

MULTIPLE CORRELATION SEQUENCES NOT APPROXIMABLE BY NILSEQUENCES

JOP BRIËT AND BEN GREEN

ABSTRACT. We show that there is a measure-preserving system (X, \mathcal{B}, μ, T) together with functions $F_0, F_1, F_2 \in L^\infty(\mu)$ such that the correlation sequence $C_{F_0, F_1, F_2}(n) = \int_X F_0 \cdot T^n F_1 \cdot T^{2n} F_2 d\mu$ is not an approximate integral combination of 2-step nilsequences.

1. INTRODUCTION

Let (X, \mathcal{B}, μ, T) be a measure-preserving system, and let $F_0, F_1, \dots, F_k \in L^\infty(\mu)$. Motivated in large part by applications in combinatorics and in particular to questions about arithmetic progressions, there has been much interest in *multiple correlation sequences*

$$C_{F_0, \dots, F_k}(n) := \int_X F_0 \cdot T^n F_1 \cdots T^{kn} F_k d\mu.$$

In fact, much more general types of correlation sequences in which the powers T, T^2, \dots, T^k appearing here are replaced by measure-preserving maps T_1, \dots, T_k have been studied, but here we restrict attention here to this special form.

In the case $k = 1$, there is a very satisfactory spectral theory of such sequences and indeed one has

$$C_{F_0, F_1}(n) = \int_0^1 e^{-2\pi i n t} d\sigma(t) \tag{1.1}$$

for some complex Borel measure σ of bounded total variation. This follows from the Herglotz theorem on positive definite sequences (which applies directly in the case $F_0 = \overline{F_1}$) and a depolarization identity.

It is natural to ask to what extent this generalises to $k \geq 2$. In the words of Frantzikinakis [7],

2000 *Mathematics Subject Classification*. Primary: 11B30; Secondary: 37A45.
 The first author is supported by the Gravitation grant NETWORKS-024.002.003 from the Dutch Research Council (NWO). The second author is supported by a Simons Investigator Award and is grateful to the Simons Foundation for their support.

“Finding a formula analogous to (1.1), with the multiple correlation sequences in place of the single correlation sequences, is a problem of fundamental importance which has been in the mind of experts for several years. A satisfactory solution is going to give us new insights and significantly improve our ability to deal with multiple ergodic averages.”

A result of Bergelson, Host and Kra [2] describes the structure of multiple correlation sequences up to an error in ℓ^1 or ℓ^2 . To state their result, we need to recall the notion of a nilsequence.

Definition 1.1 (Nilsequence). Let $k \geq 1$ be an integer. A k -step nilsequence is a sequence $(\phi(g^n x_0))_{n \in \mathbb{Z}}$. Here, $\phi : G \rightarrow \mathbb{C}$ is a continuous function satisfying the automorphy¹ condition $\phi(x\gamma) = \phi(x)$ for all $x \in G$ and all $\gamma \in \Gamma$, where G is a simply-connected k -step nilpotent Lie group with discrete and cocompact subgroup Γ , and g, x_0 are fixed elements of G .

A careful discussion of this notion may be found in many places, for instance [2]. The following result is [2, Theorem 1.9].

Theorem 1.1. *Suppose that (X, \mathcal{B}, μ, T) is a measure-preserving system and that $F_0, F_1, \dots, F_k \in L^\infty(\mu)$. Suppose that $\|F_i\|_\infty \leq 1$. Then we have a decomposition*

$$C_{F_0, F_1, \dots, F_k}(n) = a(n) + b(n),$$

where $a(n)$ is a uniform limit of k -step nilsequences with $\|a\|_\infty \leq 1$, and b is small in the sense that

$$\lim_{|I| \rightarrow \infty} \frac{1}{|I|} \sum_{n \in I} |b(n)| = 0$$

as I ranges over all subintervals of \mathbb{N} .

For applications involving the behaviour of correlation sequences at a sparse sequence of n , the error term here is too big. Frantzikinakis [7, Problem 1] has suggested, in the context of seeking a generalisation of (1.1), that a variant of Theorem 1.1 should hold with an ℓ^∞ error term. Note that in (1.1), we have not just one nilsequence $(e^{2\pi i n t})_{n \in \mathbb{N}}$, but an integral combination of (1-step) nilsequences. Frantzikinakis’s formulation generalises this concept to higher-step nilsequences.

¹Essentially equivalently, ϕ is a function on the *nilmanifold* G/Γ .

Definition 1.2. An integral combination of k -step nilsequences is a sequence of the form

$$a(n) = \int_M a_m(n) d\sigma(m).$$

Here, M is a compact metric space, σ is a complex Borel measure of bounded variation, and the a_m are k -step nilsequences, and with the map $m \mapsto a_m(n)$ being measurable for each n .

