

Ax’s theorem with an additive character

Ehud Hrushovski

Abstract. Motivated by Emmanuel Kowalski’s exponential sums over definable sets in finite fields, we generalize Ax’s theorem on pseudo-finite fields to a continuous-logic setting allowing for an additive character. The role played by Weil’s Riemann hypothesis for curves over finite fields is taken by the ‘Weil bound’ on exponential sums. Subsequent model-theoretic developments, including simplicity and the Chatzidakis–Van den Dries–Macintyre definable measures, also generalize.

Analytically, we have the following consequence: consider the algebra of functions $\mathbb{F}_p^n \rightarrow \mathbb{C}$ obtained from the additive characters and the characteristic functions of subvarieties by pre- or post-composing with polynomials, applying min and sup operators to the real part, and averaging over subvarieties. Then any element of this class can be approximated, uniformly in the variables and in the prime p , by a polynomial expression in $\Psi_p(\xi)$ at certain algebraic functions ξ of the variables, where $\Psi(n \bmod p) = \exp(2\pi i n/p)$ is the standard additive character.

1. Introduction

The first-order theory of the class of finite fields was determined in a fundamental paper of James Ax ([1]). He understood that the basic definable sets must be taken to be, not just the points of algebraic varieties, but also their finite projections; in line with the étale theory that Grothendieck and his school were developing in these years. Once this is granted, Ax showed that Weil’s Riemann hypothesis for curves leads to a quantifier-elimination result. Finally, having reduced to quantifier-free formulas, the set of sentences true in every (or in almost every) finite field is determined by the Chebotarev density theorem.

Ax’s student Kieffe studied zeta functions in this connection. A definable measure, corresponding to the leading coefficient of the asymptotic expression for the cardinality of a definable set, was shown to exist and have good properties by Chatzidakis, van den Dries and Macintyre in [9].

Now Weil wrote in 1948 about a “connection between various types of exponential sums, occurring in number-theory, and the so-called Riemann hypothesis in function-fields” [31]. An exponential sum is an expression such as $\sum_{x \in X} \chi(f(x))$, where X is (for our purposes) a subset of F^n , F a finite field, f a polynomial, and χ an additive character, i.e. a homomorphism $(\mathbb{F}_q, +) \rightarrow (\mathbb{T}, \cdot)$, where $\mathbb{T} = \{z : |z| = 1\}$ is the unit circle in \mathbb{C} . On \mathbb{F}_p , the group of characters is generated by the canonical one, $n \bmod p \mapsto e^{2\pi i n/p}$.

From his work on the Riemann hypothesis for curves, Weil was able to derive a bound on exponential sums. When $X = C(F)$ with C an absolutely irreducible curve over $F = \mathbb{F}_q$, f is not constant on C and $\chi \neq 1$, he showed

$$\frac{1}{q} \sum_{c \in C(F)} \chi(c) \leq cq^{-1/2},$$

where the constant c is bounded in algebraic families, i.e. depends only on n and the degrees of f and polynomials cutting out C . This inequality is known as the *Weil bound*. The ideas used by Weil in this connection, greatly amplified, became central in the Grothendieck–Deligne theory, and in equidistribution results; see e.g. [25].

None of this, of course, can be represented within Ax’s theory of pseudo-finite fields, since the additive character χ is not part of the structure. Nevertheless, Emmanuel Kowalski was able to demonstrate that exponential sums and definability in finite fields fit very well together. In [23], he showed how to estimate exponential sums, not only over curves or varieties but also when X above is a general definable set in the language of rings.

The contribution of the present paper is to allow the additive character to be part of the model-theoretic structure. The language thus includes the symbols of the ring language $=, +, \cdot, 0, 1$, along with an additional symbol Ψ whose interpretation is an additive character. This changes the meaning of definability: any function obtained from χ , or from characteristic functions of varieties, by min, sup and polynomial operators taken in any order is now viewed as definable. The measure arising from averaging in finite fields remains definable, and thus one may include the corresponding integration operators as well. A quantifier-elimination theorem is proved, analogous to Ax’s in that the basic functions involve a push-forward from finite covers. As stated in the abstract, this implies *a posteriori* that any of the functions obtained by sup, min or integration can be uniformly approximated by basic functions. Model-theoretic properties become meaningful; we show that the theory is simple, and indicate some interesting phenomena, in particular, a nontrivial connected component of the Kim–Pillay group, when it is embedded within the model companion ACFA of the theory of difference fields.

Adding Ψ to the language presupposes a passage to continuous logic, able to accommodate functions with real or complex values. As equality on the field is treated in the usual way (with no structural metric present) only a very basic form of continuous logic is needed; we review it below, following the statement of the main theorem. In a structure F for this logic, a formula $\phi(x_1, \dots, x_n)$ is interpreted as a function $\phi^F: F^n \rightarrow \mathbb{C}$, with bounded image (the bound depends only on ϕ). In case the image is discrete, say $\{0, 1\}$, the pullback of 1 is called a *definable subset* of F^n , or for emphasis a *discretely definable subset*; but these can be rare.

For any class of (possibly enriched) fields \mathcal{C} , we define the *characteristic zero asymptotic theory* of \mathcal{C} to be $\text{Th}(\mathcal{C}) \cup \{2 \neq 0, 3 \neq 0, \dots\}$. As we only consider characteristic 0 limits in this paper, we will also simply refer to this as the *limit theory* of \mathcal{C} .

Theorem 1.1. *Consider the characteristic zero asymptotic theory of the class of finite fields \mathbb{F}_q enriched with a nontrivial additive character Ψ .*

- (1) *The theory is decidable, and has an explicit axiomatization PF_+ .*
- (2) *PF_+ admits quantifier-elimination relative to algebraically bounded quantifiers.*
- (3) *If F is an ultraproduct of the enriched fields (\mathbb{F}_q, Ψ) , the pseudo-finite measure is definable. The image under Ψ of a definable subset of F^n with its pseudo-finite measure, is a finite union of rational affine subspaces of T^n , with rationally weighted Lebesgue measure.*
- (4) *PF_+ is a simple theory: indeed a higher amalgamation principle holds.*
- (5) *PF_+ is conservative over the pure language of rings in as far as definable sets go: any definable subset of F^n is already definable in the language of fields and falls into Ax's theory.*
- (6) *The completions of PF_+ are determined by their 'absolute numbers', i.e. by the field of algebraic numbers lying in a model of T , along with the additive character Ψ restricted to that field.*

These assertions will be explained and proved in Section 3 and Section 4. At this point we will just make a few comments.

Decidability in (1) holds in a strong sense: given a sentence S (formed using the basic relations, connectives and quantifiers), and $\epsilon > 0$, one can effectively find a sentence S' such that $|S - S'| < \epsilon$ in any model of T , and a number field L , such that S' has only quantifiers ranging L ; and the set of possible values of S' is $P(\mathbb{T}^n)$ for some explicit polynomial function P and some n .

Much of the model theory in Theorem 1.1, including a quantifier-elimination to the level of bounded quantifiers, should carry thorough for a character of any commutative algebraic group. For instance for multiplicative characters, the axioms will concern curves in Cartesian powers of the multiplicative group G_m , that do not identically satisfy a multiplicative relation $\prod_i x_i^{m_i} = c$ of bounded height $\sum_i |m_i|$. Certainly, different issues will arise in the interface with geometry. For example, some special multiplicative characters – those whose image is contained in the m th roots of unity, for fixed m – are already definable in the pseudo-finite field structure, and can be understood within the existing theory; the new limit theory will be valid for large finite fields with sufficiently general (rather than arbitrary nontrivial) multiplicative characters. In any case, for simplicity, we will restrict attention for the present to the additive character.

An alternative presentation of the results is possible, employing a new sort \widehat{F} for the whole 'definable character group', in place of naming a single specific character. \widehat{F} carries an abelian group structure, as well as an action of $G_m(F)$, transitive on $\widehat{F} \setminus (0)$. In place of Ψ the language includes a continuous-logic relation ψ on $F \times \widehat{F}$, with values in \mathbb{T} . The universal axioms assert that ψ is \mathbb{Z} -bilinear, and $\psi(cx, y) = \psi(x, cy)$. In particular, for any $d \in \widehat{F}$, $\Psi_d(x) := \psi(d, x)$ is an additive character. In addition, $\text{PF}_{\widehat{F}}$ asserts that $(F, \Psi_c) \models \text{PF}_+$. Then PF_+ interprets $\text{PF}_{\widehat{F}}$; indeed PF_+ is bi-interpretable with $\text{PF}_{\widehat{F}}$ if one

adds a constant symbol d in the sort \hat{F} to the language, with the sole axiom $d \neq 0$; the interpretation of Ψ will be Ψ_d . Then Theorem 1.1 holds for $\text{PF}_{\hat{F}}$ in full; the completions are now determined purely by their absolute numbers as fields. We will not adopt this variation in part since it is special to the additive group, and also so as not to sweep under the rug the natural question discussed in the next paragraph.

