Strategic Type Spaces*

Olivier Gossner[†] Rafael Veiel[‡]

October 16, 2024

Abstract

We provide a strategic foundation for information: in any given game with incomplete information we define strategic quotients as information representations that are sufficient for players to compute best-responses to other players. We prove 1/ existence and essential uniqueness of a minimal strategic quotient called the Strategic Type Space (STS) in which a type is given by an interim correlated rationalizability hierarchy together with the set of beliefs over other players' types and nature that rationalize this hierarchy 2/ that this minimal STS is a quotient of the universal type space and 3/ that the minimal STS has a recursive structure that is captured by a finite automaton.

1 Introduction

For games of incomplete information, Harsanyi (1967) introduced type spaces as models to describe players' information on uncertain payoff-relevant parameters (i.e. states of nature), where each type is associated to a belief on

*The authors are grateful to Marcin Pęski, Stephen Morris, Satoru Takahashi, Muhamet Yildiz and several participants of the Transatlantic Theory Workshop, the Stony Brook conference in Game Theory, the Institut Henri Poincaré Game Theory seminar, the PSE TOM seminar and One World Mathematical Game Theory seminar for useful comments. Olivier Gossner acknowledges support from the French National Research Agency (ANR), "Investissements d'Avenir" (ANR-11-IDEX-0003/LabEx Ecodec/ANR-11-LABX-0047).

†CNRS - École Polytechnique, London School of Economics

‡Massachusetts Institute of Technology

states of nature and other players' types. Mertens and Zamir (1985) show that these type spaces can be represented in canonical models (universal type spaces) of players' hierarchies of beliefs, whose descriptions do not depend on the payoff structure of a game. Universal types should then contain payoff relevant information for all games and so become very large objects.

This paper takes back the question of how to describe players' information by taking a game as fixed. We provide a universal representation of players' payoff relevant information by switching the focus from a purely informational description of types to a strategic description. For a fixed game, we introduce strategic quotients as descriptions of players' strategic information, which allow for an economical representation of every Harsanyi type space.

In our approach, an economical representation of a Harsanyi type space is obtained by mapping its types and associated beliefs into equivalence classes. A strategic quotient is a canonical set of such equivalence classes, called strategic types, which satisfies the following two conditions for every player i:

1. For any Harsanyi type space, if two of player i's types have same beliefs on nature and strategic types of other players, then these types belong to the same strategic type of player i.

2. If other player's behaviors depend only on their strategic types, there exists a best-response of player i that depends only on i's strategic type.

The first is a sufficient condition for different Harsanyi types to be merged into the same strategic type. Unlike Harsanyi (1967), we do not require the converse of this condition. We thus allow for different beliefs over states and other players' strategic types to correspond to the same strategic type, so that strategic quotients partition Harsanyi type spaces. The second condition depends on a best-reply concept and the meaning of strategic behavior. We focus on the best-reply correspondence that underlies the solution concept of interim correlated rationalizability (ICR), as defined in Dekel, Fudenberg, and Morris (2007). This condition implies that strategic quotients cannot be too coarse, as they must be closed under best-replies.

An example of a strategic quotient for any game is the universal type space of Mertens and Zamir (1985), and we show that every strategic quotient is a quotient of the universal type space. Since we are interested in economical information representations, we want to focus on strategic quotient that are the smallest, or coarsest. We prove existence and uniqueness of a minimal strategic quotient which we call the strategic type space (STS). We show this by first proving that all finite order ICR actions arising from any Harsanyi type can be recovered from any Strategic Quotient. We then provide a canonical construction of the set of best-reply hierarchies and show that it forms a STS when each type is associated to the set of beliefs that rationalize it as a best-response. We show that these hierarchies characterize all finite order ICR actions and deduce that our construction characterizes a unique STS up to isomorphisms.

We then analyze properties of the STS. Through careful exploration of the recursive structure underlying our construction of the STS, we show that ICR hierarchies exhibit a finite recursive structure. More precisely, we show the existence of a finite automaton associated to the underlying game, called the strategic automaton, such that ICR hierarchies are 1-1 with the set of paths on the automaton. This characterization allows us to further deduce that the STS is a compact and Hausdorff space when endowed with the product topology on ICR hierarchies. As an application of our approach, in a companion paper Gossner and Veiel (2024) we use the recursive structure to characterize the set of implementable distributions under interim correlated rationalizability.

1.1 Related Literature

The best-reply concept we focus on in this paper was introduced to define Interim correlated rationalizability (ICR). Rationalizability was introduced by Bernheim (1984); Pearce (1984) in games with complete information. Dekel et al. (2007) generalized this concept by introducing the concept of ICR for games of incomplete information. For every type, ICR iteratively eliminates never best-replies to that type's expectation over any state contingent, correlated beliefs over other types' actions.

Dekel et al. (2007) show that two Harsanyi types have the same ICR actions in all games if and only if they correspond to the same hierarchy of beliefs and hence the same point in the universal type space of Mertens and Zamir (1985). Therefore, ICR has been studied as a correspondence on the universal type space of Mertens and Zamir (1985). Morris, Shin, and Yildiz (2016) characterizes ICR in terms of a common belief operator on the universal type space for global games. Weinstein and Yildiz (2007) first identified critical types, i.e. points of discontinuity of ICR in the universal type space of Mertens and Zamir (1985). They provide a topological characterization of critical types. Dekel, Fudenberg, and Morris (2006) and Chen, Di Tillio, Faingold, and Xiong (2016a) characterize the coarsest topology on the universal type space, called the strategic topology, under which ICR is continuous. Chen et al. (2016a) introduce the notion of *frames* as partitions of the universal type space similar to the first property of STS. They use frames as a tool to define a strategic topology of uniform convergence over games for hierarchies of beliefs. Chen, Takahashi, and Xiong (2014) study robustness of ICR to both higher order beliefs and payoff perturbations. The authors define *curb* collections which is closely related to the second requirement of STS, i.e. strategic closure (see Section 3.2), defined in terms of the universal type space. Chen, Takahashi, and Xiong (2016b) provide an algorithm to compute hierarchies of ICR which parallels our construction of best reply hierarchies. Based on their construction, they study refinements on ICR. Finally, Ely and Peski (2011) provide a characterization of critical types in terms of common belief properties in the universal type space.

Most importantly, this paper differs from the literature described above in the following way: We fix a game and introduce a canonical language to describe strategically relevant information for this game. Unlike frames and curb collections, STS are defined as universal objects which can be characterized and constructed without reference a particular Harsanyi type space.

2 Preliminaries and Notations

We denote the cardinality of a set Y by $|Y|$. For a family of sets $(X_i)_i$ we let $X \coloneqq \prod_i X_i$ and $X_{-i} \coloneqq \prod_{j \neq i} X_j$. For a family of mappings $f_i \colon X_i \to Y_i$, f is the map from X to Y given by $f(x) = (f_i(x_i))_i$ and f_{-i} is from X_{-i} to Y_{-i} is given by $f_{-i}(x_{-i}) = (f_j(x_j))_{j \neq i}$. Similarly, if $f \colon X \to Y$ and $g : Z \to W$ are mappings we denote by $(f \times g) : X \times Z \to Y \times W$ the map given by $(f \times g)(x, z) = (f(x), g(z))$. On any set X we denote by id_X the identity mapping on X and omit the subscript X when there is no ambiguity. The set of Borel probability measures over a topological space X is written as Δ_X. We denote by $\operatorname{supp} p$ the support of a probability measure p. The marginal probability on $Y \times X_m$ of a probability measure p on a product space $X = Y \times \prod_i X_i$ is denoted $\operatorname{marg}_{Y,m}(p)$.

In commutative diagrams we describe a mapping between probability measures from Δ_X to Δ_Y which are induced by a measurable mapping from X to Y by an arrow on the subscripts as follows:

$$\begin{array}{c} \Delta_X \\ \downarrow \\ \Delta_Y \end{array}$$

Double headed arrows such as $X \twoheadrightarrow Y$ denote surjective mappings. The subscript i denotes a typical player from the finite set N of players. A finite set K of states of nature and, finite action sets $(A_i)_{i \in N}$ and a payoff function $u : A \times K \to \mathbb{R}^N$, are given.

3 Strategic Type Spaces

In this section we introduce Strategic type spaces (STS). Section 3.1 introduced the best-reply correspondence underlying the concept of Interim correlated rationalizability. In section 3.2 we introduce STS and minimal STS axiomatically. In section 3.3 we establish existence and uniqueness of a minimal STS, characterized as the space of best-reply hierarchies. Section 4 shows that the minimal STS can be represented by a finite automaton.

3.1 Interim Correlated Best-Replies

Dekel et al. (2007) show that ICR can be defined as a fixed point of a best reply correspondence, which we now state.

