Geom. Funct. Anal. Vol. 35 (2025) 1587–1637
https://doi.org/10.1007/s00039-025-00727-9

GAFA Geometric And Functional Analysis

QUADRATIC FORMS IN 8 PRIME VARIABLES

Ben Green

Abstract. We give an asymptotic for the number of prime solutions to $Q(x_1, \dots, x_8) = N$, subject to a mild non-degeneracy condition on the homogeneous quadratic form Q.

The argument initially proceeds via the circle method, but this does not suffice by itself. To obtain a nontrivial bound on certain averages of exponential sums, we interpret these sums as matrix coefficients for the Weil representation of the symplectic group $\mathrm{Sp}_8(\mathbf{Z}/q\mathbf{Z})$. Averages of such matrix coefficients are then bounded using an amplification argument and a convergence result for convolutions of measures, which reduces matters to understanding the action of certain 12-dimensional subgroups in the Weil representation. Sufficient understanding can be gained by using the basic represention theory of $\mathrm{SL}_2(k)$, k a finite field.

Contents

The author is supported by a Simons Investigator grant and is grateful to the Simons Foundation for their continued support.

1 Introduction

Let $Q : \mathbf{Z}^8 \to \mathbf{Z}$ be a homogeneous quadratic form with integer coefficients, and let $N \in \mathbf{Z}$. We will study solutions to $Q(x_1, \dots, x_8) = N$ with the x_i prime.

Throughout the paper, it is convenient to split $\mathbf{Z}^8 = \mathbf{Z}^4 \times \mathbf{Z}^4$ and to write

$$Q(x, y) = x^T a x + x^T b y + y^T c y$$

where $a, b, c \in \mathrm{Mat}_4(\mathbf{Z})$ with a, c symmetric. Equivalently,

$$Q(x, y) = (x^T, y^T) \begin{pmatrix} a & b/2 \\ b^T/2 & c \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix}.$$

Assume that b is invertible. An important role will be played by the 4-by-4 matrix

$$\Delta := 4b^{-1}ab^{-T}c - I. \tag{1.1}$$

This formally resembles the discriminant (of a form in two variables), only it is matrix-valued.

Remark. The requirement that b be invertible rules out the case that Q is diagonal. However diagonal quadratic forms can be handled (with far fewer variables) by other methods.

We turn now to the main theorem of the paper. Here, Λ is the von Mangoldt function. For q a positive integer, $\Lambda_{\mathbf{Z}/q\mathbf{Z}} : \mathbf{Z}/q\mathbf{Z} \to \mathbf{C}$ takes the value $q/\phi(q)$ when $(x, q) = 1$, and 0 otherwise. As the notation suggests, this is the natural local variant of the von Mangoldt function. The normalisation is chosen so that the average value of $\Lambda_{\mathbf{Z}/q\mathbf{Z}}$ is 1. We abuse notation by writing $\Lambda_{\mathbf{Z}/p\mathbf{Z}}$ instead of $\Lambda_{\mathbf{Z}/p^n\mathbf{Z}}$ (the domain will always be clear from context). Finally, we write $\Lambda^{\otimes 4}(x) = \Lambda(x_1)\Lambda(x_2)\Lambda(x_3)\Lambda(x_4)$ for $x \in \mathbf{Z}^4$, and $\Lambda^{\otimes 4}_{\mathbf{Z}/q\mathbf{Z}}$ is defined analogously.

Theorem 1.1. *Suppose that Q is a quadratic form for which $\det ab \neq 0$ and Δ has four distinct eigenvalues which lie in $\overline{\mathbf{Q}} \setminus \{-1, 0\}$. Then we have the expected local-global estimate for the number of solutions to $Q(x, y) = N$ in primes, namely for any real A*

$$\sum_{\substack{x,y \in [X]^4 \\ Q(x,y)=N}} \Lambda^{\otimes 4}(x)\Lambda^{\otimes 4}(y) = \mathfrak{S}(N)X^6 + O_{A,Q}(X^6 \log^{-A} X).$$

Here

$$\mathfrak{S}(N) = \beta_\infty \prod_p \beta_p(N) \tag{1.2}$$

where $\beta_p(N) = \lim_{n\to\infty} \beta_{p,n}(N)$ is the p-adic density of solutions, where

$$\beta_{p,n}(N) := p^{-7n} \sum_{\substack{x,y \in (\mathbf{Z}/p^n\mathbf{Z})^4 \\ Q(x,y) \equiv N \pmod{p^n}}} \Lambda^{\otimes 4}_{\mathbf{Z}/p\mathbf{Z}}(x)\Lambda^{\otimes 4}_{\mathbf{Z}/p\mathbf{Z}}(y), \tag{1.3}$$

and

$$\beta_\infty := \lim_{\delta\to 0} \frac{1}{2\delta}\mu_{\mathbf{R}^8}\{x, y \in [0,1]^4 : |Q(x,y) - \frac{N}{X^2}| \leqslant \delta\}$$

is an archimedean measure of the density of (*positive, real*) *solutions. Included in the statements is the fact that the limit in the definition of the p-adic density* (1.3) *exists.*

In this paper, we have restricted attention to solving equations in prime variables, since this topic is especially prominent in the literature. However, it is very likely that similar methods would allow one to handle quadratic forms in 8 variables from other arithmetic sets such as sums of two squares or smooth numbers with suitable parameters.

Throughout the paper, we will say that a form Q is *generic* if it satisfies the conditions of this theorem, that is to say if $\det ab \neq 0$ and if Δ has four distinct eigenvalues in $\overline{\mathbf{Q}}$, and that neither 0 nor -1 is one of these eigenvalues. The word generic is appropriate, since this condition holds for a Zariski-dense set of (a, b, c) in the 36-dimensional parameter space where a, c are symmetric. To see this, first note that the condition that Δ has distinct eigenvalues is (Zariski-)closed, by considering the resultant of the characteristic polynomial $\rho_\Delta(\lambda) := \det(\Delta - \lambda I)$ and its derivative $\rho'_\Delta(\lambda)$. The conditions that $\det ab = 0$, and that Δ has an eigenvalue 0 or -1, are evidently closed conditions. Finally, these conditions are nontrivial (i.e. not always satisfied) as one can see by taking a to be a diagonal matrix with distinct rational eigenvalues (not 0 or $\frac{1}{4}$) and $b = c = I$.

Note that Δ is not canonically associated to Q, being dependent on the splitting of variables into two sets of four. However, one may observe that if a is invertible then

$$\begin{pmatrix} a & b/2 \\ b^T/2 & c \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ A^T & 1 \end{pmatrix} \begin{pmatrix} B & 0 \\ 0 & C \end{pmatrix} \begin{pmatrix} 1 & A \\ 0 & 1 \end{pmatrix}$$

with $A = \frac{1}{2}a^{-1}b$, $B = a$, $C = \frac{1}{4}b^T a^{-1} b\Delta$, and so

$$\det Q := \det \begin{pmatrix} a & b/2 \\ b^T/2 & c \end{pmatrix} = 2^{-16}(\det b)^2 \det \Delta. \tag{1.4}$$

(This is also true without the assumption that a is invertible by a continuity argument.) Therefore the requirement that 0 is not an eigenvalue of Δ is essentially an invariant of Q, more-or-less equivalent to Q being non-singular (that is, $\det Q \neq 0$, also called the "regular" case in previous works such as [Zha16]). We do not expect any variant of our methods to handle the singular case $\det Q = 0$, which essentially corresponds to forms in 7 or fewer variables. In applications of the circle method to higher-degree problems, it is commonly seen that the non-singular case is easier to handle.

Previous results. Jianya Liu [Liu11] handled generic quadratic forms in 10 prime variables. Lilu Zhao [Zha16] subsequently handled all regular quadratic forms in 9 variables. These works use fairly classical forms of the Hardy-Littlewood circle method. For reasons we will go into later, 9 variables appears to be the limit of what any such method can give, and so far as I am aware results in 8 or fewer variables are known only for particular types of form with some degenerate and/or diagonal behaviour. For diagonal forms, 5 variables suffice by work of Hua [Hua38]. If one is content with almost-primes instead of primes, the number of variables can be reduced to 3: see [BGS10] and subsequent works.

Future work. In our main theorem we imposed conditions on Q, namely that $\det ab \neq 0$ and the matrix Δ has four distinct eigenvalues in $\overline{\mathbf{Q}} \setminus \{0, -1\}$. Whilst these are not especially severe restrictions, it nonetheless seems to be of interest to weaken them as far as possible, and we intend to address this in future work.

There are at least three paths to pursue in this direction. First, there are 35 essentially different ways to split 8 variables into two groups of 4, which one would expect to lead to Δs with different properties. Second, many of the arguments of the paper can be modified to work in more degenerate situations. Finally, in some highly degenerate situations Theorem 1.1 can be established by classical methods such as those in [Zha16]. One would expect this to be the case when $\det b = 0$ for all splittings of the 8 variables (the case of low "off-diagonal rank"). We anticipate this to be a somewhat tedious endeavour, with all the main ideas already present in this paper and [Zha16].

On a different matter, allowing linear terms in Q (that is, nonhomogeneous quadratics) is probably possible but seems to require a fairly significant modification of the method, and we will not attempt this here.

Notation. Most of our notation is fairly standard. We write $e(t) = e^{2\pi it}$, and for q a positive integer we write $e_q(x) := e^{2\pi ix/q}$. If S is a finite set and $F : S \to \mathbf{C}$ a function, we write $\mathbb{E}_{x\in S}F(x)$ to mean the average of F over S. We write $\mathbf{T} = \mathbf{R}/\mathbf{Z}$, and we write $\|x\|_{\mathbf{T}}$ for the distance from x to the nearest integer. We write $[X]$ to denote the discrete interval $\{1, \dots, X\}$, and $[0, X]$ for the continuous interval $\{x : 0 \leqslant x \leqslant X\} \subset \mathbf{R}$.

If G is a finite group then we write $\ell^2(G)$ for the vector space of all functions $f : G \to \mathbf{C}$ together with the inner product $\langle f_1, f_2\rangle := \mathbb{E}_{x\in G} f_1(x)\overline{f_2(x)}$ and the associated norm $\|f\|_2 := \sqrt{\langle f, f\rangle}$. Later on in the paper we will also define norms of probability measures, and we caution the reader that there we will use a different normalisation.

If V is a Hermitian inner product space (such as $\ell^2(G)$) then we write $\mathrm{U}(V)$ for the group of unitary transformations of V.

For p a prime, we will freqently encounter the group $\mathbf{Z}/p\mathbf{Z}$. When this arises as a group or a ring, we will write it $\mathbf{Z}/p\mathbf{Z}$, but when it is important that it is a field, we will write $\mathbf{F}_p$. This may seems slightly eccentric, but it does not seem stylistically correct to talk about homomorphisms from $\mathbf{Z}/q\mathbf{Z}$ to $\mathbf{F}_p$ (when $p|q$) and nor does it seem right to discuss the field $(\mathbf{Z}/p\mathbf{Z})(\theta)$ or the algebraic closure $\overline{\mathbf{Z}/p\mathbf{Z}}$. At times the distinction is somewhat arbitrary.

We use asymptotic notation $\ll$, $\gg$ and $O(\cdot)$ in the standard way. Thus, for example, $X \ll Y$ means that there is an absolute constant C such that $|X| \leqslant CY$, and $X = O(Y)$ means the same thing. Different instances of the notation may imply different constants C. If the constant depends on some underlying parameter A (as in Sect. 3 for example) then we indicate this with a subscript.

Throughout the paper we abuse notation in certain standard ways which should not cause any confusion. For example, we also write $Q(x, y)$ for the quadratic form $x^T ax + x^T by + y^T cy$ over $\mathbf{Z}/q\mathbf{Z}$, by which we mean that $x, y \in (\mathbf{Z}/q\mathbf{Z})^4$ and a, b and c are to be considered (mod q). Slightly more subtly, we also consider Δ (defined in (1.1) as an element of $\mathrm{Mat}_4(\mathbf{Q})$) as an element of $\mathrm{Mat}_4(\mathbf{Z}/p\mathbf{Z})$, which makes sense provided $p \nmid \det b$. Similarly, we consider a certain 8-by-8 symplectic matrix g (defined in (2.3) below) as an element of $\mathrm{Sp}_8(\mathbf{Z}/q\mathbf{Z})$ for squarefree q, which again will make sense provided no prime factor of q divides $\det b$.

Finally, we regard the quadratic form Q as fixed throughout the paper and will not explicitly indicate dependence on Q in asymptotic notation such as $\ll$ or $O()$.

2 Outline of the argument

The initial steps of the argument proceed in the classical fashion using the circle method, which we set up in Sect. 3. We introduce the exponential sum

$$S(\theta) := \sum_{x,y \in [X]^4} \Lambda^{\otimes 4}(x)\Lambda^{\otimes 4}(y) e(\theta Q(x, y)),$$

where $\Lambda^{\otimes 4}(x)$ is shorthand for $\Lambda(x_1)\Lambda(x_2)\Lambda(x_3)\Lambda(x_4)$, and of course Λ denotes the von Mangoldt function. Then by orthogonality we have

$$\sum_{\substack{x,y \in [X]^4 \\ Q(x,y)=N}} \Lambda^{\otimes 4}(x)\Lambda^{\otimes 4}(y) = \int_{\mathbf{T}} S(\theta) e(-\theta N) d\theta.$$

We divide $\mathbf{T}$ into the major arcs $\mathfrak{M}$ (roughly, the set of θ within distance $\sim X^{-2}$ from a rational $\frac{a}{q}$ with $q \ll \log^{O(1)} X$) and the minor arcs $\mathfrak{m}$. The major arcs give the main term in the asymptotic, and the analysis of them is entirely classical. We give this analysis in Sect. 4, referring to [Zha16] for the details when possible.

For the minor arcs, the fact that we are discussing primes is essentially irrelevant and the same arguments work with $\Lambda^{\otimes 4}(x)\Lambda^{\otimes 4}(y)$ replaced by $F_1(x)F_2(y)$ for any reasonably bounded functions F_1, F_2. We in fact divide the minor arcs into two sets $\mathfrak{m}_1$ and $\mathfrak{m}_2$, with $\mathfrak{m}_1$ being points not too close to a rational and $\mathfrak{m}_2$ being the points very close to a rational (but with moderately large denominator). The precise definitions are given at the start of Sect. 3. The treatment of the integral over $\mathfrak{m}_1$ uses diophantine approximation arguments standard in the area, and is given in Sect. 5.

The treatment of the minor arcs $\mathfrak{m}_2$ is the heart of the paper. One may reduce to considering actual rational points $\frac{r}{q}$, with $q > \log^C X$ moderately large, and one

is then led naturally led to look at exponential sums of the form

$$T_{f_1,f_2}(r) := q^2\mathbb{E}_{x,y\in(\mathbf{Z}/q\mathbf{Z})^4} f_1(x)f_2(y)e_q(rQ(x,y)), \tag{2.1}$$

where here $r \in (\mathbf{Z}/q\mathbf{Z})^*$. There is a "trivial" upper bound of 1 for such sums when $\|f_1\|_2, \|f_2\|_2 \leqslant 1$, which turns out to be (just) not good enough for the purposes of bounding the integral over $\mathfrak{m}_2$. However, any improvement of it by a factor $q^{-\delta}$ would suffice.

Unfortunately, there is no such improvement: the trivial bound is best possible. However, by a less wasteful reduction we can arrange things so that we consider instead the averages

$$\frac{1}{\phi(q)} \sum_{r\in(\mathbf{Z}/q\mathbf{Z})^*} |T_{f_1,f_2}(r)|. \tag{2.2}$$

Again, a saving of $q^{-\delta}$ over the trivial bound of 1 would be enough.

We incorporate some additional tricks which allow us to restrict attention to the case q squarefree and without very small prime factors, two features which are vital in our later arguments. The details of these reductions are given in Sect. 6.

The remainder of the paper is devoted to establishing a nontrivial bound of the required strength for averages (2.2). To make progress on this problem, we interpret the exponential sums $T_{f_1,f_2}(r)$ as matrix coefficients $\langle f_1, \rho(g^{(r)})f_2\rangle$, where here $\rho : \mathrm{Sp}_8(\mathbf{Z}/q\mathbf{Z}) \to \mathrm{U}(\ell^2((\mathbf{Z}/q\mathbf{Z})^4))$ is a certain unitary representation of the symplectic group Sp_8 over $\mathbf{Z}/q\mathbf{Z}$ called the Weil representation. After a brief introduction to the symplectic group and the Weil representation, we give this interpretation in Sect. 7. Whilst the theory of the Weil representation is well-known over $\mathbf{R}$ and somewhat well-known over finite fields, we do not know of a good source for the theory we need over $\mathbf{Z}/q\mathbf{Z}$, so we must develop some of this ourselves. This is fairly straightforward given the finite field statements, and is done in Appendix B.

The elements $g^{(r)} \in \mathrm{Sp}_8(\mathbf{Z}/q\mathbf{Z})$ are what we call "dilates" of a single element $g = g^{(1)}$ given by the formula

$$g := \begin{pmatrix} -2b^{-T}c & b^{-T} \\ 4ab^{-T}c - b & -2ab^{-T} \end{pmatrix}. \tag{2.3}$$

The *dilate* of $\left(\begin{smallmatrix} A & B \\ C & D \end{smallmatrix}\right) \in \mathrm{Sp}_8(\mathbf{Z}/q\mathbf{Z})$ by $r \in (\mathbf{Z}/q\mathbf{Z})^*$ is $\left(\begin{smallmatrix} A & r^{-1}B \\ rC & D \end{smallmatrix}\right)$; this is in fact an action of $(\mathbf{Z}/q\mathbf{Z})^*$ by automorphisms, as may be easily checked.

One is therefore led to the question of bounding an average of matrix coefficients $|\langle f_1, \rho(g^{(r)})f_2\rangle|$, where r ranges over $(\mathbf{Z}/q\mathbf{Z})^*$.

In Sect. 8 we supply a general tool for bounding averages of matrix coefficients, in principle applicable to any unitary representation ψ of any finite group G. This allows one to bound an average

$$\int_G |\langle f_1, \psi(x)f_2\rangle| d\mu(x),$$

where μ is a probability measure on G, when two conditions are satisfied:

(1) (convergence to uniform measure) Some symmetrised convolution power $\mu^\circ * \mu * \mu^\circ * \mu \cdots$ of bounded order should be close to the uniform measure on a subgroup $H \leqslant G$;
(2) (quasirandomness) $\psi|_H$ has no low-dimensional irreducible components.

We wish to apply this tool with $G = \mathrm{Sp}_8(\mathbf{Z}/q\mathbf{Z})$, $\psi = \rho$ being the Weil representation, and μ being the uniform measure on the $\phi(q)$ points $g^{(r)}$, $r \in (\mathbf{Z}/q\mathbf{Z})^*$. To do this we need to establish the convergence and quasirandomness properties.

The task (1) of showing that (symmetrised) convolution powers of μ converge to a uniform measure on a subgroup suggests the literature on the affine sieve, expanders and general measure convolutions in groups of Lie type, in particular the work of Varjú [Var12] which provides results in the appropriate generality. Some variant of this can probably be made to work in our context. However, our particular measure μ has a rather algebraic definition, being parametrised by (very simple) rational functions and we are able to offer an alternative approach using the Lang-Weil estimate. This is inspired by a blog post of Tao [Tao], giving an alternative proof (inspired by model-theoretic work of Pillay and Starchenko [PS] and unpublished notes of Hrushovski) of his own algebraic regularity lemma [Tao15]. This may be of independent interest, though we only develop it in the specific setting of interest to us here. This allows one to demonstrate rapid convergence of (symmetrised) powers of μ to the uniform measure on the group Γ_q generated by the elements $g^{-(r)}g^{(s)}$, without knowing *a priori* what this group is. The arguments may be found in Sect. 9.

The remaining task (2) is to establish the quasirandomness property for $\rho|_{\Gamma_q}$. It is easy to see that $\Gamma_q \cong \prod_{p|q} \Gamma_p$, and so by using the basic theory of tensor product representations it turns out to be enough to understand the case $q = p$ prime. First, we identify Γ_p explicitly. I was initially under the impression that the elements $g^{-(r)}g^{(s)}$ might generically generate the whole of $\mathrm{Sp}_8(\mathbf{Z}/p\mathbf{Z})$ (which has size $\sim p^{36}$), on the grounds that there is no immediately evident reason why they should not, and for the analogous situation in $\mathrm{Sp}_2(\mathbf{Z}/p\mathbf{Z})$ this is true. However, it turns out that this is not the case, and that Γ_p is (generically) a group of size $\sim p^{12}$, a conjugate (in $\mathrm{GL}_8(\mathbf{Z}/p\mathbf{Z})$) of $\mathrm{SL}_2(\mathbf{F}_p[\Delta])$, where here Δ is the matrix discriminant given in (1.1). Establishing this takes some work, involving calculations in SL_2 together with applications of lemmas of Goursat and Ribet on subgroups of direct products. A number of facts about SL_2 of a finite field are required here, and these are collated in Appendix A. These tasks are accomplished in Sect. 10.

