WITT VECTORS AND A QUESTION OF ENTIN, KEATING, AND RUDNICK

NICHOLAS M. KATZ

Abstract. This is Part II of the paper "Witt vectors and a question of Keating and Rudnick" [Ka-WVQKR]. Here we prove an independence result for tuples of character sums, formed with a variable character and its powers. In the Appendix, we prove an independence result for tuples of character sums formed with variable pairs of characters and products of the two.

1. Introduction

We work over a finite field $k = \mathbb{F}_q$ inside a fixed algebraic closure $\overline{k}$, and fix an integer $n \geq 2$. We form the k-algebra

$$B := k[X]/(X^{n+1}).$$

Following Keating and Rudnick [K-R], we say that a character

$$\Lambda : B^\times \to \mathbb{C}^\times$$

is "even" if it is trivial on the subgroup $k^\times$. The quotient $B^\times/k^\times$ is the group $(1 + Xk[[X]])/(1 + X^{n+1}k[[X]])$ of truncated "big" Witt vectors, cf. [Ka-WVQKR]. We systematically view even characters as characters of this quotient group.

We say that a character Λ is "primitive" if it is nontrivial on the subgroup $1 + kX^n$ of $B^\times$. The Swan conductor of a character Λ of $B^\times$ is the largest integer r such that Λ is nontrivial on the subgroup $1+(X^r)$. Thus a character is primitive if and only if its Swan conductor is n.

Given an even character Λ of $B^\times$, i.e. a character of $(1+Xk[[X]])/(1+X^{n+1}k[[X]])$, we can form an L-function on $\mathbb{G}_m/k$ as follows. Given an irreducible monic polynomial $P(t) \in k[t]$ with $P(0) \neq 0$, the irreducible polynomial $P(t)/P(0)$ has constant term 1, so $P(X)/P(0)$ mod X^{n+1} lies in $(1 + Xk[[X]])/(1 + X^{n+1}k[[X]])$. We define

$$\Lambda(P) := \Lambda(P(X)/P(0) \bmod X^{n+1}).$$

We then define

$$L(\mathbb{G}_m/k,\Lambda)(T) := \prod_{irred.\ monic\ P,\ P(0)\neq 0} (1-\Lambda(P)T^{deg(P)})^{-1}.$$

This L-function has a cohomological interpretation:

$$L(\mathbb{G}_m/k,\Lambda)(T) = L(\mathbb{G}_m/k,\mathcal{L}_{\Lambda(1-tX)})(T),$$

cf. [Ka-WVQKR]. This second expression, with coefficient sheaf which is lisse at 0, leads us to consider the "completed" L-function

$$L(\mathbb{A}^1/k,\mathcal{L}_{\Lambda(1-tX)})(T) = L(\mathbb{G}_m/k,\mathcal{L}_{\Lambda(1-tX)})(T)/(1-T) =$$
$$= L(\mathbb{G}_m/k,\Lambda)(T)/(1-T).$$

One knows by Weil [Weil] that so long as Λ is nontrivial, this completed L-function is a polynomial in T of degree $Swan(\Lambda)-1$, which is "pure of weight one". In other words, it is of the form $\prod_{i=1}^{Swan(\Lambda)-1}(1-\beta_i T)$ with each β_i an algebraic integer all of whose complex absolute values are $\sqrt{q}$.

For Λ primitive, we define a conjugacy class $\theta_{k,\Lambda}$ in the unitary group $U(n-1)$ in terms of its reversed characteristic polynomial by the formula

$$\det(1-T\theta_{k,\Lambda}) = L(\mathbb{A}^1/k,\mathcal{L}_{\Lambda(1-tX)})(T/\sqrt{q}).$$

In the earlier paper [Ka-WVQKR], we proved the following result.

Theorem 1.1. *Fix an integer $n \geq 4$. In any sequence of finite fields k_i (of possibly varying characteristics) whose cardinalities q_i are archimedeanly increasing to ∞, the collections of conjugacy classes*

$$\{\theta_{k_i,\Lambda}\}_{\Lambda \text{ primitive even}}$$

become equidistributed in the space $PU(n-1)^{\#}$ of conjugacy classes in the projective unitary group $PU(n-1)$ for its "Haar measure" of total mass one. We have the same result for $n=3$ if we require that no k_i have characteristic 2 or 5.

In this paper, we will prove the following independence result.

Theorem 1.2. *Fix an integer $n \geq 5$. In any sequence of finite fields k_i (of possibly varying characteristics p, none of which is 2 or 3) whose cardinalities q_i tend archimedeanly to ∞, the collections of pairs of conjugacy classes*

$$\{(\theta_{k_i,\Lambda},\theta_{k_i,\Lambda^2})\}_{\Lambda \text{ primitive even}}$$

become equidistributed in the space $PU(n-1)^{\#} \times PU(n-1)^{\#}$ of conjugacy classes in $PU(n-1) \times PU(n-1)$.

We will also prove the following more general version.

Theorem 1.3. *Let $d \geq 2$ be an integer. Fix an integer $n \geq 5$. In any sequence of finite fields k_i (of possibly varying characteristics p, subject only to $p \geq 2d+1$) whose cardinalities q_i tend archimedeanly to ∞, the collections of d-tuples of conjugacy classes*

$$\{(\theta_{k_i,\Lambda}, \theta_{k_i,\Lambda^2}, ..., \theta_{k_i,\Lambda^d})\}_{\Lambda \text{ primitive even}}$$

become equidistributed in the space $(PU(n-1)^{\#})^d$ of conjugacy classes in $(PU(n-1))^d$.

