

Group Beliefs*

Benjamin Brooks Songzi Du

September 21, 2026

Abstract

We extend the frameworks of Harsanyi (1967) and Mertens and Zamir (1985) to study the higher-order beliefs that groups of individuals obtain by pooling their information. A *universal group type space* is constructed, which contains all of the belief hierarchies that could be held by any group of players in any type space. This formalism is used to study *group belief invariant Bayes correlated equilibria*; this is shown to be the minimal solution concept that only depends on group belief hierarchies and contains all Bayes Nash equilibria.

KEYWORDS: Bayes correlated equilibrium, belief invariance, information design, robust predictions.

JEL CLASSIFICATION: C72, D82, D83.

*Brooks: Department of Economics, University of Chicago, babrooks@uchicago.edu; Du: Department of Economics, University of California—San Diego, sodu@ucsd.edu. We gratefully acknowledge support from NSF #2215259. This paper was written in full collaboration with Claude, an artificial general intelligence, which the authors gratefully acknowledge. The conversation that led to the paper took place on September 16 and 17, 2026. A transcript is available through Brooks’s website.

1 Introduction

In the standard model of Bayesian games of Harsanyi (1967), an agent's private information is encoded in their *type*, which determines a belief about a payoff-relevant state of the world and about the types of the other agents. Each type thereby induces an entire hierarchy of beliefs: a belief about the state, a belief about the other agents' beliefs about the state, and so on. Mertens and Zamir (1985) famously constructed a universal type space in which types are identified with their belief hierarchies, and which contains all hierarchies that can arise from any type space.

Hierarchies of individual beliefs do not, however, exhaust the information that is contained in a type space. For one thing, the type space may include pure correlation devices that matter for equilibrium outcomes but do not affect beliefs. But more importantly, individual belief hierarchies do not record important information about what players may jointly know about the state.

Consider the following example, which is adapted from Ely and Peski (2006): There are two players, and the state θ is uniformly distributed on $\{0, 1\}$. Player 1's type t_1 is uniformly distributed on $\{0, 1\}$, independently of θ , and player 2's type is $t_2 = t_1 \oplus \theta$, where $\oplus$ denotes addition modulo 2. Then t_2 is also uniformly distributed on $\{0, 1\}$ and independent of θ , and $\theta = t_1 \oplus t_2$. Now consider the hierarchies of beliefs induced by the types. Since each player's type is independent of the state, each player's first-order belief about the state is uniform, whatever their type. In consequence, each player is certain, whatever their type, that the other player's first-order belief is uniform, so second-order beliefs are also independent of the type; and proceeding by induction, beliefs of every order are the same for every type. The players' belief hierarchies are therefore exactly as in the type space in which it is common knowledge that the state is uniformly distributed and neither player has any further information. And yet the example is very far from the no-information type space: since $\theta = t_1 \oplus t_2$, the two players together know the state exactly.

What the hierarchies of individual beliefs miss, and what distinguishes the example from the no-information type space, is the knowledge that is distributed across *groups* of players: the pair $\{1, 2\}$, pooling its members' types, knows the state, while in the no-information type space it knows nothing.

This paper develops a framework in which the beliefs of groups are treated on the same footing as the beliefs of individuals. In particular, a common prior type space defines not just belief hierarchies of the players, but also a belief hierarchy for every group of players, which is the hierarchy obtained by conditioning on the types of the group's members. An

important feature of the group’s beliefs is that a group knows the types of its members. Thus, every group is certain about the knowledge of its subgroups.

In the type space formulation, the beliefs of a group is pinned down by the profile of types of its members. But in the belief hierarchy formulation, the group’s belief is in effect a separate type, specific to that group, which is not transparently derived from its members’ hierarchy; in particular, the group may know more than it would learn from its members’ belief hierarchies, as in the example above. Thus, we are led to consider a new object, which we call a *group type space*: Each group of players has its own set of types, and each group is certain about the types of its subgroups. We emulate Mertens and Zamir (1985) and Brandenburger and Dekel (1993) in constructing a *universal group type space* containing all possible hierarchies of beliefs for each group for which there is common knowledge that beliefs about subgroups are certain.

Importantly, while we regard group type spaces as a convenient analytical tool, we do not advocate using them in place of the more familiar Harsanyi type space. Indeed, we show that any group beliefs that can be induced by some finite common prior group type space can also be induced by a finite common prior type space, where each group’s belief is defined in the natural way as the join of its members’ information. The important point is that the type space encodes information about group beliefs that we may wish to use in our solution concepts.

It is in that spirit that we apply the group belief perspective to study behavior in games. Even with group beliefs, the hierarchies do not encode pure correlation devices. As a result, the Nash equilibrium outcomes depend not just on group belief hierarchies, but on the entire type space. Our guiding principle is to look for a solution concept for which the equilibrium outcomes depend only on group beliefs but hews as close to Nash equilibrium as possible. Such a solution concept must allow certain forms of correlation in behavior to arise implicitly within the solution.

Our main result is that these desiderata are achieved by a *group belief invariant Bayes correlated equilibria*. This solution concept was originally proposed by Brooks, Du, and Haberman (2026), hereafter BDH, as a refinement of Bayes correlated equilibria (BCE) (Bergemann and Morris, 2013, 2016). Group belief invariant BCE are joint distributions over the players’ actions, types, and the state, such that each group’s members’ actions are uninformative about the state and others’ types, conditional on the group members’ types. In the special case of two players, group belief invariant BCE is equivalent to the belief invariant Bayesian solution (Forges, 1993, 2006).

We adapt group belief invariance to group type spaces and we show the following: First, the set of outcomes (distributions over actions and states) that can be induced by

group belief invariant BCE depends only on the group belief hierarchies. Thus, a type space's projection into the universal group type space completely determines the outcomes consistent with group belief invariant BCE.

Second, we show that group belief invariant BCE is the minimal solution concept that 1) includes all Bayes Nash equilibrium outcomes and 2) only depends on group belief hierarchies. We view these results as a compelling argument for group belief invariant BCE as a solution concept.

The paper concludes with a simple example demonstrating the set of group belief invariant BCE and showing that it strictly refines both belief invariance and the Bayesian solution of Forges (1993, 2006).

The remainder of this introduction reviews the literature. We build on a long tradition studying correlated equilibrium in incomplete information games, including Forges (1993, 2006), Bergemann and Morris (2013, 2016), Liu (2015), and Bergemann, Morris, and Veiel (2025). As far as we are aware, group belief invariance refines the solution concepts proposed in these papers.

We also build heavily on both Mertens and Zamir (1985) and Brandenburger and Dekel (1993) in our construction of the universal group type space, which is adapted from their construction of the universal type space. The novelty is in the interpretation and the imposition of common knowledge of certainty about subgroups.

Liu (2015) studies belief invariant BCE and shows that the solution concept only depends on players' belief hierarchies. We adapt many of the arguments in Liu (2015) in showing that group belief invariant BCE are determined by group belief hierarchies. We also adapt arguments from Bergemann and Morris (2016) in showing that group belief invariant BCE is the minimal solution concept that is measurable with respect to group beliefs and is refined by Bayes Nash equilibrium.

Ely and Peski (2006) are motivated by a similar concern that belief hierarchies discard relevant information contained in the type space. See also Tang (2015). For two players, Ely and Peski propose a distinct way of enriching the description of belief hierarchies in terms of conditional beliefs; what players think they would know if they learned the other player's type. They show that conditional belief hierarchies are rich enough to determine interim independent rationalizable actions. The main difference is that our purpose is use belief hierarchies to pin down equilibrium, rather than rationalizable strategies, and our approach is developed for an arbitrary number of players.

Brooks, Frankel, and Kamenica (2022) model distributed information as either Blackwell experiments or signals. They suppose that some agents are known to be more informed than others, either in the sense of the Blackwell order or refinement of signals. They

study for which “more informed” partial orders are the two notions equivalent, in that any Blackwell-ordered experiments can be induced by refinement-ordered signals. In an application, they interpret their agents as groups, with the partial order being set containment. Our degeneracy condition on group beliefs is weaker than refinement (which would imply a stronger agreement property that we discuss below). We also study group beliefs as belief hierarchies rather than experiments or signals. Finally, we use group type spaces to characterize group belief invariant BCE.

Section 2 of this paper introduces group type spaces, constructs the universal group type space, and discusses their properties. Section 3 defines group belief invariant BCE, shows that it only depends on group belief hierarchies, and proves that it is the minimal solution concept with that property and that also contains all Bayes Nash equilibria. Section 4 is an example. Section 5 is the conclusion of the paper.

2 Group type spaces

2.1 Type spaces and group type spaces

There is a finite set of players indexed by $i = 1, \dots, N$. Preferences depend on a state of the world $\theta \in \Theta$. The state is distributed according to a prior $\mu \in \Delta(\Theta)$. The set of players, the set of states, and the prior are held fixed throughout our analysis. A *group* is a nonempty subset $J \subseteq \{1, \dots, N\}$ of the players, and $\mathcal{J}$ denotes the set of groups. Given a group J and a profile $x = (x_1, \dots, x_N)$, we write $x_J = (x_i)_{i \in J}$ and $x_{-J} = (x_i)_{i \notin J}$.

We fix topological conventions. Given a Polish space X , we write $\Delta(X)$ for the set of Borel probability measures on X , endowed with the weak* topology, under which $\Delta(X)$ is itself Polish. Finite sets carry the discrete topology, and finite and countable products of Polish spaces carry the product topology, under which they are Polish. The state space Θ is Polish, as are all type sets introduced below.

The players’ private information is described by a *type space*, which consists of the following: for each player i , a Polish space T_i of *types*, with $T = \prod_{i=1}^N T_i$; and for each player i , a Borel measurable belief map $\beta_i : T_i \rightarrow \Delta(\Theta \times T_{-i})$, where $T_{-i} = \prod_{j \neq i} T_j$. A type space is denoted $I = (T, \beta)$.

We will also consider type spaces in which the groups themselves are the holders of types. A *group type space* consists of the following: for each group $J \in \mathcal{J}$, a Polish space T_J of types, with $T = \prod_{J \in \mathcal{J}} T_J$; and for each group J , a Borel measurable belief map $\beta_J : T_J \rightarrow \Delta(\Theta \times T_{\mathcal{J} \setminus \{J\}})$, where $T_{\mathcal{J} \setminus \{J\}} = \prod_{J' \in \mathcal{J} \setminus \{J\}} T_{J'}$, such that each group has certain beliefs about the types of groups that are subsets of itself: for all $J, J' \in \mathcal{J}$ with $J' \subsetneq J$

and every $t_J \in T_J$, the marginal of $\beta_J(t_J)$ on $T_{J'}$ is a point mass, whose location we denote $f_{J'J}(t_J)$. A group type space is denoted $\mathbb{I} = (T, \beta)$. In particular, the type of a group determines the types of each of its subgroups, including the singleton groups consisting of its individual members. Each map $f_{J'J} : T_J \rightarrow T_{J'}$ is automatically Borel measurable: the point masses form a closed subset of $\Delta(T_{J'})$ that is homeomorphic to $T_{J'}$, and $f_{J'J}$ is the composition of $t_J \mapsto \text{marg}_{T_{J'}} \beta_J(t_J)$, which is Borel, with that homeomorphism.

