COMPOSITIO MATHEMATICA

There is an equivalence relation whose von Neumann algebra is not Connes embeddable

Aareyan Manzoor

Compositio Math. **162** (2026), 1851–1868.

doi: 10.1017/S0010437X26103431

FOUNDATION
COMPOSITIO
MATHEMATICA



LONDON
MATHEMATICAL
SOCIETY
EST. 1865



https://doi.org/10.1017/S0010437X26103431 Published online by Cambridge University Press

Compositio Math. 162 (2026) 1851–1868
doi:10.1017/S0010437X26103431

There is an equivalence relation whose von Neumann algebra is not Connes embeddable

Aareyan Manzoor

Abstract

The landmark quantum complexity result MIP* = RE has been used to prove the existence of a non-Connes-embeddable tracial von Neumann algebra. Recently, similar ideas have been used to give a negative solution to the Aldous–Lyons conjecture: there is a non-co-sofic invariant random subgroup (IRS) on any non-abelian free group. We define a notion of hyperlinearity for an IRS and show that there is a non-co-hyperlinear IRS on any non-abelian free group. As a corollary, we prove that there is a relation whose von Neumann algebra is not Connes embeddable. We do this by significantly simplifying the reduction of Aldous–Lyons to non-local games, removing the need for subgroup tests entirely.

1. Introduction

One of the biggest problems in von Neumann algebras has been the Connes embedding problem (CEP), which asks if each tracial von Neumann algebra embedded into an ultrapower R^U of the hyperfinite II_1 factors R in a trace-preserving way. This is a finite-dimensional approximation property. This problem has been famous because it turns out to be equivalent to several different problems all over operator algebras, and also in quantum information theory. For detailed discussions on different equivalent formulations of this problem, see [Gol21].

The landmark quantum complexity result MIP* = RE [JNV+21] shows a negative solution to the CEP through a long chain of equivalences. This proof is very non-constructive, essentially showing that if each tracial von Neumann algebra is Connes embeddable, then the Halting problem would be decidable. The idea is that the CEP is equivalent to a separation-of-convex-sets problem (see Proposition 2.15) and certain nice convex sets can be separated using a nice computability trick involving linear functionals (see the proof of Theorem 3.13). As this proof is non-constructive, it can be hard to say anything more about non-Connes-embeddable factors, as there is no nice counterexample on which to test properties.

However, it has been possible to use the complexity result MIP* = RE to build non-Connes-embeddable II_1 factors with additional structure. For example, Goldbring and Hart [GH24] use MIP* = RE to prove the existence of non-Connes-embeddable II_1 factors satisfying certain classes of model-theoretically nice properties. We will use the recent strengthening of

Received 24 February 2025, accepted in final form 2 April 2026, published online 30 September 2026.
2020 Mathematics Subject Classification 46L10 (primary), 20F69, 81P99 (secondary).
Keywords: Connes embedding; Aldous–Lyons; invariant random subgroup; measurable equivalence relations.

© The Author(s), 2026. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (https://creativecommons.org/licenses/by/4.0/), which permits unrestricted re-use, distribution, and reproduction in any medium, provided the original article is properly cited. *Compositio Mathematica* is © Foundation Compositio Mathematica.

https://doi.org/10.1017/S0010437X26103431 Published online by Cambridge University Press

$\mathrm{MIP}^* = \mathrm{RE}$, the so-called $\mathrm{TailoredMIP}^* = \mathrm{RE}$ [BCV24], to get the main result in Theorem 1.1. $\mathrm{TailoredMIP}^* = \mathrm{RE}$ has also been used to prove that both satisfying Lück's determinant conjecture [Man25] and Gottshack's surjunctivity conjecture [BC25] are distinct from soficity for invariant random subgroups (IRSs).

To each group Γ, one can associate a tracial von Neumann algebra $L(\Gamma)$, and one question is whether the class of group von Neumann algebras has a non-Connes-embeddable algebra. The groups whose von Neumann algebra is Connes embeddable are called hyperlinear. Another finite-approximation property for groups is soficity. Roughly speaking, hyperlinearity gives the group an approximation by finite-dimensional matrix groups, while soficity gives an approximation by finite groups. Each sofic group is, in particular, hyperlinear [ES05]. So an easier problem than asking 'Is there a non-hyperlinear group?' is 'Is there a non-sofic group?' While these computability techniques cannot solve these problems yet on the level of groups, they can solve them at the level of group actions.

Theorem 1.1. *There is some property (T) ergodic countable probability-measure-preserving (p.m.p.) Borel equivalence relation $\mathcal{R}$ so that the property (T) II_1 factor $L(\mathcal{R})$ is not Connes embeddable.*

To understand how $\mathrm{TailoredMIP}^* = \mathrm{RE}$ can be used to solve this, we first recall how $\mathrm{MIP}^* = \mathrm{RE}$ has been used to solve the CEP.

1.1 Separation of convex sets

Fix $\Gamma = \mathbb{F}_{n,2^m} := \langle u_1 \cdots u_n : u_i^{2^m} = 1 \rangle$ as the group generated by n free elements of order 2^m. Let $T(\Gamma)$ denote the compact convex set of traces (i.e. positive-definite conjugation-invariant functions) on Γ, and let $AT(\Gamma) \subseteq T(\Gamma)$ be the subset of amenable traces. By work of Kirchberg, if $AT(\Gamma) \neq T(\Gamma)$ for some (n, m), then there exists a non-Connes-embeddable II_1 factor (see Proposition 2.15).

We will witness the strict inclusion via functionals corresponding to non-local games. It is not important for this sketch to know what non-local games are, only that their strategies come from traces on some Γ (the details are discussed in § 3). So for each non-local game $\mathfrak{G}$ with question set of size n and answer set of size m, and each trace τ on Γ, we can define $\mathrm{Val}(\mathfrak{G}, \tau) \in [0, 1]$ as the winning probability under the strategy corresponding to τ. Define

$$\omega^{co}(\mathfrak{G}) = \sup_{\tau \in T(\Gamma)} \mathrm{Val}(\mathfrak{G}, \tau), \quad \omega^*(\mathfrak{G}) = \sup_{\tau \in AT(\Gamma)} \mathrm{Val}(\mathfrak{G}, \tau).$$

Thus a strict gap $\omega^{co}(\mathfrak{G}) > \omega^*(\mathfrak{G})$ produces a trace in $T(\Gamma) \setminus AT(\Gamma)$, and hence a non-Connes-embeddable factor.

Computability enters because both quantities admit effective one-sided approximations: there are computably enumerable lower bounds $\alpha_d(\mathfrak{G}) \nearrow \omega^*(\mathfrak{G})$ and computably enumerable upper bounds $\beta_d(\mathfrak{G}) \searrow \omega^{co}(\mathfrak{G})$. The theorem $\mathrm{MIP}^* = \mathrm{RE}$ [JNV+21] implies that $\omega^*(\mathfrak{G})$ is not computable as a function of $\mathfrak{G}$.

If one had $\omega^{co}(\mathfrak{G}) = \omega^*(\mathfrak{G})$ for all games, then running the procedures for α_d and β_d in parallel and waiting until $\beta_d(\mathfrak{G}) - \alpha_d(\mathfrak{G}) < 2^{-k}$ would yield an algorithm to approximate $\omega^*(\mathfrak{G})$ to precision 2^{-k}, contradicting uncomputability. Therefore, there exists some game $\mathfrak{G}$ with $\omega^{co}(\mathfrak{G}) > \omega^*(\mathfrak{G})$, and hence some (n, m) with $AT(\Gamma) \neq T(\Gamma)$.

1.2 Invariant random subgroups

An IRS of a group Γ is a Borel conjugation-invariant probability measure on its set of subgroups. These objects generalize the notion of a normal subgroup, as the Dirac measure of a normal

https://doi.org/10.1017/S0010437X26103431 Published online by Cambridge University Press

subgroup is an IRS. Since normal subgroups of the free groups correspond to all groups by taking quotients, IRSs on a free group generalize the notion of a group. The notion of a co-sofic IRS is the generalization of a sofic group in this framework (see Proposition 2.6).

The existence of a non-co-sofic IRS on a free group has been proven in [BCL+24, BCV24], disproving the Aldous–Lyons conjecture. To do this, the authors first, in [BCL+24], define subgroup tests, a type of combinatorial game that is very natural probability theoretically. IRSs on free groups correspond to strategies on these, and by the same technique discussed for the CEP, the authors reduce this problem to a computability problem. Then they reduce that computability problem to proving a stronger version of $\mathrm{MIP}^* = \mathrm{RE}$, and in the second paper they proved this computability result, $\mathrm{TailoredMIP}^* = \mathrm{RE}$ [BCV24].