With Γ_p identified explicitly, we turn to the quasirandomness property itself. It is essentially automatic from the representation theory of $\mathrm{SL}_2(\mathbf{F}_{p^n})$ that if $\rho|_{\Gamma_p}$ has an irreducible component of small degree, then this component must be the trivial representation: that is, Γ_p, acting via the Weil representation on $\ell^2((\mathbf{Z}/p\mathbf{Z})^4)$, would have a nontrivial fixed vector. The final task of the paper, then, is to rule this out. We do this in Sect. 11 using rather direct and explicit (that is, not using any representation theory) arguments.

To conclude this section we remark that classical methods only just fail for quadratic forms in 8 variables, as evidenced by the fact that it suffices to gain relatively small savings of $q^{-\delta}$ in our key arguments. To handle forms in 6 or 7 variables much larger savings would be required and so the methods of this paper do not handle those situations in their current form.

3 The circle method

In this section we describe the basic setup of the circle method. As is typical in problems of this type we will be aiming for error terms in our main theorem of $O_A(X^6 \log^{-A} X)$, for an arbitrary positive real number A. Fix such an A, without loss of generality $A \geqslant 10$, and set

$$M := \log^{C_1} X, \quad M' := \log^{C_2} X, \tag{3.1}$$

where $C_1 := 10A$ and $C_2 := 10^5 A\delta^{-1}$, where δ is the exponent appearing in Proposition 6.2 below (these choices are by no means optimal, but this is inconsequential).

Set

$$K := 8 \max_{ij} |b_{ij}|, \tag{3.2}$$

where the b_{ij} are the entries of the 4-by-4 matrix b. Thus K is a constant depending only on the quadratic form Q. For $q \in \mathbf{N}$ and for $r \in (\mathbf{Z}/q\mathbf{Z})^*$, denote

$$I_{r,q} := \Big\{\theta \in \mathbf{T} : \Big|\theta - \frac{r}{q}\Big| \leqslant \frac{M}{X^2}\Big\} \tag{3.3}$$

and

$$\tilde{I}_{r,q} := \Big\{\theta \in \mathbf{T} : \Big|\theta - \frac{r}{q}\Big| \leqslant \frac{1}{KqX}\Big\}. \tag{3.4}$$

Define the *major arcs*

$$\mathfrak{M} := \bigcup_{\substack{q \leqslant M' \\ r \in (\mathbf{Z}/q\mathbf{Z})^*}} I_{r,q}, \tag{3.5}$$

and set

$$\mathfrak{m}_1 := \bigcup_{\substack{q \leqslant KX \\ r \in (\mathbf{Z}/q\mathbf{Z})^*}} (\tilde{I}_{r,q} \setminus I_{r,q}), \quad \mathfrak{m}_2 := \bigcup_{\substack{M' < q \leqslant KX \\ r \in (\mathbf{Z}/q\mathbf{Z})^*}} I_{r,q}. \tag{3.6}$$

The sets $\mathfrak{m}_1$, $\mathfrak{m}_2$ are two different types of *minor arc*, with $\mathfrak{m}_1$ being points somewhat close to rationals with moderate denominator, but not close enough to qualify for inclusion in $\mathfrak{M}$, and $\mathfrak{m}_2$ being points very close to rationals with moderate denominator, but with the denominator not small enough to qualify for inclusion in $\mathfrak{M}$.

LEMMA 3.1. *We have* $\mathfrak{M} \cup \mathfrak{m}_1 \cup \mathfrak{m}_2 = \mathbf{T}$.

Proof. By Dirichlet's theorem on diophantine approximation,

$$\mathbf{T} = \bigcup_{\substack{q \leqslant KX \\ r \in (\mathbf{Z}/q\mathbf{Z})^*}} \tilde{I}_{r,q}.$$

The result then follows immediately. □

Set

$$S(\theta) := \sum_{x,y \in [X]^4} \Lambda^{\otimes 4}(x) \Lambda^{\otimes 4}(y) e(\theta Q(x,y)), \tag{3.7}$$

where, recall, $\Lambda^{\otimes 4}(x)$ is a convenient shorthand for $\prod_{i=1}^4 \Lambda(x_i)$. Then by orthogonality we have

$$\sum_{\substack{x,y \in [X]^4 \\ Q(x,y)=N}} \Lambda^{\otimes 4}(x) \Lambda^{\otimes 4}(y) = \int_{\mathbf{T}} S(\theta) e(-\theta N) d\theta. \tag{3.8}$$

We evaluate this by considering the contributions to the integral from $\mathfrak{M}$, $\mathfrak{m}_1$, $\mathfrak{m}_2$ separately. The major arcs $\mathfrak{M}$ give the main term in the asymptotic, as the following result shows.

PROPOSITION 3.2 (Major arcs). *Suppose that Q is regular, that is to say $\left(\begin{smallmatrix} a & b/2 \\ b^T/2 & c \end{smallmatrix}\right)$ is nonsingular. Then we have*

$$\int_{\mathfrak{M}} S(\theta) e(-\theta N) d\theta = \mathfrak{S}(N) X^6 + O_A(X^6 \log^{-A} X),$$

where the singular series $\mathfrak{S}(N)$ is as described in Theorem 1.1.

We will prove this in the next section using classical methods, referring to [Zha16] for most of the details.

Now we turn to the minor arcs $\mathfrak{m}_1$ and $\mathfrak{m}_2$. Here, as previously remarked, the fact that we are dealing with primes and the von Mangoldt function is essentially irrelevant. For any functions $F_1, F_2 : [X] \to \mathbf{C}$ we introduce the sums

$$S_{F_1,F_2}(\theta) := \sum_{x,y \in [X]^4} F_1(x) F_2(y) e(\theta Q(x,y)). \tag{3.9}$$

PROPOSITION 3.3 (Minor arcs $\mathfrak{m}_1$). *Suppose that* $\det b \neq 0$. *Then we have*

$$\int_{\mathfrak{m}_1} |S_{F_1,F_2}(\theta)| d\theta \ll_A X^6 \log^{-A-8} X,$$

uniformly for all 1-bounded functions F_1, F_2.

We will prove this in Sect. 5, using diophantine approximation arguments typical of the circle method.

PROPOSITION 3.4 (Minor arcs $\mathfrak{m}_2$). *Suppose that Q is generic (that is, Δ has four distinct eigenvalues in $\overline{\mathbf{Q}} \setminus \{0,-1\}$). Then we have*

$$\int_{\mathfrak{m}_2} |S_{F_1,F_2}(\theta)|d\theta \ll_A X^6 \log^{-A-8} X,$$

uniformly for all 1-bounded functions F_1, F_2.

This proof of this, which is a substantial undertaking, contains the new ideas of the paper and occupies the remaining sections.

Let us conclude this section by remarking that Propositions 3.2, 3.3 and 3.4 easily combine to establish Theorem 1.1. Indeed, by (3.8) and Proposition 3.2 we have

$$\Big| \sum_{\substack{x,y\in[X]^4\\ Q(x,y)=N}} \Lambda^{\otimes 4}(x)\Lambda^{\otimes 4}(y) - \mathfrak{S}(N)X^6 \Big| \leqslant O_A(X^6 \log^{-A} X) + \Big| \int_{\mathbf{T}\setminus\mathfrak{M}} S(\theta)e(-\theta N)d\theta \Big|.$$

By the triangle inequality and Lemma 3.1, the second term on the right is bounded above by

$$\int_{\mathfrak{m}_1} |S(\theta)|d\theta + \int_{\mathfrak{m}_2} |S(\theta)|d\theta.$$

By Propositions 3.3 and 3.4 (taking $F_1 = F_2 = (\log X)^{-4}\Lambda^{\otimes 4}$), both of these terms are bounded by $O_A(X^6 \log^{-A} X)$.

4 The major arcs

In this section we establish Proposition 3.2. The argument is very classical and in fact large portions of it may be simply quoted from [Zha16]. For this part of the argument, similar results hold with as few as 5 variables. Define

$$C(q,r) := \sum_{\substack{x,y\in(\mathbf{Z}/q\mathbf{Z})^4\\ (x,q)=(y,q)=1}} e_q(rQ(x,y)), \tag{4.1}$$

$$B_{Q,N}(q) := \frac{1}{\phi(q)^8} \sum_{r\in(\mathbf{Z}/q\mathbf{Z})^*} C(q,r)e_q(-rN), \tag{4.2}$$

$$\mathfrak{B}_{Q,N} := \sum_{q=1}^{\infty} B_{Q,N}(q), \tag{4.3}$$

$$I(t) := \int_{u,v\in[0,X]^4} e(tQ(u,v))dudv \tag{4.4}$$

and

$$\mathfrak{J}_{Q,N}(X) := \int_{-\infty}^{\infty} I(t)e(-Nt). \tag{4.5}$$

These are the same definitions as those in [Zha16, Sect. 3], with some notational substitutions (Zhao's f, t, β, $\mathfrak{S}(f,t)$ become our Q, N, t, $\mathfrak{B}_{Q,N}$ respectively). Also, we notate our quadratic forms using two variables x, y. Definitions like these will be familiar to anyone with knowledge of the circle method. The following is [Zha16, Lemma 3.6].

LEMMA 4.1 (Major arcs). *We have*

$$\int_{\mathfrak{M}} S(\theta)e(-\theta N)d\theta = \mathfrak{B}_{Q,N}\mathfrak{J}_{Q,N}(X) + O_A(X^6 \log^{-A} X). \tag{4.6}$$

Remarks. Our choice of C_1 and C_2 in the definition (3.1) of M, M' ensures that our major arcs are amply wide enough that the error term in [Zha16, Lemma 3.6] is $O_A(X^6 \log^{-A} X)$. The only consequence of taking the major arcs this wide is that the choice of exponent in the error term of the Siegel-Walfisz theorem towards the end of [Zha16, Sect. 3] (which is, in any case, not made explicit there) must be larger.

There is one further inconsequential difference between our setup and that in [Zha16]. In [Zha16] the major arc about $\frac{r}{q}$ has width $\frac{M}{qX^2}$, whereas we have taken the width to be $\frac{M}{X^2}$. The only other tiny change required is in (3.17), (3.19) of [Zha16] where the integrals should be taken over our slightly longer range $|\beta| \leqslant M/X^2$ (which actually helps slightly).

To reconcile this with Proposition 3.2 we must express $\mathfrak{B}_{Q.N}$ and $\mathfrak{J}_{Q,N}(X)$ in terms of the local densities β_p, β_∞, whose definitions are given in the statement of Theorem 1.1. This is again a standard endeavour, but it is not done in Zhao's paper so we give brief details now.

Recall the definition (1.3) of $\beta_{p,n}(N)$. By orthogonality, we have

$$\beta_{p,n}(N) = p^{-8n} \sum_{r \in \mathbf{Z}/p^n\mathbf{Z}} \sum_{x,y \in (\mathbf{Z}/p^n\mathbf{Z})^4} \Lambda^{\otimes 4}_{\mathbf{Z}/p\mathbf{Z}}(x)\Lambda^{\otimes 4}_{\mathbf{Z}/p\mathbf{Z}}(y)e_{p^n}(r(Q(x,y)-N)).$$

In the sum over r, write $r = p^{n-j}r'$ with $(r',p) = 1$. One may check that the contribution from a particular j is $B_{Q,N}(p^j)$, and so

$$\beta_{p,n}(N) = \sum_{j=0}^{n} B_{Q,N}(p^j).$$

Taking the limit as $n \to \infty$ gives

$$\beta_p(N) = \sum_{j=0}^{\infty} B_{Q,N}(p^j).$$

Finally, since $B_{Q,N}(q)$ is a multiplicative function of q (see [Zha16, Lemma 3.1]) we have

$$\prod_p \beta_p(N) = \sum_{q=1}^{\infty} B_{Q,N}(q) = \mathfrak{B}_{Q,N}. \tag{4.7}$$

There are, of course, convergence issues to be dealt with here, but these are fully fleshed out in [Zha16, Lemma 3.4].

To handle the archimedean factor β_∞, we proceed is as follows (we leave detailed analytic justifications to the reader). For $\varepsilon > 0$, set

$$f_\varepsilon(w) := \frac{1}{2\varepsilon}\mu_{\mathbf{R}^8}\Big\{u, v \in \mathbf{R}^4 : u, v \in [X]^4, w - \varepsilon \leqslant Q(u, v) \leqslant w + \varepsilon\Big\}.$$

Fourier inversion then gives

$$f_\varepsilon(N) = \int_{-\infty}^{\infty} \hat{f}_\varepsilon(t)e(Nt)dt.$$

However,

$$I(t) = \lim_{\varepsilon\to 0} \hat{f}_\varepsilon(-t).$$

Taking limits as $\varepsilon \to 0$ (and substituting $x := \frac{u}{X}$, $y := \frac{v}{X}$, $\delta := \frac{\varepsilon}{X^2}$ and using the homogeneity of Q) gives

$$\mathfrak{J}_{Q,N}(X) = \beta_\infty X^6. \tag{4.8}$$

Substituting (4.7) and (4.8) into Lemma 4.1 gives Proposition 3.2.

5 Minor arcs: the integral over $\mathfrak{m}_1$

In this section we prove Proposition 3.3. The reader may wish to recall the definitions of $I_{r,q}$, $\tilde{I}_{r,q}$ and $\mathfrak{m}_1$, which are (3.3), (3.4) and (3.6) respectively.

Proof of Proposition 3.3. Observe that $\tilde{I}_{r,q} \setminus I_{r,q}$ is empty if $q > X/KM$, and therefore

$$\mu_{\mathbf{T}}(\mathfrak{m}_1) \leqslant \sum_{q\leqslant X/KM} \sum_{r\in(\mathbf{Z}/q\mathbf{Z})^*} \frac{2}{KqX} < \frac{1}{M}. \tag{5.1}$$

Write

$$S_{F_1,F_2}(\theta) := \sum_{x,y\in[X]^4} F_1'(x)F_2'(y)e(\theta x^T by) \tag{5.2}$$

where $F_1'(x) := F_1(x)e(\theta x^T ax)$, $F_2'(y) := F_2(y)e(\theta y^T cy)$. By Cauchy–Schwarz,

$$|S_{F_1,F_2}(\theta)| \leqslant X^2\Big(\sum_{y,y'\in[X]^4} F_2'(y)\overline{F_2'(y')} \sum_{x\in[X]^4} e(\theta x^T b(y-y'))\Big)^{1/2}.$$

By the standard estimate

$$\Big|\sum_{x\in[X]} e(\xi x)\Big| \ll \min(X, \|\xi\|_{\mathbf{T}}^{-1})$$

and since $\|F_2\|_\infty \leqslant 1$ it follows that

$$|S_{F_1,F_2}(\theta)| \ll X^2\Big(\sum_{y,y'\in[X]^4}\prod_{j=1}^{4}\min\Big(X, \big\|\theta(b(y-y'))_j\big\|_{\mathbf{T}}^{-1}\Big)\Big)^{1/2}. \tag{5.3}$$

Now the image of $[X]^4\times[X]^4$ under the map $(y,y')\mapsto b^T(y-y')$ is contained in the box $[-\frac{1}{2}KX,\frac{1}{2}KX]^4$ and the fibres are of size at most X^4 (since b is nonsingular; recall also from the definition (3.2) that $K = 8\max_{ij}|b_{ij}|$). It follows from (5.3) that

$$|S_{F_1,F_2}(\theta)| \ll X^4\Big(\sum_{\substack{h\in\mathbf{Z}\\ |h|\leqslant KX/2}}\min\Big(X, \|\theta h\|_{\mathbf{T}}^{-1}\Big)\Big)^2. \tag{5.4}$$

Suppose now that $\theta\in\tilde{I}_{r,q}\setminus I_{r,q}$, thus

$$\theta = \frac{r}{q}+\eta \quad\text{where}\quad \frac{M}{X^2} < |\eta| \leqslant \frac{1}{KqX}.$$

Foliating into progressions modulo q we have

$$\sum_{|h|\leqslant KX/2}\min(X,\|\theta h\|_{\mathbf{T}}^{-1}) = \sum_{s(\mathrm{mod}\,q)}\sum_{\substack{|h|\leqslant KX/2\\ h\equiv s(\mathrm{mod}\,q)}}\min\Big(X, \Big\|\Big(\frac{r}{q}+\eta\Big)h)\Big\|_{\mathbf{T}}^{-1}\Big). \tag{5.5}$$

We evaluate the contributions from $s=0$ and $s\neq 0$ separately. If $s\neq 0$, $h\equiv s(\mathrm{mod}\,q)$ and $|h|\leqslant KX/2$ then

$$\Big\|\Big(\frac{r}{q}+\eta\Big)h\Big\|_{\mathbf{T}} \geqslant \Big\|\frac{rs}{q}\Big\|_{\mathbf{T}} - \frac{KX}{2}|\eta| \geqslant \Big\|\frac{rs}{q}\Big\|_{\mathbf{T}} - \frac{1}{2q}.$$

Thus, if $s\neq 0$,

$$\sum_{\substack{|h|\leqslant KX/2\\ h\equiv s(\mathrm{mod}\,q)}}\min\Big(X, \Big\|\Big(\frac{r}{q}+\eta\Big)h\Big\|_{\mathbf{T}}^{-1}\Big) \ll \frac{X}{q}\Big(\Big\|\frac{rs}{q}\Big\|_{\mathbf{T}} - \frac{1}{2q}\Big)^{-1}. \tag{5.6}$$

(Recall here that $q\leqslant KX$, so the number of terms in the sum over h is indeed $\ll X/q$.) Now as s ranges over $(\mathbf{Z}/q\mathbf{Z})\setminus\{0\}$, so does rs. Thus

$$\sum_{\substack{s(\mathrm{mod}\,q)\\ s\neq 0}}\Big(\Big\|\frac{rs}{q}\Big\|_{\mathbf{T}} - \frac{1}{2q}\Big)^{-1} = \sum_{\substack{s(\mathrm{mod}\,q)\\ s\neq 0}}\Big(\Big\|\frac{s}{q}\Big\|_{\mathbf{T}} - \frac{1}{2q}\Big)^{-1} \ll q\log q.$$

Substituting into (5.6), we see that the contribution to the right-hand side of (5.5) from the terms with $s \neq 0$ is $\ll X \log q = O(X \log X)$.

Now we look at the contribution to the right-hand side of (5.5) from $s = 0$. Making the substitution $h = kq$, this is

$$\sum_{|k| \leqslant KX/2q} \min(X, \|\eta kq\|_{\mathbf{T}}^{-1}). \tag{5.7}$$

We have

$$|\eta kq| \leqslant \frac{1}{KqX} \cdot \frac{KX}{2q} \cdot q < \frac{1}{2},$$

so $\|\eta kq\|_{\mathbf{T}} = |\eta kq|$. Therefore (5.7) is

$$\sum_{|k| \leqslant KX/2q} \min(X, |\eta kq|^{-1}) \leqslant X + \sum_{0 < |k| < KX/2q} |\eta kq|^{-1} \ll X + \frac{1}{\eta q} \log X.$$

Substituting these bounds for $s \neq 0$ and $s = 0$ into (5.5), we obtain

$$\sum_{|h| \leqslant KX/2} \min(X, \|\theta h\|_{\mathbf{T}}^{-1}) \ll \Big(X + \frac{1}{\eta q}\Big) \log X.$$

Substituting into (5.4) gives, for $\theta = \frac{r}{q} + \eta \in \tilde{I}_{r,q} \setminus I_{r,q}$,

$$|S_{F_1,F_2}(\theta)| \ll X^6 \log^2 X + \frac{X^4 \log^2 X}{\eta^2 q^2}.$$

To complete the proof of Proposition 3.3, we need to integrate this estimate over $\theta \in \mathfrak{m}_1$, that is to say over all $\tilde{I}_{r,q} \setminus I_{r,q}$ with $q \leqslant KX$ and $r \in (\mathbf{Z}/q\mathbf{Z})^*$. The contribution from the first term $X^6 \log^2 X$ is at most $X^6 M^{-1} \log^2 X$ by (5.1). The contribution from the second term is

$$\begin{aligned} &\ll X^4 \log^2 X \sum_{q \leqslant KX} \frac{1}{q^2} \sum_{r \in (\mathbf{Z}/q\mathbf{Z})^*} \int_{M/X^2}^{1/KqX} \eta^{-2} d\eta \\ &\leqslant X^4 \log^2 X \sum_{q \leqslant KX} \frac{1}{q} \frac{X^2}{M} \ll X^6 M^{-1} \log^3 X. \end{aligned}$$

Recalling that $M = \log^{C_1} X$ with $C_1 = 10A$, this completes the proof. □

6 The integral over $\mathfrak{m}_2$ – first reductions

We now begin the lengthy task of establishing Proposition 3.4. Once again, the reader may wish to begin by recalling the pertinent definitions, which are those of $I_{r,q}$ (see (3.3)), $\mathfrak{m}_2$ (see (3.6)) and $S_{F_1,F_2}(\theta)$ (given in (3.9)).

At the heart of our analysis will be certain complete exponential sums T_{f_1,f_2}. Let q be a positive integer, and suppose that $f_1, f_2 : (\mathbf{Z}/q\mathbf{Z})^4 \to \mathbf{C}$. Define

$$T_{f_1,f_2}(r) := q^2 \mathbb{E}_{x,y\in(\mathbf{Z}/q\mathbf{Z})^4} f_1(x) f_2(y) e_q(rQ(x,y)). \tag{6.1}$$

Remark. Of course, T is also depends on q, but we omit explicit mention of this from the notation. There should not be any danger of confusion. For fixed r and general f_1, f_2 we have the following bound.

