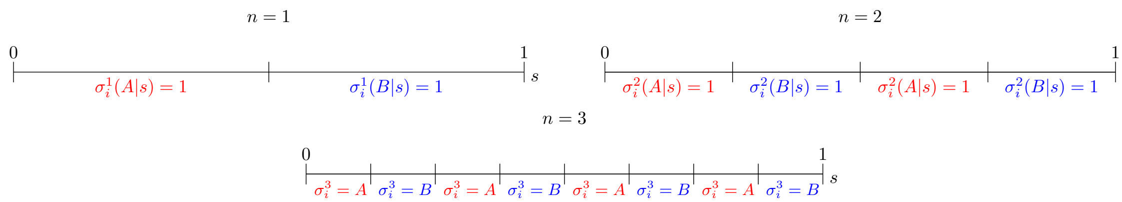

Figure 1: Sequence of strategies generating strategic entanglement.

Each player i's limit strategy σ_i^* must choose A and B with equal probability, *independently* of s.[16] However, the limit of the probability over action profiles $\sigma_1^n(\cdot|s)\cdot\sigma_2^n(\cdot|s)$ yields a perfectly correlated distribution: (A,A) and (B,B) each occur with probability $1/2$, independent of s. Since $\lim_{n\to\infty} \sigma_1^n(a_1|s)\cdot\sigma_2^n(a_2|s) \neq \sigma_1^*(a_1|s)\cdot\sigma_2^*(a_2|s)$ for $(a_1,a_2) \in \{A,B\}^2$ the players' expected payoffs, which depend on the probability over action profiles, are discontinuous with respect to their strategies. Consequently, the best response correspondence may fail to have a closed graph, invalidating the conditions required for Nash's classical fixed-point argument. ◀

Strategic entanglement has substantive implications for equilibrium existence. For instance, Simon (2003), Hellman (2014), and Simon and Tomkowicz (2018) construct one-period Bayesian games with finite actions that lack any Bayes-Nash equilibrium.

While strategic entanglement can occur in static settings, the following example illustrates a distinct form of discontinuity that emerges in multi-period games, potentially leading to equilibrium non-existence.

[16] Formally, σ_i^* is the weak limit in the space of probability measures.

Example 2.1 (Harris, Reny, and Robson (1995)). Consider the following game. In the first period, player A chooses action $a \in \mathcal{A} := \cup_{n\in\mathbb{N}}\{-1/n, 1/n\} \cup \{0\}$ and player B action $b \in \{L, R\}$, while in the second period, after observing the moves previously occurred, players C and D choose $c \in \{L, R\}$ and $d \in \{L, R\}$, respectively. Players C and D's payoff functions are identical and depend only on action a: Playing L yields a payoff of $-a$ and R of a. That is, the second period players strictly prefer to play L if $a < 0$, R if $a > 0$, and are indifferent otherwise. Player B wants to guess the future choice of player C and gets a payoff of $\mathbb{1}_{\{c=L\}} - \mathbb{1}_{\{c=R\}}$ if $b = L$ and of $2 \cdot \left(\mathbb{1}_{\{c=R\}} - \mathbb{1}_{\{c=L\}}\right)$ if $b = R$. Player A's payoff is as follows

$$-|a| \cdot \mathbb{1}_{\{b=c\}} + |a| \cdot \mathbb{1}_{\{b\neq c\}} - 10 \cdot \mathbb{1}_{\{c\neq d\}} - \tfrac{1}{2}|a|^2.$$

If B and C make the same choice, A obtains a payoff of $-|a|$, and $|a|$ otherwise; if C and D make different choices, A obtains a negative payoff of -10; A gets $-\frac{1}{2}|a|^2$.[17]

As argued by Harris et al. (1995), this game does not possess any subgame perfect equilibrium. In our case, A's action set is countable rather than uncountable, but their argument for non-existence applies unaltered. Intuitively, A would like to minimize the probability that B guesses correctly while ensuring that C and D are perfectly coordinated. This can be achieved if A randomizes uniformly between a positive and a negative number, say $a = \delta$ with probability $\frac{1}{2}$ and $a = -\delta$ with probability $\frac{1}{2}$. However, for any $\delta \in \mathcal{A}$ with $\delta \neq 0$, A incurs the cost $\frac{1}{2}|\delta|^2 > 0$, which approximates zero as $\delta \to 0$. In the limit, A's mixed strategy becomes degenerate, which does not allow C and D to coordinate their actions in a random way.[18] ◂

3 Equilibrium Existence

We now introduce sequential absolute continuity (SAC), the main assumption of our analysis. This condition restricts the transition of private signal profiles by requiring: (a) absolute continuity with respect to the product of the marginal measures of each player's private signals; (b) boundedness and continuity in the action profile history according to a novel norm, which implies continuity in the total variation norm.

Let Z be a countable set and Y a measurable space. The *strong total variation*

[17] Notice that the game falls within our framework by assuming there is no payoff-relevant state, i.e., Ω is a singleton, and second-period players receive a perfectly informative signal about the moves that previously occurred, i.e., $s_C = s_D = (a, b)$.

[18] See Supplemental Appendix B.7 for a detailed explanation.

norm of a transition measure $\xi : Z \to \mathcal{M}(Y)$ is

$$\|\xi\|_{SV} := \sup\left\{ \sum_{j\in I} |\xi(Y_j|z_j)| \;\Big|\, \{Y_j\}_{j\in I} \in \pi(Y), \{z_j\}_{j\in I} \subseteq Z \right\},$$

where $\pi(Y)$ denotes the set of finite measurable partitions of Y. For a subset $\check{Z} \subseteq Z$, let $\xi|_{\check{Z}}$ be the restriction of ξ to $\check{Z}$, and, for an element $\check{z} \in Z$, let $\xi^{\check{z}}(\cdot|z) := \xi(\cdot|\check{z})$ for every $z \in Z$. We say that ξ is *continuous in Z in the strong total variation norm* if, for every $z^* \in Z$ and $\varepsilon > 0$, there is $\delta > 0$ such that $\left\|(\xi - \xi^{z^*})|_{B(z^*,\delta)}\right\|_{SV} < \varepsilon$.

Boundedness and continuity in the strong total variation norm guarantees that whenever the transition probability admits a density with respect to some z-independent measure, such a density is also bounded by a z-independent, L^1 function on Y and continuous in Z (see Proposition 5). A density satisfying these conditions is called a *Carathéodory integrand*.[19]

Notice that ξ is *continuous in the total variation norm* if $\left\|\xi^z - \xi^{z^*}\right\|_{SV} \to 0$ as $z \to z^*$, i.e., as z approaches z^* it cannot vary across partition elements. Therefore, this form of continuity is weaker than continuity in the strong total variation norm.[20]

Assumption (Sequential absolute continuity). *The following holds:*

(*a*) *For every $t \in \mathbb{N}$, $a^{t-1} \in X^{t-1}$, $\mu_s^t(\cdot|a^{t-1})$ is absolutely continuous with respect to the product measure $\prod_{i\in N} \mu_{s_i}^t(\cdot|a^{t-1})$;*

(*b*) *For every $t \in \mathbb{N}$, μ_s^t is bounded, and continuous in X^{t-1} in the strong total variation norm.*

SAC extends the absolute continuity condition of Milgrom and Weber (1985) to dynamic games. In their static setting, signals correspond to types, and the absolute continuity condition requires the joint type distribution to be absolutely continuous with respect to the product of players' marginal type distributions, coinciding with SAC(a). SAC(b) holds trivially in one-period games as $X^0 = \emptyset$.

In a dynamic environment, sequential absolute continuity imposes two key requirements: (a) Absolute continuity must hold conditional on every possible history of play. This means that the joint distribution of players' signals, *given any sequence of past actions*, can be expressed in terms of the product of individual players' signal distributions; (b) The distribution of signals must be continuous with respect to

[19] See footnote 36 for the formal definition of a Carathéodory integrand.

[20] Example 4 in Supplemental Appendix C.1 constructs a norm that is continuous in total variation but discontinuous in strong total variation.

past play and bounded in the strong total variation norm. This condition implies continuity in both the total variation norm and set-wise continuity.[21] Consequently, any game violating these weaker forms of continuity necessarily violates continuity in the strong total variation norm. The latter is the case in Example 2.1, where the distribution of signals fails to be set-wise continuous (see Example 2.2).

SAC is always satisfied if the action and signal spaces are not too large. For instance, SAC(a) holds if the set of signals S_i is countable for all but at most one player, while SAC(b) holds if, for every player $i \in N$, the action set X_i is finite.

In general, even with finite action sets, SAC restricts the information structure of the game. For instance, when the state is drawn from a non-atomic distribution on the reals, players cannot commonly observe the state without violating SAC(a). Example 1 illustrates this violation: $\mu_s^1(D) = 1 \neq 0 = \mu_{s_1}^1 \times \mu_{s_2}^1(D)$, where $D = \{(s_1, s_2) \in [0,1]^2 | s_1 = s_2\}$ is the diagonal set, and each μ_{s_i} is uniform on $[0,1]$, $i \in \{1,2\}$.

3.1 Constrained Equilibrium

We define a constrained equilibrium that must (i) put a positive weight on each available action, and (ii) be optimal within the set of constrained strategies given the constrained strategies of the opponents.

For every $\varepsilon > 0$, a measurable function $\tilde{\varepsilon}_i : \{(h_i, a_i) | h_i \in \mathscr{H}_i, a_i \in X_i\} \to (0,1)$ is an *ε-tremble* of player $i \in N$ if $\sum_{a_i \in A_i(h_i)} \tilde{\varepsilon}_i(h_i, a_i) < \varepsilon$ for every $h_i \in \mathscr{H}_i$, and $\tilde{\varepsilon}_i((s_i^t, a_i^{t-1}), a_{i,t})$ is continuous in $(a_i^{t-1}, a_{i,t}) \in (X_i)^t$ for every $s_i^t \in (S_i)^t$. We denote by $\tilde{\varepsilon} = (\tilde{\varepsilon}_i)_{i \in N}$ an ε-tremble profile, by $\mathcal{E}(\varepsilon)$ the set of ε-tremble profiles, and by $\mathcal{E} := \cup_{\varepsilon > 0} \mathcal{E}(\varepsilon)$ the set of ε-tremble profiles for any positive ε.

Given $\tilde{\varepsilon} \in \mathcal{E}$, a strategy profile is *$\tilde{\varepsilon}$-constrained* if, at each private history $h_i \in \mathscr{H}_{A_i}$, each player $i \in N$ puts at least $\tilde{\varepsilon}_i(h_i, a_i)$ weight on each action $a_i \in A_i(h_i)$. We denote by $\Sigma_i(\tilde{\varepsilon})$ the set of player i's $\tilde{\varepsilon}$-constrained strategies, and by $\Sigma(\tilde{\varepsilon})$ the set of $\tilde{\varepsilon}$-constrained strategy profiles.

Definition 1. Let $\tilde{\varepsilon} \in \mathcal{E}$. An $\tilde{\varepsilon}$-constrained strategy profile $\sigma \in \Sigma(\tilde{\varepsilon})$ is an *$\tilde{\varepsilon}$-constrained equilibrium* if, for every player $i \in N$,

$$U_i(\sigma) \geq U_i(\sigma_i', \sigma_{-i}) \qquad \forall \sigma_i' \in \Sigma_i(\tilde{\varepsilon}).$$

[21] Formally, set-wise continuity requires that, for each $\tilde{S} \in \mathcal{M}(S^t)$ and each sequence $a^m \to a$, we have $\mu_s^t(\tilde{S} \mid a^m) \to \mu_s^t(\tilde{S} \mid a)$.

Theorem 1. *Let Γ be a dynamic game that satisfies sequential absolute continuity. For every $\tilde{\varepsilon} \in \mathcal{E}$, Γ has an $\tilde{\varepsilon}$-constrained equilibrium.*

The proof of our main theorem builds upon Balder's (1988) argument, but is technically involved and consists of the following key steps. First, we identify each player i's constrained strategy, σ_i, with its induced transition probabilities over action histories, $p_i(\cdot|\cdot, \sigma_i)$. We endow each player's strategy space, now represented as transition probabilities, with the coarsest topology that makes every expected payoff functional continuous in transition probabilities—a topology Balder (1988) terms the *weak topology.* Equipped with the product topology, the strategy space is both convex and compact, and the best response correspondence is convex-valued.

Next, we demonstrate that each player's expected payoff is continuous in strategies, a result implied by sequential absolute continuity. For every $i \in N$, $t \in \mathbb{N}$, $a^{t-1} \in X^{t-1}$, SAC(a) allows us to write

$$d\mu_s^t(s_1^t, \ldots, s_n^t | a^{t-1}) = f(s_1^t, \ldots, s_n^t, a^{t-1})\, d\nu_1(s_1^t) \times \ldots \times d\nu_n(s_n^t), \tag{3}$$

where each $\nu_i \in \Delta(S_i^t)$ is a probability measure, and f is a density function. Proposition 5 establishes that, whenever $\mu_s^t(\cdot|a^{t-1})$ can be written as in equation (3), SAC(b) holds if and only if the density f is a Carathéodory integrand. This implies that, if $\mu_s^t(\cdot|a^{t-1})$ has a density that is a Carathéodory integrand with respect to one product measure, densities with respect to any other product measure will also be Carathéodory integrands.[22] Therefore, equation (3) and our continuity condition imply that players' expected payoffs can be written as the integral of Carathéodory integrands over a product measure over players' signals. Theorem 2.5 in Balder (1988) then implies the continuity of expected payoffs, ensuring the non-emptiness and closed-graph properties of the best-response correspondence. Hence, SAC prevents the strategic entanglement of Example 1.

Notice that we define the topology over transition probabilities on action histories, rather than period-by-period behavioral strategies, to capture potential strategic entanglement of each player's strategy across periods. If the topology were defined over behavioral strategies, continuity would not follow.

With these properties established, we invoke the Kakutani-Fan-Glicksberg fixed

[22] To the best of our knowledge, we are the first to characterize the continuity properties of its supporting densities via continuity with respect to a norm. Our results show that continuity with respect to the total variation norm is not sufficient to ensure continuity of the density.

point theorem[23] to prove the existence of a fixed point of the best response correspondence over transition probabilities. Finally, we recover the equilibrium strategy from this fixed point by computing the probability of a sequence of actions up to a given period and dividing it by the probability of the sequence up to the previous period.

We next investigate the sequential rationality properties of constrained equilibria. In particular, we show that every constrained equilibrium σ prescribes an optimal course of action for each player, conditional on any set of private histories.

For every player $i \in N$, and period $t \in \mathbb{N}$, define the function $\hat{h}_i^t : \mathcal{H}^t \to \mathcal{H}_i^t$ as $\hat{h}_i^t(\omega^t, a^{t-1}) := (\gamma_i^t(\omega^t, a^{t-1}), a_i^{t-1})$. This function projects the full history $(\omega^t, a^{t-1}) \in \mathcal{H}^t$ onto player i's corresponding private history. For every $\sigma \in \Sigma$, $Z \in \mathcal{M}(\mathcal{H}_i^t)$, the probability that player i's private history belongs to Z under σ is $P_i(Z|\sigma) := P^t((\hat{h}_i^t)^{-1}(Z)|\sigma)$, where abusing notation $P^t(\cdot|\sigma)$ denotes the probability measure over $\mathcal{H}^t$ induced by σ. Whenever $P_i(Z|\sigma) > 0$, we define player i's period-t *beliefs* over $C \in \mathcal{M}(\mathcal{H}^t)$ *conditional* on Z as

$$P_i^t(C|Z,\sigma) = \frac{P^t\big(C \cap (\hat{h}_i^t)^{-1}(Z)|\sigma\big)}{P_i(Z|\sigma)}.$$

Finally, for every pair of strategy profiles $\sigma, \hat{\sigma} \in \Sigma$, and $Z \in \mathcal{M}(\mathcal{H}_i^t)$, $t \in \mathbb{N}$, player i's *expected payoff conditional on Z from σ and $\hat{\sigma}$* is defined as

$$U_i(\hat{\sigma}|Z,\sigma) := \int_{\mathcal{H}^t} U_i(\hat{\sigma}|\omega^t, a^{t-1})\, dP_i^t(\omega^t, a^{t-1}|Z,\sigma),$$

where σ and $\hat{\sigma}$ determine the conditional belief and the expected payoff, respectively.[24]

Given an ε-tremble, an $\tilde{\varepsilon}$-constrained strategy profile is an *$\tilde{\varepsilon}$-constrained conditional equilibrium* if it tests each player's rationality at every set of private histories occurring with positive probability.

Definition 2. A strategy profile $\sigma \in \Sigma(\tilde{\varepsilon})$, for $\tilde{\varepsilon} \in \mathcal{E}$, is an *$\tilde{\varepsilon}$-constrained conditional equilibrium*, if, for every $i \in N$, $t \in \mathbb{N}$ and $Z \in \mathcal{M}(\mathcal{H}_i^t)$ satisfying $P_i(Z|\sigma) > 0$,

$$U_i(\sigma|Z,\sigma) \geq U_i(\sigma_i', \sigma_{-i}|Z,\sigma) \qquad \forall \sigma_i' \in \Sigma_i(\tilde{\varepsilon}).$$

Proposition 1. *If $\tilde{\varepsilon} \in \mathcal{E}$, and σ is an $\tilde{\varepsilon}$-constrained equilibrium, then σ is an $\tilde{\varepsilon}$-constrained conditional equilibrium.*

[23] Corollary 17.55 in Aliprantis and Border (2006).

[24] Alternatively, we can define $\hat{\sigma}$ as a continuation strategy of σ after private histories in Z.

3.2 Trembling Hand Perfect Equilibrium

Inspired by Selten's (1975) seminal work on finite games, we define trembling hand perfect equilibrium as the limit of constrained strategies as their ε-trembles vanish. Consider a sequence $(\varepsilon_n)_{n\in\mathbb{N}}$ in $\mathbb{R}_+$ converging to zero, and a corresponding sequence of $\tilde{\varepsilon}_n$-constrained equilibria σ^n, where $\tilde{\varepsilon}_n \in \mathcal{E}(\varepsilon_n)$. Theorem 1 guarantees the existence of such sequences. As with constrained strategies, where we use transition probabilities to obtain continuity of payoffs, we need to construct the limit strategy from the limit of transition probabilities. However, this approach encounters difficulties as $\varepsilon_n \to 0$. While constructing strategies from transition probabilities is straightforward for ε-constrained strategies, a challenge arises when the limit strategy may assign zero probability to some action histories. In such cases, the previous construction becomes undefined along these paths.

To address this challenge, we introduce non-vanishing transition probabilities that preserve information about strategies even after zero-probability histories. These probabilities are analogous to those used in the ε-constrained case but are well-defined in the limit. We then infer the limit strategy by taking the limits of these non-vanishing probabilities, allowing us to characterize the trembling hand perfect equilibrium even in cases where some histories become infinitely unlikely.

For every $i \in N$, $t \in \mathbb{N}$, let α_i^t be a probability measure over S_i^t. Consider a sequence of transition probabilities $(\lambda^n)_{n\in\mathbb{N}}$, where $\lambda^n : S_i^t \to \Delta(X_i^t)$. We say that $(\lambda^n)_{n\in\mathbb{N}}$ *converges to* λ^* *in the weak topology of* $(S_i^t \times X_i^t, \alpha_i^t)$ if for every Carathédory integrand ϕ_i on $(S_i^t \times X_i^t, \alpha_i^t)$, the following convergence holds

$$\int_{S_i^t} \sum_{a_i^t \in X_i^t} \phi_i(s_i^t, a_i^t)\lambda^n(a_i^t|s_i^t)\, d\alpha_i^t(s_i^t) \to \int_{S_i^t} \sum_{a_i^t \in X_i^t} \phi_i(s_i^t, a_i^t)\lambda^*(a_i^t|s_i^t)\, d\alpha_i^t(s_i^t).$$

For every $i \in N$, $t \in \mathbb{N}$, define the function a_i over $(s_i^t, a_i^{t,(t-1)}) \in \mathcal{H}_{A_i}$ for any given $\sigma_i \in \Sigma_i$ as follows

$$a_i(s_i^t, a_i^t; \sigma_i) \coloneqq \left\{a_i^{t,(\tau)} | \tau = \min\left\{\hat{\tau} \le t-1 | \Pi_{\ell \ge \hat{\tau}+1}^{t-1} \sigma_i(a_{i,\ell}^t | s_i^{t,(\ell)}, a_i^{t,(\ell-1)}) > 0\right\}\right\}.$$

This function truncates a_i^t up to period τ, which is the latest period before $t-1$ where player i's action $a_{i,\tau}^t$ has zero probability under strategy σ_i, as a function of i's signals. In particular, if $\sigma_i(a_{i,\tilde{\tau}}^t \big| s_i^{t,(\tilde{\tau})}, a_i^{t,(\tilde{\tau}-1)}) > 0$, for every $\tilde{\tau} \le t-1$, then $\tau = 0$, and if $\sigma_i(a_{i,\tilde{\tau}}^t \big| s_i^{t,(\tilde{\tau})}, a_i^{t,(\tilde{\tau}-1)}) = 0$, for every $\tilde{\tau} \le t-1$, then $\tau = t-1$

For every $i \in N$, $t \in \mathbb{N}$, we define a *reference measure for player* i as $\nu_i^t \in \Delta(S_i^t)$ satisfying:

- For each $a^{t-1} \in X^{t-1}$, $\mu^t_{s_i}(\cdot|a^{t-1})$ is absolutely continuous with respect to ν^t_i;
- There exists a transition probability $\nu_{i,t} : S^{t-1}_i \to \Delta(S_i)$ such that $d\nu^t_i(s^t_i) = d\nu_{i,t}(s^t_{i,t}|s^{t,(t-1)}_i) \times d\nu^{t-1}_i(s^{t,(t-1)}_i)$.

A reference measure satisfying these conditions can always be constructed in our environment.[25]

Definition 3. A strategy $\sigma^* \in \Sigma$ is a *trembling hand perfect equilibrium* if:

(*i*) There exist sequences $(\varepsilon_n)_{n\in\mathbb{N}}$ in $\mathbb{R}_+$, $(\tilde{\varepsilon}_n)_{n\in\mathbb{N}}$ in $\mathscr{E}$, and $(\sigma^n)_{n\in\mathbb{N}}$ in Σ satisfying:

- $\tilde{\varepsilon}_n \in \mathscr{E}(\varepsilon_n)$ for each $n \in \mathbb{N}$,
- σ^n is a $\tilde{\varepsilon}_n$-constrained equilibrium,
- $\lim_{n\to\infty} \varepsilon_n = 0$;

(*ii*) For every $i \in N$, and $\tau, t \in \mathbb{N}$ with $\tau \leq t$, there exist a transition probability $p^*_i(\cdot|\cdot, a^\tau_i) : S^t_i \to \Delta(X^t_i)$, for $a^\tau_i \in X^\tau_i$, and a reference measure ν^t_i, such that:

- $p_i(\cdot|\cdot, a^\tau_i, \sigma^n_i)$ converges to $p^*_i(\cdot|\cdot, a^\tau_i)$ in the weak topology of $(S^t_i \times X^t_i, \nu^t_i)$,
- For $(s^t_i, a^{t,(t-1)}_i) \in \mathscr{H}_{A_i}$ and $a^{t,(\tau)}_i = a_i(s^t_i, a^t_i; \sigma_i)$:

$$\sigma^*_i\left(a^t_{i,t}|s^t_i, a^{t,(t-1)}_i\right) = \frac{p^*_i\left(a^t_i|s^t_i, a^{t,(\tau)}_i\right)}{p^*_i\left(a^{t,(t-1)}_i|s^{t,(t-1)}_i, a^{t,(\tau)}_i\right)} \tag{4}$$

This definition characterizes a trembling hand perfect equilibrium (THPE) as the limit of a sequence of $\tilde{\varepsilon}_n$-constrained equilibria $(\sigma^n)_{n\in\mathbb{N}}$ as the trembles ε_n vanish, extending Selten's (1975) notion of *perfect equilibrium* to infinite games. For each player $i \in N$, the limit strategy σ^*_i is constructed from the weak limits p^*_i of the transition probabilities induced by σ^n_i. Specifically, σ^*_i is defined in equation (4) as the ratio of these limiting probabilities, where the numerator is positive as it captures the probability of action histories following the last action with zero probability in the limit strategy. This construction preserves information about off-path play and, under SAC, ensures convergence of players' payoffs to the limit payoffs, thereby guaranteeing subgame perfection (Proposition 2). Our notion of convergence coincides with Selten's pointwise convergence in finite games, while extending naturally to infinite settings.

[25] Define, for instance, $\nu_{i,t} : S^{t-1}_i \to \Delta(S_i)$ as $d\nu_{i,t}(s_{i,t}|s^{t-1}_i) := \sum_{a^{t-1}\in X^{t-1}} \xi_i(a^{t-1}) \cdot d\mu^t_{s_i|s_i}(s_{i,t}|s^{t-1}_i, a^{t-1})$ where ξ_i is an arbitrary collection of strictly positive weights, i.e., for every $t \in \mathbb{N}$ and $a^{t-1} \in X^{t-1}$, $\xi_i(a^{t-1}) > 0$ and $\sum_{a^{t-1}\in X^{t-1}} \xi_i(a^{t-1}) = 1$.

Theorem 2. *Let Γ be a dynamic game that satisfies sequential absolute continuity. Then, Γ has a trembling hand perfect equilibrium.*

The following result shows that a THPE is subgame perfect and, therefore, a Nash equilibrium.

A strategy profile $\sigma^* = (\sigma_i^*)_{i \in N}$ is a *Nash equilibrium* (NE) if no player can unilaterally deviate to improve their payoff. Formally, for every $i \in N$ and $\sigma_i' \in \Sigma_i$,

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

A history $(\omega^t, a^{t-1}) \in \mathcal{H}$ is *the root of a proper subgame* if it uniquely determines each player's private history. Formally, for every $(\tilde{\omega}^t, \tilde{a}^{t-1}) \in \mathcal{H}$, $t \in \mathbb{N}$, if $(\gamma_i^t(\omega^t, a^{t-1}), a_i^{t-1}) = (\gamma_i^t(\tilde{\omega}^t, \tilde{a}^{t-1}), \tilde{a}_i^{t-1})$ for all $i \in N$, then $(\omega^t, a^{t-1}) = (\tilde{\omega}^t, \tilde{a}^{t-1})$.[26] Let $\mathcal{H}^{\varnothing} \subseteq \mathcal{H}$ denote the set of all such histories. A set $H \in \mathcal{M}(\mathcal{H})$ is *negligible* if $P^t(H \cap \mathcal{H}^t | \sigma) = 0$ for every $t \in \mathbb{N}$, $\sigma \in \Sigma$.

Finally, σ^* is a *subgame perfect equilibrium* (SPE) if it induces a Nash equilibrium in every proper subgame, except possibly for a negligible set. Formally, there exists a negligible set $H \in \mathcal{M}(\mathcal{H})$ such that for every $i \in N$, $t \in \mathbb{N}$, $\sigma_i' \in \Sigma_i$, and $(\omega^t, a^{t-1}) \in (\mathcal{H}^{\varnothing} \cap \mathcal{H}^t) \setminus H$,

$$U_i(\sigma^* | \omega^t, a^{t-1}) \geq U_i(\sigma_i', \sigma_{-i}^* | \omega^t, a^{t-1}).$$

Proposition 2. *If σ^* is a trembling hand perfect equilibrium, then it is a subgame perfect equilibrium and, a fortiori, a Nash equilibrium.*

The following example demonstrates that the game introduced by Harris et al. (1995) and described in Example 2.1, which lacks a subgame perfect equilibrium, violates SAC(b), highlighting the importance of this condition in our equilibrium existence result.

Example 2.2 (Harris et al. (1995), continued)**.** Consider any sequence of player A's actions $(a_n)_{n \in \mathbb{N}}$ in $\mathcal{A}$ where $a_n \neq 0$ for all $n \in \mathbb{N}$ and $a_n \to a_0 := 0$. Notice that players C and D's signals at the beginning of period $t = 2$ consist of players A and B's actions. For every $i \in \{C, D\}$, $n \in \mathbb{N}$, and $b \in \{L, R\}$:

- $\gamma_i(a_n, b) = (a_n, b)$;
- $\mu_s^2(\{s_{i,2}^2 = (0, b)\} | (a_n, b)) = 0$.

[26] This definition adapts Myerson and Reny's (2020) subgame notion to our setting.

Therefore,

$$\lim_{n\to\infty} \mu_s^2(\{s_{i,2}^2 = (0,b)\}|(a_n,b)) = 0 \neq 1 = \mu_s^2(\{s_{i,2}^2 = (0,b)\}|(a_0,b)).$$

This set-wise discontinuity in the signal transition implies a violation of SAC(b). ◂

3.3 Noisy Informational Asymmetries

Sequential absolute continuity and noisy informational asymmetries are closely intertwined. We show that, even in games which lack conditions for existence, adding small amounts of idiosyncratic noise to players' private signals can ensure SAC, provided this noise satisfies certain absolute continuity conditions. Specifically, SAC(a) holds when players cannot perfectly infer the original signals from their noisy observations. SAC(b) is satisfied under weaker continuity requirements on the original signal transitions, given additional regularity conditions. Both SAC(a) and SAC(b) conditions are often met when signals are real-valued and the noise is additive and absolutely continuous with respect to the Lebesgue measure. For example, additive *i.i.d.* noise following uniform or normal distributions typically satisfies the requirements for both SAC(a) and SAC(b).