2. The proof

For each integer $r \geq 1$, we have the scheme $W_r/\mathbb{F}_p$ of p-Witt vectors of length r. We fix a a faithful character $\psi_r : W_r(\mathbb{F}_p) = \mathbb{Z}/p^r\mathbb{Z} \cong \mu_{p^r}(\overline{\mathbb{Q}_\ell})$. For example, we might take $x \mapsto \exp(2\pi i x/p^r)$, so that $\psi_{r+1}^p = \psi_r$. Every character of $W_r(k)$ is of the form

$$w \mapsto \psi_r(\mathrm{Trace}_{W_r(k)/W_r(\mathbb{F}_p)}(aw))$$

for a unique element $a \in W_r(k)$. Let us denote this character $\psi_{r,a}$:

$$\psi_{r,a}(w) := \psi_r(\mathrm{Trace}_{W_r(k)/W_r(\mathbb{F}_p)}(aw)).$$

We choose a prime ℓ different from the characteristic p of k, and work with ℓ-adic cohomology. We fix an embedding of $\overline{\mathbb{Q}_\ell}$ into $\mathbb{C}$.

Attached to the character $\psi_{r,a}$ of $W_r(k)$ we have the Artin-Schreier-Witt sheaf $\mathcal{L}_{\psi_{r,a}} = \mathcal{L}_{\psi_r(aw)}$ on W_r. Given an integer $m \geq 1$ prime to p, we have the morphism of k-schemes $\mathbb{A}^1 \to W_r$ given by $t \mapsto (t^m, 0'\text{s})$. The pullback of $\mathcal{L}_{\psi_{r,a}}$ by this morphism is denoted $\mathcal{L}_{\psi_{r,a}(t^m,0'\text{s})} = \mathcal{L}_{\psi_r(a(t^m,0'\text{s}))}$. It is a lisse rank one sheaf on $\mathbb{A}^1$.

For each prime to p integer m in the range $1 \leq m \leq n$, we define

$$\ell(m,n) = 1 + \text{the largest integer } k \text{ such that } mp^k \leq n.$$

We work on the product space

$$S_n := \prod_{m \geq 1 \text{ prime to } p,\ m \leq n} W_{\ell(m,n)}.$$

As an $\mathbb{F}_p$-scheme, S_n is simply a huge affine space, with coordinates the components of all the Witt vector factors. On $\mathbb{A}^1 \times S_n$, we have the lisse rank one sheaf

$$\mathcal{L}_{univ,n} := \otimes_m \mathcal{L}_{\psi_{\ell(m,n)}(a(m)(t^m,0'\text{s}))}.$$

On S_n, we have the sheaf

$$L_{univ,n} := R^1(pr_2)_!(\mathcal{L}_{univ,n})(1/2).$$

This is a "sheaf of perverse origin" on S_n, in the terminology of [Ka-Semi], and its restriction to every lisse subscheme T of S_n is a sheaf of perverse origin on T, cf. [Ka-Semi, Cor. 6].

In the product defining S_n, there is a distinguished factor we need to single out. Write n as n_0p^{r-1} with n_0 prime to p and $r \geq 1$. Then we have a factor W_r, carrying $\mathcal{L}_{\psi_r(a_r(t^{n_0},0's))}$. Inside W_r, we have the open set $W_r^\times$ where the first component is invertible. [It is also the group of invertible elements for the ring structure on W_r.] Inside S_n, we have the open set $Prim_n \subset S_n$ where the W_r component lies in $W_r^\times$.

The sheaf $L_{univ,n}|Prim_n$ is lisse of rank $n-1$, pure of weight zero. The main theorem of [Ka-WVQKR, Thm. 5.1] is that G_{geom} for $L_{univ,n}|Prim_n$ contains $SL(n-1)$, if either $n \geq 4$ or if $n = 3$ and p is not 2 or 5. Its (already unitarized, thanks to the (1/2) Tate twist) Frobenius conjugacy classes at k-valued points of $Prim_n$ are precisely the classes

$$\{\theta_{k,\Lambda}\}_{\Lambda \text{ primitive even}},$$

cf. [Ka-WVQKR, Lemma 4.1]. [These Frobenius conjugacy classes are automatically semisimple, as they "come from curves".]

Lemma 2.1. *The complement $S_n \setminus Prim_n$ is naturally the space S_{n-1}, and under this identification the restriction of $L_{univ,n}$ to S_{n-1} is the sheaf $L_{univ,n-1}$.*

Proof. To see the complement $S_n \setminus Prim_n$ as the space S_{n-1}, we argue as follows. In this complement, all factors except the distinguished one are unchanged. In the distinguished W_r factor, the entry is required to have first component zero; that factor becomes a $W_{r-1} = W_{\ell(n_0,n-1)}$ (and is omitted entirely if $n_0 = n$). In the product with $\mathbb{A}^1$, the restriction of $\mathcal{L}_{univ,n}$ to this S_{n-1} is the sheaf $\mathcal{L}_{univ,n-1}$. The claim then results from base change for $R^1(pr_2)_!$. □

Given two distinct integers $a, b \geq 1$, both prime to the characteristic p we next consider pairs of conjugacy classes

$$\{(\theta_{k,\Lambda^a}, \theta_{k,\Lambda^b})\}_{\Lambda \text{ primitive even}}.$$

For any integer a prime to p, the classes θ_{k,Λ^a} are obtained as follows. On $\mathbb{A}^1 \times S_n$, we have the lisse rank one sheaf $\mathcal{L}^{(a)}_{univ,n}$, defined by the same recipe as for $\mathcal{L}_{univ,n}$, but replacing each chosen character ψ_r of $W_r(\mathbb{F}_p)$ by its a'th power. Then

$$L^{(a)}_{univ,n} := R(pr_2)_!(\mathcal{L}^{(a)}_{univ,n})(1/2)$$

gives rise to the classes θ_{k,Λ^a}.

Theorem 2.2. *Fix $n \geq 5$. Fix distinct integer $a, b \geq 1$ and a characteristic p not dividing $ab(a^2 - b^2)$. Then the group G_{geom} for the direct sum*

$$L^{(a)}_{univ,n} \oplus L^{(b)}_{univ,n}$$

contains the product $SL(n-1) \times SL(n-1)$.

