University of Glasgow

Adam Smith Business School

WORKING PAPER SERIES

Equilibria Existence in Bayesian Games: Climbing the Countable Borel Equivalence Relation Hierarchy

Ziv Hellman and Yehuda John Levy

Paper no. 2020-15
July 2020

EQUILIBRIA EXISTENCE IN BAYESIAN GAMES: CLIMBING THE COUNTABLE BOREL EQUIVALENCE RELATION HIERARCHY

ZIV HELLMAN & YEHUDA JOHN LEVY

ABSTRACT. The solution concept of a Bayesian equilibrium of a Bayesian game is inherently an interim concept. The corresponding ex ante solution concept has been termed Harsányi equilibrium; examples have appeared in the literature showing that there are Bayesian games with uncountable state spaces that have no Bayesian approximate equilibria but do admit Harsányi approximate equilibrium, thus exhibiting divergent behaviour in the ex ante and interim stages. Smoothness, a concept from descriptive set theory, has been shown in previous works to guarantee the existence of Bayesian equilibria. We show here that higher rungs in the countable Borel equivalence relation hierarchy can also shed light on equilibrium existence. In particular, hyperfiniteness, the next step above smoothness, is a sufficient condition for the existence of Harsányi approximate equilibria in purely atomic Bayesian games.

Keywords: Bayesian games; Equilibrium existence; Borel equivalence relations

1. INTRODUCTION

There are few concepts in contemporary game theory and decision theory as fundamental as the distinctions between the *ex ante*, interim, and *ex post* time periods. A great many theorems and results over the years, in a wide range of applications, could not have been discovered without these distinctions. It follows that sharpening our understanding of different player behaviours between the *ex ante* and interim periods, as well as delineating when these distinctions are of operative significance to begin with, are of foundational importance.

We thank Robert Simon for helpful discussions, and two anonymous reviewers. Ziv Hellman acknowledges research support by Israel Science Foundation grant 1626/18.

The *locus classicus* for the topic is a series of papers composed by John Harsányi in the 1960s. From there flowed an entire corpus of indispensible concepts in use since then, including player types, Bayesian equilibria, common priors, and much more.

In the Harsányi model of games of incomplete information, all of the details relevant to the game being played that are unknown to the players are expressed in the variety of 'types' of the game (or, alternatively, a collection of states of the world), from which hierarchies of beliefs may be formed. Borrowing concepts from the field of Bayesian inference, in the *ex ante* stage the players have prior measures over the collection of possible types. If they all share the same prior, it is said to be a common prior. The players can select strategies in this stage, using these prior measures, and form an equilibrium. This type of equilibrium has not got a uniformly standard name in the literature, but we will follow [Simon, 2003] in calling it an *Harsányi equilibrium*.

When the set of types of a Bayesian game are finite or countable, Bayesian equilibrium and Harsányi equilibrium are equivalent, hence in analysing a Bayesian game one may use either equilibrium concept.[1] When there are uncountably many types, however, the *ex ante* and interim equilibria concepts may diverge. Recent papers by the authors presented conditions guaranteeing when Bayesian equilibria and Harsányi equilibria coincide even in games with uncountably many types. Our focus here in contrast is on the opposite: understanding when such a divergence between the equilibria concepts can occur.

This goal is attained in Theorem 4.13, in which sufficient conditions for a Bayesian game to admit Harsányi ε-equilibria, but no Bayesian ε-equilibrium, are presented. This continues a series of results exhibiting relationships between the countable Borel equivalence relation hierarchy of descriptive set theory and equilibria existence in Bayesian games, and suggests further research into this topic.

1.1. **Background.** The main innovation of the Harsányi model posited that after each player receives private information (or signals), that player's specific type is revealed to him or her. Based on this, in the interim period, each player then forms a posterior probability measure over the types of the other players in a Bayesian manner. Armed

[1] Indeed, the identification of *ex ante* and interim equilibria in Bayesian games in the finite case was one of the central accomplishments of [Harsanyi, 1967].

with this posterior, each player chooses a strategy for optimal payoff in expectation. The game being played here is, as is well known, called a Bayesian game. An equilibrium of the game is a Bayesian equilibrium, which Harsányi proved always exists.

When there are a finite number of states of the world, however, the distinction between the *ex ante* and interim stages turns out to be largely immaterial. In that situation, Bayesian and Harsányi equilibria are the same and can be identified with each other [Zamir, 2008]. Furthermore, Robert Aumann's No Disagreements Theorem (later extended to a more general No Betting Theorem) showed that if players have commonly known agreement on the expected values of random variables in the *ex ante* stage then they will similarly agree on those values in the interim stage, and that the converse also holds. Given all this, some researchers opined that 'the Harsányi ex ante model is just an auxiliary construction for the analysis', essentially a fictional aspect that has no bearing on substantive aspects, and therefore 'the appropriate way of modeling an incomplete information situation is at the interim stage' [Zamir, 2008].

When the cardinality of states is extended to uncountable cardinalities[2], however, this convenient identification of the *ex ante* and interim stages from the perspective of game solutions no longer holds. [Simon, 2003] presented an example of a three-player Bayesian game with continuum many states of the world that fails to admit any measurable Bayesian equilibria, in contrast to Harsányi's finite game result always guaranteeing the existence of Bayesian equilibria.

Subsequently, [Hellman, 2014] presented an example of a two-player Bayesian game that has no measurable Bayesian ε-equilibria for sufficiently small ε. Significantly, that same example *does* admit Harsányi ε-equilibria for all ε. When $\varepsilon = 0$ the Bayesian and Harsányi equilibria can be identified (see [Simon, 2003]); however, for positive ε that identification may fail, as amply exemplified in [Hellman, 2014], drawing a sharp line of divergence between *ex ante* and interim behaviour.

In parallel, [Simon, 2000] and [Lehrer and Samet, 2011] exhibited knowledge structures in which at the *ex ante* stages the players share a common prior but in the interim stage, at each common knowledge component there is no common prior. As shown in [Hellman and Levy, 2017], these anomalies are not structural alone: building on those results

[2] When the cardinality of states is countable, Bayesian equilibria always exist and they are identical with the Harsányi equilibria ([Simon, 2003]).

one could construct examples in which the players can never agree on an acceptable bet at the *ex ante* stage but will always find acceptable bets in the interim stage. The *ex ante* stage, it turns out, is not an auxiliary construction or a convenient modeller's fiction: it can imply significant behavioural divergence from the interim.

In studying this, [Hellman and Levy, 2017] noted that all the anomalous examples shared a common element: structurally, their common knowledge components can be described as hyperfinite countable Borel equivalence classes. The designation *hyperfinite* refers to the complexity, in a topological sense, of the relation. That clue indicated a bridge between these game theoretic concepts and descriptive set theory, which includes a fundamental hierarchy of countable Borel equivalence relations. The lowest rung is that of the smooth equivalence relations, which intuitively can be described as structures in which one can measurably select a point within each equivalence class. All finite equivalence relations are smooth. The next rung up is that of the hyperfinite equivalence classes.

As shown in [Hellman and Levy, 2017] and [Hellman and Levy, 2019], if the common knowledge components of knowledge structures in games of incomplete information are restricted to be smooth countable Borel equivalence relations then all of the anomalies disappear. All the comfortable results familiar from the finite cases carry over nicely: Bayesian equilibria always exist and they are equivalent to Harsányi equilibria; common priors over the entire state space carry over to common priors in the common knowledge components. Harmony between the *ex ante* and interim stages is restored.

Here in contrast we focus on the opposite cases, studying the instances where the *ex ante* and interim strongly diverge. The starting point is the example of [Hellman, 2014]. Its knowledge structure is hyperfinite and not smooth, and it admits no Bayesian ε-equilibria but from the *ex ante* perspective it does have ε-Harsányi equilibria. In a strategy profile satisfying the latter, there may be a small mass of states in which agents can deviate for significant gains, in contrast to the strategy profiles satisfying the former, in which only small gains can be made by deviations at all but a null set of states. Is this a rare example or does it indicate a connection between hyperfinite structures and such behaviour, just as smoothness below it guarantees the existence of Bayesian equilibria? Can we identify when the *ex ante* and interim equilibrium behaviours diverge?

1.2. **Main Results.** In our main theorem here we show that the equilibrium behaviour of that [Hellman, 2014] is indeed preordained by the hyperfinite structure. Every Bayesian game whose common knowledge components structure constitutes a hyperfinite but not smooth countable Borel equivalence relation satisfies the property that it admits no Bayesian approximate equilibrium but admits Harsányi ε-equilibria for every $\varepsilon > 0$. This significantly adds our understanding of the connections between descriptive set theory and game theory.

Our study of Bayesian game equilibrium concepts is paralleled, and assisted, by our study of graphical games, introduced by ([Kearns et al., 2001]). In such games, players are arranged in a graph-like structure, and players' payoffs only depend on the actions of themselves and neighbors. Graphical games arise naturally in the agent-normal form of Bayesian games, in which each origional player is replaced by an individual agent for each possible type of that player, as an agent's payoffs' depend only on his own action and the actions of other types of other players he believes may arise.

To make use of graphical games for the study of Bayesian games, we define a parallel concept of Harsányi ε-equilibria that applies to graphical games. We also introduce a new and stronger equilibrium concept than Harsányi equilibria, which we term strong Harsányi ε-equilibria, with parallel versions for Bayesian games and graphical games. In a standard Harsányi ε-equilibrium, a deviation from equilibrium by a player may grant him at most a gain of ε, when payoffs are 'averaged' over all of his types, but it is possibile that *some* individual types will benefit by more than ε. In a strong Harsányi ε-equilibrium, a deviation from equilibrium by a player may grant him positive gain, but only in a subset of states Ω' of measure less than ε. In $\Omega \setminus \Omega'$, which is of measure $\geq 1 - \varepsilon$, a strict equilibrium holds. This is clearly a strengthening of the usual *ex ante* equilibrium concept.

Every hyperfinite graphical game of finite degree admits a strong Harsányi ε-equilibrium, and similarly every hyperfinite Bayesian game with finitely supported types posseses a strong Harsányi ε-equilibrium. (It remains an open question as to whether every hyperfinite graphical game, of countable degree, admits a strong Harsányi ε-equilibrium; and whether every hyperfinite Bayesian game, with purely atomic but not necessarily finitely supported types, possesses a strong Harsányi ε-equilibrium.)

We use this ultimately to obtain our main theorem: every hyperfinite Bayesian game with purely atomic types admits a (standard)

Harsányi ε-equilibrium. This achieves our aim of understanding when *ex ante* and interim behaviours diverge; this occurs in hyperfinite purely atomic Bayesian games, which may fail to have interim (i.e., Bayesian) ε-equilibrium but must of necessity admit *ex ante* (i.e., Harsányi) ε-equilibrium.

