TENSOR QUASI-RANDOM GROUPS

MARK SELLKE

(Communicated by Martin Liebeck)

ABSTRACT. Gowers [Combin. Probab. Comput. 17 (2008), pp. 363–387] elegantly characterized the finite groups G in which $A_1 A_2 A_3 = G$ for any positive density subsets A_1, A_2, A_3. This property, *quasi-randomness*, holds if and only if G does not admit a nontrivial irreducible representation of constant dimension. We present a dual characterization of *tensor quasi-random* groups in which multiplication of subsets is replaced by tensor product of representations.

1. INTRODUCTION

Many large finite groups G exhibit expansion and mixing phenomena. The former states that product sets $A_1 A_2 \subseteq G$ are always significantly larger than either of $A_1, A_2 \subseteq G$. The latter states that when $A_1, A_2, A_3 \subseteq G$ are fairly large, then the number of solutions to $a_1 a_2 = a_3$ for $a_i \in A_i$ is close to $\frac{|A_1 A_2 A_3|}{|G|^2}$. Such properties have attracted a lot of interest and are related to theoretical computer science via the notion of expander graphs [Mar88, LPS88]. Seminal papers have established such properties for certain simple groups; see for instance [Hel08, BV12, BGGT13, PS16, Ebe16].

The present work is inspired by a striking result of [Gow08] which qualitatively characterizes the groups G that exhibit expansion on large scales. More precisely, it characterizes the groups for which any three subsets A_1, A_2, A_3 of constant density must multiply to cover all of G. This is equivalent to stating that if $|A_1|, |A_2| = \Omega(|G|)$ then $|A_1 A_2| = (1 - o(1))|G|$. Throughout we consider the regime of large groups with $|G| \to \infty$ as all other parameters are fixed. We use the standard notations $A = O(B)$ or $B = \Omega(A)$ to indicate that $\frac{A}{B}$ is bounded, and $A = o(B)$ to indicate that $\frac{A}{B}$ tends to 0.

Theorem 1.1 ([Gow08]). *Let G be a finite group. The following are asymptotically equivalent up to dependence of constants, and define a quasi-random group:*

(1) *G has no $O(1)$-dimensional non-trivial irreducible representations.*
(2) *If $A_1, A_2, A_3 \subseteq G$ each have size $\Omega(|G|)$, then $A_1 \cdot A_2 \cdot A_3 = G$.*
(3) *Fix any constant $m \geq 3$. If $A \subseteq G$ has $|A| = \Omega(|G|)$, then $A^m = A \cdot A \cdots \cdot A = G$.*
(4) *G has neither an $O(1)$-size nor an abelian non-trivial quotient.*

We explain in Appendix A why the above equivalence follows from [Gow08] as it is not stated directly. It follows from criterion (4) that large non-abelian simple groups are quasi-random. This implies for instance that given a subset $S \subseteq G$ for

G simple, to show that $S^k = G$ for k not too large, it suffices to show that a small power of S has macroscopic $\Omega(|G|)$ size. Such an argument was used in [BNP08] to simplify the proof of [Hel08]. Of course the result above is interesting in its own right; the original application was that constant-density product-free sets do not exist in general finite groups.

Our purpose is to study a dual problem: given large G-representations V_i, when must $V_1 \otimes V_2 \otimes V_3$ contain all irreducible G-representations as subrepresentations? In such a case we say this tensor product *covers* $\mathrm{Irrep}(G)$, the set of (isomorphism classes of) irreducible G-representations. Another line of working on covering $\mathrm{Irrep}(G)$ is the Saxl conjecture, which asserts that $\mathrm{Irrep}(S_n)$ can be covered by a tensor square for n large enough - see [PPV16, Ike15, Li18, LS17]. The work [HSTZ13] establishes such a result in groups of Lie type, and [LST20, Sel20, LST21] study the number of tensor powers of a fixed irreducible representation needed to cover $\mathrm{Irrep}(G)$ in various cases. As we allow our tensor factors V_i to be reducible, it is not obvious how to best measure their size. We will use the Plancherel measure.

Definition 1.2. For G a finite group, the Plancherel measure M_G is a probability distribution on $\mathrm{Irrep}(G)$ which assigns the irreducible representation $\boldsymbol{\lambda}$ probability $M_G(\boldsymbol{\lambda}) = \frac{\dim(\boldsymbol{\lambda})^2}{|G|}$. For an arbitrary finite-dimensional G-representation V, let $M_G(V)$ denote the Plancherel measure of the set of distinct (up to isomorphism) irreducible subrepresentations $\boldsymbol{\lambda} \subseteq V$.

Our main result, a dual version of Theorem 1.1, characterizes which groups exhibit good tensor product expansion on large scales.

