Glasgow Math. J. **53** (2011) 1–50. © Glasgow Mathematical Journal Trust 2010.
doi:10.1017/S0017089510000546.

AN INVERSE THEOREM FOR THE GOWERS U^4-NORM

BEN GREEN
Centre for Mathematical Sciences, Wilberforce Road, Cambridge CB3 0WA, UK
e-mail: b.j.green@dpmms.cam.ac.uk

TERENCE TAO
UCLA Department of Mathematics, Los Angeles, CA 90095-1555, USA
e-mail: tao@math.ucla.edu

and TAMAR ZIEGLER
Department of Mathematics, Technion - Israel Institute of Technology, Haifa 32000, Israel
e-mail: tamarzr@tx.technion.ac.il

(Received 30 November 2009; accepted 18 June 2010; first published online 25 August 2010)

Abstract. We prove the so-called *inverse conjecture for the Gowers U^{s+1}-norm* in the case $s = 3$ (the cases $s < 3$ being established in previous literature). That is, we show that if $f : [N] \to \mathbb{C}$ is a function with $|f(n)| \leqslant 1$ for all n and $\|f\|_{U^4} \geqslant \delta$ then there is a bounded complexity 3-step nilsequence $F(g(n)\Gamma)$ which correlates with f. The approach seems to generalise so as to prove the inverse conjecture for $s \geqslant 4$ as well, and a longer paper will follow concerning this.

By combining the main result of the present paper with several previous results of the first two authors one obtains the generalised Hardy–Littlewood prime-tuples conjecture for any linear system of complexity at most 3. In particular, we have an asymptotic for the number of 5-term arithmetic progressions $p_1 < p_2 < p_3 < p_4 < p_5 \leqslant N$ of primes.

2010 *Mathematics Subject Classification.* Primary: 11B30; Secondary: 37A17.

NOTATION . *By a* 1*-bounded function on a set X we mean a function $f : X \to \mathbb{C}$ with $|f(x)| \leqslant 1$ for all $x \in X$. If the cardinality $|X|$ of X is finite and non-zero, we write $\mathbb{E}_{x\in X} f(x)$ for $|X|^{-1} \sum_{x\in X} f(x)$. Throughout the paper the letter M will refer to a large positive 'complexity' quantity, normally introduced in each statement of a lemma, proposition or theorem. The letters c and C are reserved for absolute constants with $0 < c < 1 < C$; different instances of the notation will generally denote different absolute constants. If $x \in \mathbb{R}$ we will write $\lfloor x \rfloor$ for the greatest integer less than or equal to x, and $\{x\} := x - \lfloor x \rfloor$. If N is a positive integer then we write $[N] := \{1, \ldots, N\}$.*

1. Introduction. This paper concerns a special case of a family of conjectures named the *Inverse Conjectures for the Gowers norms* by the first two authors. For each integer $s \geqslant 1$ the inverse conjecture GI(s), whose statement we recall shortly, describes the structure of 1-bounded functions $f : [N] \to \mathbb{C}$ whose $(s + 1)$st Gowers norm $\|f\|_{U^{s+1}}$ is large. These conjectures together with a good deal of motivation and background to them are discussed in [**10, 11, 13**]. The conjectures GI(1) and GI(2) are already known, the former being straightforward application of Fourier analysis and the latter being the main result of [**11**]. The aim of the present paper is to establish the first unknown

case, that of GI(3), using what is in essence a method which seems to generalise to prove GI(s) in general.

We have taken advantage of some shortcuts and explicit calculations that are specific to the $s = 3$ case, hoping that this will render the paper somewhat appetising as an *hors d'œuvres* for the general case. The general case will, furthermore, be phrased in the language of non-standard analysis since this provides a very effective framework in which to manage the complicated hierarchies of parameters that appear here. We offer the present paper to those readers who are not immediately comfortable with the non-standard language; it also serves as an illustration of the point, to be made in the longer paper to follow, that our arguments may be taken out of the choice-dependent realm of non-standard analysis and, in particular, can lead to effective bounds (albeit extremely weak ones).

We begin by recalling the definition of the Gowers norms. If G is a finite abelian group and if $f : G \to \mathbb{C}$ is a function then we define

$$\|f\|_{U^k(G)} := \left(\mathbb{E}_{x,h_1,\dots,h_k \in G} \Delta_{h_1} \dots \Delta_{h_k} f(x)\right)^{1/2^k},$$

where $\Delta_h f$ is the multiplicative derivative

$$\Delta_h f(x) := f(x+h)\overline{f(x)}.$$

In this paper, we will be concerned with functions on $[N]$, which is not quite a group. To define the Gowers norms of a function $f : [N] \to \mathbb{C}$, set $G := \mathbb{Z}/\tilde{N}\mathbb{Z}$ for some integer $\tilde{N} \geqslant 2^k N$, define a function $\tilde{f} : G \to \mathbb{C}$ by $\tilde{f}(x) = f(x)$ for $x = 1, \dots, N$ and $\tilde{f}(x) = 0$ otherwise, and set $\|f\|_{U^k[N]} := \|\tilde{f}\|_{U^k(G)} / \|1_{[N]}\|_{U^k(G)}$, where $1_{[N]}$ is the indicator function of $[N]$. It is easy to see that this definition is independent of the choice of $\tilde{N}$, and so for definiteness one could take $\tilde{N} := 2^k N$. Henceforth we shall write simply $\|f\|_{U^k}$, rather than $\|f\|_{U^k[N]}$, since all Gowers norms will be on $[N]$. One can show that $\|\cdot\|_{U^k}$ is indeed a norm for any $k \geqslant 2$, though we shall not need this here.

The *Inverse conjecture for the Gowers U^{s+1}-norm* posits an answer to the following question.

QUESTION 1.1. Suppose that $f : [N] \to \mathbb{C}$ is a 1-bounded function and let $\delta > 0$ be a positive real number. What can be said if $\|f\|_{U^{s+1}} \geqslant \delta$?

The conjecture made in [**13**] is that f must correlate with a certain rather algebraic object called an s-step nilsequence. In the light of subsequent work [**14, 15**] it seems natural to work with a somewhat more general object called a degree s *polynomial* nilsequence. We recall now the bald definition; for much more motivation and examples, see the introduction to [**14**].

DEFINITION 1.2 (Polynomial nilsequence). Let G be a connected, simply-connected nilpotent Lie group. By a *filtration* $G_\bullet$ of degree s we mean a nested sequence $G = G_{(0)} = G_{(1)} \supseteq G_{(2)} \supseteq \dots \supseteq G_{(s+1)} = \{\mathrm{id}\}$ with the property that $[G_{(i)}, G_{(j)}] \subseteq G_{(i+j)}$. By a polynomial sequence adapted to $G_\bullet$ we mean a map $g : \mathbb{Z} \to G$ such that $\partial_{h_i} \dots \partial_{h_1} g \in G_{(i)}$ for all $h_1, \dots, h_i \in \mathbb{Z}$, where $\partial_h \psi(n) := \psi(n+h)\psi(n)^{-1}$. Let $\Gamma \leqslant G$ be a discrete and cocompact subgroup, so that the quotient G/Γ is a nilmanifold, and assume that each of the $G_{(i)}$ are *rational* subgroups[1]. If $F : G/\Gamma \to \mathbb{C}$ is a 1-bounded,

[1]One may define rationality topologically, by stipulating that the $G_{(i)}$ are connected Lie subgroups of G and that $\Gamma \cap G_{(i)}$ is a cocompact subgroup of $G_{(i)}$. Some readers may wish to think more concretely, in terms of the existence of a Mal'cev basis as in [**14**, Definition 2.1].

https://doi.org/10.1017/S0017089510000546 Published online by Cambridge University Press

Lipschitz function then the sequence $(F(g(n)\Gamma))_{n\in\mathbb{Z}}$ is called a polynomial nilsequence of degree s.

REMARK. An important example of a filtration of a nilpotent group is the *lower central series* $G_0 \supseteq G_1 \supseteq G_2 \supseteq \ldots$, in which $G_0 = G_1 = G$, and $G_{i+1} := [G, G_i]$ for $i \geqslant 1$. It is classical (see, for example, [**3**]) that this is a filtration of degree s whenever G is s-step nilpotent. This is the minimal example of a filtration, since for any other filtration $G_\bullet$ one has $G_i \subseteq G_{(i)}$.

REMARK. An important fact about polynomial sequences adapted to a filtration $G_\bullet$ is that they form a group under pointwise multiplication: see [**20**] or [**14**, Proposition 6.2]. A polynomial sequence $g : \mathbb{Z} \to G$ can also be uniquely expressed as a Taylor expansion $g(n) = g_0^{\binom{n}{0}} g_1^{\binom{n}{1}} \ldots g_s^{\binom{n}{s}}$ for some $g_i \in G_{(i)}$ for $i = 0, 1, \ldots, s$, where $\binom{n}{i}$ is the usual binomial coefficient; see [**14**, Section 6].

REMARK. If G admits a filtration of degree s then, as we remarked above, G must be s-step nilpotent. On the other hand, the degree can exceed the step by an arbitrary amount. For instance, if $P : \mathbb{Z} \to \mathbb{R}/\mathbb{Z}$ is a polynomial of degree $d \geqslant 1$, then the function $e(P(n)) := e^{2\pi i P(n)}$ is a polynomial nilsequence of degree d, despite being associated to a nilmanifold $G/\Gamma = \mathbb{R}/\mathbb{Z}$ of step just 1.

Roughly speaking, the inverse conjecture GI(s) asserts that a 1-bounded function f has large U^{s+1}-norm if and only if it correlates with a degree s nilsequence. However, every aspect of this statement must be quantified in order to make a precise statement. The key issue here lies in defining the *complexity* of a nilsequence, a matter which was addressed in some detail in [**14**, Sec 2]. In this paper (fortunately) we can take a much rougher approach. If $\delta > 0$ is some parameter we shall simply say that the complexity of a polynomial nilsequence $(F(g(n)\Gamma))_{n\in\mathbb{Z}}$ is $O_\delta(1)$ if the following list of objects are bounded in a way that depends only on δ:

- $\dim G$;
- The rationality of some Mal'cev basis $\mathcal{X}$ for G/Γ (see [**14**, Definition 2.4]);
- The rationality of each subgroup $G_{(i)}$ in the filtration (see [**14**, Definition 2.5]);
- The Lipschitz norm of F, measured using the metric defined in [**14**, Definition 2.2].

We do not encourage the reader to read those definitions in detail at this stage. The important thing to note is that nothing is said about the polynomial sequence g, other than that it is adapted to the filtration $G_\bullet$.

We may now state the Inverse Conjecture for the Gowers U^{s+1}-norm, GI(s), properly.

CONJECTURE 1.3 (GI(s)). Suppose that $f : [N] \to \mathbb{C}$ is a 1-bounded function and that $\|f\|_{U^{s+1}} \geqslant \delta$. Then there is a degree s polynomial nilsequence $(F(g(n)\Gamma))_{n\in\mathbb{Z}}$ of complexity $O_\delta(1)$ such that $|\mathbb{E}_{n\in[N]} f(n)\overline{F(g(n)\Gamma)}| \gg_\delta 1$.

As hinted earlier, this is not quite the formulation of GI(s) originally given in [**13**, Section 8]. There, it was posited that f correlates with a *linear*[2] nilsequence $(F(g^n\Gamma))_{n\in\mathbb{Z}}$. One might now relabel this the *strong* inverse conjecture. In the longer paper to come we will show how this in fact follows from Conjecture 1.3. In the special case of the U^4-norm under consideration here, it is possible to verify the strong inverse conjecture quite directly by inspection, and we sketch this in Appendix F. We would, however,

[2]We remark that a linear nilsequence is not the same thing as a degree 1 nilsequence; a typical linear nilsequence on an s-step nilmanifold will have degree s.

https://doi.org/10.1017/S0017089510000546 Published online by Cambridge University Press

like to impress upon the reader our opinion that Conjecture 1.3 is the most natural one, a viewpoint that became apparent to the first two authors in the light of our paper [**14**]. Unfortunately [**13**] was written before that paper and hence operates under the assumption of the strong inverse conjecture. Relatively simple changes would be required to make all of the arguments there work under the assumption of Conjecture 1.3 however, the key issue being Section 11 of that paper.

The evidence for the inverse conjectures prior to the present work was a 'local version' due to Gowers [**8**], its truth in the cases $s = 1$ and 2 (see [**11**]) as well as the truth of analogues of the conjecture in both ergodic theory [**18, 28**] and in the 'finite field model' in which $[N]$ is replaced by $\mathbb{F}^n$ for some small prime field $\mathbb{F}$ [**1, 26**].

It is also known that this conjecture is necessary, in the following sense.

PROPOSITION 1.4 (Necessity of inverse conjecture). *Suppose that $f : [N] \to \mathbb{C}$ is a 1-bounded function, that $(F(g(n)\Gamma))_{n\in\mathbb{Z}}$ is a polynomial nilsequence of degree s and complexity $O_\delta(1)$, and that $|\mathbb{E}_{n\in[N]} f(n)\overline{F(g(n)\Gamma)}| \geqslant \delta$. Then $\|f\|_{U^{s+1}} \gg_\delta 1$.*

There is currently no proof of this written in the literature. In the case of *linear nilsequences* $(F(g^n\Gamma))_{n\in\mathbb{Z}}$ there are two different (albeit related) proofs in the literature: one in [**11**, Proposition 12.6] and the other in [**13**, Section 11]. The second of these proofs would generalise rather easily to the more general setting of degree s polynomial nilsequences $(F(g(n)\Gamma))_{n\in\mathbb{Z}}$, the key issue being to note that [**13**, Lemma E.4] is true for the values $(g(n + \omega \cdot h)\Gamma)_{\omega\in\{0,1\}^{s+1}}$, this being essentially [**14**, Proposition 6.5]. The reader will doubtless be relieved to hear that we recently discovered a very short proof of Proposition 1.4, and we give this in Appendix G. Note, however, that this proposition is included for motivation and interest only, and is not actually required in this paper.

Here, then is the main result of our paper.

THEOREM 1.5 (GI(3)). *The inverse conjecture for the U^4-norm,* GI(3)*, is true.*

As already remarked, in Appendix F we will also establish the strong form of the inverse conjecture for the U^4-norm, in the form given in [**13**, Section 8].

By combining this result with the previous results in [**13, 15**] we obtain a proof of what was referred to in [**13**] as the generalised Hardy–Littlewood conjecture for linear systems of complexity at most 3. In particular, we have the following.

THEOREM 1.6. *The number of quintuples of primes $p_1 < p_2 < p_3 < p_4 < p_5 \leqslant N$ in arithmetic progression is asymptotic to $\gamma N^2/\log^5 N$, where*

$$\gamma = \frac{27}{16}\prod_{p\geqslant 5}\frac{p^3(p-4)}{(p-1)^4}.$$

We refer the reader to [**13**] for further discussion. Several further applications of the GI(s) conjectures will be given in a forthcoming paper of the first two authors [**16**].

2. An outline of the proof. In this section we outline the argument we use to establish the inverse conjecture for the U^4-norm.

It is easy to show, and well-known, that if $\|f\|_{U^4} \geqslant \delta$ then there are $\gg \delta^C N$ values of h for which $\Delta_h f(n) := f(n + h)\overline{f(n)}$ has U^3-norm at least δ^C. Applying GI(2), it

https://doi.org/10.1017/S0017089510000546 Published online by Cambridge University Press

follows that for all these h we have

$$|\mathbb{E}_n \Delta_h f(n)\overline{\chi_h(n)}| \gg 1, \tag{2.1}$$

where $\chi_h(n)$ is a 2-step nilsequence (with complexity bounded uniformly in h).

Very roughly speaking, the aim is to show that these 2-step nilsequences 'line up' in such a way that they may be interpreted as the derivatives of a single 3-step object. To make this work and for ease of exposition it is convenient to assume that $\chi_h(n)$ is in fact equal to $e(\psi_h(n))$, where $\psi_h(n)$ is a bracket quadratic phase: a sum of terms of the form $\alpha_1 n\lfloor\alpha_2 n\rfloor$, $\alpha_3 n^2$ and $\alpha_4 n$. The link between these objects and 2-step nilsequences was explored in [**11**, Section 10] and will be recalled later in this paper. For the purposes of this discussion let us suppose that $\psi_h(n) = \alpha_h n\lfloor\beta_h n\rfloor$; this is something of a simplification of the true situation.

Here is a rough outline of the main steps we shall be taking to control the dependence of α_h and β_h on h. Suppose that (2.1) holds with $\chi_h(n) = e(\alpha_h n\lfloor\beta_h n\rfloor)$.

Step 1 (Reducing the h-dependence) We may assume (possibly after refining the set of h and modifying α_h and β_h somewhat) that β_h does not depend on h.

Step 2 (Approximate linearity of h-dependent frequency) We may assume (possibly after refining the set of h again) that α_h is approximately equal to a bracket linear form $\theta_1\{\theta_1' h\} + \cdots + \theta_d\{\theta_d' h\} + \theta h$.

Step 3 (Symmetry argument) Following *Step 2*, $\chi_h(n)$ is essentially $e(\psi_h(n))$ with the phase $\psi_h(n)$ being of the form $T(h, n, n)$, where $T(n_1, n_2, n_3)$ is a sum of terms of the form $\{\theta_1 n_1\}\theta_2 n_2\lfloor\theta_3 n_3\rfloor$. Not every such function $T(h, n, n)$ can be obtained as the 'derivative' of a 3-step object, however, and in order to make this assertion we need some additional symmetry properties of the 'generalised trilinear form' $T(n_1, n_2, n_3)$.

It may be of some interest to make a comparison between this strategy and that used in the proof of the U^3-inverse theorem [**11**]. If $\|f\|_{U^3} \geqslant \delta$ then for many h we have, once again,

$$|\mathbb{E}_n \Delta_h f(n)\overline{\chi_h(n)}| \gg 1,$$

but now $\chi_h(n)$ may be assumed to be nothing more complicated than a linear phase $e(\alpha_h n)$. The argument runs roughly as follows:

Step 2' (Approximate linearity of frequencies) At the possible expense of passing to a subset of the h, the frequencies α_h are approximately 'bracket-linear' in h, as above;

Step 3' (Symmetry argument) Following *Step 2*, $\chi_h(n)$ is essentially $e(\psi_h(n))$ with the phase $\psi_h(n)$ being of the form $T(h, n)$ where $T(n_1, n_2)$ is a sum of terms of the form $\{\theta_1 n_1\}\theta_2 n_2$. Not every such function $T(h, n)$ can be obtained as the 'derivative' of a 2-step object, however, and in order to make this assertion we need some additional symmetry properties of the form $T(n_1, n_2)$.

We note that *Step 2'* is essentially due to Gowers [**8**, Chapter 7], although one must apply a little extra geometry of numbers to get the precise conclusion we hint at here. *Step 3'* is due to the first two authors and is the main new result of [**11**], specifically Lemma 9.4 of that paper. Note that *Step 1* in the outline above did not feature at all in the proof of the U^3-inverse theorem and it is new to this paper.

Let us say a few words about how *Steps 1, 2* and *3* are accomplished. The key to almost all of our analysis is a straightforward adaption of a fundamental idea of

https://doi.org/10.1017/S0017089510000546 Published online by Cambridge University Press

Gowers [7], which proceeds from the assumption that

$$|\mathbb{E}_n \Delta_h f(n)\overline{\chi_h(n)}| \gg 1 \tag{2.2}$$

for many h and draws a conclusion involving just the $\chi_h(n)$, and not the function f. This argument is valid for any bounded functions $\chi_h(n)$ and we give it in Section 6.

The conclusion of that argument is that

$$\mathbb{E}_{n\in[N]}\chi_{h_1}(n)\chi_{h_2}(n+h_1-h_4)\overline{\chi_{h_3}(n)\chi_{h_4}(n+h_1-h_4)} \gg 1 \tag{2.3}$$

for many additive quadruples h_1, h_2, h_3, h_4, that is to say quadruples satisfying $h_1 + h_2 = h_3 + h_4$.

Steps 1,2 and *3* all involve interpreting this in the case that $\chi_h(n)$ is a 2-step object such as a bracket quadratic phase. One way to do this is to visualise

$$\chi_{h_1}(n)\chi_{h_2}(n+h_1-h_4)\overline{\chi_{h_3}(n)\chi_{h_4}(n+h_1-h_4)}$$

as a certain h_1, h_2, h_3, h_4-dependent nilsequence on a product of four nilmanifolds (one for each of the h_i), in which case (2.3) states that the underlying polynomial sequence $g_{h_1,h_2,h_3,h_4}(n)$ is far from equidistributed. This situation may then be studied using the distributional results on nilsequences contained in [**14**] in order to draw conclusions concerning the dependence on h of 'leading order' terms in the $\chi_h(n)$.

Steps 1 and *2* really only use the 'top-order' structure of (2.3) – that is to say the shifts $h_1 - h_4$ are not relevant. To handle *Step 3* these shifts cannot be ignored. In the general case the treatment of *Step 3* will involve another appeal to the distributional results on nilmanifolds in [**14**], but in the case of the U^4-norm a much more hands-on approach involving Bohr sets may be employed, and it is this argument that we give here.

The following deliberately vague discussion may perhaps be helpful. Suppose that $|f(n)| = 1$ for all n and that $\Delta_h f \sim \chi_h$ (where we are not attaching any real meaning to $\sim$). Then we have the 'cocycle identity' $\Delta_{h+k} f(n) = \Delta_h f(n+k)\Delta_k f(n)$, which translates to $\chi_{h+k}(n) \sim \chi_h(n+k)\chi_k(n)$. Imagining that the shift $n \mapsto n+k$ does not affect the 'top-order structure' of $\chi_h(n+k)$, we have the approximate linearity condition

$$\chi_{h+k} \sim \chi_h \chi_k \quad \text{to top order.}$$

Roughly speaking, *Steps 1* and *2* are concerned with exploiting this rigorously. On the other hand we also have the symmetry relation $\Delta_h \Delta_k f = \Delta_k \Delta_h f$, which suggests that $\Delta_h \chi_k \sim \Delta_k \chi_h$; *Step 3* may be thought of in terms of exploiting this kind of information.

3. Almost nilsequences.

In this paper we will be dealing with various objects which are 'almost' nilsequences but not quite. They can invariably be represented as $F(g(n)\Gamma)$ for some function F which is only *piecewise* Lipschitz, the discontinuities being on sets which are somehow 'polynomial'. Rather than formalise these notions, we instead introduce the notion of an approximate nilsequence, give some examples, and point out a number of consequences of the definition.

DEFINITION 3.1 (Almost nilsequences). Suppose that $\Psi : [N] \to \mathbb{C}$ is a 1-bounded function and that $M > 1$ is a complexity parameter. Then we say that Ψ is a *degree s almost polynomial nilsequence* of complexity $O_M(1)$ if, for any $\varepsilon > 0$, there is a genuine

https://doi.org/10.1017/S0017089510000546 Published online by Cambridge University Press

degree s polynomial nilsequence Ψ_ε with complexity $O_{s,\varepsilon,M}(1)$ such that $\mathbb{E}_{n\in[N]}|\Psi(n) - \Psi_\varepsilon(n)| \leqslant \varepsilon$.

REMARK . That is, Ψ can be approximated arbitrarily well, in L^1, by genuine nilsequences. We will not specify the function $O_{s,\varepsilon,M}(1)$ exactly (and indeed it does not make sense to do so, in view of the loose manner in which we have defined complexity). The reader should just imagine that there is *some* fixed function which may be taken in this definition and which makes all statements that we make later on true. Let us also remark that the non-standard analogue of this definition, which will feature in our forthcoming paper on the general case GI(s), is much cleaner and does not involve any unspecified complexity parameters $O_{s,\varepsilon,M}(1)$.

We make the following easily verified, but rather useful, claim:

LEMMA 3.2 (Algebra properties). *If Φ, Ψ are degree s almost polynomial nilsequences, then their sum $\Phi + \Psi$ and product $\Phi\Psi$, and complex conjugate $\overline{\Phi}$ are also degree s almost polynomial nilsequences* (*with a slightly different complexity bound $O_{s,\varepsilon,M}(1)$ on the approximants, of course*).

The utility of Definition 3.1 is made clear by the following lemma, which states that correlation with almost nilsequences is essentially the same thing as correlation with genuine nilsequences.

LEMMA 3.3. *Suppose that $f : [N] \to \mathbb{C}$ is a 1-bounded function and that*

$$|\mathbb{E}_{n\in[N]} f(n)\Psi(n)| \geqslant \delta$$

for some degree s almost polynomial nilsequence Ψ of complexity $O_\delta(1)$. Then there is a genuine degree s polynomial nilsequence $F(g(n)\Gamma)$ of complexity $O_{s,\delta}(1)$ such that $|\mathbb{E}_{n\in[N]} f(n)F(g(n)\Gamma)| \geqslant \delta/2$.

Proof. Simply take $\varepsilon = \delta/2$ in Definition 3.1 and set $F(g(n)\Gamma) = \Psi_\varepsilon(n)$. □

A particular consequence, which we shall make use of later, is that it suffices to establish Conjecture 1.3 with almost nilsequences instead of genuine ones.

For 1-step nilsequences there is a further, very helpful, reduction that can be made.

LEMMA 3.4 (1-step correlation). *Suppose that $f : [N] \to \mathbb{C}$ is a 1-bounded function and that $|\mathbb{E}_{n\in[N]} f(n)\Psi(n)| \geqslant \delta$ for some degree 1 almost nilsequence Ψ of complexity $O_\delta(1)$. Then there is a $\theta \in \mathbb{R}/\mathbb{Z}$ such that $|\mathbb{E}_{n\in[N]} f(n)e(\theta n)| \gg_\delta 1$. (The implied constants here depend of course on the implied constants in the definition of an almost nilsequence.)*

Proof. By the previous lemma we may assume that Ψ is a genuine degree 1 nilsequence of complexity $O_\delta(1)$, that is to say a sequence of the form $(F(n\alpha))_{n\in\mathbb{Z}}$ where $\alpha \in (\mathbb{R}/\mathbb{Z})^k$ for some $k = O_\delta(1)$ and $F : (\mathbb{R}/\mathbb{Z})^k \to \mathbb{C}$ is a function with Lipschitz constant $O_\delta(1)$. Standard Fourier analysis (see, for example, [**12**, Lemma A.9]) implies that we may expand

$$F(t) = \sum_{m\in\mathbb{Z}^k, |m|\leqslant O_\delta(1)} c_m e(m \cdot t) + O(\delta/10),$$

where the c_m are complex numbers with $|c_m| = O_\delta(1)$. The result follows quickly from this. □

https://doi.org/10.1017/S0017089510000546 Published online by Cambridge University Press

The next two lemmas collect together various examples of almost nilsequences. The proofs, which are somewhat technical and tedious, are given in Appendix E.

LEMMA 3.5. *Suppose that $\alpha, \beta \in [0,1]$ and that $M > 1$ is a complexity parameter. The following are all examples of almost nilsequences of degree 1 and complexity $O_M(1)$:*

(i) *the set of 1-step Lipschitz nilsequences of complexity at most M;*
(ii) *the set of characteristic functions 1_P, where $P \subseteq [N]$ is a progression of length at least N/M;*
(iii) *the set of functions of the form $n \mapsto e(\alpha\{\beta n\})$, with $\alpha \in \mathbb{R}$ and $\beta \in \mathbb{R}/\mathbb{Z}$;*
(iv) *the set of functions of the form $n \mapsto e(\{\alpha n\}\{\beta n\})$, with $\alpha, \beta \in \mathbb{R}/\mathbb{Z}$;*
(v) *the set of functions of the form $n \mapsto e(\alpha n \lfloor \beta n \rfloor)$, where $\|\beta\|_{\mathbb{R}/\mathbb{Z}} \leqslant M/N$.*

In particular, (by Lemma 3.4), if $f : [N] \to \mathbb{C}$ is a 1-bounded function such that $|\mathbb{E}_{n\in[N]} f(n)\Psi(n)| \geqslant \delta$, where Ψ is one of the functions on the above list, then there exists $\theta \in \mathbb{R}/\mathbb{Z}$ such that $|\mathbb{E}_{n\in[N]} f(n)e(\theta n)| \gg_{\delta,M} 1$.