We end with a conjecture. A long series of research works have revealed deep connections between the countable Borel equivalence relation hierarchy and game theory: smoothness corresponds to the existence of Bayesian equilibria, hyperfiniteness (without smoothness) to a lack of Bayesian equilibria but Harsányi approximate equilibria. There are rungs above hyperfiniteness (see Section 6.1). If it can be shown that the the treeable level of the countable Borel equivalence relation hierarchy corresponds to the categorical non-existence of Harsányi approximate equilibria,[3] the picture will be complete. We conjecture that this is true. Given the close relationship between the distinction between hyperfiniteness and treeability and the distinction between amenable and non-amenable group actions, such a result would likely open new research horizons.

2. Mathematical Preliminaries

2.1. **Countable Borel Equivalence Relations.** A Polish space is a separable and completely metrisable space. The metrisability implies a topology, from which a Borel σ-algebra is derived. Measurability without further qualification in this paper, in the context of a Polish space Ω, will be understood to mean measurability with respect to the Borel σ-algebra of Ω.

When we work with the space of probability measures over Ω, denoted $\Delta(\Omega)$, we will suppose that the σ-algebra on $\Delta(\Omega)$ is generated by the standard weak* topology of $\Delta(\Omega)$, which is the weakest topology such that for each continuous bounded $f : \Omega \to \mathbb{R}$ the mapping $\mu \mapsto \int_\Omega f d\mu$ is continuous.

A relation $\mathcal{E}$ on a Polish space Ω is said to be *Borel* if the set $\{(x, y) \in \Omega \times \Omega \mid x \, \mathcal{E} \, y\}$ is a Borel subset of $\Omega \times \Omega$. An equivalence relation is said to be *countable* if each equivalence class, referred to as a *class* or an *atom*, is countable. We will abbreviate *countable Borel equivalence relation* as *CBER*. Similarly, an equivalence relation is *finite* if each of

[3] [Simon and Tomkowicz, 2018] presents an example of a Bayesian game without Harsányi approximate equilibria, but the game there does not decompose into countable common knowledge components and hence is not directly comparable.

its equivalence classes is finite. A Borel function is a function whose graph is a Borel set.[4]

2.2. **Smoothness.**

Definition 2.1. *A Borel equivalence relation $\mathcal{E}$ on a Polish space Ω is smooth if there is a Polish space Z and a Borel function $\psi : \Omega \to Z$ such that for all $x, y \in \Omega$*

$$x \,\mathcal{E}\, y \Longleftrightarrow \psi(x) = \psi(y), \tag{2.1}$$

(i.e., the classes of $\mathcal{E}$ are precisely the level sets of ψ.) ♦

If $\mathcal{E}$ is a common knowledge relation, a Borel function ψ witnessing the smoothness of the relation can be thought of as an auxiliary tool that enables us to ascertain when x and y are in the same common knowledge component; that occurs if and only if $\psi(x) = \psi(y)$.

Example 2.2. On $\Omega = \mathbb{R}^N$, the relation $x \sim_{\mathcal{E}} y$ if and only if $x - y \in \mathbb{Z}^N$ is a smooth relation. To see why, note that if one defines $\psi : \mathbb{R}^N \to [0,1)^N$ by

$$\psi(x_1, \ldots, x_N) = (x_1 - \lfloor x_1 \rfloor, \ldots, x_N - \lfloor x_N \rfloor)$$

where $\lfloor a \rfloor = \max\{k \in \mathbb{Z} \mid k \geq a\}$ denotes the integer part of a, then $\psi(x) = \psi(y)$ if and only if $x \sim_{\mathcal{E}} y$. ♦

Example 2.3. On the unit circle $S^1 = \{z \in \mathbb{C} \mid |z| = 1\}$, let $T : S^1 \to S^1$ be an irrational rotation given by $T(z) = z \cdot e^{2\pi i \alpha}$ for a fixed irrational $\alpha \in \mathbb{R}$. It is well-known that if $\psi : S^1 \to Z$ is Borel for some Borel space Z, and $\psi = \psi \circ T$, then ψ is constant. For this reason, the equivalence relation on S^1 defined by $x \sim_{\mathcal{E}} y$ if and only if x and y are in the same T-orbit[5] is not smooth, as no Borel function can distinguish the classes of this equivalence relation by assigning different values to elements of disparate classes. ♦

A *transversal* of $\mathcal{E}$ is a set $T \subseteq X$ which intersects each $\mathcal{E}$ equivalence class at exactly one point.

Example 2.4. In Example 2.2, the set $[0,1)^N$ is a Borel transversal. ♦

It is easy to see that if a Borel equivalence relation $\mathcal{E}$ has a Borel transversal then it is smooth: the mapping $\psi(x)$ which picks out the

---

[4] Continuous functions are Borel, but not every Borel function is continuous.

[5] The T-orbit of x is the set $\{T^n(x) \mid n \in \mathbb{Z}\}$.

sole element of T that is $\mathcal{E}$-equivalent to x witnesses the smoothness of $\mathcal{E}$. For CBER's, the converse is true as well: smoothness implies the existence of a transversal (see Proposition 2.5). Hence every smooth CBER admits measurable selection of an element within each equivalence class.

From this one can show that if every equivalence class of $\mathcal{E}$ is finite then $\mathcal{E}$ is smooth.[6] However, for countable Borel equivalence relations, which are the focus of much of the material of this paper, matters are not as simple.

A set $B \subseteq \Omega$ is said to be *saturated* with respect to an equivalence relation $\mathcal{E}$ if it is the union of $\mathcal{E}$-equivalence classes. The collection of all the Borel $\mathcal{E}$-saturated sets of a Borel equivalence relation $\mathcal{E}$ forms a σ-algebra, denoted $\sigma(\mathcal{E})$.

Given a Polish state space Ω and a CBER $\mathcal{E}$, we let $\Omega/\mathcal{E}$ denote the quotient space whose elements are the equivalence classes by $\mathcal{E}$, and the induced σ-algebra consists of precisely the images of the $\mathcal{E}$-saturated Borel sets in Ω under the quotient map $\iota : \Omega \to \Omega/\mathcal{E}$; this is the finest σ-algebra on $\Omega/\mathcal{E}$ such that ι is measurable.

The following proposition follows from Propositions 6.3 and 6.4 of [Kechris and Miller, 2004] and the discussion preceding them.

Proposition 2.5. *The following conditions are equivalent for a CBER $\mathcal{E}$ on a Polish space Ω:*

- *(a) $\mathcal{E}$ is smooth.*
- *(b) There is a Borel transversal for $\mathcal{E}$.*
- *(c) There is a Borel set intersecting each class of $\mathcal{E}$ in a finite non-empty set.*
- *(d) The quotient space $\Omega/\mathcal{E}$ is standard Borel.*[7]

2.3. **Hyperfiniteness.** A CBER $\mathcal{E}$ on a Polish space Ω is *hyperfinite* if it is generated by a Borel $\mathbb{Z}$-action; that is, $\mathcal{E}$ is hyperfinite if there is a Borel bijection $T : \Omega \to \Omega$ satisfying the property that for any $x, y \in \Omega$, the relation $x\mathcal{E}y$ holds if and only if there exists $n \in \mathbb{Z}$ such that

[6] Consider the set

$$T = \{x \in X \mid \forall y \in X,\, x \,\mathcal{E}\, y \Longrightarrow x \leq y\},$$

i.e., the set of the $\leq$-elements of the $\mathcal{E}$ equivalence classes, for any Borel linear ordering on the domain of $\mathcal{E}$; such an ordering exists by a theorem of Kuratowski (see Section 15.B of [Kechris, 1995]). This T is seen to be Borel and a transversal of $\mathcal{E}$, hence finiteness of the $\mathcal{E}$-classes is sufficient for smoothness (see, e.g., Example 6.1 of [Kechris and Miller, 2004]).

[7] That is, there is a Borel-measurable bijection between it and a Polish space.

$T^n(x) = y$. Alternatively, a CBER is hyperfinite if and only if there exists an increasing sequence of Borel finite equivalence relations on X, $\mathcal{E}_1 \subseteq \mathcal{E}_2 \subseteq \cdots$, such that $\mathcal{E} = \cup_n \mathcal{E}_n$. In a sense, a hyperfinite equivalence relation is the limit of finite equivalence relations; see [Dougherty et al., 1994]. Every smooth CBER is hyperfinite, but the converse does not hold.

Example 2.6. Going back to Example 2.3, the orbit equivalence relation induced by irrational angled rotation on the circle is not smooth (as explained above in that example), but it is hyperfinite by definition: its equivalence classes are the orbits of an irrational rotation operator.

2.4. **Graphing of CBERs.** A *Borel graph* $G \subset \Omega \times \Omega$ on a Polish space Ω is an irreflexive and symmetric Borel relation on Ω (that is, $(x, x) \notin G$ for all $x \in \Omega$ and $(x, y) \in G$ implies $(y, x) \in G$). G is a *graphing* of a CBER $\mathcal{E}$ if $\mathcal{E}$ is the transitive closure of a Borel graph G, meaning that $(x, y) \in \mathcal{E}$ if and only if there exist $z_1, \ldots, z_n \in \Omega$ such that $(x, z_1) \in G, \ldots, (z_i, z_{i+1}) \in G, \ldots (z_n, y) \in G$. When this holds we will say that G *generates* $\mathcal{E}$.

Example 2.7. In Example 2.2, if

$$G = \{(u, v) \in \mathbb{R}^2 \times \mathbb{R}^2 \mid u = v \pm e_1 \text{ or } u = v \pm e_2\}$$

where e_1, e_2 are the standard unit vectors $(1, 0), (0, 1)$, respectively, then G is a graphing for $\mathcal{E}$.

2.5. **Proper Regular Conditional Distributions.** Game theorists are used to working with priors and posteriors on partitioned spaces. The appropriate generalisation to the context of the structures in this paper makes use of the concept of proper regular conditional distributions.

If $(\Omega, \mathcal{B})$ is a measurable space, $\mu \in \Delta(\Omega)$, and $\mathcal{F}$ is a sub-σ-algebra of $\mathcal{B}$, then (see [Blackwell and Ryll-Nardzewski, 1963]) a *proper regular conditional distribution* (henceforth, proper RCD) of μ given $\mathcal{F}$ is a mapping $t : \Omega \times \mathcal{B} \to [0, 1]$ such that for each $B \in \mathcal{B}$, $\omega \mapsto t_\omega(B)$ is $\mathcal{F}$-measurable and Borel, and such that

$$\mu(B) = \int_\Omega t_\omega(B)\, d\mu(\omega), \text{ for all } B \in \mathcal{B}, \tag{2.2}$$

and

$$t_\omega(A) = 1, \text{ for } \mu\text{-a.e. } \omega \in A \in \mathcal{F}.$$

It can be shown that Equation (2.2) implies that for every $T \in \mathcal{B}$

$$t_\omega(T) = E_\mu[1_T \mid \mathcal{F}](\omega), \ \mu\text{-a.e. } \omega \in \Omega.$$

In terms that may be more familiar for game theorists, a proper RCD t of a probability measure μ may be thought of as the posterior t of a prior μ with respect to a knowledge structure $\mathcal{F} = \sigma(\mathcal{E})$ (recalling that $\sigma(\mathcal{E})$ is the σ-algebra defined by the collection of all the Borel $\mathcal{E}$-saturated sets).