LEMMA 6.1. *Suppose that* $f_1, f_2 : (\mathbf{Z}/q\mathbf{Z})^4 \to \mathbf{C}$. *Suppose that* $\det b \neq 0$. *Then for any* $r \in (\mathbf{Z}/q\mathbf{Z})^*$ *we have*

$$|T_{f_1,f_2}(r)| \leqslant (q, \det b)^2 \|f_1\|_2 \|f_2\|_2. \tag{6.2}$$

Proof. Modifying $f_1(x)$ to $f_1(x)e_q(rx^T ax)$ and $f_2(y)$ to $f_2(y)e_q(ry^T cy)$, it suffices to show that

$$\mathbb{E}_{x,y\in(\mathbf{Z}/q\mathbf{Z})^4} f_1(x) f_2(y) e_q(rx^T by) \leqslant (q, \det b)^2 q^{-2} \|f_1\|_2 \|f_2\|_2.$$

By Cauchy-Schwarz, it suffices to show that

$$\mathbb{E}_{y,y'\in(\mathbf{Z}/q\mathbf{Z})^4} f_2(y)\overline{f_2(y')} \mathbb{E}_{x\in(\mathbf{Z}/q\mathbf{Z})^4} e_q\Big(rx^T b(y-y')\Big) \leqslant (q, \det b)^4 q^{-4} \|f_2\|_2^2.$$

By orthogonality, and since $(r, q) = 1$, the left-hand side is

$$\begin{aligned}\mathbb{E}_{y,y'\in(\mathbf{Z}/q\mathbf{Z})^4} f_2(y)\overline{f_2(y')} 1_{b(y-y')\equiv 0 (\mathrm{mod}\, q)} \\ = q^{-4} \sum_{\substack{h\in(\mathbf{Z}/q\mathbf{Z})^4 \\ bh\equiv 0(\mathrm{mod}\, q)}} \mathbb{E}_{y\in(\mathbf{Z}/q\mathbf{Z})^4} f_2(y)\overline{f_2(y+h)}.\end{aligned}$$

By Cauchy–Schwarz this is at most

$$q^{-4} \#\Big\{h \in (\mathbf{Z}/q\mathbf{Z})^4 : bh \equiv 0 (\mathrm{mod}\ q)\Big\} \|f_2\|_2^2,$$

and so it is enough to show that

$$\#\Big\{h \in (\mathbf{Z}/q\mathbf{Z})^4 : bh \equiv 0 (\mathrm{mod}\ q)\Big\} \leqslant (q, \det b)^4. \tag{6.3}$$

Now if $bh \equiv 0 (\mathrm{mod}\ q)$ then, multiplying on the left by $\mathrm{adj}(b)$, we have $(\det b)h \equiv 0 (\mathrm{mod}\ q)$, i.e. if $h = (h_1, h_2, h_3, h_4)$ then $q | (\det b) h_i$. The number of choices of each h_i is therefore $(q, \det b)$ and so (6.3) follows. This concludes the proof of (6.2). □

The bound in Lemma 6.1 is best possible, at least when $(q, \det b) = 1$. To see this, let $\psi : (\mathbf{Z}/q\mathbf{Z})^4 \to \mathbf{C}$ be any function with $\|\psi\|_2 = 1$, and take

$$f_1(x) := q^2 \mathbb{E}_{y\in(\mathbf{Z}/q\mathbf{Z})^4} \overline{\psi}(y) e_q(-rQ(x,y)), \quad f_2(x) := \psi(x).$$

Then one may check using the orthogonality relations that

$$|T_{f_1,f_2}(r)| = \|f_1\|_2 = \|f_2\|_2 = 1. \tag{6.4}$$

A more conceptual explanation of this is as follows. First note that

$$T_{f_1,f_2}(r) = \overline{\langle \overline{f}_1, \Phi f_2\rangle}, \tag{6.5}$$

where the map $\Phi : \ell^2((\mathbf{Z}/q\mathbf{Z})^4) \to \ell^2((\mathbf{Z}/q\mathbf{Z})^4)$ is given by

$$\Phi f(x) := q^2 \mathbb{E}_{y\in(\mathbf{Z}/q\mathbf{Z})^4} f(y) e_q(rQ(x,y)).$$

One may then make the key observation that Φ is unitary (being a composition of invertible dilations, quadratic modulations and Fourier transform). Then we have $f_1 := \overline{\Phi\psi}$, $f_2 = \psi$, and the relations (6.4) are clear from (6.5) and the unitary nature of Φ.

This also allows a very short (albeit ultimately equivalent) proof of Lemma 6.1 in the case $(q, \det b) = 1$. Indeed, by Cauchy–Schwarz and unitarity we have

$$|T_{f_1,f_2}(r)| = |\langle \overline{f}_1, \Phi f_2\rangle| \leqslant \|f_1\|_2 \|\Phi f_2\|_2 = \|f_1\|_2 \|f_2\|_2.$$

One may, using arguments similar to those below, use the bound obtained in Lemma 6.1 to show that (roughly speaking)

$$\int_{\theta\in\bigcup_{r\in(\mathbf{Z}/q\mathbf{Z})^*} I_{r,q}} |S_{F_1,F_2}(\theta)| d\theta \lessapprox X^6 q^{-1}.$$

Unfortunately the sum over q does not converge and so this (just) fails to give the desired estimate Proposition 3.4. It is this, and the sharpness of Lemma 6.1, which ultimately explain the failure of the classical circle method to handle the problem of quadratic forms in 8 prime variables.

To get around this issue we introduce the following improvement on 6.1 when an average over r is included (at least when q is squarefree and has no small prime factors, and Q is generic).

PROPOSITION 6.2. *There is an absolute constant $\delta > 0$ with the following property. Suppose that Q is generic. Then there is $p_0(Q)$ such that if q is squarefree and with all prime factors greater than $p_0(Q)$, then we have*

$$\frac{1}{\phi(q)} \sum_{r\in(\mathbf{Z}/q\mathbf{Z})^*} |T_{f_1,f_2}(r)| \ll q^{-\delta} \|f_1\|_2 \|f_2\|_2$$

for any $f_1, f_2 \in \ell^2((\mathbf{Z}/q\mathbf{Z})^4)$.

The proof of this proposition occupies most of the rest of the paper. The remainder of this section is devoted to deriving Proposition 3.4 from it.

First we observe that Lemma 6.1 and Proposition 6.2 have a fairly straightforward application to the sums $S_{F_1,F_2}(\theta)$ for $\theta = \frac{r}{q}$, which we record now.

COROLLARY 6.3. *Suppose that $F_1, F_2 : [X]^4 \to \mathbf{C}$ are 1-bounded. Suppose that $\det b \neq 0$. Then we have the pointwise bound*

$$\left|S_{F_1,F_2}\left(\frac{r}{q}\right)\right| \ll X^8 q^{-2}. \tag{6.6}$$

Suppose additionally that q is squarefree and has no prime factors of size $\leqslant p_0(Q)$, and that Q is generic. Suppose that $q \leqslant KX$. Then

$$\sum_{r\in(\mathbf{Z}/q\mathbf{Z})^*} \left|S_{F_1,F_2}\left(\frac{r}{q}\right)\right| \ll X^8 q^{-1-\delta}. \tag{6.7}$$

Proof. Extend F_1, F_2 to functions on $\mathbf{Z}$ by defining them to be 0 outside of $[X]$. Fix $x_0, y_0 \in q\mathbf{Z}^4$. Then

$$\sum_{\substack{x\in x_0+[q]^4\\ y\in y_0+[q]^4}} F_1(x)F_2(y)e_q(rQ(x,y))$$
$$= \sum_{x',y'\in[q]^4} F_1(x_0+x')F_2(y_0+y')e_q(rQ(x_0+x',y_0+y')).$$

Since x_0, y_0 are both multiples of q we have $Q(x_0+x',y_0+y') \equiv Q(x',y')(\mathrm{mod}\ q)$ and therefore, recalling the definition (6.1),

$$\sum_{\substack{x\in x_0+[q]^4\\ y\in y_0+[q]^4}} F_1(x)F_2(y)e_q(rQ(x,y)) = q^6 T_{f_1,f_2}(r) \tag{6.8}$$

where

$$f_1(x') := F_1(x_0+x') \quad \text{and} \quad f_2(y') := F_2(y_0+y').$$

By Lemma 6.1 we therefore have

$$\Big|\sum_{\substack{x\in x_0+[q]^4\\ y\in y_0+[q]^4}} F_1(x)F_2(y)e_q(rQ(x,y))\Big| \leqslant q^6(\det b)^2.$$

Covering the range $[X]^4\times[X]^4$ by $\leqslant (\frac{X}{q}+1)^8 \ll X^8q^{-8}$ boxes of the form $(x_0+[q]^4)\times(y_0+[q]^4)$ gives (6.6).

To obtain (6.7), we instead apply Proposition 6.2 to (6.8), obtaining

$$\sum_{r\in(\mathbf{Z}/q\mathbf{Z})^*}\Big|\sum_{x\in x_0+[q]^4, y\in y_0+[q]^4} F_1(x)F_2(y)e_q(rQ(x,y))\Big| \ll q^{6-\delta}\phi(q) \leqslant q^{7-\delta}.$$

Dividing into $O(X^8q^{-8})$ boxes as before gives (6.7). □

Proof of Proposition 3.4, assuming Proposition 6.2. Recall that

$$\mathfrak{m}_2 = \bigcup_{\substack{M'\leqslant q\leqslant KX\\ r\in(\mathbf{Z}/q\mathbf{Z})^*}} I_{r,q},$$

with $I_{r,q}$ as defined in (3.3), $M' = \log^{C_2} X$ (with C_2 as described in (3.1)), and $K = 8\max_{ij}|b_{ij}|$ being a constant associated to the form Q.

Therefore the bound we are trying to prove is

$$\sum_{M'<q\leqslant KX}\sum_{r\in(\mathbf{Z}/q\mathbf{Z})^*}\int_{|\eta|\leqslant M/X^2}\Big|S_{F_1,F_2}\Big(\frac{r}{q}+\eta\Big)\Big|d\eta \ll_A X^6\log^{-A-8}X. \tag{6.9}$$

We begin by using some Fourier analysis to handle the inner integral over η. Recall the definition (3.9) of S_{F_1,F_2}, that is to say

$$S_{F_1,F_2}(\theta):=\sum_{x,y\in[X]^4}F_1(x)F_2(y)e(\theta Q(x,y)).$$

Let $w:\mathbf{R}^2\to\mathbf{R}$ be some fixed smooth compactly-supported function with $w(u,v)=1$ for $(u,v)\in[0,1]^2$ and set, for any real parameter λ,

$$W_\lambda(u,v):=w(u,v)e(\lambda Q(u,v)). \tag{6.10}$$

Then

$$S_{F_1,F_2}\Big(\frac{r}{q}+\eta\Big)=\mathbb{E}_{x,y\in[X]^4}F_1(x)F_2(y)e_q(rQ(x,y))W_{\eta X^2}\Big(\frac{x}{X},\frac{y}{X}\Big). \tag{6.11}$$

Now by integration by parts and Leibniz's rule we have

$$|\hat{W}_\lambda(\xi,\xi')|\ll|\xi|^{-2}|\xi'|^{-2}\Big\|\frac{\partial^4W_\eta}{\partial^2u\partial^2v}\Big\|_1\ll\max(1,|\lambda|^4)|\xi|^{-2}|\xi'|^{-2}.$$

Since we also have the trivial bound

$$|\hat{W}_\lambda(\xi,\xi')|\leqslant\|W_\lambda\|_1=\|w\|_1\ll 1$$

it follows that

$$\|\hat{W}_\lambda\|_1\ll\max(1,|\lambda|)^4. \tag{6.12}$$

By Fourier inversion

$$W_\lambda(u,v)=\int_{\mathbf{R}}\hat{W}_\lambda(\xi,\xi')e(\xi u+\xi' v)d\xi d\xi';$$

substituting into (6.11) gives

$$S_{F_1,F_2}\Big(\frac{r}{q}+\eta\Big)=\int\hat{W}_{\eta X^2}(\xi,\xi')S_{F_{1,\xi},F_{2,\xi'}}\Big(\frac{r}{q}\Big)d\xi d\xi',$$

where

$$F_{1,\xi}(x):=F_1(x)e(\xi x/X),\quad F_{2,\xi'}(y):=F_2(y)e(\xi' y/X).$$

Therefore

$$\sum_{M'<q\leqslant KX}\sum_{r\in(\mathbf{Z}/q\mathbf{Z})^*}\Big|S_{F_1,F_2}\Big(\frac{r}{q}+\eta\Big)\Big|\leqslant \int_{\mathbf{R}}|\hat{W}_{\eta X^2}(\xi,\xi')|\sum_{M'<q\leqslant KX}\sum_{r\in(\mathbf{Z}/q\mathbf{Z})^*}\Big|S_{F_{1,\xi},F_{2,\xi'}}\Big(\frac{r}{q}\Big)\Big|d\xi d\xi'. \tag{6.13}$$

We claim the estimate

$$\sum_{M'<q\leqslant KX}\sum_{r\in(\mathbf{Z}/q\mathbf{Z})^*}\Big|S_{F_1',F_2'}\Big(\frac{r}{q}\Big)\Big|\ll X^8(M')^{-\delta/100}, \tag{6.14}$$

uniformly for all 1-bounded F_1', F_2', where δ is the exponent appearing in Proposition 6.2. Assuming this claim, (6.12) and (6.13) then imply that

$$\sum_{M'<q\leqslant KX}\sum_{r\in(\mathbf{Z}/q\mathbf{Z})^*}\Big|S_{F_1,F_2}\Big(\frac{r}{q}+\eta\Big)\Big|\ll X^8(M')^{-\delta/100}\max(1,|\eta|X^2)^4.$$

Then, integrating over $|\eta|\leqslant M/X^2$ we obtain

$$\int_{\mathfrak{m}_2}|S_{F_1,F_2}(\theta)|d\theta\ll X^6M^5(M')^{-\delta/100}\ll_A X^6(\log X)^{-A-8}.$$

For the last step, we recall that we chose $M=\log^{C_1}X$, $M'=\log^{C_2}X$ with the particular choice of C_1, C_2 specified at the start of Sect. 3. This completes the proof of Proposition 3.4, assuming the claim (6.14).

Now we must establish (6.14). The pointwise bound (6.6) is not good enough, but we do have the improved average bound (6.7), albeit only for squarefree q with no small prime factors. Most q do not have this form, and so we need the following lemma to allow us to reduce matters to the consideration of to those that do.

Lemma 6.4. *Suppose that $q=q_0q_1$ with $(q_0,q_1)=1$. Then*

$$\sup_{F_1,F_2}\sum_{r\in(\mathbf{Z}/q\mathbf{Z})^*}\Big|S_{F_1,F_2}\Big(\frac{r}{q}\Big)\Big|\ll q_1^9\sup_{F_1,F_2}\sum_{r_0\in(\mathbf{Z}/q_0\mathbf{Z})^*}\Big|S_{F_1,F_2}\Big(\frac{r_0}{q_0}\Big)\Big|,$$

where in both cases the sup *is over 1-bounded functions $F_1,F_2:[X]^4\to\mathbf{C}$.*

Proof. By the Chinese remainder theorem we have

$$\sum_{r\in(\mathbf{Z}/q\mathbf{Z})^*}\Big|S_{F_1,F_2}\Big(\frac{r}{q}\Big)\Big|=\sum_{\substack{r_0\in(\mathbf{Z}/q_0\mathbf{Z})^*\\ r_1\in(\mathbf{Z}/q_1\mathbf{Z})^*}}\Big|S_{F_1,F_2}\Big(\frac{r_0}{q_0}+\frac{r_1}{q_1}\Big)\Big|. \tag{6.15}$$

Foliating into residue classes modulo q_1, we have

$$S_{F_1,F_2}\Big(\frac{r_0}{q_0}+\frac{r_1}{q_1}\Big) = \sum_{u,v\in(\mathbf{Z}/q_1\mathbf{Z})^4} e_{q_1}(r_1Q(u,v)) \sum_{x,y\in[X]^4} F_{1,u}(x)F_{2,v}(y)e_{q_0}(r_0Q(x,y))$$

where $F_{1,u}(x) := F_1(x)1_{x\equiv u(\mathrm{mod}\, q_1)}$, $F_{2,v}(y) := F_2(y)1_{y\equiv v(\mathrm{mod}\, q_1)}$ and so by (6.15)

$$\begin{aligned}\sum_{r\in(\mathbf{Z}/q\mathbf{Z})^*}&\Big|S_{F_1,F_2}\Big(\frac{r}{q}\Big)\Big| \\ &\leqslant q_1^9 \sup_{u,v} \sum_{r_0\in(\mathbf{Z}/q_0\mathbf{Z})^*}\Big|\sum_{x,y\in[X]^4} F_{1,u}(x)F_{2,v}(y)e_{q_0}(r_0Q(x,y))\Big| \\ &= q_1^9 \sup_{u,v} \sum_{r_0\in(\mathbf{Z}/q_0\mathbf{Z})^*}\Big|S_{F_{1,u},F_{2,v}}\Big(\frac{r_0}{q_0}\Big)\Big|.\end{aligned}$$

The lemma follows. □

Let us turn to the actual proof of (6.14). Let $p_0 = p_0(Q)$ be the threshold appearing in Proposition 6.2. For any q, write q_0 for the product of all primes $p > p_0$ which divide q precisely once, and set $q_1 := q/q_0$, thus q_1 is the product of all prime powers $p^j \| q$ with $p \leqslant p_0$ or $j \geqslant 2$. Note that q_0, q_1 are coprime. By Lemma 6.4 and (6.6), (6.7) we have

$$\sum_{r\in(\mathbf{Z}/q\mathbf{Z})^*}\Big|S_{F_1,F_2}\Big(\frac{r}{q}\Big)\Big| \ll X^8 \min(q^{-1}, q_1^9 q_0^{-1-\delta}) \leqslant X^8 \min(q^{-1}, q_1^{11} q^{-1-\delta}).$$

It therefore suffices to prove that

$$\sum_{q>M'} \min(q^{-1}, q_1^{11} q^{-1-\delta}) \ll (M')^{-\delta/100}. \tag{6.16}$$

The contribution from q with $q_1 < q^{\delta/22}$ is acceptable (using the second term in the min).

Suppose now that $q_1 \geqslant q^{\delta/22}$. Let the prime factorisation of q_1 be $\prod_p p^{v_p}$, thus we have $v_p \geqslant 2$ for $p \geqslant p_0$. Set $q_2 := \prod_p p^{\lfloor v_p/2\rfloor}$. Then $q_2^2 | q$. Moreover, if $v_p \geqslant 2$ then $p^{\lfloor v_p/2\rfloor} \geqslant p^{v_p/3}$, so $q_2 \geqslant cq^{\delta/66}$ (with $c > 0$ depending only on p_0). Thus the contribution of these q to (6.16) can be bounded by

$$\sum_{d>c(M')^{\delta/66}} \sum_{\substack{q\ll d^{66/\delta}\\ d^2|q}} q^{-1} \ll \sum_{d>c(M')^{\delta/66}} \frac{\log d}{d^2} \ll (M')^{-\delta/100}.$$

This concludes the proof. □

7 Exponential sums as matrix coefficients on $\mathrm{Sp}_8(\mathbf{Z}/q\mathbf{Z})$

The remainder of the paper is occupied with the proof of Proposition 6.2, the statement of which the reader may wish to recall at this point.

We remarked after the proof of Lemma 6.1 that we can write

$$T_{f_1,f_2}(r) = \overline{\langle \overline{f}_1, \Phi_r f_2\rangle}, \tag{7.1}$$

where

$$\Phi_r f = q^2 \mathbb{E}_y f(y) e_q(rQ(x,y))$$

is unitary. The crucial observation which drives our whole argument is that the subgroup of $\mathrm{U}(\ell^2((\mathbf{Z}/q\mathbf{Z})^4))$ (the group of all unitary operators on $\ell^2((\mathbf{Z}/q\mathbf{Z})^4)$) generated by operators Φ of this type (over all Q) is rather small. Indeed, as we shall shortly see, it has size $q^{36+o(1)}$ when q is squarefree. This means that the specific operators Φ_r (with Q fixed but r allowed to vary over $(\mathbf{Z}/q\mathbf{Z})^*$) already occupy a reasonable portion of this group.

This group turns out to be the symplectic group $\mathrm{Sp}_8(\mathbf{Z}/q\mathbf{Z})$. Let us recall what these groups are, in a very concrete way. Let $R = \mathbf{Z}/q\mathbf{Z}$ with q odd. Then $\mathrm{Sp}_8(R)$ is a group of 8×8 matrices over R, which we will write in 2×2 block form with each block being a 4×4 matrix.

Definition 7.1. We define the symplectic group $\mathrm{Sp}_8(R)$ to be the group consisting of all 8×8 block matrices $g = \left(\begin{smallmatrix} A & B \\ C & D\end{smallmatrix}\right)$ with entries in R and $A^TC = C^TA$, $B^TD = D^TB$ and $A^TD - C^TB = I$.

