Introduction

It is now some thirty years since Deligne first proved his general equidistribution theorem [De-Weil II, Ka-GKM, Ka-Sar-RMFEM], thus establishing the fundamental result governing the statistical properties of suitably "pure" algebro-geometric families of character sums over finite fields (and of their associated L-functions). Roughly speaking, Deligne showed that any such family obeys a "generalized Sato-Tate law", and that figuring out which generalized Sato-Tate law applies to a given family amounts essentially to computing a certain complex semisimple (not necessarily connected) algebraic group, the "geometric monodromy group" attached to that family.

In our earlier books [Ka-GKM], [Ka-ESDE], and [Ka-TLFM], computations of geometric monodromy groups were carried out either directly on an open curve as parameter space, or by restriction to a well-chosen open curve in the parameter space. Once on an open curve, our main tool was to compute, when possible, the local monodromy at each of the missing points. This local monodromy information told us that our sought-after semisimple group contained specific sorts of elements, or specific sorts of subgroups. We typically also had a modicum of global information, e.g., we might have known that the sought-after group was an irreducible subgroup of GL(N), or of the orthogonal group O(N), or of the symplectic group Sp(N). It was often then possible either to decide either exactly which group we had, or to show that our group was on a very short list of possibilities, and then to distinguish among those possibilities by some ad hoc argument.

In this book, we introduce new techniques, which are resolutely global in nature. They are sufficiently powerful that we can sometimes prove that a geometric monodromy group is, say, the symplectic group Sp(N), without knowing the value of N; cf. Theorem 3.1.2 for an instance of this. The price we pay is that these new techniques apply only to families which depend on very many parameters, and thus our work here is nearly disjoint from our earlier "local monodromy" methods of analyzing one-parameter families. However, it is not entirely disjoint, because the new techniques will often leave us knowing, say, that our group is either SO(N) or O(N), but not knowing which. In such cases, we sometimes prove that the group is in fact O(N) by restricting to a suitable curve in the parameter space and then proving that the local monodromy at a particular missing point of this curve is a reflection: since SO(N) contains no reflections, we must have O(N).

Our work is based on two vital ingredients, neither of which yet existed at the time of Deligne's original work on equidistribution. The first of these ingredients is the theory of perverse sheaves, pioneered by Goresky and MacPherson in the topological setting, and then brilliantly transposed to algebraic geometry by Beilinson, Bernstein, Deligne, and Gabber. The second is Larsen's Alternative, discovered by Larsen ten odd years ago, which very nearly characterizes classical groups by their fourth moments.

This book has two goals, one "applied" and one "theoretical". The applied goal is to calculate the geometric monodromy groups attached to some quite specific universal families of (L-functions attached to) character sums over finite fields. The theoretical goal is to develop general techniques, based on combining a diophantine analysis of perverse sheaves and their higher moments with Larsen's Alternative and other group-theoretic results, which can be used to achieve the applied goal, and which are of interest in their own right.

Let us begin by describing some of the universal families we have in mind. Grosso modo, they are of three sorts:

families of additive character sums,

families of multiplicative character sums, and

Weierstrass (and other) families of L-functions of elliptic curves over function fields in one variable.

In the additive character case, we fix a finite field k and a nontrivial $\mathbb{C}$-valued additive character ψ of k. For any finite extension E/k, we denote by ψ_E the additive character of E defined by

$$\psi_E(x) := \psi(\mathrm{Trace}_{E/k}(x)).$$

Fix a pair of integers $n \geq 1$ and $e \geq 3$. We denote by $\mathcal{P}(n,e)(E)$ the space of polynomials over E in n variables of degree $\leq e$. We are concerned with the families of sums, parameterized by f in $\mathcal{P}(n,e)(E)$, given by

$$\mathrm{Sum}(E, f, \psi) := \Sigma_{x_1, ..., x_n \text{ in } E}\ \psi_E(f(x_1, ..., x_n)).$$