On the prime fields \mathbb{F}_p , one can define a canonical additive character

$$\Psi_p: \mathbb{F}_p \rightarrow \mathbb{T}, \quad a \bmod p \mapsto \exp\left(2\pi \frac{a}{p}\right).$$

The limit theory of the class of finite fields (\mathbb{F}_p, Ψ_p) with this specific choice of additive character contains PF_+ , and is generated over it by some quantifier-free sentences. All the assertions of Theorem 1.1 hold for this theory with the possible exception of the first, decidability. An explicit axiomatization PP_+ is suggested, that would imply decidability; but whether or not it suffices to axiomatize the limit theory hinges on a certain open question in number theory, close to a well-known problem (and theorem) of Duke, Friedlander, Iwaniec in [11]. This will be discussed in Section 5.

The last section contains some remarks, conjectures and questions regarding extensions to p -adic additive characters, and (especially) to difference fields and ACFA.

For a field F , F^a will denote the algebraic closure.

1.2. Continuous logic

A *continuous model theory* quite able to deal with real- or complex-valued functions was already available at the time of Ax's paper; see [5]. But it was much more recently, following sustained work of many, that the boundaries with discrete first-order logic were dissolved, extending the scope of model theory, almost seamlessly, beyond pure algebra; we refer to [4]. Much of the effort in developing the general theory arises from a metric, that replaces equality in the discrete case. For our present purposes this metric is unnecessary, and even impossible, i.e. the only metric one could consistently impose is the discrete one. Thus, the framework is a very small variation on the familiar case; we briefly present it here.

Function symbols and the formation of terms are treated in the usual way. Basic relations R come with a compact subset $V_R \subset \mathbb{C}$, their intended range of values. (We treat \mathbb{C} as having a distinguished $\sqrt{-1}$, so that a complex-valued relation is equivalent, if desired, to two real-valued ones.) It is best to think of ϕ not as a yes-no question with a smeared out set of possible answers, but simply as a question that has a range of possible answers in the first place. In case $V_R = \{0, 1\}$, we can call V_R two-valued; in particular, this is the case for the equality symbol. Formulas are formed using connectives and quantifiers, then closed under uniform limits (uniform over all evaluations in all structures for the language). The connectives are complex conjugation and complex polynomial operations $+$, $-$, \cdot ; however by Stone–Weierstrass we can also throw in any continuous function from \mathbb{C}^n to \mathbb{C} , without changing the set of formulas. Quantifiers are replaced by continuous \mathbb{C} -valued maps on the space of compact subsets of a given compact of \mathbb{C} ; suprema and infima of

real-valued functions suffice. A structure \underline{A} is a set A along with a function $R^{\underline{A}}: A^n \rightarrow V_R$, for each basic n -ary relation ϕ . The interpretation $\phi^{\underline{A}}$ of an formula ϕ , along with the compact set V_ϕ in which it takes values, is then defined in the obvious way.

All the usual definitions and notions of basic model theory generalize readily, once one gets used to the transposition. For instance, a sentence is a formula with no variables. For each structure \underline{A} , $\text{Th}(\underline{A})$ is the map $\phi \mapsto \phi^{\underline{A}}$ from the set of sentences ϕ to points in V_ϕ . The theory of a class \mathcal{C} of structures is the assignment of a closed subset to each V_ϕ , namely the closure of $\{\phi^{\underline{A}} : \underline{A} \in \mathcal{C}\}$. A theory is *complete* if it gives each sentence ψ a definite value (in V_ψ). A complete theory T is *decidable* if there exists an algorithm that given a sentence ψ , and any $\epsilon > 0$, is guaranteed to terminate and output the value of ψ up to a possible error below ϵ . Likewise, a possibly incomplete theory is decidable if given a sentence ψ , and any $\epsilon > 0$, an algorithm can output a finite subset F of V_ψ , such that for any $\underline{A} \models T$, $\psi^{\underline{A}}$ is at a distance at most ϵ from F ; and conversely any point of F is at a distance at most ϵ from some such $\psi^{\underline{A}}$. When T is complete and has an explicit (or recursive) axiomatization, it is automatically decidable.

The compactness theorem holds: if Φ is a set of pairs (ϕ, B) consisting of a sentence ϕ and a closed subset B of V_ϕ , we say Φ is *satisfiable* if there exists a structure \underline{A} with $\phi^{\underline{A}} \in B$ for each pair $(\phi, B) \in \Phi$. Φ is *finitely satisfiable* if each finite subset is satisfiable. Then any finitely satisfiable set is satisfiable.

There is also an ultraproduct construction of \underline{A} . The theory of saturated models generalizes, etc. In fact, deeper theories including stability and simplicity generalize too, [3], but we will not use them here.

Let us recall also the definition of quantifier-elimination in this setting:

Definition 1.3. T admits quantifier-elimination if for any formula ψ and any $\epsilon > 0$ there exist atomic formulas ϕ_1, \dots, ϕ_k and a continuous function C such that whenever $M \models T$ and $a \in M^x$, we have

$$|\psi(a) - C(\phi_1(a), \dots, \phi_k(a))| < \epsilon.$$

The usual criteria for QE go through from the discrete 1st-order logic case. When T admits quantifier-elimination, a type is determined by a quantifier-free type (and only then). (*Proof.* the continuous map restricting complete types to qf types will under these circumstances be a bijection; as the two spaces are compact Hausdorff, it is a homeomorphism.) It follows likewise that T admits QE provided, for any two λ^+ -saturated models, $\lambda \geq |L|$, any isomorphism between substructures of cardinality $\leq \lambda$ extends to an another whose domain includes a prescribed element.

Exponential sums in a model theoretic setting were discussed by Tomasic in [27], who noted that in positive characteristic the additive character is definable, and used equidistribution to determine the theory of certain reducts of pseudo-finite fields. Indeed if $(F, +)$ has exponent p , a character $\chi: F \rightarrow \mathbb{T}$ takes only p values, the p th roots of 1; the kernel of χ can be taken to be the image of the Artin-Schreier map $\wp(x) = x^p - x$; and so Ax's

discrete first-order theory already interprets the p -element group $\text{Hom}(F/\wp(F), \mathbb{C}^*)$. For this reason, we will consider only the characteristic zero limit in this paper.

Chieu Minh Tran [29] used the exponential sum estimates to determine the (discrete) first order theory of \mathbb{F}_p^a with ‘multiplicative intervals’.

Boris Zilber [32] considers exponential sums over finite rings, with a view to quantum mechanical integrals; he takes different limits than we do here. See also [2].

2. A wrong turn

This section is merely a “no through way” sign on a sideroad; the reader may pass it by without loss, except perhaps of some contrast: what seems like a slight variation leads inexorably to undecidability.

Let us identify the universe of \mathbb{F}_p with $0, 1, \dots, p - 1$, with addition and multiplication modulo p . Define

$$\Lambda_p: \mathbb{F}_p \rightarrow [0, 1] \quad \text{by } \Lambda_p(m) = m/p;$$

it is the inverse of the function

$$[0, 1] \cap \frac{\mathbb{Z}}{p} \rightarrow \mathbb{F}_p, \quad x \mapsto px \text{ mod } p.$$

Taking an ultraproduct of these structures, we obtain a pseudo-finite field F with a function $\Lambda: F \rightarrow [0, 1]$. We have $\Psi = e \circ \Lambda$, where $e(x) = e^{2\pi i x}$. From the point of view of continuous logic, this is an interpretation (e being continuous) but not a bi-interpretation, since e is not quite injective. We will soon see that $(F, +, \cdot, \Psi)$ does not interpret $(F, +, \cdot, \Lambda)$ in any way, though it does of course interpret $(F, +, \cdot, \sin \circ \Lambda)$, and is in fact bi-interpretable with it.

On first seeing Kowalski’s results, I was struck by this statement. (I understand Lou Van den Dries made similar observations and questions.)