The reduction from subgroup tests in [BCL+24] to non-local games is very long and technical. Fortunately, our approach avoids this entirely, and directly disproves Aldous–Lyons by non-local games. Taking [BCV24] as a black box, we are able to reprove the main result of [BCL+24]: we only use some minor results from that work to prove the computable upper bounds for $\omega_{\mathrm{IRS}}(\mathfrak{G})$ in § 3.2; otherwise, this paper is independent of [BCL+24].

The idea is that each IRS on a free group naturally corresponds to a trace on it:

$$\tau_H(g) = \mathbb{P}(g \in H).$$

This then naturally corresponds to strategies on non-local games. This allows us to define $\omega_{\mathrm{IRS}}(\mathfrak{G})$ to be the best winning strategy over traces coming from IRSs (with a minor technicality). This directly converts Aldous–Lyons to a problem about non-local games, without the need to go through subgroup tests of [BCL+24].

Our approach also allows us to upgrade soficity to hyperlinearity. We define, for each IRS H on a group Γ, a von Neumann algebra $L(\Gamma/H)$ generalizing the notion of quotienting against a normal subgroup. It is then natural to call an IRS H of a free group co-hyperlinear if $L(\Gamma/H)$ is Connes embeddable. We prove the following.

Theorem 1.2. *There is some free group with a non-co-hyperlinear IRS.*

In particular, each $L(\Gamma/H)$ is a subalgebra of $L(\mathcal{R})$ (see the discussion after Definition 2.12) for some countable Borel equivalence relation $\mathcal{R}$. This will allow us to recover Theorem 1.1 from Theorem 1.2.

2. Invariant random subgroups

This subsection will set up the terminology we use for IRSs. Most of these are standard definitions; compare with [BCL+24].

Definition 2.1. Let Γ be a discrete group. Let $\mathfrak{sub}(\Gamma)$ denote the space of subgroups of Γ topologized as a subset of $\{0,1\}^\Gamma$. Define a *random subgroup* as a Borel probability measure on $\mathfrak{sub}(\Gamma)$. Furthermore, if the measure is invariant under the conjugation action of Γ on $\mathfrak{sub}(\Gamma)$, then we call it an *invariant random subgroup (IRS)*. We will denote by $\mathrm{IRS}(\Gamma)$ the convex set of IRSs on Γ.

Note that $\mathfrak{sub}(\Gamma)$ is compact, so we can talk about the weak* topology on $\mathrm{IRS}(\Gamma)$. The topology on $\mathfrak{sub}(\Gamma)$ can be understood as follows: a net of subgroups H_λ converging to H means that

$$g \in H \iff g \in H_\lambda \quad \text{eventually}.$$

https://doi.org/10.1017/S0010437X26103431 Published online by Cambridge University Press

The topology on IRS(Γ) can be understood as follows: a net μ_λ converges to μ if for each continuous $f : \mathfrak{sub}(\Gamma) \to \mathbb{R}$, we have

$$\int f \, d\mu_\lambda \to \int f \, d\mu.$$

Note, in particular, that this is compact by the Banach–Alaoglu theorem.

Here are some examples of IRSs.

Example 2.2. Let Γ be a discrete group. Let $N \subset \Gamma$ be a normal subgroup. Then the point mass δ_N is an IRS.

So we should think of an IRS as a generalization of a normal subgroup of Γ. We will only really care about IRSs of free groups, and since normal subgroups of free groups correspond to all finitely generated groups by looking at the quotients, we will consider these IRSs as a generalization of the notion of a group by some generalized notion of a quotient (see Definition 2.10).

Example 2.3. Let (X, μ) be a probability space, and let $\alpha : \Gamma \to \mathrm{Aut}(X, \mu)$ be a p.m.p. action. We can then describe an IRS as follows: first sample a $x \in X$ according to μ, and then look at

$$\mathrm{Stab}(x) = \{g \in \Gamma : gx = x\}.$$

We will call the induced probability measure on $\mathfrak{sub}(\Gamma)$ by $\mathrm{Stab}(\alpha)$. Formally, if we look at the map $\mathrm{Stab} : X \to \mathfrak{sub}(\Gamma)$ sending a point to its stabilizer, $\mathrm{Stab}(\alpha)$ is the pushforward of μ under this map.

Note that the action α being p.m.p. gives conjugation invariance of the $\mathrm{Stab}(\alpha)$.

A remarkable result is that all IRSs are of the following form.

Theorem 2.4 [AGV14, Proposition 14]. *Let Γ be a countable discrete group, and let $H \in$ IRS(Γ). Then there is a p.m.p. action α of Γ on some compact space X with Borel probability measure μ such that $H = \mathrm{Stab}(\alpha)$. If H was ergodic, then α can be picked to be ergodic.*

We can now define the following.

Definition 2.5. Let Γ be a discrete group, and let α be a p.m.p. action of it on a finite set. Then the IRS $\mathrm{Stab}(\alpha)$ will be called *finitely described*. A weak* limit of finitely described IRSs will be called co-sofic.

The name is motivated by the following.

Proposition 2.6 [AGN17, Lemma 16]. *Let Γ be a free group, and let N be a normal subgroup. Then δ_N is co-sofic as an IRS if and only if Γ/N is sofic as a group.*

So if we think of IRSs of free groups as an extension of the notion of a group, we should think of co-sofic IRSs of the free group as an extension of the notion of a sofic group. One of the biggest open problems in this area is: Is every group sofic? What was proven in [BCL+24] is as follows.

Theorem 2.7. *Let Γ be a non-abelian free group. Then there is a non-co-sofic IRS on it.*

This, of course, does not prove the existence of non-sofic groups: indeed, this non-co-sofic IRS does not have to be a point-mass.

https://doi.org/10.1017/S0010437X26103431 Published online by Cambridge University Press

A slightly harder problem than finding a non-sofic group is that of finding a non hyperlinear group. The rest of the paper is dedicated to setting up what co-hyperlinearity should mean for IRSs and proving that there is a non-co-hyperlinear IRS on a free group.

2.1 The algebra of an IRS

We will formalize the notion of quotienting against an IRS in this section. The idea is to go through traces on the universal group C^* algebra $C^*(\Gamma)$ of Γ. This is a C^*-algebra densely containing the group ring $\mathbb{C}[\Gamma]$ that encodes all the representation-theoretic data of Γ. More precisely, there is a correspondence between representations of $C^*(\Gamma)$ on Hilbert spaces and unitary representations of Γ on the same Hilbert space (for exposition, see [BO08, § 2.5]).

Definition 2.8. Let Γ be a discrete group, and let $H \in \mathrm{IRS}(\Gamma)$. Then one can define a trace τ_H on $C^*(\Gamma)$ by

$$\tau_H(g) := \mathbb{P}(g \in H).$$

Note that invariance gives τ_H conjugation invariance, but one still needs to check positivity. We get this from Proposition 2.11. If $H = \mathrm{Stab}(\alpha)$, where α is a p.m.p. action of Γ on some Borel probability space (X, μ), then

$$\tau_H(g) = \mu(x : gx = x).$$

In group theory, these objects are called characters of Γ (positive-definite conjugation-invariant functions). Naimark's dialation theorem gives this equivalence (for exposition, see [Pau03, Theorem 4.8]).

Denote the trace complex of Γ by $T(\Gamma)$, and equip it with the weak* topology. This is the topology of pointwise convergence on $C^*(\Gamma)$, which is the same as pointwise convergence when restricted to Γ. Then note that there is an affine map $\mathrm{IRS}(\Gamma) \to T(\Gamma)$ given by $\mathrm{IRS}(\Gamma) \ni H \mapsto \tau_H \in T(\Gamma)$. This map is neither injective nor surjective generally.

Not being surjective is clear, as a trace $\tau \in T(\Gamma)$ can be negative on a group element, which is not the case for τ_H for $H \in \mathrm{IRS}(\Gamma)$. For non-injectivity consider S_3, which has four conjugacy classes of subgroups but $T(S_3)$ has three irreducible characters. An IRS on a finite group is simply a probability measure on the set of conjugacy class of subgroups, while a character on a finite group is in the convex hull of the irreducible characters. So a dimension argument forces this map to not be injective. This map is, however, $w^* - w^*$ continuous, as $\{H : g \in H\}$ is a clopen set making its indicator continuous. So the image of this, call it $T_{\mathrm{IRS}}(\Gamma)$, is compact.

One of the main themes in this paper will be to look at the von Neumann algebra that corresponds to a trace. This is done by looking at the Gelfand–Naimark–Segal (GNS) representation of a trace.