We introduce the set $\mathcal{B}_i := 2^{A_i}$ of action subsets and define conjectures as maps $\sigma : K \times \mathcal{B}_{-i} \to \Delta(A_{-i})$. The probability $\langle \sigma, p \rangle$ over $K \times A_{-i}$ induced by a belief $p \in \Delta_{K \times \mathcal{B}_{-i}}$ and a conjecture σ is given by the relation:

$$\langle \sigma, p \rangle(k, a_{-i}) := \sum_{b_{-i} \in \mathcal{B}_{-i}} \sigma(k, b_{-i})(a_{-i}) \; p(k, b_{-i}).$$

Player i's best-reply correspondence $\mathrm{BR}_i : \Delta_{K \times \mathcal{B}_{-i}} \to \mathcal{B}_i$ is given by:

$$\mathrm{BR}_i(p) := \bigcup_{\sigma : \text{supp } \sigma(k, b_{-i}) \subseteq b_{-i}, \ \forall k, b_{-i}} \left\{ \arg\max_{a_i \in A_i} \sum_{k, a_{-i}} u_i(a_i, a_{-i}, k) \langle \sigma, p \rangle (k, a_{-i}) \right\}.$$

A *Harsanyi type space* $\mathscr{H}$ consists of a family of topological spaces $(\Theta_i)_{i\in N}$ and of continuous mappings $\pi_i : \Theta_i \to \Delta_{K\times\Theta_{-i}}$, where $\pi_i(\theta_i)$ represents type θ_i's belief over types of other players and states of nature. As in Dekel et al. (2007), we rely on the best-reply correspondence BR_i to define ICR on any Harsanyi type space $(\Theta_i, \pi_i)_i$ as follows: ICR of a type θ_i is given by $\mathrm{ICR}_i(\theta_i) = \bigcap_m \mathrm{ICR}_i^m(\theta_i)$, where $\mathrm{ICR}_i^0(\theta_i) = A_i$ and $\mathrm{ICR}_i^m(\theta_i)$ is i's best response to the $\pi_i(\theta_i)$-mixtures (i.e. an expectation $\int_{\Theta_{-i}} \sigma(k,\theta_{-i})\pi_i(\theta_i)(k, \mathrm{d}\theta_{-i})$) of all measurable, state and type profile contingent conjectures $\sigma(k,\theta_{-i}) \in \Delta(A_{-i})$ whose support is contained in $\mathrm{ICR}_{-i}^{m-1}(\theta_{-i})$ for all θ_{-i}. We call the sequence $(\mathrm{ICR}^m(\theta))_{m\geq 0}$ the ICR-*hierarchy* of θ.

3.2 Strategic Type Spaces

We define a Strategic Quotient (for ICR) as a pair $\mathscr{S} = (\mathcal{S}_i, \psi_i)_i$ consisting of an N-tuple of topological spaces $\mathcal{S}_i$ and continuous maps $\psi_i : \Delta_{K\times\mathcal{S}_{-i}} \to \mathcal{S}_i$ which satisfy both a type space quotient requirement and a strategic requirement.

Definition 3.1 (Type Space Quotient)**.** *A space* $\mathscr{S} = (\mathcal{S}_i, \psi_i)_i$ *is a* Type Space Quotient *if, for every Harsanyi type space* $\mathscr{H} = (\Theta_i, \pi_i)_i$ *there exist a family of maps* $(\eta_i)_i$ *for which the following diagram commutes:*

$$\begin{array}{ccc} \Theta_i & \xrightarrow{\pi_i} & \Delta_{K\times\Theta_{-i}} \\ \downarrow{\scriptstyle \eta_i} & & {\scriptstyle id}\downarrow \quad \downarrow{\scriptstyle \eta_{-i}} \\ \mathcal{S}_i & \xleftarrow[\psi_i]{} & \Delta_{K\times\mathcal{S}_{-i}} \end{array}$$

Definition 3.1 imposes a sufficient condition for two types of player i to have the same representation in $\mathcal{S}_i$. The two downward pointing arrows on the right of the diagram coarsen the sigma algebra of every type's beliefs. The commutativity of the diagram then requires the following: If the beliefs of two types θ_i, θ_i' coincide on $K \times \mathcal{S}_{-i}$, then η_i maps θ_i and θ_i' to the same point in S_i. Note that the reverse implication is not required by the diagram. That is, two types with distinct beliefs on $K \times \mathcal{S}_{-i}$ could also be mapped to the same point in $\mathcal{S}_i$.

Thus, in our model, types partition beliefs[1]. This contrasts with Harsanyi

[1] Chen et al. (2016a) introduce the notion of "frames" which are partitions of type spaces that are compatible with the belief structure of the types. Frames are thus a special case of what we call type space quotients.

types spaces, where a type is associated uniquely to a belief, and with universal type spaces, where types and beliefs are homeomorphic. The universal type space together with the canonical belief maps that associate to each canonical type the corresponding belief on thet state of nature as well as on other player's canonical types is thus a special case of a type space quotient.

Note also that the universal type space is the smallest object onto which information can be projected through faithful transformations in the sense of Gossner (2000). By partitioning belief spaces we allow for type space quotients to capture coarser information than universal type spaces.

The point, $\mathcal{S}_i = \{*\}$ and constant map $\psi_i \colon \Delta_{K\times\mathcal{S}_{-i}} \to \{*\}$ is another example of a type space quotient, for any game. This second example shows that type space quotients may fail to capture strategically relevant information. In our way to introduce a minimality requirement that spaces capture such information, we now define strategically closed families of behaviors.

Definition 3.2 (Strategic Closure)**.** *A strategically closed family of behaviors for* (S, ψ) *is a family* $\mathcal{A}_i$ *for each player* i *of continuous mappings* $\alpha_i : \mathcal{S}_i \to \mathcal{B}_i$ *such that,*

1. $\mathcal{A}_i$ *contains the constant map equal to* A_i
2. *for every* $\alpha_{-i} \in \mathcal{A}_{-i}$*, there exists* $\alpha_i \in \mathcal{A}_i$ *such that the following diagram commutes:*

$$\begin{array}{ccc} \Delta_{K\times\mathcal{S}_{-i}} & \xrightarrow{\psi_i} & \mathcal{S}_i \\ {\scriptstyle id}\downarrow \;\; \downarrow{\scriptstyle \alpha_{-i}} & & \downarrow{\scriptstyle \alpha_i} \\ \Delta_{K\times\mathcal{B}_{-i}} & \xrightarrow{\mathrm{BR}_i} & \mathcal{B}_i \end{array}$$

For a given pair $(\mathcal{S}, \psi)$, a set $\mathcal{A}_i$ consists of correspondences α_i which map points in $\mathcal{S}_i$ to action sets. As a minimality requirement on $\mathcal{A}_i$ point 1 of the definition imposes that each $\mathcal{A}_i$ contains the correspondence $s_i \mapsto A_i$ that precludes no action, for any $s_i \in \mathcal{S}_i$.

In point 2 of the definition, commutativity of the diagram imposes two requirements. First, for a family $\mathcal{A}$ to be strategically closed, the diagram imposes a measurability requirement on $\mathcal{S}$: It requires beliefs that induce different best replies to a behavior in $\mathcal{A}_{-i}$ to be associated to distinct points in $\mathcal{S}_i$. That is, given any profile $\alpha_{-i} \in \mathcal{A}_{-i}$, player i's best-response correspondence to this profile, seen from $\Delta_{K\times\mathcal{S}_{-i}}$ to A_i, is in fact S_i-measurable.

Second, any strategically closed family $\mathcal{A}$ must be closed under best replies: A player's best reply to a profile in $\mathcal{A}_i$, viewed as a correspondence from $\mathcal{S}_i$ to A_i is in $\mathcal{A}_i$.

Definition 3.3 (Strategic Quotient)**.** *A Strategic Quotient is a type space quotient* $(\mathcal{S}, \psi)$ *that admits a strategically closed family of behaviors.*

As an example, the universal type space together with the canonical maps is a strategic quotient, for any game: If other players' strategies are measurable wrt. their canonical types, so is a best-response. The type space consisting of a single point for each player together with the constant map is a strategic quotient when there is no uncertainty on nature, as in those games there is always a constant best-response to constant strategies of the other players. It is, however, not a strategic quotient for general incomplete information games. Take for instance the classical *electronic mail game* (Rubinstein, 1989) or the coordination game of Section 4.2. In that game, the best-response of a player to the other player using either of possible strategies depends on their belief on nature, hence cannot be captured by a constant map.

The next definition formalizes the idea that one quotient is smaller than another one.