It turns out that these sums are, up to sign, the local traces of a perverse sheaf, say $M(n, e, \psi)$, on $\mathcal{P}(n, e)/k$. On some dense open set, say $U(n, e, \psi)$ of $\mathcal{P}(n, e)/k$, this perverse sheaf is a [shift and a Tate twist of a] single lisse sheaf, say $\mathcal{M}(n, e, \psi)$, which is pure of weight zero. [When the degree e is prime to char(k), we can take $U(n, e, \psi)$ to be the open set $\mathcal{D}(n, e)$ consisting of "Deligne polynomials" of degree e in n variables, those whose leading forms define smooth, degree e hypersurfaces in $\mathbb{P}^{n-1}$.] It is the geometric monodromy of this lisse sheaf $\mathcal{M}(n, e, \psi)$ on $U(n, e, \psi)$ which we wish to calculate.

In the multiplicative character case, we fix a finite field k and a nontrivial $\mathbb{C}$-valued multiplicative character χ of $k^\times$, extended to all of k by $\chi(0) := 0$. For any finite extension E/k, we denote by χ_E the multiplicative character of $E^\times$ defined by

$$\chi_E(x) := \chi(\mathrm{Norm}_{E/k}(x)),$$

again extended to all of E by $\chi_E(0) := 0$. Fix a pair of integers $n \geq 1$ and $e \geq 3$. We are concerned with the families of sums, parameterized by f in $\mathcal{P}(n,e)(E)$, given by

$\mathrm{Sum}(E, f, \chi) := \Sigma_{x_1, \ldots, x_n \text{ in } E}\ \chi_E(f(x_1, \ldots, x_n))$.

It turns out that these sums are, up to sign, the local traces of a perverse sheaf, say $M(n, e, \chi)$, on $\mathcal{P}(n, e)/k$. On some dense open set, say $U(n, e, \chi)$ of $\mathcal{P}(n, e)/k$, this perverse sheaf is a [shift and a Tate twist of a] single lisse sheaf, say $\mathcal{M}(n, e, \chi)$. [When the degree e is prime to char(k), we can take $U(n, e, \chi)$ to be the open set $\mathcal{SD}(n, e)$ consisting of "strong Deligne polynomials" of degree e in n variables, those which define smooth hypersurfaces in $\mathbb{A}^n$ and whose leading forms define smooth, degree e hypersurfaces in $\mathbb{P}^{n-1}$.] When χ^e is nontrivial, $\mathcal{M}(n, e, \chi)$ is pure of weight zero, and it is the geometric monodromy of $\mathcal{M}(n, e, \chi)$ we wish to calculate. When χ^e is trivial, $\mathcal{M}(n, e, \chi)$ is mixed of weight ≤ 0, and it is the weight zero quotient of $\mathcal{M}(n, e, \chi)$ whose geometric monodromy we wish to calculate.

In the simplest instance of Weierstrass families of L-functions of elliptic curves, we fix a finite field k of characteristic $p \geq 5$. We denote by χ_2 the quadratic character of $k^\times$. Fix a pair of integers $d_2 \geq 3$ and $d_3 \geq 3$. For each finite extension E/k, we have the product space $(\mathcal{P}(1,d_2)\times\mathcal{P}(1,d_3))(E)$ of pairs $(g_2(t), g_3(t))$ of one-variable polynomials over E of degrees at most d_2 and d_3 respectively. We are concerned with the families of sums, parameterized by (g_2, g_3) in $(\mathcal{P}(1,d_2)\times\mathcal{P}(1,d_3))(E)$,

$\mathrm{Sum}(E, g_2, g_3) := \Sigma_{x,\, t \text{ in } k}\ \chi_{2,E}(4x^3 - g_2(t)x - g_3(t))$.