Proof. Because G_{geom} for each of the summands contains $SL(n-1)$, the group G_{geom} for the direct sum has identity component either $SL(n-1)$ or the product $SL(n-1) \times SL(n-1)$. To show that it is the latter, it suffices to show that there is no geometric isomorphism of either $L^{(b)}_{univ,n}$ or its dual $L^{(-b)}_{univ,n}$ with any sheaf of the form $L^{(a)}_{univ,n} \otimes \mathcal{L}$ for any lisse, rank one sheaf $\mathcal{L}$ on $Prim_n$, cf. [Ka-ESDE, 1.8.1 and 1.8.2].

We first treat the case when n is odd, and argue by contradiction. The assumption that p does not divide $ab(a^2 - b^2)$ insures that $p \geq 5$. Inside $Prim_n$ we have an $\mathbb{A}^2$ with coordinates (A_1, A_3) over which the sheaf $\mathcal{L}^{(b)}_{univ}$ on $\mathbb{A}^1 \times \mathbb{A}^2$ (coordinates (t, A_1, A_3)) is the sheaf $\mathcal{L}_{\psi_r^b((t^{n_0}, 0's))} \otimes \mathcal{L}_{\psi_1(bA_1t + bA_3t^3)}$, and over which the sheaf $\mathcal{L}^{(a)}_{univ}$ is $\mathcal{L}_{\psi_r^a((t^{n_0}, 0's))} \otimes \mathcal{L}_{\psi_1(aA_1t + aA_3t^3)}$. [For simplicity, suppose that each ψ_r is chosen so that $\psi_r^{p^{r-1}} = \psi_1$. Use the $m = 1$, $m = 3$, and $m = n_0$ factors. In the $m = n_0$ factor, freeze the $W_r = W_{\ell(n_0,n)}$ component to be $(1, 0's)$. In the $m = 1$ and $m = 3$ factors, take the component to be respectively $(0's, A_1)$ and $(0's, A_3)$. In all other factors, freeze the component to be 0.]

Let us denote by $\mathcal{F}^{(b)}_n$ and by $\mathcal{F}^{(a)}_n$ the restrictions of $L^{(b)}_{univ,n}$ and $L^{(a)}_{univ,n}$ to this $\mathbb{A}^2$:

$$\mathcal{F}^{(b)}_n := L^{(b)}_{univ,n}|\mathbb{A}^2, \quad \mathcal{F}^{(a)}_n := L^{(a)}_{univ}|\mathbb{A}^2.$$

The sheaves $\mathcal{F}^{(b)}_n$ and $\mathcal{F}^{(a)}_n$ are geometrically irreducible (because they are Fourier transforms), pure of weight one, and self dual (because their trace functions are $\mathbb{R}$-valued (use $t \mapsto -t$). In fact, the duality is symplectic, compare [Ka-MMP, 3.10.3] where a result of this type is proved. Therefore G_{geom} for each of $\mathcal{F}^{(b)}_n$ and $\mathcal{F}^{(a)}_n$ is an irreducible subgroup of $Sp(n-1)$. Moreover, a moment calculation based on [Ka-MMP, 3.11.4] shows that for each of $\mathcal{F}^{(b)}_n$ and $\mathcal{F}^{(a)}_n$, the fourth moment is 3. This means precisely that in the decomposition of the tensor square of each as

$$Sym^2 \oplus \Lambda^2/\mathbb{1} \oplus \mathbb{1},$$

each of the three summands is G_{geom}-irreducible. What is key here is that the trivial factor $\mathbb{1}$ is the **only** one-dimensional component of $(\mathcal{F}^{(b)}_n)^{\otimes 2}$ and the only one-dimensional component of $(\mathcal{F}^{(a)}_n)^{\otimes 2}$.

Suppose now that there is a geometric isomorphism of either $L^{(b)}_{univ,n}$ or its dual $L^{(-b)}_{univ,n}$ with a sheaf of the form $L^{(a)}_{univ,n} \otimes \mathcal{L}$ for some lisse, rank one sheaf $\mathcal{L}$ on $Prim_n$, then after restriction to our $\mathbb{A}^2$ we get a geometric isomorphism of $\mathcal{F}^{(b)}_n$ with $\mathcal{F}^{(a)}_n \otimes \mathcal{L}$ for some lisse, rank one sheaf $\mathcal{L}$ on $\mathbb{A}^2$. From this isomorphism, we see that the tensor square

$$(\mathcal{F}^{(a)}_n \otimes \mathcal{L})^{\otimes 2} = (\mathcal{F}^{(a)}_n)^{\otimes 2} \otimes \mathcal{L}^{\otimes 2}$$

has, geometrically, a one dimensional component, which is to say that $(\mathcal{F}^{(a)}_n)^{\otimes 2}$ admits $\mathcal{L}^{\otimes -2}$ as a quotient. But the only one-dimensional quotient of $(\mathcal{F}^{(a)}_n)^{\otimes 2}$ is $\mathbb{1}$.

Hence the lisse rank one $\mathcal{L}$ is geometrically of order dividing 2. But for an affine space in odd characteristic, there are no nontrivial homomorphisms of its geometric π_1 to ± 1; i.e., $H^1(\mathbb{A}^2 \otimes \overline{\mathbb{F}_p}, \mu_2)$ vanishes. [Use the Kummer sequence. In it, $H^0(\mathbb{A}^2 \otimes \overline{\mathbb{F}_p}, \mathbb{G}_m)$ is $\overline{\mathbb{F}_p}^{\times}$, and the H^1, which is $\mathrm{Pic}(W_d \otimes \overline{\mathbb{F}_p})$, vanishes.] So we would have a geometric isomorphism of either $\mathcal{F}^{(b)}_n$ or $\mathcal{F}^{(-b)}_n$ with $\mathcal{F}^{(a)}_n$. To fix ideas, suppose it is the former. As both are geometrically irreducible and self dual, the group $\mathrm{Hom}_{geom}(\mathcal{F}^{(b)}_n, \mathcal{F}^{(a)}_n) = H^0(\mathbb{A}^2 \otimes \overline{\mathbb{F}_p}, \mathcal{F}^{(-b)}_n \otimes \mathcal{F}^{(a)}_n)$ would be nonzero, in fact one-dimensional. By Poincaré duality, the group $H^4_c(\mathbb{A}^2 \otimes \overline{\mathbb{F}_p}, \mathcal{F}^{(b)}_n \otimes \mathcal{F}^{(-a)}_n)$ would be one-dimensional. The coefficient group $\mathcal{F}^{(b)}_n \otimes \mathcal{F}^{(-a)}_n$ is pure of weight zero, so this H^4_c is pure of weight four, and one-dimensional. The lower cohomology groups are of lower weight. So for variable finite extensions $L/\mathbb{F}_p$, with $\#L := q_L$, we would have the estimate