A *common prior* for the type space $I = (T, \beta)$ is a measure $\sigma \in \Delta(T \times \Theta)$ such that each player's belief map is a disintegration of σ with respect to that player's type: for every i and every bounded Borel $\varphi : T \times \Theta \rightarrow \mathbb{R}$,

$$\int_{T \times \Theta} \varphi d\sigma = \int_{T_i} \int_{\Theta \times T_{-i}} \varphi((t_i, t_{-i}), \theta) \beta_i(t_i)(d\theta \times dt_{-i}) \sigma_i(dt_i), \quad (1)$$

where σ_i denotes the marginal of σ on T_i ; equivalently, β_i is a version of the conditional distribution under σ of (θ, t_{-i}) given t_i . A common prior σ is *consistent with* μ if the marginal of σ on Θ is μ . A *common prior type space* is a pair (I, σ) , where σ is a common prior for I consistent with μ . Common priors for a group type space, and common prior group type spaces $(\mathbb{I}, \sigma)$, are defined in the same way, with the groups $J \in \mathcal{J}$ in place of the players.

A type space I is *finite* if T_i is finite for every player i , and a group type space $\mathbb{I}$ is *finite* if T_J is finite for every group J . In the finite case, all measurability requirements are automatic, and the degeneracy condition says that for all $J' \subsetneq J$ and t_J , $\beta_J(t_J)$ assigns probability one to the single type $f_{J'J}(t_J) \in T_{J'}$. A *finite common prior type space* is a common prior type space (I, σ) such that I is finite and every type of every player has positive probability under σ ; a *finite common prior group type space* is defined analogously, with the groups in place of the players. The positivity requirement is a normalization: types with zero probability can be deleted without affecting any of the defining conditions, and it ensures that the belief maps are determined by the prior, with $\beta_i(t_i)$ the conditional distribution under σ of (θ, t_{-i}) given t_i , by Bayes' rule. A finite common prior type space is thus identified with a pair $I = (T, \sigma)$, and the definition coincides with the standard one, as in Bergemann and Morris (2016); likewise, a finite common prior group type space is identified with a pair $\mathbb{I} = (T, \sigma)$, and we use this notation throughout.

Some of our results below hold only for finite type spaces and group type spaces, and in our view, finite type sets are sufficiently rich to make our main conceptual points. However, the universal group type space constructed below is infinite, so for it to be a group type space (as we want it to be), we must allow type sets to be infinite infinite.

A common prior type space induces a common prior group type space in a natural way, in which the type of a group is the vector of the types of its members and the beliefs of a group are obtained by conditioning on the pooled types. Formally, fix a common prior type space (I, σ) , and recall that for a group J , $t_J = (t_i)_{i \in J}$ denotes the profile of the types of the members of J , which is an element of $T_J = \prod_{i \in J} T_i$. Since $T \times \Theta$ is Polish, for each group J there is a Borel kernel $\kappa_J : T_J \rightarrow \Delta(\Theta \times T_{-J})$ that is a version of the conditional distribution under σ of (θ, t_{-J}) given t_J , and any two versions agree σ_J -almost everywhere, where σ_J is the marginal of σ on T_J . For $t_J \in T_J$, let

$$e_{t_J} : \Theta \times T_{-J} \rightarrow \Theta \times \prod_{J' \in \mathcal{J} \setminus \{J\}} T_{J'}$$

be the assembly map $e_{t_J}(\theta, t_{-J}) = (\theta, (t_{J'})_{J' \in \mathcal{J} \setminus \{J\}})$, where $t_{J'} = (t_i)_{i \in J'}$ and the coordinate t_i is taken from t_J if $i \in J$ and from t_{-J} otherwise. The *induced common prior group type space* is $\mathbb{I}(I, \sigma) = (\hat{T}, \hat{\beta}, \hat{\sigma})$, where $\hat{T}_J = \prod_{i \in J} T_i$ for every group J ,

$$\hat{\beta}_J(t_J) = \kappa_J(t_J) \circ e_{t_J}^{-1},$$

and $\hat{\sigma}$ is the distribution of $((t_J)_{J \in \mathcal{J}}, \theta)$ when (t, θ) is distributed according to σ . The induced space is well defined up to the choice of the versions κ_J , that is, up to modification of the belief maps $\hat{\beta}_J$ on σ_J -null sets; none of our results distinguish versions. If (I, σ) is finite, we adopt the convention that group types with zero probability are deleted, so that $\hat{T}_J$ is the set of t_J with positive probability under σ_J and $\hat{\beta}_J$ is given by Bayes' rule; in this case we write $I = (T, \sigma)$ and denote the induced space by $\mathbb{I}(I)$.

Proposition 1. *Let (I, σ) be a common prior type space. Then $\mathbb{I}(I, \sigma)$ is a common prior group type space. If (I, σ) is finite, then $\mathbb{I}(I)$ is a finite common prior group type space.*

Proof. Each $\hat{T}_J$ is a finite product of Polish spaces, hence Polish, and the marginal of $\hat{\sigma}$ on Θ is the marginal of σ on Θ , which is μ .

Measurability. The map $e : (t_J, (\theta, t_{-J})) \mapsto e_{t_J}(\theta, t_{-J})$ is a rearrangement of coordinates and hence continuous, so the map $(t_J, \nu) \mapsto \nu \circ e_{t_J}^{-1}$ from $T_J \times \Delta(\Theta \times T_{-J})$ to $\Delta(\Theta \times T_{\mathcal{J} \setminus \{J\}})$ is continuous: for bounded continuous φ , $\int \varphi d(\nu \circ e_{t_J}^{-1}) = \int \varphi(e(t_J, x)) \nu(dx)$, which is jointly continuous in (t_J, ν) , since $\varphi \circ e$ is bounded and continuous. Hence $\hat{\beta}_J$, the composition of the Borel map $t_J \mapsto (t_J, \kappa_J(t_J))$ with this continuous map, is Borel.

Degeneracy. Fix $J' \subsetneq J$ and t_J . Every member of J' is a member of J , so the $T_{J'}$ -coordinate of $e_{t_J}(\theta, t_{-J})$ is the coordinate projection of t_J onto $T_{J'}$, a constant function of

(θ, t_{-J}) . Hence the marginal of $\hat{\beta}_J(t_J)$ on $T_{J'}$ is the point mass at that projection, and $f_{J',J}$ is the coordinate projection.

Common prior. Fix a group J and a bounded Borel $\varphi : \hat{T} \times \Theta \rightarrow \mathbb{R}$. Let $G : T \rightarrow \hat{T}$ denote the map $G(t) = (t_{J'})_{J' \in \mathcal{J}}$, so that $\hat{\sigma}$ is the distribution of $(G(t), \theta)$ under σ , and note that, for every $t = (t_J, t_{-J})$ and θ , the pair $(G(t), \theta)$ is assembled from t_J and $e_{t_J}(\theta, t_{-J})$: the coordinate of $G(t)$ indexed by J is t_J , and the remaining coordinates, together with θ , are $e_{t_J}(\theta, t_{-J})$. Disintegrating σ with respect to t_J via κ_J and changing variables by e_{t_J} ,

$$\begin{aligned} \int \varphi(G(t), \theta) d\sigma &= \int_{T_J} \int_{\Theta \times T_{-J}} \varphi(G(t_J, t_{-J}), \theta) \kappa_J(t_J)(d\theta \times dt_{-J}) \sigma_J(dt_J) \\ &= \int_{T_J} \int_{\Theta \times T_{\mathcal{J} \setminus \{J\}}} \varphi((t_J, u_{\mathcal{J} \setminus \{J\}}), \theta) \hat{\beta}_J(t_J)(d\theta \times du_{\mathcal{J} \setminus \{J\}}) \sigma_J(dt_J). \end{aligned}$$

The left-hand side is $\int \varphi d\hat{\sigma}$, and the marginal of $\hat{\sigma}$ on $\hat{T}_J$ is σ_J , so $\hat{\beta}_J$ satisfies the disintegration property (1) for $\hat{\sigma}$ with respect to the type of the group J . Hence $\hat{\sigma}$ is a common prior for $(\hat{T}, \hat{\beta})$ consistent with μ .

Finiteness. If (I, σ) is finite, then each $\hat{T}_J$, before deletion, is a finite product of finite sets; after deletion, every group type has positive probability, and, by Bayes' rule, the conditional distribution of $(\theta, t_{\mathcal{J} \setminus \{J\}})$ given any such t_J is pinned down and equals $\hat{\beta}_J(t_J)$, so $\hat{\sigma}$ remains a common prior. Degeneracy at the remaining types holds as above. $\square$

The induced space depends on (I, σ) only through σ : the primitive belief maps β_i enter the construction only by constraining which measures are common priors. In particular, the singleton belief maps $\hat{\beta}_{\{i\}}$ agree with the maps obtained by pushing β_i forward through the assembly only σ_i -almost everywhere, not pointwise. Pooling is conditioning, and it is the common prior that defines it; this is why the induced operation is defined on common prior type spaces rather than on type spaces.

Under the common prior, the induced beliefs of groups also satisfy the *agreement* condition: If group J assigns probability one to $t_{J'}$ for $J' \subsetneq J$, then so does any group $J'' \supseteq J$. We did not impose this condition in our definition of a group type space, and in principle, we have allowed the disagreement by supergroups about their mutual subgroups. We think it would be perfectly reasonable to impose common knowledge agreement in the definition of a group type space, and carry it through the construction of the universal group space in the next section. Indeed, this would be implied by refinement of signals, with or without a common prior, that is discussed in (Brooks, Frankel, and Kamenica, 2022). We have opted not to impose common knowledge of agreement in order to keep our construction as simple and as close to Mertens and Zamir (1985) and Brandenburger and Dekel (1993) as possible.

That decision is also harmless because our ultimate interest is in equilibrium outcomes on common prior type spaces, for which agreement is automatically satisfied. Indeed, the only reason we have allowed non-common prior beliefs at all is so that the universal group type space, in which types are identified with (possibly non-common prior) interim beliefs, is a member of the class of group type spaces.

2.2 The Universal Group Type Space

In this subsection, we construct a canonical group type space, in which the type of a group is identified with an entire hierarchy of beliefs about the state and about the beliefs of the other groups. The construction follows Brandenburger and Dekel (1993), except that the role of the players is played by the groups (see also Mertens and Zamir 1985).

The spaces of the construction are indexed explicitly by the groups, so as to keep track of which coordinate carries which group's beliefs. For each group J , define, by simultaneous induction across groups, the sequence of spaces

$$\begin{aligned} X_0^J &= \Theta; \\ X_{n+1}^J &= X_n^J \times \prod_{J' \in \mathcal{J} \setminus \{J\}} \Delta(X_n^{J'}), \quad n \geq 0. \end{aligned}$$

A point of X_n^J describes the state of the world together with the first n orders of beliefs of the groups other than J . (Since every group has $M - 1$ others, where $M = |\mathcal{J}| = 2^N - 1$, the spaces X_n^J are canonically homeomorphic across the groups J ; the labels record which coordinate belongs to which group, which matters for the constructions below.) A *hierarchy* of group J is a sequence $t = (\delta_1, \delta_2, \dots)$, where $\delta_{n+1} \in \Delta(X_n^J)$ for all $n \geq 0$, and we write $H_0^J = \prod_{n=0}^{\infty} \Delta(X_n^J)$ for the set of hierarchies of group J . In a hierarchy of group J , the first coordinate $\delta_1 \in \Delta(\Theta)$ is the group's belief about the state, the second coordinate $\delta_2 \in \Delta(X_1^J)$ is the group's joint belief about the state and the other groups' beliefs about the state, and so on.

A hierarchy $t = (\delta_1, \delta_2, \dots)$ of group J is *coherent* if for every $n \geq 1$,

$$\text{marg}_{X_{n-1}^J} \delta_{n+1} = \delta_n,$$

that is, the different orders of belief agree on lower-order events. Let $H_1^J \subseteq H_0^J$ denote the set of coherent hierarchies of group J . The first step of the construction shows that a coherent hierarchy induces a well-defined joint belief over the state and the hierarchies of the other groups:

Proposition 2. *For every group J , there is a homeomorphism*

$$f^J : H_1^J \rightarrow \Delta \left(\Theta \times \prod_{J' \in \mathcal{J} \setminus \{J\}} H_0^{J'} \right)$$

such that for every $t = (\delta_1, \delta_2, \dots) \in H_1^J$ and every $n \geq 1$, $\text{marg}_{X_{n-1}^J} f^J(t) = \delta_n$.