LEMMA 3.6. *Suppose that $\alpha, \beta, \gamma \in [0,1]$. Then the following are all examples of almost nilsequences of degree $s \geqslant 2$ and complexity $O(1)$:*

(i) *$n \mapsto e(\{\alpha n\}\beta n)$, of degree 2;*
(ii) *$n \mapsto e(\{\alpha n\}\beta n^2)$, of degree 3;*
(iii) *$n \mapsto e(\{\alpha n\}\{\beta n\}\gamma n)$, of degree 3.*

Although the proof of this last lemma is little tedious, it is also important in the sense that this is the only place in our paper where a 3-step nilsequence is actually constructed.

4. Distributional results concerning nilsequences. We will rely heavily on the quantitative distribution results concerning polynomial nilsequences $(g(n)\Gamma)_{n\in[N]}$ established by the first two authors in [**14**]. There were two main results in that paper, the first of which was used in the proof of the second. We are aware that the paper [**14**] is long and somewhat difficult. However, the reader wishing to understand the present paper need only be *au fait* with the statements of the results there, which means that she need only read Chapters 1 and 2 of the paper. We will assume familiarity with those chapters throughout this paper, and, in particular, will use notation from them without further comment. We will also revisit these results in a non-standard setting in the sequel to this paper, in which we will give more detailed proofs.

The first result we refer to gives a criterion for $(g(n)\Gamma)_{n\in[N]}$ being equidistributed. This is [**14**, Theorem 2.9]. This theorem is a quantitative version of a polynomial equidistribution theorem for nilmanifolds. The qualitative version basically claims that equidistribution of polynomial sequences is determined on the abelianisation $G/[G,G]\Gamma$. For linear sequences this is a classical result, and for polynomial sequences the result is due to Leibman [**21**].

THEOREM 4.1 (Quantitative Leibman dichotomy). *[**14**, Theorem 2.9] Let $m, s \geqslant 0$, $0 < \delta < 1/2$ and $N \geqslant 1$. Suppose that G/Γ is an m-dimensional nilmanifold together with a filtration $G_\bullet$ of degree s and that $\mathcal{X}$ is a $\frac{1}{\delta}$-rational Mal'cev basis adapted to $G_\bullet$. Suppose that $g : \mathbb{Z} \to G$ is a polynomial sequence adapted to $G_\bullet$. If $(g(n)\Gamma)_{n\in[N]}$ is not δ-equidistributed, then there is some $k \in \mathbb{Z}^{m'}$, where $m' := \dim G - \dim[G,G]$ is the*

https://doi.org/10.1017/S0017089510000546 Published online by Cambridge University Press

dimension of the horizontal torus of G/Γ, with $|k| \ll \delta^{-O_{m,s}(1)}$ and

$$\|k \cdot (\pi \circ g)\|_{C^\infty[N]} \ll \delta^{-O_{m,s}(1)}, \tag{4.1}$$

where $\pi : G \to G/[G,G]\Gamma \equiv (\mathbb{R}/\mathbb{Z})^{m'}$ is projection onto the horizontal torus of G/Γ.

The second result we allude to, proved in Sections 9 and 10 of [**14**] by iterating the preceding theorem, is a certain factorisation result. We will need a variant of it in the present paper involving an arbitrary growth function $\omega : \mathbb{R}^+ \to \mathbb{R}^+$; this may be established[3] by exactly the same iterative argument that is used in the proof of [**14**, Theorem 10.2].

THEOREM 4.2 (Factorisation result). *Let $s, N \geqslant 0$ be integers, let $M \geqslant 1$ be a real number, and let $\omega : \mathbb{R}^+ \to \mathbb{R}^+$ be an arbitrary growth function. Suppose that G/Γ is a nilmanifold of complexity at most M together with a filtration $G_\bullet$ of degree s. Suppose that $\mathcal{X}$ is an M-rational Mal'cev basis adapted to $G_\bullet$ and that $g : \mathbb{Z} \to G_\bullet$ is a polynomial map adapted to $G_\bullet$. Then there is an integer M_0 with $M \leqslant M_0 = O_{M,s,\omega}(1)$, a rational subgroup $G' \subseteq G$, a Mal'cev basis $\mathcal{X}'$ for G'/Γ' in which each element is an M_0-rational combination of the elements of $\mathcal{X}$, and a decomposition $g = \varepsilon g' \gamma$ into polynomial sequences $\varepsilon, g', \gamma : \mathbb{Z} \to G$ adapted to $G_\bullet$ with the following properties:*

(i) *$\varepsilon : \mathbb{Z} \to G$ is (M_0, N)-smooth;*
(ii) *$g' : \mathbb{Z} \to G'$ takes values in G', and the finite sequence $(g'(n)\Gamma'')_{n\in[N]}$ is totally $1/\omega(M_0)$-equidistributed in G'/Γ'', whenever Γ'' is a sublattice of Γ' of index at most $\omega(M_0)$, and using the metric $d_{\mathcal{X}}$ on G'/Γ'';*
(iii) *$\gamma : \mathbb{Z} \to G$ is M_0-rational, and $(\gamma(n)\Gamma)_{n\in\mathbb{Z}}$ is periodic with period at most M_0.*

REMARK. The terms 'smooth' and 'totally equidistributed' in this sort of context will not feature elsewhere in the paper, as we shall rely only on this theorem to prove Theorem 4.3 below. An extremely similar deduction was utilised (and proved in some detail) in Section 2 of [**15**].

Sketch proof. The main idea is to iterate Theorem 4.1 using the following 'dimension reduction argument'. At any given stage of the argument, one has an initial factorisation $g = \varepsilon g' \gamma$ obeying all the properties claimed in the theorem for some M_0, except for the equidistribution conclusions on g'. (Note that one can trivially obtain such an initial factorisation by setting ε and γ to be the identity, and $G' = G$.) If g' obeys the stated equidistribution properties, then we are done. Otherwise, by appealing to Theorem 4.1 and refining to a finite sublattice of Γ' if necessary, the horizontal coefficients of $g'(n)$ will contain an approximate linear dependence in the sense of (4.1). One can then use this, following the arguments used to prove [**14**, Theorem 10.2], in order to factorise $g' = \varepsilon' g'' \gamma'$, where ε', γ' satisfy similar properties to ε, γ but with a worse value of M_0, and g'' takes values in a connected subgroup G'' of G' of strictly lower dimension. We then absorb the ε' and γ' factors to ε, γ, replace G' by G'', increase M_0 to a larger quantity depending on M_0 and ω, and continue the argument. Since one cannot have an infinite descent of connected subgroups of G, the argument must eventually terminate with a factorisation with the desired properties. □

[3] We feel rather sorry for our readers at this point. One particular advantage of the non-standard analysis approach to be taken in the more general paper to follow is that the need for arbitrary growth functions ω is eliminated.

https://doi.org/10.1017/S0017089510000546 Published online by Cambridge University Press

The theorem below is a quantitative version of an equidistribution result of Leibman [**21**] stating that the orbit closure of a polynomial sequence $(g(n)\Gamma)_{n\in\mathbb{N}}$ is a finite union of subnilmanifolds Y_j, each a closed orbit of a connected closed subgroup H_j of G; moreover the polynomial sequence visits each Y_j periodically, and is well distributed there with respect to the normalised Haar measure. Much the same argument (with more details) is given in Section 2 of [**15**].

THEOREM 4.3 ('Quantitative Ratner' result). *Let $s, N \geqslant 0$ be integers, let $M \geqslant 1$ be a real number, and let $\omega : \mathbb{R}^+ \to \mathbb{R}^+$ be an arbitrary growth function. Suppose that G/Γ is a nilmanifold of complexity at most M together with a filtration $G_\bullet$ of degree s. Suppose that $\mathcal{X}$ is an M-rational Mal'cev basis adapted to $G_\bullet$ and that $g : \mathbb{Z} \to G_\bullet$ is a polynomial map adapted to $G_\bullet$. Then there is an integer M_0 with $M \leqslant M_0 = O_{M,s,\omega}(1)$ and a decomposition of $[N]$ into subprogressions P_j, each of length at least N/M_0, together with M_0-rational connected subgroups $H_j \leqslant G$ and elements $x_j \in G$ with coordinates at most M_0 such that $(g(n)\Gamma)_{n\in P_j}$ is $1/\omega(M_0)$-equidistributed on $x_jH_j\Gamma/\Gamma$ for each j.*

Sketch proof. In Theorem 4.2, take a growth function $\omega' : \mathbb{R}^+ \to \mathbb{R}^+$ even more rapidly growing than the ω in the statement here. Let $g = \varepsilon g'\gamma$ be the resulting decomposition. Take the progressions P_j to have common difference q, the period of $\gamma(n)\Gamma$, and length sufficiently small that the smooth term $\varepsilon(n)$ is almost constant on each P_j. Choose y_j, γ_j such that $\varepsilon(n) \approx y_j$ and $\gamma(n)\Gamma = \gamma_j\Gamma$ for $n \in P_j$. Then the theorem holds with $H_j := \gamma_j^{-1}G'\gamma_j$ and $x_j := y_j\gamma_j$. Note that the action of conjugation by γ_j moves Γ to a slightly different subgroup of G, but this new group intersects Γ in a subgroup of index $O_{M_0}(1)$, and so one can proceed by using the fact that g' is assumed equidistributed with respect to such subgroups also. □

5. Free nilpotent Lie groups and free nilcharacters. In previous papers in additive combinatorics in which nilsequences have been discussed, such as [**11, 14**], the Heisenberg nilmanifold has been the central example and readers have been encouraged to think of upper triangular matrix groups as the archetypal nilpotent Lie groups. A key innovation in this paper and the sequel [**17**], strongly inspired by the recent work of Leibman on bracket polynomials [**22**], is a shift away from this viewpoint. Instead, it seems that *free* nilpotent Lie groups and certain functions on them play a crucial rôle.

In this section we give some basic definitions in this regard in the 2-step case. In Appendix E we will briefly meet an example of the 3-step case, but for the most part we will be working with 2-step objects in which case it is not a particularly onerous task to proceed very explicitly. The definitions in the higher step case are similar but necessarily require some more general discussion of bases in free nilpotent Lie algebras.

DEFINITION 5.1 (Free 2-step nilpotent Lie group and nilmanifold). By the free 2-step nilpotent Lie group on generators $e_1, \dots, e_k$ we mean

$$G := \left\{e_1^{t_1} \dots e_k^{t_k} e_{[2,1]}^{t_{[2,1]}} \dots e_{[k,k-1]}^{t_{[k,k-1]}} : t_1, \dots, t_k, t_{[2,1]}, \dots, t_{[k,k-1]} \in \mathbb{R}\right\},$$

subject to the relations $e_i^{-1}e_j^{-1}e_ie_j = [e_i, e_j] = e_{[i,j]}$ for $1 \leqslant j < i \leqslant k$. By the *standard filtration* $G_\bullet$ we mean simply the lower central series filtration with $G_{(0)} = G_{(1)} = G$,

https://doi.org/10.1017/S0017089510000546 Published online by Cambridge University Press

$G_{(2)} = [G, G]$ and $G_{(3)} = \{\mathrm{id}\}$. Inside G we take the standard lattice

$$\Gamma := \{e_1^{m_1} \dots e_k^{m_k} e_{[2,1]}^{m_{[2,1]}} \dots e_{[k,k-1]}^{m_{[k,k-1]}} : m_1, \dots, m_k, m_{[2,1]}, \dots, m_{[k,k-1]} \in \mathbb{Z}\}.$$

The quotient G/Γ is then called the free 2-step nilmanifold on k generators.

A Mal'cev basis for G/Γ consists of the elements $X_i = \log e_i$ and $X_{[i',i]} = \log e_{[i',i]}$; the Mal'cev coordinates of an element of G are simply the elements

$$(t_1, \dots, t_k, t_{[2,1]}, \dots, t_{[k,k-1]}).$$

As in [**14**], such a basis may be used to coordinatise G/Γ by identifying $[0, 1]^{k+\binom{k}{2}}$ as a fundamental domain for the right action of Γ on G. Let us perform a calculation. In Mal'cev coordinates it is easy to check that the multiplication law on G corresponds to the operation

$$(t_i, t_{[i',i]}) * (u_i, u_{[i',i]}) = (t_i + u_i, t_{[i',i]} + u_{[i',i]} + t_{i'}u_i).$$

For a given element $g \in G$ with coordinates $(t_i, t_{[i',i]})$ we may pick some $\gamma \in \Gamma$ such that $g\gamma$ has coordinates in the fundamental domain $\mathcal{F} = [0, 1)^{k+\binom{k}{2}} \in \mathbb{R}^{k+\binom{k}{2}}$. Possible coordinates for γ are

$$u_i = -[t_i], u_{[i',i]} = -[t_{[i',i]} - t_{i'}[t_i]],$$

where $[\ ]$ is the floor function. These coordinates are unique if $g\gamma$ lies in the interior of the fundamental domain $\mathcal{F}$. The coordinates of $g\gamma$ are then

$$(\{t_i\}, \{t_{[i',i]} - t_{i'}[t_i]\}).$$

Definition 5.2 (Coordinates). Suppose that G/Γ is the free 2-step nilmanifold on k generators. Suppose that an element $g \in G$ has Mal'cev coordinates $t_i, t_{[i',i]}$. Then the coordinates of $g\Gamma \in G/\Gamma$ are the entries of the vector

$$\big(\{t_i\}, \{t_{[i',i]} - t_{i'}[t_i]\}\big).$$

We write them as $(t_i, t_{[i',i]})$.

Definition 5.3 (Coordinate functions). By the *basic coordinate functions* $F_i, F_{[i',i]} : G/\Gamma \to \mathbb{C}$ we mean the functions $F_i(t) = e(t_i)$ and $F_{[i',i]}(t) = e(t_{[i',i]})$. The *top order* basic coordinate functions $F_{[i',i]}$ will have a particularly important role to play.

Consider now a polynomial sequence $g : \mathbb{Z} \to G$ of the form

$$g(n) := (\xi_1 n, \dots, \xi_k n, q_{[2,1]}(n), \dots, q_{[k,k-1]}(n)), \tag{5.1}$$

where the $q_{[i',i]}$ are quadratic polynomials. By the theory developed towards the end of Section 6 of [**14**] (or simply by a short direct calculation), these are degree two polynomial sequences adapted to the standard filtration $G_\bullet$ based on the lower central series. The objects $F_i(g(n)\Gamma)$, $F_{[i',i]}(g(n)\Gamma)$ are then called *free 2-step nilcharacters*, and they will be basic building blocks in this paper. The top-order nilcharacters involving $F_{[i',i]}$ will play a particularly crucial role. In the light of the above computations these

https://doi.org/10.1017/S0017089510000546 Published online by Cambridge University Press

top-order free 2-step nilcharacters may be computed quite explicitly, and indeed we have

$$F_{[i',i]}(g(n)\Gamma) = e(\xi_i n \lfloor \xi_{i'} n \rfloor) e(\alpha_{[i',i]} n^2 + \beta_{[i',i]} n), \tag{5.2}$$

for some $\alpha_{[i',i]}, \beta_{[i',i]} \in \mathbb{R}$. By altering the quadratics $q_{[i',i]}(n)$ we may make the coefficients $\alpha_{[i',i]}, \beta_{[i',i]}$ arbitrary. These quadratic phases $e(\alpha n^2 + \beta n)$ should be thought of as essentially 1-step objects, albeit of degree 2, and the most important feature of our 2-step nilcharacters are the bracket monomials $\xi_{i'} n \lfloor \xi_i n \rfloor$. We will often use explicit bracket-quadratics in this paper. In the longer paper to come, dealing with the general case, it will not be possible to proceed so explicitly and indeed the main new innovation of that paper (following the work of Leibman) is to develop a kind of 'calculus' of bracket polynomials.

Let us note that $F_{[i',i]}(g(n)\Gamma)$ is not actually a 2-step nilsequence, because the function $F_{[i',i]}$ is only piecewise Lipschitz. From the explicit form given above and Lemma 3.6, however, one sees that it *is* an almost 2-step nilsequence.

We now give a variant of the U^3 inverse theorem involving 2-step free nilcharacters.

THEOREM 5.4 (Inverse theorem for U^3, variant). *Suppose that $f : [N] \to \mathbb{C}$ is a 1-bounded function with $\|f\|_{U^3} \geqslant \delta$. Then we have $|\mathbb{E}_{n \in [N]} f(n)\overline{\chi(n)}| \gg_\delta 1$, where*

$$\chi(n) = e(\alpha n^2 + \beta n) \prod_{1 \leqslant i < i' \leqslant k} F_{[i',i]}^{m_{[i',i]}}(g(n)\Gamma)$$

is the product of some free 2-step nilcharacters with a quadratic phase. Here, $k = O_\delta(1)$ and the $m_{[i',i]}$ are integers bounded by $O_\delta(1)$.

Proof. In [**11**, Theorem 10.9] it is shown that a function f with $\|f\|_{U^3} \geqslant \delta$ has inner product $\gg_\delta 1$ with a function which is the product of $O_\delta(1)$ bracket quadratics $e(\alpha_i n \lfloor \beta_i n \rfloor)$, $i = 1, \ldots, m$ and a quadratic phase $e(\alpha n^2 + \beta n)$. But such a function already has the form χ given in the statement of the theorem, simply by taking $k = 2m$ and horizontal frequencies $\xi_{2i-1} = \beta_i$ and $\xi_{2i} = \alpha_i$, $i = 1, 2, \ldots, m$. □

REMARK. The proof of [**11**, Theorem 10.9] was actually a stepping stone on the way to the proof of the U^3 inverse theorem itself, which requires these 2-step nilcharacters to be assembled into a *Lipschitz* Heisenberg nilsequence.

REMARK . It is possible to proceed directly from the U^3 inverse theorem, that is to say from the formulation given in Conjecture 1.3, although – as the previous remark suggests – it would be a little perverse to do so. To do this requires one to do a slightly odd kind of Fourier decomposition in the coordinate space $(t_i, t_{[i',i]}) = [0, 1]^{k+\binom{k}{2}}$,mapped onto the torus $(\mathbb{R}/\mathbb{Z})^{k+\binom{k}{2}}$, but there is an issue because a function F which is Lipschitz on G/Γ need not even be continuous on this torus. We have a way around this difficulty involving the introduction of a random shift to the fundamental domain $\mathcal{F}$. However we do not believe this argument will be necessary even in the more general paper to come, since our plan is to first prove a variant form of Conjecture 1.3, akin to Theorem 5.4, by induction and only then to deduce Conjecture 1.3 itself.

To conclude this section we give some crucial identities involving bracket quadratics. It is the proper understanding and generalisation of these that we referred to above when we talked about the development of a 'calculus' of bracket polynomials in the forthcoming longer paper.

https://doi.org/10.1017/S0017089510000546 Published online by Cambridge University Press

The key identity we shall rely on is

$$X[Y] = XY - \{X\}\{Y\} - [X]Y + [X][Y], \tag{5.3}$$

valid for all $X, Y \in \mathbb{R}$. This implies that the map $\phi : (X, Y) \mapsto X[Y] (\mathrm{mod}\, 1)$ is 'antisymmetric and bilinear modulo lower order terms'. Specifically, $\phi(X_1 + X_2, Y) - \phi(X_1, Y) - \phi(X_2, Y) = 0$, whilst $\phi(X, Y_1 + Y_2) - \phi(X, Y_1) - \phi(X, Y_2) = \{X\}\{Y_1\} + \{X\}\{Y_2\} - \{X\}\{Y_1 + Y_2\}$, and $\phi(X, Y) = \phi(Y, X) + XY - \{X\}\{Y\}$.

Let us say a clarify to some extent what we mean by 'lower order'. We shall be applying these identities when $X_i = \alpha_i n$ and $Y_j = \beta_j n$, and we shall also be considering $e(\phi)$ rather than ϕ itself. Then these obstructions to antisymmetric bilinearity take the form $e(\{\alpha n\}\{\beta n\})$, an almost 1-step nilsequence (cf. Lemma 3.5 (iv)) and $e(\theta n^2)$, another 1-step object (but of degree 2).

Let us record these observations in the form of a lemma.

LEMMA 5.5 (Bracket quadratic identities). *Suppose that* $\alpha, \alpha_1, \alpha_2, \beta, \beta_1, \beta_2, \gamma \in \mathbb{R}$. *Then*

(i) $e((\alpha_1 + \alpha_2)n\lfloor \beta n\rfloor) = e(\alpha_1 n\lfloor \beta n\rfloor)e(\alpha_2 n\lfloor \beta n\rfloor)$;
(ii) $e(\alpha n\lfloor (\beta_1 + \beta_2)n\rfloor) = e(\alpha n\lfloor \beta_1 n\rfloor)e(\alpha n\lfloor \beta_2 n\rfloor)$ *up to a product of terms of the form* $e(\{\theta n\}\{\theta' n\})$;
(iii) $e(\alpha n\lfloor \beta n\rfloor) = e(-\beta n\lfloor \alpha n\rfloor)$ *up to a product of terms of the form* $e(\theta n^2)$ *and* $e(\{\theta' n\}\{\theta'' n\})$.
(iv) $e(\gamma n\lfloor \gamma n\rfloor)$ *is a product of terms of the form* $e(\theta n^2)$ *and* $e(\{\theta' n\}\{\theta'' n\})$.

Proof. The first three of these follow immediately from (5.3) and the subsequent discussion. Part (iv) perhaps requires some comment: to prove it, first choose γ' so that $2\gamma' = \gamma (\mathrm{mod}\, 1)$. Then take $\alpha = \beta = \gamma'$ in (iii) to obtain the fact that $e(\gamma n\lfloor \gamma' n\rfloor)$ is a product of terms of the required type. Now apply (ii) to conclude the same thing for $e(\gamma n\lfloor \gamma n\rfloor)$. □

6. Some arguments of Gowers. In this section we give the observation of Gowers[7] described in Section 2, whereby one proceeds from the assumption that $|\mathbb{E}_n \Delta_h f(n)\overline{\chi_h(n)}| \geqslant \delta$ for many f to get (2.3), a kind of weak linearity statement concerning the map $h \mapsto \chi_h$. Here is a more precise statement.

PROPOSITION 6.1 (Gowers). *Suppose that* $f : [N] \to \mathbb{C}$ *is a 1-bounded function, that* $H \subseteq [N]$ *is a set with cardinality* ηN *and that for each* $h \in H$ *we have a function* $\chi_h : [N] \to \mathbb{C}$ *with* $|\chi_h(n)| \leqslant 1$ *for all n, and that*

$$|\mathbb{E}_{n\in[N]} \Delta_h f(n)\overline{\chi_h(n)}| \geqslant \delta \tag{6.1}$$

for all $h \in H$. *Then for at least* $\eta^8\delta^4 N^3/2$ *of the quadruples* (h_1, h_2, h_3, h_4) *such that* $h_1 + h_2 = h_3 + h_4$ *we have*

$$|\mathbb{E}_n \chi_{h_1}(n)\chi_{h_2}(n + h_1 - h_4)\overline{\chi_{h_3}(n)\chi_{h_4}(n + h_1 - h_4)}| \geqslant c\eta^4\delta^2.$$

REMARK. In the original paper [7], attention is restricted to the linear case $\chi_h(n) = e(\xi_h n)$, but the argument extends without difficulty to the general case, as we shall see in the proof.

https://doi.org/10.1017/S0017089510000546 Published online by Cambridge University Press

Proof. As in many arguments of analytic number theory and additive combinatorics in which a function that one does not wish to understand is to be eliminated, our main tool is the Cauchy–Schwarz inequality. Two applications of that inequality give that

$$|\mathbb{E}_{n,m} a_n b_m \Phi(n,m)|^4 \leqslant \mathbb{E}_{n,n',m,m'} \Phi(n,m)\overline{\Phi(n',m)\Phi(n,m')}\Phi(n',m') \tag{6.2}$$

whenever $(a_n)_{n\in X}, (b_m)_{m\in Y}, (\Phi(n,m))_{n\in X, m\in Y}$ are 1-bounded sequences of complex numbers.

Returning to the proposition itself, the assumptions imply that

$$\mathbb{E}_h |\mathbb{E}_n \Delta_h f(n)\overline{\chi_h(n)}|^2 \gg \eta\delta^2,$$

where we have taken the expectation over some group $\mathbb{Z}/N'\mathbb{Z}$ with $N' \sim 2N$ (say) and define all functions to be zero outside of $[N]$ and χ_h to be identically zero if $h \notin H$. Expanding out and making some obvious substitutions this yields

$$\mathbb{E}_k \mathbb{E}_{n,m} f(m)\overline{f(m+k)f(n)}f(n+k)\Delta_k \chi_{m-n}(n) \gg \eta\delta^2.$$

Applying Hölder's inequality this means that

$$\mathbb{E}_k |\mathbb{E}_{n,m} f(m)\overline{f(m+k)f(n)}f(n+k)\Delta_k \chi_{m-n}(n)|^4 \gg \eta^4\delta^8.$$

Applying (6.2) for each k, we obtain

$$\mathbb{E}_k \mathbb{E}_{n,n',m,m'} \Delta_k \chi_{m-n}(n)\overline{\Delta_k \chi_{m'-n}(n)\Delta_k \chi_{m-n'}(n')}\Delta_k \chi_{m'-n'}(n') \gg \eta^4\delta^8.$$

This is more suggestively written as

$$\mathbb{E}_{h_1+h_2=h_3+h_4} \mathbb{E}_{n,k} \Delta_k \chi_{h_1}(n)\Delta_k \chi_{h_2}(n+h_1-h_4)\overline{\Delta_k \chi_{h_3}(n)\Delta_k \chi_{h_4}(n+h_1-h_4)} \gg \eta^4\delta^8,$$

which is the same as

$$\mathbb{E}_{h_1+h_2=h_3+h_4} |\mathbb{E}_n \chi_{h_1}(n)\chi_{h_2}(n+h_1-h_4)\overline{\chi_{h_3}(n)\chi_{h_4}(n+h_1-h_4)}|^2 \gg \eta^4\delta^8.$$

This immediately implies the stated result by a trivial averaging argument. □

We now give a corollary of this in the specific case that the $\chi_h(n)$ have the form appearing in the statement of Theorem 5.4, that is to say

$$\chi_h(n) := e(\alpha_h n^2 + \beta_h n) \prod_{1\leqslant i<i'\leqslant k} F_{[i',i]}^{m_{[i',i](h)}}(g_h(n)\Gamma). \tag{6.3}$$

COROLLARY 6.2. *Suppose that $f : [N] \to \mathbb{C}$ is a 1-bounded function and that*

$$|\mathbb{E}_{n\in[N]} \Delta_h f(n)\overline{\chi_h(n)}| \geqslant \delta$$

for all h in some set H, $|H| \geqslant \delta N$. Suppose now that the functions $\chi_h(n)$ have the specific form (6.3). *Then for at least $\delta^C N^3$ additive quadruples $h_1+h_2 = h_3+h_4 \in H^4$ there are frequencies $\alpha_{h_1,h_2,h_3,h_4}, \beta_{h_1,h_2,h_3,h_4} \in \mathbb{R}/\mathbb{Z}$ such that*

$$|\mathbb{E}_{n\in[N]} \chi_{h_1}(n)\chi_{h_2}(n)\overline{\chi_{h_3}(n)\chi_{h_4}(n)}e(\alpha_{h_1,h_2,h_3,h_4}n^2 + \beta_{h_1,h_2,h_3,h_4}n)| \gg_\delta 1.$$

https://doi.org/10.1017/S0017089510000546 Published online by Cambridge University Press

Proof. Apply Proposition 6.1 and then use the bracket identities in Lemma 5.5 to expand out terms such as $\chi_{h_4}(n+h_1-h_4)$. This exhibits

$$\chi_{h_1}(n)\chi_{h_2}(n+h_1-h_4)\overline{\chi_{h_3}(n)\chi_{h_4}(n+h_1-h_4)}$$

as a product of

$$\chi_{h_1}(n)\chi_{h_2}(n)\overline{\chi_{h_3}(n)\chi_{h_4}(n)}$$

times various (possibly h_i-dependent) terms of the form $e(\alpha n^2)$ or $e(\{\alpha n\}\{\beta n\})$. By Lemma 3.5 (iv) the latter are almost 1-step nilsequences; the conclusion then follows from Lemma 3.4. □

REMARK. That this computation worked was no accident. In fact from the general theory in [**14**] one knows that if $\chi(n)$ is a Lipschitz s-step nilsequence with a vertical character then $\chi(n+k)\overline{\chi(n)}$ is an $(s-1)$-step nilsequence. We did not apply this general theory here, since we are being forced to deal with the coordinate functions $F_{[i,i']}$ which are not Lipschitz.