Definition 3.4. *A space* $\mathscr{S} = (\mathcal{S}_i, \psi_i)_i$ *is* smaller *than another space* $\tilde{\mathscr{S}} = (\tilde{\mathcal{S}}_i, \tilde{\psi}_i)_i$ *if there exist a continuous surjection from* $\tilde{\mathcal{S}}$ *to* $\mathcal{S}$ *so that the following diagram commutes:*

$$\begin{array}{ccc} \tilde{\mathcal{S}}_i & \xleftarrow{\tilde{\psi}_i} & \Delta_{K \times \tilde{\mathcal{S}}_{-i}} \\ \downarrow\!\!\downarrow & & {\scriptstyle id}\downarrow \quad \downarrow\!\!\downarrow \\ \mathcal{S}_i & \xleftarrow[\psi_i]{} & \Delta_{K \times \mathcal{S}_{-i}} \end{array}$$

In this definition a type space quotient is smaller than another if the latter admits a representation of the former. That is, all types in the former quotient can be obtained by merging types of the latter. The diagram above requires the following sufficient condition for merging types: If the beliefs of types in the latter quotient coincide on the smaller quotient then these types are merged to the same point in the smaller quotient. The definition below then identifies the minimal strategic quotient according to Definition 3.4[2].

[2]Formally, Strategic Quotients form a category whose objects are given by the pairs $\mathscr{S}$

Definition 3.5 (STS). *A strategic quotient is called the strategic type space (STS) if it is smaller than every strategic quotient.*

By Definition 3.2, all strategic quotients must distinguish types which have different best replies to some strategic behavior. Hence the STS should merge players' types whenever these types have identical best replies to all strategic behaviors from a strategically closed family. It also follows from the discussion above that a STS must be at most as large as the univeral type space.

3.3 The Strategic Type Space

In this section we establish existence and essential uniqueness of the STS. We prove this result by characterizing the STS in terms of ICR hierarchies: First, we show the ICR hierarchies can be recovered from any Strategic Quotient (Lemma 3.1). We then provide a construction of $\mathcal{S}_i$, the set of best reply hierarchies for a game. This construction is canonical as it makes no reference to any Harsanyi type space. We show that these hierarchies coincide with all ICR hierarchies that can arise in all types in all Harsanyi type spaces (Lemma 3.2). We then construct a map ψ_i, which associates beliefs to best-reply hierarchies and prove that the pair $(\mathcal{S}, \psi)$ is a Strategic Quotient (Lemma 3.3). We deduce that $(\mathcal{S}, \psi)$ is a STS and show that it is essentially unique (Theorem 3.1).

Our first theorem states that every Strategic Quotient allows to recover the ICR hierarchies from any Harsanyi type.

Lemma 3.1 (Factorization of ICR). *For every strategic quotient $(\mathcal{S}_i, \psi_i)_i$ and every $m \in \mathbb{N}$, there exists continuous $\alpha_i^m : \mathcal{S}_i \to \mathcal{B}_i$ so that for every Harsanyi type space $(\Theta_i, \pi_i)_i$ and associated maps $(\eta_i)_i$ satisfying the diagram of Definition 3.1,*

$$\mathrm{ICR}_i^m(\theta_i) = \alpha_i^m \circ \eta_i(\theta_i), \ \forall \theta_i \in \Theta_i, \ \forall i \in N$$

The proof of this result, as well as all others, is in the appendix.

We now construct the set $\mathcal{S}$ of all hierarchies of best replies. The first level of the hierarchy is given by a player's best replies to beliefs on nature and any opponents' play. Every subsequent level of a best reply hierarchy is

satisfying strategic closure and whose morphisms are given by diagrams as in Definition 3.4. A minimal strategic quotient is thus a terminal object in the category.

then obtained by computing best replies to beliefs on nature and lower levels of best reply hierarchies.

We construct inductively the sets of m-order best reply hierarchies S_i^m as m-fold sequences of action set profiles. Let $\mathcal{S}_i^0 := \{A_i\}$ for every i. Given $\mathcal{S}_i^{m-1}$ for every i, we define the subset $\mathcal{S}_i^m \subseteq \mathcal{S}_i^{m-1} \times \mathcal{B}_i$ of sequences of the form $s_m^i = (A_i, b_i^1, \ldots, b_i^m)$ for which there exists a probability distribution $p_i \in \Delta_{K \times \mathcal{S}_{-i}^{m-1}}$ satisfying

$$\mathrm{BR}_i(\mathrm{marg}_{K,l}(p_i)) = b_i^{l+1}, \ \forall\, l < m, \tag{3.1}$$

where $\mathrm{marg}_{K,l}(p_i)$ is the marginal probability of p_i on $K \times \prod_{j \neq i} \mathrm{proj}_l(\mathcal{S}_j^m)$, where proj_l denotes the projection on the l-th coordinate. We define the set of player i's best reply hierarchies as

$$\mathcal{S}_i := \{s_i \in \mathcal{B}_i^{\mathbb{N}} : s_i^m \in \mathcal{S}_i^m, \ \forall\, m \in \mathbb{N}\}.$$

Lemma 3.2 states that the best reply hierarchies $\mathcal{S}$ characterize all ICR hierarchies that can arise in any Harsanyi type space.

Lemma 3.2 (Best-Reply Hierarchies are ICR Hierarchies)**.**

(i) Let $s^m \in \mathcal{B}^m$, then $s^m \in \mathcal{S}^m$ if and only if there exists a Harsanyi type space (Θ, π) and a type profile $\theta \in \Theta$ so that $s^m = (\mathrm{ICR}^l(\theta))_{l \leq m}$.

(ii) Let $s \in \mathcal{B}^{\mathbb{N}}$, then $s \in \mathcal{S}$ if and only if there exists a Harsanyi type space (Θ, π) and a type profile $\theta \in \Theta$ so that $s = (\mathrm{ICR}^l(\theta))_{l \geq 0}$.

For every $m \in \mathbb{N}$, we define a beliefs map $\psi_i^m : \Delta_{K \times \mathcal{S}_{-i}^{m-1}} \to \mathcal{S}_i^m$ by

$$\psi_i^m(p_i) := (A_i, \mathrm{BR}_i(\mathrm{marg}_{K,1}(p_i)), \ldots, \mathrm{BR}_i(\mathrm{marg}_{K,m-1}(p_i))).$$

Any belief p_i on $K \times S_{-i}$ induces, through the projection on the first m coordinates of S_i, a belief p_i^m on $K \times S_{-i}^{m-1}$, thus an element $\psi_i^m(p_i^m) \in S_i^m$. By definition of ψ_i^m, for every $l \leq m$, the first l elements of $\psi_i^m(p_i^m)$ coincide with $\psi_i^l(p_i^l)$. Thus, the sequence $(\psi_i^m(p_i^m))_i$ defines a unique element of S_i, which we denote $\psi_i(p_i)$.

Note that once the set $\mathcal{S}_{-i}$ of all other players' best-reply hierarchies is known, player i's best-reply hierarchies are fully characterized by marginal beliefs and do not depend on correlations across different levels of $\mathcal{S}_{-i}$.

We now specify the topology on the set $\mathcal{S}$. Recall that strategic closure requires the Strategic Quotient to admit a closed family of continuous strategic behaviors. By construction, the coordinates of a best-reply hierarchy correspond to a closed family of strategic behaviors. We thus endow $\mathcal{S}$ with its product topology, i.e. the coarsest topology so that all coordinate projections are continuous. Lemma 3.3 below states that $(\mathcal{S}, \psi)$ is a Strategic Quotient and that $\mathcal{S}$ is a topological quotient of the universal type space of Mertens and Zamir (1985).

Lemma 3.3 (ICR Hierarchies form the STS)**.** *$(\mathcal{S}, \psi)$ is a strategic quotient. Moreover, the maps η from the universal type space to $\mathcal{S}$ are quotient maps, i.e. continuous open surjections.*

By Lemma 3.1 any finite order ICR hierarchy can be recovered continuously from any quotient space. By Lemma 3.2 the set $\mathcal{S}$ coincides with all ICR hierarchies. Then by Lemma 3.3, $(\mathcal{S}, \psi)$ is a Strategic Quotient which can be recovered from all Strategic Quotients. The product topology then ensures that $(\mathcal{S}, \psi)$ is a quotient which is minimal. As the property of minimality is universal, every STS is homeomorphic to $\mathcal{S}$. Theorem 3.1 thus states existence and essential uniqueness of the STS:

Theorem 3.1 (Existence and Essential Uniqueness of STS)**.**

(i) $(\mathcal{S}, \psi)$ is a STS.

(ii) If $(\mathcal{S}', \psi')$ and $(\mathcal{S}'', \psi'')$ are STS then $\mathcal{S}''$ and $\mathcal{S}'$ are homeomorphic.