It turns out that these sums are, up to sign, the local traces of a perverse sheaf, say $W(d_2, d_3)$, on $\mathcal{P}(1,d_2)\times\mathcal{P}(1,d_3)/k$. On the dense open set $U(d_2, d_3)$, of $\mathcal{P}(1,d_2)\times\mathcal{P}(1,d_3)/k$, defined by the condition that $(g_2)^3 - 27(g_3)^2$ has $\mathrm{Max}(3d_2, 2d_3)$ distinct zeroes in $\overline{k}$, this perverse sheaf is a [shift and a Tate twist of a] single lisse sheaf, say $\mathcal{W}(d_2, d_3)$, which is mixed of weight ≤ 0. The weight zero quotient $\mathrm{Gr}^0\mathcal{W}(d_2, d_3)$ of $\mathcal{W}(d_2, d_3)$ is related to L-functions of elliptic curves over function fields as follows. For (g_2, g_3) in $U(d_2, d_3)(E)$, the Weierstrass equation

$$y^2 = 4x^3 - g_2(t)x - g_3(t)$$

defines an elliptic curve over the rational function field E(t), and its (unitarized) L-function is the local L-function of $\mathrm{Gr}^0\mathcal{W}(d_2, d_3)$ at the point (g_2, g_3) in $U(d_2, d_3)(E)$. It is the geometric monodromy of $\mathrm{Gr}^0\mathcal{W}(d_2, d_3)$ we wish to calculate.

This concludes our quick overview of the sorts of universal families we wish to treat. These families have in common some essential features.

The first feature is that, in each case, the parameter space is itself a large linear space of $\mathbb{A}^m$-valued functions (m = 1 in the first two sorts of families, m = 2 in the third sort) on some fixed variety V, i.e., in each case our parameter space is a large linear subspace $\mathcal{F}$of the space $\mathrm{Hom}_{k\text{-scheme}}(V, \mathbb{A}^m)$. [It happens that V is itself an affine space in the examples we have given above ($\mathbb{A}^n$ in the first two sorts, $\mathbb{A}^1$ in the third sort), but this turns out to be a red herring.]

The second feature is that our family of sums has the following structure: for each E/k, we are given a function

$$K(E, \ ): \mathbb{A}^m(E) \to \mathbb{C},$$
$$x \mapsto K(E, x),$$

on the E-valued points of the target $\mathbb{A}^m$, and our family of sums is

$$f \text{ in } \mathcal{F}(E) \subset \mathrm{Hom}_{E\text{-scheme}}(V, \mathbb{A}^m) \mapsto \Sigma_{v \text{ in } V(E)} K(E, f(v)).$$

In the additive character case, we have m=1, and x ↦ K(E, x) is the function $x \mapsto \psi_E(x)$ on $\mathbb{A}^1(E) = E$. In the multiplicative character case, we have m=1, and x ↦ K(E, x) is the function $x \mapsto \chi_E(x)$ on $\mathbb{A}^1(E) = E$. In the case of L-functions of elliptic curves over function fields, we have m=2, and the function (a, b) ↦ K(E, a, b) on $\mathbb{A}^2(E) = E\times E$ is the function

$$(a, b) \mapsto \Sigma_{x \text{ in } E}\ \chi_{2,E}(4x^3 -ax -b).$$

[In these cases, the function x ↦ K(E, x) on $\mathbb{A}^m(E)$ also satisfies in addition the "integral zero" condition

$$\Sigma_{x \text{ in } \mathbb{A}^m(E)} K(E, x) = 0,$$

as the reader will easily check. This turns out to be an important condition, but one that can be somewhat relaxed.]

The third feature is that, in each case, the collection of functions

$$K(E, \ ): \mathbb{A}^m(E) \to \mathbb{C},$$
$$x \mapsto K(E, x),$$

is, up to sign, the trace function of a perverse sheaf K on $\mathbb{A}^m$.