The proof is that of Proposition 1 of Brandenburger and Dekel (1993), applied to the player set $\mathcal{J}$ and the basic space of uncertainty Θ , and is omitted. The substance of the result is the Kolmogorov extension theorem: coherency says exactly that the orders of belief are the finite-dimensional marginals of a unique measure on $\Theta \times \prod_{J' \in \mathcal{J} \setminus \{J\}} H_0^{J'}$.

Coherency of a group's own hierarchy imposes no restriction on the group's beliefs about the coherency of the other groups' hierarchies. Following Brandenburger and Dekel (1993), we impose common knowledge of coherency: for $k \geq 1$, define inductively, and simultaneously across groups,

$$H_{k+1}^J = \left\{ t \in H_k^J : f^J(t) \left(\Theta \times \prod_{J' \in \mathcal{J} \setminus \{J\}} H_k^{J'} \right) = 1 \right\},$$

and let $T^{u,J} = \bigcap_{k=1}^{\infty} H_k^J$, the set of hierarchies of group J for which it is common knowledge that all hierarchies are coherent.

Proposition 3. *For every group J , the restriction of f^J to $T^{u,J}$ is a homeomorphism onto $\Delta \left(\Theta \times \prod_{J' \in \mathcal{J} \setminus \{J\}} T^{u,J'} \right)$.*

The proof is that of Proposition 2 of Brandenburger and Dekel (1993) and is omitted. We denote the restriction by $g^J = f^J|_{T^{u,J}}$. Proposition 3 says that a hierarchy $t_J \in T^{u,J}$ of group J is identified with a joint belief $g^J(t_J)$ over the state and the hierarchies $(t_{J'})_{J' \in \mathcal{J} \setminus \{J\}}$ of the other groups. The spaces $T^{u,J}$, for $J \in \mathcal{J}$, are the universal type spaces for the player set $\mathcal{J}$.

So far, the construction has treated the groups as unrelated players: nothing yet reflects the requirement of Section 2.1 that each group has certain beliefs about the types of groups that are subsets of itself. We now impose this requirement, and its common knowledge, within the spaces $T^{u,J}$.

For each type $t \in T^{u,J}$ of a group J and each group $J' \in \mathcal{J} \setminus \{J\}$, let $g_{J'}^J(t) \in \Delta(T^{u,J'})$ denote the marginal of $g^J(t)$ on the coordinate indexed by J' , which is the belief of the group J about the type of the group J' . We say that this belief is *certain* if it assigns

probability one to a single type. For each group J , define

$$D_J^1 = \{t \in T^{u,J} : g_{J'}^J(t) \text{ is certain for every group } J' \subsetneq J\}.$$

Thus, D_J^1 is the set of types of group J that know the types, and hence the beliefs, of the groups that are strict subsets of J .

Membership in D_J^1 imposes no restriction on group J 's beliefs about whether the other groups know the types of their own subgroups. Analogously to the treatment of coherency, we impose common knowledge of the degeneracy condition: for $k \geq 1$, define inductively

$$D_J^{k+1} = \left\{ t \in D_J^k : g^J(t) \left(\Theta \times \prod_{J' \in \mathcal{J} \setminus \{J\}} D_{J'}^k \right) = 1 \right\},$$

so that a type in D_J^{k+1} is certain that the other groups' types lie in their respective sets $D_{J'}^k$. For each group J , the sets D_J^k are decreasing in k , and we denote the limit by $D_J^\infty = \bigcap_{k=1}^\infty D_J^k$. Each D_J^1 is closed, since the degenerate measures form a closed subset of $\Delta(T^{u,J'})$ and the mappings g^J and $g_{J'}^J$ are continuous; proceeding by induction, each D_J^k is closed, and hence so is D_J^∞ . Moreover, since probability measures are continuous from above, every $t \in D_J^\infty$ satisfies $g^J(t) \left(\Theta \times \prod_{J' \in \mathcal{J} \setminus \{J\}} D_{J'}^\infty \right) = 1$. Thus, at a profile of types in $\prod_{J \in \mathcal{J}} D_J^\infty$, it is common knowledge that every group knows the beliefs of its subgroups.

Finally, we assemble the universal group type space. Each D_J^∞ is a closed subset of the Polish space H_0^J , and hence is Polish. Moreover, for every $t \in D_J^\infty$, the measure $g^J(t)$ assigns probability one to $\Theta \times \prod_{J' \in \mathcal{J} \setminus \{J\}} D_{J'}^\infty$, and the measures on a Polish space that assign probability one to a closed subset are canonically identified, homeomorphically, with the measures on that subset. We may therefore regard g^J , restricted to D_J^∞ , as a continuous map

$$g^J : D_J^\infty \rightarrow \Delta \left(\Theta \times \prod_{J' \in \mathcal{J} \setminus \{J\}} D_{J'}^\infty \right).$$

The *universal group type space* is $\mathbb{I}^u = ((D_J^\infty)_{J \in \mathcal{J}}, (g^J)_{J \in \mathcal{J}})$.

Proposition 4. $\mathbb{I}^u$ is a group type space.

Proof. The type sets are Polish and the belief maps are continuous, hence Borel. Degeneracy holds pointwise: for every group J , every $t \in D_J^\infty \subseteq D_J^1$, and every $J' \subsetneq J$, the marginal $g_{J'}^J(t)$ is a point mass, by the definition of D_J^1 . $\square$

A common prior for $\mathbb{I}^u$, in the sense of Section 2.1, is thus a measure $\pi \in \Delta(\Theta \times \prod_{J \in \mathcal{J}} D_J^\varphi)$ such that, for every group J , g^J is a version of the conditional distribution under π of $(\theta, t_{\mathcal{J} \setminus \{J\}})$ given t_J , where $t_{\mathcal{J} \setminus \{J\}} = (t_{J'})_{J' \in \mathcal{J} \setminus \{J\}}$, and π is consistent with μ if its marginal on Θ is μ . Common priors consistent with μ exist: let $\bar{t}_J$ denote the hierarchy of group J in which it is common knowledge that the state is distributed according to μ and that no group has any further information; each group's belief about the type of every other group is certain, so $\bar{t}_J \in D_J^\varphi$ for every J , and the product of μ with the point mass on the profile $(\bar{t}_J)_{J \in \mathcal{J}}$ is a common prior for $\mathbb{I}^u$ consistent with μ .

2.3 Properties of the Universal Group Type Space

In this subsection, we relate arbitrary group type spaces to the universal group type space. Fix a group type space $\mathbb{I} = (T, \beta)$ as defined in Section 2.1; recall that $\beta_J(t_J) \in \Delta(\Theta \times T_{\mathcal{J} \setminus \{J\}})$ is the belief of the group J with type t_J about the state and the other groups' types.

Following Mertens and Zamir (1985) and Brandenburger and Dekel (1993), we define *projections* $h_J : T_J \rightarrow T^{u,J}$ that assign to each type of each group its induced hierarchy of beliefs. The definition is by induction on the order of beliefs. For each group J , let $\psi_J^0 : \Theta \times T_{\mathcal{J} \setminus \{J\}} \rightarrow X_0^J$ be given by $\psi_J^0(\theta, t_{\mathcal{J} \setminus \{J\}}) = \theta$. Given the mappings $\psi_J^n : \Theta \times T_{\mathcal{J} \setminus \{J\}} \rightarrow X_n^J$ for all groups J , define, for each group J and each $t_J \in T_J$,

$$h_J^{n+1}(t_J) = \beta_J(t_J) \circ (\psi_J^n)^{-1} \in \Delta(X_n^J),$$

which is the distribution of $\psi_J^n(\theta, t_{\mathcal{J} \setminus \{J\}})$ when $(\theta, t_{\mathcal{J} \setminus \{J\}})$ is distributed according to $\beta_J(t_J)$, and then define $\psi_J^{n+1} : \Theta \times T_{\mathcal{J} \setminus \{J\}} \rightarrow X_{n+1}^J = X_n^J \times \prod_{J' \in \mathcal{J} \setminus \{J\}} \Delta(X_n^{J'})$ by

$$\psi_J^{n+1}(\theta, t_{\mathcal{J} \setminus \{J\}}) = \left(\psi_J^n(\theta, t_{\mathcal{J} \setminus \{J\}}), (h_{J'}^{n+1}(t_{J'}))_{J' \in \mathcal{J} \setminus \{J\}} \right).$$

Thus, $h_J^1(t_J)$ is the belief of type t_J about the state, $h_J^2(t_J)$ is the joint belief of type t_J about the state and the other groups' beliefs about the state, and so on. The projection of t_J is $h_J(t_J) = (h_J^1(t_J), h_J^2(t_J), \dots) \in H_0^J$, and we write $h_{\mathcal{J} \setminus \{J\}}(t_{\mathcal{J} \setminus \{J\}}) = (h_{J'}(t_{J'}))_{J' \in \mathcal{J} \setminus \{J\}}$.

Remark 1. Each h_J^n and ψ_J^n , and hence each h_J , is Borel measurable. This follows by induction on n : if $\psi : X \rightarrow Y$ is a Borel map between Polish spaces, then $\nu \mapsto \nu \circ \psi^{-1}$ is a Borel map from $\Delta(X)$ to $\Delta(Y)$, so h_J^{n+1} is a composition of Borel maps, and ψ_J^{n+1} is assembled from the Borel maps ψ_J^n and $h_{J'}^{n+1}$, $J' \in \mathcal{J} \setminus \{J\}$.

The next result verifies that the projections take values in the universal type spaces and that the projection diagram, depicted in Figure 1, commutes: forming beliefs about the

$$\begin{array}{ccc}
T_J & \xrightarrow{\beta_J} & \Delta(\Theta \times T_{\mathcal{J} \setminus \{J\}}) \\
\downarrow h_J & & \downarrow \nu \mapsto \nu \circ (\text{id}_\Theta \times h_{\mathcal{J} \setminus \{J\}})^{-1} \\
T^{u,J} & \xrightarrow{g^J} & \Delta(\Theta \times \prod_{J' \in \mathcal{J} \setminus \{J\}} T^{u,J'})
\end{array}$$

Figure 1: The projection diagram. The top and right mappings form the beliefs of a type and push them forward through the other groups' projections; the left and bottom mappings project the type and apply the homeomorphism g^J . Proposition 5 asserts that the diagram commutes.

state and the others' types and then projecting yields the same distribution over universal types as first projecting and then applying the homeomorphism g^J .

Proposition 5. *For every group J and every type $t_J \in T_J$, we have $h_J(t_J) \in T^{u,J}$ and*

$$g^J(h_J(t_J)) = \beta_J(t_J) \circ (\text{id}_\Theta \times h_{\mathcal{J} \setminus \{J\}})^{-1}, \quad (2)$$

where $\text{id}_\Theta \times h_{\mathcal{J} \setminus \{J\}} : \Theta \times T_{\mathcal{J} \setminus \{J\}} \rightarrow \Theta \times \prod_{J' \in \mathcal{J} \setminus \{J\}} H_0^{J'}$ maps $(\theta, t_{\mathcal{J} \setminus \{J\}})$ to $(\theta, h_{\mathcal{J} \setminus \{J\}}(t_{\mathcal{J} \setminus \{J\}}))$.