$$|\sum_{(A_1,A_3)\in\mathbb{A}^2(L)} \mathrm{Trace}(Frob_{L,(A_1,A_3)}|\mathcal{F}^{(b)}_n)\mathrm{Trace}(Frob_{L,(A_1,A_3)}|\mathcal{F}^{(-a)}_n)| =$$

$$= q_L^2 + O(q_L^{3/2}).$$

We will show that this sum is precisely q_L, thus arriving at the desired contradiction.

This sum is

$$\sum_{A_1,A_3\in L} \sum_{x\in L} \psi_r^b((x^{n_0}, 0's))\psi_1(bA_3x^3 + bA_1x) \sum_{y\in L} \psi_r^{-a}((y^{n_0}, 0's))\psi_1(-aA_3y^3 - aA_1y) =$$

$$= (1/q_L) \sum_{x,y\in L} \psi_r^b((x^{n_0}, 0's))\psi_r^{-a}((y^{n_0}, 0's)) \sum_{A_1,A_3\in L} \psi_1(A_3(bx^3 - ay^3) + A_1(bx - ay)).$$

The $1/q_L$ factor comes from the (1/2) Tate twists in the definitions of $\mathcal{F}^{(b)}_n$ and $\mathcal{F}^{(-a)}_n$.

The innermost sum

$$\sum_{A_1,A_3\in L} \psi_1(A_3(bx^3-ay^3)+A_1(bx-ay))$$

vanishes unless both $bx^3 = ay^3$ and $bx = ay$, and if both vanish the sum is q_L^2. Now if $bx = ay$ then either $x = y = 0$ or x and y are both nonzero and $x/y = a/b$. If (x, y) is not $(0, 0)$, then from $bx^3 = ay^3$ we get $(x/y)^3 = a/b$. Comparing with $x/y = a/b$, we get $(a/b)^3 = a/b$, so $a^2 = b^2$. Our hypothesis that $a^2 - b^2 \neq 0 \bmod p$ tells us that the innermost sum vanishes unless $x-y = 0$. Thus the entire sum has only the $x = y = 0$ term, and the sum is q_L, as asserted.

This contradiction concludes the proof that for n odd, the group G_{geom} for the direct sum

$$L^{(b)}_{univ,n} \oplus L^{(a)}_{univ,n}$$

contains the product $SL(n-1) \times SL(n-1)$.

Suppose now that $n \geq 5$ is even. Then $n-1$ is odd, and ≥ 5. So we know that G_{geom} for the direct sum

$$L^{(b)}_{univ,n-1} \oplus L^{(a)}_{univ,n-1}$$

contains $SL(n-2) \times SL(n-2)$.

Recall that $L^{(b)}_{univ,n} \oplus L^{(a)}_{univ,n}$ is a sheaf of perverse origin on S_n, lisse on $Prim_n$ and lisse on the dense open set $Prim_{n-1}$ of $S_{n-1} = S_n \setminus Prim_n$. We now apply [Ka-Semi, Cor. 10, (2)]to this situation. The rank of G_{geom} for $L^{(b)}_{univ,n} \oplus L^{(a)}_{univ,n}$ on $Prim_n$ is at least the rank of G_{geom} for $L^{(b)}_{univ,n-1} \oplus L^{(a)}_{univ,n-1}$ on $Prim_{n-1}$. This last rank is $2(n-3)$. If G^0_{geom} for $L^{(b)}_{univ,n} \oplus L^{(a)}_{univ,n}$ on $Prim_n$ were $SL(n-1)$, we would get the inequality

$$n-2 \geq 2(n-3),$$

which is false for $n \geq 5$. Given the dearth of choice for this G^0_{geom}, it must be $SL(n-1) \times SL(n-1)$. □

Corollary 2.3. *Let $d \geq 2$ and $n \geq 5$ be integers. In any characteristic $p \geq 2d+1$, the group G_{geom} for the d-fold direct sum*

$$\bigoplus_{a=1}^{d} L^{(a)}_{univ,n}$$

contains the d-fold product $(SL(n-1))^d$.

Proof. The hypothesis that $p \geq 2d+1$ insures that for any integer a, b with $1 \leq a < b \leq d$, p does not divide $ab(a^2 - b^2)$. The corollary then follows from Theorem 2.2 by Goursat-Kolchin-Ribet, cf. [Ribet, pp. 790-791]. □

Theorem 2.4. *Let $d \geq 2$ be an integer. Fix an integer $n \geq 5$. In any sequence of finite fields k_i (of possibly varying characteristics p, subject only to $p \geq 2d+1$) whose cardinalities q_i tend archimedeanly to ∞, the collections of d-tuples of conjugacy classes*

$$\{(\theta_{k_i,\Lambda}, \theta_{k_i,\Lambda^2}, ..., \theta_{k_i,\Lambda^d})\}_{\Lambda \text{ primitive even}}$$

become equidistributed in the space $(PU(n-1)^{\#})^d$ of conjugacy classes in $(PU(n-1))^d$.

