Some combinatorial problems in group theory

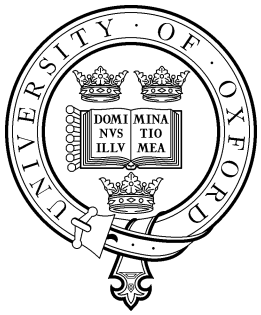

Sean Eberhard
Magdalen College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy

Hilary 2016

Some combinatorial problems in group theory

Sean Eberhard

Magdalen College
University of Oxford

A thesis submitted for the degree of Doctor of Philosophy

Hilary 2016

We study a number of problems of a group-theoretic origin or nature, but from a strongly additive-combinatorial or analytic perspective. Specifically, we consider the following particular problems.

- Given an arbitrary set of n positive integers, how large a subset can you be sure to find which is *sum-free*, i.e., which contains no two elements x and y as well as their sum $x + y$? More generally, given a linear homogeneous equation E, how large a subset can you be sure to find which contains no solutions to E?
- Given a finite group G, suppose we measure the degree of abelianness of G by its *commuting probability* $\Pr(G)$, i.e., the proportion of pairs of elements $x, y \in G$ which commute. What are the possible values of $\Pr(G)$? What is the set of all possible values like as a subset of $[0, 1]$?
- What is the probability that a random permutation $\pi \in \mathcal{S}_n$ has a fixed set of some predetermined size k? Particularly, how does this probability change as k grows? This problem is also related to the following one. Suppose we pick a few permutations $\pi_1, \dots, \pi_r \in \mathcal{S}_n$ at random. It is well known that $\pi_1, \dots, \pi_r$ will generate at least $\mathcal{A}_n$ with high probability as long as $r \geqslant 2$, but what happens if we are allowed to replace $\pi_1, \dots, \pi_r$ by arbitary conjugates $\pi'_1, \dots, \pi'_r$?
- Pick two bijections $\pi_1, \pi_2 : \{1, \dots, n\} \to \mathbf{Z}/n\mathbf{Z}$ uniformly at random. What is the probability that the pointwise sum $\pi_1 + \pi_2$ is also a bijection? This problem affords a fun interpretation in terms of queens on a toroidal chessboard.
- How big is the largest subset of the alternating group $\mathcal{A}_n$ which is *product-free*, i.e., which contains no two elements x and y as well as their product xy?

We give satisfactory answers to each of these questions, using a range of methods. More detailed abstracts are included at the beginning of each chapter.

Acknowledgements

In the first 5/3 years of his studies the author was supported by a Cambridge International Scholarship provided by the Cambridge Commonwealth Trust and the Cambridge Overseas Trust. In the subsequent 2 years he was supported by a Departmental Studentship provided by the Mathematical Institute, Oxford.

Contents

Chapter 1

Erdős's sum-free subset problem

Abstract. A clever probabilistic argument of Erdős shows that every set A of positive integers has a subset $B \subset A$ of size at least $|A|/3$ which is *sum-free*, i.e., which contains no solutions to $x + y = z$. For years it was unknown whether the constant $1/3$ here could be replaced by a larger constant. Our main theorem is that $1/3$ is actually best possible: for every $\varepsilon > 0$ there is a set A having no sum-free subset of size greater than $(1/3 + \varepsilon)|A|$. In fact we prove this for A a multiplicative Følner set.

We also study the variations of this problem arising from replacing the equation $x + y = z$ with some other equation $E_{k,l} : x_1 + \cdots + x_k = y_1 + \cdots + y_l$, where $k > l$. In this case Erdős's argument produces an $E_{k,l}$-free subset of size at least $|A|/(k + l)$. We show that $1/(k + l)$ cannot be improved if either (i) $k \leqslant 3l$, or (ii) $l = 1$. Many interesting questions remain open.

This chapter combines and slightly expands on the two papers [EGM14] and [Ebe15]. The first of these is joint work with Ben Green and Freddie Manners.

1.1 Introduction

A set B is called *sum-free* if there do not exist $x, y, z \in B$ such that $x + y = z$. Suppose we are given an arbitrary set A of positive integers. How large a sum-free subset can we find in A? For example, the set $A = \{1, \ldots, n\}$ has no sum-free subset of size greater than $\lceil n/2 \rceil$ (consider the largest element), so in general we cannot hope to find a sum-free subset B of size greater than $|A|/2$, but a now-classical argument of Erdős shows that one may always find a sum-free subset B of size at least $|A|/3$.

Theorem 1.1.1 (Erdős [Erd65])**.** *Let A be a finite set of positive integers. Then there is a sum-free subset $B \subset A$ of size $|B| \geqslant |A|/3$.*

Proof. Let $\mathbf{T} = \mathbf{R}/\mathbf{Z}$ be the additive group of real numbers modulo 1. For $x \in \mathbf{T}$, let $B_x = \{k \in A : kx \in S\}$ where $S \subset \mathbf{T}$ is the middle third $(1/3, 2/3)$. For every x the set B_x is sum-free (for if $x, y \in B_x$ then $kx, ky \in S$, so $k(x+y) = kx + ky \notin S$, so $x + y \notin B_x$), and on average the size of B_x is

$$\mathbf{E}|B_x| = \sum_{k \in A} \mathbf{P}(kx \in S) = |A|/3.$$

Thus there is at least one x such that $|B_x| \geqslant |A|/3$. □

Despite considerable effort, this silly argument has hardly been improved at all. The best lower bound we know, due to Bourgain [Bou97], is that every set of n positive integers has a sum-free subset of size at least $(n+2)/3$, and even this requires a rather careful Fourier-analytic argument. Further results in Bourgain's paper vaguely suggest a possible route to a lower bound of the form $n/3 + (\log n)^c$, but the remaining challenges are truly formidable.

On the other hand improving the simple upper bound of $n/2$ mentioned earlier is rather easy. For example, consider the set $A = \{1, 2, 3, 4, 5, 6, 8, 9, 10, 18\}$. By simple inspection, A has size 10 and has no sum-free subset of size greater than 4. Thus the set

$$A' = A \cup (100 \cdot A) \cup \cdots \cup (100^{m-1} \cdot A)$$

has size $10m$ and has no sum-free subset of size greater than $4m$. This shows that in general we cannot expect to find a sum-free subset of an arbitrary set A of density greater than $2/5$. Incremental improvements of this form have been made by several authors [Erd65; Erd92; AK90; Mal94; Lew10]. More recently, Alon [Alo13] showed that incremental improvements like this are *always* possible: whenever you have a set A with no sum-free subset of density greater than δ, then there is some $\delta' < \delta$ and another set A' such that A' has no sum-free subset of density greater than some δ'.

Our main theorem is that indeed the constant $1/3$ is best possible: we construct a set A of n integers having no sum-free subset of size greater than $n/3 + o(n)$. Thus Erdős's simple argument is actually surprisingly sharp.

Theorem 1.1.2. *For every $\varepsilon > 0$ there is a set of A of positive integers having no sum-free subset of size greater than $(1/3 + \varepsilon)|A|$.*

In fact there is no need to be so mysterious. We will show that it suffices to take the multiplicative Følner set

$$A = F_m = \{p_1^{e_1} \cdots p_m^{e_m} : 0 \leqslant e_i < m \text{ for each } i\}$$

for sufficiently large m, where $p_1, p_2, \ldots$ is the sequence of primes.

Apart from some averaging trickery to do with the Følner sets (F_m), the main ingredient of our argument is a rough structure theorem for sets A satisfying conditions of the form $|A - A| < 4|A|$, which may be of independent interest. Specifically, if A is a set of integers with $|A - A| \leqslant (4 - \varepsilon)|A|$ then A has density at least $\frac{1}{2} + c\varepsilon$ on some arithmetic progression of length $\gg_\varepsilon |A|$.

Later in the chapter we will ask how well the method generalizes to other equations. At this point it is convenient to introduce some notation. For a homogeneous linear equation E with integer coefficients in any number of variables, let δ_E be the largest constant such that every set of n positive integers has an E-free subset (that is, a subset having no solutions to E) of size at least $\delta_E n$. Thus Theorems 1.1.1 and 1.1.2 together assert that $\delta_{x+y=z} = 1/3$. Consider now more generally the equation

$$E_{k,l} : x_1 + \cdots + x_k = y_1 + \cdots + y_l,$$

where $k > l \geqslant 1$. A straightforward modification of Erdős's argument shows that $\delta_{E_{k,l}} \geqslant 1/(k + l)$, and the question now is whether $\delta_{E_{k,l}} = 1/(k + l)$.

It turns out that the method we use for sum-free sets (i.e., $k = 2, l = 1$) adapts reasonably easily as long as $k \leqslant 3l$. Since this excludes the especially natural case $l = 1$ (the case of so-called *k-sum-free* sets) except when $k \leqslant 3$, we develop an alternative approach based on a theorem of Łuczak and Schoen on the structure of large k-sum-free sets which works in this case. Thus in summary we know $\delta_{E_{k,l}} = 1/(k + l)$ whenever either $k \leqslant 3l$ or $l = 1$.

Theorem 1.1.3. *If $k > l$ and either $k \leqslant 3l$ or $l = 1$ then $\delta_{E_{k,l}} = 1/(k + l)$. In other words, for every $\varepsilon > 0$ there is a set A with no $E_{k,l}$-free subset of size greater than $(1/(k + l) + \varepsilon)|A|$. In fact as before it suffices to take $A = F_m$ for sufficiently large m.*

We conjecture that $\delta_{E_{k,l}} = 1/(k + l)$ for all $k > l$. In fact we conjecture for a general equation E that

$$\delta_E = \sup\{\mu(S) : \text{measurable } E\text{-free } S \subset \mathbf{T}\}.$$

For more idle speculation refer to Section 1.7.

1.2 The role of Følner sets

Følner sets play a universality role in the sum-free subset problem, in the following sense. Fix a Følner sequence (F_m) in $(\mathbf{N}, \cdot)$, that is a sequence of subsets $F_m \subset \mathbf{N}$ such that, for every $a \in \mathbf{N}$,

$$\frac{|(a \cdot F_m) \triangle F_m|}{|F_m|} \longrightarrow 0 \quad \text{as } m \to \infty,$$

where $\triangle$ denotes symmetric difference. For example one could take

$$F_m = \{p_1^{e_1} \cdots p_m^{e_m} : 0 \leqslant p_i < m \text{ for each } i\},$$

where $p_1, p_2, \ldots$ is the sequence of primes in some order. Then if there is any set A which has no sum-free subset of size greater than $(1/3 + \varepsilon)|A|$, then in fact for some m the set F_m has no sum-free subset of size greater than $(1/3 + \varepsilon)|F_m|$: this is part of Lemma 1.2.1 below. Thus in proving Theorem 1.1.2, we could, if we so wished, restrict our search to the sets in some Følner sequence (F_m).

In fact this universality property extends beyond just sets. If there is any probability measure μ on $\mathbf{N}$ such that no sum-free set B has measure $\mu(B) \geqslant 1/3 + \varepsilon$, then again there is some m such that the set F_m has no sum-free subset of size greater than $(1/3 + \varepsilon)|F_m|$ (and so we could have used the measure μ defined by $\mu(B) = |B \cap F_m|/|F_m|$).

The Følner property will be used in the following way. If $x \in F_m$ is uniformly random and $\mathcal{E}(x)$ is some event depending on x, then for every fixed $a \in \mathbf{N}$ we have

$$|\mathbf{P}(\mathcal{E}(ax)) - \mathbf{P}(\mathcal{E}(x))| \leqslant \frac{|(a \cdot F_m) \triangle F_m|}{|F_m|} \longrightarrow 0 \quad \text{as } m \to \infty.$$

Thus the uniform distribution on F_m is approximately dilation-invariant. Universality now follows from imitating Erdős's argument with (F_m) in place of $\mathbf{T}$.

Lemma 1.2.1. *Fix a homogeneous linear equation E (e.g., $x + y = z$ in the case of sum-free setes), and let $\delta > 0$. Fix also some Følner sequence (F_m) in $(\mathbf{N}, \cdot)$. Then the following four statements are equivalent.*

1. *For every probability measure μ on $\mathbf{N}$, there is a E-free set $B \subset \mathbf{N}$ such that $\mu(B) \geqslant \delta$.*
2. *For every finite set $A \subset \mathbf{N}$, there is a E-free subset $B \subset A$ of size at least $\delta|A|$.*
3. *For every m, there is a E-free subset $B \subset F_m$ of size at least $\delta|F_m|$.*

4. For infinitely many m, there is a E-free subset $B \subset F_m$ of size at least $\delta|F_m|$.

Consequently, each statement is true for $\delta = \delta_E$, but no larger constant.

Proof. Clearly $1 \implies 2 \implies 3 \implies 4$, so it suffices to prove $4 \implies 1$. Suppose F_m has an E-free subset S_m of size at least $\delta|F_m|$. For $x \in F_m$, let A_x be the set of all $a \in \mathbf{N}$ such that $ax \in S_m$. Then A_x is E-free, as any solution in A_x would induce a solution in S_m, and if we choose $x \in F_m$ uniformly at random then the expectation of $\mu(A_x)$ is

$$\begin{aligned}
\mathbf{E}(\mu(A_x)) &= \int \mathbf{P}(a \in A_x)\, d\mu(a) \\
&= \int \mathbf{P}(ax \in S_m)\, d\mu(a) \\
&\geqslant \int \left(\mathbf{P}(x \in S_m) - \frac{|(a \cdot F_m) \triangle F_m|}{|F_m|} \right) d\mu(a) \\
&\geqslant \delta - \int \frac{|(a \cdot F_m) \triangle F_m|}{|F_m|}\, d\mu(a).
\end{aligned}$$

We now send $m \to \infty$, and we deduce from the Følner property and the dominated convergence theorem that for every n there is an E-free set $B_n \subset \mathbf{N}$ such that $\mu(B_n) \geqslant \delta - 1/n$. By passing to a subsequence we may assume that 1_{B_n} converges pointwise to 1_B for some set B. This set B must also be E-free, and by another application of the dominated convergence theorem we have $\mu(B) \geqslant \delta$. □

Lemma 1.2.1 is interesting even in the case of translation-invariant equations (provided of course that we appropriately interpret "solution-free" as "nontrivial-solution-free"). For example let E be $x + y = 2z$. Then by Lemma 1.2.1 if infinitely many F_m contain an E-free subset of density at least δ, then every interval $\{1, \ldots, n\}$ also contains an E-free subset of density at least δ, in contradiction to Roth's theorem. In fact the same proof works for longer progressions, replacing Roth's theorem with Szmerédi's theorem, so every set of positive *multiplicative* upper density must contain arbitrarily long arithmetic progressions. This was first proved by Bergelson [Ber05], who used a more complicated method but also proved a lot more.

For us the real utility of Lemma 1.2.1 is not that it gives us the insight to look at Følner sets, but rather that it enables us to consider arbitrary measures μ. Our strategy now is to construct a measure μ on $\mathbf{N}$ such that every solution-free set has small measure. We will achieve this in the case of sum-free sets by first undertanding something about the structure of large sum-free subsets of $\{1, \ldots, N\}$, and by then carrying out a contraction-mapping-type iterative process to construct an appropriate

measure μ. (In [EGM14], the conversion of the measure to a set was achieved instead by an application of the regularity lemma. The method used here is simpler, though undoubtedly more specialized.)

So much suffices for the equation $x + y = z$, as well as some of its closer relatives, but it does not suffice for all equations, notably $x_1 + \cdots + x_k = y$ for $k \geqslant 4$. Later in this chapter we will deal with these equations with a different method, essentially by further elaborating Lemma 1.2.1 in a way which will enable us to consider still more general types of density, such as upper density, multiplicative upper density, etc. These extensions are less elementary, however, so we leave them for later.

1.3 The structure of sets of doubling $4 - \varepsilon$

1.3.1 Discussion

The purpose of this section is to prove the following weak structure theorem for sets $A \subset \mathbf{Z}$ satisfying $|A - A| \leqslant (4 - \varepsilon)|A|$. We will need this theorem in order to establish a corresponding weak structure theorem for large sum-free sets, and in turn to construct a measure μ satisfying $\mu(B) \leqslant 1/3 + \varepsilon$ for all sum-free sets B.

Theorem 1.3.1. *Let $\varepsilon > 0$, and suppose $A \subset \{1, \dots, N\}$ satisfies $|A-A| \leqslant 4|A|-\varepsilon N$. Then there is an arithmetic progression P of length $\gg_\varepsilon N$ on which the density of A is at least $1/2 + \varepsilon/5$.*

Really this a structure theorem for *dense* sets of doubling $4-\varepsilon$, since the hypothesis $|A - A| \leqslant 4|A| - \varepsilon N$ actually implies $|A| \geqslant (\varepsilon/4)N$. We can however deduce, using standard tricks from additive combinatorics, a structure theorem for arbitrary sets of doubling $4 - \varepsilon$, though we will not actually need it.

Corollary 1.3.2. *Let $\varepsilon > 0$, and suppose $A \subset \mathbf{Z}$ satisfies $|A - A| \leqslant (4 - \varepsilon)|A|$. Then there is an arithmetic progression P of length $\gg_\varepsilon |A|$ on which the density of A is at least $1/2 + c\varepsilon$, for some absolute constant c.*

Proof. By [GR06, Theorem 1.4], every set $A \subset \mathbf{Z}$ with $|A - A| \leqslant 4|A|$ is Freiman 18-isomorphic to a subset of $\{1, \dots, N\}$ for some $N \ll |A|$. Let $\pi : A \to A' \subset \{1, \dots, N\}$ be this Freiman isomorphism. Then $|A' - A'| \leqslant 4|A'| - \varepsilon' N$ with $\varepsilon' \gg \varepsilon$. Applying Theorem 1.3.1 to A', there is a progression $P' \subset \{1, \dots, N\}$ of length $|P'| \gg_\varepsilon N$ such that $X = A' \cap P'$ has size at least $(\frac{1}{2} + \frac{1}{5}\varepsilon')|P'|$. Now it follows from [Lev97, Lemma 1] with $k = 2$ and a short computation that $P' \subset 5A' - 4A'$. Now it follows from basic facts about Freiman homomorphisms (see [TV10, Section 5.2] for example)

that $\pi^{-1} : A' \to A$ induces a Freiman 2-homomorphism $\tilde{\pi}^{-1} : P' \to \mathbf{Z}$ coinciding with π^{-1} on $A' \cap P'$. The image $P = \tilde{\pi}^{-1}(P')$ is therefore a progression of length $\gg_\varepsilon |A|$ on which A has density at least $\frac{1}{2} + c\varepsilon$. □

Under stronger conditions such as $|A + A| < 3|A|$ or $|A - A| < 3|A|$, more precise information can be obtained: see [Fre73, Theorem 1.9] or [LS95]. Statements of the same form as Corollary 1.3.2, but with $\frac{1}{2}+c\varepsilon$ replaced by some small quantity $f(\varepsilon) > 0$, follow from versions of Freiman's theorem: see for example [Fre73, Theorem 2.8] and [Bil99, Theorem 1.2]. These statements come with more effective lower bounds on the length of P.

The proof of Theorem 1.3.1 depends on several nontrivial ingredients. First, we need the $s = 1$ case of the arithmetic regularity lemma of Green and Tao [GT10]. This is a rather technical lemma, but the $s = 1$ case is basically elementary and, the author believes, a central tool of modern additive combinatorics, so we provide a self-contained exposition in Appendix 1.A. Second, we use Macbeath's theorem that $\mu(X + Y) \geqslant \min(1, \mu(X) + \mu(Y))$ for compact sets $X, Y \subset \mathbf{T}^d$, or rather a more robust version of this theorem due to Tao [Tao, Section 3.2] which guarantees not just many distinct sums $x + y$ but in fact many sums with reasonably large multiplicity. Finally, and most essentially, we will use the Brunn–Minkowski theorem. In fact it is reasonable to think of Theorem 1.3.1 as the two-dimensional case of Brunn–Minkowski adapted to the integers: the theorem states that the set A exhibits either two-dimensional-type doubling or some strongly one-dimensional substructure.

Let us fix the uniform measure on each of the various spaces $[0, 1]$, $\mathbf{Z}/q\mathbf{Z}$, $\mathbf{T}$, $\{1, \dots, N\}$, as well as products of these, and in each case let us denote the uniform measure simply by μ. In the case of $\{1, \dots, N\}$ we extend the measure to all of $\mathbf{Z}$, but we continue to give each point mass $1/N$. If X is a space endowed with a measure μ (one of the above) then, as usual, we define the convolution of two sufficiently nice functions $f_1, f_2 : X \to \mathbf{C}$ by

$$f_1 * f_2(x) = \int f_1(y) f_2(x - y) d\mu(y).$$

Similarly we define L^p norms with reference to μ.

1.3.2 Proof

Let now A be a set as in Theorem 1.3.1. Let $\mathcal{F} : \mathbf{R}^+ \to \mathbf{R}^+$ be a growth function depending on ε to be chosen later. Let $\tilde{\varepsilon} = \min(\varepsilon, 1/1000)$. Then by Theorem 1.A.7

there is some $M \ll_{\varepsilon,\mathcal{F}} 1$ such that

$$1_A = f_{\mathrm{str}} + f_{\mathrm{sml}} + f_{\mathrm{unf}},$$

where $\|f_{\mathrm{sml}}\|_2 \leqslant \tilde{\varepsilon}^{100}$, $\|f_{\mathrm{unf}}\|_{U^2} \leqslant 1/\mathcal{F}(M)$, and

$$f_{\mathrm{str}} = F(n/N, n \bmod q, \theta n)$$

for some $F : [0,1] \times \mathbf{Z}/q\mathbf{Z} \times \mathbf{T}^d \to [0,1]$ such that $q, d, \|F\|_{\mathrm{Lip}} \leqslant M$ and for some $(\mathcal{F}(M), N)$-irrational $\theta \in \mathbf{T}^d$.

Let $\tilde{M} = \lceil \tilde{\varepsilon}^{-100} M \rceil$ and consider for $i \in \{1, \dots, \tilde{M}\}$ and $a \in \mathbf{Z}/q\mathbf{Z}$ the progressions

$$I_{i,a} = \left\{ n \in \left(\frac{(i-1)N}{\tilde{M}}, \frac{iN}{\tilde{M}} \right] : n \equiv a \pmod{q} \right\}.$$

Define $F_{a,i} : \mathbf{T}^d \to [0,1]$ by $F_{a,i}(x) = F(i/\tilde{M}, a, x)$. Then because F is M-Lipschitz, $F_{a,i}$ is M-Lipschitz and f_{str} differs by at most $\tilde{\varepsilon}^{100}$ from the function f'_{str} which we define by

$$f'_{\mathrm{str}}(n) = \sum_{i=1}^{\tilde{M}} \sum_{a \bmod q} 1_{I_{i,a}}(n) F_{a,i}(\theta n).$$

Absorbing the error of $\tilde{\varepsilon}^{100}$ into f_{sml}, we obtain the decomposition

$$1_A = f'_{\mathrm{str}} + f'_{\mathrm{sml}} + f_{\mathrm{unf}} \tag{1.1}$$

where $\|f'_{\mathrm{sml}}\|_2 \leqslant 2\tilde{\varepsilon}^{100}$. Given an arbitrary growth function $\tilde{\mathcal{F}}$, we may choose $\mathcal{F}$ to grow sufficiently rapidly depending on $\mathcal{F}$ and ε so that $\mathcal{F}(M) \geqslant \tilde{\mathcal{F}}(\tilde{M})$, whence $\|f_{\mathrm{unf}}\|_{U^2} \leqslant 1/\tilde{\mathcal{F}}(\tilde{M})$ and θ is $(\tilde{\mathcal{F}}(\tilde{M}), N)$-irrational. For simplicity let us now rename $\tilde{M}$ as M, $\tilde{\mathcal{F}}$ as $\mathcal{F}$, f'_{str} as f_{str}, and f'_{sml} as f_{sml}.

Write $\alpha(i,a)$ for the density of A on $I_{i,a}$. We will show that $\alpha(i,a) \geqslant 1/2 + \varepsilon/5$ for some (i,a). Note that while $|I_{a,i}|$ need not be exactly N/qM, at worst it differs from N/qM by 2. We may deal with this small wrinkle by assuming that N is sufficiently large depending on ε. This is acceptable, for if $N < N_0(\varepsilon)$ then Theorem 1.3.1 is trivially satisfied by taking P to be a suitable singleton.

We proceed by examining how the behaviour of 1_A is modelled by the more "structured" functions $F_{a,i}(\theta n)$, which in view of the decomposition 1.1 involves estimating the effect of f_{sml} and f_{unf}. The term f_{sml} is the more troublesome of the two. The following simple lemma is useful in this connection.

Lemma 1.3.3. *For all $(i,a) \in \{1, \dots, M\} \times \mathbf{Z}/q\mathbf{Z}$ outside an exceptional subset E of size at most $\tilde{\varepsilon}^4 qM$ we have $\mathbf{E}_{n \in I_{i,a}} |f_{\mathrm{sml}}(n)| \leqslant \tilde{\varepsilon}^{20}$.*

Proof. If this were not the case we would have

$$\mathbf{E}_{n\leqslant N}|f_{\mathrm{sml}}(n)| > \frac{1}{N}\left(\frac{N}{qM}-2\right)qM\tilde{\varepsilon}^{24} \geqslant \tilde{\varepsilon}^{25},$$

whence by Cauchy–Schwarz $\|f_{\mathrm{sml}}\|_2 \geqslant \tilde{\varepsilon}^{25}$, a contradiction. □

Lemma 1.3.4. *Let E be as in the preceding lemma, and assume that $\mathcal{F}$ grows sufficiently rapidly depending on ε. Then for all $(i,a) \in \{1,\dots,M\}\times \mathbf{Z}/q\mathbf{Z}\backslash E$ we have $\int_{\mathbf{T}^d} F_{i,a} \geqslant \alpha(i,a) - \tilde{\varepsilon}^{10}$.*

Proof. By Lemma 1.A.9 the average of f_{unf} over any progression $I_{i,a}$ is less than $\frac{1}{3}\tilde{\varepsilon}^{10}$ provided that $\mathcal{F}$ grows sufficiently rapidly, and by Lemma 1.3.3 for all $(i,a)\notin E$ the average of f_{sml} on $I_{i,a}$ is also at most $\frac{1}{3}\tilde{\varepsilon}^{10}$. Thus if $(i,a)\notin E$ we have

$$\alpha(i,a) = \mathbf{E}_{n\in I_{a,i}}1_A(n) \leqslant \mathbf{E}_{n\in I_{i,a}}F_{i,a}(\theta n) + \frac{2}{3}\tilde{\varepsilon}^{10} \leqslant \int_{\mathbf{T}^d} F_{i,a} + \tilde{\varepsilon}^{10},$$

where the last step last step follows from the $(\mathcal{F}(M),N)$-irrationality of θ and Lemma 1.A.8, again assuming that $\mathcal{F}$ grows sufficiently rapidly. □

We need a slightly technical lemma concerning level sets of Lipschitz functions.

Lemma 1.3.5. *Let $\eta > 0$. If $\mathcal{F}$ grows sufficiently quickly depending on η then the following is true. If $F:\mathbf{T}^d\to[0,1]$ is M-Lipschitz, θ is $(\mathcal{F}(M),N)$-irrational and $I\subset\{1,\dots,N\}$ is any progression of length at least N/M^2, then the proportion of $n\in I$ such that $F(\theta n)>\eta$ is at least $\mu(\{x\in\mathbf{T}^d : F(x)>2\eta\})-\eta$.*

Proof. We want to compute $\mathbf{E}_{n\in I}\chi\circ F(n\theta)$, where χ is the cutoff $1_{x\geqslant\eta}$. Replace χ by a function $\tilde{\chi}$ with $\|\tilde{\chi}\|_{\mathrm{Lip}} \ll 1/\eta$ such that $\tilde{\chi}(x)=0$ for $x<\eta$ and $\tilde{\chi}(x)=1$ for $x\geqslant 2\eta$. Then $\mathbf{E}_{n\in I}\chi\circ F(n\theta)\geqslant \mathbf{E}_{n\in I}\tilde{\chi}\circ F(n\theta)$. However the function $\tilde{\chi}\circ F$ is Lipschitz with $\|\tilde{\chi}\circ F\|_{\mathrm{Lip}}\ll M/\eta$ and so, if $\mathcal{F}$ grows sufficiently rapidly, since θ is so irrational, Lemma 1.A.8 implies that $\mathbf{E}_{n\in I}\tilde{\chi}\circ F(n\theta)\geqslant\int_{\mathbf{T}^d}\tilde{\chi}\circ F-\eta$. On the other hand the integral here is at least the measure of $\{x : F(x)\geqslant 2\eta\}$. □

The following lemma has more meat to it and is a crucial ingredient of our argument. It encodes the fact that if X,Y are compact subsets of a torus then $\mu(X+Y)\geqslant\min(\mu(X)+\mu(Y),1)$, a theorem originally due to Macbeath [Mac53]. In fact we require a more robust version of this result which was recently given the following elegant formulation by Tao [Tao, Section 3.2]: if $X,Y\subset\mathbf{T}$ are compact and $\mu(X),\mu(Y)\geqslant t\geqslant 0$ then

$$\int_{\mathbf{T}^d}\min(1_X * 1_Y, t)\,d\mu \geqslant t\min(\mu(X)+\mu(Y)-t,1). \tag{1.2}$$

Note that we recover Macbeath's theorem in the limit $t\to 0$.

Lemma 1.3.6. *Let $0 < \eta < 1$ and suppose that $F_1, F_2 : \mathbf{T}^d \to [0,1]$ are M-Lipschitz functions such that $\int F_1, \int F_2 \geqslant 2\eta^{1/6}$. Then the measure of the set of x for which $F_1 * F_2(x) \geqslant \eta$ is at least $\min\left(\int F_1 + \int F_2, 1\right) - 4\eta^{1/6}$.*

Proof. Let $S_i = \{x : F_i(x) > \eta^{1/3}\}$ for $i = 1, 2$. Clearly $\mu(S_i) \geqslant \int F_i - \eta^{1/3}$, so in particular $\mu(S_1), \mu(S_2) \geqslant \eta^{1/6}$. By (1.2) we therefore have

$$\int_{\mathbf{T}^d} \frac{\min(1_{S_1} * 1_{S_2}(x), \eta^{1/6})}{\eta^{1/6}} dx \geqslant \min(\mu(S_1) + \mu(S_2) - \eta^{1/6}, 1).$$

Writing X for the set of $x \in \mathbf{T}^d$ such that $1_{S_1} * 1_{S_2}(x) \geqslant \eta^{1/3}$, the left-hand side here is bounded by $\mu(X) + \eta^{1/6}$, so $\mu(X) \geqslant \min(\int F_1 + \int F_2, 1) - 4\eta^{1/6}$. On the other hand, for $x \in X$ we certainly have $F_1 * F_2(x) \geqslant \eta^{2/3} 1_{S_1} * 1_{S_2}(x) \geqslant \eta$. □

Lemma 1.3.7. *Suppose $\mathcal{F}$ grows sufficiently quickly depending on ε. Then as long as $(i,a), (i',a') \notin E$ and $\alpha(i,a)$, $\alpha(i',a') \geqslant 2\tilde{\varepsilon}^2$ then*

$$|(A - A) \cap I_{i-i',a-a'}| \geqslant \frac{N}{qM} \left(\min(\alpha(i,a) + \alpha(i',a'), 1) - 10\tilde{\varepsilon}^2\right).$$

Moreover the same bound holds for $|(A - A) \cap I_{i-i'+1,a-a'}|$.

If f is a function on an abelian group we write f° for the function $f^\circ(x) = f(-x)$.

Proof. Dealing with $I_{i-i',a-a'}$ and $I_{i-i'+1,a-a'}$ are similar, so we focus on the former. By Lemma 1.3.4, then, it suffices to prove

$$|(A - A) \cap I_{i-i',a-a'}| \geqslant \frac{N}{qM} \left(\min\left(\int F_{i,a} + \int F_{i',a'}, 1\right) - 8\tilde{\varepsilon}^2\right)$$

for (i,a) and (i',a') outside E and such that $\int F_{i,a}, \int F_{i',a'} \geqslant \tilde{\varepsilon}^2$.

For all except maybe $2\tilde{\varepsilon}^2 N/qM$ values of $d \in I_{a-a',i-i'}$ (those near the left ends),

$$|I_{i,a} \cap (d + I_{i',a'})| \geqslant \frac{\tilde{\varepsilon}^2 N}{qM},$$

and for any such d we have, if $\mathcal{F}$ grows sufficiently rapidly,

$$\begin{aligned} f_{\mathrm{str}}|_{I_{i,a}} * f^\circ_{\mathrm{str}}|_{I_{i',a'}}(d) &= \sum_{n \in I_{i,a} \cap (d + I_{i',a'})} F_{i,a}(\theta n)\, F_{i',a'}(\theta(n-d)) \\ &\geqslant \frac{1}{N} |I_{i,a} \cap (d + I_{i',a'})| \left(F_{i,a} * F^\circ_{i',a'}(\theta d) - \frac{1}{4}\tilde{\varepsilon}^{12}\right). \end{aligned}$$

Here we used the $(\mathcal{F}(M), N)$-irrationality of θ, Lemma 1.A.8 and the fact that the product of two M-Lipschitz functions, each of which is bounded pointwise by 1, is

$2M$-Lipschitz. The function $F_{a,i} * F^{\circ}_{a',i'}$ is also M-Lipschitz, so again by the $(\mathcal{F}(M), N)$-irrationality of θ and by Lemma 1.3.5 the proportion of $d \in I_{i-i',a-a'}$ such that $F_{i,a} * F^{\circ}_{i',a'}(\theta d) \geqslant \frac{1}{2}\tilde{\varepsilon}^{12}$ is at least $\mu(Y) - \tilde{\varepsilon}^{12}$, where

$$Y = \{y : F_{i,a} * F^{\circ}_{i',a'}(y) \geqslant \tilde{\varepsilon}^{12}\}.$$

But by Lemma 1.3.6 with $\eta = \tilde{\varepsilon}^{12}$, $\mu(Y) \geqslant \min\left(\int F_{a,i} + \int F_{a',i'}, 1\right) - 4\tilde{\varepsilon}^2$. Putting this all together,

$$f_{\mathrm{str}}|_{I_{i,a}} * f^{\circ}_{\mathrm{str}}|_{I_{i',a'}}(d) \geqslant \frac{\tilde{\varepsilon}^{14}}{4qM}$$

for a set of $d \in I_{i-i',a-a'}$ of size at least

$$\frac{N}{qM}\left(\min\left(\int F_{a,i} + \int F_{a',i'},\, 1\right) - 7\tilde{\varepsilon}^2\right).$$

Now by definition of the set E we can absorb the contribution of f_{sml} and conclude that

$$(f_{\mathrm{str}} + f_{\mathrm{sml}})|_{I_{i,a}} * (f_{\mathrm{str}} + f_{\mathrm{sml}})^{\circ}|_{I_{i',a'}}(d) \geqslant \frac{\tilde{\varepsilon}^{14}N}{5qM}$$

for these same values of d. Finally we add in the contribution of f_{unf}. Recalling from (1.1) that $1_A = f_{\mathrm{str}} + f_{\mathrm{sml}} + f_{\mathrm{unf}}$, Lemma 1.A.10 implies that

$$1_A|_{I_{a,i}} * 1_{-A}|_{I_{a',i'}}(d) \geqslant \frac{\tilde{\varepsilon}^{14}}{8qM}$$

for all d in a subset of $I_{i-i',a-a'}$ of size at least

$$\frac{N}{qM}\left(\min\left(\int F_{a,i} + \int F_{a',i'},\, 1\right) - 8\tilde{\varepsilon}^2\right),$$

again provided that $\mathcal{F}$ grows sufficiently rapidly. Since all these d lie in $(A - A) \cap I_{i-i',a-a'}$, the lemma follows. □

To use the bound supplied by the preceding lemma we apply the Brunn–Minkowski theorem, which states that if $X, Y \subset \mathbf{R}^d$ are compact then

$$\mu(X + Y)^{1/d} \geqslant \mu(X)^{1/d} + \mu(Y)^{1/d}.$$

We require the case $d = 2$. For a wider discussion and proof, refer to [TV10, Chapter 3.4].

Lemma 1.3.8. *Given a function* $\alpha : \{1, \dots, M\} \times \mathbf{Z}/q\mathbf{Z} \to [0, 1]$ *and* $(x, y) \in \{-M, \dots, M\} \times \mathbf{Z}/q\mathbf{Z}$*, define*

$$\tilde{\alpha}(x, y) = \max(\alpha(i, a) + \alpha(i', a')),$$

where the maximum is taken over all $(i,a),(i',a') \in \{1,\dots,M\} \times \mathbf{Z}/q\mathbf{Z}$ *such that* $\alpha(i,a)$, $\alpha(i',a') > 0$, *either* $i - i' = x$ *or* $i - i' + 1 = x$, *and* $a - a' = y$. *Then*

$$\sum_{x,y} \tilde{\alpha}(x,y) \geqslant 4 \sum_{i,a} \alpha(i,a).$$

Proof. Consider the compact sets $X, X' \subset \mathbf{R}^2 \times \mathbf{Z}/q\mathbf{Z}$ defined by

$$X = \bigcup_{\substack{(i,a) \in \{1,\dots,M\} \times \mathbf{Z}/q\mathbf{Z} \\ \alpha(i,a) > 0}} [0, \alpha(i,a)] \times [i-1, i] \times \{a\},$$

$$X' = \bigcup_{\substack{(i',a') \in \{1,\dots,M\} \times \mathbf{Z}/q\mathbf{Z} \\ \alpha(i',a') > 0}} [-\alpha(i',a'), 0] \times [i'-1, i'] \times \{a'\},$$

and note that

$$\begin{aligned} X - X' &= \bigcup_{\substack{(i,a),(i',a') \\ \alpha(i,a),\alpha(i',a') > 0}} [0, \alpha(i,a) + \alpha(i',a')] \times [i - i' - 1, i - i' + 1] \times \{a - a'\} \\ &= \bigcup_{(x,y) \in \{-M,\dots,M\} \cap \mathbf{Z}/q\mathbf{Z}} [0, \tilde{\alpha}(x,y)] \times [x-1, x] \times \{y\}. \end{aligned}$$

Thus, if ν is the product of Lebesgue measure λ on $\mathbf{R}^2$ and counting measure on $\mathbf{Z}/q\mathbf{Z}$, we have $\nu(X) = \nu(X') = \sum_{i,a} \alpha(i,a)$ and $\nu(X - X') = \sum_{x,y} \tilde{\alpha}(x,y)$. It therefore suffices to show that

$$\nu(X - X') \geqslant 4\nu(X).$$

The case $q = 1$ here is of course immediate from the Brunn–Minkowski inequality. A simple argument allows us to extend this to general q. Indeed, let X_a, X'_a be the fibres of X, X' respectively above $a \in \mathbf{Z}/q\mathbf{Z}$. Then X_a, X'_a are compact subsets of $\mathbf{R}^2$. Pick a_* such that $\lambda(X_{a_*}) = \lambda(X'_{a_*})$ is largest. If $X_a \neq \varnothing$ then the Brunn–Minkowski inequality implies that

$$\lambda(X_a - X'_{a_*}) \geqslant \big(\lambda(X_a)^{1/2} + \lambda(X'_{a_*})^{1/2}\big)^2 \geqslant 4\lambda(X_a).$$

However the sets $(X_a - X'_{a_*}) \times \{a - a_*\}$ are disjoint as a ranges over $\mathbf{Z}/q\mathbf{Z}$, since each lies in a different fibre over $\mathbf{Z}/q\mathbf{Z}$. Therefore

$$\nu(X - X') \geqslant \sum_{a : X_a \neq \varnothing} \lambda(X_a - X'_{a_*}) \geqslant \sum_{a : X_a \neq \varnothing} 4\lambda(X_a) = 4\lambda(X). \qquad \square$$

In fact we need the following more robust variant of the above, easily deduced from it.

Lemma 1.3.9. *Let $\eta > 0$. Given a function $\alpha : \{1, \dots, M\} \times \mathbf{Z}/q\mathbf{Z} \to [0,1]$ and $(x,y) \in \{-M, \dots, M\} \times \mathbf{Z}/q\mathbf{Z}$, define*

$$\tilde{\alpha}(x,y) = \max(\alpha(i,a) + \alpha(i',a')),$$

where the maximum is taken over all $(i,a),(i',a') \in \{1,\dots,M\} \times \mathbf{Z}/q\mathbf{Z}$ such that $\alpha(i,a)$, $\alpha(i',a') > \eta$, either $i - i' = x$ or $i - i' + 1 = x$, and $a - a' = y$. Then

$$\sum_{x,y} \tilde{\alpha}(x,y) \geqslant 4 \sum_{i,a} \alpha(i,a) - 4\eta q M.$$

Proof. Let

$$\alpha^\dagger(i,a) = \begin{cases} \alpha(i,a) & \text{if } \alpha(i,a) > \eta, \\ 0 & \text{otherwise.} \end{cases}$$

Then if we define, as in Lemma 1.3.8,

$$\tilde{\alpha}^\dagger(x,y) = \max(\alpha^\dagger(i,a) + \alpha^\dagger(i',a')),$$

where the maximum is taken over all $(i,a),(i',a')$ such that $\alpha^\dagger(i,a), \alpha^\dagger(i',a') > 0$, either $i - i' = x$ or $i - i' + 1 = x$, and $a - a' = y$, then $\tilde{\alpha}^\dagger = \tilde{\alpha}$. It thus follows Lemma 1.3.8 that

$$\sum_{x,y} \tilde{\alpha}(x,y) = \sum_{x,y} \tilde{\alpha}^\dagger(x,y) \geqslant 4 \sum_{a,i} \alpha^\dagger(a,i) \geqslant 4 \sum_{a,i} \alpha(a,i) - 4\eta q M. \qquad \square$$

Now we are ready to put everything together and complete the proof of Theorem 1.3.1. Let $\delta = \tilde{\varepsilon}^{20}/10M^2$. Then certainly $\delta \gg_\varepsilon 1$. Recall that $\alpha(i,a)$ is the density of A on $I_{i,a}$. Define

$$\alpha'(i,a) = \begin{cases} \alpha(i,a) & \text{if } (i,a) \notin E, \\ 0 & \text{if } (i,a) \in E. \end{cases}$$

Then Lemma 1.3.7 may be rephrased as follows: if $\alpha'(i,a), \alpha'(i',a') \geqslant 2\tilde{\varepsilon}^2$ then

$$|(A-A) \cap I_{i-i',a-a'}| \geqslant \frac{N}{qM} \left(\min(\alpha'(i,a) + \alpha'(i',a'), 1) - 10\tilde{\varepsilon}^2 \right),$$

and the same bound holds for $|(A-A) \cap I_{i-i'+1,a-a'}|$. It follows that

$$|A - A| \geqslant \frac{N}{qM} \sum_{x,y} \min(\tilde{\alpha}'(x,y), 1) - 20\tilde{\varepsilon}^2 N,$$

where $\tilde{\alpha}'$ is derived from α' as in Lemma 1.3.9 with $\eta = 2\tilde{\varepsilon}^2$. Recalling that $\tilde{\varepsilon} = \min(\varepsilon, 1/1000)$, this implies

$$\begin{aligned} |A - A| &\geqslant \frac{N}{qM} \sum_{x,y} \min(\tilde{\alpha}'(x,y), 1) - \frac{1}{10}\varepsilon N \\ &\geqslant \frac{N}{qM} \sum_{x,y} \min(\tilde{\alpha}'(x,y), 1 + 2\varepsilon/5) - \frac{9}{10}\varepsilon N. \end{aligned}$$

Supposing that $\tilde{\alpha}'(x,y) < 1 + 2\varepsilon/5$ for all (x,y), Lemma 1.3.9 then implies that

$$|A - A| > \frac{4N}{qM} \sum_{i,a} \alpha'(i,a) - \frac{99}{100}\varepsilon N > \frac{4N}{qM} \sum_{i,a} \alpha(i,a) - \frac{999}{1000}\varepsilon N > 4|A| - \varepsilon N.$$

Thus if $|A - A| \leqslant 4|A| - \varepsilon N$, there must be some (x,y) such that $\tilde{\alpha}'(x,y) \geqslant 1 + 2\varepsilon/5$, whence for some (i,a) we must have $\alpha(i,a) \geqslant 1/2 + \varepsilon/5$. This completes the proof of Theorem 1.3.1.

1.3.3 Asymmetric version

Actually, virtually the same proof gives the following asymmetric version, which even more closely resembles the Brunn–Minkowski theorem.

Theorem 1.3.10. *Let $\varepsilon > 0$, and suppose $A, B \subset \{1, \dots, N\}$ satisfy*

$$|A + B|^{1/2} \leqslant |A|^{1/2} + |B|^{1/2} - \varepsilon N^{1/2}.$$

Then there are arithmetic progressions P_1 and P_2 of the same common difference, each of length $\gg_\varepsilon N$, such that

$$|A \cap P_1|/|P_1| + |B \cap P_2|/|P_2| \geqslant 1 + c\varepsilon,$$

for some absolute constant $c > 0$.

Sketch. One wrinkle is that we need a simultaneous regularity decomposition of A and B. This can be achieved for example by applying the regularity lemma to the set $A \cup (B + N)$ as a subset of $\{1, \dots, 2N\}$.

The only other part of the proof above which needs to be seriously adapted is Lemma 1.3.8. We need to prove that if ν is the product of Lebesgue measure λ on $\mathbf{R}^2$ and counting measure on $\mathbf{Z}/q\mathbf{Z}$, then the two-dimensional Brunn–Minkowski inequality holds for ν, i.e., if $X, Y \subset \mathbf{R}^2 \times \mathbf{Z}/q\mathbf{Z}$ are compact then

$$\nu(X + Y)^{1/2} \geqslant \nu(X)^{1/2} + \nu(Y)^{1/2}.$$

In the asymmetric case there does not seem to be such a simple proof available (involving picking a largest fibre), but we can adapt the proof of the Brunn–Minkowski theorem based on the Prékopa–Leindler inequality (see [TV10, Chapter 3.4]).

We may assume that $\nu(X), \nu(Y) > 0$, as otherwise the assertion is obvious. For $x, y \in \mathbf{R}^2$, write $X_x, Y_y \subset \mathbf{Z}/q\mathbf{Z}$ for the fibres of X and Y over x and y respectively. Then whenever X_x and Y_y are both nonempty we have

$$|X_x + Y_y| \geqslant \max(|X_x|, |Y_y|).$$

Thus for every θ in the range $0 < \theta < 1$ we have

$$|X_x + Y_y| \geqslant |X_x|^\theta |Y_y|^{1-\theta}$$

for all $x, y \in \mathbf{R}^2$ (whether or not X_x or Y_y is empty). Hence by the Prékopa–Leindler inequality we have

$$\int_{\mathbf{R}^2} \max_{x+y=z} |X_x + Y_y| \, d\lambda(z) \geqslant \frac{1}{(\theta^\theta (1-\theta)^{1-\theta})^2} \left(\int_{\mathbf{R}^2} |X_x| \, d\lambda(x) \right)^\theta \left(\int_{\mathbf{R}^2} |Y_y| \, d\lambda(y) \right)^{1-\theta}.$$

This implies in particular

$$\nu(X+Y) \geqslant \frac{1}{(\theta^\theta (1-\theta)^{1-\theta})^2} \nu(X)^\theta \nu(Y)^{1-\theta}.$$

Thus by optimizing θ (i.e., taking $\theta = \nu(X)^{1/2}/(\nu(X)^{1/2} + \nu(Y)^{1/2})$ we have

$$\nu(X+Y)^{1/2} \geqslant \nu(X)^{1/2} + \nu(Y)^{1/2},$$

as required. □

1.4 Contraction-mapping-type iteration

1.4.1 The case of $x + y = z$

The weak structure theorem proved in the previous section for sets of doubling $4-\varepsilon$ implies a corresponding weak structure theorem for large sum-free subsets of $\{1, \dots, N\}$. We start with this.

Theorem 1.4.1. *Suppose $A \subset \{1, \dots, N\}$ is a sum-free set of size at least $(1/3+\varepsilon)N$. Then there is a progression $P = \{d, 2d, \dots, Md\} \subset \{1, \dots, N\}$ with $d \ll_\varepsilon 1$ and $M \gg_\varepsilon N$ such that $A \cap P = \varnothing$.*

Thus there is a progression $P = \{d, 2d \dots, Md\} \subset \{1, \dots, N\}$ with $d \ll_\varepsilon 1$ and $M \gg_\varepsilon N$ which is disjoint from every sum-free set $A \subset \{1, \dots, N\}$ of size at least $(1/3 + \varepsilon)N$.

Proof. Since A is sum-free, the sets A, $-A$, and $A-A$ are pairwise disjoint. Since $A \cup (-A) \cup (A-A) \subset \{-N+1, \dots, N-1\}$, we deduce that

$$|A-A| \leqslant 2N - 2|A| \leqslant (4/3)N - 2\varepsilon N \leqslant 4|A| - 6\varepsilon N.$$

Thus by Theorem 1.3.1 there is a progression P' of common difference $d \ll_\varepsilon 1$ and of length $\gg_\varepsilon N$ such that A has density at least $1/2 + c\varepsilon$ on P'. It follows that $A-A$ contains dx for all x such that $0 \leqslant x \leqslant M$, where $M = 2c\varepsilon|P'|$. Since $A \cap (A-A) = \varnothing$ we must have $A \cap \{d, 2d, \dots, Md\} = \varnothing$.

For the second part take the intersection of all the progressions P which might appear in the first part. □

We are now only a short step away from proving Theorem 1.1.2. It remains only to iteratively build a measure μ on $\{1, \dots, N\}$ such that μ is a convex combination of the uniform measure on $\{1, \dots, N\}$ and another measure, supported on P (with P as in Theorem 1.4.1), which again looks a bit like μ.

Proposition 1.4.2. *For every integer $t \geqslant 0$ there is a finitely supported probability measure μ on $\mathbf{N}$ such that for every sum-free set $B \subset \mathbf{N}$ we have*

$$\mu(B) \leqslant 1/3 + (9/10)^t.$$

Proof. We use induction on t. For $t = 0$ we can take μ to be anything, say the point mass δ_1. For $t > 0$, by induction there is a measure μ' such that for every sum-free set $B \subset \mathbf{N}$ we have

$$\mu'(B) \leqslant 1/3 + (9/10)^{t-1}.$$

By taking N sufficiently large and dilating μ' if necessary we may assume that μ' is supported on the progression $P = \{d, 2d \dots, Md\}$ appearing in the second part of Theorem 1.4.1 with $\varepsilon = (1/2)(9/10)^{t-1}$. Let ν be the uniform measure on $\{1, \dots, N\}$, and put

$$\mu = (4/5)\mu' + (1/5)\nu.$$

Now suppose that $B \subset \mathbf{N}$ is sum-free. By our hypothesis about μ' we have

$$\mu'(B) \leqslant 1/3 + (9/10)^{t-1}.$$

Thus if $\nu(B) \leqslant 1/3 + (9/10)^{t-1}/2$ we have

$$\mu(B) \leqslant 1/3 + (4/5)(9/10)^{t-1} + (1/5)(9/10)^{t-1}/2 = 1/3 + (9/10)^t.$$

On the other hand if $\nu(B) \geqslant 1/3 + (9/10)^{t-1}/2$ then by our hypothesis about P we have $B \cap P = \varnothing$, so $\mu'(B) = 0$, so

$$\mu(B) = (1/5)\nu(B) \leqslant 1/5.$$

This proves the proposition. □

Theorem 1.1.2 is immediate from this proposition (taking t sufficiently large) and Lemma 1.2.1.

1.4.2 The case of $E_{k,l}$ for $k \leqslant 3l$

We now briefly sketch how the method used so far generalizes. Let $k > l$ be positive integers, and consider the equation $E_{k,l}$ as in the introduction. By Erdős's argument we know $\delta_{E_{k,l}} \geqslant 1/(k+l)$. We can prove $\delta_{E_{k,l}} = 1/(k+l)$ using a method similar to the above as long as $k \leqslant 3l$.

Theorem 1.4.3. *Let k, l be positive integers satisfying $l < k \leqslant 3l$. Then $\delta_{E_{k,l}} = 1/(k+l)$. In other words, for every $\varepsilon > 0$ there is a set $A \subset \mathbf{N}$ for which every $E_{k,l}$-free subset $B \subset A$ has size at most $(1/(k+l) + \varepsilon)n$.*

Sketch. Let N be a large positive integer divisible by both l and $k - l$. A set $B \subset \{1, \ldots, N\}$ has no solutions to $E_{k,l}$ if and only if

$$(k-l)B \cap (lB - lB) = \varnothing. \tag{1.3}$$

By Theorem 1.3.10 we may assume that both $(k-l)B$ and $lB - lB$ satisfy two-dimensional-like doubling estimates, unless B has density at least

$$\min\left(\frac{1}{k-l}, \frac{1}{2l}\right) + \varepsilon = \frac{1}{2l} + \varepsilon$$

on some progression of length $\gg_\varepsilon |B|$. (The equality above is the only place where we use the hypothesis $3l \geqslant k$.) In fact we may assume that we have such doubling estimates for $(k-l)B'$ and $lB'' - lB''$ for arbitrary subsets $B', B'' \subset B$, unless B is so dense on some such progression. We apply this with

$$B' = B \cap \{1, \ldots, N/(k-l)\}$$

and

$$B'' = B \cap \{1, \ldots, N/l\}.$$

If we have the two-dimensional-like doubling estimates

$$|(k-l)B'| \geqslant ((k-l)^2 - \varepsilon)|B'|$$

and

$$|lB'' - lB''| \geqslant (4l^2 - \varepsilon)|B''|,$$

then from (1.3) we deduce as in the proof of Theorem 1.4.1 that

$$2((k-l)^2 - \varepsilon)|B'| + (4l^2 - \varepsilon)|B''| \leqslant 2N. \tag{1.4}$$

Define the probability measure ν by

$$\nu(X) = \frac{(k-l)^2|X \cap \{1, \dots, N/(k-l)\}| + 2l^2|X \cap \{1, \dots, N/l\}|}{(k+l)N}.$$

Then (1.4) implies

$$\nu(B) \leqslant \frac{1}{k+l} + O(\varepsilon).$$

On the other hand suppose B has density at least $1/(2l) + \varepsilon$ on some long progression. Then $lB - lB$ contains a long progression through 0, and hence from (1.3) it must be that $(k-l)B$ and thus B itself is disjoint from some such progression. We therefore deduce the following weak structure theorem for large $E_{k,l}$-free sets:

> If B is $E_{k,l}$-free then either $\nu(B) \leqslant 1/(k+l) + O(\varepsilon)$, or B is disjoint from a progression through 0 of length $\gg_\varepsilon N$.