Introducing noise to players' observations mitigates discontinuities arising from perfect information about other players' signals or actions. In cases of strategic entanglement, as illustrated in Example 1, noisy observations prevent players from finely tuning their strategies based on their opponents' private signals. This noise effectively smooths out the joint distribution of signals, rendering it absolutely continuous with respect to the product of players' marginal distributions. Similarly, in games like Example 2.1, adding noise to the observation of previous actions allows the signal distribution to vary more smoothly with players' moves. This prevents the abrupt changes that occur when actions are perfectly observable.

We say that a dynamic game has *decomposable noisy signals* if, for every $i \in N$ and $t \in \mathbb{N}$, each player i's period-t private signal can be represented as

$$s_{i,t} = m_i(\hat{s}_{i,t}, \epsilon_{i,t}),$$

where $\hat{s}_{i,t} \in \hat{S}_i$ denotes a fundamental signal component and $\epsilon_{i,t} \in \mathcal{E}_i$ represents a *noisy* random variable. The spaces $\hat{S}_i$ and $\mathcal{E}_i$ are Polish spaces endowed with their Borel σ-algebras.[27]

[27] We use the notation: $\epsilon_t = (\epsilon_{1,t}, \ldots, \epsilon_{n,t})$, $\hat{s}_t = (\hat{s}_{1,t}, \ldots, \hat{s}_{n,t})$, $m(\hat{s}_t, \epsilon_t) = (m_1(\hat{s}_{1,t}, \epsilon_{1,t}), \ldots, m_n(\hat{s}_{n,t}, \epsilon_{n,t}))$, $m^t(\hat{s}^t, \epsilon^t) = (m(\hat{s}_1^t, \epsilon_1^t), \ldots, m(\hat{s}_t^t, \epsilon_t^t))$, $\epsilon^t = (\epsilon_1, \ldots, \epsilon_t)$, and $\hat{s}^t = (\hat{s}_1, \ldots, \hat{s}_t)$.

For every $t \in \mathbb{N}$, denote by $\mu^t_{\hat{s},\epsilon}(\cdot|a^{t-1})$ and $\mu^t_{\epsilon}(\cdot|a^{t-1})$ the distribution of $(\hat{s}^t, \epsilon^t)$ and ϵ^t conditional on $a^{t-1} \in X^{t-1}$, respectively. Let $\mu^t_{\epsilon_i}(\cdot|a^{t-1})$ be each marginal of $\mu^t_{\epsilon}(\cdot|a^{t-1})$ for $i \in N$.

Assumption (Noisy observability). *The following conditions hold:*

(a) *For every $t \in \mathbb{N}$ and $a^{t-1} \in X^{t-1}$, $\mu^t_{\hat{s},\epsilon}(\cdot|a^{t-1})$ is absolutely continuous with respect to $\mu^t_{\hat{s}}(\cdot|a^{t-1}) \times \prod_{i\in N} \mu^t_{\epsilon_i}(\cdot|a^{t-1})$;*

(b) *If $\mu^t_s(\tilde{S}|a^{t-1}) > 0$ for $\tilde{S} \in \mathcal{M}(S^t)$, then*

$$\hat{S}(\tilde{S}) := \left\{ \hat{s} \in \hat{S}^t \mid \exists \epsilon \in \mathcal{E}^t,\, (\hat{s}, \epsilon) \in (m^t)^{-1}(\tilde{S}) \right\}$$

satisfies $(\prod_{i\in N} \mu^t_{\hat{s}_i}(\cdot|a^{t-1}))(\hat{S}(\tilde{S})) > 0$.

Noisy observability imposes two requirements: (a) The joint measure of the fundamental signal component $\hat{s}^t$ and the noise term ϵ^t is absolutely continuous with respect to the product of the measure of the fundamental signal component and the product of the marginal measures over each player's noise. This holds for every period $t \in \mathbb{N}$ and history of actions $a^{t-1} \in X^{t-1}$; (b) If a set of histories of signal profiles $\tilde{S}$ has positive measure, then its projection onto the space of histories of fundamental signals, denoted by $\hat{S}(\tilde{S})$, has positive measure according to the product measure of the marginal distributions of each player's fundamental signal component.

The following lemma provides a sufficient condition for noisy observability (b).

Lemma 1. *If noisy observability (a) holds and, for every $i \in N$, $t \in \mathbb{N}$, $s^t_i \in S^t_i$, $\mu^t_{\hat{s}_i}(\{\hat{s}^t_i \in \hat{S}^t_i | \exists \epsilon^t_i \in \mathcal{E}^t_i,\, m^t_i(\hat{s}^t_i, \epsilon^t_i) = s^t_i\}) > 0$, then noisy observability (b) holds.*

Notably, when signals are real-valued vectors and noise is additive, i.e., $S_i \subseteq \mathbb{R}^{\ell_i}$ and $m_i(\hat{s}_i, \epsilon_i) = \hat{s}_i + \epsilon_i$ for $i \in N$, $\ell_i \in \mathbb{N}$, noisy observability holds under certain conditions. For instance, it is satisfied if ϵ_t is *i.i.d.* and

$$d\mu^t_{\epsilon}(\epsilon_{1,t}, \ldots, \epsilon_{n,t}|a^{t-1}) = f^t_{\epsilon}(\epsilon_{1,t}, \epsilon_{2,t}, \ldots, \epsilon_{n,t}, a^{t-1}) \cdot d\lambda^{\ell_1}(\epsilon_{1,t}) \times \ldots \times d\lambda^{\ell_n}(\epsilon_{n,t}), \quad (5)$$

where: λ^{ℓ_i} is the Lebesgue measure over $\mathbb{R}^{\ell_i}$; $f^t_{\epsilon}(\epsilon, a^{t-1}) > 0$ implies $f^t_{\epsilon}(\hat{\epsilon}, a^{t-1}) > 0$ for every $\hat{\epsilon}$ in a neighborhood of ϵ and every $a^{t-1} \in X^{t-1}$; every $\hat{s}^t \in \hat{S}^t$ has a neighborhood with positive measure.

The next result shows that "adding some noise" to the players' observations, satisfying noisy observability, yields SAC(a).

Proposition 3. *Noisy observability implies SAC(a) in any game with decomposable noisy signals.*

Next, we investigate which noise structure implies SAC(b) by characterizing a class of signal profile distributions that are bounded and continuous in the strong total variation norm. For every $t \in \mathbb{N}$, let $\mu^t_{\hat{s},s}(\cdot|a^{t-1})$ be the distribution of $(\hat{s}^t, s^t)$ conditional on $a^{t-1} \in X^{t-1}$.

Assumption (Sequentially continuous noise)**.** *The following holds:*

(a) *For every $t \in \mathbb{N}$, $a^{t-1} \in X^{t-1}$, $\mu^t_{\hat{s}}(\cdot|a^{t-1})$ is continuous in a^{t-1} with respect to the topology of weak convergence of probability measures.*

(b) *For every $t \in \mathbb{N}$, $a^{t-1} \in X^{t-1}$, $d\mu^t_{\hat{s},s}(s^t, \hat{s}^t|a^{t-1}) = f(s^t, \hat{s}^t, a^{t-1}) d\mu^t_{\hat{s}}(\hat{s}^t|a^{t-1}) d\tilde{\mu}^t(s^t)$ where $\tilde{\mu}^t \in \Delta(S^t)$ and f is a bounded density such that, for a $\tilde{\mu}$-full-measure set $\check{S} \subset S^t$, $f(s^t, \cdot, \cdot)$ is continuous for every $s^t \in \check{S}$.*

Sequentially continuous noise imposes two conditions: (a) The measure of the fundamental signal component $\mu^t_{\hat{s}}(\cdot|a^{t-1})$ must be continuous in the history of action profiles a^{t-1} with respect to the topology of weak convergence of probability measures. Notice that this is a weaker continuity requirement compared to SAC(b). In particular, it holds whenever $\hat{s}^t$ is a deterministic and continuous function of a^{t-1}, as in Example 2.1; (b) The joint measure of the signal and fundamental signal $\mu^t_{\hat{s},s}(\cdot|a^{t-1})$ must be absolutely continuous with respect to the product of the fundamental signal measure $\mu^t_{\hat{s}}(\cdot|a^{t-1})$ and a probability measure over the space of signals $\tilde{\mu}^t(\cdot)$ that does not depend on the history of actions. Furthermore, the corresponding density function $f(s^t, \hat{s}^t, a^{t-1})$ must be bounded and almost surely continuous in $\hat{s}^t$ and a^{t-1}.

Proposition 4. *Sequentially continuous noise implies SAC(b) in any game with decomposable noisy signals.*

For real-valued vector signals with additive noise that is absolutely continuous with respect to the Lebesgue measure, we derive a simpler condition implying sequentially continuous noise (b). This condition is satisfied even by arbitrarily small independent noise added to the fundamental signal.

Lemma 2. *If the following conditions hold then sequentially continuous noise (b) is satisfied:*

1. *Noisy observability (a) holds; $S_i \subseteq \mathbb{R}^{\ell_i}$, and $m_i(\hat{s}_i, \epsilon_i) = \hat{s}_i + \epsilon_i$ for $i \in N$, $\ell_i \in \mathbb{N}$.*

2. *For every* $t \in \mathbb{N}$, $a^{t-1} \in X^{t-1}$, $\mu^t_\epsilon(\cdot|a^{t-1})$ *is absolutely continuous with respect to the Lebesgue measure* λ^t. *The density* $f^t : \mathcal{E}^t \times \hat{S}^t \times X^{t-1} \to \mathbb{R}$ *with respect to the product of* $\mu^t_{\hat{s}}(\cdot|a^{t-1})$ *and* λ^t *is bounded and there is a full measure set* $\check{S}^t \in \mathcal{M}(S^t)$ *s.t., for* $s^t \in \check{S}^t$, $\tilde{f}(s^t, \hat{s}^t, a^{t-1}) := f^t(s^t - \hat{s}^t, \hat{s}^t, a^{t-1})$ *is continuous in* $\hat{s}^t$ *and* a^{t-1}.

Thus, by Proposition 4 and Lemma 2, if the density function f^t_ϵ in equation (5) is additionally almost surely continuous in ϵ, then sequentially continuous noise (b) holds.

Lemmas 1 and 2 enable us to apply Propositions 3 and 4 to the class of games introduced in Application 1, thereby establishing equilibrium existence.

Corollary 1. *Let* Γ *be a dynamic game with Lebesgue signals as in Application 1. Then,* Γ *satisfies sequential absolute continuity.*

Example 2.1 features countable signals, ensuring SAC(a) holds. Proposition 4 then implies that adding independent, uniformly distributed noise to the observation of player A's actions yields equilibrium existence. This result follows from fundamental signals being deterministic and continuous functions of past actions, and thus satisfying continuity in the weak convergence of probability measures. We apply this result to Example 2.1 by characterizing an SPE.

Example 2.3 (Harris et al. (1995), continued)**.** Assume that, instead of perfectly observing (a, b), each player $i \in \{C, D\}$ observes $s_i = (a + u_i, b + \tilde{u}_i)$ where $u_i, \tilde{u}_i \sim$ U$[-\delta, \delta]$, for $\delta \in \mathcal{A}$, are independent, uniformly distributed private signals. Notably, the noise range δ can be arbitrarily close to zero.

The following strategy is an SPE. Player A randomizes uniformly between the actions δ and $-\delta$. The other players' best responses to A's strategy are as follows: each player $i \in \{C, D\}$ plays L if $a + u_i < 0$ and R otherwise, while B sets $\beta = 1/2$.[28] ◂

Harris et al. (1995) restores existence in this example, and more generally in games with almost perfect information, by introducing a *public* correlation device at the start of the second period. Our Corollary 1 demonstrates an alternative approach: existence can be ensured by incorporating idiosyncratic noise into players' *private* signals, a result that extends to a wider class of dynamic games with informational asymmetries.

[28] In Supplemental Appendix B.7 we show that this strategy profile is an SPE.

4 Markov Games

We now introduce *Markov* games with informational asymmetries. In this environment, the set of the states of the world can be written as $\Omega = \Omega^R \times \Omega^{\not{R}}$, which represent the *payoff-relevant* and *payoff-irrelevant states*, respectively. For every $i \in N$, players' payoffs and state transition depend only on the current payoff-relevant state and action profile, $g_i : \Omega^R \times X \to \mathbb{R}$ and $\mu : \Omega^R \times X \to \Delta(\Omega)$. Additionally, in the discounted payoffs case we can allow g_i to depend on the period without including it as part of the payoff-relevant states, i.e., $g_i(t, \omega_t^R, a_t) = \delta^t \cdot u_i(\omega_t^R, a_t)$, $\delta \in (0,1)$. Each player receives information about the current payoff-relevant and irrelevant state, which we term payoff-relevant and irrelevant private signals, respectively. That is, $S_i = S_i^R \times S_i^{\not{R}}$ and $\gamma_i = (\gamma_i^R, \gamma_i^{\not{R}})$ such that $\gamma_i^R : \Omega^R \to S_i^R$ and $\gamma_i^{\not{R}} : \cup_{t \in \mathbb{N}} \mathscr{H}^t \to S_i^{\not{R}}$. In particular, the payoff irrelevant signal may contain information about the past that is no longer payoff relevant. The available actions depend on the current payoff-relevant signals, $A_i : S_i^R \to X_i$.

For ease of exposition, we assume that regular conditional measures exist.[29] In particular, for every $\sigma \in \Sigma$, $t \in \mathbb{N}$, let $\mu_{\omega,t}^R(\cdot|\cdot, \sigma) : \cup_{i \in N} \mathscr{H}_i \to \Delta(\Omega^R)$ be the transition probability from private history $h_i \in \mathscr{H}_i$, $i \in N$, to period-t payoff-relevant states $\omega_t \in \Omega^R$ conditional on strategy σ.[30]

For every $i \in N$, player i's strategy $\sigma_i \in \Sigma_i$ is *stationary Markov* if it conditions only on the current payoff-relevant private signal. Formally, $\sigma_i(\cdot|h_i) = \sigma_i(\cdot|\mathscr{s}_i^R(h_i))$, for each $h_i \in \mathscr{H}_{A_i}$, where the measurable function $\mathscr{s}_i^R(h_i)$ projects each private history h_i onto the current payoff-relevant signal in S_i^R. Denote by $\Sigma_i^M \subset \Sigma_i$ and Σ^M the set of player i's stationary Markov strategies and strategy profiles, respectively.

We require Markov games to satisfy two additional conditions:

(*i*) *Markov information.* For every $i \in N$, $t \in \mathbb{N}$, $\sigma^M \in \Sigma^M$, $\mu_{\omega,t}^R(\cdot|h_i, \sigma^M) = \mu_{\omega,t}^R(\cdot|\tilde{h}_i, \sigma^M)$ for each $h_i, \tilde{h}_i \in \mathscr{H}_i$ such that $\mathscr{s}_i^R(h_i) = \mathscr{s}_i^R(\tilde{h}_i)$;

(*ii*) *Markov payoff.* For every $i \in N$, $t \in \mathbb{N}$, $s^t, \tilde{s}^t \in S^t$ and $a^t \in X^t$, $\hat{g}_{i,t}(s^t, a^t) = \hat{g}_{i,t}(\tilde{s}^t, a^t)$ if $s^t = ((s_\ell^{R,t})_{\ell \le t}, (s_\ell^{\not{R},t})_{\ell \le t})$, $\tilde{s}^t = ((s_\ell^{R,t})_{\ell \le t}, (\tilde{s}_\ell^{\not{R},t})_{\ell \le t})$.

Markov information requires that if all players follow stationary Markov strategies, then each player's beliefs about the current payoff-relevant state, conditional on her

[29]A sufficient condition for the existence of regular conditional measures is to assume that Ω and S are Souslin spaces. A subset of a Hausdorff space is *Souslin* if it is a continuous image of a Polish space, i.e., a complete, separable metric space. See Corollary 10.4.6 in Bogachev (2007).

[30]See Supplemental Appendix B.4 for a formal definition.

private history, depend only on the current payoff-relevant signal. This assumption implies that the additional information contained in the private history is payoff irrelevant for future decisions, and it formalizes the notion that *bygones are bygones* in an environment with informational asymmetries. Markov payoff complements Markov information by stating that payoff-relevant signal profiles determine payoffs, justifying why stationary Markov strategies condition only on them.

Markov information and Markov payoff conditions are satisfied in various games of interest. In standard stochastic games, these conditions hold naturally: the current state of the world, observed by all players, serves as the payoff-relevant signal, and flow payoffs depend solely on this signal and the current action profile. Notably, these conditions can also be met in certain games with asymmetric information. Examples include stochastic games with public and private shocks (Balbus et al., 2013), dynamic cheap talk games (Renault et al., 2013), and asynchronous games, such as revision games with or without observation of previous moves (Kamada and Kandori, 2020). We discuss these Markov games in detail in Applications 2 and 3.

We now define Markov absolute continuity (MAC), which replaces SAC in Markov games. For every $t \in \mathbb{N}$, set $S^{R,t} := (S^R)^t$, and let $\gamma^{R,t} : \Omega^t \to S^{R,t}$ be the projection of γ^t onto $S^{R,t}$. For every $a^{t-1} \in X^{t-1}$, this projection, in conjunction with $\mu^t_\omega(\cdot|a^{t-1})$, induces a measure, $\mu^{R,t}_s(C|a^{t-1}) := \mu^t_\omega(\{\omega^t \in \Omega^t : \gamma^{R,t}(\omega^t) \in C\}|a^{t-1})$, for $C \in \mathcal{M}(S^{R,t})$, over payoff-relevant signal profile histories. Let $\mu^M_{s_i,t}(\cdot|a^{t-1})$ be the marginal of $\mu^{R,t}_s$ on player i's period-t payoff-relevant signals in S^R_i.

Assumption (Markov absolute continuity). *The following holds:*

(a) *For every $t \in \mathbb{N}$, $a^{t-1} \in X^{t-1}$, $\mu^{R,t}_s(\cdot|a^{t-1})$ is absolutely continuous with respect to the product measure $\prod_{i\in N}\prod_{\ell\le t}\mu^M_{s_i,\ell}(\cdot|a^{t,(\ell-1)})$;*

(b) *For every $t \in \mathbb{N}$, $\mu^{R,t}_s$ is bounded, and continuous in X^{t-1} in the strong total variation norm.*

MAC neither weakens nor strengthens SAC, but rather modifies it in two ways. First, it applies exclusively to *payoff-relevant signal profiles*. Second, MAC requires that absolute continuity holds with respect to the product of marginal measures over payoff-relevant private signals, taken not only across players but also across periods.

The following definition adapts stationary Markov perfect equilibria (MPE) to games with informational asymmetries.

Definition 4. A stationary Markov strategy profile $\sigma^M \in \Sigma^M$ is a *stationary Markov perfect equilibrium* if it is a trembling hand perfect equilibrium.

In standard stochastic games, where players observe the history of the game before moving, an MPE is defined as a strategy profile that maximizes continuation payoffs for each player. However, in settings featuring informational asymmetries, defining continuation payoffs requires specifying beliefs over possible histories that are not fully observed. To circumvent this issue, we define an MPE as a THPE in Markov strategies. In light of Propositions 1 and 2, our notion exhibits desirable sequential rationality properties and coincides with the usual one in standard stochastic games.[31]

Theorem 3. *Let Γ be a Markov game that satisfies Markov absolute continuity. Then, Γ has a stationary Markov perfect equilibrium.*

The proof of Theorem 3 builds upon the ideas used in the proof of Theorem 2, with two main novelties. First, the MAC condition ensures absolute continuity across periods, which prevents strategic entanglement across periods and allows strategies to retain their Markov properties in the limit. Second, we need to show that in a Markov environment, when players employ Markov strategies, their opponents have best responses within the class of Markov strategies. The latter is guaranteed by the Markov information and payoff conditions.

Games in the following class are Markov and satisfy MAC.

Application 2 (Stochastic games with public and private shocks)**.** This class of games enriches standard stochastic games by incorporating independent, payoff-relevant private shocks into the state space. Such a framework is well-suited to model dynamic oligopolies, where firms set prices and the current public and private shocks represent demand and firm-specific cost characteristics, respectively.

Formally, let $\Omega^R = \hat{\Omega} \times \prod_{i \in N} \Theta_i$ and $\Omega^{\not R} = X$, where: $\hat{\Omega}$ is a set of countable, commonly observed, payoff-relevant states, interpreted as public shocks; each Θ_i represents a, potentially uncountable, set of private shocks which are only observed by player i; X contains the previous period's action profile which is also commonly observed. Payoff-relevant private shocks display the following conditionally independent structure, $\mu((\hat{\omega}', \theta')|(\hat{\omega}, \theta), a) = f(\hat{\omega}', \hat{\omega}, a) \cdot \prod_{i \in N} \mu_i(\theta_i'|\hat{\omega}')$ for every $a \in X$, and

[31] Constrained equilibria in Markov strategies converge to MPE in the weak topology "period by period". This convergence implies the convergence for THPE. See Supplemental Appendix B.4 for details.

$(\hat{\omega}, \theta), (\hat{\omega}', \theta') \in \Omega$, where $\mu_i : \hat{\Omega} \to \Delta(\Theta_i)$ is a transition probability, and f is interpreted as a density function.

For every $i \in N$, MAC(a) holds since $\hat{\Omega}$ is countable and each Θ_i is conditionally independent; MAC(b) holds, by Proposition 5, if f is uniformly bounded and continuous for every $(\hat{\omega}', \hat{\omega}, \theta)$. Furthermore, Markov information and Markov payoff are satisfied (see Supplemental Appendix B.8). Therefore, by Theorem 3, there exists a stationary Markov perfect equilibrium where each player's strategy depends only on the current public and private shocks.

An equilibrium exists even though players observe the history of action profiles, which may appear at odds with the non-existence Example 2.1. Existence holds since MAC restricts only the transition of payoff-relevant signal profiles, i.e., of public and private shocks. ◂

5 Relaxing Discounted Payoffs

As discussed in Section 3, our payoff boundedness condition not only accommodates but also generalizes the discounted payoffs assumption commonly employed by the literature. This greater generality is substantive: we show that our payoff boundedness holds for a class of non-discounted games, which we call *games with stochastic move opportunities*, where players receive opportunities to move at random times.

Consider the following class of non-discounted games where players draw opportunities to move at random periods. Opportunities to move are drawn from an interval $[0, T)$, with $T \in \mathbb{R}_+ \cup \{\infty\}$, and states of the world takes the form $\Omega = [0, T) \times \hat{\Omega}$, where $\hat{\Omega}$ is a set of underlying states that contains an absorbing state $\hat{\omega}^{end}$, such that for any $t \in [0, T)$, the game ends if $(t, \hat{\omega}^{end})$ is reached.

For every $\ell \in \mathbb{N}$, period ℓ represents the ℓ'th opportunity to move for any player, and t_ℓ is the corresponding timing. For $\omega_\ell = (t_\ell, \hat{\omega}_\ell)$ and $\omega_{\ell'} = (t_{\ell'}, \hat{\omega}_{\ell'})$, the state transition probability is such that, with probability one, $t_\ell < t_{\ell'}$ whenever $\ell < \ell'$, ensuring that later timings are assigned to later opportunities. Furthermore, for $\omega^t \in \Omega^t$, if $\omega^t_\tau = (\cdot, \hat{\omega}^{end})$ for some $\tau < t$ then $g_i(\omega^t, \cdot) = 0$ for all $i \in N$. That is, after visiting the state $\hat{\omega}^{end}$ for the first time, players do not incur flow payoffs.

Denote by $H^{end,t} := \{(\omega^t, a^t) \in \Omega^t \times X^t | \omega^t_t = (\cdot, \hat{\omega}^{end}), \omega^t_\ell \neq (\cdot, \hat{\omega}^{end}), \ell < t\}$ the subset of $\Omega^t \times X^t$ in which the game ends in period t. The following lemma provides a condition which implies payoff boundedness.

Lemma 3. *If* $\sum_{t\in\mathbb{N}} t \cdot \sup_{\sigma\in\Sigma} P^t(H^{end,t}|\sigma) < \infty$ *and* $\sum_{t\in\mathbb{N}} P^t(H^{end,t}|\sigma) = 1$*, for every* $\sigma \in \Sigma$*, then payoff boundedness holds.*

In games with stochastic move opportunities, Lemma 3 states that payoff boundedness holds whenever (i) the sequence in which the t-th element is t times an upper bound on the period-t probability that the game ends is summable, and (ii) the game ends in finite time with probability 1.

A consequence of this result is that, when actions do not affect the length of the game, i.e., $P^t(H^{end,t}|\sigma)$ is independent of σ, payoff boundedness is satisfied as long as the expected number of opportunities to move is finite, and (ii) holds. Notably, these conditions hold in revision games (Kamada and Kandori, 2020), where opportunities to move are drawn at exogenous Poisson rates independently of the previous play.[32]

6 Imperfect Observation of Moving Periods

Theorems 1, 2, and 3, and Propositions 1 and 2 hold in a broader class of games that allows players to be *inactive* in certain periods. During these inactive periods, players neither make moves nor record information in their private histories. Due to the additional complexity this generalization entails, we present the complete details in Supplemental Appendix B.1. This section outlines the main elements of this extension and illustrates an application to Markov games.

Whenever *active*, a player receives informative signals about the history of states and action profiles before moving; when *inactive*, a player neither observes signals nor takes non-trivial actions. Private histories exclude information from inactive periods, implying players may be oblivious to the number of moves that have occurred. We assume that at least one player is active in each period, and without loss, that players receive non-zero flow payoffs only when active.

Allowing for active and inactive players extends our analysis to settings where players may not be informed whether moves have taken place in the past, thus introducing uncertainty about the multi-stage structure of the game. Such games have been studied in various contexts: Kreps and Ramey (1987) provide an early example where players lack a sense of calendar time; Matsui (1989) considers an espionage

[32] Revision games can be represented as follows. Move opportunities arrive prior to an exogenous deadline, $T < \infty$, at a constant Poisson rate. The underlying state space $\hat{\Omega}$ contains two elements $\{\omega^0, \hat{\omega}^{end}\}$. The game starts in state ω^0 and switches to $\hat{\omega}^{end}$ as soon as a time larger than the deadline is drawn. There is a finite constant set of available actions, and the players' payoffs realize once the game ends.

game with private revision opportunities; Kamada and Moroni (2018) examine outcomes in coordination games with private timing; Doval and Ely (2020) investigate the range of equilibrium outcomes that can arise across different information structures and extensive forms for a given base game.

We apply the inactive player framework to the class of Markov games to show the existence of stationary Markov perfect equilibria. See Supplemental Appendix B.8 for a formal description.

Application 3 (Asynchronous games)**.** Assume that only one player is active at each period, who perfectly observes the current payoff-relevant state and, possibly imperfectly, the payoff-irrelevant one. This ensures that the Markov information condition holds. By Proposition 5, MAC holds if $d\mu(\omega_t^R|\omega_{t-1}^R, a_{t-1}) = f(\omega_t^R, \omega_{t-1}^R, a_{t-1})d\nu(\omega_t^R)$ for some bounded density f that is continuous in a_{t-1} for each $(\omega_t^R, \omega_{t-1}^R)$ and some measure $\nu \in \Delta(\Omega^R)$.

Asynchronous revision games with finite actions, where moving opportunities are drawn at Poisson rates (Kamada and Kandori, 2020), fall within this application. Notice that MAC still holds if we ascribe previous moving times to the payoff-irrelevant states, even though these timings belong to an uncountable set. Thus, Theorem 3 guarantees existence of a stationary MPE, even when Theorem 2 may not apply due to the possible violation of SAC.

Furthermore, Theorem 3 implies the existence of a stationary MPE in a class of dynamic cheap talk games (Renault et al., 2013). These games model communication between an informed sender and a receiver controlling state transitions.[33] Assuming finite message and action spaces ensures that payoff boundedness, continuity, and MAC hold, while allowing for an uncountable set of exogenous states observed only by the sender. Theorem 3 guarantees an equilibrium where the sender's strategy depends on the receiver's previous action and the current exogenous state, while the receiver's strategy depends on her previous action and the current sender's message. ◂

A Appendix

The following lemma shows that each player's expected payoff can be written in a simple way as a function of $p(\cdot|\cdot, \sigma)$.

[33] To satisfy the Markov information condition, the state transition must be either independent or controlled solely by the receiver's action, without dependence on the current state.

Lemma 4. *The function $p(\cdot|\cdot,\sigma)$ is measurable, the set Σ is non-empty, and, for each $i \in N$, player i's expected payoff when players follow strategy $\sigma \in \Sigma$ is given by*

$$U_i(\sigma) = \sum_{t\in\mathbb{N}} U_{i,t}(\sigma), \tag{6}$$

where $U_{i,t}(\sigma) := \sum_{a^t\in X^t} \int_{\Omega^t} g_i(\omega^t, a^t) \cdot p(a^t|\omega^t,\sigma)\, d\mu^t_\omega(\omega^t|a^{t,(t-1)})$.