Proof. Let us denote by $\Theta_{k_i,\Lambda}$ the d-tuple

$$\Theta_{k_i,\Lambda} := (\theta_{k_i,\Lambda}, \theta_{k_i,\Lambda^2}, ..., \theta_{k_i,\Lambda^d}).$$

By the Weyl criterion, we must show that for each fixed irreducible nontrivial representation Ξ of $(PU(n-1))^d$, the normalized Weyl sums

$$(1/\#Prim_n(k_i)) \sum_{\Lambda/k_i \text{ primitive even}} \mathrm{Trace}(\Xi(\Theta_{k_i,\Lambda}))$$

tend to 0 as $\#k_i$ grows. In each characteristic $p \geq 2d+1$, this sum is bounded in absolute value by

$$C(p,n,\Xi)/\sqrt{\#k_i}, \quad \text{for } C(p,n,\Xi) := 2\sum_i h^i_c(Prim_n \otimes_{\mathbb{F}_p} \overline{\mathbb{F}_p}, \Xi_d),$$

for Ξ_d the lisse sheaf on $Prim_n/\mathbb{F}_p$ formed by "pushing out" the direct sum $\bigoplus_{a=1}^d L^{(a)}_{univ,n}$ along the representation Ξ, now viewed as a representation of $(GL(n-1))^d$ which factors through the quotient group $(PGL(n-1))^d$.

As was the case in [Ka-WVQKR, 8.1], we do not know uniform bounds for these sums of Betti numbers $C(p,n,\Xi)$ as p varies (n and Ξ fixed). But we can bypass this problem, if we can, by other means, show that for fixed n and $p > 2n-1$, we have a bound of the form $D(n,\Xi)/\sqrt{\#k_i}$ with a constant $D(n,\Xi)$ which is independent of p. Then we use this constant for $p > 2n-1$, and we use the constant $C(p,n,\Xi)$ for the finitely many primes $p \leq 2n-1$. We will show that in fact we can take $D(n,\Xi) := 3\dim(\Xi)/(n-1)$. This will then prove the theorem.

Let us recall the argument from [Ka-WVQKR, 8.2]. Fix a nontrivial additive character ψ (formerly our ψ_1) of $\mathbb{F}_p$. For $p > n$, the space $Prim_n$ is the space of polynomials f of degree n with vanishing constant term, and for each a prime to p, the sheaf $\mathcal{L}^{(a)}_{univ,n}$ on $\mathbb{A}^1 \times Prim_n$ with coordinates (t, f) is $L_{\psi(af(t))}$. The sheaf $L^{(a)}_{univ,n}$ on $Prim_n$ has trace function at k-valued points $f \in Prim_n(k)$ given by

$$\mathrm{Trace}(Frob_{k,f}|L^{(a)}_{univ,n}) = (-1/\sqrt{q_k})\sum_{t\in k}\psi_k(af(t)).$$

The idea is to break up $Prim_n(k)$ into equivalence classes, f and g equivalent if $f - g$ has degree ≤ 1. Thus each equivalence class is a line $\{f(t) + \lambda t\}_{\lambda\in k}$. On each line, $L^{(a)}_{univ,n}$ is the Tate-twisted Fourier Transform $FT_{\psi^a}(\mathcal{L}_{\psi(af(t))})(1/2)$. These are lisse sheaves on $\mathbb{A}^1$ of rank $n - 1$, pure of weight zero and with all ∞-breaks $n/(n - 1)$. Their direct sum

$$\mathbb{V} := \bigoplus_{a=1}^{d} FT_{\psi^a}(\mathcal{L}_{\psi(af(t))})(1/2)$$

is thus lisse of rank d(n-1), with all all ∞-breaks $n/(n - 1)$. Let us denote by $\Xi_d(f)$ the lisse sheaf on $\mathbb{A}^1$ which is the pushout of this direct sum along Ξ. [Equivalently, $\Xi_d(f)$ is the restriction of Ξ_d on $Prim_n$ to the line which is the equivalence class of f.] Its rank is $\dim(\Xi)$, and all its ∞-breaks are at most $n/(n - 1)$ (because this pushout is a direct factor of some tensor power $\mathbb{V}^{\otimes e} \otimes (\mathbb{V}^{\vee})^{\otimes f}$. The sheaf $\Xi_d(f)$ will be irreducible and nontrivial provided that G_{geom} for the direct sum $\mathbb{V}$ contains the d-fold product $(SL(n - 1))^d$.

Lemma 2.5. *For $p > 2n - 1$ and f of degree n with $f(0) = 0$, we have the following results.*

1. *If n is even and the coefficient of t^2 in f is nonzero, then G_{geom} for $\mathbb{V}$ contains the d-fold product $(SL(n - 1))^d$.*
2. *We have the same result for n odd if, in addition, no translate $f(t + c) - f(c)$ of f is an odd polynomial.*

Let us temporarily admit the truth of this lemma. We then compute the "raw" Weyl sum, i.e., $\#Prim_n(k)$ times the normalized Weyl sum, for an irreducible nontrivial Ξ as the sum over equivalence classes of f's in $Prim_n(k)$. Each such sum is

$$\sum_{\lambda\in k}\mathrm{Trace}(Frob_{k,\lambda}|\Xi_d(f)) =$$

$$\mathrm{Trace}(Frob_k|H^2_c(\mathbb{A}^1\otimes\overline{\mathbb{F}_p}, \Xi_d(f)) - \mathrm{Trace}(Frob_k|H^1_c(\mathbb{A}^1\otimes\overline{\mathbb{F}_p}, \Xi_d(f)),$$

the equality by the Lefschetz trace formula.

Suppose first that the hypotheses of either part (1) or of part (2) of the above lemma apply. Then the H^2_c vanishes, and the dimension of the H^1_c is minus the Euler characteristic, itself given by

$$-\chi(\mathbb{A}^1 \otimes \overline{\mathbb{F}_p}, \Xi_d(f)) = Swan_\infty(\Xi_d(f)) - \mathrm{rank}(\Xi_d(f)) \leq$$

$$\leq (n/(n-1))\dim(\Xi) - \dim(\Xi) = \dim(\Xi)/(n-1).$$

So when the above lemma applies, the sum is bounded by

$$\dim(\Xi)\sqrt{q_k}/(n-1).$$

When the lemma does not apply, we look at the first expression for the sum to see that we have the trivial bound

$$\dim(\Xi)q_k.$$

When n is even, of the $(q_k-1)q_k^{n-2}$ equivalence classes of f's, only $1/q_k$ of them fail to satisfy the hypothesis of part (1) of the lemma. When n is odd, at most $(q_k-1)q_k^{(n-1)/2} + (q_k-1)q^{n-3}$ classes fail to satisfy the hypothesis of part (2) of the lemma (the translates of odd f's, and those with no t^2 term).