4 Finite Representation of STS

For every truncated sequence $s^m \in \mathcal{S}^m$, define its orbit as:

$$O^m(s^m) \coloneqq \{s^m_m, s) \in \mathcal{B} \times \mathcal{B}^{\mathbb{N}} : (s^m, s) \in \mathcal{S}\}. \tag{4.1}$$

Denote the collection of all orbits

$$\Omega \coloneqq \{O^m(s^m) : s^m \in \mathcal{S}^m, m \in \mathbb{N}\}. \tag{4.2}$$

Theorem 4.1. *Ω is a finite set.*

We break down the proof into six Claims, all proven in Appendix A.2. Here we provide an overview of the arguments. In order to prove Theorem 4.1, we will exploit a monotonicity property of BR_i according to which beliefs with smaller supports (according to stochastic dominance wrt set inclusion) admit smaller best-response sets: $\mathrm{BR}_i(p) \subseteq \mathrm{BR}_i(p')$ whenever p can be obtained from p' by shifting probability mass from action sets to subsets (see Lemma A.2 in Appendix A.1). In preparation for using this property, we introduce the *best-response map on sets of sequences* for player i, $B_i \colon 2^{\mathcal{B}_i^{\mathbb{N}}} \to 2^{\mathcal{B}_i^{\mathbb{N}}}$ as follows

$$B_i(X) := \{\psi_i(p_i) : p_i \in \Delta(K \times X_{-i})\}. \tag{4.3}$$

Let $B(X) := \prod_i B_i(X)$. Note that by construction, $B(\mathcal{S}) = \mathcal{S}$ and so for every subset $S \subseteq \mathcal{S}$, $B(S) \subseteq \mathcal{S}$. B inherits the following monotonicity property from BR.

Monotonicity. For any $X, X' \subseteq \mathcal{B}^{\mathbb{N}}$ write $X \ll X'$ when the following two properties hold: 1) For every $x \in X$ there exists $x' \in X'$ so that for all $m \in \mathbb{N}$, $x_m \subseteq x'_m$. 2) For every $x' \in X'$ there exists $x \in X$ so that for all $m \in \mathbb{N}$, $x_m \subseteq x'_m$. Lemma A.2 implies:

Claim 4.1. $X \ll X' \implies B(X) \ll B(X')$.

A sequence $s \in \mathcal{S}$ is *maximal after round* $m \in \mathbb{N}$ if there does not exist $\hat{s} \in \mathcal{S} \setminus \{s\}$ so that the following two properties hold: 1) $s^m = \hat{s}^m$ and 2) for all $n > m$, $s_n \subseteq \hat{s}_n$. A sequence $s \in \mathcal{S}$ is *maximal at round* $m \in \mathbb{N}$ if there does not exist $\hat{s} \in \mathcal{S} \setminus \{s\}$ so that the following two properties hold: 1) $s^{m-1} = \hat{s}^{m-1}$ and 2) $s_m \subseteq \hat{s}_m$. We use Claim 4.1 to establish that all maximal sequences are maximal at every round:

Claim 4.2. *The sequence $s \in \mathcal{S}$ is maximal after round m if and only if it is maximal at round n for all $n > m$.*

Let $\bar{\mathcal{S}}^m(s^m)$ denote the collection of sequences $\tilde{s} \in \mathcal{S}$ which are maximal after round m and satisfy $\tilde{s}^m = s^m$. Claim 4.2 then implies the following finiteness property of maximal sequences:

Claim 4.3. *There exists $L \in \mathbb{N}$ so that for every $m \in \mathbb{N}$ and every $s \in \mathcal{S}$,*

$$|\bar{\mathcal{S}}^m(s^m)| < L. \tag{4.4}$$

Say that $X \subseteq \mathcal{B}^{\mathbb{N}}$ has *converged at round* $m \in \mathbb{N}$ if for all $x \in X$ and all $n, l \geq m$, $x_n = x_l$. Claim 4.3 and the monotonicity of B imply the following convergence property of maximal sequences after any finite history:

Claim 4.4. *There exists M so that for every $m \in \mathbb{N}$ and $s \in \mathcal{S}$, the set $\bar{\mathcal{S}}^m(s^m)$ converged at round $m + M$.*

Finite Generation of the STS. We present a construction of the STS through iterated applications of the best-response operator. At each step of the construction, we allow for at most one more non-maximal transition than in the previous step. In turn, this construction will allow us to conclude on the finiteness of Ω.

We denote by $\bar{T}^{0,0}$ the set of maximal sequences $\bar{\mathcal{S}}^0(s^0)$. For every $m > 0$, let $\bar{T}^{m,0}$ denote the set of m-maximal sequences, i.e., the set of sequences that admit at most m non-maximal transitions. We also let $\bar{T}^{m,n}$ be the subset of $\bar{T}^{m,0}$ consisting of sequences that admit at most $m-1$ non-maximal transitions before round n and that are maximal after round n. Sequences in $\bar{T}^{m,n}$ may or not have a non-maximal transition at round n. Finally, let $\bar{T}^m = \cup_n \bar{T}^{m,n}$.

Our next lemma shows that the sets $\bar{T}^{m,n}$ are constructed iteratively through the best-response operator.

Claim 4.5. *For every $m \geq 0, n > 0$, $\bar{T}^{m,n}$ is the set of sequences s that are maximal after round n, admit at most $m-1$ non-maximal transitions before round n, and such that:*

$$s \in B(\bar{T}^{m,n-1}).$$

Claim 4.6 below establishes that $\mathcal{S}$ can be obtained as a finite union of the sets $(\bar{T}^{m,0})_m$. The result follows from the fact that every sequence $s \in \mathcal{S}$ only makes a finite number of non-maximal transitions.

Claim 4.6. *There exists $N \in \mathbb{N}$ so that*

$$\bigcup_{m=1}^{N} \bar{T}^{m,0} = \mathcal{S}. \tag{4.5}$$

Cyclicality and Finiteness. We now show that for every m the infinite sequence of sets $(\bar{T}^{m,n})_n$ is eventually cyclic, which will establish the finiteness of Ω. First, we conclude from Claim 4.4 and the construction of each set $\bar{T}^{m,n}$ that the sequences in $\bar{T}^{m,n}$ converge within a bounded number of rounds after round n:

Corollary 4.1. *Let N satisfy (A.15). For every $m \leq N$ there exists $M_m \in \mathbb{N}$ so that for every $n \in \mathbb{N}$, $\bar{T}^{m,n}$ has converged at round $n + M_m$.*

By monotonicity of B, sequences in $\bar{T}^{m,n}$ are best-replies to sequences in $\bar{T}^{m,n-1}$ and by Corollary 4.1 all sequences in $\cup_{l\geq M_m}\bar{T}^{m,n-l}$ have converged before round n; We conclude that for any two rounds $n>\tilde{n}$ so that

$$\{O(s^l):s\in\bar{T}^{m,l}\}=\{O(s^{l+n-\tilde{n}}):s\in\bar{T}^{m,l+n-\tilde{n}}\}, \tag{4.6}$$

for every $l\in\{\tilde{n}-M_m,\dots,\tilde{n}\}$, we also have that

$$\{O(s^{\tilde{n}+1}):s\in\bar{T}^{m,\tilde{n}+1}\}=\{O(s^{n+1}):s\in\bar{T}^{m,n+1}\}. \tag{4.7}$$

From Claim 4.3 we conclude that for every m, the number of sequences that are contained in $\bar{T}^{m,n}\setminus\bar{T}^{m,n-1}$ is bounded uniformly over all $n\in\mathbb{N}$.

Corollary 4.2. *Let N satisfy (A.15). For every $m\leq N$ there exists $L_m\in\mathbb{N}$ so that for every $n\in\mathbb{N}$,*

$$|\bar{T}^{m,n}\setminus\bar{T}^{m,n-1}|\leq L_m. \tag{4.8}$$

We conclude from both corollaries that for every $m\in\mathbb{N}$, there exists $n_m\in\mathbb{N}$ and $\tilde{n}_m<n_m$ so that (4.6) holds and so we deduce the result below:

Claim 4.7. *Let N satisfy (A.15). For every $m<N$ there exists $z_m\in\mathbb{N}$ and $n_m\in\mathbb{N}$ so that for all $s\in\bar{T}^{m+1,0}$*

$$O(s^n)=O(s^{n+z_m}),\ \forall\ n\geq n_m. \tag{4.9}$$

Claims 4.7 and 4.6 then imply that the set of tails is finite, which is what we needed to show.