Proof. Define the correspondence $\tilde{A}_i : \mathscr{H}_i \rightrightarrows \Delta(X_i)$ as $\tilde{A}_i(h_i) := \Delta(A_i(h_i))$, where $\Delta(X_i)$ is endowed with the weak topology of probability measures. For each $h_i \in \mathscr{H}_i$, $\tilde{A}_i(h_i)$ is the set of probability measures with support in $A_i(h_i)$. By Himmelberg and Van Vleck (1975), $\tilde{A}_i(h_i)$ is weakly measurable. By Theorems 15.19 and 18.13 in Aliprantis and Border (2006), $\tilde{A}_i$ admits a measurable selector. This shows that the set of strategies is non-empty. By Fubini- Tonelli theorem,[34] since g_i is bounded, we can exchange the order of integration of the counting measure over actions and the measure over states repeatedly to obtain that player i's period-t expected payoff, given σ, is $U_{i,t}(\sigma)$. □

Following Lemma 4, for $t \in \mathbb{N}$, we rewrite player i's period-t continuation expected payoff from $\sigma \in \Sigma$ after history $(\omega^\tau, a^{\tau-1}) \in \mathscr{H}^\tau$, $\tau \leq t$, as $U_{i,t}(\sigma|\omega^\tau, a^{\tau-1}) = \sum_{a^t\in X^t} \int_{\Omega^t} g_i(\omega^t, a^t) \cdot p(a^t|\omega^t, a^{\tau-1}, \sigma)\, d\mu^t_\omega(\omega^t|\omega^\tau, a^{t,(t-1)})$.

Next, we provide a simplified proof of Theorem 1; see Supplemental Appendix B.3 for a complete proof.

Proof of Theorem 1. Assume that the game has a finite length T, X is infinite, and $A_i(h_i) = X_i$ for each $h_i \in \mathscr{H}_i$, $i \in N$.

Let $\tilde{\varepsilon} \in \mathscr{E}$. We start by defining the relevant topology on the set of $\tilde{\varepsilon}$-constrained strategy profiles. For every $i \in N$, $t \leq T$, we can construct what we call a *reference measure* $\nu^t_i \in \Delta(S^t_i)$ such that, for each $a^{t-1} \in X^{t-1}$, $\mu^t_{s_i}(\cdot|a^{t-1})$ is absolutely continuous with respect to ν^t_i.[35] Define $I_{i,\varphi^t_i} : \Sigma_i \to \mathbb{R}$ as

$$I_{i,\varphi^t_i}(\sigma_i) := \int_{S^t_i} \sum_{a^t_i\in X^t_i} \varphi^t_i(s^t_i, a^t_i) \cdot p_i\left(a^t_i|s^t_i, \sigma_i\right) d\nu^t_i(s^t_i).$$

[34] Theorems 11.27 and 11.28 in Aliprantis and Border (2006).
[35] Footnote 25 exemplifies the construction of one such reference measure.

The *weak topology* on $\Sigma_i(\tilde{\varepsilon})$ is the coarsest topology such that I_{i,φ_i^t} is continuous for every $\varphi_i^t \in CI(S_i^t \times X_i^t, \nu_i^t)$, the set of *Carathéodory integrands* on $(S_i^t \times X_i^t, \nu_i^t)$.[36] The *weak topology* on $\Sigma(\tilde{\varepsilon}) = \times_{i\in N}\Sigma_i(\tilde{\varepsilon})$ is the product topology, where each $\Sigma_i(\tilde{\varepsilon})$ is endowed with its weak topology.

Consider the set of induced transition probabilities $\mathscr{P}_i(\tilde{\varepsilon}) := \{p_i(\cdot|\cdot,\sigma_i)|\sigma_i \in \Sigma_i(\tilde{\varepsilon})\}$ and endow it with the *weak topology* defined in the same manner as above. Similarly, let the *weak topology* on $\mathscr{P}(\tilde{\varepsilon}) := \times_{i\in N}\mathscr{P}_i(\tilde{\varepsilon})$ be the product topology.[37] For every $i \in N$, the resulting topological spaces $\Sigma(\tilde{\varepsilon})$ and $\mathscr{P}(\tilde{\varepsilon})$ are homeormorphic. This holds since, for every $\tilde{p}_i(\cdot|\cdot) \in \mathscr{P}_i(\tilde{\varepsilon})$, there is a strategy, defined recursively as

$$\sigma_i\big(a_{i,t}^t|s_i^t, a_i^{t,(t-1)}\big) = \frac{\tilde{p}_i(a_i^t|s_i^t)}{\tilde{p}_i(a_i^{t,(t-1)}|s_i^{t,(t-1)})}, \tag{7}$$

for $t \in \mathbb{N}$, $(s_i^t, a_i^{t,(t-1)}) \in \mathscr{H}_i$, where $\tilde{p}_i(a_i^{t,(0)}|s_i^{t,(0)}) := 1$, such that $p_i(\cdot|\cdot,\sigma_i) = \tilde{p}_i(\cdot|\cdot)$.

For every $(p_i)_{i\in N} \in \mathscr{P}(\tilde{\varepsilon})$, where $p_i = p_i(\cdot|\cdot,\sigma_i)$ for some $\sigma_i \in \Sigma_i(\tilde{\varepsilon})$, $i \in N$, let $\tilde{U}_i((p_i)_{i\in N}) := U_i((\sigma_i)_{i\in N})$ be player i's expected payoff when players follow strategy $(\sigma_i)_{i\in N}$. Player i's *$\tilde{\varepsilon}$-constrained best response correspondence* $r_i : \mathscr{P}_{-i}(\tilde{\varepsilon}) \rightrightarrows \mathscr{P}_i(\tilde{\varepsilon})$ is

$$r_i(p_{-i}) \in \arg\max\{\tilde{U}_i(\tilde{p}_i, p_{-i})|\tilde{p}_i \in \mathscr{P}_i(\tilde{\varepsilon})\},$$

and let $r : \mathscr{P}(\tilde{\varepsilon}) \rightrightarrows \mathscr{P}(\tilde{\varepsilon})$ be the Cartesian product of the r_i for each $i \in N$.

Relying on the homeomorphism introduced, every fixed point of r is mapped to a strategy profile that, by definition, forms an $\tilde{\varepsilon}$-constrained equilibrium. By applying the Kakutani-Fan-Glicksberg fixed point theorem, we establish that such a fixed point exists. In particular, we show that: **(1)** $\mathscr{P}(\tilde{\varepsilon})$ is compact, and **(2)** convex; **(3)** r has closed graph, **(4)** is non-empty, and **(5)** convex valued.[38] As the arguments establishing points **(2)** and **(5)** are standard, we omit them from the proof.

(1) $\mathscr{P}(\tilde{\varepsilon})$ is compact. By Theorem 2.3 in Balder (1988), for each $t \leq T$ the set $\{p_i^t(\cdot|\cdot,\sigma_i)|\sigma_i \in \Sigma_i(\tilde{\varepsilon})\}$ is *relatively compact*,[39] implying $\mathscr{P}(\tilde{\varepsilon})$ is relatively compact. To

[36] A measurable function $f : Y \times Z \to \mathbb{R}$ is a *Carathéodory integrand* on $(Y \times Z, \beta)$, where $\beta \in \Delta(Y)$ and Z is a compact metric space endowed with its Borel σ-algebra, if: (i) for every $y \in Y$, $f(y,\cdot)$ is continuous in Z; (ii) there exists $\psi \in L^1(Y,\beta)$ such that $|f(y,\cdot)| \leq \psi(y)$ for every $y \in Y$. For every $\psi \in L^1(Y,\beta)$, ψ is $\mathcal{M}(Y)$-measurable and $\int_Y |\psi| d\beta < +\infty$.

[37] We view $p_i(\cdot|\cdot,\sigma_i)$ as a transition probability from $\cup_{t\leq T}S_i^t$ to $\cup_{t\leq T}X_i^t$, where for each $t \leq T$ and $s_i^t \in S_i^t$, $p_i(\cdot|s_i^t,\sigma_i)$ has support in X_i^t.

[38] As discussed by Balder (1988), the requirement that $\mathscr{P}(\tilde{\varepsilon})$ is a subset of a locally convex Hausdorff space is satisfied by considering equivalent classes of transition probabilities that induce the same expected payoff. See Supplemental Appendix B.3 for further details.

[39] A set is *relatively compact* if its closure is compact.

see that it is closed, let $(\sigma_i^\alpha)_{\alpha\in\Lambda}$ be a net of player i's strategies, for some directed set Λ, and suppose that $p_i(\cdot|\cdot,\sigma_i^\alpha)$ converges weakly to a transition probability represented by the measurable function $p_i^*(\cdot|\cdot)$. Then $p_i^*(\cdot|\cdot) = p_i(\cdot|\cdot,\sigma_i^*)$, for σ_i^* defined recursively using equation (7). Lemma 6 in Supplemental Appendix B.3 shows that $\sigma_i^* \in \Sigma_i(\tilde{\varepsilon})$. Therefore, $\mathscr{P}(\tilde{\varepsilon})$ is compact.

(3) and (4) r has closed graph and is non-empty. By Radon-Nikodym theorem, SAC(a) implies the existence of a density function $\tilde{f}^t : S^t \times X^{t-1} \to \mathbb{R}$ such that $d\mu_s^t(s^t|a^{t-1}) = \tilde{f}^t(s^t, a^{t-1}) \cdot d\mu_{s_1}^t(s_1^t|a^{t-1}) \times \ldots \times d\mu_{s_n}^t(s_n^t|a^{t-1})$ for every $t \in \mathbb{N}$. For $t \in \mathbb{N}$, $s^t \in S^t$, let $\nu^t(s^t) = \prod_{i\in N} \nu_i^t(s_i^t)$. Applying again Radon-Nikodym, we can write

$$d\mu_s^t(s^t|a^{t-1}) = f^t(s^t, a^{t-1}) \cdot d\nu_1^t(s_1^t) \times \ldots \times d\nu_n^t(s_n^t)$$

where $f^t : S^t \times X^{t-1} \to \mathbb{R}$ denotes another density function. Thus, by combining Lemma 4, payoff continuity and SAC(a), we obtain

$$\tilde{U}_i(p) = \sum_{t=1}^{T} \int_{S^t} \sum_{a^t\in X^t} \hat{g}_{i,t}(s^t, a^t) \cdot f^t(s^t, a^{t,(t-1)}) \cdot p(a^t|s^t, \sigma)\, d\nu_1^t(s_1^t) \times \ldots \times d\nu_n^t(s_n^t).$$

Proposition 5 shows that there is a version of $f^t(s^t, a^{t,(t-1)})$ belonging to $CI(S^t, X^{t-1}, \nu^t)$ if and only if SAC(b) holds. Thus, by Theorem 2.5 in Balder (1988), $\tilde{U}_i(p)$ is continuous in p, since, by payoff continuity, $\hat{g}_{i,t}(s^t, a^t) \cdot f^t(s^t, a^{t,(t-1)})$ is a Carathédory integrand for every $t \le T$. For every $i \in N$, as the transition probability space $\mathscr{P}(\tilde{\varepsilon})$ is compact, the continuity of $\tilde{U}_i$ implies that r has closed graph and that r_i is non-empty. □

We present next the proofs of the results that appear in Section 3.3.

Proof of Proposition 3. Suppose that noisy observability holds. Let $\mu_s^{P,t} := \prod_{i\in N} \mu_{s_i}^t$ be the product measure induced by the players marginals over their private signals up to period t.

Let $\tilde{S} \in \mathcal{M}(S^t)$ such that $\mu_s^t(\tilde{S}) > 0$, and set $B := (m^t)^{-1}(\tilde{S})$. By Lemma 4.46 in Aliprantis and Border (2006), since a Polish space is second countable, there exist sets $\hat{B} \in \mathcal{M}(\hat{S}^t)$ and $E(\hat{s}) \in \mathcal{M}(\mathcal{E}^t)$ for each $\hat{s} \in \hat{B}$ such that we can write

$$B = \{(\hat{s}, \epsilon) | \hat{s} \in \hat{B}, \epsilon \in E(\hat{s})\}.$$

Therefore, by noisy observability (a) we have

$$\mu_s^t(\tilde{S}) = \mu_{s,\epsilon}^t(B) = \int_{\hat{B}} \int_{E(\hat{s}^t)} f(\hat{s}^t, \epsilon^t, a^{t-1})\, d\mu_{\epsilon_1}^t(\epsilon_1^t|a^{t-1}) \ldots d\mu_{\epsilon_n}^t(\epsilon_n^t|a^{t-1}) d\mu_{\hat{s}}^t(\hat{s}^t|a^{t-1}) > 0,$$

for some measurable density $f : \hat{S}^t \times \mathcal{E}^t \times X^{t-1} \to \mathbb{R}$. Hence, there is a measurable $\hat{B}_0 \subset \hat{B}$ such that $\mu^t_{\hat{s}}(\hat{B}_0) > 0$, and, for each $\hat{s} \in \hat{B}_0$,

$$\int_{E(\hat{s}^t)} f(\hat{s}^t, \epsilon^t, a^{t-1})\, d\mu^t_{\epsilon_1}(\epsilon^t_1|a^{t-1}) \dots d\mu^t_{\epsilon_n}(\epsilon^t_n|a^{t-1}) > 0.$$

Furthermore, $\mu^t_s(\tilde{S}_0) > 0$, where $\tilde{S}_0 := \{m^t(\hat{s}, \epsilon)|\hat{s} \in \hat{B}_0, \epsilon \in E(\hat{s})\}$. At the same time, we have

$$\mu^P_s(\tilde{S}_0) = \int_{\hat{B}_0} \int_{E(\hat{s}^t)} f(\hat{s}^t, \epsilon^t, a^{t-1})\, d\mu^t_{\epsilon_1}(\epsilon^t_1|a^{t-1}) \dots d\mu^t_{\epsilon_n}(\epsilon^t_n|a^{t-1}) d\mu^t_{\hat{s}_1}(\hat{s}^t_1|a^{t-1}) \dots d\mu^t_{\hat{s}_n}(\hat{s}^t_n|a^{t-1}) > 0,$$

since by noisy observability (b), $\mu^{P,t}_{\hat{s}}(\hat{B}_0) > 0$. □

Proof of Proposition 4 and Lemma 2. For every $t \in \mathbb{N}$, $B \in \mathcal{M}(S^t)$, $a^{t-1} \in X^{t-1}$,

$$\mu^t_s(B|a^{t-1}) \;= \int_B \int_{\hat{S}^t} f^t(s^t, \hat{s}^t, a^{t-1})\, d\mu^t_{\hat{s}}(\hat{s}^t|a^{t-1}) d\tilde{\mu}^t(s^t)$$

By sequentially continuous noise, $\int_{\hat{S}^t} f^t(s^t - \hat{s}^t, \hat{s}^t, a^{t-1})\, d\mu^t_{\hat{s}}(\hat{s}^t|a^{t-1}) \in CI(S^t \times X^{t-1}, \lambda^t)$. By Proposition 5, μ^t_s is bounded and continuous in the strong total variation norm. Similarly, to show Lemma 2, notice that for $\hat{B} \in \mathcal{M}(\hat{S}^t)$ we can write

$$\begin{aligned}
\mu^t_s(B \times \hat{B}|a^{t-1}) &= \int_{\hat{B}} \int_{B - \hat{s}^t} f^t(\epsilon^t, \hat{s}^t, a^{t-1})\, d\lambda^t(\epsilon^t) d\mu^t_{\hat{s}}(\hat{s}^t|a^{t-1}) \\
&= \int_{\hat{B}} \int_B f^t(s^t - \hat{s}^t, \hat{s}^t, a^{t-1})\, d\lambda^t(s^t) d\mu^t_{\hat{s}}(\hat{s}^t|a^{t-1}) \\
&= \int_B \int_{\hat{B}} f^t(s^t - \hat{s}^t, \hat{s}^t, a^{t-1})\, d\mu^t_{\hat{s}}(\hat{s}^t|a^{t-1}) d\lambda^t(s^t),
\end{aligned}$$

where the first equality holds by absolute continuity, the second equality holds by translation invariance of the Lebesgue measure λ^t, and the third by Fubini-Tonelli. □

We conclude by showing a measure theoretical result of independent interest which equates the continuity and boundedness of measures in the strong total variation norm to Carathéodory integrand densities. Let $(Y, \mathcal{M}(Y), \beta)$ be a measure space, Z a countable metric space endowed with the σ-algebra of all subsets of Z and the counting measure, $\mathcal{M}(Y)_0$ a sub σ-algebra of $\mathcal{M}(Y)$, $\mu(\cdot|\cdot) : Z \to \mathcal{M}(Y)$ a transition probability from Z to Y, and $\pi(Y)$ the set of finite measurable partitions of Y.

Proposition 5. *Suppose that there is a measure $\hat{\beta}$ over Y and a function $\hat{\varphi} : Y \times Z \to \mathbb{R}_+$, that is $\mathcal{M}(Y)_0 \otimes \mathcal{M}(Z)$ measurable and such that for every $B \in \mathcal{M}(Y)_0$ and $z \in Z$*

$$\mu(B|z) = \int_B \hat{\varphi}(y, z)\, d\hat{\beta}(y).$$

Then:

1. *There is a set $\hat{Y}$ with $\hat{\beta}(Y \setminus \hat{Y}) = 0$ such that $\hat{\varphi}(y, \cdot)$ is continuous for every $y \in \hat{Y}$ if and only if μ is continuous in Z in the strong total variation norm;*
2. *Suppose $\hat{\varphi}(\cdot, y)$ is continuous for every $y \in Y$. Then there is a function $\psi \in L^1(Y, \hat{\beta})$ such that $\hat{\varphi}(y, z) \le \psi(y)$ for every $y \in Y$ if and only if μ is bounded in the strong total variation norm.*

Proof. Part 1. **Necessity.** We argue by contradiction. Suppose there is a $B \in \mathcal{M}(Y)_0$ with $\hat{\beta}(B) \in (0, \infty)$, such that for each $y \in B$, there is $z^*(y) \in Z$, $\varepsilon(y) > 0$ and a sequence $(z^n(y))_{n \in \mathbb{N}}$ such that $d\left(z^n(y), z^*(y)\right) < \frac{1}{n}$ and $|\hat{\varphi}(y, z^n(y)) - \hat{\varphi}(y, z^*(y))| > \varepsilon(y)$. Define the correspondence $\xi : B \rightrightarrows Z$

$$\xi(y) = \left\{z \in Z \middle| \exists \varepsilon \in (0,1),\, \forall n \in \mathbb{N},\, \exists z_n \in Z,\, d(z_n, z) < 1/n,\, |\hat{\varphi}(y, z_n) - \hat{\varphi}(y, z)| > \varepsilon \right\}.$$

Since $z^*(y) \in \xi(y)$ for each $y \in B$, ξ is non-empty valued. Furthermore, let $z \in Z$ and let $\{\varepsilon_j\}_{j \in \mathbb{N}}$ be a countable dense subset of $(0, 1)$. We have

$$\begin{aligned}
&\{y \in B | z \in \xi(y)\} \\
&\quad = \left\{y \in B \middle| \exists \varepsilon \in (0,1),\, \forall n \in \mathbb{N},\, \exists z_n \in Z,\, d(z_n, z) < 1/n,\, |\hat{\varphi}(y, z_n) - \hat{\varphi}(y, z)| > \varepsilon \right\} \\
&\quad = \left\{y \in B \middle| \exists j \in \mathbb{N},\, \forall n \in \mathbb{N},\, \exists z_n \in Z,\, d(z_n, z) < 1/n,\, |\hat{\varphi}(y, z_n) - \hat{\varphi}(y, z)| > \varepsilon_j \right\} \\
&\quad = \bigcup_{j \in \mathbb{N}} \bigcap_{n \in \mathbb{N}} \bigcup_{z_n \in Z} \left\{y \in B \middle| d(z_n, z) < 1/n,\, |\hat{\varphi}(y, z_n) - \hat{\varphi}(y, z)| > \varepsilon_j \right\},
\end{aligned}$$

where the second equality follows from $\{\varepsilon_j\}_{j \in \mathbb{N}}$ dense in $(0, 1)$. Since $\hat{\varphi}$ is $\mathcal{M}(Y)_0 \otimes \mathcal{M}(Z)$ measurable, each set $\{y \in B |\, d(z, z^*) < 1/n,\, |\hat{\varphi}(y, z) - \hat{\varphi}(y, z^*)| > \varepsilon_j\}$ is in $\mathcal{M}(Y)_0$. Therefore, $\{y \in B | z \in \xi(y)\}$ is $\mathcal{M}(Y)_0$-measurable and, by Lemma 15, ξ has a $\mathcal{M}(Y)_0$-measurable selection $\hat{z}^*(y)$. Define the correspondence

$$\xi_\varepsilon(y) = \left\{\varepsilon \in \{\varepsilon_j\}_{j \in \mathbb{N}} \middle| \forall n \in \mathbb{N}, \exists z_n \in Z,\, d(z_n, \hat{z}^*(y)) < 1/n,\, |\hat{\varphi}(y, z_n) - \hat{\varphi}(y, \hat{z}^*(y))| > \varepsilon \right\}.$$

Since $\hat{z}^*(y) \in \xi(y)$ for each $y \in B$, ξ_ε is non-empty valued. As before, By Lemma 15, ξ_ε has a $\mathcal{M}(Y)_0$-measurable selection, $\hat{\varepsilon}(y)$. Analogously, define the correspondence

$$\xi_n(y) = \left\{z \in Z \middle| d(z, z^*(y)) < 1/n,\, |\hat{\varphi}(y, z) - \hat{\varphi}(y, z^*(y))| > \hat{\varepsilon}(y) \right\}.$$

Since $\hat{z}^*(y) \in \xi(y)$ and $\hat{\varepsilon}(y) \in \xi_\varepsilon(y)$, ξ_n is non-empty valued. As before, $\xi_n(y)$ has a $\mathcal{M}(Y)_0$-measurable selection, $\hat{z}_n(y)$.

Now, for each $z \in Z$, define the set $B_z := \{y \in B | \hat{z}^*(y) = z\}$. Since $\cup_{z\in Z} B_z = B$ and $\hat{\beta}(B) > 0$ then there is $\hat{z}$ such that $\hat{\beta}(B_{\hat{z}}) > 0$. Define $\hat{B}^{n,+}_z := \{y \in B_{\hat{z}} | \hat{z}_n(y) = z,\, \hat{\varphi}(y,z) - \hat{\varphi}(y,\hat{z}) > 0\}$, and $\hat{B}^{n,-}_z := \{y \in B_{\hat{z}} | \hat{z}_n(y) = z,\, \hat{\varphi}(y,z) - \hat{\varphi}(y,\hat{z}) < 0\}$. Let $\breve{\varepsilon} := \int_{B_{\hat{z}}} \hat{\varepsilon}(y)\, d\hat{\beta}(y)$. By construction $\breve{\varepsilon} > 0$.

We have, for every $n \in \mathbb{N}$, there is a finite subset of Z, denoted Z_n, such that

$$\left\|(\mu - \mu^{\hat{z}})|_{B(\hat{z},1/n)}\right\|_{SV} \geq \sum_{z\in Z_n} \left(\int_{\hat{B}^{n,+}_z} \left(\hat{\varphi}(y,z) - \hat{\varphi}(y,\hat{z})\right) d\hat{\beta}(y) + \int_{\hat{B}^{n,-}_z} \left(\hat{\varphi}(y,\hat{z}) - \hat{\varphi}(y,z)\right) d\hat{\beta}(y) \right)$$
$$\geq \int_{B_{\hat{z}}} \left|\hat{\varphi}(y,\hat{z}_n(y)) - \hat{\varphi}(y,\hat{z})\right| d\hat{\beta}(y) - \breve{\varepsilon}/2 \geq \int_{B_{\hat{z}}} \hat{\varepsilon}(y)\, d\hat{\beta}(y) - \breve{\varepsilon}/2 = \breve{\varepsilon}/2.$$

However, since μ is continuous in z in the strong total variation norm, $\left\|(\mu - \mu^{\hat{z}})|_{B(\hat{z},1/n)}\right\|_{SV}$ ought to converge to zero as $n \to \infty$. This is a contradiction.

Sufficiency. For each $n \in \mathbb{N}$, $\hat{\hat{\varepsilon}}_n(y) := \sup_{z\in B(z^*,1/n)} |\varphi(y,z) - \varphi(y,z^*)|$ is a measurable function and converges to zero almost surely in y. The conclusion follows by the dominated convergence theorem as $\left\|(\mu - \mu^{z^*})|_{B(z^*,1/n)}\right\|_{SV} \leq \int_Y \hat{\hat{\varepsilon}}_n(y)\, d\beta(y)$.

Part 2. Notice that $\bar{\varphi}(y) := \sup_{z\in Z} \varphi(z,y)$ is measurable and belongs to $L^1(Y,\beta)$ if and only if $\|\mu\|_{SV}$ is finite. □

References

Aliprantis, C. D. and K. C. Border (2006): *Infinite Dimensional Analysis: A Hitchhikers Guide*, 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.

Angeletos, G.-M., C. Hellwig, and A. Pavan (2007): "Dynamic global games of regime change: Learning, multiplicity, and the timing of attacks," *Econometrica*, 75, 711–756.

Aoyagi, M. (2003): "Bid rotation and collusion in repeated auctions," *Journal of Economic Theory*, 112, 79–105.

Ash, R. B. (2014): *Real analysis and probability*, Academic press.

Athey, S. and K. Bagwell (2008): "Collusion with persistent cost shocks," *Econometrica*, 76, 493–540.

Athey, S., K. Bagwell, and C. Sanchirico (2004): "Collusion and price rigidity," *The Review of Economic Studies*, 71, 317–349.

Aumann, R. J., M. Maschler, and R. E. Stearns (1995): *Repeated games with incomplete information*, MIT press.

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.

Bogachev, V. I. (2007): *Measure theory*, vol. 2, Springer.

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

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

Doval, L. and J. C. Ely (2020): "Sequential information design," *Econometrica*, 88, 2575–2608.

Duffie, D., J. Geanakoplos, A. Mas-Colell, and A. McLennan (1994): "Stationary markov equilibria," *Econometrica*, 745–781.

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

Ely, J. C. (2017): "Beeps," *American Economic Review*, 107, 31–53.

Fudenberg, D. and D. Levine (1983): "Subgame-perfect equilibria of finite- and infinite-horizon games," *Journal of Economic Theory*, 31, 251–268.

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*, 507–544.

He, W. and Y. Sun (2017): "Stationary Markov perfect equilibria in discounted stochastic games," *Journal of Economic Theory*, 169, 35–61.

——— (2020): "Dynamic games with (almost) perfect information," *Theoretical Economics*, 15, 811–859.

Hellman, Z. (2014): "A game with no Bayesian approximate equilibria," *Journal of Economic Theory*, 153, 138–151.

Himmelberg, C. and F. Van Vleck (1975): "Multifunctions with values in a space of probability measures," *Journal of Mathematical Analysis and Applications*, 50, 108–112.

Jofre-Bonet, M. and M. Pesendorfer (2003): "Estimation of a dynamic auction game," *Econometrica*, 71, 1443–1489.

Kamada, Y. and M. Kandori (2020): "Revision games," *Econometrica*, 88, 1599–1630.

Kamada, Y. and S. Moroni (2018): "Games with Private Timing," *Working paper*.

Kreps, D. and G. Ramey (1987): "Structural Consistency, Consistency, and Sequential Rationality," *Econometrica*, 55, 1331–1348.

Levy, Y. (2013): "Discounted stochastic games with no stationary Nash equilibrium: two examples," *Econometrica*, 81, 1973–2007.

Levy, Y. J. and A. McLennan (2015): "Corrigendum to "Discounted stochastic games with no stationary Nash equilibrium: two examples"," *Econometrica*, 83, 1237–1252.

Manelli, A. M. (1996): "Cheap talk and sequential equilibria in signaling games," *Econometrica*, 917–942.

Matsui, A. (1989): "Information leakage forces cooperation," *Games and Economic Behavior*, 1, 94–115.

Milgrom, P. R. and R. J. Weber (1985): "Distributional strategies for games with incomplete information," *Mathematics of Operations Research*, 10, 619–632.

Morris, S. and H. S. Shin (1998): "Unique equilibrium in a model of self-fulfilling currency attacks," *American Economic Review*, 587–597.

Myerson, R. B. and P. J. Reny (2020): "Perfect Conditional ε-Equilibria of Multi-Stage Games With Infinite Sets of Signals and Actions," *Econometrica*, 88, 495–531.

Nowak, A. S. (2003): "On a new class of nonzero-sum discounted stochastic games having stationary Nash equilibrium points," *International Journal of Game Theory*, 32, 121.

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.

Parthasarathy, T. and S. Sinha (1989): "Existence of stationary equilibrium strategies in non-zero sum discounted stochastic games with uncountable state space and state-independent transitions," *International Journal of Game Theory*, 18, 189–194.

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

Renault, J., E. Solan, and N. Vieille (2013): "Dynamic sender–receiver games," *Journal of Economic Theory*, 148, 502–534.