Thus when n is even, the raw Weyl sum for an irreducible nontrivial Ξ is bounded in absolute value by

$$(\#Prim_n(k)/q_k)(1-1/q_k)\dim(\Xi)\sqrt{q_k}/(n-1)+(\#Prim_n(k)/q_k^2)q_k\dim(\Xi) \leq$$

$$\leq (\#Prim_n(k))2\dim(\Xi)/((n-1)\sqrt{q_k})$$

as soon as $\sqrt{q_k} \geq n-1$.

When n is odd, the fraction of failures to satisfy the hypotheses of part (2) is at most $1/q_k + 1/q_k^{(n-3)/2}$, so always at most $2/q_k$. In this case, the raw Weyl sum will have the same bound, as soon as $\sqrt{q_k} \geq 2(n-1)$. □

It remains to prove the lemma. Under the hypothesis that either n is even or that no translate $f(t+c)-f(c)$ of f is an odd polynomial, it was proven in [Ka-MG, Thm. 19, stated there for $p \geq 2n+1$, but $p > 2n-1$ is enough], that G_{geom} for each $FT_{\psi^a}(\mathcal{L}_{\psi(af(t))})$ contains $SL(n-1)$. By Goursat-Kolchin-Ribet [Ka-ESDE, 1.8.1 and 1.8.2] and [Ribet, pp. 790-791], it suffices to treat the case of each pair of factors.

Thus we must show that if, in addition, the coefficient of t^2 in f is nonzero, then for p not dividing $ab(a^2-b^2)$, there is no geometric isomorphism of either $FT_{\psi^b}(\mathcal{L}_{\psi(bf(t))})$ or its dual $FT_{\psi^b}(\mathcal{L}_{\psi(-bf(-t))})$ with $FT_{\psi^a}(\mathcal{L}_{\psi(af(t))}) \otimes \mathcal{L}$ for any lisse rank one $\mathcal{L}$ on $\mathbb{A}^1 \otimes \overline{\mathbb{F}_p}$.

Again we argue by contradiction. Both Fourier transforms have all ∞ breaks $n/(n-1) < 2$, so $\mathcal{L}$, whose ∞-break is an integer (Hasse-Arf), must have ∞-break either 0 or 1. Thus $\mathcal{L}$ is geometrically $\mathcal{L}_{\psi(c\lambda)}$ for some $c \in \overline{\mathbb{F}_p}$. If we write $FT_{\psi^a}(\mathcal{L}_{\psi(af(t))})$ as $FT_\psi(\mathcal{L}_{\psi(af(t/a))})$, and write

$FT_{\psi^b}(\mathcal{L}_{\psi(bf(t))})$ as $FT_\psi(\mathcal{L}_{\psi(bf(t/b))})$, we find that either $FT_\psi(\mathcal{L}_{\psi(bf(t/b))})$ or its dual $FT_\psi(\mathcal{L}_{\psi(-bf(-t/b))})$ is geometrically isomorphic to

$$FT_{\psi^a}(\mathcal{L}_{\psi(af(t))}) \otimes \mathcal{L}_{\psi(c\lambda)} \cong FT_\psi(\mathcal{L}_{\psi(af(t/a)+c)})).$$

By Fourier inversion, we get that either $\mathcal{L}_{\psi(bf(t/b))}$ or $\mathcal{L}_{\psi(-bf(-t/b))}$ is geometrically isomorphic to $\mathcal{L}_{\psi(af(t/a)+b)})$. This in turn implies that either $bf(t/b)$ or $-bf(-t/b)$ is equal to $af(t/a)$. Comparing coefficients of t^2, we see that either $1/b$ or $-1/b$ is $1/a$. This is excluded by the hypothesis that p does not divide $ab(a^2-b^2)$. This concludes the proof of Lemma 2.4, and, with it, the proof of Theorem 2.3.

3. Appendix: Pairs of characters

Let us say that an ordered pair of primitive even characters (χ, Λ) is a primitive pair if the product character $\chi\Lambda$ is primitive. We denote by $\mathrm{PrimPair}(k)$ the set of primitive pairs.

Theorem 3.1. *Fix an integer $n \geq 5$. In any sequence of finite fields k_i (of possibly varying characteristics p, subject only to $p \geq 7$) whose cardinalities q_i tend archimedeanly to ∞, the collections of triples of conjugacy classes*

$$\{(\theta_{k_i,\chi}, \theta_{k_i,\Lambda}, \theta_{k_i,\chi\Lambda})\}_{(\chi,\Lambda)\in \mathrm{PrimPair}_(k_i)}$$

become equidistributed in the space $(PU(n-1)^\#)^3$ of conjugacy classes in $(PU(n-1))^3$.

More generally, given an integer $d \geq 1$, let us say that a pair of primitive even characters (χ, Λ) is a d-fold primitive pair if the $(d+1)^2-1$ characters $\chi^a\Lambda^b$, with $0 \leq a, b \leq d$ and $(a,b) \neq (0,0)$, are each primitive. [Each is necessarily even.] We denote by $\mathrm{PrimPair}_d(k)$ the set of d-fold primitive pairs. Fix an ordering on the index set

$$I_d := [0,d] \times [0,d] \setminus (0,0),$$

e.g., order first by the sum $a+b$, and within pairs of given sum order by b.

Theorem 3.2. *Fix an integer $d \geq 1$ and an integer $n \geq 5$. In any sequence of finite fields k_i (of possibly varying characteristics p, subject only to $p \geq 2(d+1)^2-1$) whose cardinalities q_i tend archimedeanly to ∞, the collections of $(d+1)^2-1$-tuples of conjugacy classes*

$$\{(\theta_{k_i,\chi^a\Lambda^b})_{(a,b)\in I_d}\}_{(\chi,\Lambda)\in \mathrm{PrimPair}_d(k_i)}$$

become equidistributed in the space $(PU(n-1)^\#)^{(d+1)^2-1}$ of conjugacy classes in $(PU(n-1))^{(d+1)^2-1}$.