Let $\mathcal{A}$ be a C^*-algebra and let τ be a trace on it. Then the GNS representation of τ is a representation $\pi_\tau : \mathcal{A} \to B(H_\tau)$ and a vector $\xi \in H_\tau$ with $\tau(a) = \langle \xi | \pi(a) | \xi \rangle$. The von Neumann algebra of τ will be $\pi_\tau(\mathcal{A})'' \subset B(H_\tau)$, and note that τ extends to this algebra by

$$\tilde{\tau}(a) = \langle \xi | \pi(a) | \xi \rangle, \quad a \in \pi_\tau(\mathcal{A})''.$$

This algebra is well defined due to uniqueness of GNSs: any two cyclic representations inducing τ will be unitarily equivalent. We will denote this von Neumann algebra as $\tau(\mathcal{A})''$. If $\mathcal{A} = C^*(\Gamma)$, we will denote this as $\tau(\Gamma)''$. It will be implicit that we are considering this von Neumann algebra with the extension of τ as its trace.

Consider the left regular representation $\lambda : \Gamma \to U(\ell^2(\Gamma))$, where $(\lambda(g)f)(h) = f(g^{-1}h)$ for $g, h \in \Gamma$. Then we define the group von Neumann algebra to be $\lambda(\Gamma)'' \subset B(\ell^2(\Gamma))$. We consider it with the trace $\tau(g) = 1$ if $g = e$ and 0 otherwise on group elements.

https://doi.org/10.1017/S0010437X26103431 Published online by Cambridge University Press

Before defining our IRS algebras, let us make one observation.

Proposition 2.9. *Let Γ be a discrete group and let N be a normal subgroup. If τ_N is the trace on Γ corresponding to the point mass $\delta_N \in \mathrm{IRS}(\Gamma)$, then*

$$\tau_N(\Gamma)'' = L(\Gamma/N),$$

where $L(\Gamma/N)$ is the group von Neumann algebra of Γ/N with its standard trace.

Proof. First note that $\tau_N(g) = 1$ if $g \in N$ and 0 otherwise. Now take the representation:

$$\pi : \Gamma \to U(\ell^2(\Gamma/N)),$$

where $(\pi(g)f)([h]) = f([g^{-1}h])$ for $g \in \Gamma, [h] \in \Gamma/N$. Note that all elements of N act trivially, so this is actually a homomorphism $\tilde{\pi} : \Gamma/N \to U(\ell^2(\Gamma/N))$. Note that this is just the left regular representation of Γ/N.

Fix the vector $\xi \in \ell^2(\Gamma/N)$ that is 1 on the identity coset and 0 everywhere else. Then it is clear that

$$\tau_N(g) = \langle \xi | \pi(g) | \xi \rangle.$$

So the corresponding algebra is $\pi(\Gamma)'' = \tilde{\pi}(\Gamma/N)''$ and by definition this is the group von Neumann algebra $L(\Gamma/N)$. Note that the extension of the trace is 1 on elements of N and 0 on every other group element, i.e. the standard trace on the group von Neumann algebra. □

Based on this, we can now define the quotient of a group by an IRS.

Definition 2.10. Let Γ be a discrete group and let $H \in \mathrm{IRS}(\Gamma)$. Then define the *quotient of* Γ *by* H as the tracial von Neumann algebra:

$$L(\Gamma/H) := \tau_H(\Gamma)''.$$

A natural question now is this: Can we write the GNS representation of τ_H explicitly? The answer is yes.

To see this, let $H \in \mathrm{IRS}(\Gamma)$ correspond to the p.m.p. action α of Γ on some (x, μ). Let $\mathcal{R}_\alpha \subset X \times X$ be the relation defined by the orbits, i.e. $x \equiv gx$. Let $\pi_0 : \mathcal{R}_\alpha \to X$ be given by $(x, y) \mapsto x$ and let Γ act on $\mathcal{R}_\alpha$ on the first coordinate, i.e. $g(x, y) = (gx, y)$. Define

$$\nu(E) := \int |\pi_0^{-1}(x) \cap E| \, d\mu(x).$$

This is a σ-finite measure on $\mathcal{R}_\alpha$ [AP17, § 1.5].

Proposition 2.11 [Ver11, Theorem 9]. *Let H be an IRS of Γ. Suppose that it corresponds to p.m.p. action α of Γ on (X, μ), i.e $H = \mathrm{Stab}(\alpha)$. Let $\pi : \Gamma \to U(L^2(\mathcal{R}_\alpha, \nu))$ correspond to the action described above. Let $\Delta \subset \mathcal{R}_\alpha$ be the diagonal of X. Then*

$$\tau_H(g) = \langle 1_\Delta | \pi(g) | 1_\Delta \rangle.$$

In particular, the algebra $L(\Gamma/H) = \pi(\Gamma)''$.

Proof. Note that $\pi(g)1_\Delta$ is the characteristic function on $g^{-1}\Delta = \{(g^{-1}x, x) : x \in X\}$. So

$$\begin{aligned}\langle 1_\Delta | \pi(g) | 1_\Delta \rangle &= \int_{\mathcal{R}} 1_\Delta 1_{g^{-1}\Delta} \, d\nu \\ &= \int_{\mathcal{R}} 1_{\{(x,x) : x = gx\}} \, d\nu \\ &= \nu\{(x, x) : x = gx\}\end{aligned}$$

https://doi.org/10.1017/S0010437X26103431 Published online by Cambridge University Press

$$= \int_X |\pi_0^{-1}(x) \cap \{(x,x) : x = gx\}| \, d\mu$$
$$= \int_X 1_{x=gx} \, d\mu$$
$$= \mu(x : gx = x). \qquad \square$$

This can be extended to a larger algebra.

Definition 2.12. Let α be a p.m.p. action of Γ on (X, μ), and let $(\mathcal{R}_\alpha, \nu)$ be as in Proposition 2.11. For $f \in L^\infty(X, \mu)$, define $M_f \in B(L^2(\mathcal{R}_\alpha, \nu))$ as

$$(M_f \xi)(x, y) = f(x)\xi(x, y).$$

Then the von Neumann algebra generated by $\pi(\Gamma)$ and $\{M_f : f \in L^\infty(X, \mu)\}$ is defined to be von Neumann algebra $L(\mathcal{R}_\alpha)$ of the relation $\mathcal{R}_\alpha$. We consider this with the trace

$$\tau_\alpha(a) = \langle 1_\Delta | a | 1_\Delta \rangle, \quad a \in L(\mathcal{R}_\alpha)$$

This can be defined for any measurable relation on a Borel space, but essentially reduces to this case. These algebras are useful as the inclusion $L^\infty(X) \subset L(\mathcal{R}_\alpha)$ is a complete invariant for the action α (up to orbit equivalence). These are also canonical examples of Cartan Inclusions. Note, in particular, $L(\Gamma/\mathrm{Stab}(\alpha)) \subset L(\mathcal{R}_\alpha)$. Also note that if α is ergodic, then $L(\mathcal{R}_\alpha)$ is a II_1 factor. For a discussion of all of these, see [AP17, §§ 1.5 and 12.1].

2.2 Co-hyperlinear IRS

Now that we can talk about IRSs algebraically, it is easy to define co-hyperlinearity for them. Recall that a group Γ is called hyperlinear if its von Neumann algebra $L(\Gamma)$ is Connes embeddable, i.e. there is a trace-preserving embedding $L(\Gamma) \hookrightarrow R^\omega$ of the group algebra into an ultrapower of the hyperfinite II_1 factor. So we want co-hyperlinearity of a $H \in \mathrm{IRS}(\Gamma)$ to line up with Connes embeddability of $L(\Gamma/H)$.

A crucial point for us is that a tracial von Neumann algebra being Connes embeddable is something we can express in terms of traces.

Definition 2.13. A trace τ on a C* algebra $\mathcal{A}$ is called amenable if there are contractively completely positive (c.c.p.) maps $\phi_n : \mathcal{A} \to M_{k_n}(\mathbb{C})$ with

$$\|\phi_n(ab) - \phi_n(a)\phi_n(b)\|_2 \to 0, \quad \forall a, b \in \mathcal{A},$$

so that $\tau(a) = \lim \mathrm{tr}_{k_n} \circ \phi_n(a)$. Here, $\|A\|_2^2 = \mathrm{tr}_n(A^*A)$ is the normalized Hilbert–Schmidt norm.

This is a finite-dimensional approximation property for the trace, and it turns out that the corresponding finite-dimensional approximation property for the algebra is Connes embeddability. We state the following proposition for groups satisfying either the local lifting property or Hilbert–Schmidt stability. The important fact is that free groups have both of these properties (see [BO08, Theorem 13.1.3] and [HS18, Theorem 1]).

