Ergod. Th. & Dynam. Sys., page 1 of 22

doi:10.1017/etds.2024.40

Uniform syndeticity in multiple recurrence

ASGAR JAMNESHAN[†] and MINGHAO PAN[‡]

[†] *Department of Mathematics, Koç University, İstanbul, Türkiye*
(e-mail: ajamneshan@ku.edu.tr)

[‡] *The Division of Physics, Mathematics and Astronomy, Caltech, Pasadena, California, USA*
(e-mail: minghaopan@ucla.edu)

(*Received* 28 *November* 2023 *and accepted in revised form* 1 *May* 2024)

Abstract. The main theorem of this paper establishes a uniform syndeticity result concerning the multiple recurrence of measure-preserving actions on probability spaces. More precisely, for any integers $d, l \geq 1$ and any $\varepsilon > 0$, we prove the existence of $\delta > 0$ and $K \geq 1$ (dependent only on d, l, and ε) such that the following holds: Consider a solvable group Γ of derived length l, a probability space (X, μ), and d pairwise commuting measure-preserving Γ-actions $T_1, \ldots, T_d$ on (X, μ). Let E be a measurable set in X with $\mu(E) \geq \varepsilon$. Then, K many (left) translates of

$$\big\{\gamma \in \Gamma \colon \mu(T_1^{\gamma^{-1}}(E) \cap T_2^{\gamma^{-1}} \circ T_1^{\gamma^{-1}}(E) \cap \cdots \cap T_d^{\gamma^{-1}} \circ T_{d-1}^{\gamma^{-1}} \circ \cdots \circ T_1^{\gamma^{-1}}(E)) \geq \delta\big\}$$

cover Γ. This result extends and refines uniformity results by Furstenberg and Katznelson. As a combinatorial application, we obtain the following uniformity result. For any integers $d, l \geq 1$ and any $\varepsilon > 0$, there are $\delta > 0$ and $K \geq 1$ (dependent only on d, l, and ε) such that for all finite solvable groups G of derived length l and any subset $E \subset G^d$ with $m^{\otimes d}(E) \geq \varepsilon$ (where m is the uniform measure on G), we have that K-many (left) translates of

$$\begin{aligned}\{g \in G \colon m^{\otimes d}(\{(a_1, \ldots, a_n) \in G^d \colon \\ (a_1, \ldots, a_n), (ga_1, a_2, \ldots, a_n), \ldots, (ga_1, ga_2, \ldots, ga_n) \in E\}) \geq \delta\}\end{aligned}$$

cover G. The proof of our main result is a consequence of an ultralimit version of Austin's amenable ergodic Szemerédi theorem.

Key words: multiple recurrence, uniform syndeticity, sated extensions, ultraproducts
2020 Mathematics Subject Classification: 37A30 (Primary); 37A15 (Secondary)

1. *Introduction*

A subset of a group is called syndetic if the union of finitely many translates of it cover the whole group. Throughout, Følner nets are left Følner nets, and syndetic sets are left syndetic sets in non-commutative groups.

One version [15] of Furstenberg's multiple recurrence theorem [14] is as follows.

THEOREM 1.1. *For every abelian group Γ, each probability space (X, μ) with finitely many pairwise commuting measure-preserving Γ-actions $T_i \curvearrowright (X, \mu)$, $i = 1, \ldots, d$, and all measurable sets E in X with positive measure, the return set*

$$\{\gamma \in \Gamma \colon \mu(E \cap T_1^{-\gamma}(E) \cap T_2^{-\gamma}(E) \cap \cdots \cap T_d^{-\gamma}(E)) > 0\}$$

is syndetic.

In this paper, our objective is to explore the uniformity of multiple recurrence theorems. There exist two directions of uniformity. First, can we establish a uniform lower bound for the measure of the multiple recurrence event, denoted as $\mu(E \cap T_1^{\gamma}(E) \cap T_2^{\gamma}(E) \cap \cdots \cap T_d^{\gamma}(E))$, keeping it away from zero? Second, can we assert that the return set is uniformly not too small? Increasing the value of d or shrinking the measure of the set E might lead to a reduction in both the multiple recurrence event and the return set. Nevertheless, our aspiration is for these measures to remain independent of certain factors: the group Γ, the probability space (X, μ), the commuting measure-preserving Γ-actions $T_1, \ldots, T_d$, and the choice of measurable set E—as long as we fixed d and $\mu(E)$.

To pursue the second aspect of uniformity, it becomes necessary to establish a method for quantifying the size of a subset within a group. In light of the statement of Theorem 1.1, a natural choice is to use the concept of K-syndeticity: given a group Γ and an integer $K \geq 1$, a subset $S \subset \Gamma$ is said to be *K-syndetic* if K many translates of S cover Γ. (In a previous version of this paper, we quantified syndeticity using the size of the lower Banach density of a subset. We are indebted to the anonymous referee for suggesting the use of the more natural (and seemingly stronger) concept of K-syndicity, which also had the benefit of significantly simplifying the proof of our main uniform syndeticity result.)

Numerous findings pertaining to uniform syndeticity are available within the existing literature. Among these, a notable contribution was made by Furstenberg and Katznelson, who demonstrated the prevalence of uniform syndeticity across all $\mathbb{Z}$-actions.

THEOREM 1.2. (Uniform syndeticity, $\mathbb{Z}$-case) *For every integer $d \geq 1$ and any $\varepsilon > 0$, there are $\delta > 0$ and $K \geq 1$ (only depending on ε, d) such that for any probability space (X, μ), every d many pairwise commuting measure-preserving transformations $T_i \colon X \to X$, $i = 1, \ldots, d$, and all measurable sets E in X with $\mu(E) \geq \varepsilon$, it holds that*

$$\{n \in \mathbb{Z} \colon \mu(E \cap T_1^{-n}(E) \cap T_2^{-n}(E) \cap \cdots \cap T_d^{-n}(E)) \geq \delta\}$$

is K-syndetic.

Proof. This result can be deduced from [5, Theorem 2.1(iii)]. □

The following weaker assertion is established for a fixed arbitrary countable abelian group by Furstenberg and Katznelson in their work [15]. In this version of uniform

https://doi.org/10.1017/etds.2024.40 Published online by Cambridge University Press

syndeticity, the probability of the multiple recurrence event is *not* shown to be uniformly bounded away from zero, as observed in Theorem 1.2 or later shown in Theorem 1.5.

THEOREM 1.3. (Weak uniform syndeticity, countable abelian case) *Let Γ be a countable abelian group. For every integer $d \geq 1$ and any $\varepsilon > 0$, there exists $K \geq 1$ (only depending on ε, d, and Γ) such that for any probability space (X, μ), every d many pairwise commuting measure-preserving actions $T_i : \Gamma \curvearrowright (X, \mu)$, $i = 1, \ldots, d$, and every measurable set E in X with $\mu(E) \geq \varepsilon$, it holds that*

$$\big\{\gamma \in \Gamma : \mu(E \cap T_1^{-\gamma}(E) \cap T_2^{-\gamma}(E) \cap \cdots \cap T_d^{-\gamma}(E)) > 0\big\}$$

is K-syndetic.

Proof. The claim follows from combining the results in [15, §10], see the last remark therein. □

Remark 1.4. In fact, in [15], Furstenberg and Katznelson establish that the return set

$$\big\{\gamma \in \Gamma : \mu(E \cap T_1^{-\gamma}(E) \cap T_2^{-\gamma}(E) \cap \cdots \cap T_d^{-\gamma}(E)) > 0\big\}$$

satisfies stronger notions of largeness than syndeticity such as IP* or even IP_r^*. However, we will focus on strengthening and generalizing the slightly weaker consequence stated in Theorem 1.3.

Our main result establishes a new proof and a joint generalization and strengthening of Theorems 1.2 and 1.3 by relaxing the dependence of δ and K on the acting group.

THEOREM 1.5. *For all integers $d, l \geq 1$ and any $\varepsilon > 0$, there exist $\delta > 0$ and $K \geq 1$ (only depending on ε, d, l) such that for any solvable group Γ of derived length l, any probability space (X, μ), every d many pairwise commuting measure-preserving actions $T_i : \Gamma \curvearrowright (X, \mu)$, $i = 1, \ldots, d$, and every measurable set E in X with $\mu(E) \geq \varepsilon$, it holds that*

$$\{\gamma \in \Gamma : \mu(E \cap T_1^{\gamma^{-1}}(E) \cap (T_{[1,2]}^{\gamma})^{-1}(E) \cap \cdots \cap (T_{[1,d]}^{\gamma})^{-1}(E)) \geq \delta\}$$

is K-syndetic, where $T_{[a,b]}^{\gamma} := T_a^{\gamma} \circ T_{a+1}^{\gamma} \circ \cdots \circ T_b^{\gamma}$.

We recall that the derived length n of a solvable group Γ is the least n for which $\Gamma^{(n)} = 1$, where $\Gamma^{(i)}$ is recursively defined by $\Gamma^{(0)} = \Gamma$ and $\Gamma^{(i+1)}$ is the commutator subgroup $[\Gamma^{(i)}, \Gamma^{(i)}]$ of $\Gamma^{(i)}$.

Remark 1.6. In fact, we establish a more general version of Theorem 1.5 where we can consider any uniformly amenable class of groups, of which a class of solvable groups of fixed derived length is an example. See §2 for the definition of uniform amenability and Theorem 3.2 for the general statement.

Remark 1.7. It is important to discern the variance in the articulation of the multiple recurrence event in the abelian case in Theorems 1.1, 1.2, and 1.3, where we consider T_i^{γ} rather than the composite actions $T_{[1,i]}^{\gamma}$ as presented in the formulation of Theorem 1.5. The possibility of attaining an analogous formulation to Theorems 1.1, 1.2, and 1.3 does

https://doi.org/10.1017/etds.2024.40 Published online by Cambridge University Press

indeed arise for nilpotent groups. However, it is crucial to note, as observed by Bergelson and Leibman in [6], that such a formulation fails to hold universally for solvable groups.

Remark 1.8. Quantitatively stronger results in the form of Khintchine-type bounds are available in more specific situations, as seen in [1, 4, 10–12, 22]. To the best of our knowledge, Theorem 1.5 is the first result of its kind to establish the existence of uniform bounds for arbitrary d, independent of Γ (within a large class of groups), and without requiring the hypothesis of ergodicity.

Remark 1.9. In [13], the first author and coauthors established a slightly weaker formulation of Theorem 1.5 (in fact, of the general Theorem 3.2) for two commuting actions involving, instead of K-syndeticity, a uniformly lower bound on the lower Banach density of the return set. The proof in [13] relied on certain technical lemmas about the interplay of Hahn–Banach type extensions for finitely additive invariant means and ultralimits of lower Banach densities. Using the stronger K-syndeticity formulation, our proof not only establishes a stronger generalization of the result in [13] to finitely many commuting actions, but also significantly simplifies (in particular, no Hahn–Banach type theorems and ultralimits of lower Banach densities are required anymore) and basically follows from an ultralimit construction of Austin's amenable ergodic Szemerédi theorem as stated next.

A pivotal step in demonstrating Theorem 1.5 hinges on employing ultraproducts of measure-preserving dynamical systems. However, the resultant ultraproduct groups are often not countable and the corresponding Loeb probability spaces lack separability. Addressing these challenges introduces certain measure-theoretic subtleties, discussed comprehensively in [20]. To navigate around these intricacies, an abstract category of probability algebra dynamical systems $\mathbf{PrbAlg}_\Gamma$ has been identified (see [18–21]). The abstract system is obtained from a concrete probability space by abstracting away the intrinsic point structure and exclusively focusing on the relationships between measurable sets, considering operations such as intersections, unions, and complementations. $\mathbf{PrbAlg}_\Gamma$, along with the tools to work with its objects, is gathered in §2.1

Equipped with these tools in the domain of uncountable ergodic theory, we extend Austin's amenable multiple recurrence theorem [2] to encompass the actions of uncountable amenable groups on inseparable probability spaces in the following theorem. This uncountable variant of Austin's theorem is then applied to the ultraproduct systems, playing a key role in proving Theorem 1.5.