Theorem 1.3. *Let G be a finite group. The following are asymptotically equivalent up to dependence of constants, and define a tensor quasi-random or TQR group:*
  (1) *G contains no $O(1)$ sized non-trivial conjugacy class.*
  (2) *If V_1, V_2, V_3 are G-representations with $M_G(V_i) = \Omega(1)$ for all i, then $V_1 \otimes V_2 \otimes V_3$ covers $\mathrm{Irrep}(G)$.*
  (3) *Fix $m \geq 3$. If V is a G-representation with $M_G(V) = \Omega(1)$, then $M_G(V^{\otimes m}) > \frac{1}{2}$.*
  (4) *G contains neither an $O(1)$ size nontrivial normal subgroup, nor an $O(1)$ index normal subgroup (possibly equal to G) with non-trivial center.*

While [Gow08] turns subsets $A \subseteq G$ into representation-theoretic data via the Fourier transform, we turn representations into class functions via their characters. The crucial lowest dimension of a non-trivial irreducible representation is for us replaced by the smallest non-trivial conjugacy class of G. In both situations, a key insight is that if all nontrivial irreducible representations or conjugacy classes are large, then an appropriate ℓ^∞ norm must be quite small - see the proof of Theorem 2.3.

We remark that from the fourth criterion in Theorem 1.3 it follows that large simple groups are also TQR. In [LST20] it was conjectured that for simple groups, any irreducible representation $\boldsymbol{\lambda}$ requires only $O\left(\frac{\log |G|}{\log \dim \boldsymbol{\lambda}}\right)$ tensor powers to cover $\mathrm{Irrep}(G)$ - this is an easy lower bound since covering $\mathrm{Irrep}(G)$ requires large dimension. This conjecture was proved there in bounded rank groups of Lie type, and then for S_n in [Sel20] and subsequently A_n in [LST21]. Because simple groups are TQR, to establish the conjecture it suffices to show that a small tensor power of $\boldsymbol{\lambda}$ has Plancherel measure $\Omega(1)$ (or even a bit smaller depending on the group in

question), at which point one could apply Theorem 1.3 to finish. This idea was used in [BNP08] to simplify the landmark result of [Hel08] on covering $SL_2(\mathbb{Z}/p\mathbb{Z})$ by products of subsets. As in [BNP08], our methods do not seem to be helpful for showing growth at small scales, but are only able to show that products of large representations quickly cover everything.

In Section 4 we specialize our results to the *tensor product Markov chains* studied in [Ful08, Ful04, Ful10, BDLT19]. The result is that the tensor quasi-randomness of a group G characterizes whether certain tensor product chains mix in constant time.

Corollary 1.4. *Associate with a G-representation V its reduced representation $\widetilde{V}$ as in Definition* 1.6. *For any $\varepsilon > 0$, if G is large and TQR, then when $M_G(V) = \Omega(1)$ the uniform ε-mixing time of the tensor product Markov chain given by $\cdot \otimes \widetilde{V}$ is at most 3. Conversely if G is large and not TQR then the total variation $\frac{1}{4}$-mixing times of the chains $\cdot \otimes \widetilde{V}$ are arbitrarily large for suitable G-representations V.*

1.1. Preliminaries.

Definition 1.5. Direct sum and tensor product of class functions are defined by element-wise sum/product (the notation is chosen to emphasize the underlying representations). For any character χ we denote by χ_0 the function $\chi_0(g) = 1_{g \neq e} \cdot \chi(g)$, so that $\chi(g) = \chi(e) \cdot 1_{g=e} + \chi_0(g)$.

In [Gow08], subsets $S \subseteq G$ correspond to their characteristic functions $1_S(x) = 1_{x \in S}$, and Fourier analysis is performed on these functions. For us, the corresponding "right" version of a G-representation V is the *reduced representation*. Throughout the paper we identify G-representations with their set of isomorphism classes of irreducible representations. We will use ℓ^p norms $|\cdot|_p$ on functions $f : G \to \mathbb{C}$ with counting measure on G, so that $|f|_p = \left(\sum_{g \in G} |f(g)|^p \right)^{1/p}$.

Definition 1.6. For V a G-representation, the corresponding reduced representation is given by

$$\widetilde{V} = \bigoplus_{\boldsymbol{\lambda} \in V} \dim(\boldsymbol{\lambda}) \cdot \boldsymbol{\lambda}.$$

Hence $\widetilde{V}$ depends only on the set of distinct irreducibles contained in V, and takes $\mathrm{Irrep}(G)$ to the regular representation. We also define the reduced character function $\widetilde{\chi}^V : G \to \mathbb{C}$ via:

$$\widetilde{\chi}^V(g) = \frac{1}{|G|} \cdot \chi^{\widetilde{V}}(g) = \frac{1}{|G|} \sum_{\boldsymbol{\lambda} \in \mathrm{Irrep}(G)} 1_{\boldsymbol{\lambda} \in V} \cdot \dim(\boldsymbol{\lambda}) \cdot \chi^{\boldsymbol{\lambda}}(g).$$

The regular representation has reduced character $1_{g=e}$. In general $\widetilde{\chi}^V(e) = M_G(V)$ and $|\widetilde{\chi}^V(e)|_2 = \sqrt{M_G(V)}$. The following covering criterion shows that if the ℓ^1 mass of a character χ is concentrated on $\chi(e)$, then the corresponding representation must contain all irreducibles.