Selten, R. (1975): "Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games," *International Journal of Game Theory*, 4, 25–55.

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*, 1171–1214.

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

Simon, R. S. and G. Tomkowicz (2018): "A Bayesian Game without ε-equilibria," *Israel Journal of Mathematics*, 227, 215–231.

Skrzypacz, A. and H. Hopenhayn (2004): "Tacit collusion in repeated auctions," *Journal of Economic Theory*, 114, 153–169.

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

B Supplemental Appendix — Omitted Proofs

B.1 Incorporating Active and Inactive Players

In the remainder of the paper, we enrich the class of games we consider by allowing for active and inactive players. In particular, the function $\mathcal{T}_i : \cup_{t\in\mathbb{N}}\Omega^t \times X^{t-1} \to \{0,1\}$ is measurable and determines if player $i \in N$ is *active* or *inactive* in period $t \in \mathbb{N}$, $\mathcal{T}_i(\cdot) = 1$ and $\mathcal{T}_i(\cdot) = 0$, respectively, as a function of the history of the states of the world and action profiles. We assume that at least one player is active at each period.

Whenever *active*, players move after receiving informative signals about the history of the state of the world and action profiles. Specifically, for $t \in \mathbb{N}$, after every history $(\omega^t, a^{t-1}) \in \Omega^t \times X^{t-1}$ with $\mathcal{T}_i(\omega^t, a^{t-1}) = 1$, player $i \in N$ observes a signal $s_{i,t} \in S_i$ and chooses an action among the available ones according to A_i. When $\mathcal{T}_i(\omega^t, a^{t-1}) = 0$, player i is *inactive* and does not observe private signals nor chooses any non-trivial action. We assume S_i contains the null signal s_* and X_i the null action a_*, the latter an isolated point of X_i, and, when $\mathcal{T}_i(\omega^t, a^{t-1}) = 0$, we assume that $\gamma_i(\omega^t, a^{t-1}) = s_*$ and $A_i((s_i^{t-1}, s_*), a_i^{t-1}) = \{a_*\}$, for any history of private signals $s_i^{t-1} \in S_i^{t-1}$.

Without loss, we assume player $i \in N$ receives a non-zero flow payoff only when active, i.e., $g_i(\omega^t, (a^{t-1}, a_t)) = 0$ if $\mathcal{T}_i(\omega^t, a^{t-1}) = 0$ for every $(\omega^t, a^{t-1}, a_t) \in \Omega^t \times X^t$.

Private histories. To accommodate for the presence of active and inactive players, we modify our definition of private history. We start by defining an auxiliary notion.

For every period $t \in \mathbb{N}$, $(s_i^t, a_i^{t-1}) \in S_i^t \times X_i^{t-1}$ is a *signal-action history* if $a_{i,\ell}^{t-1} = a_*$ whenever $s_{i,\ell}^t = s_*$ for $\ell \leq t-1$. The set of signal-action histories contained in $S_i^t \times X_i^{t-1}$ is denoted by $\mathcal{H}_{i,*}^t$ and the set of all signal-action histories is $\mathcal{H}_{i,*} := \cup_{t\in\mathbb{N}}\mathcal{H}_{i,*}^t$.

A *generalized private history* of player i is a function of a signal-action history $(s_i^t, a_i^{t-1}) \in \mathcal{H}_{i,*}$ as follows

$$\hbar_i(s_i^t, a_i^{t-1}) := \big((s_{i,\ell}^t)_{\ell\leq t,\, s_{i,\ell}^t \neq s_*}, (a_{i,\ell}^{t-1})_{\ell\leq t-1,\, s_{i,\ell}^t\neq s_*}\big).$$

Generalized private history $\hbar_i(s_i^t, a_i^{t-1})$ collects player i's signals up to t and actions up to $t-1$, describing the information available to player i before playing at period t. Notice that, private histories do not record any information from periods in which players are inactive. For this reason, a player may be oblivious to the number of moves that have occurred, e.g., for any $(s_i^t, a_i^{t-1}) \in \mathcal{H}_{i,*}$, we have $\hbar_i(s_i^t, a_i^{t-1}) = \hbar_i((s_i^t, s_*), (a_i^{t-1}, a_*))$.

For the remainder of the paper, we redefine $\mathscr{H}_i := \cup_{h_{i,*}\in\mathscr{H}_{i,*}} \hbar_i(h_{i,*})$ as the set of player i's generalized private histories instead of the set of private histories. The *length* of a generalized private history $h_i = \hbar_i(s_i^t, a_i^{t-1})$ is $|h_i| := |\{\ell \le t |\, s_{i,\ell}^t \neq s_*\}|$; for $\ell \le |h_i|$, $h_i^{(\ell)}$ is the *truncation* of h_i up to and including player i's ℓ'th active period, so that $|h_i^{(\ell)}| = \ell$. Let $\mathscr{H}_i^t := \{h_i \in \mathscr{H}_i || h_i| = t\}$ be the set of i's generalized private histories of length $t \in \mathbb{N}$.

We assume the correspondence A_i is measurable with respect to $\mathcal{M}(\mathscr{H}_{i,*}|\hbar_i)$, the sub σ-algebra of $\mathcal{M}(\mathscr{H}_{i,*})$ induced by $\hbar_i$, and we write $A_i : \mathscr{H}_i \rightrightarrows X_i$.[1]

Strategies. For every $i \in N$, $\mathscr{H}_{A_i}$ and $\sigma_i : \mathscr{H}_{A_i} \to \Delta(X_i)$ are now formally defined in terms of generalized private histories.

Conditional measures. The definition of conditional measures is generalized as follows. Conditional on player i's actions $a_i^\tau \in X_i^\tau$ and signals $s_i^t \in S_i^t$, for $t, \tau \in \mathbb{N} \cup \{0\}$ with $\tau \le t$, the strategy σ_i *induces a transition probability* over player i's action history $a_i^t \in X_i^t$ as

$$p_i(a_i^t | s_i^t, a_i^\tau, \sigma_i) := \prod_{\substack{\ell=\tau+1,\dots,t,\\ s_{i,\ell}^t \neq s_*}} \sigma_i\big(a_{i,\ell}^t | \hbar_i(s_i^{t,(\ell)}, a_i^{t,(\ell-1)})\big) \tag{8}$$

if $(s_i^t, a_i^{t,(t-1)}) \in (\hbar_i)^{-1}(\mathscr{H}_{A_i})$, $a_i^{t,(\tau)} = a_i^\tau$; zero otherwise.

Expected payoffs. The definitions of player's $i \in N$ expected payoff and continuation expected payoff apply verbatim to this more general environment. Furthermore, Lemma 4 still holds.

Constrained equilibria. We adapt the notions of $\tilde{\varepsilon}$-constrained and $\tilde{\varepsilon}$-constrained conditional equilibria by defining the objects in Section 3.1 in terms of generalized private histories instead of private histories.

For every $i \in N$, $\varepsilon > 0$, $\tilde{\varepsilon}_i : \{(h_i, a_i) | h_i \in \mathscr{H}_i, a_i \in X_i\} \to (0,1)$ is now formally defined in terms of generalized private histories.

Let $\mathscr{H}^{\mathbb{N}} := \Omega^{\mathbb{N}} \times X^{\mathbb{N}}$ be the set of histories with infinite length.[2] Define $\hat{\hbar}^t : \mathscr{H}^{\mathbb{N}} \to \mathscr{H}^t$ and $\hat{\hbar}_i^t : \mathscr{H}^{\mathbb{N}} \to \mathscr{H}_i$, for $i \in N$, functions from infinite histories to period-t histories and player's i private histories, respectively, by setting $\hat{\hbar}^t(\bar{\omega}, \bar{a}) = (\bar{\omega}^{(t)}, \bar{a}^{(t-1)})$ and $\hat{\hbar}_i^t(\bar{\omega}, \bar{a}) = \hbar_i(\gamma_i^t(\bar{\omega}^{(t)}, \bar{a}^{(t-1)}), \bar{a}_i^{(t-1)})$, for $(\bar{\omega}, \bar{a}) \in \mathscr{H}^{\mathbb{N}}$. Notice that $\hat{\hbar}_i^t$ is now defined over histories of infinite length.

[1] The sub σ-algebra of $\mathcal{M}(\mathscr{H}_{i,*})$ induced by $\hbar_i$ is $\mathcal{M}(\mathscr{H}_{i,*}|\hbar_i) := \{\hbar_i^{-1}(A) | A \in \mathcal{M}(\mathscr{H}_i)\}$. For $i \in N$, we can write $A_i : \mathscr{H}_i \rightrightarrows X_i$ by Theorem 4.41 in Aliprantis and Border (2006).

[2] For any set Y, denote $Y^{\mathbb{N}} := \{(y_t)_{t\in\mathbb{N}} | y_t \in Y, \forall t \in \mathbb{N}\}$.

We define $P(\cdot|\sigma) \in \Delta(\mathcal{H}^{\mathbb{N}})$ as the measure over infinite histories induced by the strategy profile $\sigma \in \Sigma$.[3] For every $i \in N$, $\sigma \in \Sigma$, $\ell \in \mathbb{N}$, and $Z \in \mathcal{M}(\mathcal{H}_i^\ell)$, the probability that player i's private history belongs to Z under σ is defined in the general case as $P_i(Z|\sigma) := P(\cup_{t\in\mathbb{N}}(\hat{h}_i^t)^{-1}(Z)|\sigma)$. If $P_i(Z|\sigma) > 0$, player i's period-t *beliefs* over $C \in \mathcal{M}(\mathcal{H}^t)$ *conditional* on Z can be defined as

$$P_i^t(C|Z,\sigma) = \frac{P\big((\hat{h}^t)^{-1}(C) \cap (\hat{h}_i^t)^{-1}(Z)|\sigma\big)}{P_i(Z|\sigma)}.$$

For every pair of strategy profiles $\sigma, \hat{\sigma} \in \Sigma$, and $Z \in \mathcal{M}(\mathcal{H}_i^\ell)$, for $\ell \in \mathbb{N}$, player i's *expected payoff conditional on Z from σ and $\hat{\sigma}$* is defined as

$$U_i(\hat{\sigma}|Z,\sigma) := \sum_{\tau\in\mathbb{N}} \int_{\mathcal{H}^\tau} U_i(\hat{\sigma}|\omega^\tau, a^{\tau-1})\, dP_i^\tau(\omega^\tau, a^{\tau-1}|Z,\sigma).$$

Finally, a strategy profile $\sigma \in \Sigma(\tilde{\varepsilon})$, for $\tilde{\varepsilon} \in \mathcal{E}$, is an *$\tilde{\varepsilon}$-constrained conditional equilibrium*, if, for every $i \in N$, $\ell \in \mathbb{N}$ and $Z \in \mathcal{M}(\mathcal{H}_i^\ell)$ satisfying $P_i(Z|\sigma) > 0$,

$$U_i(\sigma|Z,\sigma) \geq U_i(\sigma_i', \sigma_{-i}|Z,\sigma) \qquad \forall \sigma_i' \in \Sigma_i(\tilde{\varepsilon}).$$

Trembling hand perfect equilibrium. THPE adapts to this more general environment with minor adjustments. For every $i \in N$, $t \in \mathbb{N}$, the function a_i is now defined over $(s_i^t, a_i^{t,(t-1)}) \in h_i^{-1}(\mathcal{H}_{A_i})$, for any given $\sigma_i \in \Sigma_i$ as follows.

$$a_i(s_i^t, a_i^t; \sigma_i) = \big\{a_i^{t,(\tau)} | \tau = \min\big\{\hat{\tau} \leq t-1 | \Pi_{\ell \geq \hat{\tau}+1}^{t-1} \sigma_i(a_{i,\ell}^t | h_i(s_i^{t,(\ell)}, a_i^{t,(\ell-1)})) > 0\big\}\big\}.$$

Furthermore, the limit strategy σ^* in Definition 3(ii) is retrieved as follows

$$\sigma_i^*\big(a_{i,t}^t | h_i(s_i^t, a_i^{t,(t-1)})\big) = \frac{p_i^*\big(a_i^t | s_i^t, a_i^{t,(\tau)}\big)}{p_i^*\big(a_i^{t,(t-1)} | s_i^{t,(t-1)}, a_i^{t,(\tau)}\big)}$$

for $(s_i^t, a_i^{t,(t-1)}) \in h_i^{-1}(\mathcal{H}_{A_i})$, and $a_i^{t,(\tau)} = a_i(s_i^t, a_i^t; \sigma_i)$, for every $i \in N$, $t \in \mathbb{N}$.

We generalize the notions of root of a proper subgame and negligible histories as well. We say that a history $(\omega^t, a^{t-1}) \in \mathcal{H}$ is *the root of a proper subgame* if, for every $(\tilde{\omega}^\tau, \tilde{a}^{\tau-1}) \in \mathcal{H}$, if $h_i(\gamma_i^t(\omega^t, a^{t-1}), a_i^{t-1}) = h_i(\gamma_i^\tau(\tilde{\omega}^\tau, \tilde{a}^{\tau-1}), \tilde{a}_i^{\tau-1})$ for all $i \in N$, then $(\omega^t, a^{t-1}) = (\tilde{\omega}^\tau, \tilde{a}^{\tau-1})$. Let $\mathcal{H}^\varnothing \subseteq \mathcal{H}$ denote the set of all such histories. A set $H \in \mathcal{M}(\mathcal{H})$ is *negligible* if $P(\cup_{t\in\mathbb{N}}(\hat{h}^t)^{-1}(H \cap \mathcal{H}^t)|\sigma) = 0$ for every $\sigma \in \Sigma$.

[3]For each $\sigma \in \Sigma$, there is a probability measure $P(\cdot|\sigma) : \mathcal{H}^{\mathbb{N}} \to [0,1]$ such that for every $t \in \mathbb{N}$, and measurable, bounded functions $r^t : \mathcal{H}^t \to \mathbb{R}$, $\hat{r} : \mathcal{H}^{\mathbb{N}} \to \mathbb{R}$, defined as $\hat{r}(\bar{\omega}, \bar{a}) := r^t(\hat{h}^t(\bar{\omega}, \bar{a}))$, where $(\bar{\omega}, \bar{a}) \in \mathcal{H}^{\mathbb{N}}$, we have that $\mathbb{E}_{P(\cdot|\sigma)}(\hat{r}) = \mathbb{E}_{P^t(\cdot|\sigma)}(\hat{r})$. See Theorem 4.49 in Pollard (2002).

Markov games. For every $i \in N$, active and inactive players are determined by the current payoff-relevant state, $\mathcal{T}_i : \Omega^R \to \{0,1\}$. Furthermore, for every $\sigma \in \Sigma$, $t \in \mathbb{N}$, the objects $\mu^R_{\omega,t}(\cdot|\cdot,\sigma) : \cup_{i\in N}\mathcal{H}_i \to \Delta(\Omega^R)$, $\sigma_i \in \Sigma^M_i$, and $\mathcal{s}^R_i$, and the Markov information condition are formally defined in terms of generalized private histories.

B.2 Additional Notation

For any set Y and Z, define $Y^\infty := \cup_{t\in\mathbb{N}} Y^t$ and $(X \times Z)^\infty := \cup_{t\in\mathbb{N}} X^t \times Z^t$.

For $i \in N, t \in \mathbb{N}$, $(s^t_i, a^t_i) \in S^t_i \times X^t_i$ is a signal-action history *up to and including period t-action* if $a^t_{i,\ell} = a_*$ whenever $s^t_{i,\ell} = s_*$ for $\ell \le t$. The set of signal-action histories up to and including period t is $\bar{\mathcal{H}}^t_{i,*}$, and $\bar{\mathcal{H}}_{i,*} := \cup_{t\in\mathbb{N}}\bar{\mathcal{H}}^t_{i,*}$. A private history of player i *up to and including period t-action* is the function of $(s^t_i, a^t_i) \in \bar{\mathcal{H}}_{i,*}$ defined by $\bar{h}_i(s^t_i, a^t_i) := ((s^t_{i,\ell})_{\ell\le t,\, s^t_{i,\ell}\neq s_*}, (a^t_{i,\ell})_{\ell\le t,\, s^t_{i,\ell}\neq s_*})$. The length of private history $\bar{h}_i = \bar{h}_i(s^t_i, a^t_i)$ is $|\bar{h}_i| := |h_i(s^t_i, a^t_i)|$. We define $\bar{\mathcal{H}}_i := \cup_{\bar{h}_{i,*}\in\bar{\mathcal{H}}_{i,*}}\bar{h}_i(\bar{h}_{i,*})$, and $\bar{\mathcal{H}}^\ell_i := \{\bar{h}_i \in \bar{\mathcal{H}}_i | |\bar{h}_i| = \ell\}$. We denote by $\bar{h}^{(\ell)}_i$ the truncation of $\bar{h}_i$ up to and including the ℓ'th active period of player i. The sub σ-algebra $\mathcal{M}(\mathcal{H}_{i,*}|\bar{h}_i)$ is defined analogously to $\mathcal{M}(\mathcal{H}_{i,*}|h_i)$. Let $\bar{\mathcal{H}}_{A_i} := \{(s^t_i, a^t_i) \in \bar{\mathcal{H}}_i |\, a^t_{i,\ell} \in A_i(s^{t,(\ell)}_i, a^{t,(\ell-1)}_i),\ \ell \le t\}$, be the set of private histories in $\bar{\mathcal{H}}_i$ that are available according to A_i.

For $i \in N$, let $\bar{X}_i := X_i \cup \{\bar{a}_\star\}$ for some $\bar{a}_\star \notin X_i$, isolated point in $\bar{X}_i$. Define $A_{i,*} : \cup_{t\in\mathbb{N}} S^t_i \times X^{t-1}_i \rightrightarrows X_i$ as $A_i(h_i(s^t_i, a^{t-1}_i))$ if $(s^t_i, a^{t-1}_i) \in \mathcal{H}_{i,*}$ and $s^t_{i,t} \neq s_*$; as $\{a_*\}$ if $(s^t_i, a^{t-1}_i) \in \mathcal{H}_{i,*}$ and $s^t_{i,t} = s_*$; as $A_{i,*}(s^t_i, a^{t-1}_i) = \{\bar{a}_\star\}$ otherwise. Since A_i is weakly measurable, $A_{i,*}$ is weakly measurable. Define $\mathcal{H}_{A_{i,*}} := (h_i)^{-1}(\mathcal{H}_{A_i})$ and $\bar{\mathcal{H}}_{A_{i,*}} := (\bar{h}_i)^{-1}(\bar{\mathcal{H}}_{A_i})$, and, for $t \in \mathbb{N}$, $\mathcal{H}^t_{A_{i,*}} := \mathcal{H}_{A_{i,*}} \cap \mathcal{H}^t_{i,*}$ and $\bar{\mathcal{H}}^t_{A_{i,*}} := \bar{\mathcal{H}}_{A_{i,*}} \cap \bar{\mathcal{H}}^t_{i,*}$.

Denote by $\mathcal{H}^t_* := \{(\omega^t, a^{t-1}) \in \mathcal{H}^t : (\gamma^t_i(\omega^t, a^{t-1}), a^{t-1}_i) \in \mathcal{H}_{i,*}, \forall i \in N\}$ period-$t \in \mathbb{N}$ histories whose projection onto $S^t_i \times X^{t-1}_i$ is a signal-action history for every $i \in N$, and let $\mathcal{H}_* := \cup_{t\in\mathbb{N}}\mathcal{H}^t_*$. For $i \in N$, $h^h_i : \mathcal{H}_* \to \mathcal{H}_i$ maps histories to player's i private histories, by setting $h^h_i(\omega^t, a^{t-1}) = h_i(\gamma^t_i(\omega^t, a^{t-1}), a^{t-1}_i)$, for $t \in \mathbb{N}$, $(\omega^t, a^{t-1}) \in \mathcal{H}^t_*$.

For every $t \in \mathbb{N}$, by combining the counting measure on action profiles in X^{t-1} and μ^t_ω, we define the unconditional measure $\mu^t_{\omega,a}$ over $\mathcal{H}^t$ by setting $\mu^t_{\omega,a}(B) := \sum_{a^{t-1}\in X^{t-1}} \int_{\Omega^t} \mathbb{1}\{(\omega^t, a^{t-1}) \in B\}\, d\mu^t_\omega(\omega^t|a^{t-1})$ for $B \in \mathcal{M}(\mathcal{H}^t_*)$.[4] We define the conditional measure $\mu^t_{\omega,a}(\cdot|\omega^{t-1}, a^{t-2})$ over $\mathcal{H}^t$ and the measure $\mu^t_{\omega,a}$ over $\Omega^t \times X^t$ analogously.

[4]The counting measure "counts one" for every element of a set. Therefore, if p is a counting measure over a countable set B and $f : B \to \mathbb{R}$, $\int_B f(b)dp(b) = \sum_{b\in B} f(p)$.

B.3 Proofs of Theorem 1 and 2

SAC, in conjunction with Proposition 5, allows us to write the signal profile transition for every $t \in \mathbb{N}, s^t \in S^t$, $a^{t-1} \in X^{t-1}$ as

$$d\mu_s^t(s^t|a^{t-1}) = f^t(s^t, a^{t-1}) \cdot d\nu^t(s^t), \tag{9}$$

where $f^t \in CI(S^t \times X^{t-1}, \nu^t)$, and $\nu^t(s^t) = \prod_{i \in N} \nu_i^t(s_i^t)$ is the product of reference measures (see footnote 25) defined as $\nu_i^t(s_i^t) := \prod_{\ell \leq t} \nu_{i,\ell}(s_{i,\ell}^t | s_i^{t,(\ell-1)})$.

Let $\delta < 1$. For each $i \in N$, we define $\alpha_i \in \Delta(S_i^\infty)$ as

$$\alpha_i(B) := \frac{1-\delta}{\delta} \sum_{t \in \mathbb{N}} \delta^t \cdot \nu_i^t(B \cap S_i^t) \quad \text{for every } B \in \mathcal{M}(S_i^\infty).$$

Recall, player i's strategy $\sigma_i \in \Sigma_i$ is a transition probability from $\mathcal{H}_{A_i}$ to probabilities over available actions. We define $\hat{\sigma}_i : \cup_{t \in \mathbb{N}} S_i^t \times \bar{X}_i^{t-1} \to \Delta(\bar{X}_i)$, the extension of σ_i to $\cup_{t \in \mathbb{N}} S_i^t \times \bar{X}_i^{t-1}$, as the strategy such that $\hat{\sigma}_i(\cdot|s_i^t, a_i^{t-1})$ has support in $A_{i,*}(s_i^t, a_i^{t-1})$, and coincides with σ_i whenever $(s_i^t, a_i^{t-1}) \in \mathcal{H}_{A_i,*}$. Since the set $\mathcal{H}_{A_i,*}$ is measurable,[5] an extended strategy is a transition probability from $\cup_{t \in \mathbb{N}} S_i^t \times \bar{X}_i^{t-1}$ to $\bar{X}_i$. We denote the set of extended strategies by $\hat{\Sigma}_i$, and set of i's extended strategies that extend an $\tilde{\varepsilon}$-constrained strategy is $\hat{\Sigma}_i(\tilde{\varepsilon})$.

For $i \in N$, let $\mathcal{P}_i$ be the set of transition probabilities from S_i^∞ to $\bar{X}_i^\infty$. For each $\hat{p} \in \mathcal{P}_i$, we define the functional

$$I_{i,\varphi_i}(\hat{p}) := \int_{S_i^\infty} \sum_{\bar{a}_i \in \bar{X}_i^\infty} \varphi_i(\bar{s}_i, \bar{a}_i) \cdot \hat{p}(\bar{a}_i|\bar{s}_i), d\alpha_i(\bar{s}_i),$$

where $\varphi_i \in CI(S_i^\infty \times \bar{X}_i^\infty, \alpha_i)$. The *weak topology* on $\mathcal{P}_i$ is the coarsest topology such that for each $\varphi_i \in CI(S_i^\infty \times \bar{X}_i^\infty, \alpha_i)$ the functional $I_{i,\varphi_i}(\hat{p})$ is continuous in $\hat{p} \in \mathcal{P}_i$.

Lemma 5. *For each $j \in 1, 2$, let $(\phi_j^\lambda)_{\lambda \in \Lambda}$ be a net of transition probabilities in $\mathcal{P}_i$ converging to ϕ_j^*. If there exist functions $\varphi_1, \varphi_2 \in CI(S_i^\infty \times \bar{X}_i^\infty, \alpha_i)$ such that $\phi_1^\lambda \cdot \varphi_1 + \phi_2^\lambda \cdot \varphi_2 \geq 0$ for every $\lambda \in \Lambda$, then for any closed set $C \in \mathcal{M}(\bar{X}_i^\infty)$:*

$$\sum_{\bar{a}_i \in C} (\phi_1^*(\bar{s}_i, \bar{a}_i) \cdot \varphi_1(\bar{s}_i, \bar{a}_i) + \phi_2^*(\bar{s}_i, \bar{a}_i) \cdot \varphi_2(\bar{s}_i, \bar{a}_i)) \geq 0,$$

α_i-almost surely.

[5] The set $\mathcal{H}_{A_i}$ is measurable by Theorem 18.6 in Aliprantis and Border (2006).

Proof. Let $\hat{\phi} : S_i^\infty \times \bar{X}_i^\infty \to \mathbb{R}$ be of the form $\hat{\phi}(\bar{s}_i, \bar{a}_i) = \mathbb{1}\{\bar{s}_i \in \tilde{S}_i\} \cdot \tilde{\phi}(\bar{a}_i)$ for some $\tilde{S}_i \in \mathcal{M}(\bar{S}_i^\infty)$ and some bounded and continuous function $\tilde{\phi} : \bar{X}_i^\infty \to \mathbb{R}$. Since $\hat{\phi} \in CI(S_i^\infty \times \bar{X}_i^\infty, \alpha_i)$, we have

$$\int_{\tilde{S}_i} \sum_{\bar{a}_i \in \bar{X}_i^\infty} \tilde{\phi}(\bar{a}_i) \cdot \big(\phi_1^\lambda(\bar{s}_i, \bar{a}_i) \cdot \varphi_1(\bar{s}_i, \bar{a}_i) + \phi_2^\lambda(\bar{s}_i, \bar{a}_i) \cdot \varphi_2(\bar{s}_i, \bar{a}_i)\big)\, d\alpha_i(\bar{s}_i)$$

$$\to \int_{\tilde{S}_i} \sum_{\bar{a}_i \in \bar{X}_i^\infty} \tilde{\phi}(\bar{a}_i) \cdot \big(\phi_1^*(\bar{s}_i, \bar{a}_i) \cdot \varphi_1(\bar{s}_i, \bar{a}_i) + \phi_2^*(\bar{s}_i, \bar{a}_i) \cdot \varphi_2(\bar{s}_i, \bar{a}_i)\big)\, d\alpha_i(\bar{s}_i) \geq 0.$$

Let $U_n = \cup_{c \in C} B(c, 1/n)$ where be $B(c, 1/n)$ is the ball centered at $c \in C$ with radius $1/n$. Since X_i is a compact Hausdorff space, by Theorem 2.46 (Urysohn's lemma) and 2.48 in Aliprantis and Border (2006), for each $n \in \mathbb{N}$, there is a continuous function $f_n : X_i \to [0, 1]$ that is equal to 1 in C and it is equal to zero in the complement of U_n. Since f_n converges to $\mathbb{1}\{\bar{a}_i \in C\}$ almost surely in the counting measure in X_i, by setting $\tilde{\phi} = f_n$ in the right hand side of the previous expression, taking the limit in n, and applying the dominated convergence theorem, we obtain $\int_{\tilde{S}_i} \sum_{\bar{a}_i \in C} \big(\phi_1^*(\bar{s}_i, \bar{a}_i) \cdot \varphi_1(\bar{s}_i, \bar{a}_i) + \phi_2^*(\bar{s}_i, \bar{a}_i) \cdot \varphi_2(\bar{s}_i, \bar{a}_i)\big)\, d\alpha_i(\bar{s}_i) \geq 0$, $\tilde{S}_i \in \mathcal{M}(\bar{S}_i^\infty)$. Finally, since $\tilde{S}_i$ is arbitrary, this implies the α_i-almost surely inequality. □

For every $i \in N$, $\hat{\sigma}_i \in \hat{\Sigma}_i$, $t \in \mathbb{N}$, $\tau \in \mathbb{N} \cup \{0\}$, $a_i^\tau \in \bar{X}_i^\tau$ we define the transition probability $p_i(\cdot|\cdot, a_i^\tau, \hat{\sigma}_i) : S_i^\infty \to \Delta(\bar{X}_i^\infty)$ as

$$p_i(a_i^t | s_i^t, a_i^\tau, \hat{\sigma}_i) := \begin{cases} \prod_{j=\tau+1,\ldots,t,} \hat{\sigma}_i(a_{i,j}^t | s_i^{t,(j)}, a_i^{t,(j-1)}) & \text{if } \tau < t, a_i^\tau = a_i^{t,(\tau)} \text{ or } \tau \geq t, a_i^{\tau,(t)} = a_i^t \\ 0 & \text{otherwise.} \end{cases}.$$

For $\hat{\sigma}_i \in \hat{\Sigma}_i$, we define $a_i : S_i^t \times X_i^t \to \cup_{\tau \leq t} X_i^\tau$ as $a_i(s_i^t, a_i^t; \hat{\sigma}_i) := \{a_i^{t,(\tau)} | \tau = \min\{\hat{\tau} \leq t - 1 | \Pi_{\ell \geq \hat{\tau}+1}^{t-1} \hat{\sigma}_i(a_{i,\ell}^t | s_i^{t,(\ell)}, a_i^{t,(\ell-1)}) > 0\}\}$.