2.6. **Purely Atomic Knowledge Spaces, Type Spaces, and Priors.** Let I be a non-empty, finite set of players and Ω a Polish space of *states*. With each player $i \in I$ we associate a Borel equivalence relation over Ω denoted $\mathcal{E}^i$, called i's *knowledge relation*.[8] Adopting the convention that $\mathscr{E}$ stands for a profile of knowledge relations $(\mathcal{E}^i)_{i \in I}$, a *knowledge space* is then a triple $(\Omega, I, \mathscr{E})$.

Given a knowledge space $(\Omega, I, \mathscr{E})$, the equivalence relation *induced* by $\mathscr{E}$, which will be denoted by $\mathcal{E}$ and called the *common knowledge relation*, is the transitive closure of the union $\cup_{i \in I} \mathcal{E}^i$; in other words, it is the minimal equivalence relation containing each element in $\mathscr{E}$. Intuitively, the elements of $\sigma(\mathcal{E}^i)$ form the set of Borel events that Player i can identify and know, in the sense that if $\omega \in A \in \sigma(\mathcal{E}^i)$ then player i 'knows' at ω that event A occurs; the elements of $\mathscr{E}$ then represent events that all the players can know, e.g., $\omega \in B \in \sigma(\mathcal{E})$ then every player knows at at ω that event B occurs.

Definition 2.8. *A knowledge space such that for all $i \in I$ each equivalence class of $\mathcal{E}^i$ is finite or countably infinite will be called a* purely atomic *knowledge space.*[9] ♦

[8] Intuitively, the unions of classes of $\mathcal{E}^i$ represent the events that Player i can identify; hence, $\sigma(\mathcal{E}^i)$ is the set of Borel events that Player i can identify. Most game theory models of incomplete information with types either explicitly or implicitly work with partitionally generated type spaces. In such models, where Ω is finite or countable, each player i has a partition Π^i of Ω. This approach suffers from a difficulty in the case of a continuum of states, since the partition has to 'agree' with the measurable structure. In addition, in the continuum case one cannot work with arbitrary unions of partitions elements; only Borel unions are admissible. Our approach differs from the more classical approach given in [Nielsen, 1984] and [Brandenburger and Dekel, 1987] in favour of defining knowledge via relations (instead of σ-algebras), which is better suited for the class of purely atomic types that will concern us. Our approach also differs from the 'types' approach of [Milgrom and Weber, 1985]. See the discussion in [Hellman and Levy, 2017] for a comparison.

[9] If the equivalence relation thus defined by the equivalence classes is Borel, then a purely atomic knowledge space satisfies the conditions for being a CBER.

Purely atomic knowledge spaces were studied, as mentioned, in multiple works, e.g., [Simon, 2000], [Simon, 2003], [Lehrer and Samet, 2011], [Hellman, 2014], [Hellman and Levy, 2017], [Hellman and Levy, 2019].

Fix a knowledge space $(\Omega, I, \mathscr{E})$. For each $i \in I$, a *type function* t^i is a mapping $t^i : \Omega \to \Delta(\Omega)$ that is $\sigma(\mathcal{E}^i)$-measurable[10] and satisfies $t^i_\omega(A) = 1$ whenever $\omega \in A \in \sigma(\mathcal{E}^i)$.

Note that t^i is a type function, not a type. To define types in this model, for a given state $\omega \in \Omega$ we will sometimes write t^i_ω to denote the element of $Image(t^i) \subseteq \Delta(\Omega)$ defined by $t^i(\omega)$. Intuitively, $t^i_\omega(S)$ is the probability player i associates to a set S at state ω. Each t^i_ω is then a *type* of player i. We may also sometimes denote a generic type of player i by τ^i.

We will say that a type $t^i_\omega \in \Delta(\Omega)$ is *finitely supported* if the support of t^i_ω, as a probability measure over Ω, is concentrated on a finite subset of Ω. If the support is countably infinite, we will say that the type is countably supported. Since we concentrate in this paper on purely atomic knowledge spaces (see Definition 2.8), all types will be finitely or countable supported, unless otherwise noted.

Adopting the convention that t stands for a tuple of type functions $(t^i)_{i \in I}$, a triple (Ω, I, t) is called a *type space*. A type space implicitly defines the knowledge relations $\mathcal{E}^i$ underlying the type functions: $\omega \mathcal{E}^i \omega'$ (i.e., $(\omega, \omega') \in \mathcal{E}^i$) if and only if $t^i_\omega = t^i_{\omega'}$.

Definition 2.9. *A type space such that for all $i \in I$ and all $\omega \in \Omega$ the type t^i_ω is purely atomic will be called a* purely atomic *type space. A purely atomic type space is* smooth *(respectively* hyperfinite*) if the state space Ω along with the common knowledge equivalence relation $\mathcal{E}$ is smooth (respectively hyperfinite).* ⧫

Any type space on a purely atomic knowledge space is also purely atomic. We will henceforth always assume that knowledge spaces (and hence type spaces) are purely atomic.

A measure $\mu^i \in \Delta(\Omega)$ such that t^i is a proper RCD for μ^i given $\sigma(\mathcal{E}^i)$ is a *prior*[11] for t^i. A prior μ^i for a player i naturally induces a distribution $t^i_*(\mu^i)$ on $\Delta(\Omega)$, concentrated on $Image(t^i)$ – the space of

[10] Meaning that for any Borel $A \subseteq \Omega$, the mapping $\omega \to t^i_\omega(A)$ is $\sigma(\mathcal{E}^i)$-measurable.

[11] That a type t^i is a proper RCD for a prior may add to the intuitive explanation of the characteristics of a proper RCD. The characteristic that $\omega \mapsto t^i_\omega(B)$ is Borel translates into a statement that the type $t^i_\omega(\cdot)$ defines a probability measure over $(\Omega, \mathcal{B})$ for each ω, the condition $t^i_\omega(A) = 1$ for a.a. $\omega \in A$ states that the event A is always 'known' at state ω if it includes ω, and $\mu^i(B) = \int_\Omega t^i_\omega(B)\, d\mu^i(\omega)$ is a

player i's types – where t^i_* is the push-forward to the quotient space, i.e., $t^i_*(\mu) = \mu \circ (t^i)^{-1}$. A *common prior* is a measure μ that is a prior for the type functions of all the players $i \in I$.

3. Bayesian and Graphical Games

3.1. Bayesian Games and Bayesian Equilibria.

A *Bayesian game* $B := (\Omega, I, t, A, r)$ consists of the following components:

- (Ω, I, t) forms a type space (with knowledge relations $\mathcal{E}^i$ understood implicitly[12] as generated by the types t).
- $A = (A^i)_{i \in I}$ is a tuple consisting of a finite action set A^i for each player $i \in I$.
- $r : \Omega \times \prod_{i \in I} A^i \to \mathbb{R}^I$ is a bounded measurable payoff function, with r^i then being the resulting payoff to player i. The payoff function r extends multi-linearly to mixed actions in the usual manner.

Definition 3.1. *A Bayesian game is* purely atomic *(respectively,* smooth, *respectively* hyperfinite*) if the underlying type space is purely atomic (respectively, smooth, respectively hyperfinite).* ♦

Examples of Bayesian games that are hyperfinite but not smooth appear prominently in [Simon, 2003] and [Hellman, 2014], and are discussed extensively in [Hellman and Levy, 2017].

From here, we will presume a purely atomic Bayesian game (even though all of the definitions extend to more general Bayesian games). A *strategy* of a player $i \in I$ is a mapping $\sigma^i : \Omega \to \Delta(A^i)$ which is constant on each player's knowledge component, i.e., which is $\sigma(\mathcal{E}^i)$-measurable (recall that $\mathcal{E}^i$ is player i's knowledge relation, and $\sigma(\mathcal{E}^i)$ is the σ-algebra of Borel $\mathcal{E}^i$-saturated sets). Intuitively, the definition of a strategy must satisfy this criterion because a player can only choose a strategy based on the knowledge available.

form of 'Bayesian consistency'. In game theoretic models, this construction fits the standard interpretations in which player i has a prior measure over all the state of the world, and after receiving information, updates to a posterior belief.

[12] A type for player i was defined as a mapping from Ω to $\Delta(\Omega)$ that is $\sigma(\mathcal{E}^i)$-measurable, hence the type function depends on the knowledge relation. However, given a type space one can reconstruct the knowledge relations of the players: $\mathcal{E}^i = \sigma(t^i)$ is the coarsest σ-algebra with respect to which t^i is a Borel mapping to $\Delta(\Omega)$. This is what we mean when we say that the knowedlge relations are understood implicitly.

Given a strategy profile $\sigma = (\sigma^i)_{i\in I}$, a player i, and a type τ^i of player i, the *expected payoff* of player i is given by

$$r^i(\sigma \mid \tau^i) = \sum_{\{\omega \mid t^i(\omega)=\tau^i\}} r^i(\omega, \sigma(\omega)) \cdot \tau^i(\omega). \tag{3.1}$$

Similarly, if x is an action of Player i,

$$r^i(x, \sigma^{-i} \mid \tau^i) = \sum_{\{\omega \mid t^i(\omega)=\tau^i\}} r^i(\omega, x, \sigma^{-i}(\omega)) \cdot \tau^i(\omega). \tag{3.2}$$

A *Bayesian equilibrium* is a profile of strategies $\sigma = (\sigma^i)_{i\in I}$ such that for each $i \in I$, every type τ^i, and each alternative[13] mixed action $x \in \Delta(A^i)$ of player i,

$$r^i(\sigma \mid \tau^i) \geq r^i(x, \sigma^{-i} \mid \tau^i) \tag{3.3}$$

holds. A *Bayesian ε-equilibrium*, for $\varepsilon > 0$, is a profile of strategies $\sigma = (\sigma^i)_{i\in I}$ such that for each $i \in I$, every type τ^i, and each alternative mixed action $x \in \Delta(A^i)$ of Player i,

$$r^i(\sigma \mid \tau^i) + \varepsilon \geq r^i(x, \sigma^{-i} \mid \tau^i) \tag{3.4}$$

holds. If, in addition, each σ^i is Borel measurable, the strategy profile σ is said to be a *measurable*[14] Bayesian (ε-)equilibrium. (Hence, σ^i is $\sigma(\mathcal{E}^i)$-measurable.)

The following theorem is proved in both [Hellman and Levy, 2017] and [Hellman and Levy, 2019].

Theorem 3.2. *A smooth purely atomic Bayesian game admits a measurable Bayesian equilibrium.*

3.2. **Graphical Games of Countable Degree.** Let (Ω, G) be a graph (recall this means that the edge relation $G \subseteq \Omega \times \Omega$ is an irreflexive and symmetric relation). For finite graphs, [Kearns et al., 2001] defines a *graphical game* to be a strategic form game such that Ω is the set of players and the payoff of each player depends only on his actions and those of his neighbours under the edge relation G. Formally, suppose the (finite) action set of each agent is A. Letting $N(g)$ denote the set of the neighbours of $g \in \Omega$ under the edge relation G,

[13] Keep in mind that the strategy of a type of player i is fixed for all states ω in each of that player's knowledge components, hence the alternative mixed action does not need to be presented as a function of ω.