We can therefore now repeat the iterative process of Proposition 1.4.2 but with the measure ν in place of the uniform measure. The result is a measure μ such that $\mu(B) \leqslant 1/(k+l) + \varepsilon$ for every $E_{k,l}$-free $B \subset \mathbf{N}$. We then appeal to Lemma 1.2.1 as before. □

This proves the $k \leqslant 3l$ case of Theorem 1.1.3. The need for the restriction $k \leqslant 3l$ can be understood loosely as follows. If we rewrite the $E_{k,l}$-free condition as

$$(k-l)B \cap (lB - lB) = \varnothing,$$

then for our argument to work we need the difference set $lB - lB$ to "fill out" at least as fast as the sumset $(k-l)B$, and this amounts to requiring that $2l \geqslant k - l$.

1.5 An alternative approach based on transference

In this section our target is the equation

$$E_{k,1} : x_1 + \cdots + x_k = y$$

for arbitrary k. We call a set B with no solutions to this equation *k-sum-free*. Our goal is to prove that for every $\varepsilon > 0$ there is a set $A \subset \mathbf{N}$ having no k-sum-free subset B of size greater than $(1/(k+1) + \varepsilon)|A|$.

Theorem 1.5.1. *Let $k \geqslant 2$. Then $\delta_{E_{k,1}} = 1/(k+1)$. In other words, for every $\varepsilon > 0$ there is a set $A \subset \mathbf{N}$ having no k-sum-free-subset of size greater than $(1/(k+1)+\varepsilon)|A|$.*

Let us recall how Erdős's argument works, and in particular where the constant $1/(k+1)$ comes from. First we identify a solution-free subset $S \subset \mathbf{T}$, specifically the open interval of length $1/(k+1)$ centered at the point $1/(2k-2)$. Then we choose $x \in \mathbf{T}$ uniformly at random and put $B_x = \{k \in A : kx \in S\}$. We then observe that, for every $k \in \mathbf{N}$, the probability that $kx \in S$ is exactly $\mu(S)$, and so by linearity of expectation B_x has expected size $\mu(S)|A|$.

The starting point of this section is the observation that the only thing which makes this proof tick is the dilation-invariance of the measure μ. Here we call a measure μ on an abelian group G *dilation-invariant* if the push-forward of μ under the map $\times k : x \mapsto kx$, which we assume is measurable, is again μ for every positive integer k. At this level of abstraction we can prove a converse to Erdős's theorem, and this turns out to be oddly useful. The following lemma extends Lemma 1.2.1.

Lemma 1.5.2. *To the list of equivalent conditions in Lemma 1.2.1 we may add the following condition.*

5. *There is an abelian group G, a σ-algebra Σ of subsets of G, a dilation-invariant probability measure μ on Σ, and an E-free set $S \in \Sigma$ such that $\mu(S) \geqslant \delta$.*

Consequently δ_E is the largest constant for which we have such (G, Σ, μ) and S.

Proof. The proof of the implication $5 \implies 2$ has been suggested already, but let us repeat the argument. For $x \in G$ let $B_x = \{k \in A : kx \in S\}$. Then B_x is E-free, and by dilation-invariance we have

$$\int_G |B_x| \, d\mu(x) = \sum_{k \in A} \int_G 1_{kx \in S} \, d\mu(x) = \sum_{k \in A} (\times k)_* \mu(S) = \mu(S)|A| \geqslant \delta |A|.$$

Thus there is at least one $x \in G$ such that $|B_x| \geqslant \delta |A|$.

We will prove 3 $\implies$ 5 using an ultraproduct argument. The reader needing an introduction to ultrafilters and in particular Loeb measure should refer to Appendix 2.A.

Suppose that every set F_m has an E-free subset B_m of size at least $\delta|F_m|$. Let $p \in \beta\mathbf{N}\backslash\mathbf{N}$ be a nonprincipal ultrafilter, let G be the ultraproduct $\prod_{m\to p}\mathbf{Z}$, and let Σ be the Loeb σ-algebra on G. Defining μ_m on subsets of $\mathbf{Z}$ by

$$\mu_m(A) = |A \cap F_m|/|F_m|,$$

let μ be the Loeb measure induced by the sequence (μ_m). Let S be the internal set $\prod_{m\to p} B_m$.

To verify dilation-invariance, note that $\times k$ is measurable, as it sends internal sets to internal sets, and that $\times k$ approximately preserves μ_m by the Følner property, so it exactly preserves μ. By the basic properties of ultrafilters we know that S is E-free, and by definition of μ we have

$$\mu(S) = \mathrm{st}\left(\lim_{m\to p}\mu_m(B_m)\right) \geqslant \delta. \qquad \square$$

We do not actually need the next lemma, but we include it for independent interest. Recall that the *multiplicative upper density* of a set $B \subset \mathbf{N}$, with respect to our Følner sequence (F_m), is

$$\overline{d}_{(F_m)}(B) = \limsup_{m\to\infty}\frac{|B \cap F_m|}{|F_m|}.$$

Lemma 1.5.3. *To the equivalent conditions of Lemmas 1.2.1 and 1.5.2 we may also add the following condition.*

6. There is an E-free subset $B \subset \mathbf{N}$ of multiplicative upper density at least δ.

Consequently, δ_E is the largest possible multiplicative upper density of an E-free set $B \subset \mathbf{N}$.

Proof. To prove 5 $\implies$ 6 we use Fatou's lemma. For $x \in G$ let $B_x = \{k \in \mathbf{N} : kx \in S\}$. Then by Fatou's lemma we have

$$\begin{aligned}\int_G \overline{d}_{(F_m)}(B_x)\,d\mu(x) &\geqslant \limsup_{m\to\infty}\int_G \frac{|B_x \cap F_m|}{|F_m|}\,d\mu(x)\\ &= \limsup_{m\to\infty}\frac{1}{|F_m|}\sum_{k\in F_m}(\times k)_*\mu(S)\\ &\geqslant \delta,\end{aligned}$$

so as before there is at least one $x \in G$ such that $\overline{d}_{(F_m)}(B_x) \geqslant \delta$.

The implication 6 $\implies$ 4 is obvious, at least if we are happy to weaken δ to $\delta - \varepsilon$ for arbitrarily small $\varepsilon > 0$. We can then easily upgrade $\delta - \varepsilon$ to δ using, say, 4 $\implies$ 3. □

We can finish the proof of Theorem 1.5.1 by appealing to the following theorem of Łuczak and Schoen on the structure of infinite k-sum-free sets of large upper density.

Theorem 1.5.4 (Łuczak–Schoen [ŁS97])**.** *Let $B \subset \mathbf{N}$ be a k-sum-free set of upper density greater than $1/(k+1)$. Then B is contained in a periodic k-sum-free set.*

In order to take advantage of this theorem we need one more gadget. We define the *upper density on multiples* $\widetilde{d}(B)$ of a set $B \subset \mathbf{N}$ by

$$\widetilde{d}(B) = \limsup_{\forall p:\ \nu_p(N)\to\infty} \limsup_{n\to\infty} \frac{|B \cap \{n, 2n, \ldots, Nn\}|}{n},$$

where the outer lim sup is taken over all sequences "tending to infinity in the divisibility order", i.e., such that the p-valuation $\nu_p(N)$ tends to infinity for each prime p. In other words,

$$\widetilde{d}(B) = \lim_{m\to\infty} \sup \left\{ \limsup_{n\to\infty} \frac{|B \cap \{n, 2n, \ldots, Nn\}|}{n} : \nu_p(N) \geqslant m \text{ for all primes } p \leqslant m \right\}.$$

Then the following is a simple corollary of Theorem 1.5.4.

Corollary 1.5.5. *If $B \subset \mathbf{N}$ is k-sum-free then $\widetilde{d}(B) \leqslant 1/(k+1)$.*

Proof. Suppose $B \subset \mathbf{N}$ is k-sum-free and $\widetilde{d}(B) > 0$. Then necessarily B contains a multiple of every natural number. Thus for every $N \in \mathbf{N}$ the set

$$B/N = \{n : Nn \in B\}$$

also contains a multiple of every natural number, and thus is not contained in any periodic k-sum-free set. Thus by Theorem 1.5.4 the set B/N has upper density at most $1/(k+1)$. Since this holds for every N we have $\widetilde{d}(B) \leqslant 1/(k+1)$. □

On the other hand the following is a consequence of Lemma 1.5.2.

Lemma 1.5.6. *Fix some homogeneous linear equation E. Then there is an E-free set $B \subset \mathbf{N}$ such that $\widetilde{d}(B) \geqslant \delta_E$.*

Proof. By Lemma 1.5.2 there is an abelian group G, a σ-algebra Σ of subsets of G, a dilation-invariant probability measure μ on Σ, and an E-free set $S \in \Sigma$ such that $\mu(S) \geqslant \delta_E$. For $x \in G$ let $B_x = \{k \in \mathbf{N} : kx \in S\}$. Then by two applications of Fatou's lemma we have

$$\begin{aligned}\int_G \widetilde{d}(B_x)\, d\mu(x) &= \limsup_{\forall p:\ \nu_p(N)\to\infty} \limsup_{n\to\infty} \int_G \frac{|B_x \cap \{n, 2n, \ldots, Nn\}|}{n} \\ &= \limsup_{\forall p:\ \nu_p(N)\to\infty} \limsup_{n\to\infty} \frac{1}{n}\sum_{k=1}^{n} \int_G 1_{kx\in S}\, d\mu(x) \\ &= \limsup_{\forall p:\ \nu_p(N)\to\infty} \limsup_{n\to\infty} \frac{1}{n}\sum_{k=1}^{n} (\times k)_*\mu(S) \\ &\geqslant \delta_E.\end{aligned}$$

Thus for at least one $x \in G$ we have $\widetilde{d}(B_x) \geqslant \delta_E$. □

Theorem 1.5.1 follows immediately from the previous two results.

Incidentally, it is also possible to give a purely finitary proof of Theorem 1.5.1, by finitizing the proof of Theorem 1.5.4: see [Ebe15] for full details.

1.6 A question of Bergelson, Hindman, and Jordan

The sum-free subset problem has also appeared in the literature in a somewhat different guise, as we now explain. Inspired by a question of Bergelson about density versions of Schur's theorem, Hindman and Jordan [HJ07] asked the following question.

Question 1.6.1. What is the largest constant β such that there is a sequence (S_n) of measurable subsets of $[0, 1]$ such that $\mu(S_n) \geqslant \beta$ for each n and such that

$$S_m \cap S_n \cap S_{m+n} = \varnothing$$

for all m, n?

As observed by Bergelson, the example

$$S_n = \{x : nx \bmod 1 \in (1/3, 2/3)\}$$

shows that $\beta \geqslant 1/3$. Hindman and Jordan showed that $\beta \leqslant 0.36667$, and asked whether $\beta = 1/3$. This follows from Theorem 1.1.2.

Proposition 1.6.2. $\beta = 1/3$.

Proof. Let (S_n) be a sequence of measurable subsets of $[0,1]$ such that $\mu(S_n) \geqslant \beta$ for each n and such that $S_m \cap S_n \cap S_{m+n} = \varnothing$ for all m, n. For $A \subset \mathbf{N}$ and $x \in [0,1]$, it follows from the hypothesis $S_m \cap S_n \cap S_{m+n} = \varnothing$ that the subset $B_x = \{m \in A : x \in S_m\}$ of A is sum-free. Moreover if we pick $x \in [0,1]$ uniformly at random then

$$\mathbf{E}|B_x| = \sum_{m \in A} \mu(S_m) \geqslant \beta|A|,$$

so there is at least one x such that $|B_x| \geqslant \beta|A|$. By Theorem 1.1.2, this implies $\beta \leqslant 1/3$. □

Suppose we now ask the general version of Question 1.6.1. Given an equation E say in k variables, let β_E be the largest constant such that there is a sequence (S_n) of subsets of $[0,1]$ such that $\mu(S_n) \geqslant \beta_E$ for each n and such that $\bigcap_{i=1}^k S_{m_i} = \varnothing$ for all $m_1, \ldots, m_k \in \mathbf{N}$ satisfying E.

Proposition 1.6.3. $\beta_E = \delta_E$.

Proof. The proof that $\beta_E \leqslant \delta_E$ is much like the proof of Proposition 1.6.2, so it suffices to prove $\beta_E \geqslant \delta_E$. By Lemma 1.5.2 there is an abelian group G, a σ-algebra Σ of subsets of G, a dilation-invariant probability measure μ on Σ, and a solution-free set $S \in \Sigma$ such that $\mu(S) \geqslant \delta$. The sets $S_n = (\times n)^{-1}(S)$ are essentially what we want, except that we prefer subsets of $[0,1]$ rather than of an arbitrary probability space. This can be dealt with in a standard fashion. (Let Σ' be the sub-σ-algebra generated by $\{(\times k)^{-1}(S) : k \in \mathbf{N}\}$. Then Σ' is a countably generated σ-algebra, so by [Roy88, Theorem 15.3.4] there is a measure algebra isomorphism between the completion of (G, Σ', μ) and $[0,1]$. Now for the required set S_n take the image in $[0,1]$ of the set $(\times n)^{-1}(S)$, and remove null sets as necessary.) □

1.7 Open questions

In this chapter we proved that for every $\varepsilon > 0$ there is a set of positive integers A having no sum-free subset B of size greater than $(1/3 + \varepsilon)|A|$. This shows that in Erdős's theorem that every set A has a sum-free subset B of size at least $|A|/3$, the constant $1/3$ is best possible. More generally we proved that $\delta_{E_{k,l}} = 1/(k+l)$ whenever $k > l$ and either $k \leqslant 3l$ or $l = 1$.

Naturally we conjecture the restrictions on k and l are unnecessary, and that the same result is true for all $k > l$. In light of Section 1.5, to prove this it would suffice to show that an $E_{k,l}$-free subset of $\mathbf{N}$ of upper density larger than $1/(k+l)$ must

be contained in a periodic $E_{k,l}$-free subset. This problem might be more tractable, and in particular might be amenable to the method used by Łuczak and Schoen for k-sum-free sets.

Suppose however we now try to generalize the problem to an arbitrary linear equation. That is, consider the equation

$$E: a_1x_1 + \cdots + a_rx_r = 0,$$

where $a_1, \ldots, a_r \in \mathbf{Z}$. To avoid trivialities assume that at least one a_i is positive and at least one a_i is negative. Then before even making a conjecture we run up against a problem: what constant should we have in place of $1/(k+l)$? To run Erdős's proof we should look for the largest E-free subset of $\mathbf{T}$, but the problem is that even for such simple equations as $2x + y = z$ this is an interesting and challenging problem itself.

However, we could still conjecture an abstract transference theorem. For an arbitary compact abelian group G, let

$$d_E(G) = \sup\{\mu_G(S) : \text{measurable } E\text{-free } S \subset G\}.$$

Then we conjecture the following.

Conjecture 1.7.1. $\delta_E = d_E(\mathbf{T})$.

Applied to $E = E_{k,l}$, this conjecture would in particular remove our unnecessary restrictions on k, l, since indeed it follows from Macbeath's theorem mentioned in Section 1.3 that $d_{E_{k,l}}(\mathbf{T}) = 1/(k+l)$. Thus this conjecture generalizes Theorems 1.1.2 and 1.1.3, and of course goes much further.

We do not know whether Conjecture 1.7.1 is true, but the following proposition perhaps adds some credibility to it.

Proposition 1.7.2. *Suppose we have a compact abelian group G, a dilation-invariant Borel probability measure μ on G, and an E-free subset $S \subset G$ of measure $\mu(G) \geqslant \delta$. Then $d_E(\mathbf{T}) \geqslant \delta$.*

This proposition should be compared with Lemma 1.5.2. The point is that if we could somehow arrange for the group G appearing in that lemma to be compact, then Conjecture 1.7.1 would follow.

Proof. First we claim that the set of dilation-invariant Borel probability measures on G is the weak* closed convex hull of the set of normalized Haar measures μ_H of connected closed subgroups H. To see this, we view (G, μ) as a measure-preserving system for the action of $(\mathbf{N}, \cdot)$. By the ergodic decomposition (just the weak form based on the Krein–Milman theorem) it suffices to show that if μ is ergodic for this action then in fact $\mu = \mu_H$ for some connected closed subgroup H of G.

Let $\gamma \in \widehat{G}$ be a character (that is, a continuous homomorphism $\gamma : G \to \mathbf{T}$), and consider the push-forward $\gamma_*\mu$ of μ. This measure $\gamma_*\mu$ on $\mathbf{T}$ is again ergodic for the action of $(\mathbf{N}, \cdot)$, and it's easy to see therefore by Fourier analysis on $\mathbf{T}$ (and the fact that the function $f(x) = 1_{x>0}$ on $\mathbf{Z}$ is not the Fourier transform of any signed measure) that either $\gamma_*\mu = \delta$ or $\gamma_*\mu = \mu_{\mathbf{T}}$. Thus

$$\widehat{\mu}(\gamma) = \int_G e(\gamma(x))\, d\mu(x) = \int_{\mathbf{T}} e(t)\, d\gamma_*\mu(t) = \begin{cases} 1 & \text{if } \gamma_*\mu = \delta, \\ 0 & \text{if } \gamma_*\mu = \mu_{\mathbf{T}}. \end{cases}$$

Thus $\widehat{\mu}$ is the indicator of the subgroup $\Gamma \leqslant \widehat{G}$ defined by

$$\Gamma = \{\gamma \in \widehat{G} : \gamma(x) = 0 \ \mu\text{-a.e.}\},$$

so μ is the normalized Haar measure of $H = \Gamma^\perp$. Moreover if $k\gamma \in \Gamma$ then $k\gamma(x) = \gamma(kx) = 0$ for μ-almost-all x, so by dilation-invariance in fact $\gamma(x) = 0$ for μ-almost-all x, so $\gamma \in \Gamma$. Thus $\widehat{G}/\Gamma \cong \widehat{H}$ is torsion-free, so H is connected. This proves our first claim.

It follows immediately by considering the map $\mu \mapsto \mu(S)$ that for some connected closed subgroup H of G we have $\mu_H(S) \geqslant \delta$. Thus the conclusion of the proposition holds at least for some connected compact abelian group H in place of $\mathbf{T}$.

Finally, to get from H to $\mathbf{T}$ we use transference ideas of Candela and Sisask [Sis08; CS11]. Specifically, given $1_S \in L^2(H)$ we find a function $g \in L^2(\mathbf{T})$ such that the large spectra of 1_S and g are similar in structure. (In the technical language of [CS11], we are using the fact that, because $\widehat{H}$ is torsion-free, $\widehat{\mathbf{T}} \cong \mathbf{Z}$ can Freiman (n, k)-model $\widehat{H}$ for all n, k.) We then approximate g by the indicator of a set and use a removal lemma to conclude. Refer to [CS11] for full details about this procedure. □

One could even more generally formulate a version of Conjecture 1.7.1 for systems F of more than one equation, such as the system

$$F : \begin{cases} x + y = z, \\ x - y = w. \end{cases}$$

A set B is called F-free if there are no tuples of points satisfying each of the equations in the system. Almost all of the general results proved in this chapter carry over without change to the case of arbitrary systems. The following conjecture seems reasonable.

Conjecture 1.7.3. $\delta_F = \lim_{p\to\infty} d_F(\mathbf{Z}/p\mathbf{Z})$.

Candela and Sisask [CS14] have shown that the limit appearing in this conjecture exists, and equals $d_F(\mathbf{T})$ whenever F consists of a single equation. In general $\lim_{p\to\infty} d_F(\mathbf{Z}/p\mathbf{Z}) \geqslant d_F(\mathbf{T})$. It is actually not known whether there are systems F for which $\lim_{p\to\infty} d_F(\mathbf{Z}/p\mathbf{Z}) > d_F(\mathbf{T})$, but this is strongly suspected.

1.A Appendix: U^2 regularity

In the course of the proof of Theorem 1.1.2 we need the $s = 1$ (a.k.a. abelian, or U^2) case of the arithmetic regularity lemma of Green and Tao [GT10]. This case is considerably simpler than the general case, as it relies only on the inverse theorem for the U^2 norm, which is elementary both to state and to prove. Reading [GT10] however can be rather hard-going if one only cares about the case $s = 1$ and not the higher-order theory. The purpose of this chapter therefore is to provide a brief, self-contained treatment of this case. This chapter may also serve as an introduction to the general case.

1.A.1 The regularity lemma and its proof

The arithmetic regularity lemma states, roughly speaking, that an arbitrary function $f : [N] \to [0, 1]$ is the sum of a structured part f_{str}, a small part f_{sml}, and a Gowers-uniform part f_{unf}. Moreover we can buy higher-order uniformity of f_{unf} at the cost of more involved structure of f_{str}, but here we will only be able to afford U^2 uniformity.

We start with the inverse theorem for the U^2 norm. We define the $U^2(\mathbf{Z}/M\mathbf{Z})$ norm of a function $f : \mathbf{Z}/M\mathbf{Z} \to \mathbf{C}$ as

$$\|f\|_{U^2(\mathbf{Z}/M\mathbf{Z})} = \Big(\mathbf{E}_{a,h_1,h_2\in\mathbf{Z}/M\mathbf{Z}} f(a)\overline{f(a+h_1)f(a+h_2)}f(a+h_1+h_2)\Big)^{\frac{1}{4}},$$

and then the $U^2([N])$ norm of a function $f : [N] \to \mathbf{C}$ as

$$\|f\|_{U^2([N])} = \frac{\|f\|_{U^2(\mathbf{Z}/M\mathbf{Z})}}{\|1_{[N]}\|_{U^2(\mathbf{Z}/M\mathbf{Z})}},$$

where $M \geqslant 2N$ and we define $f(x) = 0$ if $x \notin \{1, \dots, N\}$: one easily checks that this definition is independent of the choice of M. We will often abbreviate $U^2([N])$ to U^2 when no confusion can arise.

Given $f : \mathbf{Z}/M\mathbf{Z} \to \mathbf{C}$ we define the Fourier transform $\widehat{f}$ of f by

$$\widehat{f}(r) = \mathbf{E}_{x \in \mathbf{Z}/M\mathbf{Z}} f(x) e_M(-rx)$$

for $r \in \mathbf{Z}/M\mathbf{Z}$, where $e_M(x) = e(x/M)$. The Fourier inversion formula then states

$$f(x) = \sum_{r \in \mathbf{Z}/M\mathbf{Z}} \widehat{f}(r) e_M(rx).$$

Using these formulae one easily proves

$$\|f\|_{U^2(\mathbf{Z}/M\mathbf{Z})} = \left(\sum_{r \in \mathbf{Z}/M\mathbf{Z}} |\widehat{f}(r)|^4 \right)^{\frac{1}{4}}.$$

Lemma 1.A.1 (Inverse theorem for the U^2 norm)**.** *If $f : [N] \to [-1, 1]$ is a function such that $\|f\|_{U^2} \geqslant \delta$, then there exists $\theta \in \mathbf{T}$ such that*

$$\left| \mathbf{E}_{n \in [N]} f(n) e(-\theta n) \right| \gg_\delta 1.$$

Proof. The condition $\|f\|_{U^2([N])} \geqslant \delta$ implies that $\|f\|_{U^2(\mathbf{Z}/M\mathbf{Z})} \gg \delta$, where $M = 2N$ and as usual we extend f by zero to the rest of $\mathbf{Z}/M\mathbf{Z}$. We therefore have

$$\sum_{r \in \mathbf{Z}/M\mathbf{Z}} |\widehat{f}(r)|^4 \gg \delta^4.$$

From Parseval's theorem and the hypothesis $|f| \leqslant 1$ it then follows that

$$\delta^4 \ll \sup |\widehat{f}|^2 \left(\sum_{r \in \mathbf{Z}/M\mathbf{Z}} |\widehat{f}(r)|^2 \right) = \sup |\widehat{f}|^2 \left(\mathbf{E}_{x \in \mathbf{Z}/M\mathbf{Z}} |f(x)|^2 \right) \leqslant \sup |\widehat{f}|^2.$$

Thus $|\widehat{f}(r)| \gg \delta^2$ for at least one $r \in \mathbf{Z}/M\mathbf{Z}$, so we may take $\theta = r/M$. □

We need a slightly modified form of the above lemma in order to apply an energy increment argument, but first we need some language. Let us say that $f : [N] \to \mathbf{R}$ has 1-*complexity* at most M if $f(n) = F(\theta n)$ for some $F : \mathbf{T}^d \to \mathbf{R}$ and $\theta \in \mathbf{T}^d$ such that $d, \|F\|_{\mathrm{Lip}} \leqslant M$. Here we take the Euclidean metric

$$d(x, y) = \min_{z \in \mathbf{Z}^d} \|x - y - z\|_2$$

on $\mathbf{T}^d$, and we define the *Lipschitz norm* $\|F\|_{\mathrm{Lip}}$ of $F : \mathbf{T}^d \to \mathbf{R}$ by

$$\|F\|_{\mathrm{Lip}} = \sup_x |F(x)| + \sup_{x \neq y} \frac{|F(x) - |F(y)|}{d(x,y)}.$$

The Fourier inversion formula shows that every $f : [N] \to \mathbf{C}$ has finite 1-complexity, but functions of bounded 1-complexity are special.

Our results from now on will be quantified by an arbitrary *growth function*, by which we mean simply an increasing function $\mathcal{F} : \mathbf{R}^+ \to \mathbf{R}^+$. By $\mathcal{F} \ll_X 1$ we will mean that $\mathcal{F}$ is bounded by a function $\mathbf{R}^+ \to \mathbf{R}^+$ depending only on the parameter X; in other words $\mathcal{F} \ll_X 1$ means $\mathcal{F}(M) \ll_{X,M} 1$.

We say f is 1-*measurable* with growth $\mathcal{F}$ if for every $M > 0$ there is some function $f_{\mathrm{str}} : [N] \to \mathbf{R}$ of 1-complexity at most $\mathcal{F}(M)$ such that

$$\|f - f_{\mathrm{str}}\|_2 \leqslant \frac{1}{M},$$

where the $L^2([N])$ norm of a function $f : [N] \to \mathbf{C}$ is defined by

$$\|f\|_2 = \left(\mathbf{E}_{x \in [N]} |f(x)|^2\right)^{\frac{1}{2}}.$$

A set $E \subset [N]$ is called 1-measurable with growth $\mathcal{F}$ if 1_E is so. Note that if f and g are 1-measurable with growth $\mathcal{F}$ then $f + g$ and fg are 1-measurable with growth $\ll_{\mathcal{F}} 1$, so if E and F are 1-measurable with growth $\mathcal{F}$ then $E \cup F$, $E \cap F$, $E \backslash F$, and so on, are all 1-measurable with growth $\ll_{\mathcal{F}} 1$.

Lemma 1.A.2 (U^2 inverse theorem, alternative formulation)**.** *If $f : [N] \to [-1,1]$ is a function such that $\|f\|_{U^2} \geqslant \delta$, then there is a 1-measurable set $E \subset [N]$ with growth $\ll_\delta 1$ such that*

$$\left|\mathbf{E}_{n \in [N]} f(n) 1_E(n)\right| \gg_\delta 1.$$

Proof. By the previous lemma there is some $\theta \in \mathbf{T}$ such that $\phi(n) = e(-\theta n)$ satisfies

$$\left|\mathbf{E}_{n \in [N]} f(n) \phi(n)\right| \gg_\delta 1.$$

Now by replacing ϕ with its the real or imaginary part, and then with its positive or negative part, we may assume that ϕ is real and nonnegative (e.g., if we take the real and then positive parts, then $\phi(n) = (\Re\, e(-\theta n))^+$).

For $0 \leqslant t \leqslant 1$, let

$$E_t = \{n \in [N] : \phi(n) \geqslant t\}.$$

Noting that

$$\phi(n) = \int_0^1 1_{E_t}(n)\, dt,$$

it follows that

$$\int_0^1 \left|\mathbf{E}_{n\in[N]} f(n) 1_{E_t}(n)\right| \, dt \gg_\delta 1,$$

and so

$$\left|\mathbf{E}_{n\in[N]} f(n) 1_{E_t}(n)\right| \gg_\delta 1$$

for all t in a set $\Omega \subset [0,1]$ of measure $|\Omega| \gg_\delta 1$.

Among these sets E_t with $t \in \Omega$ there must be some E_t which is approximately invariant under small changes in t. Indeed, if

$$M(t) = \sup_{r>0} \frac{1}{2r} \frac{1}{N} |\{n \in [N] : |\phi(n) - t| \leqslant r\}|$$

then the Hardy–Littlewood maximal inequality (see any standard reference, such as [Rud87]) states

$$|\{t \in [0,1] : M(t) \geqslant \lambda\}| \ll \frac{1}{\lambda}.$$

Since $|\Omega| \gg_\delta 1$ there is some $t \in \Omega$ such that $M(t) \ll_\delta 1$.

For any such t, E_t is 1-measurable with growth $\ll_\delta 1$. Indeed, note for any $r > 0$ that

$$|\{n \in [N] : |\phi(n) - t| \leqslant r\}| \ll_\delta rN.$$

Choosing $\eta : \mathbf{R} \to \mathbf{R}^+$ of Lipschitz norm $\|\eta\|_{\mathrm{Lip}} \ll 1/r$ such that $\eta(x) = 0$ if $x < t - r$ and $\eta(x) = 1$ if $x > t + r$, it follows that $\|1_{E_t} - \eta \circ \phi\|_2 \ll_\delta \sqrt{r}$. Since ϕ is a function of θn of Lipschitz norm $\ll 1$, this implies that 1_{E_t} is 1-measurable with growth $\ll_\delta 1$. □

A *factor* $\mathcal{B}$ of $[N]$ is a subalgebra of $2^{[N]}$, or equivalently a partition of $[N]$ into cells. We say a factor $\mathcal{B}'$ *refines* another $\mathcal{B}$ if every cell of $\mathcal{B}$ is a union of cells of $\mathcal{B}'$. We call $\mathcal{B}$ a 1-*factor* with complexity at most M and growth $\mathcal{F}$ if $\mathcal{B}$ has M cells, each of which is 1-measurable with growth $\mathcal{F}$. Note in this case that every $\mathcal{B}$-measurable (i.e., constant on each cell of $\mathcal{B}$) function $f : [N] \to [-1,1]$ is 1-measurable with growth $\ll_{M,\mathcal{F}} 1$.

For $x \in [N]$ we define $\mathcal{B}(x)$ to be the unique cell containing x, and we define the *conditional expectation* $\mathbf{E}(f|\mathcal{B})$ of a function $f : [N] \to \mathbf{C}$ by

$$\mathbf{E}(f|\mathcal{B})(x) = \frac{1}{|\mathcal{B}(x)|} \sum_{y \in \mathcal{B}(x)} f(y).$$

Equivalently, the function $\mathbf{E}(f|\mathcal{B})$ is the orthogonal projection of f onto the subspace of $\mathcal{B}$-measurable functions. Finally, with respect to a fixed function $f : [N] \to \mathbf{C}$, the *energy* of $\mathcal{B}$ is $\mathcal{E}(\mathcal{B}) = \|\mathbf{E}(f|\mathcal{B})\|_2^2$.

Corollary 1.A.3 (Lack of uniformity allows energy increment)**.** *Suppose $\mathcal{B}$ is a 1-factor of complexity $\leqslant M$ and growth $\mathcal{F}$ and $f : [N] \to [-1,1]$ is a function such that $\|f - \mathbf{E}(f|\mathcal{B})\|_{U^2([N])} \geqslant \delta$. Then there exists a refinement $\mathcal{B}'$ of $\mathcal{B}$ of complexity $\leqslant 2M$ and growth $\ll_{M,\delta,\mathcal{F}} 1$ such that*

$$\mathcal{E}(\mathcal{B}') - \mathcal{E}(\mathcal{B}) \gg_\delta 1.$$

Proof. By the previous corollary there is a 1-measurable set $E \subset [N]$ with growth $\ll_\delta 1$ such that

$$|\langle f - \mathbf{E}(f|\mathcal{B}), 1_E\rangle| \gg_\delta 1.$$

Let $\mathcal{B}'$ be the factor generated by $\mathcal{B}$ and E. Then $\mathcal{B}'$ is a 1-factor of complexity $\leqslant 2M$ and growth $\ll_{M,\delta,\mathcal{F}} 1$, and since 1_E is $\mathcal{B}'$-measurable we have

$$|\langle \mathbf{E}(f|\mathcal{B}') - \mathbf{E}(f|\mathcal{B}), 1_E\rangle| \gg_\delta 1.$$

Now Cauchy–Schwarz and the Pythagorean theorem imply that

$$\mathcal{E}(\mathcal{B}') - \mathcal{E}(\mathcal{B}) = \|\mathbf{E}(f|\mathcal{B}') - \mathbf{E}(f|\mathcal{B})\|_2^2 \gg_\delta 1. \qquad \square$$

We can now deduce a weak form of the regularity lemma, occasionally referred to as the Koopman–von Neumann theorem.

Corollary 1.A.4 (Weak regularity)**.** *Let $\mathcal{B}$ be a 1-factor of complexity M and growth $\mathcal{F}$, and let $f : [N] \to [-1,1]$ be a function. Then there exists a refinement $\mathcal{B}'$ of $\mathcal{B}$ of complexity $\ll_{\delta,M} 1$ and growth $\ll_{\delta,M,\mathcal{F}} 1$ such that*

$$\|f - \mathbf{E}(f|\mathcal{B}')\|_{U^2([N])} \leqslant \delta.$$

Proof. Repeatedly apply the previous corollary to refine the 1-factor $\mathcal{B}$. Since $0 \leqslant \mathcal{E}(\mathcal{B}) \leqslant 1$, this process must end after $\ll_\delta 1$ steps. $\square$

Finally, by iterating *this* result, we deduce full regularity.

Theorem 1.A.5 (The U^2 regularity lemma)**.** *Let $f : [N] \to [0,1]$ be a function, $\mathcal{F}$ a growth function, and $\varepsilon > 0$. Then there is a quantity $M \ll_{\varepsilon,\mathcal{F}} 1$ and a decomposition*

$$f = f_{\mathrm{str}} + f_{\mathrm{sml}} + f_{\mathrm{unf}}$$

of f into functions $f_{\mathrm{str}}, f_{\mathrm{sml}}, f_{\mathrm{unf}} : [N] \to [-1,1]$ such that

1. *f_{str} has 1-complexity at most M,*

2. f_{sml} *has* $L^2([N])$ *norm at most* ε,

3. f_{unf} *has* $U^2([N])$ *norm at most* $1/\mathcal{F}(M)$,

4. f_{str} *and* $f_{\mathrm{str}} + f_{\mathrm{sml}}$ *take values in* $[0,1]$.

Proof. Starting with $M_0 = 1$ and $\mathcal{B}_0 = \{\varnothing, [N]\}$, suppose inductively that $\mathcal{B}_i$ is a 1-factor of complexity and growth $\ll_{i,M_i,\mathcal{F}} 1$. Then there is a function $f_{\mathrm{str}}^{(i)} : [N] \to \mathbf{R}$ of 1-complexity $M_{i+1} \ll_{\varepsilon,i,M_i,\mathcal{F}} 1$ such that $M_{i+1} \geqslant M_i$ and

$$\|\mathbf{E}(f|\mathcal{B}_i) - f_{\mathrm{str}}^{(i)}\|_2 \leqslant \varepsilon/2.$$

Moreover, by truncating $f_{\mathrm{str}}^{(i)}$ above and below (which doesn't increase 1-complexity) we may assume that $f_{\mathrm{str}}^{(i)} : [N] \to [0,1]$. By the previous corollary there is a refinement $\mathcal{B}_{i+1}$ of $\mathcal{B}_i$ of complexity and growth $\ll_{i,M_{i+1},\mathcal{F}} 1$ such that

$$\|f - \mathbf{E}(f|\mathcal{B}_{i+1})\|_{U^2([N])} \leqslant 1/\mathcal{F}(M_{i+1}).$$

Note in the end that $M_i \ll_{\varepsilon,i,\mathcal{F}} 1$, and since $(\mathcal{E}(\mathcal{B}_i))$ is an increasing sequence in $[0,1]$ there is some $i \ll_\varepsilon 1$ such that

$$\mathcal{E}(\mathcal{B}_{i+1}) - \mathcal{E}(\mathcal{B}_i) = \|\mathbf{E}(f|\mathcal{B}_{i+1}) - \mathbf{E}(f|\mathcal{B}_i)\|_2^2 \leqslant \varepsilon^2/4.$$

Let $M = M_{i+1}$ and let

$$\begin{aligned} f_{\mathrm{str}} &= f_{\mathrm{str}}^{(i)}, \\ f_{\mathrm{sml}} &= \mathbf{E}(f|\mathcal{B}_{i+1}) - f_{\mathrm{str}}^{(i)}, \\ f_{\mathrm{unf}} &= f - \mathbf{E}(f|\mathcal{B}_{i+1}). \end{aligned}$$

□

It will be convenient to make the structure of f_{str} a little more explicit. Specifically, we would like $\theta \in \mathbf{T}^d$ to be (A,N)-*irrational* for some large A, meaning that if $q \in \mathbf{Z}^d \backslash \{0\}$ and $\|q\|_1 \leqslant A$ (where if $q = (q_1, \dots, q_d)$ then $\|q\|_1 = |q_1| + \cdots + |q_d|$) then $\|q \cdot \theta\|_{\mathbf{T}} \geqslant A/N$: this will be important in several counting lemmas, which relate a sum involving f_{str} to an integral over $\mathbf{T}^d$. Of course, there are other possible behaviours of $q \cdot \theta$: it may be that θ itself is small, in which case $q \cdot \theta$ moves slowly away from 0, or it may be that θ is rational, in which case $q \cdot \theta$ frequently returns to 0, or there may be a combination of these behaviours. Nevertheless, it turns out that once these two pollutants are boiled off, the remnant is highly irrational in the above sense.

We say a subtorus T of $\mathbf{T}^d$ of dimension d' has *complexity* at most M if there is some $L \in \mathrm{SL}_d(\mathbf{Z})$, all of whose coefficients have size at most M, such that $L(T) = \mathbf{T}^{d'} \times \{0\}^{d-d'}$. In this case we implicitly identify T with $\mathbf{T}^{d'}$ using L. For instance, we say $\theta \in \mathbf{T}^d$ is (A,N)-*irrational in* T if $L(\theta)$ is (A,N)-irrational in $\mathbf{T}^{d'}$.

Theorem 1.A.6. *Given $\theta \in \mathbf{T}^d$, a positive integer N, and a growth function $\mathcal{F}$, there is a quantity $M \ll_{d,\mathcal{F}} 1$ and a decomposition*

$$\theta = \theta_{\text{smooth}} + \theta_{\text{rat}} + \theta_{\text{irrat}}$$

such that

1. *θ_{smooth} is (M, N)-smooth, meaning $d(\theta_{\text{smooth}}, 0) \leqslant \frac{M}{N}$,*
2. *θ_{rat} is M-rational, meaning $q\theta_{\text{rat}} = 0$ for some $q \leqslant M$, and*
3. *θ_{irrat} is $(\mathcal{F}(M), N)$-irrational in a subtorus of complexity $\leqslant M$.*

Proof. Starting with $M_0 = 1$, $\theta^{(0)}_{\text{smooth}} = \theta^{(0)}_{\text{rat}} = 0$, $\theta^{(0)}_{\text{irrat}} = \theta$, and $T_0 = \mathbf{T}^d$, suppose inductively that

$$\theta = \theta^{(i)}_{\text{smooth}} + \theta^{(i)}_{\text{rat}} + \theta^{(i)}_{\text{irrat}},$$

where $\theta^{(i)}_{\text{smooth}}$ is (M_i, N)-smooth, $\theta^{(i)}_{\text{rat}}$ is M_i-rational, and $\theta^{(i)}_{\text{irrat}}$ lies in a subtorus T_i of dimension $d - i$ and complexity $\leqslant M_i$.

If $\theta^{(i)}_{\text{irrat}}$ is $(\mathcal{F}(M_i), N)$-irrational in T_i then we are done, so suppose that $L \in \mathrm{SL}_d(\mathbf{Z})$ is a linear map of complexity $\leqslant M_i$ identifying T with $\mathbf{T}^{d-i}$ and such that

$$\|q \cdot L(\theta^{(i)}_{\text{irrat}})\|_{\mathbf{T}} \leqslant \frac{\mathcal{F}(M_i)}{N}$$

for some $q \in \mathbf{Z}^{d-i} \backslash \{0\}$ such that $\|q\|_1 \leqslant \mathcal{F}(M_i)$. Choose $\theta^{(i)\prime}_{\text{smooth}} \in T$ so that

$$\|q \cdot L(\theta^{(i)}_{\text{irrat}} - \theta^{(i)\prime}_{\text{smooth}})\|_{\mathbf{T}} = 0$$

and such that $d(L(\theta^{(i)\prime}_{\text{smooth}}), 0) \leqslant \mathcal{F}(M_i)/N$, so $d(\theta^{(i)\prime}_{\text{smooth}}, 0) \ll_{M_i,d,\mathcal{F}} 1/N$. Let $q = mq'$ where $m \in \mathbf{Z}^+$ and q' is primitive in $\mathbf{Z}^{d-i}$. Then

$$q' \cdot L(\theta^{(i)}_{\text{irrat}} - \theta^{(i)\prime}_{\text{smooth}}) \in \tfrac{1}{m}\mathbf{Z}.$$

Now using the Euclidean algorithm, choose $\theta^{(i)\prime}_{\text{rat}} \in T$ so that

$$q' \cdot L(\theta^{(i)}_{\text{irrat}} - \theta^{(i)\prime}_{\text{smooth}} - \theta^{(i)\prime}_{\text{rat}}) \in \mathbf{Z}$$

and such that $mL(\theta^{(i)\prime}_{\text{rat}}) = 0$, so that $m\theta^{(i)\prime}_{\text{rat}} = 0$. Finally, let

$$\begin{aligned}
\theta^{(i+1)}_{\text{smooth}} &= \theta^{(i)}_{\text{smooth}} + \theta^{(i)\prime}_{\text{smooth}}, \\
\theta^{(i+1)}_{\text{rat}} &= \theta^{(i)}_{\text{rat}} + \theta^{(i)\prime}_{\text{rat}}, \\
\theta^{(i+1)}_{\text{irrat}} &= \theta^{(i)}_{\text{irrat}} - \theta^{(i)\prime}_{\text{smooth}} - \theta^{(i)\prime}_{\text{rat}},
\end{aligned}$$

and choose $M_{i+1} \ll_{M_i,d,\mathcal{F}} 1$ so that $\theta^{(i+1)}_{\text{smooth}}$ is (M_{i+1}, N)-smooth, $\theta^{(i+1)}_{\text{rat}}$ is M_{i+1}-rational, and the subtorus $T_{i+1} = \{x \in T_i : q' \cdot L(x) = 0\}$ has complexity $\leqslant M_{i+1}$.

In the end note that $M_i \ll_{i,d,\mathcal{F}} 1$, and since T_i has dimension $d - i$ we can iterate this argument no more than d times, so for some $i \leqslant d$ we must have that $\theta^{(i)}_{\text{irrat}}$ is $(\mathcal{F}(M_i), N)$-irrational in T_i. □

We can now state and prove the version of the regularity lemma which will be most useful to us. This version improves on Theorem 1.A.5 by giving f_{str} the structure

$$f_{\text{str}}(n) = F(n/N, n \bmod q, \theta n),$$

where

$$F : [0,1] \times \mathbf{Z}/q\mathbf{Z} \times \mathbf{T}^d \to \mathbf{R},$$

$q, d, \|F\|_{\text{Lip}} \leqslant M$, and θ is $(\mathcal{F}(M), N)$-irrational. Here we take the usual Euclidean metrics on $[0,1]$ and $\mathbf{T}^d$, the discrete metric on $\mathbf{Z}/q\mathbf{Z}$, the sum of these metrics on $[0,1] \times \mathbf{Z}/q\mathbf{Z} \times \mathbf{T}^d$, and then define $\|F\|_{\text{Lip}}$ as before.

Theorem 1.A.7 (U^2 regularity, version 2)**.** *Let $f : [N] \to [0,1]$ be a function, $\mathcal{F}$ a growth function, and $\varepsilon > 0$. Then there is a quantity $M \ll_{\varepsilon,\mathcal{F}} 1$ and a decomposition*

$$f = f_{\text{str}} + f_{\text{sml}} + f_{\text{unf}}$$

of f into functions $f_{\text{str}}, f_{\text{sml}}, f_{\text{unf}} : [N] \to [-1,1]$ such that

1. *$f_{\text{str}}(n) = F(n/N, n \bmod q, \theta n)$, where*

$$F : [0,1] \times \mathbf{Z}/q\mathbf{Z} \times \mathbf{T}^d \to [0,1],$$

   $q, d, \|F\|_{\text{Lip}} \leqslant M$, and $\theta \in \mathbf{T}^d$ is $(\mathcal{F}(M), N)$-irrational,

2. *f_{sml} has $L^2([N])$ norm at most ε,*

3. *f_{unf} has $U^2([N])$ norm at most $1/\mathcal{F}(M)$,*

4. *f_{str} and $f_{\text{str}} + f_{\text{sml}}$ take values in $[0,1]$.*

Proof. Let $\mathcal{F}_1$ and $\mathcal{F}_2$ be growth functions depending on ε and $\mathcal{F}$ in a manner to be determined. By Theorem 1.A.5 there exists $M_1 \ll_{\varepsilon,\mathcal{F}_1} 1$ and a decomposition

$$f = f_{\text{str}} + f_{\text{sml}} + f_{\text{unf}}$$

of f into functions $f_{\text{str}}, f_{\text{sml}}, f_{\text{unf}} : [N] \to [-1,1]$ such that

1. $f_{\mathrm{str}}(n) = F(\theta n)$, where $F : \mathbf{T}^d \to [0,1]$, $d, \|F\|_{\mathrm{Lip}} \leqslant M_1$, and $\theta \in \mathbf{T}^d$,
2. f_{sml} has $L^2([N])$ norm at most ε,
3. f_{unf} has $U^2([N])$ norm at most $1/\mathcal{F}_1(M_1)$, and
4. f_{str} and $f_{\mathrm{str}} + f_{\mathrm{sml}}$ take values in $[0,1]$.

Now by the previous theorem we can find $M_2 \ll_{M_1,\mathcal{F}_2} 1$ such that $M_2 \geqslant M_1$ and such that θ decomposes as

$$\theta = \theta_{\mathrm{smooth}} + \theta_{\mathrm{rat}} + \theta_{\mathrm{irrat}},$$

where

1. θ_{smooth} is (M_2, N)-smooth, meaning $d(\theta_{\mathrm{smooth}}, 0) \leqslant \frac{M_2}{N}$,
2. θ_{rat} is M_2-rational, meaning $q\theta_{\mathrm{rat}} = 0$ for some $q \leqslant M_2$, and
3. θ_{irrat} is $(\mathcal{F}_2(M_2), N)$-irrational in a subtorus of complexity $\leqslant M_2$.

Then

$$F(\theta n) = F(\theta_{\mathrm{smooth}} n + \theta_{\mathrm{rat}} n + \theta_{\mathrm{irrat}} n) = \tilde{F}(n/N, n \bmod q, nL(\theta_{\mathrm{irrat}})),$$

where $\tilde{F} : [0,1] \times \mathbf{Z}/q\mathbf{Z} \times \mathbf{T}^{d'} \to [0,1]$ is defined by

$$\tilde{F}(x,y,z) = F(N\theta_{\mathrm{smooth}} x + \theta_{\mathrm{rat}} y + L^{-1}(z)).$$

Noting that $\|\tilde{F}\|_{\mathrm{Lip}} \ll_{M_2} 1$, we can find $M \ll_{M_2} 1$ exceeding both M_2 and $\|\tilde{F}\|_{\mathrm{Lip}}$. But since $M \ll_{M_2} 1$, if $\mathcal{F}_2$ is sufficiently large depending on $\mathcal{F}$ then $\mathcal{F}_2(M_2) \geqslant \mathcal{F}(M)$, and similarly $M_2 \ll_{M_1,\mathcal{F}_2} 1$, so if $\mathcal{F}_1$ is sufficiently large depending on $\mathcal{F}_2$ then $\mathcal{F}_1(M_1) \geqslant \mathcal{F}_2(M_2) \geqslant \mathcal{F}(M)$. After all these dependencies are fixed we have $M \ll_{\varepsilon,\mathcal{F}} 1$, and the conclusion of the theorem holds. □

1.A.2 Counting lemmas

If $\theta \in \mathbf{T}^d$ is highly irrational (i.e., (A,N)-irrational for large A), then the sequence θn is highly equidistibuted over $\mathbf{T}^d$ as n ranges over long progressions. Moreover, if $\mathcal{F}$ is grows sufficiently rapidly and f_{str} is as in the conclusion of Theorem 1.A.7, then the triple $(n/N, n \bmod q, n\theta)$ is highly equidistributed over $[0,1] \times \mathbf{Z}/q\mathbf{Z} \times \mathbf{T}^d$ as n ranges over $\{1, \dots, N\}$. This allows us to relate counts weighted by f_{str} to integrals of F.

Lemma 1.A.8. *Suppose that $\theta \in \mathbf{T}^d$ is (A, N)-irrational, and let $F : \mathbf{T}^d \to \mathbf{C}$ be a function with Lipschitz constant at most M. Suppose that $P \subset \{1, \dots, N\}$ is an arithmetic progression of length at least ηN. Then, provided that $A > A_0(M, d, \eta, \delta)$ is large enough,*

$$\left| \mathbf{E}_{n \in P} F(\theta n) - \int F \, d\mu \right| \leqslant \delta.$$

Proof. The key here (as usual in equidistribution theory) is to take a Fourier expansion of F and truncate it. In particular, we may find $M_0 = O_{M,d,\delta}(1)$ and coefficients c_m with $c_0 = \int F$ and $c_m = O_{M,d}(1)$ for $m \neq 0$ such that

$$\left| F(x) - \sum_{\|m\|_1 \leqslant M_0} c_m e(m \cdot x) \right| \leqslant \delta/2$$

uniformly in x. For a proof, see for example [GT08, Lemma A.9]. It follows, of course, that

$$\left| \mathbf{E}_{n \in P} F(\theta n) - \int F \, d\mu \right| \leqslant \sum_{\|m\|_1 \leqslant M_0, m \neq 0} |c_m| |\mathbf{E}_{n \in P} e(m \cdot \theta n)| + \frac{\delta}{2}.$$

Thus we need only show that

$$\mathbf{E}_{n \in P} e(m \cdot \theta n) = o_{m,\eta;A \to \infty}(1),$$

and then take A sufficiently large. If the common difference of the arithmetic progression P is h, then by summing the geometric progression we have the bound

$$\mathbf{E}_{n \in P} e(m \cdot \theta n) \ll \frac{1}{\eta N \|(m \cdot \theta) h\|_{\mathbf{T}}}.$$

But if $A > |h| \|m\|_1$ then, by the definition of (A, N)-irrationality, $\|(m \cdot \theta) h\|_{\mathbf{T}} \geqslant A/N$. Since $h \leqslant 2\eta^{-1}$, the result follows immediately. □

Now that we know how to relate counts weighted by f_{str} to counts weighted by F, we need to know how to relate counts weighted by f to counts weighted by f_{str}. That is, we need to control the effect of the uniform part f_{unf}, as well as that of the small part f_{sml}. In general controlling the effect of the small part f_{sml} is delicate, and depends on the application. Controlling the effect of the uniform part tends to be more straightforward. The following lemma is typical.

Lemma 1.A.9. *Suppose that $f : \{1, \dots, N\} \to \mathbf{C}$ is a function. Then $|\mathbf{E}_{n \leqslant N} f(n)| \ll \|f\|_{U^2}$. More generally suppose that $P \subset \{1, \dots, N\}$ is a progression of length at least ηN. Then $|\mathbf{E}_{n \in P} f(n)| \ll \eta^{-1} \|f\|_{U^2}$.*

Proof. We establish the second statement, the first being a special case of it. Fix a prime $M \in [4N, 8N]$ as in the definition of the U^2 norm and write $G = \mathbf{Z}/M\mathbf{Z}$. Note that the $U^2([N])$-norm and the $U^2(G)$-norm are comparable up to an absolute constant. We use the inequality

$$\begin{aligned}|\mathbf{E}_{x\in G} f(x)g(x)| &= \left|\sum_r \widehat{f}(r)\widehat{g}(r)\right| \\ &\leqslant \left(\sum_r |\widehat{f}(r)|^4\right)^{1/4} \left(\sum_r |\widehat{g}(r)|^{4/3}\right)^{3/4} \\ &= \|f\|_{U^2(G)} \left(\sum_r |\widehat{g}(r)|^{4/3}\right)^{3/4}.\end{aligned}$$

Taking $g = 1_P$, the characteristic function of the progression P, it therefore suffices to show that $\sum_r |\widehat{g}(r)|^{4/3} = O(1)$. Dilating, we may assume that the common difference of P is 1. But then we have, upon summing the geometric progression, the bound $|\widehat{g}(r)| \ll \min(1, |r|^{-1})$, from which the result follows immediately. □

Another important fact is that f_{unf} does not substantially affect the value of convolutions $f * g(n) = \mathbf{E}_{m\leqslant n} f(m)g(n-m)$.

Lemma 1.A.10. *Let $f, \tilde{f}, g : \{1, \dots, N\} \to [-1, 1]$ be functions such that $\|\tilde{f} - f\|_{U^2} \leqslant \delta$. Then $|\tilde{f} * g(d) - f * g(d)| \leqslant \delta$ for all except at most $O(\delta N)$ values of d.*

Proof. Regard the functions $f, \tilde{f}, g$ as functions on $G = \mathbf{Z}/M\mathbf{Z}$ in the usual way, where $M \in [4N, 8N]$ is prime. Let $h = \tilde{f} - f$. Then

$$\begin{aligned}\mathbf{E}_{x\in G}|\mathbf{E}_{y\in G} h(y)g(x-y)|^2 &= \sum_r |\widehat{h}(r)|^2|\widehat{g}(r)|^2 \\ &\leqslant \left(\sum_r |\widehat{h}(r)|^4\right)^{1/2} \left(\sum_r |\widehat{g}(r)|^4\right)^{1/2} \\ &= \|h\|_{U^2(G)}^2 \|g\|_{U^2(G)}^2 \leqslant \|h\|_{U^2(G)}^2 \leqslant 10\delta^2.\end{aligned}$$

By returning from G to $\{1, \dots, N\}$ we have

$$\frac{1}{N}\sum_x |h * g(x)|^2 \ll \delta^2.$$

Thus there cannot be more than $O(\delta N)$ values of x for which $|h * g(x)| \geqslant \delta$. □

References

[AK90] N. Alon and D. J. Kleitman. "Sum-free subsets". *A tribute to Paul Erdős*. Cambridge: Cambridge Univ. Press, 1990, pp. 13–26.