Proof. We work first in a given characteristic p. The scheme $S_n/\mathbb{F}_p$, a product of Witt groups, is a commutative group scheme, with componentwise operations. For $k/\mathbb{F}_p$ a finite extension, the points of $S_n(k)$ are precisely the even characters of $(k[X]/(X^{n+1}))^\times$. The points of $Prim_n(k)$ are precisely the primitive even characters. For an integer a prime to p, the a-fold addition map $[a] : w \mapsto aw$ is an automorphism of each Witt group factor, so an automorphism of S_n which preserves $Prim_n$ and induces an automorphism of $Prim_n$. The sheaves $L^{(a)}_{univ,n}$ on $Prim_n$, for a prime to p, are simply the pullbacks of $L_{univ,n}$ by the map $[a] : Prim_n \to Prim_n$:

$$L^{(a)}_{univ,n}|Prim_n = [a]^\star(L_{univ,n}|Prim_n).$$

For each pair of integers (a, b), we have the weighted addition maps

$$[a, b] : S_n \times S_n \to S_n, \quad , (s, t) \mapsto as + bt.$$

Lemma 3.3. *Suppose $p > 2d$. Define $U_d := \cap_{(a,b)\in I_d}[a, b]^{-1}(Prim_n)$. Then U_d is a dense open set of $Prim_n \times Prim_n$. For each finite extension $k/\mathbb{F}_p$, the k-valued points $U_d(k)$ are precisely the points* $\mathrm{PrimPair}_d(k)$, *i.e. the d-fold primitive pairs.*

Proof. Since $Prim_n$ is open in S_n, each inverse image $[a, b]^{-1}(Prim_n)$ is open in $S_n \times S_n$. Already intersecting inverse images by the two projections $[1, 0]$ and $[0, 1]$ shows that U_d lies in $Prim_n \times Prim_n$. To see that U_d is nonempty, observe that it contains the diagonal of $Prim_n \times Prim_n$; this is merely the statement that if $s \in Prim_n$, then $(a + b)s$ lies in $Prim_n$ for any (a, b) in I_d, simply because $p > 2d \geq a + b$. That $U_d(k) = \mathrm{PrimPair}_d(k)$ is a tautology. □

On the space U_d, we have for each (a, b) in I_d the sheaf

$$\mathcal{F}^{(a,b)} := [a, b]^\star(L_{univ,n}|Prim_n).$$

Theorem 3.4. *Suppose $n \geq 5$, $d \geq 1$, and $p \geq 2(d + 1)^2 - 1$. Then the group G_{geom} for the direct sum sheaf*

$$\bigoplus_{(a,b)\in I_d} \mathcal{F}^{(a,b)}$$

on U_d contains the product $(SL(n - 1))^{(d+1)^2-1}$.

Proof. After pullback, G_{geom} can only get smaller. So it suffices to exhibit a pullback on which G_{geom} contains $(SL(n - 1))^{(d+1)^2-1}$. For this, we use the morphism

$$Prim_n \to U_d, \chi \mapsto (\chi, \chi^{d+1}).$$

The pullback of $\mathcal{F}^{(a,b)}$ by this morphism is the sheaf $L^{a+b(d+1)}_{univ,n}|Prim_n$. So the entire direct sum pulls back to

$$\bigoplus_{c=1}^{(d+1)^2-1} L^{(c)}_{univ,n}$$

on $Prim_n$. By Corollary 2.3, its G_{geom} contains $(SL(n-1))^{(d+1)^2-1}$. □

To conclude the proof of Theorem 3.2, we proceed exactly as the deduction of Theorem 2.4 from Corollary 2.3. For $(\chi,\Lambda) \in \mathrm{PrimPair}_d(k)$, we denote by $\Theta_{k,\chi,\Lambda}$ the $(d+1)^2-1$-tuple of conjugacy classes

$$\Theta_{k,\chi,\Lambda} := (\theta_{k_i,\chi^a\Lambda^b})_{(a,b)\in I_d}.$$

Fix an irreducible nontrivial representation Ξ of $(PU(n-1))^{(d+1)^2-1}$. Denote by

$$\Xi_{I_d}$$

the lisse sheaf on U_d obtained from pushing out lisse sheaf $\bigoplus_{(a,b)\in I_d}\mathcal{F}^{(a,b)}$ by Ξ, now viewed as a representation of $(GL(n-1))^{(d+1)^2-1}$ which factors through $(PGL(n-1))^{(d+1)^2-1}$.

For $p \geq 2(d+1)^2-1$ but $p \leq 2n-1$, and $k/\mathbb{F}_p$ a finite extension with $q := \#k$, we have the estimate

$$|\sum_{(\chi,\Lambda)\in \mathrm{PrimPair}_d(k)=U_d(k)} \mathrm{Trace}(\Xi(\Theta_{k,\chi,\Lambda}))| \leq (\sum_i h^i_c(U_d \otimes_{\mathbb{F}_p} \overline{\mathbb{F}_p}, \Xi_{I_d}))q^{2n-1/2}.$$

In each of these finitely many characteristics, for q large enough we will have $\#U_d(k) \geq q^{2n}/2$ (simply by Lang-Weil, as $U_d/\mathbb{F}_p$ is smooth and geometrically connected of dimension $2n$), and for such q we then have

$$|(1/\#U_d(k)) \sum_{(\chi,\Lambda)\in \mathrm{PrimPair}_d(k)=U_d(k)} \mathrm{Trace}(\Xi(\Theta_{k,\chi,\Lambda}))| \leq \frac{2(\sum_i h^i_c(U_d \otimes_{\mathbb{F}_p} \overline{\mathbb{F}_p}, \Xi_{I_d})}{\sqrt{q}}.$$

It remains to treat the case of characteristic $p > 2n-1$ (and $p \geq 2(d+1)^2-1$). Here we will need the following lemma.