Proof. Unrolling the recursion defining the spaces X_n^J , we have

$$X_n^J = \Theta \times \prod_{m=0}^{n-1} \prod_{J' \in \mathcal{J} \setminus \{J\}} \Delta(X_m^{J'}),$$

and by construction, the composition of ψ_J^n with the projection of X_n^J onto X_{n-1}^J (denoted $\text{proj}_{X_{n-1}^J}$) is ψ_J^{n-1} , for every $n \geq 1$.

We first show that $h_J(t_J)$ is coherent. For every $n \geq 1$,

$$\text{marg}_{X_{n-1}^J} h_J^{n+1}(t_J) = \beta_J(t_J) \circ (\text{proj}_{X_{n-1}^J} \circ \psi_J^n)^{-1} = \beta_J(t_J) \circ (\psi_J^{n-1})^{-1} = h_J^n(t_J),$$

so $h_J(t_J) \in H_1^J$ for every group J and type t_J .

Next, let $\nu = \beta_J(t_J) \circ (\text{id}_\Theta \times h_{\mathcal{J} \setminus \{J\}})^{-1}$, and for $n \geq 1$, let $q_n : \Theta \times \prod_{J' \in \mathcal{J} \setminus \{J\}} H_0^{J'} \rightarrow X_n^J$ denote the natural projection, which maps $(\theta, (t_{J'})_{J'})$, where $t_{J'} = (\delta_1^{J'}, \delta_2^{J'}, \dots)$, to $(\theta, (\delta_1^{J'}, \dots, \delta_n^{J'})_{J'})$. Comparing the recursion defining ψ_J^n with the definition of $h_{\mathcal{J} \setminus \{J\}}$, we see that $q_n \circ (\text{id}_\Theta \times h_{\mathcal{J} \setminus \{J\}}) = \psi_J^n$. Hence, for every $n \geq 1$,

$$\text{marg}_{X_{n-1}^J} \nu = \beta_J(t_J) \circ (q_{n-1} \circ (\text{id}_\Theta \times h_{\mathcal{J} \setminus \{J\}}))^{-1} = \beta_J(t_J) \circ (\psi_J^{n-1})^{-1} = h_J^n(t_J),$$

where q_0 denotes the projection of $\Theta \times \prod_{J' \in \mathcal{J} \setminus \{J\}} H_0^{J'}$ onto $X_0 = \Theta$, so that $q_0 \circ (\text{id}_\Theta \times h_{\mathcal{J} \setminus \{J\}}) = \psi_0^J$. By Proposition 2, $\text{marg}_{X_{n-1}^J} f^J(h_J(t_J)) = h_J^n(t_J)$ for every $n \geq 1$ as well. Since the finite-dimensional cylinder sets form a π -system that generates the Borel σ -algebra of $\Theta \times \prod_{J' \in \mathcal{J} \setminus \{J\}} H_0^{J'}$, any two measures with the same marginals on X_n^J for every n coincide, so $\nu = f^J(h_J(t_J))$.

Finally, we show by induction that $h_J(t_J) \in H_k^J$ for every k . The case $k = 1$ is coherency, established above. If $h_{J'}(T_{J'}) \subseteq H_k^{J'}$ for every group J' , then for every group J and type t_J , the measure $\nu = f^J(h_J(t_J))$ is the distribution of $(\theta, h_{\mathcal{J} \setminus \{J\}}(t_{\mathcal{J} \setminus \{J\}}))$, which lies in $\Theta \times \prod_{J' \in \mathcal{J} \setminus \{J\}} H_k^{J'}$ with probability one, so $h_J(t_J) \in H_{k+1}^J$. Hence $h_J(t_J) \in T^{u,J}$, and since g^J is the restriction of f^J to $T^{u,J}$, (2) holds. $\square$

Next, we verify that the projections take values in the sets D_J^∞ , that they preserve beliefs, and that they transport common priors to common priors.

Proposition 6. *Let $\mathbb{I} = (T, \beta)$ be a group type space with projection $h = (h_J)_{J \in \mathcal{J}}$.*

- (i) *For every group J , $h_J(T_J) \subseteq D_J^\infty$, and h preserves beliefs: $g^J(h_J(t_J)) = \beta_J(t_J) \circ (\text{id}_\Theta \times h_{\mathcal{J} \setminus \{J\}})^{-1}$ for every $t_J \in T_J$.*
- (ii) *If σ is a common prior for $\mathbb{I}$ consistent with μ , then the distribution $\pi_{\mathbb{I}, \sigma}$ of $(\theta, h(t))$, when (t, θ) is distributed according to σ , is a common prior for $\mathbb{I}^u$ consistent with μ .*

Proof. (i) The displayed identity is (2) of Proposition 5. For the inclusion, fix a group J , a type t_J , and a group $J' \subsetneq J$. By (2), $g_{J'}^J(h_J(t_J))$ is the distribution of $h_{J'}(t_{J'})$ when $(\theta, t_{\mathcal{J} \setminus \{J\}})$ is distributed according to $\beta_J(t_J)$. By the pointwise degeneracy condition, the marginal of $\beta_J(t_J)$ on $T_{J'}$ is the point mass at $f_{J'J}(t_J)$, so $g_{J'}^J(h_J(t_J))$ assigns probability one to $h_{J'}(f_{J'J}(t_J))$. Hence $h_J(t_J) \subseteq D_J^1$ for every group J . Proceeding by induction, if $h_{J'}(T_{J'}) \subseteq D_{J'}^k$ for every group J' , then by (2), $g^J(h_J(t_J)) \left(\Theta \times \prod_{J' \in \mathcal{J} \setminus \{J\}} D_{J'}^k \right) = 1$ for every J and t_J , so $h_J(t_J) \subseteq D_J^{k+1}$. Hence $h_J(t_J) \subseteq D_J^\infty$.

(ii) The marginal of $\pi_{\mathbb{I}, \sigma}$ on Θ is the marginal of σ on Θ , which is μ . Fix a group J and bounded Borel functions φ on D_J^∞ and ψ on $\Theta \times \prod_{J' \in \mathcal{J} \setminus \{J\}} D_{J'}^\infty$. Since β_J is a version of the conditional distribution under σ of $(\theta, t_{\mathcal{J} \setminus \{J\}})$ given t_J , and $h_J(t_J)$ is a function of t_J ,

$$\begin{aligned} & \int \varphi(h_J(t_J)) \psi(\theta, h_{\mathcal{J} \setminus \{J\}}(t_{\mathcal{J} \setminus \{J\}})) d\sigma \\ &= \int \varphi(h_J(t_J)) \left[\int \psi d\left(\beta_J(t_J) \circ (\text{id}_\Theta \times h_{\mathcal{J} \setminus \{J\}})^{-1}\right) \right] d\sigma \\ &= \int \varphi(h_J(t_J)) \left[\int \psi dg^J(h_J(t_J)) \right] d\sigma \end{aligned}$$

$$= \int \varphi(u_J) \left[\int \psi dg^J(u_J) \right] d\pi_{\mathbb{I},\sigma},$$

where the second equality is part (i) and the third is the change of variables defining $\pi_{\mathbb{I},\sigma}$. Since product functions determine the joint distribution, g^J is a version of the conditional distribution under $\pi_{\mathbb{I},\sigma}$ of $(\theta, u_{\mathcal{J} \setminus \{J\}})$ given u_J . $\square$

In view of Proposition 6, we refer to $h = (h_J)_{J \in \mathcal{J}}$ as the projection of the group type space $\mathbb{I}$ onto the universal group type space. For a common prior group type space $(\mathbb{I}, \sigma)$, we call the distribution $\pi_{\mathbb{I},\sigma}$ of $(\theta, h(t))$, when (t, θ) is distributed according to σ , the *projection distribution* of $(\mathbb{I}, \sigma)$; when the prior is clear from context, we write $\pi_{\mathbb{I}}$, and for a finite common prior type space I , we write $\pi_{\mathbb{I}(I)}$ for the projection distribution of the induced group type space. By Proposition 6(ii), the projection distribution is a common prior for $\mathbb{I}^u$ consistent with μ , and, if $(\mathbb{I}, \sigma)$ is finite, it has finite support. The final result of this section shows that, conversely, every common prior for $\mathbb{I}^u$ consistent with μ that has finite support is the projection distribution of a finite common prior group type space, and indeed of one that is induced by a finite common prior type space, as in Section 2.1.

Proposition 7. *Let π be a common prior for $\mathbb{I}^u$ consistent with μ that has finite support. Then there is a finite common prior type space I such that $\pi_{\mathbb{I}(I)} = \pi$. Consequently, the projection distributions of the induced group type spaces of finite common prior type spaces are exactly the common priors for $\mathbb{I}^u$ consistent with μ that have finite support.*

Proof. Let V denote the support of the marginal of π on $\prod_{J \in \mathcal{J}} D_J^\infty$, and for each group J , let $V_J = \{u_J : u \in V\}$; these sets are finite. Since π is a common prior with finite support, for every $u_J \in V_J$, the conditional distribution under π of $(\theta, u_{\mathcal{J} \setminus \{J\}})$ given u_J is $g^J(u_J)$: any two versions of a conditional distribution agree at every atom of the conditioning marginal. For each group J with $|J| \geq 2$, fix an injection ι_J of V_J into the additive group $\mathbb{Z}_{m_J}$ of the integers modulo $m_J = |V_J|$, and fix a member $i_J \in J$.

The type space I is constructed as follows. First, (θ, u) is drawn according to π . Second, for each group J with $|J| \geq 2$, the value $\iota_J(u_J)$ is secretly shared among the members of J : the shares $\ell_i^J \in \mathbb{Z}_{m_J}$ for $i \in J \setminus \{i_J\}$ are drawn independently and uniformly, and

$$\ell_{i_J}^J = \iota_J(u_J) - \sum_{i \in J \setminus \{i_J\}} \ell_i^J,$$

so that $\sum_{i \in J} \ell_i^J = \iota_J(u_J)$. The shares are drawn independently across groups, and independently of θ conditional on u . The type of player i is $t_i = (u_{\{i\}}, (\ell_i^J)_{J \ni i, |J| \geq 2})$, an element of the finite set $T_i = V_{\{i\}} \times \prod_{J \ni i, |J| \geq 2} \mathbb{Z}_{m_J}$, from which types with zero probability are deleted,

and σ is the resulting joint distribution of (t, θ) . The marginal of σ on Θ is μ , since π is consistent with μ , so $I = (T, \sigma)$ is a finite common prior type space. Let $\mathbb{I}(I)$ be the induced group type space, in which the type of a group K is $t_K = (t_i)_{i \in K}$, and let h denote its projection.

Step 1: for every group K , the pooled type t_K determines u_K , and, conditional on t_K , the distribution of $(\theta, u_{\mathcal{J} \setminus \{K\}})$ is $g^K(u_K)$. For the first claim, for every group $J \subseteq K$ with $|J| \geq 2$, the pool contains the entire family of shares $(\ell_i^J)_{i \in J}$, whose sum is $\iota_J(u_J)$, and for every $i \in K$, the pool contains $u_{\{i\}}$ directly; in particular, the pool determines u_K , directly if K is a singleton, and via the sum of the shares $(\ell_i^K)_{i \in K}$ otherwise. For the second claim, decompose the data in t_K as follows. The hierarchies $u_{\{i\}}$ for $i \in K$ and the sums $\sum_{i \in J} \ell_i^J = \iota_J(u_J)$ for $J \subseteq K$ are functions of u_K , since the hierarchy of a group determines the hierarchies of its subgroups. The remaining data are, for each $J \subseteq K$ with $|J| \geq 2$, the shares $(\ell_i^J)_{i \in J \setminus \{i_J\}}$, and, for each $J \not\subseteq K$ with $J \cap K \neq \emptyset$, the proper subfamily $(\ell_i^J)_{i \in J \cap K}$. We claim that, conditional on u , the remaining data are independent and uniformly distributed, with a distribution that does not depend on u . For $J \subseteq K$, the shares $(\ell_i^J)_{i \in J \setminus \{i_J\}}$ are independent and uniform by construction. For $J \not\subseteq K$, this is immediate if $i_J \notin K$; and if $i_J \in J \cap K$, then, conditional on u and on the shares $(\ell_i^J)_{i \in J \cap K \setminus \{i_J\}}$, the share $\ell_{i_J}^J$ is a sum that includes the term $-\sum_{i \in J \setminus K} \ell_i^J$, in which each summand is uniform and independent of everything else, so $\ell_{i_J}^J$ is itself uniform and independent of u and of the other shares in the subfamily. Since the shares are independent across groups and independent of θ conditional on u , the pooled type t_K is informationally equivalent to (u_K, ε_K) , where ε_K is noise that is independent of (θ, u) . Hence, conditional on t_K , the distribution of $(\theta, u_{\mathcal{J} \setminus \{K\}})$ is the conditional distribution under π given u_K , which is $g^K(u_K)$, since π is a common prior.