[Alo13] N. Alon. "Paul Erdős and probabilistic reasoning". *Erdős centennial*. Vol. 25. Bolyai Soc. Math. Stud. János Bolyai Math. Soc., Budapest, 2013, pp. 11–33. DOI: 10.1007/978-3-642-39286-3_1.

[Ber05] V. Bergelson. "Multiplicatively large sets and ergodic Ramsey theory". *Israel J. Math.* 148 (2005). Probability in mathematics, pp. 23–40. ISSN: 0021-2172. DOI: 10.1007/BF02775431.

[Bil99] Y. Bilu. "Structure of sets with small sumset". *Astérisque* 258 (1999). Structure theory of set addition, pp. xi, 77–108. ISSN: 0303-1179.

[Bou97] J. Bourgain. "Estimates related to sumfree subsets of sets of integers". *Israel J. Math.* 97 (1997), pp. 71–92. ISSN: 0021-2172. DOI: 10.1007/BF02774027.

[CS11] P. Candela and O. Sisask. "On the asymptotic maximal density of a set avoiding solutions to linear equations modulo a prime". *Acta Math. Hungar.* 132.3 (2011), pp. 223–243. ISSN: 0236-5294. DOI: 10.1007/s10474-011-0124-0.

[CS14] P. Candela and O. Sisask. "Convergence results for systems of linear forms on cyclic groups and periodic nilsequences". *SIAM J. Discrete Math.* 28.2 (2014), pp. 786–810. ISSN: 0895-4801. DOI: 10.1137/130935677.

[Ebe15] S. Eberhard. "Følner sequences and sum-free sets". *Bull. Lond. Math. Soc.* 47.1 (2015), pp. 21–28. ISSN: 0024-6093. DOI: 10.1112/blms/bdu091.

[EGM14] S. Eberhard, B. Green, and F. Manners. "Sets of integers with no large sum-free subset". *Ann. of Math. (2)* 180.2 (2014), pp. 621–652. ISSN: 0003-486X. DOI: 10.4007/annals.2014.180.2.5.

[Erd65] P. Erdős. "Extremal problems in number theory". *Proc. Sympos. Pure Math., Vol. VIII*. Providence, R.I.: Amer. Math. Soc., 1965, pp. 181–189.

[Erd92] P. Erdős. Letter to Klarner. 1992. URL: http://www.plambeck.org/oldhtml/mathematics/klarner/ep/index.htm.

[Fre73] G. A. Freiman. *Foundations of a structural theory of set addition*. Translated from the Russian, Translations of Mathematical Monographs, Vol 37. Providence, R. I.: American Mathematical Society, 1973, pp. vii+108.

[GR06] B. Green and I. Z. Ruzsa. "Sets with small sumset and rectification". *Bull. London Math. Soc.* 38.1 (2006), pp. 43–52. ISSN: 0024-6093. DOI: 10.1112/S0024609305018102.

[GT08] B. Green and T. Tao. "Quadratic uniformity of the Möbius function". *Ann. Inst. Fourier (Grenoble)* 58.6 (2008), pp. 1863–1935. ISSN: 0373-0956. URL: http://aif.cedram.org/item?id=AIF_2008__58_6_1863_0.

[GT10] B. Green and T. Tao. "An arithmetic regularity lemma, an associated counting lemma, and applications". *An irregular mind*. Vol. 21. Bolyai Soc. Math. Stud. Budapest: János Bolyai Math. Soc., 2010, pp. 261–334. DOI: 10.1007/978-3-642-14444-8_7.

[HJ07] N. Hindman and H. Jordan. "Measures of sum-free intersecting families". *New York J. Math.* 13 (2007), pp. 97–106. ISSN: 1076-9803. URL: http://nyjm.albany.edu:8000/j/2007/13_97.html.

[Lev97] V. F. Lev. "Optimal representations by sumsets and subset sums". *J. Number Theory* 62.1 (1997), pp. 127–143. ISSN: 0022-314X. DOI: 10.1006/jnth.1997.2012.

[LS95] V. F. Lev and P. Y. Smeliansky. "On addition of two distinct sets of integers". *Acta Arith.* 70.1 (1995), pp. 85–91. ISSN: 0065-1036.

[Lew10] M. Lewko. "An improved upper bound for the sum-free subset constant". *J. Integer Seq.* 13.8 (2010), Article 10.8.3, 15. ISSN: 1530-7638.

[ŁS97] T. Łuczak and T. Schoen. "On infinite sum-free sets of natural numbers". *J. Number Theory* 66.2 (1997), pp. 211–224. ISSN: 0022-314X. DOI: 10.1006/jnth.1997.2169.

[Mac53] A. M. Macbeath. "On measure of sum sets. II. The sum-theorem for the torus". *Proc. Cambridge Philos. Soc.* 49 (1953), pp. 40–43.

[Mal94] J. L. Malouf. "Combinatorial approaches to integer sequences". PhD thesis. University of Illinois at Urbana-Champaign, 1994.

[Roy88] H. L. Royden. *Real analysis*. Third. Macmillan Publishing Company, New York, 1988, pp. xx+444. ISBN: 0-02-404151-3.

[Rud87] W. Rudin. *Real and complex analysis*. Third. McGraw-Hill Book Co., New York, 1987, pp. xiv+416. ISBN: 0-07-054234-1.

[Sis08] O. Sisask. "Combinatorial properties of large subsets of abelian groups". PhD thesis. University of Bristol, 2008.

[Tao] T. Tao. *Spending Symmetry*. URL: https://terrytao.wordpress.com/books/spending-symmetry/.

[TV10] T. Tao and V. H. Vu. *Additive combinatorics*. Vol. 105. Cambridge Studies in Advanced Mathematics. Paperback edition [of MR2289012]. Cambridge: Cambridge University Press, 2010, pp. xviii+512. ISBN: 978-0-521-13656-3.

Chapter 2

Commuting probability

ABSTRACT. The *commuting probability* (or *commutativity degree*) of a finite group is defined to be the probability that two randomly chosen group elements commute. A little known theorem of P. Neumann from 1989 determines the structure of finite groups whose commuting probability is bounded away from zero. As our main result we use this theorem to investigate the structure the set $\mathcal{P} \subset (0, 1]$ of possible commuting probabilities of finite groups. In particular we prove two conjectures of Joseph from 1977: all limit points of $\mathcal{P}$ are rational, and $\mathcal{P}$ is well ordered by $>$. A third conjecture is left open. Partially answering another question of Joseph, we show that the order type of $\mathcal{P}$ is either ω^ω or ω^{ω^2}.

Additionally we give a simple proof of a theorem of Hofmann and Russo on the structure of compact groups with positive commuting probability. We also briefly touch on 2-*step nilpotency degree*, and we show that the naïve analogue of Neumann's theorem fails.

The material in this chapter relating to Joseph's conjectures appears as the paper [Ebe15].

2.1 Introduction

A reasonable measure of how close a finite group G is to abelian is the proportion of pairs of elements which commute: we call this the *commuting probability* (or *commutativity degree*) of G, and we denote it

$$\Pr(G) = \mathbf{P}_{x,y \in G}(xy = yx) = \frac{1}{|G|^2} |\{(x, y) \in G^2 : xy = yx\}|.$$

The commuting probability has been extensively studied (see for example [ET68; Gus73; Neu89; Rus79]), not least because $|G| \Pr(G)$ is just the number of conjugacy classes of G. To highlight a few particular facts about $\Pr(G)$,

- if $\Pr(G) > 5/8$ then G is abelian (a popular undergraduate exercise);
- if $\Pr(G) \geqslant \delta$ then G has a large abelian section (P. Neumann [Neu89]);
- always $\Pr(G) \geqslant (\log_2 \log_2 |G|)/|G|$ (Erdős and Turán [ET68]).

Of these results we are most interested in Neumann's theorem that a group with large commuting probability must have a large abelian section, where a *section* of a group G is a quotient H/K of a subgroup $H \leqslant G$, and a *large section* is a section H/K where $|G : H|$ and $|K|$ are bounded in terms of δ, so $|H/K| \geqslant \delta'|G|$ for some δ' depending only on δ. Conversely it is easy to see that any group G with such an abelian section must have $\Pr(G) \geqslant \delta''$ for some δ'' depending only on δ', so this theorem provides an if-and-only-if characterization of groups with large commuting probability.

Our main result is that Neumann's theorem can be used to settle two conjectures of Keith Joseph from 1977 about the structure of the set

$$\mathcal{P} = \{\Pr(G) : G \text{ a finite group}\}.$$

We know that $\mathcal{P}$ contains the point 1, and that if $\Pr(G) < 1$ then $\Pr(G) \leqslant 5/8$. Moreover it is not too hard to see there is another gap following 5/8: if $\Pr(G) < 5/8$ then $\Pr(G) \leqslant 17/32$. The first limit point is 1/2, which happens to be the commuting probability of S_3, and then again if $\Pr(G) < 1/2$ then $\Pr(G) \leqslant 1/2 - \varepsilon$ for some $\varepsilon > 0$. For all this see Rusin [Rus79]. Following observations of this sort, Joseph [Jos69; Jos77] made the following three conjectures.

Conjecture 2.1.1 (Joseph's conjectures)**.**

J1. All limit points of $\mathcal{P}$ are rational.

J2. $\mathcal{P}$ is well ordered by $>$.

J3. $\{0\} \cup \mathcal{P}$ is topologically closed.

In other words if $p_n \in \mathcal{P}$ and $p_n \to p$ as $n \to \infty$ then $p \in \mathbf{Q}$, $p_n \geqslant p$ for all but at most finitely many n, and in fact if $p > 0$ then $p \in \mathcal{P}$.

Progress on these conjectures has been slow, however. Prior to the present work, the best partial result, due to Hegarty [Heg13], states that J1 and J2 hold for the set $\mathcal{P} \cap (2/9, 1]$. Our main result is that Neumann's theorem can be used to prove J1 and J2 in their generality.

Theorem 2.1.2. *All limit points of $\mathcal{P}$ are rational, and $\mathcal{P}$ is well ordered by $>$.*

This theorem follows from a somewhat more technical statement to the effect that every $p \in \mathcal{P}$ is approximately an Egyptian fraction of bounded complexity: see Theorem 2.3.2.

A useful model case for Theorem 2.1.2 is the case in which the group G is 2-step nilpotent. In this case the commutator map induces a bilinear map

$$[,] : (G/Z(G)) \times (G/Z(G)) \to [G, G],$$

and the theorem then classifies the possible zero probabilities of this bilinear map. Here the *zero probability* of a bilinear map $\phi : A \times B \to C$ of abelian groups is

$$\Pr(\phi) = \mathbf{P}_{a \in A, b \in B}(\phi(a, b) = 0) = \frac{1}{|A||B|}|\{(a, b) \in A \times B : \phi(a, b) = 0\}|.$$

Let $\mathcal{P}_{\mathrm{b}}$ be the set of all $\Pr(\phi)$, where ϕ is such a bilinear map. The argument is so much easier to motivate and understand in this context that we find it useful to run the two arguments in parallel (there is no formal dependence).

Theorem 2.1.3. *All limit points of $\mathcal{P}_{\mathrm{b}}$ are rational, and $\mathcal{P}_{\mathrm{b}}$ is well ordered by $>$.*

Next we investigate the possible order types of $(\mathcal{P}, >)$, partially answering another question of Joseph. By examining the proof of Theorem 2.1.2 more carefully we reduce the number of possibilities for the order type to two.

Theorem 2.1.4. *The order type of $(\mathcal{P}, >)$ is either ω^ω or ω^{ω^2}.*

The same theorem holds for $\mathcal{P}_{\mathrm{b}}$.

In the final two sections of this chapter we will investigate two related questions. First we will study commuting probabilities of compact groups. A theorem of Hofmann and Russo [HR12], extending an earlier theorem of Lévai and Pyber [LP00] in the profinite case, states that a compact group with positive commuting probability must have finite-index FC-center. This is a heavy restriction which in particular forces

the commuting probability to be rational; in fact it forces the commuting probability to be the same as that of some finite group. We give a simple proof of this theorem in Section 2.6 based on the proof of Neumann's theorem.

Finally in the last section we briefly study 2-*step nilpotency degree*

$$d_2(G) = \mathbf{P}([[x,y],z] = 1) = \frac{1}{|G|^3}|\{(x,y,z) \in G^3 : [[x,y],z] = 1\}|.$$

A naïve analogue of Neumann's theorem might assert that if $d_2(G)$ is bounded away from zero then G has a large 2-step nilpotent section. We present a counterexample to this tempting conjecture.

2.2 Neumann's theorem and variations

First, we need the following simple but quite useful lemma which asserts that whenever X is a symmetric subset of positive density in a group G, then X generates a subgroup in a bounded number of steps. Lemmas of this sort are familiar in additive combinatorics: see for example Hamidoune [Ham92].

Lemma 2.2.1. *Let G be a finite group and X a symmetric subset of G containing the identity. Then $\langle X \rangle = X^{3r}$ provided $(r+1)|X| > |G|$.*

Proof. Suppose $x \in X^{i+1} \backslash X^i$. Then

$$xX \subset X^{i+2} \backslash X^{i-1}.$$

Indeed, $xX \subset X^{i+2}$ is clear, and if $xX \cap X^{i-1} \neq \varnothing$ then because $X = X^{-1}$ we would have $x \in X^{i-1}X^{-1} = X^i$, contrary to hypothesis.

Thus if $x_i \in X^{3i+1} \backslash X^{3i}$ for each $i = 0, \ldots, r$ then the sets $x_0X, \ldots, x_rX$ are disjoint, as $x_iX \subset X^{3i+2} \backslash X^{3i-1}$. Since each of these sets has size $|X|$ we must have

$$(r+1)|X| \leqslant |G|.$$

Thus if $(r+1)|X| > |G|$ we must have $X^{3i+1} = X^{3i}$ for some $i \leqslant r$, so we must have $\langle X \rangle = X^{3i} = X^{3r}$. □

2.2.1 Bilinear maps

For $\phi : A \times B \to C$ a bilinear map and $A' \leqslant A$ and $B' \leqslant B$ subgroups, we denote by $\phi(A', B')$ the group generated by the values $\phi(a', b')$ with $a' \in A'$, $b' \in B'$. The following is our analogue of Neumann's theorem for bilinear maps of large zero probability.

Theorem 2.2.2 (Neumann's theorem for bilinear maps)**.** *Let $\varepsilon > 0$, and let $\phi : A \times B \to C$ be a bilinear map of finite abelian groups such that $\Pr(\phi) \geqslant \varepsilon$. Then there are subgroups $A' \leqslant A$ and $B' \leqslant B$ such that $|A/A'|$, $|B/B'|$ and $|\phi(A', B')|$ are each ε-bounded.*

Proof. Let $X \subset A$ be the set of $x \in A$ such that $|\ker \phi(x, \cdot)| \geqslant (\varepsilon/2)|B|$, and let A' be the group generated by X. Then $|X| \geqslant (\varepsilon/2)|A|$, so A' has index at most $2/\varepsilon$ in A, and by the lemma every $a \in A'$ is a sum of at most $6/\varepsilon$ elements of X, so for every $a \in A'$ we have $|\ker \phi(a, \cdot)| \geqslant (\varepsilon/2)^{6/\varepsilon}|B|$. Similarly, there is a subgroup B' of B of index at most $2/\varepsilon$ such that for every $b \in B'$ we have $|\ker \phi(\cdot, b)| \geqslant (\varepsilon/2)^{6/\varepsilon}|A|$. Then for every $a \in A'$ the subgroup $\ker \phi(a, \cdot) \cap B'$ has index at most $(2/\varepsilon)^{6/\varepsilon}$ in B' and for every $b \in B'$ the subgroup $\ker \phi(\cdot, b) \cap A'$ has index at most $(2/\varepsilon)^{6/\varepsilon}$ in A'.

Now consider any value c of ϕ on $A' \times B'$, say $c = \phi(a, b)$. If we replace a by any element a' of $a + (\ker \phi(\cdot, b) \cap A')$ and then b by any element b' of $b + (\ker \phi(a', \cdot) \cap B')$ then we still have $\phi(a', b') = c$, so

$$|\{(a', b') \in A' \times B' : \phi(a', b') = c\}| \geqslant (\varepsilon/2)^{12/\varepsilon}|A'||B'|,$$

so ϕ takes at most $(2/\varepsilon)^{12/\varepsilon}$ different values on $A' \times B'$. But every element of $\phi(A', B')$ is a sum of distinct values of ϕ on $A' \times B'$, since if say

$$c = \sum_{i=1}^{m} \phi(a_i, b_i)$$

with the term $\phi(a_j, b_j)$ appearing twice then we can reduce the total number of terms by replacing $\phi(a_j, b_j) + \phi(a_j, b_j)$ with $\phi(2a_j, b_j)$. Thus $|\phi(A', B')| \leqslant 2^{(2/\varepsilon)^{12/\varepsilon}}$. □

We need a stronger variant of the above theorem which asserts the existence of subgroups A' and B' such that $\phi(A', B')$ is small and such that $A' \times B'$ contains almost all pairs $(a, b) \in A \times B$ such that $\phi(a, b) \in \phi(A', B')$, in particular almost all pairs such that $\phi(a, b) = 0$. The precise formulation is the following.

Theorem 2.2.3 (Neumann's theorem for bilinear maps, amplified)**.** *For every decreasing function $\eta : \mathbf{N} \to (0, 1)$ there is some $M = M(\eta)$ such that the following holds. For every bilinear map $\phi : A \times B \to C$ there are subgroups $A' \leqslant A$ and $B' \leqslant B$ such that*

1. *$|\phi(A', B')| \leqslant M$,*
2. *with at most $\eta(|\phi(A', B')|)|A||B|$ exceptions, every pair $(a, b) \in A \times B$ such that $\phi(a, b) \in \phi(A', B')$ is contained in $A' \times B'$.*

Remark. We have not stated a bound on $|A/A'|$ or $|B/B'|$, but such a bound is automatic if $\Pr(\phi) \geqslant \varepsilon$, since then

$$\varepsilon \leqslant \Pr(\phi) \leqslant \frac{1}{|A/A'||B/B'|} + \eta(|\phi(A', B')|).$$

Thus by ensuring $\eta(1) \leqslant \varepsilon/2$ one automatically has $|A/A'||B/B'| \leqslant 2\varepsilon^{-1}$.

Proof. If $\Pr(\phi) \leqslant \eta(1)$ then we can just take $A' = B' = \{0\}$, so assume otherwise. Then we can apply Theorem 2.2.2 with $\varepsilon = \eta(1)$. Let $A_1 \leqslant A$ and $B_1 \leqslant B$ be the resulting subgroups, let $C_1 = \phi(A_1, B_1)$, and suppose that more than $\eta(|C_1|)|A||B|$ pairs $(a, b) \in (A \times B) \backslash (A_1 \times B_1)$ satisfy $\phi(a, b) \in C_1$. Then by the pigeonhole principle there is some $(a, b) \in (A \times B) \backslash (A_1 \times B_1)$, say with $a \notin A_1$ (a similar argument applies in the other case), such that at least $\eta(|C_1|)|A_1||B_1|$ pairs $(a', b') \in A_1 \times B_1$ satisfy

$$\phi(a + a', b + b') \in C_1.$$

By replacing (a, b) by $(a + a', b + b')$ for some appropriate $(a', b') \in A_1 \times B_1$ and using the pigeonhole principle again, we may assume that

$$\phi(a, b) \in C_1$$

and that

$$\phi(a, b + b') \in C_1$$

for at least $\eta(|C_1|)|B_1|$ elements $b' \in B_1$. By linearity this implies that

$$\phi(a, b') \in C_1$$

for all these $b' \in B_1$. But this implies that $\phi(a, \cdot) : B_1 \to C/C_1$ has kernel of index at most $\eta(|C_1|)^{-1}$, and hence that

$$|(C_1 + \phi(a, B_1))/C_1| \leqslant \eta(|C_1|)^{-1}.$$

Thus if we put $A_2 = A_1 + \langle a \rangle$, $B_2 = B_1$, $C_2 = \phi(A_2, B_2) = C_1 + \phi(a, B_1)$, then $|C_2| \leqslant \eta(|C_1|)^{-1}|C_1|$, $|B/B_2| \leqslant |B/B_1|$, and $|A/A_2| < |A/A_1|$.

Now repeat the argument with A_2, B_2, C_2 in place of A_1, B_1, C_1, and so on. Since $|A/A_1||B/B_1|$ is an $\eta(1)$-bounded integer and $|A/A_2||B/B_2| < |A/A_1||B/B_1|$, this process must end after an $\eta(1)$-bounded number of steps, at which time we will have the conclusion of the theorem. □

2.2.2 The commutator map

We now turn our attention to the commutator map on groups, which behaves enough like a bilinear map for the above arguments to be emulated. In an arbitrary group G we write $[x, y]$ for the commutator $x^{-1}y^{-1}xy$ of two elements $x, y \in G$. We also write x^y for the conjugate $y^{-1}xy$, and we will use the relation $[x, y] = x^{-1}x^y$. For $H, K \leqslant G$ we write $[H, K]$ for the group generated by all commutators $[h, k]$ with $h \in H, k \in K$. The following is essentially the original version of Neumann's theorem [Neu89], except that we have made the 2-step nilpotency consequence explicit and we have supplied a somewhat simpler proof. First, however, we need a theorem of *B.* Neumann:

Theorem 2.2.4 (B. Neumann [Neu54])**.** *Let G be a group. Then the elements of G have boundedly many conjugates if and only if $[G, G]$ is bounded.*

Proof. The reverse implication is trivial, since for every $x, y \in G$ we have $x^y = x^{-1}[x, y] \in x^{-1}[G, G]$. Now assume that every $x \in G$ has at most n conjugates. Then for every $x \in G$ the subgroup $C_G(x)$ has index at most n. Consider a commutator $c = [x, y]$ of two elements $x, y \in G$. If we replace x by any element x' of $C_G(y)x$ and then y by any element y' of $C_G(x')y$ then we still have $[x', y'] = c$, so

$$|\{(x', y') \in G^2 : [x', y'] = c\}| \geqslant 1/n^2.$$

Thus there are at most n^2 distinct commutators of elements of G. Now a classical theorem of Schur (see [Rob96, Theorem 10.1.4]) implies that $|[G, G]|$ is n-bounded. □

Theorem 2.2.5 (P. Neumann's theorem)**.** *Let $\varepsilon > 0$, and let G be a finite group such that $\Pr(G) \geqslant \varepsilon$. Then G has a normal 2-step nilpotent subgroup H of ε-bounded index such that $|[H, H]|$ is ε-bounded.*

Proof. Let $X \subset G$ be the set of all $x \in G$ such that $|C_G(x)| \geqslant (\varepsilon/2)|G|$, where $C_G(x)$ is the centraliser of x in G, and let K be the group generated by X. Then $|K| \geqslant (\varepsilon/2)|G|$, so K has index at most $2/\varepsilon$ in G, and by Lemma 2.2.1 every $k \in K$ is the product of at most $6/\varepsilon$ elements of K, so for every $k \in K$ we have $|C_G(k)| \geqslant (\varepsilon/2)^{6/\varepsilon}|G|$. Thus also $|C_K(k)| \geqslant (\varepsilon/2)^{6/\varepsilon}|K|$. Thus the elements of K have boundedly many conjugates, so by B. Neumann's theorem $[K, K]$ is ε-bounded.

To finish let $H = C_K([K, K])$. Then H has ε-bounded index in K, hence ε-bounded index in G, and since $[H, H] \subset [K, K]$ we see that H is 2-step nilpotent and $|[H, H]|$ is ε-bounded. □

Now as in the case of bilinear maps we can prove a stronger variant which asserts the existence of a normal subgroup H such that $[H, H]$ is small and such that $H \times H$ contains almost all $(x, y) \in G \times G$ such that $[x, y] \in [H, H]$, in particular almost all commuting pairs. We will need the following generalisation of Schur's theorem due to Baer (see [Rob96, Theorem 14.5.2]).

Lemma 2.2.6 (Baer's theorem)**.** *If M and N are normal subgroups of a group G then $|[M, N]|$ is bounded by a function of $|M/C_M(N)|$ and $|N/C_N(M)|$.*

By analogy with the bilinear case, we will also use the commutator expansion formula

$$[ab, cd] = [a, d]^b[b, d][a, c]^{bd}[b, c]^d, \tag{2.1}$$

which can be verified directly.

Theorem 2.2.7 (Neumann's theorem, amplified)**.** *For every decreasing function $\eta : \mathbf{N} \to (0, 1)$ there is some $M = M(\eta)$ such that the following holds. Every finite group G has a normal subgroup H such that*

1. *$|[H, H]| \leqslant M$,*
2. *with at most $\eta(|[H, H]|)|G|^2$ exceptions, every pair $(x, y) \in G^2$ such that $[x, y] \in [H, H]$ is contained in H^2.*

Proof. If $\Pr(G) \leqslant \eta(1)$ then we can just take $H = 1$, so assume otherwise. Then we can apply Theorem 2.2.5 with $\varepsilon = \eta(1)$. Let $K \leqslant G$ be the resulting subgroup. Now by analogy with the bilinear case we will define an iterative process starting with $K_1 = L_1 = K$. Suppose that more than $\eta(|[K_1, L_1]|)/2 \cdot |G|^2$ pairs $(x, y) \in G^2 \backslash (K_1 \times L_1)$ satisfy $[x, y] \in [K_1, L_1]$. Then by the pigeonhole principle there must be some $(x, y) \in G^2 \backslash (K_1 \times L_1)$, say with $x \notin K_1$ (again, a similar argument applies in the other case), such that at least $\eta(|[K_1, L_1]|)/2 \cdot |K_1||L_1|$ pairs $(k, l) \in K_1 \times L_1$ satisfy

$$[xk, yl] \in [K_1, L_1].$$

By replacing (x, y) with (xk, yl) for appropriate $(k, l) \in K_1 \times L_1$ and applying the pigeonhole principle again, we may assume that

$$[x, y] \in [K_1, L_1]$$

and that

$$[x, yl] \in [K_1, L_1]$$

for at least $\eta(|[K_1, L_1]|)/2 \cdot |L_1|$ elements $l \in L_1$. But by (2.1) we have

$$[x, yl] = [x, l][x, y]^l,$$

so we deduce that

$$[x, l] \in [K_1, L_1]$$

for all these $l \in L_1$.

Thus the subgroup $N_0 \leqslant L_1$ defined by

$$N_0 = \{l \in L_1 : [x, l] \in [K_1, L_1]\}$$

has index at most $2\eta(|[K_1, L_1]|)^{-1}$ in L_1, thus index at most

$$2\eta(|[K_1, L_1]|)^{-1}|G/L_1|$$

in G. If N is the largest normal subgroup of G contained in N_0 then it follows that

$$|G/N| \leqslant (2\eta(|[K_1, L_1]|)^{-1}|G/L_1|)!.$$

But note that if K_2 is the normal subgroup of G generated by K_1 and x then in fact

$$N = \{l \in L_1 : [K_2, l] \subset [K_1, L_1]\},$$

so

$$N/[K_1, L_1] = C_{L_1/[K_1,L_1]}(K_2/[K_1, L_1]).$$

Since trivially

$$K_1/[K_1, L_1] \leqslant C_{K_2/[K_1,L_1]}(L_1/[K_1, L_1]),$$

Lemma 2.2.6 implies that the size of

$$[K_2/[K_1, L_1], L_1/[K_1, L_1]] = [K_2, L_1]/[K_1, L_1]$$

is bounded by a function of $|L_1/N| \leqslant |G/N|$ and $|K_2/K_1| \leqslant |G/K_1|$, and thus the size of $[K_2, L_1]$ is bounded by a function of $\eta(|[K_1, L_1]|)$.

Now repeat the argument with K_2 and $L_2 = L_1$ in place of K_1 and L_1, and so on. Since $|G/K_1||G/L_1|$ is an $\eta(1)$-bounded integer and $|G/K_2||G/L_2| < |G/K_1||G/L_1|$, this process must end after an $\eta(1)$-bounded number of steps. When the process ends we will have normal subgroups $K, L \leqslant G$ such that

1. $|[K, L]| \leqslant M$,

2. with at most $\eta(|[K,L]|)/2 \cdot |G|^2$ exceptions, every pair $(x,y) \in G^2$ such that $[x,y] \in [K,L]$ is contained in $K \times L$.

But 2 implies that with at most $\eta(|[K,L]|)|G|^2$ exceptions every pair $(x,y) \in G^2$ such that $[x,y] \in [K,L]$ is contained in both $K \times L$ and $L \times K$, and hence in $(K \cap L)^2$, so because

$$[K \cap L, K \cap L] \subset [K,L]$$

the conclusion of the theorem is satisfied by $H = K \cap L$. □

2.3 Joseph's conjectures

2.3.1 Approximation by Egyptian fractions

An *Egyptian fraction* is a series of the form $q = 1/n_1 + \cdots + 1/n_m$, where we usually think of m as small. Define the *Egyptian complexity* $\mathcal{E}(q)$ of $q > 0$ to be the smallest $m \in \mathbf{N}$ for which there exist $n_1, \ldots, n_m \in \mathbf{N}$ such that $q = 1/n_1 + \cdots + 1/n_m$. Our main technical result is that zero probabilities and commuting probabilities admit good approximations from below by Egyptian fractions of bounded complexity.

Theorem 2.3.1. *For every decreasing function $\eta : \mathbf{N} \to (0,1)$ there is some $M = M(\eta) \in \mathbf{N}$ such that every bilinear zero probability $\Pr(\phi)$ has the form $q + \varepsilon$, where $\mathcal{E}(q) \leqslant M$ and $0 \leqslant \varepsilon \leqslant \eta(\mathcal{E}(q))$.*

Theorem 2.3.2. *For every decreasing function $\eta : \mathbf{N} \to (0,1)$ there is some $M = M(\eta) \in \mathbf{N}$ such that every commuting probability $\Pr(G)$ has the form $q + \varepsilon$, where $\mathcal{E}(q) \leqslant M$ and $0 \leqslant \varepsilon \leqslant \eta(\mathcal{E}(q))$.*

We start with Theorem 2.3.1. For an abelian group A we denote by $\widehat{A}$ the group of characters $\gamma : A \to S^1$. Recall the size relation $|\widehat{A}| = |A|$ and the orthogonality relations

$$\mathbf{E}_{a \in A} \gamma(a) = 1_{\gamma = 1},$$
$$\mathbf{E}_{\gamma \in \widehat{A}} \gamma(a) = 1_{a=0}.$$

Lemma 2.3.3. *Let A, B, C be finite abelian groups and $\phi : A \times B \to C$ a bilinear map. Then $\mathcal{E}(\Pr(\phi)) \leqslant |C|$.*

Proof. By two applications of orthogonality of characters we have

$$\begin{aligned}\Pr(\phi) &= \mathbf{E}_{a\in A}\mathbf{E}_{b\in B}1_{\phi(a,b)=0}\\ &= \mathbf{E}_{a\in A}\mathbf{E}_{b\in B}\mathbf{E}_{\gamma\in\widehat{C}}\gamma(\phi(a,b))\\ &= \mathbf{E}_{a\in A}\mathbf{E}_{\gamma\in\widehat{C}}1_{\gamma(\phi(a,B))=1}\\ &= \mathbf{E}_{\gamma\in\widehat{C}}\left(\frac{1}{|A|}|\{a\in A:\gamma(\phi(a,B))=1\}|\right).\end{aligned}$$

But for fixed $\gamma\in\widehat{C}$ the set $\{a\in A:\gamma(\phi(a,B))=1\}$ is a subgroup of A, so the above formula expresses $\Pr(\phi)$ as a sum of $|C|$ terms of the form $1/n$ with n a positive integer. □

Proof of Theorem 2.3.1. Fix $\eta:\mathbf{N}\to(0,1)$ and $\phi:A\times B\to C$. Applying Theorem 2.2.3, we find some $M=M(\eta)$ and subgroups $A'\leqslant A$ and $B'\leqslant B$ such that $|\phi(A',B')|\leqslant M$ and such that no more than $\eta(|\phi(A',B')|)|A||B|$ pairs $(a,b)\in(A\times B)\backslash(A'\times B')$ satisfy $\phi(a,b)=0$. Thus

$$\Pr(\phi)=\frac{1}{|A/A'||B/B'|}\Pr(\phi_{A'\times B'})+\varepsilon,$$

where by the lemma

$$\mathcal{E}(\Pr(\phi_{A'\times B'}))\leqslant|\phi(A',B')|\leqslant M,$$

and

$$\begin{aligned}\varepsilon &= \frac{|\{(a,b)\in(A\times B)\backslash(A'\times B'):\phi(a,b)=0\}|}{|A||B|}\\ &\leqslant \eta(|\phi(A',B')|)\\ &\leqslant \eta(\mathcal{E}(\Pr(\phi_{A'\times B'}))).\end{aligned}$$

□

The proof of Theorem 2.3.2 is similar, but in order to get a suitable analogue of Lemma 2.3.3 we need the following theorem of P. Hall [Hal56].

Lemma 2.3.4 (Hall's theorem)**.** *In any group G the index of the second centre*

$$Z_2(G)=\{g\in G:[g,G]\subset Z(G)\}$$

is bounded by a function of $|[G,G]|$.

Lemma 2.3.5. *Let G be a finite group. Then $\mathcal{E}(\Pr(G))\leqslant|G/Z_2(G)|\cdot|[G,G]|$. In particular by Hall's theorem $\mathcal{E}(\Pr(G))$ is bounded by a function of $|[G,G]|$.*

Proof. Let A be the abelian group $[G,G] \cap Z(G)$, and let Z_2 be the second centre of G. Then by the orthogonality relations we have

$$\begin{aligned}\Pr(G) &= \mathbf{E}_{x\in G}\mathbf{E}_{y\in G}1_{[x,y]=1}\\ &= \mathbf{E}_{x\in G}\mathbf{E}_{y\in G}\mathbf{E}_{z\in Z_2}1_{[x,yz]=1}\\ &= \mathbf{E}_{x\in G}\mathbf{E}_{y\in G}\mathbf{E}_{z\in Z_2}\mathbf{E}_{\gamma\in\widehat{A}}1_{[x,yz]\in A}\gamma([x,yz])\\ &= \mathbf{E}_{x\in G}\mathbf{E}_{y\in G}\mathbf{E}_{z\in Z_2}\mathbf{E}_{\gamma\in\widehat{A}}1_{[x,y]\in A}\gamma([x,yz]),\end{aligned}$$

since, by (2.1), $[x,yz] = [x,z][x,y]^z \in A$ if and only if $[x,y] \in A$. Moreover, if $[x,y] \in A$ then $[x,yz] = [x,z][x,y]$, so by orthogonality again we have

$$\begin{aligned}\Pr(G) &= \mathbf{E}_{x\in G}\mathbf{E}_{y\in G}\mathbf{E}_{z\in Z_2}\mathbf{E}_{\gamma\in\widehat{A}}1_{[x,y]\in A}\gamma([x,z])\gamma([x,y])\\ &= \mathbf{E}_{x\in G}\mathbf{E}_{y\in G}\mathbf{E}_{\gamma\in\widehat{A}}1_{[x,y]\in A}\left(\mathbf{E}_{z\in Z_2}\gamma([x,z])\right)\gamma([x,y])\\ &= \mathbf{E}_{x\in G}\mathbf{E}_{y\in G}\mathbf{E}_{\gamma\in\widehat{A}}1_{[x,y]\in A}1_{\gamma([x,Z_2])=1}\gamma([x,y]),\end{aligned}$$

since $z \mapsto [x,z]$ defines a homomorphism $Z_2 \to A$. For fixed $y \in G$, $\gamma \in \widehat{A}$, let

$$G_{y,\gamma} = \{x \in G : [x,y] \in A, \gamma([x,Z_2]) = 1\}.$$

Then, again by (2.1), $G_{y,\gamma}$ is a subgroup of G and $x \mapsto [x,y]$ defines a homomorphism $G_{y,\gamma} \to A$, so

$$\Pr(G) = \mathbf{E}_{y\in G}\mathbf{E}_{\gamma\in\widehat{A}}\frac{1}{|G/G_{y,\gamma}|}1_{\gamma([G_{y,\gamma},y])=1}.$$

Finally, the integrand here depends on y only through yZ_2, so we can replace the expectation over $y \in G$ by an expectation over $yZ_2 \in G/Z_2$, so

$$\mathcal{E}(\Pr(G)) \leqslant |G/Z_2| \cdot |A| \leqslant |G/Z_2| \cdot |[G,G]|. \qquad \square$$

Proof of Theorem 2.3.2. Fix $\eta : \mathbf{N} \to (0,1)$ and G. By the lemma we can find another decreasing function $\eta' : \mathbf{N} \to (0,1)$ such that

$$\eta'(|[G,G]|) \leqslant \eta(\mathcal{E}(\Pr(G)))$$

for all finite groups G. Applying Theorem 2.2.7 with η', we find some $M = M(\eta)$ and a subgroup $H \leqslant G$ such that $|[H,H]| \leqslant M$ and such that no more than $\eta'(|[H,H]|)|G|^2$ pairs $(x,y) \in G^2 \backslash H^2$ satisfy $[x,y] = 1$. Thus

$$\Pr(G) = \frac{1}{|G/H|^2}\Pr(H) + \varepsilon,$$

where by the lemma $\mathcal{E}(\Pr(H))$ is bounded by a function of $|[H,H]| \leqslant M$, and

$$\begin{aligned}\varepsilon &= \frac{|\{(x,y) \in G^2 \backslash H^2 : [x,y] = 1\}|}{|G|^2} \\ &\leqslant \eta'(|[H,H]|) \\ &\leqslant \eta(\mathcal{E}(\Pr(H)))\end{aligned}$$

by the choice of η'. □

2.3.2 Deduction of Joseph's conjectures

The following lemma is well known.

Lemma 2.3.6. *For every $x > 0$ and $m \in \mathbf{N}$ the supremum of the set of $q < x$ such that $\mathcal{E}(q) \leqslant m$ is strictly less than x.*

Proof. Suppose for contradiction that n_{1i}, ..., n_{mi} are m sequences of positive integers such that for all i

$$1/n_{1i} + \cdots + 1/n_{mi} < x$$

and

$$1/n_{1i} + \cdots + 1/n_{mi} \to x.$$

After rearranging and passing to a subsequence we may assume $n_{1i} = n_1, \ldots, n_{ki} = n_k$ are constants while $n_{k+1,i}, \ldots, n_{mi} \to \infty$. But then

$$1/n_{1i} + \cdots + 1/n_{mi} \to 1/n_1 + \cdots + 1/n_k < x,$$

a contradiction. □

Proof of Theorem 2.1.2. Let $x > 0$ be a limit point of $\mathcal{P} = \{\Pr(G) : G \text{ a finite group}\}$. We will prove that x is rational, and that if $p_n \to x$ then $p_n \geqslant x$ for all but finitely many n.

For $m \in \mathbf{N}$ let

$$Q(m,x) = \sup\{q < x : \mathcal{E}(q) \leqslant m\}$$

and define

$$\eta_x(m) = (x - Q(m,x))/2.$$

By the lemma $\eta_x(m) > 0$ for every m, so by Theorem 2.3.2 there is some $M = M(\eta_x)$ such that every $p \in \mathcal{P}$ has the form $q + \varepsilon$, where $\mathcal{E}(q) \leqslant M$ and $0 \leqslant \varepsilon \leqslant \eta_x(\mathcal{E}(q))$.

Fix some such $p = q + \varepsilon$ and suppose $q < x$. Then

$$\varepsilon \leqslant \eta_x(\mathcal{E}(q)) \leqslant (x - q)/2,$$

so

$$p = q + \varepsilon \leqslant (q + x)/2 \leqslant (Q(M, x) + x)/2 = x - \eta_x(M),$$

so p is bounded away from x. Thus if $p_n = q_n + \varepsilon_n \to x$ then we must have $p_n \geqslant q_n \geqslant x$ for all but finitely many n. In particular $q_n \to x$, but the set of Egyptian fractions of complexity at most M is closed, so this implies $\mathcal{E}(x) \leqslant M$, so $x \in \mathbf{Q}$. □

Theorem 2.1.3 is proved in exactly the same way.

2.4 The order type of $\mathcal{P}$

Having shown in previous sections that $\mathcal{P}$ is well ordered by $>$, our goal in this section is to show that $\mathcal{P}$ has order type either ω^ω or ω^{ω^2}. The proof consists of two parts: first, the mere fact that $\mathcal{P}$ is a semigroup forces its order type to have the form ω^{ω^β} for some β; second, we can extract from the proof of Theorem 2.1.2 a bound of ω^{ω^2}.

2.4.1 Consequences of the semigroup property

First we need a standard definition. For X a closed subset of $[0, 1]$ let $X' \subset X$ be the set of limit points of X. Iterating this operation, define X^α for ordinals α as follows:

$$\begin{aligned} X^0 &= X, \\ X^{\alpha+1} &= (X^\alpha)', \\ X^\alpha &= \bigcap_{\beta<\alpha} X^\beta \quad \text{if } \alpha \text{ is a limit ordinal.} \end{aligned}$$

If X is countable then there is a unique countable ordinal α for which X^α is finite and nonempty; we call α the *Cantor–Bendixson rank* of X. If X happens to be well ordered by $>$ then its order type is at most $\omega^\alpha + 1$, and if $X^\alpha = \{0\}$ and $\alpha > 0$ then in fact the order type of X is exactly $\omega^\alpha + 1$. (For a detailed introduction to Cantor–Bendixson rank see Dasgupta [Das14, Chapter 16].)

Lemma 2.4.1. *Let X be a countably infinite closed subset of $[0, 1]$ closed under multiplication, and let α be the Cantor–Bendixson rank of X. Then $X^\alpha = \{0\}$ and $\alpha = \omega^\beta$ for some ordinal β.*

Proof. By induction on γ if $x \in X^\gamma$ and $y \in X$ and $y > 0$ then

$$xy \in X^\gamma.$$

Hence by induction on δ if $x \in X^\gamma$ and $x > 0$ and $y \in X^\delta$ and $y > 0$ then

$$xy \in X^{\gamma+\delta}.$$

Suppose $x \in X^\gamma$ and $x > 0$. Fix $y \in X \cap (0,1)$. Then for all n we have

$$xy^n \in X^\gamma,$$

and $xy^n \to 0$, so

$$0 \in X^{\gamma+1}.$$

Hence we must have $X^\alpha = \{0\}$.

Now suppose $\gamma < \alpha$. Since $0 \in X^\alpha \subset X^{\gamma+1}$ there must be some $x \in X^\gamma \cap (0,1)$. But then for all n we have

$$x^n \in X^{\gamma \cdot n},$$

so

$$0 \in X^{\gamma \cdot \omega}.$$

We deduce that

$$\alpha \geqslant \sup_{\gamma<\alpha}(\gamma \cdot \omega).$$

Let ω^β be the largest power of ω such that $\omega^\beta \leqslant \alpha$. If $\omega^\beta < \alpha$ then

$$\alpha \geqslant \sup_{\gamma<\alpha}(\gamma \cdot \omega) \geqslant \omega^\beta \cdot \omega = \omega^{\beta+1},$$

a contradiction, so we must have $\alpha = \omega^\beta$. □

Let $\overline{\mathcal{P}}$ be the closure of $\mathcal{P}$ in $[0,1]$. By the formula

$$\Pr(G \times H) = \Pr(G)\Pr(H)$$

we know that $\mathcal{P}$, and hence $\overline{\mathcal{P}}$, is closed under multiplication, so if α is the Cantor–Bendixson rank of $\overline{\mathcal{P}}$ then the lemma and the previous discussion imply that $\overline{\mathcal{P}}$ has order type $\omega^\alpha + 1$ and that $\alpha = \omega^\beta$ for some β. It follows that $\mathcal{P}$ has order type $\omega^\alpha = \omega^{\omega^\beta}$. Since for instance $1/2 \in \overline{\mathcal{P}}'$ we know that $\beta > 0$. Next we will prove that $\alpha \leqslant \omega^2$, which then implies $\beta \in \{1,2\}$.

2.4.2 The bound ω^2 for the Cantor–Bendixson rank

For $n \in \mathbf{N}$ let $\mathcal{E}_n = \{q : \mathcal{E}(q) \leqslant n\}$ be the set of Egyptian fractions of complexity at most n. The following lemma follows from the proof of Theorem 2.1.2.

Lemma 2.4.2. *For every $\varepsilon_0 > 0$ there exist $k \in \mathbf{N}$ and a function $m : (0,1] \to \mathbf{N}$ such that for all $\varepsilon_1, \ldots, \varepsilon_k > 0$ the set*

$$\overline{\mathcal{P}} \cap [0, \varepsilon_0]^c \cap \bigcap_{i=0}^{k-1} (\mathcal{E}_{m(\varepsilon_i)} + [0, \varepsilon_{i+1}])^c$$

is finite.

Proof. Define $\eta_x(m) = (x - Q(m,x))/2$ as in the proof of Theorem 2.1.2. By inspecting the proofs of Theorem 2.3.2 and Theorem 2.2.7, we see that if $x > \varepsilon_0$ the constant $M = M(\eta_x)$ can be taken to be the result of iterating some function

$$t \mapsto b(\eta_x(h(t)))$$

some $n(\varepsilon_0)$ times starting with $n(\varepsilon_0)$, where

- $b : (0,1] \to \mathbf{N}$ is a decreasing function (from Baer's theorem, Lemma 2.2.6),
- $h : \mathbf{N} \to \mathbf{N}$ is an increasing function (from Hall's theorem, Lemma 2.3.4),
- $n : (0,1] \to \mathbf{N}$ is a decreasing function (from Neumann's theorem, Theorem 2.2.5).

Let $k = n(\varepsilon_0) + 1$ and $m(\varepsilon) = \max(h(b(\varepsilon)), h(n(\varepsilon_0)))$, and suppose

$$x \in \mathcal{P} \cap [0, \varepsilon_0]^c \cap \bigcap_{i=0}^{k-1} (\mathcal{E}_{m(\varepsilon_i)} + [0, \varepsilon_{i+1}])^c. \tag{2.2}$$

Define the sequence $t_0, t_1, \ldots, t_k$ by

$$\begin{aligned} t_0 &= n(\varepsilon_0), \\ t_{i+1} &= b(\eta_x(h(t_i))) \quad \text{for } 0 \leqslant i < k. \end{aligned}$$

Then inductively

$$\begin{aligned} h(t_i) &\leqslant m(\varepsilon_i), \\ \eta_x(h(t_i)) &\geqslant \varepsilon_{i+1}, \\ t_{i+1} &\leqslant b(\varepsilon_{i+1}), \\ h(t_{i+1}) &\leqslant m(\varepsilon_{i+1}) \end{aligned}$$

for all i in the range $0 \leqslant i \leqslant k-1$, so

$$\eta_x(M) = \eta_x(t_{k-1}) \geqslant \eta_x(h(t_{k-1})) \geqslant \varepsilon_k.$$

But from the proof of Theorem 2.1.2 we know that

$$(x - \eta_x(M), x) \cap \mathcal{P} = \varnothing,$$

so there can be at most $1/\varepsilon_k$ elements x in the set (2.2), and the lemma follows. □

Proof of Theorem 2.1.4. From the discussion in the previous subsection it suffices to prove that $\alpha \leqslant \omega^2$. By the lemma we have

$$\overline{\mathcal{P}} \subset [0, \varepsilon_0] \cup \bigcup_{i=0}^{k-1} (\mathcal{E}_{m(\varepsilon_i)} + [0, \varepsilon_{i+1}]) \cup F$$

for some finite set F. From the rule $(X \cup Y)' = X' \cup Y'$ we have

$$\overline{\mathcal{P}}' \subset [0, \varepsilon_0] \cup \bigcup_{i=0}^{k-1} (\mathcal{E}_{m(\varepsilon_i)} + [0, \varepsilon_{i+1}]).$$

But since this holds for all $\varepsilon_k > 0$ we have

$$\overline{\mathcal{P}}' \subset [0, \varepsilon_0] \cup \bigcup_{i=0}^{k-2} (\mathcal{E}_{m(\varepsilon_i)} + [0, \varepsilon_{i+1}]) \cup \mathcal{E}_{m(\varepsilon_{k-1})}.$$

Now using $\mathcal{E}_n' = \mathcal{E}_{n-1}$ and the rule $(X \cup Y)' = X' \cup Y'$ again we have

$$\overline{\mathcal{P}}^{1+m(\varepsilon_{k-1})} \subset [0, \varepsilon_0] \cup \bigcup_{i=0}^{k-2} (\mathcal{E}_{m(\varepsilon_i)} + [0, \varepsilon_{i+1}]).$$

In particular

$$\overline{\mathcal{P}}^{\omega} \subset [0, \varepsilon_0] \cup \bigcup_{i=0}^{k-2} (\mathcal{E}_{m(\varepsilon_i)} + [0, \varepsilon_{i+1}]).$$

Repeating this argument another $k-1$ times, we have

$$\begin{aligned}
\overline{\mathcal{P}}^{\omega \cdot 2} &\subset [0, \varepsilon_0] \cup \bigcup_{i=0}^{k-3} (\mathcal{E}_{m(\varepsilon_i)} + [0, \varepsilon_{i+1}]), \\
&\vdots \\
\overline{\mathcal{P}}^{\omega \cdot (k-1)} &\subset [0, \varepsilon_0] \cup (\mathcal{E}_{m(\varepsilon_0)} + [0, \varepsilon_1]), \\
\overline{\mathcal{P}}^{\omega \cdot k} &\subset [0, \varepsilon_0].
\end{aligned}$$

Thus

$$\overline{\mathcal{P}}^{\omega^2} \subset [0, \varepsilon_0]$$

for all $\varepsilon_0 > 0$, so

$$\overline{\mathcal{P}}^{\omega^2} \subset \{0\},$$

as claimed. □

The same argument applies unchanged in the case of $\mathcal{P}_{\mathrm{b}}$.

2.5 Ultrafinite groups (optional)

Throughout the proof of Joseph's conjectures we needed to quantify dependencies in terms of an arbitrary decreasing function η, which ultimately we took to describe the quality of under-approximations by Egyptian fractions. Consequently the proof requires pretty careful epsilon management. Alternatively we could use ultrafilters to give a cleaner proof, as we explain in this optional section. We did not do so above for two reasons: first, many readers are uninterested in ultrafilters; second, we need the finitary proof in order to effectively bound the order type of $\mathcal{P}$.

For all the necessary background on ultrafinite groups and Loeb measure refer to Appendix 2.A. See in particular Theorem 2.A.2 for a version of Fubini's theorem adapted to Loeb measure: we will use it throughout this section without special reference.

Given a group G, the *FC-center* of G is the subgroup of all $g \in G$ having only finitely many conjugates.

Theorem 2.5.1 (Neumann's theorem, ultrafinitary version)**.** *Let G be an ultrafinite group such that* $\operatorname{st}\Pr(G) > 0$*. Let F be the FC-center of G. Then F is a finite-index internal subgroup such that $[F, F]$ is finite.*

Proof. Let $X_n \subset G$ be the set of all $x \in G$ with at most n conjugates. Since

$$\operatorname{st}\Pr(G) = \int \operatorname{st}(1/|x^G|)\, d\mu(x) \leqslant \mu(X_n) + 1/n,$$

we see that X_n has positive measure for large enough n, and thus the group F_n generated by X_n has finite index. Since (F_n) is therefore an increasing sequence of (eventually) finite-index subgroups, it must terminate with some subgroup F, which must be the FC-center. Moreover, by Lemma 2.2.1 we have

$$F_n = X_n^r \subset X_{n^r}$$

for $r = \lfloor 6\mu(X_n)^{-1} \rfloor$, so in fact $F = X_n$ for some n. Thus F is internal, and B. Neumann's theorem (Theorem 2.2.4) implies that $[F, F]$ is finite. □

The proof of the next lemma is exactly that of Lemma 2.3.5: the ultrafinitary perspective adds little here.

Lemma 2.5.2. *Let G be an ultrafinite group such that $[G, G]$ is finite. Then* $\Pr(G)$ *has the form*

$$\text{(standard rational)} + \text{(nonnegative infinitesimal)}.$$

Sketch. As in the proof of Lemma 2.3.5 we have

$$\begin{aligned}\Pr(G) &= \mathbf{E}_{x\in G}\mathbf{E}_{y\in G}1_{[x,y]=1} \\ &= \mathbf{E}_{y\in G/Z_2}\mathbf{E}_{\gamma\in\widehat{A}}\frac{1}{|G/G_{y,\gamma}|}1_{\gamma([G_{y,\gamma},y])=1},\end{aligned}$$

where Z_2 is the second centre of G, $A = [G, G] \cap Z(G)$, and $G_{y,\gamma}$ is the subgroup of all $x \in G$ such that $[x, y] \in A$ and $\gamma([x, Z_2]) = 1$. Since A is finite, and Z_2 has finite index by Hall's theorem (Lemma 2.3.4), we have expressed $\Pr(G)$ in the desired form. □

Now let G_n be a sequence of finite groups, let $p \in \beta\mathbf{N}\backslash\mathbf{N}$, and let G be the ultrafinite group $\prod_{n\to p} G_n$. Let F be the FC-center of G. Then

$$\begin{aligned}\mu_{G^2}\left(\{(x, y) \in G^2\backslash F^2 : [x, y] = 1\}\right) &\leqslant 2\mu_{G^2}\left(\{(x, y) \in G^2 : x \notin F, [x, y] = 1\}\right) \\ &= 2\int_{G\backslash F} \mathrm{st}(1/|x^G|)\, d\mu(x) \\ &= 0.\end{aligned}$$

Thus if $\mathrm{st}\Pr(G) > 0$, Theorem 2.5.1 implies that

$$\Pr(G) = \frac{1}{|G/F|^2}\Pr(F) + \varepsilon,$$

where ε is nonnegative and infinitesimal. It now follows from Lemma 2.5.2 that $\Pr(G)$ has the form

$$\text{(standard rational)} + \text{(nonnegative infinitesimal)}.$$

But this is nothing more than a neat restatement of J1 and J2.