[14] It is possible for a game to have Bayesian ε-equilibria that are not measurable as in, for example, [Simon, 2003] or [Hellman, 2014].

the payoff function of a player g is a function $u_g : A \times A^{N(g)} \to \mathbb{R}$. Payoffs extend multi-linearly in the usual way.[15]

The generalisation to infinite Ω adds the requirement that the payoff function $u_g : A \times A^{N(g)} \to \mathbb{R}$ is continuous in the Tychonoff topology. If for all $g \in \Omega$ the set $N(g)$ is countable (respectively, finite) we will call (Ω, G) a graphical game of countable (respectively, finite) degree.

The existence of Nash equilibria in such games was proven in [Peleg, 1969]; an example showing the necessity of continuity with respect to the Tychonoff topology is also presented in that paper.[16] However, even if Ω has a Borel structure and the payoffs obey some natural measurability requirements, the resulting equilibria need not be measurable; see [Levy, 2013, Sec. 8].

This framework generalises straight-forwardly to the case in which Ω is a Polish space, $G \subseteq \Omega \times \Omega$ is Borel, the degree of G is countable (i.e., for all $g \in \Omega$, $N(g)$ is countable), and the payoffs are Borel. Specifically, the payoff function here is a function $u : \Omega \times A^{\Omega} \to \mathbb{R}$, understood as $u(g, a)$ being the payoff of $g \in \Omega$ when the profile a is played, such that $u(g, \cdot)$ depends only on the actions of g and his neighbours. (Formally, if $g \in \Omega$ and $a, a' : \Omega \to A$ are two action profiles such that $a(g) = a'(g)$, and $a(h) = a'(h)$ for each $h \in N(g)$, then $u(g, a) = u(g, a')$.) A *strategy profile* is then a Borel mapping $\sigma : \Omega \to \Delta(A)$.

Given these definitions, we will assume that:

- u_g is continuous in the Tychonoff topology on $A \times A^{N(g)}$,
- u is bounded,
- whenever $\alpha : \Omega \times \Omega \to A$ is Borel, $g \to u(g, \alpha(g, \cdot))$ is Borel.

To understand the latter condition, observe that for each fixed $g \in \Omega$, $\alpha(g, \cdot)$ defines a strategy profile, hence the Borel requirement.

Definition 3.3. *For a graph (Ω, G) over a Polish space of countable degree with a Borel edge relation $\mathcal{E}$, the equivalence relation induced by the edges G (i.e., the equivalence relation whose classes are the connected components of G) straightforwardly satisfies the conditions for being a CBER. We will call such a game a* purely atomic graphical game.

[15] For the sake of clarity, we will usually denote the payoff function of a graphical game by u, as opposed to r which will be used to denote the payoff function of a Bayesian game.

[16] Continuity implies that $N(g)$ must be countable for each $g \in G$ (cf. [Peleg, 1969]).

Definition 3.4. *Let (Ω, G) be a purely atomic graph whose Borel edge relation $\mathcal{E}$ induces a smooth (resp. hyperfinite) CBER. We will call such a graph a* smooth graph *(resp.* hyperfinite graph*). A graphical game over such a graph will be called a* smooth (resp. hyperfinite) graphical game*.*

A (measurable) Nash equilibrium of a graphical game is naturally a (measurable) mapping $\sigma : G \to \Delta(A)$ such that $u(g, \sigma) \geq u(g, (b, (\sigma(h))_{h \neq g}))$ for each $g \in G$ and each $b \in A$, i.e., replacing the mixed action chosen by player g by any pure action b (while the mixed actions of the other players remain fixed) does not increase g's payoff.

The following theorem is proved in [Hellman and Levy, 2019].

Theorem 3.5. *A smooth graphical game admits a measurable Nash equilibrium.*

4. Main Theorem

Our main theorem appears in Theorem 4.13: every hyperfinite Bayesian game with purely atomic types admits an Harsányi ε-equilibrium. Recall, as discussed in the introduction and defined formally below, Harsányi ε-equilibrium implies that no deviating strategy for a player would improve payoff by more than ε when aggregated over all states w.r.t. that player's prior. To prove our main theorem, we make use of graphical games, essentially converting a hyperfinite Bayesian game into a graphical game, proving statements about the associated graphical game, and then using that to draw conclusions about the Bayesian game. This is presented in greater detail in the following sketch, which also makes use of our notion of *strong* Harsányi ε-equilibrium, in which no deviating player can improve his payoff *at all* at all but an ε-measure set of states.

Sketch of the main argument: Given a hyperfinite Bayesian game with purely atomic types B, we first form an associated induced graphical game Γ (using Proposition 4.5) that encodes in its vertices and edge relations many of the properties of the Bayesian game B; most importantly, Γ is also hyperfinite.

Next we consider a graphical game of finite degree Γ' that is ε-close to Γ (existence guaranteed by Lemma 4.8). By Proposition 4.10, Γ', being of finite degree, admits a strong Harsányi ε- equilibrium.

We then 'roll' this back: the existence of a strong Harsányi ε-equilibrium of Γ' implies the existence of an Harsányi ε-equilibrium of Γ (Proposition 4.11). This in turn yields the conclusion of the existence of an

Harsányi ε-equilibrium of the associated Bayesian game B (Lemma 4.6), completing the argument.

4.1. **Harsányi Equilibrium and Strong Harsányi ε-Equilibrium.** Bayesian equilibrium is an interim concept; it asks which strategies players will choose after they have been informed of their types. There is a corresponding *ex ante* equilibrium concept, also known as Harsányi equilibrium (cf. [Simon, 2003]), in which each player simultaneously chooses strategies for all of his types.

To proceed, assume each player i has a prior μ^i. We define notions of Harsányi equilibrium, and then contrast an Harsányi equilibrium with a Bayesian equilibrium. Recall that the notation of expected payoff to each type given in Equation (3.2).

Definition 4.1. *Let $(\Omega, (t^i), (\mu^i))$ with actions $A_1, \ldots, A_I$ and payoffs $r : \Omega \times \prod A_i \to \mathbb{R}^I$ be a Bayesian game, and let $\varepsilon > 0$. A strategy profile $\sigma = (\sigma^i)_{i \in I}$, with $\sigma^i : \Omega \to \Delta(A_i)$ measurable with respect to Player i's knowledge, is an* Harsányi ε-equilibrium*, if for each $i \in I$ and each alternative mixed action $a \in \Delta(A^i)$ of player i,*

$$\int_{\omega \in \Omega} \Big[\max_{a \in A^i} r^i(a, \sigma^{-i} \mid \tau^i) - r^i(\sigma \mid \tau^i) \Big] d\mu^i(\tau^i) \leq \varepsilon. \tag{4.1}$$

Equivalently, a strategy profile $\sigma = (\sigma^i)_{i \in I}$ is an Harsányi ε-equilibrium*, if for each $i \in I$ and each strategy $\tilde{\sigma}^i$ of Player i,*

$$r^i(\tilde{\sigma}^i, \sigma^{-i}) \leq r^i(\sigma) + \varepsilon. \tag{4.2}$$

♦

The equivalence of the two definitions in Equations (4.1) and (4.2) above follows by application of a standard measurable selection theorem to select for each type τ^i an action maximising $r^i(\cdot, \sigma^{-i} \mid \tau^i)$. Equation (4.2) may seem more intuitive, but as our techniques make use of reductions to graphical games, Equation (4.1) is more practical for our purposes here.

Recall that a Bayesian ε-equilibrium, for $\varepsilon > 0$, is a profile of strategies $\sigma = (\sigma^i)_{i \in I}$ such that for each $i \in I$, every type τ^i, and each alternative mixed action $x \in \Delta(A^i)$ of Player i,

$$r^i(\sigma \mid \tau^i) + \varepsilon \geq r^i(x, \sigma^{-i} \mid \tau^i) \tag{4.3}$$

holds. If $\varepsilon = 0$ in the inequalilty in Equation (4.3) then a Bayesian equilibrium holds.

The difference between a Bayesian ε-equilibrium and an Harsányi ε-equilibrium can be seen in contrasting Equation (4.3) with Equation (4.2). In the Bayesian case, the test of whether or not a deviation to an alternative strategy satisfies Equation (4.3) needs to be applied separately for each type τ^i of each player i, as expressed in the conditional payoff expression $r^i(\sigma \mid \tau^i)$; if the defining inequality fails even for one type τ^i, the condition fails and a Bayesian equilibrium has not been defined.

In contrast, the test of whether a strategy satisfies the conditions of being an Harsányi equilibrium involves the entire state space Ω, without conditioning on individual types; note that Equation (4.1) involves an integral over Ω. A rough analogy would be to consider the Bayesian equilibrium as requiring optimal best-reply strategies for all the players with respect to the payoff at each individual type, with the Harsányi equilibrium requiring best-reply strategies with respect to the average payoff of each player. In other words, in an Harsányi equilibrium a player may theoretically choose to accept a sub-optimal payoff at one of his types in exchange for a high payoff at another type in such a way that the average payoff over all types is optimal.

A measurable Bayesian ε-equilibrium satisfies the conditions of being also an Harsányi ε-equilibrium. For $\varepsilon = 0$, the converse also holds: by a theorem in [Simon, 2003], any measurable Harsányi equilibrium generates a measurable Bayesian equilibrium (up to a null set).[17] For $\varepsilon > 0$, however, this may not be true; the main example in [Hellman, 2014] admits a measurable Harsányi ε-equilibrium but no measurable Bayesian ε-equilibrium.

The concepts of Harsányi equilibrium and Harsányi ε-equilibrium have been well studied as far back as Harsányi's seminal papers on equilibria in games of incomplete information. [Harsanyi, 1967] showed that in the finite state case the *ex ante* Harsányi equilibrium and the interim Bayesian equilibrium concepts are equivalent.

Definition 4.2 below introduces a stronger *ex ante* equilibrium concept than the standard one in the literature as presented in Definition 4.1.

Definition 4.2. *Let* $(\Omega, (t^i), (\mu^i))$ *with actions* $A_1, \ldots, A_I$ *and payoffs* $u : \Omega \times \prod A_i \to \mathbb{R}^I$ *be a Bayesian game, and let* $\varepsilon > 0$*. A strategy profile* $\sigma = (\sigma^i)_{i \in I}$*, such that for each* i*,* $\sigma^i : \Omega \to \Delta(A_i)$ *is measurable with*

[17] That is, in an Harsányi equilibrium, each player may fail to optimise against the others' strategies on a set of states which is null with respect to his own prior.

respect to player i's knowledge, is a strong Harsányi ε-equilibrium *if for each player i there is a subset $\Omega' \subseteq \Omega$ with $\mu^i(\Omega') < \varepsilon$ such that for all $\omega \notin \Omega'$, $i \in I$, for $\tau^i = t^i(\omega)$,*

$$r^i(\sigma \mid \tau^i) = \max_{a \in A^i} r^i(a, \sigma^{-i} \mid \tau^i).$$

♦

It is important for the results in this paper to note carefully the distinctions between the definitions of Harsányi ε-equilibrium (Definition 4.1) and our new concept of strong Harsányi ε-equilibrium (Definition 4.2). In words, a standard Harsányi ε-equilibrium requires that a deviation from equilibrium by player i grant him at most a gain of ε when his payoffs are averaged over the payoffs of all of his types; this, however, still leaves open the possibility that some types will benefit by more than ε (as long as this is balanced by other types getting less than average in such a way that the overall average does not exceed ε).