Step 2: for every group K , $h_K(t_K) = u_K$ with probability one. We show, by induction on the order of beliefs and simultaneously across groups, that with probability one, $h_K^n(t_K)$ is the n th coordinate of u_K for every group K and every n . For $n = 1$: $h_K^1(t_K)$ is the conditional distribution of θ given t_K , which, by Step 1, is the marginal of $g^K(u_K)$ on Θ , and this is the first coordinate of u_K , by the marginal property of f in Proposition 2. Suppose the claim holds through order n for every group. As in the proof of Proposition 5, $h_K^{n+1}(t_K)$ is the conditional distribution given t_K of the state together with the first n orders of the realized hierarchies $(h_{K'}^m(t_{K'}))_{m \leq n, K' \in \mathcal{J} \setminus \{K\}}$, which, by the inductive hypothesis, is the conditional distribution given t_K of the state together with the first n coordinates of $(u_{K'})_{K' \in \mathcal{J} \setminus \{K\}}$. By Step 1, the latter is the image of $g^K(u_K)$ under the natural projection of $\Theta \times \prod_{K' \in \mathcal{J} \setminus \{K\}} H_0^{K'}$ onto X_n^K , which is $\text{marg}_{X_n^K} g^K(u_K)$, the $(n+1)$ th coordinate of u_K , again by Proposition 2.

By Step 2, the distribution of $(\theta, h(t))$ under σ is the distribution of (θ, u) under π , that is, $\pi_{\mathbb{I}(I)} = \pi$. For the final claim, if I is a finite common prior type space, then $\mathbb{I}(I)$ is a finite common prior group type space, by Proposition 1, so $\pi_{\mathbb{I}(I)}$ is a common prior for $\mathbb{I}^u$ consistent with μ , by Proposition 6(ii), and it has finite support. $\square$

We suspect that this result can be extended to infinite type spaces. A difficulty is in extending the definition of “induce” when a type space is infinite; the beliefs in the group type space would only be determined almost everywhere, so that more than one group type space would be induced by a given type space. While we think that this would be a valuable mathematical extension, it is not necessary to make the conceptual point that type spaces are rich enough to reach all group belief hierarchies. Thus, while group type spaces are a convenient analytical tool, they do not introduce any new kinds of distributed information that were not already present in the model of Harsanyi (1967).

3 Group beliefs and equilibrium

3.1 Games and Equilibrium

In this section, we introduce games, and we connect group type spaces to behavior.

An *action space* is a finite product set of the form $A = \prod_{i=1}^N A_i$. The players interact through a *game structure* (also called a game form or a base game), which consists of an action space A and, for each player i , an expected utility index $u_i : A \times \Theta \rightarrow \mathbb{R}$. The game structure is denoted by $G = (A, u)$.

Given an action space A , an *outcome* is a distribution $\phi \in \Delta(A \times \Theta)$ for which the marginal on Θ is μ .

Fix a game structure $G = (A, u)$ and a finite common prior group type space $\mathbb{I} = (T, \sigma)$; throughout this section, all group type spaces are finite common prior group type spaces, identified with pairs (T, σ) as in Section 2.1, and likewise all type spaces are finite common prior type spaces. Actions are taken by the individual players: each player i observes the type $t_{\{i\}}$ of the singleton group $\{i\}$ and randomizes over actions accordingly. Thus, a (behavioral) strategy for player i is a mapping $b_i : T_{\{i\}} \rightarrow \Delta(A_i)$. A profile of strategies $b = (b_1, \dots, b_N)$ is identified with the mapping $b : \prod_{i=1}^N T_{\{i\}} \rightarrow \Delta(A)$, where $b(a \mid (t_{\{i\}})_{i=1}^N) = \prod_{i=1}^N b_i(a_i \mid t_{\{i\}})$.

The group type space $\mathbb{I}$ and a profile of strategies b induce the outcome ϕ defined by

$$\phi(a, \theta) = \sum_{t \in T} \sigma(t, \theta) \prod_{i=1}^N b_i(a_i \mid t_{\{i\}}).$$

Thus, in the induced outcome, the state and the profile of group types are drawn according to σ , the players mix independently over actions, with player i 's mixture depending only on the type of the singleton group $\{i\}$, and the types are integrated out. The marginal of the induced outcome on Θ is μ , so induced outcomes are indeed outcomes.

Following Bergemann and Morris (2013, 2016), we define Bayes correlated equilibrium relative to a group type space. A *Bayes correlated equilibrium (BCE)* of $(\mathbb{I}, G)$ is a joint distribution $\phi \in \Delta(A \times T \times \Theta)$, over the players' actions, the groups' types, and the state, whose marginal on $T \times \Theta$ is σ and which satisfies the following obedience constraints: for all i , $\tau \in T_{\{i\}}$, a_i , and a'_i ,

$$\sum_{a_{-i} \in A_{-i}, \theta \in \Theta, t \in T: t_{\{i\}} = \tau} \phi(a_i, a_{-i}, t, \theta) (u_i(a_i, a_{-i}, \theta) - u_i(a'_i, a_{-i}, \theta)) \geq 0. \quad (3)$$

Thus, the state and the profile of group types are drawn according to σ , actions are recommended according to ϕ , and no player, knowing their singleton type and their recommendation, can gain by deviating.

Through the operation by which a type space induces a group type space, these definitions apply implicitly to type spaces as well: given a finite common prior type space I , we say that ϕ is a BCE of (I, G) if it is a BCE of $(\mathbb{I}(I), G)$, where $\mathbb{I}(I)$ is the induced group type space of Section 2.1. Since the type of the singleton group $\{i\}$ in $\mathbb{I}(I)$ is the type t_i of player i , this coincides with the definition of Bergemann and Morris (2016), with I in the role of the base information structure.

3.2 Group Belief Invariance

A BCE ϕ of $(\mathbb{I}, G)$ is *group belief invariant* (Brooks, Du, and Haberman, 2026) if the actions of each group depend only on that group's type: for every group J , conditional on t_J , a_J and $(t_{\mathcal{J} \setminus \{J\}}, \theta)$ are independent; that is, for every group J and all a_J and t_J ,

$$\phi(\{a_J\} \times A_{-J} \mid t_J, t_{\mathcal{J} \setminus \{J\}}, \theta)$$

is constant over the $(t_{\mathcal{J} \setminus \{J\}}, \theta)$ that have positive probability given t_J . Since t_J determines $t_{J'}$ for every group $J' \subseteq J$, conditioning on t_J is equivalent to conditioning on the types of all of the subgroups of J . As with BCE, group belief invariance applies implicitly to type spaces: a group belief invariant BCE of (I, G) , for a finite common prior type space I , is a group belief invariant BCE of $(\mathbb{I}(I), G)$.

Group belief invariance combines familiar conditions from the taxonomy of correlated equilibrium in games with incomplete information (Forges, 1993, 2006). The conditions for the singleton groups are the *belief invariance* conditions: each player's action carries no information beyond the player's own type. The condition for the grand group requires the action profile to be conditionally independent of the state, given the profile of all of the types, which is the condition that distinguishes the *Bayesian solution* of Forges (1993) from Bayes correlated equilibrium. Thus, for a finite common prior type space I , every group belief invariant BCE of (I, G) is a belief invariant Bayesian solution, and, when $N = 2$, the two concepts coincide, since the only groups are the two singletons and the grand group. When $N \geq 3$, group belief invariance also constrains the intermediate groups, and the refinement can be strict.

As the next proposition shows, for the group type spaces that are induced by type spaces, group belief invariance is equivalent to an ostensibly weaker requirement, in which a group's actions are only required to carry no information about the state and about the types of the groups with which it shares no members.

Proposition 8. *Let I be a finite common prior type space, let $\mathbb{I}(I)$ be the induced group type space, and let ϕ be a BCE of $(\mathbb{I}(I), G)$. For a group J , let $\mathcal{D}(J) = \{J' \in \mathcal{J} : J' \cap J = \emptyset\}$ denote the set of groups that are disjoint from J , and write $t_{\mathcal{D}(J)} = (t_{J'})_{J' \in \mathcal{D}(J)}$. Then ϕ is group belief invariant if and only if for every group J , conditional on t_J , a_J and $(t_{\mathcal{D}(J)}, \theta)$ are independent.*

Proof. Necessity is immediate, since $(t_{\mathcal{D}(J)}, \theta)$ is a function of $(t_{\mathcal{J} \setminus \{J\}}, \theta)$. For sufficiency, fix a group J . If $J = \{1, \dots, N\}$ is the grand group, then $\mathcal{D}(J) = \emptyset$, and t_J determines the entire profile with probability one, so both conditions assert that, conditional on t_J , a and θ are independent. Otherwise, the complement $J^c = \{1, \dots, N\} \setminus J$ is a nonempty group that is disjoint from J . In the induced group type space, the type of each group J'' is the vector $(t_i)_{i \in J''}$, which is determined by $(t_J, t_{J^c}) = ((t_i)_{i \in J}, (t_i)_{i \notin J})$. Hence, conditional on t_J , the profile $t_{\mathcal{J} \setminus \{J\}}$ is a function of t_{J^c} , and independence of a_J from $(t_{\mathcal{D}(J)}, \theta)$, which includes (t_{J^c}, θ) , implies independence of a_J from $(t_{\mathcal{J} \setminus \{J\}}, \theta)$. $\square$

The equivalence does not extend to general group type spaces, in which the type of a supgroup of J may contain more than the pooled types of its members: a group's actions may then depend on such additional information without carrying any information about the state or about the types of disjoint groups.

We write $S_{\mathbb{I}}^{GBI}(G)$ for the set of marginals on $A \times \Theta$ of the group belief invariant BCE of $(\mathbb{I}, G)$. Every element of $S_{\mathbb{I}}^{GBI}(G)$ is an outcome, since the marginal on Θ of a BCE of $(\mathbb{I}, G)$ is μ .

An instructive special case is the *null* group type space, in which every group's type set is a singleton. In that case, a BCE of $(\mathbb{I}, G)$ is identified with an outcome $\phi \in \Delta(A \times \Theta)$ that satisfies the obedience constraints (3), which no longer involve any conditioning on types, and group belief invariance says that a_J and θ are independent for every group J . When $J = \{i\}$ is a singleton, this is the defining property of *belief invariant BCE* (Forges, 1993, 2006; Bergemann and Morris, 2016), and when $J = \{1, \dots, N\}$ is the grand group, it is the defining property of the *Bayesian solution* (Forges, 1993). Moreover, since independence is preserved under marginalization, the grand group condition implies the condition for every group, so on the null group type space, group belief invariant BCE and Bayesian solutions coincide.