Our main theorem states that, even in the case $k = 2$, one cannot hope for a version of Theorem 1.1 in which the error b is small in ℓ^∞ , even if one allows a to be an integral combination of nilsequences.

Theorem 1.2. *There is a measure-preserving system (X, \mathcal{B}, μ, T) , functions $F_0, F_1, F_2 \in L^\infty(\mu)$ and an $\varepsilon > 0$ such that the correlation sequence*

$$C_{F_0, F_1, F_2}(n) := \int_X F_0 \cdot T^n F_1 \cdot T^{2n} F_2 d\mu$$

cannot be written as $a(n) + b(n)$, where $\|b\|_\infty \leq \varepsilon$ and a is an integral combination of 2-step nilsequences.

This theorem casts some serious doubt on the existence of a formula generalising (1.1).

Theorem 1.2 does not provide a negative answer to [7, Problem 1], because Frantzikinakis allows the automorphic functions ϕ in the definition of a nilsequence to be merely Riemann-integrable, rather than continuous. He calls these *generalised* nilsequences. An explanation of why our construction does not allow one to establish an analogue of Theorem 1.2 for generalised nilsequences is given in Appendix A. Note, however, that the Riemann-integrable functions ϕ appearing in Appendix A are very singular and we certainly do not expect that the corresponding generalised nilsequences have any important role to play in the theory.

One reason for considering Riemann-integrable functions rather than just continuous ones is that there is a somewhat natural and well-studied class of nilsequences in which ϕ is not continuous, namely the bracket polynomial phases [3]. In this case, the corresponding ϕ have only mild discontinuities, and our argument adapts easily to show that Theorem 1.2 remains true even if one allows a to be an integral combination of this more general class of nilsequences. We sketch the argument at the end of Section 3.

A key motivation for Frantzikinakis in formulating [7, Problem 1] was that it provides a potential route to understanding Szemerédi's theorem

with common difference in a sparse random set, a problem for which our current understanding is extremely incomplete for progressions of length 3 or longer (see [4] for recent progress). Whilst Theorem 1.2 seems to rule this out as a viable strategy, our example unfortunately does not give any new information about Szemerédi's theorem with common differences from a random set, which remains a tantalising open problem.

Notation. Our notation is standard. We will occasionally write $\mathbb{E}_{x \in A}$ for $\frac{1}{|A|} \sum_{x \in A}$, where A is a finite set. We write $[N] = \{1, 2, \dots, N\}$ as usual, and sometimes we will write $[0, N - 1] = \{0, 1, 2, \dots, N - 1\}$. For real t , we write $e(t) = e^{2\pi it}$.

Acknowledgements. JB would like to thank Xuancheng Shao for helpful discussions and pointers to the literature. The authors would like to thank Bryna Kra for pointing them to a reference for Proposition 5.1, and Nikos Frantzikinakis for helpful comments on the first draft of the paper.

2. OUTLINE OF THE ARGUMENT

Our argument is part deterministic and part random. It is random in the sense that we do not explicitly construct a system (X, \mathcal{B}, μ, T) and functions F_0, F_1, F_2 for which the correlation sequence $C_{F_0, F_1, F_2}(n)$ is not approximable by an integral combination of nilsequences, but rather we show there are too many possibilities for the correlation functions $C_{F_0, F_1, F_2}(n)$ for this to be so.

To do this, we first explicitly construct a certain infinite sequence $\mathcal{S} \subset \mathbb{N}$ whose growth is slower than exponential in the sense that

$$\lim_{N \rightarrow \infty} \frac{|\mathcal{S}[N]|}{\log N} = \infty, \quad (2.1)$$

where $\mathcal{S}[N] := \#\{n \in \mathcal{S} : n \leq N\}$.

We show that for any choice of function $\eta : \mathcal{S} \rightarrow \{1, -\frac{1}{3}\}$ there is a system (X, \mathcal{B}, μ, T) and functions F_0, F_1, F_2 such that $C_{F_0, F_1, F_2}(n) = \eta(n)$ for $n \in \mathcal{S}$.

For a random choice of η , such a function will almost surely not be approximable by an integral combination of nilsequences. We give the details of this deduction, which uses nothing about \mathcal{S} other than the growth property (2.1), in Proposition 3.1.

The heart of the argument, then, is the construction of the system (X, \mathcal{B}, μ, T) and the functions F_0, F_1, F_2 , given $\eta : \mathcal{S} \rightarrow \{1, -\frac{1}{3}\}$. This is assembled from a sequence of finitary examples, via a (well-known)

variant of Furstenberg's correspondence principle, and here the specific nature of \mathcal{S} is critical.