Define

$$J := \begin{pmatrix} 0 & I \\ -I & 0\end{pmatrix}. \tag{7.2}$$

It is an simple exercise to check that $g \in \mathrm{Sp}_8(R)$ if and only if $g^TJg = J$. In fact, this is the more usual definition of the sympletic group, but it suits us to be more explicit.

Note that if $g = \left(\begin{smallmatrix} A & B \\ C & D\end{smallmatrix}\right) \in \mathrm{Sp}_8(R)$ then g is left-invertible with left-inverse $\left(\begin{smallmatrix} D^T & -B^T \\ -C^T & A^T\end{smallmatrix}\right)$. This is then, of course, also a right-inverse for g, and this gives us the additional relations

$$AB^T = BA^T,\ CD^T = DC^T \quad \text{and} \quad AD^T - BC^T = I$$

for any symplectic matrix.

We will also need the fact that $(\mathbf{Z}/q\mathbf{Z})^*$ acts on $\mathrm{Sp}_8(\mathbf{Z}/q\mathbf{Z})$ by "dilation" automorphisms. If $g = \left(\begin{smallmatrix} A & B \\ C & D\end{smallmatrix}\right)$ and if $r \in (\mathbf{Z}/q\mathbf{Z})^*$ then we define

$$g^{(r)} = \begin{pmatrix} A & r^{-1}B \\ rC & D\end{pmatrix}. \tag{7.3}$$

It is then easy to see that this gives an action of $(\mathbf{Z}/q\mathbf{Z})^*$ on $\mathrm{Sp}_8(\mathbf{Z}/q\mathbf{Z})$ by automorphisms.

Finally, we note for future reference (see, for example, [Neu02]) that

$$|\mathrm{Sp}_8(\mathbf{F}_p)| = p^{16}(p^8-1)(p^6-1)(p^4-1)(p^2-1) = (1+o(1))p^{36}. \tag{7.4}$$

(From this, the earlier claim that $|\mathrm{Sp}_8(\mathbf{Z}/q\mathbf{Z})| = q^{36+o(1)}$ when q is squarefree follows from the fact that, by the Chinese remainder theorem, $\mathrm{Sp}_8(\mathbf{Z}/q\mathbf{Z}) \cong \prod_{p|q} \mathrm{Sp}_8(\mathbf{Z}/p\mathbf{Z})$.)

Now we come to the key link between Sp_8 and operators such as Φ_r in (7.1), which stems from [Wei64] and is thus known as the Weil representation (or, depending on the context, the Segal-Shale-Weil representation or the oscillator representation).

Proposition 7.2 (Weil representation). *Let q be squarefree and odd. Then there is a unitary representation*

$$\rho : \mathrm{Sp}_8(\mathbf{Z}/q\mathbf{Z}) \to \mathrm{U}(\ell^2((\mathbf{Z}/q\mathbf{Z})^4))$$

and a function $\xi : \mathrm{Sp}_8(\mathbf{Z}/q\mathbf{Z}) \to \{z \in \mathbf{C} : |z| = 1\}$ satisfying the following:

- (Dilations) *If $g = s(E)$ where $s(E) := \left(\begin{smallmatrix} E & 0 \\ 0 & E^{-T} \end{smallmatrix}\right)$ with E invertible then*
$$\rho(g)f(x) = \xi(s(E))f(E^{-1}x);$$
- (Fourier transform) *If $g = J$ then*
$$\rho(g)f(x) = \xi(J)q^2\mathbb{E}_{y\in(\mathbf{Z}/q\mathbf{Z})^4} f(y)e(x^Ty);$$
- (Quadratic modulations) *If $g = l(W)$ where $l(W) := \left(\begin{smallmatrix} I & 0 \\ W & I \end{smallmatrix}\right)$ with W symmetric then*
$$\rho(g)f(x) = \xi(l(W))e_q(-\frac{1}{2}x^TWx)f(x).$$

Remarks. 1. There is nothing special about Sp_8 here; similar results hold for Sp_{2m} for any positive integer m. The Weil representation is well-known over $\mathbf{R}$ (where one needs to pass to the double cover of the symplectic group), but in finite situations it seems to me that it is only at all widely discussed over finite fields. In this case, the construction is given in detail in (for example) the paper [Neu02] of Neuhauser. The analogue of this in the lower-dimensional setting of $\mathrm{Sp}_2(\mathbf{Z}/p\mathbf{Z}) = \mathrm{SL}_2(\mathbf{Z}/p\mathbf{Z})$ already contains the key ideas, and a very nice description of this may be found in the notes of Charlotte Chan [Cha], which I found helpful in preparing this material. It is not difficult to derive the case q squarefree from the prime case, and we do this in Appendix B.

2. The phase ξ can be given explicitly if desired. When $q = p$ is an odd prime, we can take $\xi(l(W)) = \xi(J) = 1$ and $\xi(s(E)) = \left(\frac{\det E}{p}\right)$, and the general squarefree case can then be deduced from the arguments in Appendix B. For details of these calculations in the case $q = p$ (which are somewhat involved) see [Neu02]. In this

paper, we will not need explicit values of ξ, and the mere existence is a much easier statement to prove, this being [Neu02, Theorem 4.3].

3. If desired one can also add in the translations $f(x) \mapsto f(x-v)$ and the linear modulations $f(x) \mapsto e(-t^T x)f(x)$, getting an action by the "Jacobi group" $\mathrm{Sp}_8(\mathbf{Z}/q\mathbf{Z}) \ltimes \mathrm{H}_8(\mathbf{Z}/q\mathbf{Z})$, where H_8 is the Heisenberg group on $(\mathbf{Z}/q\mathbf{Z})^8 \times \mathbf{Z}/q\mathbf{Z}$.

4. It is not really correct to call ρ "the" Weil representation. In the case $q = p$ a prime, further representations $\tilde{\rho}$ of the same dimension can be obtained by twisting with the dilation $\sigma(g) = g^{(r)}$, that is to say $\tilde{\rho}(g) := \rho(g^{(r)})$. When r is not a square in $\mathbf{F}_p^*$, the dilation is an outer automorphism and it is known that $\tilde{\rho} \not\cong \rho$. (In the literature this would be described in terms of different central characters on the Heisenberg group giving different Weil representations, see [Neu02, Sect. 7] or [Sze98, Proposition 4]). Thus there are two Weil representations of $\mathrm{Sp}_8(\mathbf{Z}/p\mathbf{Z})$, and the one we are considering is a concrete realisation of one of them.

When $q = p_1 \cdots p_n$, one may obtain 2^n non-isomorphic representations of $\mathrm{Sp}_8(\mathbf{Z}/q\mathbf{Z})$ by taking tensor products. The representation whose existence we assert in Proposition 7.2 is one of these. However, it turns out *not* to be simply the tensor product of the ρs associated to each p_i: we must first apply some twists. The details are given in Appendix B.

5. Even when $q = p$, the Weil representation is not irreducible. It splits into its actions on odd and even functions, which *are* irreducible representations of degrees $\frac{1}{2}(p^4 \pm 1)$. We will not need this fact here.

Now that we have defined the Weil representation, we can interpret the exponential sums $T_{f_1,f_2}(r)$ as matrix coefficients. We begin with an important definition which will be relevant for the rest of the paper.

Definition 7.3 (Symplectic element). Suppose that $Q(x,y) = x^T ax + x^T by + y^T cy$ is a quadratic form and that $\det b \neq 0$. Then we associate to Q the element $g = g(Q) \in \mathrm{Sp}_8(\mathbf{Q})$ defined by

$$g := \begin{pmatrix} -2b^{-T}c & b^{-T} \\ 4ab^{-T}c - b & -2ab^{-T} \end{pmatrix}.$$

We call this the *symplectic element* associated to Q.

Remark. This is not a standard term and we are not aware of any other appearance of this matrix in the literature.

As mentioned in the introduction, we will abuse notation by regarding g as an element of $\mathrm{Sp}_8(\mathbf{Z}/q\mathbf{Z})$ for squarefree odd q, coprime to $\det b$. Here is the promised interpretation of exponential sums as matrix coefficients.

Proposition 7.4. *Suppose that q is odd, squarefree and coprime to $\det b$. Then for any $f_1, f_2 \in \ell^2((\mathbf{Z}/q\mathbf{Z})^4)$ we have*

$$|T_{f_1,f_2}(r)| = |\langle \overline{f}_1, \rho(g^{(r)})f_2 \rangle|,$$

where $g^{(r)}$ is the dilate of g (regarded as an element of $\mathrm{Sp}_8(\mathbf{Z}/q\mathbf{Z})$) by r as defined in (7.3) and ρ is the Weil representation described in Proposition 7.2.

Proof. It suffices to establish the case $r = 1$, since then the formula for general r follows by applying that case with Q replaced by rQ (or, to be pedantic, $\overline{r}Q$ for some $r \in \mathbf{Z}$ projecting to $r (\mathrm{mod}\ q)$).

To handle the case $r = 1$, note that $q^2\mathbb{E}_{y\in(\mathbf{Z}/q\mathbf{Z})^4} f(y)e(Q(x,y))$ may be built up as a composition of four (unitary) operations, as follows:

(1) A quadratic modulation $f(y) \mapsto f(y)e_q(y^T cy)$;
(2) Fourier transform $f \mapsto q^2\mathbb{E}_{y\in(\mathbf{Z}/q\mathbf{Z})^4} f(y)e_q(x^T y)$;
(3) Dilation $f(x) \mapsto f(b^T x)$;
(4) Quadratic modulation $f(x) \mapsto f(x)e_q(x^T ax)$.

In the Weil representation these four operations correspond, up to scalar multiplication by unit complex numbers, respectively, to the following elements of $\mathrm{Sp}_8(\mathbf{Z}/q\mathbf{Z})$: $l(-2c)$, J, $s(b^{-T})$ and $l(-2a)$. Therefore by Proposition 7.2 (since ρ is a homomorphism!) we have

$$q^2\mathbb{E}_{y\in(\mathbf{Z}/q\mathbf{Z})^4} f(y)e(Q(x,y)) = z\rho(g)f(x)$$

for some unit complex number $z = z(Q)$ where

$$g = l(-2a) \cdot s(b^{-T}) \cdot J \cdot l(-2c)$$

is the product of the four elements just written down. A short computation confirms that g is the symplectic element of Q as defined in Definition 7.3.

Finally, we have

$$\begin{aligned} T_{f_1,f_2}(1) &= \mathbb{E}_{x\in(\mathbf{Z}/q\mathbf{Z})^4} f_1(x)\mathbb{E}_{y\in(\mathbf{Z}/q\mathbf{Z})^4} f_2(y)e(Q(x,y)) \\ &= z\mathbb{E}_{x\in(\mathbf{Z}/q\mathbf{Z})^4} f_1(x)(\rho(g)f_2)(x) \\ &= z\overline{\langle \overline{f}_1, \rho(g)f_2\rangle}. \end{aligned}$$

This completes the proof. □

The following definition will play a key role in what follows.

Definition 7.5. Fix $Q(x,y) = x^T ax + x^T by + y^T cy$, a quadratic form over $\mathbf{Z}$ with $\det b \neq 0$. Let g be the symplectic element of Q (see Definition 7.3). Then for every odd squarefree q coprime to $\det b$ we associate a probability measure μ_q on $\mathrm{Sp}_8(\mathbf{Z}/q\mathbf{Z})$, which puts weight $\frac{1}{\phi(q)}$ on each of the points $g^{(r)}$, $r \in (\mathbf{Z}/q\mathbf{Z})^*$.

We are now in a position to rephrase Proposition 6.2 in terms of matrix coefficients.

Proposition 7.6. *Suppose that Q is generic. Then there is $p_0(Q)$ such that the following is true. Suppose that q is squarefree with all prime factors greater than $p_0(Q)$. Then*

$$\int |\langle f_1, \rho(x)f_2\rangle| d\mu_q(x) \ll q^{-\delta}\|f_1\|_2\|f_2\|_2$$

for all $f_1, f_2 \in \ell^2((\mathbf{Z}/q\mathbf{Z})^4)$.

Remarks. For a discussion of the integral notation, see the start of the next section. By taking $p_0(Q)$ larger than any prime factor of $\det b$, we can ensure that the conditions of Proposition 7.4 are satisfied. For notational simplicity we switched $\overline{f}_1$ to f_1, which makes no difference since these functions have the same ℓ^2-norm.

8 Averages of matrix coefficients

In this section we give a general bound for averages of matrix coefficients. Whilst we do not know of a reference for quite this result, the first part of the argument is related to earlier work, particularly Bourgain [Bou12] and Skhredov [Shk21, Sect. 4]. The second idea, of using quasirandomness (no small-dimensional representations apart from the trivial one), is also by no means new. It is exploited in related ways in many works starting with Sarnak–Xue [SX91] and continuing with, for instance, Bourgain–Gamburd [BG08] and Gowers [Gow08].

Probability measures. We begin by recalling some basic notions about probability measures on (finite) groups. Let G be a finite group. A probability measure μ on G is simply a function $\mu : G \to [0,1]$ with $\sum_{x\in G}\mu(x) = 1$. The opposite measure μ° is defined by $\mu^\circ(x) = \mu(x^{-1})$. If $\mu^\circ = \mu$ then we say that μ is *symmetric*. If μ_1, μ_2 are two probability measures then their convolution $\mu_1 * \mu_2$ is defined by $\mu_1 * \mu_2(x) = \sum_{g_1g_2=x}\mu_1(g_1)\mu_2(g_2)$. This is also a probability measure. If μ is a probability measure and $2m$ a positive even integer, we write $\mu^{(2m)}$ for the $2m$-fold *symmetrised* convolution power $\mu^\circ * \mu * \mu^\circ * \cdots * \mu$. This is slightly non-standard, but very convenient as these are the only types of convolution power we will consider in this paper. At one place in Sect. 9 we will use a similar notation with an odd power, thus $\mu^{(2m-1)}(x) = \mu^\circ * \mu * \mu^\circ * \cdots * \mu^\circ$. This, of course, is not necessarily a symmetric measure. We have $\mu^{(2m)} = \mu^{(2m-1)} * \mu$. If $H \leqslant G$ is a subgroup then we write μ_H for the uniform probability measure on H, that is to say $\mu_H(x) = |H|^{-1}1_{x\in H}$.

If μ is a probability measure on a finite group G then we write $\|\mu\| = \left(\sum_x \mu(x)^2\right)^{1/2}$. Note that this is normalised differently to the ℓ^2-norm of functions which has appeared in previous sections: to reduce the potential for confusion, we omit any subscript from the norm. We extend this notion to differences of measures in the obvious way, thus $\|\mu - \nu\| = \left(\sum_x(\mu(x) - \nu(x))^2\right)^{1/2}$.

If $F : G \to \mathbf{C}$, we will adopt the fairly standard convention in this context of writing $\int F(x)d\mu(x)$ instead of $\sum_x F(x)\mu(x)$.

We will need the following consequence of Schur's lemma which is standard but cannot be reliably found in every textbook.

Lemma 8.1. *Let $\psi : G \to \mathrm{U}(V)$ be an irreducible representation of a finite group G. Suppose that $v, w \in V$. Then we have*

$$\int |\langle v, \psi(x)w\rangle|^2 d\mu_G(x) = \frac{1}{\dim\psi}\|v\|^2\|w\|^2.$$

Proof. See [Kow14, Proposition 4.3.5]. Note that here $\dim\psi$ is defined to be $\dim V$. □

PROPOSITION 8.2. *Let G be a finite group, and let $\rho : G \to \mathrm{U}(V)$ be a finite-dimensional unitary representation of G. Let μ be a probability measure on G. Let $H \leqslant G$ be the group generated by $\mathrm{Supp}(\mu^{(2)}) = \mathrm{Supp}(\mu^\circ * \mu)$. Suppose that*

- (Almost uniform distribution of convolution powers) *For some real number $K \geqslant 1$ and for some power of two m we have*
$$\mu^{(m)}(x) \leqslant K\mu_H(x) \tag{8.1}$$
pointwise;
- (Quasirandomness) *If $\rho|_H = \bigoplus_i \psi_i$ as a sum of irreducible representations (of H) then $\dim\psi_i \geqslant D$ for all i.*

Then we have the bound
$$\int |\langle v, \rho(x)w\rangle| d\mu(x) \leqslant K^{1/m} D^{-1/2m} \tag{8.2}$$
for all $v, w \in V$ with $\|v\| = \|w\| = 1$.

Remark. Note that the trivial bound is 1 (by the unitary nature of ρ and Cauchy–Schwarz). If $K \approx 1$ and D is somewhat large, (8.2) is therefore an appreciable improvement of the trivial bound.

Proof. Set
$$\eta := \int |\langle v, \rho(x)w\rangle| d\mu(x).$$
For each $x \in G$, let $\xi(x) = e(\arg\langle v, \rho(x)w\rangle)$, so $\xi(x)$ is a unit complex number and
$$\Big\langle v, \int \xi(x)\rho(x)w d\mu(x)\Big\rangle = \int \overline{\xi(x)}\langle v, \rho(x)w\rangle d\mu(x) = \eta.$$
By Cauchy–Schwarz,
$$\Big\| \int \xi(x)\rho(x)w d\mu(x)\Big\| \geqslant \eta.$$
Squaring and expanding out gives
$$\int \xi(x)\overline{\xi(x')}\langle \rho(x)w, \rho(x')w\rangle d\mu(x) d\mu(x') \geqslant \eta^2,$$
thus
$$\int |\langle \rho(x)w, \rho(x')w\rangle| d\mu(x) d\mu(x') \geqslant \eta^2.$$
Since ρ is a unitary representation, this implies
$$\int |\langle w, \rho(x^{-1}x')w\rangle| d\mu(x) d\mu(x') \geqslant \eta^2,$$

or in other words

$$\int |\langle w, \rho(x)w\rangle| d\mu^{(2)}(x) \geqslant \eta^2.$$

We may now apply the same argument again repeatedly, noting that $\mu^{(2)}, \mu^{(4)}, \dots$ are symmetric, to obtain

$$\int |\langle w, \rho(x)w\rangle| d\mu^{(m)}(x) \geqslant \eta^m$$

for any power of two m. By the almost uniform distribution assumption (8.1), this implies (with m as in (8.1)) that

$$\int |\langle w, \rho(x)w\rangle| d\mu_H(x) \geqslant K^{-1}\eta^m. \tag{8.3}$$

Now decompose $V = \bigoplus_{i=1}^n V_i$ as a sum of orthogonal $\rho(H)$-invariant subspaces, irreducible for $\rho|_H$. Let w_i be the projection of w to V_i, so $w = \sum_{i=1}^n w_i$ and

$$\sum_{i=1}^n \|w_i\|^2 = 1. \tag{8.4}$$

By Lemma 8.1 we have for $i = 1, \dots, n$

$$\int |\langle w_i, \rho(x)w_i\rangle|^2 d\mu_H(x) = \frac{1}{\dim V_i}\|w_i\|^4,$$

so by Cauchy–Schwarz and the quasirandomness assumption

$$\int |\langle w_i, \rho(x)w_i\rangle| d\mu_H(x) \leqslant D^{-1/2}\|w_i\|^2. \tag{8.5}$$

Since (by orthogonality)

$$\langle w, \rho(x)w\rangle = \sum_{i=1}^n \langle w_i, \rho(x)w_i\rangle$$

for all $x \in H$, it follows from (8.4) and (8.5) that

$$\int |\langle w, \rho(x)w\rangle| d\mu_H(x) \leqslant D^{-1/2}.$$

Comparing this with (8.3) gives the claimed bound. □

We now outline the rest of the paper. Recall that we have reduced the proof of our main theorem to the task of proving Proposition 7.6. We now have a tool, Proposition 8.2, to use on this problem. However, we must verify the two requirements, the uniform distribution property (8.1) and the quasirandomness property, in our setting.

The formal statements are Propositions 8.4 and 8.5 below. First, we give a definition which will play an important role in the rest of the paper.

Definition 8.3. Suppose that q is odd, squarefree and coprime to $\det b$. Let $\Gamma_q \leqslant \mathrm{Sp}_8(\mathbf{Z}/q\mathbf{Z})$ be the group generated by the elements $g^{-(r)}g^{(s)}$, $r,s \in (\mathbf{Z}/q\mathbf{Z})^*$, or equivalently by the support of $\mathrm{Supp}(\mu_q^{(2)})$.

Proposition 8.4. *Let Q be a quadratic form. Then there is some $p_0(Q)$ such that the following is true. Suppose that q is squarefree and has all prime factors greater than $p_0(Q)$. Let μ_q be the measure described in Definition 7.5. Then there is power of two $m = O(1)$ such that $\mu_q^{(m)} \ll \mu_{\Gamma_q}$ pointwise, with the implied constant being absolute.*

We will prove this in Sect. 9.

Proposition 8.5. *Let Q be a generic quadratic form. Then there is some $p_0(Q)$ such that the following is true. Let ρ be the Weil representation on $\mathrm{Sp}_8(\mathbf{Z}/q\mathbf{Z})$ (as given in Proposition 7.2). Suppose that q is squarefree and has all prime factors greater than $p_0(Q)$. Then $\rho|_{\Gamma_q}$ splits into irreducible subrepresentations of dimensions $\geqslant q^{1-o(1)}$.*

We will prove this in Sects. 10 and 11.