Proposition 2.1 (Kowalski). *For an affine algebraic variety $V \subset \mathbb{A}^n$, defined over \mathbb{F}_p , the number of points of V in $[0, p/2]^n$ is approximately $p^{\dim(V)} 2^{-n}$, unless V is contained in some linear hyperplane of \mathbb{A}^n .*

The caveat on linear hyperplanes has an innocent ring at first hearing, and the statement sounds very much like the geometric input needed for a quantifier-elimination theorem, in the style of Ax, for the theory of pseudo-finite fields enriched with Λ . However, this cannot be the case:

Proposition 2.2. *Let $I_p = \lambda(\mathbb{F}_p)$ be the image in \mathbb{F}_p of an interval in $\{0, \dots, p - 1\}$. Assume $|I_p|$ and $|\mathbb{F}_p \setminus I_p|$ are unbounded. Let*

$$T = \{ \phi : (\exists n)(\forall p > n)(\mathbb{F}_p, +, \cdot, I_p) \models \phi \}, \quad T^\perp = \{ \phi : \neg \phi \in T \}.$$

Then T, T^\perp are recursively inseparable.

Proof. Let $<$ be the image of the archimedean ordering on $\{0, \dots, p - 1\}$. We may assume

$$a_p := |I_p| < p/2.$$

By intersecting the interval I_p with translates, one obtains a uniformly definable family of convex definable sets including all intervals $[0, a]$ with $a \leq a_p$. So the restriction of $<$ to $[0, a_p]$ is definable. Note that $\lceil a_p/m \rceil$ (the integer part of a_p/m) is definable for each m , and that $\lceil \sqrt{a_p} \rceil$ is definable since $[0, \sqrt{a_p}/2]$ is the largest segment contained in

$$[0, a_p/4] \cap \{x : x^2 \in [0, a_p/2]\};$$

similarly for higher roots. Now consider polynomials F, G in variables $x = (x_1, \dots, x_k)$ and of some given degree d , with non-negative integer coefficients. Let D be a set with $T \subset D$ and $D \cap T^\perp = \emptyset$. Then $\mathbb{N} \models (\exists x)(F(x) = G(x))$ if and only if

$$\lceil (\exists x_1, \dots, x_k \in [0, \sqrt[d+1]{a_p}]) F = G \rceil \in D$$

Thus, we can decide existence of solutions of integer polynomials if we have access to D , so by Matiyasevich's theorem D is not recursive. ■

Remark 2.3. (1) Following Gödel, using the Chinese remainder theorem, we could define the exponential function $x \bmod p \mapsto 2^x \bmod p$ on an unbounded interval contained in $[0, \log_2(p)]$; then we could appeal to Davis–Putnam–Robinson instead, or (with more quantifiers) even directly to Gödel's theorem, in place of Matiyasevich.

(2) T is in fact Σ_2^0 -complete: one can reduce to T the question of whether a given r.e. set E is finite. Say again that E is defined by $(\exists x)(F(x) = G(x))$, with F, G of degree $d \geq 2$. It follows from [20] that for large n , there exists a prime p with $n^{d+1} < p < (n + 1)^{d+1}$; so $n = \lceil \sqrt[d+1]{p} \rceil$. Thus, E is finite iff for all but finitely many primes p ,

$$(\mathbb{F}_p, +, \cdot, I_p) \models \neg (F(\lceil \sqrt[d+1]{p} \rceil) = G(\lceil \sqrt[d+1]{p} \rceil));$$

iff this sentence lies in T . (Here Ingham's theorem was used to obtain a reasonable large 'tally'. We could use the more basic Chebyshev's theorem / Bertrand's postulate instead, obtaining arithmetic up to $\log(p)$ in place of $\sqrt[d+1]{p}$; or a trivial bound on prime gaps, reaching approximately $\log \log(p)$; all sufficient of course for our present purposes.)

Let us note two consequences of Proposition 2.2. The first is a curious, purely negative alternative proof of another result of Kowalski's, that he drew as a corollary of his positive results.

Corollary 2.4. *Intervals cannot be defined uniformly in \mathbb{F}_p , unless they are bounded or co-bounded.*

Proof. Ax showed that the theory of pseudo-finite fields is decidable, while according to Proposition 2.2 they cannot be uniformly defined in any decidable theory. ■

We include one further statement in this vein.

Corollary 2.5. *Let $X_p \subset \mathbb{F}_p$ be a definable set in the theory of pseudo-finite fields of characteristic 0, or any decidable expansion. Assume $|\mathbb{F}_p \setminus X_p|$ is unbounded with p . Let $k(p)$ be the minimal archimedean gap in $\mathbb{F}_p \setminus X_p$, i.e. the minimum $k \geq 1$ such that for some $a \in \mathbb{F}_p \setminus X_p$ we have $a + 1, \dots, a + k - 1 \in X_p$, $a + k \notin X_p$. Then $k(p)$ is bounded.*

Proof. Let $k = k(p)$. By taking $u = a$ we see that

$$\{x : (\forall u)(u \notin X, u + 1 \in X \Rightarrow u + x \in X)\}$$

is contained in $[1, k]$; it must be equal to $[1, k]$ by minimality of k . Thus, an interval of length $k(p)$ is definable. So by Proposition 2.2 $k(p)$ is bounded, or else the interval is cobounded, and then $F \setminus X$ is bounded. ■

Thus, one cannot expect to have a good quantifier elimination for

$$\text{Th}(\{(\mathbb{F}_p, +, \cdot, I_p) : p\}),$$

and hence not for $\text{Th}(F, +, \cdot, \lambda)$.

I found this situation puzzling; it seems to deprive us of a natural logical setting for Proposition 2.1, playing a role analogous to the theory [1] of pseudo-finite fields for [9].

A part of the solution is the use of continuous logic; When intervals are defined in continuous logic, via a map $[0, p - 1] \rightarrow [0, 1]$ or more conveniently to the unit circle $\mathbb{T} \subset \mathbb{C}$, the nature of the logic blurs the endpoints of intervals, and thus removes at least the apparent source of undecidability. The price one pays, of course, is that while the theory accounts well for intervals of length p/m for any given m , it cannot access any intervals of length of order $o(p)$.

But even the continuous logic theory of $(F, +, \cdot, \Lambda)$ is undecidable. Recall that a \wedge -definable is called *definable* if its complement is also \wedge -definable. On $\Lambda^{-1}[1/3, 2/3]$, the relation $\Lambda(x) < \Lambda(y)$ is ∞ -definable: since $\Lambda(x) < \Lambda(y)$ iff $\Lambda(x - y) \geq 2/3$. Thus, $\Lambda^{-1}(2/5, 3/5)$ is \wedge -definable (the conjunction of $2/5 < \Lambda(x)$ and $x < 3/5$, both taken on $[1/3, 2/3]$; but clearly so is the complement $\Lambda^{-1}[0, 2/5] \cup \Lambda^{-1}[3/5, 1]$; so $\Lambda^{-1}(2/5, 3/5)$ is definable, and $\Lambda^{-1}(<)$ is definable on it.

Corollary 2.6. *The continuous logic theory of the class of enriched prime fields $(\mathbb{F}_p, +, \cdot, \Lambda)$ is undecidable.*

In the next section, we begin studying the theory PF_+ , with Λ replaced by $e \circ \Lambda$. For measure-theoretic or distribution statements (see Section 3.24) it is indistinguishable from the theory of $(F, +, \cdot, \Lambda)$ and, in particular, does provide a framework including Proposition 2.1, while remaining in the same world.

3. The theory of pseudo-finite fields with an additive character

We work in continuous logic, as described above.

3.1. The language

The language L_+ has a sort F for the field, with equality treated in the usual (discrete) way; with the ring operations. And there is one additional unary relation

$$\Psi: F \rightarrow \mathbb{T},$$

where $\mathbb{T} = \{z \in \mathbb{C} : |z| = 1\}$ is the complex unit circle. The function

$$(x_1, \dots, x_n) \mapsto (\Psi(x_1), \dots, \Psi(x_n)): F^n \rightarrow \mathbb{T}^n$$

will be denoted by $\Psi^{(n)}$, or just Ψ when no confusion can arise.

3.2. Defined terms

Let $P(u_1, \dots, u_n; x)$ and $Q(u_1, \dots, u_n; x)$ be integral polynomials. In any field F , we define $\kappa(b_1, \dots, b_n) = c$ if there exists at least one root in F of $P(b_1, \dots, b_n)$, and for any such root d we have $Q(b_1, \dots, b_n, d) = c$. (If no such c exists, let $\kappa_{P,Q}(b_1, \dots, b_n) = 0$.) The functions κ are algebraic functions, and at the same time are well-defined functions in the usual sense; we will view them as basic terms in the language.

With these terms in the language, Ax's theory PF (in the usual, discretely valued first order setting) admits quantifier-elimination, and any substructure is definably closed.