The idea behind the construction of these finitary examples ultimately comes from coding theory, and in particular a construction of Yekhanin [9]. We will only need the most basic form of these ideas; for instance, we can replace all the finite-field theory in Yekhanin's work with the simple observation that the function $\psi : \mathbb{Z} \rightarrow \{-1, 1\}$ defined by $\psi(0) = 1$, $\psi(1) = \psi(2) = -1$, and periodic mod 3 has the property that

$$\psi(x)\psi(x+d)\psi(x+2d) = \begin{cases} \psi(x) & d \equiv 0 \pmod{3} \\ 1 & d \not\equiv 0 \pmod{3}. \end{cases}$$

The idea of using Yekhanin's construction to give interesting examples in the additive combinatorics of higher-order correlations first arose in the finite field setting, in joint work of the first author and Labib [5]. Those ideas have inspired the present work.

3. ENTROPY AND NILSEQUENCES

Proposition 3.1. *Let \mathcal{S} be an increasing sequence of natural numbers such that*

$$\lim_{N \rightarrow \infty} \frac{|\mathcal{S}[N]|}{\log N} = \infty. \quad (3.1)$$

Then there is a function $\eta : \mathcal{S} \rightarrow \{1, -\frac{1}{3}\}$ such that

$$\lim_{N \rightarrow \infty} \frac{1}{|\mathcal{S}[N]|} \sum_{n \in \mathcal{S}[N]} \eta(s)a(s) = 0 \quad (3.2)$$

for all nilsequences a .

Proof. The space of C^∞ -functions on G/Γ is dense in the space of continuous functions; to approximate a continuous function by a smooth function, average with respect to a smooth kernel supported near the identity on G . It therefore suffices to verify (3.2) for $a(n) = \phi(g^n x)$ with $\phi \in C^\infty(G/\Gamma)$. Now we use the fact that there is a map

$$\text{Complexity} : \{\text{smooth nilsequences}\} \rightarrow (0, \infty)$$

and a function $M : (0, \infty) \times (0, 1) \rightarrow (0, \infty)$ such that the set

$$\{a : \text{Complexity}(a) \leq C\}$$

can be covered by $N^{M(C, \varepsilon)}$ balls of radius ε in $\ell^\infty[N]$.

Results of this type were first observed by Frantzikinakis [6, Proposition 6.2], and in fact Proposition 3.1 and its proof are very closely related to [6, Theorem 1.4]. A discussion which gives what we need here

is in the appendix of Altman [1] (note that $(g^n x_0)_{n \in \mathbb{Z}}$ is a particular example of a polynomial sequence as considered by Altman).

We will pick the values of $\eta(n)$ at random, choosing $\eta(n) = -\frac{1}{3}$ with probability $\frac{3}{4}$, and $\eta(n) = 1$ with probability $\frac{1}{4}$, these choices being independent for different values of $n \in \mathcal{S}$. Then $\mathbb{E}\eta(n) = 0$. By well-known large deviation estimates (Hoeffding's inequality), for any fixed 1-bounded function b , and for any distinct n_1, \dots, n_m ,

$$\mathbb{P}\left(\left|\sum_{i=1}^m \eta(n_i)b(n_i)\right| \geq t\right) \ll e^{-ct^2/m}, \quad (3.3)$$

where $c > 0$ is absolute.

Let $\omega : \mathbb{N} \rightarrow (0, \infty)$ be some function tending to infinity, to be specified later.

For each N , let E_N be the following event: for all 1-bounded nilsequences a of complexity $\leq \omega(N)$,

$$\left|\sum_{n \in \mathcal{S}[N]} \eta(n)a(n)\right| \leq \frac{1}{\omega(N)}|\mathcal{S}[N]|. \quad (3.4)$$

We estimate $\mathbb{P}(E_N)$ as follows. Pick some collection $\{a_1, \dots, a_J\}$, $J \leq N^{M(\omega(N), 1/2\omega(N))}$ of functions such that, for every 1-bounded nilsequence a of complexity at most $\omega(N)$, there is some a_i with $\|a - a_i\|_{\ell^\infty[N]} \leq 1/2\omega(N)$. Note that we do not need to assume that the a_i are nilsequences (though this could be arranged if desired) and they are automatically 2-bounded.