Proposition 2.14. *Let Γ be a group such that $C^*(\Gamma)$ has the local lifting property, and let τ be a trace on it. The following are equivalent by [Bro06, Proposition 6.3.4]:*

(1) *$\tau(\Gamma)''$ has a trace-preserving embedding into R^ω; and*
(2) *τ is amenable.*

In addition, if Γ is Hilbert–Schmidt stable, then the above are also equivalent to the following by [HS18, Theorem 3].

https://doi.org/10.1017/S0010437X26103431 Published online by Cambridge University Press

(3) *There is a sequence of finite-dimensional representations* $\pi_n : \Gamma \to U_n(\mathbb{C})$, *with* τ *being the weak* limit of* $\mathrm{tr}_n \circ \pi_n$.

(4) *There is a sequence of traces* τ_n *with finite-dimensional GNS representation whose weak* limit is* τ.

Here, R^ω *is an ultrapower of the hyperfinite* II_1 *factor.*

Note that the above proposition is not true for a general C^*-algebra $\mathcal{A}$. In general, amenability of the trace means that $\tau(A)''$ has an embedding into R^ω with a completely positive lift onto $\prod_{\mathbb{N}} R$. The local lifting property gives this lift for free. For a detailed discussion of the local lifting property, see [BO08, § 13.1]. We will use this property in our proof of Theorem 3.5 in a minor way. Hilbert–Schmidt stability allows the completely positive lift of the embedding of $\tau(A)'' \hookrightarrow R^\omega$ to actually be a $*$-homomorphism. This gives the equivalence of (1) and (3).

Let us call the collection of traces with finite-dimensional GNS representation the finite-dimensional traces of Γ, and let us denote this set by $T_{f.d}(\Gamma) \subset T(\Gamma)$.

We also remark that while the traces are called amenable, the corresponding algebras are not amenable von Neumann algebras. For that, one would need to restrict to a class of traces called uniformly amenable traces.

Of course, one of the nice things about the free group is that it is an universal object. It turns out that every tracial von Neumann algebra arises from a trace on some free group (the extreme points correspond to factors). So the fact that there is a non-Connes-embeddable tracial von Neumann algebra can be stated as follows.

Proposition 2.15. *There is a non-amenable trace on each non-abelian free group. That is, if* Γ *is free, then* $T(\Gamma) \neq \overline{T_{\mathrm{f.d.}}(\Gamma)}^{w^*}$.

For more details, see [Bro06, Proposition 6.3.5]. It is now clear how to generalize the notion of hyperlinearity to IRSs.

Definition 2.16. Let Γ be a discrete group, and let $H \in \mathrm{IRS}(\Gamma)$. Then call H *co-hyperlinear* if the corresponding trace τ_H is amenable.

Proposition 2.17. *Let* Γ *be a free group, and let* $H \in \mathrm{IRS}(\Gamma)$. *Then:*

(1) H *is co-hyperlinear if and only if* $L(\Gamma/H)$ *is Connes embeddable; and*

(2) $H = \delta_N$ *is co-hyperlinear if and only if* Γ/N *is hyperlinear as a group.*

Let us make explicit what the difference is between co-hyperlinear and co-sofic.

Lemma 2.18. *Let* Γ *be a discrete group, and let* $H \in \mathrm{IRS}(\Gamma)$ *be ergodic.*

Then H *being finitely described implies that* τ_H *is a finite-dimensional trace.*

Furthermore, if Γ *is finitely generated, then* τ_H *being a finite-dimensional trace implies that* H *is finitely described.*

Proof. ($\Longrightarrow$) Let $H = \mathrm{Stab}(\alpha)$, where α is the action of Γ on a finite set with the uniform measure. Then the orbit relation $\mathcal{R}_\alpha \subset X \times X$ is also a finite set and hence by Proposition 2.11 there is a finite-dimensional representation $\pi : \Gamma \to U(L^2(\mathcal{R}_\alpha))$ inducing τ_H.

($\Longleftarrow$) We can find an ergodic action α of Γ on some probability space (X, μ) with $H = \mathrm{Stab}(\alpha)$.

Let $F_n \subset X$ denote the set of points with orbits of size less than n. Note that a point having orbit of size n is the same as the stabilizer K having index n. Since the action Γ on Γ/K uniquely determines K (by looking at the stabilizer of the identity coset), the corresponding

https://doi.org/10.1017/S0010437X26103431 Published online by Cambridge University Press

homomorphism $\Gamma \to \mathrm{Sym}(n)$ uniquely determines K. In particular, since Γ is finitely generated, there are finitely many subgroups of index n of Γ. So

$$F_n = \bigcup_{K:|\Gamma/K|\leqslant n} \{x : gx = x \,\forall g \in K\}.$$

This is a finite union of measurable sets, and hence measurable. F_n is Γ invariant and hence has measure 0 or 1. If it has measure 1, then H is supported on subgroups of index $\leqslant n$, and this is a finite set; call it $\mathrm{Supp}(H)$. Now ergodicity of H forces $\mathrm{Supp}(H)$ to be the set of conjugates of a given subgroup K', and this is the IRS induced by the action of Γ on the finite set of cosets Γ/K'.

If all F_n have measure 0, then the action has infinite orbits almost everywhere. By [PT16, Proposition 3.1], the representation $\Gamma \to U(L^2(\mathcal{R}_\alpha, \nu))$ is weakly mixing, i.e. it has no finite-dimensional sub-representations. In particular, τ_H cannot be finite dimensional, ruling this case out. □

This immediately gives us the following.

Proposition 2.19. *Let Γ be a countable discrete group, and let $H \in \mathrm{IRS}(\Gamma)$ be co-sofic. Then H is co-hyperlinear.*

Proof. Let H be co-sofic; then there is a sequence of finitely described IRS (H_n) with $H_n \to H$. Since the map sending an IRS to its trace is $w^* - w^*$ continuous, we have $\tau_{H_n} \to \tau_H$. By Lemma 2.18, we get that each τ_{H_n} is finite dimensional, and so τ_H is a w^* limit of finite-dimensional traces. By Definition 2.13, since these finite-dimensional traces factor through matrix algebras by $*$-homomorphisms (which are, in particular, c.c.p. maps that are approximately multiplicative), we have that τ_H is amenable. □

Let Γ be a free group. Note that showing a separation between co-sofic IRSs and all IRSs is implied by showing a separation between $T_{\mathrm{IRS}}(\Gamma)$ and $\overline{T_{\mathrm{IRS}}(\Gamma) \cap T_{\mathrm{f.d.}}(\Gamma)}$. On the other hand, showing a separation between co-hyperlinear IRSs and all IRSs is showing a separation between $T_{\mathrm{IRS}}(\Gamma)$ and $T_{\mathrm{IRS}}(\Gamma) \cap \overline{T_{\mathrm{f.d.}}(\Gamma)}$. The latter is an *a priori* harder problem, as one might be able to approximate an IRS trace with finite-dimensional traces, but not finite-dimensional traces arising from IRSs. Actually, these notions agreeing would imply that hyperlinearity and soficity agree for groups, by specializing to Dirac measures.

We also remark that the set of co-hyperlinear IRSs of free group Γ is a w^* closed convex subset of $\mathrm{IRS}(\Gamma)$ containing the co-sofic IRSs. This is because it is the pre-image of $T_{\mathrm{IRS}}(\Gamma) \cap \overline{T_{\mathrm{f.d.}}(\Gamma)}$ under the affine continuous $H \mapsto \tau_H$.

3. Non-local games

To each trace on a certain group C^* algebra to be defined below, one can associate a strategy on non-local games. Since we have associated a trace to each IRS, we can associate a strategy to each IRS for some non-local game. For a less bare-bones discussion of the topic, see [Gol21, §§ 4 and 5].

Definition 3.1. A non-local game $\mathfrak{G}$ has the following parameters: a finite question set Q, a probability distribution q on $Q \times Q$, an integer m describing the length of the answer, and a decider function $D : (\{0,1\}^m)^2 \times Q^2 \to \{0,1\}$. We will write this as $D(a, b|x, y)$ for $a, b \in \{0,1\}^m$ and $x, y \in Q$.

https://doi.org/10.1017/S0010437X26103431 Published online by Cambridge University Press

We will also associate to this game a set

$$S(\mathfrak{G}) = \{u_{x,i} : x \in Q, 1 \leqslant i \leqslant m\}.$$

Each $u_{x,i}$ is to be regarded as a formal variable that encodes the value of the ith bit in the answer a corresponding to question x. We will simply write S when there is no chance for confusion. We will write S_x for the variables corresponding to question x.

One should think of this as an interactive proof system. There are two players, Alice and Bob, and a verifier. The verifier samples a pair of questions (x, y) according to q from $Q \times Q$, then sends x to Alice and y to Bob. Alice and Bob then both send m-bit answers a and b to the verifier, respectively. The verifier evaluates $D(a, b|x, y)$ and Alice and Bob win when this is 1.

We are using the notation of [BCV24] here, specifically with regard to the variable set. A general game can have different question sets for Alice and Bob, and different answer sets for them too. However, we will not need that level of generality.