Propositions 8.4 and 8.5 are precisely what is needed in order to apply Proposition 8.2, and the conclusion is precisely Proposition 7.6. Therefore we have, as the remaining outstanding tasks, the proofs of these two propositions.

9 Uniform distribution of convolution powers

In this section we establish Proposition 8.4. Our proof of this statement has a hint of model theory about it, though we will not use that language. As remarked in the introduction, it is somewhat related to Tao's argument in [Tao]. Here is a rough plan of the proof.

(1) (Step 1) Consider first the case $q = p$ a sufficiently large prime. We argue that the sequence $\|\mu_p^{(2^j)}\|$, $j = 0,1,2,\dots$ (which is non-increasing by Young's inequality) stabilises at some time $t = O(1)$, in the sense that $\|\mu_p^{(2^{t+1})}\| = (1+O(p^{-1/2}))\|\mu_p^{(2^t)}\|$. This uses the Lang-Weil bound from algebraic geometry.
(2) (Step 2) By standard arguments from additive combinatorics (recalled in Appendix C), this implies that $\|\mu_p^{(2^t)} - \mu_H\| \ll p^{-c}\|\mu_H\|$, for some subgroup $H \leqslant \mathrm{Sp}_8(\mathbf{Z}/p\mathbf{Z})$.
(3) (Step 3) By some group-theoretic arguments, H must in fact be Γ_p.
(4) (Step 4) Taking a few further convolution powers, we upgrade the estimate to a much stronger bound $\|\mu^{(s)} - \mu_{\Gamma_p}\|_\infty \ll p^{-100}$.
(5) (Step 5) We deduce the general squarefree case of Proposition 8.4.

Lang-Weil estimate. We keep algebro-geometric terminology to an absolute minimum. A good down-to-earth account of what we need may be found in [Tao15, Chap. 9]. Let M be a real parameter. Then for the purposes of this paper, by a

variety of complexity $\leqslant M$ defined over $\mathbf{F}_p$ we mean a set of points of the form

$$V := \{x \in \overline{\mathbf{F}}_p^n : P_1(x) = \cdots = P_m(x) = 0\},$$

where $P_1, \dots, P_m \in \mathbf{F}_p[X_1, \dots, X_n]$ are polynomials all of degree $\leqslant M$, and $m, n \leqslant M$. Denote by $V(\mathbf{F}_p)$ the $\mathbf{F}_p$-points of V, that is to say the points of V all of whose coordinates lie in $\mathbf{F}_p$.

PROPOSITION 9.1 (Lang-Weil). *We have*

$$|V(\mathbf{F}_p)| = (c(V) + O_M(p^{-1/2}))p^{\dim V},$$

for some integer $c(V) = O_M(1)$.

Remark. In fact, $c(V)$ is the number of top-dimensional components of V which are definable over $\mathbf{F}_p$, but we shall not need this description. Nor will we need to really know what dimension means, other than that it is an integer in the range $0 \leqslant \dim V \leqslant n$. All we need is the fact that the quantities $|V(\mathbf{F}_p)|$ are restricted to a rather discretised set of values. This kind of application of Lang-Weil has appeared in several model-theoretic works and is related to the concept of *stability*.

The Lang-Weil estimate has the following consequence for convolution powers of our measures μ_p.

LEMMA 9.2. *There are functions* $\alpha, \beta : \mathbf{N} \times \{\text{primes}\} \to \mathbf{Z}_{\geqslant 0}$ *and non-decreasing functions* $\alpha_* : \mathbf{N} \to \mathbf{Z}_{\geqslant 0}$ *and* $p_0 : \mathbf{N} \to \mathbf{N}$ *such that for all* $j \geqslant 1$ *and for all primes* p *we have*

$$\|\mu_p^{(2^j)}\|^2 = \Big(\alpha(j,p) + O_j(p^{-1/2})\Big)p^{-\beta(j,p)}, \tag{9.1}$$

where $\alpha(j,p)$, $\beta(j,p)$ *are integers with* $0 < \alpha(j,p) \leqslant \alpha_*(j)$. *Moreover, if* $p \geqslant p_0(j)$ *then we have*

$$0 \leqslant \beta(j,p) \leqslant 36; \tag{9.2}$$

$$\beta(j+1,p) \geqslant \beta(j,p) \tag{9.3}$$

and

$$\alpha(j+1,p) \leqslant \alpha(j,p) \quad \text{if} \quad \beta(j+1,p) = \beta(j,p). \tag{9.4}$$

Proof. The key point is to interpret $\|\mu_p^{(2^j)}\|^2$ in terms of the number of $\mathbf{F}_p$-points on a variety of bounded complexity. To this end, we have

$$\begin{aligned}\|\mu_p^{(2^j)}\|^2 &= \sum_x \mu_p^{(2^j)}(x)^2 \\ &= (p-1)^{-2^{j+1}} \#\{(r,r') \in (\mathbf{F}_p^*)^{2^j} \times (\mathbf{F}_p^*)^{2^j} : g^{-(r_1)}g^{(r_2)} \cdots g^{(r_{2^j})} = \\ &\qquad = g^{-(r'_1)}g^{(r'_2)} \cdots g^{(r'_{2^j})}\}\end{aligned}$$

$$= (p-1)^{-2^{j+1}} \#\{(r, r', y, y') \in \mathbf{F}_p^{2^j} \times \mathbf{F}_p^{2^j} \times \mathbf{F}_p \times \mathbf{F}_p : r_1 \cdots r_{2^j} y = \\ = r'_1 \cdots r'_{2^j} y' = 1,\ g^{-(r_1)} g^{(r_2)} \cdots g^{(r_{2^j})} = g^{-(r'_1)} g^{(r'_2)} \cdots g^{(r'_{2^j})}\}$$

$$= (p-1)^{-2^{j+1}} |V_j(\mathbf{F}_p)| = (1 + O_j(\frac{1}{p}))p^{-2^{j+1}} |V_j(\mathbf{F}_p)|,$$

where $V_j \subset \overline{\mathbf{F}}_p^{2^{j+1}+2}$ is some variety of complexity $O_j(1)$, defined over $\mathbf{F}_p$. Note here that, although (for instance) $g^{(r_1)} = \begin{pmatrix} -2b^{-T}c & r_1^{-1}b^{-T} \\ r_1(4ab^{-T}c - b) & -2ab^{-T} \end{pmatrix}$ is not *a priori* given by polynomials, we can express

$$g^{(r_1)} = \begin{pmatrix} -2b^{-T}c & r_2 \cdots r_{2^j} y b^{-T} \\ r_1(4ab^{-T}c - b) & -2ab^{-T} \end{pmatrix},$$

and this *is* given by polynomials. (Alternatively, one could talk about quasiprojective varieties, but the trick of introducing y, y' avoids the need to do that.)

This immediately implies, by the Lang-Weil bound, the first statement (9.1) (with $\beta(j,p) = 2^{j+1} - \dim V_j$). We now proceed to derive the additional statements (9.2), (9.3) and (9.4), which we do by combinatorial means (with reference to (9.1)).

For (9.2), note that *any* probability measure ν on a finite group G satisfies $|G|^{-1} \leqslant \|\nu\|^2 \leqslant 1$; the lower bound is Cauchy–Schwarz, and the upper bound is the trivial bound (with equality only if ν is concentrated at one point). Since (see (7.4)) $|\operatorname{Sp}_8(\mathbf{Z}/p\mathbf{Z})| = (1 + o(1))p^{36}$, (9.2) follows if p is large enough.

For items (9.3) and (9.4) we use Young's inequality (Lemma C.1), which implies that $\|\mu_p^{(2^j)}\|^2$ is a non-increasing function of j. Therefore we have, by (9.1),

$$\Big(\alpha(j+1,p) + O_j(p^{-1/2})\Big)p^{-\beta(j+1,p)} \leqslant \Big(\alpha(j,p) + O_j(p^{-1/2})\Big)p^{-\beta(j,p)}.$$

Here, $\alpha(j+1,p)$, $\alpha(j,p)$ are positive integers of size $O_j(1)$ and so by taking p sufficiently large in terms of j both (9.3) and (9.4) follow.

Finally, note that $\alpha(j,p)$ is bounded above by $O_{M_j}(1)$, where $M_j = O_j(1)$ is an upper bound for the complexity of V_j. Thus $\alpha(j,p) \leqslant \alpha_*(j)$ for some function α_*, which we may clearly assume to be non-decreasing (else replace it by $\sup_{i \leqslant j} \alpha_*(i)$). For the same reasons, we may also assume that p_0 is a non-decreasing function. □

Remark. The inequalities (9.3), (9.4) can also be seen in a more purely algebro-geometric manner by noting that the varieties V_j are defined over $\mathbf{Q}$ and that (appropriately embedding them into a common underlying affine space) we have $V_j \subseteq V_{j+1}$. For p sufficiently large in terms of j, $\alpha(j,p)$ then counts the number of irreducible components of $V_j(\overline{\mathbf{Q}})$, which makes the inequality $\alpha(j,p) \leqslant \alpha(j+1,p)$ clear.

We now turn to the realisation of Step 1 of the outline.

PROPOSITION 9.3. *Suppose that p is sufficiently large. Then there is some $t = O(1)$ such that $\|\mu_p^{(2^{t+1})}\| = (1 + O(p^{-1/2}))\|\mu_p^{(2^t)}\|$.*

Proof. We use Lemma 9.2 and the notation there. It clearly suffices to show that, for some $t = O(1)$, we have

$$\alpha(t+1,p) = \alpha(t,p) \quad \text{and} \quad \beta(t+1,p) = \beta(t,p). \tag{9.5}$$

Define a sequence $T_1, T_2, \dots, T_{37}$ as follows. Set $T_1 := \alpha_*(1)$, and then inductively define $T_{i+1} := T_i + \alpha_*(T_i)$ for $i = 1, 2, 3, \dots, 36$. If $p \geqslant p_0(T_{37})$ we have the bounds (9.2), (9.3) and (9.4), for all $j \leqslant T_{37}$.

We claim that there is some $t \leqslant T_{37}$ such that (9.5) holds. Suppose not. Then, by (9.3), (9.4) we have that for all $j \leqslant T_{37}$ either

(1) $\beta(j+1,p) > \beta(j,p)$ or
(2) $\beta(j+1,p) = \beta(j,p)$ and $\alpha(j+1,p) < \alpha(j,p)$.

By (9.2), there are at most 36 values of j for which (1) occurs; suppose they are $T_1', \dots, T_m'$, $m \leqslant 36$.

For $j = 1, 2, \dots, T_1' - 1$ we must have (2), which means that $\alpha(T_1', p) \leqslant \alpha(1,p) - T_1' \leqslant \alpha_*(1) - T_1' + 1 = T_1 - T_1' + 1$. Since $\alpha(j,p)$ is always positive, this implies that $T_1' \leqslant T_1$.

Now for $j = T_1' + 1, \dots, T_2' - 1$ we must also have (2), which means that

$$\begin{aligned} \alpha(T_2', p) &\leqslant \alpha(T_1', p) - (T_2' - T_1' - 1) \\ &\leqslant \alpha_*(T_1') - (T_2' - T_1' - 1) \leqslant \alpha_*(T_1) + T_1 - T_2' + 1. \end{aligned}$$

Since $\alpha(j,p)$ is always positive, this implies that $T_2' \leqslant \alpha_*(T_1) + T_1 = T_2$. Continuing in this manner we see inductively that $T_m' \leqslant T_m \leqslant T_{36}$. Continuing now with $j = T_m' + i$, $i = 1, 2, \dots$, only (2) can occur, and so

$$\alpha(T_m' + i, p) \leqslant \alpha(T_m', p) - i \leqslant \alpha_*(T_m') - i \leqslant \alpha_*(T_{36}) - i.$$

Since (yet again) $\alpha(j,p)$ is always positive, this can only continue as far as $i = \alpha_*(T_{36}) - 1$ before we get a contradiction. Note that then $j \leqslant T_m' + \alpha_*(T_{36}) \leqslant T_{37}$, so all the appeals we made to (9.2), (9.3) and (9.4) were indeed valid.

This contradiction shows that we were wrong to assume that there is no $t \leqslant T_{37}$ for which (9.5) holds. □

Step 2. The conclusion of Step 1 (Proposition 9.3) is that for some t, $1 \leqslant t \ll O(1)$, we have, for the symmetic measure $\nu := \mu^{(2^t)}$ the very strong "flattening"

$$\|\nu * \nu\| = (1 + O(p^{-1/2}))\|\nu\|.$$

It is well-known that any probability measure satisfying this kind of property is close to uniform on a subgroup. The precise statement we need is Corollary C.3 in Appendix C, from which we conclude that there is some subgroup H such

that

$$\|\mu_p^{(2^t)} - \mu_H\| \ll p^{-c}\|\mu_H\|, \tag{9.6}$$

$$\mu_p^{(2^t)}(H) \geqslant 1 - O(p^{-c}), \tag{9.7}$$

and

$$|\operatorname{Supp}(\mu_p^{(2^t)})| > (1 - O(p^{-c}))|H|, \tag{9.8}$$

where μ_H is the uniform measure on H. It follows from (9.7) that

$$1 - O(p^{-c}) \leqslant \mu_p^{(2^t)}(H) = (\mu_p^{(2^t-1)} * \mu_p)(H) = \\ = \sum_x \mu_p^{(2^t-1)}(x)\mu_p(x^{-1}H) \leqslant \sup_x \mu_p(xH).$$

Thus there is some coset xH such that

$$\mu_p(xH) \geqslant 1 - O(p^{-c}). \tag{9.9}$$

Step 3. In this step we use a group-theoretic argument, making use of some slightly specific features of the problem, to upgrade the statement (9.9) to $\mu(xH) = 1$, or in other words (recalling Definitions 7.3 and 7.5) to show that all $g^{(r)}$, $r \in (\mathbf{Z}/p\mathbf{Z})^*$, lie in xH.

Let $R := \{r \in (\mathbf{Z}/p\mathbf{Z})^* : g^{(r^{-1})} \in xH\}$. Thus, by (9.9),

$$|R| \geqslant (1 - O(p^{-c}))(p-1). \tag{9.10}$$

Perform the following algorithm to generate distinct elements $r_1, r_2, \dots$ of R as long as possible. Write $S_j := \bigcap_{i \leqslant j} (xH)^{(r_i)}$ (where $A^{(r)}$ means $\{a^{(r)} : a \in A\}$). Each S_j is a coset (of some subgroup of $\mathrm{Sp}_8(\mathbf{Z}/p\mathbf{Z})$) and, no matter how we choose the r_i, we have the nesting

$$S_1 \supseteq S_2 \supseteq \cdots$$

If, at step j of the construction, it is possible to choose $r_{j+1} \in R$ so that S_{j+1} is a *proper* subset of S_j then do so; otherwise, stop.

Note that, as long as the algorithm continues, we have $|S_{j+1}| \leqslant \frac{1}{2}|S_j|$ (since the S_j are all cosets of subgroups). Therefore, the algorithm stops in at most $O(\log p)$ steps.

When the algorithm finishes, we have $r_1, \dots, r_m \in R$, $m = O(\log p)$, and a coset $S := S_m = \bigcap_{i=1}^m (xH)^{(r_i)}$ (of some subgroup). Note that, since $r_i \in R$, we have $g^{(r_i^{-1})} \in xH$ and so $g \in (xH)^{(r_i)}$, and therefore $g \in S$.

Now set

$$R' := r_1^{-1}R \cap \cdots \cap r_m^{-1}R,$$

and suppose that $r \in R'$. Since the algorithm we described stopped at the mth stage, we have

$$S \cap (xH)^{(r_i r)} = S \qquad \text{for } i = 1, \dots, m,$$

since otherwise we could take $r_{m+1} := r_i r$ (which would be an element of R by the definition of R'). It follows that

$$S \cap S^{(r)} = S \cap \bigcap_{i=1}^{m} (xH)^{(r_i r)} = S.$$

That is, if $r \in R'$ then $S = S^{(r)}$. It follows that S is invariant under the entire subgroup of $(\mathbf{Z}/p\mathbf{Z})^*$ generated by R'. However,

$$|R'| \geqslant 1 - m|(\mathbf{Z}/p\mathbf{Z})^* \setminus R| > \frac{1}{2}(p-1),$$

by (9.10) and the fact that $m = O(\log p)$. Therefore the group generated by R' is the whole of $(\mathbf{Z}/p\mathbf{Z})^*$, and so we have that $S = S^{(r)}$ for all $r \in (\mathbf{Z}/p\mathbf{Z})^*$.

We showed earlier that $g \in S$. It now follows that $g^{(rr_1)} \in S$ for all $r \in (\mathbf{Z}/p\mathbf{Z})^*$. In particular, $g^{(rr_1)} \in (xH)^{(r_1)}$, which implies that $g^{(r)} \in xH$ for all $r \in (\mathbf{Z}/p\mathbf{Z})^*$, which is what we wanted to show.

Consequently, all the elements $g^{-(r)}g^{(s)}$ lie in H. By definition (Definition 8.3), we have $\Gamma_p \leqslant H$.

It follows that $\mathrm{Supp}(\mu_p^{(2^t)}) \subseteq H$. However, we showed in (9.8) that $|\mathrm{Supp}(\mu_p^{(2^t)})| > (1 - O(p^{-c}))|H| > \frac{1}{2}|H|$, and therefore the group generated by $\mathrm{Supp}(\mu_p^{(2^t)})$ is all of H. However, the group generated by $\mathrm{Supp}(\mu_p^{(2^t)})$ is the group generated by $\mathrm{Supp}(\mu_p^{(2)})$ which, as we remarked earlier, is precisely Γ_p. Finally, we may conclude that $H = \Gamma_p$. Therefore (9.6) may be rewritten as

$$\|\mu_p^{(2^t)} - \mu_{\Gamma_p}\| \ll p^{-c}\|\mu_{\Gamma_p}\|. \tag{9.11}$$

Step 4. In this step of the argument we upgrade (9.11) to a highly uniform estimate by taking a few extra convolution powers. By Cauchy–Schwarz (and using that $\mu_{\Gamma_p} * \nu = \mu_{\Gamma_p}$ for any ν with support contained in Γ_p) is we have

$$\begin{aligned}\|\mu_p^{(2^{t+1})} - \mu_{\Gamma_p}\|_\infty &= \|(\mu_p^{(2^t)} - \mu_{\Gamma_p}) * (\mu_p^{(2^t)} - \mu_{\Gamma_p})\|_\infty \\ &\leqslant \|\mu_p^{(2^t)} - \mu_{\Gamma_p}\|^2 \\ &\ll p^{-2c}\|\mu_{\Gamma_p}\|^2 = p^{-2c}|\Gamma_p|^{-1}. \end{aligned} \tag{9.12}$$

However, if ν is some probability measure on a finite group Γ of size N and if

$$\|\nu - \mu_\Gamma\|_\infty \leqslant \frac{\varepsilon}{N}$$

then

$$\Big|\nu * \nu(x) - \frac{1}{N}\Big| = |(\nu - \mu_\Gamma) * (\nu - \mu_\Gamma)(x)|$$
$$\leqslant \sum_y \Big|\nu(y) - \frac{1}{N}\Big|\Big|\nu(y^{-1}x) - \frac{1}{N}\Big| \leqslant \frac{\varepsilon^2}{N},$$

that is to say

$$\|\nu^{(2)} - \mu_\Gamma\|_\infty \leqslant \frac{\varepsilon^2}{N}.$$

Applying this s times to (9.12) gives

$$\|\mu_p^{(2^{t+1+s})} - \mu_{\Gamma_p}\|_\infty \leqslant (Cp^{-2c})^{2^s}|\Gamma_p|^{-1},$$

and so, taking a suitably large s, there is some power of two $m = 2^{t+1+s} = O(1)$ such that

$$\|\mu_p^{(m)} - \mu_{\Gamma_p}\|_\infty < p^{-38}, \tag{9.13}$$

provided (as always) p is sufficiently large. Since we are free to choose any sufficiently large $s = O(1)$, we can make the choice so that m is independent of p.

Note that (9.13) is a much stronger version of Proposition 8.4 in the case $q = p$ a sufficiently large prime.

Since $|\Gamma_p| \leqslant |\mathrm{Sp}_8(\mathbf{Z}/p\mathbf{Z})| \ll p^{36}$ it follows from (9.13) that

$$\mu_p^{(m)} \leqslant (1 + O(p^{-2}))|\Gamma_p|^{-1} \tag{9.14}$$

pointwise. This one-sided estimate is the only one we will need subsequently.

Step 5. Finally, we turn to the deduction of Proposition 8.4 itself. That is, we pass from the case $q = p$ a prime to the general case. Suppose then that q is squarefree, and that all its prime factors are sufficiently large (larger than $p_0(T_{37})$, the quantity appearing in Step 1, is enough). We have a natural homomorphism

$$\pi : \mathrm{Sp}_8(\mathbf{Z}/q\mathbf{Z}) \to \prod_{p|q} \mathrm{Sp}_8(\mathbf{Z}/p\mathbf{Z}).$$

By the Chinese remainder theorem, the measure μ_q pushes forward under π to the product $\times_{p|q}\mu_p$: we have $\pi(g^{(r)}) = (g^{(r(\mathrm{mod}\, p))})_{p|q}$, and the tuple $(r(\mathrm{mod}\, p))_{p|q}$ takes all values in $\prod_{p|q}(\mathbf{Z}/p\mathbf{Z})^*$ as r ranges over $(\mathbf{Z}/q\mathbf{Z})^*$.