If we are not in E_N , there is some a_i such that

$$\left|\sum_{n \in \mathcal{S}[N]} \eta(n)a_i(n)\right| \geq \frac{1}{2\omega(N)}|\mathcal{S}[N]|. \quad (3.5)$$

By (3.3), the probability of (3.5) happening, for some fixed i , is bounded above by $e^{-c'|\mathcal{S}[N]|/\omega(N)^2}$ for some $c' > 0$. Summing over i , it follows that

$$\mathbb{P}(\neg E_N) \leq N^{M(\omega(N), 1/2\omega(N))} e^{-c'|\mathcal{S}[N]|/\omega(N)^2}.$$

Choose ω (with $\omega(N) \rightarrow \infty$) so that

$$\frac{|\mathcal{S}[N]|}{\log N} > \frac{\omega(N)^2}{c'}(10 + M(\omega(N), 1/2\omega(N)))$$

for N sufficiently large. (Here, of course, we have used the assumption on \mathcal{S}). This then means that

$$\mathbb{P}(\neg E_N) \leq N^{-10}$$

for large N . In particular, $\sum_N \mathbb{P}(\neg E_N) < \infty$ which, by Borel-Cantelli, implies that almost surely only finitely many of the $\neg E_N$ occur. In particular, there is some particular choice of η such that (3.4) holds for all sufficiently large N . Since every nilsequence has finite complexity, this implies the result. \square

Remark. There is of course nothing special about $\{1, -\frac{1}{3}\}$; any set containing both positive and negative numbers would do.

To conclude this section, let us quickly sketch how one could extend Proposition 3.1 to include the case where $a()$ is a bracket polynomial or a product of such (and hence not a nilsequence with a *continuous* automorphic function ϕ). Write $\chi_{\alpha,\beta}(n) := e(\alpha n \lfloor \beta n \rfloor)$. The key point is that the set of functions $\chi_{\alpha,\beta}(n)$, like the set of nilsequences of fixed complexity, has polynomially-bounded covering numbers in $\ell^\infty[N]$.

To see why this is so, first note that $\chi_{\alpha,\beta}$ depends only on $\alpha \pmod{1}$, so we may assume $0 \leq \alpha < 1$. Next, replacing β by $\beta + k$ for $k \in \mathbb{Z}$ has the effect of multiplying by a quadratic phase $e(\gamma n^2)$ (where $\gamma = \alpha k$). However, the set of all quadratic phases $e(\gamma n^2)$ is covered by $\ll_\varepsilon N^2$ balls of radius ε in $\ell^\infty[N]$, since we may assume $0 \leq \gamma < 1$ and changing γ by $\frac{\varepsilon}{N^2}$ only changes $e(\gamma n^2)$ by $O(\varepsilon)$, uniformly for $n \leq N$.

It therefore suffices to show that the covering numbers of the set $\Xi := \{\chi_{\alpha,\beta} : 0 \leq \alpha, \beta < 1\}$ are polynomially bounded in $\ell^\infty[N]$. Now, restricted to $n \leq N$, there are only polynomially many functions $\lfloor \beta n \rfloor$ as β ranges in $[0, 1)$. Indeed, the map $\beta \mapsto (\lfloor \beta n \rfloor)_{n \leq N}$ is only discontinuous at the points where $\beta n \in \mathbb{Z}$ for some $n \leq N$, of which there are no more than N^2 with $0 \leq \beta < 1$. Thus $\chi_{\alpha,\beta} = \chi_{\alpha,\beta'}$, with β' varying in a set of size N^2 . Changing α by $\frac{\varepsilon}{N^2}$ only changes $\chi_{\alpha,\beta}(n)$ by $O(\varepsilon)$, uniformly for $n \leq N$. Therefore the covering number of Ξ in $\ell^\infty[N]$ is $\ll_\varepsilon N^4$.

It follows immediately that, for fixed C , the set of functions of type $e(\sum_{i=1}^k \alpha_i n \lfloor \beta_i n \rfloor)$, where $k \leq C$, is covered by $N^{M(C,\varepsilon)}$ balls of radius ε in $\ell^\infty[N]$. One could include various types of 1-step nilsequence or bracket polynomial and obtain a similar result.

Bounds on covering numbers were all we needed to know about nilsequences, and the rest of the argument goes over verbatim.

4. THE HEART OF THE CONSTRUCTION

Define $\psi : \mathbb{Z} \rightarrow \{-1, 1\}$ to be the function with $\psi(0) = 1$, $\psi(1) = \psi(2) = -1$, and periodic mod 3. The crucial property of this function

we will use is the following, which is easily checked:

$$\psi(x)\psi(x+d)\psi(x+2d) = \psi(x) \quad (4.1)$$

if $d \equiv 0 \pmod{3}$, and 1 if $d \not\equiv 0 \pmod{3}$.

Fix, once and for all, a sequence $M_1 < M_2 < \dots$ be a sequence of positive integers such that

- (1) Each M_i is a multiple of 3;
- (2) $\lim_{n \rightarrow \infty} k^{-2} \sum_{i=1}^k \log M_i = 0$;
- (3) $\prod_{i=1}^{\infty} (1 - \frac{3}{M_i}) = \gamma > 0$.