We also include some defined terms using Ψ . Let

$$\Psi_{\text{sym}}^n(c_1, \dots, c_n) := \sum \{\Psi(x) : x^n + c_1 x^{n-1} + \dots + c_n = 0\}$$

(To be precise, if $x^n + c_1 x^{n-1} + \dots + c_n = \prod_{i=1}^n (x - \alpha_i)$ with $\alpha_i \in F^a$, we define $\Psi_{\text{sym}}^n(c_1, \dots, c_n)^F = \sum_{i=1}^n \Psi(\alpha_i)$; where $\Psi(x) = 0$ for $x \in F^a \setminus F$.) Ψ_{sym}^n is clearly definable; it is our analogue of the algebraically bounded quantifiers required for quantifier-elimination in the case of Ax.

Thus, a basic formula has the form $\Psi_{\text{sym}}^n(g_1(u), \dots, g_n(u))$ with g_i a basic PF-definable function as above.

Remark 3.3. We note a few closure properties of these terms.

- (i) $\Psi_{\text{sym}}^1(-c) = \Psi(c)$. In particular, $1 = \Psi_{\text{sym}}^1(0)$.
- (ii) $\Psi_{\text{sym}}^n(-c_1, c_2, \dots, (-1)^n c_n)$ is the complex conjugate of $\Psi_{\text{sym}}^n(c_1, \dots, c_n)$.
- (iii) $\Psi_{\text{sym}}^n(c_1, \dots, c_n) + \Psi_{\text{sym}}^m(d_1, \dots, d_m) = \Psi_{\text{sym}}^{n+m}(e_1, \dots, e_{n+m})$, where the e_i are the coefficients of $(\sum c_i x^i)(\sum d_j x^j)$.
- (iv) Let $\alpha(x, u)$ be any PF formula, such that $PF \models (\forall u)(\exists^{\leq n} x)\alpha(x, u)$. Then $\sum \{\Psi(x) : \alpha(x, u)\}$ can be expressed as a basic formula. (Let t be a variable and write $\Pi\{(t - c) : \alpha(c, u)\} = \sum_{i=0}^n g_i(u)t^{n-i}$; then the $g_i(u)$ ($i = 0, \dots, n$) are PF-definable functions, and $\Psi^\alpha(u) = \Psi_{\text{sym}}^n(g_1(u), \dots, g_n(u))$.)

- (v) More generally, if $\beta(x, u)$ is a function into \mathbb{N} with finite image, whose level sets are PF-definable, then $\sum_x \{\beta(x, u)\Psi(x) : \alpha(x, u)\}$ can be expressed as a basic formula. (Split α into a disjoint union of α_i on which β has constant value $v_i \in \mathbb{N}$, and apply the previous two items.)
- (vi) The product $\Psi_{\text{sym}}^n(c_1, \dots, c_n)\Psi_{\text{sym}}^m(d_1, \dots, d_m)$ can be expressed as a basic formula. (Take β to give the number of ways that an element can be written as a sum of a root of $\sum c_i x^i$ and a root of $\sum d_j x^j$.)
- (vii) Hence, the expressions $\frac{1}{m}\Psi_{\text{sym}}^n(u)$ form a ring, closed under complex conjugation. By Stone–Weierstrass, an arbitrary continuous function composed with basic formulas can be uniformly approximated by a basic formula.

3.4. Axioms for PF₊

A hyperplane $Y \subset \mathbb{A}^n$ is said to have *height* $\leq m$ if it can be defined by a linear equation

$$\sum A_i X_i = b$$

with $A_i \in \mathbb{Z}$, $|A_i| \leq m$.

- (1) F is a field containing \mathbb{Q} , with a unique Galois extension of order n for each n ;
- (2) $\Psi: (F, +) \rightarrow \mathbb{T}$ is a homomorphism;
- (3) Let $h \in \mathbb{Q}[z_1, z_1^{-1}, \dots, z_n, z_n^{-1}]$ be a finite Fourier series (Laurent polynomial) with degrees $\leq m$, real-valued on \mathbb{T}^n , with no constant term. For any absolutely irreducible curve $C \subset \mathbb{A}^n$, not contained in any hyperplane of height at most m ,

$$\sup \{h(\Psi^{(n)}(x)) : x \in C\} \geq 0;$$

- (∞) Of course, if we add the defined terms $\kappa, \Psi_{\text{sym}}$ to the language, their definition must be included among the axioms; in particular for κ ,

$$(\forall u_1, \dots, u_n, x)(P(u_1, \dots, u_n, x) = 0 \rightarrow Q(u_1, \dots, u_n, x) = \kappa_{P,Q}(u_1, \dots, u_n)).$$

The essential content of the AE axiom group (3) is that $\Psi(C)$ is dense in \mathbb{T}^n , provided C is not contained in any *rational* affine subspace of F^n ; see Lemma 3.5 below. One could more easily formulate axioms for C not contained in *any* subspace of F^n ; but that would not suffice.

Lemma 3.5. *Let F be a field of characteristic zero, and $\Psi: (F, +) \rightarrow \mathbb{T}$ a homomorphism. Then $F \models \text{PF}_+(3)$ if and only if the following condition holds:*

- (3*) *Let $C \subset \mathbb{A}^n$ be an absolutely irreducible curve over F , not contained in any rational hyperplane of \mathbb{A}^n . Then $\Psi^{(n)}(C(F))$ is dense in \mathbb{T}^n .*

Here a *rational hyperplane* is one defined by an equation $\sum_{i=1}^n m_i X_i = b$, with $m_i \in \mathbb{Z}$ and $b \in F$.

Proof. Assume PF_+ (3) holds. Let U be a nonempty open subset of \mathbb{T}^n . Find a continuous $g \geq 0$, whose support is contained in U , and with $\int g \, d\lambda = 3$, where λ denotes normalized Haar measure on \mathbb{T}^n . Using Stone–Weierstrass, find a polynomial $h(x)$ in $z_1, \dots, z_n, \bar{z}_1, \dots, \bar{z}_n$ with rational coefficients, of degree say m , with $\|g - h\| < 1$ in the uniform norm; replacing h by the average of h with its complex conjugate, we may assume h takes real values; in particular, the constant term $h_0 \in \mathbb{R}$. So,

$$h_0 = \int_{\mathbb{T}^n} h \geq \left(\int_{\mathbb{T}^n} g - 1 \right) > 1.$$

By the axiom, there exists $c \in C(F)$ with $|h(\Psi(c))| \geq 1$. Thus, $\Psi(c) \in U$. As U was arbitrary, $\Psi(C(F))$ is dense in \mathbb{T}^n .

Conversely, assume (3*). Let C, h be as in PF_+ (3). Let (a_1, \dots, a_n) be a generic point of C over F . Then $\{a_1, \dots, a_n\}$ generate a torsion-free, hence free, subgroup of $\mathbb{A}/\mathbb{A}(F)$; let $b_1, \dots, b_l \in \langle a_1, \dots, a_n \rangle$ represent a free generating set. Writing $a = (a_1, \dots, a_n)$, $b = (b_1, \dots, b_l)$, we have $b = Ba$ for some $l \times n$ -integral matrix B ; and $a = Mb + a_0$ for some $n \times l$ -matrix M and $a_0 \in F^n$. The matrix B defines a linear transformation $B: \mathbb{A}^n \rightarrow \mathbb{A}^l$, and $z \mapsto Mz + a_0$ an affine transformation $\mathbb{A}^l \rightarrow \mathbb{A}^n$; their composition is the identity on C , showing that B restricts to an isomorphism $C \rightarrow C' := B(C)$.

Write $h = \sum c_i h_i$ as a finite linear combination of Laurent monomials. For each monomial h_i we have

$$h_i \circ \Psi^{(n)} \circ M = \Psi \circ H_i,$$

where $H_i(x) = \sum k_i x_i$ is a linear form with $|k_i| \leq m$; H_i is nonzero since h , by assumption, has no constant term. The linear form $H_i \circ M$ is also nonzero: otherwise H_i would be constant on the image of L , but this image contains C , so C is contained in a translate of $\ker(H_i)$, contradicting the assumption that C is not contained in any hyperplane of height $\leq m$. It follows that

$$h \circ \Psi^{(n)} \circ L = h' \circ \Psi^{(m)}$$

for an appropriate real-valued Laurent polynomial; the monomials of h' are (constant multiples of) $z^{H_i \circ M}$, so h' also has no constant term. Integrating h' with respect to the Haar measure on \mathbb{T}^m , we have $\int h' = 0$ (since each monomial has integral 0 on \mathbb{T}^m) and so $\sup h' \geq 0$. By (3*) applied to C' , $\Psi^{(m)}(C')$ is dense in \mathbb{T}^m , so $\sup_{C'} h' \circ \Psi^{(m)} \geq 0$. This proves (3). ■

Lemma 3.6. *Let χ_q be any nontrivial additive character on \mathbb{F}_q , and let F be any ultra-product of (\mathbb{F}_q, χ_q) of characteristic zero. Then $F \models \text{PF}_+$.*

Proof. Let h be a Laurent polynomial (or finite Fourier series) of degree $\leq m$ in variables $z_1, \dots, z_n, z_1^{-1}, \dots, z_n^{-1}$, and taking real values on T^n . By subtracting the constant term, we may assume it is zero.