Recall that Γ_q is the group generated by the $g^{-(r)}g^{(s)}$, $r, s \in (\mathbf{Z}/q\mathbf{Z})^*$. These groups also behave nicely under projection, as the following lemma shows.

Lemma 9.4. *Suppose that q is squarefree, and let $\pi : \mathrm{Sp}_8(\mathbf{Z}/q\mathbf{Z}) \to \prod_{p|q} \mathrm{Sp}_8(\mathbf{Z}/p\mathbf{Z})$ be the natural isomorphism. Then $\pi(\Gamma_q) = \prod_{p|q}\Gamma_p$.*

Proof. It is easy to see that $\pi(\Gamma_q) \subseteq \prod_{p|q} \Gamma_p$. To see that the two are in fact equal, suppose we have elements $\gamma_p \in \Gamma_p$. For some N, we may write $\gamma_p = g^{-(r_{1,p})} g^{(s_{1,p})} \dots g^{-(r_{N,p})} g^{(s_{N,p})}$ for elements $r_{i,p}, s_{i,p} \in (\mathbf{Z}/p\mathbf{Z})^*$, that is to say as a word consisting of a product of N of the generators. Note that we can use the same N for each p by padding with exponents $r_{i,p} = s_{i,p}$, if necessary, each of which contributes the identity, and we have also taken advantage of the fact that $(g^{-(r)} g^{(s)})^{-1} = g^{-(s)} g^{(r)}$, which means we do not need to worry about including inverses separately. By the Chinese remainder theorem there are $r_i, s_i \in (\mathbf{Z}/q\mathbf{Z})^*$ such that $r_i (\mathrm{mod}\ p) = r_{i,p}$, and similarly for the s_i. Setting $\gamma := g^{-(r_1)} g^{(s_1)} \dots g^{(-r_N)} g^{(s_N)} \in \mathrm{Sp}_8(\mathbf{Z}/q\mathbf{Z})$, we see that $\pi(\gamma) = (\gamma_p)_{p|q}$, as desired. □

Remark. We caution that this lemma is a rather specific result. If, for example, we defined $\tilde{\Gamma}_q$ to be the group generated by the $g^{(r)}$, $r \in (\mathbf{Z}/q\mathbf{Z})^*$, the same argument would not work (consider, for example, the question of how to find an element whose projection to $\mathrm{Sp}_8(\mathbf{Z}/p\mathbf{Z})$ is g, and whose projection to $\mathrm{Sp}_8(\mathbf{Z}/p'\mathbf{Z})$ is g^{-1}).

With these facts in hand, we may now complete the proof. Let m, a power of two, be as in (9.14). Then

$$\|\mu_q^{(m)}\|_\infty = \prod_{p|q} \|\mu_p^{(m)}\|_\infty \leqslant \prod_{p|q} (1 + O(p^{-2})) |\Gamma_p|^{-1} \ll \prod_{p|q} |\Gamma_p|^{-1} = |\Gamma_q|^{-1}.$$

This bound, coupled with the fact that $\mathrm{Supp}(\mu_q^{(m)}) \subset \Gamma_q$, implies Proposition 8.4.

10 Identifying Γ_q

We turn now to the task of proving Proposition 8.5. The first stage is to actually identify the group Γ_q that is to say the subgroup of $\mathrm{Sp}_8(\mathbf{Z}/q\mathbf{Z})$ generated by the $g^{-(r)} g^{(s)}$, $r, s \in (\mathbf{Z}/q\mathbf{Z})^*$, in an explicit algebraic form. Recall that the symplectic element g is a particular element of $\mathrm{Sp}_8(\mathbf{Q})$ associated to the quadratic form Q (see Definition 7.3). In view of Lemma 9.4, it is enough to consider the prime case Γ_p.

To understand this group, it seems best to think of $\mathrm{Sp}_8(\mathbf{Z}/p\mathbf{Z})$ as embedded in the full group $\mathrm{GL}_8(\mathbf{Z}/p\mathbf{Z})$ (in the obvious way). It then turns out that Γ_p is a conjugate (in $\mathrm{GL}_8(\mathbf{Z}/p\mathbf{Z})$) of the group $\mathrm{SL}_2(\mathbf{F}_p[\Delta]) \leqslant \mathrm{GL}_8(\mathbf{Z}/p\mathbf{Z})$, where (recall) $\Delta = 4b^{-1}ab^{-T}c - I$ is the matrix discriminant of our form Q. We will discuss such groups at much greater length in due course. For now, note that $\mathbf{F}_p[\Delta]$ is an algebra over $\mathbf{F}_p$ of dimension at most 4 (by the Cayley-Hamilton theorem). Generically, this algebra will have dimension exactly 4 and in this case $\mathrm{SL}_2(\mathbf{F}_p[\Delta])$ (and hence Γ_p) will be a group of size $\sim p^{12}$.

The precise structure of $\mathrm{SL}_2(\mathbf{F}_p[\Delta])$ will depend on how the characteristic polynomial ρ_Δ for Δ splits over $\mathbf{F}_p$, but it will be a direct product of some groups $\mathrm{SL}_2(\mathbf{F}_{p^n})$ for various n.

Suppose henceforth that Q is generic, and that p is large enough that a, b, Δ are invertible over $\mathbf{F}_p$. In order to examine the elements $g^{-(r)} g^{(s)}$ (which, by definition,

generate Γ_p), one eventually hits upon the idea of looking at what might be called a "block DUL factorisation" of g. Namely, there is an invertible A and symmetric B, C (all 4-by-4 matrices) such that

$$g = \begin{pmatrix} A & 0 \\ 0 & A^{-T} \end{pmatrix} \begin{pmatrix} I & B \\ 0 & I \end{pmatrix} \begin{pmatrix} I & 0 \\ C & I \end{pmatrix}. \tag{10.1}$$

To see that such a factorisation exists is simple: writing $g = \left(\begin{smallmatrix} P & Q \\ R & S \end{smallmatrix}\right)$ where $P = -2b^{-T}c$, $Q = b^{-T}$, $R = 4ab^{-T}c - b$, $S = -2ab^{-T}$, we can take

$$A = S^{-T},\ B = S^T Q,\ C = S^{-1}R. \tag{10.2}$$

Note that S is well-defined and invertible (over $\mathbf{Q}$ and over $\mathbf{F}_p$) since both a and b are invertible.

The matrices B and C here will be somewhat important in their own right. We calculate

$$B = -2b^{-1}ab^{-T}, \quad C = -2c + \frac{1}{2}b^T a^{-1} b, \quad BC = \Delta. \tag{10.3}$$

We note that the appearance of Δ here is the reason for its definition. Since Δ is assumed invertible over $\mathbf{F}_p$, both B and C are invertible over $\mathbf{F}_p$.

The purpose of looking at the DUL factorisation (10.1) is that it renders the action of dilation easy to understand. Indeed,

$$g^{-(r)}g^{(s)} = \begin{pmatrix} I & 0 \\ -rC & I \end{pmatrix} \begin{pmatrix} I & -r^{-1}B \\ 0 & I \end{pmatrix} \begin{pmatrix} I & s^{-1}B \\ 0 & I \end{pmatrix} \begin{pmatrix} I & 0 \\ sC & I \end{pmatrix}. \tag{10.4}$$

Now set

$$\tau := \begin{pmatrix} I & 0 \\ 0 & B \end{pmatrix} \in \mathrm{GL}_8(\mathbf{Q}). \tag{10.5}$$

This is invertible over $\mathbf{F}_p$ and so may considered as an element of $\mathrm{GL}_8(\mathbf{Z}/p\mathbf{Z})$. Now observe that for $\lambda \in \mathbf{F}_p^*$ we have

$$\begin{pmatrix} I & \lambda B \\ 0 & I \end{pmatrix} = \tau^{-1} \begin{pmatrix} I & \lambda I \\ 0 & I \end{pmatrix} \tau, \quad \begin{pmatrix} I & 0 \\ \lambda C & I \end{pmatrix} = \tau^{-1} \begin{pmatrix} I & 0 \\ \lambda \Delta & I \end{pmatrix} \tau. \tag{10.6}$$

It follows from this and (10.4) that $g^{-(r)}g^{(s)}$ takes values in the subgroup $\tau^{-1}\,\mathrm{SL}_2(\mathbf{F}_p[\Delta])\tau \leqslant \mathrm{Sp}_8(\mathbf{Z}/p\mathbf{Z})$.

Remark. It is important to note that neither τ nor $\mathrm{SL}_2(\mathbf{F}_p[\Delta])$ are contained in $\mathrm{Sp}_8(\mathbf{Z}/p\mathbf{Z})$ in general, although of course the conjugate $\tau^{-1}\,\mathrm{SL}_2(\mathbf{F}_p[\Delta])\tau$ is. It turns out that (assuming Q is generic, and for sufficiently large p) this is the group Γ_p. This is the first key result of the section.

Proposition 10.1. *Suppose that Q is generic and that p is sufficiently large in terms of Q. Then $\Gamma_p = \tau^{-1}\,\mathrm{SL}_2(\mathbf{F}_p[\Delta])\tau$.*

Before turning to the proof, we assemble some lemmas. We will need the following two facts about polynomials.

LEMMA 10.2. *Let $\eta > 0$. Suppose that $F \in \overline{\mathbf{F}}_p[X,Y]$ has total degree D and that $p > 2D/\eta$. Then*

(1) *If $F(x,y) = 0$ for at least a proportion η of all pairs $(x,y) \in \mathbf{F}_p \times \mathbf{F}_p$, then F is identically zero*;
(2) *If $F(x,y)$ takes values in some subfield $k_0 < \overline{\mathbf{F}}_p$, for at least a proportion η of all pairs $(x,y) \in \mathbf{F}_p \times \mathbf{F}_p$, then all the coefficients of F lie in k_0.*

Proof. (i) This is an immediate consequence of the Schwartz-Zippel lemma, which states that if F is not identically zero then the number of solutions to $F(x,y) = 0$ with $x, y \in \mathbf{F}_p$ is at most Dp, which is less than ηp^2 under the assumptions of the lemma.

(ii) We begin with a 1-variable version. If $f(x)$ is a polynomial of degree D which takes values in k_0 for $D+1$ different values of $x \in \mathbf{F}_p$ then it follows from Lagrange interpolation that all the coefficients of f lie in k_0.

Turning to the 2-variable statement we actually want, write $F(X) = \sum_{i=0}^{D} f_i(X) \times Y^i$, where $\deg f_i \leqslant D$. Let $S \subset \mathbf{F}_p \times \mathbf{F}_p$ be the set of pairs (x,y) for which $F(x,y) \in k_0$. For each x, let $S_x := \{y : (x,y) \in S\}$. Then there is a set $A \subset \mathbf{F}_p$, $|A| \geqslant \eta p/2$, such that $|S_x| \geqslant \eta p/2$ for all $x \in A$. If $x \in A$, the 1-variable result implies (since $p > 2D/\eta$) that all the $f_i(x)$ lie in k_0. A second application of the 1-variable result then implies that all the coefficients of each f_i lie in k_0.

We remark that the proof technique for (ii) can also be used for (i) (in fact this is essentially the usual proof of Schwartz-Zippel by induction). □

We will also need a couple of lemmas about subgroups of direct products. Both may be found in [Ser16, Chap. 1]. The first result is well-known.

LEMMA 10.3 (Goursat's lemma). *Let G_1, G_2 be groups. Consider the direct product $G_1 \times G_2$ and let $\pi_i : G_1 \times G_2 \to G_i$ be the two projection maps. Let $H \leqslant G_1 \times G_2$ be a subgroup, and suppose that $\pi_i(H) = G_i$ for $i = 1, 2$. Then there are normal subgroups $N_i \lhd G_i$ and an isomorphism $\phi : G_1/N_1 \to G_2/N_2$ such that H has the form $\{(g_1, g_2) \in G_1 \times G_2 : \phi(\overline{g}_1) = \overline{g}_2\}$, where $\overline{g}_i$ is the image of g_i in G_i/N_i.*

Proof. See, for example, [Ser16, Proposition 1.6]. □

The second result is somewhat less well-known and is called Ribet's lemma by Serre [Ser16].

LEMMA 10.4 (Ribet's Lemma). *Let $G_1, \dots, G_n$ be perfect groups, that is to say equal to their own commutator subgroups. Let $H \leqslant G_1 \times \cdots \times G_n$ be such that the projection π_{ij} of H to $G_i \times G_j$ is surjective for every pair (i,j). Then H is the whole of $G_1 \times \cdots \times G_n$.*

Proof. See [Ser16, Proposition 1.8]. □

The next lemma, which looks a little *ad hoc*, is in some sense the scalar version of Proposition 10.1, and is the heart of the proof of it.

LEMMA 10.5. *Suppose that* $p \geqslant 5$ *and that* $\theta \in \overline{\mathbf{F}}_p \setminus \{0, -1\}$. *Then the matrices*

$$M_\theta(r,s) = \begin{pmatrix} 1 & 0 \\ -r\theta & 1 \end{pmatrix} \begin{pmatrix} 1 & -r^{-1} \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & s^{-1} \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ s\theta & 1 \end{pmatrix},$$

as r, s *range over* $\mathbf{F}_p^*$, *generate* $\mathrm{SL}_2(\mathbf{F}_p(\theta))$.

Proof. Set $k := \mathbf{F}_p(\theta)$. We use the fact that any proper subgroup of $\mathrm{SL}_2(k)$ has a subgroup of index $C = O(1)$ which is conjugate to a subgroup of one of

(1) the group of upper triangular matrices or
(2) $\mathrm{SL}_2(k_0)$ for some proper subfield $k_0 < k$.

See Appendix A for further comments on this fact.

Suppose, then, that the $M_\theta(r,s)$ generate a proper subgroup $\Gamma < \mathrm{SL}_2(\mathbf{F}_p(\theta))$. Let $\Gamma' \leqslant \Gamma$, $[\Gamma : \Gamma'] \leqslant C$, be a subgroup conjugate to a group of type (1) or (2) above. By the pigeonhole principle there is some coset $\Gamma' x$ containing the elements $g^{-(r)}g$, $r \in R$, for some set $R \subset \mathbf{F}_p^*$ of size at least $\frac{1}{C}(p-1)$. Then Γ' contains the elements $g^{-(r)}g^{(s)} = g^{-(r)}g^{-1}(g^{-(s)}g^{-1})^{-1}$ for all $r, s \in R$, that is to say for more than ηp^2 pairs (r,s), where $\eta = (2C^2)^{-1}$.

Now we may explicitly compute that if $P = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right) \in \mathrm{SL}_2(k)$ (here a, b, c are local to the proof of Lemma 10.5, not the coefficients of Q) then the bottom-left entry of $PM_\theta(r,s)P^{-1}$ is

$$\frac{1}{rs}\Big(d^2\theta(1+\theta)(s^2r - r^2s) + c^2(s-r) + cd\theta(r^2 - s^2)\Big).$$

Consider the bracketed expression as a polynomial in r and s. If this is to be zero for at least ηp^2 pairs $(r,s) \in \mathbf{F}_p^* \times \mathbf{F}_p^*$ then the first part of Lemma 10.2 implies that (if p is sufficiently large) all the coefficients of this polynomial vanish, and so $c^2 = cd\theta = d^2\theta(1+\theta) = 0$. We cannot have $c = d = 0$ (since P is invertible) and so, since $\theta \neq 0, -1$, possibility (1) is excluded.

Turning to the subfield case (2), we proceed similarly, but now using the second part of Lemma 10.2. This implies that all of $c^2 = cd\theta = d^2\theta(1+\theta)$ lie in k_0. If $cd \neq 0$, it follows that

$$1 + \frac{1}{\theta} = \frac{c^2 \cdot d^2\theta(1+\theta)}{(cd\theta)^2} \in k_0,$$

and so $\theta \in k_0$. This, however, is impossible since $k = \mathbf{F}_p(\theta)$ and k_0 is assumed to be a *proper* subfield of k.

If $c = 0$, we consider additionally the top-left entry of $PM(r,s)P^{-1}$, which (when $c = 0$) is

$$\frac{1}{r}\Big(ad(1+\theta)r - ad\theta s + bd(rs - r^2)\theta(1+\theta)\Big).$$

Note that $ad = \det P = 1$. Therefore, by the second part of Lemma 10.2, $\theta \in k_0$. This is again a contradiction.

Finally if $d = 0$, the top-left entry of $PM(r,s)P^{-1}$ is

$$\frac{1}{rs}\Big(ac(s-r) - bc(1+\theta)rs + bc\theta r^2\Big).$$

Note that $bc = \det P = -1$, and so again we get $\theta \in k_0$. □

Proof of Proposition 10.1. We must show that the elements $\gamma(r,s) := \tau g^{-(r)} g^{(s)} \tau^{-1}$ generate $\mathrm{SL}_2(\mathbf{F}_p[\Delta])$. From (10.4), (10.6) we have

$$\gamma(r,s) = \begin{pmatrix} I & 0 \\ -r\Delta & I \end{pmatrix} \begin{pmatrix} I & -r^{-1}I \\ 0 & I \end{pmatrix} \begin{pmatrix} I & s^{-1}I \\ 0 & I \end{pmatrix} \begin{pmatrix} I & 0 \\ s\Delta & I \end{pmatrix}. \tag{10.7}$$

This should be compared with the definition of $M_\theta(r,s)$ in Lemma 10.5.

Now we are assuming that Q is generic which, by definition, means that the characteristic polynomial $\rho_\Delta(\lambda) = \det(\Delta - \lambda I)$ has four distinct roots in $\overline{\mathbf{Q}} \setminus \{0,-1\}$. Consequently, it will also be the minimum polynomial of Δ over $\mathbf{Q}$.

Let $\overline{\rho}_\Delta \in \mathbf{F}_p[X]$ be the reduction of ρ_Δ modulo p. If p is sufficiently large, $\overline{\rho}_\Delta$ will have four distinct roots in $\overline{\mathbf{F}}_p \setminus \{-1,0\}$. (The resultant $\mathrm{Res}(\rho_\Delta, \rho'_\Delta) \in \mathbf{Q}[X]$ is not the zero polynomial, by assumption, and so it is also not the zero polynomial when reduced mod p, for p sufficiently large; also $\rho_\Delta(0), \rho_\Delta(-1) \neq 0$ in $\mathbf{Q}$, and hence in $\mathbf{F}_p$ for p sufficiently large.) Since it has distinct roots, $\overline{\rho}_\Delta$ will be the minimal polynomial of Δ over $\mathbf{F}_p$.

Henceforth, suppose that $p \geqslant p_0(Q)$. Suppose that the factorisation of $\overline{\rho}_\Delta$ into irreducibles polynomials over $\mathbf{F}_p$ is $f_1 \cdots f_n$. Then, since $\overline{\rho}_\Delta$ has distinct roots in $\overline{\mathbf{F}}_p$, the f_i will be coprime. For each i, let $\alpha_i \in \overline{\mathbf{F}}_p$ be a root of f_i and consider the map

$$\Phi : \mathbf{F}_p[\Delta] \to \prod_{i=1}^{n} \mathbf{F}_p(\alpha_i)$$

given by

$$\Phi(F(\Delta)) = (F(\alpha_1), \dots, F(\alpha_n))$$

for any $F \in \mathbf{F}_p[X]$. This is a well-defined ring homomorphism: if $F_1(\Delta) = F_2(\Delta)$ then $\overline{\rho}_\Delta | F_1 - F_2$ (since $\overline{\rho}_\Delta$ is the minimal polyomial of Δ) and hence $f_i | F_1 - F_2$ for each i, whence $F_1(\alpha_i) = F_2(\alpha_i)$. We claim that Φ is injective. If $\Phi(F(\Delta)) = 0$ then for $i = 1, \dots, n$ we have $F(\alpha_i) = 0$ which implies $f_i | F$. Since the f_i are coprime, $\overline{\rho}_\Delta | F$ and so $F(\Delta) = 0$. Both the domain and range of Φ have size p^4 and so it is in fact a ring isomorphism.

Therefore there are five possible isomorphism types for the ring $\mathbf{F}_p[\Delta]$, namely $\mathbf{F}_{p^4}$ (if $\overline{\rho}_\Delta$ is irreducible over $\mathbf{F}_p$), $\mathbf{F}_{p^3} \times \mathbf{F}_p$, $\mathbf{F}_{p^2} \times \mathbf{F}_{p^2}$, $\mathbf{F}_{p^2} \times \mathbf{F}_p \times \mathbf{F}_p$, or $\mathbf{F}_p \times \mathbf{F}_p \times \mathbf{F}_p \times \mathbf{F}_p$ (if $\overline{\rho}_\Delta$ splits completely over $\mathbf{F}_p$.) Note that by standard algebraic number theory we can expect all of these possibilities to occur as p varies over primes.