4.1 Automaton Representation

Define an automaton as a tuple $\hat{\mathscr{A}}=(\hat{\Omega},\hat{\beta},\hat{\preceq},\hat{\omega}^0)$, where $\hat{\Omega}$ is a finite set of automaton states, $\hat{\beta}\colon\hat{\Omega}\to\mathcal{B}$ assigns an action set profile to every state, $\hat{\preceq}$ is a binary successor relation on $\hat{\Omega}$ and $\hat{\omega}^0\in\hat{\Omega}$ is an initial state. A path on the automaton is a sequence of states $(\omega^0,\omega^1,\omega^2,\dots)$ so that $\omega^0=\hat{\omega}^0$ and for every $m\in\mathbb{N}$,

$$\omega^m\hat{\preceq}\omega^{m+1}. \tag{4.10}$$

Let $P_{\hat{\mathscr{A}}}$ denote the set of paths. From Theorem 4.1 we obtain a finite automaton representation of all ICR-hierarchies: Let $\beta\colon\Omega\to\mathcal{B}$ recover the first coordinate from sequences in each orbit. Define the shift operator on sequences $\gamma\colon \mathcal{B}^{\mathbb{N}} \to \mathcal{B}^{\mathbb{N}}$, which for each sequence $(s^1, s^2, \dots)$ removes the first coordinate

$$\gamma\colon (s^1, s^2, \dots) \mapsto (s^2, \dots). \tag{4.11}$$

Define the successor relation $\preceq$ on Ω, where for every $\omega, \hat{\omega} \in \Omega$, $\omega \preceq \hat{\omega}$ if and only if

$$\hat{\omega} \subseteq \{\gamma(s) : s \in \omega\}. \tag{4.12}$$

Letting $\omega^0 = O^0(s^0)$, we obtain an automaton.

Theorem 4.2. *The tuple* $\mathscr{A} = (\Omega, \beta, \preceq, \omega^0)$ *is an automaton so that*

$$\{(\beta(\omega^0), \beta(\omega^1), \dots) : (\omega^0, \omega^1, \dots) \in P_{\mathscr{A}}\} = \mathcal{S}. \tag{4.13}$$

We deduce that the STS is countable, compact and Hausdorff.

4.2 Example

In this section we illustrate our results and construct the STS in an two players coordination Game.

Consider the following two-player game: $N = \{1, 2\}$, $K = \{-1, 1\}$ and $A_i = A = \{a, b\}$ where payoffs are given by:

	a	b
a	$k,\ k$	$-1,\ 0$
b	$0, -1$	$0,\ 0$

If player i (row player) believes that $k = 1$ with probability less than $\frac{1}{2}$ then b is a dominant action. Otherwise, neither action dominates the other. First order hierarchies of best replies in this game, $\mathcal{S}_i^1$, are thus given by $\{(A, b), (A, A)\}$. The first pair corresponds to beliefs which put less than half of the probability on $k = 1$. Indeed, recall that $\mathcal{S}_{-i}^0 = \{a, b\}$ from Section 3.3 and consider any belief $p \in \Delta_{K \times \mathcal{S}_{-i}^0}$. Player i thus forms best replies to p-mixtures of state-contingent conjectures $\sigma : K \to \Delta_A$. In the simplex $\Delta_{K \times A}$, these p-mixtures over conjectures form geometric rectangles - the set of probabilities on $K \times A$ with constant marginal belief on K given by (p_1, p_{-1}). The right panel of Figure 1 illustrates these rectangles for $p_1 < \frac{1}{2}$ and $p_1 \geq \frac{1}{2}$. The left of Figure 1 plots the simplex $\Delta_{K \times A}$, where the shaded triangle with dashed contour marks the boundary of the partition induced by the best response correspondence of player i. When $p_1 < \frac{1}{2}$, the mixture of the conjectures is entirely included in the region where b_i is the unique best-response. When $p_1 \geq \frac{1}{2}$, the conjectures cross regions where a and b, or both are best-responses. Hence $\mathcal{S}_i^1 = \{(A, A), (A, b)\}$.

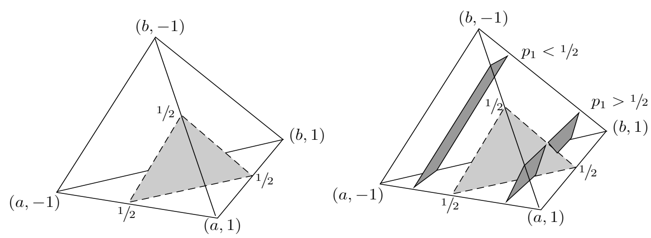

Figure 1: For all beliefs in the region between the shaded triangle (excluded) and the sub simplex spanned by $(a, -1)$, $(b, -1)$ and $(b, 1)$, player i's best response is always b. For all beliefs in the region between the shaded triangle (excluded) and $(a, 1)$, player i's best response is always a and on the shaded triangle all beliefs induce both actions a and b as best response.

We repeat the same procedure on $\mathcal{S}_{-i}^1$. For a belief p on $K \times \mathcal{S}_{-i}^1$ of player i, let p_b denote the probability put on hierarchies ending at b, let p_1 denote the probability put on $k = 1$ and $p_{k,b}$ be the joint probability on state $k \in \{-1, 1\}$ and hierarchies ending at b. As can be seen in Figure 1, $2p_1 - p_b + (p_{-1,b} - p_{1,b}) < 1$ describes the portion of a rectangle associated to p_1 where b is a unique best reply for player i. Hence the set of beliefs on $K \times \mathcal{S}_{-i}^1$ so that BR_i maps to A is given by $p_1 \geq \frac{1}{2}$ and $2p_1 - p_b + (p_{-1,b} - p_{1,b}) \geq 1$. We deduce that $\mathcal{S}_i^2 = \{(A, A, A), (A, A, b), (A, b, b)\}$ corresponds to the following partition of $\Delta_{K \times \mathcal{S}_{-i}^1}$:

(1) $2p_1 - p_b + (p_{-1,b} - p_{1,b}) \geq 1$ and $p_1 \geq 1/2$, for (A, A, A)

(2) $2p_1 - p_b + (p_{-1,b} - p_{1,b}) < 1$ and $p_1 \geq 1/2$, for (A, A, b)

(3) $p_1 < 1/2$, for (A, b, b)

Note that these conditions only depend i's beliefs on K and on the last coordinate in $\mathcal{S}_{-i}^1$. As the last coordinates of $\mathcal{S}^1$ are the same as the last coordinates of $\mathcal{S}^1$ we deduce that the game is indeed simple. we argue that all transitions in $\mathcal{S}$ are described by the three rules above. The STS automaton in Figure 2 below illustrates the transition for coordinates in $\mathcal{S}_i$ and $\mathcal{S}_{-i}$ in this game:

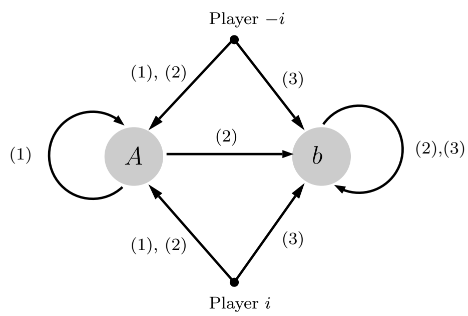

Figure 2: STS Automaton.

The automaton represented in Figure 2 above describes the following transitions: If player i's m-th coordinate is A, then it must be that $p_1 \geq \frac{1}{2}$ and i's $(m+1)$-th coordinate in any strategic type must be one of $\beta_i(A) \coloneqq \{A, b\}$. Moreover, the $(m+1)$-th coordinate is A if i also believes in strategic types of $-i$ whose m-th coordinate is b with low enough probability. That is, i believes that $-i$ plays b with low enough probability in round m (i.e. condition (1)). Otherwise, the $(m+1)$-th coordinate must be b (condition (2)). However, if i's m-th coordinate was b, then i's $(m+1)$-th coordinate must be $\beta_i(b) = b$. In this case, i's beliefs satisfy either condition (3) or condition (2).

For this example, note first that the only possible change in the last coordinate when going from $\mathcal{S}_i^1$ to $\mathcal{S}_i^2$ is to move from A to b. A probability on $K \times \mathcal{S}_{-i}^2$ must therefore put at least as much probability on sequences ending with b than its marginal on $K \times \mathcal{S}_{-i}^1$. But under this constraint, third order types can also only move from A to b or stay unchanged. Hence the automaton above generates all the sequences in $\mathcal{S}$.