In a strong Harsányi ε-equilibrium, a deviation by player i may grant him positive gain, but only in a subset of states Ω' such that $\mu^i(\Omega') < \varepsilon$. At all states other than those in Ω', player i has no incentive at all to deviate because any deviation from equilibrium can at most give him the same payoff and perhaps worsen his payoff. In other words, a strong Harsányi ε-equilibrium requires each player i to be precisely optimising up to a μ^i-measure ε set – a very strict requirement in comparison with standard Harsányi ε-equilibrium.

Graphical games are not amenable to *ex ante* versus interim period analysis, in contrast to Bayesian games (because in graphical games there is no *ex ante* stage before signals are revealed; the players in graphical games choose their strategies without the benefit of information provided by a signal indicating their types). However we may, in analogy with Definitions 4.1 and 4.2, define what we will (for terminological consistency) term Harsányi ε-equilibrium and strong Harsányi ε-equilibrium for graphical games.

Definition 4.3. *Let $\Gamma := (\Omega, G, u)$, along with action set A, be a graphical game of countable degree. Let μ be a measure over Ω.*

A strategy profile $\sigma : \Omega \to \Delta(A)$ is an Harsányi ε-equilibrium *of Γ if*

$$\int_\Omega \big[\max_{a \in A} u(g, (a, \sigma_{-g})) - u(g, \sigma)\big] d\mu(g) \leq \varepsilon. \tag{4.4}$$

A strategy profile σ is a strong Harsányi ε-equilibrium *of Γ if there is a subset $\Omega' \subseteq \Omega$ with $\mu(\Omega') \leq \varepsilon$ such that for all $g \notin \Omega'$,*

$$\max_{a \in A} u(g,(a,\sigma_{-g})) = u(g,\sigma).$$

In words, a strategy profile of a graphical game is an Harsányi ε-equilibrium if deviating from the equilibrium strategy yields the players an average gain of at most ε when we integrate over the payoffs of all the players. It is a strong Harsányi ε-equilibrium if it is a strict Nash equilibrium ε μ-almost everywhere; in other words, we permit a measure ε of players to gain by deviating, but $1-\varepsilon$ of the players have no incentive at all to deviate.

The concept of strong Harsányi ε-equilibrium of a graphical game with bounded payoffs is stronger than that of the standard Harsányi ε-equilibrium. This can be seen directly: if M is a bound on the payoffs, then every strong Harsányi ε-equilibrium is an Harsányi ε-equilibrium.

Lemma 4.4. *A strong Harsányi ε-equilibrium of a graphical game (respectively a Bayesian game) with the absolute value of payoffs bounded by $M > 0$ is an Harsányi $2M \cdot \varepsilon$-equilibrium of the game.*

Proof. We prove the version for graphical games; the version of the proof for Bayesian games can be copied from this proof almost verbatim. Let $\Gamma := (\Omega, G, u)$ be a graphical game, let μ be a measure over Ω, let $\varepsilon > 0$, and let σ be strong Harsányi ε-equilibrium of Γ. Let $\Omega' \subset \Omega$ be a set with $\mu(\Omega') < \varepsilon$ such that all the agents in $\Omega \setminus \Omega'$ are optimising by using σ. By definition,

$$\int_{\Omega\setminus\Omega'} \big[\max_{a\in A} u(g,(a,\sigma_{-g})) - u(g,\sigma)\big]\, d\mu(g) = 0.$$

Hence,

$$\begin{aligned}
&\int_{\Omega} \big[\max_{a\in A} u(g,(a,\sigma_{-g})) - u(g,\sigma)\big]\, d\mu(g)\\
&\qquad = \int_{\Omega'} \big[\max_{a\in A} u(g,(a,\sigma_{-g})) - u(g,\sigma)\big]\, d\mu(g)\\
&\qquad \leq 0 + \int_{\Omega'} 2M\, d\mu(g) = 2M\mu(\Omega') < 2M\varepsilon.
\end{aligned}$$

□

4.2. **Graphical Games from Bayesian Games.** We will prove results pertaining to Bayesian games by converting them to graphical games, where we will perform most of the analysis, and then pivot back to apply those results to Bayesian games. For this we need a way to construct a graphical game associated with any given Bayesian game.

Let $\mathcal{B} := (\Omega, (t^i)_{i\in I}, (\mu^i)_{i\in I}, A, (r^i)_{i\in I})$ be a Bayesian game with common action space A for each player (for simplicity, assume all the players share the same action space) and payoffs $r^i : \Omega \times \prod A_i \to \mathbb{R}^I$ for each player. Then the *agent-normal form graphical game associated with B*, (based on a concept first presented, in relation to extensive-form games, in [Selten, 1975]), is the graphical game $\Gamma_{\mathcal{B}} := (\Omega_T, G_T)$ with payoff function u_T and measure μ_T over Ω_T defined as follows:

- For each i, set $T^i := \Omega/t^i \sim Image(t^i)$ as a quotient space. T^i is the type space of player i, i.e., the collection of all of that player's types.
- The set of players of $\Gamma_{\mathcal{B}}$ is the disjoint union $\Omega_T := \sqcup_{i=1}^I T^i$. In words, each type of each player of $\mathcal{B}$ becomes a player of $\Gamma_{\mathcal{B}}$.
- For each player $\tau_i \in T^i \in \Omega_T$ and strategy profile $\sigma : \Omega_T \to \Delta(A)$, a payoff u_T is defined by the expected payoff:

  $$u_T(\tau_i, \sigma) = r^i(\sigma \mid \tau_i).$$

  It is immediate to verify that such u_T is measurable in the sense we require.
- The graph G_T on the vertices $\Omega_T = \sqcup_{i=1}^I T^i$ is such that there is an edge between vertices τ^i and τ^j if and only if there is a state $\omega \in \Omega$ in the original Bayesian game B such that $t^i(\omega) = \tau^i$ and $t^j(\omega) = \tau^j$; i.e., simultaneously, τ^i, τ^j can be the types of players i,j respectively.
- What remains is defining the measure μ_T induced on Ω_T. When $\mu_1 = \cdots = \mu_N = \mu$ is a common prior, the natural candidate is $\mu_T := \frac{1}{|I|}\sum_{i=1}^I t_*^i(\mu)$, where t_*^i is the push-forward to the quotient space, i.e., $t_*^i(\mu) = \mu \circ (t^i)^{-1}$. More generally, define

  $$\mu_T := \frac{1}{|I|^2}\sum_{i,j=1}^I t_*^i(\mu^j), \tag{4.5}$$

  which reduces to the previous definition in the case of a common prior (the square term $|I|^2$ is needed because the sum in Equation (4.5) over $i, j \in I$ counts players twice).

Proposition 4.5. *If a Bayesian game is smooth (resp. hyperfinite) then the induced graphical game is smooth (resp. hyperfinite).*

Proof. Recall that Ω_T is the disjoint union $\sqcup_{i=1}^I T^i$. Let $\mathcal{E}_T$ be the equivalence relation in the induced graphical game of which G_T is a graphing. We observe that $(\omega, \theta) \in \mathcal{E}$ if and only if there are $\omega_0 = \omega, \omega_1, \ldots, \omega_N = \theta$ and players $i_1, \ldots, i_n$ such that $t^{i_k}(\omega_{k-1}) = \tau^{i_k}(\omega_k)$ for each $k = 1, \ldots, N$; this holds if and only if there are types $\tau_0, \ldots, \tau_K$ such that $(\tau_{k-1}, \tau_k) \in G_T$ with $\tau_0 \in \{t^i(\omega) \mid i \in I\}$ and $\tau_N \in \{t^i(\theta) \mid i \in I\}$, which in turn holds if and only if $(\tau_0, \tau_N) \in \mathcal{E}_T$ for some $\tau_0 \in \{t^i(\omega) \mid i \in I\}$ and $\tau_N \in \{t^i(\theta) \mid i \in I\}$. Hence, we see that $(\omega, \theta) \in \mathcal{E}$ if and only if $(t^i(\omega), t^j(\theta)) \in \mathcal{E}_T$ for each i, j.

The relation $\mathcal{E}'$ on $\Omega' = \Omega \times I$ is defined by $(\omega, i)\, \mathcal{E}'\, (\theta, j)$ if and only if $\omega\, \mathcal{E}\, \theta$. By the above inference, the relation $\mathcal{E}_T$ is the relation induced by $\mathcal{E}'$ and the quotient map $\iota(\omega, i) = t^i(\omega) \in T^i \subseteq \Omega_T$; that is, $\iota(\omega, i)\, \mathcal{E}_T\, \iota(\theta, j)$ if and only if $(\omega, i)\, \mathcal{E}'\, (\theta, j)$ (if and only if $\omega\, \mathcal{E}\, \theta$). In words, this is saying that $\iota(\omega, i)$ and $\iota(\theta, j)$ are connected in the induced graphical game (that is, $\iota(\omega, i)\, \mathcal{E}_T\, \iota(\theta, j)$) if and only if ω and θ are in the same common knowledge equivalence class in the Bayesian game (that is, $\omega\, \mathcal{E}\, \theta$), as we had established above.

If the Bayesian game is smooth (respectively hyperfinite), i.e., the relation $\mathcal{E}$ on Ω is smooth (respectively hyperfinite), then so is $\mathcal{E}'$. As quotient maps preserve smoothness (respectively hyperfiniteness), this completes the proof. (Since we have not found a reference in the literature to this last assertion, we state it as a separate lemma, Lemma A.1, in an appendix. Intuitively, this is because quotient maps can only simplify the structure, never make it more complex, in a topological sense.) □

Strategies in the graphical game induced by a Bayesian game naturally induce strategies in the original Bayesian game. The relation between the equilibrium concepts for graphical games and Bayesian games is given by:

Lemma 4.6. *A (strong) Harsanyi ε-equilibrium of the associated graphical game of a Bayesian game is a (strong) Harsanyi $|I|^2 \cdot \varepsilon$-equilibrium of that Bayesian game.*

A converse can be proved but we do not need it.