Lemma 6. *Let $\{\hat{\sigma}_i^\lambda\}_{\lambda \in \Lambda}$ be a net in $\hat{\Sigma}_i$ for some directed set Λ and $i \in N$. Suppose that for each $\tau \in \mathbb{N}$ and $a_i^\tau \in \bar{X}_i^\tau$, the net $\{p_i(\cdot|\cdot, a_i^\tau, \hat{\sigma}_i^\lambda)\}_{\lambda \in \Lambda}$ converges to $p_i^*(\cdot|\cdot, a_i^\tau) \in \mathcal{P}_i$ in the weak topology of $\mathcal{P}_i$. Then:*

(i) There exists $\sigma_i^* \in \Sigma_i$ that satisfies equation (4) for every $(s_i^t, a_i^{t,(t-1)}) \in \mathcal{H}_{A_i,}$.

(ii) If $\tilde{\varepsilon} \in \mathcal{E}$ and $\{\hat{\sigma}_i^\lambda\}_{\lambda \in \Lambda} \subseteq \hat{\Sigma}_i(\tilde{\varepsilon})$, then $\sigma_i^* \in \hat{\Sigma}_i(\tilde{\varepsilon})$.

Proof. We define $\sigma_i^* : S_i^t \times X_i^t \to [0, 1]$ recursively as follows:

- For $s_i \in S_i$, let $\sigma_i^*(a_i | s_i) = p^{*,1}(a_i | s_i, \emptyset)$.

- Suppose $\sigma_i^*(a_{i,\ell}^\ell|s_i^\ell, a_i^{\ell,(\ell-1)})$ has been defined for $(s_i^\ell, a_i^\ell) \in \bar{\mathscr{H}}_{A_i,*}$, for every $\ell \leq t-1$, and let $\hat{\sigma}_i^*$ denote its extension to $\cup_{\ell\leq t-1} S_i^\ell \times \bar{X}_i^{\ell-1}$. For $(s_i^t, a_i^{t,(t-1)}) \in \mathscr{H}_{A_i,*}$, define

$$\sigma_i^*(a_{i,t}^t|s_i^t, a_i^{t,(t-1)}) := \frac{p_i^*(a_i^t|s_i^t, \alpha_i(s_i^t, a_i^t; \hat{\sigma}_i^*))}{p_i(a_i^{t,(t-1)}|s_i^{t,(t-1)}, \alpha_i(s_i^t, a_i^t; \hat{\sigma}_i^*), \hat{\sigma}_i^*)}.$$

Remark 1. *Consider the following:*

1. *Notice that* $\alpha_i(s_i^t, a_i^t; \hat{\sigma}_i^*)$ *and* $p_i(a_i^{t,(t-1)}|s_i^{t,(t-1)}, \alpha_i(s_i^t, a_i^t; \hat{\sigma}_i^*), \hat{\sigma}_i^*)$ *depend only on* $\hat{\sigma}_i^*|_{S_i^\ell \times \bar{X}_i^{\ell-1}}$ *for* $\ell \leq t-1$. *Therefore,* $\hat{\sigma}_i^*|_{S_i^t \times \bar{X}_i^{t-1}}$ *is well-defined (see point 3).*
2. *By the definition of* σ_i^* *we have* $p_i^*(a_i^t|s_i^t, \alpha_i(s_i^t, a_i^t; \hat{\sigma}_i^*)) = p_i(a_i^t|s_i^t, \alpha_i(s_i^t, a_i^t; \hat{\sigma}_i^*), \hat{\sigma}_i^*)$ *for every* $(s_i^t, a_i^t) \in \bar{\mathscr{H}}_{A_i,*}^t$.
3. *By the definition of* α_i, $p_i(a_i^{t,(t-1)}|s_i^{t,(t-1)}, \alpha_i(s_i^t, a_i^t; \hat{\sigma}_i^*), \hat{\sigma}_i^*) > 0$, $\forall (s_i^t, a_i^t) \in S_i^t \times \bar{X}_i^t$.
4. *For every* $\hat{\sigma} \in \hat{\Sigma}$, *the restrictions of* $p_i(a_i^{t,(t-1)}|s_i^{t,(t-1)}, a_i^\tau, \hat{\sigma}_i)$ *and* $\alpha_i(s_i^t, a_i^t; \hat{\sigma}_i)$ *to* $\bar{\mathscr{H}}_{i,*}$, *are measurable with respect to the sub* σ*-algebra of* $\bar{\mathscr{H}}_{i,*}$ *induced by* $\hbar_i$.

Let $\tilde{\varepsilon} \in \mathcal{E}$. The extension of $\tilde{\varepsilon}_i$, $\hat{\tilde{\varepsilon}}_i : (S_i \times \bar{X}_i)^\infty \to \mathbb{R}$, is given by $\tilde{\varepsilon}_i(\hbar_i(s_i^t, a_i^{t-1}), a_i)$ if $(s_i^t, (a_i^{t-1}, a_i)) \in \bar{\mathscr{H}}_{i,*}$; 0, otherwise. Since a_* and $\bar{a}_\star$ are isolated points of $\bar{X}_i$, $\hat{\tilde{\varepsilon}}$ is continuous in (a_i^{t-1}, a_i) for each s_i^t, and, therefore, we have $\hat{\tilde{\varepsilon}}_i \in CI(S_i^\infty \times \bar{X}_i^\infty, \alpha_i)$.

Next we show that **(a.1)** If $(\hat{\sigma}_i^\lambda)_{\lambda\in\Lambda} \subseteq \hat{\Sigma}_i(\tilde{\varepsilon})$ then $\sigma_i^*(a_i|s_i^t, a_i^{t-1}) \geq \tilde{\varepsilon}_i(\hbar_i(s_i^t, a_i^{t-1}), a_i)$ for $(s_i^t, (a_i^{t-1}, a_i)) \in \bar{\mathscr{H}}_{A_i,*}^t$, **(a.2)** $\sigma_i^*(\cdot|s_i^t, a_i^{t-1}) \in \Delta(A_i(s_i^t, a_i^{t-1}))$ for every $(s_i^t, a_i^{t-1}) \in \mathscr{H}_{A_i,*}^t$, and **(b)** $\sigma_i^*|_{\bar{\mathscr{H}}_{A_i,*}}$ is $\mathcal{M}(\bar{\mathscr{H}}_{A_i,*}|\bar{\hbar}_i)$-measurable.

Proof of (a.1) and (a.2). We argue by induction. Let $t \in \mathbb{N}$, and suppose that **(a.1)** and **(a.2)** hold for $\sigma_i^*(\cdot|s_i^{\ell-1}, a_i^{\ell-2})$ with $(s_i^{\ell-1}, a_i^{\ell-2}) \in \mathscr{H}_{A_i,*}^{\ell-1}$, for $\ell \leq t$. Note first that for $\tau \leq t$, and $a_i^\tau \in \bar{X}_i^\tau$, $p_i(a_i^{t-1}|s_i^{t-1}, a_i^\tau, \hat{\sigma}_i^\alpha)$ converges to $p_i^*(a_i^{t-1}|s_i^{t-1}, a_i^\tau)$ in the weak topology of $(S_i^t \times \bar{X}_i^t, \nu_i^t)$. In fact, for any $\phi \in CI(S_i^t \times \bar{X}_i^t, \nu_i^t)$, $\int_{S_i} \sum_{a_i \in \bar{X}_i} \phi(s_i^t, a_i^{t-1}, a_i)\, d\nu_{i,t}(s_i|s_i^{t-1})$ belongs to $CI(S_i^{t-1} \times \bar{X}_i^{t-1}, \nu_i^{t-1})$. In fact, it is measurable with respect to s_i^{t-1} (Theorem 2.6.4 in Ash (2014)), and continuous in (a_i^{t-1}, a_i) for each s_i^{t-1} by the dominated convergence theorem.[6] Therefore, the convergence of $p_i(a_i^{t-1}|s_i^{t-1}, a_i^\tau, \hat{\sigma}_i^\lambda)$ to $p_i^*(a_i^{t-1}|s_i^{t-1}, a_i^\tau)$ in the weak topology of $(S_i^{t-1} \times \bar{X}_i^{t-1}, \nu_i^{t-1})$ yields convergence in the weak topology of $(S_i^t \times \bar{X}_i^t, \nu_i^t)$.

Proof of (a.1). Assume $(\hat{\sigma}_i^\lambda)_{\lambda\in\Lambda} \subseteq \hat{\Sigma}_i(\tilde{\varepsilon})$. By the induction hypothesis, $\alpha_i(s_i^t, a_i^t; \hat{\sigma}_i^*) = \emptyset$ for $(s_i^t, a_i^t) \in \bar{\mathscr{H}}_{A_i,*}$. Now, for $\lambda \in \Lambda$, and $((s_i^{t-1}, s_i), (a_i^{t-1}, a_i)) \in \bar{\mathscr{H}}_{A_i,*}$

$$p_i(a_i^{t-1}|s_i^{t-1}, \emptyset, \hat{\sigma}_i^\lambda)\big(\hat{\sigma}_i^\lambda\left(a_i|(s_i^{t-1}, s_i), a_i^{t-1}\right) - \hat{\tilde{\varepsilon}}_i\left((s_i^{t-1}, s_i, a_i^{t-1}), a_i\right)\big) \geq 0.$$

[6]Fix $a_i \in X_i$. As ϕ is continuous in $a_i^t \in X_i^t$, $(a_i^{t-1,n}, a_i) \to (a_i^{t-1,*}, a_i)$ implies $\phi((s_i^{t-1}, s_i, a_i^{t-1,n}), a_i) \to \phi((s_i^{t-1}, s_i, a_i^{t-1,*}), a_i)$ for every $(s_i^{t-1}, s_i) \in S_i^t$. Furthermore, ϕ is bounded by an $L^1(S_i^t, \nu_i^t)$ function.

By Lemma 5 with $C = \{(a_i^{t-1}, a_i)\}$, ν_i^t-almost surely for $((s_i^{t-1}, s_i), (a_i^{t-1}, a_i)) \in \bar{\mathscr{H}}_{A_i,*}$

$$p_i(a_i^{t-1}|s_i^{t-1}, \emptyset, \hat{\sigma}_i^*)\big(\hat{\sigma}_i^*(a_i|(s_i^{t-1}, s_i), a_i^{t-1}) - \hat{\bar{\varepsilon}}_i\left((s_i^{t-1}, s_i, a_i^{t-1}), a_i\right)\big) \geq 0.$$

By Remark 1.3, this yields the desired result.

Proof of (a.2). Since $\sum_{a_i \in X_i} \hat{\sigma}_i^{\lambda}(a_i|(s_i^{t-1}, s_i), a_i^{t-1}) = 1$ for all $\lambda \in \Lambda$, $((s_i^{t-1}, s_i), a_i^{t-1}) \in \mathscr{H}_{A_i,*}^t$, by Lemma 5, we obtain

$$p_i(a_i^{t-1}|s_i^{t-1}, a_i^{t,(\tau)}, \hat{\sigma}_i^*)\left(\sum_{a_i \in X_i} \hat{\sigma}_i^*(a_i|(s_i^{t-1}, s_i), a_i^{t-1}) - 1\right) = 0,$$

ν_i^t-almost surely for every $((s_i^{t-1}, s_i), a_i^{t-1}) \in \mathscr{H}_{A_i,*}^t$ s.t. $a_i^{t,(\tau)} = a_i((s_i^{t-1}, s_i), a_i^{t-1}; \hat{\sigma}_i^*)$.[7] As this is true for each $\tau \leq t-1$, this shows $\sigma_i^*(\cdot|s_i^t, a_i^{t-1}) \in \Delta(X_i)$ for $(s_i^t, a_i^{t-1}) \in \mathscr{H}_{A_i,*}^t$.

We now show that, $\operatorname{supp} \sigma_i^*(\cdot|s_i^t, a_i^{t-1}) \subseteq A_i(s_i^t, a_i^{t-1})$ for every $(s_i^t, a_i^{t-1}) \in \mathscr{H}_{A_i,*}^t$. In fact, define the function $\hat{\varphi}_i : (S_i \times \bar{X}_i)^\infty \to \mathbb{R}$ as $\hat{\varphi}_i(s_i^t, a_i^t) := \inf\{d(a_i^t, \tilde{a}_i^t)|(s_i^t, \tilde{a}_i^t) \in \bar{\mathscr{H}}_{A_i,*}\}$, where d is a distance in the metric space $\bar{X}^t$.[8] The function $\hat{\varphi}_i$ is continuous in a_i^t, as it is a "distance to a set" function, and it is measurable in s_i^t for each a_i^t.[9] The function $\hat{\varphi}_i(s_i^t, a_i^t)$ is zero for $(s_i^t, a_i^t) \in \bar{\mathscr{H}}_{A_i,*}$, and strictly positive otherwise, as A_i is closed-valued. For $\tilde{a}_i \in X_i$, by Lemma 5 for $C = \{(a_i^{t-1}, \tilde{a}_i)\}$, we obtain

$$\hat{\varphi}_i(s_i^{t-1}, s_i, a_i^{t-1}, \tilde{a}_i) \cdot p_i(a_i^{t-1}|s_i^{t-1}, a_i^{t,(\tau)}, \hat{\sigma}_i^*) \cdot \hat{\sigma}_i^*(\tilde{a}_i|(s_i^{t-1}, s_i), a_i^{t-1}) = 0,$$

ν_i^t-almost surely for $((s_i^{t-1}, s_i), a_i^{t-1}) \in \mathscr{H}_{A_i,*}^t$ such that $a_i^{t,(\tau)} = a_i((s_i^{t-1}, s_i), a_i^{t-1}, \hat{\sigma}_i^*)$. This shows that $\hat{\sigma}_i^*(\tilde{a}_i|(s_i^{t-1}, s_i), a_i^{t-1}) = 0$ if $((s_i^{t-1}, s_i), a_i^{t-1}) \in \mathscr{H}_{A_i,*}^t$, but $\tilde{a}_i \notin A_i((s_i^{t-1}, s_i), a_i^{t-1})$. Hence, $\hat{\sigma}_i^*(\cdot|(s_i^{t-1}, s_i), a_i^{t-1}) \in \Delta(A_i((s_i^{t-1}, s_i), a_i^{t-1}))$.

Proof of (b). It follows by Lemma 16 that $p^*|_{\bar{\mathscr{H}}_{i,*}}$ must be measurable with respect to the sub σ algebra $\mathcal{M}(\bar{\mathscr{H}}_{i,*}|\bar{h}_i)$. □

Proof of Theorem 1. Let the *weak topology on* $\hat{\Sigma}_i$ be the coarsest topology such that the functional $I_{\varphi_i}(p_i(\cdot|\cdot, \emptyset, \hat{\sigma}_i))$ is continuous for every Carathédory integrand $\varphi_i \in CI((S_i \times \bar{X}_i)^\infty, \alpha_i)$. The *weak topology on* $\hat{\Sigma} = \times_{i \in N} \hat{\Sigma}_i$ is the product topology,

[7] Notice that $a_i(s_i^t, a_i^t, \hat{\sigma}_i^*)$ does not depend on the value of $a_{i,t}^t$. Thus, we abuse notation slightly by writing $a_i((s_i^{t-1}, s_i), a_i^{t-1}, \hat{\sigma}_i^*)$.

[8] The distance d can be any distance that generates the product topology. Without loss of generality, we can assume $\bar{a}_\star$ is at distance 1 from every element in X^t.

[9] This follows by the measurable maximum theorem as A_i is weakly measurable and non-empty and compact valued. See Theorem 18.19 in Aliprantis and Border (2006).

where each $\hat{\Sigma}_i$ is endowed with its weak topology. By Lemma 6, for each $i \in N$, $\hat{\Sigma}_i(\tilde{\varepsilon})$ is closed in the weak topology.[10]

In the remainder, we define U_i over the set of extended strategies by setting, for every $t \in \mathbb{N}$ and $a^t \in \bar{X}^t$, $g_i(\cdot, a^t) = 0$ whenever $a^t_{j,\tau} = \bar{a}_\star$ for some $j \in N$ and $\tau \leq t$.

Lemma 7. *Under payoff boundedness, continuity, and SAC, for every $i \in N$, $U_i(\sigma)$ is continuous in the weak topology of $\hat{\Sigma}$.*

Proof. Let $(\hat{\sigma}^\lambda)_{\lambda \in \Lambda}$ be a net in $\hat{\Sigma}$ that converges to $\hat{\sigma}^*$ in the weak topology. We will show that $U_i(\hat{\sigma}^\lambda) \to U_i(\hat{\sigma}^*)$.

Let $\bar{P}^t$ be the projection of $\Omega^t \times X^t$ onto $S^t \times X^t$, i.e., $\bar{P}^t : \Omega^t \times X^t \to S^t \times X^t$ is such that $\bar{P}^t(\omega^t, a^t) = (\gamma^t(\omega^t, a^{t,(t-1)}), a^t)$ for $(\omega^t, a^t) \in \Omega^t \times X^t$. The sub σ-algebra of $\mathcal{M}\left(\Omega^t \times X^t\right)$ induced by $\bar{P}^t$ is $\mathcal{M}(\Omega^t \times X^t | \bar{P}^t) := \{(\bar{P}^t)^{-1}(B) | B \in \mathcal{M}\left(S^t \times X^t\right)\}$. For $a^t \in X^t$, let $\mathbb{E}(g_{i,t}(\cdot, a^t)|\bar{P}^t)$ denote the expectation of $g_{i,t}(\cdot, a^t) := g_i(\cdot, \cdot)|_{\Omega^t \times \{a^t\}}$ with respect to the sub-σ-algebra $\mathcal{M}(\Omega^t \times X^t | \bar{P}^t)$, where the measure in $\Omega^t \times \{a^t\}$ is $\mu^t_\omega(\cdot|a^t)$. By Theorem 4.41 in Aliprantis and Border (2006) there is a measurable function $E^{a^t}_{g_i} : (S \times X)^\infty \to \mathbb{R}$ such that $E^{a^t}_{g_i}(\bar{P}^t(\cdot, a^t)) = \mathbb{E}(g_{i,t}(\cdot, a^t)|\bar{P}^t)$.

If X is infinite, let $\hat{g}_{i,t}$ be defined as in our payoff continuity condition. Define $\tilde{g}_{i,t} : S^t \times X^t \to \mathbb{R}$ as $\tilde{g}_{i,t}(s^t, a^t) := E^{a^t}_{g_i}(s^t)$ if X is finite and as $\hat{g}_{i,t}(s^t, a^t)$, otherwise. Let $\sigma \in \hat{\Sigma}$, we have $U_{i,t}(\sigma) = \sum_{a^t \in X^t} \int_{S^t} \tilde{g}_{i,t}(s^t, a^t) \cdot f^t(s^t, a^{t,(t-1)}) \cdot p(a^t | s^t, \sigma) \, d\nu^t(s^t)$ where the equality follows by Theorems 4.41 and 13.46 in Aliprantis and Border (2006) and equation (9). By Proposition 5, SAC(b) implies $f^t(s^t, a^{t,(t-1)}) \in CI(S^t \times \bar{X}^{t-1}, \nu^t)$; payoff continuity implies there is a version of $\tilde{g}_{i,t}(s^t, a^t) \in CI(S^t \times \bar{X}^t, \nu^t)$. Therefore, there is a version of $\tilde{g}_{i,t}(s^t, a^t) \cdot f^t(s^t, a^{t,(t-1)}) \in CI(S^t \times \bar{X}^t, \nu^t)$, and by Theorem 2.5 in Balder (1988), we obtain $U_{i,t}(\hat{\sigma}^\lambda) \to U_{i,t}(\hat{\sigma}^*)$.

Now, by Lemma 4, $|U_i(\sigma)| = \left|\sum_{t=1}^\infty U_{i,t}(\sigma)\right| \leq \sum_{t \in \mathbb{N}} \sup_{\sigma \in \Sigma} |U_{i,t}(\sigma)| < \infty$, where the last inequality follows by payoff boundedness. Therefore, $U_i(\hat{\sigma}^\lambda) \to U_i(\hat{\sigma}^*)$ follows by applying the dominated convergence theorem over nets shown in Proposition 8. □

Let $\tilde{\varepsilon} \in \mathcal{E}$. Define the $\tilde{\varepsilon}$-constrained best response correspondence r as in the proof sketch of Theorem 1, but on the set of *induced extended transition probabilities* $\hat{\mathcal{P}}_i(\tilde{\varepsilon}) := \{p_i(\cdot|\cdot, \sigma_i) | \sigma_i \in \hat{\Sigma}_i(\tilde{\varepsilon})\}$.

The topological spaces $\hat{\mathcal{P}}(\tilde{\varepsilon})$, endowed with the relative weak topology, and $\hat{\Sigma}(\tilde{\varepsilon})$ are homeomorphic. To see this, notice that, by the definition of $\sigma_i \in \hat{\Sigma}(\tilde{\varepsilon})$, the ratio

[10] Endow $\hat{\Sigma}_i(\tilde{\varepsilon})$ with the relative topology inherited from the weak topology on $\hat{\Sigma}_i$, and $\hat{\Sigma}(\tilde{\varepsilon}) := \times_{i \in N} \hat{\Sigma}_i(\tilde{\varepsilon})$ with the corresponding product topology.

$$\sigma_i\big(a^t_{i,t}|s^t_i, a^{t,(t-1)}_i\big) = \frac{p_i\big(a^t_i|s^t_i, \sigma_i\big)}{p_i\big(a^{t,(t-1)}_i|s^{t,(t-1)}_i, \sigma_i\big)} \tag{10}$$

is well-defined for every $t \in \mathbb{N}$, $(s^t_i, a^{t,(t-1)}_i) \in \mathscr{H}_{A_i,*}$, where $p_i(\cdot|\cdot, \sigma_i) \in \hat{\mathscr{P}}(\tilde{\varepsilon})$ and we set $p_i(a^{t,(0)}_i|s^{t,(0)}_i, \sigma_i) := 1$. If $(s^t_i, a^{t,(t-1)}_i) \notin \mathscr{H}_{A_i,*}$ and the ratio is not well-defined, set it to 1 when $a^t_{i,t} = \bar{a}_\star$, and to 0 otherwise.

For every $i \in N$, define i's utility over the set of induced extended transition probabilities as $\tilde{U}_i(p(\cdot|\cdot, \sigma)) = U_i(\sigma)$. By Lemma 7, and the homeomorphism between $\hat{\Sigma}(\tilde{\varepsilon})$ and $\mathscr{P}(\tilde{\varepsilon})$, $\tilde{U}_i$ is continuous.

Relying on the homeomorphism introduced in equation (10), every fixed point of r is mapped to a strategy profile that, by definition, corresponds to a $\tilde{\varepsilon}$-constrained equilibrium. We establish that r has a fixed point by applying the Kakutani-Fan-Glicksberg fixed point theorem. In particular, we show that: **(1)** $\hat{\mathscr{P}}(\tilde{\varepsilon})$ is compact, and **(2)** convex; **(3)** r has closed graph, **(4)** is non-empty, and **(5)** convex valued.[11]

Point **(1)** follows by Lemma 6 and Theorem 2.3 in Balder (1988); **(4)** follows by **(1)** and the continuity of $\tilde{U}_i$; **(5)** follows by Lemma 4. We show points **(2)** and **(3)**.

(2) $\hat{\mathscr{P}}(\tilde{\varepsilon})$ is convex. Let $p', p'' \in \hat{\mathscr{P}}(\tilde{\varepsilon})$, with associated strategies σ' and σ'', respectively. To show that $\lambda \cdot p' + (1-\lambda) \cdot p'' \in \hat{\mathscr{P}}(\tilde{\varepsilon})$ for every $\lambda \in [0,1]$ it is sufficient to show $\lambda \cdot p'_i + (1-\lambda) \cdot p''_i \in \hat{\mathscr{P}}_i(\tilde{\varepsilon})$ for every $i \in N$. In particular, we prove that exists $\sigma_i \in \hat{\Sigma}_i(\tilde{\varepsilon})$ such that $(\lambda \cdot p'_i + (1-\lambda) \cdot p''_i)(a^t_i|s^t_i) = p_i(a^t_i|s^t_i, \sigma_i)$. Let σ_i be defined as in equation (10). We show that $\sigma_i \in \Sigma_i(\tilde{\varepsilon})$. Measurability of σ_i follows from Theorem 4.27 in Aliprantis and Border (2006). For every $(s^t_i, a^{t,(t-1)}_i) \in \mathscr{H}_{A_i,*}$, $\sigma_i \in [0,1]$, since $p_i(a^{t,(t-1)}_i|s^{t,(t-1)}_i) \geq p_i(a^t_i|s^t_i)$ for every $p_i \in \hat{\mathscr{P}}_i(\tilde{\varepsilon})$. Furthermore, if $a^t_{i,t} \neq \bar{a}_\star$, we have

$$\begin{aligned}\sigma_i\big(a^t_{i,t}|s^t_i, a^{t,(t-1)}_i\big) = &\frac{\sigma'_i\big(a^t_{i,t}|s^t_i, a^{t,(t-1)}_i\big) \cdot \lambda \cdot p_i(a^{t,(t-1)}_i|s^{t,(t-1)}_i, \sigma'_i)}{\lambda \cdot p_i(a^{t,(t-1)}_i|s^{t,(t-1)}_i, \sigma'_i) + (1-\lambda) \cdot p_i(a^{t,(t-1)}_i|s^{t,(t-1)}_i, \sigma''_i)} \\ &+ \frac{\sigma''_i\big(a^t_{i,t}|s^t_i, a^{t,(t-1)}_i\big) \cdot (1-\lambda) \cdot p_i(a^{t,(t-1)}_i|s^{t,(t-1)}_i, \sigma''_i)}{\lambda \cdot p_i(a^{t,(t-1)}_i|s^{t,(t-1)}_i, \sigma'_i) + (1-\lambda) \cdot p_i(a^{t,(t-1)}_i|s^{t,(t-1)}_i, \sigma''_i)} \geq \tilde{\varepsilon}((s^t_i, a^{t,(t-1)}_i), a^t_{i,t}).\end{aligned}$$

where the equality follows from equation (8), and the inequality follows as $\sigma', \sigma'' \in \hat{\Sigma}(\tilde{\varepsilon})$. Finally, $\sum_{a_i \in \bar{X}_i} \sigma_i(a_i|s^t_i, a^{t,(t-1)}_i) = 1$ follows from the equality above as well.

(3) r has closed graph. Let (p^λ) and $(\hat{p}^\lambda)$ be two nets of transition probabilities such that, for every $\lambda \in \Lambda$, $\hat{p}^\lambda \in r(p^\lambda)$, $\hat{p}^\lambda \to \hat{p}^*$ and $p^\lambda \to p^*$, where all convergences

[11] We identify $\hat{\mathscr{P}}_i(\tilde{\varepsilon})$ with its quotient space, by associating each transition probability with the equivalence class of transition probabilities yielding the same $\mathscr{I}_{i,\varphi_i}$ for every φ_i on $CI((S_i \times \bar{X}_i)^\infty, \alpha_i)$. As discussed in Balder (1988) this quotient space is a locally convex Hausdorff space.

are in the weak topology. Let us show that $\hat{p}^* \in r(p^*)$.

The condition $\hat{p}^\lambda \in r(p^\lambda)$ implies that, for every $i \in N$, $p_i \in \hat{\mathscr{P}}_i(\tilde{\varepsilon})$, $\tilde{U}_i(\hat{p}_i^\lambda, p_{-i}^\lambda) \geq \tilde{U}_i(p_i, p_{-i}^\lambda)$. Also, notice that, by the definition of weak convergence, $(\hat{p}_i^\lambda, p_{-i}^\lambda) \to (\hat{p}_i^*, p_{-i}^*)$ and $(p_i, p_{-i}^\lambda) \to (p_i, p_{-i}^*)$. Thus, as Lemma 7 implies continuity of $\tilde{U}_i$, we obtain $\tilde{U}_i(\hat{p}_i^*, p_{-i}^*) \geq \tilde{U}_i(p_i, p_{-i}^*)$, for each $i \in N$, which shows that $\hat{p}^* \in r(p^*)$. □

Proof of Theorem 2. By Remark 2.4 in Balder (1988), $\hat{\Sigma}$ is weakly sequentially compact. Therefore, for any sequences $(\varepsilon_n)_{n\in\mathbb{N}}$ in $\mathbb{R}_+$, $(\tilde{\varepsilon}_n)_{n\in\mathbb{N}}$ in $\mathscr{E}$, and $(\sigma^n)_{n\in\mathbb{N}}$ in Σ such that, for each $n \in \mathbb{N}$, $\tilde{\varepsilon}_n \in \mathscr{E}(\varepsilon_n)$, σ_n is $\tilde{\varepsilon}_n$-constrained, and $\lim_{n\to\infty} \varepsilon_n = 0$, we can construct sequences $\{p_i(\cdot|\cdot, a_i^\tau, \hat{\sigma}_i^n)\}_{n\in\mathbb{N}}$ that converge to $p_i^*(\cdot|\cdot, a_i^\tau)$ in the weak topology of $\mathscr{P}_i$ for each $\tau \in \mathbb{N}$ and $a_i^\tau \in \bar{X}_i^\tau$. Theorem 2 follows by Lemma 6. □

B.4 Proof of Theorem 3

MAC, together with Proposition 5, implies that, for every $t \in \mathbb{N}, s^t \in S^{R,t}$, $a^{t-1} \in X^{t-1}$, we can write the payoff-relevant signal profile transition as $d\mu_s^{R,t}(s^t|a^{t-1}) = f^t(s^t, a^{t-1}) \cdot d\nu^{M,t}(s^t)$, where $f^t \in CI(S^{R,t} \times X^{t-1}, \nu^{M,t})$, and $\nu^{M,t}(s^t) := \prod_{i\in N} \nu_i^{M,t}(s_i^t)$ is the product of the reference measures $\nu_i^{M,t}(s_i^t) := \prod_{\ell\leq t} \nu_{i,\ell}^M(s_{i,\ell}^t)$. For instance, each $\nu_{i,t}^M \in \Delta(S_i^R)$ can be constructed as $d\nu_{i,t}^M(s_i) := \sum_{a^{t-1}\in X^{t-1}} \xi_i(a^{t-1}) \cdot d\mu_{s_i,t}^M(s_i|a^{t-1})$ for ξ_i arbitrary collection of strictly positive weights.