Lemma 1.7. *Let V be a G-representation and let $\chi : G \to \mathbb{C}$ be a class function in the $\mathbb{C}$-linear span of $\{\chi^{\boldsymbol{\lambda}} | \boldsymbol{\lambda} \in V\}$. Suppose also that $|\chi(e)| > \sum_{g \in G \setminus \{e\}} |\chi(g)|$. Then $S = \mathrm{Irrep}(G)$.*

Proof. We show $\langle \chi, \chi^{\boldsymbol{\lambda}} \rangle \neq 0$ for any $\boldsymbol{\lambda} \in \mathrm{Irrep}(G)$. Indeed,

$$\langle \chi, \chi^{\boldsymbol{\lambda}} \rangle = \chi(e) \chi^{\boldsymbol{\lambda}}(e) + \langle \chi_0, \chi_0^{\boldsymbol{\lambda}} \rangle.$$

Now by assumption $|\chi(e)| > |\chi_0|_{\ell^1}$, and since $\boldsymbol{\lambda}$ is a genuine G-representation $|\chi^{\boldsymbol{\lambda}}(e)| \geq |\chi_0^{\boldsymbol{\lambda}}|_{\ell^\infty}$. Therefore

$$|\langle \chi_0, \chi_0^{\boldsymbol{\lambda}} \rangle| \leq |\chi_0|_1 \cdot |\chi_0^{\boldsymbol{\lambda}}|_\infty < |\chi(e) \chi^{\boldsymbol{\lambda}}(e)|.$$

Hence $\langle \chi, \chi^{\boldsymbol{\lambda}} \rangle \neq 0$ completing the proof. $\qquad \square$

2. Covering $Irrep(G)$ when all conjugacy classes are large

In this section we prove the implication $(1) \implies (2)$ of Theorem 1.3.

As a warmup we reprove a result from our previous work [Sel20] on covering $\mathrm{Irrep}(G)$ by a tensor product of two representations. We consider this a dual to the trivial statement that if $A_1, A_2 \subseteq G$ satisfy $|A_1| + |A_2| > |G|$ then $A_1 A_2 = G$.

Theorem 2.1 ([Sel20])**.** *Suppose* $M_G(V_1) + M_G(V_2) > 1$. *Then* $V_1 \otimes V_2$ *covers* $\mathrm{Irrep}(G)$.

Proof. For convenience take $a_i = M_G(V_i)$. Then

$$\left(\widetilde{\chi}^{V_1} \otimes \widetilde{\chi}^{V_2} \right)(g) = a_1 a_2 \cdot 1_{g=e} + \widetilde{\chi}_0^{V_1} \otimes \widetilde{\chi}_0^{V_2}(g).$$

The Cauchy-Schwarz inequality implies:

$$|\widetilde{\chi}_0^{V_1} \otimes \widetilde{\chi}_0^{V_2}|_1 \leq |\widetilde{\chi}_0^{V_1}|_2 |\widetilde{\chi}_0^{V_2}|_2 = \sqrt{(a_1 - a_1^2)(a_2 - a_2^2)}.$$

The assumption $a_1 + a_2 > 1$ implies $a_1 > 1 - a_2$ and $a_2 > 1 - a_1$. Therefore $a_1 a_2 > (1 - a_1)(1 - a_2)$, and so $a_1^2 a_2^2 > (a_1 - a_1^2)(a_2 - a_2^2)$. Therefore $\left(\widetilde{\chi}^{V_1} \otimes \widetilde{\chi}^{V_2} \right)(e) > 0$, concluding the proof. $\qquad \square$

Let $c(G)$ denote the minimal size of a non-trivial conjugacy class in G. We now show that tensor triple products cover $\mathrm{Irrep}(G)$ when $c(G)$ is large. The idea is to use Holder's inequality with $\frac{1}{2} + \frac{1}{2} + \frac{1}{\infty} = 1$.

Lemma 2.2. *Let* V *be a G-representation; then* $|\widetilde{\chi}_0^V|_\infty \leq c(G)^{-1/2}$.

Proof. Certainly $|\widetilde{\chi}_0^V|_2 \leq 1$. Moreover $\widetilde{\chi}_0^V$ is constant on conjugacy classes, all of which have size at least $c(G)$. This implies the claim. $\qquad \square$

Theorem 2.3. *For representations* V_1, V_2, V_3 *of G, if*

$$M_G(V_1) M_G(V_2) M_G(V_3) > c(G)^{-1/2},$$

then $V_1 \otimes V_2 \otimes V_3$ *contains all irreducible representations of G.*

Proof. We consider the character values on $e \in G$ and on all other values. $\widetilde{\chi}^{V_i}(e) = M_G(V_i)$ and so

$$\widetilde{\chi}^{V_1} \otimes \widetilde{\chi}^{V_2} \otimes \widetilde{\chi}^{V_3}(e) = M_G(V_1) M_G(V_2) M_G(V_3).$$

On the other hand, Holder's inequality and Lemma 2.2 imply:

$$|\widetilde{\chi}_0^{V_1} \otimes \widetilde{\chi}_0^{V_2} \otimes \widetilde{\chi}_0^{V_3}|_1 \leq |\widetilde{\chi}_0^{V_1}|_2 \cdot |\widetilde{\chi}_0^{V_2}|_2 \cdot |\widetilde{\chi}_0^{V_3}|_\infty \leq c(G)^{-1/2}.$$