7. Step 1: Reducing the h-dependence. The aim of this rather long and technical section is to handle *Step 1* of the outline in Section 2. Our first task is to formulate properly exactly what it is we intend to do. Recall that if $\|f\|_{U^4} \geqslant \delta$ then, from the fact that $\|\Delta_h f\|_{U^3} \gg \delta^C$ for $\gg \delta^C N$ values of h and Theorem 5.4 we have

$$\mathbb{E}_{n\in[N]}\Delta_h f(n)\overline{\chi_h(n)} \gg_\delta 1 \tag{7.1}$$

where χ_h is an object having the form (6.3), that is to say

$$\chi_h(n) = e(\alpha_h n^2+\beta_h n)\prod_{1\leqslant i<i'\leqslant k} F_{[i',i]}^{m_{[i',i]}(h)}(g_h(n)\Gamma). \tag{7.2}$$

Each term involving an $F_{[i,i']}$ is, by the calculations in Section 5 and in particular those around (5.2), essentially a bracket quadratic $e(\xi n\lfloor\xi' n\rfloor)$ involving ξ,ξ', two of the frequencies in the 'horizontal' part of the polynomial sequence $g_h(n)$.

Let us be a little more precise and write $g_h(n) = (\xi_{h,1}n,\dots,\xi_{h,k}n,\dots)$; thus the numbers $\xi_{h,i}$ are the horizontal frequencies just alluded to. Write $\Xi_h := \{\xi_{h,1},\dots,\xi_{h,k}\}$ for this set. When we outlined *Step 1* earlier on, we did little more than suggest that our aim was to show that no bracket quadratic $e(\xi_{h,i}n\lfloor\xi_{h,i'}n\rfloor)$ involving two genuinely h-dependent frequencies $\xi_{h,i}$ actually occurs in the formula for $\chi_h(n)$.

To attach meaning to this, we will split Ξ_h as a union $\Xi_*\cup\Xi_h'$ of a 'core' set $\Xi_* = \{\xi_{h,1},\dots,\xi_{h,k_*}\}$ and a 'petal' set $\Xi_h' = \{\xi_{h,k_*+1},\dots,\xi_{h,k}\}$ in such a way that the frequencies $\xi_{h,i}$, $i=1,\dots,k_*$, do not actually depend on h. Our task, then, is to show that (7.1) and (7.2) may be achieved in such a way that $m_{[i,i']}(h)=0$ when $i,i'>k_*$. In other words, no bracket quadratic $e(\xi_{h,i}n\lfloor\xi_{h,i'}n\rfloor)$ with $i,i>k_*$ actually occurs in the expression for $\chi_h(n)$.

We will not prove that any situation such as (7.1) and (7.2) has this form automatically. Rather, we will perform an inductive procedure in which the underlying frequency sets are slowly modified so that they take on more and more characteristics of the above 'sunflower' decomposition into core and petals. At the same time, the set of h for which (7.1) holds will be gradually reduced, although it will always have cardinality $\gg_\delta N$.

https://doi.org/10.1017/S0017089510000546 Published online by Cambridge University Press

Here is a precise statement.

PROPOSITION 7.1 (Step 1). *Suppose that $\|f\|_{U^4} \geqslant \delta$. Then for $\gg_\delta N$ values of h we have $\mathbb{E}_{n\in[N]}\Delta_h f(n)\overline{\chi_h(n)} \gg_\delta 1$, where where*

$$\chi_h(n) = e(\alpha_h n^2 + \beta_h n) \prod_{1\leqslant i<i'\leqslant k} F_{[i',i]}^{m_{[i',i]}(h)}(g_h(n)\Gamma)$$

with $k, |m_{[i,i']}(h)| = O_\delta(1)$. Furthermore there is a 'sunflower' decomposition of the frequency sets $\Xi_h = \{\xi_{h,1}, \dots, \xi_{h,k}\}$ of $g_h(n)$ into a 'core' $\Xi_* = \{\xi_{h,1}, \dots, \xi_{h,k_*}\}$ which does not depend on h together with 'petals' $\Xi'_h = \{\xi_{h,k_*+1}, \dots, \xi_{h,k}\}$, in such a way that $m_{[i,i']}(h) = 0$ if $i, i' > k_*$.

The last statement – that is to say the assertion that there are no bracket quadratics with two petal frequencies – is of course the beef here.

Here is a plan of the rest of this section. Proposition 7.1 is proved by a kind of induction (on the 'complexity' of the core-petal decomposition). The inductive step is stated as Proposition 7.5 below, and we give the full derivation of Proposition 7.1 shortly after the proof of that. Proposition 7.5 is itself deduced from Corollary 7.4, which is in turn an easy deduction from Lemma 7.3. This latter result is the main business of this section, and indeed is probably the hardest part of the entire argument. For that reason we will, between stating it and proving it, give a kind of model variant of the argument to illustrate the underlying algebraic structure.

Before we can begin we require a definition which will also feature later in the paper.

DEFINITION 7.2 (Approximate relations and dissociativity). Suppose that

$$\Xi = \{\xi_1, \dots, \xi_k\} \subseteq \mathbb{R}/\mathbb{Z}$$

is a finite set of frequencies. We say that this set *satisfies an M-linear relation* up to ε if there are integers $m_1, \dots, m_k$, $|m_i| \leqslant M$, not all zero, such that $\|m_1\xi_1 + \cdots + m_k\xi_k\|_{\mathbb{R}/\mathbb{Z}} \leqslant \varepsilon$. If a set Ξ satisfies no such linear relation then we say that it is (M, ε)-*dissociated*. We say that a further frequency ξ lies in the M-linear span of Ξ up to ε if there are integers $m_1, \dots, m_k$, $|m_i| \leqslant M$, such that $\|\xi - m_1\xi_1 - \cdots - m_k\xi_k\|_{\mathbb{R}/\mathbb{Z}} \leqslant \varepsilon$.

Let us now state the main lemma of this section. We remark that the hypothesis of this lemma comes from applying Proposition 6.1 to the assumption (7.1). However we shall revisit this point later on when we actually perform the inductive application of the lemma.

LEMMA 7.3. *Fix $h_1, h_2, h_3, h_4 \in H$ and suppose that for $j = 1, \dots, 4$ we have a decomposition $\Xi_{h_j} = \Xi_* \cup \Xi'_{h_j}$ of the frequency set Ξ_{h_j} into a core $\Xi_* = \{\xi_{h_j,1}, \dots, \xi_{h_j,k_*}\}$ not depending on j and a petal set $\Xi'_{h_j} = \{\xi_{h_j,k_*+1}, \dots, \xi_{h_j,k}\}$. Suppose that the functions $\chi_h(n)$ have the form* (6.3) *above, where both k and the indices $m_{[i',i]}(h)$ are bounded by M, and suppose that we have*

$$|\mathbb{E}_{n\in[N]}\chi_{h_1}(n)\chi_{h_2}(n)\chi_{h_3}(n)\chi_{h_4}(n)e(\alpha_{h_1,h_2,h_3,h_4}n^2 + \beta_{h_1,h_2,h_3,h_4}n)| \geqslant 1/M. \qquad (7.3)$$

Suppose that $m_{[i',i]}(h_1) \neq 0$ for some pair $i, i' > k_*$. Then either there is an $O_M(1)$-linear relation, up to $O_M(1/N)$, between the elements in $\Xi_* \cup \Xi'_{h_1} \cup \Xi'_{h_2} \cup \Xi'_{h_3}$, or else there is such a relation between the elements of $\Xi_* \cup \Xi'_{h_1} \cup \Xi'_{h_2} \cup \Xi'_{h_4}$.

https://doi.org/10.1017/S0017089510000546 Published online by Cambridge University Press

Proof. The main idea is to apply the distributional results on nilsequences, and, in particular, the 'Quantitative Ratner' result, Theorem 4.3, to the assumption (7.3). There is a very natural way to do this, which is to write (7.3) as

$$|\mathbb{E}_{n\in[N]}\tilde{F}(\tilde{g}(n)\tilde{\Gamma})| \geqslant 1/M, \tag{7.4}$$

where of course

$$\tilde{F}(\tilde{g}(n)\tilde{\Gamma}) := e(\alpha_{h_1,h_2,h_3,h_4}n^2 + \beta_{h_1,h_2,h_3,h_4}n) \prod_{1\leqslant i<i'\leqslant k} \prod_{j=1,2,3,4} F_{[i',i]}^{m_{[i',i]}(h_j)}(g_{h_j}(n)\Gamma).$$

We may interpret the left-hand side as one big polynomial nilsequence on the 2-step nilmanifold $\tilde{G}/\tilde{\Gamma}$, where $\tilde{G} = G \times G \times G \times G \times \mathbb{R}$ and $\tilde{\Gamma} = \Gamma \times \Gamma \times \Gamma \times \Gamma \times \mathbb{Z}$, and the polynomial sequence $\tilde{g} = \tilde{g}_{h_1,h_2,h_3,h_4}(n)$ is given by

$$\tilde{g}(n) = \tilde{g}_{h_1,h_2,h_3,h_4}(n) = g_{h_1}(n) \times g_{h_2}(n) \times g_{h_3}(n) \times g_{h_4}(n) \times (\alpha_{h_1,h_2,h_3,h_4}n^2 + \beta_{h_1,h_2,h_3,h_4}n).$$

The Quantitative Ratner results are a little complicated, and so before continuing with the proof we sketch how it goes in what might be termed the *asymptotic limit case*, in which we work not with any given scale N, but rather with the limiting behaviour as $N \to \infty$. More precisely, instead of (7.4) we assume merely that[4]

$$\lim_{N\to\infty} \mathbb{E}_{n\in[N]}\tilde{F}(\tilde{g}(n)\tilde{\Gamma}) \neq 0,$$

and instead of finding quantitative relations amongst the frequency sets we merely conclude that the frequencies in either $\Xi_* \cup \Xi_{h_1} \cup \Xi_{h_2} \cup \Xi_{h_3}$ or $\Xi_* \cup \Xi_{h_1} \cup \Xi_{h_2} \cup \Xi_{h_4}$ are rationally dependent. The main difference between the model case and the actual one is that the corresponding nilmanifold distribution results, due to Leibman [**21**], are much cleaner in this setting. For simplicity of notation (in this sketch) let us suppose that $\Xi_* = \emptyset$.

Suppose, then that $\Xi_{h_1} \cup \Xi_{h_2} \cup \Xi_{h_3}$ and $\Xi_{h_1} \cup \Xi_{h_2} \cup \Xi_{h_4}$ are both rationally independent. Consider the orbit $(\tilde{g}(n)\tilde{\Gamma})_{n\in\mathbb{N}}$ in $\tilde{G}$. Roughly speaking[5], the results of [**21**] assert that this orbit is equidistributed on a subnilmanifold of the form $H\tilde{\Gamma}/\tilde{\Gamma}$, where H is a closed connected rational subgroup of $\tilde{G}$.

If $\tilde{F}$ were continuous then this would imply that

$$\lim_{N\to\infty} \mathbb{E}_{n\in[N]}\tilde{F}(\tilde{g}(n)\tilde{\Gamma}) = \int \tilde{F}(x)dm_H(x) \neq 0 \tag{7.5}$$

where m_H is the Haar measure on $H\tilde{\Gamma}/\tilde{\Gamma}$. Unfortunately $\tilde{F}$ is not quite continuous, a further technicality we will have to handle when discussing the proof of Lemma 7.3 proper. For the purposes of this sketch, however, let us assume that (7.5) holds.

Let $\pi_1, \pi_2, \pi_3, \pi_4 : \tilde{G} \to G$ be the projections of $\tilde{G}$ onto each of the four factors of G comprising $\tilde{G}$, and by abuse of notation use the same notation for the projection maps from $\tilde{G}/\tilde{\Gamma}$ to the factors G/Γ. Now the projection $(\pi_1 \times \pi_2 \times \pi_3)(\tilde{g}(n))$ has, as its set of horizontal frequencies, $\Xi_{h_1} \cup \Xi_{h_2} \cup \Xi_{h_3}$, a set which is rationally independent.

[4]The convergence of all limits involving polynomial nilsequences was established in [**21**], at least in the case when $\tilde{F}$ is continuous.

[5]In actual fact this is only true after subdividing $\mathbb{N}$ into finitely many subprogressions, and furthermore we need to work with a translate $x_0H\tilde{\Gamma}/\tilde{\Gamma}$. Both of these points are merely technical. The finitary analogue of this result is, of course, the Quantitative Ratner Theorem, Theorem 4.3.

https://doi.org/10.1017/S0017089510000546 Published online by Cambridge University Press

But these frequencies are precisely those occurring in the projection of $(\pi_1 \times \pi_2 \times \pi_3)(\tilde{g}(n)\tilde{\Gamma})$ onto the horizontal torus (abelianisation) of $G/\Gamma \times G/\Gamma \times G/\Gamma$, and hence the orbit of this (abelian) nilsequence is dense. However Leibman's criterion[6] asserts that a polynomial nilsequence is dense if and only if its abelianisation is, and so $((\pi_1 \times \pi_2 \times \pi_3)(\tilde{g}(n)\tilde{\Gamma}))_{n\in\mathbb{N}}$ is dense in $G/\Gamma \times G/\Gamma \times G/\Gamma$.

Since $(\tilde{g}(n)\tilde{\Gamma})_{n\in\mathbb{N}}$ equidistributes in $H\tilde{\Gamma}/\tilde{\Gamma}$, we must have

$$(\pi_1 \times \pi_2 \times \pi_3)(H\tilde{\Gamma}/\tilde{\Gamma}) = G/\Gamma \times G/\Gamma \times G/\Gamma.$$

Topological arguments[7] using the fact that H is closed and connected let us lift this statement to G to conclude that

$$(\pi_1 \times \pi_2 \times \pi_3)(H) = G \times G \times G. \tag{7.6}$$

By exactly the same argument we have

$$(\pi_1 \times \pi_2 \times \pi_4)(H) = G \times G \times G. \tag{7.7}$$

We claim that as a consequence of these observations we have

$$[G, G] \times \mathrm{id} \times \mathrm{id} \times \mathrm{id} \times \mathrm{id} \subseteq H.$$

To see this, let $g, g' \in G$ be arbitrary. Then (7.6) implies that H contains an element of the form $(g, \mathrm{id}, \mathrm{id}, x, z)$, for some $x \in G$ and some $z \in \mathbb{R}$, whilst (7.7) implies that H contains an element of the form $(\mathrm{id}, g', x', \mathrm{id}, z')$ for some $x' \in G$ and some $z' \in \mathbb{R}$. The commutator of these two elements is $([g, g'], \mathrm{id}, \mathrm{id}, \mathrm{id}, \mathrm{id})$, thereby establishing the claim.

REMARK . This idea has appeared in related contexts before, for example in the work of Furstenberg and Weiss [**6**], as well as in less related contexts such as a paper of Hrushovski [**19**, Lemma 4.11].

As a special case of the above claim, we see that for each pair i, i' with $1 \leqslant i < i' \leqslant k$ and for each $t \in \mathbb{R}$ the element $z := (e^t_{[i,i']}, \mathrm{id}, \mathrm{id}, \mathrm{id}, \mathrm{id})$ lies in H. It follows that

$$\int \tilde{F}(x) dm_H(x) = \int \tilde{F}(zx) dm_H(x).$$

However a direct calculation using the definition of $\tilde{F}$ confirms that

$$\tilde{F}(zx) = e(tm_{[i,i']}(h_1))\tilde{F}(x).$$

Since t is arbitrary, the only way to reconcile this with (7.5) is to conclude that $m_{[i,i']}(h_1) = 0$. Thus in this case (in which there is no core Ξ_*) we see that either the functions $\chi_h(n)$ are somewhat trivial in the sense that all of the $m_{[i,i']}(h)$ vanish, or else we were wrong to assume that the frequencies in both $\Xi_{h_1} \cup \Xi_{h_2} \cup \Xi_{h_3}$ and $\Xi_{h_1} \cup \Xi_{h_2} \cup \Xi_{h_4}$ are rationally independent.

This concludes our sketch of the asymptotic limit case, and we now return to our original task of proving Lemma 7.3. The underlying idea is the same as in the above

[6]The finitary analogue of this is the Quantitative Leibman dichotomy, Theorem 4.1.
[7]In the finitary world these are somewhat painful and involve, for example, some quantitative linear algebra; see Appendix A.

https://doi.org/10.1017/S0017089510000546 Published online by Cambridge University Press

sketch except that everything must be made quantitative, without any recourse to limits. Furthermore there was one point in the above sketch where we treated a special case (the core set Ξ_* is empty) and others where we waived our hands somewhat (the function $\tilde{F}$ is not Lipschitz, the orbit only equidistributes on a *coset* of a nilmanifold, and then only after passing to a subprogression). These issues must, of course, be dealt with properly.

Consider the orbit $(\tilde{g}(n)\tilde{\Gamma})_{n\in[N]}$. Let $\omega : \mathbb{R}^+ \to \mathbb{R}^+$ be a growth function to be specified later. By Theorem 4.3 there is some $M_0 = O_{\omega,M}(1)$ (which we may clearly assume to be at least $\max(M, \#\tilde{G}/\tilde{\Gamma})$, since both of these quantities are $O_M(1)$) with the following property. We may partition $[N]$ into subprogressions P_j with lengths at least N/M_0, such that corresponding to each progression P_j the uniform measure

$$\mu_j := \frac{1}{|P_j|}\sum_{n\in P_j} 1_{\tilde{g}(n)\tilde{\Gamma}}$$

is $1/\omega(M_0)$-close to the Haar measure m_{H_j} on $b_j H_j\tilde{\Gamma}/\tilde{\Gamma}$, where $H_j \leqslant \tilde{G}$ is some closed, connected, M_0-rational subgroup. Namely for any Lipschitz function F on $\tilde{G}/\tilde{\Gamma}$ we have

$$|\mathbb{E}_{n\in P}F(\tilde{g}(n)\tilde{\Gamma}) - \int F dm_{H_j}| \leqslant \frac{1}{\omega(M_0)}\|F\|_{\mathrm{Lip}} \tag{7.8}$$

By a trivial averaging argument, condition (7.4) implies that there is some $P = P_j$ such that

$$|\mathbb{E}_{n\in P}\tilde{F}(\tilde{g}(n)\tilde{\Gamma})| \geqslant 1/M; \tag{7.9}$$

let $H = H_j$ be the corresponding group, and m_H the Haar measure on $bH\tilde{\Gamma}/\tilde{\Gamma}$.

Let $z \in [H, H]$ be an element, all of whose coordinates are bounded by $O_{M_0}(1)$, and let $F : \tilde{G}/\tilde{\Gamma} \to \mathbb{C}$ be a Lipschitz function. Then $F_z(x\tilde{\Gamma}) = F(zx\tilde{\Gamma})$ is also Lipschitz and $\|F_z\|_{\mathrm{Lip}} = O_{M_0}(1)\|F\|_{\mathrm{Lip}}$. Furthermore since m_H is invariant under translation by z (which lies in the centre of G) we have

$$\int F_z dm_H = \int F dm_H,$$

and thus from (7.8) we get

$$|\mathbb{E}_{n\in P}F_z(\tilde{g}(n)\tilde{\Gamma}) - \int F dm_H| = O_{M_0}(1/\omega(M_0))\|F\|_{\mathrm{Lip}}.$$

And by the triangle inequality

$$|\mathbb{E}_{n\in P}F_z(\tilde{g}(n)\tilde{\Gamma}) - \mathbb{E}_{n\in P}F(\tilde{g}(n)\tilde{\Gamma})| = O_{M_0}(1/\omega(M_0))\|F\|_{\mathrm{Lip}}, \tag{7.10}$$

thus if ω is sufficiently rapidly-growing then the error term here is negligible and thus

$$\mathbb{E}_{n\in P}F(z\tilde{g}(n)\tilde{\Gamma}) \approx \mathbb{E}_{n\in P}F(\tilde{g}(n)\tilde{\Gamma}). \tag{7.11}$$

Let ϵ_{M_0} be the quantity from the lifting Proposition A.4. Namely any element of x of $H[\tilde{G}, \tilde{G}]\tilde{\Gamma}/[\tilde{G}, \tilde{G}]\tilde{\Gamma} = \mathbb{R}^m/\mathbb{Z}^m$ whose coordinates are bounded by ϵ_{M_0} has a lift under

https://doi.org/10.1017/S0017089510000546 Published online by Cambridge University Press

the natural projection $\tilde{G} \to \tilde{G}/[\tilde{G}, \tilde{G}]\tilde{\Gamma}$ to an element in H with coordinates $O_{M_0}(1)$, whose first m coordinates are the reduced coordinates of x.

We now deal with the issue of $\tilde{F}$ not being Lipschitz. Fix $\delta_0 = \frac{1}{10}\epsilon_{M_0}$. We first need to modify the function $\tilde{F}$. We will choose two parameters δ_1, δ_2, such that δ_1 is much smaller that δ_0, and δ_2 still smaller depending on δ_1. However, both these quantities will be $\gg_{M,w} 1$. Consider the distribution of some fixed coordinate $t_{h_j,[i',i]}$ of $\tilde{g}(n)\tilde{\Gamma}$ as n varies over P. We may clearly suppose that there is no $O_{M,\omega}(1)$-linear relation, up to $O_{M,\omega}(1/N)$, amongst the frequencies $\Xi_* \cup \Xi_{h_j}$ since otherwise the conclusion of the lemma is trivially satisfied. If there is no such relation, and if the implicit constants in the $O_{M,\omega}(1)$ notation above are chosen sufficiently large, then by the quantitative Leibman dichotomy, Theorem 4.1, the sequence $(g_{h_j}(n)\Gamma)_{n\in P}$ is δ_2-equidistributed in G/Γ. Fix a $j \in \{1, 2, 3, 4\}$ and a pair i, i' with $1 \leqslant i < i' \leqslant k$. Let $\psi = \psi_{j,i,i'} : G/\Gamma \to [0, 1]$ be supported where $t_{h_j,[i',i]} \leqslant 2\delta_1$ or $t_{h_j,[i',i]} \geqslant 1 - 2\delta_1$ and be equal to 1 whenever $t_{h_j,[i',i]} \leqslant \delta_1$ or $t_{h_j,[i',i]} \geqslant 1 - \delta_1$ and have $\|\psi\|_{\mathrm{Lip}} = O_{M,\delta_1}(1)$. Let $\tilde{\psi} = \tilde{\psi}_{j,i,i'} : \tilde{G}/\tilde{\Gamma} \to [0, 1]$ be the pullback of $\psi_{j,i,i'}$ under the natural projection from $\tilde{G}$ to the jth copy of G.

Our preceding observation about the distribution of $(g_{h_j}(n)\Gamma)_{n\in P}$ implies that

$$\begin{aligned}|\mathbb{E}_{n\in P}\tilde{\psi}(\tilde{g}(n)\tilde{\Gamma})| = |\mathbb{E}_{n\in P}\psi(g_{h_j}(n)\Gamma)| &\leqslant \int_{G/\Gamma} \psi\, dm_{G/\Gamma} + \delta_2\|\psi\|_{\mathrm{Lip}} \\ &= o_{M;\delta_1\to 0}(1) + O_{M,\delta_1}(\delta_2),\end{aligned}$$

where $o_{M;\delta_1\to 0}(1)$ denotes a quantity that is bounded in magnitude by $c_M(\delta_1)$ for some $c_M(\delta_1)$ that goes to zero as $\delta_1 \to 0$ for any fixed M.

Let z be an element in $[H, H]$ with $O_{M_0}(1)$-bounded coordinates, and suppose $z = (z_{h_1}, z_{h_2}, z_{h_3}, z_{h_4}, w)$ under the decomposition of $\tilde{G}$ as $G \times G \times G \times G \times \mathbb{R}$. Then $\psi_{z_{h_j}}(g_{h_j}(n)\Gamma) = \psi(z_{h_j}g_{h_j}(n)\Gamma)$ is Lipschitz with $\|\psi_{z_{h_j}}\|_{\mathrm{Lip}} = O_{M_0}(1)$, and by the invariance of $m_{G/\Gamma}$ under multiplication by z_{h_j} we get

$$|\mathbb{E}_{n\in P}\tilde{\psi}(z\tilde{g}(n)\tilde{\Gamma})| = o_{M_0;\delta_1\to 0}(1) + O_{M_0,\delta_1}(\delta_2).$$

By adding, the same type of bounds hold for the function $\Psi := \sum_{j=1}^{4}\sum_{1\leqslant i<i'\leqslant k}\tilde{\psi}_{j,i,i'}$, that is to say

$$|\mathbb{E}_{n\in P}\Psi(\tilde{g}(n)\tilde{\Gamma})|, |\mathbb{E}_{n\in P}\Psi(z\tilde{g}(n)\tilde{\Gamma})| = o_{M_0;\delta_1\to 0}(1) + O_{M_0,\delta_1}(\delta_2).$$

Let us, at this point, fix $\delta_1 \gg_{M_0} 1$ in such a way that the $o_{M_0;\delta_1\to 0}(1)$ term here is bounded by δ_0^{10} (say), and let us then choose $\delta_2 \gg_{M_0} 1$ in such a way that the $O_{M_0,\delta_1}(\delta_2)$ term is also bounded by δ_0^{10}. Then the last displayed equation becomes

$$|\mathbb{E}_{n\in P}\Psi(\tilde{g}(n)\tilde{\Gamma})|, |\mathbb{E}_{n\in P}\Psi(z\tilde{g}(n)\tilde{\Gamma})| = O(\delta_0^{10}). \tag{7.12}$$

Note that by construction Ψ is equal to 1 in a δ_1-neighbourhood of all of the discontinuities of our function F. As a result of this it is clear that we may find a function $\tilde{F}_0$ with the property that

$$\|\tilde{F}_0\|_{\mathrm{Lip}} = O_{M_0}(1)$$

whilst

$$|\tilde{F} - \tilde{F}_0| \leqslant \Psi$$

https://doi.org/10.1017/S0017089510000546 Published online by Cambridge University Press

pointwise. By (7.10) we have

$$|\mathbb{E}_{n\in P}\tilde{F}_0(z\tilde{g}(n)\tilde{\Gamma}) - \mathbb{E}_{n\in P}\tilde{F}_0(\tilde{g}(n)\Gamma)| = O_{M_0}(1/\omega(M_0))$$

and by (7.12) we have

$$|\mathbb{E}_{n\in P}(\tilde{F}-\tilde{F}_0)(z\tilde{g}(n)\tilde{\Gamma})|, |\mathbb{E}_{n\in P}(\tilde{F}-\tilde{F}_0)(\tilde{g}(n)\tilde{\Gamma})| = O(\delta_0^{10}).$$

Adding, we obtain

$$|\mathbb{E}_{n\in P}\tilde{F}(z\tilde{g}(n)\tilde{\Gamma}) - \mathbb{E}_{n\in P}\tilde{F}(\tilde{g}(n)\tilde{\Gamma})| = O_{M_0}(1/\omega(M_0)) + O(\delta_0^{10}).$$

Recall that $\delta_0 = \frac{1}{10}\varepsilon_{M_0}$ depends only on M_0. By choosing $\omega : \mathbb{R}^+ \to \mathbb{R}^+$ to be sufficiently rapidly-growing, the whole of the right-hand side can therefore be made $O(\delta_0^{10})$, that is to say

$$\mathbb{E}_{n\in P}\tilde{F}(z\tilde{g}(n)\tilde{\Gamma}) = \mathbb{E}_{n\in P}\tilde{F}(\tilde{g}(n)\tilde{\Gamma}) + O(\delta_0^{10}). \tag{7.13}$$

Now that ω has been fixed, we have $M_0 = O_M(1)$ and $\delta_0 \gg_M 1$. As stated before, our aim now is to assume that $\Xi_* \cup \Xi_{h_1} \cup \Xi_{h_2} \cup \Xi_{h_3}$ and $\Xi_* \cup \Xi_{h_1} \cup \Xi_{h_2} \cup \Xi_{h_4}$ are highly dissociated and use this to produce an element $z \in [H, H]$ which, in conjunction with (7.4), contradicts (7.13). We shall require a further parameter $\delta_3 \gg_M 1$, much smaller than δ_0. We will specify it later on.