THEOREM 1.10. *Let Γ be an arbitrary discrete amenable group and let (X, μ, T) be a $\mathbf{PrbAlg}_{\Gamma^d}$-system, that is, there are finitely many commuting measure-preserving Γ-actions $T_i \colon \Gamma \curvearrowright (X, \mu)$, $i = 1, \ldots, d$, where (X, μ) is a probability algebra. Let $f_1, \ldots, f_d \in L^\infty(X, \mu)$ and let (Φ_κ) be a Følner net for Γ. Then the limit*

$$\lim_{\kappa} \frac{1}{|\Phi_\kappa|} \sum_{\gamma \in \Phi_\kappa} \prod_{i=1}^{d} f_i \circ T_{[1,i]}^{\gamma} \tag{1}$$

exists in $L^2(X, \mu)$ and is independent of the Følner net. Moreover, if a measurable set E in X is such that $\mu(E) > 0$, then

https://doi.org/10.1017/etds.2024.40 Published online by Cambridge University Press

$$\lim_{\kappa} \frac{1}{|\Phi_\kappa|} \sum_{\gamma \in \Phi_\kappa} \mu\bigg(\bigcap_{i=0}^{d} T_{[1,i]}^{\gamma^{-1}}(E) \bigg) > 0. \tag{2}$$

In particular, there exists $\varepsilon > 0$ such that

$$\bigg\{ \gamma \in \Gamma : \mu\bigg(\bigcap_{i=0}^{d} T_{[1,i]}^{\gamma^{-1}}(E) \bigg) > \varepsilon \bigg\} \tag{3}$$

is syndetic in Γ.

The case of two commuting transformations of Theorem 1.10 was previously established in [13] by the first author and coauthors. They generalized the proof of the amenable double recurrence theorem by Bergelson, McCutcheon, and Zhang from [8]. Zorin-Kranich [26] establishes the limit claim in equation (1) of Theorem 1.10 in full generality using an adaptation of functional analytic methods developed by Walsh [24], who established the L^2-limit in the case of finitely many actions of a nilpotent group. However, Zorin-Kranich's result does not provide information about the limit; in particular, it does not yield the multiple recurrence statements in equations (2) and (3). These multiple recurrence statements are established in the case of countable amenable groups by Austin [2] using sated extensions. Our proof of Theorem 1.10 will modify necessary steps in [2] to tailor his proof to our uncountable setting.

1.1. *Combinatorial application.*

As an immediate consequence of Theorem 1.5, we obtain the following combinatorial application. For a finite group G, we denote by m the uniform measure. On G^d, we denote by $m^{\otimes d}$ the d-fold product of the uniform measure.

COROLLARY 1.11. *Let $d, l \geq 1$ be integers and let $\varepsilon > 0$. Then there are $\delta = \delta(d, l, \varepsilon) > 0$ and $K = K(d, l, \varepsilon) \geq 1$ such that for all finite solvable groups G of derived length l and any subset $E \subset G^d$ with $m^{\otimes d}(E) \geq \varepsilon$, we have that*

$$\begin{aligned} \{g \in G : m^{\otimes d}(\{(a_1, \ldots, a_n) \in G^d :& \\ (a_1, \ldots, a_n), (ga_1, a_2, \ldots, a_n), \ldots, &(ga_1, ga_2, \ldots, ga_n) \in E\}) \geq \delta\} \end{aligned}$$

is K-syndetic.

Remark 1.12. In [9, Theorem 1.5], the authors establish K-syndeticity for the density of triangle configurations with a Khintchine-type lower bound in a fixed class of quasirandom ultraproduct groups. In [9, Corollary 1.6], they deduce a similar consequence for the class of all non-cyclic finite simple groups. Their proof relies on a convergence theorem along minimal idempotent ultrafilters for ergodic averages formed by two commuting actions of a minimally almost periodic group.

1.2. *Discussion.*

In Theorem 3.2, we prove a more general version of Theorem 1.5, establishing a strong form of uniform syndeticity within any *uniformly amenable* class of groups. We then derive Theorem 1.5 by demonstrating that the class of solvable groups with a fixed derived length is uniformly amenable. A key property in this verification is

https://doi.org/10.1017/etds.2024.40 Published online by Cambridge University Press

that the class of solvable groups with a fixed derived length is closed under countable direct products.

Since Theorem 1.10 holds for all amenable groups, a natural question arises: does Theorem 1.5 remain valid for the entire class of amenable groups (that is, K, δ only depend on ε, d and are uniform for all amenable acting groups)? The methods employed in this paper cannot address this question. Indeed, the countable direct product of arbitrary nilpotent groups is not amenable in general, which essentially undermines our strategy based on Theorem 1.10 to prove uniform syndeticity. We are grateful to Dave Benson for this observation. Additionally, an example of an amenable group that is not uniformly amenable is given in [**25**].

However, the results of [**7**] give hope that Theorem 1.10 holds for all amenable groups (or even for the class of all groups). In [**7**, Theorem 1.3], the authors establish the consequence in equation (3) of Theorem 1.10 for the action of an arbitrary countable group in the case of two commuting actions. (Actually, they establish the stronger conclusion that the return set is $\mathcal{C}^*$, we refer the interested reader to [**7**] for the definition of a central* subset of a group.) Proving an uncountable version of [**7**, Theorem 1.3] should, in principle, yield the analogue of Theorem 1.5 for the class of all groups in the case of two commuting actions by the same proof as given in §3. To our knowledge, the analogue of [**7**, Theorem 1.3] in the case of more than two commuting actions of an arbitrary countable group is unknown. A potential line of attack, suggested by Austin [**2**], is to combine the technique of stated extensions and the ultrafilter techniques in [**7**]. We hope to address these questions in future work.

2. *Tools*

2.1. *The category of probability algebra dynamical systems.* We now formalize the 'point-free' approach by introducing the category of probability algebra dynamical systems and the canonical model functor, see Figure 1. For a comprehensive background, references, and any unexplained concepts which are used in the following, we refer the interested reader to [**20**].

Definition 2.1. (Categories and functors)

(i) We denote by **CHPrb** the category of compact Hausdorff spaces equipped with a Baire–Radon probability measure and measure-preserving continuous functions.

(ii) We denote by **AbsMbl** the category of abstract measurable spaces which we define as the opposite category of the category of σ-complete Boolean algebras and σ-complete Boolean homomorphisms.

(iii) We denote by **PrbAlg** the opposite category of the category of probability algebras and measure-preserving Boolean homomorphisms. Note that the category **PrbAlg** has arbitrary inverse limits (e.g., see [**20**]), a fact which will be useful for us later.

(iv) We denote by **Hilb** the category of complex Hilbert spaces and linear isometries.

(v) Given a (discrete) group Γ, we can turn a category $\mathcal{C} =$ **Hilb**, **PrbAlg**, **CHPrb** into a dynamical category $\mathcal{C}_\Gamma$ as follows. Given an object X in $\mathcal{C}$, we can associate with X the group $\mathrm{Aut}(X)$ of its automorphisms in $\mathcal{C}$. The dynamical category $\mathcal{C}_\Gamma$ now consists of pairs (X, T), where X is a $\mathcal{C}$-object and $T\colon \Gamma \to \mathrm{Aut}(X)$ a

https://doi.org/10.1017/etds.2024.40 Published online by Cambridge University Press

$$\begin{tikzcd}
\mathbf{CHPrb} \arrow[d, shift right, "\mathtt{Alg}\circ\mathtt{Abs}"'] & \mathbf{CHPrb}_{\Gamma} \arrow[l, tail] \arrow[d, shift right, "\mathtt{Alg}\circ\mathtt{Abs}"'] \\
\mathbf{PrbAlg} \arrow[u, two heads, bend right, "\mathtt{Conc}"'] \arrow[d, tail, "L^2"'] & \mathbf{PrbAlg}_{\Gamma} \arrow[l, tail] \arrow[u, two heads, bend right, "\mathtt{Conc}"'] \arrow[d, tail, "L^2"'] \\
\mathbf{Hilb}^{\mathrm{op}} & \mathbf{Hilb}_{\Gamma}^{\mathrm{op}} \arrow[l, tail]
\end{tikzcd}$$

FIGURE 1. The main categories and functors used in this paper (op indicates the use of the opposite category). Arrows with tails are faithful functors and arrows with two heads in one direction are full. Unlabeled functors are forgetful. The diagram is not fully commutative (even modulo natural isomorphisms).

group homomorphism. A $\mathcal{C}_\Gamma$-morphism is a $\mathcal{C}$-morphism which intertwines with the Γ-actions.

(vi) The abstraction functor $\mathtt{Abs}$ maps a concrete measurable space (X, Σ_X) to the **AbsMbl**-object Σ_X and a measurable function $f : (X, \Sigma_X) \to (Y, \Sigma_Y)$ to the **AbsMbl**-morphism $f^* : \Sigma_X \to \Sigma_Y$, where f^* is the (opposite) pullback map $f^*(E) := f^{-1}(E)$, $E \in \Sigma_Y$. We apply the $\mathtt{Abs}$-functor to a concrete probability space (X, Σ_X, μ) and obtain an abstract probability space (Σ_X, μ). Let $\mathcal{I}_\mu = \{E \in \Sigma_X : \mu(E) = 0\}$ be the ideal of μ-null sets. Then the quotient Boolean algebra $X_\mu := \Sigma_X/\mathcal{I}_\mu$ is σ-complete. We can lift the measure μ to X_μ in a natural way and, by an abuse of notation, we denote this lift by μ again. The tuple (X_μ, μ) is a probability algebra and we define $\mathtt{Alg} \circ \mathtt{Abs}(X, \Sigma_X, \mu_X) := (X_\mu, \mu)$. If $f : (X, \Sigma_X, \mu) \to (Y, \Sigma_Y, \nu)$ is a measure-preserving function, then the pullback map f^* maps the ideal $\mathcal{I}_\nu$ to the ideal $\mathcal{I}_\mu$. We obtain a **PrbAlg**-morphism $\mathtt{Alg} \circ \mathtt{Abs}(f) : (X_\mu, \mu) \to (Y_\nu, \nu)$.

(vii) The canonical model functor $\mathtt{Conc}$ reverses the process described in the previous item. More precisely, if (X, μ) is a probability algebra, there exists a **CHPrb**-space $\mathtt{Conc}(X, \mu) := (Z, \mathcal{B}a(Z), \mu_Z)$ such that $\mathtt{Alg} \circ \mathtt{Abs}(Z, \mathcal{B}a(Z), \mu_Z)$ is isomorphic to (X, μ) in **PrbAlg**. A complete construction of $\mathtt{Conc}(X, \mu)$ can be found in [20]. Given a **PrbAlg**-morphism $f : (X, \mu) \to (Y, \nu)$, we define $\mathtt{Conc}(f) : Z_X \to Z_Y$ by $\mathtt{Conc}(f)(\theta) = \theta \circ \phi$.

(viii) Next we define the L^2-functor. Let (X, μ) be a **PrbAlg**-space with canonical model $\mathtt{Conc}(X, \mu) = (Z, \mathcal{B}a(Z), \mu_Z)$, as constructed previously. We define the L^2-functor on objects by $L^2(X, \mu) := L^2(Z, \mathcal{B}a(Z), \mu_Z)$. If $\pi : (X, \mu) \to (Y, \nu)$ is a **PrbAlg**-morphism, then $L^2(\pi) : L^2(Y, \nu) \to L^2(X, \mu)$ is defined by the Koopman operator $L^2(\pi)(f) := \pi^* f$, where $\pi^* f := f \circ \mathtt{Conc}(\pi)$.