5 Discussion

Our construction is topological and the minimality property implies that the STS is endowed with the product topology on sequences. However, the automaton we have constructed in Section 4.2 illustrates how strategic continuity fails in the product topology. In Figure 2, the type space where the state in K is common knowledge can be represented by the constant sequence cycling through the full action set A forever and the sequence cycling through action b forever after the first transition. The strategic type space for the Email game in Rubinstein (1989) can be described by attaching transition probabilities to the arrows in Figure 2: Transitions labeled (1), (2) and (1) have probability $1-\epsilon$ while transitions labeled (3) and (2) have probability ϵ. The paths in this type space converge to the common knowledge path in the product topology. However, their limits do not converge. Results in Chen et al. (2016a) and Chen, Di Tillio, Faingold, and Xiong (2010) suggest that the topology of uniform convergence would ensure strategic continuity for a fixed game.

A related observation is that the sequence of approximating paths requires an ever growing type space while the common knowledge type is binary. The complexity of a type, loosely defined as the minimal size of the type space required to contain this type is therefore not continuous in the product topology. We leave the study of type-complexity for future work.

In this paper we focus on Harsanyi type spaces and not on common prior models. In a companion paper, Gossner and Veiel (2024) represent common prior models as Markov chains on STS automata and use this representation to provide a finite characterization of rationalizable outcomes in common prior models.

References

Bernheim, B. D. (1984): "Rationalizable strategic behavior," *Econometrica*, 1007–1028.

Chen, Y.-C., A. Di Tillio, E. Faingold, and S. Xiong (2010): "Uniform topologies on types," *Theoretical Economics*, 5, 445–478.

——— (2016a): "Characterizing the strategic impact of misspecified beliefs," *The Review of Economic Studies*, 84, 1424–1471.

Chen, Y.-C., S. Takahashi, and S. Xiong (2014): "The robust selection of rationalizability," *Journal of Economic Theory*, 151, 448–475.

——— (2016b): "Predictions by Refining Rationalizability: Sharpness Versus Robustness," *unpublished manuscript.*

Dekel, E., D. Fudenberg, and S. Morris (2006): "Topologies on types," *Theoretical Economics*.

——— (2007): "Interim correlated rationalizability," *Theoretical Economics*, 2, 15–40.

Ely, J. C. and M. Peski (2011): "Critical types," *The Review of Economic Studies*, 78, 907–937.

Gossner, O. (2000): "Comparison of information structures," *Games and Economic Behavior*, 30, 44–63.

Gossner, O. and R. Veiel (2024): "Information Design for Rationalizability," .

Harsanyi, J. C. (1967): "Games with incomplete information played by "Bayesian" players, I–III Part I. The basic model," *Management science*, 14, 159–182.

Mertens, J.-F. and S. Zamir (1985): "Formulation of Bayesian analysis for games with incomplete information," *International Journal of Game Theory*, 14, 1–29.

Morris, S., H. S. Shin, and M. Yildiz (2016): "Common belief foundations of global games," *Journal of Economic Theory*, 163, 826–848.

Pearce, D. G. (1984): "Rationalizable strategic behavior and the problem of perfection," *Econometrica*, 1029–1050.

Rubinstein, A. (1989): "The Electronic Mail Game: Strategic Behavior Under "Almost Common Knowledge"," *The American Economic Review*, 385–391.

Weinstein, J. and M. Yildiz (2007): "A structure theorem for rationalizability with application to robust predictions of refinements," *Econometrica*, 75, 365–400.

A Appendix

A.1 Preliminaries

We introduce some additional notation. For any mapping $f : X \to Y$ we denote the image of f by $\mathrm{Im}(f)$.The following lemma is key for our results:

Lemma A.1 (BR Factorization of ICR)**.** *Let $(T_i, \pi_i)_i$ be a Harsanyi type space. Then for every m and every i, ICR_i^m admits the following factorization through BR_i,*

$$\begin{array}{ccc} T_i & \xrightarrow{\pi_i} & \Delta_{K\times T_{-i}} \\ \big\downarrow{\scriptstyle \mathrm{ICR}_i^m} & & {\scriptstyle \mathrm{id}}\big\downarrow \quad \big\downarrow{\scriptstyle \mathrm{ICR}_{-i}^{m-1}} \\ \mathcal{B}_i & \xleftarrow[\mathrm{BR}_i]{} & \Delta_{K\times \mathcal{B}_{-i}} \end{array}$$

Proof. Let $\sigma : K \times T_{-i} \to \Delta(A_{-i})$ be a $\pi_i(t_i)$-measurable conjecture. Write the t_i mixture of σ as

$$\langle \sigma, \pi_i(t_i)\rangle_{T_i}(k, a_{-i}) \coloneqq \int_{T_{-i}} \sigma(k, t_{-i})(a_{-i})\ \pi_i(t_i)(k, \mathrm{d}t_{-i}),\ \forall\ k, a_{-i}$$

Then by definition of ICR we have that

$$\mathrm{ICR}_i^m(t_i) = \left\{\mathbf{B}(\langle \sigma, t_i\rangle_{T_i}) : \sigma \text{ is } \pi_i(t_i)\text{-meas., supp } \sigma(k, t_{-i}) \subseteq \mathrm{ICR}_{-i}^{m-1}(t_{-i})\right\}$$

where $\mathbf{B}(p) \coloneqq \mathrm{argmax}_{a_i} \sum_{k,a_{-i}} u_i(k, a_i, a_{-i})p(k, a_{-i})$ for every $p \in \Delta(K \times A_{-i})$. We now show that for every $\pi_i(t_i)$-measurable conjecture $\sigma : K \times T_{-i} \to \Delta(A_{-i})$ so that supp $\sigma(k, t_{-i}) \subseteq \mathrm{ICR}_{-i}^{m-1}(t_{-i})$ we can construct a conjecture $\tilde{\sigma} : K \times \mathcal{A}_{-i} \to \Delta(A_{-i})$ so that supp $\sigma(k, b_{-i}) \subseteq b_{-i}$ and

$$\langle \tilde{\sigma}, p_i^{m-1}(t_i)\rangle(k, a_{-i}) = \langle \sigma, \pi_i(t_i)\rangle_{T_i}(k, a_{-i}),\ \forall\ k, a_{-i} \tag{A.1}$$

where $p_i^{m-1}(t_i) \coloneqq \pi_i(t_i) \circ (id \times \mathrm{ICR}_{-i}^{m-1})^{-1}$ is the push forward probability. Define the required conjecture for every k, a_{-i} and $b_{-i} \in \mathrm{Im}(\mathrm{ICR}_{-i}^{m-1})$,

$$\tilde{\sigma}(a_{-i}|k, b_{-i}) \coloneqq \int_{(\mathrm{ICR}_{-i}^{m-1})^{-1}(b_{-i})} \sigma(k, t_{-i})(a_{-i})\ \pi_i(t_i)(k, \mathrm{d}t_{-i})$$

and for $b_{-i} \notin \mathrm{Im}(\mathrm{ICR}_{-i}^{m-1})$ let $\tilde{\sigma}$ be arbitrary. This construction clearly satisfies (A.1). Conversely, $\tilde{\sigma}$ can be written as a conjecture $K \times T_{-i} \to \Delta(A_{-i})$ which is constant on the pre-image of ICR_{-i}^{m-1} and so the result follows. $\square$

A *monotone stochastic transformation* for player i is a map $\rho_i\colon K\times B_{-i}\to\Delta(B_{-i})$ so that for every $b\in B$ and $k\in K$,

$$b'_{-i}\subseteq b_{-i},\ \forall\ b'_{-i}\in\operatorname{supp}(\rho_i(k,b_{-i})). \tag{A.2}$$

Lemma A.2 (Monotonicity of BR)**.** *For any monotone stochastic transformation $\rho_i\colon K\times\mathcal{B}_{-i}\to\Delta(\mathcal{B}_{-i})$ and for any $p_i\in\Delta(K\times\mathcal{B}_{-i})$,*

$$\mathrm{BR}_i(p_i\circ\rho_i)\subseteq\mathrm{BR}_i(p_i), \tag{A.3}$$

where for all $k\in K$ and $b_{-i}\in B_{-i}$,

$$p_i\circ\rho_i(k,b_{-i}):=\sum_{b'_{-i}\in B_{-i}}\rho_i(b_{-i}|k,b'_{-i})p_i(k,b'_{-i}). \tag{A.4}$$

Proof. Consider any conjecture $\sigma_i\colon K\times B_{-i}\to\Delta(A_{-i})$ so that $\operatorname{supp}(\sigma(\cdot|k,b_{-i}))\subseteq b_{-i}$ for all $k\in K, b_{-i}\in B_{-i}$. Now define the conjecture $\sigma_i\circ\rho_i$, which for every $a_{-i}\in A_{-i}, k\in K, b'_{-i}\in B_{-i}$ is given by

$$\sigma_i\circ\rho_i(a_{-i}|k,b'_{-i}):=\sum_{b_{-i}}\sigma_i(a_{-i}|k,b_{-i})\rho_i(b_{-i}|k,b'_{-i}). \tag{A.5}$$