The map Φ induces a group isomorphism

$$\Phi : \mathrm{SL}_2(\mathbf{F}_p[\Delta]) \to \prod_{i=1}^{n} \mathrm{SL}_2(\mathbf{F}_p(\alpha_i)).$$

In view of (10.7), we have

$$\Phi(\gamma(r,s)) = (M_{\alpha_i}(r,s))_{i=1}^n,$$

where the $M_{\alpha_i}(r,s)$ are as defined in Lemma 10.5. It follows from this and Lemma 10.5 that, if $\Gamma = \langle \gamma(r,s) : r,s \in \mathbf{F}_p^* \rangle = \tau\Gamma_p\tau^{-1}$ is the group generated by the $\gamma(r,s)$, then the projection of $\Phi(\Gamma)$ on to each factor $\mathrm{SL}_2(\mathbf{F}_p(\alpha_i))$ is surjective. When $n = 1$, this is the end of the proof, but we must work a little harder in the other cases.

Let us begin by looking at H, the projection of $\Phi(\Gamma)$ to the product $\mathrm{SL}_2(\mathbf{F}_p(\alpha_1)) \times \mathrm{SL}_2(\mathbf{F}_p(\alpha_2))$ of two of the factors (without loss of generality, the first two). Write π_i, $i = 1, 2$ for projection onto each factor. As we have remarked, $\pi_i(H) = \mathrm{SL}_2(\mathbf{F}_p(\alpha_i))$. This allows us to apply Goursat's lemma (Lemma 10.3). We conclude that there are $N_i \lhd \mathrm{SL}_2(\mathbf{F}_p(\alpha_i))$, and an isomorphism $\phi : \mathrm{SL}_2(\mathbf{F}_p(\alpha_1))/N_1 \to \mathrm{SL}_2(\mathbf{F}_p(\alpha_2))/N_2$ such that $H = \{(g_1, g_2) \in \mathrm{SL}_2(\mathbf{F}_p(\alpha_1)) \times \mathrm{SL}_2(\mathbf{F}_p(\alpha_2)) : \phi(\overline{g}_1) = \overline{g}_2\}$, where $\overline{g}_i$ denotes reduction mod N_i.

Now the $\mathrm{SL}_2(\mathbf{F}_p(\alpha_i))$ are almost simple: each N_i must be either trivial, $\{\pm I\}$ or $\mathrm{SL}_2(\mathbf{F}_p(\alpha_i))$. See Appendix A. Moreover, the fact that ϕ is an isomorphism, and the fact that $\mathrm{SL}_2(\mathbf{F}_p(\alpha_i))$ is not isomorphic to $\mathrm{PSL}_2(\mathbf{F}_p(\alpha_j))$ (consider cardinalities), means that up to relabelling there are only three essentially different cases, which we consider separately below.

Case 1. $N_1 = \mathrm{SL}_2(\mathbf{F}_p(\alpha_1))$. Then $N_2 = \mathrm{SL}_2(\mathbf{F}_p(\alpha_2))$, and H is the whole of the product $\mathrm{SL}_2(\mathbf{F}_p(\alpha_1)) \times \mathrm{SL}_2(\mathbf{F}_p(\alpha_2))$.

Case 2. $N_1 = N_2 = \{I\}$. Then $H = \{(x, \phi(x)) : x \in \mathrm{SL}_2(\mathbf{F}_p(\alpha_1))\}$, where $\phi : \mathrm{SL}_2(\mathbf{F}_p(\alpha_1)) \to \mathrm{SL}_2(\mathbf{F}_p(\alpha_2))$ is some isomorphism. Note that H contains the elements $(M_{\alpha_1}(r,s), M_{\alpha_2}(r,s))$, $r,s \in \mathbf{F}_p^*$, so in this scenario we must have

$$\phi(M_{\alpha_1}(r,s)) = M_{\alpha_2}(r,s) \tag{10.8}$$

for all r, s. By looking at cardinalities, the fields $\mathbf{F}_p(\alpha_1)$ and $\mathbf{F}_p(\alpha_2)$ must be isomorphic, so to ease notation we may suppose that $\alpha_2 \in \mathbf{F}_p(\alpha_1)$.

Now it is known (see Appendix A) that the automorphism group of $\mathrm{SL}_2(k)$ is generated by conjugation by elements of $\mathrm{GL}_2(k)$ and field automorphisms. Therefore for some $P \in \mathrm{GL}_2(\mathbf{F}_p(\alpha_1))$ and for some field automorphism σ of $\mathbf{F}_p(\alpha_1)$ we have

$$\phi(M_{\alpha_1}(r,s)) = PM_{\sigma(\alpha_1)}(r,s)P^{-1}.$$

Comparing with (10.8) gives

$$PM_{\sigma(\alpha_1)}(r,s) = M_{\alpha_2}(r,s)P \tag{10.9}$$

for all r, s. Writing $P = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$ and comparing top left entries gives, writing $\theta = \sigma(\alpha_1)$ and $\theta' = \alpha_2$,

$$a\Big(1+\theta-\frac{s}{r}\theta\Big)+b(s-r)\theta(1+\theta)=a\Big(1+\theta'-\frac{s}{r}\theta'\Big)+\Big(\frac{1}{s}-\frac{1}{r}\Big)c \tag{10.10}$$

for all $r, s \in \mathbf{F}_p^*$. By Lemma 10.2 (and since $\theta \neq 0, -1$) we have $b = c = 0$, thus $a \neq 0$ and $\theta = \theta'$. That is, $\sigma(\alpha_1) = \alpha_2$, and so α_1, α_2 have the same minimal polynomial over $\mathbf{F}_p$. This is a contradiction, since we assumed that the minimal polynomials of α_1, α_2 over $\mathbf{F}_p$ (the f_i, that is to say the factors of the minimal polynomial of Δ over $\mathbf{F}_p$) are coprime.

Case 3. $N_i = N_j = \{\pm I\}$. We reprise the argument from Case 2, only now we must allow a sign error. Included in the classification of automorphisms of $\mathrm{PSL}_2(\mathbf{F}_p(\alpha_1))$ (see Appendix A) is the fact that such automorphisms lift to automorphisms of $\mathrm{SL}_2(\mathbf{F}_p(\alpha_1))$. Therefore $H = \{(x, \varepsilon(x)\phi(x)) : x \in \mathrm{SL}_2(\mathbf{F}_p(\alpha_1)), \varepsilon(x) \in \pm I\}$ for some isomorphism $\phi : \mathrm{SL}_2(\mathbf{F}_p(\alpha_1)) \to \mathrm{SL}_2(\mathbf{F}_p(\alpha_2))$. We may now proceed as before but with an additional sign error, thus (10.9) becomes

$$PM_{\sigma(\alpha_1)}(r,s) = \varepsilon_{r,s} M_{\alpha_2}(r,s)P \tag{10.11}$$

for all $r, s \in \mathbf{F}_p^*$ and for some choice of signs $\varepsilon_{r,s} \in \{\pm 1\}$. If $\varepsilon_{r,s} = 1$ for at least half of all pairs $(r,s) \in \mathbf{F}_p^* \times \mathbf{F}_p^*$ then we are done, exactly as before (taking $\eta = \frac{1}{4}$ in Lemma 10.2). If $\varepsilon_{r,s} = -1$ for at least half of all pairs $(r,s) \in \mathbf{F}_p^* \times \mathbf{F}_p^*$ then (10.10) is modified to

$$a\Big(1+\theta-\frac{s}{r}\theta\Big)+b(s-r)\theta(1+\theta)=-a\Big(1+\theta'-\frac{s}{r}\theta'\Big)-\Big(\frac{1}{s}-\frac{1}{r}\Big)c, \tag{10.12}$$

for half of all pairs $(r,s) \in \mathbf{F}_p^* \times \mathbf{F}_p^*$. From this we conclude that $b = c = 0$, hence $a \neq 0$ and so both $\theta = -\theta'$ and $1 + \theta = -(1 + \theta')$. This is impossible.

Since only Case 1 in the above analysis did not lead to a contradiction (and since we can replace $\{1,2\}$ by any pair $\{i,j\}$), we have now shown that the projection of $\Phi(\Gamma)$ to the product $\mathrm{SL}_2(\mathbf{F}_p(\alpha_i)) \times \mathrm{SL}_2(\mathbf{F}_p(\alpha_j))$ of any pair of factors is surjective. Proposition 10.1 now follows from Lemma 10.4, together with the fact (see Appendix A) that all the factors $\mathrm{SL}_2(\mathbf{F}_p(\alpha_i))$ are perfect. □

11 Quasirandomness of $\rho|_{\Gamma_q}$

We turn now to the proof of Proposition 8.5 itself. Let us begin by recalling the statement.

Proposition 11.1 (Proposition 8.5). *Let Q be a generic quadratic form. Then there is some $p_0(Q)$ such that the following is true. Let ρ be the Weil representation on $\mathrm{Sp}_8(\mathbf{Z}/q\mathbf{Z})$ (as given in Proposition 7.2). Suppose that q is squarefree and has all prime factors greater than $p_0(Q)$. Then $\rho|_{\Gamma_q}$ splits into irreducible subrepresentations of dimensions $\geqslant q^{1-o(1)}$.*

We begin by reducing to the prime case. Write $q = p_1 \cdots p_n$. The representation ρ is constructed in Appendix B as a tensor product $\otimes_{i=1}^n \tilde{\rho}_i$, where $\tilde{\rho}_i : \mathrm{Sp}_8(\mathbf{Z}/p_i\mathbf{Z}) \to \mathrm{U}(\ell^2((\mathbf{Z}/p_i\mathbf{Z})^4)$ is a twisted version of the mod p_i Weil representation ρ_i, given by $\tilde{\rho}_i(g) = \rho_i(g^{\sigma_i})$ where $\begin{pmatrix} M_1 & M_2 \\ M_3 & M_4 \end{pmatrix}^{\sigma_i} = \begin{pmatrix} M_1 & \lambda_i M_2 \\ \lambda_i^{-1} M_3 & M_4 \end{pmatrix}$ and $\lambda_i = \prod_{j \neq i} p_i$. This tensor product may be realised on $\ell^2((\mathbf{Z}/q\mathbf{Z})^4)$ by $\rho(g)f = \prod_{i=1}^n \tilde{\rho}_i(g) f_i(x)$ for "pure tensors" $f(x) = \prod_{i=1}^n f_i(x)$, where f_i factors through the projection $\pi_i : (\mathbf{Z}/q\mathbf{Z})^4 \to (\mathbf{Z}/p_i\mathbf{Z})^4$ (see Appendix B for a discussion of the notation here). Consider the restriction to $\Gamma_q \cong \prod_{i=1}^n \Gamma_{p_i}$. The decomposition of $\rho|_{\Gamma_q}$ into irreducibles is then given by decomposing each $\ell^2((\mathbf{Z}/p_i\mathbf{Z})^4)$ into irreducible $\tilde{\rho}_i|_{\Gamma_{p_i}}$-invariant subspaces V_i and taking tensor products. (Here we use the fact that if V_i is an irreducible G_i representation then $\bigotimes_{i=1}^n V_i$ is an irreducible $\times_{i=1}^n G_i$-representation, which is a standard fact of representation theory. See for instance [JL01, Theorem 19.18].)

Now each Γ_{p_i} is a direct product of groups $\mathrm{SL}_2(\mathbf{F}_{p_i^j})$, and therefore by Appendix A any irreducible representation of Γ_{p_i} is either trivial, or has dimension at least $\frac{1}{2}(p_i - 1)$. That is, if $\dim V_i \neq 1$ then $\dim V_i \geqslant \frac{1}{2}(p_i - 1)$. Consequently, if for all i the representation $\tilde{\rho}_i|_{\Gamma_{p_i}}$ has no invariant vector (that is, 1-dimensional invariant subspace) then $\dim \rho \geqslant \prod_{p|q} \frac{1}{2}(p-1) = q^{1-o(1)}$, as desired.

This reduces the task of proving Proposition 8.5 to the following, which is the final task for the main part of the paper.

Proposition 11.2. *Suppose that Q is generic and that p is sufficiently large in terms of Q. Let $\rho : \mathrm{Sp}_8(\mathbf{Z}/p\mathbf{Z}) \to \mathrm{U}(\ell^2((\mathbf{Z}/p\mathbf{Z})^4))$ be the mod p Weil representation. Let $r \in (\mathbf{Z}/p\mathbf{Z})^*$, and let $\tilde{\rho}$ be the twist of ρ by dilation by r, that is to say $\tilde{\rho}(x) = \rho(x^{(r)})$. Then $\tilde{\rho}|_{\Gamma_p}$ has no nontrivial invariant vector.*

Proof. We will show that the conclusion holds under the assumption that Δ has distinct eigenvalues and is invertible over $\mathbf{F}_p$. This includes all sufficiently large primes p. Indeed Q is generic, so by definition Δ has distinct eigenvalues and is invertible over $\mathbf{Q}$. Therefore the same is true over $\mathbf{F}_p$, p sufficiently large, for the reasons detailed at the start of the proof of Proposition 10.1.

Suppose from now on that Δ has distinct eigenvalues and is invertible over $\mathbf{F}_p$. Since Γ_p is invariant under the dilation $\gamma \mapsto \gamma^{(r)}$, it suffices to consider the case $\tilde{\rho} = \rho$. Suppose, then, that

$$\rho(\gamma) f = f \tag{11.1}$$

for all $\gamma \in \Gamma_p$. Our aim is to show that f is identically zero. To examine the condition (11.1) we will look at the following particular elements γ, where $B, C \in \mathrm{Mat}_4(\mathbf{F}_p)$ are the specific symmetric matrices described in (10.3):

(1) the upper triangular elements $u(MB) = \begin{pmatrix} I & MB \\ 0 & I \end{pmatrix}$, where $M \in \mathbf{F}_p[\Delta]$;
(2) the lower triangular elements $l(CM) = \begin{pmatrix} I & 0 \\ CM & I \end{pmatrix}$, where $M \in \mathbf{F}_p[\Delta]$;
(3) the diagonal elements $s(\lambda I) = \begin{pmatrix} \lambda I & 0 \\ 0 & \lambda^{-1} I \end{pmatrix}$, $\lambda \in \mathbf{F}_p^*$.

Now that we know from Proposition 10.1 that $\Gamma_p = \tau^{-1}\,\mathrm{SL}_2(\mathbf{F}_p[\Delta])\tau$ (where τ is defined in (10.5)), so one may easily check using (10.6) that all of these elements do lie in Γ_p.

Now we already have formulae for the actions of the elements in (2) and (3), directly from Proposition 7.2. Namely,

$$\rho(l(CM))f(x) = \xi_M e_p\Big(-\frac{1}{2}x^T CMx\Big)f(x) \tag{11.2}$$

for some unit complex number $\xi_M = \xi(l(CM))$ (we do not care exactly what this is) and

$$\rho(s(\lambda I))f(x) = \xi'_\lambda f(\lambda^{-1}x) \tag{11.3}$$

for some unit complex number ξ'_λ. If (11.1) holds, it follows from (11.2) that for each M

$$\mathrm{Supp}(f) \subset \{x : x^T CMx = t_M\}$$

for some parameters t_M (in fact satisfying $e_p(-\frac{1}{2}t_M) = \xi_M$). Now (11.3) tells us that $\mathrm{Supp}(f(x)) = \mathrm{Supp}(f(\lambda^{-1}x))$ for all $\lambda \in \mathbf{F}_p^*$, therefore (taking $\lambda \neq \pm 1$) we see that in fact all the t_M must be zero, that is

$$\mathrm{Supp}(f) \subset \{x : x^T CMx = 0\} \quad \text{for all } M \in \mathbf{F}_p[\Delta]. \tag{11.4}$$

To get a formula for the upper triangular action, we note the identity

$$u(W) = J^{-1}l(-W)J$$

for any matrix W, where as usual $J = \left(\begin{smallmatrix} 0 & I \\ -I & 0 \end{smallmatrix}\right)$. Thus if f is invariant under $\rho(u(MB))$ then $\rho(J)f$ is invariant under $\rho(l(-MB))$. As above, this implies that $\mathrm{Supp}(\rho(J)f) \subset \{x : x^T MBx = t'_M\}$. Since $J^{-1}s(\lambda I)J = s(\lambda^{-1}I)$, we see that $\rho(J)f$ is invariant (up to multiplication by a scalar) by dilation $x \mapsto \lambda^{-1}x$. Therefore all of the t'_M are in fact zero. Moreover, from Proposition 7.2 we know that $\rho(J)f$ is the (normalised) Fourier transform of f times a scalar, and so we come to the conclusion that

$$\mathrm{Supp}(\hat{f}) \subset \{x : x^T MBx = 0\} \quad \text{for all } M \in \mathbf{F}_p[\Delta]. \tag{11.5}$$

Conditions (11.4) and (11.5) seem highly incompatible and should, for example, violate the uncertainty principle under reasonable assumptions. However, a proof seems not to be completely straightforward (and indeed the statement fails in sufficiently degenerate situations, for example if we were to allow $\Delta = 0$).

Let us begin the argument. Consider the bilinear form $\phi(x, y) = x^T Cy$ on $\overline{\mathbf{F}}_p^4 \times \overline{\mathbf{F}}_p^4$. Since $BC = \Delta$ and Δ is invertible over $\mathbf{F}_p$, C is invertible over $\mathbf{F}_p$ and

so this is a non-degenerate form. We note that Δ is self-adjoint with respect to ϕ:

$$\phi(x, \Delta y) = x^T CBCy,$$

$$\phi(\Delta x, y) = x^T \Delta^T Cy = x^T CBCy$$

(note that B, C are both symmetric). By the usual argument, this means that eigenvectors of Δ with distinct eigenvalues are orthogonal with respect to ϕ. Indeed, if $\Delta v_1 = \lambda_1 v_1$ and $\Delta v_2 = \lambda_2 v_2$ then

$$\lambda_2 \phi(v_1, v_2) = \phi(v_1, \Delta v_2) = \phi(\Delta v_1, v_2) = \lambda_1 \phi(v_1, v_2).$$

Now since Δ has distinct eigenvalues, it is diagonalisable over $\overline{\mathbf{F}}_p$. Suppose that v_i are eigenvectors with (distinct) eigenvalues λ_i, $i = 1, 2, 3, 4$. These are a basis for $\overline{\mathbf{F}}_p^4$. Write $x \in \overline{\mathbf{F}}_p^4$ as $\sum_{i=1}^4 x_i v_i$. Then

$$x^T C \Delta^j x = \phi(x, \Delta^j x) = \sum_{i=1}^4 x_i^2 \lambda_i^j \phi(v_i, v_i). \tag{11.6}$$

Note that $\phi(v_i, v_i) \neq 0$ (if it was, v_i would be orthogonal with respect to ϕ to all of v_1, v_2, v_3, v_4 and hence to all of $\overline{\mathbf{F}}_p^4$, contrary to the fact that ϕ is non-degenerate).

Therefore the matrix with (i, j)-entry $\lambda_i^j \phi(v_i, v_i)$ $(1 \leqslant i \leqslant 4,\ 0 \leqslant j \leqslant 3)$ is non-singular, its determinant being a non-zero multiple $\prod_{i=1}^4 \phi(v_i, v_i)$ of a certain Vandermonde determinant. It follows from (11.6) that if $x^T C \Delta^j x = 0$ for $j = 0, 1, 2, 3$ then $x = 0$, and so any f satisfying (11.4) is supported only at zero.

Noting that $x^T \Delta^{j+1} Bx = (Bx)^T C \Delta^j (Bx)$, we can also conclude that if $x^T \Delta^{j+1} \times Bx = 0$ for $j = 0, 1, 2, 3$ then $x = 0$, and so any f satisfying (11.5) has $\hat{f}$ supported only at zero.

These two facts about f are completely incompatible, unless f is identically zero: if f is supported at zero, $\hat{f}$ is in fact constant on $(\mathbf{Z}/p\mathbf{Z})^4$.

This completes the proof of Proposition 11.2, and hence that of Proposition 8.5. □

All of the main results in the paper are now established.

Remark. It is in fact possible to show (under the assumptions on Δ in force throughout this section) that $\rho|_{\Gamma_p}$ is isomorphic to a Weil representation of $\mathrm{SL}_2(\mathbf{F}_p[\Delta])$ on $\ell^2(\mathbf{F}_p[\Delta])$, by giving an explicit intertwining map. This fact can be used to give an alternative proof of Proposition 11.2 which, while more natural than the *ad hoc* arguments presented here, requires quite a bit more setting up. We intend to give a full account in future work.

Appendix A: Facts about $\mathrm{SL}_2(k)$

We collect various well-known facts about $\mathrm{SL}_2(k)$, k a finite field, which we used in the main text. For our purposes, "rough" versions of these facts (passing to subgroups of index $O(1)$, etc) would be quite sufficient but we use the precise versions when sufficiently clean results are relatively easily-available.

PROPOSITION A.1. *Let k be a finite field of odd characteristic. Then*

(1) *The smallest nontrivial complex representation of $\mathrm{PSL}_2(k)$ has dimension at least $\frac{1}{2}(|k|-1)$.*
(2) *If k has order at least 5 then $\mathrm{SL}_2(k)$ is perfect.*
(3) *Any proper subgroup of $\mathrm{SL}_2(k)$ has a subgroup of index at most C which is conjugate to a subgroup of one of* (i) *the group of upper triangular matrices or* (ii) *$\mathrm{SL}_2(k_0)$ for some proper subfield $k_0 < k$.*
(4) *Every automorphism of $\mathrm{PSL}_2(k)$ or $\mathrm{SL}_2(k)$ is a composition of a conjugation by elements of $\mathrm{GL}_2(k)$, and a power of the Frobenius automorphism of k.*

Proof. (1) goes back well over a century, to Jordan and Schur. For a nice and easy-to-access discussion, see Prasad's notes [Pra].