$\qquad \square$

This establishes the implication $(1) \implies (2)$ of Theorem 1.3. In fact examination of the proof above shows slightly more, in analogy with Theorem 3.3 of [Gow08]. Let V_1, V_2, V_3 be representations of a TQR group G with Plancherel measures $M_G(V_i) = a_i$. Then the multiplicity of an irreducible $\boldsymbol{\lambda}$ in $\widetilde{V}_1 \otimes \widetilde{V}_2 \otimes \widetilde{V}_3$ is $|G|^2(a_1 a_2 a_3 \dim(\boldsymbol{\lambda}) + |\widetilde{\chi}_0^{V_1} \otimes \widetilde{\chi}_0^{V_2} \otimes \widetilde{\chi}_0^{V_3}|_1)$, where the last term is at most $c(G)^{-1/2}$. Hence in a TQR group, such a tensor product of reduced representations contains every irreducible $\boldsymbol{\lambda}$ a number of times approximately proportional to $\dim(\boldsymbol{\lambda})$. The L^∞ mixing time result is a slight generalization of this argument.

3. Remaining proofs for Theorem 1.3

The implication $(2) \implies (3)$ of Theorem 1.3 is clear. We now show the implication $(4) \implies (1)$ of Theorem 1.3, and in Section 3.1 show $(3) \implies (4)$.

Proof of Theorem 1.3, implication $(4) \implies (1)$. We go by contradiction and assume G contains neither a small non-trivial conjugacy class C at most $k = O(1)$ in size nor a normal subgroup at most $k!$ in size. We will show that G contains a constant-index normal subgroup with non-trivial center.

First, if $k = 1$ then G itself has nontrivial center so we may assume $k > 1$. Then G acts on C by conjugation; this defines a non-trivial homomorphism $\phi : G \to S_k$. Therefore $N = \ker \phi \trianglelefteq G$ is a normal subgroup of G commuting with each element of C. Its index is $|G/N| \leq k!$. The subgroup H generated by C is normal (as C is a conjugacy class) and commutes with N. Since H is a normal subgroup of G, by assumption $|H| > k!$. Let $K = H \cap N$. We have

$$|K| \geq \frac{|H||N|}{|G|} \geq \frac{|H|}{k!} > 1.$$

Since H and N commute, K is central in N. Therefore G contains a constant index subgroup N with nontrivial center K. This shows $(4) \implies (1)$. $\square$

3.1. **Tensor power condition implies normal subgroup condition.** To prove the implication $(3) \implies (4)$ we begin with two preparatory lemmas in additive combinatorics. Throughout, if n is a positive integer and B is a subset of an abelian group, we denote by nB the n-fold sum $B + B + \cdots + B$.

Lemma 3.1. *Let k, m, n be positive integers and let $B \subseteq K$ be a subset of an abelian group K with size $|B| = k + 1$. Then mnB can be covered by $(10km)^k$ translates of nB. In particular $\frac{|mnB|}{|nB|} \leq (10km)^k$.*

Proof. It suffices to show the result when $B = \{0, e_1, \ldots, e_k\}$ consists of zero and standard basis vectors in $K = \mathbb{Z}^k$. Indeed given any such identification, the resulting homomorphism from $\mathbb{Z}^k \to K$ transfers a covering inside $\mathbb{Z}^k$ to a covering inside K. Observe that $nB \supseteq \{0, \ldots, j-1\}^k$ for $j = 1 + \lfloor \frac{n}{k} \rfloor$, while $mnB \subseteq \{0, \ldots, J-1\}^k$ for $J = mn + 1$. As $J \leq 10kmj$ the result follows. $\square$

Lemma 3.2. *Let K be a finite non-trivial abelian group admitting an action of automorphisms by a group L of size k. Then there exists a subset $A \subseteq K$ with $|A| = \Omega_{k,m}(|K|)$ and $|mA| \leq \frac{|K|}{2}$.*

Proof. Fix a small, positive value $\varepsilon < \frac{1}{(10km)^{k+1}}$. We construct A using the following iterative algorithm initialized with $A = \{e\}$ and $a \in K \backslash \{e\}$. Throughout, for $a \in K$,

we denote by $La \subseteq K$ the orbit of a under the action of L.

(1) If $A \neq A + La$, update $A \leftarrow A \cup (A + La)$.

(2) If $A = A + La$, update a to any element of $K \setminus A$ and return to step (1).

(3) At the first time that $|A| \geq \varepsilon |K|$, terminate and output the current set $\widehat{A}$.

This algorithm repeats multiple iterations of step (1) interrupted by single iterations of step (2) until terminating in step (3). To show that the outputed set $\widehat{A}$ satisfies the conditions of the lemma it remains to show that $mA \subsetneq K$. First observe that the size $|A|$ grows by a factor of at most $k+1$ at each iteration because $|La| \leq |L| \leq k$ for each $a \in K$.

We claim that all times, $|mA| \leq (10km)^k |A|$ holds. Indeed let $a_1, a_2, \ldots, a_t$ be the sequence of values of a so far in the algorithm, and let A_{t-1} be the set A at the time that $a \leftarrow a_t$ was updated. Then A_{t-1} is exactly the subgroup generated by the sets La_i for $i \leq t-1$. Letting $B = La_t \cup e$, the current set A may be written as

$$A = A_{t-1} + nB$$

for some non-negative integer n, where $B = e \cup La$ has size $|B| = k + 1$. Since $A_{t-1} \subseteq K$ is a subgroup, we may set $\psi : K \to K/A_{t-1}$ to be the natural quotient map. Then Lemma 3.1 yields $\frac{|mA|}{|A|} = \frac{|nm\psi(B)|}{|\psi(nB)|} \leq (10km)^k$ proving the claim.