Lemma 3.5. *Suppose that (a,b) and (c,d) are two different elements of I_d. Then for $p \geq 2d+1$, the two linear forms*

$$(aX+bY)/(a+b)^2 \pm (cX+dY)/(c+d)^2$$

are both nonzero in $\mathbb{F}_p[X,Y]$.

Proof. If the two vectors (a,b) and (c,d) are linearly independent, then the two linear forms $(aX+bY)/(a+b)^2$ and $(cX+dY)/(c+d)^2$ are linearly independent, and the assertion is obvious. If they are $\mathbb{F}_p$-linearly dependent, say $(a,b) = t(c,d)$ for some $t \in \mathbb{F}_p^\times$, then

$$t(aX+bY)/(a+b)^2 = (cX+dY)/(c+d)^2.$$

In this case,

$$(aX+bY)/(a+b)^2 \pm (cX+dY)/(c+d)^2 = (1\pm t)(cX+dY)/(c+d)^2,$$

which is nonzero so long as $t \neq \pm 1$. We cannot have $t = -1$, for then $a+c = b+d = 0$, which is impossible because $p \geq 2d+1$, while both $a+c, b+d$ lie in $[0, 2d]$ and are not both zero. We cannot have $t = 1$ by the assumption that (a,b) and (c,d) are two different elements of I_d. □

For $p > n$, the space $U_d/\mathbb{F}_p$ is the space of pairs of polynomials $(f = \sum_{i=1}^n A_i x^i, g = \sum_{i=1}^n B_i x^i)$, such that for all (a,b) in I_d, the polynomial $af + bg$ has its leading coefficient $aA_n + bB_n$ invertible.

Suppose now that $p \geq 2(d+1)^2 - 1$ and $p \geq 2n-1$. For unordered each pair of distinct points (a,b) and (c,d) in I_d, we have the hyperplanes $H_{\pm,(a,b),(c,d)}$ in U_d consisting of those points $(f = \sum_{i=1}^n A_i x^i, g = \sum_{i=1}^n B_i x^i)$ for which the vector (A_2, B_2) of their x^2-coefficients satisfies

$$(aA_2 + bB_2)/(a+b)^2 = \pm(cA_2 + dB_2)/(c+d)^2.$$

We denote by $H \subset U_d$ the union of these $((d+1)^2-1)((d+1)^2-2)$ hyperplanes.

The idea is to break up $U_d(k)$ into equivalence classes, with $(f(x), g(x)) \cong (f(x)+tx, g(x)+tx)$. The pullback of

$$\bigoplus_{(a,b)\in I_d} \mathcal{F}^{(a,b)}$$

to this t-line is, geometrically, the direct sum

$$\bigoplus_{(a,b)\in I_d} FT_{\psi^{a+b}}(\mathcal{L}_{\psi(af(x)+bg(x))}) =$$

$$\bigoplus_{(a,b)\in I_d} FT_{\psi}(\mathcal{L}_{\psi(af(\frac{x}{a+b})+bg(\frac{x}{a+b}))}).$$

Exactly as in Lemma 2.5, we have the following lemma, with essentially the same proof.

Lemma 3.6. *Suppose $p \geq 2(d+1)^2 - 1$ and $p \geq 2n-1$. Suppose $(f,g) \in U_d(k) \setminus H(k)$. Then we have the following results.*

(1) *If n is even, then G_{geom} for*

$$\bigoplus_{(a,b)\in I_d} FT_\psi(\mathcal{L}_{\psi(af(\frac{x}{a+b})+bg(\frac{x}{a+b}))})$$

contains $(SL(n-1))^{(d+1)^2-1}$.

(2) *We have the same result for n odd if in addition no translate $F(X+c)-F(c)$ is odd, for any of the polynomials $af(X)+bg(X)$, $(a,b)\in I_d$.*

Exactly as in the deduction of Theorem 2.4 from Lemma 2.5, we find the following: for $n \geq 5$, in any characteristic $p > 2n-1$ (and $p \geq 2(d+1)^2-1$), then for q large enough that $\sqrt{q} > n(d+1)^4$, we will have the bound

$$|(1/\#U_d(k)) \sum_{(\chi,\Lambda)\in \mathrm{PrimPair}_d(k)=U_d(k)} \mathrm{Trace}(\Xi(\Theta_{k,\chi,\Lambda}))| \leq \frac{2\dim(\Xi)}{(n-1)\sqrt{q}}.$$

This concludes the proof of Theorem 3.2.

□

References

[De-Weil II] Deligne, P., La conjecture de Weil II. Publ. Math. IHES 52 (1981), 313-428.

[Ka-ESDE] Katz, N., Exponential sums and differential equations. Annals of Mathematics Studies, 124. Princeton Univ. Press, Princeton, NJ, 1990. xii+430 pp.

[Ka-MG] Katz, N., On the monodromy groups attached to certain families of exponential sums. Duke Math. J. 54 (1987), no. 1, 41-56.

[Ka-MMP] Katz, N., Moments, monodromy, and perversity: a Diophantine perspective. Annals of Mathematics Studies, 159. Princeton University Press, Princeton, NJ, 2005. viii+475 pp.

[Ka-Semi] Katz, N., A semicontinuity result for monodromy under degeneration. Forum Math. 15 (2003), no. 2, 191-200.

[Ka-WVQKR] Katz, N., Witt vectors and a question of Keating and Rudnick. Int. Math. Res. Not. IMRN 2013, no. 16, 3613-3638.

[K-R] Keating, J.P., and Rudnick, Z., The variance of the number of prime polynomials in short intervals and in residue classes. Int. Math. Res. Not. IMRN 2014, no. 1, 259-288.

[Ribet] Ribet, K., Galois action on division points of Abelian varieties with real multiplications. Amer. J. Math. 98 (1976), no. 3, 751-804.

[Weil] Weil, A., Variétés abéliennes et courbes algébriques. Actualités Sci. Ind., no. 1064 = Publ. Inst. Math. Univ. Strasbourg 8 (1946). Hermann & Cie., Paris, 1948. 165 pp.

Princeton University, Mathematics, Fine Hall, NJ 08544-1000, USA
E-mail address: nmk@math.princeton.edu