Let $C \subset \mathbb{A}^n$ be an absolutely irreducible curve over \mathbb{F}_p , not contained in any subspace of height at most m . For a nonzero $k \in \mathbb{Z}^n$ of height $\leq m$, we consider

$$f(x) = k \cdot x = \sum_i k_i x_i$$

as a function on C . Then f is not constant on C , so the Weil bound applies. Thus,

$$\left| \sum_{x \in C(\mathbb{F}_p)} \chi_p(k \cdot x) \right| \leq bp^{1/2}$$

for some b independent of p and of any parameters needed to define $C \subset \mathbb{A}^n$ within an algebraic family. Now for each monomial $\prod_i z_i^{k_i}$ of h , we have

$$\prod_i \chi_p(x_i)^{k_i} = \chi_p\left(\sum_i k_i x_i\right).$$

It follows that

$$\left| \sum_{x \in C(\mathbb{F}_p)} h(\chi_p^{(n)}(x)) \right| \leq b' p^{1/2}$$

for an appropriate $b' > 0$ (namely b times the number of monomials in h). Thus,

$$|C(\mathbb{F}_p)| \max_{x \in C(\mathbb{F}_p)} h(\chi_p^{(n)}(x)) \geq \sum_{x \in C(\mathbb{F}_p)} h(\chi_p^{(n)}(x)) \geq -b' p^{1/2}.$$

So,

$$\max_{x \in C(\mathbb{F}_p)} h(\chi_p^{(n)}(x)) \geq -b' p^{1/2} / |C(\mathbb{F}_p)|.$$

Letting $p \rightarrow \infty$, taking into account that $|C(\mathbb{F}_p)| \sim p$, we obtain the result. The proof for \mathbb{F}_q is identical. ■

3.7. Quantifier elimination

Lemma 3.8. *Let $F, F' \models \text{PF}_+$, let A, A' be substructures of F, F' , respectively, for the language enriched with function symbols $\kappa, \Psi_{\text{sym}}^{(n)}$ as in Section 3.2. Let $\alpha: A \rightarrow A'$ be an isomorphism. Then α extends to an isomorphism $\text{acl}(A) \rightarrow \text{acl}(A')$; here $\text{acl}(A)$ denotes the relative algebraic closure of A in F .*

Proof. Since A, A' are closed under the κ -functions, they are definably closed as substructures of the pseudo-finite fields F, F' . We can extend α further to a *field* isomorphism

$$\text{acl}(A) \rightarrow \text{acl}(A').$$

Thus, we may assume

$$\text{acl}(A) = \text{acl}(A') =: C$$

as fields, $A = \text{dcl}(A) \leq C$; F, F' have two additive characters Ψ, Ψ' , agreeing on A . We also know that $\Psi_{\text{sym}}^n(c_1, \dots, c_n)$ has the same value, for $c_1, \dots, c_n \in A$, whether computed with respect to Ψ or to Ψ' . We have to find in this situation a field automorphism taking Ψ to Ψ' . By compactness of $\text{Aut}(C/A)$, it suffices to find, given a finite set β of elements of C , an automorphism σ of C over A such that $\Psi'(b) = \Psi(\sigma(b))$ for $b \in \beta$. We may enlarge β so as to be $\text{Aut}(C/A)$ -invariant. Let $B = A(\beta)$, and $G = \text{Aut}(B/A)$. Introduce

variables T_b for $b \in \beta$, and obtain a G -action on $\mathbb{C}[T_b : b \in \beta]$ (fixing \mathbb{C} ; $(g, T_b) \mapsto T_{g(b)}$). Consider the linear polynomials

$$f = \sum_{b \in \beta} \Psi(b)T_b, \quad f' = \sum_{b \in \beta} \Psi'(b)T_b,$$

where T_i are variables. Let

$$\phi = \prod_{g \in G} g \cdot f, \quad \phi' = \prod_{g \in G} g \cdot f'.$$

These are products of linear polynomials; if we show $\phi = \phi'$, it will follow that f divides ϕ' , so $f = g \cdot f'$ for some $g \in G$; then g^{-1} would be precisely the σ we are after.

In order to show that $\phi = \phi'$, we compute ϕ explicitly in terms of Ψ_{sym} over A . Write

$$\phi = \sum \phi_v T^v.$$

Then ϕ_v is a finite linear combination, with combinatorial coefficients, of products

$$\Psi(b_1) \cdots \Psi(b_n), \quad n = |G|,$$

with $b_1, \dots, b_n \in \beta$; we rewrite such products as $\Psi(b_1 + \cdots + b_n)$. Let $n\beta$ be the set of all n -fold sums of elements of β ; G acts naturally on $n\beta$. If $b = (b_1 + \cdots + b_n) \in n\beta$ and $\Psi(b_1 + \cdots + b_n)$ occurs in ϕ_v , then so will $\Psi(hb_1 + \cdots + hb_n)$, for $h \in G$; we can thus express ϕ_v as an integral linear combination of certain sums $\sum_{h \in G} \Psi(hb)$, $b \in \beta$. Write

$$\prod_{h \in G} (t - hb) = t^n + d_1(b)t^{n-1} + \cdots + d_n(b).$$

Then each $d_i(b) \in A$, and

$$\sum_{h \in G} \Psi(hb) = \Psi_{\text{sym}}^n(d_1(b), \dots, d_n(b)).$$

This expresses ϕ_v as a sum of Ψ_{sym}^n applied to elements of A , and the same expression is valid for ϕ'_v . ■

Proposition 3.9. PF_+ admits quantifier elimination, in the language including the defined terms $\Psi^{(n)}, \kappa$ of Section 3.2.

Proof. We will use the criterion mentioned at the end of Section 1, constructing an isomorphism between saturated models M, N by an inductive back-and-forth procedure, starting with an arbitrary isomorphism between substructures; see [6].

At a given stage we have small substructures A of M and B of N , and an isomorphism between them, preserving the field structures and the κ and Ψ ; we are to extend it to an isomorphism of M and N . By Lemma 3.8, we may take A, B to be relatively algebraically

closed, as the isomorphism extends. It suffices to show how to extend the isomorphism by one step, so that the domain is a small extension A' of A within M , containing a given element a . We may take

$$A' = \text{acl}(A, a),$$

the relative algebraic closure of $A \cup \{a\}$ in M . Since A' is relative algebraically closed, it suffices to find a field embedding $A' \rightarrow N$ over A , preserving Ψ , and whose image B' is also relatively algebraically closed; for then the terms κ and Ψ_{sym}^n can be computed correctly within A' and B' . We may assume $A = B$ and the given isomorphism is the identity.

Consider finitely generated substructures, say $A(a_1, \dots, a_n)$ where $a = (a_1, \dots, a_n)$ now denotes a generic point of a curve C in affine n -space; since $a \in M^n$ and A is relatively algebraically closed in M , C is absolutely irreducible. By compactness, it suffices to find a generic $b \in C(N)$ with $\Psi^{(n)}(b) = \Psi^{(n)}(a)$. Genericity of b means here that $b \notin A$; by another use of compactness it suffices to avoid finitely many points $\alpha_1, \dots, \alpha_k \in C(A)$; replacing C by the affine curve $C \setminus \{\alpha_1, \dots, \alpha_k\}$, we see that it suffices to find any $b \in C(N)$ with $\Psi^{(n)}(b) = \Psi^{(n)}(a)$.

We can replace (a_1, \dots, a_n) by \mathbb{Q} -linearly independent elements b_1, \dots, b_m generating the same additive group as (a_1, \dots, a_n) . (This uses the fact that if $a_1 = \sum m_j b_j$ with $m_j \in \mathbb{Z}$ then $\Psi(a_1) = \prod \Psi(b_j)^{m_j}$ in the circle group.) Now by Lemma 3.5, there exists an m -tuple (c_1, \dots, c_m) of $C(N)$ with the same values $\Psi(b_i) = \Psi(c_i)$. By saturation of N , we can indeed find an embedding of A' into N , preserving Ψ . If we knew that the image of A' is relatively algebraically closed, it would follow that Ψ_{sym} and κ are preserved too. At this point we have no basis to claim this, however.