2.6 Compact groups

It is natural to extend the definition of commuting probability to infinite compact groups G, each of which is endowed with its unique normalized Haar measure μ_G. Thus we define

$$\Pr(G) = \mu_{G^2}\left(\{(x,y) \in G^2 : xy = yx\}\right).$$

The thesis of this section however is that this is not actually that interesting, because if $\Pr(G) > 0$ then the structure of G is harshly restricted.

Let us hasten to add that this is quite distinct from our consideration in Section 2.5 of the commuting probability of ultrafinite groups: there ultrafinite groups served merely as a logical tool for epsilon management while we were still effectively studying finite groups; here the compact group G is intrinsically infinite. There are strong parallels, however, as we will see.

The structure of compact groups of positive commuting probability is determined by following theorem of Hofmann and Russo [HR12], which extends earlier work of Lévai and Pyber [LP00] in the profinite case. We give a much simpler proof, very similar to that of Theorem 2.5.1.

As in Section 2.5 we define the *FC-center* of a group G to be the subgroup $F(G)$ of all elements $g \in G$ with finitely many conjugates.

Theorem 2.6.1 (Hofmann–Russo [HR12])**.** *Let G be a compact group such that $\Pr(G) > 0$. Let F be the FC-center of G. Then F is a finite-index open subgroup such that $[F,F]$ is finite. Moreover, $Z(F)$ is a finite-index open subgroup of F.*

Proof. Let $X_n \subset G$ be the set of all $x \in G$ with at most n conjugates. Then X_n is closed, since any element x with at least $n+1$ distinct conjugates x^{g_i} has a neighbourhood U such that for all $u \in U$ the points u^{g_i} are distinct. Since

$$\Pr(G) = \int 1/|x^G| \, d\mu_G(x) \leqslant \mu_G(X_n) + 1/n,$$

we see that X_n has positive measure for large enough n, and thus the group F_n generated by X_n has finite index. Since (F_n) is therefore an increasing sequence of (eventually) finite-index subgroups, it must terminate with some subgroup F, which must be the FC-center. Moreover, by Lemma 2.2.1 we have

$$F_n = X_n^r \subset X_{n^r}$$

for $r = \lfloor 6\mu(X_n)^{-1} \rfloor$, so in fact $F = X_n$ for some n. Thus F is closed and finite-index, hence open, and B. Neumann's theorem (Theorem 2.2.4) implies that $[F,F]$ is finite.

In particular in its own right F is a compact group such that $[F,F]$ is finite. Since the commutator map $[,] : F \times F \to [F,F]$ is then a continuous map to a discrete set satisfying $[F,1] = 1$ there must be a neighbourhood U of 1 such that $[F,U] = 1$. This implies that $Z(F)$ is open, hence of finite-index in F. □

As a corollary we can prove that for every compact group G such that $\Pr(G) > 0$ there is a finite group H such that $\Pr(G) = \Pr(H)$. We need a simple lemma first.

Lemma 2.6.2. *For every $n \in \mathbf{N}$ there is a finite group K_n such that $\Pr(K_n) = 1/n$.*

Proof. We need one calculation: If n is odd, then in the dihedral group D_n there are precisely $1 + 1 + (n-1)/2 = (n+3)/2$ conjugacy classes: one for the identity, one for the set of reflections, and one for each pair of mutually inverse rotations. Thus

$$\Pr(D_n) = \frac{\text{number of conjugacy classes}}{|D_n|} = \frac{n+3}{4n}.$$

Now we define the groups K_n inductively. Put $K_1 = 1$ and $K_2 = D_3$. If $n > 2$ is even then put $K_n = K_2 \times K_{n/2}$. If $n = 4k+1 > 2$ then put $K_n = K_{k+1} \times D_n$. If $n = 4k+3 > 2$ then put $K_n = K_{k+1} \times D_{3n}$. □

Corollary 2.6.3. *For every compact group G such that $\Pr(G) > 0$ there is a finite group H such that $\Pr(G) = \Pr(H)$.*

Proof. Suppose $\Pr(G) > 0$. Then by Theorem 2.6.1 the FC-center F of G is a finite-index open subgroup of G, $[F,F]$ is finite, and $Z(F)$ is a finite-index open subgroup of F. Moreover since

$$\begin{aligned}\mu_{G^2}\left(\{(x,y) \in G^2 \backslash F^2 : [x,y] = 1\}\right) &\leqslant 2\mu_{G^2}\left(\{(x,y) \in G^2 : x \notin F, [x,y] = 1\}\right) \\ &= 2\int_{G\backslash F} 1/|x^G|\, d\mu_G(x) \\ &= 0,\end{aligned}$$

we have

$$\Pr(G) = \frac{1}{|G/F|^2}\Pr(F).$$

Now a theorem of Hall [Hal40] asserts that F is *isoclinic* to a group E satisfying $Z(E) \leqslant [E,E]$, i.e., there exist isomorphisms $f : F/Z(F) \to E/Z(E)$ and $g : [F,F] \to [E,E]$ making the diagram

$$\begin{array}{ccc} F/Z(F) \times F/Z(F) & \xrightarrow{[,]} & [F,F] \\ \downarrow{\scriptstyle f\times f} & & \downarrow{\scriptstyle g} \\ E/Z(E) \times E/Z(E) & \xrightarrow{[,]} & [E,E] \end{array}$$

commute. But since $E/Z(E) \cong F/Z(F)$ and $[E,E] \cong [F,F]$ are finite and $Z(E) \leqslant [E,E]$, it follows that E is a finite group, and the above diagram implies that $\Pr(F) = \Pr(E)$. Thus it suffices to take

$$H = K_{|G/F|^2} \times E. \qquad \square$$

The corollary can be compared with Joseph's third conjecture (which is still open), which can be stated as follows: For every ultrafinite group G such that $\operatorname{st}\Pr(G) > 0$ there is a finite group H such that $\operatorname{st}\Pr(G) = \Pr(H)$.

2.7 Nilpotency degree

The *s-step nilpotency degree* of a finite group G is the quantity

$$d_s(G) = \frac{1}{|G|^{s+1}} |\{(x_1, \ldots, x_{s+1}) \in G^{s+1} : [x_1, \ldots, x_{s+1}] = 1\}|,$$

where we inductively define $[x_1, \ldots, x_{s+1}] = [[x_1, \ldots, x_s], x_{s+1}]$. In this notation the commuting probability $\Pr(G)$ is the same as the 1-step nilpotency degree $d_1(G)$. Motivated by Neumann's theorem, it is natural and tempting to ask whether every group G such that $d_s(G)$ is bounded away from zero has a large s-step nilpotent section. We show that this need not be the case for $s = 2$: we will give an example of a group G with 2-step nilpotency degree at least $1/4$ and no large 2-step nilpotent section. We do not know whether G must at least have a large 3-step nilpotent section.

Fix a prime p (take $p = 2$ to get $1/4$ as claimed above) and let $F_{n,3,p}$ be the group with generators

$$\begin{aligned} \alpha_i &\qquad (1 \leqslant i \leqslant n), \\ \beta_{ij} &\qquad (1 \leqslant j < i \leqslant n), \\ \gamma_{ijk} &\qquad (1 \leqslant j < i \leqslant n, j \leqslant k \leqslant n), \end{aligned}$$

and relators

$$\begin{aligned} [\alpha_i, \alpha_j] &= \beta_{ij} && (1 \leqslant j < i \leqslant n), \\ [\beta_{ij}, \alpha_k] &= \gamma_{ijk} && (1 \leqslant j < i \leqslant n, j \leqslant k \leqslant n), \\ [\gamma_{ijk}, \alpha_l] &= 1 && (1 \leqslant j < i \leqslant n, j \leqslant k \leqslant n, 1 \leqslant l \leqslant n), \\ \alpha_i^p = \beta_{ij}^p = \gamma_{ijk}^p &= 1 && (\text{all valid } i, j, k). \end{aligned}$$

Every element of $F_{n,3,p}$ has a unique representation of the form

$$\prod_{1\leqslant i\leqslant n}\alpha_i^{e_i}\prod_{1\leqslant j<i\leqslant n}\beta_{ij}^{e_{ij}}\prod_{\substack{1\leqslant j<i\leqslant n\\ j\leqslant k\leqslant n}}\gamma_{ijk}^{e_{ijk}},$$

where each $e_i, e_{ij}, e_{ijk} \in \{0, 1, \ldots, p-1\}$ and the factors in the products are taken in some fixed but arbitrary order. For more details on this standard material see M. Hall [Hal59].

Our group $G = G_{n,p}$ is a quotient of $F_{n,3,p}$ by a certain central subgroup. Specifically, to the generators above we add

$$\gamma_i \qquad (1 \leqslant i \leqslant n)$$

and to the relators we add

$$\begin{aligned}
\gamma_{ijk} &= 1 && (1 \leqslant j < i \leqslant n, j \leqslant k \leqslant n, k \notin \{j, i\}),\\
\gamma_{ijj} &= \gamma_i && (1 \leqslant j < i \leqslant n),\\
\gamma_{iji} &= \gamma_j^{-1} && (1 \leqslant j < i \leqslant n).
\end{aligned}$$

Every element of G has a unique representation of the form

$$\prod_{1\leqslant i\leqslant n}\alpha_i^{e_i}\prod_{1\leqslant j<i\leqslant n}\beta_{ij}^{e_{ij}}\prod_{1\leqslant i\leqslant n}\gamma_i^{e'_i},$$

where as before each $e_i, e_{ij}, e'_i \in \{0, 1, \ldots, p-1\}$ and the factors are taken in some fixed but arbitrary order. The order of $G = Z_3(G) = \langle \alpha_i \rangle$ is $p^{2n+\binom{n}{2}}$, the order of $\Gamma_2(G) = Z_2(G) = \langle \beta_{ij} \rangle$ is $p^{n+\binom{n}{2}}$, and the order of $\Gamma_3(G) = Z(G) = \langle \gamma_i \rangle$ is p^n.

From the relations defining G it follows that

$$[\alpha_i, \alpha_j, \alpha_k] = \gamma_i^{\delta_{jk}}\gamma_j^{-\delta_{ik}} \qquad (\text{all } i, j, k),$$

where δ_{ab} is defined as 1 if $a = b$ and 0 otherwise. Thus by trilinearity of the triple commutator every triple commutator in G has the form

$$\left[\prod \alpha_i^{e_i}, \prod \alpha_i^{f_i}, \prod \alpha_i^{g_i}\right] = \prod \gamma_i^{e_i(f\cdot g)-f_i(e\cdot g)}. \tag{2.3}$$

In particular if e and f are fixed then the above triple commutator is trivial provided $f \cdot g = e \cdot g = 0$, so $d_2(G) \geqslant 1/p^2$.

Since $\Gamma_3(G) \cong \mathbf{F}_p^n$ is large, plainly G has no large 2-step nilpotent quotient. Moreover if $H \leqslant G$ has bounded-index then its image $H\Gamma_2/\Gamma_2$ under the abelianization map is a bounded-codimension subspace of $G/\Gamma_2 \cong \mathbf{F}_p^n$, so from (2.3) one can see that $\Gamma_3(H)$ is a bounded-codimension subspace of $\Gamma_3(G) \cong \mathbf{F}_p^n$ (fix f and g such that $f \cdot g \neq 0$ and consider varying e), so H also has no large 2-step nilpotent quotient.

2.A Appendix: A crash course on ultraproducts and Loeb measure

For the convenience of the reader we provide here a rapid introduction to ultraproducts and Loeb measure. See also Bergelson and Tao [BT14, Section 2] for similar coverage.

2.A.1 What is $\beta\mathbf{N}$?

A *filter* p on $\mathbf{N}$ is a collection of subsets of $\mathbf{N}$ with the following properties:

1. $\mathbf{N} \in p$ and $\emptyset \notin p$;
2. if $A \in p$ and $A \subset B$ then $B \in p$;
3. if $A, B \in p$ then $A \cap B \in p$.

An *ultrafilter* is a maximal filter. Equivalently we could add the fourth condition

4. for every $A \subset \mathbf{N}$ either $A \in p$ or $A^c \in p$.

There are many ultrafilters. For example for any $x \in \mathbf{N}$ one could take the ultrafilter p_x such that $A \in p_x$ if and only if $x \in A$: these are the *principal* ultrafilters. More interestingly, by Zorn's lemma there is at least one ultrafilter extending the the filter of cofinite sets. In fact by an easy exercise every nonprincipal ultrafilter extends the filter of cofinite sets.

We often use the notation $\beta\mathbf{N}$ to denote the set of ultrafilters on $\mathbf{N}$, in which case we think of $\mathbf{N}$ as a subset of $\beta\mathbf{N}$ by mapping each $x \in \mathbf{N}$ to the corresponding principal ultrafilter p_x. The set $\beta\mathbf{N}\backslash\mathbf{N}$ therefore denotes the set of nonprincipal ultrafilters. We call $\beta\mathbf{N}$ the *Stone–Čech* compactification of $\mathbf{N}$; it carries a natural topology, as well as a semicontinuous semigroup structure, neither of which need concern us here.

Given $p \in \beta\mathbf{N}\backslash\mathbf{N}$ and some event $E(x)$ depending on $x \in \mathbf{N}$, we say that $E(x)$ holds "for p-almost-all x" if $\{x : E(x)\} \in p$. We might also write $\forall_p x : E(x)$. It follows from the properties of ultrafilters that for arbitrary events E and F we have

$$\begin{aligned} \forall_p x : (E(x) \wedge F(x)) &\iff (\forall_p x : E(x)) \wedge (\forall_p x : F(x)), \\ \forall_p x : (E(x) \vee F(x)) &\iff (\forall_p x : E(x)) \vee (\forall_p x : F(x)), \\ \neg \forall_p x : E(x) &\iff \forall_p x : (\neg E(x)). \end{aligned}$$

This "logical Freshman's dream" is one of the reasons ultrafilters are so useful for quantification and epsilon management. Note also that if p is nonprincipal then if $\forall_p x : E(x)$ then in particular $E(x)$ holds for infinitely many x.

2.A.2 Ultraproducts and the nonstandard reals

Fix from now on some particular $p \in \beta\mathbf{N}\backslash\mathbf{N}$. Given a sequence of structures (X_n) (sets, groups, fields, ...), the *ultraproduct* $\prod_{n\to p} X_n$ is defined to be the set of all sequences $(x_n) \in \prod X_n$ modulo p-almost-everywhere equality. Thus $(x_n) = (y_n)$ in $\prod_{n\to p} X_n$ if and only if $\{n : x_n = y_n\} \in p$.

For example suppose we take X_n to be the field of real numbers $\mathbf{R}$. Then the ultralimit ${}^*\mathbf{R} = \prod_{n\to p} \mathbf{R}$ is called the field of *hyperreals* or *nonstandard reals*. It is indeed a field, where we define $(x_n) + (y_n) = (x_n + y_n)$ and $(x_n)(y_n) = (x_n y_n)$. It follows from the properties of ultrafilters that these operations are well defined and that the field axioms are satisfied. For example, suppose $(x_n) \neq 0$. Then if we define $y_n = x_n^{-1}$ whenever $x_n \neq 0$ and, say, $y_n = 1$ otherwise, then indeed $x_n y_n = 1$ for p-almost-all n, so (x_n) has multiplicative inverse (y_n).

In general operations tend to extend in an obvious way to the ultraproduct. Moreover a first-order property is true in $\prod_{n\to p} X_n$ if and only if it is true in X_n for p-almost-all n. The buzzword for this phenomenon is "Łos's theorem".

The reals embed naturally into the nonstandard reals via the diagonal embedding $r \mapsto (r)$. We will abuse notation and denote (r) simply by r. The nonstandard reals are much bigger than this though: there are positive nonstandard reals (exercise: define *positive*) smaller than every positive standard real, e.g., the nonstandard real $(1/n)$, and there are likewise nonstandard reals larger than every standard real, e.g., the nonstandard real (n). We call $x = (x_n) \in {}^*\mathbf{R}$ *bounded* if $|x| \leqslant M$ for some $M \in \mathbf{R}$, equivalently if $|x_n| \leqslant M$ for p-almost-all n. In this case we can define the *standard part* $\operatorname{st} x$ of x as the supremum of all $r \in \mathbf{R}$ such that $r < x$. Thus $\operatorname{st} x$ is a real number such that $|x - \operatorname{st} x|$ is smaller than every positive real (i.e., such that $x - \operatorname{st} x$ is *infinitesimal*).

Commonly we consider an ultralimit $G = \prod_{n\to p} G_n$ of groups G_n. The group operation extends naturally to G, just as the field operations of $\mathbf{R}$ extended naturally to ${}^*\mathbf{R}$. The properties of $G = \prod_{n\to p} G_n$ tend to reflect the asymptotic properties of the sequence of groups (G_n). For example, G has an element of finite order p if and only if G_n has such an element for p-almost-all n. If each group G_n is finite, then we call G an *ultrafinite* group.

2.A.3 Loeb measure

Given a sequence of countable sets X_n, consider the ultraproduct $X = \prod_{n\to p} X_n$. A subset S of X is called *internal* if it is defined by subsets $S_n \subset X_n$ in the same way

as X, namely if $(s_n) \in S$ if and only if $s_n \in S_n$ for p-almost-all n; in this case we write $S = \prod_{n \to p} S_n$. The colletion of internal subsets is a ring, and the σ-algebra they generate is called the *Loeb σ-algebra* $\mathcal{L}_X$.

Now suppose for each n that we have some measure μ_n defined on the set of all subsets of X_n. Then (μ_n) defines a function on the internal subsets of X by mapping $S = \prod_{n \to p} S_n$ to the hyperreal $(\mu_n(S_n))$. We define a new function μ on the internal sets by taking standard parts:

$$\mu(S) = \mathrm{st}(\mu_n(S_n)).$$

It turns out that μ defines a premeasure.

Lemma 2.A.1. *The function μ is a premeasure on the ring of internal sets.*

Proof. Finite additivity is easy to verify, as

$$\left(\prod_{n \to p} S_n\right) \cup \left(\prod_{n \to p} S'_n\right) = \prod_{n \to p} (S_n \cup S'_n).$$

We thus need only check that

$$\mu(S) = \lim_{m \to \infty} \mu(S^m)$$

whenever $S^1 \subset S^2 \subset \cdots$ is a sequence of internal sets whose union S is also internal (we use superscripts to distinguish from the standard components S_n). We will accomplish this by surprisingly trivial means: we claim that under these hypotheses in fact $S = S^m$ for some m. (The buzzword for this type of argument is "countable saturation".)

Write $S = \prod_{n \to p} S_n$ and $S^m = \prod_{n \to p} S_n^m$, and suppose $S = \bigcup_{m=1}^{\infty} S^m$. By replacing S_n^m with $S_n^1 \cup \cdots \cup S_n^m$ for each m and n if necessary, we may assume that

$$S_n^1 \subset S_n^2 \subset \cdots$$

for each n. For each n, write m_n for the largest index such that

$$S_n \notin S_n^{m_n},$$

or 0 if none exists, or n if there is no largest. If $m_n = 0$ for p-almost-all n then $S = S^1$ and we're done. Otherwise for p-almost-all n pick

$$x_n \in S_n \backslash S_n^{m_n}$$

and let $x = (x_n)$. Then clearly $x \in S = \bigcup_{m=1}^{\infty} S^m$, so $x \in S^m$ for some m, so for p-almost-all n we have $x \in S_n^m$. Since $x_n \notin S_n^{m_n}$ this implies $m_n < m$ for p-almost-all n. Thus by definition of m_n for p-almost-all n we have $S_n \subset S_n^m$, whence we deduce that $S = S^m$. □

It follows by usual measure theory that μ extends to a countably additive measure on $\mathcal{L}_X$. We call μ the *Loeb measure*, and the space $(X, \mathcal{L}_X, \mu)$ the *Loeb space*. It is something of a paradox that the even though the above proof is so trivial the resulting measure μ is still so interesting.

Actually it was not important for the sets X_n to be countable, and we coud have started each X_n with its own σ-algebra Σ_n and let μ_n be a measure on Σ_n. Of course we would then only consider the σ-algebra generated by ultraproducts of measurable sets. We could even allow Σ_n to be only an algebra and μ_n only finitely additive, and the resulting Loeb space would still be a bona fide measure space.

2.A.4 Fubini's theorem

It is a sad fact of life that, given Loeb spaces $(X, \mathcal{L}_X, \mu_X)$ and $(Y, \mathcal{L}_Y, \mu_Y)$, based on ultraproducts of countable sets as above, the product $\mathcal{L}_X \times \mathcal{L}_Y$ of the two Loeb σ-algebras is not the same as the Loeb σ-algebra $\mathcal{L}_{X\times Y}$ of the product: the latter, being generated by the internal subsets of

$$X \times Y = \prod_{n\to p} X_n \times Y_n,$$

is much larger. This is potentially annoying because the traditional form of Fubini's theorem only applies to functions which are $\mathcal{L}_X \times \mathcal{L}_Y$-measurable, whereas functions on $X \times Y$ that we want to apply Fubini's theorem to are typically only $\mathcal{L}_{X\times Y}$-measurable. Fortunately, however, we can prove a version of Fubini's theorem for functions which are only $\mathcal{L}_{X\times Y}$-measurable.

Theorem 2.A.2 (Fubini's theorem for Loeb measure)**.** *Fix Loeb spaces $(X, \mathcal{L}_X, \mu_X)$ and $(Y, \mathcal{L}_Y, \mu_Y)$. If $f : X \times Y \to \mathbf{R}$ is a nonnegative $\mathcal{L}_{X\times Y}$-measurable function, then for all $x \in X$ the function $y \mapsto f(x, y)$ is $\mathcal{L}_Y$-measurable, the function $x \mapsto \int_Y f(x, y)\, d\mu_Y(y)$ is $\mathcal{L}_X$-measurable, and*

$$\int_{X\times Y} f\, d\mu_{X\times Y} = \int_X \int_Y f(x, y)\, d\mu_Y(y)\, d\mu_X(x). \tag{2.4}$$

In general if $f \in L^1(\mu_{X\times Y})$ then for almost all $x \in X$ the function $y \mapsto f(x,y)$ is in $L^1(\mu_Y)$, the function defined almost everywhere by $x \mapsto \int f(x,y)\,d\mu_Y(y)$ is in $L^1(\mu_X)$, and again (2.4) holds.

The same result holds of course with the roles of X and Y reversed, so (2.4) implies that

$$\int_X \int_Y f(x,y)\,d\mu_Y(y)\,d\mu_X(x) = \int_Y \int_X f(x,y)\,d\mu_X(x)\,d\mu_Y(y)$$

if f is $\mathcal{L}_{X\times Y}$-measurable and either nonnegative or integrable.

Again, the advantage here over the usual form of Fubini's theorem is that we only assume that f is $\mathcal{L}_{X\times Y}$-measurable, which is much weaker than assuming f is $\mathcal{L}_X \times \mathcal{L}_Y$-measurable.

Sketch. As in the proof of the usual form of Fubini's theorem, modulo a few applications of the monotone convergence theorem it suffices to check the theorem when f is the indicator of some $Q \in \mathcal{L}_{X\times Y}$. Moreover, modulo an application of the monotone class theorem (or the π-λ theorem) we may assume that Q is internal, say $Q = \lim_{n\to p} Q_n$, where $Q_n \subset X_n \times Y_n$ for each n. Now we appeal to the trivial form of Fubini's theorem on $X_n \times Y_n$. See [BT14, Theorem 19] for more details. □

If X_n and Y_n are not necessarily countable but instead start life as measurable spaces in their own rights then result still holds: at the end of the proof above we would just appeal to the usual form of Fubini's theorem on $X_n \times Y_n$. If on the other hand μ_{X_n} and μ_{Y_n} are only finitely additive then the theorem need not hold.

The fundamental reason we need the above theorem is that the group operation $(x,y) \mapsto xy^{-1}$ on an ultrafinite group G is not $\mathcal{L}_G \times \mathcal{L}_G$-measurable, but only $\mathcal{L}_{G^2}$-measurable. See Tserunyan [Tse14] for further discussion about this, and a general definition of groups obeying such a rule.

References

[BT14] V. Bergelson and T. Tao. "Multiple recurrence in quasirandom groups". *Geom. Funct. Anal.* 24.1 (2014), pp. 1–48. ISSN: 1016-443X. DOI: 10.1007/s00039-014-0252-0.

[Das14] A. Dasgupta. *Set theory: with an introduction to real point sets*. Birkhäuser/Springer, New York, 2014, pp. xvi+444. ISBN: 978-1-4614-8853-8; 978-1-4614-8854-5. DOI: 10.1007/978-1-4614-8854-5.

[Ebe15] S. Eberhard. "Commuting probabilities of finite groups". *Bull. Lond. Math. Soc.* 47.5 (2015), pp. 796–808. ISSN: 0024-6093. DOI: 10.1112/blms/bdv050.

[ET68] P. Erdős and P. Turán. "On some problems of a statistical group-theory. IV". *Acta Math. Acad. Sci. Hungar* 19 (1968), pp. 413–435. ISSN: 0001-5954.

[Gus73] W. H. Gustafson. "What is the probability that two group elements commute?" *Amer. Math. Monthly* 80 (1973), pp. 1031–1034. ISSN: 0002-9890.

[Hal59] M. Hall Jr. *The theory of groups*. The Macmillan Co., New York, N.Y., 1959, pp. xiii+434.

[Hal40] P. Hall. "The classification of prime-power groups". *J. Reine Angew. Math.* 182 (1940), pp. 130–141. ISSN: 0075-4102.

[Hal56] P. Hall. "Finite-by-nilpotent groups". *Proc. Cambridge Philos. Soc.* 52 (1956), pp. 611–616.

[Ham92] Y. O. Hamidoune. "On a subgroup contained in some words with a bounded length". *Discrete Math.* 103.2 (1992), pp. 171–176. ISSN: 0012-365X. DOI: 10.1016/0012-365X(92)90267-J. URL: http://dx.doi.org/10.1016/0012-365X(92)90267-J.

[Heg13] P. Hegarty. "Limit points in the range of the commuting probability function on finite groups". *J. Group Theory* 16.2 (2013), pp. 235–247. ISSN: 1433-5883. DOI: 10.1515/jgt-2012-0040.

[HR12] K. H. Hofmann and F. G. Russo. "The probability that x and y commute in a compact group". *Math. Proc. Cambridge Philos. Soc.* 153.3 (2012), pp. 557–571. ISSN: 0305-0041. DOI: 10.1017/S0305004112000308.

[Jos77] K. S. Joseph. "Research Problems: Several Conjectures on Commutativity in Algebraic Structures". *Amer. Math. Monthly* 84.7 (1977), pp. 550–551. ISSN: 0002-9890. DOI: 10.2307/2320020.

[Jos69] K. S. Joseph. "Commutativity in non-abelian groups". PhD thesis. University of California, Los Angeles, 1969.

[LP00] L. Lévai and L. Pyber. "Profinite groups with many commuting pairs or involutions". *Arch. Math. (Basel)* 75.1 (2000), pp. 1–7. ISSN: 0003-889X. DOI: 10.1007/s000130050466.

[Neu54] B. H. Neumann. "Groups covered by permutable subsets". *J. London Math. Soc.* 29 (1954), pp. 236–248. ISSN: 0024-6107.

[Neu89] P. M. Neumann. "Two combinatorial problems in group theory". *Bull. London Math. Soc.* 21.5 (1989), pp. 456–458. ISSN: 0024-6093. DOI: 10.1112/blms/21.5.456.

[Rob96] D. J. S. Robinson. *A course in the theory of groups*. Second. Vol. 80. Graduate Texts in Mathematics. Springer-Verlag, New York, 1996, pp. xviii+499. ISBN: 0-387-94461-3. DOI: 10.1007/978-1-4419-8594-1.

[Rus79] D. J. Rusin. "What is the probability that two elements of a finite group commute?" *Pacific J. Math.* 82.1 (1979), pp. 237–247. ISSN: 0030-8730. URL: http://projecteuclid.org/euclid.pjm/1102785075.

[Tse14] A. Tserunyan. *Mixing and triple recurrence in probability groups*. 2014. URL: http://arxiv.org/abs/1405.5629.

Chapter 3

Fixed sets of permutations

Abstract. We consider two related problems about permutations.

First, we count the number of permutations having a fixed set of size k, or in other words failing to induce a derangement in the action on k-sets. Let $i(n,k)$ be the proportion of such permutations in $\mathcal{S}_n$. By adapting arguments of Ford in connection with the multiplication table problem from analytic number theory, we prove that $i(n,k) \asymp k^{-\delta}(1+\log k)^{-3/2}$ uniformly for $1 \leqslant k \leqslant n/2$, where $\delta = 1 - \frac{1+\log\log 2}{\log 2}$. As an application we show that the proportion of $\pi \in \mathcal{S}_n$ contained in a transitive subgroup not containing $\mathcal{A}_n$ is at least $n^{-\delta+o(1)}$ if n is even, partially answering a question of Cameron.

Second, we consider invariable generation in $\mathcal{S}_n$. Following Dixon, we say that permutations $\pi_1, \dots, \pi_r \in \mathcal{S}_n$ *invariably generate* $\mathcal{S}_n$ if, no matter how one chooses conjugates $\pi_1', \dots, \pi_r'$ of these permutations, $\pi_1', \dots, \pi_r'$ generate $\mathcal{S}_n$. We show that if π_1, π_2, π_3 are chosen uniformly at random from $\mathcal{S}_n$ then, with probability tending to 1 as $n \to \infty$, π_1, π_2, π_3 do not invariably generate $\mathcal{S}_n$, because they have fixed sets of a common size. By contrast it was shown recently by Pemantle, Peres and Rivin that four random elements do invariably generate $\mathcal{S}_n$, with probability bounded away from zero. We include a proof of this statement which, while sharing many features with their argument, is short and completely combinatorial.

This chapter combines the papers [EFG15b] and [EFG15a], both of which are joint work with Kevin Ford and Ben Green.

3.1 Introduction

In this chapter we are concerned with two related problems about fixed sets of random permutations $\pi \in \mathcal{S}_n$. First, given an integer k in the range $1 \leqslant k \leqslant n$, what is the probability $i(n,k)$ that π has a fixed set of size k? Second, given a few (three or four, to be precise) random permutations $\pi_1, \dots, \pi_r \in \mathcal{S}_n$, what is the probability that for some k in the range $0 < k < n$ each of $\pi_1, \dots, \pi_r$ has a fixed set of size k?

Both these questions are much less about group theory than they may appear to be at first glance, because they both depend just on the set of cycle lengths of π, the distribution of which we understand well. Let $c_j = c_j(\pi)$ be the number of cycles of length j in π, let $\mathbf{c} = (c_1, \dots, c_n)$, and let

$$\mathscr{L}(\mathbf{c}) = \left\{ \sum_{j=1}^{n} j x_j : 0 \leqslant x_j \leqslant c_j \text{ for each } j \right\}.$$

Then the first question above concerns $i(n,k) = \mathbf{P}(k \in \mathscr{L}(\mathbf{c}))$, while the second question concerns $\mathscr{L}(\mathbf{c}^{(1)}) \cap \dots \cap \mathscr{L}(\mathbf{c}^{(r)})$, where $\mathbf{c}^{(1)}, \dots, \mathbf{c}^{(r)}$ are independent copies of $\mathbf{c}$.

Now it is well known that $c_j(\pi)$ is approximately Poisson with parameter $1/j$, independently of other values of j. In particular we have the following lemma, which is fundamental for this chapter. See for example Arratia and Tavaré [AT92].

Lemma 3.1.1. *Let $\mathbf{X} = (X_1, X_2, \dots)$ be a sequence of independent Poisson random variables, where X_j has parameter $1/j$. If $c_j = c_j(\pi)$ is the number of cycles of length j in a random permutation $\pi \in \mathcal{S}_n$, then if k is fixed and $n \to \infty$ the distribution of $(c_1, \dots, c_k)$ converges to that of $\mathbf{X}_k = (X_1, \dots, X_k)$.*

Thus at least in the limit $n \to \infty$ we are effectively just asking questions about the random sumset $\mathscr{L}(\mathbf{X})$, as well as its self-intersections $\mathscr{L}(\mathbf{X}^{(1)}) \cap \dots \cap \mathscr{L}(\mathbf{X}^{(r)})$. Thus these questions have much more of an additive-combinatorial flavour than a group-theoretic flavour.

3.1.1 Permutations fixing a k-set

As explained already we have $i(n,k) = \mathbf{P}(k \in \mathscr{L}(\mathbf{c}))$, where $\mathbf{c}$ is the cycle type of a random permutation, and in particular by Lemma 3.1.1 we have

$$\lim_{n \to \infty} i(n,k) = \mathbf{P}(k \in \mathscr{L}(\mathbf{X}_k));$$

in particular the limit exists for all k. This makes it easy to compute $\lim_{n\to\infty} i(n,k)$ for small values of k. For example we have the well known result that

$$\lim_{n\to\infty} i(n,1) = \mathbf{P}(X_1 \geqslant 1) = 1 - e^{-1},$$

as well as the less well known result that

$$\lim_{n\to\infty} i(n,2) = 1 - \mathbf{P}(X_1 = X_2 = 0) - \mathbf{P}(X_1 = 1, X_2 = 0) = 1 - 2e^{-3/2},$$

and so on.

The behaviour of $i(n,k)$ as k grows, however, is much less clear. In fact, somewhat surprisingly even the behaviour of $\lim_{n\to\infty} i(n,k)$ as a function of k has only recently been at all well understood. The timeline is roughly as follows.

1. Dixon [Dix92] defined $i(n,k)$ while studying invariable generation and proved $i(n,k) \ll \exp(-c(\log k)^2(\log n)^{-1})$ for some constant $c > 0$.

2. The upper bound $i(n,k) \ll k^{-0.01}$ $(1 \leqslant k \leqslant n/2)$ was proved by Łuczak and Pyber [ŁP93]. (These authors did not make any special effort to optimize the exponent 0.01, but their method does not lead to a sharp bound.)

3. The asymptotic lower bound $\lim_{n\to\infty} i(n,k) \gg k^{-1}\log k$ is contained in a paper of Diaconis, Fulman, and Guralnick [DFG08].

4. Recently, Pemantle, Peres, and Rivin [PPR14, Theorem 1.7] established the asymptotic estimate
$$\lim_{n\to\infty} i(n,k) = k^{-\delta+o(1)},$$
where
$$\delta = 1 - \frac{1+\log\log 2}{\log 2} = 0.08607....$$

In this chapter we will prove the following estimate, which establishes the order of $i(n,k)$ up to constant factors uniformly for all n and k.

Theorem 3.1.2. *$i(n,k) \asymp k^{-\delta}(\log k)^{-3/2}$ uniformly for $2 \leqslant k \leqslant n/2$.*

Here and throughout this chapter we use the analytic-number-theoretic notation $X \ll Y$ to mean that $X \leqslant CY$ for some constant $C > 0$, and similarly the notation $X \asymp Y$ to mean that $c_1 Y \leqslant X \leqslant c_2 Y$ for some constants c_1, c_2, i.e., $X \ll Y$ and $X \gg Y$. We also use $O(\cdot)$ and $o(\cdot)$ notation as usual.

Theorem 3.1.2 has implications for a conjecture of Cameron related to random generation of the symmetric group. Cameron conjectured that the proportion of

$\pi \in \mathcal{S}_n$ contained in a transitive subgroup not containing $\mathcal{A}_n$ tends to zero: this was proved by Łuczak and Pyber [ŁP93] using their bound $i(n,k) \ll k^{-0.01}$. Cameron further guessed that this proportion might decay as fast as $n^{-1/2+o(1)}$ (see [ŁP93, Section 5]). However, Theorem 3.1.2 has the following corollary.

Corollary 3.1.3. *Suppose that $n > 2$ is even. Then the proportion of $\pi \in \mathcal{S}_n$ contained in a transitive subgroup not containing $\mathcal{A}_n$ is $\gg n^{-\delta}(\log n)^{-3/2}$.*

Proof. By Theorem 3.1.2 the proportion of $\pi \in \mathcal{S}_n$ fixing a set B_1 of size $n/2$ is $\asymp n^{-\delta}(\log n)^{-3/2}$. Such a permutation π must also fix the set $B_2 = \{1,\dots,n\}\backslash B_1$, and thus preserve the partition $\{B_1, B_2\}$ of $\{1,\dots,n\}$. Since $|B_1| = |B_2|$, the set of all τ preserving this partition is a transitive subgroup not containing $\mathcal{A}_n$. □

A matching upper bound $O(n^{-\delta}(\log n)^{-3/2})$ holds in Corollary 3.1.3, and in fact for odd n there is an upper bound of the form $O(n^{-\delta'})$ for some $\delta' > \delta$. This will be proved in a future paper [EFK16].

As we explained already our method is founded on the equality

$$\lim_{n\to\infty} i(n,k) = \mathbf{P}(k \in \mathscr{L}(\mathbf{X}_k)).$$

Our theorem, however, asserts a uniform estimate for $i(n,k)$. Thus what we really need is an approximate analogue of the above identity which asserts that

$$i(n,k) \asymp \mathbf{P}(k \in \mathscr{L}(\mathbf{X}_k))$$

uniformly for $k \leqslant n/2$. Instead of directly estimating the probability of k lying in $\mathscr{L}(\mathbf{X}_k)$, however, it will be more convenient to apply a kind of local-to-global principle to reduce the problem to studying the *size* of $\mathscr{L}(\mathbf{X}_k)$. To see roughly how this works, note that with significant probability there is some such j with $j > k/10$, in which case at least half of the sums $x_1 + 2x_2 + \cdots + kx_k$ will be at least $k/10$ (those with $x_j > 0$). At the same time it is reasonably likely that all *all* elements of $\mathscr{L}(\mathbf{X}_k)$ are at most $10k$. Assuming this heuristic is reasonable, we should expect that

$$i(n,k) \asymp \mathbf{P}(k \in \mathscr{L}(\mathbf{X}_k)) \asymp \frac{1}{k}\mathbf{E}|\mathscr{L}(\mathbf{X}_k)|.$$

This estimate the topic of Section 3.3, in which we establish the following proposition.

Proposition 3.1.4. $i(n,k) \asymp \frac{1}{k}\mathbf{E}|\mathscr{L}(\mathbf{X}_k)|$ *uniformly for* $1 \leqslant k \leqslant n/2$.

Our main theorem follows immediately from this and the next proposition, whose proof occupies Sections 3.4 and 3.5. We emphasize that we are now working just with the sequence $\mathbf{X}_k = (X_1, X_2, \ldots, X_k)$ of genuinely independent random variables, which is independent of n.

Proposition 3.1.5. $\mathbf{E}|\mathscr{L}(\mathbf{X}_k)| \asymp k^{1-\delta}(\log k)^{-3/2}$ *for* $k \geqslant 2$.

To briefly explain the origin of the exponent δ, first observe the simple inequality

$$|\mathscr{L}(\mathbf{X}_k)| \leqslant \min\left(2^{X_1+\cdots+X_k}, 1 + X_1 + 2X_2 + \cdots + kX_k\right). \tag{3.1}$$

Assume this is close to being sharp with reasonably high probability, and condition on $Y = X_1 + \cdots + X_k$. Following our earlier heuristic, the second term on the right side of (3.1) is $\asymp k$ most of the time, and so there is a change of behaviour around $Y = \log_2 k + O(1)$. Since Y is Poisson with parameter $\log k + O(1)$, a short calculation demonstrates that

$$\mathbf{E}\min(2^Y, k) \asymp k^{1-\delta}(\log k)^{-1/2}.$$

We err in the logarithmic term because (3.1) is only sharp with probability about $1/\log k$, a fact related to order statistics.

Let us finally mention two open questions.

Question 3.1.6. Is there some constant C such that as $k \to \infty$ we have

$$\lim_{n\to\infty} i(n,k) \sim Ck^{-\delta}(\log k)^{-3/2}?$$

It would be surprising if this were not the case, but establishing it seems to require a significantly more involved analysis.

Question 3.1.7. Is $\lim_{n\to\infty} i(n,k)$ monotonically decreasing in k?

Data collected by Britnell and Wildon [BW15] shows that this is so at least as far as $k = 30$, and of course a positive answer is plausible just from the fact that $\lim_{n\to\infty} i(n,k) \to 0$ as $k \to \infty$.

3.1.2 Invariable generation

Permutations $\pi_1, \ldots, \pi_r$ are said to have a property P *invariably* if $\pi_1', \ldots, \pi_r'$ have property P whenever π_i' is conjugate to π_i for each i. Thus $\pi_1, \ldots, \pi_r$ *invariably generate* $\mathcal{S}_n$ if $\pi_1', \ldots, \pi_r'$ generate $\mathcal{S}_n$ whenever π_i' is conjugate to π_i for each i.

By a theorem of Dixon [Dix69], if $\pi_1, \pi_2 \in \mathcal{S}_n$ are chosen uniformly at random then, with high probability (i.e., with probability tending to 1 as $n \to \infty$), the group they generate is either $\mathcal{S}_n$ or $\mathcal{A}_n$. Thus

$$\mathbf{P}(\langle \pi_1, \pi_2 \rangle = \mathcal{S}_n) \to 3/4,$$
$$\mathbf{P}(\langle \pi_1, \pi_2 \rangle = \mathcal{A}_n) \to 1/4,$$

as $n \to \infty$. But what about invariable generation? We cannot expect such a strong result, because given any constant r, with probability $i(n,1)^r \approx (1-1/e)^r$ each of $\pi_1, \dots, \pi_r$ has a fixed point, thus they have conjugates having the same fixed point, and thus they do not invariably generate. However it is reasonable to ask the following question.

Question 3.1.8. Does there exist a constant r such that $\pi_1, \dots, \pi_r$ invariably generate with probability bounded away from zero?

There is strong motivation for this question coming from computational Galois theory, and consequently the question has been studied by several different authors [Mus78; Dix92; ŁP93; DS00; KZ12; PPR14]. To briefly explain this connection, suppose we are given a polynomial $f \in \mathbf{Z}[x]$ of degree n with no repeated factors. Information about the Galois group can be gained by reducing f modulo random primes p and factorizing the reduced polynomial $\bar{f}$ over $\mathbf{Z}/p\mathbf{Z}$. By classical Galois theory, if $\bar{f}$ has irreducible factors of degrees $n_1, \dots, n_r$ then the Galois group G of f over $\mathbf{Q}$ has an element with cycle lengths $n_1, \dots, n_r$. Moreover by Frobenius's density theorem, if $G = \mathcal{S}_n$ then the frequency with which a given cycle type arises is equal to the proportion of elements in $\mathcal{S}_n$ with that cycle type. Thus if we suspect that $G = \mathcal{S}_n$ then the number of times we expect to have to iterate this procedure before being able to deduce that $G = \mathcal{S}_n$ is controlled by the expected number of random elements required to invariably generate $\mathcal{S}_n$.

Łuczak and Pyber [ŁP93] were the first to answer Question 3.1.8 positively. This follows from their bound $i(n,k) \ll k^{-0.01}$: Note by a union bound that the probability that $\pi_1, \dots, \pi_r$ each have a fixed set of the same size k for some k is bounded by

$$\sum_{k=1}^{n/2} i(n,k)^r \ll_r \sum_{k=1}^{n/2} k^{-0.01r},$$

so if r is sufficiently large then the probability that $\pi_1, \dots, \pi_r$ can be conjugated to each fix a common set is at most $1/10$, independent of n. In other words, $\pi_1, \dots, \pi_r$ invariably generate a transitive subgroup with probability at least $9/10$. Moreover, by

another theorem of Łuczak and Pyber (another consequence of their bound $i(n,k) \ll k^{-0.01}$), the union of all transitive subgroups besides $\mathcal{S}_n$ and $\mathcal{A}_n$ covers only $o(n!)$ of $\mathcal{S}_n$, so with probability $1 - o(1)$ our first permutation π_1 has no conjugate lying in a transitive subgroup smaller than $\mathcal{A}_n$. Finally, the probability that each of $\pi_1, \dots, \pi_r$ is contained in $\mathcal{A}_n$ is $(1/2)^r \leqslant 1/10$. Thus by ruling out all other possibilities we deduce that with probability at least $1/2$ the permutations $\pi_1, \dots, \pi_r$ invariably generate $\mathcal{S}_n$.

This proof does not readily yield a reasonable value of r, however, and in any case because $i(n,k) = k^{-\delta+o(1)}$ it seems unlikely to prove that any r smaller than 12 could suffice. However, Pemantle, Peres, and Rivin [PPR14] recently proved that we may take $r = 4$.

Theorem 3.1.9 (Pemantle–Peres–Rivin [PPR14])**.** *If $\pi_1, \pi_2, \pi_3, \pi_4 \in \mathcal{S}_n$ are chosen uniformly at random then the probability that $\pi_1, \pi_2, \pi_3, \pi_4$ invariably generate $\mathcal{S}_n$ is bounded away from zero.*

We give a somewhat simplified proof of this theorem in Section 3.6. Incidentally, Pemantle, Peres, and Rivin only prove that $\pi_1, \pi_2, \pi_3, \pi_4$ invariably generate a transitive subgroup of $\mathcal{S}_n$, but it is only a little more work to prove the theorem as stated above. Indeed, as in the Łuczak–Pyber proof, almost all $\pi_1 \in \mathcal{S}_n$ are contained in no transitive subgroup besides $\mathcal{S}_n$ or $\mathcal{A}_n$, so we need only worry about $\mathcal{A}_n$. Thus as long as we are careful in the proof to arrange that at least one of $\pi_1, \pi_2, \pi_3, \pi_4$ is odd, then no proper transitive subgroup will trap $\pi_1, \pi_2, \pi_3, \pi_4$. For full details see the proof of Theorem 3.1.9 in Section 3.6.

Our main contribution however is the lower bound $r > 3$, which will be proved in Section 3.7. Thus we finally have a complete answer: r can be taken as small as 4, but no smaller.

Theorem 3.1.10. *If $\pi_1, \pi_2, \pi_3 \in \mathcal{S}_n$ are chosen uniformly at random then the probability that π_1, π_2, π_3 invariably generate a transitive subgroup (or, in particular, all of $\mathcal{S}_n$) tends to zero as $n \to \infty$. Equivalently, with probability tending to 1 there is some k in the range $0 < k < n$ such that π_1, π_2, π_3 each have a fixed set of size k.*

As with estimating $i(n,k)$, these assertions are more or less equivalent to analogous assertions about our model $\mathscr{L}(\mathbf{X})$, where now it is more natural to consider the entire sequence $\mathbf{X}$ and the entire sumset

$$\mathscr{L}(\mathbf{X}) = \left\{ \sum_{j \geqslant 1} j x_j : 0 \leqslant x_j \leqslant X_j \text{ for each } j \right\}. \tag{3.2}$$

Thus unsurprisingly the main task in proving Theorem 3.1.9 is to show that

$$\mathscr{L}(\mathbf{X}) \cap \mathscr{L}(\mathbf{X}') \cap \mathscr{L}(\mathbf{X}'') \cap \mathscr{L}(\mathbf{X}''') = \{0\}$$

with positive probability, where $\mathbf{X}', \mathbf{X}'', \mathbf{X}'''$ are independent copies of $\mathbf{X}$. Similarly, excepting the details already mentioned to do with transitive subgroups, Theorem 3.1.10 follows almost immediately from the assertion that

$$\mathscr{L}(\mathbf{X}) \cap \mathscr{L}(\mathbf{X}') \cap \mathscr{L}(\mathbf{X}'')$$

is almost surely infinite.

In a sentence, the essential reason for the difference in behavior between the 3- and 4-fold intersections is that $\mathscr{L}(\mathbf{X}) \cap [k, 2k]$ typically has size $k^{\log 2 + o(1)}$, and

$$2/3 < \log 2 < 3/4.$$

3.1.3 The analogy with number theory

Let X be large, and choose an integer n uniformly at random from $\{1, \ldots, X\}$. Then for all small primes p, the probability that p divides n is close to $1/p$, and these events are approximately independent for different values of p. Thus if we fix an interval (e^t, e^{t+1}), with t much smaller than x, and ask how many primes $p \in (e^t, e^{t+1})$ divide n, then by the law of rare events the answer will be approximately Poisson with parameter

$$\begin{aligned}\sum_{p \in (e^t, e^{t+1})} 1/p &= \log\log(e^{t+1}) - \log\log(e^t) + O\left(1/t^2\right) \\ &= 1/t + O\left(1/t^2\right),\end{aligned}$$

these variables being approximately independent for disjoint intervals.

Compare this heuristic with Lemma 3.1.1, which asserts that if π is a random permutation then the number of cycles of length j in π is approximately Poisson with parameter $1/j$, these variables being approximately independent for different values of j. The similarity of these two heuristics is the foundation of a powerful analogy with analytic number theory, briefly outlined in Table 3.1.3.

Estimating the proportion $i(n, k)$ of permutations π having a fixed set of a given size k is thus roughly analogous to estimating the number $H(X, Y)$ of integers $n \leqslant X$ having a divisor in a given dyadic interval $(Y, 2Y]$. This is more or less equivalent to the well known *multiplication table problem*, which asks for the number of distinct

permutations	integers
$\pi \in \mathcal{S}_n$ (n large)	$n \leqslant X$ (X large)
cycles σ of π	primes $p \mid n$
length of σ	$\log p$
fixed sets A of π	divisors $d \mid n$
size of A	$\log d$

Table 3.1: The dictionary between permutations and integers

integers appearing in the standard $N \times N$ multiplication table, or in other words the quantity

$$A(N) = |\{n \leqslant N^2 : n = d_1 d_2 \text{ for some } d_1, d_2 \leqslant N\}|.$$

These problems have a long history, originating as a problem of Besicovitch [Bes34], and were solved up to constant factors only recently by Ford [For08b; For08a], who proved the following theorem.

Theorem 3.1.11 (Ford [For08b; For08a])**.** *Uniformly for $3 \leqslant Y \leqslant X^{1/2}$ we have*

$$H(X, Y) \asymp X(\log Y)^{-\delta}(\log\log Y)^{-3/2}.$$

Thus for $N \geqslant 3$ we have

$$A(N) \asymp N^2(\log N)^{-\delta}(\log\log N)^{-3/2}.$$

The similarity between Theorems 3.1.2 and 3.1.11 should now not be surprising. Unfortunately the analogy is not so perfect that it is possible to easily deduce either theorem from the other, but the reader will not be surprised to hear that our proof of Theorem 3.1.2 strongly parallels Ford's proof of Theorem 3.1.11. Even before the appearance of Theorem 3.1.2 this connection was noted by Pemantle, Peres, and Rivin [PPR14], as well as by Diaconis and Soundararajan (see [Sou13, page 14]).

Similarly, the probability that $\pi_1, \dots, \pi_r$ have fixed sets of the same size is roughly analogous the probability that $n_1, \dots, n_r \leqslant X$ have divisors $d_1, \dots, d_r$ all within a small factor of one another. This is reminiscent of the following well known theorem of Maier and Tenenbaum [MT84].

Theorem 3.1.12 (Maier–Tenenbaum [MT84])**.** *A random integer $n \leqslant X$ has, with probability tending to 1 as $X \to \infty$, two distinct divisors d_1, d_2 such that $d_1 < d_2 \leqslant 2d_1$.*

The corresponding assertion for permutations is similar in spirit to what we want: it would assert that with high probability a single permutation π has for some k with

$0 < k < n$ at least two distinct fixed sets of size k (a true statement, but not one we establish here). We prove Theorem 3.1.10 by building on the method of Maier and Tenenbaum.

The analogy between permutations and integers is sufficiently strong that in [EFG15a] we made the following conjecture, which is analogous to Theorems 3.1.9 and 3.1.10, and we said we would return to it in a future paper.

Conjecture 3.1.13. Choose n_1, n_2, n_3, n_4 at random from $1, \dots, X$. Then with probability tending to 1 as $X \to \infty$ we can find

$$d_1 | n_1, d_2 | n_2, d_3 | n_3$$

such that $\max d_i \leqslant 1.001 \min d_i$, but with probability bounded away from zero we cannot find

$$d_1 | n_1, d_2 | n_2, d_3 | n_3, d_4 | n_4$$

such that $\max d_i \leqslant 1000 \min d_i$.

Following the submission of our paper it was brought to our attention that this has actually already been proved! See Raouj and Stef [RS99].

The analogy we are exploting between permutations and integers has been written about by many authors, most notably by Andrew Granville, who even wrote a play about it with his sister Jennifer Granville, called *The Anatomy of Integers and Permutations*: see [GG09; Gra09].

3.2 A permutation sieve

As mentioned in the introduction, the asymptotic distribution (as $n \to \infty$ with k fixed) of the cycle lengths $(c_1(\pi), \dots, c_k(\pi))$ of a random $\pi \in \mathcal{S}_n$ is that of $\mathbf{X}_k = (X_1, \dots, X_k)$, where the X_i are independent with $X_i \stackrel{d}{=} \mathrm{Pois}(1/i)$. In the nonasymptotic regime, where n may be as small as $2k$, this property is lost. We do, however, have the following substitute which will suffice for this paper.

Proposition 3.2.1. *Let $1 \leqslant m < n$ and $c_1, \dots, c_m$ be nonnegative integers satisfying*

$$c_1 + 2c_2 + \cdots + mc_m \leqslant n - m - 1.$$

Suppose that $\pi \in \mathcal{S}_n$ is chosen uniformly at random. Then

$$\frac{1}{(2m+2)\prod_{i=1}^m c_i! i^{c_i}} \leqslant \mathbf{P}(c_1(\pi) = c_1, \dots, c_m(\pi) = c_m) \leqslant \frac{1}{(m+1)\prod_{i=1}^m c_i! i^{c_i}}.$$

We will prove this shortly, but first let us fix some notation. As every permutation $\pi \in \mathcal{S}_n$ factors uniquely as a product of disjoint cycles, in keeping with the analogy with integers we say that any product of these cycles, including the empty product, is a *factor* or *divisor* of π. We may think of the factors as partially defined permutations. The domains of these factors are precisely the invariant sets of π. We make the following further definitions:

- $\mathcal{C}_{k,n}$ is the set of cycles of length k on $\{1, \dots, n\}$;
- $|\sigma|$ is the length of any factor σ (of some permutation in $\mathcal{S}_n$);
- $\tau | \pi$ means that τ is an invariant set or divisor of π.

The following lemma is a slight generalization of the well known formula of Cauchy.