Proof. Suppose σ is a strong Harsanyi ε-equilibrium of the associated graphical game; then by definition there is a set $\Omega'_T \subseteq \Omega$ with

$\mu_T(\Omega'_T) < \varepsilon$ such that for all $i \in I$, each $\tau^i \notin T_i \cap \Omega'_T$ satisfies

$$u_T(\tau^i, \sigma) = \max_{a \in A} u_T(a, \sigma_{-\tau_i}), \text{ i.e. }, r^i(\sigma \mid \tau_i) = \max_{a \in A} r^i(a, \sigma_{-i} \mid \tau_i).$$

Hence, if $\Omega' = \bigcup_i (t^i)^{-1}(\Omega'_T)$, then for all $\omega \notin \Omega'$ and for each player i,

$$r^i(\sigma \mid t^i(\omega)) = \max_{a \in A} r^i(a, \sigma_{-i} \mid t^i(\omega)).$$

Furthermore, for each player i, recalling the definition of μ_T in Equation (4.5)

$$\mu^i(\Omega') < \sum_{j=1}^{|I|} \mu^i((t^j)^{-1}(\Omega'_T)) \leq |I|^2 \cdot \mu_T(\Omega'_T) < |I|^2 \cdot \varepsilon.$$

Suppose σ is an Harsányi ε-equilibrium of the associated graphical game. Then by definition,

$$\int_{\Omega_T} \big[\max_{a \in A} u(g, (a, \sigma_{-g})) - u(g, \sigma) \big] d\mu_T(g) < \varepsilon,$$

and in particular, since $T^i \subseteq \Omega_T$ and $t^i_*(\mu^i) \leq |I|^2 \cdot \mu_T$,

$$\int_{T^i} \big[\max_{a \in A} u(g, (a, \sigma_{-g})) - u(g, \sigma) \big] d(t^i_*(\mu^i))(g) < |I|^2 \cdot \varepsilon,$$

which is re-written by the definition of the payoffs in the associated graphical game as

$$\int_{\Omega} \big[\max_{a \in A} r^i(a, \sigma^{-i} \mid \tau^i) - r^i(\sigma \mid \tau^i) \big]] d\mu^i(\tau^i) < |I|^2 \cdot \varepsilon. \tag{4.6}$$

Comparing Equation (4.6) with Equation (4.1) gives the desired conclusion. □

4.3. **Hyperfiniteness and Equilibria.** Lemma 4.8 shows that every graphical game in which vertices have infinite neighbours is 'near' a graphical game with vertices that all have finite neighbours.

Definition 4.7. *Let $\Gamma = (\Omega, G, u)$ and $\Gamma' = (\Omega, G', u')$ be two graphical games. Let $\varepsilon > 0$. Then Γ' is ε-*close *to Γ if the difference of the payoffs, $|u(\omega) - u'(\omega)| \leq \varepsilon$ holds at all states ω except on a μ-null set.*

Lemma 4.8. *Let $\Gamma = (\Omega, G, u)$ be a graphical game. Let μ be a measure over Ω and let $\varepsilon > 0$. Then there is a graphical game $\Gamma' = (\Omega, G', u')$, with $G' \subseteq G$, such Γ' is ε-close to Γ, and furthermore each node of G' has finite degree.*

Proof. We may suppose that every vertex of G has infinitely many neighbours (if some vertices of G have finitely many neighbours we can simply leave those alone and apply the argument here only to the remaining vertices).

For each integer j, let $\phi_j : \Omega \to \Omega$ be a measurable function, such that the collection $\{\phi_j\}_{j\in\mathbb{N}}$ satisfies the property that for each $\omega \in \Omega$ the set $\bigcup_{j\in\mathbb{N}}\{\phi_j(\omega)\}$ equals $\{\omega\} \cup \{\theta \mid (\omega, \theta) \in G\}$ (i.e., every neighbour of ω, and ω itself, eventually appears in $\bigcup_{j\in\mathbb{N}}\{\phi_j(\omega)\}$), and furthermore $\phi_j(\omega) \neq \phi_k(\omega)$ if $j \neq k$. Such a collection $\{\phi_j\}_{j\in\mathbb{N}}$ can be shown to exist by standard arguments by the Lusin–Novikov theorem, e.g., [Kechris, 1995, Thm. 18.10].

From the perspective of player ω, the collection $\{\phi_j\}_{j\in\mathbb{N}}$ can serve as a way to enumerate the set of neighbours of ω and itself in the graph, and furthermore enumerate the actions that they choose under any strategy profile. In more detail, let σ be a strategy profile. If $\phi_j(\omega) = \theta$, denote the action that player θ plays under σ by a_j, i.e., $a_j := \sigma(\theta)$. This yields a list $(a_1, a_2, \ldots)$ of actions of ω and all of his immediate neighbours. Since the utility ω derives from the implementation of strategy profile σ, namely $u(\omega, \sigma)$, is a function solely of $(a_1, a_2, \ldots)$, we may denote:

$$v(\omega, a_1, a_2, \ldots) := u(\omega, \sigma), \text{ where } \sigma(\theta) = a_j \text{ if } \phi_j(\omega) = \theta.$$

We claim that the function $v : \Omega \times A^{\mathbb{N}} \to \mathbb{R}$ is Borel measurable. It suffices to show that for fixed ω, $v(\omega, \cdot) : A^{\mathbb{N}} \to \mathbb{R}$ is continuous (this is true by assumption) and that for fixed $a = (a_1, a_2, \ldots)$, $v(\cdot, a) : \Omega \to \mathbb{R}$ is Borel.[18] Indeed, fix such $a = (a_1, a_2, \ldots)$, and define the mapping $\alpha : \Omega \times \Omega \to A$ given by

$$\alpha(\omega, \theta) = a_j, \text{ if } \phi_j(\omega) = \theta,$$

i.e., $\alpha = a_j$ on $Gr(\phi_j)$. The mapping α is clearly Borel, and $v(\omega, a) = u(\omega, \alpha(\omega, \cdot))$ is Borel by assumption.

Next, for each $N \in \mathbb{N}$, define mappings $\overline{\psi} : \Omega \times A^N \to \mathbb{R}$ and $\underline{\psi} : \Omega \times A^N \to \mathbb{R}$ by

$$\overline{\psi}_N(\omega, a_1, \ldots, a_N) = \max\{v(\omega, a_1, a_2, \ldots, a_N, a_{N+1}, \ldots) \mid a_{N+1}, a_{N+2}, \ldots \in A\},$$

$$\underline{\psi}_N(\omega, a_1, \ldots, a_N) = \min\{v(\omega, a_1, a_2, \ldots, a_N, a_{N+1}, \ldots) \mid a_{N+1}, a_{N+2}, \ldots \in A\}.$$

[18] Functions defined on a product of Polish spaces, continuous in one variable and Borel measurable in the other, are called Caratheodory functions, and are known to be jointly Borel measurable (see [Himmelberg, 1975, Thm. 6.1]).

Next, define a mapping $N : \Omega \to \mathbb{N}$ by

$$N(\omega) := \min\{N \mid \forall a \in A^N,\ \overline{\psi}_N(a,\omega) - \underline{\psi}_N(a,\omega) \leq \varepsilon\}.$$

By continuity and compactness, such a mapping $N(\omega)$ is guaranteed to exist, and is Borel. Set[19]

$$u'(\omega,\sigma) := \overline{\psi}_{N(\omega)}(\omega, \sigma(\phi_1(\omega)), \ldots, \sigma(\phi_N(\omega))).$$

Finally, define a graph

$$G' := \{(\omega,\theta) \mid \theta \neq \omega, \text{ and } \exists j = 1, \ldots, N(\omega) \text{ such that } \theta = \phi_j(\omega)\}.$$

This satisfies all that we set out to attain. By construction, the graph $G' \subseteq G$ has the property that that each vertex in G' has finite degree, each player's payoff in u' depends only on his or her own action and those of his or her neighbours in G', and $|u - u'| \leq \varepsilon$. □

Proposition 4.9. *Let (Ω, G) be a hyperfinite graph of finite degree. Then there exists an increasing sequence $\Omega_1 \subseteq \Omega_2 \subseteq \Omega_3 \subseteq \cdots \subseteq \Omega$ satisfying $\Omega = \cup_n \Omega_n$ such that for each integer $n \in \mathbb{N}$, the equivalence classes of the relation on Ω_n generated by $G \cap (\Omega_n \times \Omega_n)$ are finite.*

Proof. By one of the definitions of hyperfiniteness, there exists an increasing sequence of equivalence relations on Ω with finite classes, $\mathcal{E}_1 \subseteq \mathcal{E}_2 \subseteq \cdots$, such that $\mathcal{E} = \cup_n \mathcal{E}_n$. Define $\Omega_n = \{x \in \Omega \mid y\, G\, x \to y\, \mathcal{E}_n\, x\}$, i.e., all those points in Ω which are in the same $\mathcal{E}_n$ equivalence class as all their neighbours. Then for all n, $\Omega_n \subseteq \Omega_{n+1}$. Since G is of finite degree, for each $\omega \in \Omega$ for sufficiently large n we are guaranteed that $\omega \in \Omega_n$. Let $S_n = \Omega \setminus \Omega_n$.

To complete the proof, we contend that the equivalence relation on Ω_n generated by the graph $G \cap (\Omega_n \times \Omega_n)$ coarsens the restriction of $\mathcal{E}_n$ to Ω_n, i.e., is contained in $\mathcal{E}_n \cap (\Omega_n \times \Omega_n)$. Indeed, suppose (x, y) is in the transitive closure of $G \times (\Omega_n \times \Omega_n)$, hence there are $z_1, \ldots, z_n \in \Omega_n$ with $(x, z_1) \in G, \ldots, (z_i, z_{i+1}) \in G, \ldots (z_n, y) \in G$. But since $x, z_1, \ldots, z_n, y \in \Omega_n$, that implies that $(x, z_1), (z_1, z_2), \ldots, (z_{n-1}, z_n), (z_n, y) \in \mathcal{E}_n$. □

Proposition 4.10. *Every hyperfinite graphical game of finite degree admits a strong Harsányi ε-equilibrium.*

Proof. Let (Ω, G, u) be a hyperfinite graphical game of finite degree. Let μ be a measure on Ω, and let $\varepsilon > 0$. Let $M > 0$ be a bound on the absolute value of the payoffs, i.e. $|u| \leq M$, and let $\delta = \frac{\varepsilon}{2M}$.

[19] Alternatively, we could have used $\underline{\psi}$, or anything between $\underline{\psi}$ and $\overline{\psi}$.

By Proposition 4.9 there exists a sequence of subsets $\Omega_1 \subseteq \Omega_2 \subseteq \Omega_3 \subseteq \cdots \subseteq \Omega$ satisfying $\Omega = \cup_n \Omega_n$ such that for each n the equivalence relation induced by $G_n := G \cap (\Omega_n \times \Omega_n)$ is finite, and hence smooth (since every finite equivalence relation is smooth).

Fix N such that $\mu(\Omega_N) \geq 1 - \delta$ and denote $\Omega' = \Omega_N$, $G' = G_N$, and $\mu' = \mu(\cdot \mid \Omega')$. We can now consider graphical games over (Ω', G').