Let $\pi_1, \pi_2, \pi_3, \pi_4 : \tilde{G} \to G$ be the projections from $\tilde{G}$ onto the four copies of G (recall, of course, that $\tilde{G} = G \times G \times G \times G \times \mathbb{R}$). Once again we abuse notation and use the same notation for the corresponding projections from $\tilde{G}/\tilde{\Gamma}$ to G/Γ. Suppose that $\Xi_* \cup \Xi_{h_1} \cup \Xi_{h_2} \cup \Xi_{h_3}$ is $O_M(1)$-dissociated up to $O_M(1/N)$. Let us examine the abelian part of $(\pi_1 \times \pi_2 \times \pi_3)(\tilde{g}(n)\Gamma)_{n\in P}$, that is to say the image of $(\tilde{g}(n)\tilde{\Gamma})_{n\in P}$ under the projection

$$\pi_{123}^{\mathrm{ab}} : \tilde{G}/\tilde{\Gamma} \to G/[G,G]\Gamma \times G/[G,G]\Gamma \times G/[G,G]\Gamma \cong (\mathbb{R}/\mathbb{Z})^k \times (\mathbb{R}/\mathbb{Z})^k \times (\mathbb{R}/\mathbb{Z})^k.$$

This image takes the form

$$\big((\xi_{h_1,i}n)_{i=1}^{k_*}, (\xi_{h_1,i}n)_{i=k_*+1}^{k}, (\xi_{h_2,i}n)_{i=1}^{k_*}, (\xi_{h_2,i}n)_{i=k_*+1}^{k}, (\xi_{h_3,i}n)_{i=1}^{k_*}, (\xi_{h_3,i}n)_{i=k_*+1}^{k}\big) (\mathrm{mod}\ 1)$$

Recalling that $\Xi_* = \{\xi_{h,1}, \dots, \xi_{h,k_*}\}$ and that $\Xi_h = \{\xi_{h,k_*+1}, \dots, \xi_{h,k}\}$, it follows from the asserted dissociativity (assuming the implicit $O_M(1)$ terms are large enough) and Kronecker's theorem in quantitative form (cf. Lemma D.2), that this image is δ_3-equidistributed in the subtorus

$$\{(t, u_1, t, u_2, t, u_3) : t \in (\mathbb{R}/\mathbb{Z})^{k_*}, u_1, u_2, u_3 \in (\mathbb{R}/\mathbb{Z})^{k-k_*}\} \subseteq (\mathbb{R}/\mathbb{Z})^k \times (\mathbb{R}/\mathbb{Z})^k \times (\mathbb{R}/\mathbb{Z})^k.$$

In particular, there is an element in $(\pi_{123}^{\mathrm{ab}}(\tilde{g}(n)\tilde{\Gamma}))_{n\in P}$ within $O(\delta_3)$ of

$$(0, (0, \dots, \delta_0, \dots, 0), 0, 0, 0, 0) + \pi_{123}^{\mathrm{ab}}(b),$$

where the δ_0 lies in the ith position (note that $i > k_*$ by assumption). Now since the uniform probability measure on $(\tilde{g}(n)\tilde{\Gamma})_{n\in P}$ is $O_{M_0}(1/\omega(M_0)) = \delta_0^{10}$-close to the Haar measure on $bH\tilde{\Gamma}/\tilde{\Gamma}$, the projection $(\pi_{123}^{\mathrm{ab}}(\tilde{g}(n)\tilde{\Gamma}))_{n\in P}$ is δ_0^{10}-equidistributed in $\pi_{123}^{\mathrm{ab}}(bH\tilde{\Gamma}/\tilde{\Gamma})$. This means that there is an element x of $\pi_{123}^{\mathrm{ab}}(H\tilde{\Gamma}/\tilde{\Gamma})$ within $O(\delta_3)$ of

$$(0, (0, \dots, \delta_0, \dots, 0), 0, 0, 0, 0).$$

https://doi.org/10.1017/S0017089510000546 Published online by Cambridge University Press

Recall that we chose δ_0 so that the lifting property A.4 holds. Recalling the relationship between distance in coordinates and distance in $\tilde{G}$ (cf. [**14**, Lemma A.4]) we can thus find an element in $(\pi_1 \times \pi_2 \times \pi_3)(H)$ at distance $O_M(\delta_3)$ from $e_i^{\delta_0} z_1 \times z_2 \times z_3$ where $z_1, z_2, z_3 \in [G, G]$ are arbitrary (with coordinates bounded by $O_{M_0}(1)$). It follows that we can find an $g \in H$ with

$$d_{\tilde{G}}(g, e_i^{\delta_0} z_1 \times z_2 \times z_3 \times w_4 \times u) = O_M(\delta_3)$$

where $w_4 \in G$ and $u \in \mathbb{R}$ are arbitrary.

Similarly, if $\Xi_* \cup \Xi_{h_1} \cup \Xi_{h_2} \cup \Xi_{h_4}$ is $O_M(1)$-dissociated up to $O_M(1/N)$ then we may locate inside H an element g' with

$$d_{\tilde{G}}(g', e_{i'}^{\delta_0} z_1' \times z_2' \times w_3' \times z_4' \times u') = O_M(\delta_3),$$

where $z_1', z_2', z_4' \in [G, G]$, and $w_3' \in G$, $u' \in \mathbb{R}$ are arbitrary.

We then take for our element $z \in [H, H]$ the commutator $[g, g']$. Noting that[8]

$$\begin{aligned}&[e_{i'}^{\delta_0} z_1 \times z_2 \times z_3 \times w_4 \times u, e_i^{\delta_0} z_1' \times z_2' \times w_3' \times z_4' \times u'] \\ &\quad = [e_{i'}^{\delta_0}, e_i^{\delta_0}] = [e_{i'}, e_i]^{\delta_0^2} \times \mathrm{id} \times \mathrm{id} \times \mathrm{id} \times \mathrm{id}\end{aligned}$$

and that the maps $g \mapsto [g, g_0]$ are uniformly Lipschitz for g_0 in any bounded set, we have

$$d_{\tilde{G}}(z, [e_{i'}, e_i]^{\delta_0^2} \times \mathrm{id} \times \mathrm{id} \times \mathrm{id} \times \mathrm{id}) = O_{M_0}(\delta_3). \tag{7.14}$$

Now, as we have remarked, the coordinate functions $F_{[l',l]} : G/\Gamma \to \mathbb{C}$ are not Lipschitz. However, they *are* $O_M(1)$-Lipschitz when restricted to $[G, G]\Gamma/\Gamma$, as an easy computation confirms. It follows from this observation, (7.14) and the definition of the functions $F_{[l',l]}$ that for $j = 1, 2, 3, 4$ and for any $x \in G/\Gamma$ we have

$$F_{[l',l]}(zx)^{m_{[l',l]}(h_j)} = F_{[l',l]}(x)^{m_{[l',l]}(h_i)} + O_{M_0}(\delta_3)$$

unless $l = i$, $l' = i'$ and $j = 1$ in which case

$$F_{[i',i]}(zx)^{m_{[i',i]}(h_1)} = e(\delta_0^2 m_{[i',i]}(h_1)) F_{[i',i]}(x)^{m_{[i',i]}(h_1)} + O_{M_0}(\delta_3).$$

Taking products over all choices of i, i', it follows that

$$\tilde{F}(zx) = e(\delta_0^2 m_{[i',i]}(h_1))\tilde{F}(x) + O_{M_0}(\delta_3),$$

from which it of course follows that

$$\mathbb{E}_{n \in P} \tilde{F}(z\tilde{g}(n)\tilde{\Gamma}) = e(\delta_0^2 m_{[i',i]}(h_1)) \mathbb{E}_{n \in P} \tilde{F}(\tilde{g}(n)\Gamma) + O_{M_0}(\delta_3).$$

Choosing δ_3 so small that the error term here is $O(\delta_0^{10})$, we obtain upon comparison with (7.13) that

$$|1 - e(\delta_0^2 m_{[i',i]}(h_1))| |\mathbb{E}_{n \in P} \tilde{F}(\tilde{g}(n)\Gamma)| = O(\delta_0^{10}).$$

[8] Here we have used the fact, specific to the 2-step case, that $[x^t, y^{t'}] = [x, y]^{tt'}$. One way to check this would be to verify it for $t, t' \in \mathbb{Z}$ and use the fact that both sides are polynomials in a suitable coordinate system. In the higher step case, the more general *Baker-Campbell-Hausdorff formula* could be used instead.

https://doi.org/10.1017/S0017089510000546 Published online by Cambridge University Press

Recalling that $m_{[i',i]}(h_1)$ is an integer bounded in magnitude by M, that

$$|\mathbb{E}_{n\in P}\tilde{F}(\tilde{g}(n)\tilde{\Gamma})| \geqslant 1/M,$$

and that δ_0 may certainly be assumed to be much smaller than $1/M$, we are forced to conclude (at last!) that $m_{[i',i]}(h_1) = 0$. □

We may put Lemma 7.3 together with Corollary 6.2 in a straightforward manner.

COROLLARY 7.4. *Suppose that* $|\mathbb{E}_{n\in[N]}\Delta_h f(n)\overline{\chi_h(n)}| \geqslant 1/M$ *for all* $h \in H$, *where* $H \subseteq [N]$, $|H| \geqslant N/M$ *and* $\chi_h(n)$ *has the form* (6.3) *with complexity at most* M, *and with decompositions of the frequency sets* $\Xi_h = \{\xi_{h,1}, \dots, \xi_{h,k}\}$ *into cores* $\Xi_* = \{\xi_{h,1}, \dots, \xi_{h,k_*}\}$ *which do not depend on* h *and petal sets* $\Xi'_h = \{\xi_{h,k_*+1}, \dots, \xi_{h,k}\}$. *Then one of the following two alternatives holds true:*

(i) *There is a set* $H' \subseteq H$, $|H'| \gg_M |H|$, *such that* $m_{[i,i']}(h) = 0$ *whenever* $i, i' > k_*$ *and* $h \in H'$;

(ii) *For* $\gg_M N^3$ *triples* $h, h', h'' \in H^3$ *the set* $\Xi_* \cup \Xi'_h \cup \Xi'_{h'} \cup \Xi'_{h''}$ *fails to be* $O_M(1)$*-dissociated up to* $O_M(1/N)$.

Proof. By Corollary 6.2 there are $\gg_M N^3$ additive quadruples $h_1 + h_2 = h_3 + h_4$ such that there are α_{h_1,h_2,h_3,h_4}, β_{h_1,h_2,h_3,h_4} for which

$$|\mathbb{E}_{n\in[N]}\chi_{h_1}(n)\chi_{h_2}(n)\chi_{h_3}(n)\chi_{h_4}(n)e(\alpha_{h_1,h_2,h_3,h_4}n^2 + \beta_{h_1,h_2,h_3,h_4}n)| \gg_M 1.$$

By pigeonhole there must either be $\gg_M N^3$ of these quadruples such that $m_{[i',i]}(h_1) = 0$ for all $i, i' > k_*$, in which case we are clearly in alternative (i), or else there must be some choice of $i, i' > k_*$ such that there are $\gg_M N^3$ quadruples with $m_{[i',i]}(h_1) \neq 0$. By Lemma 7.3 it follows that for each of these quadruples at least one of the sets $\Xi_* \cup \Xi'_{h_1} \cup \Xi'_{h_2} \cup \Xi'_{h_3}$ or $\Xi_* \cup \Xi'_{h_1} \cup \Xi'_{h_2} \cup \Xi'_{h_4}$ fails to be $O_M(1)$-dissociated up to $O_M(1/N)$. It follows immediately that we are in case (ii). □

Now if alternative (i) holds in this last corollary then *Step 1* is complete (that is, Proposition 7.1 is proven). If alternative (ii) holds, then it is possible to replace the core-petal decomposition $\Xi_h = \Xi_* \cup \Xi'_h$ by one in which some of the petal behaviour is absorbed into the core. The precise statement of this, which follows now, is slightly long:

PROPOSITION 7.5. *Let* $H \subseteq [N]$ *be a set with* $|H| \geqslant N/M$. *Suppose that*

$$|\mathbb{E}_{n\in[N]}\Delta_h f(n)\overline{\chi_h(n)}| \geqslant 1/M$$

for all $h \in H$, *where* $|H| \geqslant N/M$ *and the nilcharacter* $\chi_h(n)$ *has the form* (6.3) *with complexity at most* M *and there is a decomposition of the underlying frequency set* $\Xi_h = \{\xi_{h,1}, \dots, \xi_{h,k}\}$ *into*

- *a core component* $\Xi_* = \{\xi_{h,1}, \dots, \xi_{h,k_*}\}$ *which does not depend on* h *and*
- *a petal component* $\Xi'_h = \{\xi_{h,k_*+1}, \dots, \xi_{h,k}\}$.

Then either

- *there is a set* $H' \subseteq H$, $|H'| \gg_M |H|$, *such that* $m_{[i',i]}(h) = 0$ *for all* $i, i' > k_*$ *and for all* $h \in H'$, *or*
- *there is a set* $\tilde{H} \subseteq H$, $|\tilde{H}| \gg_M |H|$ *and nilcharacters* $\tilde{\chi}_h(n)$ *of complexity* $O_M(1)$, $h \in H''$, *such that*

$$\mathbb{E}_{n\in[N]}\Delta_h f(n)\overline{\tilde{\chi}_h(n)} \gg_M 1 \tag{7.15}$$

https://doi.org/10.1017/S0017089510000546 Published online by Cambridge University Press

for all $h \in \tilde{H}$. Here the nilcharacters $\tilde{\chi}_h(n)$ have the form

$$\tilde{\chi}_h(n) = e(\alpha_h n^2 + \beta_h n) \prod_{1 \leqslant i < i' \leqslant \tilde{k}} F_{[i',i]}^{\tilde{m}_{[i',i]}(h)}(\tilde{g}_h(n)\tilde{\Gamma}),$$

where $\tilde{g}_h(n) = (\tilde{\xi}_{h,1}, \dots, \tilde{\xi}_{h,\tilde{k}}, 0, \dots, 0)$. Furthermore writing

$$\tilde{\Xi}_h := \{\tilde{\xi}_{h,1}, \dots, \tilde{\xi}_{h,\tilde{k}}\}$$

we have a decomposition $\tilde{\Xi}_h = \tilde{\Xi}_* \cup \tilde{\Xi}'_h$, where either

(i) (core decreases) $|\tilde{\Xi}_*| < |\Xi_*|$ *and* $|\tilde{\Xi}'_h| = |\Xi'_h|$ *or*
(ii) (petals decrease) $|\tilde{\Xi}_*| \leqslant |\Xi_*| + 1$ *and* $|\tilde{\Xi}'_h| < |\Xi'_h|$.

Proof of Proposition 7.1, a.k.a. Step 1. Before embarking on the proof of this last proposition, we remark how a simple iteration of it leads to Proposition 7.1. One starts with the trivial decomposition $\Xi_h = \Xi_* \cup \Xi'_h$ where $\Xi_* = \emptyset$ and $\Xi'_h = \Xi_h$, and with the initial value of M being $O_\delta(1)$. It is not hard to see that there cannot be more than $O_M(1)$ iterations of alternatives (i) (core decreases) or (ii) (petals decrease). □

Proof of Proposition 7.5. By Corollary 7.4 we may assume that there are $\gg_M N^3$ triples $h, h', h'' \in H$ such that $\Xi_* \cup \Xi_h \cup \Xi_{h'} \cup \Xi_{h''}$ fails to be $O_M(1)$-dissociated up to $O_M(1/N)$. To each such triple is associated a $k_* + 3(k - k_*)$ tuple

$$q_{*,1}, \dots, q_{*,k_*}, q_{h,k_*+1}, \dots, q_{h,k}, q_{h',k_*+1}, \dots, q_{h',k}, q_{h'',k_*+1}, \dots, q_{h'',k}$$

of integers, all at most $O_M(1)$ in magnitude, such that

$$\begin{aligned}\|q_{*,1}\xi_{*,1} + \dots + q_{*,k_*}\xi_{*,k_*} + q_{h,k_*+1}\xi_{h,k_*+1} + \dots + q_{h,k}\xi_{h,k} + q_{h',k_*+1}\xi_{h,k_*+1} + \dots \\ + q_{h',k}\xi_{h',k} + q_{h'',k_*+1}\xi_{h'',k_*+1} + \dots + q_{h'',k}\xi_{h'',k}\|_{\mathbb{R}/\mathbb{Z}} = O_M(1/N).\end{aligned}$$

By pigeonholing we may pass to a further subcollection of triples h, h', h'' for which these integers $q_{h,j}, q_{h',j'}, q_{h'',j''}$ have no h, h', h''-dependence. If at least one of these latter quantities (with $j > k_*$) is non-zero then by relabeling we may assume it is $q_{h,k}$. All this having been done, let us fix h' and h'' appearing in $\gg_M N$ of these triples. We then have integers $q_1, \dots q_k = O_M(1)$, not all zero, and some frequency ξ_0 such that

$$\|\xi_0 + q_1\xi_{*,1} + \dots + q_{k_*}\xi_{*,k_*} + q_{k_*+1}\xi_{h,k_*+1} + \dots + q_k\xi_{h,k}\|_{\mathbb{R}/\mathbb{Z}} = O_M(1/N)$$

for at least $\gg_M N$ values of h. Furthermore (case 1) we have $\xi_0 = 0$ if $q_{k_*+1} = \dots = q_k = 0$; otherwise (case 2) we have $q_k \neq 0$.

Suppose we are in case 1 and that, without loss of generality, we have $q_{k_*} \neq 0$. Then ξ_{*,k_*} is in the $O_M(1)$-linear span, up to $O_M(1/N)$, of the set $\tilde{\Xi}_* := \{\frac{1}{Q}\xi_{*,1}, \dots, \frac{1}{Q}\xi_{*,k_*-1}\}$, where Q is the lowest common multiple of the integers up to $O_M(1)$. Taking $\tilde{\Xi}'_h = \Xi'_h$, we see that (i) is satisfied and also that Ξ_h is in the $O_M(1)$-linear span, up to $O_M(1/N)$, of $\tilde{\Xi}_* \cup \tilde{\Xi}'_h$. Suppose now that we are in case 2; then take $\tilde{\Xi}_* = \frac{1}{Q}\Xi_* \cup \{\frac{1}{Q}\xi_0\}$ and $\tilde{\Xi}'_h = \frac{1}{Q}\Xi'_h \setminus \{\frac{1}{Q}\xi_{h,k}\}$. Now condition (ii) is satisfied, and once again Ξ_h is in the $O_M(1)$-linear span, up to $O_M(1/N)$, of $\tilde{\Xi}_* \cup \tilde{\Xi}'_h$.

https://doi.org/10.1017/S0017089510000546 Published online by Cambridge University Press

The treatment of the two cases is, henceforth, the same and at this point we revert to the bracket quadratic expressions

$$F_{[i',i]}^{m_{[i',i]}(h)}(g_h(n)\Gamma) = e(m_{[i',i]}(h)\xi_{h,i'}n\lfloor\xi_{h,i}n\rfloor).$$

For each $\xi_{h,i}$ we substitute in the expression for this frequency as an $O_M(1)$-linear combination of the frequencies in $\tilde{\Xi}'_* \cup \tilde{\Xi}'_h$, plus an error which is $O_M(1/N)$. To simplify this we use the bracket identities of Lemma 5.5 repeatedly to express the whole product $\chi_h(n)$ as a product of terms $e(\tilde{m}_{[i',i]}(h)\tilde{\xi}_{h,i'}n\lfloor\tilde{\xi}_{h,i}n\rfloor)$ with $i < i'$, where the exponents $\tilde{m}_{[i',i]}(h)$ are still $O_M(1)$, together with various terms of the form $e(\theta n^2)$, $e(\{\alpha n\}\{\beta n\})$ and $e(\alpha n\lfloor\beta n\rfloor)$ with $\beta = O_M(1/N)$.

Now we may use Lemma 3.5 (ii), (iii) and (iv) repeatedly, bearing in mind the assumption $|\mathbb{E}_{n\in[N]}\Delta_h f(n)\overline{\chi_h(n)}| \geqslant 1/M$, to remove all terms of these last two types and replace them by a single linear term $e(\theta' n)$. Doing this and then taking the new bracket quadratics $e(\tilde{m}_{[i',i]}(h)\tilde{\xi}_{h,i'}n\lfloor\tilde{\xi}_{h,i}n\rfloor)$ and writing them as nilcoordinate functions $F_{[i',i]}^{\tilde{m}_{[i',i]}(h)}(\tilde{g}_h(n)\tilde{\Gamma})$, we obtain precisely the desired conclusion (7.15). □

8. Step 2: Approximate linearity. In this section we address *Step 2* of the outline in Section 2. In the last section we decomposed the underlying frequency sets $\Xi_h = \{\xi_{h,1}, \dots, \xi_{h,k}\}$ into a core set Ξ_* and a petal set Ξ'_h, in such a way that no nilcharacter $F_{[i,i']}(g_h(n))$ corresponding to two petal frequencies $\xi_{h,i}, \xi_{h,i'}$ appears in the expression for $\chi_h(n)$. Our task now is to proceed from here to show that, at least for many h, the petal set Ξ'_h has a weak linear structure. There follows a precise statement of what we shall prove. By a *bracket-linear form* of complexity M we mean a function $\psi : \mathbb{Z} \to \mathbb{R}/\mathbb{Z}$ of the form

$$\psi(h) = \beta_0 + \alpha_1\{\beta_1 h\} + \cdots + \alpha_m\{\beta_m h\} + \theta h,$$

where the α_j, β_j, β lie in $\mathbb{R}$ and $m \leqslant M$.

PROPOSITION 8.1. *Suppose that $f : [N] \to \mathbb{C}$ is a 1-bounded function with $\|f\|_{U^4} \geqslant \delta$. Then there is a set $H \subseteq [N]$, $|H| \gg_\delta N$, such that for all $h \in H$ we have*

$$|\mathbb{E}_{n\in[N]}\Delta_h f(n)\overline{\chi_h(n)}| \gg_\delta 1.$$

Here we have

$$\chi_h(n) = e(\alpha_h n^2 + \beta_h n) \prod_{1\leqslant i<i'\leqslant k} F_{[i',i]}^{m_{[i',i]}(h)}(g_h(n)) \tag{8.1}$$

with k, $|m_{[i,i']}(h)| = O_\delta(1)$, where $g_h(n) = (\xi_{h,1}n, \dots, \xi_{h,k}n, 0, \dots, 0)$, $m_{[i',i]}(h) = 0$ if $i, i' > k_*$, the frequency set Ξ_h decomposes as $\Xi_* \cup \Xi'_h$ with $\Xi_* = \{\xi_{h,1}, \dots, \xi_{h,k_*}\}$ independent of h, and every frequency $\xi_{h,i}$, $i > k_*$, in the petal set Ξ'_h is a bracket linear form in h of complexity $O_\delta(1)$.

We shall establish this proposition inductively in a manner not too dissimilar to that in the last section. The inductive step which drives Proposition 8.1 is the following; it might be compared to Proposition 7.5 in the last section.

https://doi.org/10.1017/S0017089510000546 Published online by Cambridge University Press

PROPOSITION 8.2. *Suppose that $H \subseteq [N]$ is a set with $|H| \geqslant N/M$. Suppose that for all $h \in H$ we have*

$$|\mathbb{E}_{n\in[N]}\Delta_h f(n)\overline{\chi_h(n)}| \geqslant 1/M,$$

where $\chi_h(n)$ has the form (8.1) *and the frequency set Ξ_h is decomposed as $\Xi_* \cup \Xi_h^{\mathrm{struct}} \cup \Xi_h^{\mathrm{unstruct}}$, where the frequencies in Ξ_* do not depend on h and those in Ξ_h^{struct} are bracket-linear in h with complexity at most M. Then there is a set $\tilde{H} \subseteq H$, $|\tilde{H}| \gg_M 1$, such that*

$$|\mathbb{E}_{n\in[N]}\Delta_h f(n)\overline{\tilde{\chi}_h(n)}| \gg_M 1,$$

where $\tilde{\chi}_h(n)$ has the form

$$\tilde{\chi}_h(n) = e(\tilde{\alpha}_h n^2 + \tilde{\beta}_h n) \prod_{1\leqslant i<i'\leqslant \tilde{k}} F_{[i',i]}^{\tilde{m}_{[i',i]}(h)}(\tilde{g}_h(n)\tilde{\Gamma}),$$

a nilcharacter with complexity $O_M(1)$ in which the frequency set $\tilde{\Xi}_h$ decomposes as $\tilde{\Xi}_* \cup \tilde{\Xi}_h^{\mathrm{struct}} \cup \tilde{\Xi}_h^{\mathrm{unstruct}}$ where either

(i) (core decreases) $|\tilde{\Xi}_*| < |\Xi_*|$, $|\tilde{\Xi}_h^{\mathrm{struct}}| \leqslant |\Xi_h^{\mathrm{struct}}|$, $|\tilde{\Xi}_h^{\mathrm{unstruct}}| \leqslant |\tilde{\Xi}_h^{\mathrm{unstruct}}|$;

(ii) (unstructured part decreases) $|\tilde{\Xi}_*| = O_{M,|\Xi_*|}(1)$, $|\tilde{\Xi}_h^{\mathrm{struct}}| = |\Xi_h^{\mathrm{struct}}| + 1$, $|\tilde{\Xi}_h^{\mathrm{unstruct}}| = |\Xi_h^{\mathrm{unstruct}}| - 1$.