Lemma 3.2.2. *Let $1 \leqslant m \leqslant n$, and let $c_1, \dots, c_m$ be nonnegative integers with $t = c_1 + 2c_2 + \cdots + mc_m \leqslant n$. Then the number of ways of choosing $c_1 + \cdots + c_m$ disjoint cycles consisting of c_i cycles in $\mathcal{C}_{i,n}$ for $1 \leqslant i \leqslant m$ is*

$$\frac{n!}{(n-t)!} \prod_{j=1}^{m} \frac{1}{c_j! j^{c_j}}.$$

Proof. First count the number of ways of choosing the subsets that make up the cycles, and then multiply by the number of ways to arrange the elements of these subsets into cycles. The result is

$$\binom{n}{\underbrace{1 \cdots 1}_{c_1} \underbrace{2 \cdots 2}_{c_2} \cdots \underbrace{m \cdots m}_{c_m} n-t} \frac{1}{c_1! \cdots c_m!} \times \prod_{j=1}^{m} (j-1)!^{c_j},$$

which simplifies to the claimed expression. □

Our next lemma is an analogue for permutations of a basic lemma from sieve theory.

Lemma 3.2.3. *Suppose that m, n are integers with $1 \leqslant m \leqslant n$. Let $\pi \in \mathcal{S}_n$ be chosen uniformly at random. Then*

$$\frac{1}{2m} \leqslant \mathbf{P}(\pi \text{ has no cycle of length} < m) \leqslant \frac{1}{m}.$$

Remark. Both upper and lower bounds are best possible, since trivially the probability in question is exactly $1/n$ when $n/2 < m \leqslant n$ (if a permutation has no cycle of length $< m$, with m in this range, then it must be an n-cycle). In fact, it is not difficult to prove an asymptotic formula $\sim \omega(n/m)/m$ ($n \to \infty$, $m \to \infty$, $m \leqslant n$) for the probability in question, where ω is Buchstab's function and $\omega(u) \to e^{-\gamma}$ as $u \to \infty$: see Granville [Gra06, Theorem 2.2].

Proof. (See the proof of [Gra06, Theorem 2.2]). We phrase the proof combinatorially rather than probabilistically; thus let $c(n, m)$ be the number of permutations of $\mathcal{S}_n$ that have no cycles of length $< m$. We proceed by induction on n, the result being trivial when $n = 1$. Let $\sum^*$ denote a sum over permutations with no cycle of length $< m$. Using the fact that the sum of lengths of cycles in a permutation in $\mathcal{S}_n$ is n, we get

$$\begin{aligned} nc(n,m) &= \sum_{\pi\in\mathcal{S}_n}{}^* n = \sum_{\pi\in\mathcal{S}_n}{}^* \sum_{\substack{\sigma|\pi \\ \sigma \text{ a cycle}}} |\sigma| = \sum_{k\geqslant m} k \sum_{\sigma\in\mathcal{C}_{k,n}} \sum_{\substack{\pi\in\mathcal{S}_n \\ \sigma|\pi}}{}^* 1 \\ &= \sum_{\sigma\in\mathcal{C}_{n,n}} n + \sum_{m\leqslant k\leqslant n-m} k \sum_{\sigma\in\mathcal{C}_{k,n}} c(n-k,m) \\ &= n! + \sum_{m\leqslant k\leqslant n-m} \frac{n!}{(n-k)!} c(n-k,m). \end{aligned}$$

If $\frac{n}{2} < m \leqslant n$, then $c(n, m) = \frac{n!}{n}$ and the result follows. Otherwise, by the induction hypothesis,

$$nc(n,m) \leqslant n! + \sum_{m\leqslant k\leqslant n-m} \frac{n!}{m} = n!\left(1 + \frac{n-2m+1}{m}\right) \leqslant \frac{n!\cdot n}{m}$$

and

$$nc(n,m) \geqslant n! + \sum_{m\leqslant k\leqslant n-m} \frac{n!}{2m} = n!\left(1 + \frac{n-2m+1}{2m}\right) \geqslant \frac{n!\cdot n}{2m}. \qquad \square$$

It is now a simple matter to establish Proposition 3.2.1.

Proof of Proposition 3.2.1. Let $t = c_1 + 2c_2 + \cdots + mc_m$. For each choice of the $c_1 + \cdots + c_m$ disjoint cycles consisting of c_j cycles from $\mathcal{C}_{j,n}$ ($1 \leqslant j \leqslant m$), there are $c(n - t, m + 1)$ permutations $\pi \in \mathcal{S}_n$ containing these cycles as factors and no other cycles of length at most m, where $c(n - t, m + 1)$ is the number of permutations on $n-t$ letters with no cycle of length $< m+1$, as in the proof of Lemma 3.2.3. Applying Lemmas 3.2.2 and 3.2.3 completes the proof. $\square$

Later in this chapter we will need another lemma involving the same sort of argument. This time we need to count permutations with a given number of cycles of length at most k. By the Poisson model Lemma 3.1.1, if k is fixed and $n \to \infty$, this statistic has distribution approaching that of a Poisson variable with parameter $h_k = 1 + \frac{1}{2} + \cdots + \frac{1}{k}$. The next result tells us that the distribution is still approximately Poisson uniformly over all choices of the parameters k and n.

We record here the basic inequalities

$$\log(k+1) \leqslant h_k \leqslant 1 + \log k, \qquad (k \geqslant 1) \tag{3.3}$$

which follow from the obvious inequalities $\frac{1}{n+1} \leqslant \int_n^{n+1} dt/t \leqslant \frac{1}{n}$ by taking a sum.

Lemma 3.2.4. *Let n, k, ℓ be integers with $1 \leqslant k \leqslant n$ and $\ell \geqslant 0$. Select $\pi \in \mathcal{S}_n$ at random. Then*

$$\mathbf{P}(\pi \text{ has exactly } \ell \text{ cycles with length } \leqslant k) \leqslant \frac{e}{k} \frac{h_k^\ell}{\ell!} \left(1 + \frac{\ell}{h_k}\right).$$

In particular if $\ell \ll \log k$ then this is $O\left(k^{-1} h_k^\ell / \ell!\right)$, while if $\ell \gg \log k$ then this is $O\left(k^{-1} h_k^{\ell-1}/(\ell-1)!\right)$.

Proof. Denote by $\mathcal{S}_n(k, \ell)$ the set of $\pi \in \mathcal{S}_n$ containing exactly ℓ cycles of length at most k. Evidently

$$n|\mathcal{S}_n(k,\ell)| = \sum_{\pi \in \mathcal{S}_n(k,\ell)} \sum_{\substack{\sigma | \pi \\ \sigma \text{ a cycle}}} |\sigma|.$$

Write $\pi = \sigma\pi'$, and observe that π' has either $\ell - 1$ or ℓ cycles of length at most k, depending on whether $|\sigma| \leqslant k$ or not. Thus $\pi' \in S_{n-|\sigma|}(k, m)$, where $m = \ell - 1$ or $m = \ell$, so

$$\begin{aligned} n|\mathcal{S}_n(k,\ell)| &\leqslant \sum_{h=1}^{n} \sum_{m=\ell-1}^{\ell} \sum_{\pi' \in \mathcal{S}_{n-h}(k,m)} |\mathcal{C}_{h,n}| h \\ &= \sum_{h=1}^{n} \sum_{m=\ell-1}^{\ell} \sum_{\pi' \in \mathcal{S}_{n-h}(k,m)} \frac{n!}{(n-h)!}. \end{aligned}$$

Now rearrange the sum according the cycle type $\mathbf{c} = (c_1, \ldots, c_n)$ of the permutation π', i.e., suppose π' has c_j cycles of length j for each j. By Cauchy's formula the

number of $\pi' \in \mathcal{S}_{n-h}$ with the cycle type $\mathbf{c}$ is $(n-h)!/\prod_j c_j! j^{c_j}$. It follows that

$$\begin{aligned}
n|\mathcal{S}_n(k,\ell)| &\leqslant n! \sum_{h=1}^{n} \sum_{\substack{c_1,\dots,c_n \geqslant 0 \\ c_1+2c_2+\cdots+nc_n=n-h \\ c_1+\cdots+c_k \in \{\ell-1,\ell\}}} \frac{1}{\prod_j c_j! j^{c_j}} \\
&\leqslant n! \sum_{\substack{c_1,\dots,c_n \geqslant 0 \\ c_1+\cdots+c_k \in \{\ell-1,\ell\}}} \frac{1}{\prod_j c_j! j^{c_j}} \\
&= n! \left(\sum_{\substack{c_1,\dots,c_k \geqslant 0 \\ c_1+\cdots+c_k=\ell-1}} \frac{1}{\prod_j c_j! j^{c_j}} + \sum_{\substack{c_1,\dots,c_k \geqslant 0 \\ c_1+\cdots+c_k=\ell}} \frac{1}{\prod_j c_j! j^{c_j}} \right) \\
&\quad \times \sum_{c_{k+1},\dots,c_n \geqslant 0} \frac{1}{\prod_{j=k+1}^{n} c_j! j^{c_j}} \\
&= n! \left(\frac{h_k^{\ell-1}}{(\ell-1)!} + \frac{h_k^{\ell}}{\ell!} \right) \prod_{k<i\leqslant n} e^{1/i},
\end{aligned}$$

where in the last line we used the multinomial theorem. The claimed bound now follows from

$$\sum_{k<i\leqslant n} \frac{1}{i} = h_n - h_k \leqslant \log n - \log k + 1. \qquad \square$$

3.3 The local-to-global principle

Our aim in this section is to prove our local-to-global principle Proposition 3.1.4. As in the introduction, let $X_1, X_2, \dots$ be independent random variables with distribution $X_j \stackrel{d}{=} \mathrm{Pois}(1/j)$. We record here the basic equality

$$\begin{aligned}
\mathbf{E}|\mathscr{L}(\mathbf{X}_k)| &= \sum_{c_1,\dots,c_k \geqslant 0} |\mathscr{L}(\mathbf{c})| \mathbf{P}(X_1 = c_1) \cdots \mathbf{P}(X_k = c_k) \\
&= e^{-h_k} \sum_{c_1,\dots,c_k \geqslant 0} \frac{|\mathscr{L}(\mathbf{c})|}{\prod_{i=1}^{k} c_i! i^{c_i}}. \qquad (3.4)
\end{aligned}$$

Lemma 3.3.1. *Let $k \in \mathbf{N}$, $c_1, \dots, c_k \geqslant 0$, $I \subset \{1, \dots, k\}$ and $c_i' = c_i$ for $i \notin I$, $c_i' = 0$ for $i \in I$. then*

$$|\mathscr{L}(\mathbf{c})| \leqslant |\mathscr{L}(\mathbf{c}')| \prod_{i \in I} (c_i + 1).$$

Proof. Clearly $\mathscr{L}(\mathbf{c})$ is the union of $\prod_{i\in I}(c_i+1)$ translates of $\mathscr{L}(\mathbf{c}')$. $\square$

Lemma 3.3.2. *Suppose that $\ell' \leqslant \ell$. Then*

$$\frac{1}{\ell} \mathbf{E}|\mathscr{L}(\mathbf{X}_\ell)| \leqslant \frac{1}{\ell'} \mathbf{E}|\mathscr{L}(\mathbf{X}_{\ell'})|.$$

Proof. By Lemma 3.3.1, $|\mathscr{L}(\mathbf{X}_\ell)| \leqslant (1+X_{\ell'+1})\cdots(1+X_\ell)|\mathscr{L}(\mathbf{X}_{\ell'})|$. Thus by independence,

$$\mathbf{E}|\mathscr{L}(\mathbf{X}_\ell)| \leqslant \left(\prod_{i=\ell'+1}^{\ell} \mathbf{E}(1+X_i)\right)\mathbf{E}|\mathscr{L}(\mathbf{X}_{\ell'})| = \frac{\ell+1}{\ell'+1}\mathbf{E}|\mathscr{L}(\mathbf{X}_{\ell'})| \leqslant \frac{\ell}{\ell'}\mathbf{E}|\mathscr{L}(\mathbf{X}_{\ell'})|. \quad \square$$

We also need to compute the mixed moments of $|\mathscr{L}(\mathbf{X}_k)|$ with powers of some X_j. Write B_m for the mth moment $\mathbf{E}X^m$, if $X \stackrel{d}{=} \mathrm{Pois}(1)$ (B_m is the mth Bell number).

Lemma 3.3.3. *Suppose that $j_1, \ldots, j_h \leqslant k$ are distinct integers and that $a_1, \ldots, a_h$ are positive integers. Then*

$$\mathbf{E}|\mathscr{L}(\mathbf{X}_k)|X_{j_1}^{a_1}\cdots X_{j_h}^{a_h} \leqslant \frac{C_{a_1,\ldots,a_h}}{j_1\ldots j_h}\mathbf{E}|\mathscr{L}(\mathbf{X}_k)|.$$

We may take $C_{a_1,\ldots,a_h} = \prod_{i=1}^{h}(B_{a_i}+B_{a_i+1})$. In particular we may take $C_1 = 3$.

Proof. Define $\mathbf{X}'_k$ by putting $X'_{j_1} = \cdots = X'_{j_h} = 0$ and $X'_j = X_j$ for all other j. By Lemma 3.3.1, we have

$$|\mathscr{L}(\mathbf{X}_k)| \leqslant |\mathscr{L}(\mathbf{X}'_k)|(1+X_{j_1})\cdots(1+X_{j_h}).$$

Thus by independence

$$\mathbf{E}|\mathscr{L}(\mathbf{X}_k)|X_{j_1}^{a_1}\cdots X_{j_h}^{a_h} \leqslant \mathbf{E}|\mathscr{L}(\mathbf{X}'_k)|\prod_{i=1}^{h}(\mathbf{E}X_{j_i}^{a_i}+\mathbf{E}X_{j_i}^{a_{i+1}}). \tag{3.5}$$

For $X \stackrel{d}{=} \mathrm{Pois}(\lambda)$ we have $\mathbf{E}X^m = \phi_m(\lambda)$, where $\phi_m(\lambda)$ is the mth Touchard (or Bell) polynomial, a polynomial with positive coefficients and zero constant coefficient. If $\lambda \leqslant 1$, it follows that $\mathbf{E}X^m \leqslant \lambda B_m$ for $m \geqslant 1$. The result follows immediately from this, (3.5), and the trivial observation that $\mathbf{E}|\mathscr{L}(\mathbf{X}'_k)| \leqslant \mathbf{E}|\mathscr{L}(\mathbf{X}_k)|$. $\square$

We turn now to the proof of Proposition 3.1.4. In what follows, if $\mathbf{c} = (c_1, \ldots, c_\ell)$, we write

$$S(\mathbf{c}) = c_1 + 2c_2 + \cdots + \ell c_\ell = \max \mathscr{L}(\mathbf{c}).$$

We will treat the lower bound and upper bound in Proposition 3.1.4 separately, the former being somewhat more straightforward than the latter.

3.3.1 The lower bound in Proposition 3.1.4

If $k < 40$ then for trivial reasons

$$i(n,k) \asymp 1 \asymp \frac{1}{k}\mathbf{E}|\mathscr{L}(\mathbf{X}_k)|,$$

so we may assume $k \geqslant 40$. Let $r = \lfloor k/20 \rfloor$ (so $r \geqslant 2$), and consider the permutations

$$\pi = \alpha\sigma_1\sigma_2\beta \in \mathcal{S}_n,$$

where σ_1 and σ_2 are cycles, $|\alpha| \leqslant 4r < |\sigma_1| < |\sigma_2| < 16r$, all cycles in α have length $\leqslant r$, all cycles in β have length at least $16r$, and $\alpha\sigma_1\sigma_2$ has a fixed set of size k. Because of the size restrictions on $\alpha, \sigma_1, \sigma_2$, if α is of type $\mathbf{c} = (c_1, \ldots, c_r)$, with c_i cycles of length i for $1 \leqslant i \leqslant r$, then the last condition is equivalent to

$$k - |\sigma_1| - |\sigma_2| \in \mathscr{L}(\mathbf{c}).$$

In particular $|\sigma_1| + |\sigma_2| \leqslant k$, and hence

$$n - |\alpha| - |\sigma_1| - |\sigma_2| \geqslant \frac{4}{5}k \geqslant 16r.$$

Fix $\mathbf{c}$ and ℓ_1, ℓ_2 with $4r < \ell_1 < \ell_2 < 16r$ such that $k - \ell_1 - \ell_2 \in \mathscr{L}(\mathbf{c})$. By Proposition 3.2.1, the probability that a random $\pi \in \mathcal{S}_n$ has c_i cycles of length i ($1 \leqslant i \leqslant r$), one cycle each of length ℓ_1, ℓ_2 and no other cycles of length $< 16r$ is at least

$$\frac{1}{32r\ell_1\ell_2 \prod_{i=1}^r c_i! i^{c_i}} \geqslant \frac{1}{2^{13}r^3 \prod_{i=1}^r c_i! i^{c_i}}.$$

For any ℓ_1 satisfying $4r + 1 \leqslant \ell_1 \leqslant 8r - 1$, there are $|\mathscr{L}(\mathbf{c})|$ admissible values of $\ell_2 > \ell_1$ for which $k - \ell_1 - \ell_2 \in \mathscr{L}(\mathbf{c})$, since $\max \mathscr{L}(\mathbf{c}) \leqslant 4r \leqslant k/5$. We conclude that

$$i(n,k) \gg \frac{4r-1}{2^{13}r^3} \sum_{\substack{c_1,\cdots,c_r \geqslant 0 \\ S(\mathbf{c}) \leqslant 4r}} \frac{|\mathscr{L}(\mathbf{c})|}{\prod_{i=1}^r c_i! i^{c_i}}.$$

As in (3.4), the sum above equals $e^{h_r}\mathbf{E}|\mathscr{L}(\mathbf{X}_r)|1_{S(\mathbf{X}_r)\leqslant 4r}$. Hence by (3.3) we have

$$i(n,k) \gg \frac{1}{r}\mathbf{E}|\mathscr{L}(\mathbf{X}_r)|1_{S(\mathbf{X}_r)\leqslant 4r}.$$

To estimate this, we use the inequality

$$1_{S(\mathbf{X}_r)\leqslant 4r} \geqslant 1 - \frac{S(\mathbf{X}_r)}{4r}.$$

By Lemma 3.3.3 we have

$$\mathbf{E}|\mathscr{L}(\mathbf{X}_r)|S(\mathbf{X}_r) = \sum_{j=1}^{r} \mathbf{E}|\mathscr{L}(\mathbf{X}_r)|jX_j \leqslant 3r\mathbf{E}|\mathscr{L}(\mathbf{X}_r)|.$$

It follows that

$$\mathbf{E}|\mathscr{L}(\mathbf{X}_r)|1_{S(\mathbf{X}_r)\leqslant 4r} \geqslant \mathbf{E}|\mathscr{L}(\mathbf{X}_r)|\left(1 - \frac{S(\mathbf{X}_r)}{4r}\right) \geqslant \frac{1}{4}\mathbf{E}|\mathscr{L}(\mathbf{X}_r)|,$$

and hence that

$$i(n,k) \gg \frac{1}{r}\mathbf{E}|\mathscr{L}(\mathbf{X}_r)|.$$

The lower bound in Proposition 3.1.4 now follows from Lemma 3.3.2.

Remark. Strictly for the purposes of proving our main theorem, the final appeal to Lemma 3.3.2 in the above proof is unnecessary. However, that lemma is straightforward and it is more aesthetically pleasing to have $\mathbf{E}|\mathscr{L}(\mathbf{X}_k)|$ in the lower bound for $i(n,k)$ rather than $\mathbf{E}|\mathscr{L}(\mathbf{X}_r)|$.

3.3.2 The upper bound in Proposition 3.1.4

Suppose $\pi \in \mathcal{S}_n$ has an invariant set of size k. Let $k_1 = k$ and $k_2 = n - k$. Then $\pi = \pi_1\pi_2$, where π_j is a product of cycles which, all together, have total length k_j, for $j = 1, 2$. For each j let σ_j be the longest cycle of π_j (or one of them if there are ties). Fix $j \in \{1, 2\}$ so that $|\sigma_j| \leqslant |\sigma_{3-j}|$, and let $\ell = |\sigma_j|$. Note then that

$$\ell \leqslant \min(k_1, k_2) = k.$$

Now decompose π as

$$\pi = \alpha\sigma_1\sigma_2\beta,$$

where α is the product of all cycles of π of length at most ℓ other than σ_1, σ_2 and β is a (possibly empty) product of cycles of length greater than ℓ.

By definition of σ_1 and α, $\alpha\sigma_j$ has a divisor of size k_j. Thus if α has type $\mathbf{c} = (c_1, c_2, \dots, c_\ell)$ then $k_j - \ell \in \mathscr{L}(\mathbf{c})$. For ℓ and $\mathbf{c}$ satisfying this condition, the number of choices for $\alpha\sigma_1$ is at most (by Lemma 3.2.2)

$$\frac{n!}{(n - S(\mathbf{c}) - \ell)!}\prod_{i<\ell}\frac{1}{c_i!i^{c_i}} \times \frac{1}{(c_\ell + 1)!\ell^{c_\ell+1}} \leqslant \frac{n!}{\ell(n - S(\mathbf{c}) - \ell)!}\prod_{i\leqslant\ell}\frac{1}{c_i!i^{c_i}}.$$

Having fixed $\alpha\sigma_1$, since $n - |\alpha\sigma_1| = |\sigma_2\beta| \geqslant \ell$, Lemma 3.2.3 implies that the number of choices for $\sigma_2\beta$ is at most $(n - |\alpha\sigma_1|)!/\ell$. Thus the number of choices for $\pi = \alpha\sigma_1\sigma_2\beta$ is at most

$$\frac{n!}{\ell^2}\prod_{i\leqslant\ell}\frac{1}{c_i!i^{c_i}}.$$

We deduce that

$$i(n,k) \leqslant \sum_{j=1}^{2}\sum_{\ell=1}^{k} \sum_{\substack{c_1,\dots,c_\ell \geqslant 0 \\ k_j-\ell \in \mathscr{L}(\mathbf{c})}} \frac{1}{\ell^2} \prod_{i\leqslant \ell} \frac{1}{c_i! i^{c_i}} = \sum_{j=1}^{2} \sum_{c_1,\dots,c_k\geqslant 0} \prod_{i\leqslant k} \frac{1}{c_i! i^{c_i}} \sum_{\substack{m(\mathbf{c})\leqslant \ell \leqslant k \\ k_j-\ell\in\mathscr{L}(\mathbf{c})}} \frac{1}{\ell^2},$$

where $m(\mathbf{c}) = \max\{i : c_i > 0\} \cup \{1\}$. With $\mathbf{c}$ fixed, note that $\ell \geqslant \max(m(\mathbf{c}), k_j - S(\mathbf{c}))$. Also, the number of ℓ such that $k_j - \ell \in \mathscr{L}(\mathbf{c})$ is at most $|\mathscr{L}(\mathbf{c})|$. Thus, the innermost sum on the right side above is at most

$$\frac{|\mathscr{L}(\mathbf{c})|}{\max(m(\mathbf{c}), k_j - S(\mathbf{c}))^2} \leqslant \frac{|\mathscr{L}(\mathbf{c})|}{\max(m(\mathbf{c}), k - S(\mathbf{c}))^2}.$$

Thus as in (3.4), using (3.3) we thus see that

$$i(n,k) \leqslant 2ek\,\mathbf{E}\,\frac{|\mathscr{L}(\mathbf{X}_k)|}{\max(m(\mathbf{X}_k), k - S(\mathbf{X}_k))^2}. \tag{3.6}$$

To bound this we use the inequality

$$\frac{1}{\max(m, k-S)^2} \leqslant \frac{4}{k^2}\left(1 + \frac{S^2}{m^2}\right),$$

which can be checked in the cases $S \geqslant k/2$ and $S \leqslant k/2$ separately. It follows from this and (3.6) that

$$i(n,k) \leqslant 8e\frac{1}{k}\mathbf{E}|\mathscr{L}(\mathbf{X}_k)| + 8e\frac{1}{k}\mathbf{E}\frac{|\mathscr{L}(\mathbf{X}_k)|S(\mathbf{X}_k)^2}{m(\mathbf{X}_k)^2}. \tag{3.7}$$

The first of these two terms is what we want, but the second requires a keener analysis. By conditioning on $m = m(\mathbf{X}_k)$ we have

$$\begin{aligned}
\mathbf{E}\frac{|\mathscr{L}(\mathbf{X}_k)|S(\mathbf{X}_k)^2}{m(\mathbf{X}_k)^2} &= \sum_{m=1}^{k} \frac{1}{m^2} \sum_{\substack{c_1,\dots,c_m\geqslant 0 \\ c_m \geqslant 1}} |\mathscr{L}(\mathbf{c})|S(\mathbf{c})^2 \mathbf{P}(\mathbf{X}_k = (c_1,\dots,c_m,0,\dots,0)) \\
&= \sum_{m=1}^{k} \frac{1}{m^2}\mathbf{E}|\mathscr{L}(\mathbf{X}_m)|S(\mathbf{X}_m)^2 1_{X_m\geqslant 1} \exp\left(-\sum_{j=m+1}^{k}\frac{1}{j}\right) \\
&\leqslant \frac{e}{k}\sum_{m=1}^{k}\frac{1}{m}\mathbf{E}|\mathscr{L}(\mathbf{X}_m)|S(\mathbf{X}_m)^2 X_m,
\end{aligned}$$

where in the last step we used the crude inequality $1_{X_m\geqslant 1} \leqslant X_m$. Expanding $S(\mathbf{X}_m)^2 = (X_1 + 2X_2 + \cdots + mX_m)^2$ and using (3.7), we thus arrive at

$$i(n,k) \ll \frac{1}{k}\mathbf{E}|\mathscr{L}(\mathbf{X}_k)| + \frac{1}{k^2}\sum_{m=1}^{k}\frac{1}{m}\sum_{i,i'=1}^{m} ii'\mathbf{E}|\mathscr{L}(\mathbf{X}_m)|X_iX_{i'}X_m. \tag{3.8}$$

The innermost sum is estimated using Lemma 3.3.3, splitting into various cases depending on the set of distinct values among i, i', m. Write $\mathbf{Y}_m = |\mathscr{L}(\mathbf{X}_m)|$ for brevity.

Case 1 i, i', m all distinct. Then $ii'\mathbf{EY}_m X_i X_{i'} X_m \leqslant C_{1,1,1} m^{-1} \mathbf{EY}_m = \frac{27}{m}\mathbf{EY}_m$.

Case 2 $i = i' \neq m$. Then $ii'\mathbf{EY}_m X_i X_{i'} X_m \leqslant C_{1,2} i m^{-1} \mathbf{EY}_m \leqslant C_{1,2}\mathbf{EY}_m = 21\mathbf{EY}_m$.

Case 3 $i = i' = m$. Then $ii'\mathbf{EY}_m X_i X_{i'} X_m \leqslant C_3 m \mathbf{EY}_m = 20m\mathbf{EY}_m$.

Case 4 $i \neq i' = m$ or $i' \neq i = m$. In both cases $ii'\mathbf{EY}_m X_i X_{i'} X_m \leqslant 21\mathbf{EY}_m$.

By taking a sum over all cases we find that

$$\sum_{i,i'=1}^{m} ii'\mathbf{EY}_m X_i X_{i'} X_m \ll m\mathbf{EY}_m.$$

Since clearly $\mathbf{EY}_m \leqslant \mathbf{EY}_k$ for every $m \leqslant k$ the result follows from this and (3.8).

3.4 The lower bound in Proposition 3.1.5

In this section we prove the lower bound in Proposition 3.1.5, and hence the lower bound in Theorem 3.1.2. We begin by noting that from (3.4) and (3.3) follows

$$\mathbf{E}|\mathscr{L}(\mathbf{X}_k)| \geqslant \frac{1}{ek} \sum_{c_1,\dots,c_k \geqslant 0} \frac{|\mathscr{L}(\mathbf{c})|}{\prod_{i=1}^k c_i! i^{c_i}}. \tag{3.9}$$

If we fix $r = c_1 + \cdots + c_k$, which we may think of as the number of cycles in a random permutation, then

$$\sum_{c_1+\cdots+c_k=r} \frac{|\mathscr{L}(\mathbf{c})|}{\prod_{i=1}^k c_i! i^{c_i}} = \frac{1}{r!} \sum_{a_1,\dots,a_r=1}^{k} \frac{|\mathscr{L}^*(\mathbf{a})|}{a_1 \cdots a_r}, \tag{3.10}$$

where

$$\mathscr{L}^*(\mathbf{a}) = \left\{ \sum_{i \in I} a_i : I \subset \{1, \dots, r\} \right\}. \tag{3.11}$$

The equality is most easily seen by starting from the right side and setting

$$c_i = |\{j : a_j = i\}|$$

for each i: then $\mathscr{L}(\mathbf{c}) = \mathscr{L}^*(\mathbf{a})$, $\prod_{i=1}^k i^{c_i} = a_1 \cdots a_k$, and each $\mathbf{c} = (c_1, \dots, c_k)$ comes from $\frac{r!}{c_1! \cdots c_k!}$ different choices of $a_1, \dots, a_k$. If one thinks of c_j as the number of cycles of length j in a random permutation, then one may think of $a_1, \dots, a_r$ as the unordered cycle lengths, in this case conditioned so that there are r total cycles.

Now let $J = \left\lfloor \frac{\log k}{\log 2} \right\rfloor$ and suppose that $b_1, \dots, b_J$ are arbitrary nonnegative integers with sum r. Consider the part of the sum in which

$$b_i = \sum_{j=2^{i-1}}^{2^i-1} c_j \ \ (i = 1, 2, \dots, J), \qquad c_j = 0 \ (j > 2^J - 1).$$

Equivalently, suppose there are exactly b_i of the a_j in each interval $[2^{i-1}, 2^i - 1]$. Writing $\mathscr{D}(\mathbf{b}) = \prod_{i=1}^{J} \{2^{i-1}, \dots, 2^i - 1\}^{b_i}$, we have

$$\frac{1}{r!} \sum_{a_1,\dots,a_r=1}^{2^J-1} \frac{|\mathscr{L}^*(\mathbf{a})|}{a_1 \cdots a_r} = \sum_{b_1,\dots,b_J} \frac{1}{b_1! \cdots b_J!} \sum_{\mathbf{d} \in \mathscr{D}(\mathbf{b})} \frac{|\mathscr{L}^*(\mathbf{d})|}{d_1 \cdots d_r}. \tag{3.12}$$

To see this, fix $b_1, \dots, b_J$ and observe that there are $\frac{r!}{b_1! \cdots b_J!}$ ways to choose which b_i of the variables $a_1, \dots, a_r$ lie in $[2^{i-1}, 2^i - 1]$ for $1 \leqslant i \leqslant J$.

Combining (3.9), (3.10) and (3.12) gives

$$\mathbf{E}|\mathscr{L}(\mathbf{X}_k)| \gg \frac{1}{k} \sum_{r} \sum_{b_1 + \dots + b_J = r} \frac{1}{b_1! \cdots b_J!} \sum_{\mathbf{d} \in \mathscr{D}(\mathbf{b})} \frac{|\mathscr{L}^*(\mathbf{d})|}{d_1 \cdots d_r}. \tag{3.13}$$

In the light of this, the motivation for the following lemma is clear.

Lemma 3.4.1. *For any $\mathbf{b} = (b_1, \dots, b_J)$ with $b_1 + \dots + b_J = r$ we have*

$$\sum_{\mathbf{d} \in \mathscr{D}(\mathbf{b})} \frac{|\mathscr{L}^*(\mathbf{d})|}{d_1 \cdots d_r} \gg \frac{(2 \log 2)^r}{\sum_{i=1}^{J} 2^{b_1 + \dots + b_i - i}}.$$

Proof. Given $\ell \in \mathbf{N}$, let $R(\mathbf{d}, \ell)$ be the number of $I \subset \{1, \dots, r\}$ with $\ell = \sum_{i \in I} d_i$ (one can think of this as the number of sets of cycles whose lengths add up to precisely ℓ). Then $\sum_\ell R(\mathbf{d}, \ell) = 2^r$. Also, define

$$\lambda_i = \sum_{j=2^{i-1}}^{2^i-1} 1/j$$

for $1 \leqslant i \leqslant J$ (thus $\lambda_i \approx \log 2$). Then by Cauchy–Schwarz,

$$\begin{aligned} 2^{2r} \prod_{j=1}^{J} \lambda_j^{2b_j} &= \left(\sum_{\mathbf{d} \in \mathscr{D}(\mathbf{b})} \frac{1}{d_1 \cdots d_r} \sum_{\ell} R(\mathbf{d}, \ell) \right)^2 \\ &= \left(\sum_{\mathbf{d} \in \mathscr{D}(\mathbf{b})} \frac{1}{d_1 \cdots d_r} \sum_{\ell \in \mathscr{L}^*(\mathbf{d})} R(\mathbf{d}, \ell) \right)^2 \\ &\leqslant \left(\sum_{\mathbf{d} \in \mathscr{D}(\mathbf{b}), \ell} \frac{R(\mathbf{d}, \ell)^2}{d_1 \cdots d_r} \right) \left(\sum_{\mathbf{d} \in \mathscr{D}(\mathbf{b})} \frac{|\mathscr{L}^*(\mathbf{d})|}{d_1 \cdots d_r} \right). \end{aligned} \tag{3.14}$$

To bound the first sum on right side of (3.14), observe that

$$\sum_{\mathbf{d}\in\mathscr{D}(\mathbf{b}),\ell} \frac{R(\mathbf{d},\ell)^2}{d_1\cdots d_r} = \sum_{I_1,I_2\subset\{1,\dots,r\}} S(I_1,I_2), \tag{3.15}$$

where

$$S(I_1,I_2) = \sum_{\substack{\mathbf{d}\in\mathscr{D}(\mathbf{b})\\ \sum_{i\in I_1} d_i=\sum_{i\in I_2} d_i}} \frac{1}{d_1\cdots d_r}.$$

If $I_1 = I_2$, then evidently $S(I_1,I_2) = \lambda_1^{b_1}\cdots\lambda_J^{b_J}$. If I_1 and I_2 are distinct, let $j = \max(I_1\triangle I_2)$ be the largest coordinate at which I_1 and I_2 differ. With all of the quantities d_i fixed except for d_j, we see that d_j is uniquely determined by the relation $\sum_{i\in I_1} d_i = \sum_{i\in I_2} d_i$. If we define $e(j)\in\{1,\dots,J\}$ uniquely by

$$b_1+\cdots+b_{e(j)-1}+1 \leqslant j \leqslant b_1+\cdots+b_{e(j)},$$

then $d_j \geqslant 2^{e(j)-1}$, regardless of the choice of $d_1,\dots,d_{j-1},d_{j+1},\dots,d_r$ and thus

$$S(I_1,I_2) \leqslant \prod_{\substack{1\leqslant i\leqslant r\\ i\neq j}} \left(\sum_{d_i}\frac{1}{d_i}\right)\frac{1}{2^{e(j)-1}} = \frac{\lambda_1^{b_1}\cdots\lambda_J^{b_J}\lambda_{e(j)}^{-1}}{2^{e(j)-1}} \ll \frac{\lambda_1^{b_1}\cdots\lambda_J^{b_J}}{2^{e(j)}},$$

where the sums over d_i go over the dyadic intervals determined by $\mathbf{d}\in\mathscr{D}(\mathbf{b})$. Here we used the fact that $\lambda_i \asymp 1$; in fact one may note that $\lambda_i \geqslant \lambda_{i+1}$ for all i (since $\frac{1}{n} \geqslant \frac{1}{2n}+\frac{1}{2n+1}$) and that $\lim_{i\to\infty}\lambda_i = \log 2$, so in fact $\lambda_i \geqslant \log 2$ for all i.

Since the number of pairs of subsets $I_1,I_2\subset\{1,\dots,r\}$ with $\max(I_1\triangle I_2) = j$ is exactly 2^{r+j-1}, we deduce from (3.15) that

$$\begin{aligned}
\prod_{j=1}^{J}\lambda_j^{-b_j}\sum_{d\in\mathscr{D}(\mathbf{b}),\ell}\frac{R(\mathbf{d},\ell)^2}{d_1\cdots d_r} &\ll 2^r + 2^r\sum_{j=1}^{r} 2^{j-e(j)}\\
&= 2^r + 2^r\sum_{i=1}^{J}2^{-i}\sum_{j:e(j)=i}2^j\\
&\ll 2^r + 2^r\sum_{i=1}^{J}2^{b_1+\cdots+b_i-i}\\
&\ll 2^r\sum_{i=1}^{J}2^{b_1+\cdots+b_i-i}.
\end{aligned}$$

Compare with (3.14) and use again the fact that $\lambda_i \geqslant \log 2$ to complete the proof. □

Combining Lemma 3.4.1 and (3.13), we obtain

$$\begin{aligned}\mathbf{E}|\mathscr{L}(\mathbf{X}_k)| &\gg \frac{1}{k}\sum_r \sum_{b_1+\cdots+b_J=r} \frac{1}{b_1!\cdots b_J!}\frac{(2\log 2)^r}{\sum_{i=1}^J 2^{b_1+\cdots+b_i-i}} \\ &\geqslant \frac{(2\log 2)^J}{k}\sum_{b_1+\cdots+b_J=J} \frac{1}{b_1!\cdots b_J!}\frac{1}{\sum_{i=1}^J 2^{b_1+\cdots+b_i-i}}. \end{aligned} \tag{3.16}$$

Here we threw away the terms with $r \neq J$. (This is not as wasteful as it appears. A more careful but unnecessary analysis would demonstrate that the main contribution comes from terms with $r = J + O(1)$.) Somewhat surprisingly, the right hand side here can be evaluated explicitly, as in [For08a], by (redundantly) averaging over the J cyclic permutations of $b_1, \dots, b_J$.

Lemma 3.4.2. *Let $x_1, \dots, x_J$ be positive reals such that $x_1 \cdots x_J = 1$. Then the average of $\left(\sum_{i=1}^J x_1 \cdots x_i\right)^{-1}$ over cyclic permutations of $x_1, \dots, x_J$ is exactly $1/J$.*

Proof. In this proof we read all indices modulo J. This makes sense even for the products such as $x_1 \cdots x_i$, because of the assumption $x_1 \cdots x_J = 1$. Now observe that if $1 \leqslant t \leqslant J$ then

$$\sum_{i=1}^J x_{t+1}\cdots x_{t+i} = \frac{\sum_{i=1}^J x_1\cdots x_{t+i}}{x_1\cdots x_t} = \frac{\sum_{i=1}^J x_1\cdots x_i}{x_1\cdots x_t}.$$

Thus

$$\sum_{t=1}^J \frac{1}{\sum_{i=1}^J x_{t+1}\cdots x_{t+i}} = \sum_{t=1}^J \frac{x_1\cdots x_t}{\sum_{i=1}^J x_1\cdots x_i} = 1. \qquad \square$$

Applying the lemma with $x_i = 2^{b_i-1}$ gives that

$$\sum_{b_1+\cdots+b_J=J} \frac{1}{b_1!\cdots b_J!\sum_{i=1}^J 2^{b_1+\cdots+b_i-i}} = \frac{1}{J}\sum_{b_1+\cdots+b_J=J}\frac{1}{b_1!\cdots b_J!} = \frac{1}{J}\frac{J^J}{J!},$$

the second equality being a consequence of the multinomial theorem.

Substitute this into (3.16) and recall that $J = \frac{\log k}{\log 2} + O(1)$. The lower bound in Proposition 3.1.5 now follows from Stirling's formula.

3.5 The upper bound in Proposition 3.1.5

In this section we turn to the upper bound in Proposition 3.1.5, i.e., the bound

$$\mathbf{E}|\mathscr{L}(\mathbf{X}_k)| \ll k^{1-\delta}(\log k)^{-3/2}.$$

As with the lower bound, we condition on $r = X_1 + \cdots + X_k$. Recall from (3.11) the definition of $\mathscr{L}^*(\mathbf{a})$:

$$\mathscr{L}^*(\mathbf{a}) = \left\{ \sum_{i \in I} a_i : I \subset \{1, \dots, r\} \right\}.$$

From (3.4), (3.3) and (3.10) we have

$$\mathbf{E}|\mathscr{L}(\mathbf{X}_k)| \leqslant \frac{1}{k} \sum_r \frac{1}{r!} \sum_{a_1, \dots, a_r = 1}^{k} \frac{|\mathscr{L}^*(\mathbf{a})|}{a_1 \cdots a_r}. \tag{3.17}$$

The most common cause for $|\mathscr{L}^*(\mathbf{a})|$ to be small is for many of the a_i to be small. To capture this, let $\tilde{a}_1, \tilde{a}_2, \dots$ be the increasing rearrangement of the sequence $\mathbf{a}$, so that $\tilde{a}_1 \leqslant \tilde{a}_2 \leqslant \cdots$ (the *order statistics* of $\mathbf{a}$). For any j satisfying $0 \leqslant j \leqslant r$, we then have

$$\mathscr{L}^*(\mathbf{a}) \subset \left\{ m + \sum_{i \in I} \tilde{a}_i : 0 \leqslant m \leqslant \sum_{i=1}^{j} \tilde{a}_i,\ I \subset \{j+1, \dots, r\} \right\}.$$

Thus we have

$$|\mathscr{L}^*(\mathbf{a})| \leqslant G(\mathbf{a}),$$

where

$$G(\mathbf{a}) = \min_{0 \leqslant j \leqslant r} 2^{r-j} \left(\tilde{a}_1 + \cdots + \tilde{a}_j + 1 \right). \tag{3.18}$$

It is reasonable to expect that

$$\begin{aligned} \sum_{a_1, \dots, a_r = 1}^{k} \frac{G(\mathbf{a})}{a_1 \cdots a_r} &\sim \int_1^k \cdots \int_1^k \frac{G(\mathbf{t})}{t_1 \cdots t_r} d\mathbf{t} \\ &= (\log k)^r \int_0^1 \cdots \int_0^1 G(e^{\xi_1 \log k}, \dots, e^{\xi_r \log k}) d\boldsymbol{\xi}. \end{aligned} \tag{3.19}$$

Here we have enlarged the domain of G to include r-tuples of positive real numbers. However, G is not an especially regular function, so (3.19) is perhaps too much to hope for, but G is at least increasing in every coordinate, and we may exploit this to prove an approximate version of (3.19).

Lemma 3.5.1. *For any $r \geqslant 1$, we have*

$$\sum_{a_1, \dots, a_r = 1}^{k} \frac{|\mathscr{L}^*(\mathbf{a})|}{a_1 \cdots a_r} \ll (2h_k)^r r! \int_{\Omega_r} \min_{0 \leqslant j \leqslant r} 2^{-j} (k^{\xi_1} + \cdots + k^{\xi_j} + 1) d\boldsymbol{\xi},$$

where $\Omega_r = \{(\xi_1, \dots, \xi_r) : 0 \leqslant \xi_1 \leqslant \xi_2 \leqslant \ldots \leqslant \xi_r \leqslant 1\}$.

Proof. Motivated by the fact that $1/a = \int_{\exp(h_{a-1})}^{\exp(h_a)} dt/t$, define the product sets

$$R(\mathbf{a}) = \prod_{i=1}^{r} \left[\exp\left(h_{a_i-1}\right), \exp\left(h_{a_i}\right)\right].$$

By (3.18), we have

$$\sum_{a_1,\dots,a_r=1}^{k} \frac{|\mathscr{L}^*(\mathbf{a})|}{a_1\cdots a_r} \leqslant \sum_{a_1,\dots,a_r=1}^{k} \frac{G(\mathbf{a})}{a_1\cdots a_r} = \sum_{a_1,\dots,a_r=1}^{k} G(\mathbf{a}) \int_{R(\mathbf{a})} \frac{d\mathbf{t}}{t_1\cdots t_r}.$$

Consider some $\mathbf{t} \in R(\mathbf{a})$. Writing $\tilde{t}_1 \leqslant \tilde{t}_2 \leqslant \ldots \leqslant \tilde{t}_r$ for the increasing rearrangement of $\mathbf{t}$, we have

$$\exp\left(h_{\tilde{a}_i-1}\right) \leqslant \tilde{t}_i \leqslant \exp\left(h_{\tilde{a}_i}\right) \quad \text{for } 1 \leqslant i \leqslant r.$$

From (3.3) we see that $\tilde{t}_i \geqslant \tilde{a}_i$ for all i. Hence

$$G(\mathbf{a}) \leqslant \min_{0\leqslant j\leqslant r} 2^{r-j}(\tilde{t}_1 + \cdots + \tilde{t}_j + 1) = G(\mathbf{t}).$$

for all $\mathbf{t} \in R(\mathbf{a})$. Thus

$$\begin{aligned}\sum_{a_1,\dots,a_r=1}^{k} G(\mathbf{a}) \int_{R(\mathbf{a})} \frac{d\mathbf{t}}{t_1\cdots t_r} &\leqslant \sum_{a_1,\dots,a_r=1}^{k} \int_{R(\mathbf{a})} \frac{G(\mathbf{t})}{t_1\cdots t_r} d\mathbf{t} \\ &= \int_1^{\exp(h_k)} \cdots \int_1^{\exp(h_k)} \frac{G(\mathbf{t})}{t_1\cdots t_r} d\mathbf{t}.\end{aligned}$$

The integrand on the right is symmetric in $t_1, \dots, t_r$. Making the change of variables $t_i = e^{\xi_i h_k}$ yields

$$\sum_{a_1,\dots,a_r=1}^{k} \frac{|\mathscr{L}^*(\mathbf{a})|}{a_1\cdots a_r} \leqslant (2h_k)^r r! \int_{\Omega_r} \min_{0\leqslant j\leqslant r} 2^{-j} \left(e^{\xi_1 h_k} + \cdots + e^{\xi_j h_k} + 1\right) d\boldsymbol{\xi}.$$

The lemma now follows from the upper bound in (3.3), namely $h_k \leqslant 1 + \log k$. □

With Lemma 3.5.1 established, we can finish the proof of the upper bound in Proposition 3.1.5 by quoting [For08a, Lemma 3.6]. Indeed, in the notation of that paper

$$\int_{\Omega_r} \min_{0\leqslant j\leqslant r} 2^{-j} \left(k^{\xi_1} + \cdots + k^{\xi_j} + 1\right) d\boldsymbol{\xi} = U_r(\log_2 k),$$

and thus by (3.17) and Lemma 3.5.1 we have

$$\mathbf{E}|\mathscr{L}(\mathbf{X}_k)| \ll \frac{1}{k} \sum_r (2h_k)^r U_r(\log_2 k). \tag{3.20}$$

Now [For08a, Lemma 3.6] provides the bound

$$U_r(\log_2 k) \ll \frac{1+|\log_2 k - r|^2}{(r+1)!(2^{r-\log_2 k}+1)},$$

uniformly for $0 \leqslant r \leqslant 10\log_2 k$. Put $r_* = \lfloor \log_2 k \rfloor$. If $r = r_* + m$ with $m \leqslant 9\log_2 k$ then we have

$$\begin{aligned}(2h_k)^r U_r(\log_2 k) &\ll \frac{(2h_k)^r}{(r+1)!}\frac{1+m^2}{2^m} \\ &\leqslant \frac{(2h_k)^{r_*}}{(r_*+1)!}\left(\frac{h_k}{r_*+1}\right)^m (1+m^2).\end{aligned}$$

This series is of course rapidly convergent because $h_k/(r_*+1) \approx \log 2 < 1$, and thus the contribution to (3.20) from this range of r is acceptable. Similarly if $r = r_* - m$ then we have

$$\begin{aligned}(2h_k)^r U_r(\log_2 k) &\ll \frac{(2h_k)^r}{(r+1)!}(1+m^2) \\ &\leqslant \frac{(2h_k)^{r_*}}{(r_*+1)!}\left(\frac{2h_k}{r_*+1}\right)^{-m} (1+m^2),\end{aligned}$$

and this contribution is acceptable as well because $2h_k/(r_*+1) \approx 2\log 2 > 1$. Finally, if $r > 10\log_2 k$ then we use the trivial bound $U_r(\log_2 k) \leqslant 1/r^2$ to deduce that

$$\sum_{r>10\log_2 k} (2h_k)^r U_r(\log_2 k) \ll \sum_{r>10\log_2 k} \frac{(2h_k)^r}{r!} \ll k^{-10}.$$

Thus from (3.20) we have

$$\mathbf{E}|\mathscr{L}(\mathbf{X}_k)| \ll \frac{1}{k}\frac{(2h_k)^{r_*}}{(r_*+1)!} \asymp k^{1-\delta}(\log k)^{-3/2}.$$

Remark. It is obvious from this analysis and the lower bound in our main theorem that a proportion $\geqslant 1-\varepsilon$ of all permutations fixing some set of size k have $\log_2 k + O(\log(1/\varepsilon))$ cycles of length at most k. It is most probably also true that for a proportion $\geqslant 1-\varepsilon$ of all permutations fixing some set of size k we have $\log \tilde{a}_j \geqslant j\log 2 - O_\varepsilon(1)$ for $j \leqslant \log_2 k - O_\varepsilon(1)$, where the $\tilde{a}_j$ are the (ordered) cycle lengths of the permutation. To establish this would require opening up some of the arguments used to bound the quantities U_k in [For08a]. We plan to return to this and other issues in a future paper.

3.6 Four generators are enough

We turn now to the problem of invariable generation, starting with the four-generator theorem, Theorem 3.1.9. The principal result needed for the proof is the following proposition.

Proposition 3.6.1. *Let $k \leqslant n/2$. If $\pi_1, \pi_2, \pi_3, \pi_4 \in \mathcal{S}_n$ are chosen uniformly at random, then the probability that there is some $\ell \in (k/2, k]$ such that $\pi_1, \pi_2, \pi_3, \pi_4$ each fix some set of size ℓ is $O(k^{-c})$ for some $c > 0$.*

Over the past several sections we showed that the probability that a random permutation $\pi \in \mathcal{S}_n$ fixes some set of size k is $k^{-\delta+o(1)}$, where $\delta = 1 - \frac{1+\log\log 2}{\log 2} \approx 0.086$. As noted there, the main contribution to this estimate comes from rather exceptional permutations with an unexpectedly large number of cycles of length at most k, roughly $(\log 2)^{-1} \log k$ such cycles. By contrast a typical permutation has roughly $\log k$ cycles of length at most k. By restricting to this "quenched" regime we can establish a much stronger bound.

Lemma 3.6.2. *Suppose $k \leqslant n/2$, $0 < \varepsilon \leqslant 1/2$, and choose $\pi \in \mathcal{S}_n$ uniformly at random. Then the probability that π both (a) fixes a set of size k and (b) has at most $(1+\varepsilon)\log k$ cycles of length at most k is at most $O(k^{\log 2 - 1 + 2\varepsilon})$.*

Proof. Fix $\ell \leqslant (1+\varepsilon)\log k$ and consider permutations π with exactly ℓ cycles of length at most k and having some fixed set X of size k. Write $\pi = \pi_1\pi_2$, where π_1 is supported on X and π_2 is supported on X^c, and suppose that π_1 has ℓ_1 cycles of length at most k and π_2 has ℓ_2 cycles of length at most k, where $\ell_1 + \ell_2 = \ell$. By Lemma 3.2.4 the number of such π, for a given choice of X and ℓ_1, ℓ_2, is bounded by a constant times

$$\frac{h_k^{\ell_1}}{k\ell_1!} k! \cdot \frac{h_k^{\ell_2}}{k\ell_2!} (n-k)!.$$

Thus by multiplying by the number of choices of X and summing over $\ell_1 + \ell_2 = \ell$, the probability we are interested in is bounded by a constant times

$$\sum_{\ell_1+\ell_2=\ell} \frac{1}{k^2} \frac{h_k^\ell}{\ell_1!\ell_2!} = \frac{(2h_k)^\ell}{k^2\ell!}.$$

Now by summing over all $\ell \leqslant \ell_0 = \lfloor (1+\varepsilon)\log k \rfloor$ we get the bound

$$\frac{1}{k^2} \sum_{\ell \leqslant (1+\varepsilon)\log k} \frac{(2h_k)^\ell}{\ell!} \ll \frac{1}{k^2} \frac{(2h_k)^{\ell_0}}{\ell_0!}.$$

The result now follows from $h_k = \log k + O(1)$ and Stirling's formula, as usual. □

Proof of Proposition 3.6.1. Let $\varepsilon > 0$ be a small. We will first use Lemma 3.2.4 to bound the probability that one of $\pi_1, \pi_2, \pi_3, \pi_4$ has more than $\ell_0 = \lfloor (1+\varepsilon)\log k \rfloor$ cycles of length at most k. By that lemma, for each $\ell \geqslant \ell_0$, the probability that π_1 has ℓ cycles of length at most k is bounded by

$$O\left(\frac{h_k^{\ell-1}}{k(\ell-1)!}\right),$$

so the probability that π_1 has more than ℓ_0 cycles is bounded by a constant times

$$\sum_{\ell>\ell_0} \frac{h_k^{\ell-1}}{k(\ell-1)!} \ll \frac{h_k^{\ell_0-1}}{k(\ell_0-1)!} \ll \frac{1}{k}\left(\frac{eh_k}{\ell_0-1}\right)^{\ell_0-1} \ll \frac{1}{k}\left(\frac{e}{1+\varepsilon}\right)^{(1+\varepsilon)\log k}.$$

By a Taylor expansion of $-1+(1+\varepsilon)\log(e/(1+\varepsilon))$, this is bounded by $O(k^{-\varepsilon^2/3})$ if $\varepsilon \leqslant \frac{1}{2}$. Thus the probability that one of $\pi_1, \pi_2, \pi_3, \pi_4$ has more than ℓ_0 cycles of length at most k is bounded by $O(k^{-\varepsilon^2/3})$.

On the other hand, by Lemma 3.6.2, for each $\ell \in (k/2, k]$ the probability that π_i has at most $(1+\varepsilon)\log k$ cycles of length at most k and fixes a set of size ℓ is at most $k^{\log 2 - 1 + 2\varepsilon}$. Thus the probability that $\pi_1, \pi_2, \pi_3, \pi_4$ each have at most $(1+\varepsilon)\log k$ cycles of length at most k and each fix a set of the same size ℓ for some $\ell \in (k/2, k]$ is at most $k^{1+4(\log 2-1+2\varepsilon)}$. Since $1+4(\log 2-1)<0$, we have $1+4(\log 2-1+2\varepsilon)<0$ if ε is small enough ($\varepsilon = 1/40$ works), and so the theorem holds with

$$c = \min(\varepsilon^2/3, -1-4(\log 2 - 1 + 2\varepsilon)). \qquad \square$$

An immediate corollary of Proposition 3.6.1 is obtained by fixing k, letting $n \to \infty$, and recalling the Lemma 3.1.1 and the definition (3.2) of $\mathscr{L}(\mathbf{X})$.