Let τ be an arbitrary Borel map $\tau : \Omega \setminus \Omega' \to \Delta(A)$, which we will use to extend strategy profiles from $\Omega \setminus \Omega'$ to all of Ω. More specifically, recall that a strategy profile of a graphical game over Ω is a Borel map from Ω to $\Delta(A)$. Then, for each strategy profile $\sigma : \Omega' \to \Delta(A)$ over Ω', define a strategy profile $\hat{\sigma}$ in the original game by setting $\hat{\sigma}(\omega) = \sigma(\omega)$ for $\omega \in \Omega'$ and $\hat{\sigma}(\omega) = \tau(\omega)$ for $\omega \in \Omega \setminus \Omega'$. Along with this, define a payoff function u' on (Ω', G') by

$$u'(\omega, \sigma) = u(\omega, \hat{\sigma}).$$

Hence we have a well-defined graphical game (Ω', G', u') with measure μ'. This game is finite, hence smooth. Appealing to [Hellman and Levy, 2019], by smoothness (Ω', G', u') admits a 0-Harsányi-equilibrium σ. But then $\hat{\sigma}$ is a strong Harsányi $2 \cdot M \cdot \delta = \varepsilon$-equilibrium over our original game (Ω, G, u), by a brief calculation very similar to the one in the proof of Lemma 4.4. □

Proposition 4.11. *Every hyperfinite graphical game with bounded payoffs admits an Harsányi ε-equilibrium.*

Proof. Let Γ be a hyperfinite graphical game and fix $\varepsilon > 0$. Let M be a bound on the absolute values of the payoffs in the game, i.e, $|u| \leq M$.

Applying Lemma 4.8, approximate Γ by a finite-degree graphical game Γ' which is ε-close to Γ. By Proposition 4.10, Γ' admits a strong Harsányi ε-equilibrium σ, which is an Harsányi $2M\varepsilon$-equilibrium σ of Γ' by Lemma 4.4. Since $|u - u'| < \varepsilon$ by the definition of closeness, σ is a $(2M + 1)\varepsilon$-equilibrium of Γ. □

Theorem 4.12. *Every hyperfinite Bayesian game with finitely supported types possesses a strong Harsányi ε-equilibrium.*

Proof. Let B be a hyperfinite Bayesian game with finitely supported types. Assume first that all players have the same action space. The graphical game associated with B satisfies the property that every vertex has finite degree; hence the conclusion we seek follows from Proposition 4.10 and Lemma 4.6.

More generally, if the players have different action spaces, one can replace the profile of action sets $(A^i)_{i \in I}$ with a superset of all of them, e.g. $A = \cup A^i$, which can serve as a common action set for all the players. We can then inductively define payoffs so that for each player i, any actions outside of A^i in the original game are strictly dominated in the resulting game, which is guaranteed to have a strong Harsányi ε-equilibrium $\sigma' = (\sigma'_i)_{i \in I}$. We may heuristically change σ'_i for i to σ_i in the original game in a way that all the probaility weight not put on A^i in σ'_i is shifted to A^i in σ_i. The measure of types which must be shifted in such a way is 'small'; we omit the formal development of this proof idea. □

We finally have all the ingredients needed for the proof of our main result in Theorem 4.13.

Theorem 4.13. *Every hyperfinite Bayesian game with purely atomic types admits an Harsányi ε-equilibrium.*

Proof. Let B be a hyperfinite Bayesian game with purely atomic types. As in the proof of Theorem 4.12, we may without loss of generality assume a common action space for all the players. By constructing the graphical game associated with B, as above, we may apply Proposition 4.11 (existence of Harsányi ε-equilibrium in graphical games) and Lemma 4.6 (which transfers ε-equilibria from graphical games to Bayesian games) to yield Theorem 4.13. □

A summary of the lemmata and propositions leading to the main theorems of this paper, and their interrelationships, appears in graphical form in Figure 2.

Remark 4.14. As remarked in the introduction, it is an open question whether Bayesian games with purely atomic (but not necessarily finite) types possess strong Harsányi ε-equilibria, and similarly it is an open question whether graphical games of countable (but not necessarily finite) degree possess strong Harsányi ε-equilibria.

5. Extensions

The results of this paper, both with respect to Bayesian games and graphical games, extend readily to the case of players with continuous and compact action spaces. Indeed, the proofs in this paper carry through verbatim in that case. One must only be careful when relying on the results of [Hellman and Levy, 2019] (as we do in the proof of Proposition 4.10), which establishes the existence of (0)-equilibrium in smooth hyperfinite games. The results of that paper do use finite action spaces, but the results there also extend easily.

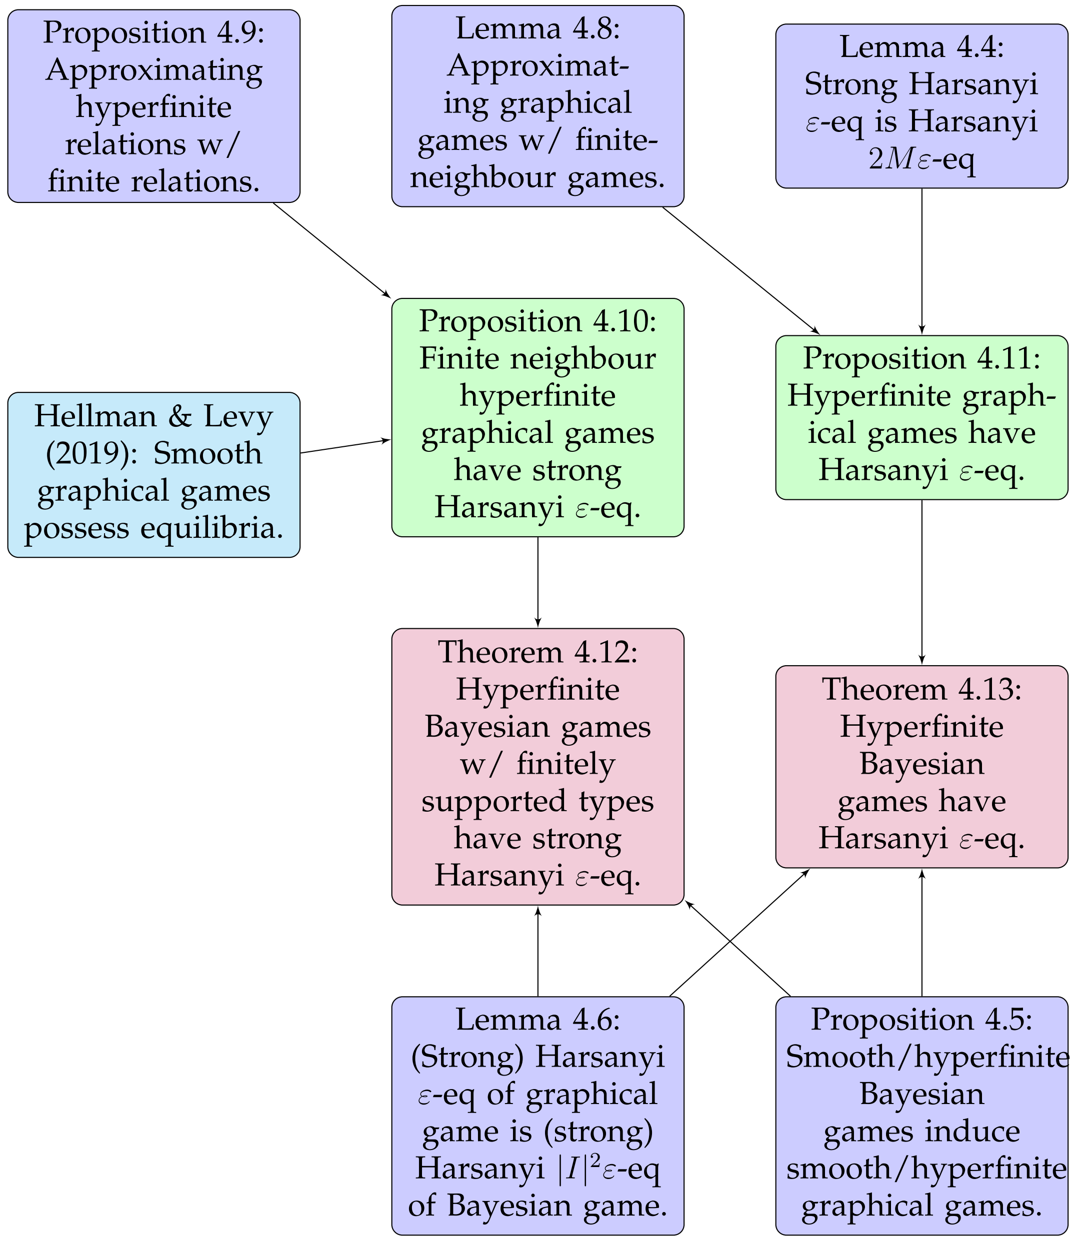

Figure 1. Logical Flow of Results

One also need not assume that payoffs are bounded to prove Theorem 4.12 and Theorem 4.13; it is sufficient to assume that the payoffs are integrably bounded, that is,

$$\int_\Omega \sup_{a \in \prod_i A^i} |r^i(\omega, a)| d\mu(\omega) < \infty.$$

Similarly, Proposition 4.10 and Corollary 4.11 extend to integrably bounded payoffs,

$$\int_{\Omega} \sup |u(g, \cdot)| d\mu(g) < \infty.$$

with the supremum taken over strategy profiles.

6. EQUILIBRIA AND THE COUNTABLE BOREL EQUIVALENCE RELATION HIERARCHY

6.1. **The Countable Borel Equivalence Relation Hierarchy.** Smoothness, as presented in Definition 2.1, is but one element in a fundamental hierarchy of countable Borel equivalence relations that has been studied intensely in recent years (see, for example, [Thomas and Schneider, 2012]).

The basic building block of the hierarchy is the Borel reducibility order. If A, B are standard Borel spaces with countable Borel equivalences $\mathcal{E}$ and $\mathcal{F}$ respectively, then $\mathcal{E}$ is *Borel reducible* to $\mathcal{F}$, denoted $\mathcal{E} \leq_B \mathcal{F}$, if there is a Borel map $\phi : A \to B$ such that $x\,\mathcal{E}\,y \Longleftrightarrow \phi(x)\,\mathcal{F}\,\phi(y)$ for $x \in A, y \in B$. In these terms, writing $id\,\mathbb{R}$ to denote the identity equivalence class of the set $\mathbb{R}$, a CBER $\mathcal{E}$ is smooth if $\mathcal{E} \leq_B id\,\mathbb{R}$.

But smoothness constitutes only the lowest rung of the hierarchy of countable Borel equivalence relations. It represents the 'simplest' equivalence relations. For the next rung, we define the CBER $\mathcal{E}_0$ over $2^{\mathbb{N}}$ by

$$x\,\mathcal{E}_0\,y \Longleftrightarrow \exists m \in \mathbb{N},\ \ \forall n > m,\ \ x(n) = y(n),$$

where $x(n)$ stands for the n-th element of x considered as a sequence in $2^{\mathbb{N}}$.

The CBER $\mathcal{E}_0$ is called the relation of eventual agreement on $2^{\mathbb{N}}$, and it is a fundamental building block of the CBER hierarchy. Hyperfiniteness[20] is the rung above smoothness in the hierarchy: as shown in Example 2.6, every smooth CBER is hyperfinite. The converse, however, does not hold: $\mathcal{E}_0$ itself is the canonical example of a CBER that is hyperfinite but not smooth. Any CBER $\mathcal{E}$ is hyperfinite if and only if $\mathcal{E} \leq_B \mathcal{E}_0$, and any two non-smooth hyperfinite equivalence relations are reducible to each other [Dougherty et al., 1994].