Since $|A|$ grows by a factor of at most $k + 1$ each iteration, as long as $|K| \geq \varepsilon^{-1}$ the output set $\widehat{A}$ satisfies $\frac{|\widehat{A}|}{|K|} < \frac{1}{10 \cdot (10km)^k}$. Therefore $|m\widehat{A}| < \frac{|K|}{10}$ as desired. On the other hand if $|K| \leq \varepsilon^{-1}$ then step (3) immediately outputs $\widehat{A} = \{e\}$. In this case $|m\widehat{A}| = 1$, which also concludes the proof as K is assumed to be non-trivial. $\square$

We now complete the proof of Theorem 1.3.

Proof of Theorem 1.3, *implication* $(3) \implies (4)$. It is easy to see that if G has a constant size normal subgroup N, then condition (3) fails - simply take V to be the regular representation on the quotient G/N, viewed as a G-representation. Then in $V^{\otimes m}$, N will still act trivially so $V^{\otimes m}$ cannot cover $\mathrm{Irrep}(G)$. Moreover $M_G(V) = \frac{1}{|N|} = \Omega(1)$ because irreducible G/N representations are also irreducible G-representations. Therefore V contradicts condition (3).

For the main part of the proof, take $N \trianglelefteq G$ with $|G/N| = O(1)$ and with non-trivial center $K \subseteq N$. We will construct a G-representation V with $M_G(V) = \Omega(1)$ such that $M_G(V^{\otimes m}) \leq \frac{1}{2}$. The idea is to construct G-representations from K-characters via induction. Letting $\theta : K \to \mathbb{C}^{\times}$ be a multiplicative character, we define $V_\theta = \mathrm{Ind}_K^N \theta$. From the Mackey formula and centrality of K in N,

$$\chi^{V_\theta}(g) = \left\{ \begin{array}{ll} 0, & \text{for } g \notin K \\ \frac{|N|}{|K|} \cdot \theta(g), & \text{for } g \in K \end{array} \right\}.$$

It follows that the induced representations V_θ remain mutually orthogonal and satisfy $\bigoplus_{\theta \in K^*} V_\theta = \chi_N^{Reg}$. This implies that they are of the form

$$V_\theta := \bigoplus_{\boldsymbol{\lambda} \in \Lambda_\theta \subseteq \mathrm{Irrep}(N)} (\dim \boldsymbol{\lambda}) \boldsymbol{\lambda},$$

where $\{\Lambda_\theta : \theta \in K^*\}$ is a partition of $\mathrm{Irrep}(N)$. Hence their N-Plancherel measures satisfy

$$M_N(V_\theta) = \frac{\dim(V_\theta)}{|N|} = \frac{1}{|K|}.$$

We induct again from N to G, obtaining G-representations

$$W_\theta := \operatorname{Ind}_K^G \theta = \operatorname{Ind}_N^G V_\theta.$$

Observe that any inner automorphism on G restricts to an automorphism on N, hence preserves its center K. Therefore conjugation defines a group homomorphism $\varphi_0 : G \to Aut(K)$. $\varphi_0(n)$ acts trivially for any $n \in N$ so we obtain a quotient map $\varphi : G/N \to Aut(K)$. The Mackey formula again takes a simple form. Writing $\varphi_g(\theta)(\cdot) = \theta(\varphi(g)(\cdot))$ we may view φ as a G/N action on K^*. Then

$$\chi^{W_\theta}(k) = \frac{|N|}{|K|} \sum_{g_i \in G/N} \varphi_{g_i}(\theta)(k). \tag{1}$$

Partition the set K^* of K-characters into equivalence classes $(S_i)_{i=1}^j$ by the action of G/N. (1) implies that if θ, θ' are in the same equivalence class then $W_\theta = W^{\theta'}$, while if not then $W_\theta, W^{\theta'}$ have orthogonal characters. Let $W^{S_i} = W_\theta$ if $\theta \in S_i$. Next observe that $\bigoplus_{\theta \in K^*} \theta \simeq V_K^{Reg}$ and so

$$\bigoplus_{i \leq j} |S_i| \cdot W^{S_i} = \bigoplus_{\theta \in K^*} W_\theta = \operatorname{Ind}_K^G V_K^{Reg} = V_G^{Reg}.$$