For instance, one could take $M_i = 3i^2$.

Define

$$\Omega_k := \{(x_1, x_2, \dots) : 0 \leq x_i < M_i, x_{k+1} = x_{k+2} = \dots = 0\}.$$

Later on we will need the technical variant

$$\tilde{\Omega}_k := \{(x_1, x_2, \dots) : 0 \leq x_i < M_i - 3, x_{k+1} = x_{k+2} = \dots = 0\}.$$

Define also Σ_k to consist of all sequences (x_1, x_2, \dots) with precisely two nonzero entries x_a, x_b , both of which equal 1, and with $x_{k+1} = x_{k+2} = \dots = 0$. Write

$$\Omega := \bigcup_k \Omega_k, \quad \tilde{\Omega} := \bigcup_k \tilde{\Omega}_k, \quad \Sigma := \bigcup_k \Sigma_k.$$

We have a bijective map

$$\beta : \Omega \rightarrow \mathbb{Z}_{\geq 0}$$

defined by

$$\beta(x_1, x_2, \dots) = x_1 + M_1 x_2 + M_1 M_2 x_3 + \dots.$$

Let $\mathcal{S} = \beta(\Sigma)$. Thus \mathcal{S} consists of the sums of two distinct elements of the sequence $\{1, M_1, M_1 M_2, M_1 M_2 M_3, \dots\}$. We claim that \mathcal{S} satisfies the hypothesis (3.1) of Lemma 3.1, that is to say $\lim_{N \rightarrow \infty} \frac{|\mathcal{S}[N]|}{\log N} = \infty$.

To see this, let k be maximal so that $M_1 \cdots M_k \leq N/2$. Then $|\mathcal{S}[N]| \geq \binom{k}{2}$, whilst $\log(N/2) \leq \sum_{i=1}^{k+1} \log M_i$. Therefore it is enough that

$$\lim_{k \rightarrow \infty} k^{-2} \sum_{i=1}^{k+1} \log M_i = 0,$$

which follows immediately from assumption (2) above.

We now apply Lemma 3.1 to get a function $\eta : \mathcal{S} \rightarrow \{1, -\frac{1}{3}\}$ satisfying (3.2). Define

$$\Sigma_k^+ := \{x \in \Sigma_k : \eta(\beta(x)) = 1\} \quad \text{and} \quad \Sigma_k^- := \{x \in \Sigma_k : \eta(\beta(x)) = -\frac{1}{3}\}.$$

Thus $\Sigma_k = \Sigma_k^- \cup \Sigma_k^+$.

We introduce one more piece of notation. If $z \in \Sigma_k$ and if $x \in \Omega_k$ then we write

$$\sigma_z(x) := \sum_{i \in [k]: z_i=0} x_i.$$

Now we come to the crucial definition. Let $k \in \mathbb{N}$. For $x \in \Omega_k$ define

$$f_k(\beta(x)) = \prod_{z \in \Sigma_k^-} \psi(\sigma_z(x)). \quad (4.2)$$

Note that $\beta(\tilde{\Omega}_k) = [0, N_k - 1]$, where

$$N_k := M_1 \cdots M_k, \quad (4.3)$$

and so f_k is a well-defined function on $[0, N_k - 1]$, taking values in $\{-1, 1\}$. Define also the technical variant

$$\tilde{f}_k(\beta(x)) := 1_{x \in \tilde{\Omega}_k} f_k(\beta(x)). \quad (4.4)$$

Thus \tilde{f}_k is defined on $[0, N_k - 1]$ and takes values in $\{-1, 0, 1\}$. Extend both f_k and \tilde{f}_k to functions on all of $\mathbb{Z}_{\geq 0}$ by defining $f_k(n) = \tilde{f}_k(n) = 0$ for $n \geq N_k$.

The following lemma is the heart of the argument. Here, recall that $\gamma > 0$ is just a positive constant (appearing in point (3) of the list of properties satisfied by the M_i).

Lemma 4.1. *For $d \in \mathbb{Z}_{\geq 0}$, write*

$$S_k(d) := \frac{1}{N_k} \sum_{n \in [0, N_k - 1]} \tilde{f}_k(n) f_k(n + d) f_k(n + 2d).$$

Then for $d \in \mathcal{S}$ we have $\lim_{k \rightarrow \infty} S_k(d) = \gamma \eta(d)$.

Proof. Let $d \in \mathcal{S} = \beta(\Sigma)$. For k large enough, $d \in \beta(\Sigma_k)$, and we will assume this is so in what follows.