(ix) Similarly, we define the dynamical version of the functors $\mathtt{Abs} \circ \mathtt{Alg}$, $\mathtt{Conc}$, and L^2.

A significant feature of the canonical model $\mathtt{Conc}(X) = (Z_X, \mathcal{B}a(Z_X), \mu_{Z_X})$ of a **PrbAlg**-space $X = (X, \mu)$ is the strong Lusin property (cf. [20, §7]), which states that the commutative von Neumann algebra $L^\infty(Z_X, \mathcal{B}a(Z_X), \mu_{Z_X})$ is isomorphic to the commutative C^*-algebra $C(Z_X)$ of continuous functions on Z_X in the category of unital

https://doi.org/10.1017/etds.2024.40 Published online by Cambridge University Press

C^*-algebras. In **CHPrb**-spaces with the strong Lusin property, every equivalence class of bounded measurable functions has a continuous representative. A very useful consequence of this property is that it comes with a canonical disintegration of measures.

THEOREM 2.2. *Let Γ be a discrete group. Let $\pi : (X, \mu, T) \to (Y, \nu, S)$ be a $\mathbf{PrbAlg}_\Gamma$-morphism. Then there exists a unique Baire–Radon probability measure μ_y on Z_X for each $y \in Z_Y$, which depends continuously on y in the vague topology in the sense that $y \mapsto \int_{Z_X} f\, d\mu_y$ is continuous for every f in the space of continuous functions $C(Z_X)$, and such that*

$$\int_{Z_X} f(x) g(\mathtt{Conc}(\pi)(x))\, d\mu_{Z_X}(x) = \int_{Z_Y} \bigg(\int_{Z_X} f\, d\mu_y \bigg) g\, d\mu_{Z_Y} \tag{4}$$

for all $f \in C(Z_X)$, $g \in C(Z_Y)$. Furthermore, for each $y \in Z_Y$, μ_y is supported on the compact set $\mathtt{Conc}(\pi)^{-1}(\{y\})$, in the sense that $\mu_y(E) = 0$ whenever E is a measurable set disjoint from $\mathtt{Conc}(\pi)^{-1}(\{y\})$. (Note that this conclusion does not require the fibers $\mathtt{Conc}(\pi)^{-1}(\{y\})$ to be Baire measurable.) Moreover, we have $\mu_{S^\gamma_{Z_Y}(y)} = (T^\gamma_{Z_X})^* \mu_y$ for all $y \in Z_Y$ and $\gamma \in \Gamma$.

Let $\pi : (X, \mu, T) \to (Y, \nu, S)$ be a $\mathbf{PrbAlg}_\Gamma$-morphism and let $\mathtt{Conc}(\pi) : Z_X \to Z_Y$ be its canonical representation. For every $f \in L^2(\mathtt{Conc}(Y))$, the composition $\mathtt{Conc}(\pi)^* f$ is an element of $L^2(\mathtt{Conc}(X))$ since $\mathtt{Conc}(\pi)$ is measure-preserving. In fact, $\{\mathtt{Conc}(\pi)^* f : f \in L^2(\mathtt{Conc}\ (Y))\}$ is a closed Γ-invariant subspace of $L^2(\mathtt{Conc}(X))$. Thus, we can identify $L^2(\mathtt{Conc}(Y))$ with the closed subspace $\mathtt{Conc}(\pi)^*(L^2(\mathtt{Conc}(Y)))$ in $L^2(\mathtt{Conc}(X))$. Using this identification, we can define a conditional expectation operator $\mathbb{E}(\cdot|Y)$ from $L^2(\mathtt{Conc}(X))$ to $L^2(\mathtt{Conc}(Y))$. Since L^∞ is dense in L^2 in the L^2 topology and by Theorem 2.2, we obtain the disintegration of measures

$$\mathbb{E}(f|Y)(y) = \int_{Z_X} f\, d\mu_y$$

almost surely for all $f \in L^2(\mathtt{Conc}(X))$.

An important application is a canonical construction of relatively independent products. Indeed, let $\pi_1 : (X_1, \mu_1, T_1) \to (Y, \nu, S)$, $\pi_2 : (X_2, \mu_2, T_2) \to (Y, \nu, S)$ be two $\mathbf{PrbAlg}_\Gamma$-morphisms. Let $(\mu^1_y)_{y \in Z_Y}$ and $(\mu^2_y)_{y \in Z_Y}$ be the corresponding canonical disintegration of measures. Define the probability measure

$$\mu_{Z_{X_1}} \times_{Z_Y} \mu_{Z_{X_2}}(E) := \int_{Z_Y} \mu^1_y \times \mu^2_y(E)\, d\mu_{Z_Y}$$

for all $E \in \mathcal{B}a(Z_{X_1}) \otimes \mathcal{B}a(Z_{X_2})$. Then

$$(Z_{X_1} \times Z_{X_2}, \mathcal{B}a(Z_{X_1}) \otimes \mathcal{B}a(Z_{X_2}), \mu_{Z_{X_1}} \times_{Z_Y} \mu_{Z_{X_2}}, T_{Z_{X_1}} \times T_{Z_{X_2}})$$

is a $\mathbf{CHPrb}_\Gamma$-object coming with two $\mathbf{CHPrb}_\Gamma$-morphisms $\psi_1 : Z_{X_1} \times Z_{X_2} \to Z_{X_1}$ and $\psi_2 : Z_{X_1} \times Z_{X_2} \to Z_{X_2}$. Applying the functor $\mathtt{Alg} \circ \mathtt{Abs}$, we obtain a $\mathbf{PrbAlg}_\Gamma$-object

https://doi.org/10.1017/etds.2024.40 Published online by Cambridge University Press

$(X_1 \times_Y X_2, \mu_1 \times_Y \mu_2, T_1 \times T_2)$ and the two $\mathbf{PrbAlg}_\Gamma$-morphisms $\mathtt{Alg} \circ \mathtt{Abs}(\psi_1)$, $\mathtt{Alg} \circ \mathtt{Abs}(\psi_2)$ satisfying the following commutative diagram in $\mathbf{PrbAlg}_\Gamma$:

$$\begin{tikzcd}
& X_1 \times_Y X_2 \arrow[dl, "\mathtt{Alg}\circ\mathtt{Abs}(\psi_1)"'] \arrow[dr, "\mathtt{Alg}\circ\mathtt{Abs}(\psi_2)"] & \\
X_1 \arrow[dr, "\pi_1"'] & & X_2 \arrow[dl, "\pi_2"] \\
& Y &
\end{tikzcd}$$

2.2. *Ultrafilters and non-standard analysis.* A filter on a set X is a non-empty collection f of subsets of X satisfying the following properties.

(i) $\emptyset \notin f$.

(ii) If $A, B \in f$, then $A \cap B \in f$.

(iii) If $A \in f$, $B \subset X$, and $A \subset B$, then $B \in f$.

An ultrafilter on X is a maximal element in the set of filters on X with respect to set inclusion. A *non-principal ultrafilter* is an ultrafilter such that none of its elements is finite.

We are concerned with non-principal ultrafilters on the set of natural numbers $\mathbb{N}$. The *Fréchet filter* consists of all subsets $A \subset \mathbb{N}$ for which there is $n \in \mathbb{N}$ such that A contains the tail $\{n, n+1, \ldots\}$. By definition, the Fréchet filter is contained in any non-principal ultrafilter on $\mathbb{N}$.

Fix a non-principal ultrafilter p on $\mathbb{N}$. The *ultraproduct* of a sequence $\{X_n\}$ of sets with respect to p is the quotient set $X^* = \prod_{n\to p} X_n$ of the Cartesian product $\prod_{n\in\mathbb{N}} X_n$ with respect to the equivalence relation $(x_n) \sim (y_n)$ if and only if $\{n \colon x_n = y_n\} \in p$. The equivalence class of (x_n) is denoted by $\lim_{n\to p} x_n$ and called an *ultralimit* of the elements x_n.

A subset A of X^* is said to be *internal* if it is of the form $A = \prod_{n\to p} A_n$ for some subset $A_n \subset X_n$ for each n. One can check that the collection of internal subsets of X^* forms an algebra of sets, in particular, $\bigcup_{i=1}^K A_i = \prod_{n\to p} \bigcup_{i=1}^K A_{i,n}$ for finitely many internal subsets A_i of X^*.

One can also verify that the ultraproduct of a sequence of groups is a group. Moreover, a class $\mathfrak{G}$ of groups is said to be *uniformly amenable* if the ultraproduct group Γ^* is amenable for any sequence of groups $\Gamma_n \in \mathfrak{G}$.

Let (r_n) be a bounded sequence of real numbers. By the Bolzano–Weierstrass theorem, there is a unique real number r such that $\lim_{n\to p}(r_n - r)$ is infinitesimal, that is, an ultralimit real number in the equivalence class of a null sequence. We define the standard part function $\mathrm{st}(\lim_{n\to p} r_n) \coloneqq r$.

3. *Proof of Theorem 1.5*

We prove Theorem 1.5 via an ultralimit construction from Theorem 1.10. The following simple lemma is key. Throughout, fix a non-principal ultrafilter p on $\mathbb{N}$.

LEMMA 3.1. *Let $\{\Gamma_n\}$ be a sequence of groups and let $\Gamma^* = \prod_{n\to p} \Gamma_n$ be their ultraproduct. Let $S = \prod_{n\to p} S_n$ be an internal subset of Γ^* which is a K-syndetic subset of Γ^* for*

https://doi.org/10.1017/etds.2024.40 Published online by Cambridge University Press

some $K \geq 1$. Then

$$\{n\colon S_n \text{ is a } K\text{-syndetic set in } \Gamma_n\} \in p.$$

Proof. There are $\gamma_1, \ldots, \gamma_K \in \Gamma$ such that $\Gamma^* = \bigcup_{i=1}^K \gamma_i \cdot S$. We have $\gamma_i = \lim_{n\to p} \gamma_{i,n}$ for some choice of $\gamma_{i,n} \in \Gamma_n$ for each n and every $1 \leq i \leq K$ such that

$$\Gamma^* = \bigcup_{i=1}^K \gamma_i \cdot S = \prod_{n\to p} \bigcup_{i=1}^K \gamma_{i,n} \cdot S_n,$$

and thus $\{n\colon \bigcup_{i=1}^K \gamma_{i,n} \cdot S_n = \Gamma_n\} \in p$. □

At the end of this section, we show how to derive Theorem 1.5 from the following generalization of it.

THEOREM 3.2. *For every uniformly amenable class $\mathfrak{G}$ of groups, for all integer $d \geq 1$, and any $\varepsilon > 0$, there exist an integer $K \geq 1$ and $\delta > 0$ (only depending on $\mathfrak{G}, \varepsilon, d$) such that for any group $\Gamma \in \mathfrak{G}$, any probability space (X, μ), every finitely many pairwise commuting measure-preserving actions $T_i\colon \Gamma \curvearrowright (X, \mu)$, $i = 1, \ldots, d$, and every measurable set E in X with $\mu(E) \geq \varepsilon$, it holds that*

$$\big\{\gamma \in \Gamma\colon \mu(T_1^{\gamma^{-1}}(E) \cap (T_{[1,2]}^{\gamma})^{-1}(E) \cap \cdots \cap (T_{[1,d]}^{\gamma})^{-1}(E)) \geq \delta\big\}$$

is K-syndetic, where $T_{[a,b]}^{\gamma} := T_a^{\gamma} \circ T_{a+1}^{\gamma} \circ \cdots \circ T_b^{\gamma}$.