3.1 Strategies

Note that the only thing Alice and Bob can do, when asked questions $(x, y) \in Q^2$, is give answers $(a, b) \in (\{0, 1\}^m)^2$. So a strategy for Alice and Bob will be some conditional distribution $(p(a, b|x, y))_{(\{0,1\}^m)^2 \times Q^2}$. We will call these *correlations*. The idea in these games is that Alice and Bob can discuss a strategy beforehand but cannot communicate during the game: so Alice has no idea what question Bob got and vice versa. Therefore, a strategy in which Alice's answer is independent of Bob's answer would be a classical strategy.

However, if Alice and Bob share quantum entanglement, then there can be correlation between their choices of their answer. This works precisely in the following way.

Definition 3.2. Let $\mathfrak{G}$ be a non-local game. A *quantum strategy* is a map $\rho : S(\mathfrak{G}) \to U_n(\mathbb{C})$ for some n with:

(i) $\rho(u)^2 = 1$ for each $u \in S$; and

(ii) $\rho(u)$ and $\rho(v)$ commuting whenever $u, v \in S_x$, i.e. different bits corresponding to the same answer should commute.

Define

$$e_x^a = \prod_{i \leqslant m} \frac{1 + (-1)^{a_i} \rho(u_{x,i})}{2}.$$

This is the product of the spectral projections of each bit onto $(-1)^{a_i}$. Then, the corresponding correlation is defined as

$$p(a, b|x, y) = \mathrm{tr}_n(e_x^a e_y^b).$$

A *quantum approximate correlation* is a pointwise limit of quantum correlations.

The value of a game $\mathfrak{G}$ on strategy ρ, $\mathrm{Val}(\mathfrak{G}, \rho)$ is just the expected value of winning if Alice and Bob answer (a, b) on questions (x, y) with probability $p(a, b|x, y)$. We define the *quantum value* of the game as

$$\omega^*(\mathfrak{G}) := \sup_{\rho} \mathrm{Val}(\mathfrak{G}, \rho)$$

the best probability of winning with a quantum strategy.

Actually, in the literature these are called synchronous quantum approximate strategies, but we will not consider any non-synchronous strategies, so we omit that word. This is also defined

https://doi.org/10.1017/S0010437X26103431 Published online by Cambridge University Press

slightly differently than the standard literature, but Theorem 3.3 shows the equivalence. Note that if the dimension was 1, this would just be a classical correlation, since we would have $p(a,b|x,y) = p(a|x)p(b|y)$.

There are two perspectives: one of them is that we can associate to each bit of the answer a formal variable, to be interpreted as a matrix. This is what we are doing. In this case, the spectral projections of these variables correspond to what bit that variable outputs. On the other hand, one can associate a variable v_x to each question $x \in Q$ with $v_x^{2^m} = 1$. So now v_x will have 2^m spectral projections, and they will correspond to the 2^m different outputs on question x. These perspectives are equivalent.

Let $\mathfrak{G}$ be a non-local game with question set Q, answers of size m, and variable set S. By $\mathcal{F}(S,2)$, we denote the group generated by elements of S with the condition that they have order 2. By $*_Q \mathbb{Z}_2^m$, we will denote the quotient of $\mathcal{F}(S,2)$ by $\{[u_{x,i}, u_{x,j}] : x \in X, i,j \leqslant m\}$. That is, we are imposing the condition that the variables corresponding to the same question commute. Let $U_{x,i}$ denote the group unitaries corresponding to $u_{x,i}$ in $C^*(*_Q \mathbb{Z}_2^m)$, and let e_x^a be the spectral projections as above, for $a \in \mathbb{Z}_2^m$. That is,

$$e_x^a = \prod_{i \leqslant m} \frac{1 + (-1)^{a_i} U_{x,i}}{2}.$$

Note that these projections generate the algebra and have no relation between each other except that $e_x^a e_x^b = \delta_{ab}$ and $\sum_a e_x^a = 1$. Hence we can find an isomorphism $C^*(*_Q \mathbb{Z}_2^m) \cong C^*(\mathcal{F}(Q, 2^m))$ by looking at the 2^m spectral projections of the defining unitaries in the latter. In particular, this algebra has the local lifting property (see [ES24, Corollary 4.6]).

Any set of projections satisfying the same properties as e_x^a are called *projective valued measurements (PVMs)*. One should think of these as 'quantum probabilities' of answering a on question x. To get actual probabilities out of these, one needs to use traces.

Theorem 3.3 [KPS18, Theorem 3.6]. *Let $\mathfrak{G}$ be a non-local game with question set Q and answers of size m. Then p is a quantum approximate correlation if and only if there is an amenable trace τ on $C^*(*_Q \mathbb{Z}_2^m) \cong C^*(\mathcal{F}(Q, 2^m))$ so that*

$$p(a,b|x,y) = \tau(e_x^a e_y^b).$$

If we remove the restriction to amenable traces and allow any trace, then the resulting classes of correlations are called quantum commuting. The idea of $\mathsf{MIP}^* = \mathsf{RE}$ was separating the quantum commuting value and quantum value of some game, and that would give a non-amenable trace on this group.

Of course, we do not want just any old trace to induce our correlations, but specifically the ones coming from IRSs. So let us define the following.

Definition 3.4. Let $\mathfrak{G}$ be a non-local game with question set Q and answers of size m. An *IRS strategy* is defined by the following data: a map $\sigma : S \cup \{J\} \to \mathrm{Aut}(X, \mu)$, where (X, μ) is a Borel probability space. We will require (equality here means that they agree μ-almost everywhere):

- $\mu(x : \sigma(J)x = x) = 0$ and $\sigma(J)^2 = 1$;
- for $u \in \sigma(S)$, we have $u^2 = 1$ and $u\sigma(J) = \sigma(J)u$; and
- for $u, v \in \sigma(S_x)$ we have $uv = vu$.

If μ was the uniform measure on a finite set, call this a *permutation strategy*.

Note that σ extends to a homomorphism $\tilde{\sigma} : *_Q \mathbb{Z}_2^m \times \mathbb{Z}_2 \to \mathrm{Aut}(X, \mu)$. That is, it is a p.m.p. action of the group $\mathbb{Z}_2^m \times \mathbb{Z}_2$ on (X, μ). So we can consider the IRS $\mathrm{Stab}(\tilde{\sigma})$ and then the

https://doi.org/10.1017/S0010437X26103431 Published online by Cambridge University Press

corresponding trace τ on $C^*(*_Q \mathbb{Z}_2^m \times \mathbb{Z}_2) \cong C^*(\mathcal{F}(Q, 2^m)) \otimes C^*(\mathbb{Z}_2)$ has

$$\tau(g) = \mu(x : gx = x).$$

So we can define the correlation

$$p(a, b|x, y) = \tau((1 - J)e_x^a e_y^b).$$

Note that this is positive as $\tau((1 - J)a) = \frac{1}{2}\tau((1 - J)a(1 - J))$ and these sum up to 1 as $\tau(J) = 0$. We will define $\omega_{\text{IRS}}(\mathfrak{G})$ to be the best winning probability over all IRS strategies.

We remark that this definition agrees with that of [BCL+24, Definition 6.11] for permutation strategies. The authors defined permutation strategies as maps $\sigma : S \cup \{J\} \to \text{Sym}(2n)$ with the same conditions, and if W^- is the negative eigenspace of the image of J in $U(2n)$, then they defined the correlation as

$$p(a, b|x, y) = \text{tr}_n(e_x^a e_y^b|_{W^-}).$$

Since $(1 - J)/2$ is the projection onto W^-, this is the same correlation as our definition. This is why we have had to add this $(1 - J)$ to our construction, as without it we would not be able to use the resulting TailoredMIP$^* =$ RE of [BCV24]. What we have done is, in essence, to define the dual of the permutation strategies of [BCL+24], which will allow us a direct separation of the convex sets in question without going through some intermediary-like subgroup tests.

Now, if we had a game with an IRS strategy that was not quantum approximate, then intuitively that should arise from an IRS trace that is not amenable, which is what we want. Formally, we have the following.

Theorem 3.5. *Suppose that there is a game $\mathfrak{G}$ with*

$$\omega_{\text{IRS}}(\mathfrak{G}) > \omega^*(\mathfrak{G}).$$

Then there is a non-co-hyperlinear IRS on some free group.

Proof. Let $\mathfrak{G}$ be a game with $\omega_{\text{IRS}}(\mathfrak{G}) > \omega^*(\mathfrak{G})$. Let $\sigma : S \cup \{J\} \to \text{Aut}(X, \mu)$ be an IRS strategy with $\text{Val}(\mathfrak{G}, \sigma) > \omega^*(\mathfrak{G})$. Let τ be the corresponding trace on $C^*(\mathcal{F}(Q, 2^m)) \otimes C^*(\mathbb{Z}_2)$ so that

$$p(a, b|x, y) = \tau((1 - J)e_x^a e_y^b)$$

is the correlation.