From the definition of \tilde{f}_k , we see that the sum over n ranges over $n = \beta(x)$, $x \in \tilde{\Omega}_k$. Now for n of this form and for $d = \beta(y)$, $y \in \Sigma_k$, we have $x + y, x + 2y \in \Omega_k$ and moreover

$$\begin{aligned} \beta(x + y) &= \beta(x) + \beta(y) = n + d, \\ \beta(x + 2y) &= \beta(x) + 2\beta(y) = n + 2d. \end{aligned}$$

Note that this “lack of carries” was precisely the reason for defining the set $\tilde{\Omega}_k$. It follows that

$$S_k(d) = \mathbb{E}_{x \in \Omega_k} \tilde{f}_k(\beta(x)) f_k(\beta(x+y)) f_k(\beta(x+2y)),$$

for $d = \beta(y)$, $y \in \Sigma_k$. Substituting the definitions of f_k, \tilde{f}_k (and noting that σ_z is linear), we see that

$$S_k(d) = \mathbb{E}_{x \in \Omega_k} 1_{x \in \tilde{\Omega}_k} \prod_{z \in \Sigma_-^k} \psi(\sigma_z(x)) \psi(\sigma_z(x) + \sigma_z(y)) \psi(\sigma_z(x) + 2\sigma_z(y)).$$

From (4.1) it follows that

$$S_k(d) = \mathbb{E}_{x \in \Omega_k} 1_{x \in \tilde{\Omega}_k} \prod_{z \in \Sigma_-^k : \sigma_z(y) \equiv 0 \pmod{3}} \psi(\sigma_z(x)).$$

Now both y and z here are vectors with only two nonzero entries and so $\sigma_z(y)$ takes only the values 0, 1, 2 with $\sigma_z(y) = 0$ iff $y = z$. Therefore

$$S_k(d) = \begin{cases} \mathbb{E}_{x \in \Omega_k} 1_{x \in \tilde{\Omega}_k} \psi(\sigma_y(x)) & \text{if } y \in \Sigma_-^k \\ \mathbb{E}_{x \in \Omega_k} 1_{x \in \tilde{\Omega}_k} & \text{if } y \in \Sigma_+^k. \end{cases} \quad (4.5)$$

The second expression is

$$\mathbb{E}_{x \in \Omega_k} 1_{x \in \tilde{\Omega}_k} = \frac{|\tilde{\Omega}_k|}{|\Omega_k|} = \prod_{i=1}^k \left(1 - \frac{3}{M_i}\right) \rightarrow \gamma$$

as $k \rightarrow \infty$. The first expression in (4.5) may be written explicitly as

$$\frac{|\tilde{\Omega}_k|}{|\Omega_k|} \mathbb{E}_{x \in \tilde{\Omega}_k} \psi(x_1 + \cdots + \hat{x}_i + \cdots + \hat{x}_j + \cdots + x_k), \quad (4.6)$$

where y has nonzero coordinates at i, j and the hat means that \hat{x}_i does not appear in the sum. Note, however, that $\tilde{\Omega}_k$ is a box with sidelengths $M_i - 3$, each of which is a multiple of 3. Therefore $x_1 + \cdots + \hat{x}_i + \cdots + \hat{x}_j + \cdots + x_k$ is uniformly distributed mod 3, as x ranges uniformly over $\tilde{\Omega}_k$, and the average in (4.6) is

$$\frac{|\tilde{\Omega}_k|}{|\Omega_k|} \cdot \left(-\frac{1}{3}\right) = -\frac{1}{3} \prod_{i=1}^k \left(1 - \frac{3}{M_i}\right) \rightarrow -\frac{\gamma}{3}.$$

This completes the proof. \square

5. PUTTING EVERYTHING TOGETHER

Our final task is to build a measure-preserving system from the functions constructed in the last section. For this we will need a slight variant of the usual Furstenberg correspondence principle, proven in a very similar way. An essentially equivalent statement may be found, for instance, in [8, Proposition 3.3].

Proposition 5.1. *Let $A \subset \mathbb{R}$ be a finite set. Suppose that for each $k \in \mathbb{N}$ we have functions $f_{0,k}, \dots, f_{r,k} : \mathbb{Z}_{\geq 0} \rightarrow A$, and that $(N_k)_{k=1}^{\infty}$ is an increasing sequence of positive integers. Then there is a measure-preserving system (X, \mathcal{B}, μ, T) and functions $F_0, F_1, \dots, F_r \in L^\infty(\mu)$ such that the following is true: if (d_1, \dots, d_r) is a tuple of distinct positive integers such that*

$$S(d_1, \dots, d_r) := \lim_{k \rightarrow \infty} \frac{1}{N_k} \sum_{n \in [0, N_k - 1]} f_{0,k}(n) f_{1,k}(n + d_1) \cdots f_{r,k}(n + d_r)$$

exists, then

$$S(d_1, \dots, d_r) = \int_X F_0 \cdot T^{d_1} F_1 \cdot T^{d_2} F_2 \cdots T^{d_r} F_r d\mu.$$

We will apply this with the functions constructed in the last section, taking $r = 2$, $f_{0,k} := \tilde{f}_k$, $f_{1,k} = f_{2,k} = f_k$, and $N_k = M_1 \cdots M_k$ as before.