Proof. Toward a contradiction, assume there is a uniformly amenable class of groups $\mathfrak{G}$, $d \geq 1$, $\varepsilon > 0$ such that for every $n \geq 1$, there is a group $\Gamma_n \in \mathfrak{G}$, d pairwise commuting measure-preserving Γ_n-actions $T_{n,1}, \ldots, T_{n,d}$ on a probability space $(X_n, \mathcal{X}_n, \mu_n)$, and $E_n \in \mathcal{X}_n$ with $\mu_n(E_n) \geq \varepsilon$ such that

$$A_n := \big\{\gamma \in \Gamma_n\colon \mu(T_{n,1}^{\gamma^{-1}}(E_n) \cap (T_{n,[1,2]}^{\gamma})^{-1}(E_n) \cap \cdots \cap (T_{n,[1,d]}^{\gamma})^{-1}(E_n) \geq 1/n\big\} \tag{5}$$

is not n-syndetic.

Let $X^* = \prod_{n\to p} X_n$ be the ultraproduct of the sets X_n and denote by

$$\mathcal{A} = \bigg\{ \prod_{n\to p} D_n : (D_n) \in \prod_{n\in\mathbb{N}} \mathcal{X}_n \bigg\}$$

the algebra of internal subsets of X^*. We define the Loeb premeasure

$$\mu^* : \mathcal{A} \to [0, 1], \quad \mu^*\bigg(\prod_{n\to p} D_n \bigg) := \mathrm{st}\Big(\lim_{n\to p} \mu_n(D_n) \Big).$$

By Carathéodory's extension and uniqueness theorem, μ^* extends to a unique countably additive probability measure μ on the σ-algebra of sets generated by $\mathcal{A}$. By a slight abuse of notation, we let (X_μ, μ) denote the probability algebra associated to $(X^*, \sigma(\mathcal{A}), \mu)$.

https://doi.org/10.1017/etds.2024.40 Published online by Cambridge University Press

We let $\Gamma^* = \prod_{n\to p} \Gamma_n$. Since $\mathfrak{G}$ is uniformly amenable, Γ^* is an amenable group. For each $i = 1, \ldots, d$, and for every $\gamma^* = \lim_{n\to p} \gamma_n \in \Gamma^*$ and $\prod_{n\to p} D_n \in \mathcal{A}$, define

$$(T_i^*)^{\gamma^*}\bigg(\prod_{n\to p} D_n\bigg) := \prod_{n\to p} T_{n,i}^{\gamma_n}(D_n).$$

One checks that $(T_i^*)^{\gamma^*}$ is a well-defined Γ^*-action by Boolean automorphism of $\mathcal{A}$ that preserves the probability measure μ. Since μ is a finite measure, by a standard approximation result in measure theory (see, e.g., [**3**, Theorem 5.7]), for any $D \in X_\mu$, there is a sequence $D_n \in \mathcal{A}$ such that $\mu(D\Delta D_n) \to 0$ as n tends to ∞, where Δ denotes symmetric set difference. Thus, we can extend actions T_i^* to abstract $\mathbf{PrbAlg}_\Gamma$-actions $T_i : \Gamma^* \to \mathrm{Aut}(X_\mu, \mu)$. We obtain an abstract $(\Gamma^*)^d$-system (X_μ, μ, T).

By construction, we have $\mu(E^*) \geq \varepsilon$, where $E^* := \prod_{n\to p} E_n$ and the E_n are as in equation (5). By Theorem 1.10, there exist $\delta > 0$ and $K \geq 1$ such that

$$B = \big\{\gamma \in \Gamma^* \colon \mu((T_1^*)^{\gamma^{-1}}(E^*) \cap (T_{[1,2]}^*)^{\gamma^{-1}}(E^*) \cap \cdots \cap (T_{[1,d]}^*)^{\gamma^{-1}}(E^*)) \geq \delta\big\}$$

is K-syndetic. Let

$$B_n = \{\gamma \in \Gamma_n \colon \mu_n(T_{n,1}^{\gamma^{-1}}(E_n) \cap (T_{n,[1,2]}^{\gamma})^{-1}(E_n) \cap \cdots \cap (T_{n,[1,d]}^{\gamma})^{-1}(E_n)) \geq \delta\}.$$

By construction, we have $B = \prod_{n\to p} B_n$. Then Lemma 3.1 gives

$$\{n \colon B_n \text{ is } K\text{-syndetic in } \Gamma_n\} \in p.$$

Moreover, $B_n \subset A_n$ as long as $n > 1/\delta$, and since the Fréchet filter is contained in any non-principal ultrafilter, we must have $\{n \colon B_n \subset A_n\} \in p$. Since a filter is intersection closed and p does not contain finite sets as a non-principal ultrafilter, there are infinitely many A_n which are K-syndetic, which contradicts our assumptions. $\square$

As a corollary, we obtain Theorem 1.5.

Proof. By Theorem 3.2, it suffices to verify that the class of solvable groups of derived length at most l is uniformly amenable for some fixed $l \geq 1$.

Let Γ^* be the ultraproduct of a sequence $\{\Gamma_n\}$ of solvable groups of derived length at most l. We claim that Γ^* is solvable (and amenable in particular). The direct product $\prod_n \Gamma_n$ is a solvable group as $\{\Gamma_n\}$ has uniformly bounded length. The ultraproduct $\Gamma^* = \prod_{n\to p} \Gamma_n$ is a quotient group of $\prod_n \Gamma_n$ so Γ^* is solvable as well. $\square$

4. *Proof of Theorem 1.10*

Austin [**2**] proved a subcase of Theorem 1.10 when the acting group Γ is countable and the space (X, μ) is standard Lebesgue by constructing characteristic spaces on stated extensions. Our proof of Theorem 1.10 aims to facilitate Austin's proof [**2**] to be carried out in a setup where spaces may not be standard Lebesgue and groups may not be countable. We will reuse most of the arguments of Austin [**2**] and only modify the steps in which the assumptions about the space and the group are substantially used. We will now use the notation introduced in §2.1.

https://doi.org/10.1017/etds.2024.40 Published online by Cambridge University Press

Throughout this section, we let Γ be a (discrete) amenable group and (Φ_κ) be a left Følner net for Γ. Let (X, μ) be a **PrbAlg**-space and $T_1, T_2, \ldots, T_d$ be commuting group homomorphisms $T_i \colon \Gamma \to \mathrm{Aut}(X, \mu)$, that is, $T_i^\gamma \circ T_j^\eta = T_j^\eta \circ T_i^\gamma$ for all $\gamma, \eta \in \Gamma$ and $1 \leq i < j \leq d$ (the automorphism group $\mathrm{Aut}(X, \mu)$ is taken in the category **PrbAlg**). We denote by $(\tilde{X}, \tilde{\mu}, \tilde{T})$ the corresponding canonical model, so that $\tilde{X}$ is a compact Hausdorff space and each $\tilde{T}_i$ acts on $(\tilde{X}, \tilde{\mu})$ by measure-preserving homeomorphisms. This canonical model has the advantage that tools such as disintegration of measure and relative independent product are available as discussed in §2.1. For the most part, we can then adopt Austin's arguments on the concrete space $(\tilde{X}, \tilde{\mu}, \tilde{T})$. Then the corresponding results for the probability algebra (X, μ) follow by applying the functor $\mathtt{Alg} \circ \mathtt{Abs}$.

We begin with Zorin-Kranich's convergence theorem [**26**] which holds in the generality that we have just set up.

THEOREM 4.1. (Zorin-Kranich's convergence theorem) *Let $f_1, f_2, \ldots, f_d \in L^\infty(X, \mu)$. Then the limit*

$$\lim_\kappa \frac{1}{|\Phi_\kappa|} \sum_{\gamma \in \Phi_\kappa} \prod_{i=1}^d f_i \circ T_{[1,i]}^\gamma$$

exists in $L^2(X, \mu)$ and is independent of the choice of the (left) Følner net (Φ_κ).

This is exactly the claim in equation (1) in Theorem 1.10. It remains to prove the claim in equation (2) in the same theorem. This is achieved as follows. Consider the set of d-fold couplings on $\tilde{X}^d$, which is the collection of Baire probability measures on $(\tilde{X}^d, \mathcal{B}a(\tilde{X}^d))$ all of whose coordinate projections are $\tilde{\mu}$. By Theorem 4.1, for each $x \in \tilde{X}$, the averages

$$\frac{1}{|\Phi_\kappa|} \sum_{\gamma \in \Phi_\kappa} \delta_{\tilde{T}_1^\gamma x, \tilde{T}_{[1,2]}^\gamma x, \ldots, \tilde{T}_{[1,d]}^\gamma x}$$

converge weakly to a measure λ in the set of d-fold couplings. This weak convergence implies that

$$\frac{1}{|\Phi_\kappa|} \sum_{\gamma \in \Phi_\kappa} \tilde{\mu}(\tilde{T}_1^{\gamma^{-1}}(E) \cap \tilde{T}_{[1,2]}^{\gamma^{-1}}(E) \cap \cdots \cap \tilde{T}_{[1,d]}^{\gamma^{-1}}(E)) \to \lambda(E^d)$$

for all $E \in \mathcal{B}a(\tilde{X})$. To establish multiple recurrence, it suffices to show that

$$\lambda(E_1 \times E_2 \times \cdots \times E_d) = 0 \Rightarrow \tilde{\mu}(E_1 \cap E_2 \cap \cdots \cap E_d) = 0. \tag{6}$$

Austin first introduces the notion of satedness and proves that any space has an extension that is sated. It is then not a loss of generality to assume that $(\tilde{X}, \tilde{\mu})$ is a sated space in the first place. The advantage of a sated space is that it constrains how a certain relevant σ-subalgebra of the space is lifted to any of its extensions. More about satedness can be found in §4.1.

Define

$$H_{i,j} := \{\gamma \in \Gamma^d : \gamma_{i+1} = \gamma_{i+2} = \cdots = \gamma_j\},$$
$$L_{i,j} := \{\gamma \in H_{i,j} : \gamma_l = 1 \text{ for all } l \notin (i, j]\}.$$

https://doi.org/10.1017/etds.2024.40 Published online by Cambridge University Press

Austin then constructs recursively a tower of Γ^d-spaces, as a variant of Host–Kra self-joinings [16, 17]:

$$(\tilde{Y}^{(d)}, \tilde{\nu}^{(d)}, S^{(d)}) \to (\tilde{Y}^{(d-1)}, \tilde{\nu}^{(d-1)}, \tilde{S}^{(d-1)}) \to \cdots \to (\tilde{Y}^{(0)}, \tilde{\nu}^{(0)}, \tilde{S}^{(0)}) = (\tilde{X}, \tilde{\mu}, \tilde{T}).$$

Assume that the tower has already been constructed up to some $j \leq d-1$. Define an $H_{d-j-1,d}$-action $\tilde{R}^{(j)}$ on $(\tilde{Y}^{(j)}, \tilde{\nu}^{(j)})$ by setting

$$\tilde{R}_i^{(j)} = \begin{cases} \tilde{S}_i^{(j)} & \text{for } i < d-j-1, \\ \tilde{S}_{[d-j-1,d]}^{(j)} & \text{for } i = d-j-1, \\ \mathrm{id} & \text{for } i = [d-j, d-1]. \end{cases}$$

An $H_{d-j-1,d}$-space is then constructed by the relative product

$$(\tilde{Z}^{(j+1)}, \tilde{\theta}^{(j+1)}, \tilde{R}^{(j+1)}) = (\tilde{Y}^{(j)} \times \tilde{Y}^{(j)}, \tilde{\nu}^{(j)} \otimes_{\Sigma_{\tilde{Y}^{(j)}}^{L_{d-j-1,d}}} \tilde{\nu}^{(j)}, (\tilde{S}^{(j)})_{H_{d-j-1,d}} \times \tilde{R}^{(j)}),$$

where for a subgroup H of Γ^d, we let $(\tilde{S})_H$ be the restriction of an action $\tilde{S}$ by Γ^d to an action by H, and for a subgroup L of Γ^d, we denote by $\Sigma_{\tilde{Y}}^{L}$ the σ-subalgebra of $\Sigma_{\tilde{Y}}$ of invariant sets with respect to the restriction of an action $\tilde{S}$ by Γ^d to an action by L.