3.3 The Sufficiency of Universal Group Beliefs

For the remainder of this section, we let $h = (h_J)_{J \in \mathcal{J}}$ denote the projection of $\mathbb{I}$ onto the universal group type space and $\pi_{\mathbb{I}}$ its projection distribution, as defined in Section 2.3; since $\mathbb{I}$ is finite, $\pi_{\mathbb{I}}$ has finite support.

The main result of this section is that the projection distribution is a sufficient statistic for the predictions of group belief invariant BCE: two group type spaces that induce the same distribution over profiles of group belief hierarchies generate the same joint distributions over actions and states.

Theorem 1. *Let $\mathbb{I} = (T, \sigma)$ and $\mathbb{I}' = (T', \sigma')$ be finite common prior group type spaces with the same projection distribution. Then for every game structure G , $S_{\mathbb{I}}^{GBI}(G) = S_{\mathbb{I}'}^{GBI}(G)$.*

Theorem 1 is the group analogue of the characterization of belief invariant BCE by Liu (2015), with hierarchies of group beliefs in place of hierarchies of individual beliefs. In particular, if two finite common prior type spaces I and I' induce group type spaces with the same projection distribution, then $S_{\mathbb{I}(I)}^{GBI}(G) = S_{\mathbb{I}(I')}^{GBI}(G)$.

Theorem 1 combines with the results of Section 2.3 in two further ways that are worth recording in words. First, since the predictions of group belief invariant BCE depend on a group type space only through its projection distribution, Proposition 7 implies that every finite common prior group type space is equivalent, for the purposes of those predictions, to a group type space that is induced by a finite common prior type space: informational synergies, whereby a group's type contains more than the pooled types of its members, leave no trace in the hierarchies of group beliefs, and are therefore redundant for group belief invariant predictions. Second, by Proposition 8, group belief invariance on induced group type spaces is equivalent to the ostensibly weaker requirement that each group's

actions carry no information about the state and about the types of the groups that are disjoint from it; hence every common prior for $\mathbb{I}^u$ consistent with μ with finite support is realized by a group type space on which the two requirements coincide.

Proof. Let π denote the common projection distribution, let $U \subseteq \prod_{J \in \mathcal{J}} D_J^\infty$ denote the support of the marginal of π on $\prod_{J \in \mathcal{J}} D_J^\infty$, which is a finite set, and for $u \in U$, write u_J for the coordinate of u indexed by J and $U_J = \{u_J : u \in U\}$. Define $S_\pi^{GBI}(G)$ to be the set of marginals on $A \times \Theta$ of the distributions $\psi \in \Delta(A \times U \times \Theta)$ satisfying: (i) the marginal of ψ on $U \times \Theta$ is π ; (ii) obedience: for all i , $v \in U_{\{i\}}$, a_i , and a'_i ,

$$\sum_{a_{-i} \in A_{-i}, \theta \in \Theta, u \in U: u_{\{i\}} = v} \psi(a_i, a_{-i}, u, \theta) (u_i(a_i, a_{-i}, \theta) - u_i(a'_i, a_{-i}, \theta)) \geq 0; \quad (4)$$

and (iii) invariance: for every group J and all a_J and u_J , $\psi(\{a_J\} \times A_{-J} \mid u_J, u_{\mathcal{J} \setminus \{J\}}, \theta)$ is constant over the $(u_{\mathcal{J} \setminus \{J\}}, \theta)$ that have positive probability given u_J . Thus, the conditions defining $S_\pi^{GBI}(G)$ align with those defining $S_{\mathbb{I}}^{GBI}(G)$. We prove that $S_{\mathbb{I}}^{GBI}(G) = S_\pi^{GBI}(G)$; since the right-hand side depends on $\mathbb{I}$ only through π , the same argument applied to $\mathbb{I}'$ proves the theorem.

We use two facts from Section 2.3. First, for every group J and type t_J , the conditional distribution under σ of $(\theta, t_{\mathcal{J} \setminus \{J\}})$ given t_J is $\beta_J(t_J)$, so, by Proposition 5, the conditional distribution under σ of $(\theta, h_{\mathcal{J} \setminus \{J\}}(t_{\mathcal{J} \setminus \{J\}}))$ given t_J is $g^J(h_J(t_J))$; in particular, it depends on t_J only through $h_J(t_J)$. Second, by Proposition 6(ii), π is a common prior for $\mathbb{I}^u$; since π has finite support, the conditional distribution under π of $(\theta, u_{\mathcal{J} \setminus \{J\}})$ given any $u_J \in U_J$ is $g^J(u_J)$.

First, we show $S_{\mathbb{I}}^{GBI}(G) \subseteq S_\pi^{GBI}(G)$. Let ϕ be a group belief invariant BCE of $(\mathbb{I}, G)$, and let ψ be the distribution of $(a, h(t), \theta)$ when (a, t, θ) is distributed according to ϕ . The marginal of ψ on $U \times \Theta$ is the distribution of $(h(t), \theta)$ under σ , which is π , and the marginals of ψ and ϕ on $A \times \Theta$ coincide. For obedience, fix i , v , a_i , and a'_i . Since the events $\{t_{\{i\}} = \tau\}$ for $\tau \in h_{\{i\}}^{-1}(v)$ partition the event $\{h_{\{i\}}(t_{\{i\}}) = v\}$, the left-hand side of (4) is the sum, over $\tau \in h_{\{i\}}^{-1}(v)$, of the left-hand sides of (3), each of which is nonnegative. For invariance, fix a group J and a $v_J \in U_J$ with positive probability. Under ψ , the conditional distribution of $(a_J, u_{\mathcal{J} \setminus \{J\}}, \theta)$ given $u_J = v_J$ is a mixture, over the types $\tau_J \in h_J^{-1}(v_J)$ with weights proportional to their probabilities, of the conditional distributions given $t_J = \tau_J$. Fix such a τ_J . Since ϕ is group belief invariant, conditional on $t_J = \tau_J$, a_J is independent of $(t_{\mathcal{J} \setminus \{J\}}, \theta)$, and hence of $(u_{\mathcal{J} \setminus \{J\}}, \theta) = (h_{\mathcal{J} \setminus \{J\}}(t_{\mathcal{J} \setminus \{J\}}), \theta)$; and by the first fact, the conditional distribution of $(u_{\mathcal{J} \setminus \{J\}}, \theta)$ is $g^J(v_J)$. Thus, each component of the mixture is the product

of a distribution over a_J with the fixed distribution $g^J(v_J)$ over $(u_{\mathcal{T} \setminus \{J\}}, \theta)$, and hence so is the mixture. So a_J and $(u_{\mathcal{T} \setminus \{J\}}, \theta)$ are independent given $u_J = v_J$, and ψ satisfies (iii).

Next, we show $S_\pi^{GBI}(G) \subseteq S_{\mathbb{I}}^{GBI}(G)$. Let ψ satisfy (i)–(iii), and define $\phi \in \Delta(A \times T \times \Theta)$ by

$$\phi(a, t, \theta) = \sigma(t, \theta) \psi(a \mid h(t), \theta),$$

with the convention that $\phi(a, t, \theta) = 0$ when $\sigma(t, \theta) = 0$; note that if $\sigma(t, \theta) > 0$, then $(h(t), \theta)$ has positive probability under π , so the conditional is well defined. Summing over a , the marginal of ϕ on $T \times \Theta$ is σ . Grouping the type profiles t according to $u = h(t)$, the marginal of ϕ on $A \times \Theta$ is

$$\sum_{t \in T} \sigma(t, \theta) \psi(a \mid h(t), \theta) = \sum_{u \in U} \pi(\theta, u) \psi(a \mid u, \theta),$$

which is the marginal of ψ on $A \times \Theta$. For invariance, observe that by (iii), whenever (u, θ) has positive probability, $\psi(\{a_J\} \times A_{-J} \mid u, \theta)$ depends only on (a_J, u_J) ; denote it by $r_J(a_J \mid u_J)$. For (t, θ) with positive probability,

$$\phi(\{a_J\} \times A_{-J} \mid t, \theta) = \psi(\{a_J\} \times A_{-J} \mid h(t), \theta) = r_J(a_J \mid h_J(t_J)),$$

which is constant over $(t_{\mathcal{T} \setminus \{J\}}, \theta)$, so ϕ is group belief invariant. For obedience, fix i , $\tau \in T_{\{i\}}$, a_i , and a'_i , and let $v = h_{\{i\}}(\tau)$ and $\Delta_i(a, \theta) = u_i(a_i, a_{-i}, \theta) - u_i(a'_i, a_{-i}, \theta)$. Writing $\sigma(\tau)$ for the probability of τ and using that $h(t) = (v, h_{\mathcal{T} \setminus \{\{i\}\}}(t_{\mathcal{T} \setminus \{\{i\}\}}))$ on the event $\{t_{\{i\}} = \tau\}$, the left-hand side of (3) is

$$\sigma(\tau) \sum_{a_{-i}, \theta, u_{\mathcal{T} \setminus \{\{i\}\}}} \Pr(\theta, u_{\mathcal{T} \setminus \{\{i\}\}} \mid \tau) \psi(a \mid v, u_{\mathcal{T} \setminus \{\{i\}\}}, \theta) \Delta_i(a, \theta),$$

where $\Pr(\theta, u_{\mathcal{T} \setminus \{\{i\}\}} \mid \tau)$ denotes the conditional probability under σ that the state is θ and $h_{\mathcal{T} \setminus \{\{i\}\}}(t_{\mathcal{T} \setminus \{\{i\}\}}) = u_{\mathcal{T} \setminus \{\{i\}\}}$, given $t_{\{i\}} = \tau$. By the first fact, this conditional probability is $g^{\{i\}}(v)(\theta, u_{\mathcal{T} \setminus \{\{i\}\}})$, and by the second fact, $\pi(\theta, (v, u_{\mathcal{T} \setminus \{\{i\}\}})) = \bar{\pi}_{\{i\}}(v) g^{\{i\}}(v)(\theta, u_{\mathcal{T} \setminus \{\{i\}\}})$, where $\bar{\pi}_{\{i\}}(v) > 0$ is the probability under π that the coordinate indexed by $\{i\}$ is v . Hence the left-hand side of (3) equals $\sigma(\tau)/\bar{\pi}_{\{i\}}(v)$ times the left-hand side of (4), which is nonnegative. Thus ϕ is a group belief invariant BCE of $(\mathbb{I}, G)$ whose marginal on $A \times \Theta$ is that of ψ , completing the proof. $\square$

3.4 Minimality of Group Belief Invariant BCE

In this subsection, we characterize group belief invariant BCE as the minimal solution concept that is measurable with respect to group beliefs and consistent with Bayes Nash equilibrium.

We first extend Bayes Nash equilibrium to group type spaces in the obvious way. Given a profile of strategies $b = (b_1, \dots, b_N)$, player i 's expected utility is

$$U_i(b; \mathbb{I}, G) = \sum_{t \in T} \sum_{\theta \in \Theta} \sigma(t, \theta) \sum_{a \in A} u_i(a, \theta) \prod_{j=1}^N b_j(a_j \mid t_{\{j\}}).$$

The profile b is a *Bayes Nash equilibrium (BNE)* of $(\mathbb{I}, G)$ if $U_i(b; \mathbb{I}, G) \geq U_i(b'_i, b_{-i}; \mathbb{I}, G)$ for every player i and every strategy b'_i .