Since the representations $|S_i| \cdot W^{S_i}$ have pairwise orthogonal characters and their direct sum is the regular representation, they must be of the form

$$|S_i| \cdot W^{S_i} = \sum_{\boldsymbol{\lambda} \in \Lambda_i \subseteq \operatorname{Irrep}(G)} (\dim \boldsymbol{\lambda})\boldsymbol{\lambda}$$

for some partition $\{\Lambda_i : i \in [j]\}$ of $\operatorname{Irrep}(G)$. Therefore $M_G(W^{S_i}) = \frac{|S_i| \dim(W^{S_i})}{|G|} = \frac{|S_i|}{|K|}$. Altogether for any φ-invariant subset $S \subseteq K^*$, letting $V^S = \bigoplus_{\theta \in S} \theta$ we have

$$M_G(\operatorname{Ind}_K^G(V^S)) = M_K(V^S).$$

By Lemma 3.2, there exists $A \subseteq K^*$ invariant under the action φ of G/N of size $|A| = \Omega_{k,m}(|K|)$, and which also satisfies $|mA| \leq \frac{|K^*|}{2}$. We take $V = \bigoplus_{\theta \in A} W_\theta$. It is not difficult to see by φ-invariance of A that $\widetilde{\chi}^V(g)$ vanishes outside of K, and that for any $k \in K$:

$$\widetilde{\chi}^V(k) = \frac{|N|}{|G|} \sum_{\theta \in A} \theta(k).$$

Hence the m-th tensor power of this character is in the $\mathbb{C}$-linear span of $\{\chi^{W_\theta} | \theta \in mA\}$. As a result, $M_G(V^{\otimes m}) \leq \frac{1}{2}$, completing the proof. $\qquad\square$

4. Application to tensor product Markov chains

Here we prove Corollary 1.4 on tensor product Markov chains. We consider the chains $\cdot \otimes \widetilde{V}$ for V a G-representation, where the reduced representation $\widetilde{V}$ is as in Definition 1.6. To take a step in these Markov chains, from a starting representation $\boldsymbol{\lambda}$ one samples from the irreducible subrepresentations of $\boldsymbol{\lambda} \otimes \widetilde{V}$, weighted by the product of their multiplicity and dimension. The Plancherel measure is a stationary distribution for any such Markov chain, and the distribution $p_t(\boldsymbol{\lambda}, \cdot)$ after t steps can be generated by sampling in the same way from the irreducible subrepresentations of $\boldsymbol{\lambda} \otimes (\widetilde{V})^{\otimes t}$. Recall also that the ℓ^∞ (or uniform) ε-mixing time is given by $\inf \left\{ t \geq 0 : \max_{\boldsymbol{\lambda}, \boldsymbol{\mu}} \left| \frac{p_t(\boldsymbol{\lambda}, \boldsymbol{\mu})}{M_G(\boldsymbol{\mu})} - 1 \right| \leq \varepsilon \right\}$ while the ℓ^1 (or total variation) ε-mixing

time is $\inf\{t \geq 0 : \max_{\boldsymbol{\lambda},\boldsymbol{\mu}} |p_t(\boldsymbol{\lambda}, \boldsymbol{\mu}) - M_G(\boldsymbol{\mu})| \leq \varepsilon\}$. There are several other related definitions for mixing times, such as the ℓ^p mixing time and the separation distance mixing time. Among all of these choices, uniform mixing time is the largest while total variation mixing time is the smallest. Therefore in the setting of Corollary 1.4, when G is TQR all of these chains mix within 3 steps, while when G is not TQR none of these chains are guaranteed to mix within $O(1)$ steps.

Proof of Corollary 1.4. We begin with the first statement. Let $\boldsymbol{\lambda}, V$ be G-representations with $\boldsymbol{\lambda}$ irreducible and V (without loss of generality) reduced. Then:

$$|\widetilde{\chi}_0^{V^{\otimes 3}} \otimes \widetilde{\chi}_0^{\boldsymbol{\lambda}_i}|_1 \leq M_G(\boldsymbol{\lambda}) \cdot |\widetilde{\chi}_0^{V^{\otimes 3}}|_1 \leq |\widetilde{\chi}_0^V|_2 \cdot |\widetilde{\chi}_0^V|_2 \cdot |\widetilde{\chi}_0^V|_\infty \leq M_G(\boldsymbol{\lambda}_i) \cdot c(G)^{-1/2}.$$

On the other hand,

$$\widetilde{\chi}^{V^{\otimes 3}}(e)\widetilde{\chi}^{\boldsymbol{\lambda}}(e) = M_G(\boldsymbol{\lambda}) \cdot M_G(V)^3.$$

Let $p_3(\boldsymbol{\lambda}, \boldsymbol{\mu})$ denote the probability to reach $\boldsymbol{\mu}$ from $\boldsymbol{\lambda}$ in exactly 3 steps. Then with $\propto$ indicating proportionality as $\boldsymbol{\mu}$ varies over $\mathrm{Irrep}(G)$,

$$p_3(\boldsymbol{\lambda}, \boldsymbol{\mu}) \propto \langle \widetilde{\chi}^{\boldsymbol{\mu}}, \widetilde{\chi}^{V^{\otimes 3}} \otimes \chi^{\boldsymbol{\lambda}} \rangle \propto M_G(\boldsymbol{\mu})M_G(\boldsymbol{\lambda})M_G(V)^3 + \langle \widetilde{\chi}_0^{\boldsymbol{\mu}}, \widetilde{\chi}_0^{V^{\otimes 3}} \otimes \chi_0^{\boldsymbol{\lambda}} \rangle.$$

As $M_G(V) = \Omega(1)$ and G is TQR, it follows that $c(G)^{-1/2} = o(M_G(V)^3)$. Recall:

$$\langle \widetilde{\chi}_0^{\boldsymbol{\mu}}, \widetilde{\chi}_0^{V^{\otimes 3}} \otimes \chi_0^{\boldsymbol{\lambda}} \rangle \leq |\widetilde{\chi}_0^{\boldsymbol{\mu}}|_\infty |\widetilde{\chi}_0^{V^{\otimes 3}} \otimes \chi_0^{\boldsymbol{\lambda}}|_1 \leq M_G(\boldsymbol{\mu})M_G(\boldsymbol{\lambda}) \cdot c(G)^{-1/2}.$$