Proof of Proposition 8.1 given Proposition 8.2. To prove Proposition 8.1 one first, of course, applies *Step 1*. With that in hand one may pick $M = O_\delta(1)$ and initialise the inductive use of Proposition 8.2 by taking $\Xi_h^{\mathrm{unstruct}}$ to equal to the entire petal frequency Ξ_h' and $\Xi_h^{\mathrm{struct}} = \emptyset$. It is not hard to see that this repeated application of Proposition 8.2 terminates in time $O_M(1)$, at which point the unstructured component $\Xi_h^{\mathrm{unstruct}}$ is empty. □

It remains, of course, to prove Proposition 8.2, and this will be the main business of this section. Once again the key tool is Proposition 6.1, of which we require the following variant.

LEMMA 8.3. *Suppose that $H \subseteq [N]$ is a set with $|H| \geqslant N/M$ and that*

$$|\mathbb{E}_{n\in[N]}\Delta_h f(n)\overline{\chi_h(n)}| \geqslant 1/M$$

for all $h \in H$, where $\chi_h(n)$ has the form (8.1) *with $m_{[i,i']}(h) = 0$ if $i, i' > k_*$ and the underlying frequency set Ξ_h has been decomposed as $\Xi_* \cup \Xi_h^{\mathrm{struct}} \cup \Xi_h^{\mathrm{unstruct}}$, where the core Ξ_* does not depend on h and Ξ_h^{struct} consists of bracket linear forms of complexity at most M. Write $\chi_h(n) = \chi_h^{\mathrm{struct}}(n)\chi_h^{\mathrm{unstruct}}(n)$, where the two parts here correspond to the structured and unstructured frequencies in Ξ_h. Then there is a set $\tilde{H} \subseteq H$, $|\tilde{H}| \gg_M |H|$, and frequencies $\alpha_{h_1,h_2,h_3,h_4}, \beta_{h_1,h_2,h_3,h_4} \in \mathbb{R}/\mathbb{Z}$ such that*

$$\mathbb{E}_{n\in[N]}\chi_{h_1}^{\mathrm{unstruct}}(n)\chi_{h_2}^{\mathrm{unstruct}}(n)\chi_{h_3}^{\mathrm{unstruct}}(n)\chi_{h_4}^{\mathrm{unstruct}}(n)e(\alpha_{h_1,h_2,h_3,h_4}n^2 + \beta_{h_1,h_2,h_3,h_4}n) \gg_M 1$$

for $\gg_M N^3$ additive quadruples $h_1 + h_2 = h_3 + h_4 \in \tilde{H}$.

Proof. The idea is to apply Proposition 6.1 and then simply observe that the contribution from the structured parts $\chi_h^{\mathrm{struct}}(n)$ can be made to cancel out. Bracket linear forms are not quite genuinely linear, but if $\psi(h) = \alpha_1\{\beta_1 h\} + \cdots + \alpha_m\{\beta_m h\} + \theta h$

https://doi.org/10.1017/S0017089510000546 Published online by Cambridge University Press

then we have $\psi(h_1)+\psi(h_2)=\psi(h_3)+\psi(h_4)$ whenever the tuple $(\beta_1 h,\dots,\beta_m h)(\mathrm{mod}\ 1)$ lies in some cube $\prod_{j=1}^m [i_j/10, \gamma_j+(i_j+1)/10]$ (say), where the i_j are integers between 0 and 9. By pigeonholing we may pass to a set $\tilde{H}\subseteq H$ such that for each bracket-linear form $\psi(h)$ in Ξ_h^{struct}, and for all $h\in\tilde{H}$, the corresponding tuple always lies in a cube of this form depending only on ψ, and not on h.

By Proposition 6.1 there are $\gg_M N^3$ additive quadruples $h_1+h_2=h_3+h_4\in\tilde{H}$ and frequencies $\alpha_{h_1,h_2,h_3,h_4}, \beta_{h_1,h_2,h_3,h_4}\in\mathbb{R}/\mathbb{Z}$ such that

$$|\mathbb{E}_n \chi_{h_1}(n)\chi_{h_2}(n)\overline{\chi_{h_3}(n)\chi_{h_4}(n)}e(\alpha_{h_1,h_2,h_3,h_4}n^2+\beta_{h_1,h_2,h_3,h_4}n)|\gg_M 1.$$

Now the contribution to this from the structured parts,

$$\chi_{h_1}^{\mathrm{struct}}(n)\chi_{h_2}^{\mathrm{struct}}(n)\overline{\chi_{h_3}^{\mathrm{struct}}(n)\chi_{h_4}^{\mathrm{struct}}(n)},$$

is a product of bracket quadratic terms of the form

$$e(\psi(h_1)n\lfloor\theta n\rfloor+\psi(h_2)n\lfloor\theta n\rfloor-\psi(h_3)n\lfloor\theta n\rfloor-\psi(h_4)n\lfloor\theta n\rfloor)$$

or

$$e(\theta n\lfloor\psi(h_1)n\rfloor+\theta n\lfloor\psi(h_2)n\rfloor-\theta n\lfloor\psi(h_3)n\rfloor-\theta n\lfloor\psi(h_4)n\rfloor).$$

For the quadruples h_1,h_2,h_3,h_4 under consideration we have $\psi(h_1)+\psi(h_2)=\psi(h_3)+\psi(h_4)$, and so the first of these expressions is identically 1. The second is not, but by applying Lemma 5.5 we see that it is merely a combination of terms of the form $e(\theta n^2)$, $e(\{\alpha n\}\{\beta n\})$ and $e(\alpha n\lfloor\beta n\rfloor)$ with $\|\beta\|_{\mathbb{R}/\mathbb{Z}}=O_M(1/N)$, where α, β and θ depend on h_1,h_2,h_3,h_4. Applying Lemma 3.5, it follows that we may completely ignore the contribution from these structured parts, although we may need to modify the frequencies α_{h_1,h_2,h_3,h_4}, β_{h_1,h_2,h_3,h_4}. □

The next task is to use a similar (but much simpler) argument to that used for Lemma 7.3 to study the conclusion of Lemma 8.3 for a particular quadruple $h_1+h_2=h_3+h_4$.

LEMMA 8.4. *Let h_1,h_2,h_3,h_4 be fixed and suppose that nilcharacters $\chi_{h_j}(n)$ have the form* (8.1). *Suppose that for each $j=1,2,3,4$ the underlying frequency set Ξ_{h_j} is decomposed as $\Xi_*\cup\Xi_{h_j}^{\mathrm{struct}}\cup\Xi_{h_j}^{\mathrm{unstruct}}$, where the core set Ξ_* does not depend on h_j and each element of $\Xi_{h_j}^{\mathrm{struct}}$ is a bracket linear form $\psi(h_j)$, again not depending on h_j. Suppose that*

$$|\mathbb{E}_{n\in[N]}\chi_{h_1}^{\mathrm{unstruct}}(n)\chi_{h_2}^{\mathrm{unstruct}}(n)\overline{\chi_{h_3}^{\mathrm{unstruct}}(n)\chi_{h_4}^{\mathrm{unstruct}}(n)}e(\alpha_{h_1,h_2,h_3,h_4}n^2+\beta_{h_1,h_2,h_3,h_4}n)|\geqslant 1/M.$$

Suppose that not all of the integers $m_{[i,i']}(h_j)$ corresponding to frequencies $\xi_{h_j,i},\xi_{h_j,i'}$, one of which is in $\Xi_{h_j}^{\mathrm{unstruct}}$, vanish. Then some there is some $O_M(1)$-rational relation, up to $O_M(1/N)$, amongst the elements of $\Xi_*\cup\Xi_{h_1}^{\mathrm{unstruct}}\cup\Xi_{h_2}^{\mathrm{unstruct}}\cup\Xi_{h_3}^{\mathrm{unstruct}}\cup\Xi_{h_4}^{\mathrm{unstruct}}$.

Proof. Once again we interpret the assumption as an assertion about a 2-step nilsequence. Perhaps the 'correct' way to do this (and the manner more amenable to generalisation) would be to mimic the construction of the last section and apply the Quantitative Ratner theorem once again. However in the special case of the U^4-norm being addressed by this paper a shortcut in which only the (simpler) quantitative Leibman dichotomy, Theorem 4.1, is needed and we give this here. Let

https://doi.org/10.1017/S0017089510000546 Published online by Cambridge University Press

us take G to be the free 2-step nilpotent Lie group on the ordered generating set $\{e_\xi : \xi \in \Xi_* \cup \Xi_{h_1}^{\text{unstruct}} \cup \Xi_{h_2}^{\text{unstruct}} \cup \Xi_{h_3}^{\text{unstruct}} \cup \Xi_{h_4}^{\text{unstruct}}\}$. As in Section 5 we identify the 'coordinate' functions $F_{\xi,\xi'} : G/\Gamma \to \mathbb{C}$, and we take a polynomial sequence $g : \mathbb{Z} \to G$ whose coordinate at e_ξ is ξn, for all ξ in the above indexing set, and all of whose other coordinates are zero except for that at $[e_\xi, e_{\xi'}]$ for some arbitrary pair of frequencies ξ, ξ' in the above set, where the coordinate of g is some quadratic $q = q_{h_1,h_2,h_3,h_4}(n)$ to be specified shortly. Inside G take Γ to be the lattice of integer points in the free 2-step nilpotent Lie group. Finally, take

$$F := \prod_{i=1}^{4} F_{\xi_{*,j},\xi_{h_i,j'}}^{m_{j,j'}(h_i)}.$$

By an appropriate choice of the quadratic term q we may ensure that

$$F(g(n)\Gamma) = \chi_{h_1}^{\text{unstruct}}(n) \dots \chi_{h_4}^{\text{unstruct}}(n) e(\alpha_{h_1,h_2,h_3,h_4} n^2 + \beta_{h_1,h_2,h_3,h_4} n).$$

Note that we have $\int_{G/\Gamma} F = 0$. Although F is only piecewise Lipschitz, it is nonetheless the case that if $(g(n)\Gamma)_{n\in P}$ is δ-equidistributed for an appropriate $\delta \gg_M$ then $|\mathbb{E}_{n\in[N]} F(g(n)\Gamma)| \leqslant 1/10M$, contrary to assumption. This is because, as in the last section, we may decompose F as a sum $F_0 + F_1$ where $\|F_0\|_{\text{Lip}} = O_{M,\varepsilon}(1)$ and $|F_1|$ is bounded above pointwise by a function Ψ with $\int_{G/\Gamma} \Psi = O(\varepsilon)$ and $\|\Psi\|_{\text{Lip}} = O_{M,\varepsilon}(1)$.

Thus we are forced to conclude that $(g(n)\Gamma)_{n\in P}$ is not δ-equidistributed on G/Γ, for some $\delta \gg_M 1$. By the quantitative Leibman dichotomy, Theorem 4.1, this implies that there is some $k \in \mathbb{Z}^{\dim(G:[G,G])}$, $0 < |k| = O_M(1)$, such that $\|k \cdot (\pi \circ g)\|_{C^\infty[N]} = O_M(1/N)$. In view of the way that $\pi \circ g$ was constructed, namely the fact that the horizontal part $\pi \circ g$ contains only the terms ξn with $\xi \in \Xi_* \cup \Xi_{h_1}^{\text{unstruct}} \cup \Xi_{h_2}^{\text{unstruct}} \cup \Xi_{h_3}^{\text{unstruct}} \cup \Xi_{h_4}^{\text{unstruct}}$, this is precisely the result claimed. □

The conclusion of Lemma 8.4 looks rather weak, but using the tools of additive combinatorics pioneered in this context by Gowers (particularly in [**8**, Ch. 7]) it turns out to be enough for us to be able to impose some bracket linear behaviour on some of the unstructured sets Ξ_h^{unstruct}. The following result concerning approximate homomorphisms is our key tool. We know of no source for this precise result in the literature, though we feel it should be somehow be regarded as 'known'. It is appropriate to associate the names of Freĭman, Ruzsa and Gowers with results of this kind.

PROPOSITION 8.5 (Approximate homomorphisms). *Let $\delta, \varepsilon \in (0,1)$ be parameters and suppose that $f_1, f_2, f_3, f_4 : S \to \mathbb{R}/\mathbb{Z}$ are functions defined on some subset $S \subseteq [N]$ such that there are at least δN^3 quadruples $(x_1, x_2, x_3, x_4) \in S^4$ with $x_1 + x_2 = x_3 + x_4$ and $\|f_1(x_1) + f_2(x_2) - f_3(x_3) + f_4(x_4)\|_{\mathbb{R}/\mathbb{Z}} \leqslant \varepsilon$. Then there is a bracket linear phase $\psi : \mathbb{Z} \to \mathbb{R}/\mathbb{Z}$ of complexity $O_\delta(1)$ and a set $S' \subseteq S$, $|S'| \gg_\delta N$, such that $f_1(x) = \psi(x) + O(\varepsilon)$ for all $x \in S'$.*

Proof. See Appendix C. □

LEMMA 8.6. *Let $H \subseteq [N]$ be a set of size at least N/M, and suppose that we have a core set Ξ_* and, for each $h \in H$, sets Ξ_h^{unstruct}. Suppose that $|\Xi_*|, |\Xi_h^{\text{unstruct}}| \leqslant M$. Suppose that for at least N^3/M additive quadruples $h_1 + h_2 = h_3 + h_4$ in H there is an M-linear relation, up to $O(M/N)$, in $\Xi_* \cup \Xi_{h_1}^{\text{unstruct}} \cup \Xi_{h_2}^{\text{unstruct}} \cup \Xi_{h_3}^{\text{unstruct}} \cup \Xi_{h_4}^{\text{unstruct}}$. Then either*

https://doi.org/10.1017/S0017089510000546 Published online by Cambridge University Press

(i) *There is some element of the core* Ξ_* *which lies in the* $O_M(1)$*-span of the others, up to* $O_M(1/N)$*, or*

(ii) *There is a bracket linear form* ψ *of degree* $O_M(1)$ *and a set* $H' \subseteq H$, $|H'| \gg_M |H|$, *such that* $\psi(h)$ *lies in the* $O_M(1)$*-linear span up to* $O_M(1/N)$ *of* $\Xi_h^{\mathrm{unstruct}}$ *for all* $h \in H'$.

Proof. Let the elements of the core set Ξ_* be $\{\xi_{*,1}, \dots, \xi_{*,M}\}$ and those of the petal set Ξ_h be $\{\xi_{h,1}, \dots, \xi_{h,M}\}$. Suppose, for a given quadruple $h_1 + h_2 = h_3 + h_4$, that the approximate linear relation between the elements of $\Xi_* \cup \Xi_{h_1} \cup \Xi_{h_2} \cup \Xi_{h_3} \cup \Xi_{h_4}$ is

$$\begin{aligned}\|&q_{*,1}(h_1,h_2,h_3,h_4)\xi_{*,1} + \cdots + q_{*,M}(h_1,h_2,h_3,h_4)\xi_{*,M}\\ &+ q_{1,1}(h_1,h_2,h_3,h_4)\xi_{h_1,1} + \cdots + q_{1,M}(h_1,h_2,h_3,h_4)\xi_{h_1,M}\\ &+ \cdots + q_{4,1}(h_1,h_2,h_3,h_4)\xi_{h_4,1} + \cdots + q_{4,M}(h_1,h_2,h_3,h_4)\xi_{h_4,M}\|_{\mathbb{R}/\mathbb{Z}} = O(M/N),\end{aligned}$$

where each integer q has magnitude at most M. There are only $(2M+1)^{5M}$ choices for these integers and so we may pass to a subcollection of $\gg_M N^3$ quadruples for which there is such a relation and for which none of the q's depend on h_1, h_2, h_3, h_4. Since Ξ_* is M-dissociated, at least one of the $q_{i,j}$ must be non-zero, $i = 1, 2, 3, 4$; without loss of generality, suppose that $q_{1,1} \neq 0$.

Writing $f_1(h_1) := q_{1,1}\xi_{h_1,1} + \cdots + q_{1,M}\xi_{h_1,M}$, we see that we have found functions $f_2, f_3, f_4 : H \to \mathbb{R}/\mathbb{Z}$ such that

$$\|f_1(h_1) + f_2(h_2) - f_3(h_3) - f_4(h_4)\|_{\mathbb{R}/\mathbb{Z}} \leqslant 1/N'$$

for $\gg_M N^3$ additive quadruples $h_1 + h_2 = h_3 + h_4 \in H$, for some $N' \gg N/M$. Now we apply Proposition 8.5 to conclude that there is a bracket linear phase ψ of complexity $O_M(1)$ such that $f_1(h) = \psi(h) + O_M(1/N)$ for all h in some set $H' \subseteq H$, $|H'| \gg_M N$. This concludes the proof of the lemma. □

We are now in a position to prove Proposition 8.2 which, recall, was the inductive step driving the main result of this section, namely Proposition 8.1. The argument is very similar to that employed in the proof of Proposition 7.5, hingeing on repeated use of the bracket identities of Lemma 5.5 to expand out linear combinations of frequencies.

Proof of Proposition 8.2. The assumption that $|\mathbb{E}_{n\in P}\Delta_h f(n)\overline{\chi_h(n)}| \geqslant 1/M$ may be fed into Lemma 8.3 to conclude the existence of a set $H' \subseteq H$ with $|H'| \gg_M |H|$ such that

$$\mathbb{E}_n \chi_{h_1}^{\mathrm{unstruct}}(n)\chi_{h_2}^{\mathrm{unstruct}}(n)\chi_{h_3}^{\mathrm{unstruct}}(n)\chi_{h_4}^{\mathrm{unstruct}}(n)e(\alpha_{h_1,h_2,h_3,h_4}n^2 + \beta_{h_1,h_2,h_3,h_4}n)) \gg_M 1$$

for $\gg_M N^3$ additive quadruples $h_1 + h_2 = h_3 + h_4$ in H'. This in turn may be fed into Lemma 8.4, which allows us to conclude that for each of these additive quadruples there is an $O_M(1)$ linear relation, up to $O_M(1/N)$, between the elements of $\Xi_* \cup \Xi_{h_1}^{\mathrm{unstruct}} \cup \Xi_{h_2}^{\mathrm{unstruct}} \cup \Xi_{h_3}^{\mathrm{unstruct}} \cup \Xi_{h_4}^{\mathrm{unstruct}}$. There is one other possibility here, namely that in the attempt to apply Lemma 8.4 we find that, for many quadruples $h_1 + h_2 = h_3 + h_4$, all of the integers $m_{[i,i']}(h_j)$ corresponding to frequencies $\xi_{h_j,i}, \xi_{h_j,i'}$, one of which is in $\Xi_{h_j}^{\mathrm{unstruct}}$, are zero. This is a rather trivial case, however, for we may then pass to the set $\tilde{H}$ of h_1 (say) appearing here, and simply delete the unstructured frequencies $\Xi_h^{\mathrm{unstruct}}$,

https://doi.org/10.1017/S0017089510000546 Published online by Cambridge University Press

which play no actual role in the expression for $\chi_h(n)$. The conclusion of Proposition 8.2 is then immediate in this case.

Returning to the main line of the argument, we may then apply Lemma 8.6 to conclude that either

(i) There is some element $\xi \in \Xi_*$ which lies in the $O_M(1)$-linear span of the others, up to $O_M(1/N)$, or
(ii) There is a bracket linear form ψ of degree $O_M(1)$ and a set $\tilde{H} \subseteq H'$, $|\tilde{H}| \gg_M |H|$, so that $\psi(h)$ lies in the $O_M(1)$-linear span of $\Xi_h^{\mathrm{unstruct}}$ for all $h \in \tilde{H}$.

These two possibilities will correspond to alternatives (i) and (ii) respectively in Proposition 8.2. To see this we proceed rather as in the proof of Proposition 7.5, making use once again of Lemma 5.5 as well as extensive use of Lemma 3.5 to handle the somewhat annoyingly non-Lipschitz 1-step objects which arise. The treatment of (i) is exactly analogous to the aforementioned argument, so we only describe (ii) in any detail.

Assume that the sets $\Xi_h^{\mathrm{unstruct}}$ are ordered as $\xi_{h,k_0+1}, \dots, \xi_{h,k}$. We are assuming that there is a bracket-linear form $\psi(h)$ having the form $q_{h,k_0+1}\xi_{h,k_0+1} + \cdots + q_{h,k}\xi_{h,k} + O_M(1/N)$, for all $h \in H'$. Here the integers $q_{h,j}$ are all bounded in magnitude by $O_M(1)$ and so we may, by passing to a further subset $H'' \subseteq H'$, assume that they do not depend on h. Without loss of generality let us suppose that $q_{h,k} \neq 0$. Then we may write $\xi_{h,k}$ as an $O_M(1)$-linear combination of $\frac{1}{Q}\psi(h)$ and the frequencies $\frac{1}{Q}\xi_{h,k_0+1}, \dots, \frac{1}{Q}\xi_{h,k-1}$, plus an error of $O_M(1/N)$, where Q is the lcm of the numbers up to $O_M(1)$. Now we replace Ξ_h^{struct} by $\Xi_h^{\mathrm{struct}} \cup \{\frac{1}{Q}\psi(h)\}$ and $\Xi_h^{\mathrm{unstruct}}$ by $\{\frac{1}{Q}\xi_{h,k_0+1}, \dots, \frac{1}{Q}\xi_{h,k-1}\}$, and then proceed to rewrite the bracket quadratics $e(\xi n\lfloor \xi' n\rfloor)$ using these new sets of frequencies by means of Lemma 5.5 and Lemma 3.5 exactly as we did at the end of Section 7. □

Before moving onto the next section we apply one additional piece of analysis to Proposition 8.1. This allows us to conclude that the quadratic frequency α_h varies bracket-linearly in h as well. Thus, once this is done, only the linear term $e(\beta_h n)$ does not have a rigid structure imposed upon it.

PROPOSITION 8.7. *In the statement of Proposition 8.1, we may assume that the quadratic frequency α_h varies bracket-linearly in h.*

Proof. We may, of course, take for granted the conclusion of Proposition 8.1. We apply Proposition 6.1 once again, using the same argument we employed at the start of the proof of Lemma 8.3 to first pass to a subset $H' \subseteq H$, $|H'| \gg_M N$, on which all the bracket linear forms ψ in the petals Ξ_h' are linear in the sense that $\psi(h_1) + \psi(h_2) = \psi(h_3) + \psi(h_4)$ whenever $h_1 + h_2 = h_3 + h_4$ with $h_1, h_2, h_3, h_4 \in H'$. This gives

$$\mathbb{E}_{n\in[N]}\chi_{h_1}(n)\chi_{h_2}(n+h_1-h_4)\chi_{h_3}(n)\chi_{h_4}(n+h_1-h_4) \gg_M 1.$$

As in Corollary 6.2, this implies that $\chi_{h_1}(n)\chi_{h_2}(n)\chi_{h_3}(n)\chi_{h_4}(n)$ correlates with a quadratic phase $e(\alpha_{h_1,h_2,h_3,h_4}n^2 + \beta_{h_1,h_2,h_3,h_4}n)$. Moreover a careful analysis of the proof of that corollary, looking at the decomposition $\chi_h(n) = \chi_h'(n)e(\alpha_h n^2 + \beta_h n)$, where

$$\chi_h'(n) = \prod_{1\leqslant i<i'\leqslant k} F_{[i',i]}^{m_{[i',i]}(h)}(g_h(n)\Gamma),$$

reveals that we can take $\alpha_{h_1,h_2,h_3,h_4} = \alpha_{h_1} + \alpha_{h_2} - \alpha_{h_3} - \alpha_{h_4}$. That is, the genuinely bracket-quadratic objects comprising $\chi_h'(n)$ only give rise to *linear* terms.

https://doi.org/10.1017/S0017089510000546 Published online by Cambridge University Press

The term $\chi'_{h_1}(n)\chi'_{h_2}(n)\chi'_{h_3}(n)\chi'_{h_4}(n)$ arising from the genuinely bracket quadratic parts is a product of terms of the form $e(\alpha n\lfloor\psi(h_1)n\rfloor + \alpha n\lfloor\psi(h_2)n\rfloor - \alpha n\lfloor\psi(h_3)n\rfloor - \alpha n\lfloor\psi(h_4)n\rfloor)$ where, recall, $\psi(h_1)+\psi(h_2)=\psi(h_3)+\psi(h_4)$. Using Lemma 5.5 (iii) to move the ψ terms to the outside of the brackets and applying Lemma 3.5 repeatedly, we conclude that

$$\mathbb{E}_{n\in[N]}e((\alpha'_{h_1}+\alpha'_{h_2}-\alpha'_{h_3}-\alpha'_{h_4})n^2+\theta_{h_1,h_2,h_3,h_4}n) \gg_M 1$$

for all these quadruples $h_1+h_2=h_3+h_4$, where $\alpha'_h-\alpha_h$ is a bracket-linear form of complexity $O_M(1)$. By Lemma D.1 it follows that there is some $q=O_M(1)$ such that

$$\|q(\alpha'_{h_1}+\alpha'_{h_2}-\alpha'_{h_3}-\alpha'_{h_4})\|_{\mathbb{R}/\mathbb{Z}} = O_M(1/N^2).$$

By Proposition 8.5 there is a further subset $H''\subseteq H$, $|H''|\gg_M |H|$, together with a bracket linear form $\psi'(h)$ of complexity $O_M(1)$, such that

$$q\alpha'_h = \psi'(h) + O_M(1/N^2)$$

for all $h\in H''$. This means that

$$\alpha_h = \psi''(h) + \frac{r_h}{q} + O_M(1/N^2),$$

where $\psi''(h)$ is another bracket linear form and r_h takes integer values. Refining $[N]$ into progressions of common difference q and length $\gg_M N$ small enough to make the $O_M(1/N^2)$ error negligible, and then applying Lemma 3.5 (ii), we obtain the claim. □

9. Step 3: The symmetry argument. Finally we turn to *Step 3* of the programme outlined in Section 2, the so-called *symmetry argument*. Here we shall take an approach somewhat different to the one we shall employ in the general case of the U^{s+1}-norm, $s\geqslant 4$, where further use is made of the nilmanifold distribution results of Section 4 and there are slightly complicated issues concerning the keeping-track of the complexity of various bracket expressions.

In the special case of the U^4-norm that this paper is concerned with, a rather direct argument using Bohr sets is possible. Let $S=\{\theta_1,\dots,\theta_d\}\subseteq\mathbb{R}/\mathbb{Z}$ be a set of frequencies and suppose that $\rho\in(0,1)$. Then we set

$$B(S,\rho,N) := \{n\in[\rho N] : \|n\theta_j\|_{\mathbb{R}/\mathbb{Z}}\leqslant\rho \quad \text{for all } j=1,\dots,d\}.$$

We shall need a small amount of the theory of such sets, particularly pertaining to the notion of *regularity* – the idea that there is a plentiful supply of ρ for which the size of $B(S,\rho',N)$ is nicely controlled for $\rho'\approx\rho$. The need to introduce this idea in additive combinatorics was first appreciated in [**5**] and it has now appeared in several places, for example [**11**] where the notion is defined in Definition 2.6 and discussed in more detail in Chapter 8.

For our purposes here we say that a value ρ is *regular* if we have

$$|B(S,(1+\kappa)\rho,N)| = |B(S,\rho,N)|(1+O(d|\kappa|))$$

https://doi.org/10.1017/S0017089510000546 Published online by Cambridge University Press

uniformly for $|\kappa| \leqslant 1/d$. We shall need the following facts about regular Bohr sets. It would be possible to obtain much more precise statements but we shall not need to do so here.