We say that player i's strategy is *non-stationary Markov* if in each period it only depends on the payoff-relevant component of i's signal but may condition on the length of the i's private history at the time of play.

Denote by $\hat{\Sigma}_i^{nsM}$ the set of player's i extended non-stationary Markov strategies and by $\hat{\Sigma}^{nsM}$ the corresponding set of strategy profiles. Define the *weak topology on* $\hat{\Sigma}_i^{nsM}$ as the coarsest topology such that, for $t \in \mathbb{N}$, the functional $\mathscr{G}_{i,\varphi_i}^t : \hat{\Sigma}_i^{nsM} \to \mathbb{R}$ $\mathscr{G}_{i,\varphi_i}^t(\hat{\sigma}_i^{nsM}) := \int_{S_i^R} \sum_{a_i\in\bar{X}_i} \varphi_i(s_i, a_i) \cdot \hat{\sigma}_i^{nsM}(a_i|s_i, t)\, d\nu_{i,t}^M(s_i)$, is continuous for every $\varphi_i \in CI(S_i^R \times \bar{X}_i, \nu_{i,t}^M)$. The *weak topology* on $\hat{\Sigma}^{nsM}$ is the product topology, where each $\hat{\Sigma}_i^{nsM}$ is endowed with its weak topology.

Denote by $\hat{\Sigma}_i^{nsM}(\tilde{\varepsilon})$ and $\hat{\Sigma}^{nsM}(\tilde{\varepsilon})$ the sets of non-stationary Markov $\tilde{\varepsilon}$-constrained strategies, and endow them with the relative and product topology, respectively. By Lemma 6 with $\tau = t-1$, $\hat{\Sigma}_i^{nsM}(\tilde{\varepsilon})$ is closed in the weak topology, for each $i \in N$. The corresponding sets of stationary Markov strategies are denoted by the superscript M. Note that, $\hat{\Sigma}^M$ and $\hat{\Sigma}^M(\tilde{\varepsilon})$ are closed subsets of $\hat{\Sigma}^{nsM}$ and $\hat{\Sigma}^{nsM}(\tilde{\varepsilon})$, respectively.

Define U_i over the set of extended non-stationary Markov strategies by setting, for $t \in \mathbb{N}$, $a^t \in \bar{X}^t$, $g_i(\cdot, a^t) = 0$ whenever $a_{j,\tau}^t = \bar{a}_\star$ for $j \in N$ and $\tau \leq t$.

Lemma 8. *Under payoff boundedness, continuity, and MAC, for every $i \in N$, $U_i(\sigma)$ is continuous in the weak topology of $\hat{\Sigma}^{nsM}$.*

Proof. Let $(\hat{\sigma}^\lambda)_{\lambda\in\Lambda}$ be a net in $\hat{\Sigma}^{nsM}$ that converges to $\hat{\sigma}^*$ in the weak topology. We will show that $U_i(\hat{\sigma}^\lambda) \to U_i(\hat{\sigma}^*)$.

For every $\sigma^{nsM} \in \hat{\Sigma}^{nsM}$, by applying Lemma 4, payoff boundedness, payoff continuity, and MAC, together with Theorems 4.41 and 13.46 in Aliprantis and Border (2006), we have

$$\begin{aligned} U_{i,t}(\sigma^{nsM}) &= \sum_{a^t\in X^t}\int_{S^t} \hat{g}_i(s^t,a^t)\cdot p(a^t|s^t,\sigma^{nsM})\, d\mu_s^t\left(s^t|a^t\right) \\ &= \sum_{a^t\in X^t}\int_{S^{R,t}} \hat{g}_{i,t}(s^t,a^t)\cdot f^t(s^t,a^{t,(t-1)})\cdot \prod_{i\in N,\ell\leq t}\sigma_i^{nsM}(a_{i,\ell}^t|s_{i,\ell}^t,\ell)\, d\nu^{M,t}(s^t) \end{aligned}$$

where the second equality follows as $\hat{g}_{i,t}(s^t,a^t) = \hat{g}_{i,t}(s^{R,t},a^t)$ is constant on payoff-relevant signals by Markov payoff, and $p(a^t|s^t,\sigma^{nsM}) = p(a^t|s^{R,t},\sigma^{nsM})$ by definition of non-stationary Markov strategy, respectively.

By Proposition 5, MAC(b) implies there is a version of $f^t(s^t,a^{t,(t-1)}) \in CI(S^{R,t}\times \bar{X}^{t-1},\nu^{M,t})$; payoff continuity implies there is a version of $\hat{g}_{i,t}(s^t,a^t)\cdot f^t(s^t,a^{t,(t-1)}) \in CI(S^{R,t}\times\bar{X}^t,\nu^{M,t})$. Thus, by Theorem 2.5 in Balder (1988), we obtain $U_{i,t}(\hat{\sigma}^\lambda) \to U_{i,t}(\hat{\sigma}^*)$.

By Lemma 4, and payoff boundedness, $|U_i(\sigma)| = \left|\sum_{t=1}^\infty U_{i,t}(\sigma)\right| \leq \sum_{t\in\mathbb{N}}\sup_{\sigma\in\Sigma}|U_{i,t}(\sigma)| < \infty$. Therefore, by applying Proposition 8 we obtain $U_i(\hat{\sigma}^\lambda) \to U_i(\hat{\sigma}^*)$. □

Let $\tilde{\varepsilon} \in \mathcal{E}$. Define the $\tilde{\varepsilon}$-constrained best response correspondence $r^M : \hat{\Sigma}^M(\tilde{\varepsilon}) \to \hat{\Sigma}^M(\tilde{\varepsilon})$ as usual. We establish r^M has a fixed point by applying the Kakutani-Fan-Glicksberg fixed point theorem. In particular, we show: **(1)** $\hat{\Sigma}^M(\tilde{\varepsilon})$ is compact, and **(2)** convex; **(3)** r^M has closed graph, **(4)** is non-empty, and **(5)** convex valued.[12] Point **(1)** follows by Lemma 6 and Theorem 2.3 in Balder (1988); point **(2)** is straightforward; point **(3)** holds as argued in Theorem 1, applying Lemma 8; point **(4)** follows by **(1)** and the continuity of U_i; point **(5)** follows by the next lemma.

[12] As before, we identify $\hat{\mathcal{P}}_i^M(\tilde{\varepsilon})$ with its quotient space of transition probabilities yielding the same $\mathcal{I}_{i,\varphi_i}^t$ for every $\varphi_i^t \in CI(S_i^t\times\bar{X}_i^t,\nu_i^{M,t})$, $t\in\mathbb{N}$. By Balder (1988) this quotient space is locally convex Hausdorff.

Lemma 9. *If $\hat{\sigma}_i, \tilde{\sigma}_i \in r_i^M(\sigma)$ for $\sigma \in \hat{\Sigma}^M(\tilde{\varepsilon})$ then $\sigma_i^\lambda := \lambda \cdot \hat{\sigma}_i + (1-\lambda) \cdot \tilde{\sigma}_i \in r_i^M(\sigma)$ for every $\lambda \in (0,1)$.*

Proof. Let $(\sigma_i^{\lambda,t}, \sigma_i)$ denote the strategy that coincides with $\lambda \cdot \hat{\sigma}_i + (1-\lambda) \cdot \tilde{\sigma}_i$ in i's first t active periods and equals $\sigma_i \in \hat{\Sigma}_i^M$ in later active periods. We will show, inductively in t, that $U_i((\sigma_i^{\lambda,t}, \hat{\sigma}_i), \sigma_{-i}^M) = U_i((\sigma_i^{\lambda,t}, \tilde{\sigma}_i), \sigma_{-i}^M) = U_i(\hat{\sigma}_i, \sigma_{-i}^M) = U_i(\tilde{\sigma}_i, \sigma_{-i}^M)$ for every $t \in \mathbb{N} \cup \{0\}$. The statement is clearly true for $t = 0$.

Suppose the induction hypothesis holds for $t \leq \tau$. We show it holds for $t = \tau + 1$. We can write $U_i((\sigma_i^{\lambda,\tau+1}, \hat{\sigma}_i), \sigma_{-i}^M) = \lambda \cdot U_i((\sigma_i^{\lambda,\tau}, \hat{\sigma}_i), \sigma_{-i}^M) + (1-\lambda) \cdot U_i((\sigma_i^{\lambda,\tau}, \tilde{\sigma}_{i,\tau+1}, \hat{\sigma}_i), \sigma_{-i}^M)$, where $(\sigma_i^{\lambda,\tau}, \tilde{\sigma}_{i,\tau+1}, \hat{\sigma}_i)$ is the strategy that is equal to $\sigma_i^{\lambda,\tau}$ in the first τ active periods, equal to $\tilde{\sigma}_i$ in active period $\tau+1$ and equal to $\hat{\sigma}_i$ afterwards.

We can write $U_i(\sigma) = \int_{\mathcal{H}_i^{\tau+1}} U_i(\sigma|h_i, \sigma_{-i}) p_i(a_i(h_i)|s_i(h_i), \sigma_i) dP_{-i}(h_i|\sigma_{-i})$, for any $\sigma \in \hat{\Sigma}$, where the measurable functions $a_i(h_i)$ and $s_i(h_i)$ project the private history $h_i \in \mathcal{H}_i$ onto $X_i^{|h_i|}$ and $S_i^{|h_i|}$, respectively, and $U_i(\sigma|h_i, \sigma_{-i})$ is defined as in equation (11). Therefore, condition (13) in Lemma 11 implies that $U_i((\sigma_i^{\lambda,\tau}, \tilde{\sigma}_{i,\tau+1}, \hat{\sigma}_i), \sigma_{-i}^M) = U_i((\sigma_i^{\lambda,\tau}, \tilde{\sigma}_i), \sigma_{-i}^M)$, which concludes our induction argument. The lemma follows by payoff boundedness, as it implies that $U_i((\sigma_i^{\lambda,t}, \hat{\sigma}_i), \sigma_{-i}^M)$ converges to $U_i(\sigma_i^\lambda, \sigma_{-i}^M)$. □

Remark 2. *In a Markov game with discounted payoffs an argument analogous to the one above establishes the existence of a stationary Markov equilibrium.*

For every $i \in N$, $\sigma \in \Sigma$, $\ell \in \mathbb{N}$ and $Z \in \mathcal{M}(S_i^R)$, define the probability that player i's payoff relevant signal belongs to Z and yields a private history of length $\hat{t}$ under σ as $P_i^{\hat{t},R}(Z|\sigma) := P_i(\cup_{t\in\mathbb{N}}(\hat{h}_i^t)^{-1}((s_i^R)^{-1}(Z) \cap \mathcal{H}_i^{\hat{t}})|\sigma)$. Furthermore, define the transition probability $P_i^{\hat{t}}(\cdot|\cdot,\sigma) : S_i^R \to \Delta(\mathcal{H}_i^{\hat{t}})$ as $dP_i(h_i|\sigma) = dP_i^{\hat{t}}(h_i|s_i^R, \sigma) \times dP_i^{\hat{t},R}(s_i^R|\sigma)$.

Lemma 10. *If σ^M is a fixed point of r^M for $\tilde{\varepsilon} \in \mathcal{E}$, then $U_i(\sigma^M) \geq U_i(\sigma_i, \sigma_{-i}^M)$ for every $i \in N$, $\sigma_i \in \hat{\Sigma}(\tilde{\varepsilon})$.*

Proof. Suppose that σ^M is a fixed point of r^M and that there is a player $i \in N$, $\beta > 0$ and a strategy $\sigma_i' \in \hat{\Sigma}_i(\tilde{\varepsilon})$ such that $U_i(\sigma_i', \sigma_{-i}^M) - U_i(\sigma^M) = \beta > 0$. By payoff boundedness, there is $\hat{t} \in \mathbb{N}$ such that the strategy σ_i coinciding with σ_i' in the first $\hat{t}$ active periods and with σ_i^M, otherwise, satisfies $U_i(\sigma_i, \sigma_{-i}^M) - U_i(\sigma^M) \geq \beta/2 > 0$.

We show that we can construct a strategy that is non-stationary Markov in active period $\hat{t}$ that is also a profitable deviation. Arguing recursively we then show that σ_i^M must have a profitable deviation that is non-stationary Markov.

Let $\hat{\mathscr{H}}^{i,\hat{t}} := \cup_{t\in\mathbb{N}}(\hbar_i \circ \gamma^t)^{-1}(\mathscr{H}_i^{\hat{t}})$ be the set of histories in $\mathscr{H}$ that yield a player i private history of length $\hat{t}$. For every $\sigma \in \Sigma$ define $\hat{t}$'th active period strategy as

$$\tilde{\sigma}_i^{M,\hat{t}}(a_i|s_i^R) := \int_{\mathscr{H}_i^{\hat{t}}} \sigma_i(a_i|h_i)dP_i^{\hat{t}}(h_i|s_i^R, \sigma_i, \sigma_{-i}^M),$$

for every $s_i^R \in \mathscr{s}_i^R(\mathscr{H}_i^{\hat{t}})$. Define the strategy $\tilde{\sigma}_i$ to correspond to σ_i in the first $\hat{t}-1$ active periods, to $\tilde{\sigma}_i^{M,\hat{t}}$ at $\hat{t}$, and to σ_i^M thereafter.

Denote by $\bar{p}^R$ the projection of histories in $\mathscr{H}_*$ onto $\cup_{t\in\mathbb{N}}\Omega^R \times X^{t-1}$. The transition $\mu_{\omega,t}^R(\cdot|\cdot,\sigma) : \cup_{i\in N}\mathscr{H}_i \to \Delta(\Omega^R)$ conditional on $\sigma \in \Sigma$ is defined by

$$\int_{(\Omega^R)^{t-1}\times\hat{\Omega}\times X^{t-1}\cap(\bar{p}^R\circ\hbar_i^h)^{-1}(Z)} p(a^{t-1}|\omega^{t-1},\sigma)d\mu_\omega^{R,t}(\omega_t,\omega^{t-1}|a^{t-1}) = \int_Z d\mu_{\omega,t}^R(\hat{\Omega}|h_i,\sigma)\,dP_i(h_i|\sigma),$$

for $\hat{\Omega} \in \mathcal{M}(\Omega^R)$, $Z \in \mathscr{H}_i^\ell$ for some $i \in N, \ell \in \mathbb{N}$, where $\mu_\omega^{R,t}$ denotes the transition probability for payoff-relevant states.

Let $\sigma = (\sigma_i, \sigma_{-i}^M)$, $t, \tau \in \mathbb{N}$, $t > \tau$. The measure $P^{h,t}(\cdot|\omega_\tau, a_\tau, \sigma)$ over time-t states and actions in $\Omega^R \times X$ induced by σ after state action pair $(\omega_\tau, a_\tau) \in \Omega^R \times X$ is defined by

$$P^{h,t}(\hat{\Omega}\times\hat{X}|\omega_\tau,a_\tau,\sigma) := \sum_{a^t\in X^{t-1}\times\hat{X}} \int_{\Omega^{R,t-(\tau+1)}\times\hat{\Omega}^R} p(a^t|\omega^t,a_0^\tau,\sigma) \prod_{k=\tau+1}^{t} d\mu_\omega^R(\omega_k^t|\omega_{k-1}^t,a_{k-1}^t),$$

for each $\hat{\Omega}^R \times \hat{X} \in \mathcal{M}(\Omega^R \times X)$, and where $a_0^\tau \in X^\tau$, with $a_{0,\tau}^\tau = a_\tau$. Notice that $P^{h,t}$ is independent of the choice of a_0^τ when σ is a Markov strategy. The sum of player i's expected flow payoffs over periods $t \geq \hat{t}$ and σ is given by

$$\begin{aligned} U_{i,\hat{t}\to\infty}(\sigma) = \sum_{\tau\geq\hat{t},t\geq\hat{t}} \int_{S_i^R}\int_{\mathscr{H}_i^{\hat{t}}}\int_{\Omega^R\times X}\int_{\Omega^R\times X} g_i(\omega_t,a_t)\cdot\sigma_i(a_{i,\tau}|h_i)\cdot\sigma_{-i}^M(a_{-i,\tau}|\omega_\tau) \\ dP^{h,t}(\omega_t,a_t|\omega_\tau,a_\tau,\sigma)\,d\mu_{\omega,\tau}^R(\omega_\tau|h_i,\sigma)\,dP_i(h_i|s_i^R,\sigma)\,dP_i^R(s_i^R|\sigma) = U_{i,\hat{t}\to\infty}(\tilde{\sigma}_i,\sigma_{-i}^M) \end{aligned}$$

where the last equality follows from σ_i Markov after i's active period $\hat{t}$, and the Markov payoff assumption, as $d\mu_{\omega,\tau}^R(\omega_\tau|h_i,\sigma)$ depends on h_i only through $\mathscr{s}_i^R(h_i)$. Since player i's flow payoffs from active periods 1 through $\hat{t}-1$ are equal under σ_i and $\tilde{\sigma}_i$, we can conclude that $U_i(\sigma) = U_i(\tilde{\sigma}_i, \sigma_{-i}^M)$.

We now show that σ_{-i}^M has a stationary Markov best response. Let us first show that σ_i^M does not have a profitable one-shot deviation, i.e., a strategy that yields a higher payoff but only differs from σ_i^M in one active period. In fact, suppose σ_i^M has a profitable one-shot deviation $\tilde{\sigma}_i^M$ that differs from σ_i^M in active period t. Therefore, the strategy $\sigma_i^{t,t'}$ that coincides with $\tilde{\sigma}_i^M$ from periods t through t' for $t' > t$, $t' \in \mathbb{N}$, and with σ_i^M in all other periods must also be a profitable deviation. However, by payoff boundedness, as t' converges to ∞ the expected payoff of $\sigma_i^{t,t'}$ converges to that of $\sigma_i^{t,\infty}$—the strategy that coincides with σ_i^M up to period $t-1$ and coincides with $\tilde{\sigma}_i^M$ afterwards. Hence, $\sigma_i^{t,\infty}$ must also be a profitable deviation. However, since, by Markov information, we can write $U_i(\sigma^M|h_i, \sigma_{-i}^M) = U_i(\sigma^M|\mathscr{s}_i^R(h_i), \sigma_{-i}^M)$, the fact that $\sigma_i^{t,\infty}$ is a profitable deviation contradicts the optimality of σ_i^M implied by condition (13) in Lemma 11.

Let σ_i^{nSM} be a Markov non-stationary best response, and $\sigma_i^{M,t}$ a best response among the strategies that coincide with σ_i^M from the $t+1$'th active period on. These best responses exist by compactness and Lemma 8. By payoff boundedness, for every ε, there is a t, $|U_i(\sigma_i^{nSM}, \sigma_{-i}^M) - U_i(\sigma_i^{M,t}, \sigma_{-i}^M)| < \varepsilon$. For every $t \in \mathbb{N}$, we will argue that $U_i(\sigma_i^{M,t}, \sigma_{-i}^M) = U_i(\sigma^M)$, which, in light of the previous inequality, establishes the desired conclusion. In fact, since σ_i^M does not have one-shot deviations, the strategy that coincides with $\sigma_i^{M,t}$ in the first $t-1$ active periods and with σ_i^M from period t on, must give i the same expected payoff as $\sigma_i^{M,t}$ from active periods $t+1$ on. Thus, applying this argument recursively, we obtain $U_i(\sigma_i^M) = U_i(\sigma_i^{M,t}, \sigma_{-i}^M)$. $\square$

The previous arguments show that, for every $\tilde{\varepsilon} \in \mathcal{E}$, there exist a fixed point of r^M, and, by Lemma 10, is an $\tilde{\varepsilon}$-constrained equilibrium.

Let $(\tilde{\varepsilon}_n)_{n\in\mathbb{N}}$ in $\mathcal{E}$ with $\tilde{\varepsilon}_n \in \mathcal{E}(\varepsilon_n)$, and $\varepsilon_n \to 0$, and σ^n be a stationary Markov $\tilde{\varepsilon}_n$-constrained equilibrium. Passing to a subsequence, there is a strategy σ^* such that σ^n converges to σ^* in the weak topology. Notice that this convergence also implies convergence in the sense of THPE.

B.5 Proof of Propositions 1 and 2

For $Z \in \mathcal{M}(\mathcal{H}_i^\ell)$, $\ell \in \mathbb{N}$, let $\mathcal{H}_i(Z) := \{h_i \in \mathcal{H}_i | h_i^{(\ell)} \in Z\}$ be the set of private histories that follow private histories in Z. We say that a strategy $\sigma_i' \in \Sigma_i$ is a *continuation of strategy* $\sigma_i \in \Sigma_i$ *after* Z if σ_i coincides with σ_i' in every $h_i \notin \mathcal{H}_i(Z)$. The set of strategies that are a continuation of σ_i after Z is denoted $\Sigma_i(\sigma_i, Z)$. For a pair of strategies $\sigma_i, \sigma_i' \in \Sigma_i$ we write $(\sigma_i'|_Z, \sigma_i)$ for the strategy in $\Sigma_i(\sigma_i, Z)$ that coincides with σ_i' at private histories in $\mathcal{H}_i(Z)$, and is equal to σ_i otherwise.

Let $\ell \in \mathbb{N}$, $a_i^\ell \in (X_i \setminus \{a_*\})^\ell$. Define $\mathcal{H}_i(a_i^\ell) := \{h_i \in \mathcal{H}_i^\ell | h_i = (s_i^\ell, a_i^\ell),\ s_i^\ell \in S_i^\ell\}$. Let $\sigma_i(a_i^\ell)$ denote the strategy that takes action $a_{i,t}^\ell$ in active period t when it is feasible and takes action $\bar{a}_\star$, otherwise. Define $P_{-i}(Z|\sigma_{-i}) = P_i(Z|(\sigma_i(a_i^\ell), \sigma_{-i}))$ for $Z \in \mathcal{M}(\mathcal{H}_i(a_i^\ell))$, for some $a_i^\ell \in (X_i \setminus \{a_*\})^\ell$.

Let $\sigma_{-i} \in \Sigma_{-i}$, and the transition probability $P_{\omega|h_i}(\cdot|\cdot, \sigma_{-i}) : \mathcal{H}_i \to \Delta(\mathcal{H})$ defined by

$$\int_{C\cap(\hbar_i^h)^{-1}(Z)} p_{-i}(a^{t-1} \mid \omega^{t,(t-1)}, \sigma) d\mu_{\omega,a}^t(\omega^t, a^{t-1}) \quad = \int_Z P_{\omega|h_i}\left(C \mid h_i, \sigma_{-i}\right) dP_{-i}(h_i \mid \sigma_{-i}),$$

for $Z \in \mathcal{M}(\mathcal{H}_i(a_i^\ell))$, $a_i^\ell \in (X_i \setminus \{a_*\})^\ell$ and $C \in \mathcal{M}(\mathcal{H})$. Define

$$U_i(\hat{\sigma}|h_i, \sigma_{-i}) := \sum_{\tau\in\mathbb{N}} \int_{\mathcal{H}^\tau} U_i(\hat{\sigma}|\omega^\tau, a^{\tau-1})\, dP_{\omega|h_i}(\omega^\tau, a^{\tau-1}|h_i, \sigma_{-i}). \tag{11}$$

Lemma 11. *Let $(\sigma_i, \sigma_{-i}) \in \Sigma$, and $Z \in \mathcal{M}(\mathcal{H}_i^\ell)$, $\ell \in \mathbb{N}$, with $P_i(Z|\sigma) > 0$. Let the set $\tilde{\Sigma}_i \subseteq \Sigma_i$ for $i \in N$, be such that $(\sigma_i'|_Z, \sigma_i) \in \tilde{\Sigma}_i$ for every $\sigma_i' \in \tilde{\Sigma}_i$. If*

$$U_i(\sigma_i, \sigma_{-i}) \geq U_i(\sigma_i', \sigma_{-i}) \qquad \forall \sigma_i' \in \tilde{\Sigma}_i, \tag{12}$$

then

$$U_i(\sigma_i, \sigma_{-i}|Z, \sigma) \geq U_i(\sigma_i', \sigma_{-i}|Z, \sigma) \qquad \forall \sigma_i' \in \tilde{\Sigma}_i.$$

Furthermore, if $P_{\omega|h_i}$ exists then, for $h_i \in \mathcal{H}_i \setminus N$ for some N with $P_{-i}(N|\sigma_{-i}) = 0$,

$$U_i(\sigma_i, \sigma_{-i}|h_i, \sigma_{-i}) \geq U_i(\sigma_i', \sigma_{-i}|h_i, \sigma_{-i}) \qquad \forall \sigma_i' \in \tilde{\Sigma}_i. \tag{13}$$

Proof. We argue by contradiction. Suppose that σ_i satisfies equation (12), but there are $\ell \in \mathbb{N}$, $Z \in \mathcal{M}(\mathcal{H}_i^\ell)$ with $P_i(Z|\sigma) > 0$, and $\sigma_i' \in \tilde{\Sigma}_i$ such that

$$U_i(\sigma|Z, \sigma) < U_i(\sigma_i', \sigma_{-i}|Z, \sigma).$$

By the definition of P_i^t, for every $C \in \mathcal{M}(\mathcal{H}^t)$, we have that

$$P_i^t(C|Z, \sigma) = \frac{\int_{C\cap(\hbar_i^h)^{-1}(Z)} p(a^{t-1}|\omega^{t,(t-1)}, \sigma) d\mu_{\omega,a}^t(\omega^t, a^{t-1})}{P_i(Z|\sigma)}.$$

Therefore, multiplying both sides by $P_i(Z|\sigma) > 0$, we obtain

$$\sum_{\substack{\tau\in\mathbb{N},\\ t\geq\tau}} \int_{(\hbar_i^h)^{-1}(Z)\cap\mathcal{H}^\tau} U_{i,t}(\sigma|\omega^\tau, a^{\tau-1}) p(a^{\tau-1}|\omega^{\tau,(\tau-1)}, \sigma)\, d\mu_{\omega,a}^\tau(\omega^\tau, a^{\tau-1})$$
$$< \sum_{\substack{\tau\in\mathbb{N},\\ t\geq\tau}} \int_{(\hbar_i^h)^{-1}(Z)\cap\mathcal{H}^\tau} U_{i,t}(\sigma_i', \sigma_{-i}|\omega^\tau, a^{\tau-1}) p(a^{\tau-1}|\omega^{\tau,(\tau-1)}, \sigma)\, d\mu_{\omega,a}^\tau(\omega^\tau, a^{\tau-1}). \tag{14}$$

By the definition of $U_{i,t}(\sigma|\omega^\tau, a^{\tau-1})$, the previous expression yields,

$$\sum_{t\in\mathbb{N}}\int_{\Omega^t} \mathbb{1}\{\hbar_i^h(\omega^t, a^{t,(t-1)}) \in \mathcal{H}_i(Z)\} \cdot g_i(\omega^t, a^t) \cdot p(a^t|\omega^t, \sigma)\, d\mu^t_{\omega,a}(\omega^t, a^{t,(t-1)})$$
$$< \sum_{t\in\mathbb{N}}\int_{\Omega^t} \mathbb{1}\{\hbar_i^h(\omega^t, a^{t,(t-1)}) \in \mathcal{H}_i(Z)\} \cdot g_i(\omega^t, a^t) \cdot p(a^t|\omega^t, (\sigma_i'', \sigma_{-i}))\, d\mu^t_{\omega,a}(\omega^t, a^{t,(t-1)}),$$

where $\sigma_i'' = (\sigma_i'|_Z, \sigma_i)$. Since σ_i'' belongs to $\tilde{\Sigma}_i$, and σ_i'' coincides with σ_i at $h_i \notin \mathcal{H}_i(Z)$, this contradicts that σ_i is a best response, as required by condition (12).