Corollary 3.6.3. *There exists $c > 0$ such that for any $k \geqslant 1$ the probability that*

$$\mathscr{L}(\mathbf{X}) \cap \mathscr{L}(\mathbf{X}') \cap \mathscr{L}(\mathbf{X}'') \cap \mathscr{L}(\mathbf{X}''') \cap (k/2, k] \neq \varnothing$$

is $O(k^{-c})$.

Remark. If one wished to prove only this, we could replace Lemma 3.2.4 with a corresponding bound for $\mathbf{P}(X_1+\cdots+X_k \leqslant \ell)$, which is trivial given that $X_1+\cdots+X_k$ is Poisson with parameter h_k.

Corollary 3.6.4. *$\mathscr{L}(\mathbf{X}) \cap \mathscr{L}(\mathbf{X}') \cap \mathscr{L}(\mathbf{X}'') \cap \mathscr{L}(\mathbf{X}''')$ is almost surely finite, and equal to $\{0\}$ with positive probability.*

Proof. Let F_k be the event that

$$\mathscr{L}(\mathbf{X}) \cap \mathscr{L}(\mathbf{X}') \cap \mathscr{L}(\mathbf{X}'') \cap \mathscr{L}(\mathbf{X}''') \cap (k, \infty)$$

is nonempty. By applying Corollary 3.6.3 with k replaced by $2^j k$, $j \in \mathbf{N}$, and summing the geometric series, we obtain $\mathbf{P}(F_k) \ll k^{-c}$ for $k \geqslant 1$. In particular $\mathbf{P}(F_k) \to 0$, so $\mathbf{P}\left(\bigcap F_k\right) = 0$, so the first part of the corollary holds. For the second part, fix k_0 such that $\mathbf{P}(F_{k_0}) < 1$. Then

$$\begin{aligned}\mathbf{P}(\mathscr{L}(\mathbf{X}) \cap \mathscr{L}(\mathbf{X}') \cap \mathscr{L}(\mathbf{X}'') \cap \mathscr{L}(\mathbf{X}''') = \{0\}) &\geqslant \mathbf{P}(X_j = 0 \text{ for all } j \leqslant k_0, \text{ and } F_{k_0}^c) \\ &\geqslant \mathbf{P}(X_j = 0 \text{ for all } j \leqslant k_0)\, \mathbf{P}(F_{k_0}^c) \\ &> 0.\end{aligned}$$

The second inequality here is a simple case of the FKG inequality (see [TV10, Theorem 1.19]). To see the inequality directly, define $\mathbf{X}^* = (X_1^*, X_2^*, \dots)$ by putting $X_j^* = X_j$ if $j > k_0$ and $X_j^* = 0$ if $j \leqslant k_0$, and let $F_{k_0}^*$ be the event that

$$\mathscr{L}(\mathbf{X}^*) \cap \mathscr{L}(\mathbf{X}') \cap \mathscr{L}(\mathbf{X}'') \cap \mathscr{L}(\mathbf{X}''') \cap (k_0, \infty)$$

is nonempty. Clearly $F_{k_0}^*$ implies F_{k_0}, so

$$\begin{aligned}\mathbf{P}(X_j = 0 \text{ for all } j \leqslant k_0, \text{ and } F_{k_0}^c) &= \mathbf{P}(X_j = 0 \text{ for all } j \leqslant k_0, \text{ and } F_{k_0}^{*c}) \\ &= \mathbf{P}(X_j = 0 \text{ for all } j \leqslant k_0)\, \mathbf{P}(F_{k_0}^{*c}) \\ &\geqslant \mathbf{P}(X_j = 0 \text{ for all } j \leqslant k_0)\, \mathbf{P}(F_{k_0}^c). \qquad \square\end{aligned}$$

Shortly we will complete the proof of Theorem 3.1.9. In the proof, we will need a trick to deal with the possibility that $\pi_1, \pi_2, \pi_3, \pi_4 \in \mathcal{A}_n$. The following lemma is helpful in this regard. It shows that random even and random odd permutations have the same small-cycle structure as random permutations with unconstrained parity (Lemma 3.1.1).

Lemma 3.6.5. *Let $\pi \in \mathcal{S}_n$ be a random even permutation, and let $c_j(\pi)$ be the number of cycles of length j. Fix $k \in \mathbf{N}$. Then as $n \to \infty$ the distribution of $(c_1(\pi), \dots, c_k(\pi))$ converges to that of $(X_1, \dots, X_k)$. The same is true if π is a random odd permutation.*

Proof. Choose $\pi \in \mathcal{S}_n$ uniformly at random, and define σ by putting $\sigma = 1$ if π is even and $\sigma = (12)$ if π is odd. Then $\pi\sigma$ is uniformly distributed over $\mathcal{A}_n$. By Lemma 3.1.1, as $n \to \infty$, the number of cycles in π of length at most $2k$ approaches a Poisson distribution with parameter $h_{2k} \leqslant 1 + \log 2k$. Thus, with high probability (as $n \to \infty$) the total number of points in cycles of π of length at most $2k$ is at most

$2k \log n$, so with high probability each of these cycles is disjoint from (12). That is, the points 1 and 2 are both contained in cycles of π of length at least $2k+1$ with high probability.

Now consider the probability that 1 and 2 are both contained in the same cycle and are close together. For each $\ell \geqslant 2k+1$, the number of cycles of length ℓ containing both 1 and 2, which are a distance $\leqslant k$ from each other, equals $\binom{n-2}{\ell-2} 2k(\ell-2)!$. Hence, the number of permutations π containing such a cycle is at most

$$\sum_{2k+1 \leqslant \ell \leqslant n} 2k(n-2)! \leqslant 2k(n-1)!.$$

Hence, with high probability, if 1 and 2 are in the same cycle they are a distance at least $k+1$ from each other. Thus, with high probability, $c_j(\pi\sigma) = c_j(\pi)$ for each $j \leqslant k$.

Similarly $\pi\sigma(12)$ is uniformly distributed over odd permutations, and with high probability $c_j(\pi\sigma(12)) = c_j(\pi)$. □

We will also need the following important theorem of Łuczak and Pyber, mentioned several times in the introduction.

Theorem 3.6.6 (Łuczak–Pyber [ŁP93])**.** *The union of all transitive subgroups $H \leqslant \mathcal{S}_n$ other than $\mathcal{S}_n$ and $\mathcal{A}_n$ has size $o(n!)$. In other words, the probability that a random permutation π is contained in some transitive subgroup other than $\mathcal{S}_n$ or $\mathcal{A}_n$ tends to zero.*

We can now finally complete the proof of Theorem 3.1.9.

Proof of Theorem 3.1.9. Let $\pi_1, \pi_2, \pi_3, \pi_4 \in \mathcal{S}_n$ be random permutations with π_1 odd. Let $E_{n,k}$ be the event that $\pi_1, \pi_2, \pi_3, \pi_4$ each fix a set of size ℓ for some ℓ in the range $1 \leqslant \ell \leqslant k$, and let $F_{n,k}$ be the event that $\pi_1, \pi_2, \pi_3, \pi_4$ each fix a set of size ℓ for some ℓ in the range $k < \ell \leqslant n/2$. By Proposition 3.6.1 (and summing a geometric series as in the proof of Corollary 3.6.4) we have

$$\mathbf{P}(F_{n,k}) \ll k^{-c}$$

uniformly for $1 \leqslant k \leqslant n/2$, while by Corollary 3.6.4 and Lemma 3.6.5 we have

$$\lim_{n \to \infty} \mathbf{P}(E_{n,k}) \leqslant 1 - \delta$$

for all k, for some constant $\delta > 0$. Fix k_0 such that $\mathbf{P}(F_{n,k_0}) \leqslant \delta/3$ for all $n \geqslant 2k_0$. Then

$$\mathbf{P}(E_{n,k_0}) + \mathbf{P}(F_{n,k_0}) \leqslant 1 - \delta/3$$

for all sufficiently large n. Hence with probability bounded away from zero the four permutations $\pi_1, \pi_2, \pi_3, \pi_4$ do not share any fixed-set size $\ell \in [1, n/2]$.

Thus with probability bounded away from zero $\pi_1, \pi_2, \pi_3, \pi_4$ invariably generate a transitive subgroup of $\mathcal{S}_n$. By Theorem 3.6.6, π_1 is, with high probability, not contained in any transitive subgroup other $\mathcal{S}_n$ and possibly $\mathcal{A}_n$. Since $\pi_1 \notin \mathcal{A}_n$, it must be that with probability bounded away from zero $\pi_1, \pi_2, \pi_3, \pi_4$ invariably generate $\mathcal{S}_n$. □

3.7 Three generators are not enough

Finally, we now present our main new contribution, Theorem 3.1.10, which asserts that three random permutations with high probability fail to invariably generate. Theorem 3.1.10 follows from the following more specific proposition.

Proposition 3.7.1. *For every $\varepsilon > 0$ there exists $k_0 = k_0(\varepsilon)$ and $n_0 = n_0(\varepsilon)$ such that if $n \geqslant n_0$ then with probability at least $1 - \varepsilon$ there is some $\ell \leqslant k_0$ such that π_1, π_2, π_3 each fix a set of size ℓ.*

Let $\mathbf{X}$ be defined as before, and let $\mathbf{Y}$ and $\mathbf{Z}$ be independent copies of $\mathbf{X}$. For I an interval in $\mathbf{N}$ let

$$\mathscr{L}(I, \mathbf{X}) = \left\{ \sum_{j \in I} j x_j : 0 \leqslant x_j \leqslant X_j \text{ for each } j \right\},$$

and define $\mathscr{L}(I, \mathbf{Y})$ and $\mathscr{L}(I, \mathbf{Z})$ analogously.

Lemma 3.7.2. *Let $I = \{1, \ldots, k\}$ and let $\varepsilon > 0$. Then with probability at least $1 - \varepsilon$ we have $\mathscr{L}(I, \mathbf{X}), \mathscr{L}(I, \mathbf{Y}), \mathscr{L}(I, \mathbf{Z}) \subset [0, 3\varepsilon^{-1}k]$.*

Proof. Since $\mathbf{E} \sum_{j \in I} j X_j = |I| = k$, by Markov's inequality we have $\sum_{j \in I} j X_j \leqslant 3\varepsilon^{-1}k$ with probability at least $1 - \varepsilon/3$. Similarly $\sum_{j \in I} j Y_j \leqslant 3\varepsilon^{-1}k$ and $\sum_{j \in I} j Z_j \leqslant 3\varepsilon^{-1}k$ each with probability at least $1 - \varepsilon/3$, and the lemma follows. □

Lemma 3.7.3. *Fix $\varepsilon > 0$. There is a constant $C(\varepsilon)$ so that with probability at least $1 - \varepsilon$ we have*

$$\sum_{m < j \leqslant k} X_j \geqslant 0.99 \log(k/m) - C(\varepsilon),$$
$$\sum_{m < j \leqslant k} Y_j \geqslant 0.99 \log(k/m) - C(\varepsilon), \text{ and}$$
$$\sum_{m < j \leqslant k} Z_j \geqslant 0.99 \log(k/m) - C(\varepsilon).$$

for every nonnegative integer $m \leqslant k$.

Proof. Let $C = C(\varepsilon)$ be a constant whose properties will be specified later. It suffices to show that the first inequality holds for all $m \leqslant k$ with probability at least $1 - \varepsilon/3$. There is nothing to prove if $m \geqslant e^{-C}k$, so we may suppose $m \leqslant e^{-C}k$. We may also suppose that $C \geqslant 1$.

Let E be the event that

$$\sum_{m<j\leqslant k} X_j \geqslant 0.99\log(k/m) - 1$$

for all $m \leqslant e^{-C}k$. Suppose E fails, say

$$\sum_{m<j\leqslant k} X_j < 0.99\log(k/m) - 1$$

for some $m \leqslant e^{-C}k$. Writing m' for the smallest power of 2 with $m' > m$, we thus have

$$\sum_{m'<j\leqslant k} X_j \leqslant \sum_{m<j\leqslant k} X_j \leqslant 0.99\log(k/m) - 1 \leqslant 0.99\log(k/m').$$

Thus

$$1_{E^c} \leqslant \sum_{\substack{m'\leqslant 2e^{-C}k \\ \text{dyadic}}} 0.99^{\left(\sum_{m'<j\leqslant k} X_j - 0.99\log(k/m')\right)}.$$

Whenever P is Poisson of parameter λ and $a > 0$ we have $\mathbf{E}a^P = e^{(a-1)\lambda}$, and the sum $\sum_{m'<j\leqslant k} X_j$ is Poisson with parameter $\sum_{m'<j\leqslant k} 1/j = \log(k/m') + O(1)$, so

$$\begin{aligned}\mathbf{P}(E^c) &\ll \sum_{\substack{m'\leqslant 2e^{-C}k \\ \text{dyadic}}} \exp\left((0.99 - 1 - 0.99\log(0.99))\log(k/m')\right) \\ &\leqslant \sum_{\substack{m'\leqslant 2e^{-C}k \\ \text{dyadic}}} (k/m')^{-0.00005} \\ &\ll e^{-0.00005C}.\end{aligned}$$

Therefore, $\mathbf{P}(E^c) \leqslant \varepsilon/3$ if C is taken large enough. □

We need a standard estimate for the partial sums of the Fourier series

$$\sum_{j=1}^{\infty} \frac{\cos(2\pi j\theta)}{j} = -\log|2\sin(\pi\theta)|.$$

For $\theta \in \mathbf{R}$ denote by $\|\theta\|$ be the distance from θ to $\mathbf{Z}$.

Lemma 3.7.4.

$$\sum_{j\leqslant m} \frac{\cos(2\pi j\theta)}{j} = \log\min\left(\frac{1}{\|\theta\|}, m\right) + O(1) \quad \textit{for } \theta \in \mathbf{R}.$$

Proof. We may assume that $0 < \theta \leqslant \frac{1}{2}$. Using the approximation $\cos(2\pi j\theta) = 1 + O(j^2\theta^2)$, we have

$$\sum_{j \leqslant \min(m,1/\theta)} \frac{\cos(2\pi j\theta)}{j} = \log \min(m, 1/\theta) + O(1).$$

This proves the lemma if $\|\theta\| \leqslant 1/m$. Suppose, then, that $\|\theta\| > 1/m$. Set

$$S_j = \sum_{n=0}^{j} e^{2\pi i n\theta},$$

and note that by summing the geometric series we have

$$S_j = \frac{e^{2\pi ij\theta} - 1}{e^{2\pi i\theta} - 1} \ll \frac{1}{\theta}. \tag{3.21}$$

Thus (by "Abel summation"),

$$\begin{aligned}\sum_{1/\theta < j \leqslant m} \frac{\cos(2\pi j\theta)}{j} &= \Re \sum_{1/\theta < j \leqslant m} \frac{e^{2\pi ij\theta}}{j} = \Re \sum_{1/\theta < j \leqslant m} \frac{S_j - S_{j-1}}{j} \\ &= \Re \sum_{1/\theta < j \leqslant m-1} \frac{S_j}{j(j+1)} + \frac{S_m}{m} - \frac{S_{\lceil 1/\theta \rceil - 1}}{\lceil 1/\theta \rceil}.\end{aligned}$$

The latter two terms here are $O(1)$ by the trivial bound $|S_j| \leqslant j$, while from (3.21) the sum is bounded by a constant times

$$\frac{1}{\theta} \sum_{j > 1/\theta} \frac{1}{j^2} \ll 1. \qquad \square$$

Let $\mathbf{T} = \mathbf{R}/\mathbf{Z}$ be the unit torus, and let $e(z) = e^{i2\pi z}$. Given $I, \mathbf{X}, \mathbf{Y}, \mathbf{Z}$ define $F : \mathbf{T}^2 \to \mathbf{C}$ by

$$F(\boldsymbol{\theta}) = \prod_{j \in I} \left(\frac{1 + e(j\theta_1)}{2}\right)^{X_j} \left(\frac{1 + e(j\theta_2)}{2}\right)^{Y_j} \left(\frac{1 + e(j(-\theta_1 - \theta_2))}{2}\right)^{Z_j}.$$

By expanding the product we see that $\widehat{F} : \mathbf{Z}^2 \to \mathbf{C}$ is supported on the set

$$S(I, \mathbf{X}, \mathbf{Y}, \mathbf{Z}) = \{(n_1 - n_3, n_2 - n_3) : n_1 \in \mathscr{L}(I, \mathbf{X}), n_2 \in \mathscr{L}(I, \mathbf{Y}), n_3 \in \mathscr{L}(I, \mathbf{Z})\}. \tag{3.22}$$

Since $\sum_{a \in \mathbf{Z}^2} \widehat{F}(a) = F(0) = 1$, by Cauchy–Schwarz we have

$$1 = \left(\sum_{a \in \mathbf{Z}^2} \widehat{F}(a)\right)^2 \leqslant \left(\sum_{a : \widehat{F}(a) \neq 0} 1\right) \sum_{a \in \mathbf{Z}^2} |\widehat{F}(a)|^2 \leqslant |S(I, \mathbf{X}, \mathbf{Y}, \mathbf{Z})| \sum_{a \in \mathbf{Z}^2} |\widehat{F}(a)|^2.$$

Applying Parseval, we get

$$|S(I, \mathbf{X}, \mathbf{Y}, \mathbf{Z})| \geqslant \left(\sum_{a \in \mathbf{Z}^2} |\widehat{F}(a)|^2\right)^{-1} = \left(\int_{\mathbf{T}^2} |F(\boldsymbol{\theta})|^2 \, d\boldsymbol{\theta}\right)^{-1}. \tag{3.23}$$

Lemma 3.7.5. *Let*

$$\beta = 1 - \frac{2}{3\log 2} - 0.02 \approx 0.0182,$$

and let $I = (k^\beta, k]$. Fix $\varepsilon \in (0, 1/2)$, and let $E = E(\varepsilon)$ be the event from Lemma 3.7.3. Then both of the bounds

$$\mathbf{E} 1_E |F(\boldsymbol{\theta})|^2 \ll_\varepsilon (k\|\theta_1\|^{1/3}\|\theta_2\|^{1/3}\|\theta_3\|^{1/3})^{-2.02} \tag{3.24}$$

and

$$\mathbf{E} 1_E |F(\boldsymbol{\theta})|^2 \ll_\varepsilon (k\|\theta_i\|^{1/2}\|\theta_j\|^{1/2})^{-1.3} \quad (\{i,j\} \subset \{1,2,3\}) \tag{3.25}$$

hold uniformly for $\boldsymbol{\theta} \in \mathbf{T}^2$, where $\theta_3 = -\theta_1 - \theta_2$. The expectation is over $\mathbf{X}, \mathbf{Y}, \mathbf{Z}$.

Proof. Define, for $i \in \{1,2,3\}$,

$$k_i = \begin{cases} k^\beta & \text{if } \|\theta_i\| \geqslant k^{-\beta}, \\ 1/\|\theta_i\| & \text{if } 1/k < \|\theta_i\| < k^{-\beta}, \\ k & \text{if } \|\theta_i\| \leqslant 1/k. \end{cases}$$

It is useful to note the crude bound

$$k_i \leqslant \frac{k^\beta}{\|\theta_i\|^{1-\beta}} \tag{3.26}$$

which follows by an analysis of the three cases in the definition of k_i. If E holds then

$$\sum_{k_1 < j \leqslant k} X_j + \sum_{k_2 < j \leqslant k} Y_j + \sum_{k_3 < j \leqslant k} Z_j \geqslant 0.99 \log(k^3/(k_1k_2k_3)) - 3C(\varepsilon),$$

so

$$1_E |F(\boldsymbol{\theta})|^2 \ll_\varepsilon (k^3/k_1k_2k_3)^{-0.99\log 2} |F(\boldsymbol{\theta})|^2 2^{\sum_{k_1<j\leqslant k} X_j + \sum_{k_2<j\leqslant k} Y_j + \sum_{k_3<j\leqslant k} Z_j}.$$

From (3.26) and the inequality $3 \times 0.99 \log 2 \times (1-\beta) > 2.02$, we deduce that

$$1_E |F(\boldsymbol{\theta})|^2 \ll_\varepsilon (k\|\theta_1\|^{1/3}\|\theta_2\|^{1/3}\|\theta_3\|^{1/3})^{-2.02} |F(\boldsymbol{\theta})|^2 2^{\sum_{k_1<j\leqslant k} X_j + \sum_{k_2<j\leqslant k} Y_j + \sum_{k_3<j\leqslant k} Z_j}.$$

Thus (3.24) will follow if we can prove

$$\mathbf{E} |F(\boldsymbol{\theta})|^2 2^{\sum_{k_1<j\leqslant k} X_j + \sum_{k_2<j\leqslant k} Y_j + \sum_{k_3<j\leqslant k} Z_j} \ll 1. \tag{3.27}$$

Similarly, from (3.26) for $i = 1, 2$ and the trivial bound $k_3 \leqslant k$, and using $2 \times 0.99 \log 2 \times (1-\beta) > 1.3$, we deduce that

$$1_E |F(\boldsymbol{\theta})|^2 \ll_\varepsilon (k\|\theta_1\|^{1/2}\|\theta_2\|^{1/2})^{-1.3} |F(\boldsymbol{\theta})|^2 2^{\sum_{k_1<j\leqslant k} X_j + \sum_{k_2<j\leqslant k} Y_j + \sum_{k_3<j\leqslant k} Z_j},$$

and similarly for other permutations of the indices $1, 2, 3$, so (3.25) will also follow from (3.27).

It remains only to prove (3.27). We have a factorization

$$\begin{aligned}\mathbf{E}|F(\boldsymbol{\theta})|^2 2^{\sum_{k_1<j\leqslant k} X_j + \sum_{k_2<j\leqslant k} Y_j + \sum_{k_3<j\leqslant k} Z_j} \\ = \prod_{k^\beta<j\leqslant k_1} \mathbf{E}\left|\frac{1+e(j\theta_1)}{2}\right|^{2X_j} \prod_{k_1<j\leqslant k} \mathbf{E}\left(2\left|\frac{1+e(j\theta_1)}{2}\right|^2\right)^{X_j} \\ \times \prod_{k^\beta<j\leqslant k_2} \mathbf{E}\left|\frac{1+e(j\theta_2)}{2}\right|^{2Y_j} \prod_{k_2<j\leqslant k} \mathbf{E}\left(2\left|\frac{1+e(j\theta_2)}{2}\right|^2\right)^{Y_j} \\ \times \prod_{k^\beta<j\leqslant k_3} \mathbf{E}\left|\frac{1+e(j\theta_3)}{2}\right|^{2Z_j} \prod_{k_3<j\leqslant k} \mathbf{E}\left(2\left|\frac{1+e(j\theta_3)}{2}\right|^2\right)^{Z_j}.\end{aligned}$$

By using again the calculation $\mathbf{E}a^P = e^{(a-1)\lambda}$ for P Poisson with parameter λ, we get

$$\begin{aligned}&\mathbf{E}|F(\boldsymbol{\theta})|^2 2^{\sum_{k_1<j\leqslant k} X_j + \sum_{k_2<j\leqslant k} Y_j + \sum_{k_3<j\leqslant k} Z_j} \\ &= \exp\sum_{i=1}^{3}\left(\sum_{k^\beta<j\leqslant k_i}\frac{1}{j}\left(\left|\frac{1+e(j\theta_i)}{2}\right|^2-1\right)+\sum_{k_i<j\leqslant k}\frac{1}{j}\left(2\left|\frac{1+e(j\theta_i)}{2}\right|^2-1\right)\right) \\ &= \exp\sum_{i=1}^{3}\left(\sum_{k^\beta<j\leqslant k_i}\frac{\cos(2\pi j\theta_i)-1}{2j}+\sum_{k_i<j\leqslant k}\frac{\cos(2\pi j\theta_i)}{j}\right) \\ &= \exp\sum_{i=1}^{3}\left(\frac{1}{2}\log\frac{\min(k_i,1/\|\theta_i\|)}{\min(k^\beta,1/\|\theta_i\|)}-\frac{1}{2}\log\frac{k_i}{k^\beta}+\log\frac{\min(k,1/\|\theta_i\|)}{\min(k_i,1/\|\theta_i\|)}+O(1)\right)\end{aligned}$$

by Lemma 3.7.4. Checking the three cases in the definition of k_i separately, it can be confirmed that this is always $O(1)$. □

Corollary 3.7.6. *With notation as in Lemma 3.7.5, we have*

$$\int_{\mathbf{T}^2} \mathbf{E}1_E|F(\boldsymbol{\theta})|^2\,d\boldsymbol{\theta} \ll_\varepsilon k^{-2}.$$

Proof. Divide $\mathbf{T}^2$ into three regions R_1, R_2, R_3 as follows:

$$\begin{aligned}R_1 &= \{\boldsymbol{\theta}\in\mathbf{T}^2 : \|\theta_i\|\geqslant 1/k \text{ for all three } i\in\{1,2,3\}\},\\ R_2 &= \{\boldsymbol{\theta}\in\mathbf{T}^2 : \|\theta_i\|\geqslant 1/k \text{ for exactly two } i\in\{1,2,3\}\},\\ R_3 &= \{\boldsymbol{\theta}\in\mathbf{T}^2 : \|\theta_i\|\geqslant 1/k \text{ for at most one } i\in\{1,2,3\}\}.\end{aligned}$$

We will bound the integral differently in each region.

Further subdivide R_1 according to which of $\|\theta_1\|, \|\theta_2\|, \|\theta_3\|$ is largest. In the subregion R_1' in which say $\|\theta_1\|$ is largest we have $\|\theta_1\| \geqslant \|\theta_2\|^{1/2}\|\theta_3\|^{1/2}$, so by (3.24) we have

$$\begin{aligned}\int_{R_1'} \mathbf{E}1_E|F(\boldsymbol{\theta})|^2\,d\boldsymbol{\theta} &\ll_\varepsilon \int_{R_1} (k\|\theta_2\|^{1/2}\|\theta_3\|^{1/2})^{-2.02}\,d\boldsymbol{\theta}\\ &= \left(\int_{\|\theta\|\geqslant 1/k} (k\|\theta\|)^{-1.01}\,d\theta\right)^2\\ &\asymp k^{-2}.\end{aligned}$$

We can bound the integral over the other subregions in the same way, so the integral over R_1 is indeed $\ll_\varepsilon k^{-2}$.

Similarly, subdivide R_2 according to the relative order of $\|\theta_1\|, \|\theta_2\|, \|\theta_3\|$, and focus for the moment on the subregion R_2' in which $\|\theta_1\| \leqslant \|\theta_2\| \leqslant \|\theta_3\|$. This implies in particular that $\|\theta_1\| \leqslant 1/k$ while $\|\theta_2\| \geqslant 1/k$. Thus by (3.25) with $i = 2$ and $j = 3$ we have

$$\int_{R_2'} \mathbf{E}1_E|F(\boldsymbol{\theta})|^2\,d\boldsymbol{\theta} \ll_\varepsilon \int_{R_2'} (k\|\theta_2\|)^{-1.3}\,d\boldsymbol{\theta} \asymp k^{-2}.$$

Again we can bound the integral over the other subregions in the same way, so the integral over R_2 is also $\ll_\varepsilon k^{-2}$.

Finally, in the region R_3 note that because $\theta_1+\theta_2+\theta_3 = 0$ we must have $\|\theta_i\| < 2/k$ for each i. Thus from the trivial bound $|F(\boldsymbol{\theta})| \leqslant 1$ we have

$$\int_{R_3} \mathbf{E}1_E|F(\boldsymbol{\theta})|^2\,d\boldsymbol{\theta} \leqslant \int_{R_3} 1 \asymp k^{-2}. \qquad \square$$

Recall the definition of $S(I, \mathbf{X}, \mathbf{Y}, \mathbf{Z})$, given in (3.22).

Proposition 3.7.7. *Let $I = (k^\beta, k]$. There is a constant $c > 0$ such that with probability at least $1/2$ we have $S(I, \mathbf{X}, \mathbf{Y}, \mathbf{Z}) \subset [-10k, 10k]^2$ and $|S(I, \mathbf{X}, \mathbf{Y}, \mathbf{Z})| \geqslant ck^2$.*

Proof. Apply Lemma 3.7.3 with $\varepsilon = 0.01$, and let E be the resulting event. By Corollary 3.7.6 (and interchanging the order of integration and expectation) we have

$$\mathbf{E}1_E \int_{\mathbf{T}^2} |F(\boldsymbol{\theta})|^2\,d\boldsymbol{\theta} \ll k^{-2}.$$

Thus by Markov's inequality there is a constant C such that $1_E \int_{\mathbf{T}^2} |F(\boldsymbol{\theta})|^2\,d\boldsymbol{\theta} \leqslant Ck^{-2}$ with probability at least 0.99. Since $\mathbf{P}(E) \geqslant 0.99$ we deduce that $\int_{\mathbf{T}^2} |F(\boldsymbol{\theta})|^2\,d\boldsymbol{\theta} \leqslant Ck^{-2}$ with probability at least 0.98. Applying (3.23), we have $|S(I, \mathbf{X}, \mathbf{Y}, \mathbf{Z})| \geqslant C^{-1}k^2$ with probability at least 0.98.

On the other hand, by Lemma 3.7.2 with $\varepsilon = 1/3$ we have $S(I, \mathbf{X}, \mathbf{Y}, \mathbf{Z}) \subset [-10k, 10k]^2$ with probability at least $2/3$, so we must have both $S(I, \mathbf{X}, \mathbf{Y}, \mathbf{Z}) \subset [-10k, 10k]^2$ and $|S(I, \mathbf{X}, \mathbf{Y}, \mathbf{Z})| \geqslant C^{-1}k^2$ with probability at least $1 - 1/3 - 0.02 \geqslant 1/2$. □

Proposition 3.7.8. *Let $I = (k^\beta, 60k]$. Then with probability bounded away from zero we can find $(x_j)_{j\in I}, (y_j)_{j\in I}, (z_j)_{j\in I}$ not all zero such that $0 \leqslant x_j \leqslant X_j$, $0 \leqslant y_j \leqslant Y_j$, and $0 \leqslant z_j \leqslant Z_j$ for each $j \in I$ and*

$$\sum_{j\in I} jx_j = \sum_{j\in I} jy_j = \sum_{j\in I} jz_j.$$

Proof. Let $I' = (k^\beta, k]$. By Proposition 3.7.7, with probability at least $1/2$ we have

$$S(I', \mathbf{X}, \mathbf{Y}, \mathbf{Z}) \subset [-10k, 10k]^2$$

and $|S(I', \mathbf{X}, \mathbf{Y}, \mathbf{Z})| \gg k^2$. This event depends only on X_j, Y_j, Z_j for $j \leqslant k$, so independently with probability at least $1/2$ we can find $j_3 \in (20k, 50k]$ such that $Z_{j_3} > 0$, as

$$\mathbf{P}(Z_j = 0 \text{ for all } j \in (20k, 50k]) = \prod_{j\in(20k,50k]} e^{-1/j} \leqslant 1/2. \tag{3.28}$$

Given such a j_3 the set T of pairs of integers (j_1, j_2) such that $10k < j_1, j_2 \leqslant 60k$ and for which

$$j_1(1, 0) + j_2(0, 1) - j_3(1, 1) \in -S(I', \mathbf{X}, \mathbf{Y}, \mathbf{Z})$$

has size $|T| \gg k^2$. In particular there is a set T_1 of integers j_1 in the range $10k < j_1 \leqslant 60k$ of size $|T_1| \gg k$ such that for each $j_1 \in T_1$ there are $\gg k$ integers j_2 in the same range $10k < j_2 \leqslant 60k$ such that $(j_1, j_2) \in T$. Thus by two further computations along the lines of (3.28), independently with probability $\gg 1$ we can find $j_1 \in T_1$ such that $X_{j_1} > 0$, and then j_2 such that $(j_1, j_2) \in T$ and such that $Y_{j_2} > 0$.

But then by definition of $S(I', \mathbf{X}, \mathbf{Y}, \mathbf{Z})$ we can find $(x_j)_{j\in I'}, (y_j)_{j\in I'}, (z_j)_{j\in I'}$ such that $0 \leqslant x_j \leqslant X_j$, $0 \leqslant y_j \leqslant Y_j$, and $0 \leqslant z_j \leqslant Z_j$ for all $j \in I'$ and such that

$$j_1 + \sum_{j\in I'} jx_j = j_2 + \sum_{j\in I'} jy_j = j_3 + \sum_{j\in I'} jz_j.$$

Thus the proposition follows from putting $x_{j_1} = y_{j_2} = z_{j_3} = 1$, and putting all other x_j, y_j, z_j with $j > k$ equal to 0. □

Corollary 3.7.9. *$\mathscr{L}(\mathbf{X}) \cap \mathscr{L}(\mathbf{Y}) \cap \mathscr{L}(\mathbf{Z})$ is almost surely infinite.*

Proof. Define k_1 to be sufficiently large, and thereafter $k_{i+1} = (60k_i)^{1/\beta}$. Then the intervals $I_i = (k_i^\beta, 60k_i]$ are pairwise disjoint and by the proposition for each the probability that we can find $(x_j)_{j\in I_i}, (y_j)_{j\in I_i}, (z_j)_{j\in I_i}$ not all zero such that $0 \leqslant x_j \leqslant X_j$, $0 \leqslant y_j \leqslant Y_j$, and $0 \leqslant z_j \leqslant Z_j$ for each $j \in I_i$ and

$$\sum_{j\in I_i} jx_j = \sum_{j\in I_i} jy_j = \sum_{j\in I_i} jz_j$$

is bounded away from zero. Since these events are independent for different values of i the corollary follows. □

Proof of Proposition 3.7.1. By Corollary 3.7.9 there is some $k_0 = k_0(\varepsilon)$ such that

$$\mathscr{L}(\mathbf{X}) \cap \mathscr{L}(\mathbf{Y}) \cap \mathscr{L}(\mathbf{Z}) \cap [1, k_0]$$

is nonempty with probability at least $1-\varepsilon/2$. Thus by Lemma 3.1.1 there is some $n_0 = n_0(\varepsilon)$ such that if $n \geqslant n_0$ and $\mathbf{c}_1, \mathbf{c}_2, \mathbf{c}_3$ are the cycle types of random permutations $\pi_1, \pi_2, \pi_3 \in \mathcal{S}_n$ then

$$\mathscr{L}(\mathbf{c}_1) \cap \mathscr{L}(\mathbf{c}_2) \cap \mathscr{L}(\mathbf{c}_3) \cap [1, k_0]$$

is nonempty with probability at least $1 - \varepsilon$. But this means that π_1, π_2, π_3 each fix a set of size ℓ for some ℓ in the range $1 \leqslant \ell \leqslant k_0$, as claimed. □

References

[AT92] R. Arratia and S. Tavaré. "The cycle structure of random permutations". *Ann. Probab.* 20.3 (1992), pp. 1567–1591. ISSN: 0091-1798. URL: http://links.jstor.org/sici?sici=0091-1798(199207)20:3%3C1567:TCSORP%3E2.0.CO;2-T&origin=MSN.

[Bes34] A. S. Besicovitch. "On the density of certain sequences of integers". *Math. Ann.* 110 (1934), pp. 336–341.

[BW15] J. Britnell and M. Wildon. *Computing derangement probabilities of the symmetric group acting on k-sets*. 2015. URL: http://arxiv.org/abs/1511.04106.

[DS00] J. H. Davenport and G. C. Smith. "Fast recognition of alternating and symmetric Galois groups". *J. Pure Appl. Algebra* 153.1 (2000), pp. 17–25. ISSN: 0022-4049. DOI: 10.1016/S0022-4049(99)00078-X.

[DFG08] P. Diaconis, J. Fulman, and R. Guralnick. "On fixed points of permutations". *J. Algebraic Combin.* 28.1 (2008), pp. 189–218. ISSN: 0925-9899. DOI: 10.1007/s10801-008-0135-2.

[Dix92] J. D. Dixon. "Random sets which invariably generate the symmetric group". *Discrete Math.* 105.1-3 (1992), pp. 25–39. ISSN: 0012-365X. DOI: 10.1016/0012-365X(92)90129-4.

[Dix69] J. D. Dixon. "The probability of generating the symmetric group". *Math. Z.* 110 (1969), pp. 199–205. ISSN: 0025-5874.

[EFG15a] S. Eberhard, K. Ford, and B. Green. *Invariable generation of the symmetric group*. 2015. URL: http://arxiv.org/abs/1508.01870.

[EFG15b] S. Eberhard, K. Ford, and B. Green. "Permutations fixing a k-set". *IMRN* (2015). DOI: 10.1093/imrn/rnv371.

[EFK16] S. Eberhard, K. Ford, and D. Koukoulopoulos. *Permutations contained in transitive subgroups*. In preparation. 2016.

[For08a] K. Ford. "Integers with a divisor in $(y, 2y]$". *Anatomy of integers*. Vol. 46. CRM Proc. Lecture Notes. Amer. Math. Soc., Providence, RI, 2008, pp. 65–80.

[For08b] K. Ford. "The distribution of integers with a divisor in a given interval". *Ann. of Math. (2)* 168.2 (2008), pp. 367–433. ISSN: 0003-486X. DOI: 10.4007/annals.2008.168.367.

[Gra06] A. Granville. "Cycle lengths in a permutation are typically Poisson". *Electron. J. Combin.* 13.1 (2006), Research Paper 107, 23. ISSN: 1077-8926. URL: http://www.combinatorics.org/Volume_13/Abstracts/v13i1r107.html.

[Gra09] A. Granville. *The Anatomy of Integers and Permutations*. 2009. URL: http://www.dms.umontreal.ca/~andrew/PDF/Anatomy.pdf.

[GG09] A. Granville and J. Granville. *The Anatomy of Integers and Permutations*. 2009. URL: http://www.dms.umontreal.ca/~andrew/PDF/MSIProgram.pdf.

[KZ12] E. Kowalski and D. Zywina. "The Chebotarev invariant of a finite group". *Exp. Math.* 21.1 (2012), pp. 38–56. ISSN: 1058-6458. DOI: 10.1080/10586458.2011.565261.

[ŁP93] T. Łuczak and L. Pyber. "On random generation of the symmetric group". *Combin. Probab. Comput.* 2.4 (1993), pp. 505–512. ISSN: 0963-5483. DOI: 10.1017/S0963548300000869.

[MT84] H. Maier and G. Tenenbaum. "On the set of divisors of an integer". *Invent. Math.* 76.1 (1984), pp. 121–128. ISSN: 0020-9910. DOI: 10.1007/BF01388495.

[Mus78] D. R. Musser. "On the efficiency of a polynomial irreducibility test". *J. Assoc. Comput. Mach.* 25.2 (1978), pp. 271–282. ISSN: 0004-5411. DOI: 10.1145/322063.322071.

[PPR14] R. Pemantle, Y. Peres, and I. Rivin. *Four random permutations conjugated by an adversary generate S_n with high probability*. 2014. URL: http://arxiv.org/abs/1412.3781.

[RS99] A. Raouj and A. Stef. "Sur la proximité des diviseurs des entiers". *J. Number Theory* 76.1 (1999), pp. 66–93. ISSN: 0022-314X. DOI: 10.1006/jnth.1998.2355. URL: http://dx.doi.org/10.1006/jnth.1998.2355.

[Sou13] K. Soundararajan. "Ramanujan and the anatomy of integers, Srinivasa Ramanujan: going strong at 125, Part II". *Notices Amer. Math. Soc.* 60.1 (2013). Krishnaswami Alladi, Editor, pp. 10–23. ISSN: 0002-9920. DOI: 10.1090/noti926.

[TV10] T. Tao and V. H. Vu. *Additive combinatorics*. Vol. 105. Cambridge Studies in Advanced Mathematics. Paperback edition [of MR2289012]. Cambridge: Cambridge University Press, 2010, pp. xviii+512. ISBN: 978-0-521-13656-3.

Chapter 4

Additive triples of bijections

Abstract. We prove an asymptotic for the number of additive triples of bijections $\{1,\dots,n\} \to \mathbf{Z}/n\mathbf{Z}$, that is, the number of pairs of bijections

$$\pi_1, \pi_2 : \{1,\dots,n\} \to \mathbf{Z}/n\mathbf{Z}$$

such that the pointwise sum $\pi_1 + \pi_2$ is also a bijection. Specifically we show that the number of such pairs is

$$(e^{-1/2} + o(1))n!^3/n^{n-1}.$$

This problem appears in the literature under several different names: counting the number of orthomorphisms or complete mappings of $\mathbf{Z}/n\mathbf{Z}$, counting the number of arrangements of n mutually nonattacking semi-queens on an $n \times n$ toroidal chessboard, and counting the number of transversals in a cyclic Latin square. Our asymptotic supersedes several previous bounds, and settles conjectures of Vardi and Wanless.

The method of proof is a version of the Hardy–Littlewood circle method from analytic number theory, adapted to the group $(\mathbf{Z}/n\mathbf{Z})^n$. By Fourier analysis the problem is equivalent to estimating the sum of the cubes of the Fourier coefficients of the set of bijections, thought of as a subset of $(\mathbf{Z}/n\mathbf{Z})^n$. We split the universe of Fourier coefficients into four different regions and use a range of techniques to estimate the total contribution from each region.

This chapter reproduces the paper [EMM15], which is joint work with Freddie Manners and Rudi Mrazović.

4.1 Introduction

Let S be the set of bijections $\{1, \dots, n\} \to \mathbf{Z}/n\mathbf{Z}$, thought of as a subset of the group $(\mathbf{Z}/n\mathbf{Z})^n$. We are interested in counting the number s_n of additive triples in S, that is, the number of pairs of bijections $\pi_1, \pi_2 : \{1, \dots, n\} \to \mathbf{Z}/n\mathbf{Z}$ such that the pointwise sum $\pi_1 + \pi_2$ is also a bijection.

The number s_n has been studied somewhat extensively, but under a different guise. Since S is invariant under precomposition with permutations of $\{1, \dots, n\}$ it is easy to see by arbitrarily identifying $\{1, \dots, n\}$ with $\mathbf{Z}/n\mathbf{Z}$ that $s_n/n!$ is number of permutations π of $\mathbf{Z}/n\mathbf{Z}$ such that $x \mapsto \pi(x) - x$ is also a permutation. Such maps are called orthomorphisms or complete mappings, and afford the following fun interpretation. Define a semiqueen to be a chess piece which can move any distance horizontally, vertically, or diagonally in the northeast-southwest direction. Then orthomorphisms represent ways of arranging n mutually nonattacking semiqueens on an $n \times n$ toroidal chessboard. Thus $s_n/n!$ is the number of such arrangements.

In another guise this problem is that of counting the number of transversals of a cyclic Latin square. Here a transversal of an $n \times n$ Latin square is a set of n squares with no two sharing the same row, column, or symbol, and the cyclic Latin square is the Latin square with (i, j) entry given by $i + j \pmod n$. Then as above the number of such transversals is $s_n/n!$.

If n is even then $s_n = 0$. Indeed, in this case for any bijection $\pi : \{1, \dots, n\} \to \mathbf{Z}/n\mathbf{Z}$ we have $\sum_{i=1}^{n} \pi(i) = n/2$, so the sum of two bijections is never again a bijection. If n is odd then it is easy to see that $s_n > 0$, but estimating s_n is not easy.

In 1991, Vardi [Var91] made a conjecture equivalent to the following.

Conjecture 4.1.1. There are constants $c_1 > 0$ and $c_2 < 1$ such that for any large enough odd number n we have

$$c_1^n n!^2 \leqslant s_n \leqslant c_2^n n!^2.$$

The upper bound in this conjecture is known, and various authors have made incremental improvements to the constant c_2. Cooper and Kovalenko [CK96] showed that $c_2 = e^{-0.08854}$ is acceptable, and this was later improved by Kovalenko [Kov96] to $c_2 = 1/\sqrt{2}$ and by McKay, McLeod and Wanless [MMW06] to $c_2 = 0.614$. More recently, Taranenko [Tar15] proved that one can take $c_2 = 1/e + o(1)$. Glebov and Luria [GL15] proved the same bound using a somewhat simpler method based on entropy.

There has been much less progress on the lower bound, though some nontrivial lower bounds for s_n have been proved under various arithmetic assumptions about n. For example, Cooper [Coo00] proved that $s_n \geqslant e^{\frac{1}{2}\sqrt{n}\log n} n!$ under the hypothesis that n has a divisor of size roughly $\sqrt{n}$, while if, say, n is prime then a lower bound of Rivin, Vardi, and Zimmermann [RVZ94] for the torodial queens problem gives $s_n \geqslant 2^{\sqrt{(n-1)/2}} n!$. These lower bounds were superseded by work of Cavenagh and Wanless [CW10], which showed that $s_n \geqslant 3.246^n n!$ for all odd n. However, this lower bound is still a long way from the one in Conjecture 4.1.1.

Some researchers have also tried investigating the growth rate of s_n numerically: see Cooper, Gilchrist, Kovalenko, and Novaković [Coo+99] and Kuznetsov [Kuz07; Kuz08; Kuz09]. Based on this numerical evidence, Wanless [Wan11] conjectured that we can take $c_1, c_2 = 1/e + o(1)$; specifically he conjectured that

$$\lim_{n\to\infty} \frac{1}{n} \log(s_n/n!^2) = -1.$$

Finally, we should mention that $s_n/(n \cdot n!)$ is sequence A003111 in [Slo].

In this chapter we prove Vardi's conjecture with the optimal values $c_1, c_2 = 1/e + o(1)$, thus also confirming Wanless's conjecture. In fact our estimate is much more precise: we compute s_n up to a factor of $1 + o(1)$.

Theorem 4.1.2. *Let n be an odd integer. Then*

$$s_n = (e^{-1/2} + o(1)) n!^3 / n^{n-1}.$$

Perhaps the most surprising feature of this estimate is the appearance of the constant $e^{-1/2}$. This constant arises as the sum of the singular series in an argument resembling the circle method from analytic number theory, but it can be rationalized heuristically as follows. If $\pi_1, \pi_2 : \{1, \dots, n\} \to \mathbf{Z}/n\mathbf{Z}$ are random bijections then the sum $f = \pi_1 + \pi_2$ is something like a random function subject to $\sum_{x \in \mathbf{Z}/n\mathbf{Z}} f(x) = 0$, and so we might guess that the probability $\pi_1 + \pi_2$ is a bijection is about $n \cdot n!/n^n$. However if we fix two elements $x, y \in \mathbf{Z}/n\mathbf{Z}$ then the difference

$$f(x) - f(y) = (\pi_1(x) - \pi_1(y)) + (\pi_2(x) - \pi_2(y))$$

is the sum of two uniformly random nonzero elements, so $f(x) = f(y)$ with probability $1/(n-1)$, not $1/n$. Thus the probability that $f(x) \neq f(y)$ is smaller than we previously suggested by a factor of

$$\frac{1 - 1/(n-1)}{1 - 1/n} = 1 - 1/n^2 + O(1/n^3).$$

There are a total of $n^2/2 + O(n)$ pairs x, y, so we might guess that the probability $\pi_1 + \pi_2$ is a bijection is more like

$$\left(1 - 1/n^2\right)^{n^2/2} n \cdot n!/n^n \approx e^{-1/2} n!/n^{n-1}.$$

Our proof applies with only notational modifications when $\mathbf{Z}/n\mathbf{Z}$ is replaced by any abelian group G of odd order. To be precise, given a finite abelian group G, let $s(G)$ be the number of pairs of bijections $\pi_1, \pi_2 : \{1, \ldots, |G|\} \to G$ such that $\pi_1 + \pi_2$ is also a bijection. Then by the same method which proves Theorem 4.1.2, we have the following theorem.

Theorem 4.1.3. *Let G be a finite abelian group of odd order n. Then*

$$s(G) = (e^{-1/2} + o(1))n!^3/n^{n-1}.$$

There are additional complications however when G is allowed to have even order, say when $G = (\mathbf{Z}/2\mathbf{Z})^d$, and we do not know whether the same method can be made to work in this case.

Like the toroidal semiqueens problem, the toroidal queens problem—the problem of counting the number of ways of arranging n mutually nonattacking queens on an $n \times n$ toroidal chessboard—is equivalent to counting the number of solutions to a particular linear system in the set of bijections, namely the system

$$\pi_1 + \pi_2 = \pi_3,$$
$$\pi_1 - \pi_2 = \pi_4.$$

Linear systems of complexity 2 like this one are known not to be controlled by Fourier analysis alone, but there is some hope that one could use the higher-order theory pioneered by Gowers. This is an interesting avenue which has not yet been fully explored.

Notation. All the groups we work with will be finite, and we equip each of them with either the uniform probability measure or the counting measure, depending on whether the group is considered to be on the physical side or the frequency side, respectivey. Of course, this convention is respected in our definitions of convolution, inner product, and L^2-norm. When G has the uniform measure we use the expectation notation $\mathbf{E}_{x \in G} f(x)$ to denote the average of f over G.

Occasionally we will use a primed sum $\sum'_{x_1, \ldots, x_k \in G}$ to denote the sum over all k-tuples of distinct elements $x_1, \ldots, x_k \in G$. Finally, we will use the standard notation $e(x) = e^{i2\pi x}$.

4.2 Outline of the proof

Our approach to proving Theorem 4.1.2 is Fourier-analytic. Write $G = \mathbf{Z}/n\mathbf{Z}$ for n an odd integer, and write $S \subset G^n$ for the set of bijections $\{1, \dots, n\} \to G$. Our goal is to compute the quantity

$$\mathbf{E}_{x,y \in G^n} 1_S(x) 1_S(y) 1_S(x+y) = \langle 1_S * 1_S, 1_S \rangle.$$

A standard application of Parseval's identity shows that this quantity can be expressed in terms of the Fourier transform of 1_S as

$$\sum_{\chi \in \widehat{G}^n} |\widehat{1_S}(\chi)|^2 \widehat{1_S}(\chi). \tag{4.1}$$

Here we identify $\widehat{G}$ with $\frac{1}{n}\mathbf{Z}/\mathbf{Z}$ in the usual way, so that for $\chi = (r_1, \dots, r_n) \in (\frac{1}{n}\mathbf{Z}/\mathbf{Z})^n$ we have explicitly

$$\widehat{1_S}(r_1, \dots, r_n) = \frac{1}{n^n} \sum_{x_1, \dots, x_n}{}' e(-r_1 x_1 - \cdots - r_n x_n), \tag{4.2}$$

where, as mentioned in the previous section, we use $\sum'$ to denote the sum over distinct $x_1, \dots, x_n \in G$, i.e., the sum over S. In fact it is clear that $\widehat{1_S}$ is real-valued, since $S = -S$, and hence we may drop the absolute value signs in (4.1).

The form of the proof can then be viewed as an analogue of the Hardy–Littlewood circle method from analytic number theory, adapted to the group G^n rather than $\mathbf{Z}$. We borrow some nomenclature from the classical setting.

- There are a small number of characters $\chi \in \widehat{G}^n$, namely the characters $\chi = (r_1, \dots, r_n)$ for which almost all of the r_i are equal, that make a substantial contribution to the sum (4.1). We call the totality of these χ the *major arcs*. We compute explicitly the contribution of these χ to (4.1), up to small errors: this is an analogue of the singular series, which accounts for the main term in Theorem 4.1.2, including the constant $e^{-1/2}$. For all this see Section 4.3.

- We bound the contribution from all other characters using the triangle inequality; that is, we obtain an upper bound for

$$\sum_{\text{all other } \chi} |\widehat{1_S}(\chi)|^3$$

that is smaller than the main term. We call such χ collectively the *minor arcs*.

A large part of this chapter is therefore devoted to obtaining good bounds for Fourier coefficients $\widehat{1_S}(\chi)$, either pointwise or on average in a suitable sense. In fact we will need to combine several different arguments which are effective in different regimes.

4.2.1 Preliminaries on $\widehat{1_S}$

In order to describe how the characters $\widehat{G}^n$ are divided up into different pieces in which different approaches will be effective, we will need some preliminary remarks.

First, note that the function $\widehat{1_S}(r_1, \dots, r_n)$ is invariant under permutation of the r_i, i.e.,

$$\widehat{1_S}(r_1, \dots, r_n) = \widehat{1_S}(r_{\sigma(1)}, \dots, r_{\sigma(n)})$$

for any bijection $\sigma : \{1, \dots, n\} \to \{1, \dots, n\}$. This is immediate from the definition (4.2). Hence it makes sense to specify a Fourier coefficient of interest in the form $\widehat{1_S}(r_1^{a_1}, \dots, r_k^{a_k})$, where $r_i \in \widehat{G}$ are distinct, the a_i are positive integers such that $\sum a_i = n$, and the notation r^a means r repeated a times.

Second, observe that $\widehat{1_S}(r_1, \dots, r_n) = 0$ unless $\sum r_i = 0$. This follows from the fact that S is invariant under global shifts $(x_i) \mapsto (x_i + t)$, so one can compute

$$\begin{aligned}
\widehat{1_S}(r_1, \dots, r_n) &= \frac{1}{n^n} \sum_{x_1, \dots, x_n}{}' e(-r_1 x_1 - \cdots - r_n x_n) \\
&= \frac{1}{n^n} \sum_{x_1, \dots, x_n}{}' e(-r_1(x_1 + t) - \cdots - r_n(x_n + t)) \\
&= \widehat{1_S}(r_1, \dots, r_n) e(-(r_1 + \cdots + r_n)t).
\end{aligned}$$

Dually, note that $\widehat{1_S}$ is invariant under global shifts. This follows from the fact that $\sum x_i = 0$ for any $(x_i) \in S$. Hence,

$$\begin{aligned}
\widehat{1_S}(r_1, \dots, r_n) &= \frac{1}{n^n} \sum_{x_1, \dots, x_n}{}' e(-r_1 x_1 - \cdots - r_n x_n) \\
&= \frac{1}{n^n} \sum_{x_1, \dots, x_n}{}' e(-(r_1 + t)x_1 - \cdots - (r_n + t)x_n) \\
&= \widehat{1_S}(r_1 + t, \dots, r_n + t).
\end{aligned}$$

Finally, note the trivial bound

$$|\widehat{1_S}(\chi)| \leqslant \widehat{1_S}(0) = \frac{n!}{n^n}.$$

4.2.2 Entropy ranges for minor arcs

We now explain a straightforward but still fairly powerful bound on $|\widehat{1_S}(\chi)|$ that follows directly from these elementary considerations.