Finally, a *lifting lemma* proves the existence of a Γ^d-space extension

$$\pi : (\tilde{Y}^{(j+1)}, \tilde{\nu}^{(j+1)}, \tilde{S}^{(j+1)}) \to (\tilde{Y}^{(j)}, \tilde{\nu}^{(j)}, \tilde{S}^{(j)}),$$

which admits a commutative diagram of $H_{d-j-1,d}$-spaces

$$\begin{tikzcd} & \tilde{Z}^{(j+1)} \arrow[dr] & \\ (\tilde{Y}^{(j+1)})_{H_{\{d-j-1,d\}}} \arrow[ur] \arrow[rr, "\pi"] & & (\tilde{Y}^{(j)})_{H_{\{d-j-1,d\}}} \end{tikzcd},$$

where Y_H is the H-space with the same probability space but with the action restricted to H.

The advantage of the sequence of extensions is a variant of Host–Kra inequality: the asymptotic behavior of the ergodic average of f_i (the term inside the limit of equation (1)) is governed by an integral of a product of those f_i lifted to $\tilde{Y}^{(d)}$.

These Host–Kra-like self-joinings admit characteristic subspaces. A closed subspace $V \leq L^2(\tilde{\mu})$ is *partially characteristic in position i* if the ergodic averages of $f_1, \ldots, f_d$ are asymptotically the same as those of $f_1, \ldots, f_{i-1}, P^V f_i, f_{i+1}, \ldots, f_d$, where P^V denotes the orthogonal projection onto the space V. In a sated space $(\tilde{X}, \tilde{\mu})$, it can be proven that

$$L^2\bigg(\tilde{\mu}\bigg| \bigvee_{l=0}^{i-1} \Sigma_{\tilde{X}}^{\tilde{T}_{[l;i]}} \vee \bigvee_{l=i+1}^{d} \Sigma_{\tilde{X}}^{\tilde{T}_{[i;l]}}\bigg)$$

is partially characteristic in position i. The significance of these characteristic subspaces is that $P^V f_d$ can be approximated by a finite sum of products of the form $h_0 h_1 \ldots h_{d-1}$, where each h_i is $\Sigma_{\tilde{X}}^{\tilde{T}_{[i;d]}}$-measurable. This then allows us to reduce an ergodic average of d functions to an ergodic average of $d-1$ functions.

Satedness helps to prove that some spaces related to the characteristic subspaces are relatively orthogonal. As an illustrative example, let us say we want to prove that $L^2(\tilde{\mu}|\Phi_1)$ and $L^2(\tilde{\mu}|\Phi_2)$ are relatively independent over $V(\tilde{X})$, where Φ_1, Φ_2 are σ-algebras over $\tilde{X}$ and $V(\cdot)$ is a functorial L^2-subspace ($V(\tilde{Z})$ is an L^2-subspace of $L^2(\tilde{Z})$; see Definition 4.3 for details). We assume that $\tilde{X}$ is a Ψ-sated space.

Let $f \in L^2(\tilde{\mu}|\Phi_1)$ and $g \in L^2(\tilde{\mu}|\Phi_2)$. We construct a relative product measure

$$(\tilde{Y}, \tilde{\nu}) = (\tilde{X}^2, \tilde{\mu} \otimes_{\Phi_1} \tilde{\mu})$$

and carefully define a Γ-action on the space. Let β_1 and β_2 be the projections of $\tilde{Y} = \tilde{X}^2$ onto the first and second coordinate, respectively. Since f is Φ_1-measurable, we have

$$\int_{\tilde{X}} fg\,d\tilde{\mu} = \int_{\tilde{Y}} (f \circ \beta_2)(g \circ \beta_2)\,d\tilde{\nu} = \int_{\tilde{Y}} (f \circ \beta_1)(g \circ \beta_2)\,d\tilde{\nu}. \tag{7}$$

Then we use V-satedness of $\tilde{X}$, which gives that $L^2(\tilde{\mu}) \circ \beta_1$ and $V(\tilde{Y})$ are relatively orthogonal over $V(\tilde{X}) \circ \beta_1$. We will need $g \circ \beta_2 \in V(\tilde{Y})$, so equation (7) equals

$$\int_{\tilde{Y}} (P^{V(\tilde{X})} f \circ \beta_1) \cdot (g \circ \beta_2)\,d\tilde{\nu}.$$

By the same line of reasoning as in equation (7), we have this equals

$$\int_{\tilde{X}} P^{V(\tilde{X})} f \cdot g\,d\tilde{\mu} = \int_{\tilde{X}} P^{V(\tilde{X})} f \cdot P^{V(\tilde{X})} g\,d\tilde{\mu},$$

as desired.

From here, Austin's ergodic version of Tao's removal lemma [**23**] yields equation (6).

The modifications we need to extend these arguments to our uncountable setup are as follows.

(i) Construction of sated extensions for **PrbAlg**-spaces.

(ii) Lifting lemma: extending a factor map relative to a subgroup to the whole group for **PrbAlg**-spaces and for uncountable groups.

These modifications are carried out in the following two subsections.

4.1. **PrbAlg**$_\Gamma$-*sated extensions.* In this section, we verify that probability algebra dynamical systems admit sated extensions.

Following the standard notation, if $\pi : (Y, \nu, S) \to (X, \mu, T)$ is an extension, we let $\pi^* f := f \circ \pi$ on $L^2(X)$. If $\mathcal{H}$ is a closed subspace of a Hilbert space, we denote by $P_{\mathcal{H}}$ the orthogonal projection onto $\mathcal{H}$.

Recall that if $\mathcal{H}_1, \mathcal{H}_2$, and $\mathcal{I}$ are closed subspaces of a Hilbert space, then $\mathcal{H}_1$ and $\mathcal{H}_2$ are said to be *relatively orthogonal* over $\mathcal{I}$ if for any $u \in \mathcal{H}_1$ and $v \in \mathcal{H}_2$, we have $\langle u, v\rangle = \langle P_{\mathcal{I}} u, P_{\mathcal{I}} v\rangle$. The following simple characterization of relative orthogonality will be useful.

LEMMA 4.2. *Suppose $\mathcal{H}_1, \mathcal{H}_2$, and $\mathcal{I}$ are closed subspaces of a Hilbert space. If in addition $\mathcal{I} \subset \mathcal{H}_2$, then $\mathcal{H}_1$ and $\mathcal{H}_2$ are relatively orthogonal over $\mathcal{I}$ if and only if for any $u \in \mathcal{H}_1$, $P_{\mathcal{I}}(u) = P_{\mathcal{H}_2}(u)$.*

Proof. Suppose $\mathcal{H}_1$ and $\mathcal{H}_2$ are relatively orthogonal over $\mathcal{I}$. Fix $u \in \mathcal{H}_1$. For any $v \in \mathcal{H}_2$, $\langle u, v\rangle = \langle P_{\mathcal{I}}u, P_{\mathcal{I}}v\rangle = \langle P_{\mathcal{I}}u, v\rangle$. Since $P_{\mathcal{I}}u \in \mathcal{H}_2$, we have $P_{\mathcal{H}_2}u = P_{\mathcal{H}_2}(P_{\mathcal{I}}u) = P_{\mathcal{I}}u$.

Conversely, suppose for any vector in $\mathcal{H}_1$, its projection onto $\mathcal{I}$ and $\mathcal{H}_2$ are the same. For any $u \in \mathcal{H}_1$ and $v \in \mathcal{H}_2$, we have

$$\langle P_{\mathcal{I}}u, P_{\mathcal{I}}v\rangle = \langle P_{\mathcal{I}}u, v\rangle = \langle P_{\mathcal{H}_2}u, v\rangle = \langle u, v\rangle,$$

as desired. □

Definition 4.3. (**PrbAlg**$_\Gamma$-sated extensions) A *functorial L^2-subspace* of **PrbAlg**$_\Gamma$-spaces is a composition of functors $V = W \circ L^2$, where W is a functor from the category **Hilb**$_\Gamma$ to **Hilb** sending any object $\mathcal{H}$ to a closed subspace $V(\mathcal{H})$ of $\mathcal{H}$ and any morphism ϕ from a **Hilb**$_\Gamma$-object $\mathcal{H}$ to a **Hilb**$_\Gamma$-object $\mathcal{K}$ to the restriction $V(\phi) : V(\mathcal{H}) \to V(\mathcal{K})$.

Let V be a functorial L^2-subspace of **PrbAlg**$_\Gamma$-spaces. A **PrbAlg**$_\Gamma$-space $X = (X, \mu)$ is said to be *V-sated* if for any **PrbAlg**$_\Gamma$-morphism $\pi : (Y, \nu) \to (X, \mu)$, the subspaces $\pi^*(L^2(X))$ and $V(Y)$ of $L^2(Y)$ are relatively orthogonal over their common further subspace $\pi^*(V(X))$. The condition is equivalent to $\pi^*(P_{V(X)}(h)) = P_{V(Y)}(\pi^*(h))$ for any $h \in L^2(X)$, by Lemma 4.2 and the inclusion relation $\pi^*(V(X)) \subset V(Y)$. Moreover, a **PrbAlg**$_\Gamma$-morphism $\pi : (Y, \nu) \to (X, \mu)$ is said to be *relatively V-sated* if for any further **PrbAlg**$_\Gamma$-morphism $\psi : (Z, \lambda) \to (Y, \nu)$, the subspaces $(\pi \circ \psi)^*(L^2(X))$ and $V(Z)$ are relatively orthogonal over $\psi^*(V(Y))$.

For the remainder of this section, we fix a functorial L^2-subspace V.

LEMMA 4.4. *Suppose that $\pi : Y \to X$ is a relatively V-sated* **PrbAlg**$_\Gamma$*-morphism and $\phi : Z \to Y$ is a* **PrbAlg**$_\Gamma$*-morphism. Then $\pi \circ \phi : Z \to X$ is relatively V-sated.*

Proof. Let $\psi : W \to Z$ be a **PrbAlg**$_\Gamma$-morphism. Fix $f \in L^2(X)$ and $g \in V(W)$. We want to show

$$\langle(\pi \circ \phi \circ \psi)^* f, g\rangle_{L^2(W)} = \langle\psi^* P_{V(Z)}((\pi \circ \phi)^* f), g\rangle_{L^2(W)}$$

(see Figure 2). Applying the definition of relative satedness of $Y \xrightarrow{\pi} X$ to the further extension $W \xrightarrow{\phi\circ\psi} Y$, we have

$$\langle(\pi \circ \phi \circ \psi)^* f, g\rangle_{L^2(W)} = \langle(\phi \circ \psi)^* P_{V(Y)}(\pi^* f), g\rangle_{L^2(W)}.$$

It remains to prove $P_{V(Z)}((\pi \circ \phi)^* f) = \phi^* P_{V(Y)}(\pi^* f)$. By the relative satedness of $Y \xrightarrow{\pi} X$ applied to the further extension $Z \xrightarrow{\phi} Y$, $V(Z)$ and $(\pi \circ \phi)^*(L^2(X))$ are relatively orthogonal over $\phi^*(V(Y))$. Lemma 4.2 gives the desired result. □

LEMMA 4.5. *Let $(A, \leq)$ be a directed set with no maximal element and $((X_\alpha)_{\alpha\in A}, (\pi_{\alpha_1,\alpha_2})_{\alpha_1,\alpha_2\in A,\alpha_1\leq\alpha_2})$ be an inverse system of* **PrbAlg**$_\Gamma$*-spaces with inverse limit $(X, (\pi_\alpha)_{\alpha\in A})$. Further, assume that for all $\alpha_1, \alpha_2 \in A$ with $\alpha_1 < \alpha_2$, the* **PrbAlg**$_\Gamma$*-morphism π_{α_1,α_2} is relatively V-sated. Then the inverse limit X is V-sated.*