Although not apparent from these examples, there is also interest in introducing, in addition to our perverse sheaf K on $\mathbb{A}^m$, a perverse sheaf L on the source variety V, with trace function

$$L(E, \ ): V(E) \to \mathbb{C},$$
$$v \mapsto L(E, v),$$

and considering the family of sums

$$f \text{ in } \mathcal{F}(E) \subset \mathrm{Hom}_{E\text{-scheme}}(V, \mathbb{A}^m) \mapsto \Sigma_{v \text{ in } V(E)} K(E, f(v))L(E, v).$$

Slightly more generally, one might fix a single function

$$h \text{ in } \mathrm{Hom}_{k\text{-scheme}}(V, \mathbb{A}^m),$$

and consider the family of sums "with an offset of h", namely

$$f \text{ in } \mathcal{F}(E) \subset \mathrm{Hom}_{E\text{-scheme}}(V, \mathbb{A}^m)$$
$$\mapsto \Sigma_{v \text{ in } V(E)} K(E, h(v) + f(v))L(E,v).$$

Let us now turn to a brief description of the "theoretical" aspects of this book, which are mainly concentrated in the first two chapters.

In the first chapter, we show that, under very mild hypotheses on K and L, these sums are, for any fixed h, the trace function (up to sign) of a perverse sheaf Twist(L,K,$\mathcal{F}$,h) on the function space $\mathcal{F}$. This general construction is responsible for the perverse sheaves M(n, e, ψ) and M(n, e, χ) on the space $\mathcal{P}$(n, e)/k discussed in the additive and multiplicative character cases, and it is responsible for the perverse sheaf W(d_2, d_3), on $\mathcal{P}(1,d_2)\times\mathcal{P}(1,d_3)/k$ discussed in the Weierstrass family case. We then formulate in diophantine terms a general orthogonality theorem for pure perverse sheaves, which is formally analogous to the orthogonality theorem for the characters of finite-dimensional representations of a compact Lie group. Proceeding along the same lines, we formulate in diophantine terms the theory of the Frobenius-Schur indicator for geometrically irreducible pure lisse sheaves. This theory is formally analogous to that of the Frobenius-Schur indicator for irreducible representations of a compact Lie group, which tells us whether a given irreducible representation is self dual or not, and tells us, in the autodual case, whether the autoduality is symplectic or orthogonal We then show that, given these diophantine invariants for suitable input perverse sheaves K on $\mathbb{A}^m$ and L on V, there is a simple rule for calculating them for (a suitable quotient of) the perverse sheaf Twist(L,K,$\mathcal{F}$,h) on the function space $\mathcal{F}$.

Up to this point in our theoretical analysis, we require relatively little of our space of functions $\mathcal{F}$, only that it contain the constant functions and that it separate points. We then formulate the notion of "higher moments" for pure perverse sheaves. [The orthogonality theorem is concerned with the "second moment".] To get results on the higher moments, we must require that the function space $\mathcal{F}$ be suitably large, more precisely, that it be "d-separating" for some d ≥ 4. Here d-separating means that given any field extension E/k, and any d distinct points v_1, ..., v_d in V(E), the E-linear map "simultaneous evaluation" at the points v_1, ..., v_d",

$$\mathcal{F}\otimes_k E \to (\mathbb{A}^m(E))^d,$$
$$f \mapsto (f(v_1), f(v_2), ..., f(v_d))$$

is surjective. [In the examples, the degrees ("e" in the first two cases, "d_2" and "d_3" in the Weierstrass case) are taken to be at least 3 in order to insure that our function spaces are at least 4-separating.]

We end the first chapter by proving a quite general "Higher Moment Theorem" . We suppose that the function space $\mathcal{F}$ is d-separating for some d ≥ 4. Then we get control of the even moments M_{2k}, for every positive even integer 2k ≤ d, of (a suitable quotient of) the perverse sheaf Twist(L,K,$\mathcal{F}$,h) on the function space $\mathcal{F}$. An immediate consequence of this control is the fact that the support of (this suitable quotient of) the perverse sheaf Twist(L,K,$\mathcal{F}$,h) is the entire space $\mathcal{F}$. Its restriction to an open dense set of $\mathcal{F}$ is a (shift of a single) lisse sheaf, whose geometric monodromy is what we wish to calculate, and whose higher moments we now control.