By Proposition 5.1 and Lemma 4.1, there is a measure-preserving system (X, \mathcal{B}, μ, T) together with functions $F_0, F_1, F_2 \in L^\infty(\mu)$ such that, writing $C_{F_0, F_1, F_2}(d) := \int_X F_0 \cdot T^d F_1 \cdot T^{2d} F_2 d\mu$, we have

$$C_{F_0, F_1, F_2}(d) = \eta(d) \quad \text{for } d \in \mathcal{S}. \quad (5.1)$$

(Note it is clearly possible to scale the F_i to remove γ factor appearing in Lemma 4.1.) We claim that it is impossible to write

$$C_{F_0, F_1, F_2}(n) = a(n) + b(n)$$

with a an integral combination of 2-step nilsequences and $\|b\|_\infty \leq \frac{1}{100}$. Suppose that this were possible. Then, from (5.1) and the fact that η takes values in $\{1, -\frac{1}{3}\}$, we would have $(a(d) + b(d))\eta(d) \in \{\frac{1}{9}, 1\}$ for all $d \in \mathcal{S}$. However, $|b(d)\eta(d)| \leq \frac{1}{100}$, and therefore

$$\Re(a(d)\eta(d)) \geq \frac{1}{9} - \frac{1}{100} > \frac{1}{10} \quad (5.2)$$

for all $d \in \mathcal{S}$.

Suppose that

$$a(n) = \int_M a_m(n) d\sigma(m).$$

Here, M is a compact metric space, σ is a complex Borel measure of bounded variation and the a_m are nilsequences, with the map $m \mapsto a_m(n)$ being in $L^\infty(\sigma)$.

Then (5.2) implies that

$$\left| \frac{1}{|\mathcal{S}[N]|} \sum_{n \in \mathcal{S}[N]} a(n)\eta(n) \right| \geq \frac{1}{10}.$$

On the other hand we have

$$\left| \frac{1}{|\mathcal{S}[N]|} \sum_{n \in \mathcal{S}[N]} a(n)\eta(n) \right| \leq \int_M \left| \frac{1}{|\mathcal{S}[N]|} \sum_{n \in \mathcal{S}[N]} a_m(n)\eta(n) \right| d|\sigma|$$

However, by the choice of η (Lemma 3.1) we have

$$\lim_{N \rightarrow \infty} \frac{1}{|\mathcal{S}[N]|} \sum_{n \in \mathcal{S}[N]} a_m(n)\eta(n) = 0$$

for all m . Therefore, by the dominated convergence theorem,

$$\lim_{N \rightarrow \infty} \int_M \left| \frac{1}{|\mathcal{S}[N]|} \sum_{n \in \mathcal{S}[N]} a_m(n)\eta(n) \right| d|\sigma| = 0.$$

Putting these statements together gives a contradiction, and this completes the proof of Theorem 1.2.

APPENDIX A. GENERALISED NILSEQUENCES

In this appendix we explain why our example does not seem to give a negative solution to [7, Problem 1]. That is, we explain why our example (or similar ones) do not seem to be able to rule out the possibility that $C_{F_0, F_1, F_2}(n)$ is an approximate integral combination of *generalised* 2-step nilsequences, in which the automorphic function ϕ is allowed to be merely Riemann-integrable. In fact, our examples agree with 1-step generalised nilsequences on the crucial set \mathcal{S} .

Recall that $\mathcal{S} = \mathcal{A} \hat{+} \mathcal{A}$, where

$$\mathcal{A} = \{N_0, N_1, N_2, \dots\} \quad \text{and} \quad N_i := \prod_{j \leq i} M_j$$

(thus $N_0 = 1$, $N_1 = M_1$, $N_2 = M_1 M_2$ and so on). Here, $\mathcal{A} \hat{+} \mathcal{A}$ means the restricted sumset of \mathcal{A} with itself, that is to say the set of sums of two distinct elements of \mathcal{A} .