Note that there is a natural inclusion $C^*(\mathcal{F}(Q, 2^m)) \subset C^*(\mathcal{F}(Q, 2^m)) \otimes C^*(\mathbb{Z}_2)$ induced by the natural group inclusion [BO08, Proposition 2.5.8]. Note that $a \mapsto \tau((1 - J)a)$ for $a \in C^*(\mathcal{F}(Q, 2^m)) \otimes C^*(\mathbb{Z}_2)$ is a trace as J commutes with everything. Call this trace τ'. Then $\tau'|_{C^*(\mathcal{F}(Q,2^m))}$ is also a trace, and has

$$p(a, b|x, y) = \tau'|_{C^*(\mathcal{F}(Q,2^m))}(e_x^a e_y^b).$$

Since p is not a quantum approximate correlation, $\tau'|_{C^*(\mathcal{F}(Q,2^m))}$ is not an amenable trace. Since amenability of traces is preserved under restriction [BO08, Proposition 6.3.5], we get that τ' itself is not amenable. We now have

$$\tau(a) = \tfrac{1}{2}\tau((1 - J)a) + \tfrac{1}{2}\tau(Ja)$$

and both are traces, and since the amenable traces are a face [BO08, Proposition 6.3.7], we have that τ is not amenable.

Note that $C^*(*_Q \mathbb{Z}_2^m \times \mathbb{Z}_2) \cong C^*(\mathcal{F}(Q, 2^m)) \otimes C^*(\mathbb{Z}_2)$ has the local lifting property [BO08, Exercise 13.2]. Note that there is a natural quotient map $q : \mathcal{F}(S \cup \{J\}) \to *_Q \mathbb{Z}_2^m \times \mathbb{Z}_2$ and it induces a quotient on the C^* level. Here, $\mathcal{F}(S \cup \{J\})$ is the free group generated by $S \cup \{J\}$, so we can lift up τ to $\tilde{\tau} = \tau \circ q$ on $C^*(\mathcal{F}(S \cup \{J\}))$. Note that $\tilde{\tau}$ corresponds to an action of

https://doi.org/10.1017/S0010437X26103431 Published online by Cambridge University Press

$\mathcal{F}(S \cup \{J\})$ on (X, μ), so it is an IRS trace. We also have

$$\tilde{\tau}(\mathcal{F}(S \cup \{J\}))'' = \tau(\underset{Q}{*} \mathbb{Z}_2^m \times \mathbb{Z}_2)''.$$

Since the latter has the local lifting property, the non-amenability of τ implies that the algebra is not Connes embeddable, and hence that $\tilde{\tau}$ is not amenable. □

Since this implies the negation to Aldous–Lyons by Proposition 2.19, this proof allows us to skip [BCL+24]. That is, we reduce to non-local games without having to go through subgroup tests. In that paper, the authors also reduced to the *a priori* harder problem of finding this separation for a restricted class of games called tailored games, which have a sort of linearization baked into them, allowing translation as group-like objects. However, we do not require this, as these IRS strategies already bake in a group-like structure.

This highlights two different approaches when one wants to use these non-local games: restrict the class of strategy or restrict the class of games. In [BCL+24], the authors ended up doing both but their ideas were based off of the latter. If there was a linear constraint game with a perfect quantum commuting strategy but $\omega^* < 1$, then there would be a non-hyperlinear group. Based off of this, the tailored games of [BCL+24] were designed to be 'semi-linear'.

3.2 NPA hierarchy

To get the separation, we will need an analogue of the NPA hierarchy for these IRS strategies. That is, we will need the following.

Theorem 3.6. *Let $\mathfrak{G}$ be a non-local game; then there is a computably enumerable sequence α_n that is monotonically decreasing and converges to $\omega_{\mathrm{IRS}}(\mathfrak{G})$.*

Recall that this means there is a Turing machine that takes as input the description of some non-local game, and enumerates the sequence α_n.

To see this, we will recall some notions from [BCL+24, § 2]. Our NPA hierarchy for IRS value of a game is inspired by the one in [BCL+24] for subgroup tests, and is using some of the authors' lemmas leading up to that. Throughout this section, fix a non-local game $\mathfrak{G}$ and its variable set S. Then let $\mathcal{F} = \mathcal{F}(S \cup \{J\})$ be the free group generated by $S \cup \{J\}$.

Definition 3.7. Let $B \subset C \subset \mathcal{F}$. Then define the restriction map $R_{B \subset C} : \{0,1\}^C \to \{0,1\}^B$ as the map that sends $A \subset C$ to $A \cap B \subset B$. If $C = \mathcal{F}$, then call this map R_B.

Definition 3.8. A pseudo-subgroup of $B \subset \mathcal{F}$ is a subset $A \subset B$ that is the restriction of some subgroup. That is, there is some subgroup $H \subset \mathcal{F}$ so that $A = R_B(H) = H \cap B$.

Definition 3.9. A *random pseudo-subgroup* of $B \subset \mathcal{F}$ is a probability measure π on $\{0,1\}^B$ supported on the pseudo-subgroups of B.

A random pseudo-subgroup π is, in addition, invariant if for each $K, L \subset B$ so that $sKs^{-1}, sLs^{-1} \subset B$ for each $s \in S \cup S^{-1}$, we have

$$\pi\{A \subset B : K \subset A,\ A \cap L = \varnothing\} = \pi\{A \subset B : sKs^{-1} \subset A,\ A \cap sLs^{-1} = \varnothing\}$$

Denote the set of invariant random pseudo-subgroups of B as $\mathcal{Q}_B$, and let $\tilde{\mathcal{Q}}_B = R_{B*}^{-1}(\mathcal{Q}_B)$: that is, the set of measures on $\{0,1\}^{\mathcal{F}}$ whose pushforward to $\{0,1\}^B$ is an invariant random pseudo-subgroup.

For each $\pi \in \tilde{\mathcal{Q}}_B$, we will extend the value of a non-local game to it. Here is the idea: if $\sigma : S \cup \{J\} \to \mathrm{Aut}(X, \mu)$ is an IRS strategy, then it corresponds to an IRS $H = \mathrm{Stab}(\sigma)$ on $\mathcal{F}$.

https://doi.org/10.1017/S0010437X26103431 Published online by Cambridge University Press

Now the correlations with respect to this are

$$p(a,b|x,y)=\tau_H((1-J)e_x^a e_y^b),\quad e_x^a=\prod_{i\leqslant m}\frac{1+(-1)^{a_i}U_{x,i}}{2}.$$

In particular, each $p(a,b|x,y)$ only requires the data of τ on a finite number of group elements. One can see this by expanding out e_x^a and e_y^b. The only group elements appearing in this will be those that are a product of elements $J,U_{x,i},U_{y,i}$ of length $\leqslant m+1$.

So we can find a finite set K with

$$p(a,b|x,y)=\sum_{g\in K}c_g^{abxy}\mathbb{P}(g\in H).$$

We can assume that K contains $S\cup\{J\}$, all of the squares from $S\cup\{J\}$, and all the commutators from it, simply by unioning this finite set with K.

Now, for a probability distribution π on $\{0,1\}^B$ for some set $B\supset K$, we can define $\mathrm{Val}(\mathfrak{G},\pi)$ as the expected value of winning with correlation

$$p(a,b|x,y)=\sum_{g\in K}c_g^{abxy}\pi(A:g\in A).$$

Note that this will typically not actually be a conditional probability distribution, as these can become negative. This is because π is not actually invariant and so we are not dealing with traces. But we can still find the value of the game on it. That is, if q is the distribution on the question set $Q\times Q$ and D is the decider, then define

$$\mathrm{Val}(\mathfrak{G},\pi)=\sum_{x,y}\sum_{a,b}q(x,y)D(a,b|x,y)\left(\sum_{g\in K}c_g^{abxy}\pi(A:g\in A)\right).$$

Now we prove the following.

Lemma 3.10. *Let $K\subset B\subset C\subset\mathcal{F}$, and let π be a probability measure on $\{0,1\}^C$. Then*

$$\mathrm{Val}(\mathfrak{G},\pi)=\mathrm{Val}(\mathfrak{G},R_{B\subset C*}\pi).$$

Proof. Note that

$$R_{B\subset C*}(\pi)(A\subset B:g\in A)=\pi(A\subset C:g\in R_{B\subset C}(A))=\pi(A\subset C:g\in A)$$

for $g\in K$. Note that we have used that $g\in B$ always. Since these determine the value of the game, we are done. □

Only one thing is missing: not all IRSs will define a strategy for our game. For example, we want all members of S to square to identity in the action. But this is easy to enforce.