Proof. Each **PrbAlg**$_\Gamma$-morphism $\pi_\alpha : X \to X_\alpha$ is relatively V-sated because we can factorize $\pi_\alpha = \pi_{\alpha,\alpha'} \circ \pi_{\alpha'}$ for some $\alpha < \alpha'$ and apply Lemma 4.4. Let $\psi : Y \to X$ be an

https://doi.org/10.1017/etds.2024.40 Published online by Cambridge University Press

$$
\begin{tikzcd}
& L^2(W) \arrow[dl] \arrow[dr] & & & \\
V(W) \arrow[dr] & & L^2(Z) \arrow[dl] \arrow[dr] & & \\
& V(Z) \arrow[dr] & & L^2(Y) \arrow[dl] \arrow[dr] & \\
& & V(Y) \arrow[dr] & & L^2(X) \arrow[dl] \\
& & & V(X) &
\end{tikzcd}
$$

FIGURE 2. The subspace relations in Lemma 4.4.

arbitrary further $\mathbf{PrbAlg}_\Gamma$-morphism. For any $g \in V(Y)$ and $f \in \bigcup_{\alpha\in A} \pi_\alpha^*(L^2(X_\alpha))$, we have

$$\langle \psi^* f, g\rangle_{L^2(Y)} = \langle \psi^* P_{V(X)}(f), g\rangle_{L^2(Y)}.$$

Since $\bigcup_{\alpha\in A} \pi_\alpha^*(L^2(X_\alpha))$ is dense in $L^2(X)$, f in the last equation can be replaced by any function in $L^2(X)$. Thus, $V(Y)$ and $\psi^*(L^2(X))$ are relatively orthogonal over $\psi^*(V(X))$, and so X is V-sated. $\square$

LEMMA 4.6. *Every* $\mathbf{PrbAlg}_\Gamma$*-space X has a relatively V-sated extension.*

Proof. We write all elements of $L^2(X)$ as $\{f_{\beta'}\}_{\beta'<\alpha'}$ for some limit ordinal number α'. For each ordinal $\gamma < \alpha'$, we let $A_\gamma = \{f_{\beta'} : \beta' \le \gamma\}$. Define a well ordering on the set $\bigcup_{\gamma<\alpha'}\{\gamma\} \times A_\gamma$ by the relation

$$(\gamma_1, f_{\beta_1'}) < (\gamma_2, f_{\beta_2'}) \Longleftrightarrow \gamma_1 < \gamma_2 \quad \text{or} \quad (\gamma_1 = \gamma_2 \text{ and } \beta_1' < \beta_2').$$

Since $\bigcup_{\gamma<\alpha'}\{\gamma\} \times A_\gamma$ is well ordered, there exists an ordinal α and an order-preserving bijection Ψ from $\{\beta : \beta < \alpha\}$ to $\bigcup_{\gamma<\alpha'}\{\gamma\} \times A_\gamma$. Since α' is a limit ordinal, for each $\gamma < \alpha'$, there is some $\gamma < \gamma' < \alpha'$. As a result, there is no maximal element in $\bigcup_{\gamma<\alpha'}\{\gamma\} \times A_\gamma$; in other words, α is a limit ordinal as well.

Let $\Phi = \Pi \circ \Psi$, where Π is the projection mapping to the second coordinate. For each $\gamma < \alpha$ and $f_{\beta'} \in L^2(X)$, we claim that there exists $\tau > \gamma$ such that $\Phi(\tau) = f_{\beta'}$. Suppose $\Psi(\gamma) = (\beta, g)$. We let $\beta_{\max} := \max\{\beta, \beta'\} + 1 < \alpha'$ and then $\Psi^{-1}(\beta_{\max}, f_{\beta'})$ is the desired τ. The interpretation is, when enumerating $L^2(X)$ by Φ, each function appears not only infinitely many times, but also arbitrarily late.

Resorting to transfinite induction, we construct a $\mathbf{PrbAlg}_\Gamma$-extension X_β of X for every $\beta < \alpha$ and a $\mathbf{PrbAlg}_\Gamma$-morphism ϕ_γ^β from X_β to X_γ for every $\gamma < \beta < \alpha$. Set $X_\emptyset := X$. Suppose $\epsilon \le \alpha$ is an ordinal and for each $\gamma < \beta < \epsilon$, X_β and ϕ_γ^β have been constructed.

https://doi.org/10.1017/etds.2024.40 Published online by Cambridge University Press

Case 1: ϵ is the successor of $\epsilon - 1$. For any $\mathbf{PrbAlg}_\Gamma$-extension $\eta : Z \to X_{\epsilon-1}$, we have

$$\begin{aligned}\|P_{V(Z)}((\phi_\emptyset^{\epsilon-1} \circ \eta)^* \Phi(\epsilon))\|_2 - \|P_{V(X_{\epsilon-1})}((\phi_\emptyset^{\epsilon-1})^* \Phi(\epsilon))\|_2 \\ \leq \|\Phi(\epsilon)\|_2 - \|P_{V(X_{\epsilon-1})}((\phi_\emptyset^{\epsilon-1})^* \Phi(\epsilon))\|_2\end{aligned}$$

since every orthogonal projection is a contraction. Hence, we find a $\mathbf{PrbAlg}_\Gamma$-extension $\phi_{\epsilon-1}^{\epsilon} : X_\epsilon \to X_{\epsilon-1}$ such that the difference

$$\|P_{V(X_\epsilon)}((\phi_\emptyset^{\epsilon-1} \circ \phi_{\epsilon-1}^{\epsilon})^* \Phi(\epsilon))\|_2 - \|P_{V(X_{\epsilon-1})}((\phi_\emptyset^{\epsilon-1})^* \Phi(\epsilon))\|_2$$

is at least half its supremum value over all extensions $\eta : Z \to X_{\epsilon-1}$. For any $\gamma < \epsilon - 1$, we set $\phi_\gamma^\epsilon := \phi_\gamma^{\epsilon-1} \circ \phi_{\epsilon-1}^\epsilon$.

Case 2: ϵ is a limit ordinal. Let $(Z_\epsilon, (\psi_\beta^\epsilon)_{\beta<\epsilon})$ be the inverse limit of the inverse system $((X_\beta)_{\beta<\epsilon}, (\phi_\gamma^\beta)_{\gamma\leq\beta<\epsilon})$. Let $\psi_\epsilon : X_\epsilon \to Z_\epsilon$ be a $\mathbf{PrbAlg}_\Gamma$-extension such that the difference

$$\|P_{V(X_\epsilon)}((\psi_\epsilon \circ \psi_\emptyset^\epsilon)^* \Phi(\epsilon))\|_2 - \|P_{V(Z_\epsilon)}((\psi_\emptyset^\epsilon)^* \Phi(\epsilon))\|_2$$

is at least half its supremum possible value over all extensions of Z_ϵ. Set $\phi_\gamma^\epsilon := \psi_\gamma^\epsilon \circ \psi_\epsilon$.

We now show that $\phi_\emptyset^\alpha : X_\alpha \to X$ is relatively V-sated. Let $\pi : Y \to X_\alpha$ be an arbitrary further extension. By Lemma 4.2, it is equivalent to showing that for any $f \in L^2(X)$,

$$P_{V(Y)}((\phi_\emptyset^\alpha \circ \pi)^* f) = \pi^* P_{V(X_\alpha)}((\phi_\emptyset^\alpha)^* f).$$

Since $\pi^*(V(X_\alpha)) \subset V(Y)$, it suffices to show

$$\|P_{V(Y)}((\phi_\emptyset^\alpha \circ \pi)^* f)\|_2 \leq \|P_{V(X_\alpha)}((\phi_\emptyset^\alpha)^* f)\|_2.$$

Suppose for contradiction, $\|P_{V(Y)}((\phi_\emptyset^\alpha \circ \pi)^* f)\|_2 > \|P_{V(X_\alpha)}((\phi_\emptyset^\alpha)^* f)\|_2$. We know $\|P_{V(X_\gamma)}((\phi_\emptyset^\gamma)^* f)\|_2$ is increasing in γ and bounded above by $\|f\|_2$. By the construction of Φ, f appears in the image of Φ infinitely many times. There exists an ordinal γ large enough such that $\Phi(\gamma) = f$ and one of the following holds:

(i) γ is a successor and

$$\begin{aligned}\|P_{V(X_\gamma)}((\phi_\emptyset^\gamma)^* f)\|_2 - \|P_{V(X_{\gamma-1})}((\phi_\emptyset^{\gamma-1})^* f)\|_2 \\ < \tfrac{1}{2}(\|P_{V(Y)}((\phi_\emptyset^\alpha \circ \pi)^* f)\|_2 - \|P_{V(X_\alpha)}((\phi_\emptyset^\alpha)^* f)\|_2);\end{aligned}$$

(ii) γ is a limit ordinal and

$$\begin{aligned}\|P_{V(X_\gamma)}((\phi_\emptyset^\gamma)^* f)\|_2 - \|P_{V(Z_\gamma)}((\psi_\emptyset^\gamma)^* f)\|_2 \\ < \tfrac{1}{2}(\|P_{V(Y)}((\phi_\emptyset^\alpha \circ \pi)^* f)\|_2 - \|P_{V(X_\alpha)}((\phi_\emptyset^\alpha)^* f)\|_2).\end{aligned}$$

Since $\|P_{V(X_\alpha)}((\phi_\emptyset^\alpha)^* f)\|_2 \geq \|P_{V(X_{\gamma-1})}((\phi_\emptyset^{\gamma-1})^* f)\|_2$ when γ is a successor and $\|P_{V(X_\alpha)}((\phi_\emptyset^\alpha)^* f)\|_2 \geq \|P_{V(Z_\gamma)}((\psi_\emptyset^\gamma)^* f)\|_2$ when γ is a limit ordinal, we have

$$\begin{aligned}\|P_{V(X_\gamma)}((\phi_\emptyset^\gamma)^* f)\|_2 - \|P_{V(X_{\gamma-1})}((\phi_\emptyset^{\gamma-1})^* f)\|_2 \\ < \tfrac{1}{2}(\|P_{V(Y)}((\phi_\emptyset^\alpha \circ \pi)^* f)\|_2 - \|P_{V(X_{\gamma-1})}((\phi_\emptyset^{\gamma-1})^* f)\|_2)\end{aligned}$$

https://doi.org/10.1017/etds.2024.40 Published online by Cambridge University Press

or

$$\begin{aligned}&\|P_{V(X_\gamma)}((\phi_\emptyset^\gamma)^* f)\|_2 - \|P_{V(Z_\gamma)}((\psi_\emptyset^\gamma)^* f)\|_2\\ &\quad < \tfrac{1}{2}(\|P_{V(Y)}((\phi_\emptyset^\alpha \circ \pi)^* f)\|_2 - \|P_{V(Z_\gamma)}((\psi_\emptyset^\gamma)^* f)\|_2),\end{aligned}$$

either of which contradicts our choice of X_γ. Thus, we have $\phi_\emptyset^\alpha : X_\alpha \to X$ is relatively V-sated. □

Lemmas 4.5 and 4.6 give the following theorem.

THEOREM 4.7. *If V is a functorial L^2-subspace of $\mathbf{PrbAlg}_\Gamma$-spaces, then for every $\mathbf{PrbAlg}_\Gamma$-space $X = (X, \mu)$, there is a $\mathbf{PrbAlg}_\Gamma$-morphism $\pi : Y \to X$ such that $Y = (Y, \nu)$ is V-sated.*

4.2. *Extending factors relative to subgroups.* In this section, we show how to extend a factor map relative to a subgroup to a factor map of the whole group. The corresponding result for countable groups and standard Lebesgue spaces is [2, Theorem 2.1].