LEMMA 9.1 (Regular Bohr sets – Basic Facts). *Fix a set $S = \{\theta_1, \dots, \theta_d\}$ of frequencies, and write $B := B(S, \rho, N)$. We have the following facts.*

(i) (Ubiquity of regular values) *For any $\rho_0 \in (0, 1/2)$ there is a regular value of ρ in the interval $[\rho_0, 2\rho_0]$.*

(ii) (Fourier expansion of Bohr cutoffs) *Suppose that ρ is regular, and that $\varepsilon > 0$ is a parameter. Then we may decompose the cutoff $1_B(n)$ as $\psi_1(n) + \psi_2(n)$, where $\psi_1(n) = \int_0^1 \widehat{\psi}_1(\theta) e(\theta n)\, d\theta$ with $\|\widehat{\psi}_1\|_1 := \int_0^1 |\widehat{\psi}_1(\theta)|\, d\theta \leqslant C_{\epsilon,\rho}$ and $\sum_n |\psi_2(n)| \leqslant \varepsilon N$.*

(iii) (Large generalised Fourier coefficients) *Suppose that ρ is regular and that $\phi : B(S, 2\rho, N) \to \mathbb{R}/\mathbb{Z}$ is locally linear on B in the sense that $\phi(x + y) = \phi(x) + \phi(y)$ whenever $x, y \in B$. Suppose that $|\mathbb{E}_{x \in B} e(\phi(x))| \geqslant \eta$. Then there is a regular value of ρ', $\rho' \gg_{\varepsilon,\eta,\rho} 1$, such that $\|\phi(x)\|_{\mathbb{R}/\mathbb{Z}} \leqslant \varepsilon$ for all $x \in B(S, \rho')$.*

Sketch Proof. The definition of Bohr set we are using here is very slightly different to that used in [**11**], in that our Bohr sets are contained in $[N]$ and not in $\mathbb{Z}/N\mathbb{Z}$. Nonetheless, the proofs of the above statements are so close to those in $\mathbb{Z}/N\mathbb{Z}$ that we simply refer to the relevant sections of the aforementioned paper. Statement (i) is [**11**, Lemma 8.2]. Statement (ii) is not explicitly mentioned in [**11**]. To prove it, take $\psi_1(n) = \frac{1}{|B'|} 1_B * 1_{B'}(n)$, where $B' := B(S, \rho', N)$ for a suitably small $\rho' \gg_{\varepsilon,\rho} 1$. The bound on $\|\widehat{\psi}_1\|_1$ follows from Plancherel, whilst the bound on $\|\psi_2\|_1$ is a consequence of the regularity of B and the observation that $\frac{1}{|B'|} 1_B * 1_{B'}(n) = 1_B(n)$ provided that $n \notin B(S, \rho + \rho', N) \setminus B(S, \rho - \rho', N)$. Finally, (iii) is [**11**, Lemma 8.4]. □

Let us return to the main business of this section, which is to conclude the proof of Theorem 1.5. The main result of the last section, Proposition 8.7, took us from the assumption that $\|f\|_{U^4} \geqslant \delta$ to the conclusion that

$$|\mathbb{E}_{n \in [N]} f(n) \overline{f(n+h)} \chi_h(n)| \gg_\delta 1 \tag{9.1}$$

for a set H of size $\gg_\delta N$, where $\chi_h(n)$ is a product of terms of the form $e(\{\alpha h\}\beta n \lfloor \gamma n \rfloor)$, $e(\alpha\{\beta h\} n^2)$ and $e(\theta_h n)$. Using the fact that $\lfloor \gamma n \rfloor = \gamma n - \{\gamma n\}$, we may assume that

$$\chi_h(n) = e\Big(\sum_{j=1}^{k} \{\alpha_j h\}\beta_j n \{\gamma_j n\}\Big) e\Big(\sum_{j=1}^{k} \alpha'_j \{\beta'_j h\} n^2\Big) e(\theta_h n).$$

Later on it will be convenient to assume that

$$\text{For all } h \in H \text{ we have } \|\theta h\|_{\mathbb{R}/\mathbb{Z}} \geqslant \rho_1 \text{ for all } \theta \in \{\alpha_1, \dots, \alpha_k, \beta'_1, \dots, \beta'_k\}, \tag{9.2}$$

for some small parameter $\rho_1 > 0$ to be specified later. This can be achieved at the expense of thinning out H somewhat to a set of size merely $\gg_{\rho_1,\delta} N$, as we now show.

To demonstrate the last claim we distinguish two types of such θ. We say that θ is *good* if the number of $h \in H$ such that $\|\theta h\|_{\mathbb{R}/\mathbb{Z}} < \rho_1$ is at most $10\rho_1 N$. By refining H to a set $H' \subseteq H$ with $|H'| \geqslant |H| - 20\rho_1 kN$, we may assume that $\|\theta h\|_{\mathbb{R}/\mathbb{Z}} \geqslant \rho_1$ for all $h \in H'$ and for all good θ. Note that $|H'| \geqslant |H|/2$ if ρ_1 is chosen small enough as a function of δ, as it will be later on. If θ is not good then the sequence $\{n\theta (\mathrm{mod}\ \mathbb{Z})\}_{n \in [N]}$ is

https://doi.org/10.1017/S0017089510000546 Published online by Cambridge University Press

not ρ_1-equidistributed, and by well-known results of diophantine approximation (see, for example, [**14**, Proposition 3.1]) there is some $q \ll \rho_1^{-C}$ such that $\|q\theta\|_{\mathbb{R}/\mathbb{Z}} \ll \rho_1^{-C}/N$. This means that the bracket $\{\theta h\}$ takes on only ρ_1^{-2C} values as h ranges over $[N]$, and so there is a subset $H'' \subseteq H'$, $|H''| \gg \rho_1^{Ck}|H|$, on which all these brackets are constant. This means that the corresponding terms in $\chi_h(n)$ may be ignored, for the purpose of (9.1), since they depend just on n and not on h. Replacing H by H'' gives the claim, and henceforth we assume that (9.2) holds, remembering that we now only have the weaker bound $|H| \gg_{\rho_1,\delta} N$.

Write

$$T(x, y, z) := \sum_{j=1}^{k} \{\alpha_j x\} \frac{\beta_j}{3} y \{\gamma_j z\} + \sum_{j=1}^{k} \frac{\alpha_j'}{3} \{\beta_j' x\} yz,$$

so that $3T(h, n, n)$ is the form appearing in the definition of $\chi_h(n)$. Here, there are three possible choices for each $\beta_j/3$, $\alpha_j'/3$ and it does not matter which we take; the reason for introducing these 3's will become apparent later. Then $T(x, y, z)$ is trilinear on the Bohr set $B := B(S, \rho_0, N)$, where $S = \{\alpha_1, \dots, \alpha_k, \gamma_1, \dots, \gamma_k, \beta_1', \dots, \beta_k'\}$ and the parameter $\rho_0 \in [\frac{1}{20}, \frac{1}{10}]$ is chosen so that B is regular. By stating that T is trilinear we mean that, for example, $T(x_1 + x_2, y, z) = T(x_1, y, z) + T(x_2, y, z)$ when all of $x_1, x_2, x_1 + x_2, y, z$ lie in B. We begin by symmetrising T in the last two variables, a straightforward task. For each j pick some $\tilde{\beta}_j$ such that $2\tilde{\beta}_j = \beta_j/3$ (there are two choices) and set

$$\tilde{T}(x, y, z) := \sum_{j=1}^{k} \{\alpha_j x\} \tilde{\beta}_j y \{\gamma_j z\} + \sum_{j=1}^{k} \{\alpha_j x\} \tilde{\beta}_j z \{\gamma_j y\} + \sum_{j=1}^{k} \frac{\alpha_j'}{3} \{\beta_j' x\} yz.$$

Then of course $\tilde{T}(h, n, n) = T(h, n, n)$, but now $\tilde{T}(x, y, z)$ is symmetric in the last two variables. Dropping the tildes, we assume henceforth that T itself is symmetric in the last two variables.

Our assumption, then, is that

$$|\mathbb{E}_{n \in [N]} f(n)\overline{f(n+h)} e(3T(h, n, n)) e(\theta_h n)| \gg_\delta 1$$

for all h lying in some set H of size at least $\gg_{\rho_1,\delta} N$, where H additionally satisfies (9.2). Our immediate goal is to localise the variables h and n to small Bohr sets so that we may properly exploit the trilinearity of T.

Let us briefly reprise the heuristic mentioned in the Section 2 to recall why it is that we expect T to be symmetric in the first two coordinates as well (on a'nice set'). Suppose we knew that $f(n)\overline{f(n+h)} = \chi_h(n) = e(3T(h, n, n))e(\theta_h n)$ for all n, h. Then we get

$$\chi_h(n+k)\chi_k(n) = f(n)\overline{f(n+h+k)} = \chi_k(n+h)\chi_h(n).$$

Using the trilinearity of T and symmetry in the last two coordinates we get

$$6T(h, k, n) = 6T(k, h, n).$$

Now of course we do not have proper equations but only correlations, we don't have correlation for all h but only for 'many', and we have trilinearity only when the variables are restricted to Bohr sets, so we must work much harder.

https://doi.org/10.1017/S0017089510000546 Published online by Cambridge University Press

We start with the h variable. Set $B_1 := B(S, \rho_1, N)$, where ρ_1 is the (as yet unspecified) quantity appearing in (9.2). Modifying ρ_1 by at most a factor of two, we may assume that B_1 is regular. We claim that it is possible to find an $h_0 \in H$ such that the intersection $H \cap (h_0 + B_1)$ has size $\gg_{\rho_1,\delta} N$. A slight trick is necessary to establish this: consider

$$\sum_{n\in[N]} 1_H * 1_{B'} * 1_{B'}(n)1_H(n),$$

where $B' := B(S, \rho_1/2, N)$. On the one hand this equals $\sum_n 1_H * 1_{B'}(n)^2$ which, by the Cauchy–Schwarz inequality, is $\gg_{\rho_1,\delta} N^3$. On the other hand we have $1_{B'} * 1_{B'}(n) \leqslant |B_1|1_{B_1}(n)$ for $n \in [N]$, and from these two inequalities the claim follows immediately.

Our assumption now implies that

$$|\mathbb{E}_{n\in[N]} f(n)\overline{f(n + h_0 + h')}e(3T(h_0 + h', n, n))e(\theta_{h_0+h'}n)| \gg_\delta 1$$

for all h' lying in some set $H' \subseteq B_1 = B(S, \rho_1, N)$, $|H'| \gg_{\rho_1,\delta} N$. By the careful construction of H (cf. (9.2)) and the fact that $h_0 \in H$ we have $\{\alpha_j(h_0 + h')\} = \{\alpha_j h_0\} + \{\alpha_j h'\}$, and similarly for the β'_j, and hence we obtain the linearity property $T(h_0 + h', n, n) = T(h_0, n, n) + T(h', n, n)$. After relabelling we hence have

$$|\mathbb{E}_{n\in[N]} f_1(n)\overline{f_2(n + h)}e(3T(h, n, n))e(\theta_h n)| \gg_\delta 1$$

for all $h \in H$, where $H \subseteq B_1$, $|H| \gg_{\rho_1,\delta} N$, $f_1(n) := f(n)e(T(h_0, n, n))$ and $f_2(n) := f(n + h_0)$.

We must now localise the n variable, and for this we use a somewhat different trick. By averaging there is some n_0 such that

$$\mathbb{E}_{n\in[N]} f_1(n_0 + n)\overline{f_2(n_0 + n + h)}e(3T(h, n_0 + n, n_0 + n))e(\theta_h n)1_{B_1}(n) \gg_\delta 1.$$

Now we have

$$\begin{aligned}\beta(n + n_0)\{\gamma(n + n_0)\} = {} & \beta n_0\{\gamma(n + n_0)\} + \beta n\{\gamma n\} + \beta n\{\gamma n_0\} \\ & + \beta n(\{\gamma(n + n_0)\} - \{\gamma n\} - \{\gamma n_0\}).\end{aligned}$$

Substituting into the expression for $e(3T(h, n + n_0, n + n_0))$ and expanding, we see that the contribution from the term $e(\beta n_0\{\gamma(n + n_0)\})$ may be absorbed into the linear term $e(\theta_h n)$ (by Lemma 3.5), as may the term $e(\beta n\{\gamma n_0\})$ (trivially). The term $\{\gamma(n + n_0)\} - \{\gamma n\} - \{\gamma n_0\}$ takes values in $\{-1, 0, 1\}$ according to whether $\gamma n \pmod 1$ lies in certain intervals $I_\gamma^{-1}, I_\gamma^0, I_\gamma^{+1}$, and so we obtain

$$\begin{aligned}&\mathbb{E}_{n\in[N]} f'_1(n)\overline{f'_2(n + h)}e(3T(h, n, n))e(\theta_h n) \\ &\quad\times \prod_{j=1}^{k} \big(1_{\gamma_j n\in I_{\gamma_j}^{-1}} e(-\{\alpha_j h\}\beta_j n) + 1_{\gamma_j n\in I_{\gamma_j}^0} + 1_{\gamma_j n\in I_{\gamma_j}^{+1}} e(\{\alpha_j h\}\beta_j n)\big)1_{B_1}(n) \gg 1,\end{aligned}$$

where $f'_1(n) = f_1(n + n_0)$ and $f'_2(n) = f_2(n + n_0)$. It follows that there is a choice of $\varepsilon_j \in \{-1, 0, 1\}$ and a θ'_h such that

$$\mathbb{E}_{n\in[N]} f'_1(n)\overline{f'_2(n + h)}e(3T(h, n, n))e(\theta'_h n)\prod_{j=1}^{k} 1_{\gamma_j n\in I_{\gamma_j}^{\varepsilon_j}} 1_{B_1}(n) \gg 1.$$

https://doi.org/10.1017/S0017089510000546 Published online by Cambridge University Press

By Lemma 3.5 we may remove the last term at the expense of changing θ_h' again. Removing the dashes for notational convenience we now obtain

$$\mathbb{E}_{n\in[N]}f_1(n)f_2(n+h)e(3T(h,n,n))e(\theta_h n)1_{B_1}(n) \gg 1.$$

Here, $f_1(n) = f(n+n_0)e(T(h_0, n_0+n, n_0+n))$ and $f_2(n) = f(n+h_0+n_0)$, and we recall once more that this is known to hold for $\gg_{\rho_1,\delta} N$ values of $h \in B_1$.

Set $\chi_h(n) := e(3T(h,n,n))e(\theta_h n)1_{B_1}(n)$. Applying Proposition 6.1, we obtain

$$\mathbb{E}_{n\in[N]}\chi_{h_1}(n)\chi_{h_2}(n+h_1-h_4)\overline{\chi_{h_3}(n)\chi_{h_4}(n+h_1-h_4)} \gg_{\rho_1,\delta} 1 \tag{9.3}$$

for at least $c_{\rho_1,\delta}N^3$ additive quadruples $(h_1,h_2,h_3,h_4) \in B_1$ with $h_1+h_2 = h_3+h_4$. We have already, in previous sections, extracted 'top order' information from statements like this and our task here is to exploit the additional structure inherent in (9.3), particularly that present in the terms $h_1 - h_4$.

Parametrising these by $h_1 = h$, $h_2 = h+a+b$, $h_3 = h+a$, $h_4 = h+b$ we obtain

$$\mathbb{E}_{n\in[N]}\chi_h(n)\chi_{h+a+b}(n+b)\overline{\chi_{h+a}(n)\chi_{h+b}(n+b)} \gg_{\rho_1,\delta} 1$$

for at least $c_\delta N^3$ triples h, a, b with $h \in t+B'$ and $a, b \in B(S, 3\rho_1, N)$. Substituting in the definition of $\chi_h(n)$, and using the trilinearity of T we obtain

$$\mathbb{E}_n e(6T(a,b,n) + (\theta_h + \theta_{h+a+b} - \theta_{h+a} - \theta_{h+b})n)1_{B_1}(n)1_{B_1}(n+b) \gg_{\rho_1,\delta} 1 \tag{9.4}$$

for at least $c_{\rho_1,\delta}N^3$ triples h, a, b with $h \in B_1$ and $a, b \in B(S, 3\rho_1, N)$. Pigeonholing in h, one sees that there is some fixed h such that this holds for at least $c_{\rho_1,\delta}N^2$ pairs $a, b \in B(S, 3\rho_1, N)$. Let $\varepsilon = \varepsilon(\rho_1, \delta)$ be a small positive quantity to be specified very shortly. By Lemma 9.1 (ii) and the regularity of B_1 we may expand

$$1_{B_1}(n+b) = \int_0^1 \widehat{\psi}_1(\theta)e(\theta(n+b))\,d\theta + \psi_2(n),$$

where $\|\widehat{\psi}_1\|_1 \leqslant C_{\varepsilon,\rho_1,\delta}$ and $\sum_n |\psi_2(n)| \leqslant \varepsilon N$. Choosing ε so that the contribution to (9.4) from $\psi_2(n)$ is negligible, we see using the triangle inequality that there is some $\theta \in [0,1]$ such that

$$\mathbb{E}_{n\in B_1}e(6T(a,b,n) + (\theta_h + \theta_{h+a+b} - \theta_{h+a} - \theta_{h+b} + \theta)n) \gg_{\rho_1,\delta} 1 \tag{9.5}$$

for the same fixed h and many pairs a, b as before. For each a, b write $\phi_{a,b}(n)$ for the phase appearing here, thus

$$\phi_{a,b}(n) := 6T(a,b,n) + (\eta_a + \eta_b' + \eta_{a+b}'')n$$

where $\eta_a := \theta_h - \theta_{h+a} + \theta$, $\eta_b' := \theta_{h+b}$ and $\eta_{a+b}'' := \theta_{h+a+b}$. Equation (9.5) implies that

$$|\mathbb{E}_{n\in B_1}e(\phi_{a,b}(n))| \gg_{\rho_1,\delta} 1.$$

Let $\varepsilon = \varepsilon(\delta, \rho_1)$ be a small positive parameter to be specified later. By Lemma 9.1 (iii) there is some $\rho_2 = \rho_2(\varepsilon, \rho_1, \delta)$ such that we have

$$\|\phi_{a,b}(n)\|_{\mathbb{R}/\mathbb{Z}} \leqslant \varepsilon$$

https://doi.org/10.1017/S0017089510000546 Published online by Cambridge University Press

for all $n \in B_2 := B(S, \rho_2, N)$ and for these same pairs a, b, that is to say for at least $c_{\rho_1,\delta}N^2$ pairs $a, b \in B(S, 3\rho_1)$. Thus

$$6T(a, b, n) = (\eta_a + \eta'_b + \eta''_{a+b})n + O(\varepsilon)$$

for at least $c_{\rho_1,\delta}N^2$ choices of $a, b \in B(S, 3\rho_1, N)$ and for all $n \in B_2$. For at least $c_\delta N^3$ triples a, b, b' we thus have

$$6T(a, b - b', n) = (\eta'_b - \eta'_{b'} + \eta''_{a+b} - \eta''_{a+b'})n + O(\varepsilon)$$

for all $n \in B_2$. Writing $c := a + b + b'$ it follows that

$$6T(c - b - b', b - b', n) = (\eta'_b - \eta'_{b'} + \eta''_{c-b'} - \eta''_{c-b})n + O(\varepsilon)$$

for at least $c_{\rho_1,\delta}N^3$ triples $c, b, b' \in B(S, 9\rho_1, N)$ and for all $n \in B_2$. Fix some c for which this holds for at least $c_{\rho_1,\delta}N^2$ pairs b, b'; then by trilinearity of T we have

$$6(T(b, b', n) - T(b', b, n)) = \kappa_b(n) + \kappa'_{b'}(n) + O(\varepsilon)$$

for all these pairs b, b' and for all $n \in B_2$, where $\kappa_b := \eta'_b n - T(c, b, n) + T(b, b, n) - \eta''_{c-b}n$ and $\kappa'_{b'} = -\eta'_{b'}n + \eta''_{c-b'}n - \psi(c, b', n) + \psi(b', b', n)$. The exact form of these expressions is not relevant, as we shall very shortly see.

Indeed for at least $c_{\rho_1,\delta}N^3$ triples $b_1, b_2, b' \in B(S, 3\rho_1)$ we have

$$6T(b_1 - b_2, b', n) - T(b', b_1 - b_2, n) = \kappa_{b_1}(n) - \kappa_{b_2}(n) + O(\varepsilon)$$

for all $n \in B_2$, and hence for at least $c_{\rho_1,\delta}N^4$ quadruples $b_1, b_2, b'_1, b'_2 \in B(S, 3\rho_1)$ we have

$$6T(b_1 - b_2, b'_1 - b'_2, n) - T(b'_1 - b'_2, b_1 - b_2, n) = O(\varepsilon)$$

for all $n \in B_2$. There are at least $c_{\rho_1,\delta}N^2$ different pairs $x, y \in B(S, 6\rho_1)$ represented as $x = b_1 - b_2, y = b'_1 - b'_2$, and for each of them

$$6(T(x, y, n) - T(y, x, n)) = O(\varepsilon)$$

for all $n \in B_2$. Write $A \subseteq [N]^2$ for the set of these pairs, thus $|A| \geqslant c_\delta N^2$. Let us write $A \oplus A$ for the set of all pairs $(x, y_1 \pm y_2)$ where both (x, y_1) and (x, y_2) lie in A, together with all pairs $(x_1 \pm x_2, y)$ where both (x_1, y) and (x_2, y) lie in A. By bilinearity we see that

$$6(T(x, y, n) - T(y, x, n)) = O(k\varepsilon)$$

for all pairs (x, y) in the k-fold bilinear sumset $A \oplus A \oplus \cdots \oplus A$ and for all $n \in \tilde{B}$.

Now by Lemma B.2 this k-fold bilinear sumset $A' := A \oplus A \cdots \oplus A$ contains a product $P \times P$ provided that $k \geqslant C_\delta$, where P is an arithmetic progression which contains 0 and has length N and common differences $d = O_\delta(1)$. Thus for all triples $x, y, z \in P \cap \tilde{B}$ we have

$$T(x, y, z) - T(y, x, z) = O(k\varepsilon) + \sigma_{x,y,z}, \tag{9.6}$$

where $\sigma_{x,y,z}$ takes values in $\mathbb{Z}/6\mathbb{Z}$.

https://doi.org/10.1017/S0017089510000546 Published online by Cambridge University Press

Recall that we have

$$|\mathbb{E}_{n\in[N]}f_1(n)f_2(n+h)e(3T(h,n,n))e(\theta_h n)1_{B_1}(n)| \gg 1$$

for many $h \in B_1$. By the pigeonhole principle, there are h_1, n_1 such that

$$|\mathbb{E}_{n\in[N]}f_1(n+n_1)f_2(n+h_1+h)e(3T(h_1+h,n_1+n,n_1+n))1_{P\cap B_1}(n)1_{B_1}(n_1+n)| \gg 1$$

for many $h \in P \cap B_1$. Obviously $n_0 \in B(S, 2\rho_1, N)$, and so we may expand $T(h_1 + h, n_1 + n, n_1 + n)$ using trilinearity. Doing this, absorbing the linear terms into $e(\theta_h n)$ using Lemma 3.5 and expanding the cutoff $1_{B_1}(n_1 + n)$ as a Fourier series using Lemma 9.1 (ii), we obtain

$$|\mathbb{E}_{n\in[N]}f_1'(n)f_2'(n+h)e(3T(h,n,n))1_{P\cap B_1}(n)e(\theta_h' n)| \gg 1$$

for may $h \in P \cap B_1$. Here, $f_1'(n) = f_1(n + n_1)e(T(h_1, n, n)) = f(n + n_0 + n_1)e(T(h_0, n_0 + n_1 + n, n_0 + n_1 + n) + T(h_1, n, n))$ whilst $f_2'(n) = f(n + h_0 + h_1 + n_0 + n_1)$. Once again we drop the dashes in what follows for notational convenience.

By the trilinearity of T and the approximate symmetry (9.6) of T in the first two variables, the genuine symmetry in the last two and another application of Lemma 3.5 to handle the terms which are linear in n, it follows that

$$|\mathbb{E}_{n\in[N]}f_1(n)e(-T(n,n,n))f_2(n+h)e(T(n+h,n+h,n+h))1_{P\cap B_1}(n)e(\tilde{\theta}_h n)e(2\sigma_{n,h,n})| \gg 1,$$

provided that ε was chosen sufficiently small in terms of δ. This, recall, is for many $h \in P \cap B_1$.

Now from (9.6) and the smallness of ε we see that $\sigma : (P \cap B_1)^3 \to \mathbb{Z}/6\mathbb{Z}$ is trilinear. Thus $\sigma_{n,h,n}$ is constant as n, h vary over any translate of $Q := 6 \cdot (P \cap B_1) := \{6x : x \in P \cap B_1\}$. Since $P \cap B_1$ may be covered by $O_\delta(1)$ such translates, we may pigeonhole yet again to conclude the existence of h_2, n_2 such that

$$|\mathbb{E}_{n\in[N]}f_1(n+n_2)e(-T(n_2+n,n_2+n,n_2+n))F_2(n+h)1_Q(n+n_2)e(\theta_h' n)| \gg 1.$$

for many h, where

$$\begin{aligned} F_1(n) &:= f_1(n+n_2)e(-T(n_2+n,n_2+n,n_2+n)) \\ &= f(n+n_0+n_1+n_2)e(-T(n_2+n,n_2+n,n_2+n) \\ &\quad + T(h_0,n_0+n_1+n_2+n,n_0+n_1+n_2+n) + T(h_1,n_2+n,n_2+n)) \end{aligned}$$

and F_2 is a 1-bounded function whose precise nature is unimportant. It follows from this and an expansion of $1_Q(n + n_2)$ as a Fourier series that

$$\mathbb{E}_h\|F_1(n)F_2(n+h)\|_{U^2}^4 \gg 1.$$

Expanding out implies that the Gowers inner product $\langle F_1, F_1, F_1, F_1, F_2, F_2, F_2, F_2\rangle_{U^3}$ is $\gg 1$. By the Gowers–Cauchy–Schwarz inequality we see that $\|F_1\|_{U^3} \gg 1$ which, by the inverse theorem for the U^3 norm, implies that

$$\mathbb{E}_{n\in[N]}F_1(n)\Psi(n) \gg 1$$

for some 2-step nilsequence $\Psi(n)$.

https://doi.org/10.1017/S0017089510000546 Published online by Cambridge University Press

Now $F_1(n)$ is equal to $f(n+n_0+n_1+n_2)$ times a variety of bracket terms. By Lemma 3.6, each of those bracket terms is a product of almost nilsequences of degree at most 3. Thus f itself has inner product $\gg_\delta 1$ with a degree 3 almost nilsequence on $[N]$. As we observed in Lemma 3.3, this is enough to establish (at last!) the inverse conjecture for the U^4-norm, that is to say Theorem 1.5. □

Appendix A. Lifting results for nilmanifolds. In this section we establish some slightly technical results concerning the relationship between points on a connected, simply-connected nilpotent Lie group G and points in the nilmanifold G/Γ. These results were necessary in Section 7.

We begin with a folklore result of quantitative linear algebra type.

LEMMA A.1 (Bounded equations have bounded solutions). *Suppose that A is an $m \times n$ matrix and that $b \in \mathbb{C}^m$. Suppose that all of the entries of A are rational numbers of complexity at most M, and that the entries of b are bounded by M. Then if the equation $Ax = b$ has a solution over $\mathbb{C}^n$, it has a solution in which each coordinate is bounded by $O_{M,m,n}(1)$.*

Sketch proof. By removing rows of A if necessary we may assume that the rows of A are linearly independent. One may then augment A to a non-singular $n \times n$ matrix $\tilde{A}$ by adding appropriate basis vectors e_i. Augment b to a vector $\tilde{b} \in \mathbb{C}^n$ by simply adding $n-m$ zeros to b. Then the equation $\tilde{A}\tilde{x} = \tilde{b}$ has a solution given by $\tilde{x} = \tilde{A}^{-1}\tilde{b}$. All entries of $\tilde{x}$ are bounded by $O_{M,m,n}(1)$ by the construction of $\tilde{A}^{-1}$, the key point here being to note that $|\det \tilde{A}|$ is bounded below by $\Omega_{M,m,n}(1)$ since it is a non-zero rational number of complexity $O_{M,m,n}(1)$. □

We record the following special case.