Call a substructure D of M *full* in M if the restriction homomorphism

$$\text{Gal}(M) \rightarrow \text{Gal}(D)$$

is bijective, where Gal denotes the absolute Galois group. If D is full in N , and $D \leq D'$, the composition $\text{Gal}(N) \rightarrow \text{Gal}(D') \rightarrow \text{Gal}(D)$ is injective, hence so is the restriction $\text{Gal}(N) \rightarrow \text{Gal}(D')$. In case D' is relatively algebraically closed it is also surjective, so D' is full. Assume A', B' contain full substructures A_0, B_0 . Let B'' be the relative algebraic closure of B' in N . Then,

$$\widehat{\mathbb{Z}} \cong \text{Gal}(M) \cong \text{Gal}(A') \cong \text{Gal}(B').$$

Consider the homomorphisms

$$\text{Gal}(N) \rightarrow \text{Gal}(B') \rightarrow \text{Gal}(B_0).$$

Since the composition $\text{Gal}(N) \rightarrow \text{Gal}(B_0)$ is an isomorphism, $\text{Gal}(B') \rightarrow \text{Gal}(B_0)$ is surjective. But a surjective self-map of $\widehat{\mathbb{Z}}$ is an isomorphism. Thus, $\text{Gal}(B') \rightarrow \text{Gal}(B_0)$ is bijective; and since $\text{Gal}(N) \rightarrow \text{Gal}(B_0)$ is surjective, $\text{Gal}(N) \rightarrow \text{Gal}(B')$ must be surjective too. Thus, B' is relatively algebraically closed in N .

We have shown so far, combining the smallness of the Galois group with Lemma 3.5, that if $D \leq M$ is small (compared to the saturation) and full, and D' is an extension of transcendence degree in which D is full too, then D' embeds into M as a full (and relatively algebraically closed) substructure. Thus, we will be done as soon as we extend the original isomorphism $A \rightarrow B$ to one between full substructures (that we took to be the identity).

By Lemma 3.11, there exists a full subfield A_0 of M of transcendence degree 1 over \mathbb{Q} , and likewise B_0 in N . Let A' be the relative algebraic closure in M of A_0 , and B' of B_0 . We note in passing that the Chebotarev field crossing argument does not immediately work for fields with additive character, since the Ψ -structure on a field amalgam is not uniquely determined. However the argument does show that *some* amalgam A'' of A', B' over $A = B$ exists, with Galois group \widehat{Z} , and A', B' relatively algebraically closed in A'' . By the previous paragraph, since A' is full in M , A'' embeds into M over A' . Likewise A'' embeds into M over B' . Thus, while A', B' need not be isomorphic, we can skip them and obtain an isomorphism $A'' \rightarrow B''$ with $B'' \leq N$, so that we have an isomorphism between full substructures as required. ■

Remark 3.10. Lemma 3.5 can be extended as follows: Let M be an \aleph_1 -saturated model of PF_+ , $M_0 \leq M$ countable, relatively algebraically closed, and p any type over M_0 in the language of pseudo-finite fields. Assume $p(x_1, \dots, x_n)$ implies the \mathbb{Q} -linear independence of the images of x_1, \dots, x_n in M/M_0 . Let $\alpha_i \in T$. Then there exists a realization $a = (a_1, \dots, a_n) \in M$ with $\Psi(a_i) = \alpha_i$.

Namely, let M'_1 be a model of PF extending M_0 , and with $a = (a_1, \dots, a_n)$ realizing p . Extend Ψ to a homomorphism on M'_1 with $\Psi(a_i) = \alpha_i$; this gives a structure M_1 for the language of PF_+ . Extend M_1 iteratively to a structure satisfying $\text{PF}_+ (1,2,3^*)$ (Lemma 3.5). Then, $M_1 \models \text{PF}_+$. By quantifier elimination over relatively algebraically closed substructures, M_1 embeds elementarily into M over M_0 .

Lemma 3.11. *Let T be a theory of pseudo-finite fields of characteristic 0. Then there exists $F \models T$ and a subfield F_1 of F of transcendence degree 1 over \mathbb{Q} , such that the restriction $\text{Gal}(F) \rightarrow \text{Gal}(F_1)$ is an isomorphism.*

Proof. It is easy to give a direct proof, but let us use Ax's theorem instead. Let $M \models T$ and $F_0 = \mathbb{Q}^a \cap M$. Then F_0 is the fixed field of some $\sigma_0 \in \text{Aut}(\mathbb{Q}^a)$. Let $L = \lim_n \mathbb{Q}^a((t^{1/n}))$ be the field of Puiseux series over \mathbb{Q}^a ; so L is algebraically closed. Lift σ_0 to a (continuous) automorphism of L fixing $t^{1/n}$ for each n . Let σ_1 be a topological generator of $\text{Aut}(L/\mathbb{Q}_a(t))$. Let $\sigma = \sigma_0\sigma_1$, and let F_1 be the fixed field of σ . $\text{Gal}(F_1)$ is generated by σ .

Since \widehat{Z} is free, there is a unique homomorphism $\widehat{Z} \rightarrow \text{Gal}(F_1)$ with $1 \mapsto \sigma$, and we denote it by $m \mapsto \sigma^m$. Since σ_0, σ_1 commute, if $m \in \widehat{Z}$ and $\sigma^m = 1$, then

$$\sigma_0^m \sigma_1^m = (\sigma_0\sigma_1)^m = 1,$$

so σ_1^m fixes each $t^{1/n}$; from this it follows that $m = 0$. Thus,

$$\text{Gal}(F_1) = \widehat{\mathbb{Z}}.$$

Let $(L', \sigma') \models ACFA$ with $(L, \sigma) \leq (L', \sigma')$. Let $F = \text{Fix}(\sigma') \leq L'$. Then F is pseudo-finite and $\text{Gal}(F) \rightarrow \text{Gal}(F_1)$ is an isomorphism. According to Ax, $F \models T$. ■

3.12. Completions

As in Ax's case, the completions are determined by the 'absolute numbers' as a field with additive character. In more detail, let $G = \text{Aut}(\mathbb{Q}^a/\mathbb{Q})$ be the absolute Galois group of \mathbb{Q} . It follows from quantifier-elimination that the completions of PF_+ are determined by pairs (σ, ψ) , where $\sigma \in G$ and

$$\psi: \text{Fix}(\sigma) \rightarrow \mathbb{T}$$

is a homomorphism, extending $x \mapsto \exp(2\pi i x)$ on \mathbb{Z} , and with $\sigma(\omega) = \omega^k$,

$$\psi(k/n) = \exp(2\pi i/n)$$

for $\omega^n = 1$. Two such pairs are equivalent iff they are isomorphic as fields with additive characters; and any pair (σ, ψ) can occur. (It is easy to find by an inductive limit argument, a model of PF_+ containing the given structure on $\text{Fix}(\sigma)$.)

3.13. Universal theory

PF_+ is the model completion of the theory described in axioms (1,2); it is easy to extend any model of (the universal part of) this theory to a model of PF_+ , using Lemma 3.5 (3*).

3.14. Decidability

PF_+ is decidable. Given a sentence ψ and $\epsilon > 0$, one can first look for a quantifier-free sentence ψ' and a proof that the values of ψ, ψ' are within $\epsilon/2$. Now ψ' concerns a number field L , that one can take to be Galois over \mathbb{Q} , with Galois group G ; and is determined by an element $g \in G$, and a homomorphism $\psi: \text{Fix}(g) \rightarrow \mathbb{C}^*$. Actually only the values of ψ on finitely many elements e_1, \dots, e_k are concerned. Now e_1, \dots, e_k generate a finitely generated subgroup of $(L, +)$, isomorphic to \mathbb{Z}^l for some l , and we may replace them by a lattice basis for the group they generate; so we may assume they are \mathbb{Q} -linearly independent. In this case, $\psi(e_i) \in \mathbb{T}$ can be chosen arbitrarily and independently. Using this, we can determine the set of possible values of ψ' , to $\epsilon/2$ -accuracy.