(2) See Lang [Lan02, Chapter XIII, Theorem 8.3].

(3) The rough statement given here, which suffices for our purposes is [Tao15, Theorem 5.2.7]. As one would expect, a detailed classification of maximal subgroups of $\mathrm{SL}_2(k)$ has been known for more than a century. It is somewhat complicated; the details, as well as references to the original papers, may be found in [Kin].

(4) This is certainly well-known. A standard reference is [Ste60, 3.2]. See also [Ste16]. The MathOverflow post [Wha] is helpful in navigating these papers. □

Appendix B: Weil representation of $\mathrm{Sp}_8(\mathbf{Z}/q\mathbf{Z})$

In the section we construct the representation of $\rho : \mathrm{Sp}_8(\mathbf{Z}/q\mathbf{Z}) \to \mathrm{U}(\ell^2((\mathbf{Z}/q\mathbf{Z})^4)$ which we have been using throughout the paper, and whose properties are detailed in Proposition 7.2. Suppose that $q = p_1 \cdots p_n$ is squarefree and odd. We assume the existence, for each i, of the Weil representations $\rho_i : \mathrm{Sp}_8(\mathbf{Z}/p_i\mathbf{Z}) \to \mathrm{U}(\ell^2((\mathbf{Z}/p_i\mathbf{Z})^4))$, the construction of which is given in detail in [Neu02] and shown to satisfy the properties of Proposition 7.2 (and in fact that paper gives details of the multiplier ξ, whose precise properties are unimportant in this paper). Since $\mathrm{Sp}_8(\mathbf{Z}/q\mathbf{Z}) \cong \prod_{i=1}^n \mathrm{Sp}_8(\mathbf{Z}/p_i\mathbf{Z})$, one thinks of looking at the (exterior) tensor product $\rho = \otimes_{i=1}^n \rho_i$. However, this turns out to need a small modification.

For each i, denote by $\sigma_i : \mathrm{Sp}_8(\mathbf{Z}/p_i\mathbf{Z}) \to \mathrm{Sp}_8(\mathbf{Z}/p_i\mathbf{Z})$ the automorphism defined by $\left(\begin{smallmatrix} M_1 & M_2 \\ M_3 & M_4 \end{smallmatrix}\right)^{\sigma_i} := \left(\begin{smallmatrix} M_1 & \lambda_i M_2 \\ \lambda_i^{-1} M_3 & M_4 \end{smallmatrix}\right)$, where $\lambda_i := \prod_{i \neq j} p_i$. Note that this is, in fact, one of the dilates we considered earlier (see (7.3)), but with parameter $r = \lambda_i^{-1}$. It may not be an inner automorphism (this depends on whether or not λ_i is a square mod p_i). Each ρ_i may be twisted by σ_i to give a representation $\tilde{\rho}_i : \mathrm{Sp}_8(\mathbf{Z}/p_i\mathbf{Z}) \to \mathrm{U}(\ell^2((\mathbf{Z}/p_i\mathbf{Z})^4))$,

defined by $\tilde{\rho}_i(g)f := \rho(g^{\sigma_i})f$. This will be isomorphic to ρ_i if λ_i is a square mod p_i, but not otherwise. However, we will not need this last fact.

We now construct ρ as the tensor product $\bigotimes_{i=1}^n \tilde{\rho}_i$, which we shall shortly show how to realise concretely in $\ell^2((\mathbf{Z}/q\mathbf{Z})^4)$. From here on, we will abuse notation by omitting explicit notation for projection maps from $(\mathbf{Z}/q\mathbf{Z})^4$ to $(\mathbf{Z}/p_i\mathbf{Z})^4$, from $\mathrm{Sp}_8(\mathbf{Z}/q\mathbf{Z})$ to $\mathrm{Sp}_8(\mathbf{Z}/p_i\mathbf{Z})$, or from $\mathrm{Mat}_4(\mathbf{Z}/q\mathbf{Z})$ to $\mathrm{Mat}_4(\mathbf{Z}/p_i\mathbf{Z})$ when the domain is clear from context. Thus, for example, for functions $f_i \in \ell^2((\mathbf{Z}/p_i\mathbf{Z})^4)$ we define the "pure tensor" $f(x) := \prod_{i=1}^n f_i(x)$, but it would be more correct, though cumbersome, to write $\prod_{i=1}^n f_i(\pi_i(x))$ where $\pi_i : (\mathbf{Z}/q\mathbf{Z})^4 \to (\mathbf{Z}/p_i\mathbf{Z})^4$ is the natural projection.

For a pure tensor f as above define

$$\rho(g)f(x) = \prod_{i=1}^n \tilde{\rho}_i(g)f_i(x) = \prod_{i=1}^n \rho_i(g^{\sigma_i})f_i(x). \tag{B.1}$$

This is well-defined by the universal property of tensor products.

We now turn to the verification of the properties stated in Proposition 7.2. Recall that the properties to be established are as follows (for some unit complex numbers $\xi(\cdot)$):

$$\rho(s(E))f(x) = \xi(s(E))f(E^{-1}x) \tag{B.2}$$

for $s(E) := \left(\begin{smallmatrix} E & 0 \\ 0 & E^{-T} \end{smallmatrix}\right)$ with E invertible;

$$\rho(J)f(x) = \xi(J)q^2\mathbb{E}_{y\in(\mathbf{Z}/q\mathbf{Z})^4} f(y)e(x^Ty) \tag{B.3}$$

where $J := \left(\begin{smallmatrix} 0 & I \\ -I & 0 \end{smallmatrix}\right)$, and

$$\rho(l(W))f(x) = \xi(l(W))e_q\Big(-\frac{1}{2}x^TWx\Big)f(x) \tag{B.4}$$

for $l(W) := \left(\begin{smallmatrix} I & 0 \\ W & I \end{smallmatrix}\right)$ with W symmetric.

We will also, of course, be using the corresponding properties for the prime case ρ_i. To avoid confusion, we write $\xi_i(\cdot)$ for the corresponding unit complex numbers. It is enough to check (B.2), (B.3) and (B.4) for pure tensors f.

Proof of (B.2). We have

$$\begin{aligned}\rho(s(E))f(x) &= \prod_{i=1}^n \rho_i(s(E)^{\sigma_i})f_i(x) = \prod_{i=1}^n \rho_i(s(E))f_i(x)\\ &= \prod_{i=1}^n \xi_i(s(E))f_i(E^{-1}x) = \xi(s(E))f(E^{-1}x),\end{aligned}$$

where $\xi(s(E)) := \prod_{i=1}^n \xi_i(s(E))$. This establishes (B.2).

For the remaining two parts, we will need the relation

$$e_q(\lambda t) = \prod_{i=1}^n e_p(t), \tag{B.5}$$

where $\lambda = \sum_{i=1}^n \lambda_i$.

Proof of (B.3). We have

$$\rho(J)f(x) = \prod_{i=1}^{n} \rho_i(J^{\sigma_i}) f_i(x) = \prod_{i=1}^{n} \rho_i(s(\lambda_i I)J) f_i(x)$$
$$= q^2 \xi(J) \prod_{i=1}^{n} \mathbb{E}_{y_i \in (\mathbf{Z}/p_i\mathbf{Z})^4} f_i(y_i) e_{p_i}(\lambda_i^{-1} x^T y_i),$$

where $\xi(J) := \prod_{i=1}^{n} \xi_i(s(\lambda_i I))\xi_i(J)$. Applying (B.5) gives

$$\prod_{i=1}^{n} e_{p_i}(\lambda_i^{-1} x^T y_i) = e_q(x^T y),$$

where $y \in (\mathbf{Z}/q\mathbf{Z})^4$ reduces to y_i in $(\mathbf{Z}/p_i\mathbf{Z})^4$, and the claim follows.

Proof of (B.4). We have

$$\rho(l(W))f(x) = \prod_{i=1}^{n} \rho_i(l(W)^{\sigma_i}) f_i(x) = \prod_{i=1}^{n} \rho_i(l(\lambda_i^{-1} W)) f_i(x)$$
$$= \xi(l(W)) \prod_{i=1}^{n} e_{p_i}\Big(-\frac{1}{2} x^T \lambda_i^{-1} W x \Big) f_i(x),$$

where $\xi(l(W)) := \prod_{i=1}^{n} \xi_i(l(\lambda_i^{-1} W))$. The claim again follows using (B.5).

Appendix C: Almost-invariant measures

We invite the reader to recall the notation concerning probability measures on finite groups G as described at the start of Sect. 8. In particular, $\|\cdot\|$ means the 2-norm with respect to the counting measure on G, $\|\mu\| = \big(\sum_{x \in G} \mu(x)^2\big)^{1/2}$. We extend the notation μ_H for the uniform measure on a subgroup H to arbitrary sets: thus if $A \subset G$ is a finite set then we write μ_A for the uniform measure on A, that is to say the measure which puts weight $|A|^{-1}$ on each point of A.

We have the following instance of Young's inequality.

Lemma C.1 (Young's inequality). *Let μ, ν be two probability measures on G. Then we have $\|\mu * \nu\| \leqslant \|\mu\|$.*

Proof. We have

$$\mu * \nu(x) = \sum_{y \in G} \Big(\mu(y)^2 \nu(y^{-1}x)\Big)^{1/2} \nu(y^{-1}x)^{1/2}.$$

By Cauchy–Schwarz and the fact that ν is a probability measure,

$$\mu * \nu(x)^2 \leqslant \sum_{y \in G} \mu(y)^2 \nu(y^{-1}x).$$

Finally, summing over $x \in G$ gives the result. □

In the main text we required a statement about almost equality here in the case that $\mu = \nu$ and both are symmetric. The actual statement we quoted in the main text is Corollary C.3 below, but the heart of it is Lemma C.2. This result should be thought of as "well-known", but it is hard to give a precise reference. The basic idea of the proof goes back to Fournier [Fou77]; see [ET12, Proposition 5.4] for the abelian case or [B+15, Appendix A] for a closely related result.

LEMMA C.2. *Let μ be a symmetric probability measure on a finite group G. Let ε be sufficiently small positive constant. Suppose that $\|\mu * \mu\| \geqslant (1-\varepsilon)\|\mu\|$. Then there is a subgroup $H \leqslant G$ such that $\|\mu - \mu_H\| \ll \varepsilon^c \|\mu\|$.*

Proof. We assume throughout the proof that ε is sufficiently small. Set

$$A := \{x : \mu * \mu(x) \geqslant (1 - 2\varepsilon^{1/2})\|\mu\|^2\}. \tag{C.1}$$

Then A is a symmetric set, and we have (using the pointwise bound $\mu * \mu(x) \leqslant \|\mu\|^2$)

$$\begin{aligned}(1-2\varepsilon)\|\mu\|^2 \leqslant \|\mu * \mu\|^2 &= \sum_x \mu * \mu(x)^2 \\ &\leqslant \|\mu\|^2 \mu * \mu(A) + (1 - 2\varepsilon^{1/2})\|\mu\|^2 \mu * \mu(A^c).\end{aligned}$$

Writing $\mu * \mu(A) = 1 - \delta$, so that $\mu * \mu(A^c) = \delta$, this rearranges to give $\delta \leqslant \varepsilon^{1/2}$, that is to say $\mu * \mu(A) \geqslant 1 - \varepsilon^{1/2}$. Since

$$\mu * \mu(A) = \sum_y \mu(y)\mu(Ay^{-1}),$$

it follows that there is some $B := Ay^{-1}$ such that $\mu(B) \geqslant 1 - \varepsilon^{1/2}$. Since μ is symmetric, $\mu(B^{-1}) \geqslant 1 - \varepsilon^{1/2}$. Therefore, setting $S := B \cap B^{-1}$, we see that S is symmetric and $\mu(S) \geqslant 1 - 2\varepsilon^{1/2}$. From (C.1) we have

$$|S| \leqslant |A| \leqslant (1 + 4\varepsilon^{1/2})\|\mu\|^{-2}.$$

Now we have

$$\|\mu - \mu_S\|^2 = \|\mu\|^2 - \frac{1}{|S|}(2\mu(S) - 1) \leqslant \|\mu\|^2\Big(1 - \frac{1 - 4\varepsilon^{1/2}}{1 + 4\varepsilon^{1/2}}\Big) < 8\varepsilon^{1/2}\|\mu\|^2$$

and so

$$\|\mu - \mu_S\| = O(\varepsilon^{1/4})\|\mu\|. \tag{C.2}$$

Note that this implies

$$\frac{1}{2}\|\mu\| \leqslant \|\mu_S\| \leqslant 2\|\mu\|. \tag{C.3}$$

Writing $\mu_S = \mu + (\mu_S - \mu)$ and expanding and using the triangle inequality, we have

$$\|\mu_S * \mu_S\| \geqslant \|\mu * \mu\| - 2\|\mu * (\mu_S - \mu)\| - \|(\mu_S - \mu) * (\mu_S - \mu)\|.$$

By Young's inequality, (C.2), (C.3) and the assumption of the lemma it follows that

$$\|\mu_S * \mu_S\| \geqslant (1 - O(\varepsilon^{1/4})\|\mu\| \geqslant (1 - O(\varepsilon^{1/4}))\|\mu_S\|. \tag{C.4}$$

At this point we have essentially reduced the proof of the lemma to the case of a uniform measure on a set. Equation (C.4) is equivalent to the statement that the number of multiplicative quadruples $s_1 s_2 = s_3 s_4$ in S is $(1 - O(\varepsilon^{1/4})|S|^3$. This is a well-known situation and (for example) Fournier [Fou77] implies that there is a subgroup H such that $|S \triangle H| = O(\varepsilon^c)\min(|H|, |S|)$. Therefore

$$\|\mu_S - \mu_H\|^2 = \frac{|S \triangle H|}{|S||H|} \ll \varepsilon^c \|\mu_S\|^2 \ll \varepsilon^c \|\mu\|^2.$$

The result follows from this, (C.2) and the triangle inequality. □

Finally we give the result actually quoted in the main text.

Corollary C.3. *Let μ be a symmetric probability measure on a finite group G. Let ε be a sufficiently small constant. Suppose that $\|\mu * \mu\| \geqslant (1 - \varepsilon)\|\mu\|$. Then there is a subgroup $H \leqslant G$ such that $\|\mu - \mu_H\| \ll \varepsilon^c \|\mu_H\|$, $\mu(H) \geqslant 1 - O(\varepsilon^c)$ and $|\operatorname{Supp}(\mu)| \geqslant (1 - O(\varepsilon^c))|H|$.*

Proof. Let H be as in the conclusion of Lemma C.2. The first statement is just the conclusion of Lemma C.2. From it we deduce

$$\|\mu\| = (1 + O(\varepsilon^c))\|\mu_H\| = (1 + O(\varepsilon^c))|H|^{-1/2}. \tag{C.5}$$

In particular,

$$\langle \mu, \mu_H \rangle = \frac{1}{2}\big(\|\mu\|^2 + \|\mu_H\|^2 - \|\mu - \mu_H\|^2\big) = (1 + O(\varepsilon^c))|H|^{-1}. \tag{C.6}$$

Now we have $\langle \mu, \mu_H \rangle = |H|^{-1}\mu(H)$, and so the second statement follows.

For the third statement, write $A := \operatorname{Supp}(\mu) \cap H$. By Cauchy–Schwarz we have

$$\langle \mu, \mu_H \rangle = \frac{1}{|H|}\sum_x \mu(x) 1_A(x) \leqslant \frac{1}{|H|}\|\mu\||A|^{1/2}.$$

The required bound now follows by comparing this with (C.5) and (C.6). □

Acknowledgements

It is a pleasure to thank Emmanuel Breuillard, Charlotte Chan, Tom Fisher and Balázs Szendrői for helpful correspondence related to this work and earlier versions of it, and Roger Baker and James Maynard for discussions which introduced me to the problem in around 2014. I thank the anonymous referee for a careful reading of the paper. The author is a Simons Investigator and is very grateful to the Simons Foundation for their continued support.

References

[Bou12] Bourgain, J.: A modular Szemerédi-Trotter theorem for hyperbolas. *C. R. Math. Acad. Sci. Paris* **350**(17–18), 793–796 (2012).

[BG08] Bourgain, J., Gamburd, A.: Uniform expansion bounds for Cayley graphs of $\mathrm{SL}_2(\mathbf{F}_p)$. *Ann. Math.* **167**(2), 625–642 (2008).

[BGS10] Bourgain, J., Gamburd, A., Sarnak, P.: Affine linear sieve, expanders, and sum-product. *Invent. Math.* **179**(3), 559–644 (2010).

[B+15] Breuillard, E., Green, B.J., Guralnick, R., Tao, T.: Expansion in finite simple groups of Lie type. *J. Eur. Math. Soc.* **17**(6), 1367–1434 (2015).

[Cha] Chan, C.: The Weil Representation. Stanford Senior Honors Thesis available at http://www-personal.umich.edu/~charchan/TheWeilRepresentation.pdf.

[ET12] Eisner, T., Tao, T.: Large values of the Gowers-Host-Kra seminorms. *J. Anal. Math.* **117**, 133–186 (2012).

[Fou77] Fournier, J.J.: Sharpness in Young's inequality for convolution. *Pac. J. Math.* **72**, 383–397 (1977).

[Gow08] Gowers, W.T.: Quasirandom groups. *Comb. Probab. Comput.* **17**(3), 363–387 (2008).

[Hua38] Hua, L.K.: Some results in additive prime number theory. *Q. J. Math.* **9**, 60–80 (1938).

[JL01] James, G., Liebeck, M.: Representations and Characters of Groups, 2nd edn. Cambridge University Press, Cambridge (2001).

[Lan02] Lang, S.: Algebra, 3rd edn. Springer, Berlin (2002).

[Liu11] Liu, J.: Integral points on quadrics with prime coordinates. *Monatshefte Math.* **164**, 439–465 (2011).

[Kin] King, O.: The subgroup structure of finite classical groups in terms of geometric configurations. In: Surveys in Combinatorics 2005. London Math. Soc. Lecture Notes, vol. 327, pp. 29–56.

[Kow14] Kowalski, E.: An Introduction to the Representation Theory of Groups. Graduate Studies in Mathematics, vol. 155. Am. Math. Soc., Providence (2014).

[PS] Pillay, A., Starchenko, S.: Remarks on Tao's algebraic regularity lemma. Unpublished note available at arXiv:1310.7538.

[Neu02] Neuhauser, M.: An explicit construction of the metaplectic representation over a finite field. *J. Lie Theory* **12**, 15–30 (2002).

[Pra] Prasad, A.: Representations of $\mathrm{GL}_2(\mathbf{F}_q)$ and $\mathrm{SL}_2(\mathbf{F}_q)$, and some remarks about $\mathrm{GL}_n(\mathbf{F}_q)$. arXiv:0712.4051.

[SX91] Sarnak, P., Xue, X.X.: Bounds for multiplicities of automorphic representations. *Duke Math. J.* **64**(1), 207–227 (1991).

[Ser16] Serre, J.-P.: Finite Groups: An Introduction. Surveys of Modern Mathematics, vol. 10. Int. Press, Somerville (2016).

[Shk21] Shkredov, I.: Modular hyperbolas and bilinear forms of Kloosterman sums. *J. Number Theory* **220**, 182–211 (2021).

[Ste60] Steinberg, R.: Automorphisms of finite linear groups. *Can. J. Math.* **12**, 606–615 (1960).

[Ste16] Steinberg, R.: Lectures on Chevalley Groups. University Lecture Series, vol. 66. Am. Math. Soc., Providence (2016). Revised and corrected edition of the 1968 original.

[Sze98] Szechtman, F.: Weil representations of the symplectic group. *J. Algebra* **208**, 662–686 (1998).

[Tao15] Tao, T.: Expanding polynomials over finite fields of large characteristic, and a regularity lemma for definable sets. *Contrib. Discrete Math.* **10**(1), 22–98 (2015).

[Tao] Tao, T.: A spectral theory proof of the algebraic regularity lemma. Blog post available at https://terrytao.wordpress.com/2013/10/29/a-spectral-theory-proof-of-the-algebraic-regularity-lemma/.

[Tao15] Tao, T.: Expansion in Finite Simple Groups of Lie Type. Graduate Studies in Mathematics, vol. 164. Am. Math. Soc., Providence (2015).

[Var12] Varjú, P.P.: Expansion in $\mathrm{SL}_d(\mathcal{O}_K/I)$, I squarefree. *J. Eur. Math. Soc.* **14**(1), 273–305 (2012).

[Wei64] Weil, A.: Sur certains groupes d'opérateurs unitaires. *Acta Math.* **111**, 143–211 (1964).

[Zha16] Zhao, L.: The quadratic form in nine prime variables. *Nagoya Math. J.* **223**(1), 21–65 (2016).

[Wha] What is the outer automorphism group of $\mathrm{SL}(2,\mathbf{F}_q)$? Mathoverflow discussion available at https://mathoverflow.net/questions/348440/what-is-the-outer-automorphism-group-of-operatornamesl2-mathbbf-q.

Ben Green
Mathematical Institute, Radcliffe Observatory Quarter, Woodstock Road, Oxford OX2 6GG, England, UK.

ben.green@maths.ox.ac.uk

Received: 29 September 2025
Accepted: 18 November 2025