In chapter 2, we bring to bear some very important ideas of Michael Larsen, about the determination of classical groups through their higher moments. The idea which we exploit most extensively is "Larsen's Alternative", in which we are given an integer $N \geq 2$ and a reductive subgroup H of one of the classical groups $GL(N, \mathbb{C})$ or $O(N, \mathbb{C})$, or, when N is even and at least 4, $Sp(N, \mathbb{C})$, and we are told that H has the same fourth moment as the ambient group in the given N-dimensional representation (namely 2, 3, 3 in the three successive cases). Larsen's Alternative is the marvelous statement that either H is finite, or that, in the three successive cases, we have

H contains SL(N), in the GL(N) case,
H is either SO(N) or O(N), in the O(N) case,
H is Sp(N), in the Sp(N) case.

This very nearly reduces us to ruling out the possibility that H is finite. [We say very nearly, because we must still compute determinants, i.e., we must still distinguish between SO(N) and O(N), and we must still distinguish among the various groups between SL(N) and GL(N).] Fortunately, there is a great deal known about the possible finite groups which could arise in this context. For $N \geq 3$, any such finite group is, because of its low fourth moment, automatically a primitive subgroup of GL(N). We can then apply the plethora of known results on finite primitive irreducible subgroups of GL(N), due (in chronological order) to Blichfeldt, Mitchell, Huffman-Wales, Zalesskii, and Wales. We can apply all this theory to an H which is the geometric monodromy group attached to (a suitable quotient of the restriction to a dense open set of) the perverse sheaf $\mathrm{Twist}(L,K,\mathcal{F},h)$, thanks to the control over moments gained in the first chapter. For such an H, there are further tools we can bring to bear, both algebro-geometric (the theory of "sheaves of perverse origin") and diophantine in nature. All of this is explained in the second chapter.

A further idea of Michael Larsen is his unpublished "Eighth Moment Conjecture". Suppose $N \geq 8$, and suppose we are given a reductive subgroup H of one of $GL(N, \mathbb{C})$ or $O(N, \mathbb{C})$, or, when N is even, $Sp(N, \mathbb{C})$. Suppose H has the same eighth moment as the ambient group in the given N-dimensional representation. Then Larsen conjectured that, in the successive cases, we have

H contains SL(N), in the GL(N) case,
H is either SO(N) or O(N), in the O(N) case,
H is Sp(N), in the Sp(N) case.

In other words, if we have the correct eighth moment (which implies that the lower even moments are also "correct"), then the "H finite" case of Larsen's Alternative cannot arise. Larsen's Eighth Moment Conjecture has recently been proven by Guralnick and Tiep. Combining their result and the Higher Moment Theorem, we avoid the need to rule out the "H finite" case, provided only that our space of functions $\mathcal{F}$ is at least 8-separating. To see what this means in practice, consider the three examples of universal families we considered above. To have $\mathcal{F}$ 8-separating, we need to take the degree $e \geq 7$ in the cases of additive and multiplicative character sums, and we need to take the degrees d_2 and d_3 both ≥ 7 in the Weierstrass case. [But we still face the earlier mentioned problem of computing the determinant.]

With these tools at hand, we get down to concrete applications. Chapters 3 and 4 are devoted to additive character sums, first on $\mathbb{A}^n$ and then on more general varieties. In chapter 5, we study multiplicative character sums. The results we obtain in these chapters are nearly complete, except that in a number of cases we cannot distinguish whether we have SO(N) or O(N). In chapter 6, we apply the theory of middle additive convolution on the additive group $\mathbb{G}_a = \mathbb{A}^1$ to both additive and multiplicative character sums on $\mathbb{A}^n$. This theory allows us in many cases to compute determinants, and thus distinguish between the O(N) and SO(N) cases. It is in this use of middle additive convolution that we are falling back on the method of restricting to a suitable curve and then computing local monodromies, in order to show that our group contains pseudoreflections of specified determinant. In an appendix to chapter 6, we further develop some technical themes which appeared in the proof of a key technical result, Theorem 6.2.11, which was worked out jointly with Eric Rains.