Proposition 4.2.1 ("Entropy bound")**.** *We have*

$$|\widehat{1_S}(r_1^{a_1}, \dots, r_k^{a_k})| \leqslant \binom{n}{a_1, \dots, a_k}^{-1/2} \left(\frac{n!}{n^n}\right)^{1/2}.$$

Proof. Let $\mathcal{O} \subset \widehat{G}^n$ denote the set of characters obtained by permuting the elements of $(r_1^{a_1}, \dots, r_k^{a_k})$. Hence, $|\mathcal{O}| = \binom{n}{a_1,\dots,a_k}$ and by permutation invariance, $\widehat{1_S}$ takes a constant value on $\mathcal{O}$. Thus by Parseval,

$$\frac{n!}{n^n} = \sum_{\chi \in \widehat{G}^n} |\widehat{1_S}(\chi)|^2 \geqslant \sum_{\chi \in \mathcal{O}} |\widehat{1_S}(\chi)|^2 = |\mathcal{O}||\widehat{1_S}(r_1^{a_1}, \dots, r_k^{a_k})|^2,$$

and the result follows. □

The term "entropy bound" refers to the fact that the quantity

$$H(\chi) = \frac{1}{n} \log \binom{n}{a_1, \dots, a_k}$$

is roughly the entropy $\sum_{i=1}^k (a_i/n) \log(n/a_i) \in [0, \log n]$ of a random variable taking the value r_i with probability a_i/n. In a slight abuse of nomenclature, we refer to this first quantity $H(\chi)$ as the *entropy of* χ. In this language, the bound of Proposition 4.2.1 is precisely $\exp(-Hn/2)(n!/n^n)^{1/2}$ or roughly $\exp(-Hn/2 - n/2)$. This is already sufficient to control the contribution to (4.1) for most characters χ.

Corollary 4.2.2. *We have*

$$\sum_{\chi : H(\chi) \geqslant R} |\widehat{1_S}(\chi)|^3 \leqslant \exp((3-R)n/2) \left(\frac{n!}{n^n}\right)^3.$$

Proof. By Parseval and Proposition 4.2.1,

$$\begin{aligned}
\sum_{\chi : H(\chi) \geqslant R} |\widehat{1_S}(\chi)|^3 &\leqslant \left(\sum_{\chi : H(\chi) \geqslant R} |\widehat{1_S}(\chi)|^2\right) \sup_{\chi : H(\chi) \geqslant R} |\widehat{1_S}(\chi)| \\
&\leqslant \left(\sum_{\chi \in \widehat{G}^n} |\widehat{1_S}(\chi)|^2\right) \sup_{\chi : H(\chi) \geqslant R} \exp(-H(\chi)n/2) \left(\frac{n!}{n^n}\right)^{1/2} \\
&\leqslant \exp(-Rn/2) \left(\frac{n!}{n^n}\right)^{3/2} \\
&\leqslant \exp((3-R)n/2) \left(\frac{n!}{n^n}\right)^3
\end{aligned}$$

as required, where we have used the estimate $n! > (n/e)^n$. □

This is smaller than the main term for any fixed $R > 3$, so from now on such high-entropy characters need not concern us. We refer to such characters as the *high-entropy minor arcs*. Since a typical character χ has entropy comparable to $\log n$, the high-entropy case comprises almost all characters in some sense, so this is a good first step.

On the opposite extreme we have characters χ with entropy $o(1)$, such as $\chi = \left(r^k, -r^k, 0^{n-2k}\right)$ for $k = o(n)$. The characters with entropy $O\left(\frac{\log n}{n}\right)$ are the major arcs, but between $O\left(\frac{\log n}{n}\right)$ and $o(1)$ we have the *low-entropy minor arcs*. Such characters are necessarily of the form $(\dots, r^{n-o(n)})$, i.e., they take a single value almost all the time. By shift-invariance we may assume without loss of generality that this value r is zero. Thus the low-entropy minor arcs are closely related to the set of *sparse* characters, characters χ comprised almost entirely of zeros.

That leaves the characters χ with $o(1) \leqslant H(\chi) \leqslant 3$, which form the *medium-entropy minor arcs*. A good model case are characters such as $\chi = (r^{n/3}, -r^{n/3}, 0^{n/3})$, which are particularly troublesome.

There are two fundamental areas of inefficiency in the arguments of Proposition 4.2.1 and Corollary 4.2.2 that prevent them from giving good bounds for low- or medium-entropy minor arcs. First, the bound on $|\widehat{1_S}|$ given by Proposition 4.2.1 is less effective for smaller H, and in particular worse than trivial when $H < 1$. Clearly there is no hope for this programme unless we can find a bound on $|\widehat{1_S}|$ that is nontrivial throughout the range $0 < H < 1$, and moreover obtains a substantial exponential saving for most of that range.

Second, in the proof of Corollary 4.2.2 we made use of the convenient bound

$$\sum_{\chi \in X} |\widehat{1_S}(\chi)|^3 \leqslant \left(\sum_{\chi \in \widehat{G}} |\widehat{1_S}(\chi)|^2 \right) \sup_{\chi \in X} |\widehat{1_S}(\chi)|$$

for some appropriate set $X \subset \widehat{G}^n$. It is fairly clear that we cannot afford this luxury for small entropies: since the main term has order $(n!/n^n)^3$, we need $|\widehat{1_S}(\chi)| = o((n!/n^n)^2)$ for all $\chi \in X$ for this to be effective. This is a saving of about $\exp(-n)$ over the trivial bound of $n!/n^n$, and it is simply not reasonable to expect such a bound to hold if, say, $H(\chi) = 1/1000$. One way around this is to pigeonhole the characters χ into various sets X_i, and apply such an $L^2 \cdot L^\infty$ bound on each set. To make this work, we would need a bound on

$$\sum_{\chi \in X_i} |\widehat{1_S}(\chi)|^2$$

which makes a significant saving over the crude estimate

$$\sum_{\chi \in X_i} |\widehat{1_S}(\chi)|^2 \leqslant \sum_{\chi \in \widehat{G^n}} |\widehat{1_S}(\chi)|^2$$

employed in Corollary 4.2.2. Alternatively, one could try to bound the L^3 sum

$$\sum_{\chi \in X_i} |\widehat{1_S}(\chi)|^3$$

directly, using an L^∞ bound and an estimate for the size of X_i.

We address both of these inefficiencies in order to reach our final bound. Specifically, we will do each of the following.

- Improving on Proposition 4.2.1, we obtain a good general-purpose bound for $|\widehat{1_S}(\chi)|$ that is nontrivial for values of $H(\chi)$ approaching zero. Roughly speaking, where Proposition 4.2.1 gives the bound $e^{-Hn/2-n/2}$ we will prove the bound $e^{-Hn/2-n}$. This appears in Section 4.4.

- We use this bound to estimate the contribution to (4.1) from χ with $10^{-10} < H(\chi) < 10$, say, by using a dyadic decomposition and a simple estimate for the number of characters of a given entropy. Thus we dispatch the medium-entropy minor arcs. This is covered in Section 4.7.

- Separately, we obtain a good estimate for

  $$\sum_{m\text{-sparse } \chi} |\widehat{1_S}(\chi)|^2,$$

  i.e., an improvement to Parseval when summing only over those characters χ having exactly m nonzero terms. Here we might imagine $m \leqslant n/1000$. We call this a *sparseval* bound, and prove this in Section 4.5.

- Finally, we obtain a slightly different L^∞ bound specialized to the case of sparse characters χ. This appears in Section 4.6. The total contribution from the low-entropy minor arcs is controlled by combining this with the L^2 bound above. (Incidentally, this is the only part of the argument in which we really use the odd order assumption—elsewhere we only use the assumption that the sum of the elements in G is zero.)

In summary, we split the universe of all χ into several slightly overlapping regions—major arcs, low-entropy minor arcs, medium-entropy minor arcs, and high-entropy minor arcs—and apply a cocktail of different bounds and explicit computations adapted to each region.

4.2.3 Interpretations of minor arc bounds

As we have said, a significant part of our effort will be spent proving nontrivial bounds on Fourier coefficients $|\widehat{1_S}(\chi)|$ for χ in the minor arcs. Although such bounds are at first glance fairly esoteric, in fact they have natural interpretations. For instance, they are intimately related to questions of the following flavour.

Question 4.2.3. Suppose $G = \mathbf{Z}/n\mathbf{Z}$ is partitioned into k sets $A_1, \dots, A_k$, each of a pre-determined size $a_i \approx n/k$, at random. Consider the random variables

$$T_i = \sum_{r \in A_i} r \in G$$

for $i = 1, \dots, k$. The vector $T = (T_1, \dots, T_k)$ is a random variable taking values in the subgroup $H = \{(x_i) \in G^k : \sum x_i = 0\}$ of G^k. How close is T to being equidistributed in H, in a quantitative sense? In other words how close is the law of T to the uniform distribution on H?

Because of the connection of this question to the size of the Fourier coefficients $\widehat{1_S}(r_1^{a_1}, \dots, r_k^{a_k})$, our bounds give a strong answer to this question in many regimes. It is not inconceivable that these kind of bounds may find applications elsewhere.

4.3 Major arcs

In this section we compute $\widehat{1_S}(r_1, \dots, r_m, 0, \dots, 0)$ explicitly for bounded m. To this end observe that

$$\widehat{1_S}(r_1, \dots, r_m, 0, \dots, 0) = \frac{(n-m)!}{n^n} \sum_{x_1, \dots, x_m}{}' e(-r_1 x_1 - \dots - r_m x_m). \tag{4.3}$$

The sum over distinct $x_1, \dots, x_m$ can be related to a sum over partitions of $\{1, \dots, m\}$ using a type of Möbius inversion: for any function $F(x_1, \dots, x_m)$ we have

$$\sum_{x_1, \dots, x_m}{}' F(x_1, \dots, x_m) = \sum_{\mathcal{P}} \mu(\mathcal{P}) \sum_{\substack{x_1, \dots, x_m \\ x_i = x_j \text{ whenever } i \overset{\mathcal{P}}{\sim} j}} F(x_1, \dots, x_m),$$

where the outer sum runs over partitions $\mathcal{P}$ of $\{1, \dots, m\}$, and

$$\mu(\mathcal{P}) = (-1)^{m - |\mathcal{P}|} \prod_{P \in \mathcal{P}} (|P| - 1)!.$$

See, e.g., Bóna [Bón15] for more information. To apply this to $\widehat{1_S}$, say that a partition $\mathcal{P}$ of $\{1,\dots,m\}$ *kills* $(r_1,\dots,r_m)$ if $\sum_{i\in P} r_i = 0$ for each $P \in \mathcal{P}$, and observe that

$$\sum_{\substack{x_1,\dots,x_m \\ x_i=x_j \text{ whenever } i\overset{\mathcal{P}}{\sim}j}} e(-r_1x_1 - \cdots - r_mx_m) = \begin{cases} n^{|\mathcal{P}|} & \text{if } \mathcal{P} \text{ kills } (r_1,\dots,r_m), \\ 0 & \text{otherwise,} \end{cases}$$

so

$$\widehat{1_S}(r_1,\dots,r_m,0,\dots,0) = \frac{(n-m)!}{n^n} \sum_{\mathcal{P} \text{ killing } (r_1,\dots,r_m)} \mu(\mathcal{P}) n^{|\mathcal{P}|}.$$

The following two observations are immediate from this formula:

- Suppose every killing partition of $(r_1,\dots,r_m)$ has at most k parts. Then

$$|\widehat{1_S}(r_1,\dots,r_m,0,\dots,0)| \leqslant O_m\left(\frac{1}{n^{m-k}}\frac{n!}{n^n}\right). \tag{4.4}$$

- Suppose that m is even, that $r_1,\dots,r_m \in \widehat{G}$ are nonzero, and that $(r_1,\dots,r_m)$ is killed by a unique partition with $m/2$ parts. In other words suppose that $r_1,\dots,r_m \in \widehat{G}$ are distinct and nonzero, and that, up to a permutation, $r_{2j} = -r_{2j-1}$ for each $j = 1,\dots,m/2$. Then

$$\widehat{1_S}(r_1,\dots,r_m,0,\dots,0) = \frac{(-1)^{m/2}}{n^{m/2}}\frac{n!}{n^n} + O_m\left(\frac{1}{n^{m/2+1}}\frac{n!}{n^n}\right). \tag{4.5}$$

Proposition 4.3.1. *If m is even then*

$$\sum_{m\text{-sparse } \chi} \widehat{1_S}(\chi)^3 = \frac{(-1)^{m/2}}{2^{m/2}(m/2)!}\frac{n!^3}{n^{3n}} + O_m\left(\frac{1}{n}\frac{n!^3}{n^{3n}}\right),$$

while if m is odd then

$$\sum_{m\text{-sparse } \chi} \widehat{1_S}(\chi)^3 = O_m\left(\frac{1}{n}\frac{n!^3}{n^{3n}}\right).$$

Proof. First of all note by permutation-invariance that

$$\sum_{m\text{-sparse } \chi} \widehat{1_S}(\chi)^3 = \binom{n}{m} \sum_{r_1,\dots,r_m \neq 0} \widehat{1_S}(r_1,\dots,r_m,0,\dots,0)^3. \tag{4.6}$$

For each $(r_1,\dots,r_m)$ choose a maximal-size partition $\mathcal{P}$ which kills $(r_1,\dots,r_m)$, and split the sum up according to $\mathcal{P}$. Suppose $\mathcal{P}$ has k parts. Since $\mathcal{P}$ cannot have singletons we have $k \leqslant m/2$, and by (4.4) we have

$$|\widehat{1_S}(r_1,\dots,r_m,0,\dots,0)| \leqslant O_m\left(\frac{1}{n^{m-k}}\frac{n!}{n^n}\right).$$

Since the number of $(r_1, \dots, r_m)$ killed by $\mathcal{P}$ is bounded by n^{m-k}, the contribution to (4.6) from these $(r_1, \dots, r_m)$ is bounded by

$$\binom{n}{m} n^{m-k} \left(\frac{O_m(1)}{n^{m-k}} \frac{n!}{n^n} \right)^3 = O_m \left(\frac{1}{n^{m-2k}} \frac{n!^3}{n^{3n}} \right),$$

which is satisfactory unless $k = m/2$. Moreover the number of $(r_1, \dots, r_m)$ killed by at least two partitions $\mathcal{P}$ with $m/2$ parts is bounded by $n^{m/2-1}$, so the contribution from these $(r_1, \dots, r_m)$ is bounded by

$$\binom{n}{m} n^{m/2-1} O_m \left(\frac{1}{n^{m/2}} \frac{n!}{n^n} \right)^3 = O_m \left(\frac{1}{n} \frac{n!^3}{n^{3n}} \right),$$

which is again satisfactory. Thus we may restrict our attention to those $(r_1, \dots, r_m)$ which are killed by a unique partition $\mathcal{P}$ with $m/2$ parts, and for these (4.5) gives

$$\widehat{1_S}(r_1, \dots, r_m, 0, \dots, 0) = \frac{(-1)^{m/2}}{n^{m/2}} \frac{n!}{n^n} + O_m \left(\frac{1}{n^{m/2+1}} \frac{n!}{n^n} \right).$$

For a fixed such $\mathcal{P}$ the number of such $(r_1, \dots, r_m)$ is $n^{m/2} + O_m(n^{m/2-1})$, and the number of such $\mathcal{P}$ is

$$(m-1)(m-3) \cdots 1 = \frac{m!}{2^{m/2}(m/2)!},$$

so the total contribution from these $(r_1, \dots, r_m)$ is

$$\begin{aligned} \binom{n}{m} \frac{m!}{2^{m/2}(m/2)!} (n^{m/2} + O_m(n^{m/2-1})) \left(\frac{(-1)^{m/2}}{n^{m/2}} \frac{n!}{n^n} + O_m \left(\frac{1}{n^{m/2+1}} \frac{n!}{n^n} \right) \right)^3 \\ = \frac{(-1)^{m/2}}{2^{m/2}(m/2)!} \frac{n!^3}{n^{3n}} + O_m \left(\frac{1}{n} \frac{n!^3}{n^{3n}} \right). \end{aligned}$$

This proves the proposition. □

4.4 Square-root cancellation for general Fourier coefficients

The aim of this section is to prove the following refinement of Proposition 4.2.1.

Theorem 4.4.1. *Suppose $\chi = (r_1^{a_1}, \dots, r_k^{a_k})$, where $\sum_{i=1}^k a_i = n$. Then*

$$|\widehat{1_S}(\chi)| \leqslant \binom{n+k-1}{k-1}^{1/2} \binom{n}{a_1, \dots, a_k}^{-1/2} \frac{n!}{n^n}.$$

Before proving this theorem we make a few remarks in the way of motivation and explanation. First, note that the factor $\binom{n+k-1}{k-1}^{1/2}$ is small in the regimes we care about, for if $H(\chi) = O(1)$ then $k = O(n/\log n) = o(n)$, so this factor is of size $\exp(o(n))$. Thus where Proposition 4.2.1 provides the bound $e^{-Hn/2-n/2}$, this theorem provides the bound $e^{-Hn/2-n+o(n)}$. This saving of around $e^{-n/2+o(n)}$ may be uninspiring if H is large, but if H is small it is decisive.

In general the bound in this bound is not sharp, often by a substantial amount. However, it is sometimes attained asymptotically for very special, arithmetically structured classes of characters. For instance, if we temporarily allow n to be even, then one can compute directly that

$$\log\left[|\widehat{1_S}((1/2)^a, 0^{n-a})|/(n!/n^n)\right] \sim -\frac{1}{2}\log\binom{n}{a}$$

for large n. Similarly, numerical evidence strongly suggests that

$$\log\left[|\widehat{1_S}((1/3)^a, (-1/3)^a, 0^{n-2a})|/(n!/n^n)\right] \sim -\frac{1}{2}\log\binom{n}{a, a, n-2a}$$

if n is divisible by 3; for example, when $n = 3003$,

$$\log\left[|\widehat{1_S}((1/3)^{1001}, (-1/3)^{1001}, 0^{1001})|/(n!/n^n)\right] = -1649.01782245...,$$

while

$$-\frac{1}{2}\log\binom{3003}{1001, 1001, 1001} = -1645.46757758....$$

But, for example, the evidence also suggests that

$$\log\left[|\widehat{1_S}((1/3)^a, 0^{n-a})|/(n!/n^n)\right] \sim -\frac{2}{3}\log\binom{n}{a}.$$

Thus we might hope to improve on Theorem 4.4.1, but only by ruling out the particular characters discussed above and others similar to them. For instance, a strengthening is almost certainly available under the assumption that n is prime. Thus one might think of Theorem 4.4.1 as the best general-purpose bound available.

We should also comment briefly on the term "square-root cancellation". If $\chi = (r_1^{a_1}, \dots, r_k^{a_k})$ then note that

$$\widehat{1_S}(\chi) = \frac{n!}{n^n}\mathbf{E}_{\mathcal{P}} e\left(-r_1\left(\sum_{x\in P_1} x\right) - \cdots - r_k\left(\sum_{x\in P_k} x\right)\right),$$

where the expectation is over all ordered partitions $\mathcal{P} = (P_1, \dots, P_k)$ of G into k pieces with $|P_i| = a_i$. The number of such partitions is precisely $\binom{n}{a_1,\dots,a_k}$. Assuming

that the phases in the average behave randomly then we should expect square-root cancellation, and we heuristically recover the bound in Theorem 4.4.1 up to lower-order terms.

Proof of Theorem 4.4.1. Let γ be the Gaussian measure on $\mathbf{C}^k$ defined with respect to Lebesgue measure λ by

$$\frac{d\gamma}{d\lambda} = \frac{1}{\pi^k} \exp\left(- \sum_{i=1}^{k} |z_i|^2 \right).$$

We claim that

$$n^n \widehat{1_S}(\chi) = \int_{\mathbf{C}^k} z_1^{a_1} \cdots z_k^{a_k} \prod_{x \in G} \left(\sum_{i=1}^{k} e(-r_i x) \overline{z_i} \right) d\gamma. \tag{4.7}$$

One can check this by expanding the product and using the identity

$$\int_{\mathbf{C}^k} \prod_{i=1}^{k} z_i^{a_i} \overline{z_i}^{b_i} \, d\gamma = \begin{cases} \prod_{i=1}^{k} a_i! & \text{if } a_i = b_i \text{ for each } i, \\ 0 & \text{otherwise.} \end{cases} \tag{4.8}$$

We can now proceed by bounding the integral in (4.7) using various techniques. Applying the Cauchy–Schwarz inequality we bound the right-hand side by

$$\left(\int_{\mathbf{C}^k} |z_1|^{2a_1} \cdots |z_k|^{2a_k} \, d\gamma \right)^{1/2} \left(\int_{\mathbf{C}^k} \prod_{x \in G} \left| \sum_{i=1}^{k} e(-r_i x) \overline{z_i} \right|^2 d\gamma \right)^{1/2}.$$

The first factor here is exactly $\left(\prod_{i=1}^{k} a_i! \right)^{1/2}$ by (4.8). If we now apply the AM–GM inequality to the second term and evaluate the resulting integral, we get

$$\begin{aligned} \int_{\mathbf{C}^k} \prod_{x \in G} \left| \sum_{i=1}^{k} e(-r_i x) \overline{z_i} \right|^2 d\gamma &\leqslant \int_{\mathbf{C}^k} \left(\frac{1}{n} \sum_{x \in G} \left| \sum_{i=1}^{k} e(-r_i x) \overline{z_i} \right|^2 \right)^n d\gamma \\ &= \int_{\mathbf{C}^k} \left(\sum_{i=1}^{k} |z_i|^2 \right)^n d\gamma \\ &= \sum_{s_1 + \cdots + s_k = n} \binom{n}{s_1, \ldots, s_k} \int_{\mathbf{C}^k} |z_1|^{2s_1} \cdots |z_k|^{2s_k} \, d\gamma \\ &= n! \sum_{s_1 + \cdots + s_k = n} 1 \\ &= n! \binom{n + k - 1}{k - 1}. \end{aligned}$$

Thus the theorem follows by dividing by n^n. □

Remark. Though the identity (4.7) is easily verified directly, we briefly sketch here a possible route by which one might be motivated to write down such a formula in the first place.

Let $\{X_x : x \in G\}$ be indeterminates, and observe that the left-hand side of (4.7) is precisely the coefficient of $\prod_{x \in G} X_x$ in the multivariate polynomial

$$\prod_{i=1}^{k} \left(\sum_{x \in G} X_x e(-r_i x) \right)^{a_i},$$

and so equivalently the value of the derivative

$$\left(\prod_{x \in G} \frac{\partial}{\partial X_x} \right) \prod_{i=1}^{k} \left(\sum_{x \in G} X_x e(-r_i x) \right)^{a_i} \tag{4.9}$$

evaluated at zero (although in fact this evaluation is redundant as this expression is a constant function).

In some generality, it is possible to equate a quantity of the form

$$\frac{\partial^{b_1}}{\partial Y_1^{b_1}} \cdots \frac{\partial^{b_k}}{\partial Y_k^{b_k}} f \big|_{Y_1 = \cdots = Y_k = 0},$$

where $f(Y_1, \ldots, Y_k)$ is a suitable holomorphic function of k complex variables, with the corresponding integral expression

$$\int_{\mathbf{C}^k} \overline{Y_1^{b_1}} \ldots \overline{Y_k^{b_k}} \, f(Y_1, \ldots, Y_k) \, d\gamma,$$

where again γ denotes Gaussian measure. This can be verified by applying the case $k = 1$ iteratively, which in turn follows by averaging the usual Cauchy integral formula over different radii. This formula can be considered a particular higher-dimensional variant of the Cauchy integral formula.

The identity (4.7) follows by applying this identity to (4.9) and making an orthogonal change of variables.

4.5 Sparseval

Recall from Section 4.2 that one of our goals is to obtain good bounds on

$$\sum_{m\text{-sparse } \chi} |\widehat{1_S}(\chi)|^2, \tag{4.10}$$

where the sum is over all characters $\chi = (r_1, \ldots, r_n)$ with precisely m nonzero entries.

First consider the related sum

$$\sum_{\leqslant m\text{-sparse } \chi} |\widehat{1_S}(\chi)|^2,$$

i.e., the sum over characters with *at most* m nonzero entries. This can be bounded by applying Parseval to the set S_m of injections $\{1,\dots,m\} \to G$ as a subset of the group G^m. Indeed, letting $N(\chi)$ be the set of indices i for which r_i is nonzero, by permutation-invariance we have

$$\begin{aligned}
\sum_{\leqslant m\text{-sparse } \chi} |\widehat{1_S}(\chi)|^2 &\leqslant \sum_{|N|=m} \sum_{N(\chi)\subset N} |\widehat{1_S}(\chi)|^2 \\
&= \binom{n}{m} \sum_{r_1,\dots,r_m} |\widehat{1_S}(r_1,\dots,r_m,0,\dots,0)|^2 \\
&= \binom{n}{m} \sum_{r_1,\dots,r_m} \frac{(n-m)!^2}{n^{2n-2m}} |\widehat{1_{S_m}}(r_1,\dots,r_m)|^2 \qquad (4.11) \\
&= \binom{n}{m} \frac{(n-m)!^2}{n^{2n-2m}} \frac{n!}{n^m(n-m)!} \\
&= \frac{n^m}{m!} \frac{n!^2}{n^{2n}}.
\end{aligned}$$

Here we used

$$\widehat{1_S}(r_1,\dots,r_m,0,\dots,0) = \frac{(n-m)!}{n^{n-m}} \widehat{1_{S_m}}(r_1,\dots,r_m),$$

which follows from (4.3).

This is our first nontrivial *sparseval* bound. It turns out that overcounting as above is too inefficient in the context of the wider argument, when, say, m/n is a small constant. The purpose of the rest of this section therefore is to improve the bound on (4.10) as much as possible.

Theorem 4.5.1. *If $m \leqslant n/2$ then*

$$\sum_{m\text{-sparse } \chi} |\widehat{1_S}(\chi)|^2 \leqslant O(m^{1/4}) e^{O(m^{3/2}/n^{1/2})} \binom{n}{m}^{1/2} \frac{n!^2}{n^{2n}}$$

Proof. The starting point of our strategy is to apply inclusion-exclusion to obtain an expression for (4.10). For $m \leqslant n$ let

$$Q(m,n) = \frac{n^{2n}}{n!^2} \sum_{m\text{-sparse } \chi} |\widehat{1_S}(r)|^2,$$

and observe as in (4.11) that

$$\begin{aligned}\frac{n^m}{m!} &= \frac{n^{2n}}{n!^2}\sum_{|N|=m}\sum_{N(\chi)\subset N}|\widehat{1_S}(\chi)|^2\\ &= \frac{n^{2n}}{n!^2}\sum_{k=0}^{m}\sum_{|N(\chi)|=k}|\widehat{1_S}(\chi)|^2\binom{n-k}{m-k}\\ &= \sum_{k=0}^{m}\binom{n-k}{m-k}Q(k,n).\end{aligned}$$

By inverting this relation we have

$$\begin{aligned}Q(m,n) &= \sum_{k=0}^{m}(-1)^{m-k}\binom{n-k}{m-k}\frac{n^k}{k!}\\ &= \frac{1}{m!}\sum_{k=0}^{m}\binom{m}{k}(-1)^k(n-m+1)\cdots(n-k)n^k. \qquad (4.12)\end{aligned}$$

This expression exhibits a vast amount of cancellation. To capture this, we use a generating function argument, followed by some complex analysis in the spirit of the saddle-point method (see Flajolet and Sedgewick [FS09, Chapter VIII] for background); see also the remark following the proof for some further motivation.

Observe from (4.12) that $Q(m,n)$ is exactly the coefficient of X^m in

$$f(X) = \frac{n^{n+1}e^X}{(n+X)^{n-m+1}}. \qquad (4.13)$$

Thus, by Cauchy's formula,

$$Q(m,n) = \frac{n^{n+1}}{i2\pi}\oint_{|z|=r}\frac{e^z}{(n+z)^{n-m+1}z^{m+1}}\,dz$$

for any r in the range $0 < r < n$, so

$$Q(m,n) \leqslant \frac{n^{n+1}}{r^m}\max_{|z|=r}\left|\frac{e^z}{(n+z)^{n-m+1}}\right|.$$

We claim that $|e^z/(n+z)^{n-m+1}|$ has only two local maxima on $|z| = r$: one at $z = +r$ and the other at $z = -r$. To see this note that if $|z| = r$ and $\Re z = t$ then

$$\left|\frac{e^z}{(n+z)^{n-m+1}}\right|^2 = \frac{e^{2t}}{(n^2+r^2+2nt)^{n-m+1}},$$

so

$$\frac{d}{dt}\log\left|\frac{e^z}{(n+z)^{n-m+1}}\right|^2 = 2 - \frac{2n(n-m+1)}{n^2+r^2+2nt}.$$

This function has a unique pole at $t = -(n + r^2/n)/2$, which is to the left of the physical region $-r \leqslant t \leqslant r$, and it has a unique zero somewhere corresponding to a minimum, not a maximum, so the claim holds. Thus $|e^z/(n+z)^{n-m+1}|$ is bounded on $|z| = r$ by

$$\frac{e^{\pm r}}{(n \pm r)^{n-m+1}},$$

so

$$Q(m,n) \leqslant \frac{n^{n+1}e^{\pm r}}{r^m(n \pm r)^{n-m+1}} = \frac{e^{\pm r}}{(1 \pm r/n)^{n-m+1}}\frac{n^m}{r^m}.$$

Here

$$\begin{aligned}\log\left(\frac{e^{\pm r}}{(1 \pm r/n)^{n-m+1}}\right) &= \pm r - (n-m+1)(\pm r/n - r^2/2n^2 + O(r^3/n^3)) \\ &= r^2/2n + O(r^3/n^2 + rm/n).\end{aligned}$$

Taking $r = (mn)^{1/2}$ we thus have, by Stirling's formula,

$$Q(m,n) \leqslant e^{(1/2+O((m/n)^{1/2}))m}\frac{n^{m/2}}{m^{m/2}} \leqslant O(m^{1/4})e^{O(m^{3/2}/n^{1/2})}\binom{n}{m}^{1/2}. \qquad \square$$

Remark. Although it is straightforward to check that $Q(m,n)$ is indeed the coefficient of X^m in $f(X)$ by expanding (4.13), the proof is unsatisfying in that one needs to know the formula for f in advance. Here is an alternative proof which does not have this fault.

Let ∂ be the discrete derivative operator defined on functions $g : \mathbf{N} \to \mathbf{R}$ by

$$\partial g(k) = g(k+1) - g(k),$$

and observe that

$$Q(m,n) = \frac{1}{m!}\partial^m g_{m,n}(0),$$

where $g_{m,n}$ is the function defined by

$$g_{m,n}(k) = n^k(n-k)(n-k-1)\cdots(n-m+1).$$

Intuitively, the reason the sum (4.12) exhibits such enormous cancellation is that it is a high derivative of a very "smooth" function $g_{m,n}$. To capture this smoothness we start by investigating the first derivative $\partial g_{m,n}$. By directly computing we find that

$$\begin{aligned}\partial g_{m,n}(k) &= kn^k(n-k-1)(n-k-2)\cdots(n-m+1) \\ &= \frac{k}{n}g_{m,n}(k+1).\end{aligned}$$

Thus from the Leibniz-type formula

$$\partial(gh)(k) = \partial g(k)h(k+1) + g(k)\partial h(k)$$

we derive the recurrence

$$\begin{aligned}\partial^\ell g_{m,n}(k) &= \partial^{\ell-1}\left(\frac{k}{n}g_{m,n}(k+1)\right)\\ &= \frac{k}{n}\partial^{\ell-1}g_{m,n}(k+1) + \frac{\ell-1}{n}\partial^{\ell-2}g_{m,n}(k+2),\end{aligned}$$

which holds for all $\ell \geqslant 2$. In particular, letting

$$a_\ell = \frac{1}{\ell!}\partial^\ell g_{m,n}(m-\ell),$$

we have the recurrence

$$\ell a_\ell = \frac{m-\ell}{n}a_{\ell-1} + \frac{1}{n}a_{\ell-2} \tag{4.14}$$

for $(a_\ell)_{\ell\geqslant 0}$. Observe that $a_0 = n^m$, $a_1 = (m-1)n^{m-1}$, and that $Q(m,n) = a_m$.

In the next section we will bound solutions to recurrences like (4.14) in a direct fashion, but for (4.14) itself it is more convenient to use a generating function technique. For an indeterminate X put

$$f(X) = \sum_{\ell=0}^{\infty} a_\ell X^\ell,$$

and observe by multiplying (4.14) by $X^{\ell-1}$ and summing for $\ell \geqslant 2$ that

$$f'(X) - a_1 = \frac{m-1}{n}(f(X) - a_0) - \frac{1}{n}Xf'(X) + \frac{1}{n}Xf(X)\ .$$

By inputting $a_0 = n^m$ and $a_1 = (m-1)n^{m-1}$ and rearranging we arrive at

$$(n-X)f'(X) = (m-1+X)f(X).$$

The formula (4.13) is obtained by solving this differential equation.

4.6 An L^∞ bound for low-entropy minor arcs

The main result of this section is the following proposition.

Proposition 4.6.1. *Let χ be a character with exactly m nonzero coordinates, where $m \leqslant n/3$. Then*

$$|\widehat{1_S}(\chi)| \leqslant e^{O(m^{3/2}/n^{1/2}+m^{1/2})}\cdot 2^{-m/2}\binom{n}{m}^{-1/2}\cdot\frac{n!}{n^n}.$$

It is instructive to compare this bound to Theorem 4.4.1. If χ has exactly m nonzero coordinates and takes precisely k distinct nonzero values, i.e., if χ has the form $(r_1^{a_1}, \ldots, r_k^{a_k}, 0^{n-m})$ where $\sum a_i = m$, then Theorem 4.4.1 gives us a bound of

$$|\widehat{1_S}(\chi)| \leqslant \binom{n+k}{k}^{1/2} \binom{n}{a_1, \ldots, a_k, n-m}^{-1/2} \frac{n!}{n^n}.$$

This bound is worst approximately for characters of the form $\chi = (r^m, 0^{n-m})$, for which we get the bound

$$|\widehat{1_S}(\chi)| \leqslant O(n^{1/2}) \binom{n}{m}^{-1/2} \frac{n!}{n^n}.$$

Moreover if we allow n to be even and set $r = 1/2$ then this bound is approximately attained, as remarked in the comments following Theorem 4.4.1. Here, however, we are asserting an improvement to this bound by a significant factor of $2^{-m/2}$. Such an improvement is only possible under the assumption that n is odd, and thus our proof of Proposition 4.6.1 must exploit this assumption in a fundamental way.

Since we do not know how to adapt the proof of Theorem 4.4.1 to exploit the absence of 2-torsion, we employ a more specialized approach that suffices in the sparse regime.

Incidentally we observe that this is the unique stage of the proof in which we need the full strength of the hypothesis that G has odd order. Elsewhere we only need to assume that $\sum_{x \in G} x = 0$.

Proof of Proposition 4.6.1. Let $\chi = (r_1, \ldots, r_m, 0^{n-m})$. The key tool in our proof is an exact recursive formula for Fourier coefficients $\widehat{1_S}(\chi)$, in terms of related Fourier coefficients for which $m' < m$. Specifically, as long as r_m is nonzero we have

$$\begin{aligned}
\widehat{1_S}(\chi) &= \frac{(n-m)!}{n^n} \sum_{x_1, \ldots, x_m}{}' e(-r_1 x_1 - \cdots - r_m x_m) \\
&= \frac{(n-m)!}{n^n} \sum_{x_1, \ldots, x_{m-1}}{}' \sum_{x_m \neq x_1, \ldots, x_{m-1}} e(-r_1 x_1 - \cdots - r_m x_m) \\
&= -\frac{(n-m)!}{n^n} \sum_{x_1, \ldots, x_{m-1}}{}' \sum_{x_m = x_1, \ldots, x_{n-1}} e(-r_1 x_1 - \cdots - r_m x_m) \\
&= -\frac{1}{n-m+1} \sum_{i=1}^{m-1} \widehat{1_S}(r_1, \ldots, r_i + r_m, \ldots, r_{m-1}, 0, \ldots, 0). \qquad (4.15)
\end{aligned}$$

(As an aside, we remark that this formula can be used for efficient explicit computation of certain kinds of Fourier coefficient.)

Let U_m be the maximal value of the $|\widehat{1_S}|$ taken over all characters with exactly m nonzero coordinates. If we apply the triangle inequality to the right-hand side of (4.15) and bound each $|\widehat{1_S}(\chi')|$ appearing there by the appropriate value $U_{m'}$ (where $m' < m$), we obtain recursive bounds on the U_m.

However, there is a subtlety in this process: the number of nonzero coordinates of

$$(r_1, \dots, r_i + r_m, \dots, r_{m-1}, 0, \dots, 0)$$

might be either $(m-1)$ (if $r_i + r_m \neq 0$) or $(m-2)$ (if $r_i + r_m = 0$), and we do not have much control over which occurs. Since we expect $U_{m-1} < U_{m-2}$, it is in our interests to land in the former case as much as possible. We have the freedom to reorder the r_i before applying (4.15), so our goal is to optimize the recursive bound over all choices of r_m.

For each $i \leqslant m$ let $\mathcal{N}_i = \{j \leqslant m \colon r_j = -r_i\}$. By reordering, we may assume without loss of generality that $|\mathcal{N}_i|$ is minimized when $i = m$. Suppose $j \in \mathcal{N}_m$. Then $r_j = -r_m$, so since $|G|$ is odd we have $r_j \neq r_m$, so $\mathcal{N}_m \cap \mathcal{N}_j = \varnothing$. Thus by minimality $|\mathcal{N}_m| \leqslant m/2$.

Hence, by applying (4.15) we deduce the bound

$$\begin{aligned} |\widehat{1_S}(\chi)| &\leqslant \frac{|\mathcal{N}_m| U_{m-2} + (m - 1 - |\mathcal{N}_m|) U_{m-1}}{n - m + 1} \\ &\leqslant \frac{1}{n-m+1} \max\left\{(m-1)U_{m-1},\ (m/2)U_{m-2} + (m/2-1)U_{m-1}\right\} \end{aligned}$$

where the former term in the maximum covers the case $U_{m-1} > U_{m-2}$ and the latter the (inuitively more likely) case $U_{m-1} \leqslant U_{m-2}$. Hence

$$U_m \leqslant \frac{1}{n-m+1} \max\left\{(m-1)U_{m-1},\ (m/2)U_{m-2} + (m/2-1)U_{m-1}\right\}. \tag{4.16}$$

Our task is now to obtain bounds on U_m by solving this recurrence, which, ignoring the maximum, resembles a kind of time-dependent Fibonacci sequence. In Section 4.5 we dealt with a similar recurrence using generating function methods. Here we will use a more hands-on method. Specifically, we will phrase (4.16) in terms of a product of 2×2 matrices (as one might for the Fibonacci sequence), and control the L^2 norm of the result by bounding the $L^2 \to L^2$ operator norms of each matrix in the product.

Write $\alpha_m = \frac{m}{2(n-m+1)}$, $\beta_m = \frac{m-2}{2(n-m+1)}$, and $\gamma_m = \frac{m-1}{n-m+1}$. Additionally, we work with a rescaled version $V_m = (\alpha_1 \alpha_2 \dots \alpha_m)^{-1/2} U_m$ of U_m, for $m \geqslant 1$. From (4.16) we deduce that

$$V_m \leqslant \max\left\{\gamma_m \alpha_m^{-1/2} V_{m-1},\ \alpha_m^{1/2} \alpha_{m-1}^{-1/2} V_{m-2} + \beta_m \alpha_m^{-1/2} V_{m-1}\right\} \tag{4.17}$$

holds for any $m \geqslant 3$.

If we define matrices

$$M_m = \begin{pmatrix} \gamma_m \alpha_m^{-1/2} & 0 \\ 1 & 0 \end{pmatrix} \quad \text{and} \quad N_m = \begin{pmatrix} \beta_m \alpha_m^{-1/2} & \alpha_m^{1/2} \alpha_{m-1}^{-1/2} \\ 1 & 0 \end{pmatrix},$$

and vectors $v_m = (V_m, V_{m-1})^T$, we can write (4.17) equivalently as

$$v_m \leqslant \max\{M_m v_{m-1}, N_m v_{m-1}\}.$$

Using the easily obtainable bounds $\alpha_m \alpha_{m-1}^{-1} = 1 + O(1/m)$ and $\beta_m^2 \alpha_m^{-1} \leqslant m/n$, we get that

$$\operatorname{tr}\left(N_m^T N_m\right) \leqslant 2 + m/n + O(1/m)$$

and

$$\det\left(N_m^T N_m\right) = 1 + O(1/m),$$

and so by considering the singular values of N_m, we get the operator norm bound

$$\|N_m\|_{L^2 \to L^2} \leqslant 1 + O((m/n)^{1/2} + m^{-1/2}).$$

Completely analogously, using the easy bound $\gamma_m^2 \alpha_m^{-1} \leqslant 4m/n$, we deduce

$$\|M_m\|_{L^2 \to L^2} \leqslant 1 + O(m/n).$$

Thus we have

$$\begin{aligned} V_m &\leqslant \sqrt{V_m^2 + V_{m-1}^2} \\ &\leqslant \left(\prod_{r=3}^{m} \max\{\|M_r\|_{L^2 \to L^2}, \|N_r\|_{L^2 \to L^2}\} \right) \cdot \sqrt{V_2^2 + V_1^2} \\ &\leqslant \left(\prod_{r=3}^{m} \left(1 + O\left((r/n)^{1/2} + r^{-1/2}\right)\right) \right) \cdot \sqrt{V_2^2 + V_1^2}. \end{aligned}$$

We already know from Section 4.2.1 that $U_1 = 0$ and it is evident from (4.15) that $U_2 = O(n!/n^{n+1})$. Hence, we can bound the term $\sqrt{V_2^2 + V_1^2}$ in the above inequality by $O(n!/n^n)$, which gives

$$V_m \leqslant e^{O\left(m^{3/2}/n^{1/2} + m^{1/2}\right)} \cdot \frac{n!}{n^n},$$

and the claim stated in the proposition follows easily from this. □

4.7 End of the argument

To estimate the sum of cubes

$$\sum_{\chi \in \widehat{G}^n} \widehat{1_S}(\chi)^3$$

we divide the set of all $\chi \in \widehat{G}^n$ into three regions depending on the value of $H(\chi)$: a high-entropy range $H \geqslant 10$, a medium-entropy range $\varepsilon \leqslant H \leqslant 10$, and a low-entropy range $H \leqslant \varepsilon$. Here ε is a small positive constant which will be chosen at the end of the argument.

In the high-entropy range $H \geqslant 10$ it is enough to use the bound from Corollary 4.2.2:

$$\sum_{\chi:\, H(\chi) \geqslant 10} |\widehat{1_S}(\chi)|^3 \leqslant e^{-3.5n} n!^3 / n^{3n}.$$

In the medium-entropy range $\varepsilon \leqslant H \leqslant 10$, we first need a bound for the number of characters of a given entropy.

Lemma 4.7.1. *If $H \leqslant 10$ then the number of characters of entropy at most H is bounded by $e^{Hn+o(n)}$.*

Proof. Every character of entropy at most H has an orbit under the permutation action of size at most e^{Hn}, so it suffices to show that there are only $e^{o(n)}$ such orbits of characters of entropy at most 10. Every orbit is uniquely specified by giving the number a_r of appearances of r for each $r \in \widehat{G}$, so we must show that the number of nonnegative integer vectors $(a_r)_{r \in \widehat{G}}$ with $\sum_r a_r = n$ and satisfying

$$\frac{n!}{\prod_{r \in \widehat{G}} a_r!} \leqslant e^{10n}$$

is $e^{o(n)}$. Let $\delta > 0$ be a small constant, and let t be the sum of the a_r for which $a_r \leqslant e^{-11/\delta} n$. Then

$$e^{-11n} n^n \leqslant e^{-10n} n! \leqslant \prod_{r \in \widehat{G}} a_r! \leqslant \prod_{r \in \widehat{G}} a_r^{a_r} \leqslant (e^{-11/\delta} n)^t n^{n-t} = e^{-11t/\delta} n^n,$$

so $t \leqslant \delta n$. Since at most $e^{11/\delta}$ of the a_r are bigger than $e^{-11/\delta} n$, the number of $(a_r)_{r \in \widehat{G}}$ is bounded by the number of ways of choosing the set B of at most $e^{11/\delta}$ indices r for which a_r is big, times the number of ways of choosing these a_r, times the number of ways of choosing $(a_r)_{r \notin B}$ so that $\sum_{r \notin B} a_r = n - \sum_{r \in B} a_r \leqslant \delta n$. This is bounded by

$$n^{e^{11/\delta}} n^{e^{11/\delta}} \binom{n + \delta n - 1}{\delta n - 1}.$$

The claimed estimate now follows by applying Stirling's formula and choosing δ appropriately. $\square$

We conclude by combining Lemma 4.7.1 with Theorem 4.4.1. Note that if the entropy of χ is $O(1)$ then the number of distinct coordinates of χ is $O(n/\log n)$, so Theorem 4.4.1 implies

$$|\widehat{1_S}(\chi)| \leqslant e^{-H(\chi)n/2+o(n)}n!/n^n.$$

Thus for any $H \leqslant 10$ we have

$$\begin{aligned}\sum_{\chi\colon 9H/10<H(\chi)\leqslant H} |\widehat{1_S}(\chi)|^3 &\leqslant |\{\chi\colon H(\chi) \leqslant H\}| \max_{\chi\colon H(\chi)\geqslant 9H/10} |\widehat{1_S}(\chi)|^3 \\ &\leqslant e^{-7Hn/20+o(n)}n!^3/n^{3n}.\end{aligned}$$

Decomposing the range $[\varepsilon, 10]$ into ranges of this type and bounding crudely, we obtain

$$\sum_{\chi\colon \varepsilon\leqslant H(\chi)\leqslant 10} |\widehat{1_S}(\chi)|^3 \leqslant O(\log(1/\varepsilon))e^{-7\varepsilon n/20+o(n)}n!^3/n^{3n}.$$

In the low-entropy range $H \leqslant \varepsilon$, one can easily verify that $\chi = (r_1, \dots, r_n)$ must repeat some coordinate $r \in \widehat{G}$ at least $(1-\varepsilon)n$ times. Thus by global shift-invariance of $\widehat{1_S}$ we have

$$\sum_{\chi\colon H(\chi)\leqslant\varepsilon} \widehat{1_S}(\chi)^3 = n\sum_{m=0}^{\varepsilon n} \sum_{\substack{m\text{-sparse } \chi \\ H(\chi)\leqslant\varepsilon}} \widehat{1_S}(\chi)^3. \tag{4.18}$$

By combining Proposition 4.6.1 with Theorem 4.5.1 we have that for any $m \leqslant \varepsilon n$,

$$\begin{aligned}\sum_{m\text{-sparse } \chi} |\widehat{1_S}(\chi)|^3 &\leqslant \left(\max_{m\text{-sparse } \chi} |\widehat{1_S}(\chi)|\right) \sum_{m\text{-sparse } \chi} |\widehat{1_S}(\chi)|^2 \\ &\leqslant e^{O\left(m^{3/2}/n^{1/2}+m^{1/2}\right)}2^{-m/2}\frac{n!^3}{n^{3n}} \\ &\leqslant e^{O\left(\varepsilon^{1/2}m+m^{1/2}\right)}2^{-m/2}\frac{n!^3}{n^{3n}}.\end{aligned}$$

As long as ε is sufficiently small depending on the constant implicit in the $O(\varepsilon^{1/2}m)$ term, this is negligible except when m has size $O(1)$. In this range, we can apply Proposition 4.3.1. We introduce a further parameter M and split the sum (4.18) into the ranges $1 \leqslant m \leqslant 2M$ and $2M < m \leqslant \varepsilon n$, to obtain

$$\begin{aligned}\sum_{\chi\colon H(\chi)\leqslant\varepsilon} \widehat{1_S}(\chi)^3 = n\sum_{m=0}^{M} \frac{(-1)^m}{2^m m!}\frac{n!^3}{n^{3n}} + nO_M\left(\frac{1}{n}\frac{n!^3}{n^{3n}}\right) \\ + O\left(n\sum_{m=2M}^{\varepsilon n} e^{O\left(\varepsilon^{1/2}m+m^{1/2}\right)}2^{-m/2}\frac{n!^3}{n^{3n}}\right).\end{aligned}$$

Finally, by combining the three bounds we have just proved for each range $H \leqslant \varepsilon$, $\varepsilon \leqslant H \leqslant 10$, and $H \geqslant 10$, we deduce

$$\begin{aligned}\sum_{\chi\in\widehat{G^n}} \widehat{1_S}(\chi)^3 = n\sum_{m=0}^{M}\frac{(-1)^m}{2^m m!}\frac{n!^3}{n^{3n}} + nO_M\left(\frac{1}{n}\frac{n!^3}{n^{3n}}\right)\\ + O\left(n\sum_{m=2M}^{\varepsilon n} e^{O\left(\varepsilon^{1/2}m+m^{1/2}\right)}2^{-m/2}\frac{n!^3}{n^{3n}}\right)\\ + O\left(\log(1/\varepsilon)e^{-7\varepsilon n/20+o(n)}\frac{n!^3}{n^{3n}}\right)\\ + O\left(e^{-3.5n}\frac{n!^3}{n^{3n}}\right).\end{aligned}$$

Theorem 4.1.2 now follows by choosing ε and M appropriately.

References

[Bón15] M. Bóna, ed. *Handbook of enumerative combinatorics*. Discrete Mathematics and its Applications (Boca Raton). CRC Press, Boca Raton, FL, 2015, pp. xxiii+1061. ISBN: 978-1-4822-2085-8.

[CW10] N. J. Cavenagh and I. M. Wanless. "On the number of transversals in Cayley tables of cyclic groups". *Discrete Appl. Math.* 158.2 (2010), pp. 136–146. ISSN: 0166-218X. DOI: 10.1016/j.dam.2009.09.006.

[Coo00] C. Cooper. "A lower bound for the number of good permutations". *Data Recording, Storage and Processing (Nat. Acad. Sci. Ukraine)* 213 (2000), pp. 15–25.

[Coo+99] C. Cooper, R. Gilchrist, I. N. Kovalenko, and D. Novaković. "Deriving the number of "good" permutations, with applications to cryptography". *Kibernet. Sistem. Anal.* 5 (1999), pp. 10–16, 187. ISSN: 0023-1274. DOI: 10.1007/BF02733401.

[CK96] C. Cooper and I. N. Kovalenko. "The upper bound for the number of complete mappings". *Theory Probab. Math. Statist.* 53 (1996), pp. 77–83.

[EMM15] S. Eberhard, F. Manners, and R. Mrazović. *Additive triples of bijections, or the toroidal semiqueens problem*. 2015. URL: http://arxiv.org/abs/1510.05987.

[FS09] P. Flajolet and R. Sedgewick. *Analytic combinatorics*. Cambridge University Press, Cambridge, 2009, pp. xiv+810. ISBN: 978-0-521-89806-5. DOI: 10.1017/CBO9780511801655.

[GL15] R. Glebov and Z. Luria. *On the maximum number of Latin transversals*. 2015. URL: http://arxiv.org/abs/1506.00983.

[Kov96] I. N. Kovalenko. "On an upper bound for the number of complete mappings". *Kibernet. Sistem. Anal.* 1 (1996), pp. 81–85, 188. ISSN: 0023-1274. DOI: 10.1007/BF02366583.

[Kuz07] N. Y. Kuznetsov. "Applying fast simulation to find the number of good permutations". *Cybernet. Systems Anal.* 43 (2007), pp. 830–837.

[Kuz08] N. Y. Kuznetsov. "Estimating the number of good permutations by a modified fast simulation method". *Cybernet. Systems Anal.* 44 (2008), pp. 547–554.

[Kuz09] N. Y. Kuznetsov. “Estimating the number of latin rectangles by the fast simulation method”. *Cybernet. Systems Anal.* 45 (2009), pp. 69–75.

[MMW06] B. D. McKay, J. C. McLeod, and I. M. Wanless. “The number of transversals in a Latin square”. *Des. Codes Cryptogr.* 40.3 (2006), pp. 269–284. ISSN: 0925-1022. DOI: 10.1007/s10623-006-0012-8.

[RVZ94] I. Rivin, I. Vardi, and P. Zimmermann. “The n-queens problem”. *Amer. Math. Monthly* 101.7 (1994), pp. 629–639. ISSN: 0002-9890. DOI: 10.2307/2974691.

[Slo] N. J. A. Sloane. *The On-Line Encyclopedia of Integer Sequences*. URL: https://oeis.org/.

[Tar15] A. A. Taranenko. “Multidimensional permanents and an upper bound on the number of transversals in Latin squares”. *J. Combin. Des.* 23.7 (2015), pp. 305–320. ISSN: 1063-8539. DOI: 10.1002/jcd.21413.

[Var91] I. Vardi. *Computational recreations in Mathematica*. Addison-Wesley Publishing Company, Redwood City, CA, 1991, pp. xviii+286. ISBN: 0-201-52989-0.

[Wan11] I. M. Wanless. “Transversals in Latin squares: a survey”. *Surveys in combinatorics 2011*. Vol. 392. London Math. Soc. Lecture Note Ser. Cambridge Univ. Press, Cambridge, 2011, pp. 403–437.

Chapter 5

Product mixing in the alternating group

Abstract. We prove the following one-sided product-mixing theorem for the alternating group: Given subsets $X, Y, Z \subset A_n$ of densities α, β, γ satisfying

$$\min(\alpha\beta, \alpha\gamma, \beta\gamma) \gg n^{-1}(\log n)^7,$$

there are at least

$$(1 + o(1))\alpha\beta\gamma|A_n|^2$$

solutions to $xy = z$ with $x \in X, y \in Y, z \in Z$. One consequence is that the largest product-free subset of A_n has density at most $n^{-1/2}(\log n)^{7/2}$, which is best possible up to logarithms and improves the best previous bound of $n^{-1/3}$ due to Gowers. The main tools are a Fourier-analytic reduction noted by Ellis and Green to a problem just about the standard representation, a Brascamp–Lieb-type inequality for the symmetric group due to Carlen, Lieb, and Loss, and a concentration of measure result for rearrangements of inner products.

This chapter reproduces the paper [Ebe16].

5.1 Introduction

Product mixing for a group G generally refers to any estimate of the following form: whenever subsets $X, Y, Z \subset G$ have densities α, β, γ above some threshold, the number of solutions to $xy = z$ with $x \in X, y \in Y, z \in Z$ is $(1 + o(1))\alpha\beta\gamma|G|^2$. The following foundational theorem proved by Gowers [Gow08] (and expanded by Babai, Nikolov, and Pyber [BNP08]) explains this idea further.