COROLLARY A.2 (Linear lifting). *Suppose that $V \leqslant \mathbb{R}^n$ is a vector subspace given by the vanishing of linear forms over $\mathbb{Z}$ with coefficients of magnitude at most M. Let $\pi : \mathbb{R}^n \to \mathbb{R}^m$ be projection onto the first m coordinates. Suppose that the entries of $x \in \mathbb{R}^m$ are bounded by M, and that $\pi^{-1}(x) \cap V$ is nonempty. Then $\pi^{-1}(x) \cap V$ contains a vector whose entries are bounded by $O_{M,m,n}(1)$.*

Proof. The condition that a vector y lies in $\pi^{-1}(x) \cap V$ may be encoded as $Ay = b$, where this linear system includes the equations $y_1 = x_1, \ldots, y_m = x_m$ and the equations that y must satisfy in order to lie in V. By construction the entries of A are rational numbers of complexity at most M and the entries of b are bounded. The corollary therefore follows from the preceding lemma. □

Using a little Lie theory, this last result has the following further corollary.

COROLLARY A.3. *Suppose that G is a connected, simply-connected nilpotent Lie group and let $\pi : G \to G/[G,G]$ be the natural projection. Suppose that the Lie algebra $\mathbf{g} = \log G$ has a basis $\mathcal{X} = \{X_1, \ldots, X_m, X_{m+1}, \ldots, X_n\}$, where $\pi(\mathcal{X}) := \{\pi(X_1), \ldots, \pi(X_m)\}$ is a basis for $\mathbf{g}/[\mathbf{g},\mathbf{g}] = \log(G/[G,G])$ as a vector space over $\mathbb{R}$. Suppose that H is an M-rational connected subgroup relative to $\mathcal{X}$, and that $\pi(H)$ contains an element $x \in \mathbb{R}^m$ whose entries, written in the basis $\pi(\mathcal{X})$, are bounded by M. Then there is an element $\tilde{x} \in H$ with $\pi(\tilde{x}) = x$ whose entries are bounded by $O_{M,n}(1)$.*

Proof. Let $\mathbf{h} = \log H$ be the Lie algebra of H. By standard Lie theory (see, for example, [**3**]) the exponential/logarithm maps from $\mathbf{g}$ to G and from $\mathbf{h}$ to H are

https://doi.org/10.1017/S0017089510000546 Published online by Cambridge University Press

diffeomorphisms. The result now follows from the preceding corollary upon taking $V = \mathbf{h}$.

This last corollary took place at the level of Lie groups. The actual result we required in Section 7 concerned lifting from nilmanifolds. We state it now. □

PROPOSITION A.4 (Lifting from nilmanifolds). *Let G/Γ be a nilmanifold with Mal'cev basis $\mathcal{X} = \{X_1, \dots, X_m, X_{m+1}, \dots, X_n\}$ and of complexity at most M, and let $H \leqslant G$ be a closed connected M-rational subgroup giving rise to a subnilmanifold $H\Gamma/\Gamma$. Then there is a quantity $\varepsilon_M > 0$ with the following property. Suppose that $H[G, G]\Gamma/[G, G]\Gamma$, identified with the torus $\mathbb{R}^m/\mathbb{Z}^m$ using the Mal'cev basis $\mathcal{X}$, contains an element x whose reduced coordinates (those nearest 0) are all at most ε_M. Let $\psi : G \to G/[G, G]\Gamma$ be the natural projection onto the horizontal torus of G/Γ. Then there is a lift $\tilde{x} \in H$ with coordinates $O_M(1)$ whose first m coordinates are precisely the reduced coordinates of x.*

Proof. The Mal'cev coordinates give a commutative diagram

$$\begin{array}{ccc} G/[G, G] & \longrightarrow & \mathbb{R}^m \\ \downarrow & & \downarrow \\ G/\Gamma[G, G] & \longrightarrow & \mathbb{R}^m/\mathbb{Z}^m. \end{array} \tag{A.1}$$

The inclusion of $H[G, G]/[G, G]$ into $G/[G, G]$ identifies the former with a vector subspace $V \leqslant \mathbb{R}^m$ given by the vanishing of linear forms over $\mathbb{Z}$ with coefficients of magnitude $O_M(1)$, and then $H[G, G]\Gamma/[G, G]\Gamma$ becomes identified with $V\mathbb{Z}^m/\mathbb{Z}^m$. Note that this last object is not in general connected, being a union of a finite number of cosets of a subtorus of $\mathbb{R}^m/\mathbb{Z}^m$. We claim that there is an intermediate lift x' of x to $H[G, G]/[G, G]$ whose coordinates in $\mathbb{R}^m$ are the same as the reduced coordinates of x in $\mathbb{R}^m/\mathbb{Z}^m$. Once this claim is proved we may use the last corollary to lift x' again, under the map $\pi : G \to G/[G, G]$, thereby confirming the proposition.

The claim is a completely abelian statement concerning tori. To prove it, suppose that the linear relations over $\mathbb{Z}$ which define V as a subspace of $\mathbb{R}^m$ are given by $\sum_{j=1}^m k_{ij}x_j = 0, i = 1, \dots, m'$. Suppose that $\varepsilon_M < |k_{ij}|/10m$ (say) and that x, written as $(x_1, \dots, x_m)$ in reduced coordinates, lies in $V\mathbb{Z}^m/\mathbb{Z}^m$. By assumption we have $|x_j| \leqslant \varepsilon_M$ for all j. Then $\sum_{j=1}^m k_{ij}x_j$ is an integer, yet it also has magnitude at most $1/10$. It must therefore vanish, which means that element $x' \in G/[G, G]$ whose coordinates in $\mathbb{R}^m$ are precisely those of x must lie in $H[G, G]/[G, G]$, as claimed. □

Appendix B. Sárközy-type results. In this section we prove a lemma that was used in the course of the so-called symmetry argument in Section 9. It is a familiar principle in additive combinatorics that if one takes some fairly 'dense' set A in an abelian group then the sumsets $2A = A + A, 3A = A + A + A$ become progressively more structured, containing longer and longer progressions and ever larger Bohr sets. See, for example, **[2, 4, 9]**. Sárközy **[24]** was the first to observe that in *very* high-order sumsets kA, one may locate very large amounts of structure indeed. The following rather neat version of his result follows directly from a theorem of Lev ([**23**, Theorem 2′]):

THEOREM B.1 (Lev). *Suppose that $A \subseteq [N]$ is a set of size αN. Then for any $k \geqslant 2/\alpha$ the set $kA - kA := A + \cdots + A - A - \cdots - A$ contains an arithmetic progression $\{0, d, 2d, \dots, (N-1)d\}$ where $d \leqslant 1/\alpha$.*

https://doi.org/10.1017/S0017089510000546 Published online by Cambridge University Press

In Section 9 we required a kind of 'bilinear' version of this. Suppose that $A \subseteq [N]^2$ is a set. Let us write $A \oplus A$ for the set of all pairs $(x, y_1 \pm y_2)$ where both (x, y_1) and (x, y_2) lie in A, together with all pairs $(x_1 \pm x_2, y)$ where both (x_1, y) and (x_2, y) lie in A. The importance of this definition for us lies in the fact that if a bilinear form is approximately annihilated by A then it is also also approximately annihiliated by $A \oplus A$.

PROPOSITION B.2 (Bilinear Sárközy result). *Suppose that $A \subseteq [N]^2$ is a set of size αN^2. Then for $k \geqslant 128/\alpha^3$ the k-fold iterated bilinear sumset $A \oplus A \cdots \oplus A$ contains a product $P \times P'$, where $P = \{0, d, 2d, \ldots, (N-1)d\}$ and $P' = \{0, d', 2d', \ldots, (N-1)d'\}$ with $0 < d, d' \leqslant 4/\alpha^2$.*

Proof. For each $x \in [N]$ write $A_x := \{y \in [N] : (x, y) \in A\}$ for the vertical fibre of A above x. By a simple averaging argument there are at least $\alpha N/2$ values of x for which $|A_x| \geqslant \alpha N/2$. For each such x the vertical sumset $kA_x - kA_x$, where $k \geqslant 4/\alpha$, contains a progression $P = \{0, d_x, 2d_x, \ldots, (N-1)d_x\}$ with $0 < d_x \leqslant 2/\alpha$. By the pigeonhole principle we may pass to a further set $\{A_x : x \in X\}$ of vertical fibres , $|X| \geqslant \alpha^2 N/4$, which all have the same value of d_x, say d. By a further application of Lev's theorem the set $lX - lX$, $l \geqslant 8/\alpha^2$, contains a progression $P' = \{0, d', 2d', \ldots, (N-1)d'\}$ with $0 < d' \leqslant 4/\alpha^2$. □

REMARK . We believe that it ought to be possible to prove a structural result in which only some bounded sum $A \oplus A \oplus \cdots \oplus A$ is involved, where the number of summands does not depend on α (and might, for example, be 16). Such a result would deserve to be called a 'bilinear Bogolyubov theorem' by analogy with Bogolyubov's lemma [**2**]. One would not expect to find a structure as simple and rich as the product $P \times P'$; we expect the relevant structure to be, rather, a 'transverse set', the intersection of sets of the form $\{(x, y) \in [N]^2 : \|\theta xy\|_{\mathbb{R}/\mathbb{Z}} \leqslant \varepsilon\}$.

Appendix C. Structure of approximate homomorphisms. The aim of this appendix is to indicate a proof of Proposition 8.5, whose statement we recall now. As we said before, this result is somehow 'known' without being explicitly given anywhere in the literature. The forthcoming *Barbados lectures* of the first author will give a self-contained treatment of results of this type.

PROPOSITION 8.5 (Approximate homomorphisms). *Let $\delta, \varepsilon \in (0, 1)$ be parameters and suppose that $f_1, f_2, f_3, f_4 : S \to \mathbb{R}/\mathbb{Z}$ are functions defined on some subset $S \subseteq [N]$ such that there are at least δN^3 quadruples $(x_1, x_2, x_3, x_4) \in S^4$ with $x_1 + x_2 = x_3 + x_4$ and $\|f_1(x_1) + f_2(x_2) - f_3(x_3) - f_4(x_4)\|_{\mathbb{R}/\mathbb{Z}} \leqslant \varepsilon$. Then there is a bracket linear phase $\psi : \mathbb{Z} \to \mathbb{R}/\mathbb{Z}$ of complexity $O_\delta(1)$ and a set $S' \subseteq S$, $|S'| \gg_\delta N$, such that $f_1(x) = \psi(x) + O(\varepsilon)$ for all $x \in S'$.*

Proof. We begin with a 'rounding' trick to dispose of the error of ε in the range. Take $N := [1/\varepsilon]$ and for $i = 1, 2, 3, 4$ define $\tilde{f}_i : S \to \mathbb{R}/\mathbb{Z}$ by taking $\tilde{f}_i(x) = r/N$, where r, $0 \leqslant r < N$, is the integer such that r/N is nearest to $f_i(x)$ in $\mathbb{R}/\mathbb{Z}$ (ties being broken arbitrarily). Then of course $\tilde{f}_i(x) = f_i(x) + O(\varepsilon)$ for all $x \in S$ and so

$$\tilde{f}_1(x_1) + \tilde{f}_2(x_2) - \tilde{f}_3(x_3) - \tilde{f}_4(x_4) = O(\varepsilon)$$

for the set of additive quadruples $(x_1, x_2, x_3, x_4) \in S^4$ in the hypothesis of the proposition. The quantity $\|\tilde{f}_1(x_1) + \tilde{f}_2(x_2) - \tilde{f}_3(x_3) - \tilde{f}_4(x_4)\|_{\mathbb{R}/\mathbb{Z}}$ is quantised and

https://doi.org/10.1017/S0017089510000546 Published online by Cambridge University Press

restricted to integer multiples of $1/N$, and there are only $O(1)$ such numbers with magnitude $O(\varepsilon)$. It follows that there is some θ_0 such that $\|\tilde{f}_1(x_1)+\tilde{f}_2(x_2)-\tilde{f}_3(x_3)-\tilde{f}_4'(x_4)\|_{\mathbb{R}/\mathbb{Z}}=0$ for $c\delta N^3$ additive quadruples $(x_1,x_2,x_3,x_4)\in S^4$, where $\tilde{f}_4'(x)=\tilde{f}_4(x)+\theta_0$.

Writing $\Gamma_i:=\{(x,\tilde{f}_i(x)):x\in S\}\subseteq\mathbb{Z}\times\mathbb{R}/\mathbb{Z}$, $i=1,2,3$, and $\Gamma_4':=\{(x,\tilde{f}_4'(x)):x\in S\}\subseteq\mathbb{R}/\mathbb{Z}$ for the 'graphs' of $\tilde{f}_1,\tilde{f}_2,\tilde{f}_3$ and $\tilde{f}_4'$, this means that the additive energy (cf. [**25**, Chapter 2]) $E(\Gamma_1,\Gamma_2,\Gamma_3,\Gamma_4')$ is at least $c\delta N^3$. By [**25**, Corollary 2.10] (or the Cauchy–Schwarz–Gowers inequality) it follows that the additive energy $E(\Gamma_1,\Gamma_1,\Gamma_1,\Gamma_1)$ is at least $c\delta^C N^3$, or in other words that there are $\geqslant c\delta^C N^3$ additive quadruples $(x_1,x_2,x_3,x_4)\in S^4$ for which $\|\tilde{f}_1(x_1)+\tilde{f}_1(x_2)-\tilde{f}_1(x_3)-\tilde{f}_1(x_4)\|_{\mathbb{R}/\mathbb{Z}}=0$.

From this point on we give references to the paper [**11**] of the first two authors, which is reasonably well-adapted to our purposes. Most of the ideas here go back to [**8**, Chapter 7] and to earlier work of Ruzsa. Starting from the assumption that the graph Γ has large additive energy, the key steps are the following[9].

(i) [**11**, Proposition 5.4] Apply the Balog-Szemerédi-Gowers theorem followed by the Plünnecke-Ruzsa inequalities to conclude that there is a set $S_0\subseteq S$, $|S_0|\geqslant c\delta^C N$, such that the graph $\Gamma:=\{(x,\tilde{f}_1(x)):x\in S_0\}$ satisfies an iterative sumset estimate $|k\Gamma-l\Gamma|\ll_{k,l}N$ for all integers $k,l\geqslant 1$.

(ii) [**11**, Proposition 9.1] The function $\tilde{f}_1$ correlates with a function which is locally linear on a Bohr set. This means that there are is a Bohr set $B=B(\Theta,\rho,N)$ with $\Theta=\{\theta_1,\dots,\theta_d\}\subseteq\mathbb{R}/\mathbb{Z}$, $d=O_\delta(1)$ and $\rho\gg_\delta 1$ together with a function $\phi:B\to\mathbb{R}/\mathbb{Z}$ satisfying $\phi(x+y)=\phi(x)+\phi(y)$ whenever $x,y,x+y\in B(\Theta,\rho,N)$, as well as some $x_0\in[N]$ and some $\theta_0\in\mathbb{R}/\mathbb{Z}$ such that $\tilde{f}(x+x_0)=\theta_0+\phi(x)$ for $\gg_\delta N$ values of $x\in(S_0-x_0)\cap B$. The appropriate definitions here are given in full in [**11**] and are also recalled in Section 9 of the present paper.

(iii) Apply some geometry of numbers to conclude that any such linear function ϕ has the form $\phi(x)=\alpha_1\{\theta_1x\}+\cdots+\alpha_d\{\theta_dx\}+\theta x$ on some multidimensional progression $P\subseteq B$ with $|P|\gg_\delta N$. The proof of this is very similar to, but easier than, that of [**11**, Proposition 10.8], where an analogous statement is established for locally *quadratic* phase functions on Bohr sets.

It follows from all of this that we have

$$\tilde{f}_1(x)=\theta_0+\alpha_1\{\theta_1(x-x_0)\}+\cdots+\alpha_d\{\theta_d(x-x_0)\}+\theta(x-x_0)$$

for all x in some set $S_1\subseteq S_0$, $|S_1|\gg_\delta N$.

Now we have $\{\theta_j(x-x_0)\}=\{\theta_jx\}-\{\theta_jx_0\}+\tau_{j,x}$, where $\tau_{j,x}$ takes values in $\{-1,0,1\}$. By the pigeonhole principle we may pass to a further subset $S_2\subseteq S_1$ with $|S_2|\gg_\delta N$ such that, for all $x\in S_2$, each of the $\tau_{j,x}$ is independent of x.

Take $S':=S_2$. Then for $x\in S'$ we have

$$\tilde{f}_1(x)=\theta_0'+\alpha_1\{\theta_1x\}+\cdots+\alpha_d\{\theta_dx\}+\theta x,$$

a bracket linear form of complexity $d=O_\delta(1)$. Recalling that $\tilde{f}_1(x)=f_1(x)+O(\varepsilon)$, the result follows. □

[9]Strictly speaking, the tools we are applying here only apply to groups rather than to intervals such as $[N]$. However, this can be easily addressed by temporarily embedding $[N]$ in, say, $\mathbb{Z}/10N\mathbb{Z}$; we omit the details.

https://doi.org/10.1017/S0017089510000546 Published online by Cambridge University Press

REMARK. The rounding trick we used to remove the ε errors was a slightly dirty one but makes the argument quite short given known results. It would probably be possible, and more natural in some moral sense, to run through the Balog-Szemerédi-Gowers and Freiman arguments carrying an $O(\varepsilon)$ error throughout.

Appendix D. Some diophantine results. This section recalls some well-known results from Diophantine approximation which, in the context of this paper, may be naturally viewed as distributional results for abelian (1-step) nilsequences. We will use them repeatedly in the next section. Furthermore Lemma D.2 below was crucial in Section 7, and Lemma D.1 was required at the end of Section 8.

LEMMA D.1. *Let $d \geqslant 1$ be an integer, let $\varepsilon \in (0, 1/2)$ be a parameter, and suppose that $\psi(n) = \alpha_d n^d + \cdots + \alpha_0$ is a polynomial of degree d such that $(\psi(n)(\mathrm{mod}\ 1))_{n\in[N]}$ is not ε-equidistributed on $\mathbb{R}/\mathbb{Z}$. Then for all $i = d, d-1, \ldots, 1$ there are coprime integers a_i, q_i, $q_i \leqslant \varepsilon^{-C_d}$, such that*

$$\alpha_i = \frac{a_i}{q_i} + O\left(\frac{\varepsilon^{-C_d}}{N^i}\right).$$

Proof. This is actually a special case of the Quantitative Leibman Dichotomy, Theorem 4.1, although this is a somewhat misleading statement to make since it is also a crucial ingredient in the *proof* of that result. It is proven using Weyl's criterion for equidistribution and Weyl's inequality (see, for example, [**27**]), and indeed the statement that the lead coefficient α_d is close to rational is essentially equivalent to that inequality. The other coefficients $\alpha_{d-1}, \alpha_{d-2} \ldots$ may be shown to be almost rational iteratively; the argument is given in detail in [**14**, Section 4]. □

Secondly we recall a quantitative version of Kronecker's theorem, phrased in language appropriate to Section 7. Once again this is a special case of the Quantitative Leibman Dichotomy, and once again it is very well-known.

LEMMA D.2. *Let $d \geqslant 1$ be an integer, let $\varepsilon \in (0, 1/2)$ be a parameter, and let $\alpha_1, \ldots, \alpha_d \in \mathbb{R}/\mathbb{Z}$ be frequencies. Suppose that $((\alpha_1 n, \ldots, \alpha_d n)(\mathrm{mod}\ 1))_{n\in[N]}$ fails to be ε-equidistributed in the torus $(\mathbb{R}/\mathbb{Z})^d$. Then the set $\{\alpha_1, \ldots, \alpha_d\}$ satisfies an ε^{-C_d}-linear relation up to ε^{-C_d}/N (that is, there are integers $m_1, \ldots, m_d$, not all zero, with $|m_i| \leqslant \varepsilon^{-C_d}$ for all i and $\|m_1\alpha_1 + \cdots + m_d\alpha_d\|_{\mathbb{R}/\mathbb{Z}} \leqslant \varepsilon^{-C_d}/N$).*

Proof. This is discussed in detail in [**14**, Section 3]. Here is a very rough sketch: if the sequence is not ε-equidistributed, there is some Lipschitz function $F : (\mathbb{R}/\mathbb{Z})^d \to \mathbb{C}$ with

$$|\mathbb{E}_{n\in[N]} F(\alpha_1 n, \ldots, \alpha_d n) - \int_{(\mathbb{R}/\mathbb{Z})^d} F(\theta) d\theta| \geqslant \varepsilon \|F\|_{\mathrm{Lip}}.$$

Lipschitz functions are well-approximated in L^∞ by their Fourier series; exanding F into such a series, it follows that some exponential sum

$$\mathbb{E}_{n\in[N]} e(\vec{m} \cdot \vec{\alpha} n)$$

has modulus at least ε^{C_d}, where $\vec{m} = (m_1, \ldots, m_d)$ and $|m_i| \leqslant \varepsilon^{-C_d}$. The lemma now follows with an application of the formula for the sum of a geometric series. □

https://doi.org/10.1017/S0017089510000546 Published online by Cambridge University Press

Appendix E. Almost nilsequences. The aim of this section is to establish Lemmas 3.5 and 3.6, which asserted that various objects – chiefly bracket polynomials – are 1-, 2- and 3-step almost nilsequences.

LEMMA 3.5. *Suppose that* $\alpha, \beta \in [0, 1]$ *and that* $M > 1$ *is a complexity parameter. The following are all examples of almost nilsequences of degree* 1 *and complexity* $O_M(1)$*:*

(i) *the set of* 1*-step Lipschitz nilsequences of complexity at most* M*;*
(ii) *the set of characteristic functions* 1_P*, where* $P \subseteq [N]$ *is a progression of length at least* N/M*;*
(iii) *the set of functions of the form* $n \mapsto e(\alpha\{\beta n\})$*, with* $\alpha \in \mathbb{R}$ *and* $\beta \in \mathbb{R}/\mathbb{Z}$*;*
(iv) *the set of functions of the form* $n \mapsto e(\{\alpha n\}\{\beta n\})$*, with* $\alpha, \beta \in \mathbb{R}/\mathbb{Z}$*;*
(v) *the set of functions of the form* $n \mapsto e(\alpha n \lfloor \beta n \rfloor)$*, where* $\|\beta\|_{\mathbb{R}/\mathbb{Z}} \leqslant M/N$.

Proof. (i) is trivial.

To prove (ii) we first note that $1_P(n)$ can be expressed as the product of $1_I(n)$ and $1_{n \equiv a (\mathrm{mod}\, q)}$, where $I \subseteq N/M$ is an interval and $q \leqslant M$. The second object is in fact a 1-step nilsequence $F(g(n)\Gamma)$ on $\mathbb{R}/\mathbb{Z}$, the polynomial sequence $g : \mathbb{Z} \to \mathbb{R}$ being $g(n) = n/q$ and the function $F : \mathbb{R}/\mathbb{Z} \to [0, 1]$ being Lipschitz, equal to 1 at a/q and supported within $1/10q$ (say) of a/q. The first object, $1_I(n)$, is not quite a genuine 1-step nilsequence. However let us observe that any function $\psi : [N] \to \mathbb{C}$ with Lipschitz constant $O(1/N)$ is a genuine 1-step nilsequence; indeed we have $\psi(n) = F(g(n)\Gamma)$ on $\mathbb{R}/\mathbb{Z}$, where $g(n) = n/2N$ and $F : \mathbb{R}/\mathbb{Z} \to \mathbb{C}$ is defined by setting $F(n/2N) := \psi(n)$ for $n \in [N]$ and by Lipschitz extension elsewhere. Now simply note that 1_I may be approximated arbitrarily closely, in $L^1[N]$, by functions ψ of this type. Specifically, we may take a sequence of Lipschitz 'tent' functions ψ which equal 1 on I and are zero at points distance more than εN from I. The claim now follows from Lemma 3.2.

To establish (iii) we first note that if $\alpha \equiv \alpha' (\mathrm{mod}\, 1)$ then $\alpha\{\beta n\} \equiv \alpha'\{\beta n\} + (\alpha - \alpha')\beta n (\mathrm{mod}\, 1)$, and so we may assume that $0 \leqslant \alpha \leqslant 1$. Let $\varepsilon > 0$ be arbitrary and define $F : \mathbb{R}/\mathbb{Z} \to \mathbb{C}$ by $F(x) = e(\alpha\{x\})$ and divide into two cases: either $(\beta n (\mathrm{mod}\, 1))_{n \in [N]}$ is $\varepsilon/10$-equidistributed on $\mathbb{R}/\mathbb{Z}$, or it is not. In the former case we take a 100ε-Lipschitz function $\tilde{F}$ which agrees with F outside of the set $\{x \in \mathbb{R}/\mathbb{Z} : \|x\|_{\mathbb{R}/\mathbb{Z}} \leqslant \varepsilon/10\}$ and is bounded by 1 elsewhere. By the assumed equidistribution we obviously have $e(\alpha\{\beta n\}) = F(\beta n) = \tilde{F}(\beta n)$ for all except at most $\varepsilon N/2$ values of n. The result is then immediate.

If, on the other hand, the sequence $(\beta n (\mathrm{mod}\, 1))_{n \in [N]}$ fails to be $\varepsilon/10$-equidistributed then by Lemma D.1 with $d = 1$ there is an integer $q \leqslant \varepsilon^{-C}$ and an $a \in \mathbb{Z}$ such that $\|\beta - \frac{a}{q}\|_{\mathbb{R}/\mathbb{Z}} \ll \varepsilon^{-C}/N$. This in turn means that we may divide $[N]$ into progressions $P_1 \cup \cdots \cup P_m$, $m \ll \varepsilon^{-C}$, on which $n \mapsto e(\alpha\{\beta n\})$ varies by at most $\varepsilon/100$. Since (by part (ii)) functions which are constant on progressions are almost 1-step nilsequences, the result follows (using Lemma 3.2 as necessary).

To prove (iv) we use a trick. The function $(x, y) \mapsto e(xy)$ on the square $[0, 1]^2$ may be smoothly extended to a periodic function on $[0, 2]^2$. By Fourier analysis (cf. [**12**, Lemma A.9]) it may then be uniformly approximated to any desired accuracy ε by a linear combination of frequencies $e((kx + ly)/2)$, $k, l \in \mathbb{Z}$. Thus $e(\{\alpha n\}\{\beta n\})$ may be approximated uniformly by a linear combination of functions of the form $e(k\{\alpha n\}/2)e(l\{\beta n\}/2)$. But such functions are almost 1-step nilsequences by (iii), and the claim follows from Lemma 3.2.

Finally we turn to (v). The condition that $\|\beta\|_{\mathbb{R}/\mathbb{Z}} \leqslant M/N$ means that we may divide $[N]$ into subprogressions (in fact subintervals) $P_1 \cup \cdots \cup P_m$, $m = O_M(1)$, such

https://doi.org/10.1017/S0017089510000546 Published online by Cambridge University Press

that $\lfloor \beta n \rfloor$ is equal to some constant c_j for $n \in P_j$. The result then follows from (ii) and Lemma 3.2. □

Now we turn to higher degree bracket polynomial phases.