Since ρ_i is monotone, the conjecture $\sigma_i\circ\rho_i$ also satisfies the support constraint of σ_i. Hence

$$\begin{aligned}\langle\sigma_i,p_i\circ\rho_i\rangle(k,a_{-i})&=\sum_{b'_{-i}\in B_{-i}}\left(\sum_{b_{-i}\in B_{-i}}\sigma_i(a_{-i}|k,b_{-i})\rho_i(b_{-i}|k,b'_{-i})\right)p_i(k,b'_{-i})\\&=\sum_{b'_{-i}\in B_{-i}}\sigma_i\circ\rho_i(a_{-i}|k,b'_{-i})p_i(k,b'_{-i})\\&=\langle\sigma_i\circ\rho_i,p_i\rangle(k,a_{-i}).\end{aligned} \tag{A.6}$$

Now the result is immediate from the definition of BR_i. □

A.2 Strategic Type Spaces

Lemma 3.1 *For every strategic type space $(\mathcal{S}_i,\psi_i)_i$ and every $m\in\mathbb{N}$, there exists measurable $\sigma_i^m:\mathcal{S}_i\to\mathcal{A}_i$ so that for every Harsanyi type space $(T_i,\pi_i)_i$ and associated maps $(\eta_i)_i$ satisfying the diagram of Definition 3.1,*

$$\mathrm{ICR}_i^m(t_i)=\sigma_i^m\circ\eta_i(t_i),\ \forall t_i\in T_i,\ \forall i\in N \tag{A.7}$$

Proof. Let $(\mathcal{S}_i, \psi_i)_i$ be a STS and let $(\Sigma_i)_i$ be a measurable family of strategically closed behaviors. Proceed by induction. Base case: The constant map $\sigma_i^0 : s_i \mapsto A_i$ is in Σ_i and for every Harsanyi type space $(T_i, \pi_i)_i$, $\mathrm{ICR}_i^0(t_i) = A_i = \sigma_i^0 \circ \eta_i(t_i)$ for all i and $t_i \in T_i$. Inductive hypothesis: Suppose $(\sigma_i^l)_{l \le m, i}$ satisfy (A.7). Then for every player i and type t_i, $\pi_i(t_i) \circ (id \times \mathrm{ICR}_{-i}^{m-1})^{-1} = \pi_i(t_i) \circ (id \times \sigma_{-i}^{m-1} \circ \eta_{-i})^{-1}$ and so by Lemma A.1 $\mathrm{ICR}_i(t_i) = \mathrm{BR}_i(\pi_i(t_i) \circ (id \times \sigma_{-i}^{m-1} \circ \eta_{-i})^{-1})$. By the *quotient property* in Definition 3.1 $\psi_i(\pi_i(t_i) \circ (id \times \eta_{-i})^{-1}) = \eta_i(t_i)$ and by the *strategic closure property* in Definition 3.2 there exists σ_i^m so that $\sigma_i^m \circ \eta_i(t_i) = \mathrm{BR}_i(\pi_i(t_i) \circ (id \times \sigma_{-i}^{m-1} \circ \eta_{-i})^{-1})$. □

Lemma 3.2

(i) Let $s^m \in \mathcal{B}^m$, then $s^m \in \mathcal{S}^m$ if and only if there exists a Harsanyi type space (T, π) and a type profile $t \in T$ so that $s^m = (\mathrm{ICR}^l(t))_{l \le m}$.

(ii) Let $s \in \mathcal{B}^{\mathbb{N}}$, then $s \in \mathcal{S}$ if and only if there exists a Harsanyi type space (T, π) and a type profile $t \in T$ so that $s = (\mathrm{ICR}^l(t))_{l \ge 0}$.

Proof. We start with (i): We prove both directions inductively using the following base case: For every Harsanyi type space $(T_i, \pi_i)_i$ and player i we have $\mathcal{S}_i^0 = \{A_i\} = \mathrm{ICR}_i^0(t_i)$. For the "if" direction: Inductive hypothesis: Fix a Harsanyi type space $(T_i, \pi)_i$, and suppose that for any player i and type $t_i \in T_i$, $(\mathrm{ICR}_i^l(t_i))_{l \le m-1} \in \mathcal{S}_i^{m-1}$. Then by Lemma A.1, $\mathrm{ICR}_i^m(t_i) = \mathrm{BR}_i(\pi_i(t_i) \circ (id \times \mathrm{ICR}_{-i}^{m-1})^{-1})$ and so $\pi_i(t_i) \circ (id \times \prod_{l \le m-1} \mathrm{ICR}_{-i}^l)^{-1} \in \Delta(K \times \mathcal{S}_{-i}^{m-1})$ implies $(\mathrm{ICR}_i^l(t_i))_{l \le m} \in \mathcal{S}_i^m$, as required.

Prove the "only if" by constructing a type space for every $s^m \in \mathcal{S}^m$ containing a type t so that $\prod_{l \le m} \mathrm{ICR}^l(t) = s^m$. For every i, $s_i^m \in \mathcal{S}_i^m$ implies that there exists $p_i^m \in \Delta(K \times \mathcal{S}_{-i}^{m-1})$ so that $s_i^m = (\mathrm{BR}_i(\mathrm{marg}_{K, \mathcal{S}_{-i}^l}(p_i^m)))_{l \le m-1}$. We proceed by induction on m. Suppose we have picked a selection $\iota_i^m \colon \mathcal{S}_i^m \to \Delta(K \times \mathcal{S}_{-i}^{m-1})$ from $(\psi_i^m)^{-1}(\cdot)$. For every s_i^{m+1}, we may construct a selection ι_i^{m+1} so that for every $s_i^{m+1} \in \mathcal{S}_i^{m+1}$, $\mathrm{marg}_{K, \mathcal{S}_{-i}^m}(\iota_i^{m+1}(s_i^{m+1})) = \iota_i^m(s_i^m)$. Finally, by Lemma A.1 the m-th component $s_{i,m}^m$ of s_i^m is given by $s_{i,m}^m = \mathrm{ICR}_i^m(\iota_i^m(s_i^m))$, for all m, as required.

Finally, (ii) follows from the properties of inverse limits of the constructions in the proof of (i). □

Lemma 3.3 *(S, ψ) obtained in our construction of hierarchies of best replies are an STS.*

Proof. Lemmas A.1 and 3.2 (i) ensure that $(\mathcal{S}, \psi)$ is a type space quotient. The coordinate projection then yields a best-reply closed family of measurable strategic behaviors. □

Theorem 3.1

(i) $(\mathcal{S}, \psi)$ *is a minimal STS.*

(ii) If $(\mathcal{S}', \psi')$ *and* $(\mathcal{S}'', \psi'')$ *are minimal STS then* $\mathcal{S}''$ *and* $\mathcal{S}'$ *are isomorphic.*

Proof. By Lemma 3.1 any finite order ICR hierarchy can be recovered continuously from any STS. By Lemma 3.2 the set $\mathcal{S}$ coincides with all ICR hierarchies. Then by Lemma 3.3, $(\mathcal{S}, \psi)$ is a STS which can be recovered from all STS. The product sigma algebra then ensures that $(\mathcal{S}, \psi)$ is a STS which is minimal. As the property of minimality is universal, every minimal STS is isomorphic to $\mathcal{S}$. □

Claim 4.1 $X \ll X' \implies B(X) \ll B(X')$.

Proof of Claim 4.1. Since $X \ll X'$ there exists a mapping $\rho\colon X \to \Delta(X')$ so that for all $x \in X$,

$$\rho(\{x' \in X' : \forall\ n,\ x_n \subseteq x'_n\}|x) = 1. \tag{A.8}$$

For every player i and $x_i \in X_i$, we thus derive a stochastic transformation $\rho_{-i,x_i}\colon K \times X_{-i} \to \Delta(X_{-i})$, which for every k, x_{-i} satisfies

$$\rho_{-i,x_i}(x'_{-i}|k, x_{-i}) \coloneqq \sum_{x'_i \in X'_i} \rho(x'_i, x'_{-i}|k, x_{-i}, x_i). \tag{A.9}$$

Then the result follows by applying Lemma A.2 at every coordinate of x_i. □

Claim 4.2 *The sequence* $s \in \mathcal{S}$ *is maximal after round* m *if and only if it is maximal at round* n *for all* $n > m$.