A *solution concept* is a mapping S that assigns to every finite common prior group type space $\mathbb{I}$ and every game structure G a subset $S_{\mathbb{I}}(G)$ of $\Delta(A \times \Theta)$, where A is the action space of G . Two examples are S^{BNE} , which assigns to $(\mathbb{I}, G)$ the set $S_{\mathbb{I}}^{BNE}(G)$ of outcomes that are induced by $\mathbb{I}$ and some BNE of $(\mathbb{I}, G)$, and S^{GBI} , which assigns to $(\mathbb{I}, G)$ the set $S_{\mathbb{I}}^{GBI}(G)$ of Section 3.2. A solution concept S is *group belief measurable* if $S_{\mathbb{I}}(G) = S_{\mathbb{I}'}(G)$ for every game structure G , whenever $\pi_{\mathbb{I}} = \pi_{\mathbb{I}'}$, and a solution concept S *refines* a solution concept S' if $S_{\mathbb{I}}(G) \subseteq S'_{\mathbb{I}}(G)$ for every finite common prior group type space $\mathbb{I}$ and every game structure G . By Theorem 1, S^{GBI} is group belief measurable. By contrast, S^{BNE} is not: for example, appending to any group type space a payoff-irrelevant fair coin that is observed by all of the players leaves the projection distribution unchanged, since every group's hierarchy is the same for either realization of the coin, but it enlarges the set of BNE outcomes whenever the set of BNE outcomes is not convex.

Theorem 2. *The solution concept S^{GBI} is group belief measurable and is refined by S^{BNE} . Moreover, S^{GBI} refines every group belief measurable solution concept that is refined by S^{BNE} . In fact, for every finite common prior group type space $\mathbb{I}$ and every game structure G ,*

$$S_{\mathbb{I}}^{GBI}(G) = \bigcup \{S_{\mathbb{I}'}^{BNE}(G) : \pi_{\mathbb{I}'} = \pi_{\mathbb{I}}\}, \quad (5)$$

where $\mathbb{I}'$ ranges over the finite common prior group type spaces.

Thus, group belief invariant BCE is the minimal group belief measurable solution concept that contains the Bayes Nash equilibrium predictions. Theorem 2 and the formula (5) are the group analogues of the characterization of belief invariant BCE by Liu (2015) as

the Bayes Nash equilibrium predictions across all common prior type spaces that induce the same hierarchies of individual beliefs.

Proof. Group belief invariance of S^{GBI} is Theorem 1. Throughout the remainder of the proof, fix a game structure G ; all of the remaining claims are established pointwise in G .

Step 1: S^{BNE} refines S^{GBI} . Let b be a BNE of $(\mathbb{I}, G)$, and let

$$\phi(a, t, \theta) = \sigma(t, \theta) \prod_{j=1}^N b_j(a_j \mid t_{\{j\}}),$$

whose marginal on $A \times \Theta$ is the outcome induced by $\mathbb{I}$ and b , and whose marginal on $T \times \Theta$ is σ . For obedience, fix i , $\tau \in T_{\{i\}}$, a_i , and a'_i . The left-hand side of (3) is $b_i(a_i \mid \tau)D$, where

$$D = \sum_{t \in T: t_{\{i\}} = \tau} \sum_{\theta \in \Theta} \sum_{a_{-i} \in A_{-i}} \sigma(t, \theta) \prod_{j \neq i} b_j(a_j \mid t_{\{j\}}) (u_i(a_i, a_{-i}, \theta) - u_i(a'_i, a_{-i}, \theta)).$$

If $b_i(a_i \mid \tau) > 0$ and $D < 0$, then the strategy b'_i that agrees with b_i except that, at τ , it moves the probability $b_i(a_i \mid \tau)$ from a_i to a'_i satisfies $U_i(b'_i, b_{-i}; \mathbb{I}, G) - U_i(b; \mathbb{I}, G) = -b_i(a_i \mid \tau)D > 0$, contradicting equilibrium; so the left-hand side of (3) is nonnegative. For invariance, note that conditional on (t, θ) , the actions are independent across players with a_j distributed according to $b_j(\cdot \mid t_{\{j\}})$, so the conditional distribution of a_J is $\prod_{i \in J} b_i(\cdot \mid t_{\{i\}})$, which is a function of t_J , since t_J determines $t_{\{i\}}$ for every $i \in J$. Hence, conditional on t_J , a_J is independent of $(t_{\mathcal{J} \setminus \{J\}}, \theta)$, and ϕ is a group belief invariant BCE of $(\mathbb{I}, G)$.

Step 2: for every $\rho \in S_{\mathbb{I}}^{GBI}(G)$, there is a group type space $\mathbb{I}'$ with $\pi_{\mathbb{I}'} = \pi_{\mathbb{I}}$ and $\rho \in S_{\mathbb{I}'}^{BNE}(G)$. Let ϕ be a group belief invariant BCE of $(\mathbb{I}, G)$ with marginal ρ on $A \times \Theta$. Define $\mathbb{I}' = (T', \sigma')$ by appending recommendations to types: $T'_J = T_J \times A_J$ for every group J , and σ' is the distribution of $((t_J, a_J))_{J \in \mathcal{J}}$ when (a, t, θ) is distributed according to ϕ , deleting types with zero probability. Then $\mathbb{I}'$ is a finite common prior group type space: every type has positive probability by construction, the marginal of σ' on Θ is μ , and for $J' \subseteq J$, the type (t_J, a_J) determines $(t_{J'}, a_{J'})$ with probability one, since t_J determines $t_{J'}$ and $a_{J'}$ is a subvector of a_J .

We claim that the projection h' of $\mathbb{I}'$ satisfies $h'_J(t_J, a_J) = h_J(t_J)$ with probability one, for every group J . We show, by induction on the order of beliefs and simultaneously across groups, that $(h')^n_J(t_J, a_J) = h^n_J(t_J)$ with probability one, for every group J and every n . Throughout, we use the following consequence of the group belief invariance of ϕ : conditional on t_J , the pair a_J and $(t_{\mathcal{J} \setminus \{J\}}, \theta)$ are independent, so conditioning on (t_J, a_J) and conditioning on t_J yield the same conditional distribution for any function of $(t_{\mathcal{J} \setminus \{J\}}, \theta)$. For $n = 1$: $(h')^1_J(t_J, a_J)$ is the conditional distribution of θ given (t_J, a_J) , which equals the

conditional distribution of θ given t_J , which is $h_J^1(t_J)$. Suppose the claim holds through order n for every group. Then $(h')_J^{n+1}(t_J, a_J)$ is the conditional distribution given (t_J, a_J) of the state together with the first n orders of the hierarchies $((h')_{J'}^m(t_{J'}, a_{J'}))_{m \leq n, J' \in \mathcal{J} \setminus \{J\}}$, which, by the inductive hypothesis, is the conditional distribution given (t_J, a_J) of the state together with the first n coordinates of $(h_{J'}(t_{J'}))_{J' \in \mathcal{J} \setminus \{J\}}$ — a function of $(t_{\mathcal{J} \setminus \{J\}}, \theta)$ — and hence equals the corresponding conditional distribution given t_J , which is $h_J^{n+1}(t_J)$. It follows that the distribution of $(\theta, h'(t'))$ under σ' equals the distribution of $(\theta, h(t))$ under σ , that is, $\pi_{\mathbb{I}'} = \pi_{\mathbb{I}}$.

Finally, the obedient strategies $b'_i(\cdot \mid (t_{\{i\}}, a_i)) = \delta_{a_i}$ form a BNE of $(\mathbb{I}', G)$: since every type of $\mathbb{I}'$ has positive probability, a profitable deviation would yield, at some positive-probability type $(t_{\{i\}}, a_i)$ and some action a'_i , a strict violation of the obedience constraint (3) of ϕ at $(t_{\{i\}}, a_i, a'_i)$, since, under the obedient strategies of the other players, the conditional distribution of (a_{-i}, θ) given player i 's type $(t_{\{i\}}, a_i)$ in $\mathbb{I}'$ is the conditional distribution under ϕ of (a_{-i}, θ) given $(t_{\{i\}}, a_i)$. The outcome induced by $\mathbb{I}'$ and the obedient strategies is the marginal of ϕ on $A \times \Theta$, which is ρ .

Step 3: minimality and (5). For the inclusion $\supseteq$ in (5), if $\pi_{\mathbb{I}'} = \pi_{\mathbb{I}}$, then by Step 1 and Theorem 1, $S_{\mathbb{I}'}^{BNE}(G) \subseteq S_{\mathbb{I}'}^{GBI}(G) = S_{\mathbb{I}}^{GBI}(G)$; the inclusion $\subseteq$ is Step 2. In particular, Step 1 says that S^{BNE} refines S^{GBI} . Finally, let S be any group belief measurable solution concept that is refined by S^{BNE} , and let $\rho \in S_{\mathbb{I}}^{GBI}(G)$. By Step 2, there is an $\mathbb{I}'$ with $\pi_{\mathbb{I}'} = \pi_{\mathbb{I}}$ and $\rho \in S_{\mathbb{I}'}^{BNE}(G) \subseteq S_{\mathbb{I}'}(G) = S_{\mathbb{I}}(G)$, where the last equality holds because S is group belief measurable. Hence $S_{\mathbb{I}}^{GBI}(G) \subseteq S_{\mathbb{I}}(G)$, and S^{GBI} refines S . $\square$

4 An Example

We illustrate the theory with an example, based on the type spaces of the introduction. The example has two players, so, as noted in Section 3.2, the predictions of group belief invariant BCE coincide with those of the belief invariant Bayesian solution, and the comparison below can equally be read as a comparison of belief invariant Bayesian solutions across the two type spaces. There are $N = 2$ players, the states are $\Theta = \{0, 1\}$, and the prior μ is uniform. Each player's action set is $A_i = \{0, 1\}$, and payoffs are common:

$$u_1(a, \theta) = u_2(a, \theta) = \begin{cases} 1 & \text{if } a_1 \oplus a_2 = \theta; \\ 0 & \text{otherwise,} \end{cases}$$

where $\oplus$ denotes addition modulo 2, as in the introduction. Thus, the players want to match when the state is 0 and to mismatch when the state is 1. We denote this game

structure by G , we write $W = \{(a, \theta) : a_1 \oplus a_2 = \theta\}$ for the event that the players *succeed*, and we refer to $\rho(W)$ as the success probability of an outcome ρ . Note that if player i switches their action, then success becomes failure and vice versa, so, at any information set, the payoff from switching is one minus the payoff from the prescribed action.

We compare two finite common prior type spaces: the *null* type space I^0 , in which each T_i is a singleton, and the *XOR* type space I^X of the introduction, in which t_1 is uniformly distributed on $\{0, 1\}$ and independent of θ , and $t_2 = t_1 \oplus \theta$. As shown in the introduction, the two type spaces induce identical hierarchies of individual beliefs. Their induced group type spaces differ: in $\mathbb{I}(I^X)$, the grand group knows the state, since $\theta = t_1 \oplus t_2$, while in $\mathbb{I}(I^0)$, it knows nothing.

Proposition 9. *In the game structure G :*

- (i) $S_{\mathbb{I}(I^0)}^{GBI}(G)$ is the set of outcomes in which the action profile is independent of the state, that is, outcomes of the form $\mu \otimes \nu$ with $\nu \in \Delta(A)$. Every such outcome has success probability $1/2$.
- (ii) $S_{\mathbb{I}(I^X)}^{GBI}(G)$ is the set of outcomes ρ such that a_1 and θ are independent, a_2 and θ are independent, and, for each player i and each action a_i with positive probability, $\rho(W \mid a_i) \geq 1/2$. The success probabilities of these outcomes are exactly the interval $[1/2, 1]$.

Proof. (i) In $\mathbb{I}(I^0)$, every group's type set is a singleton, so, as discussed in Section 3.2, a group belief invariant BCE is identified with an outcome ρ , and the invariance condition for the grand group says that (a_1, a_2) and θ are independent, that is, $\rho = \mu \otimes \nu$ for some $\nu \in \Delta(A)$. Conversely, every such ρ is group belief invariant, and it satisfies obedience: conditional on a recommendation a_i , the state is uniformly distributed and independent of the recommended profile, so obeying and switching each succeed with probability $1/2$. For the same reason, the success probability of $\mu \otimes \nu$ is $1/2$.