LEMMA 3.6. *Suppose that* $\alpha, \beta, \gamma \in [0, 1]$. *Then the following are all examples of almost nilsequences of degree* $s \geqslant 2$ *and complexity* $O(1)$*:*

(i) $n \mapsto e(\lfloor \alpha n \rfloor \beta n)$, *of degree* 2*;*
(ii) $n \mapsto e(\lfloor \alpha n \rfloor \beta n^2)$, *of degree* 3*;*
(iii) $n \mapsto e(\lfloor \alpha n \rfloor \lfloor \beta n \rfloor \gamma n)$, *of degree* 3.

Proof. The proofs of all three parts are somewhat similar and proceed along the following lines: each object may be exhibited in a fairly obvious way as a nilsequence $F(g(n)\Gamma)$, where F is, however, only piecewise Lipschitz. If the sequence $(g(n)\Gamma)_{n \in [N]}$ is highly equidistributed then it spends sufficient time away from singularities for one to be able to approximate by $\tilde{F}(g(n)\Gamma)$, where $\tilde{F}$ is genuinely Lipschitz. If not then there must be an approximate rational relation between the horizontal frequencies of $g(n)$ (that is, the frequencies occurring in the projection to $G/[G, G]$). This may then be used to approximate the object in question by objects of lower complexity.

To exhibit these arguments as part of a more general theory is not a particularly easy matter and involves a more conceptual understanding of bracket identities such as those in Lemma 5.5 and others such as (E.3) below. The required theory is implicit in the work of Leibman [**22**] and will be introduced properly in our longer paper to come.

In this paper we can proceed in an *ad hoc* and slightly calculational way, taking advantage of one or two simplifications specific to the U^4 (3-step) case. In a sense, however, these calculations also serve as motivation for the longer paper to come. We begin by recalling the constructions of Section 5 leading up to (5.2). Specialising to the free 2-step nilpotent group on two generators (essentially the Heisenberg group) we have

$$F_{[1,2]}(g(n)\Gamma) = e(\lfloor \alpha n \rfloor \beta n)$$

and

$$F_{[1,2]}(g'(n)\Gamma) = e(\lfloor \alpha n \rfloor \beta n^2)$$

where $F_{[1,2]} : G/\Gamma \to \mathbb{C}$ is the basic coordinate function introduced in Definition 5.3 and $g, g' : \mathbb{Z} \to G$ are polynomial sequences of degree 2 and 3 respectively given in coordinates by $g(n) = (\alpha n, -\beta n, 0)$, $g'(n) = (\alpha n, -\beta n^2, 0)$. Only the first two coordinates (corresponding to the horizontal torus $G/[G, G]$) are really important.

The discontinuities of $F_{[1,2]}$ are very manageable: the key point, already exploited in Section 7, is that for any $\varepsilon > 0$ there is are ε^{-C}-Lipschitz functions $\tilde{F} : G/\Gamma \to \mathbb{C}$ and $\Psi : G/\Gamma \to [0, 1]$ such that $\int_{G/\Gamma} \Psi \leqslant \varepsilon$ and $|F(x) - \tilde{F}(x)| \leqslant \Psi(x)$ pointwise.

Fix $\varepsilon > 0$. Let us consider statement (i), for which we consider the sequence $(g(n)\Gamma)_{n \in [N]}$. If it is ε-equidistributed then, by the preceding, $e(\{\alpha n\}\beta n) = F(g(n)\Gamma)$ and $\tilde{F}(g(n)\Gamma)$ are within 2ε in $L^1[N]$. If this is not the case then, by the Quantitative Leibman Dichotomy (Theorem 4.1) there must be some $O(\varepsilon^{-C})$-linear relation, up to $O(\varepsilon^{-C}/N)$, between α and β. The rest of the argument in this case is essentially identical to that at the very end of Section 7; we may find some γ such that $\alpha = q_1\gamma + O(\varepsilon^{-C}/N)$ and $\beta = q_2\gamma + O(\varepsilon^{-C}/N)$, where q_1, q_2 are integers with magnitude at most ε^{-C}. Substituting

https://doi.org/10.1017/S0017089510000546 Published online by Cambridge University Press

into $e(\alpha n\lfloor\beta n\rfloor)$ and making repeated use of the bracket identities of Lemma 5.5 as well as Lemma 3.5, one sees that in this case $e(\alpha n\lfloor\beta n\rfloor)$ lies within ε in $L^1[N]$ of a degree 2 nilsequence (of step 1) of complexity $O_\varepsilon(1)$. Thus in either case we have approximated $e(\alpha n\lfloor\beta n\rfloor)$ within $O(\varepsilon)$ by a degree 2 polynomial nilsequence of complexity $O_\varepsilon(1)$, thereby completing the proof of (i).

The analysis of (ii) is similar but, obviously, involves consideration of the sequence $(g'(n)\Gamma)_{n\in[N]}$ instead. If the sequence $(g'(n)\Gamma)_{n\in[N]}$ is ε-equidistributed then we are done, as before. If not, the Quantitative Leibman Dichotomy implies that either $\alpha = \frac{a_1}{q_1} + O(\varepsilon^{-C}/N)$ or else $\beta = \frac{a_2}{q_2} + O(\varepsilon^{-C}/N^2)$. In the first case we may then partition $[N]$ into progressions $P_1 \cup \cdots \cup P_m$, $m = O_\varepsilon(1)$, on which $\lfloor\alpha n\rfloor$ is constant and then apply Lemma 3.5 (ii) to approximate $e(\lfloor\alpha n\rfloor\beta n^2)$ within $O(\varepsilon)$ by a degree 2 polynomial nilsequence of complexity $O_\varepsilon(1)$. In the second case we first apply the bracket identity (5.3) to write

$$e(\lfloor\alpha n\rfloor\beta n^2) = e(\alpha\beta n^3)e(-\alpha n[\beta n^2])e(-\{\alpha n\}\{\beta n^2\}). \tag{E.1}$$

The first term here is already a degree 3 polynomial nilsequence of complexity $O(1)$. In the second term we may partition $[N]$ into progressions $P_1 \cup \cdots \cup P_m$, $m = O_\varepsilon(1)$, on which $[\beta n^2]$ is constant and then apply Lemma 3.5 (ii) to approximate arbitrarily closely by a degree 1, nilsequence. The third term, $e(-\{\alpha n\}\{\beta n^2\})$, may be handled using the same trick as in the proof of Lemma 3.5 (iv). This reduces matters to handling $e(\theta\{\theta' n\})$ (already known to be a degree 1 almost nilsequence by Lemma 3.5 (iii)) and $e(\theta\{\theta' n^2\})$. By an argument almost identical to that used in the proof of Lemma 3.5 (iii), only using Lemma D.1 with $d = 2$ instead, this second object may be shown to be an degree 2 almost nilsequence. Using Lemma 3.2 to put everything together, we obtain the claim.

We turn now to the proof of (iii), which is important in the sense that it is the only place in our paper where a 3-step nilmanifold is actually constructed!

Specifically, we let $\mathfrak{g}$ be the free 3-step Lie algebra generated by three generators e_1, e_2, e_3, or equivalently

$$G := \left\{e_1^{t_1}e_2^{t_2}e_3^{t_3}e_{21}^{t_{21}}e_{211}^{t_{211}}e_{31}^{t_{31}}e_{311}^{t_{311}}e_{32}^{t_{32}}e_{322}^{t_{322}}e_{212}^{t_{212}}e_{312}^{t_{312}}e_{213}^{t_{213}}e_{313}^{t_{313}}e_{323}^{t_{323}} : t_i, t_{ij}, t_{ijk} \in \mathbb{R}, 1 \leqslant i, j, k, \leqslant 3\right\}.$$

subject to the relations $e_i^{-1}e_j^{-1}e_ie_j = [e_i, e_j] = e_{[i,j]}$ for $1 \leqslant j < i \leqslant 3$, $[[e_i, e_j], e_k] = e_{ijk}$, and the Jacobi relation $[[e_i, e_j], e_k][[e_j, e_k], e_i][[e_k, e_i], e_j] = 1$. Inside G we take the standard lattice

$$\Gamma := \left\{e_1^{n_1}e_2^{n_2}e_3^{n_3}e_{21}^{n_{21}}e_{211}^{n_{211}}e_{31}^{n_{31}}e_{311}^{n_{311}}e_{32}^{n_{32}}e_{322}^{n_{322}}e_{212}^{n_{212}}e_{312}^{n_{312}}e_{213}^{n_{213}}e_{313}^{n_{313}}e_{323}^{n_{323}} : n_i, n_{ij}, n_{ijk} \in \mathbb{Z}, 1 \leqslant i, j, k, \leqslant 3\right\}.$$

Then G/Γ is the free 3-step nilmanifold on 3 generators. We take $G_\bullet$ to be the lower central series on G

We abbreviate $e_1^{t_1}\ldots e_{323}^{t_{323}}$ as $(t_1, \ldots, t_{323})$. A computation yields the multiplication law

$$(t_1, \ldots, t_{323}) \star (u_1, \ldots, u_{323}) = (s_1, \ldots, s_{323})$$

where $s_i = t_i + u_i$ for $i = 1, 2, 3$, $s_{ij} = t_{ij} + u_{ij} + t_iu_j$ for $1 \leqslant j < i \leqslant 3$, and $s_{312} = t_{312} + u_{312} + t_{32}u_1 + t_{31}u_2 + t_3u_1u_2$; we will ignore the other coordinates, as they will not be needed in this calculation.

https://doi.org/10.1017/S0017089510000546 Published online by Cambridge University Press

Using this law, we see that for any real numbers $t_1, \dots, t_{323}$, one has

$$(t_1, \dots, t_{323})\Gamma = (s_1, \dots, s_{323})\Gamma$$

where

$$\begin{aligned} s_i &:= \{t_i\} \text{ for } i = 1, 2, 3; \\ s_{ij} &:= \{t_{ij} - t_i[t_j]\} \text{ for } 1 \leqslant j < i \leqslant 3; \\ s_{312} &:= \{t_{312} - t_{32}[t_1] - t_{31}[t_2] + t_3[t_1][t_2]\}, \end{aligned} \tag{E.2}$$

and with the other coordinates $s_{ijk} \in [0, 1]$ being explicitly computable, but not relevant for this discussion. Thus if we let

$$g(n) := e_1^{\alpha n} e_2^{\beta n} e_3^{\gamma n}$$

and let $F : G/\Gamma \to \mathbb{C}$ be the 3-step basic coordinate function function

$$F((s_1, \dots, s_{323})\Gamma) := e(s_{312})$$

for $s_1, \dots, s_{323} \in [0, 1]$, then one sees that $e(\lfloor \alpha n \rfloor \lfloor \beta n \rfloor \gamma n)$ is equal to $F(g(n)\Gamma)$ times objects already known to be almost nilsequences by earlier parts.

This concludes the argument unless $(g(n)\Gamma)_{n\in[N]}$ spends too much time near the singularities of F, which are at the points $s_j = 0$ and $s_j = 1$, $j = 1, 2, 3$. There will be no problem unless[10] one of the sequences $(\alpha n (\mathrm{mod}\, 1))_{n\in[N]}$, $(\beta n (\mathrm{mod}\, 1))_{n\in[N]}$, $(\gamma n (\mathrm{mod}\, 1))_{n\in[N]}$ fails to be ε-equidistributed. If $(\alpha n (\mathrm{mod}\, 1))_{n\in[N]}$ is not ε-equidistributed then, by the now-familiar application of Lemma D.1 with $d = 1$, we may partition $[N]$ as a union $P_1 \cup \cdots \cup P_m$ of at most ε^{-C} progressions such that $\lfloor \alpha n \rfloor$ is constant on P_i. We may then conclude using part (i) and Lemma 3.5 (ii). An identical argument works if $(\beta n (\mathrm{mod}\, 1))_{n\in[N]}$ fails to be ε-equidistributed.

The final case is when $(\gamma n (\mathrm{mod}\, 1))_{n\in[N]}$ fails to be ε-equidistributed. In this case we note that

$$\{\alpha n\}\{\beta n\}\{\gamma n\} = (\alpha n - \lfloor \alpha n \rfloor)(\beta n - \lfloor \beta n \rfloor)(\gamma n - \lfloor \gamma n \rfloor) \tag{E.3}$$

so that

$$\begin{aligned} e(\lfloor \alpha n \rfloor \lfloor \beta n \rfloor \gamma n) =& e(\{\alpha n\}\{\beta n\}\{\gamma n\}) e(-\lfloor \alpha n \rfloor \beta n \lfloor \gamma n \rfloor) \\ & \times e(-\alpha n \lfloor \beta n \rfloor \lfloor \gamma n \rfloor) e(\alpha \beta n^2 \lfloor \gamma n \rfloor) e(\alpha \gamma n^2 \lfloor \beta n \rfloor) e(\beta \gamma n^2 \lfloor \alpha n \rfloor). \end{aligned}$$

Each of the terms on the right except the first can be handled using part (ii) or by those instances of part (iii) already established. To deal with the first term $e(\{\alpha n\}\{\beta n\}\{\gamma n\})$ one may proceed exactly as in Lemma 3.5(iv) to show that this is in fact an degree 1 almost nilsequence. Applying Lemma 3.2 to collect terms, we obtain the claim. □

The main business of the paper is now concluded. The remaining two appendices were promised in the introduction but are not necssary for the proof of Theorem 1.5.

[10] This observation, which is stronger than saying that the abelianisation $((\alpha n, \beta n, \gamma n)(\mathrm{mod}\, 1))_{n\in[N]} \subseteq (\mathbb{R}/\mathbb{Z})^3$ is not equidistributed, is somewhat specific to the three-step situation we are working with and represents something of a simplification over the argument required in general.

https://doi.org/10.1017/S0017089510000546 Published online by Cambridge University Press

Appendix F. The strong inverse conjecture. We have shown, in Theorem 1.5, that a 1-bounded function $f : [N] \to \mathbb{C}$ with $\|f\|_{U^4} \geqslant \delta$ correlates with a degree 3 polynomial nilsequence $F(g(n)\Gamma)$. As we remarked after the statement of Conjecture 1.3, this does not quite establish the result used in (for example) **[13]**, where correlation with a nilsequence $F(g^n x\Gamma)$ was used. In this section we shall refer to *linear nilsequences* to distinguish objects of this last type from more general polynomial nilsequences.

In this section we indicate, very briefly, how our arguments may be modified to obtain this apparently stronger statement. In the longer paper to come we will provide a quite general proof that Conjecture 1.3 implies this strong variant. Let us recall once more, however, our view that this is the 'wrong' perspective and that **[13]** works, with rather minimal changes, in the context of polynomial nilsequences.

We need only show that large U^4-norm entails correlation with *almost* linear nilsequences, defined in exact analogy with Definition 3.1. We already have correlation with almost polynomial sequences, so it will suffice to show that the almost polynomial sequences used in the paper are also almost linear sequences of the same degree.

Clearly, any degree 1 almost nilsequence is already an almost linear 1-step nilsequence, and an inspection of the previous appendix shows that $e(\lfloor \alpha n \rfloor \beta n)$ is an almost linear 2-step nilsequence, modulo a quadratic phase $e(\gamma n^2)$, and similarly $e(\lfloor \alpha n \rfloor \lfloor \beta n \rfloor \gamma n)$ is an almost linear 3-step nilsequence modulo phases such as $e(\gamma n^3)$ and $e(\lfloor \gamma n \rfloor \delta n^2)$. As Lemma 3.2 is clearly also valid for almost linear nilsequences, one only needs to verify three remaining claims, for any real numbers α, β:

- $e(\alpha n^2)$ is an almost linear 2-step nilsequence;
- $e(\alpha n^3)$ is an almost linear 3-step nilsequence; and
- $e(\lfloor \alpha n \rfloor \beta n^2)$ is an almost linear 3-step nilsequence.

We look first at $e(\alpha n^2)$ and consider once again the 2-step nilpotent group on 2 generators (Heisenberg group); looking all the way back to (5.2) and taking $g = (2\alpha, 1, 0)$ one may compute that $F_{[1,2]}(g^n\Gamma) = e(\alpha n^2 + \theta n)$ for some $\theta \in \mathbb{R}/\mathbb{Z}$. Now $F_{[1,2]}(t_1, t_2, t_{12})$ is discontinuous when $t_{12} = 0$ or 1. If we wish to approximate $e(\alpha n^2 + \theta n)$ within ε (in $L^1[N]$) by a *Lipschitz* linear nilsequence, we must show (for example) that there are no more than $10\varepsilon N$ values of $n \in [N]$ for which $\alpha n^2 + \theta n (\mathrm{mod}\ 1)$ is within ε of 0. But if this is not the case then, by Lemma D.1, we have $\alpha = a/q + O(\varepsilon^{-C}/N^2)$, at which point we can split $[N]$ into ε^{-C} progressions on which $e(\alpha n^2 + \theta n)$ is within $O(\varepsilon)$ of a linear phase. One may then proceed using Lemma 3.5.

Now we turn to the 3-step objects $e(\alpha n^3)$ and $e(\lfloor \alpha n \rfloor \beta n^2)$, which require some slightly more careful calculations on the free 3-step nilmanifold are required. With the notation for the free 3-step nilpotent Lie group as in the preceding section, let $g = e_1^\alpha e_2^\beta e_3^\gamma$. Then one can check that

$$g^n = e_1^{n\alpha} e_2^{n\beta} e_3^{n\gamma} e_{21}^{\binom{n}{2}\alpha\beta} e_{31}^{\binom{n}{2}\alpha\gamma} e_{32}^{\binom{n}{2}\beta\gamma} \cdots e_{312}^{\alpha\beta\gamma(2\binom{n}{3}+\binom{n}{2})} \cdots$$

and hence one may compute (cf. (E.2))

$$F_{312}(g^n\Gamma) = e\big(\alpha\beta\gamma\left(2\tbinom{n}{3} + \tbinom{n}{2}\right) - \tbinom{n}{2}\beta\gamma\lfloor\alpha n\rfloor - \tbinom{n}{2}\alpha\gamma\lfloor\beta n\rfloor + n\gamma[\alpha n]\lfloor\beta n\rfloor\big) \qquad \text{(F.1)}$$

Taking $\beta = \gamma = 1$ and replacing α by 6α gives $F_{312}(g^n\Gamma) = e(\alpha n^3 + q(n))$ for some quadratic q. The discontinuities of F_{312} may be handled as for $F_{[1,2]}$ above, and so we see that $e(\alpha n^3 + q(n))$ is an almost 3-step linear nilsequence for some quadratic q. Since we can already obtain pure quadratic and linear phases as almost linear nilsequences of step less than 3, it follows that $e(\alpha n^3)$ itself is an almost 3-step linear nilsequence.

https://doi.org/10.1017/S0017089510000546 Published online by Cambridge University Press

Next, we take $\beta = 1$ and replace γ by -2γ. Taking into account objects already known to be almost linear nilsequences, we have now obtained $e(\gamma n^2[n\alpha])$ as a 3-step almost linear nilsequence. Applying (E.1), we see that to obtain the desired object $e([n\alpha]\gamma n^2)$ it suffices to examine $e(\{\alpha n\}\{\gamma n^2\})$. By the trick used in the proof of Lemma 3.5 (iv), it suffices in turn to handle $e(\theta\{\theta' n\})$ and $e(\theta\{\theta' n^2\})$. The first of these is an almost 1-step (linear) nilsequence by Lemma 3.5 (iii). To handle the second, proceed in the same way as in the proof of Lemma 3.5 (iii) but in the obvious places substitute the fact (established above of course) that pure quadratic phases are 2-step linear nilsequences, together with the case $d = 2$ of Lemma D.1.

Appendix G. Necessity of the inverse conjectures. In this appendix we sketch a rather short proof of Proposition 1.4, which asserted that functions which correlate with a degree s polynomial nilsequence must have large U^{s+1}-norm. Since *linear* nilsequences are merely special cases of polynomial ones, this kind of argument could substitute in, for example, [**13**, Sec. 10], where a rather more complicated approach was taken.

PROPOSITION 1.4. *Suppose that $f : [N] \to \mathbb{C}$ is a 1-bounded function, that $(F(g(n)\Gamma))_{n\in\mathbb{Z}}$ is a polynomial nilsequence of degree s and complexity $O_\delta(1)$, and that*

$$|\mathbb{E}_{n\in[N]} f(n)\overline{F(g(n)\Gamma)}| \geqslant \delta.$$

Then $\|f\|_{U^{s+1}} \gg_\delta 1$.

Sketch proof. The argument is only a sketch in that we do not address such issues as the complexity of the nilsequences involved. We leave this as a (not particularly interesting) exercise to the reader, most of the details of which may be found in [**14**] where these complexity issues are discussed in detail. We proceed by induction on s, the claim being obvious when $s = 0$. Let $f : [N] \to \mathbb{C}$ be a 1-bounded function, and let $g : \mathbb{Z} \to G$ be a polynomial sequence of degree s adapted to the filtration $G_\bullet$. Let $F(g(n)\Gamma)$ be a polynomial nilsequence of complexity $O_\delta(1)$. Assume that

$$|\mathbb{E}_{n\in[N]} f(n)F(g(n)\Gamma)| \gg_\delta 1. \tag{G.1}$$

By decompositing F into vertical characters as in [**14**, Lemma 3.7], we may assume that F has a vertical frequency: that is, there is some non-trivial character $\xi : G_{(s)}/G_{(s)} \cap \Gamma$ such that

$$F(g_s x) = e(\xi(g_s))F(x)$$

for all $g_s \in G_{(s)}$ and $x \in G/\Gamma$.

By taking the modulus squared of (G.1) and making the substitution $n' = n + h$ we see that

$$\mathbb{E}_{n\in[N]} \Delta_h f(n)F(g(n+h)\Gamma)\overline{F(g(n)\Gamma)} \gg_\delta 1$$

for $\gg_\delta N$ values of $h \in [N]$.

However for each fixed h the 'derivative' $n \mapsto F(g(n+h)\Gamma)\overline{F(g(n)\Gamma)}$ of the degree s-step nilsequence $F(g(n)\Gamma)$ is a Lipschitz polynomial nilsequence of degree $(s-1)$, the underlying nilmanifold being

$$\overline{(G^{\square})} = (G \times_{G_{(2)}} G)/G^{\triangle}_{(s)},$$

https://doi.org/10.1017/S0017089510000546 Published online by Cambridge University Press

where $G \times_{G_{(2)}} G = \{(g,h) : g,h \in G, gh^{-1} \in G_{(2)}\}$, and $G_s^{\Delta} = \{(g_s, g_s) : g_s \in G_{(s)}\}$. For details of this theory see Section 7 of [**14**].

We now invoke our induction hypothesis to conclude that

$$\|\Delta_h f\|_{U_s} \gg_\delta 1$$

for $\gg_\delta N$ values of h.

Noting that

$$\|f\|_{U_{s+1}}^{2^{s+1}} = \mathbb{E}_{h \in \mathbb{Z}/N'\mathbb{Z}} \|\Delta_h f\|_{U_s}^{2^s},$$

we are done. □

It is perhaps worth reiterating the main point of the above argument, since it explains the importance of nilsequences in the whole theory: the derivative of a degree s polynomial nilsequence with a vertical character is a degree $(s-1)$ polynomial nilsequence.

ACKNOWLEDGEMENTS. BG was, for some of the period during which this work was carried out, a fellow of the Radcliffe Institute at Harvard. He is very grateful to the Radcliffe Institute for providing excellent working conditions. TT is supported by NSF Research Award DMS-0649473, the NSF Waterman award and a grant from the MacArthur Foundation. TZ is supported by ISF grant 557/08, an Alon fellowship, and a Landau fellowship of the Taub foundation . All three authors are very grateful to the University of Verona for allowing them to use classrooms at Canazei during a week in July 2009 when this work was largely completed.

REFERENCES

1. V. Bergelson, T. C. Tao and T. Ziegler, An inverse theorem for uniformity seminorms associated with the action of F^{ω}, *Geom. Funct. Anal.* **19**(6) (2010), 1539–1596.

2. N. N. Bogolyubov, Sur quelques propriétés arithmétiques des presque-périodes, *Ann. Chaire Math. Phys. Kiev* **4** (1939), 185–194.

3. N. Bourbaki, Lie groups and Lie algebras. Chaps. 1–3 (translated from French), in *Elements of mathematics* (Springer-Verlag, Berlin, 1998), xviii+450 pp.

4. J. Bourgain, On arithmetic progressions in sums of sets of integers, in *A tribute to Paul Erdős* (Cambridge University Press, Cambridge, UK, 1990), 105–109.

5. J. Bourgain, On triples in arithmetic progression, *Geom. Funct. Anal.* **9**(5) (1999), 968–984.

6. H. Furstenberg and B. Weiss, A mean ergodic theorem for $1/N \sum_{n=1}^{N} f(T^n x)g(T^{n^2}x)$, in *Convergence in ergodic theory and probability* (Columbus, OH, 1993), Ohio State University Mathematical Research Institute Publications, vol. 5 (de Gruyter, Berlin, 1996), 193–227.

7. W. T. Gowers, A new proof of Szemerédi's theorem for progressions of length four, *GAFA* **8**(3) (1998), 529–551.

8. W. T. Gowers, A new proof of Szemerédi's theorem, *GAFA* **11** (2001), 465–588.

9. B. J. Green, Arithmetic progressions in sumsets, *GAFA* **12**(3) (2002), 584–597.

10. B. J. Green, Generalising the Hardy–Littlewood method for primes, in *International Congress of Mathematicians*, vol. II (European Mathematical Society, Zurich, 2006), 373–399.

11. B. J. Green and T. C. Tao, An inverse theorem for the Gowers U^3-norm, with applications, *Proc. Edinburgh Math. Soc.* **51**(1) (2008), 71–153.

12. B. J. Green and T. C. Tao, Quadratic uniformity of the Möbius function, *Ann. l'Inst. Fourier* (Grenoble) **58**(6) (2008), 1863–1935.

13. B. J. Green and T. C. Tao, Linear equations in primes, *Ann. Math.* (in press).

14. B. J. Green and T. C. Tao, The quantitative behaviour of polynomial orbits on nilmanifolds, *Ann. Math.* (in press).

15. B. J. Green and T. C. Tao, The Möbius function is strongly orthogonal to nilsequences, *Ann. Math.* (in press).

16. B. J. Green and T. C. Tao, An arithmetic regularity lemma, associated counting lemma, and applications, in *Proceedings of the conference in honour of the 70th birthday of Endre Szemerédi* (in press).

17. B. J. Green, T. C. Tao and T. Ziegler, An inverse theorem for the Gowers norms (manuscript submitted for publication).

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

19. E. Hrushovski, Totally categorical structures, *Trans. Amer. Math. Soc.* **313**(1) (1989), 131–159.

20. A. Leibman, Polynomial sequences in groups, *J. Algebra* **201** (1998), 189–206.

21. A. Leibman, Pointwise convergence of ergodic averages of polynomial sequences of translations on a nilmanifold, *Ergodic Theory Dyn. Syst.* **25**(1) (2005), 201–213.

22. A. Leibman, A canonical form and the distribution of values of generalised polynomials, *Israel J. Math*.

23. V. Lev, Optimal representations by sumsets and subset sums, *J. Number Theory* **62**(1) (1997), 127–143.

24. A. Sárközy, Finite addition theorems, I, *J. Number Theory* **32** (1989), 114–130.

25. T. C. Tao and V. Vu, Additive combinatorics, in *Cambridge studies in advanced mathematics*, vol. 105 (Cambridge University Press, Cambridge, UK, 2006).

26. T. C. Tao and T. Ziegler, The inverse conjecture for the Gowers norms over finite fields via the correspondence principle, *Analysis PDE* (in press).

27. R. C. Vaughan, The Hardy–Littlewood method, in *Cambridge tracts in mathematics*, vol. 80 (Cambridge University Press, New York, 1981).

28. T. Ziegler, Universal characteristic factors and Furstenberg averages, *J. Amer. Math. Soc.* **20** (2007), 53–97.