THEOREM 4.8. *Let Γ be an arbitrary discrete group, not necessarily countable or amenable. Let H be a subgroup of Γ. Let $X = (X, \mu, T)$ be a $\mathbf{PrbAlg}_\Gamma$-system and $Y = (Y, \nu, S)$ a $\mathbf{PrbAlg}_H$-system. Denote by $X_H = (X, \mu, T|_H)$ the $\mathbf{PrbAlg}_H$-system, where $T|_H$ is the restriction of the group homomorphism $T : \Gamma \to \mathrm{Aut}(X, \mu)$ to H. If $\beta : Y \to X_H$ is a $\mathbf{PrbAlg}_H$-morphism, then there are a $\mathbf{PrbAlg}_\Gamma$-system $Z = (Z, \theta, R)$, a $\mathbf{PrbAlg}_\Gamma$-extension $\pi : Z \to X$, and a $\mathbf{PrbAlg}_H$-extension $\alpha : Z_H \to Y$ such that the diagram*

$$\begin{tikzcd} & Y \arrow[dr, "\beta"] & \\ Z_H \arrow[ur, "\alpha"] \arrow[rr, "\pi"] & & X_H \end{tikzcd}$$

commutes in $\mathbf{PrbAlg}_H$.

Proof. We pass to the canonical models $\tilde{X} = (\tilde{X}, \mathcal{B}a(\tilde{X}), \tilde{\mu}, \tilde{T})$, $\tilde{Y} = (\tilde{Y}, \mathcal{B}a(\tilde{Y}), \tilde{\nu}, \tilde{S})$, and $\tilde{\beta}$ of X, Y, and β, respectively. We construct Z, α, and π as follows. First, we construct a $\mathbf{CH}_\Gamma$-system $(\tilde{Z}, \tilde{R})$, and $\mathbf{CH}_\Gamma$-maps $\tilde{\alpha}$ and $\tilde{\pi}$ satisfying a related commutative diagram in the dynamical category $\mathbf{CH}_\Gamma$. (We denote by $\mathbf{CH}$ the category of compact Hausdorff spaces and continuous maps, and $\mathbf{CH}_\Gamma$ denotes the category of topological dynamical Γ-systems formed on compact Hausdorff spaces and continuous factor maps, where the Γ-action is given by homeomorphisms) Second, we construct a probability measure $\tilde{\theta}$ on $(\tilde{Z}, \mathcal{B}a(\tilde{Z}))$ and show that it preserves the $\tilde{R}$-action. Finally, we verify that this $\mathbf{CHPrb}_\Gamma$-system satisfies the right commutative diagram. We can then map this diagram to the dynamical categories of probability algebras using the deletion and abstraction functors `Alg` and `Abs`.

Step 1: we build a $\mathbf{CH}_\Gamma$-system $(\tilde{Z}, \tilde{R})$. Let

$$\tilde{Z} := \{(y_\gamma)_\gamma \in \tilde{Y}^\Gamma : y_{\gamma\eta} = \tilde{S}^{\eta^{-1}} y_\gamma \text{ and } \tilde{\beta}(y_\gamma) = \tilde{T}^{\gamma^{-1}} \tilde{\beta}(y_e) \text{ for all } \gamma \in \Gamma, \eta \in H\},$$

where e is the identity of Γ. Note that $\tilde{Z}$ is a compact subspace of $\tilde{Y}^{\Gamma}$ (this basically follows from the fact that $\tilde{S}, \tilde{T}$ act by homeomorphisms and $\tilde{\beta}$ is continuous). We can define a $\mathbf{CH}_{\Gamma}$-action $\tilde{R} : \Gamma \to \mathrm{Aut}(\tilde{Z})$ by

$$\tilde{R}^{\gamma'}((y_{\gamma})_{\gamma\in\Gamma}) = (y_{(\gamma'^{-1}\gamma)})_{\gamma\in\Gamma}.$$

(One easily checks that $\tilde{Z}$ is an $\tilde{R}$-invariant set so that $\tilde{R}$ is well defined).

We set $\tilde{\alpha} : \tilde{Z} \to \tilde{Y}$ to be $\tilde{\alpha}((y_{\gamma})_{\gamma\in\Gamma}) := y_e$ and $\tilde{\pi} := \tilde{\beta} \circ \tilde{\alpha}$. By construction, the diagram

$$\begin{tikzcd} & \tilde{Y} \arrow[dr, "\tilde{\beta}"] & \\ \tilde{Z}_H \arrow[ur, "\tilde{\alpha}"] \arrow[rr, "\tilde{\pi}"] & & \tilde{X}_H \end{tikzcd}$$

commutes in the dynamical category $\mathbf{CH}_H$. By construction of $\tilde{Z}$, the map $\tilde{\pi}$ is also a $\mathbf{CH}_{\Gamma}$-factor map.

Endow the space $\tilde{Z}$ with the Baire σ-algebra $\mathcal{B}a(\tilde{Z})$, which coincides with the restriction of $\mathcal{B}a(\tilde{Y}^{\Gamma}) = \mathcal{B}a(\tilde{Y})^{\otimes\Gamma}$ to $\tilde{Z}$ by [21, Lemma 2.1]. In particular, the maps in the previous diagram preserve Baire measurability.

Step 2: we construct a probability measure $\tilde{\theta}$ on $(\tilde{Z}, \mathcal{B}a(\tilde{Z}))$. Let $\{\nu_x\}_{x\in\tilde{X}}$ be the canonical disintegration (see Theorem 2.2) of $\tilde{\nu}$ with respect to the factor map $\tilde{\beta} : \tilde{Y} \to \tilde{X}$. For each $\gamma \in \Gamma$ and $x \in \tilde{X}$, define $\tilde{\nu}_{\gamma,x}$ on $(\tilde{Y}^{\gamma H}, \mathcal{B}a(\tilde{Y}^{\gamma H}))$ by

$$\tilde{\nu}_{\gamma,x}(E) := \nu_x(\{y \in \tilde{Y} : (\tilde{S}^{\eta^{-1}} y)_{\gamma\eta\in\gamma H} \in E\}).$$

Since we can identify the Baire σ-algebra $\mathcal{B}a(\tilde{Y}^{\gamma H})$ with the product σ-algebra $\mathcal{B}a(\tilde{Y})^{\otimes\gamma H}$, it follows that $\tilde{\nu}_{\gamma,x}$ is well defined by first verifying cylinder sets and then applying the π-λ theorem.

By the axiom of choice, we pick a representative from each left coset γH an element ω and denote their collection by Ω. We identify $\tilde{Y}^{\Gamma} = \Pi_{\omega\in\Omega}\tilde{Y}^{\omega H}$ so as to define a probability measure

$$\tilde{\nu}'_x := \otimes_{\omega\in\Omega}\tilde{\nu}_{\omega,\tilde{T}^{\omega^{-1}}x}$$

on $\mathcal{B}a(\tilde{Y}^{\Gamma}) = \mathcal{B}a(\tilde{Y}^{H})^{\otimes\Omega}$. We show that the definition of $\tilde{\nu}'_x$ is independent from the choice of representatives. If $x \in X$, $A \in \mathcal{B}a(\tilde{Y}^{\gamma H})$, and $\gamma_1 = \gamma_2\eta_1$ for some $\gamma_1, \gamma_2 \in \Gamma$ and $\eta_1 \in H$, which means $\gamma_1 H = \gamma_2 H$, then we have

$$\begin{aligned}
\tilde{\nu}_{\gamma_1,\tilde{T}^{\gamma_1^{-1}}x}(E) &= \nu_{\tilde{T}^{\gamma_1^{-1}}x}(\{y : (\tilde{S}^{\eta^{-1}} y)_{\gamma_1\eta\in\gamma_1 H} \in E\}) \\
&= \nu_{\tilde{T}^{\eta_1^{-1}}\tilde{T}^{\gamma_2^{-1}}x}(\{y : (\tilde{S}^{\eta^{-1}} y)_{\gamma_2\eta_1\eta\in\gamma_2 H} \in E\}) \\
&= \tilde{S}_*^{\eta_1^{-1}}\nu_{\tilde{T}^{\gamma_2^{-1}}x}(\{y : (\tilde{S}^{\eta^{-1}} y)_{\gamma_2\eta_1\eta\in\gamma_2 H} \in E\}) \\
&= \nu_{\tilde{T}^{\gamma_2^{-1}}x}(\{y : (\tilde{S}^{\eta^{-1}}\tilde{S}^{\eta_1^{-1}} y)_{\gamma_2\eta_1\eta\in\gamma_2 H} \in E\}) \\
&= \tilde{\nu}_{\gamma_2,\tilde{T}^{\gamma_2^{-1}}x}(E).
\end{aligned}$$

https://doi.org/10.1017/etds.2024.40 Published online by Cambridge University Press

Hence, a finite product of the probability measures $\tilde{\nu}_{\gamma,\tilde{T}^{\gamma^{-1}}x}$ for γ ranging from different left cosets is independent from the choice of representatives. By the uniqueness part of Carathéodory's extension theorem, we conclude that the product probability measure $\tilde{\nu}'_x$ is independent from the choice of representatives.

Next, we define a measure $\tilde{\nu}_x$ on $(\tilde{Z}, \mathcal{B}a(\tilde{Z}))$ by

$$\tilde{\nu}_x(E \cap \tilde{Z}) := \tilde{\nu}'_x(E)$$

for each $E \in \mathcal{B}a(\tilde{Y}^\Gamma)$. Note that $\tilde{Z}$ is a closed subset of $\tilde{Y}^\Gamma$, but may not be Baire measurable. Therefore, we need to check the well definedness of $\tilde{\nu}_x$. It suffices to show that

$$E \in \mathcal{B}a(\tilde{Y}^\Gamma) \quad \text{and} \quad E \cap \tilde{Z} = \emptyset \Rightarrow \tilde{\nu}'_x(E) = 0.$$

Since $\mathcal{B}a(\tilde{Y}^\Gamma) = \mathcal{B}a(\tilde{Y})^{\otimes\Gamma}$, E depends on only countably many coordinates. Hence, there exists $\{\gamma_i\}_{i=1}^\infty$ such that $E = E' \times \otimes_{\gamma\in\Gamma\setminus\{\gamma_i\}_{i=1}^\infty} \tilde{Y}$, where $E' \in \mathcal{B}a(\tilde{Y})^{\otimes\mathbb{N}}$. Let

$$\begin{aligned}\tilde{Z}^* = \{(y_\gamma)_\gamma \in \tilde{Y}^\Gamma : y_{\gamma_i\eta} = \tilde{S}^{\eta^{-1}} y_{\gamma_i} \text{ and } \tilde{\beta}(y_{\gamma_i}) \\ = \tilde{T}^{\gamma_i^{-1}}\tilde{\beta}(y_e) \ \text{ for all } i \geq 1 \text{ and } \gamma_i\eta \in \{\gamma_i : i \geq 1\}\}.\end{aligned}$$

Since $\tilde{Z} \cap E = \emptyset$ implies $\tilde{Z}^* \cap E = \emptyset$, to show $\tilde{\nu}'_x(E) = 0$, it suffices to show $\tilde{\nu}'_x(\tilde{Z}^*) = 1$ (since $\tilde{Z}^*$ only depends on countable many coordinates, it is guaranteed to be Baire measurable). We group the γ_i according to the left cosets ωH to which they belong. So suppose $\{\gamma_i\} = \bigcup_i \{\omega_i\eta_{i,j}\}_j$, where each $\eta_{i,j} \in H$, $\omega_i \in \Omega$. For each ω_i, we have

$$\tilde{\nu}_{\omega_i,\tilde{T}^{\omega_i^{-1}}x}(\{(y_{\omega_i\eta})_{\omega_i\eta\in\omega_i H} : y_{\omega_i\eta_{i,j}} = \tilde{S}^{\eta_{i,j}^{-1}} y_{\omega_i}, \tilde{\beta}(y_{\omega_i}) = \tilde{T}^{\omega_i^{-1}}x\}) = 1.$$

Note that

$$\otimes_{i=1}^\infty\{(y_{\omega_i\eta})_{\omega_i\eta\in\omega_i H} : y_{\omega_i\eta_{i,j}} = \tilde{S}^{\eta_{i,j}^{-1}} y_{\omega_i}, \tilde{\beta}(y_{\omega_i}) = \tilde{T}^{\omega_i^{-1}}x\} \times \otimes_{\omega\neq\omega_i}\tilde{Y}^{\omega H} \subset \tilde{Z}^*.$$

Thus, $\tilde{\nu}'_x(\tilde{Z}^*) = 1$ and consequently, $\tilde{\nu}_x$ is well defined.