We conclude that for $c(G)$ large,

$$p_3(\boldsymbol{\lambda}, \boldsymbol{\mu}) \propto M_G(\boldsymbol{\mu})(1 + o(1)).$$

This implies the uniform mixing time result. For total variation non-mixing, we recall Theorem 1.3, assertion (3). Taking the contrapositive we see that if G is not TQR, then there is a G-representation V with $M_G(V) = \Omega_m(1)$ but $M_G(V^{\otimes m}) \leq \frac{1}{2}$ for any fixed m. Taking as starting point $\boldsymbol{\lambda}$ the trivial representation, after m steps at least half of $\mathrm{Irrep}(G)$ is still completely inaccessible to the tensor product Markov chain. Such a distribution must have large total variation distance from the Plancherel measure stationary distribution, which concludes the proof. $\square$

5. RELATIONS BETWEEN TQR AND QUASI-RANDOM GROUPS

Here we make some comments on the structure of TQR and quasi-random groups. We observe that any quasi-random group has a quotient which is both quasi-random and TQR, but there seems to be no analogous general way to go from a TQR group to a quasi-random group.

Corollary 5.1. *Any large center-free quasi-random group G is TQR.*

Proof. If G is not TQR, then it contains a nontrivial conjugacy class of size $k = O(1)$. As G is center-free, $k > 1$. Therefore we obtain a nontrivial conjugacy action $G \to S_k$ with nontrivial normal kernel $N \subsetneq G$. Then $|G/N| \leq k!$, so G contains a constant-sized quotient, contradicting quasi-randomness. $\square$

Corollary 5.2. *Any large quasi-random group G has a quotient H of super-constant size which is simultaneously quasi-random and TQR.*

Proof. Begin from G and repeatedly quotient out the center until a center-free quotient H of G is reached. Being quasi-random G contains no non-trivial abelian quotient, so $|H| > 1$ as reaching $|H| = 1$ requires coming from an abelian quotient in the previous step. G also contains no constant-size non-trivial quotient, implying $|H| = \omega(1)$ is super-constant. Since all quotients of H are quotients of G, condition (4) of Theorem 1.1 implies that H is also quasi-random. Moreover H is center-free; hence it is TQR as well. $\square$

Proposition 5.3. *Let $G = \mathbb{F}_p \rtimes \mathbb{F}_p^*$ be the group of affine bijections $x \to ax + b$ on $\mathbb{F}_p$. Then G is TQR, but none of its subgroups or quotients are quasi-random.*

Proof. The TQR property follows by considering conjugacy classes so we focus on the latter assertions. First we show that any non-trivial subgroup H of G has a non-trivial abelian quotient, implying G has no quasi-random subgroup. This holds because restriction of the quotient map $G \to \mathbb{F}_p^*$ defined by $(ax + b) \to a$ gives an abelian quotient unless H is contained in the set $\{x + b\}$. In the latter case H is abelian already.

Next we show that G cannot have a quasi-random quotient, and in fact that all non-trivial quotients of G are abelian. Indeed any irreducible representation of a quotient G/N pulls back to an irreducible representation of G. As G has order $p^2 - p$ and has irreducible representations of dimension only $p - 1$ and 1, simple size considerations imply only the 1-dimensional irreducible G-representations can be irreducible representations of a non-trivial quotient. Therefore any quotient of G is abelian, concluding the proof. $\square$

Appendix A. Statement of Theorem 1.1

Here we justify our statement of Theorem 1.1. In [Gow08, Theorems 3.3, 4.5, 4.6, 4.8] the following statements are shown to be equivalent for a finite group G:

  (A) G has no $O(1)$-dimensional non-trivial irreducible representations.
  (B) If $A_1, A_2, A_3 \subseteq G$ each have size $\Omega(|G|)$, then $a_1 a_2 = a_3$ has a solution in $(a_1, a_2, a_3) \in A_1 \times A_2 \times A_3$.
  (C) G has neither an $O(1)$-size nor an abelian non-trivial quotient.

It is easy to see that Assertion (B) is equivalent to assertion (2) of Theorem 1.1. As Assertion (2) of Theorem 1.1 clearly implies Assertion (3), it only remains to show that Assertion (3) of Theorem 1.1 implies one of the others. We will show it implies Assertion (1). Going by contrapositive, we suppose we are given a nontrivial homomorphism $\varphi : G \to U(k)$ for $k = O(1)$. Take the set A to be the preimage by φ of a small neighborhood of the identity in $U(k)$. Then A^m contains no element with negative trace, but $\varphi(G)$ must contain such elements as the average trace of $\varphi(g)$ for $g \in G$ uniformly random is 0. Moreover a simple volume argument on cosets of A shows that $\frac{|A|}{|G|} = \Omega_{k,m}(1)$. This proves Assertion (1) of Theorem 1.1 from Assertion (3).