Definition 3.11. Let $K\subset B\subset\mathcal{F}$. Define the set $\mathcal{T}_B$ to be the set of probability measures on $\{0,1\}^B$ that has measure 1 on the following sets:

(i) $\{A:J\notin A\}$;
(ii) $\{A:[u,v]\in A$ for $u,v\in S_x\}$ for each $x\in Q$;
(iii) $\{A:u^2\in A$ for $u\in S\cup\{J\}\}$; and
(iv) $\{A:[u,J]\in A$ for $u\in S\}$.

Note that $R_{B\subset C*}^{-1}\mathcal{T}_B=\mathcal{T}_C$.

The point of this is that $\mathrm{IRS}(\mathcal{F})\cap\mathcal{T}_{\mathcal{F}}$ correspond precisely to the IRSs that induce strategies on $\mathfrak{G}$. To see this, take any IRS H of $\mathcal{F}$, and any action σ of $\mathcal{F}$ on some Borel probability space

https://doi.org/10.1017/S0010437X26103431 Published online by Cambridge University Press

(X, μ) that has $H = \mathrm{Stab}(\alpha)$. Then having measure 1 on, e.g., $\{A : u^2 \in A \text{ for } u \in S \cup \{J\}\}$ is simply saying that u^2 stabilizes almost all points of X, i.e. $u^2 = 1$.

Now, we can finally construct the NPA hierarchy (the proof follows the structure of [BCL+24, Proof of Theorem 1.10], with the appropriate modification for our setting).

Proof of Theorem 3.6. Let K be as above and let $B \supset K$ be a finite set. We can then define

$$\mathrm{Val}_B(\mathfrak{G}) = \sup(\mathrm{Val}(\mathfrak{G}, \pi) : \pi \in \mathcal{Q}_B \cap \mathcal{T}_B) = \sup(\mathrm{Val}(\mathfrak{G}, \pi) : \pi \in \tilde{\mathcal{Q}}_B \cap \mathcal{T}_{\mathcal{F}}).$$

Here, we use that $R_{B*}^{-1}(\mathcal{T}_B) = \mathcal{T}_{\mathcal{F}}$. Let B_t be a sequence of sets, $B_1 \supset K$ and $B_{t+1} \supset B_t$. Also suppose that $\bigcup_t B_t = \mathcal{F}$. Then we prove the following.

- That $\mathrm{Val}_{B_t}(\mathfrak{G})$ is computable. To see this, first from [BCL+24, Lemma 2.16], $\mathcal{Q}_{B_t}$ can be considered as a computable polytope of $\mathbb{R}^{\{0,1\}^{B_t}}$ (i.e. assigning to each subset of B_t a real number). This means that $\mathcal{Q}_{B_t}$ can be computably defined using integer linear equations and inequalities. Note that $\mathcal{Q}_{B_t} \cap \mathcal{T}_{B_t}$ just adds finitely more linear equations to this: for fixed subsets of B_t, we are asking the sum over subsets of those to add up to 1. Hence we can do a linear program over $\mathcal{Q}_{B_t} \cap \mathcal{T}_{B_t}$ to compute $\mathrm{Val}_{B_t}(\mathfrak{G})$ in finite time.
- That $\mathrm{Val}_{B_t}(\mathfrak{G}) \geqslant \mathrm{Val}_{B_{t+1}}(\mathfrak{G})$. To see this, note that $R_{B \subset C*}(\mathcal{Q}_C) = \mathcal{Q}_B$. So
$$\tilde{\mathcal{Q}}_{B_t} \cap \mathcal{T}_{\mathcal{F}} \supset \tilde{\mathcal{Q}}_{B_{t+1}} \cap \mathcal{T}_{\mathcal{F}},$$
giving the result.
- That $\mathrm{Val}_{B_t}(\mathfrak{G}) \to \omega_{\mathrm{IRS}}(\mathfrak{G})$. The $\tilde{\mathcal{Q}}_B$ and $\mathcal{T}_{\mathcal{F}}$ are weak* closed subsets of $\mathrm{Prob}(\{0,1\}^{\mathcal{F}})$ and hence their intersection is compact. So there is some $\pi_t \in \tilde{\mathcal{Q}}_{B_t} \cap \mathcal{T}_{\mathcal{F}}$, which has
$$\mathrm{Val}_{B_t}(\mathfrak{G}) = \mathrm{Val}(\mathfrak{G}, \pi_t).$$
We can pick a weak* cluster of π_t; say, π_∞. By passing to a subsequence, we may assume π_t weak* converge to π_∞. Now note that for all $s \geqslant t$, we have $\pi_s \in \tilde{\mathcal{Q}}_{B_t} \cap \mathcal{T}_S$. Since this set is closed, we get $\pi_\infty \in \tilde{\mathcal{Q}}_{B_t} \cap \mathcal{T}_{\mathcal{F}}$ for all t, i.e.
$$\pi_\infty \in \bigcap_t \tilde{\mathcal{Q}}_{B_t} \cap \mathcal{T}_{\mathcal{F}} = \mathrm{IRS}(\mathcal{F}) \cap \mathcal{T}_{\mathcal{F}}.$$
For the proof of this intersection, see [BCL+24, Lemma 2.12].

  This means that π_∞ is an IRS strategy for the game. And now since $\mathrm{Val}(\mathfrak{G}, -)$ is a continuous functional, we get
$$\lim_{t \to \infty} \mathrm{Val}_{B_t}(\mathfrak{G}) = \mathrm{Val}(\mathfrak{G}, \pi_t) = \mathrm{Val}(\mathfrak{G}, \pi_\infty) \leqslant \omega_{\mathrm{IRS}}(\mathfrak{G}).$$
But as $\tilde{\mathcal{Q}}_{B_t} \cap \mathcal{T}_{\mathcal{F}} \supset \mathrm{IRS}(\mathcal{F}) \cap \mathcal{T}_{\mathcal{F}}$, we have $\mathrm{Val}_{B_t}(\mathfrak{G}) \geqslant \omega_{\mathrm{IRS}}(\mathfrak{G})$. We are done. □

3.3 Separation

We now get the separation directly from the main theorem of [BCV24].

Theorem 3.12. *For each Turing machine $\mathcal{M}$, there is a non-local game $\mathfrak{G}_{\mathcal{M}}$ whose description can be computed in finite time from $\mathcal{M}$ with the following.*

- *If $\mathcal{M}$ halts, then $\mathfrak{G}_{\mathcal{M}}$ has a perfect permutation strategy.*
- *If $\mathcal{M}$ does not halt, then $\omega^*(\mathfrak{G}_{\mathcal{M}}) < \frac{1}{2}$.*

Actually, the authors have proved something *a priori* stronger; that this is true when restricting to a class of games called tailored non-local games. With this, we can prove the following.

https://doi.org/10.1017/S0010437X26103431 Published online by Cambridge University Press

Theorem 3.13. *There exists a non local game* $\mathfrak{G}$ *with* $\omega_{\text{IRS}}(\mathfrak{G}) > \omega^*(\mathfrak{G})$.

Proof. First note that for each $\mathfrak{G}$, there is a computably enumerable sequence β_n that is increasing and converges to $\omega^*(\mathfrak{G})$. This is the inclusion $\mathsf{MIP}^* \subset \mathsf{RE}$ [JNV+21, Corollary 12.9]. This is by exhaustive search: a quantum strategy only requires the choice of unitaries of order 2 with some commuting properties, and we can enumerate the finite set of rational unitaries in dimension n with denominators at most n. The density of rational unitaries gives the convergence [LDC+13, Theorem 2].

Suppose that for all games we had $\omega^*(\mathfrak{G}) \geqslant \omega_{\text{IRS}}(\mathfrak{G})$. Under this assumption, we will construct an algorithm for the Halting problem.

Let $\mathcal{M}$ be a Turing machine, and let $\mathfrak{G}_{\mathcal{M}}$ be the game from Theorem 3.12. On this input, our algorithm will alternate enumerating the α_n from Theorem 3.6 and the β_n for $\mathfrak{G}_{\mathcal{M}}$. The algorithm will accept on the input if ever $\beta_n \geqslant \frac{1}{2}$. It will reject if ever $\alpha_n < 1$.

There are two cases. First suppose that $\mathcal{M}$ halts. This means that there is a perfect permutation strategy and so $\omega_{\text{IRS}}(\mathfrak{G}_{\mathcal{M}}) = 1$ and $\omega^*(\mathfrak{G}_{\mathcal{M}}) = 1$. So the α_n are always $\geqslant 1$ and the β_n are eventually $\geqslant \frac{1}{2}$. So the algorithm accepts on this input.