For any set $A = A' \cap \tilde{Z} \in \mathcal{B}a(\tilde{Z})$ where $A' \in \mathcal{B}a(\tilde{Y}^\Gamma)$, we aim to prove the mapping $x \mapsto \tilde{\nu}_x(A)$ is Baire measurable. Suppose there are $\omega_1, \ldots, \omega_m \in \Omega$, $\eta_{i,1}, \ldots, \eta_{i,n_i} \in H$ for each $i \leq m$, and $A_{i,j} \in \mathcal{B}a(\tilde{Y})$ for all $i \leq m$ and $j \leq n_i$ such that $A' = \{(y_\gamma)_\gamma : y_{\omega_i\eta_{i,j}} \in A_{i,j} \text{ for all } i \leq m, j \leq n_i\}$. Then

$$\begin{aligned}\tilde{\nu}'_x(A') &= \prod_{i=1}^m \tilde{\nu}_{\omega_i,\tilde{T}^{\omega_i^{-1}}x}(\{(y_{\omega_i\eta})_{\eta\in H} : y_{\omega_i\eta_{i,j}} \in A_{i,j} \text{ for any } j \leq n_i\}) \\ &= \prod_{i=1}^m \nu_{\tilde{T}^{\omega_i^{-1}}x}(\tilde{S}^{\eta_{i,1}}(A_{i,1}) \cap \cdots \cap \tilde{S}^{\eta_{i,n_i}}(A_{i,n_i})). \qquad (8)\end{aligned}$$

Since $x \mapsto \nu_x$ is Baire measurable and the product of finitely many Baire measurable functions is still Baire measurable, it follows that $x \mapsto \tilde{\nu}_x(A)$ is Baire measurable whenever $A' \in \mathcal{B}a(\tilde{Y}^\Gamma)$, as the cylinder sets generate $\mathcal{B}a(\tilde{Y}^\Gamma)$. As a result, we are able to define

$$\tilde{\theta} := \int_{\tilde{X}} \tilde{\nu}_x \, d\tilde{\mu}(x).$$

Observe that each $\tilde{\nu}_x(\tilde{Z}) = \tilde{\nu}'_x(\tilde{Y}^\Gamma) = 1$. Therefore, $\tilde{\theta}$ is a probability measure as well.

Step 3: we verify that $(\tilde{Z}, \mathcal{B}a(\tilde{Z}), \tilde{\theta}, \tilde{R})$ is a $\mathbf{CHPrb}_\Gamma$-system satisfying the desired diagram.

We claim that $\tilde{R}$ is a measure-preserving transformation. Suppose $\gamma' \in \Gamma$, $x \in \tilde{X}$, $\omega_1, \ldots, \omega_m \in \Omega$, $\eta_{i,1}, \ldots, \eta_{i,n_i} \in H$ for each $i \leq m$, and $A_{i,j} \in \mathcal{B}a(\tilde{Y})$ for all $i \leq m$ and $j \leq n_i$. Then we obtain

$$\begin{aligned}
&\tilde{R}_*^{\gamma'} \tilde{\nu}'_x(\{(y_\gamma)_\gamma : y_{\omega_i \eta_{i,j}} \in A_{i,j} \text{ for all } i \leq m, j \leq n_i\}) \\
&\quad = \tilde{\nu}'_x(\{(y_\gamma)_\gamma : y_{\gamma'^{-1}\omega_i \eta_{i,j}} \in A_{i,j} \text{ for all } i \leq m, j \leq n_i\}) \\
&\quad = \prod_{i=1}^m \nu'_{\tilde{T}^{\omega_i^{-1}\gamma'} x}(\tilde{S}^{\eta_{i,1}}(A_{i,1}) \cap \cdots \cap \tilde{S}^{\eta_{i,n_i}}(A_{i,n_i})) \\
&\quad = \tilde{\nu}'_{\tilde{T}^{\gamma'} x}\{(y_\gamma)_\gamma : y_{\omega_i \eta_{i,j}} \in A_{i,j} \text{ for all } i \leq m, j \leq n_i\},
\end{aligned}$$

where the last two equalities follow from two applications of equation (8) (while in the first applications, we work with a family of representatives $\gamma'^{-1}\Omega$ instead of Ω). Therefore, $\tilde{R}_*^{\gamma'} \tilde{\nu}'_x = \tilde{\nu}'_{\tilde{T}^{\gamma'} x}$. By the definition of $\tilde{\nu}_x$, $\tilde{R}_*^{\gamma'} \tilde{\nu}_x = \tilde{\nu}_{\tilde{T}^{\gamma'} x}$. Since $\tilde{\mu}$ is $\tilde{T}$ -invariant, integrating $\tilde{R}_*^{\gamma'} \tilde{\nu}_x = \tilde{\nu}_{T^{\gamma'} x}$ over $\tilde{\mu}$ gives $\tilde{R}_*^{\gamma'} \tilde{\theta} = \tilde{\theta}$.

Recall that the map $\tilde{\alpha}$ is a $\mathbf{CH}_H$-factor map. Moreover, for any $x \in \tilde{X}$, we have $\tilde{\alpha}^* \tilde{\nu}_x = \nu_x$ by observing that for any $A \in \mathcal{B}a(\tilde{Y})$,

$$\tilde{\nu}_x(\alpha^{-1} A) = \tilde{\nu}'_x(A \times \otimes_{\gamma \neq e, \gamma \in \Gamma} Y) = \nu_x(A).$$

Therefore,

$$\tilde{\alpha}_* \tilde{\theta} = \int_{\tilde{X}} \tilde{\alpha}_* \tilde{\nu}_x \tilde{\mu}(dx) = \int_{\tilde{X}} \nu_x \tilde{\mu}(dx) = \tilde{\nu}, \tag{9}$$

which shows that $\tilde{\alpha}$ is a $\mathbf{CHPrb}_H$-factor map.

It remains to show that $\tilde{\pi}$ is a $\mathbf{CHPrb}_\Gamma$-factor map, but this is a direct consequence of equation (9):

$$\tilde{\pi}_* \tilde{\theta} = \beta^* \tilde{\nu} = \tilde{\mu}. \qquad \square$$

Acknowledgments. A.J. was supported by DFG research fellowship JA 2512/3-1. We thank John Griesmer for helpful comments. We are sincerely grateful to an anonymous referee for a very constructive and detailed report, which helped to state a beautiful strengthening of our uniformity result and improve the presentation.

References

[1] E. Ackelsberg, V. Bergelson and O. Shalom. Khintchine-type recurrence for 3-point configurations. *Forum Math. Sigma* **10** (2022), Article id e107.

[2] T. Austin. Non-conventional ergodic averages for several commuting actions of an amenable group. *J. Anal. Math.* **130** (2016), 243–274.

https://doi.org/10.1017/etds.2024.40 Published online by Cambridge University Press

[3] H. Bauer. *Measure and Integration Theory (De Gruyter Studies in Mathematics, 26)*. de Gruyter, Berlin, 2001; translation from the German by R. B. Burckel.

[4] V. Bergelson, B. Host and B. Kra. Multiple recurrence and nilsequences. *Invent. Math.* **160**(2) (2005), 261–303; with an appendix by I. Ruzsa.

[5] V. Bergelson, B. Host, R. McCutcheon and F. Parreau. Aspects of uniformity in recurrence. *Colloq. Math.* **84–85** (2000), 549–576.

[6] V. Bergelson and A. Leibman. Failure of Roth theorem for solvable groups of exponential growth. *Ergod. Th. & Dynam. Sys.* **24**(1) (2004), 45–53.

[7] V. Bergelson and R. McCutcheon. Central sets and a non-commutative Roth theorem. *Amer. J. Math.* **129**(5) (2007), 1251–1275.

[8] V. Bergelson, R. McCutcheon and Q. Zhang. A Roth theorem for amenable groups. *Amer. J. Math.* **119**(6) (1997), 1173–1211.

[9] V. Bergelson, D. Robertson and P. Zorin-Kranich. Triangles in Cartesian squares of quasirandom groups. *Combin. Probab. Comput.* **26**(2) (2017), 161–182.

[10] V. Bergelson, T. Tao and T. Ziegler. Multiple recurrence and convergence results associated to F_P^{ω}-actions. *J. Anal. Math.* **127** (2015), 329–378.

[11] Q. Chu. Multiple recurrence for two commuting transformations. *Ergod. Th. & Dynam. Sys.* **31**(3) (2011), 771–792.

[12] Q. Chu and P. Zorin-Kranich. Lower bound in the Roth theorem for amenable groups. *Ergod. Th. & Dynam. Sys.* **35**(6) (2015), 1746–1766.

[13] P. Durcik, R. Greenfeld, A. Iseli, A. Jamneshan and J. Madrid. An uncountable ergodic Roth theorem and applications. *Discrete Contin. Dyn. Syst.* **42**(11) (2022), 5509–5540.

[14] H. Furstenberg. Ergodic behavior of diagonal measures and a theorem of Szemeredi on arithmetic progressions. *J. Anal. Math.* **31** (1977), 204–256.

[15] H. Furstenberg and Y. Katznelson. An ergodic Szemerédi theorem for IP-systems and combinatorial theory. *J. Anal. Math.* **45** (1985), 117–168.

[16] B. Host. Ergodic seminorms for commuting transformations and applications. *Studia Math.* **195**(1) (2009), 31–49.

[17] B. Host and B. Kra. Nonconventional ergodic averages and nilmanifolds. *Ann. of Math. (2)* **161**(1) (2005), 397–488.

[18] A. Jamneshan. An uncountable Furstenberg–Zimmer structure theory. *Ergod. Th. & Dynam. Sys.* **43**(7) (2023), 2404–2436.

[19] A. Jamneshan and T. Tao. An uncountable Mackey–Zimmer theorem. *Studia Math.* **266**(3) (2022), 241–289.

[20] A. Jamneshan and T. Tao. Foundational aspects of uncountable measure theory: Gelfand duality, Riesz representation, canonical models, and canonical disintegration. *Fund. Math.* **261**(1) (2023), 1–98.

[21] A. Jamneshan and T. Tao. An uncountable Moore–Schmidt theorem. *Ergod. Th. & Dynam. Sys.* **43**(7) (2023), 2376–2403.

[22] O. Shalom. Multiple ergodic averages in abelian groups and Khintchine type recurrence. *Trans. Amer. Math. Soc.* **375**(4) (2022), 2729–2761.

[23] T. Tao. Szemerédi's regularity lemma revisited. *Contrib. Discrete Math.* **1**(1) (2006), 8–28.

[24] N. M. Walsh. Norm convergence of nilpotent ergodic averages. *Ann. of Math. (2)* **175**(3) (2012), 1667–1688.

[25] J. Wysoczański. On uniformly amenable groups. *Proc. Amer. Math. Soc.* **102**(4) (1988), 933–938.

[26] P. Zorin-Kranich. Norm convergence of multiple ergodic averages on amenable groups. *J. Anal. Math.* **130** (2016), 219–241.