3.15. Standard exponential characters

The *standard interpretation* of Ψ for $\mathbb{F}_p = \mathbb{Z}/p\mathbb{Z}$ is by definition the character map

$$\Psi_p: n + p\mathbb{Z} \mapsto \exp\left(2\pi i \frac{n}{p}\right)$$

On \mathbb{F}_q (with q a power of p) we set

$$\Psi_q := \Psi_p(\text{Tr}(a)),$$

where Tr is the trace from \mathbb{F}_q to \mathbb{F}_p . The field \mathbb{F}_q endowed with the additive character Ψ_q will be denoted by \mathbb{F}_q^+ .

Any other additive character on \mathbb{F}_q has the form $\Psi_q(cx)$, for a unique $c \in \mathbb{F}_q$; this is a version of the statement that a finite abelian group is isomorphic to its dual. Hence the group \hat{A} of all additive characters, along with the evaluation map $\hat{A} \times A \rightarrow \mathbb{T}$, can be interpreted in the language naming only one character.

Proposition 3.16. *The characteristic 0 asymptotic theory of the standard finite fields-with-additive-character \mathbb{F}_q^+ is precisely PF_+ . It is equal to the characteristic 0 asymptotic theory of all finite fields with a nontrivial additive character; and also to the characteristic 0 asymptotic theory of all prime fields with a nontrivial additive character.*

Proof. We have seen that PF_+ eliminates quantifiers and hence becomes complete upon a description of Ψ on \mathbb{Q}^a . We thus have only to show that for any $\sigma \in \text{Aut}(\mathbb{Q}^a)$ and any homomorphism $\Psi: \text{Fix}(\sigma) \rightarrow \mathbb{T}$ vanishing on \mathbb{Z} , there exists an \mathbb{F}_q^+ approximating $(\text{Fix}(\sigma), \Psi)$.

The following explicit statement is a little stronger than what we need: *Let $f(X)$ be an irreducible monic polynomial of degree d over \mathbb{Q} ,*

$$K = \mathbb{Q}[X]/(f).$$

Assume the Galois hull L of K is cyclic of degree d' over K , and let r_1, \dots, r_b be elements of K such that $1, r_1, \dots, r_b$ are \mathbb{Q} -linearly independent; also let r_0 be the reciprocal of some integer $k > 0$. Let $g_i \in \mathbb{Q}[X]$ be a polynomial whose image modulo (f) is r_i . Let U be a nonempty open subset of \mathbb{T}^b , and let $u \in \mathbb{T}$ satisfy $u^k = 1$. Then there exist infinitely many primes p and such that f has a root a modulo p , and the splitting field of f over \mathbb{F}_p has degree d' over \mathbb{F}_p ; and such that, moreover, there exists e with

$$(\Psi_{p^e}(g_1(a)), \dots, \Psi_{p^e}(g_b(a))) \in U$$

and such that

$$\Psi_{p^e}(1/k) = u.$$

Using Chebotarev, we can find an infinite sequence of primes p_m and $a_m \in \mathbb{F}_{p_m}$, such that the splitting field of f over \mathbb{F}_{p_m} has degree d' over \mathbb{F}_{p_m} . We will let $q_m = p_m^{e_m}$ for an appropriate sequence of integers e_m . Then

$$\Psi_{q_m}(g_i(a_m)) = \Psi_{p_m}(e_m g_i(a_m)) \quad \text{and} \quad \Psi_{q_m}(1/k) = \Psi_{p_m}(e_m/k).$$

Let $c_m = (g_1(a_m), \dots, g_b(a_m))$. Then it suffices to show given $\epsilon > 0$ that for some infinite set of indices m , we can find e_m such that $\Psi_{p_m}(e_m c_m)$ lies in a prescribed open subset of \mathbb{T}^b , and $\Psi_{p_m}(e_m/k)$ is a prescribed k th root of unity. This is true by Lemma 3.17,

noting that as $m \rightarrow \infty$, the $g_i(a_m)$ avoid any given finite number of \mathbb{Z} -linear relations, even modulo 1 (since the $g_i(a)$ and 1 are linearly independent over \mathbb{Q}). At the same time, we can make sure that e_m is relatively prime to d' ; thus the splitting field of f over $\mathbb{F}_{p_m}^{e_m}$ still has degree d' over $\mathbb{F}_{p_m}^{e_m}$.

The proof for the class of prime fields with an arbitrary nontrivial additive character is completely parallel; at the end, having found the sequences p_m, e_m , in place of moving to $\mathbb{F}_{p_m}^{e_m}$ we remain with \mathbb{F}_{p_m} , but multiply the standard character by e_m . ■

Lemma 3.17. *Let $c_m \in \mathbb{F}_{p_m}^b$ be a sequence of b -tuples from prime fields of increasing size. Assume: for any given nonzero rational vector $(\alpha_1, \dots, \alpha_b)$, for almost all m it is not the case that*

$$\sum_{i=1}^b \alpha_i c_{m,i} = 0.$$

Let U be a nonempty open subset of \mathbb{T}^b . Also fix $k \in \mathbb{N}$ and $l \in (\mathbb{Z}/k\mathbb{Z})^$. Then for arbitrarily large m , for some $e_m \in \mathbb{N}$,*

$$\Psi_{p_m}(e_m c_m) \in U.$$

Moreover, we can choose $e_m = l \pmod k$.

Proof. This follows from an effective form of Weyl’s criterion for equidistribution; but it also follows from the pseudo-finite version Lemma 3.18, that we will prove separately. Note that the ‘moreover’ is obvious, since changing e_m by a multiple of p_m does not affect $\Psi_{p_m}(e_m c_m)$, and p_m is a unit mod k . ■

Lemma 3.18. *Let $(F, +, \cdot, \Psi)$ be an ultraproduct of enriched finite fields \mathbb{F}_q^+ . Let $n \in \mathbb{N}$, $c \in F^n$ and assume $m \cdot c \neq 0$ for $m \in \mathbb{Z}^n \setminus (0)$. Then, $\Psi(Fc) = \mathbb{T}^n$.*

Proof. In any case, $\Psi(Fc)$ is a closed subgroup of \mathbb{T}^n , using \aleph_0 -saturation of the ultraproduct; so if it is not all of \mathbb{T}^n then for some $m \in \mathbb{Z}^n \setminus (0)$ we have $\Psi(Fc)^m = 1$ or $\Psi(Fc') = 1$, where $c' = m \cdot c \in F$. But by assumption, $c' \neq 0$ so $Fc' = F$, and $\Psi(F) = \mathbb{T}$. ■

Note we could not simply apply Weyl’s theorem directly to $\Psi(c)$, since that may fall into some rational hyperplane.

3.19. Simplicity

Proposition 3.20. *PF_+ admits n -amalgamation for algebraically independent algebraically closed substructures. By the case $n = 3$, PF_+ is a simple theory.*

Proof. In fact, a stronger statement is true, namely n -amalgamation over PF : let A be any functor on the partially ordered set $P(n)$ of subsets of $[n] = \{1, \dots, n\}$ into algebraically closed substructures of a model of PF ; see [18]. Let $P(n)^- = \{s \subset n : |s| < n\}$. Assume

further that for $s \in P(n)^-$, a character Ψ_s is given on $A(s)$, so that $\Psi_s|_{A(s')} = \Psi_{s'}$ when $s' \subset s$. Then there exists

$$\Psi_{[n]}: A([n]) \rightarrow T$$

completing the data to a functor into models of the universal part of PF_+ .

The proof is similar to the proof for pseudo-finite fields [17], but moving first to the \mathbb{Q} -linear span of the $n - 2$ -skeleton, arguing that the values of Ψ are determined thus far, and that beyond this one has \mathbb{Q} -linear independence and thus sufficient freedom in the choice of Ψ . More precisely, for $|s| = n - 1$, let $B(s)$ be the \mathbb{Q} -span of $\cup_{s' \subset s} A(s')$. Let $(e_{s,i} : i \in I_s)$ be a subset of $A(s)$, mapping to a \mathbb{Q} -basis of $A(s)/B(s)$. We have to show consistency of $\Psi(e_{s,i}) = \alpha_{s,i}$, for any choice of $\alpha_{s,i} \in T$. By an internalization argument as in [17], we see that the elements

$$(e_{s,i} : s \in P(n), |s| = n - 1, i \in I_s)$$

are \mathbb{Q} -linearly independent as a set. E.g. any linear dependence among the elements $e_{2,3,i}$ over $A(1, 2) \cup A(1, 3)$ would imply, by internalizing $A(1)$ into the base, a similar linear dependence over $A(2) \cup A(3)$, contradicting the choice of the $E_{2,3,i}$. Hence, by Remark 3.10, the PF-type of $A([n])$ along with $\Psi(e_{s,i}) = \alpha_{s,i}$ is consistent with PF_+ . ■

For a definable group G in a saturated structure and a base structure A , G_A^0 denotes the intersection of all A -definable subgroups of finite index; whereas G_A^{00} is the smallest subgroup of G cut out by a partial type over A , and with *bounded* index, i.e. the quotient remains bounded in any elementary extension; in this case the quotient admits a natural compact topological group structure. It has long been open, for discrete first order theories that are simple, whether G_0^{00} must equal G_0^0 . PF_+ gives a natural *continuous-logic* example of a definable group G in a simple theory, where $G_A^0 = G$ for all A , but $G_0^{00} \neq G$, and where G_A^{00} does not stabilize with A . Indeed, let G be the additive group, and let $(a_i : i \in I)$ be a \mathbb{Q} -linearly independent subset of $G(A)$. Then we have a surjective homomorphism

$$G \rightarrow T^A, \quad g \mapsto (\Psi(a_i g) : i \in I);$$

the kernel is clearly A - \wedge -definable.