A CBER $\mathcal{E}$ is *treeable* if there is a Borel acyclic graph whose connected components are $\mathcal{E}$-equivalence classes. Treeable CBERs are above hyperfinite relations in the hierarchy, in the sense that every hyperfinite CBER is treeable but the converse does not hold, thus yielding a nice, clean nesting containment sequence, smooth

[20] As defined above in Section 2.3.

Common Knowledge Equivalence Relation	*Measurable Equilibria*
Treeable	*Conjecture:* Harsányi ε-equilibria may not exist
Hyperfinite	Harsányi ε-equilibria exist Bayesian equilibria may not exist
Smooth	Bayesian and Harsányi equilibria exist

Figure 2. Relationship between the CBER hierarchy, as it relates to the underlying common knowlege structure of the state space of a Bayesian game, and the existence of measurable equilibria.

within hyperfinite within treeable (see Figure 2). Above the hyperfinite level, however, the neat linearity of the hierarchy breaks down; among the treeable relations alone, there are $2^{\aleph_0}$ many treeable CBERs that are incomparable in the Borel reducibility order [Hjorth, 2012]. There is also a universal CBER, denoted $\mathcal{E}_\infty$, which satisfies $\mathcal{E} \leq_B \mathcal{E}_\infty$ for any CBER $\mathcal{E}$.

6.2. **Climbing the Hierarchy.** The results of the previous section both advance our understanding of conditions guaranteeing the existence of Harsányi equilibria and sharpen the distinctions between *ex ante* and interim solution concepts in Bayesian games. They also deepen the connections recently discovered between those solution concepts and the countable Borel equivalence relation hierarchy, while hinting at further intriguing possibilities.

As noted above in Section 6.1, the lowest rung of the countable Borel equivalence relation hierarchy is that of smoothness, which includes as a special case the finite spaces. The classic of papers of John Harsányi established that all finite Bayesian games admit both Bayesian and Harsányi equilibria. [Hellman and Levy, 2017] extended this to the class of smooth purely atomic Bayesian games: as in the finite case, they all admit both Bayesian and Harsányi equilibria.

Counter-examples to the existence of Bayesian equilibria in non-finite games appeared in [Simon, 2003] and [Hellman, 2014], with the latter not even admitting Bayesian ε-equilibria for sufficiently small ε. Both examples are non-smooth; by the results of [Hellman and Levy, 2017] this is no coincidence, as they had to be non-smooth in order to be counter-examples.

Both of those examples are, however, hyperfinite games, and they also positively admit Harsányi ε-equilibria for all ε. This, too, is not a coincidence; that is precisely what is established in this paper. By Theorem 4.13, the class of hyperfinite but not smooth Bayesian games is exactly the class of games that do not admit Bayesian equilibria but do have Harsányi ε-equilibria.

This parallelism between the smooth and hyperfinite rungs of the CBER hierarchy and results on equilibrium existence is summarised in Figure 2. What happens even higher up the hierarchy, at the treeable level or even beyond? The general theory of the upper reaches of the countable Borel hierarchy is itself not fully understood at present. [Simon and Tomkowicz, 2018] have presented a three-player example of a Bayesian game that admits no Harsányi ε-equilibria. However, that example does not satisfy the definition of a purely atomic Bayesian game (in particular, the only common knowledge component is the entire space) and hence cannot shed light on the relationship between the CBER hierarchy and equilibria of purely atomic games developed here.

It is of significant interest to see if the example of [Simon and Tomkowicz, 2018] can be adopted to the purely atomic framework. Currently, we do not know if there exists a purely atomic game without Harsányi equilibrium.

If such a game can be shown to exist, then the next question would be: given a treeable (but not hyperfinite) knowledge structure, is it always possible to construct a Bayesian game over that structure that does not admit an Harsányi equilibrium? We conjecture that this is the case.

[Hellman and Levy, 2017] includes a theorem that states a form of converse to the statement in Theorem 3.2, namely that over any non-smooth knowledge structure it is possible to construct a Bayesian game that does *not* admit any Bayesian equilibria. The above conjecture would be the exact parallel of that theorem, replacing non-smooth by non-hyperfinite and Bayesian equilibrium by Harsányi equilibrium. Proving (or disproving) this conjecture will require more careful study of the inter-relationships between purely atomic Bayesian games and the countable Borel equivalence relation hierarchy.

APPENDIX A. SMOOTHNESS & HYPERFINITENESS UNDER QUOTIENTS

Lemma A.1. *If $\mathcal{E}$ is a smooth (resp. hyperfinite) CBER on a Polish space Ω, and $\iota : \Omega \to \Theta$ is a surjective map such that $\iota(\omega_1) = \iota(\omega_2)$ implies $(\omega_1, \omega_2) \in \mathcal{E}$ for any $\omega_1, \omega_2 \in \Omega$, then the induced relation $\mathcal{D} = \{(\iota(\omega_1), \iota(\omega_2)) \mid (\omega_1, \omega_2) \in \mathcal{E}\}$ on Θ is a CBER and is smooth (resp. hyperfinite).*

Before proving the lemma, recall that a reduction from $(\Omega, \mathcal{E})$ to $(\tilde{\Omega}, \tilde{\mathcal{E}})$, denoted $\mathcal{E} \leq_B \tilde{\mathcal{E}}$, is a Borel mapping $\phi : \Omega \to \tilde{\Omega}$ with $(\omega_1, \omega_2) \in \mathcal{E} \Leftrightarrow (\phi(\omega_1), \phi(\omega_2)) \in \tilde{\mathcal{E}}$. In fact, the lemma remains true if smoothness or hyperfiniteness is replaced with any property $\mathcal{P}$ of a CBER which can be expressed in an equivalent way via reduction to a CBER $\tilde{\mathcal{E}}$ on a space $\tilde{\Omega}$. Smoothness is equivalent to the existence of a reduction to $\mathbb{R}$ with the identity CBER, denoted id, while hyperfiniteness is equivalent to the existence of a reduction to $2^{\mathbb{N}}$ with the CBER $\mathcal{E}_0$ induced by the shift. (See Section 6.1 below, or [Dougherty et al., 1994].)

Proof. That $\mathcal{D}$ is a CBER is immediate to verify. Suppose $\mathcal{E}$ is smooth (resp. hyperfinite); then there is a reduction ϕ from $(\Omega, \mathcal{E})$ to $(\mathbb{R}, id)$ (resp. $(2^{\mathbb{N}}, \mathcal{E}_0)$). By the Lusin–Novikov theorem ([Kechris, 1995, Theorem 18.10]) there is a reduction ψ from $(\Theta, \mathcal{D})$ to $(\Omega, \mathcal{E})$ such that $\iota \circ \psi = id$. Hence, $\phi \circ \psi$ is a reduction from $(\Theta, \mathcal{D})$ to $(\mathbb{R}, id)$ (resp. $(2^{\mathbb{N}}, \mathcal{E}_0)$), showing that the former is smooth (resp. hyperfinite). □

REFERENCES

[Blackwell and Ryll-Nardzewski, 1963] Blackwell, D. and Ryll-Nardzewski, C. (1963). Non-existence of everywhere proper conditional distributions. *Ann. Math. Statist.*, 34(1):223–225.

[Brandenburger and Dekel, 1987] Brandenburger, A. and Dekel, E. (1987). Common knowledge with probability 1. *Journal of Mathematical Economics*, 16(3):237 – 245.

[Dougherty et al., 1994] Dougherty, R., Jackson, S., and Kechris, A. S. (1994). The structure of hyperfinite Borel equivalence relations. *Trans. Amer. Math. Soc.*, 341(1):193–225.

[Harsanyi, 1967] Harsanyi, J. C. (1967). Games with Incomplete Information Played by Bayesian Players, I–III Part I. The Basic Model. *Management Science*, 14(3):159–182.

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

[Hellman and Levy, 2017] Hellman, Z. and Levy, Y. J. (2017). Bayesian games with a continuum of states. *Theoretical Economics*, 12:1089 – 1120.

[Hellman and Levy, 2019] Hellman, Z. and Levy, Y. J. (2019). Measurable selection for purely atomic games. *Econometrica*, 87(2):593–629.

[Himmelberg, 1975] Himmelberg, C. (1975). Measurable relations. *Fundamenta Mathematicae*, 87(1):53–72.

[Hjorth, 2012] Hjorth, G. (2012). Treeable equivalence relations. *Journal of Mathematical Logic*, 12(1):1250003.

[Kearns et al., 2001] Kearns, M. J., Littman, M. L., and Singh, S. P. (2001). Graphical models for game theory. pages 253–260.

[Kechris, 1995] Kechris, A. (1995). *Classical Descriptive Set Theory*. Graduate Texts in Mathematics. Springer Verlag, New York.

[Kechris and Miller, 2004] Kechris, A. and Miller, B. (2004). *Topics in Orbit Equivalence*. Lecture Notes in Mathematics. Springer Verlag, Berlin and Heidelberg.

[Lehrer and Samet, 2011] Lehrer, E. and Samet, D. (2011). Agreeing to agree. *Theoretical Economics*, 6(2):269–287.

[Levy, 2013] Levy, Y. (2013). A Cantor set of games with no shift-homogeneous equilibrium selection. *Mathematics of Operations Research*, 38(3):492–503.

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

[Nielsen, 1984] Nielsen, L. T. (1984). Common knowledge, communication, and convergence of beliefs. *Mathematical Social Sciences*, 8(1):1 – 14.

[Peleg, 1969] Peleg, B. (1969). Equilibrium points for games with infinitely many players. *Journal of the London Mathematical Society*, s1-44(1):292–294.

[Selten, 1975] Selten, R. (1975). Reexamination of the perfectness concept for equilibrium points in extensive games. *International Journal of Game Theory*, 4(1):25–55.

[Simon, 2000] Simon, R. S. (2000). The common prior assumption in belief spaces: An example. Technical report, The FedermannCenter For The Study of Rationality, Hebrew University of Jerusalem.

[Simon, 2003] Simon, R. S. (2003). Games of incomplete information, ergodic theory, and the measurability of equilibria. *Israel Journal of Mathematics*, 138(1):73–92.

[Simon and Tomkowicz, 2018] Simon, R. S. and Tomkowicz, G. (2018). A Bayesian game without ε-equilibria. *Israel Journal of Mathematics*, 227(1):215 – 231.

[Thomas and Schneider, 2012] Thomas, S. and Schneider, S. (2012). Countable borel equivalence relations. In Cummings, J. and Schimmerling, E., editors, *Appalachian Set Theory, 2006-2012*, pages 25–62. Cambridge University Press, Cambridge.

[Zamir, 2008] Zamir, S. (2008). Bayesian Games: Games with Incomplete Information. Discussion Paper Series dp486, The Federmann Center for the Study of Rationality, the Hebrew University, Jerusalem.

DEPARTMENT OF ECONOMICS, BAR ILAN UNIVERSITY, RAMAT GAN, ISRAEL.

ADAM SMITH BUSINESS SCHOOL, UNIVERSITY OF GLASGOW, GLASGOW, UK