Similarly, if there is a set Z with $P_{-i}(Z|\sigma_{-i}) > 0$, and a strategy $\sigma_i' \in \tilde{\Sigma}_i$ such that condition (13) does not hold, then we obtain a contradiction by multiplying equation $U_i(\sigma_i, \sigma_{-i}|h_i, \sigma_{-i}) < U_i(\sigma_i', \sigma_{-i}|h_i, \sigma_{-i})$ by the corresponding p_i and integrating by $dP_{-i}(h_i|\sigma_{-i})$ over Z, which yields equation (14). □

Proof of Proposition 1. The fact that an $\tilde{\varepsilon}$-constrained equilibrium is an $\tilde{\varepsilon}$-constrained conditional equilibrium follows by Lemma 11 by setting $\tilde{\Sigma}_i = \Sigma_i(\tilde{\varepsilon})$. □

Proof of Proposition 2. For $h \in \mathcal{H}^\varnothing$, we have that $(\hbar_i^h)^{-1}(\hbar_i^h(h))$ is a singleton. Therefore, $P_{\omega|h_i}(h|h_i, \sigma_{-i}) = \mathbb{1}\{h_i = \hbar_i^h(h)\}$, since the right hand side is a probability measure. This implies that $U_i(\hat{\sigma}|h_i, \sigma_{-i})$ is independent of σ_{-i} and equal to $U_i(\hat{\sigma}|\omega^t, a^{t-1})$ for $(\omega^t, a^{t-1}) = (\hbar_i^h)^{-1}(h_i)$. Therefore, by condition (13) in Lemma 11 we obtain the following lemma.

Lemma 12. *If σ is an $\tilde{\varepsilon}$-constrained equilibrium, for $\tilde{\varepsilon} \in \mathcal{E}$, then there exists a negligible set $H \in \mathcal{M}(\mathcal{H})$ such that $U_i(\sigma|\omega^t, a^{t-1}) \geq U_i(\sigma_i', \sigma_{-i}|\omega^t, a^{t-1})$, for every $\sigma_i' \in \Sigma_i(\tilde{\varepsilon})$, $(\omega^t, a^{t-1}) \in \mathcal{H}^\varnothing \setminus H$.*

The following lemma shows that a THPE is an SPE.

Lemma 13. *Let σ^* be a THPE. There exists a negligible set $H \in \mathcal{M}(\mathcal{H})$ such that $U_i(\sigma^*|\omega^t, a^{t-1}) \geq U_i(\sigma_i', \sigma_{-i}^*|\omega^t, a^{t-1})$, for every $\sigma_i' \in \Sigma_i$, $(\omega^t, a^{t-1}) \in \mathcal{H}^\varnothing \setminus H$.*

Proof. Let σ^* be a THPE and suppose that there is a set $B \subseteq \mathcal{H}^\varnothing \cap \mathcal{H}^t$ non-negligible such that $U_i(\sigma^*|\omega^t, a^{t-1}) < U_i(\sigma_i', \sigma_{-i}^*|\omega^t, a^{t-1})$ for $(\omega^t, a^{t-1}) \in B$. We assume, without loss, that for each $i \in N$ there is $\tau_i \in \mathbb{N}$, $a_i(\gamma_i^{t-1}(\omega^{t,(t-1)}, a^{t-1}), a_i^{t-1}; \sigma_i^*) = a_i^{t-1,(\tau_i)}$ for every $(\omega^t, a^{t-1}) \in B$, and the projection of B onto X^{t-1} is a singleton $\{a^{t-1}\}$. Therefore, we obtain

$$- \int_B U_i(\sigma^*|\omega^t, a^{t-1})\Pi_{j\in N} p_j(a^{t-1}|\omega^{t,(t-1)}, a^{t,(\tau_j)}, \sigma^*)\, d\mu^t_\omega(\omega^t|a^{t-1})$$
$$+ \int_B U_i(\sigma'_i, \sigma^*_{-i}|\omega^t, a^{t-1})\Pi_{j\in N} p_j(a^{t-1}|\omega^{t,(t-1)}, a^{t,(\tau_j)}, (\sigma'_i, \sigma^*_{-i}))\, d\mu^t_\omega(\omega^t|a^{t-1}) := \beta > 0.$$

By the definition of THPE, Lemma 5, and Lemma 7,[13] for $\tilde{\varepsilon} \in \mathcal{E}$, there is an $\tilde{\varepsilon}$-constrained equilibrium σ^ε and $\tilde{\varepsilon}$-constrained strategy σ'^{ε}_i such that

$$\Bigg| \int_B U_i(\sigma^*|\omega^t, a^{t-1})\Pi_{j\in N} p_j(a^{t-1}|\omega^{t,(t-1)}, a^{t,(\tau_j)}, \sigma^*)\, d\mu^t_\omega(\omega^t|a^{t-1})$$
$$- \int_B U_i(\sigma^\varepsilon|\omega^t, a^{t-1})\Pi_{j\in N} p_j(a^{t-1}|\omega^{t,(t-1)}, a^{t,(\tau_j)}, \sigma^\varepsilon)\, d\mu^t_\omega(\omega^t|a^{t-1})\Bigg| < \beta/3,$$

and,

$$\Bigg| \int_B U_i(\sigma'_i, \sigma^*_{-i}|\omega^t, a^{t-1})\Pi_{j\in N} p_j(a^{t-1}|\omega^{t,(t-1)}, a^{t,(\tau_j)}, (\sigma'_i, \sigma^*_{-i}))\, d\mu^t_\omega(\omega^t|a^{t-1})$$
$$- \int_B U_i(\sigma'^{\varepsilon}_i, \sigma^*_{-i}|\omega^t, a^{t-1})\Pi_{j\in N} p_j(a^{t-1}|\omega^{t,(t-1)}, a^{t,(\tau_j)}, (\sigma'^{\varepsilon}_i, \sigma^*_{-i}))\, d\mu^t_\omega(\omega^t|a^{t-1})\Bigg| < \beta/3.$$

Combining the last three expressions yields a contradiction with Lemma 12. □

B.6 Stochastic Move Opportunities

Proof of Lemma 3. Let M be s.t. $\sup\{|g_i(\omega^t, a^t)|\,|(\omega^t, a^t) \in \Omega^t \times X^t,\ t \in \mathbb{N}\} < M$. For each $t \in \mathbb{N}$, define $\bar{g}_i(\omega^t, a^t) := M$ if $g_i(\omega^t, a^t) \neq 0$ and $\bar{g}_i(\omega^t, a^t) := 0$, otherwise.

For $i \in N$, and $(\omega^t, a^t) \in \Omega^t \times X^t$, $|g_i(\omega^t, a^t)| \leq \bar{g}_i(\omega^t, a^t)$. By Lemma 4, we write

$$\begin{aligned}
|U_i(\sigma)| &\leq \sum_{t=1}^{\infty} \sup_{\sigma\in\Sigma} \left| \sum_{a^t\in X^t} \int_{\Omega^t} g_i(\omega^t, a^t)\cdot p(a^t|\omega^t, \sigma)\, d\mu^t_\omega(\omega^t|a^{t,(t-1)}) \right| \\
&\leq \sum_{t=1}^{\infty} \sup_{\sigma\in\Sigma} \sum_{a^t\in X^t} \int_{\Omega^t} \bar{g}_i(\omega^t, a^t)\cdot p(a^t|\omega^t, \sigma)\, d\mu^t_\omega(\omega^t|a^{t,(t-1)}) \\
&= \sum_{t=1}^{\infty} \sup_{\sigma\in\Sigma} \sum_{a^t\in X^t} \int_{\Omega^t} \mathbb{1}\{(\omega^{t,(\tilde{t})}, a^{t,(\tilde{t})}) \notin \bar{H}^{\tilde{t}}, \tilde{t} \leq t\} \cdot M \cdot p(a^t|\omega^t, \sigma)\, d\mu^t_\omega(\omega^t|a^{t,(t-1)}) \\
&= \sum_{t=1}^{\infty} \sup_{\sigma\in\Sigma} \big(1 - \sum_{\tilde{t}=1}^{t} P^{\tilde{t}}(\bar{H}^{\tilde{t}}|\sigma)\big)\cdot M \\
&\leq \sum_{t=1}^{\infty} \sum_{\tilde{t}=t+1}^{\infty} \sup_{\sigma\in\Sigma} P^{\tilde{t}}(\bar{H}^{\tilde{t}}|\sigma)\cdot M = M\cdot \sum_{\tilde{t}=1}^{\infty} \sup_{\sigma\in\Sigma} \big((\tilde{t}-1)\cdot P^{\tilde{t}}(\bar{H}^{\tilde{t}}|\sigma)\big),
\end{aligned}$$

[13] Notice that we can view the previous integral as the expected utility of a strategy in which each player j puts weight 1 on actions in a^{t-1} up to τ_j, and in which $g_i = 0$ for $t' < t$.

where the second equality follows from the definition of $P^t(\cdot|\sigma)$ and $\bar{H}^\tau \cap \bar{H}^{\tau'} = \emptyset$ for $\tau \neq \tau'$, the third inequality from the fact that the game has finitely many periods with probability 1,[14] and the last equality is obtained by changing the order of the summation. □

B.7 Examples 2.1, 2.2, and 2.3

Consider the game of Examples 2.1, 2.2, and 2.3 adapted from Harris et al. (1995).

We first explain why the base game does not have any subgame perfect equilibrium. Denote by β the probability that player B plays R. If $\beta > 1/2$, player A prefers to play $a < 0$ as this would imply that $c = d = L$, maximizing the probability of a wrong guess by B. Anticipating this, player B would profitably deviate to $\beta < 1/2$. The same argument applies when $\beta < 1/2$. Consider the case $\beta = 1/2$. For every $a \neq 0$, there exists $a' \neq 0$ closer to zero, i.e., $|a'| < |a|$, that yields a higher payoff to player A. Finally, the action $a = 0$ is optimal for player A if and only if players C and D are perfectly coordinated, i.e., $c = d = L$ or $c = d = R$, as a function of b. However, this implies that B cannot be indifferent between L and R, as B's payoffs from her two actions never coincide when C plays pure actions. This shows that no subgame perfect equilibrium exists. Consequently, no Markov perfect equilibrium exists either.

Next, we show that the strategy profile described in Example 2.3 is an SPE. Restrict player A to randomize uniformly between the actions $\delta - \varepsilon$ and $-\delta + \varepsilon$, for $\varepsilon \in [\delta - 1, \delta]$ such that $a_\varepsilon := 1/2 \cdot (\delta - \varepsilon) + 1/2 \cdot (\delta + \varepsilon) \in \Delta(\mathcal{A})$. It is immediate to see that players B, C, and D cannot profitably deviate; we focus on player A's best response. Player A's payoff is

$$-10 \cdot \mathbb{P}(\{c \neq d\}|a_\varepsilon) - \tfrac{1}{2}|\delta - \varepsilon|^2,$$

where $\mathbb{P}(\{c \neq d\}|a_\varepsilon)$ is the probability that C and D play different actions conditional

[14] This latter condition implies $(1 - \sum_{\tilde{t}=1}^{t} P^{\tilde{t}}(\bar{H}^{\tilde{t}}|\sigma)) = \sum_{\tilde{t}=t+1}^{\infty} P^{\tilde{t}}(\bar{H}^{\tilde{t}}|\sigma)$.

on a_ε, which equals[15]

$$\mathbb{P}(\{c \neq d\}|a_\varepsilon) = \varepsilon \cdot \frac{(2\delta - \varepsilon)}{2\delta^2}$$

for $\varepsilon > 0$ and zero otherwise.

Clearly, $\varepsilon < 0$ cannot be optimal since, compared to $\varepsilon = 0$, it does not affect $\mathbb{P}(\{c \neq d\}|a_\varepsilon)$ and increases $\frac{1}{2}|\delta - \varepsilon|^2$. Furthermore, it is easy to check that A's payoff is decreasing in ε for $\delta \in \mathcal{A}$, making it optimal to set $\varepsilon = 0$ and uniformly randomize between δ and $-\delta$.

We are left to check whether player A could profitably play actions different from the ones that take the form of a_ε. This is not the case since for every mixed strategy $\alpha \in \Delta(\mathcal{A})$ and action $\bar{a} \in \text{supp}(\alpha) \cap (-\delta, \delta) \neq \emptyset$, player A obtains a higher payoff by assigning $\alpha(\bar{a}) > 0$ weight to δ instead of $\bar{a}$.

Formally, let $\alpha \in \Delta(\mathcal{A})$ be a mixed strategy. Clearly, if $\text{supp}(\alpha) \cap [-\delta, \delta] = \emptyset$ then α is not optimal. Assume that $\bar{a} \in \text{supp}(\alpha) \cap (-\delta, \delta) \neq \emptyset$. By following steps analogous to the ones displayed in footnote 15, one can show that the probability that C and D play different actions given $\bar{a}$ equals

$$\mathbb{P}(\{c \neq d\}|\bar{a}) = \frac{\delta^2 - |\bar{a}|^2}{2\delta^2}.$$

Therefore, A's payoff from playing action $\bar{a}$ amounts at

$$U_A(\bar{a}) = -10 \cdot \frac{\delta^2 - |\bar{a}|^2}{2\delta^2} - \frac{1}{2}|\bar{a}|^2 < -\frac{1}{2}|\delta|^2 = U_A(\delta)$$

where the inequality follows from first order conditions as $\delta \in \mathcal{A} \subseteq [-1, 1]$.

B.8 Applications to Markov Games

Application 2. We are left to show that Markov information and Markov payoff are satisfied. The former holds since the unobservable component of the current payoff-relevant state, that is, each opponent's private shock, is independent from past information conditional on the current payoff-relevant signal of each player. Therefore,

[15] Applying conditioning, it follows that

$$\begin{aligned}\mathbb{P}(\{c \neq d\}|a_\varepsilon) &= 1/2 \cdot \big(\mathbb{P}(\{c \neq d\}|\delta - \varepsilon) + \mathbb{P}(\{c \neq d\}| - \delta + \varepsilon)\big) \\ &= 1/2 \cdot \big(2 \cdot \mathbb{P}(s_i \geq 0|\delta - \varepsilon) \cdot \mathbb{P}(s_{-i} < 0|\delta - \varepsilon) + 2 \cdot \mathbb{P}(s_i \geq 0| - \delta + \varepsilon) \cdot \mathbb{P}(s_{-i} < 0| - \delta + \varepsilon)\big) \\ &= \left(\int_0^{2\delta - \varepsilon} \frac{1}{2\delta} dy\right) \cdot \left(\int_{-\varepsilon}^0 \frac{1}{2\delta} dy\right) + \left(\int_0^\varepsilon \frac{1}{2\delta} dy\right) \cdot \left(\int_{-2\delta + \varepsilon}^0 \frac{1}{2\delta} dy\right) = \varepsilon \cdot \frac{(2\delta - \varepsilon)}{2\delta^2}\end{aligned}$$

for every $i \in \{C, D\}$.

the current-payoff relevant signal is as informative as the private history to predict the current payoff-relevant state. Markov payoff holds by the same argument and the fact that flow payoffs depend only on the current action profile and the payoff-relevant private signals.

Application 3. In asynchronous games, if the active player perfectly observes the current payoff-relevant state, then both Markov information and payoff are satisfied. Indeed, as private histories do not record inactive periods, the current payoff-relevant signals coincide with the payoff-relevant state, satisfying Markov information. Markov payoff follows by the fact that the payoff received during inactive periods equals zero.

Asynchronous revision games. Let $\Omega = \cup_{t\in\mathbb{N}}(X^t \times [0, T))$ with $T \in \mathbb{R}_+ \cup \{\infty\}$, representing the history of action profiles and the timing of the current move. For $i \in N$, assume $\Omega^R = \cup_{t\in\mathbb{N}}(X^t \times [0, T))$ and $\gamma_i^R(\omega^R) = \omega^R$ for each $\omega^R \in \Omega^R$. For $\ell \in \mathbb{N}, h^\ell \in \mathcal{H}^\ell$, $\gamma_i^{\not R}(h^\ell) = (t_1, \ldots, t_{\ell-1})$, that is, the previous moving times are part of the payoff-irrelevant signals. The rest of the environment is as in games with stochastic moves opportunities described in Section 5, with the additional assumption that opportunities are drawn independently at exogenous Poisson rates. Payoff boundedness and continuity are satisfied under standard restrictions concerning payoffs. For instance, payoff boundedness is satisfied by the hypotheses of Lemma 3. MAC(a) holds across periods since timing of moves are drawn at Poisson rates, and the history of actions profile is countable, while MAC(b) holds as there are finitely many actions. Markov information and payoff hold since the active player perfectly observes the current payoff-relevant state.

Dynamic cheap talk games. The sender observes the whole history of the game, including the current exogenous state $\hat{\omega} \in \hat{\Omega}$, and selects a message $m \in M$; the receiver, observing only the sender's messages, implements an action $a \in A$; the exogenous state transition $\mu : A \to \Delta(\hat{\Omega})$ depends on the current action played by the receiver; payoffs depend on the current state and receiver's action, $(\hat{\omega}, a) \in \hat{\Omega} \times A$. We assume the sets M and A are finite to ensure that payoff boundedness, payoff continuity, and MAC hold.

We can write the set of the states of the world as $\Omega^R = \hat{\Omega} \times M \times A$, incorporating the current message and previous actions played by the receiver. For $i \in \{s, r\}$, representing the sender and the receiver, respectively, the current payoff-relevant signals are: $\mathcal{s}_s^R(h_s) = (\hat{\omega}_t^t, a_{t-1}^{t-1})$ for $h_s = (\hat{\omega}^t, m^{t-1}, a^{t-1}) \in \mathcal{H}_s$; $\mathcal{s}_r^R(h_r) = (a_{t-1}^{t-1}, m_t^t)$ for $h_r = (a^{t-1}, m^t) \in \mathcal{H}_r$. Markov information is satisfied as, before moving, the sender observes the current exogenous state and the previous receiver's action, which constitute the current payoff-relevant state, and the receiver's inference about the current exogenous state depends only on her previous actions played and the current message received, which constitute her current payoff-relevant signals. Markov payoff holds by the same argument.

B.9 Conditional ν-equilibrium

We discuss the relation between our notion of $\tilde{\varepsilon}$-constrained equilibrium and the one of *conditional ν-equilibrium*, where, conditional on every set of private histories occurring with positive probability, each player's strategy is a best reply to the opponents' strategies up to ν payoff. The following definition adapts Myerson and Reny's (2020) conditional ν-equilibrium to our setting.

Definition 5. Let $\nu > 0$. A strategy profile $\sigma \in \Sigma$ is a *conditional ν-equilibrium* if, for every $i \in N$, $\ell \in \mathbb{N}$ and $Z \in \mathcal{M}(\mathcal{H}_i^\ell)$ satisfying $P_i(Z|\sigma) > 0$,

$$U_i(\sigma|Z,\sigma) + \nu \geq U_i(\sigma_i', \sigma_{-i}|Z,\sigma), \qquad \forall \sigma_i' \in \Sigma_i.$$

The following assumption requires that players' expected payoffs decrease at a sufficiently fast rate over periods. Let $\bar{U}_i(\hat{\sigma}|Z,\sigma)$ be calculated as $U_i(\hat{\sigma}|Z,\sigma)$ except that g_i replaced by $|g_i|$, and let $U_{i,t}(\hat{\sigma}|Z,\sigma) = \sum_{\tau \leq t} \int_{\mathcal{H}^\tau} U_{i,t}(\hat{\sigma}|\omega^\tau, a^{\tau-1})\, dP_i^\tau(\omega^\tau, a^{\tau-1}|Z,\sigma)$.

Assumption (Regularity conditions)**.** *The following holds:*

1. *For every $\varepsilon > 0$, there is $t(\varepsilon) \in \mathbb{N}$ such that*

$$\sup\{|\textstyle\sum_{t=t(\varepsilon)}^{\infty} U_{i,t}(\hat{\sigma}|Z,\sigma)|\,|i \in N,\, Z \in \mathcal{M}(\mathcal{H}_i^\ell),\, \ell \in \mathbb{N},\, \sigma, \hat{\sigma} \in \Sigma\} < \varepsilon,$$

   and $(1-\varepsilon)^{t(\varepsilon)} \to 1$ as $\varepsilon \to 0$.

2. *For every $Z \in \mathcal{M}(\mathcal{H}_i^\ell)$, $\ell \in \mathbb{N}$, $\sup_{\hat{\sigma},\sigma \in \Sigma} \bar{U}_i(\hat{\sigma}|Z,\sigma) < \infty$.*

3. *For every $\hat{\sigma} \in \Sigma$, $\tau \in \mathbb{N}$, and $a^{\tau-1} \in X^{\tau-1}$, there is a transition probability $\mu_{\omega,a}^\tau(\cdot, a^{\tau-1}|\cdot,\hat{\sigma}) : \mathcal{H}_i \to \Delta(\Omega^\tau)$ such that*

$$d\mu_{\omega,a}^\tau(\omega^\tau|a^{\tau-1}) \cdot p(a^{\tau-1}|\omega^{\tau,(\tau-1)},\hat{\sigma}) = d\mu_{\omega,a}^\tau(\omega^\tau, a^{\tau-1}|h_i,\hat{\sigma}) \times dP_i(h_i|\hat{\sigma}).$$

The regularity conditions (RC) hold in any game that has finitely many periods or discounted payoffs. The following proposition shows that $\tilde{\varepsilon}$-constrained equilibria are conditional ν-equilibria, relating the two concepts.

Proposition 6. *If payoff boundedness, payoff continuity, SAC, and RC hold, then, for every $\varepsilon > 0$, $\tilde{\varepsilon} \in \mathcal{E}(\varepsilon)$, and $\tilde{\varepsilon}$-constrained equilibrium σ^ε, there exists $\nu(\varepsilon) > 0$ such that $\lim_{\varepsilon\to 0} \nu(\varepsilon) = 0$, and σ^ε is a conditional $\nu(\varepsilon)$-equilibrium.*

Proof. The notation used in this proof is introduced in Supplemental Appendix B.5.

Lemma 14. *Under payoff boundedness, payoff continuity, SAC, and RC, the following statements hold:*

1. *For every $\ell \in \mathbb{N}$, $Z \in \mathcal{M}(\mathcal{H}_i^\ell)$ with $P_i(Z|\sigma) > 0$, and for every strategy $\hat{\sigma} \in \Sigma$, there exists a strategy σ_i' such that*

$$\sigma_i' \in \arg\max\{U_i(\tilde{\sigma}_i, \hat{\sigma}_{-i}|Z, \hat{\sigma}_i)|\tilde{\sigma}_i \in \Sigma_i(\hat{\sigma}_i, Z)\}. \tag{15}$$

2. *If (σ_i, σ_{-i}) satisfies condition (12) with $\tilde{\Sigma}_i = \Sigma_i$, then there is a pure strategy $a_i : \mathcal{H}_i \to X_i$ such that $U_i(\sigma_i, \sigma_{-i}) = U_i(a_i(\cdot), \sigma_{-i})$.*
3. *Let $\ell \in \mathbb{N}$ and $Z \in \mathcal{M}(\mathcal{H}_i^\ell)$. For every $\sigma_i' \in \Sigma_i(\hat{\sigma}_i, Z)$ that satisfies condition (15), there is a strategy $\sigma_i'' \in \Sigma_i(\hat{\sigma}_i, Z)$ that is pure at every $h_i \in \mathcal{H}_i(Z)$, such that $U_i(\sigma_i', \hat{\sigma}_{-i}|Z, \hat{\sigma}_i) = U_i(\sigma_i'', \hat{\sigma}_{-i}|Z, \hat{\sigma}_i)$.*

Proof. **Point 1.** It follows from the fact that $\check{\Sigma}_i := \{\sigma \in \hat{\Sigma}|\sigma_i(\cdot|s_i^t, a_i^{t-1}) = \hat{\sigma}_i(\cdot|s_i^t, a_i^{t-1}), (s_i^t, a_i^{t-1}) \in S_i^t \times \bar{X}_i^t, \hbar_i(s_i^t, a_i^{t-1}) \in \mathcal{H}_i^r, r < \ell \leq t\}$, the set of extended strategies coinciding with $\hat{\sigma}_i$ at all private histories of length less than ℓ, is closed in the weak topology, and, therefore, it is compact. Under payoff boundedness and continuity, SAC, and RC, by Lemma 7, $U_i(\cdot, \hat{\sigma}_{-i})$ is continuous and, therefore, attains its maximum, which we term σ_i'', over $\check{\Sigma}_i$.[16] Now, σ_i'' extends a strategy that satisfies the best response equation (12) with $\tilde{\Sigma}_i = \check{\Sigma}_i$ and coincides with $\hat{\sigma}_i$ in private histories of length less than ℓ. Therefore, $(\hat{\hat{\sigma}}_i|_Z, \sigma_i'') \in \check{\Sigma}_i$ for every $\hat{\hat{\sigma}}_i \in \check{\Sigma}_i$. By Lemma 11, σ_i'' satisfies condition (15). By the definition of the conditional expected payoff, $\sigma_i' = (\sigma_i''|_Z, \hat{\sigma}_i)$ satisfies condition (15) as well.

Point 3. Let $\sigma_i' \in \Sigma_i(\hat{\sigma}_i, Z)$ satisfy (15) for some $\ell \in \mathbb{N} \cup \{0\}$, and $Z \in \mathcal{M}(\mathcal{H}_i^\ell)$. Define the correspondence $\xi : Z \rightrightarrows X_i$ as $\xi(h_i) = \{a_i \in X_i|\sigma_i'(a_i|h_i) > 0\}$. By Lemma 15, ξ has a measurable selection $a_{i,\ell}(h_i)$.

[16] See Corollary 2.35 in Aliprantis and Border (2006). Notice that we abuse notation slightly by denoting strategies and their extensions by the same symbol.

We show that $U_i((a_{i,\ell}(\cdot),\sigma_i'),\hat\sigma_{-i}|Z,\hat\sigma) = U_i(\sigma_i',\hat\sigma_{-i}|Z,\hat\sigma)$, where $(a_{i,\ell}(\cdot),\sigma_i')$ denotes the strategy that that coincides with $a_{i,\ell}(h_i)$ on every $h_i \in Z$, and is equal to σ_i', otherwise. Suppose $U_i((a_{i,\ell}(\cdot),\sigma_i'),\hat\sigma_{-i}|Z,\hat\sigma) < U_i(\sigma_i',\hat\sigma_{-i}|Z,\hat\sigma)$.

For any strategy $\sigma_i'' \in \Sigma_i(\hat\sigma_i, Z)$, we can write

$$
\begin{aligned}
U_i(\sigma_i'',\hat\sigma_{-i}|Z,\hat\sigma) &= \frac{1}{P_i(Z|\hat\sigma)} \sum_{\substack{\tau\in\mathbb{N},\\ t\geq\tau}} \int_{(\hbar_i^h)^{-1}(Z)} \int_{\mathcal{H}^{\tau+1}} U_{i,t}(\sigma_i'',\hat\sigma_{-i}|\omega^{\tau+1},a^\tau)\cdot p_i(a_i^\tau|\omega^\tau,\sigma_i'') \\
&\qquad \cdot p_{-i}(a_{-i}^\tau|\omega^\tau,\hat\sigma_{-i})\, d\mu_{\omega,a}^{\tau+1}(\omega^{\tau+1},a^\tau|\omega^\tau,a^{\tau-1}) d\mu_{\omega,a}^\tau(\omega^\tau,a^{\tau-1}) \\
&= \frac{1}{P_i(Z|\hat\sigma)} \sum_{\substack{\tau\in\mathbb{N},\\ t\geq\tau}} \int_Z \int_{\mathcal{H}^\tau} \int_{\mathcal{H}^{\tau+1}} U_{i,t}(\sigma_i'',\hat\sigma_{-i}|\omega^{\tau+1},a^\tau)\cdot \sigma_i''(a_{i,\tau}^\tau|\omega^\tau,a^{\tau-1})\cdot\hat\sigma_{-i}(a_{-i,\tau}^\tau|\omega^\tau,a^{\tau-1}) \\
&\qquad d\mu_{\omega,a}^{\tau+1}(\omega^{\tau+1},a^\tau|\omega^\tau,a^{\tau-1}) d\mu_{\omega,a}^\tau(\omega^\tau,a^{\tau-1}|h_i,\hat\sigma) dP_i(h_i|\hat\sigma),
\end{aligned}
$$

where $U_{i,\tau}(\sigma_i'',\hat\sigma_{-i}|\omega^{\tau+1},a^\tau) := g_i(\omega^{\tau+1,(\tau)},a^\tau)$ and the last equality follows from the definition of P_i and RC.3. Therefore, $U_i((a_{i,\ell}(\cdot),\sigma_i'),\hat\sigma_{-i}|Z,\hat\sigma) < U_i(\sigma_i',\hat\sigma_{-i}|Z,\hat\sigma)$ implies that there is $\hat Z \in \mathcal{M}(\mathcal{H}_i^\ell)$, $\hat Z \subseteq Z$ with $P_i(\hat Z) > 0$ such that for every $h_i \in \hat Z$

$$
\begin{aligned}
&\sum_{\substack{\tau\in\mathbb{N},\\ t\geq\tau}} \int_{\mathcal{H}^{\tau+1}} U_{i,t}(\sigma_i',\hat\sigma_{-i}|\omega^{\tau+1},a^\tau)\cdot\sigma_i'(a_{i,\tau}^\tau|\omega^\tau,a^{\tau-1})\cdot\hat\sigma_{-i}(a_{-i,\tau}^\tau|\omega^\tau,a^{\tau-1})\cdot d\mu_{\omega,a}^{\tau+1}(\omega^{\tau+1},a^\tau|h_i,\hat\sigma) \\
> &\sum_{\substack{\tau\in\mathbb{N},\\ t\geq\tau}} \int_{\mathcal{H}^{\tau+1}} U_{i,t}(\hat\sigma_i,\hat\sigma_{-i}|\omega^{\tau+1},a^\tau)\cdot a_{i,\ell}(a_{i,\tau}^\tau|\omega^\tau,a^{\tau-1})\cdot\hat\sigma_{-i}(a_{-i,\tau}^\tau|\omega^\tau,a^{\tau-1})\cdot d\mu_{\omega,a}^{\tau+1}(\omega^{\tau+1},a^\tau|h_i,\hat\sigma)
\end{aligned}
$$

where $a_{i,\ell}(a_{i,\tau}^\tau|\omega^\tau,a^{\tau-1}) = 1$ if $a_{i,\tau}^\tau = a_{i,\ell}(\hbar_i^h(\omega^\tau,a^{\tau-1}))$, and $d\mu_{\omega,a}^{\tau+1}(\omega^{\tau+1},a^\tau|h_i,\hat\sigma) = d\mu_{\omega,a}^{\tau+1}(\omega^{\tau+1},a^\tau|\omega^\tau,a^{\tau-1})d\mu_{\omega,a}^\tau(\omega^\tau,a^{\tau-1}|h_i,\hat\sigma)$.