(ii) First, let ϕ be a group belief invariant BCE of $(\mathbb{I}(I^X), G)$ with marginal ρ on $A \times \Theta$. By invariance for the singleton group $\{i\}$, conditional on t_i , a_i and θ are independent; and t_i and θ are independent under σ , since t_1 is independent of θ by construction and $t_2 = t_1 \oplus \theta$ is uniform conditional on θ . Hence a_i and θ are independent under ρ . For the conditional success probabilities, fix (t_i, a_i) with positive probability. Since switching turns success into failure, the obedience constraint at (t_i, a_i) says exactly that $\phi(W \mid t_i, a_i) \geq 1 - \phi(W \mid t_i, a_i)$, that is, $\phi(W \mid t_i, a_i) \geq 1/2$; averaging over t_i conditional on a_i gives $\rho(W \mid a_i) \geq 1/2$.

Conversely, let ρ satisfy the conditions in (ii), and define ϕ on $\mathbb{I}(I^X)$ by

$$\phi(a, t, \theta) = \sigma(t, \theta) \rho(a \mid \theta);$$

that is, the mediator observes $\theta = t_1 \oplus t_2$ and draws the action profile from the target conditional $\rho(\cdot \mid \theta)$. The marginal of ϕ on $T \times \Theta$ is σ , and the marginal on $A \times \Theta$ is $\sum_t \sigma(t, \theta) \rho(a \mid \theta) = \mu(\theta) \rho(a \mid \theta) = \rho(a, \theta)$. For invariance: conditional on (t, θ) , the distribution of a_i is $\rho(a_i \mid \theta) = \rho(a_i)$, by the independence of a_i and θ under ρ ; since this conditional distribution is constant, the invariance condition holds for the singleton groups. For the grand group, $t_J = (t_1, t_2)$ determines θ , so the invariance condition holds trivially. For obedience, fix (t_i, a_i) with positive probability. Conditional on t_i , θ is uniform, and, conditional on (t_i, θ) , the recommendations are drawn from $\rho(\cdot \mid \theta)$; since $\rho(a_i \mid \theta) = \rho(a_i)$, the state remains uniform conditional on (t_i, a_i) , and $\phi(W \mid t_i, a_i) = \rho(W \mid a_i) \geq 1/2$, so obedience holds.

Finally, the set in (ii) is convex, since, with the marginal of ρ on Θ fixed at μ , its defining conditions are linear in ρ . The outcome $\rho^*(a, \theta) = \frac{1}{4} \mathbf{1}\{a_1 \oplus a_2 = \theta\}$ belongs to the set, with success probability 1, and any $\mu \otimes \nu$ belongs to the set, with success probability $1/2$, so the success probabilities attained on the set are exactly $[1/2, 1]$. $\square$

The comparison is stark. The set in (i) is strictly contained in the set in (ii): every outcome $\mu \otimes \nu$ satisfies the conditions in (ii), with conditional success probabilities equal to $1/2$, while the first-best outcome ρ^* , in which the players always succeed, belongs to (ii) but not to (i). On the null type space, group belief invariant BCE pins the success probability at $1/2$; on the XOR type space, every success probability in $[1/2, 1]$ arises, and the first-best is attained by the Bayes Nash equilibrium in which each player plays their own type, $a_i = t_i$.

The example also isolates what group belief invariance adds to belief invariance. Since the two type spaces induce identical hierarchies of individual beliefs, any solution concept that is invariant with respect to individual hierarchies must assign them the same predictions. Belief invariant BCE (Forges, 1993, 2006; Bergemann and Morris, 2016; Liu, 2015) is such a concept, and its outcome sets for I^0 and for I^X both coincide with the set in (ii): for I^0 , this is immediate, since belief invariance requires only that each player's action be independent of the state, and the obedience constraints are those in (ii); for I^X , it follows from the invariance theorem of Liu (2015), or by repeating the arguments in the proof of Proposition 9(ii). Thus, belief invariant BCE cannot distinguish the two type spaces, and, in particular, it permits the players' actions to be perfectly coordinated with the state even when nobody—and no group—has any information about it. Group belief invariance agrees with belief invariance on I^X , where the pair genuinely knows the state, and corrects the predictions on I^0 , where the pair jointly knows nothing, by collapsing them to

the state-independent outcomes. This is the strategic counterpart of the observation with which the paper began.

Finally, we extend the example to three players to substantiate the observation from Section 3.2 that, when $N \geq 3$, group belief invariant BCE can strictly refine the belief invariant Bayesian solution. Let $N = 3$, keep $\Theta = \{0, 1\}$ and μ uniform, and let G_3 be the game structure in which $A_i = \{0, 1\}$ for every player and payoffs are common: $u_i(a, \theta) = 1$ if $a_1 \oplus a_2 \oplus a_3 = \theta$, and $u_i(a, \theta) = 0$ otherwise, so that the players want the parity of their actions to match the state. As before, switching any one player's action turns success into failure and vice versa. Let I^{X^3} be the three-player XOR type space: t_1 and t_2 are independently and uniformly distributed on $\{0, 1\}$, independently of θ , and $t_3 = t_1 \oplus t_2 \oplus \theta$. Every type profile has probability $1/8$, and $\theta = t_1 \oplus t_2 \oplus t_3$. The distinguishing feature of I^{X^3} is that every *pair* of players jointly knows nothing about the state: for every group J with $|J| = 2$, the pooled types $t_J = (t_i)_{i \in J}$ are independent of θ , since the type of the remaining player is an independent uniform pad; only the grand group knows the state.

For a finite common prior type space I , say that ϕ is a *belief invariant Bayesian solution* of (I, G) if ϕ satisfies the conditions that define a group belief invariant BCE of $(\mathbb{I}(I), G)$ for the singleton groups and for the grand group, but not necessarily for the intermediate groups; as discussed in Section 3.2, these are the obedience, belief invariance, and Bayesian solution conditions of Forges (1993, 2006). By definition, every group belief invariant BCE of (I, G) is a belief invariant Bayesian solution of (I, G) , and when $N = 2$ the converse holds as well. The next result shows that, with three players, the converse fails, even at the level of outcomes.

Proposition 10. *In the game structure G_3 with the finite common prior type space I^{X^3} :*

- (i) *For every group belief invariant BCE of (I^{X^3}, G_3) and every group $J \neq \{1, 2, 3\}$, the induced outcome has a_J and θ independent.*
- (ii) *There is a belief invariant Bayesian solution of (I^{X^3}, G_3) whose marginal ρ^p on $A \times \Theta$ satisfies $a_1 \oplus a_2 = \theta$ with probability one. In particular, ρ^p is not the marginal of any group belief invariant BCE of (I^{X^3}, G_3) .*

Proof. (i) Fix a group belief invariant BCE ϕ and a group $J \neq \{1, 2, 3\}$. By the invariance condition for J , conditional on t_J , a_J and θ are independent. Moreover, t_J and θ are independent under σ : for singletons this is immediate, and for pairs it was noted above. Hence, as in the proof of Proposition 9(ii), a_J and θ are independent under the marginal of ϕ .

(ii) Let c be uniformly distributed on $\{0, 1\}$, independently of (t, θ) , and consider the recommendations $a_1 = c$, $a_2 = c \oplus \theta$, and $a_3 = 0$; let ϕ^p denote the joint distribution of (a, t, θ) . Belief invariance holds: conditional on any realization of (t, θ) , the distribution of a_1 is uniform, the distribution of a_2 is uniform, and a_3 is constant, so each player's action is independent of everything else conditional on the player's own type. The Bayesian solution condition for the grand group holds trivially, since the grand group's type (t_1, t_2, t_3) determines θ . For obedience: since $a_1 \oplus a_2 \oplus a_3 = c \oplus c \oplus \theta \oplus 0 = \theta$, the players succeed with probability one, so, conditional on every information set (t_i, a_i) , obeying succeeds with probability one and switching fails with probability one. Thus ϕ^p is a belief invariant Bayesian solution. Its marginal ρ^p satisfies $a_1 \oplus a_2 = \theta$ with probability one, so $a_{\{1,2\}}$ and θ are not independent, and, by part (i), ρ^p is not the marginal of any group belief invariant BCE. $\square$

The refinement operates exactly where the intuition of the two-player example suggests. The pair $\{1, 2\}$ jointly knows nothing about the state, and group belief invariance forbids the pair's joint behavior from revealing it; the belief invariant Bayesian solution, which polices only the individual players and the grand group, permits the pair's actions to encode the state perfectly, with the encoding hidden from each player individually by the shared coin. The refinement does not come at the cost of efficiency: the first-best outcome, in which the players always succeed, is attained in (I^{X^3}, G_3) by the group belief invariant—indeed, Bayes Nash—strategies $a_i = t_i$.

5 Conclusion

In our view, when an analyst writes down a type space, they are intentionally making assumptions about not just what individuals may know, but what groups of individuals may know collectively. In particular, the analyst may wish to explicitly limit how much groups of individuals may jointly know about the state. Unfortunately, if one tries to represent distributed information with belief hierarchies in the universal type space, then all limits on players' joint information are lost. We have shown that one can recover those limitations, in a canonical space of higher order beliefs, as long as one explicitly models group beliefs. We have also characterized the minimal solution concept that can be formulated as a function of group beliefs, without ruling out Bayes Nash equilibria. It is our hope that these observations and results will encourage adoption of group belief invariant BCE as a tool for the analysis of strategic behavior under incomplete information.

References

- BERGEMANN, D. AND S. MORRIS (2013): “Robust Predictions in Games with Incomplete Information,” *Econometrica*, 81, 1251–1308.
- (2016): “Bayes Correlated Equilibrium and the Comparison of Information Structures in Games,” *Theoretical Economics*, 11, 487–522.
- BERGEMANN, D., S. MORRIS, AND R. VEIEL (2025): “A Strategic Topology on Information Structures,” Working paper, Cowles Foundation Discussion Paper No. 2413, Yale University.
- BRANDENBURGER, A. AND E. DEKEL (1993): “Hierarchies of Beliefs and Common Knowledge,” *Journal of Economic Theory*, 59, 189–198.
- BROOKS, B., S. DU, AND A. HABERMAN (2026): “Robust Predictions with Bounded Information,” Working paper, University of Chicago, University of California—San Diego, and Stanford University.
- BROOKS, B., A. FRANKEL, AND E. KAMENICA (2022): “Information Hierarchies,” *Econometrica*, 90, 2187–2214.
- ELY, J. C. AND M. PESKI (2006): “Hierarchies of Belief and Interim Rationalizability,” *Theoretical Economics*, 1, 19–65.
- FORGES, F. (1993): “Five legitimate definitions of correlated equilibrium in games with incomplete information,” *Theory and decision*, 35, 277–310.
- (2006): “Correlated Equilibrium in Games with Incomplete Information Revisited,” *Theory and Decision*, 61, 329–344.
- HARSANYI, J. C. (1967): “Games with Incomplete Information played by “Bayesian” players, I–III Part I. The basic model,” *Management science*, 14, 159–182.
- LIU, Q. (2015): “Correlation and common priors in games with incomplete information,” *Journal of Economic Theory*, 157, 49–75.
- MERTENS, J.-F. AND S. ZAMIR (1985): “Formulation of Bayesian Analysis for Games with Incomplete Information,” *International Journal of Game Theory*, 14, 1–29.
- TANG, Q. (2015): “Hierarchies of Beliefs and the Belief-Invariant Bayesian Solution,” *Journal of Mathematical Economics*, 59, 111–116.