Proof. One direction is immediate: If s is maximal at round n for all $n > m$ then it must be maximal. Suppose now that s is maximal after round m. For any $n \in \mathbb{N}$, let $\bar{\mathcal{S}}^n \subseteq \mathcal{S}$ denote the set of sequences which are maximal after round n. Moreover, let $\tilde{\mathcal{S}}^n \subseteq \mathcal{S}$ denote the set of sequences which are maximal at round n. For every $\hat{s} \in \bar{\mathcal{S}}^n$, define the set of sequences obtained by making the n-th entry maximal and preserving the rest of the sequence:

$$\hat{\mathcal{S}}^n(\hat{s}) \coloneqq \{(\hat{s}^{n-1}, \tilde{s}_n, \tilde{s}_{n+1}, \dots) : \tilde{s} \in \tilde{\mathcal{S}}^{n+1} \text{ s.t. } \tilde{s}^n = \hat{s}^n\}. \tag{A.10}$$

Note that for every $\bar{s} \in \mathcal{S}$ and any $\hat{s} \in \hat{\mathcal{S}}^n(\bar{s})$ we have that $\hat{s}^n \in \mathcal{S}^n$. By Claim 4.1 we conclude that

$$B(\mathcal{S}) \ll B(\bar{\mathcal{S}}^n) \ll B(\cup_{\bar{s}\in\bar{\mathcal{S}}^n}\hat{\mathcal{S}}^n(\bar{s})) \tag{A.11}$$

Since $B(\mathcal{S}) = \mathcal{S}$, we conclude that $\mathcal{S} \ll B(\cup_{\bar{s}\in\bar{\mathcal{S}}^n}\hat{\mathcal{S}}^n(\bar{s}))$. Moreover, by construction of $\mathcal{S}$ we must have that for all $\hat{s} \in B(\cup_{\bar{s}\in\bar{\mathcal{S}}^n}\hat{\mathcal{S}}^n(\bar{s}))$, $\hat{s}^{n+1} \in \mathcal{S}^{n+1}$. Applying this argument for all $n' \geq n$ implies that $\hat{\mathcal{S}}^n(\hat{s}) = \bar{\mathcal{S}}^n(\hat{s})$ and so $\bar{\mathcal{S}}^n = \bigcap_{l>n}\tilde{\mathcal{S}}^l$, which concludes the proof. □

Claim 4.3 *There exists $L \in \mathbb{N}$ so that for every $m \in \mathbb{N}$ and every $s \in \mathcal{S}$,*

$$|\bar{\mathcal{S}}^m(s^m)| < L. \tag{A.12}$$

Proof. Since the sequences in $\mathcal{S}$ are weakly decreasing in the set inclusion order, we conclude from the sequential maximality property established in Claim 4.2 that for any sequence $s \in \bar{\mathcal{S}}^m$ and $\bar{s} \in \bar{\mathcal{S}}^m(s^m)$,

$$\{s_n : n \in \mathbb{N}\} = \{\bar{s}_n : n \in \mathbb{N}\} \implies s = \bar{s}. \tag{A.13}$$

Hence the result. □

Claim 4.4 *There exists M so that for every $m \in \mathbb{N}$ and $s \in \mathcal{S}$, the set $\bar{\mathcal{S}}^m(s^m)$ converged at round $m + M$.*

Proof. We proceed inductively on $m \in \mathbb{N}$. We start with the base case: $\bar{\mathcal{S}}^0(s^0)$. By Claim 4.3 there is some finite M so that $\bar{\mathcal{S}}^0(s^0)$ has converged at round M. Suppose now that $\bar{\mathcal{S}}^{m-1}(s^{m-1})$ has converged at round $m - 1 + M$ for every s^{m-1}. It follows from the definition of B that $B(\bar{\mathcal{S}}^{m-1})$ has converged at round $m + M$. By maximality of $\bar{\mathcal{S}}^m$ and monotonicity of B, we deduce that $\bar{\mathcal{S}}^m(s^m) \subseteq B(\bar{\mathcal{S}}^{m-1})$, and so the result follows. □

Claim 4.5 *For every $m \geq 0, n > 0$, $\bar{T}^{m,n}$ is the set of sequences s that are maximal after round n, admit at most $m - 1$ non-maximal transitions before round n, and such that:*

$$s \in B(\bar{T}^{m,n-1}).$$

Proof. We already established that $\bar{\mathcal{S}}^m(s^m) \subseteq B(\bar{\mathcal{S}}^{m-1})$. We deduce from the monotonicity of B that every sequence that makes $m-1$ non-maximal transitions before round n is a best-reply to beliefs supported on sequences that make at most $m-1$ non-maximal transitions before round $n-1$. Given $\bar{T}^{m,0}$ and $\bar{T}^{m,n-1}$, we thus have that

$$\bar{T}^{m,n} := \bigcup_{s \in B(\bar{T}^{m,n-1}): s^{n-1} \in (\bar{T}^{m,0})^{n-1}} \bar{\mathcal{S}}^n(s^n), \tag{A.14}$$

where $(\bar{T}^{m,0})^{n-1} := \{s^{n-1} : s \in \bar{T}^{m,0}\}$. □

Claim 4.6 *There exists $N \in \mathbb{N}$ so that*

$$\bigcup_{m=1}^{N} \bar{T}^{m,0} = \mathcal{S}. \tag{A.15}$$

Proof. We will argue that for the infinite sequence of sets $(\bar{T}^{m,0})_{m\in\mathbb{N}}$ obtained from the recursive construction above, there is $N \leq |\mathcal{B}|$ so that

$$\mathcal{T}^N := \bigcup_{m=1}^{N} \bar{T}^{m,0} = \mathcal{S}. \tag{A.16}$$

Since $\bar{T}^{1,0} \subseteq \mathcal{S}$, we conclude that $\mathcal{T}^w \subseteq \mathcal{S}$ for all w. Fix any $s \in \mathcal{S}$. Note that s makes at most $|\mathcal{B}|$ many non-maximal transitions: The set of rounds $m \in \mathbb{N}$ so that $s \notin \tilde{\mathcal{S}}^m(s^m)$ is at most $|\mathcal{B}|$, where $\tilde{\mathcal{S}}^m(s^m)$ is the set of sequences $\tilde{s} \in \mathcal{S}$ which are maximal at round m and satisfy $\tilde{s}^{m-1} = s^{m-1}$. We now show by induction on n that

$$s^n \in \{\bar{s}^n : \bar{s} \in \mathcal{T}^w\}, \tag{A.17}$$

for every $w \geq |\mathcal{B}|$. Note that by construction of $\mathcal{S}$, condition (A.17) holds for $n = 0$. Suppose then that we have shown condition (A.17) for all $s \in \mathcal{S}$ and some $n = n' - 1$. Then there exists $m \leq |\mathcal{B}|$ so that

$$s^{n'-1} \in \{\bar{s}^{n'-1} : \bar{s} \in \bar{T}^{m,n'-1}\}. \tag{A.18}$$

By the construction of $\mathcal{S}$, we conclude that

$$s^{n'} \in \{\bar{s}^{n'} : \bar{s} \in B(\bar{T}^{m,n'-1})\}, \tag{A.19}$$

and so there is $m' \leq |\mathcal{B}|$ so that

$$s^{n'} \in \{\bar{s}^{n'} : \bar{s} \in \bar{T}^{m',n'}\}. \tag{A.20}$$

We deduce that for all $w \geq |\mathcal{B}|$ we must have that $s \in \mathcal{T}^w$ and so the result follows. □

Claim 4.7 *Let N satisfy (A.15). For every $m < N$ there exists $z_m \in \mathbb{N}$ and $n_m \in \mathbb{N}$ so that for all $s \in \bar{T}^{m+1,0}$*

$$O(s^n) = O(s^{n+z_m}), \ \forall \ n \geq n_m. \tag{A.21}$$

Proof. By monotonicity of B we have that for every $n \in \mathbb{N}$,

$$\bar{\mathcal{S}}^n \subseteq B(\bar{\mathcal{S}}^{n-1}). \tag{A.22}$$

Hence $\bar{T}^{m,n} \subseteq B(\bar{T}^{m,n-1})$. For any two rounds $n > \tilde{n}$ so that

$$\{O(s^l) : s \in \bar{T}^{m,l}\} = \{O(s^{l+n-\tilde{n}}) : s \in \bar{T}^{m,l+n-\tilde{n}}\}, \tag{A.23}$$

for every $l \in \{\tilde{n} - M_m, \ldots, \tilde{n}\}$, we also have that

$$\{O(s^{\tilde{n}+1}) : s \in \bar{T}^{m,\tilde{n}+1}\} = \{O(s^{n+1}) : s \in \bar{T}^{m,n+1}\}. \tag{A.24}$$

We conclude from both corollaries that for every $m \in \mathbb{N}$, there exists $n_m \in \mathbb{N}$ and $\tilde{n}_m < n_m$ so that (A.23) holds. Hence the result follows. □

Theorem 4.1 *Ω is a finite set.*

Proof. Follows from Claims 4.6 and 4.7. □