Now, notice that we can write $\sigma_i'(a_{i,\tau}^\tau|\omega^\tau,a^{\tau-1}) = \sigma_i'(a_{i,\tau}^\tau|\omega^\tau,a^{\tau-1})(1 - a_{i,\ell}(a_{i,\tau}^\tau|\omega^\tau,a^{\tau-1})) + \sigma_i'(a_{i,\tau}^\tau|\omega^\tau,a^{\tau-1})a_{i,\ell}(a_{i,\tau}^\tau|\omega^\tau,a^{\tau-1})$. Therefore, the previous inequality implies

$$
\begin{aligned}
U_i(\sigma_i',\hat\sigma_{-i}|\hat Z,\hat\sigma)\cdot P_i(\hat Z|\hat\sigma) &< \sum_{\substack{\tau\in\mathbb{N},\\ t\geq\tau}} \int_{\hat Z} \int_{\mathcal{H}^{\tau+1}} U_{i,t}(\sigma_i',\hat\sigma_{-i}|\omega^{\tau+1},a^\tau)\cdot\sigma_i'(a_{i,\tau}^\tau|\omega^\tau,a^{\tau-1}) \\
&\qquad \cdot\Big(1 - a_{i,\ell}(a_{i,\tau}^\tau|\omega^\tau,a^{\tau-1})\Big)\cdot\hat\sigma_{-i}(a_{-i,\tau}^\tau|\omega^\tau,a^{\tau-1})\, d\mu_{\omega,a}^{\tau+1}(\omega^{\tau+1},a^\tau|h_i,\hat\sigma)dP_i(h_i|\hat\sigma) \\
&+ \sum_{\substack{\tau\in\mathbb{N},\\ t\geq\tau}} \int_{\hat Z} \int_{\mathcal{H}^{\tau+1}} U_{i,t}(\sigma_i',\hat\sigma_{-i}|\omega^{\tau+1},a^\tau)\cdot\sigma_i'(a_{i,\tau}^\tau|\omega^\tau,a^{\tau-1})\cdot\sigma_i'(a_{i,\ell}(\hbar_i^h(\omega^\tau,a^{\tau-1}))|\omega^\tau,a^{\tau-1}) \\
&\qquad \cdot\hat\sigma_{-i}(a_{-i,\tau}^\tau|\omega^\tau,a^{\tau-1})\, d\mu_{\omega,a}^{\tau+1}(\omega^{\tau+1},a^\tau|h_i,\hat\sigma)dP_i(h_i|\hat\sigma).
\end{aligned}
$$

where we used Tonelli's theorem to write the bound as a sum of two integrals.

However, this last inequality contradicts the optimality of σ_i': the payoff on the right hand side of the inequality is the payoff from the strategy that instead of putting weight $\sigma_i'(a_{i,\ell}(h_i)|h_i) > 0$ on $a_{i,\ell}(h_i)$ randomizes according to σ_i', for $h_i \in \hat{Z}$. From the inequality, this alternative strategy yields a higher expected conditional payoff. Arguing recursively, by defining a new correspondence from $\mathcal{H}_i(Z) \cap \mathcal{H}_i^\tau$ to X_i, for each $\tau > \ell$, and finding a measurable selection $a_{i,\tau}$ we can construct σ_i''.

Point 2 can be shown by an analogous argument. □

Let $\sigma^* \in \Sigma(\tilde{\varepsilon})$. Let us show that, if $\tilde{\varepsilon}$ is an ε-tremble, then it is a conditional ν-equilibrium for every $\nu > \nu(\varepsilon)$, where $\nu : [0, 1] \to (0, \infty)$ satisfies $\lim_{\varepsilon \to 0} \nu(\varepsilon) = 0$.

Suppose there is $\nu > 0$, $Z \in \mathcal{M}(\mathcal{H}_i)$, and $\sigma_i' \in \hat{\Sigma}_i$, such that

$$U_i(\sigma^*|Z, \sigma^*) + \nu < U_i(\sigma_i', \sigma_{-i}^*|Z, \sigma^*). \tag{16}$$

Let $\hat{\sigma}_i' \in \operatorname{argmax}\{U_i(\sigma_i', \sigma_{-i}^*|Z, \sigma^*)|\sigma_i' \in \Sigma_i(\sigma_i^*, Z)\}$. The strategy $\hat{\sigma}_i'$ exists by Lemma 14.1, and by Lemma 14.3 it can be chosen so that it is pure in $\mathcal{H}_i(Z)$. Let $a_i : \mathcal{H}_i(Z) \to X_i$ be the function that yields the action that is drawn with probability 1 by $\hat{\sigma}_i'(\cdot|h_i)$ for $h_i \in \mathcal{H}_i(Z)$.

For each $h_i \in \mathcal{H}_i(Z)$ define $\beta(h_i) \coloneqq \sum_{a_i \in X_i \setminus a_i(h_i)} \tilde{\varepsilon}(h_i, a_i)$. By Lemma 15, β is $\mathcal{M}(\mathcal{H}_i)$-measurable. Define the strategy

$$\tilde{\sigma}_i(a_i|h_i) \coloneqq \begin{cases} 1 - \beta(h_i) & \text{if } a_i = a_i(h_i),\ h_i \in \mathcal{H}_i(Z) \\ \tilde{\varepsilon}(h_i, a_i) & \text{if } a_i \neq a_i(h_i),\ h_i \in \mathcal{H}_i(Z) \\ \sigma_i^*(a_i|h_i) & \text{otherwise.} \end{cases}$$

Notice that, $\tilde{\sigma}_i$ is $\tilde{\varepsilon}$-constrained.

Define $H^t(Z) = \{(\omega^t, a^t) \in \Omega^t \times X^t | \hbar_i^h(\omega^t, a^{t,(t-1)}) \in \mathcal{H}_i(Z)\}$ and $\alpha(\omega^t, a^t) = \Pi_{\tau \leq |\hbar^h{}_i(\omega^t, a^{t,(t-1)})|}(1 - \beta(\hbar_i^h(\omega^{t,(\tau)}, a^{t,(\tau-1)})))$ for $(\omega^t, a^t) \in \Omega^t \times X^t$.

For $\sigma \in \Sigma$ and $\sigma_i' \in \Sigma_i(\sigma_i, Z)$, we can write

$$\begin{aligned} U_{i,t}(\tilde{\sigma}_i, \sigma_{-i}^*|Z, \sigma^*) \cdot P_i(Z|\sigma^*) &= \int_{H^t(Z)} g_i(\omega^t, a^t) \cdot p_i(a_i^t|\omega^t, \tilde{\sigma}_i) p_{-i}(a_{-i}^t|\omega^t, \sigma_{-i}^*) \, d\mu_{\omega,a}^t(\omega^t, a^t) \\ = \int_{H^t(Z)} & g_i(\omega^t, a^t) \cdot p_i(a_i^t|\omega^t, \hat{\sigma}_i') \cdot p_{-i}(a_{-i}^t|\omega^t, \sigma_{-i}^*) \cdot \alpha(\omega^t, a^t) \, d\mu_{\omega,a}^t(\omega^t, a^{t-1}) \\ &+ \int_{H^t(Z)} g_i(\omega^t, a^t) \cdot m(\tilde{\varepsilon}, \omega^t, a^t) \cdot p_{-i}(a_{-i}^t|\omega^t, \sigma_{-i}^*) \, d\mu_{\omega,a}^t(\omega^t, a^t), \end{aligned}$$

where $m(\tilde{\varepsilon}, \omega^t, a^t) > 0$ is a function of $\tilde{\varepsilon}$ such that $m(\tilde{\varepsilon}, \omega^t, a^t) < \varepsilon$ for every $(\omega^t, a^t) \in \mathcal{H}_i(Z) \cap \mathcal{H}^t$. Notice that $U_i(\sigma_i', \sigma_{-i}|Z, \sigma) = \sum_{t \in \mathbb{N}} U_{i,t}(\sigma_i', \sigma_{-i}|Z, \sigma)$.

Let $t(\varepsilon)$ be as in RC.1, then

$$U_i(\tilde\sigma_i,\sigma^*_{-i}|Z,\sigma^*) \geq \sum_{t\leq t(\varepsilon)} U_{i,t}(\tilde\sigma_i,\sigma^*_{-i}|Z,\sigma^*) - \varepsilon \geq \sum_{t\leq t(\varepsilon)} U_{i,t}(\hat\sigma'_i,\sigma^*_{-i}|Z,\sigma^*) - \varepsilon$$
$$-\frac{1}{P_i(Z|\sigma^*)}\Bigg(\sum_{t\leq t(\varepsilon)}\int_{H^t(Z)}(1-\alpha(\omega^t,a^t))|g_i(\omega^t,a^t)|\cdot\tilde p_i(a^t_i|\omega^t,\hat\sigma'_i)\cdot p_{-i}(a^t_{-i}|\omega^t,\sigma^*_{-i})\,d\mu^t_{\omega,a}(\omega^t,a^t)$$
$$\sum_{t\leq t(\varepsilon)}\int_{H^t(Z)}|g_i(\omega^t,a^t)|\cdot m(\tilde\varepsilon,\omega^t,a^t)\cdot p_{-i}(a^t_{-i}|\omega^t,\sigma^*_{-i})\,d\mu^t_{\omega,a}(\omega^t,a^t)\Bigg)$$
$$\geq U_i(\hat\sigma'_i,\sigma^*_{-i}|Z,\sigma^*) - 2\varepsilon - (1-(1-\varepsilon)^{t(\varepsilon)}+\varepsilon)\sup_{\hat\sigma,\sigma\in\Sigma}\bar U_i(\hat\sigma|Z,\sigma),$$

where in the first and third inequality we use RC.1, and in the third inequality we use $\alpha(\omega^t,a^t)\geq(1-\varepsilon)^{t(\varepsilon)}$ for $t\leq t(\varepsilon)$.

Define $\nu(\varepsilon) \coloneqq 2\varepsilon + (1-(1-\varepsilon)^{t(\varepsilon)}+\varepsilon)\sup_{\hat\sigma,\sigma\in\Sigma}\bar U_i(\hat\sigma|Z,\sigma)$. If $\nu > \nu(\varepsilon)$ then $U_i(\sigma^*|Z,\sigma^*)+\nu < U_i(\hat\sigma'_i,\sigma^*_{-i}|Z,\sigma^*) \leq U_i(\tilde\sigma_i,\sigma^*_{-i}|Z,\sigma^*)+\nu(\varepsilon)$, which is implied by equation (16), contradicts that σ^* is an $\tilde\varepsilon$-constrained equilibrium. By RC, $\lim_{\varepsilon\to 0}\nu(\varepsilon)=0$.

C Supplemental Appendix — Mathematical Results

C.1 Carathéodory integrands and measurability of weak limits

Let $(Y,\mathcal{M}(Y),\beta)$ be a measure space, Z be a countable metric space endowed with the σ-algebra of all subsets of Z and the counting measure, and $\mathcal{M}(Y)_0$ be a sub σ-algebra of $\mathcal{M}(Y)$.

The following result shows that a correspondence from Y to Z, which may not be closed-valued, has a measurable selection under a condition weaker than measurability. It relies on the countability of the set Z.

Lemma 15. *Let $\phi : Y \rightrightarrows Z$ be a non-empty valued correspondence such that for every $z\in Z$ the set $\{y\in Y|z\in\phi(y)\}$ is $\mathcal{M}(Y)_0$-measurable. Then ϕ has a $\mathcal{M}(Y)_0$-measurable selection, and for any $\mathcal{M}(Y)_0\times\mathcal{M}(Z)$-measurable, real valued function $\hat g$, $\sum_{z\in\phi(y)}\hat g(y,z)$ is $\mathcal{M}(Y)_0$-measurable.*

Proof. Let $(z_j)_{j\in\mathbb{N}}$ be an enumeration of the set Z. Define the function

$$m_j(y) \coloneqq \mathbb{1}\{z_j\notin\phi(y)\} - 1/j.$$

Then, the function $\bar m(y) \coloneqq \inf_{j\in\mathbb{N}} m_j(y)$ is strictly negative (ϕ is non-empty valued), finite for each y, and $\mathcal{M}(Y)_0$-measurable. It yields $1/j$ for the smallest j such that $z_j\in\phi(y)$.

The selection $z(y) = z_{|1/\bar{m}(y)|}$ is $\mathcal{M}(Y)_0$-measurable. In fact, for a measurable set $\hat{Z} \in \mathcal{M}(Z)$, we can write $\hat{Z} = (z_{n_j})_{n_j \in \mathbb{N}}$ for some subsequence $(n_j)_{j\in\mathbb{N}}$ of $\mathbb{N}$. Then $z^{-1}(\hat{Z}) = \{y \in Y | \, |1/\bar{m}(y)| \in (n_j)_{j\in\mathbb{N}}\}$ is in $\mathcal{M}(Y)_0$ since all countable subsets of $\mathbb{R}$ are measurable.

Let $\hat{m}_k(y) = \sum_{j=1}^{k} \hat{g}(y, z_j) \cdot \mathbb{1}\{z_j \in \phi(y)\}$, then $\hat{m}_k(y)$ is measurable and, therefore, $\lim_{k\to\infty} \hat{m}_k(y) = \sum_{z\in\phi(y)} \hat{g}(y, z)$ is $\mathcal{M}(Y)_0$-measurable. □

The following corollary is a consequence of Proposition 5.

Corollary 2. *Let $\varphi \in CI(Y \times Z, \mathcal{M}(Y), \beta)$ and let $\hat{\varphi} = \mathbb{E}(\varphi|\mathcal{M}(Y)_0 \otimes \mathcal{M}(Z))$ denote the conditional expectation of φ with respect to $\mathcal{M}(Y)_0 \otimes \mathcal{M}(Z)$. Then there is a version of $\hat{\varphi} \in CI(Y \times Z, \mathcal{M}(Y)_0, \beta)$.*

The following example shows that continuity in the strong total variation norm is stronger than continuity in the total variation norm.

Example 4. Let $\hat{\beta}$ be a the uniform distribution over $Y = (0, 1)$, and define $Z := \{z_{(n,k)} | z_{(n,k)} = \left(1/n, k/n^2\right), k, n \in \mathbb{N}, k \in [1, n-1]\} \cup (0, 0)$. Let $\varphi(y, (0, 0)) = 1$ for every $y \in Y$, and

$$\varphi(y, z_{(n,k)}) = \begin{cases} 0 & \text{if } y \in [k-1/n, k/n] \\ 2 & \text{if } y \in (k/n, (k+1)/n] \\ 1 & \text{otherwise.} \end{cases}$$

Define $\xi(\hat{Y}|z) = \int_{\hat{Y}} \varphi(y, z)\, d\hat{\beta}(y)$ for each $\hat{Y} \in \mathcal{M}(Y)$.

Now take a sequence $z_m \to (0, 0)$. We can write $z_m = z_{(n_m,k_m)}$ with $k_m/n_m \to 0$. We have

$$\sup_{\pi\in\pi(Y)} \sum_{\tilde{Y}\in\pi} \int_{\tilde{Y}} |\varphi(y, z_m) - \varphi(y, (0, 0))|\, d\hat{\beta}(y) \le 2/n_m,$$

and, hence, $\xi(\cdot|z)$ is continuous in z in the total variation norm. Now, consider the y dependent sequence $z_n(y) = z_{(n,k(y))}$ with $k(y) = \lceil y \cdot n \rceil$. Then for each $n \in \mathbb{N}, y \in Y$, $|z_{(n,k(y))} - (0, 0)| < 2/n$, while $|\varphi(y, z_{(n,k(y))}) - \varphi(y, (0, 0))| = 1$, and, therefore, $||\,(\xi - \xi^{(0,0)})|_{B((0,0),1/n)}||_{SV} \ge 1$. This shows that $\xi(\cdot|z)$ is not continuous in z in the strong total variation norm. ◀

Let Z and $\hat{Z}$ be countable sets endowed with the σ-algebra of all of their subsets and the counting measure. Let $T_y : Y \to \hat{Y}$ and $T_z : Z \to \hat{Z}$ be measurable functions, and $(Y, \mathcal{M}(Y), \beta)$ and $(\hat{Y}, \mathcal{M}(\hat{Y}), \hat{\beta} := \beta \circ T_y^{-1})$ be measure spaces.[17]

[17] The measure $\beta \circ T_y^{-1}$ is defined by $\beta \circ T_y^{-1}(\hat{B}) = \beta(T_y^{-1}(\hat{B}))$ for each $\hat{B} \in \mathcal{M}(\hat{Y})$.

We say that a net of measurable functions $(f^\lambda)_{\lambda\in\Lambda} \subseteq \mathbb{R}^{Y\times Z}$ converges to f if $I_\phi(f^\lambda) \coloneqq \int_Y \sum_{z\in Z} \phi(y,z) f^\lambda(y,z)\,dy$ converges to $I(f)$ for every $\phi \in CI(Y\times Z,\beta)$. The convergence for functions in $\mathbb{R}^{\hat{Y}\times\hat{Z}}$ is defined analogously.

Lemma 16. *Let $\mathcal{M}(Y\times Z)_0 = \sigma\{(T_y,T_z)^{-1}(\hat{B}) \cap \tilde{B} | \hat{B} \in \mathcal{M}(\hat{Y}\times\hat{Z})\}$, for a set $\tilde{B} \in \mathcal{M}(Y\times Z)$. Suppose that $(f^\lambda)_{\lambda\in\Lambda}$ is a net of functions each of which is measurable with respect to $\mathcal{M}(Y\times Z)_0$ and has support in $\tilde{B}$. If $(f^\lambda)_{\lambda\in\Lambda}$ converges to f, then f is measurable with respect to $\mathcal{M}(Y\times Z)_0$ and has support in $\tilde{B}$.*

Proof. Let $T = (T_y,T_z)$. By Theorem 4.41 in Aliprantis and Border (2006), we can write $f^\lambda(y,z) = \hat{f}^\lambda(T(y,z))$, where $\hat{f}^\lambda$ is a $\mathcal{M}(\hat{Y}\times\hat{Z})$-measurable function for each $\lambda\in\Lambda$. Let $\hat{f}$ be the limit of a subnet of $(\hat{f}^\lambda)_{\lambda\in\Lambda}$. We show that $\hat{f}(T(y,z)) = f(y,z)$ for each $y \in Y$, establishing the desired conclusion by the same theorem.

Suppose not, then there is a test function $\psi \in CI(Y\times Z,\beta)$ such that $I_\psi(f) \neq I_\psi(\hat{f}\circ T)$. Let $\bar{\psi}$ be the measurable function such that for every $B \in \mathcal{M}(Y\times Z)_0$, $\int_Y \sum_{z\in Z} \bar{\psi}(y,z)\mathbb{1}\{(y,z)\in B\}\,d\beta(y) = \int_Y \sum_{z\in Z} \psi(y,z)\mathbb{1}\{(y,z)\in B\}\,d\beta(y)$ for every $B \in \mathcal{M}(Y\times Z|T)$. Since the measure $\hat{\mu}(B) = \int_Y \sum_{z\in Z} \mathbb{1}\{(y,z)\in B\}\,d\beta(y)$ for $B \in \mathcal{M}(Y\times Z|T)$ is σ-finite, $\bar{\psi}$ exists by the Radon-Nikodym theorem.

By Theorem 4.41 in Aliprantis and Border (2006), there is a measurable function $\hat{\psi} : \hat{Y}\times\hat{Z} \to \mathbb{R}$ such that $\hat{\psi} = \bar{\psi}\circ T$, and, by Theorem 13.46, for every measurable $g : Y\times Z \to \mathbb{R}$ with support in $T(\tilde{B})$

$$\hat{I}_{\hat{\psi}}(g) \coloneqq \int_{\hat{Y}} \sum_{z\in\hat{Z}} \hat{\psi}(y,z)\cdot g(\hat{y},\hat{z})\,d\hat{\beta}(\hat{y}) = \int_Y \sum_{z\in Z} \psi(y,z)\cdot g\circ T(y,z)\,d\beta(\hat{y}) \tag{17}$$

By Corollary 2, $\hat{\psi}(y,z)$ is continuous in $\hat{Z}$, for $\hat{y}$ in a $\hat{\beta}$-full-measure subset of $\hat{Y}$.[18] This contradicts $\lim_\lambda I_\psi(f^\lambda) \neq I_\psi(\hat{f}\circ T)$ as we have $\lim_\lambda I_\psi(f^\lambda) = \lim_\lambda I_\psi(\hat{f}^\lambda\circ T) = \lim_\lambda \hat{I}_{\hat{\psi}}(\hat{f}^\lambda) = \hat{I}_{\hat{\psi}}(\hat{f}) = I_\psi(\hat{f}\circ T)$, where the third equality follows by the definition of $\hat{f}$, since $\hat{\psi} \in CI(\hat{Y}\times Z,\hat{\beta})$, and the second and last equalities follow by (17). □

C.2 A special case of dominated convergence for nets

The following results show that results analogous to Fatou's lemma and Lebesgue's dominated convergence theorem hold not just for sequences but also for nets under the counting measure. The proofs of Lemma 17 and Proposition 8 are close to the standard ones—except that one replaces sequences by nets—and use Proposition 7.

[18] By Corollary 2, $\bar{\psi}$ is continuous in Z almost surely in Y. This implies the continuity of $\hat{\psi}_i$.

We include the proof of these Lemmas for completion. Proposition 7 does not hold when one replaces the sum by a measure over an uncountable set.

Proposition 7 (Monotone convergence). *Let $(x_{\alpha,k})_{\alpha\in\mathcal{A},k\in\mathbb{N}} \subseteq \mathbb{R}_+$ be a net, with $\mathcal{A}$ a directed set, be such that (i) $\alpha \geq \alpha'$ implies $x_{\alpha,k} \geq x_{\alpha',k}$ for all $k \in \mathbb{N}$, and (ii) there is M such that $\sum_{k=1}^{\infty} x_{\alpha,k} < M$ for all $\alpha \in \mathcal{A}$ then $x_k = \lim_\alpha x_{\alpha,k}$ exists for each $k \in \mathbb{N}$ and $\lim_\alpha \sum_{k=1}^{\infty} x_{\alpha,k} = \sum_{k=1}^{\infty} x_k$.*

Proof. The net $x_{\alpha,k}$ has a limit, x_k, for each k, as it is non-decreasing and bounded by (ii).

Let $B = \{(\alpha, n) | \alpha \in \mathcal{A}, n \in \mathbb{N}\}$ be a directed set with $(\alpha, n) \geq (\alpha', n')$ if and only if $\alpha \geq \alpha'$ and $n \geq n'$. And define $y_{\alpha,n} := \sum_{k=1}^{n} x_{\alpha,k}$. The net $(y_b)_{b\in B}$ is non-decreasing and bounded and, therefore, has a limit—its supremum in $\mathbb{R}$—which we denote s. That is, for each $\varepsilon > 0$ there is $\bar{\alpha}$ and $\bar{n}$ such that $\alpha \geq \bar{\alpha}$ and $n \geq \bar{n}$ implies $\left|\sum_{k=1}^{n} x_{\alpha,k} - s\right| < \varepsilon$. For each $\alpha \geq \bar{\alpha}$, the sequence $(y_{\alpha,n})_{n\in\mathbb{N}}$ is monotone non-decreasing and bounded. Therefore, it has a limit and by the continuity of the absolute value we obtain, $\left|\sum_{k=1}^{\infty} x_{\alpha,k} - s\right| < \varepsilon$. This shows that $\lim_\alpha \sum_{k=1}^{\infty} x_{\alpha,k} = s$.

For each fixed $n \geq \bar{n}$, there is $\hat{\alpha}(n) \geq \bar{\alpha}$ such that $\left|\sum_{k=1}^{n} x_{\hat{\alpha}(n),k} - \sum_{k=1}^{n} x_k\right| < \varepsilon$. Therefore, for each $n \geq \bar{n}$ there is $\hat{\alpha}(n)$ such that $\left|\sum_{k=1}^{n} x_k - s\right| \leq \left|\sum_{k=1}^{n} x_{\hat{\alpha}(n),k} - \sum_{k=1}^{n} x_k\right| + \left|\sum_{k=1}^{n} x_{\hat{\alpha}(n),k} - s\right| < 2\varepsilon$. This shows that $\sum_{k=1}^{\infty} x_k = s$. □

Lemma 17 (Fatou's Lemma). *Let $(x_{\alpha,k})_{\alpha\in\mathcal{A},k\in\mathbb{N}} \subseteq \mathbb{R}_+$ be a net, with $\mathcal{A}$ a directed set, and $\sup_{\bar{\alpha}} \inf_{\alpha\geq\bar{\alpha}} \left(\sum_{k=1}^{\infty} x_{\alpha,k}\right) < \infty$. Then, for each $k \in \mathbb{N}$, $x_k := \sup_{\bar{\alpha}} \inf_{\alpha\geq\bar{\alpha}} x_{\alpha,k}$ exists and $\sum_{k=1}^{\infty} \sup_{\bar{\alpha}} \inf_{\alpha\geq\bar{\alpha}} x_{\alpha,k} \leq \sup_{\bar{\alpha}} \inf_{\alpha\geq\bar{\alpha}} \sum_{k=1}^{\infty} x_{\alpha,k}$.*

Proof. Define $y_{k,\alpha} := \inf\{x_{k,\hat{\alpha}} | \hat{\alpha} \geq \alpha\}$. Since $\geq$ is transitive $y_{k,\alpha}$ is non-decreasing in α. Then, $(\sum_{k=1}^{\infty} y_{k,\alpha})_{\alpha\in\mathcal{A}}$ is a non-decreasing and bounded net, and, therefore, has a limit. Furthermore, for $k \in \mathbb{N}, \alpha \in \mathcal{A}$ we have $y_{k,\alpha} \leq x_{k,\alpha}$. Therefore, $\lim_\alpha \sum_{k=1}^{\infty} y_{k,\alpha} \leq \sup_{\bar{\alpha}} \inf_{\alpha\geq\bar{\alpha}} \sum_{k=1}^{\infty} x_{k,\alpha} < \infty$. To conclude note that by Proposition 7 the left hand side of the previous expression is equal to $\sum_{k=1}^{\infty} \lim_\alpha y_{k,\alpha} = \sum_{k=1}^{\infty} \sup_{\bar{\alpha}} \inf_{\alpha\geq\bar{\alpha}} x_{\alpha,k}$. □

Proposition 8 (Dominated convergence). *If $x_k = \lim_\alpha x_{k,\alpha}$ for each $k \in \mathbb{N}$, and there is $(y_k)_{k\in\mathbb{N}}$, with $|\sum_{k=1}^{\infty} y_k| < \infty$, such that $|x_{\alpha,k}| \leq y_k$, then $\lim_\alpha \sum_{k=1}^{\infty} |x_{\alpha,k} - x_k| = 0$.*

Proof. Notice that $|x_{\alpha,k} - x_k| \leq |x_{\alpha,k}| + |x_k| \leq 2y_k$. Therefore, by Lemma 17 $\inf_{\bar{\alpha}} \sup_{\alpha\geq\bar{\alpha}} \sum_{k=1}^{\infty} |x_{\alpha,k} - x_k| \leq \sum_{k=1}^{\infty} \inf_{\bar{\alpha}} \sup_{\alpha\geq\bar{\alpha}} |x_{\alpha,k} - x_k| = 0$. This shows that $\lim_\alpha \sum_{k=1}^{\infty} |x_{k,\alpha} - x_k| = 0$. □