Theorem 5.1.1 (Gowers)**.** *Let G be a group and let m be the minimal dimension of a nontrivial representation of G. Let $X, Y, Z \subset G$ have densities α, β, γ, respectively. Then*

$$|\langle 1_X * 1_Y, 1_Z \rangle - \alpha\beta\gamma| < m^{-1/2}\alpha^{1/2}\beta^{1/2}\gamma^{1/2}.$$

In particular if $\alpha\beta\gamma \gg m^{-1}$ then

$$\langle 1_X * 1_Y, 1_Z \rangle = (1 + o(1))\alpha\beta\gamma.$$

Here and throughout this chapter we write $X \lesssim Y$ to mean that $X \leqslant O(Y)$, and we write $X \ll Y$ to mean that $X \leqslant o(Y)$. We will write $X \sim Y$ to mean $X \lesssim Y$ and $X \gtrsim Y$. This differs from the convention used for example in Chapters 1 and 3, but it will be convenient for us. The convolutions and inner products above are defined with respect to the uniform measure.

There are several immediate corollaries of Theorem 5.1.1. For example, if $\alpha\beta\gamma \geqslant m^{-1}$, then the intersection $XY \cap Z$ is nonempty, and in fact $XYZ^{-1} = G$. In particular, if $X \subset G$ is product-free (meaning that there are no solutions to $xy = z$ with $x, y, z \in X$), then X has density at most $m^{-1/3}$.

For the purpose of illustration let us assume $\alpha \sim \beta \sim \gamma$. Then Theorem 5.1.1 asserts that there is a product-mixing phenomenon for sets of density greater than $m^{-1/3}$. On the other hand Kedlaya [Ked97] proved that any group G acting transitively on a set of size n has a product-free subset of density $n^{-1/2}$. For a broad class of groups, including for example the alternating groups and special linear groups, we have $m \sim n$, so for these groups this leaves a gap between $m^{-1/3}$ and $m^{-1/2}$.

In Section 5.2 we partly explain this gap by showing that any group G acting transitively on a set of size n has a subset X of density $\sim n^{-1/3}$ for which there are significantly *more* than the expected number of solutions to $xy = z$. In groups with $m \sim n$ this shows that the density threshold for product mixing is $m^{-1/3}$, as in Gowers's theorem.

Our main purpose, however, is to demonstrate that a one-sided product-mixing phenomenon persists in the alternating group A_n for somewhat lower densities. Specifically we prove the following theorem.

Theorem 5.1.2. *If $X, Y, Z \subset A_n$ have densities α, β, γ, respectively, and*

$$\min(\alpha\beta, \alpha\gamma, \beta\gamma) \gg \frac{(\log n)^7}{n},$$

then

$$\langle 1_X * 1_Y, 1_Z \rangle \geqslant (1 + o(1))\alpha\beta\gamma.$$

As a corollary we deduce that if X has density $\gg n^{-1/2}(\log n)^{7/2}$ then X^2 has density $1 - O(n^{-1/2}(\log n)^{7/2})$. In particular if X is product-free then X has density at most $O(n^{-1/2}(\log n)^{7/2})$. This is best possible up to the logarithmic factors.

As to the methods, we first use nonabelian Fourier analysis to reduce to a problem taking place only in the standard representation, an idea due to Ellis and Green. This problem is then interpreted in terms of random rearrangements of inner products, and we tackle this problem using concentration of measure and entropy subadditivity. The backbone of our proof is a Brascamp–Lieb-type inequality for the symmetric group due to Carlen, Lieb, and Loss, which we explain in Section 5.4.

Notation. As already mentioned, in addition to the usual asymptotic notation $O(\cdot)$ and $o(\cdot)$, we write $X \lesssim Y$ to mean that $X \leqslant O(Y)$, and we write $X \ll Y$ to mean that $X \leqslant o(Y)$. We write $X \sim Y$ to mean $X \lesssim Y$ and $X \gtrsim Y$.

We write Ω throughout for the ground set $\{1, \dots, n\}$ on which S_n and A_n act. We attach the uniform measures to S_n, A_n, and Ω, and we write an unadorned integral $\int f$ to mean the integral with respect to the uniform measure on the domain of f. We also define inner products, L^p norms, and convolutions accordingly.

5.2 Examples of sets with poor product mixing

In this section we give two concrete examples of fairly dense sets with poor product-mixing properties. The first example, a relatively large product-free set, is due to Kedlaya [Ked97] and independently Edward Crane (Ben Green, personal communication), but we recall the construction here as it shows that Theorem 5.1.2 is best possible up to logarithms. The second construction is original, and shows that Theorem 5.1.1 is best possible for two-sided mixing.

5.2.1 Sets with no solutions to $xy = z$

First we give an example of a set X of density $\sim n^{-1/2}$ with no solutions to $xy = z$. Fix a set $T \subset \Omega$ of size t and a point 1 not in T and let X be the set of all $\pi \in A_n$ such that $\pi(1) \in T$ and such that $\pi(T) \subset T^c$. Then clearly X^2 is disjoint from X, as every $\pi \in X^2$ satisfies $\pi(1) \in T^c$, and it is straightforward to see that X has density

$$\frac{1}{n!} t \binom{n-t}{t} t!(n-t-1)! = \frac{t(n-t)!(n-t-1)!}{n!(n-2t)!} = \frac{t}{n} e^{O(t^2/n)}.$$

Thus if $t \sim n^{1/2}$ then X has density $\sim n^{-1/2}$. This example is due to Kedlaya [Ked97] and independently Edward Crane (Ben Green, personal communication), and it shows that Theorem 5.1.2 is best possible up to logarithms.

As explained by Kedlaya [Ked97], his construction adapts straightforwardly to any 2-transitive subgroup $G \leqslant S_n$, and in fact it adapts to any transitive subgroup $G \leqslant S_n$ through an averaging argument.

Proposition 5.2.1 (Kedlaya [Ked97])**.** *Let G be a transitive subgroup of S_n. Then there is a subset $X \subset G$ of density $\sim n^{-1/2}$ such that $X^2 \cap X = \varnothing$.*

5.2.2 Sets with too many solutions to $xy = z$

Next we give an example of a set X of density $\alpha \sim n^{-1/3}$ having many more than the expected number of solutions (namely, $\alpha^3 n!^2$) to $xy = z$. Fix a set T of size t and let X be the set of all $\pi \in A_n$ such that $\pi(T) \cap T$ is nonempty. As long as $t = o(n^{1/2})$ then X has density roughly t^2/n, and if you choose π_1, π_2 randomly from X then $\pi_1\pi_2$ is again in X with probability of order $t^2/n + 1/t$. To see this it may help to notice that X is symmetric, and that $\pi_1^{-1}\pi_2 \in X$ if and only if $\pi_1(T) \cap \pi_2(T) \neq \varnothing$. Each of $\pi_1(T)$ and $\pi_2(T)$ is required to intersect T nontrivially, so $\pi_1(T)$ and $\pi_2(T)$ intersect with probability at least $1/t$. Aside from that restriction $\pi_1(T)$ and $\pi_2(T)$ are just random sets of size t, so they intersect with probability at least t^2/n. (We can afford to be somewhat lax with this computation as we will shortly prove a more general proposition.) Note that the probability $t^2/n + 1/t$ is much larger than the expected probability t^2/n whenever t is small compared to $n^{1/3}$.

As with the previous construction, this construction adapts straightforwardly to any 2-transitive subgroup $G \leqslant S_n$, and to an arbitrary transitive subgroup $G \leqslant S_n$ through an averaging argument.

Proposition 5.2.2. *Let G be a transitive subgroup of S_n. Then there is a subset $X \subset G$ of density $\alpha \sim n^{-1/3}$ for which there are at least $100\alpha^3|G|^2$ solutions to $xy = z$ with $x, y, z \in X$.*

Proof. For $T \subset \Omega$ of size t let X_T be the set of all $g \in G$ for which $g(T) \cap T \neq \varnothing$. Also fix an arbitrary total order $<$ on Ω. Clearly $|X_T|/|G| \leqslant t^2/n$. We will bound $|X_T|$ below by the number of $g \in G$ for which there are $i, j \in T$ with $i < j$ such that $g(i) = j$. Thus by inclusion-exclusion we have

$$\frac{|X_T|}{|G|} \geqslant \sum_{\substack{i,j \in T \\ i<j}} \frac{|\{g : g(i) = j\}|}{|G|} - \sum_{\substack{i,j,i',j' \in T \\ i<j, i'<j' \\ (i,j) \neq (i',j')}} \frac{|\{g : g(i) = j, g(i') = j'\}|}{|G|}.$$

The first sum here is $\sim t^2/n$ by transitivity, for any T. The second sum can be rewritten as

$$\sum_{\substack{i,i'\in\Omega\\ i\neq i'}}\frac{1}{|G|}\sum_{\substack{g\in G\\ g(i)>i, g(i')>i'}} 1_{i\in T}1_{g(i)\in T}1_{i'\in T}1_{g(i')\in T}. \tag{5.1}$$

Now note that for any fixed $i, i' \in \Omega$ such that $i \neq i'$ and for any g satisfying $g(i) > i$ and $g(i') > i'$ we have $|\{i, g(i), i', g(i')\}| \geqslant 3$, and in fact $|\{i, g(i), i', g(i')\}| = 4$ except for a proportion at most $O(1/n)$ of $g \in G$. It follows that the average of (5.1) over $T \subset \Omega$ is bounded by

$$O(n^2(t/n)^4 + n(t/n)^3) = O(t^4/n^2).$$

Thus, by Markov's inequality, (5.1) is $O(t^4/n^2)$ with probability at least 9/10.

Similarly let us count solutions to $xy = z$ in X_T. We will bound the number N_T of solutions below by the number of pairs $(g_1, g_2) \in G^2$ for which there exists $i, j, k \in T$ with $i < j < k$ such that $g_1(i) = j$ and $g_2(j) = k$. Thus by inclusion-exclusion again we have

$$\begin{aligned}\frac{N_T}{|G|^2} \geqslant &\sum_{\substack{i,j,k\in T\\ i<j<k}}\frac{|\{(g_1,g_2)\in G^2 : g_1(i)=j, g_2(j)=k\}|}{|G|^2}\\ &-\sum_{\substack{i,j,k,i',j',k'\in T\\ i<j<k, i'<j'<k'\\ (i,j,k)\neq(i',j',k')}}\frac{|\{(g_1,g_2)\in G^2 : g_1(i)=j, g_2(j)=k, g_1(i')=j', g_2(j')=k'\}|}{|G|^2}.\end{aligned}$$

The first sum is $\sim t^3/n^2$ by transitivity. The second sum can be rewritten

$$\sum_{\substack{j,j'\in\Omega\\ j\neq j'}}\frac{1}{|G|^2}\sum_{\substack{g_1,g_2\in G\\ g_1^{-1}(j)<j<g_2(j)\\ g_1^{-1}(j')<j'<g_2(j')}} 1_{g_1^{-1}(j)\in T}1_{j\in T}1_{g_2(j)\in T}1_{g_1^{-1}(j')\in T}1_{j'\in T}1_{g_2(j')\in T}. \tag{5.2}$$

To bound this we again average over $T \subset \Omega$. For $j \neq j'$ and g_1, g_2 under the stated restrictions the set

$$S = \{g_1^{-1}(j), j, g_2(j), g_1^{-1}(j'), j', g_2(j')\}$$

always has size at least 4, has size 4 for at most a proportion $O(1/n^2)$ of $(g_1, g_2) \in G^2$, has size 5 for at most a proportion $O(1/n)$ of $(g_1, g_2) \in G^2$, and otherwise has size 6. It follows that the average of (5.2) over $T \subset \Omega$ is bounded by

$$O(n^2(t/n)^6 + n(t/n)^5 + (t/n)^4) = O(t^6/n^4).$$

Thus, by Markov's inequality, (5.2) is $O(t^6/n^4)$ with probability at least 9/10.

We deduce that there is some T for which (5.1) is $O(t^4/n^2)$ and (5.2) is $O(t^6/n^4)$. For this T it follows that

$$\frac{|X_T|}{|G|} \sim t^2/n + O(t^4/n^2)$$

and that

$$\frac{N_T}{|G|^2} \gtrsim t^3/n^2 + O(t^6/n^4).$$

Thus as long as $t = o(n^{1/2})$ we see that X_T has density $\alpha \sim t^2/n$ while there are at least $(t^3/n^2)|G|^2 \sim (n/t^3)\alpha^3|G|^2$ solutions to $xy = z$ in X. Now take $t = \lfloor cn^{1/3} \rfloor$ for a sufficiently small constant c. □

5.3 Nonabelian Fourier analysis

Here we briefly recall the fundamentals of nonabelian Fourier analysis, and then we give a short Fourier-analytic proof of Theorem 5.1.1. This proof seems to be well known among experts: see for example Wigderson [Wig10, Chapter 2.11].

Let G be a compact group endowed with the uniform measure. The Fourier transform of a function $f \in L^2(G)$ at an irreducible unitary representation $\xi : G \to U(d_\xi)$ is defined by

$$\widehat{f}(\xi) = \int_G f(x)\xi(x).$$

We then have the inversion formula

$$f(x) = \sum_\xi d_\xi \langle \widehat{f}(\xi), \xi(x) \rangle_{\mathrm{HS}},$$

and Parseval's identity

$$\langle f, g \rangle = \sum_\xi d_\xi \langle \widehat{f}(\xi), \widehat{g}(\xi) \rangle_{\mathrm{HS}}. \tag{5.3}$$

Here the sums are taken over a complete set of representatives of the irreducible representations of G up to equivalency, and the Hilbert–Schmidt inner product $\langle \cdot, \cdot \rangle_{\mathrm{HS}}$ is defined by

$$\langle R, S \rangle_{\mathrm{HS}} = \mathrm{tr}(RS^*).$$

Like classical Fourier analysis, nonabelian Fourier analysis is a powerful tool for understanding the behaviour of convolutions. Here the convolution $f * g$ of two functions $f, g \in L^2(G)$ is defined by

$$f * g(x) = \int_G f(y)g(y^{-1}x),$$

and by an application of Fubini's theorem we have the rule

$$\widehat{f * g}(\xi) = \widehat{f}(\xi)\widehat{g}(\xi). \tag{5.4}$$

For all this and more the reader might refer to Tao [Tao14, §2.8].

We can now give a short proof of Theorem 5.1.1.

Proof of Theorem 5.1.1. Suppose that G is finite, that $d_\xi \geqslant m$ for $\xi \neq 1$, and that $X, Y, Z \subset G$ have densities α, β, γ, respectively. Let $f = 1_X, g = 1_Y, h = 1_Z$. Then by the convolution rule (5.4) and Parseval (5.3) we have

$$\begin{aligned}\langle f * g, h\rangle &= \sum_\xi d_\xi \langle \widehat{f}(\xi)\widehat{g}(\xi), \widehat{h}(\xi)\rangle_{\mathrm{HS}} \\ &= \alpha\beta\gamma + \sum_{\xi\neq 1} d_\xi \langle \widehat{f}(\xi)\widehat{g}(\xi), \widehat{h}(\xi)\rangle_{\mathrm{HS}}.\end{aligned}$$

Here we have written 1 for the trivial representation of G. Now by Cauchy–Schwarz and the algebra property $\|RS\|_{\mathrm{HS}} \leqslant \|R\|_{\mathrm{HS}}\|S\|_{\mathrm{HS}}$ of the Hilbert–Schmidt norm we have

$$|\langle \widehat{f}(\xi)\widehat{g}(\xi), \widehat{h}(\xi)\rangle_{\mathrm{HS}}| \leqslant \|\widehat{f}(\xi)\widehat{g}(\xi)\|_{\mathrm{HS}}\|\widehat{h}(\xi)\|_{\mathrm{HS}} \leqslant \|\widehat{f}(\xi)\|_{\mathrm{HS}}\|\widehat{g}(\xi)\|_{\mathrm{HS}}\|\widehat{h}(\xi)\|_{\mathrm{HS}},$$

so by using Cauchy–Schwarz together with Parseval again we have

$$\begin{aligned}\sum_{\xi\neq 1} d_\xi |\langle \widehat{f}(\xi)\widehat{g}(\xi), \widehat{h}(\xi)\rangle_{\mathrm{HS}}| &\leqslant \sum_{\xi\neq 1} d_\xi \|\widehat{f}(\xi)\|_{\mathrm{HS}}\|\widehat{g}(\xi)\|_{\mathrm{HS}}\|\widehat{h}(\xi)\|_{\mathrm{HS}} \\ &\leqslant \max_{\xi\neq 1} \|\widehat{f}(\xi)\|_{\mathrm{HS}} \sum_\xi d_\xi \|\widehat{g}(\xi)\|_{\mathrm{HS}}\|\widehat{h}(\xi)\|_{\mathrm{HS}} \\ &\leqslant m^{-1/2}\|f\|_2\|g\|_2\|h\|_2 \\ &= m^{-1/2}\alpha^{1/2}\beta^{1/2}\gamma^{1/2}.\end{aligned} \tag{5.5}$$

This proves Theorem 5.1.1. □

For the rest of the chapter we specialize to the alternating group $G = A_n$. As explained in the introduction, Theorem 5.1.1 provides a satisfactory estimate for $\langle 1_X * 1_Y, 1_Z\rangle$ only if $\alpha\beta\gamma \gg 1/n$. However, as observed by Ellis and Green (personal communication), by examination of the proof above it is clear that only the standard $(n-1)$-dimensional representation σ is problematic: Again taking $f = 1_X, g = 1_Y, h = 1_Z$, we have

$$\begin{aligned}\langle f * g, h\rangle &= \sum_\xi d_\xi \langle \widehat{f}(\xi)\widehat{g}(\xi), \widehat{h}(\xi)\rangle_{\mathrm{HS}} \\ &= \alpha\beta\gamma + (n-1)\langle \widehat{f}(\sigma)\widehat{g}(\sigma), \widehat{h}(\sigma)\rangle_{\mathrm{HS}} + \sum_{\xi\neq 1,\sigma} d_\xi \langle \widehat{f}(\xi)\widehat{g}(\xi), \widehat{h}(\xi)\rangle_{\mathrm{HS}},\end{aligned} \tag{5.6}$$

and since $d_\xi \gtrsim n^2$ for $\xi \neq 1, \sigma$ (this follows from the hook formula: see for example [Ras77, Result 2]) we have, by straightforward adaptation of (5.5),

$$\sum_{\xi \neq 1,\sigma} d_\xi |\langle \widehat{f}(\xi)\widehat{g}(\xi), \widehat{h}(\xi)\rangle_{\mathrm{HS}}| \lesssim n^{-1}\alpha^{1/2}\beta^{1/2}\gamma^{1/2}.$$

This is negligible compared to the main term $\alpha\beta\gamma$ whenever $\alpha\beta\gamma \gg n^{-2}$. Thus it remains only to control $(n-1)\langle \widehat{f}(\sigma)\widehat{g}(\sigma), \widehat{h}(\sigma)\rangle_{\mathrm{HS}}$.

For each $i \in \Omega$ we have a map $S_n \to \Omega$ given by $\pi \mapsto \pi(i)$, which induces a map $L^2(\Omega) \to L^2(S_n)$ given by composition with $\pi \mapsto \pi(i)$. We denote by p_i the adjoint of this map, and we call $p_i f$ the *pushforward* of f under $\pi \mapsto \pi(i)$. Explicitly $p_i f$ is defined by

$$p_i f(\omega) = n \int_{S_n} f(\pi) 1_{\pi(i)=\omega} = \frac{1}{(n-1)!} \sum_{\substack{\pi \in S_n \\ \pi(i)=\omega}} f(\pi),$$

and for any $g \in L^2(\Omega)$ we have

$$\int_{S_n} f(\pi) g(\pi(i)) = \int_\Omega p_i f(\omega) g(\omega).$$

Now by direct computation whenever at least one of $\int f, \int g, \int h$ is zero we have

$$\begin{aligned}
(n-1)\langle \widehat{f}(\sigma)\widehat{g}(\sigma), \widehat{h}(\sigma)\rangle_{\mathrm{HS}} &= (n-1)\int_{S_n^2} (f * g)(x)\overline{h(y)} \operatorname{tr} \sigma(xy^{-1}) \\
&= (n-1)\int_{S_n^2} (f * g)(x)\overline{h(y)} \left(\sum_{i \in \Omega} 1_{x(i)=y(i)} - 1 \right) \\
&= (n-1)\sum_{i \in \Omega} \int_{S_n^2} (f * g)(x)\overline{h(y)} 1_{x(i)=y(i)} \\
&= \frac{n-1}{n} \sum_{i \in \Omega} \langle f * p_i g, p_i h\rangle \\
&\sim \sum_{i \in \Omega} \langle f * p_i g, p_i h\rangle. \qquad (5.7)
\end{aligned}$$

Here we define the convolution of functions $f \in L^2(S_n)$ and $u \in L^2(\Omega)$ by the same formula:

$$f * u(\omega) = \int_{S_n} f(\pi) u(\pi^{-1}(\omega));$$

$f * u$ is then a function defined on Ω, and one may check the relation

$$p_i(f * g) = f * p_i g.$$

Note that the assumption that one of $\int f, \int g, \int h$ is zero is innocuous, since changing f by a constant does not change $\widehat{f}(\sigma)$.

Similarly whenever $\int f = 0$ we have the following remnant of Parseval's identity:

$$\|f\|_2^2 \geqslant (n-1)\|\widehat{f}(\sigma)\|_{\mathrm{HS}}^2 \sim \sum_{i\in\Omega} \|p_i f\|_2^2. \tag{5.8}$$

We can now summarize the rest of the proof. We will prove a concentration-of-measure result for the randomly rearranged inner product

$$\langle \pi * p_i g, p_i h\rangle = \int_\Omega p_i g(\pi^{-1}(\omega)) p_i h(\omega).$$

This result will ensure that

$$\langle \pi * p_i g, p_i h\rangle \approx \textstyle\int p_i g \int p_i h = \int g \int h$$

with high probability, and with a tail depending on the variances $\|p_i g - \int g\|_2^2$ and $\|p_i h - \int h\|_2^2$ of $p_i g$ and $p_i h$ and on the entropies of $p_i g/\int g$ and $p_i h/\int h$. Crucially, when the variances are small there is rather strong concentration from below, unless one of the entropies is large. We will then apply the Parseval remnant (5.8) and a version of subadditivity of entropy to conclude.

5.4 An inequality of Carlen, Lieb, and Loss

The following inequality was proved by Carlen, Lieb, and Loss [CLL06].

Theorem 5.4.1. *Let $f_1, \ldots, f_n : \Omega \to \mathbf{C}$ be functions. Then*

$$\int_{S_n} \prod_{i=1}^n |f_i(\pi(i))| \leqslant \prod_{i=1}^n \|f_i\|_2.$$

This inequality can be viewed in at least two ways. First, as it resembles the classical Loomis–Whitney inequality, or more generally the Brascamp–Lieb inequality, it can be viewed as an inequality of Brascamp–Lieb-type for the symmetric group. In this light Theorem 5.4.1 bears a striking resemblance to another Brascamp–Lieb-type inequality proved by Carlen, Lieb, and Loss for the sphere: see [CLL04].

Theorem 5.4.1 can also be viewed as a Hadamard-type inequality for permanents. The classical Hadamard inequality states that if M is a matrix with columns $v_1, \ldots, v_n \in \mathbf{C}^n$ then

$$|\det(M)| \leqslant \prod_{i=1}^n |v_i|,$$

where $|\cdot|$ is the usual Euclidean norm on $\mathbf{C}^n$. By comparison Theorem 5.4.1 states that

$$|\mathrm{perm}(M)| \leqslant \frac{n!}{n^{n/2}} \prod_{i=1}^n |v_i|.$$

In this section we deduce two consequences of Theorem 5.4.1, neither of them original: a version of entropy subadditivity for the symmetric group, and a concentration-of-measure result for a statistic of Hoeffding.

5.4.1 Entropy subadditivity for the symmetric group

Given $f : S_n \to [0, \infty)$ with $\alpha = \int f$ we define the entropy of f to be

$$S(f) = \int_{S_n} (f/\alpha) \log(f/\alpha).$$

To be more precise we might call $S(f)$ the Kullback–Liebler divergence of $(f/\alpha)\, d\pi$ from uniform, but we will use the shorter term for simplicity. Similarly, given $g : \Omega \to [0, \infty)$ with $\beta = \int g$ we define

$$S(g) = \int_{\Omega} (g/\beta) \log(g/\beta).$$

All logarithms are of course taken to the natural base.

In the coup de grâce of our argument we will apply the following entropy-subadditivity inequality.

Theorem 5.4.2 (Subadditivity of entropy)**.** *Suppose $f : S_n \to [0, \infty)$. Then*

$$S(f) \geqslant \frac{1}{2} \sum_{i \in \Omega} S(p_i f).$$

Note that this is much stronger than what one gets from just applying usual entropy subadditivity to f as a function $[n]^n \to [0, \infty)$.

Theorems 5.4.1 and 5.4.2 are more closely related than it may appear, as shown in some generality by Carlen and Cordero-Erausquin [CC09]. We repeat the rather simple deduction of Theorem 5.4.2 from Theorem 5.4.1 here for the convenience of the reader.

Proof of Theorem 5.4.2. Put $\alpha = \int f$. Define $f' : S_n \to [0, \infty)$ by

$$f'(\pi) = \prod_{i=1}^{n} p_i f(\pi(i))^{1/2},$$

and put $\alpha' = \int f'$. Then by Jensen's inequality we have

$$\begin{aligned} 0 &\leqslant \int_{S_n} (f/\alpha) \log \left(\frac{f/\alpha}{f'/\alpha'} \right) \\ &= S(f) - \frac{1}{2} \sum_{i=1}^{n} \int_{S_n} (f/\alpha) \log p_i f(\pi(i)) + \log \alpha' \\ &= S(f) - \frac{1}{2} \sum_{i=1}^{n} \int_{\Omega} (p_i f/\alpha) \log p_i f + \log \alpha' \\ &= S(f) - \frac{1}{2} \sum_{i=1}^{n} S(p_i f) - \frac{n}{2} \log \alpha + \log \alpha'. \end{aligned} \tag{5.9}$$

On the other hand by Theorem 5.4.1 we have

$$\alpha' = \int_{S_n} \prod_{i=1}^{n} p_i f(\pi(i))^{1/2} \leqslant \prod_{i=1}^{n} \left(\int_{\Omega} p_i f \right)^{1/2} = \alpha^{n/2},$$

so $\log \alpha' \leqslant \frac{n}{2} \log \alpha$ and the theorem follows from (5.9). □

5.4.2 Concentration for Hoeffding's statistic

Given an $n \times n$ complex matrix (a_{ij}) we consider the sum

$$X = \sum_{i=1}^{n} a_{i\pi(i)},$$

where $\pi \in S_n$ is random permutation. The study of such sums goes back at least to Hoeffding [Hoe51], who proved a central limit theorem for X under suitable hypotheses, and so we refer to X as *Hoeffding's statistic*. More recently work on Hoeffding's statistic has been more or less wedded to Stein's method of exchangeable pairs, starting with Bolthausen's [Bol84] Berry–Esseen-type estimate for the error in Hoeffding's theorem, and following with the work of Chatterjee [Cha07], who proved the first nonasymptotic concentration-type result for such sums.

In the next section we will need the following Bernstein-type concentration inequality for Hoeffding's statistic, which was proved in the more general context of random matrix theory by Mackey, Jordan, Chen, Farrell, and Tropp [Mac+14, Corollary 10.3], using an extension of Chatterjee's method.

Theorem 5.4.3. *Let (a_{ij}) be an $n \times n$ matrix such that $\sum_{i,j=1}^{n} a_{ij} = 0$ and such that $|a_{ij}| \leqslant M$ for each i, j. Let $v = \frac{1}{n} \sum_{i,j=1}^{n} |a_{ij}|^2$. Let $\pi \in S_n$ be chosen uniformly at random, and let*

$$X = \sum_{i=1}^{n} a_{i\pi(i)}.$$

Then for all $t > 0$ we have

$$\mathbf{P}(|X| > t) \leqslant 2\exp\left(\frac{-ct^2}{v + Mt}\right),$$

where c is some positive constant.

The purpose of this subsection is to give another proof of the above theorem, not relying on Stein's method, but instead relying on the Carlen–Lieb–Loss inequality Theorem 5.4.1. The main value of doing so is to reduce the reliance of the present chapter on results proved elsewhere, but it may also be of independent interest.

Proof of Theorem 5.4.3. By replacing (a_{ij}) with $(a_{ij} - \frac{1}{n}\sum_{j'} a_{ij'})$ if necessary and slightly reducing the constant c we may assume that $\sum_j a_{ij} = 0$ for each i: note that this operation does not change X, it can at worst double $\max|a_{ij}|$, and it can only reduce v. We may also assume that (a_{ij}) is real, for otherwise we may just deal with the real and imaginary parts separately.

Now for $\lambda > 0$ we have, by Theorem 5.4.1,

$$\begin{aligned}\mathbf{E}\exp(\lambda X) &= \int_{S_n}\prod_{i=1}^n \exp(\lambda a_{i\pi(i)}) \\ &\leqslant \prod_{i=1}^n\left(\frac{1}{n}\sum_{j=1}^n \exp(2\lambda a_{ij})\right)^{1/2}. \end{aligned} \tag{5.10}$$

Define

$$h(x) = \frac{e^x - 1 - x}{x^2} = \sum_{k=0}^\infty \frac{x^k}{(k+2)!}.$$

Then

$$\begin{aligned}\frac{1}{n}\sum_{j=1}^n \exp(2\lambda a_{ij}) &= \frac{1}{n}\sum_{j=1}^n \left(1 + 2\lambda a_{ij} + 4\lambda^2 a_{ij}^2 h(2\lambda a_{ij})\right) \\ &= 1 + \frac{1}{n}\sum_{j=1}^n 4\lambda^2 a_{ij}^2 h(2\lambda a_{ij}) \\ &\leqslant 1 + 4\lambda^2\left(\frac{1}{n}\sum_{j=1}^n a_{ij}^2\right) h(2\lambda M) \\ &\leqslant \exp\left(4\lambda^2\left(\frac{1}{n}\sum_{j=1}^n a_{ij}^2\right) h(2\lambda M)\right),\end{aligned}$$

so from (5.10) and the simple bound

$$h(x) \leqslant \sum_{k=0}^\infty x^k = \frac{1}{1-x} \qquad (0 < x < 1),$$

we have

$$\mathbf{E}\exp(\lambda X) \leqslant \exp\left(\frac{2\lambda^2 v}{1-2\lambda M}\right)$$

for $2\lambda M < 1$. The claimed result now follows by bounding

$$\mathbf{P}(X > t) = \mathbf{P}(\exp(\lambda X) > e^{\lambda t}) \leqslant e^{-\lambda t}\mathbf{E}\exp(\lambda X) \leqslant \exp\left(-\lambda t + \frac{2\lambda^2 v}{1-2\lambda M}\right)$$

and putting

$$\lambda = \frac{t}{4v + 2Mt},$$

and similarly bounding $\mathbf{P}(-X > t)$. □

The reader familiar with the proof of the usual Bernstein inequality may recognize that from (5.10) onwards all we have done is reproduce the that proof. Indeed, if Y is the sum of n independent random variables, the ith of which takes values $a_{i1}, \ldots, a_{in}$ each with probability $1/n$, then (5.10) states that

$$\mathbf{E}\exp(\lambda X) \leqslant (\mathbf{E}\exp(2\lambda Y))^{1/2},$$

so it suffices to extract from the proof of the usual Bernstein inequality an upper bound for $\mathbf{E}\exp(2\lambda Y)$.

5.5 Refined concentration for rearrangements

In this section we prove a refined concentration estimate for Hoeffding's statistic

$$X = \sum_{i=1}^{n} a_{i\pi(i)}$$

under the hypothesis that $a_{ij} = u_i v_j$ for some (u_i) and (v_i) for which we have some sort of entropy control. Moreover we are particularly interested in the concentration from below, which in certain regimes we expect to be stronger than the concentration from above.

Theorem 5.5.1. *Let $f : S_n \to [0,1]$ be a function with $\int f = \alpha$. Let $g_1, g_2 : \Omega \to [0,1]$ be functions with $\int g_1 = \beta$ and $\int g_2 = \gamma$. Then*

$$\begin{aligned} -(\langle f * g_1, g_2\rangle - \alpha\beta\gamma) \lesssim\ & \frac{\alpha\|g_1 - \beta\|_2\|g_2 - \gamma\|_2 \log n}{n^{1/2}} \\ & + \frac{\alpha^{1/2}\beta^{1/2}\gamma^{1/2}(\beta^{1/2} + \gamma^{1/2})S(g_1)^{1/2}S(g_2)^{1/2}(\log n)^{5/2}}{n^{1/2}} \\ & + O(n^{-99}). \end{aligned}$$

Lemma 5.5.2. *Let $h_1, h_2 : \Omega \to [0,1]$ be functions such that h_i is supported on a set H_i of density δ_i, and such that $1/2 \leqslant h_i \leqslant 1$ on H_i. Let $f : S_n \to [0,1]$ be a function with $\int f = \alpha$. Then if $\delta_1\delta_2 \gtrsim n^{-1}$ we have*

$$|\langle f * h_1, h_2\rangle - \alpha \int h_1 \int h_2| \lesssim \frac{\alpha \delta_1^{1/2}\delta_2^{1/2}\log n}{n^{1/2}} + O(n^{-100}),$$

while if $\delta_1\delta_2 \lesssim n^{-1}$ we have

$$-\alpha\delta_1\delta_2 \lesssim \left(\langle f * h_1, h_2\rangle - \alpha\int h_1 \int h_2\right) \lesssim \min\left(\frac{\alpha \log n}{n} + O(n^{-100}), \delta_1\delta_2\right).$$

To explain the two cases appearing in Lemma 5.5.2, let us momentarily think of h_i as the indicator of H_i. The inner product $\langle f * h_1, h_2\rangle/\alpha$ is then the density of a random intersection $\pi(H_1) \cap H_2$, where π is chosen randomly according to f/α. If H_1 and H_2 are not too small then we expect $|\pi(H_1) \cap H_2|$ to be highly concentrated around $\delta_1\delta_2 n$ with a Gaussian-type tail: this is the first case in the lemma. However if H_1 and H_2 are small then $|\pi(H_1) \cap H_2|$ has a Poisson-type distribution, so we expect $\pi(H_1) \cap H_2$ to be nonempty with probability about $\delta_1\delta_2$, and in any case almost surely bounded in size by about $\log n$: this is the second case in the lemma. The lower bound in the second case is trivial.

Proof. Apply Theorem 5.4.3 to $a_{ij} = (h_1(i) - \int h_1)(h_2(j) - \int h_2)$, noting that $|a_{ij}| \leqslant 1$ and

$$\frac{1}{n}\sum_{i,j=1}^n a_{ij}^2 = \frac{1}{n}\sum_{i=1}^n (h_1(i) - \int h_1)^2 \sum_{j=1}^n (h_2(j) - \int h_2)^2 \lesssim \delta_1\delta_2 n.$$

The result is that

$$\mathbf{P}(n|\langle \pi * (h_1 - \int h_1), (h_2 - \int h_2)\rangle| > t) \leqslant 2\exp\left(\frac{-ct^2}{\delta_1\delta_2 n + t}\right).$$

Thus for every $t > 0$ we have

$$|\langle f * (h_1 - \int h_1), (h_2 - \int h_2)\rangle| \lesssim \frac{\alpha t}{n} + 2\exp\left(\frac{-ct^2}{\delta_1\delta_2 n + t}\right).$$

For the first part of the lemma put $t = C\delta_1^{1/2}\delta_2^{1/2}n^{1/2}\log n$ for some constant C. Then we obtain

$$|\langle f * (h_1 - \int h_1), (h_2 - \int h_2)\rangle| \lesssim_C \frac{\alpha\delta_1^{1/2}\delta_2^{1/2}\log n}{n^{1/2}} + 2\exp\left(\frac{-cC^2(\log n)^2}{1 + (\delta_1\delta_2 n)^{-1/2}C\log n}\right).$$

If $\delta_1\delta_2 \gtrsim n^{-1}$ and C is sufficiently large it follows that

$$|\langle f * (h_1 - \int h_1), (h_2 - \int h_2)\rangle| \lesssim \frac{\alpha\delta_1^{1/2}\delta_2^{1/2}\log n}{n^{1/2}} + O(n^{-100}),$$

as claimed.

For the second part of the lemma put $t = C \log n$ for some constant C. Then we obtain

$$|\langle f * (h_1 - \int h_1), (h_2 - \int h_2)\rangle| \lesssim_C \frac{\alpha \log n}{n} + 2\exp\left(\frac{-cC^2(\log n)^2}{\delta_1\delta_2 n + C\log n}\right).$$

Now if $\delta_1\delta_2 \lesssim n^{-1}$ and C is sufficiently large it follows that

$$|\langle f * (h_1 - \int h_1), (h_2 - \int h_2)\rangle| \lesssim \frac{\alpha \log n}{n} + O(n^{-100}).$$

The remaining inequalities asserted by the lemma are trivial: just note that

$$\langle f * h_1, h_2\rangle \leqslant \langle 1 * h_1, h_2\rangle = \int h_1 \int h_2 \lesssim \delta_1\delta_2,$$

and

$$\alpha \int h_1 \int h_2 \lesssim \alpha\delta_1\delta_2. \qquad \square$$

We will deduce Theorem 5.5.1 from Lemma 5.5.2 using a dyadic decomposition, but first we need two basic entropy computations.

Lemma 5.5.3. *Let $g : \Omega \to [0,1]$ be a function such that $\int g = \beta$ and such that $g \leqslant \beta - t$ on a set of density at least δ, where $t, \delta > 0$. Then*

$$S(g) \gtrsim \frac{\delta t^2}{\beta^2}.$$

Proof. We must have $t \leqslant \beta$, so by replacing t with $t/100$ if necessary we may assume that $t/\beta \leqslant 1/100$. Similarly, by reducing δ if necessary we may assume that $\delta \leqslant 1/2$ and that δn is an integer. Now by convexity $S(g)$ is minimized under the stated conditions when $g = \beta - t$ on a set of density δ and otherwise equal to $\beta + \frac{\delta}{1-\delta}t$, and in this case

$$S(g) = \delta\left(1 - \frac{t}{\beta}\right)\log\left(1 - \frac{t}{\beta}\right) + (1-\delta)\left(1 + \frac{\delta}{1-\delta}\frac{t}{\beta}\right)\log\left(1 + \frac{\delta}{1-\delta}\frac{t}{\beta}\right).$$

By inserting the Taylor expansion

$$(1+x)\log(1+x) = x + x^2/2 + O(x^3) \tag{5.11}$$

we thus have

$$S(g) \geqslant \frac{1}{2}\frac{\delta}{1-\delta}\frac{t^2}{\beta^2} + O\left(\delta\frac{t^3}{\beta^3}\right) \gtrsim \frac{\delta t^2}{\beta^2}.$$

The last inequality follows from our assumption $t/\beta \leqslant 1/100$. $\square$

Lemma 5.5.4. *Let $g : \Omega \to [0,1]$ be a function such that $\int g = \beta$ and such that $g \geqslant \beta + t$ on a set of density at least δ, where $t, \delta > 0$. Then*

$$S(g) \gtrsim \min\left(\frac{\delta t}{\beta}, \frac{\delta t^2}{\beta^2}\right) \geqslant \frac{\delta t^2}{\beta}.$$

Proof. We must have $(\beta + t)\delta \leqslant \int g = \beta$, i.e.,

$$\frac{\delta}{1-\delta}\frac{t}{\beta} \leqslant 1,$$

so by replacing t with $t/100$ if necessary we may assume that

$$\frac{\delta}{1-\delta}\frac{t}{\beta} \leqslant \frac{1}{100}.$$

As before we may also assume that $\delta \leqslant 1/2$ and that δn is an integer. Now by convexity $S(g)$ is minimized under the stated conditions when $g = \beta + t$ on a set of density δ and otherwise equal to $\beta - \frac{\delta}{1-\delta}t$, and in this case

$$S(g) = \delta\left(1 + \frac{t}{\beta}\right)\log\left(1 + \frac{t}{\beta}\right) + (1-\delta)\left(1 - \frac{\delta}{1-\delta}\frac{t}{\beta}\right)\log\left(1 - \frac{\delta}{1-\delta}\frac{t}{\beta}\right).$$

By inserting (5.11) we thus have

$$S(g) \geqslant \delta\left(1 + \frac{t}{\beta}\right)\log\left(1 + \frac{t}{\beta}\right) - \delta\frac{t}{\beta} + \frac{1}{2}\frac{\delta^2}{(1-\delta)}\frac{t^2}{\beta^2} + O\left(\frac{\delta^3 t^3}{\beta^3}\right).$$

Now we separate into cases depending on the size of t/β. If $t/\beta \geqslant 1$ then we have

$$S(g) \geqslant \delta\frac{t}{\beta}(2\log 2 - 1) + O\left(\frac{\delta^2 t^2}{\beta^2}\right) \gtrsim \frac{\delta t}{\beta}.$$

On the other hand if $t/\beta \leqslant 1$ then by reducing t if necessary we may assume that $t/\beta \leqslant 1/100$, and then by inserting (5.11) again we have

$$S(g) \geqslant \frac{1}{2}\frac{\delta}{1-\delta}\frac{t^2}{\beta^2} + O\left(\frac{\delta t^3}{\beta^3}\right) \gtrsim \frac{\delta t^2}{\beta^2}.$$

As before we used our assumption about the size of $\delta t/\beta$ or t/β to justify the absorption of the error terms. □

Proof of Theorem 5.5.1. Write

$$g_i - \int g_i = \sum_s g_i^s + O(n^{-100}) = \sum_s (g_i^s - \int g_i^s) + O(n^{-100}),$$

where s ranges over all s of the form $\pm 2^{-k}$ for which $n^{-100} \leqslant |s| \leqslant 1$, and where g_i^s is defined to be equal to $g_i - \int g_i$ where $g_i - \int g_i$ has the same sign as s and $|s|/2 < |g_i - \int g_i| \leqslant |s|$ and zero elsewhere. Then

$$\langle f * g_1, g_2\rangle - \alpha\beta\gamma = \sum_{s,t} \big(\langle f * g_1^s, g_2^t\rangle - \alpha \int g_1^s \int g_2^t\big) + O(n^{-100}).$$

For each s, t we apply Lemma 5.5.2 with $h_1 = g_1^s/s$ and $h_t = g_2^t/t$. Let δ_1^s be the density of points where $g_1 - \int g_1$ has the same sign as s and $|s|/2 < |g_1 - \int g_1| \leqslant |s|$ and let δ_2^t be the density of points where $g_2 - \int g_2$ has the same sign as t and $|t|/2 < |g_2 - \int g_2| \leqslant |t|$. If $\delta_1^s \delta_2^t \gtrsim 1/n$ then we get the bound

$$\big|\langle f * g_1^s, g_2^t\rangle - \alpha \int g_1^s \int g_2^t\big| \lesssim \frac{\alpha |s||t| (\delta_1^s)^{1/2} (\delta_2^t)^{1/2} \log n}{n^{1/2}} + O(n^{-100}),$$

and the total contribution from all such cases is bounded by

$$\sum_{s,t} \left(\frac{\alpha |s||t| (\delta_1^s)^{1/2} (\delta_2^t)^{1/2} \log n}{n^{1/2}} + O(n^{-100}) \right)$$
$$\lesssim \frac{\alpha \|g_1 - \int g_1\|_2 \|g_2 - \int g_2\|_2 \log n}{n^{1/2}} + O(n^{-99}).$$

Now consider the cases in which $\delta_1^s \delta_2^t \lesssim 1/n$ and in which s and t have the same sign. By Lemma 5.5.2 we have

$$-\big(\langle f * g_1^s, g_2^t\rangle - \alpha g_1^s g_2^t\big) \lesssim \alpha |s||t| \delta_1^s \delta_2^t \lesssim \frac{\alpha |s||t| (\delta_1^s)^{1/2} (\delta_2^t)^{1/2}}{n^{1/2}},$$

so the total contribution from these cases is again acceptable.

Finally consider the cases in which $\delta_1^s \delta_2^t \lesssim 1/n$ and in which s and t have opposite sign, say $s < 0$ and $t > 0$. By Lemmas 5.5.3 and 5.5.4 we have

$$S(g_1) \gtrsim \frac{\delta_1^s s^2}{\beta^2}$$

and

$$S(g_2) \gtrsim \frac{\delta_2^t t^2}{\gamma},$$

so

$$S(g_1)^{1/2} S(g_2)^{1/2} \gtrsim \frac{|s||t| (\delta_1^s \delta_2^t)^{1/2}}{\beta \gamma^{1/2}}.$$

Thus by Lemma 5.5.2 we can bound

$$|\langle f * g_1^s, g_2^t\rangle - \alpha \int g_1^s \int g_2^t| \lesssim |s||t| \left(\frac{\alpha \log n}{n}\right)^{1/2} \big(\delta_1^s \delta_2^t\big)^{1/2} + O(n^{-100})$$
$$\lesssim \frac{\alpha^{1/2} \beta \gamma^{1/2} S(g_1)^{1/2} S(g_2)^{1/2} (\log n)^{1/2}}{n^{1/2}} + O(n^{-100}).$$

If $s < 0$ and $t > 0$ then we get the analogous bound

$$|\langle f * g_1^s, g_2^t\rangle - \alpha \int g_1^s \int g_2^t| \lesssim \frac{\alpha^{1/2}\beta^{1/2}\gamma S(g_1)^{1/2}S(g_2)^{1/2}(\log n)^{1/2}}{n^{1/2}} + O(n^{-100}).$$

The number of choices of s and t is bounded by $(\log n)^2$, so the total contribution from all these cases is bounded by

$$\frac{\alpha^{1/2}\beta^{1/2}\gamma^{1/2}(\beta^{1/2}+\gamma^{1/2})S(g_1)^{1/2}S(g_2)^{1/2}(\log n)^{5/2}}{n^{1/2}} + O(n^{-99}). \qquad \square$$

5.6 Bounding the second term in (5.6)

Proof of Theorem 5.1.2. Let $f = 1_X$, $g = 1_Y$, and $h = 1_Z$, where $X, Y, Z \subset A_n$ have densities $\alpha, \beta, \gamma \geqslant n^{-O(1)}$ respectively. Then the first term in (5.6) is

$$\alpha\beta\gamma,$$

the third term is bounded by

$$c\alpha^{1/2}\beta^{1/2}\gamma^{1/2}/n,$$

and the second term is, by (5.7) and Theorem 5.5.1,

$$\begin{aligned}
(n-1)&\langle \widehat{f}(\sigma)\widehat{g}(\sigma), \widehat{h}(\sigma)\rangle_{\mathrm{HS}} \\
&\sim \sum_{i\in\Omega} \langle f * (p_i g - \beta), (p_i h - \gamma)\rangle \\
&\gtrsim -\frac{\alpha \log n}{n^{1/2}} \sum_{i\in\Omega} \|p_i g - \beta\|_2 \|p_i h - \gamma\|_2 \\
&\quad - \frac{\alpha^{1/2}\beta^{1/2}\gamma^{1/2}(\beta^{1/2}+\gamma^{1/2})(\log n)^{5/2}}{n^{1/2}} \sum_{i\in\Omega} S(p_i g)^{1/2} S(p_i h)^{1/2} \\
&\quad + O(n^{-98}).
\end{aligned}$$

By Cauchy–Schwarz and the Parseval remnant (5.8), the first term here is bounded in magnitude by

$$\frac{\alpha \log n}{n^{1/2}} \left(\sum_{i\in\Omega} \|p_i g - \beta\|_2^2\right)^{1/2} \left(\sum_{i\in\Omega} \|p_i h - \gamma\|_2^2\right)^{1/2} \lesssim \frac{\alpha\beta^{1/2}\gamma^{1/2}\log n}{n^{1/2}}.$$

Similarly, by Cauchy–Schwarz and subadditivity of entropy (Theorem 5.4.2) the second term is bounded in magnitude by

$$\begin{aligned}
&\frac{\alpha^{1/2}\beta^{1/2}\gamma^{1/2}(\beta^{1/2}+\gamma^{1/2})(\log n)^{5/2}}{n^{1/2}}\left(\sum_{i\in\Omega}S(p_i g)\right)^{1/2}\left(\sum_{i\in\Omega}S(p_i h)\right)^{1/2}\\
&\lesssim \frac{\alpha^{1/2}\beta^{1/2}\gamma^{1/2}(\beta^{1/2}+\gamma^{1/2})(\log n)^{5/2}}{n^{1/2}}(\log\beta^{-1})^{1/2}(\log\gamma^{-1})^{1/2}\\
&\lesssim \frac{\alpha^{1/2}\beta^{1/2}\gamma^{1/2}(\beta^{1/2}+\gamma^{1/2})(\log n)^{7/2}}{n^{1/2}}.
\end{aligned}$$

Thus we deduce that $\langle f*g,h\rangle \geqslant (1+o(1))\alpha\beta\gamma$ provided that

$$\begin{aligned}
\frac{\alpha^{1/2}\beta^{1/2}\gamma^{1/2}}{n} &\ll \alpha\beta\gamma,\\
\frac{\alpha\beta^{1/2}\gamma^{1/2}(\log n)}{n^{1/2}} &\ll \alpha\beta\gamma,\\
\frac{\alpha^{1/2}\beta\gamma^{1/2}(\log n)^{7/2}}{n^{1/2}} &\ll \alpha\beta\gamma,\\
\frac{\alpha^{1/2}\beta^{1/2}\gamma(\log n)^{7/2}}{n^{1/2}} &\ll \alpha\beta\gamma, \text{ and}\\
n^{-98} &\ll \alpha\beta\gamma.
\end{aligned}$$

In other words what we require is that

$$\min(\alpha\beta,\alpha\gamma,\beta\gamma) \gg (\log n)^7/n. \qquad \square$$

5.7 Open questions

The most obvious outstanding open question is whether the logarithms can be removed from Theorem 5.1.2. Specifically, does the largest product-free subset of A_n have density $O(n^{-1/2})$? Can you say anything about the extremal examples? It is possible that all near-extremizers look roughly like the first example in Section 5.2, or its inverse, but this may be difficult to quantify, and even more difficult to prove.

Another obvious outstanding open question is whether a one-sided product-mixing phenomenon persists in other groups for densities lower than that given by Theorem 5.1.1. For example take $G=\mathrm{SL}_2(p)$. For this group $m\sim p$. By Theorem 5.1.1 there is two-sided product mixing for sets of density at least $p^{-1/3}$, by Proposition 5.2.2 there is no two-sided product mixing for sets of density less than $p^{-1/3}$, and by Proposition 5.2.1 there is no product mixing at all below density $p^{-1/2}$. Do we have one-sided product mixing for sets of density between $p^{-1/2}$ and $p^{-1/3}$?

Another great question, which has been asked before by both Kedlaya [Ked98] and Gowers [Gow08], is about the product-mixing properties of $\mathrm{SU}(n)$. To make the question concrete, what is the measure of the largest product-free subset of $\mathrm{SU}(n)$? By straightforward adaptation of Theorem 5.1.1 it is at most $O(n^{-1/3})$, but the only lower bounds we know have the form c^n for some $c < 1$. Apart from being an interesting and natural question in its own right, answering this question may be relevant for understanding the product-mixing behaviour of groups not having a permutation representation of dimension $\sim m$.

References

[BNP08] L. Babai, N. Nikolov, and L. 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.

[Bol84] E. Bolthausen. "An estimate of the remainder in a combinatorial central limit theorem". *Z. Wahrsch. Verw. Gebiete* 66.3 (1984), pp. 379–386. ISSN: 0044-3719. DOI: 10.1007/BF00533704.

[CC09] E. A. Carlen and D. Cordero-Erausquin. "Subadditivity of the entropy and its relation to Brascamp-Lieb type inequalities". *Geom. Funct. Anal.* 19.2 (2009), pp. 373–405. ISSN: 1016-443X. DOI: 10.1007/s00039-009-0001-y.

[CLL04] E. A. Carlen, E. H. Lieb, and M. Loss. "A sharp analog of Young's inequality on S^N and related entropy inequalities". *J. Geom. Anal.* 14.3 (2004), pp. 487–520. ISSN: 1050-6926. DOI: 10.1007/BF02922101.

[CLL06] E. A. Carlen, E. H. Lieb, and M. Loss. "An inequality of Hadamard type for permanents". *Methods Appl. Anal.* 13.1 (2006), pp. 1–17. ISSN: 1073-2772. DOI: 10.4310/MAA.2006.v13.n1.a1.

[Cha07] S. Chatterjee. "Stein's method for concentration inequalities". *Probab. Theory Related Fields* 138.1-2 (2007), pp. 305–321. ISSN: 0178-8051. DOI: 10.1007/s00440-006-0029-y.

[Ebe16] S. Eberhard. "Product mixing in the alternating group". *Discrete Analysis* 2016:2 (2016). DOI: 10.19086/da.610.

[Gow08] W. T. Gowers. "Quasirandom groups". *Combin. Probab. Comput.* 17.3 (2008), pp. 363–387. DOI: 10.1017/S0963548307008826.

[Hoe51] W. Hoeffding. "A combinatorial central limit theorem". *Ann. Math. Statistics* 22 (1951), pp. 558–566. ISSN: 0003-4851.

[Ked97] K. S. Kedlaya. "Large product-free subsets of finite groups". *J. Combin. Theory Ser. A* 77.2 (1997), pp. 339–343. ISSN: 0097-3165. DOI: 10.1006/jcta.1997.2715.

[Ked98] K. S. Kedlaya. "Product-free subsets of groups". *Amer. Math. Monthly* 105.10 (1998), pp. 900–906. ISSN: 0002-9890. DOI: 10.2307/2589282.

[Mac+14] L. Mackey, M. I. Jordan, R. Y. Chen, B. Farrell, and J. A. Tropp. "Matrix concentration inequalities via the method of exchangeable pairs". *Ann. Probab.* 42.3 (2014), pp. 906–945. ISSN: 0091-1798. DOI: 10.1214/13-AOP892.

[Ras77] R. Rasala. "On the minimal degrees of characters of S_n". *J. Algebra* 45.1 (1977), pp. 132–181. ISSN: 0021-8693.

[Tao14] T. Tao. *Hilbert's fifth problem and related topics*. Vol. 153. Graduate Studies in Mathematics. American Mathematical Society, Providence, RI, 2014, pp. xiv+338. ISBN: 978-1-4704-1564-8.

[Wig10] A. Wigderson. *Lecture notes for the 22nd McGill invitational workshop on computational complexity*. 2010. URL: http://www.math.ias.edu/~avi/TALKS/additive-lectures-v2.pdf.