References

[BDLT19] Georgia Benkart, Persi Diaconis, Martin W. Liebeck, and Pham Huu Tiep, *Tensor product Markov chains*, J. Algebra **561** (2020), 17–83, DOI 10.1016/j.jalgebra.2019.10.038. MR4135538

[BGGT13] Emmanuel Breuillard, Ben Green, Robert Guralnick, and Terence Tao, *Expansion in finite simple groups of Lie type*, J. Eur. Math. Soc. (JEMS) **17** (2015), no. 6, 1367–1434, DOI 10.4171/JEMS/533. MR3353804

[BNP08] László Babai, Nikolay Nikolov, and László Pyber, *Product growth and mixing in finite groups*, Proceedings of the Nineteenth Annual ACM-SIAM Symposium on Discrete Algorithms, ACM, New York, 2008, pp. 248–257. MR2485310

[BV12] Jean Bourgain and Péter P. Varjú, *Expansion in $SL_d(\mathbf{Z}/q\mathbf{Z})$, q arbitrary*, Invent. Math. **188** (2012), no. 1, 151–173, DOI 10.1007/s00222-011-0345-4. MR2897695

[Ebe16] Sean Eberhard, *Product mixing in the alternating group*, Discrete Anal., posted on 2016, Paper No. 2, 19, DOI 10.19086/da.610. MR3533301

[Ful04] Jason Fulman, *Card shuffling and the decomposition of tensor products*, Pacific J. Math. **217** (2004), no. 2, 247–262, DOI 10.2140/pjm.2004.217.247. MR2109933

[Ful08] Jason Fulman, *Convergence rates of random walk on irreducible representations of finite groups*, J. Theoret. Probab. **21** (2008), no. 1, 193–211, DOI 10.1007/s10959-007-0102-1. MR2384478

[Ful10] Jason Fulman, *Separation cutoffs for random walk on irreducible representations*, Ann. Comb. **14** (2010), no. 3, 319–337, DOI 10.1007/s00026-010-0062-5. MR2737322

[Gow08] W. T. Gowers, *Quasirandom groups*, Combin. Probab. Comput. **17** (2008), no. 3, 363–387, DOI 10.1017/S0963548307008826. MR2410393

[Hel08] H. A. Helfgott, *Growth and generation in $\mathrm{SL}_2(\mathbb{Z}/p\mathbb{Z})$*, Ann. of Math. (2) **167** (2008), no. 2, 601–623, DOI 10.4007/annals.2008.167.601. MR2415382

[HSTZ13] Gerhard Heide, Jan Saxl, Pham Huu Tiep, and Alexandre E. Zalesski, *Conjugacy action, induced representations and the Steinberg square for simple groups of Lie type*, Proc. Lond. Math. Soc. (3) **106** (2013), no. 4, 908–930, DOI 10.1112/plms/pds062. MR3056296

[Ike15] Christian Ikenmeyer, *The Saxl conjecture and the dominance order*, Discrete Math. **338** (2015), no. 11, 1970–1975, DOI 10.1016/j.disc.2015.04.027. MR3357782

[Li18] Xin Li, *Saxl conjecture for triple hooks*, Discrete Math. **344** (2021), no. 6, Paper No. 112340, 19, DOI 10.1016/j.disc.2021.112340. MR4222286

[LPS88] A. Lubotzky, R. Phillips, and P. Sarnak, *Ramanujan graphs*, Combinatorica **8** (1988), no. 3, 261–277, DOI 10.1007/BF02126799. MR963118

[LS17] Sammy Luo and Mark Sellke, *The Saxl conjecture for fourth powers via the semigroup property*, J. Algebraic Combin. **45** (2017), no. 1, 33–80, DOI 10.1007/s10801-016-0700-z. MR3591371

[LST20] Martin W. Liebeck, Aner Shalev, and Pham Huu Tiep, *On the diameters of McKay graphs for finite simple groups*, Israel J. Math. **241** (2021), no. 1, 449–464, DOI 10.1007/s11856-021-2109-1. MR4242157

[LST21] Martin W. Liebeck, Aner Shalev, and Pham Huu Tiep, *McKay graphs for alternating and classical groups*, Trans. Amer. Math. Soc. **374** (2021), no. 8, 5651–5676, DOI 10.1090/tran/8395. MR4293783

[Mar88] G. A. Margulis, *Explicit group-theoretic constructions of combinatorial schemes and their applications in the construction of expanders and concentrators* (Russian), Problemy Peredachi Informatsii **24** (1988), no. 1, 51–60; English transl., Problems Inform. Transmission **24** (1988), no. 1, 39–46. MR939574

[PPV16] Igor Pak, Greta Panova, and Ernesto Vallejo, *Kronecker products, characters, partitions, and the tensor square conjectures*, Adv. Math. **288** (2016), 702–731, DOI 10.1016/j.aim.2015.11.002. MR3436397

[PS16] László Pyber and Endre Szabó, *Growth in finite simple groups of Lie type*, J. Amer. Math. Soc. **29** (2016), no. 1, 95–146, DOI 10.1090/S0894-0347-2014-00821-3. MR3402696

[Sel20] Mark Sellke, *Covering Irrep(S_n) with tensor products and powers*, Preprint, `arXiv:2004.05283`, 2020.

Department of Mathematics, Stanford University, Stanford, California 94305
Email address: `msellke@stanford.edu`