Proposition A.1. *There is $\theta \in \mathbb{R}/\mathbb{Z}$ such that the following is true. Let $\eta : \mathcal{S} \rightarrow [-1, 1]$ be any function. Then there is a Riemann-integrable function $\phi : \mathbb{R}/\mathbb{Z} \rightarrow [-1, 1]$ such that $\phi(\theta n) = \eta(n)$ for all $n \in \mathcal{S}$.*

Proof. Set $\theta := \sum_{i=1}^{\infty} \frac{1}{N_i}$. Since $M_1 < M_2 < \dots$, we certainly have $M_j \geq j$. As a consequence, the usual proof that e is irrational may be adapted easily to show that θ is irrational: if $\theta = \frac{p}{q}$ then $\alpha := \frac{M_1 \dots M_{qp}}{q} \in \frac{1}{q}\mathbb{Z}$, but on the other hand the fractional part of α satisfies

$$0 < \{\alpha\} = \frac{1}{M_{q+1}} + \frac{1}{M_{q+1}M_{q+2}} + \dots \leq \frac{1}{q+1} + \frac{1}{(q+1)(q+2)} + \dots < \frac{1}{q}.$$

Now define $\phi : \mathbb{R}/\mathbb{Z} \rightarrow [-1, 1]$ as follows: $\phi(\theta n) = \eta(n)$ for all $n \in \mathcal{S}$, and $\phi(x) = 0$ if $x \notin \theta\mathcal{S}$. Since θ is irrational, this is a well-defined function.

We claim that it is Riemann-integrable, with integral zero. It is enough to show that for every $\varepsilon > 0$ there is some finite collection of intervals, of total length $< \varepsilon$, whose union covers $\theta\mathcal{S}$.

Note that for every j we have

$$\|\theta N_j\|_{\mathbb{R}/\mathbb{Z}} = \frac{1}{M_{j+1}} + \frac{1}{M_{j+1}M_{j+2}} + \dots < \frac{1}{M_{j+1} - 1}. \quad (\text{A.1})$$

Moreover, condition (3) in the definition of the M_j s implies that

$$\limsup_{j \rightarrow \infty} \frac{M_j}{j} = \infty. \quad (\text{A.2})$$

In particular we may choose k so that $\frac{1}{M_{k+1}-1} < \frac{\varepsilon}{10k}$, and by (A.1) it follows that

$$\|\theta N_j\|_{\mathbb{R}/\mathbb{Z}} < \frac{\varepsilon}{10k} \quad \text{for } j \geq k.$$

It follows that

$$\theta\mathcal{A} \subseteq \{\theta N_0, \dots, \theta N_{k-1}\} \cup I,$$

where $I = (-\varepsilon/10k, \varepsilon/10k) \subseteq \mathbb{R}/\mathbb{Z}$. Therefore

$$\theta\mathcal{S} \subseteq \theta\mathcal{A} + \theta\mathcal{A} \subseteq \bigcup_{i,j < k} \{\theta(N_i + N_j)\} \cup \bigcup_{i < k} (\theta N_i + I) \cup (I + I),$$

which makes it clear that $\theta\mathcal{S}$ is contained in a finite union of intervals of length $< \varepsilon$. \square

REFERENCES

- [1] D. Altman, *On Szemerédi's theorem with differences from a random set*, Acta. Arith. **195** (2020), no. 1, 97–108.
- [2] V. Bergelson, B. Host and B. Kra, *Multiple recurrence and nilsequences*, Invent. Math. **160** (2005), no. 2, 261–303.
- [3] V. Bergelson and A. Leibman, *Distribution of values of bounded generalised polynomials*, Acta. Math. **198** (2007), no. 2, 155–230.
- [4] J. Briët and S. Gopi, *Gaussian width bounds with applications to arithmetic progressions in random settings*, Int. Math. Res. Not. (IMRN), published online October 2018.
- [5] J. Briët and F. Labib, in preparation.
- [6] N. Frantzikinakis, *Equidistribution of sparse sequences on nilmanifolds*, J. Anal. Math. **109** (2009), 353–395.
- [7] N. Frantzikinakis, *Some open problems on multiple ergodic averages*, Bull. Hellenic Math. Soc. **60** (2016), 41–90.
- [8] N. Frantzikinakis, *An averaged Chowla and Elliott conjecture along independent polynomials*, Int. Math. Res. Not. (IMRN) 2018, no. 12, 3721–3743.
- [9] S. Yekhanin, *Towards 3-query locally decodable codes of subexponential length*, J. ACM **55** (2008), no. 1, Art. 1, 16pp.

CENTRUM WISKUNDE & INFORMATICA (CWI), SCIENCE PARK 123, 1098 XG
AMSTERDAM, THE NETHERLANDS

Email address: `j.briet@cwi.nl`

MATHEMATICAL INSTITUTE, ANDREW WILES BUILDING, RADCLIFFE OBSER-
VATORY QUARTER, WOODSTOCK RD, OXFORD OX2 6QW

Email address: `ben.green@maths.ox.ac.uk`