Suppose that $\mathcal{M}$ never halts. This means that $\omega^*(\mathfrak{G}_{\mathcal{M}}) < \frac{1}{2}$ and so $\omega_{\text{IRS}}(\mathfrak{G}_{\mathcal{M}}) < \frac{1}{2}$. In particular, α_n is eventually < 1 and the β_n are never $\geqslant \frac{1}{2}$. So the algorithm rejects on this input. □

We now combine the results of the paper to prove the main theorem.

Proof of Theorems 1.1 and 1.2. By Theorems 3.5 and 3.13, we see that there is a non-co-hyperlinear IRS H on some free group Γ.

Since $\text{IRS}(\Gamma)$ and the set of co-hyperlinear IRSs are closed convex sets, a separation between them implies that some extreme point of $\text{IRS}(\Gamma)$ is not co-hyperlinear. So there is some ergodic $H \in \text{IRS}(\Gamma)$ that is not co-hyperlinear.

By Proposition 2.17, this means that $L(\Gamma/H)$ is not Connes embeddable. By the discussion preceding Definition 2.12, we have that $L(\Gamma/H)$ embeds into some $L(\mathcal{R})$, where $\mathcal{R}$ is a countable p.m.p. relation. $\mathcal{R}$ can be chosen to be ergodic as H was ergodic.

Finally, by [CDI22, Theorem D], we have that $L(\mathcal{R})$ embeds into some $L(\mathcal{R}')$, where $\mathcal{R}'$ is an ergodic countable p.m.p. relation with property (T). Since $L(\Gamma/H)$ was not Connes embeddable, $L(\mathcal{R}')$ is not either. □

4. Further directions

Here are some future research directions.

(i) Given a trace τ on a free group (or any group), is there any way to determine if τ arises from an IRS? An intrinsic characterization of this would be interesting. Maybe one could try to generalize the notion of an IRS trace to general C^* algebras if such a characterization existed.

(ii) In [BCL+24], the authors have proved the existence of a non-co-sofic IRS on free groups, and we have proved the existence of a non-co-hyperlinear IRS. So a natural question is: Are these notions the same? If they are the same, then hyperlinearity and soficity is the same for groups.

If they are distinct, a naive application of non-local games will not help. The non-local-game method for separating two convex sets $B \subset A$ requires one to be able to have computable upper bounds for A and computable lower bounds for B. The issue is that ω^* cannot have computable upper bounds; otherwise, it would be computable. So really,

https://doi.org/10.1017/S0010437X26103431 Published online by Cambridge University Press

something novel is needed for this. Such a separation would be evidence for soficity and hyperlinearity not agreeing on groups.

It is not clear that they should be distinct either: perhaps ergodic theoretic methods could be used to show they agree.

We note that the existence of a Hilbert–Schmidt stable group that is not permutation stable [ES23] implies a non-co-sofic co-hyperlinear IRS on that group, but this does not extend to free groups.

(iii) Let $\mathcal{R}$ be a p.m.p. equivalence relation on (X, μ). Consider the uniform metric d_u on the full group $[\mathcal{R}] \subset \mathrm{Aut}(X, \mu)$. That is,

$$d_u(T, S) = \mu(x : T(x) \neq S(x)).$$

It is not hard to see that $L(\mathcal{R})$ being Connes embeddable is the same as $[\mathcal{R}]$ being metrically hyperlinear. Compare with [Cor16, Proposition 2.4].

However, the issue is that a metric group can fail to be metrically hyperlinear while the underlying discrete group is hyperlinear. Can we drop the metric condition in the above? If so, we have a non-hyperlinear group.

We expect that this should not be possible: there should be a non-Connes-embeddable $L(\mathcal{R})$ with hyperlinear full group. However, it is not clear how to prove it one way or the other.

Acknowledgements

I would like to thank Michael Brannan for discussions, supervision, and very thoroughly proofreading this paper. I would also like to thank Michael Chapman for discussion, proofreading, and providing suggestions to significantly clean up the proof of Theorem 3.13. I would like to thank Stefan Frunza for pointing out the explicit GNS construction in Proposition 2.11. I would also like to thank Jesse Peterson for pointing out the proof of Lemma 2.18. I would also like to thank the anonymous referee for helpful comments that cleaned up the exposition in parts of this paper.

Conflicts of interest

None.

Journal information

Compositio Mathematica is owned by the Foundation Compositio Mathematica and published by the London Mathematical Society in partnership with Cambridge University Press. All surplus income from the publication of *Compositio Mathematica* is returned to mathematics and higher education through the charitable activities of the Foundation, the London Mathematical Society and Cambridge University Press.

References

AGN17 M. Abert, T. Gelander and N. Nikolov, *Rank, combinatorial cost and homology torsion growth in higher rank lattices*, Duke Math. J. **166** (2017), 2925–2964.

AGV14 M. Abert, Y. Glasner and B. Virag, *Kesten's theorem for invariant random subgroups*, Duke Math. J. **163** (2014), 465–488.

AP17 C. Anantharaman and S. Popa, *An introduction to II1 factors* (version 15), Preprint (2017), https://www.math.ucla.edu/∼popa/Books/IIunV15.pdf.

https://doi.org/10.1017/S0010437X26103431 Published online by Cambridge University Press

BC25 L. Bowen and M. Chapman, *Surjunctivity does not characterize cosoficity of invariant random subgroups*, Preprint (2025), arXiv:2511.06586.

BCL$^+$24 L. Bowen, M. Chapman, A. Lubotzky and T. Vidick, *The Aldous–Lyons conjecture I: Subgroup tests*, Preprint (2024), arXiv:2408.00110.

BCV24 L. Bowen, M. Chapman and T. Vidick, *The Aldous–Lyons conjecture II: Undecidability*, Preprint (2024), arXiv:2501.00173.

Bro06 N. P. Brown, *Invariant means and finite representation theory of C^*-algebras*, Memoirs of the American Mathematical Society, vol. 184 (American Mathematical Society, 2006).

BO08 N. P. Brown and N. Ozawa, *C^*-algebras and Finite-dimensional Approximations*, Graduate Studies in Mathematics, vol. 88 (American Mathematical Society, 2008).

CDI22 I. Chifan, D. Drimbe and A. Ioana, *Embedding universality for II_1 factors with property (T)*, Preprint (2022), arXiv:2205.07442.

Cor16 L. Cordeiro, *An elementary approach to sofic equivalence relations*, Preprint (2016), arXiv:1608.05608.

ES23 C. Eckhardt and T. Shulman, *On amenable Hilbert–Schmidt stable groups*, Preprint (2023), arXiv:2207.01089.

ES05 G. Elek and E. Szab'o, *Hyperlinearity, essentially free actions and L^2-invariants. The sofic property*, Math. Ann. **332** (2005), 421–441.

ES24 D. Enders and T. Shulman, *On the (local) lifting property*, Preprint (2024), arXiv:2403.12224.

Gol21 I. Goldbring, *The Connes embedding problem: A guided tour*, Preprint (2021), arXiv:2109.12682.

GH24 I. Goldbring and B. Hart, *The universal theory of the hyperfinite II_1 factor is not computable*, Bull. Symb. Log. **30** (2024), 181–198.

HS18 D. Hadwin and T. Shulman, *Stability of group relations under small Hilbert–Schmidt perturbations*, J. Funct. Anal. **275** (2018), 761–792.

JNV$^+$21 Z. Ji, A. Natarajan, T. Vidick, J. Wright and H. Yuen, *MIP* = RE*, Commun. ACM **64** (2021), 131–138.

KPS18 S.-J. Kim, V. I. Paulsen and C. Schafhauser, *A synchronous game for binary constraint systems*, J. Math. Phys. **59** (2018), 032201.

LDC$^+$13 M.-X. Luo, Y. Deng, X.-B. Chen and Y.-X. Yang, *The rational approximations of the unitary groups*, Quantum Inf. Process. **12** (2013), 3149–3166.

Man25 A. Manzoor, *Invariant random subgroups, soficity, and Lück's determinant conjecture*, Preprint (2025), arXiv:2508.15154.

Pau03 V. Paulsen, *Completely bounded maps and operator algebras*, Cambridge Studies in Advanced Mathematics, vol. 78 (Cambridge University Press, Cambridge, 2003).

PT16 J. Peterson and A. Thom, *Character rigidity for special linear groups*, J. Reine Angew. Math. **716** (2016), 207–228.

Ver11 A. Vershik, *Nonfree actions of countable groups and their characters*, J. Math. Sci. **174** (2011), 1–6.

Aareyan Manzoor a2manzoo@uwaterloo.ca
Department of Pure Mathematics, University of Waterloo, 200 University Avenue West, Waterloo, Ontario N2L 3G1, Canada

https://doi.org/10.1017/S0010437X26103431 Published online by Cambridge University Press