In chapter 7, we work systematically with "pullback to a curve from $\mathbb{A}^1$" situations. A typical example of the situation we study is this. Take a finite field k of odd characteristic, and consider the rational function field in one variable $k(\lambda)$, over which we have the Legendre curve, defined by the equation

$$y^2 = x(x-1)(x - \lambda).$$

Fix an integer $e \geq 3$. For each finite extension E/k, and each each polynomial $f(\lambda)$ in $E[\lambda]$ of degree at most e (i.e., f lies in $\mathcal{P}(1, e)(E)$), we have the pullback equation

$$y^2 = x(x-1)(x - f(\lambda)).$$

The sums

$$\mathrm{Sum}(f, E) := \Sigma_{x,\lambda \text{ in } E}\ \chi_{2,E}(x(x-1)(x - f(\lambda)))$$

are, up to sign, the trace function of a perverse sheaf on $\mathcal{P}(1, e)$. For f in the dense open set U of $\mathcal{P}(1, e)$ consisting of those polynomials f such that $f(f-1)$ has 2e distinct roots in $\overline{k}$, this perverse sheaf is a (shift and a Tate twist of a) single lisse sheaf, whose rank N is

$$2e - 2, \text{ if } e \text{ odd},$$
$$2e - 3, \text{ if } e \text{ even},$$

and whose local L-function at f in U(E) is precisely the unitarized L-function of the elliptic curve over $E(\lambda)$ defined by the pullback equation

$$y^2 = x(x-1)(x - f(\lambda)).$$

We prove that this lisse sheaf has geometric monodromy group the full orthogonal group O(N), provided that $N \geq 9$. At the very end of this chapter, we give some results on degeneration of Leray spectral sequences, which are certainly well known to the experts, but for which we know of no convenient reference.

In chapter 8, we indicate how the general theory of Twist(L,K,$\mathcal{F}$,h) developed here allows us to recover some of the results of [Ka-TLFM].

Chapters 9, 10, and 11 are devoted to a detailed study of families of L-functions of elliptic curves over function fields in one variable over finite constant fields. Chapter 9 is devoted to explaining how various classical families of elliptic curves provide appropriate input, namely a suitable perverse sheaf K on an $\mathbb{A}^m$, to the general theory. Chapter 10 works out what the general theory gives for various sorts of Weierstrass families, and Chapter 11 works it out for other, more neglected, universal families, which we call FJTwist families.

In chapter 12, we return to theoretical questions, developing some general if ad hoc methods which allow us to work "over $\mathbb{Z}$" instead of "just" over a finite field. These methods apply nicely to the case of multiplicative character sums, and to the various Weierstrass and FJTwist families. What they make possible is equidistribution statements where we are allowed to work over bigger and bigger finite fields, whose characteristics are allowed to vary, e.g., bigger and bigger prime fields, rather than the more restrictive setting of bigger and bigger finite fields of a fixed characteristic. Unfortunately, the methods do not apply at all to additive character sums. Nonetheless, we believe that the corresponding equidistribution statements, about additive character sums over bigger and bigger finite fields whose characteristics are allowed to vary, are in fact true statements. It is just that we are presently incapable of proving them.

In the final chapter 13, we make explicit the application of our results to the arithmetic of elliptic curves over function fields. We first give results on average analytic rank in our families. We then pass to the large-N limit, e.g., by taking Weierstrass families of type (d_2, d_3) as described earlier, and letting $\mathrm{Max}(3d_2, 2d_3)$ tend to infinity, and give results concerning low-lying zeroes as incarnated in the eigenvalue location measures of [Ka-Sar, RMFEM].

It is a pleasure to acknowledge the overwhelming influence on this book of the ideas and work of Deligne, Gabber, and Larsen. In the course of working on the book, I visited the Institute for Advanced Study, the University of Tokyo, the University of Minnesota, the University of Paris at Orsay, I.H.E.S., and the University of Paris VI. I thank all these institutions for their hospitality and support.