3.21. Galois group

The compact (Lascar–Kim–Pillay) absolute Galois group G_T of a theory T , whether in discrete or continuous logic, is defined as follows. A (hyper)imaginary sort has the form $S = K^n/E$, with E an ∞ -definable equivalence relation (without parameters). If (for some K , or all sufficiently saturated K) S does not grow upon replacing K by an elementary extension, we say that S is *bounded*. One then defines a compact Hausdorff topology, the *logic topology*, on S , by taking projections of definable sets (with parameters) to be the basic closed sets. G_T is by definition the automorphism group of the family of all compact sorts (the permutations that extend to an automorphism of some model). It is itself naturally a compact Hausdorff group.

If T admits 3-amalgamation over algebraically closed sets in the discrete logic sense, then G_T is totally disconnected. For PF_+ this is the case by Proposition 3.20. See, however, Section 6.3.

3.22. The theory with a space of characters

Recall the theory $\text{PF}_{\hat{\uparrow}}$ from the introduction. When a nonzero constant symbol is added to the dual sort \hat{F} , it is clear that $\text{PF}_{\hat{\uparrow}}$ becomes bi-interpretable with PF_+ . The results of Theorem 1.1 thus generalize easily to $\text{PF}_{\hat{\uparrow}}$. Let us only review the completeness statement, that becomes somewhat stronger: the completions of $\text{PF}_{\hat{\uparrow}}$ are entirely determined by the field of absolute numbers, i.e. by the restriction to the field language.

To see this, let $M, N \models \text{PF}_{\hat{\uparrow}}$ with N \aleph_1 -saturated, and assume

$$M \cap \mathbb{Q}^a = N \cap \mathbb{Q}^a = K.$$

Choose $\chi \in \hat{F}(M)$ and $\psi \in \hat{F}(N)$. Let $(b_i : i \in I)$ be a \mathbb{Q} -basis for K . Use saturation of N to find $d \in N$ such that $\psi(db_i) = \chi(b_i)$ for $i \in I$. (For any finite $\{\beta_1, \dots, \beta_k\} \subset I$, existence of such a d with respect to $\{\beta_1, \dots, \beta_k\}$ follows from the axioms, since the image of \mathbb{A}^1 under $x \mapsto (x\beta_1, \dots, x\beta_k)$ is not contained in any \mathbb{Q} -linear affine space.) Let $\psi'(x) = \psi(dx)$. Then by completeness of PF_+ we see that (N, ψ') and (M, χ) are elementarily equivalent in the language of PF_+ . It follows that M, N are elementarily equivalent in the language of $\text{PF}_{\hat{\uparrow}}$.

Corollary 3.23. *The theory of the class of prime fields, enriched by their dual groups and evaluation map, is precisely $\text{PF}_{\hat{\uparrow}}$.*

3.24. The theory in probability logic

An alternative and more conceptual axiomatization appears if one uses probability logic; specifically the expectation logic in the sense of [15]. We will only evoke it in passing.

We continue to work in continuous logic as sketched in the introduction. *Probability logic*, superimposed upon this, consists of additional syntactical operators E_x behaving formally like quantifiers; thus if $\phi(x, y_1, \dots, y_m)$ is interpreted as a function $F^{m+1} \rightarrow \mathbb{C}$ with bound N , then $E_x\phi(x, y_1, \dots, y_m)$ has free variables among y_1, \dots, y_m , and is interpreted as a function $F^m \rightarrow \mathbb{C}$ with the same bound, intended to denote the expectation of ϕ with respect to x . The probability axioms are:

- (1) $E_x(1) = 1$ for the constant function 1 (viewed as a function of any set of variables);
- (2) $E_x(\phi + \phi') = E_x(\phi) + E_x(\phi')$, and $E_x(\psi \cdot \phi) = \psi \cdot E_x(\phi)$ when x is not free in ψ ;
- (3) $E_x(|\phi|) \geq 0$;
- (4) $E_x E_{x'} = E_{x'} E_x$.

Axioms (1)–(3) ensure that a model carries a definable Keisler measure μ , and $E_x\phi$ is the μ -expectation of ϕ . Axiom (4), asserting the commutativity of the operators $E_x, E_{x'}$,

is the model-theoretic Fubini rule. We need an additional version of Fubini that we formulate for expansions of fields. Let C be an absolutely irreducible curve, and let f, g be two nonconstant regular maps on C . Also let ϕ be a \mathbb{C} -valued definable function on C . For $x \in F$, let

$$f_*\phi(x) = \sum_{y \in C, f(y)=x} \phi(y),$$

the sum being taken over F -points of C . Similarly define $g_*\phi$.

$$(5) \quad E_x f_*\phi = E_x g_*\phi.$$

This allows us to define $E_C\phi$ to be the common value of $E_x f_*\phi$, for any nonconstant regular f . This version is suitable for (pseudo-)finite fields, where no Jacobian term is needed. (This extends uniquely to an integration theory on varieties, see [19, 5.14]; but only the case of curves will be needed below.)

Now for the language of rings with an additive character, we can formulate the following axioms:

- (a) $|\Psi(x)| = 1$;
- (b) $\Psi(x + y) = \Psi(x)\Psi(y)$, $\Psi(0) = 1$;
- (c) For any absolutely irreducible affine curve C , $E_C 1 = 1$ and $E_C \Psi = 0$.

This implies the main axiom scheme of PF_+ : if $C \subset \mathbb{A}^n$ is an absolutely irreducible affine curve, not contained in any \mathbb{Q} -linear hyperplane, let $Y \subset \mathbb{T}^n$ be the image of $\Psi^{(n)}$ on a saturated model. Then Y is the support of a measure μ on \mathbb{T}^n , such that

$$\int \psi \theta = E_C \theta \circ \Psi^{(n)}$$

for any continuous function θ on PF_+ . The μ -integral of any nontrivial character $z_1^{m_1} \cdots z_k^{m_k}$ vanishes on \mathbb{T}^n , while the μ -integral of 1 is 1; so μ is the Haar measure, and $Y = \mathbb{T}^n$. Then PF_+ implies elimination of sup and min quantifiers, while Axiom (c) determines the underlying Keisler measure. We leave the details to the interested reader.

4. Definable integration, and comparison to PF

Let F be an ultraproduct of finite fields with an additive character, F_i , and let V be an absolutely irreducible variety over F . Then for almost all i , $V(F_i)$ we have the normalized counting measure μ_i on $V(F_i)$; the integral of ψ under this measure is

$$\frac{\sum_{v \in V(F_i)} \psi(v)}{|V(F_i)|}.$$

The pseudo-finite measure μ_V is obtained as an ultraproduct of these normalized counting measures.

Proposition 4.1. *The pseudo-finite counting measure μ_V is definable. In other words if $\phi(x, y)$, $y = (y_1, \dots, y_m)$ is a formula then the map*

$$y \mapsto \int \phi(x, y) d\mu_V(x)$$

is a definable map in the sense of continuous logic. Moreover, when V varies in a definable family of definable sets, μ_V is uniformly definable.

Proof. It will be more convenient to show definability for a different normalization, namely for the ultraproduct of the counting measure on V divided by $|F_i|^{\dim(V)}$. The dimension of V is definable in definable families, and one can then obtain the probability measure upon dividing by $\int 1$.

The statement reduces using Fubini to the case of affine curves $C_b \subset \mathbb{A}^n$ varying in some definable family. Using quantifier elimination and Remark 3.3 (vii), it suffices to integrate $\Psi_{\text{sym}}^n(c_1, \dots, c_k)$, where the c_i are PF-definable functions of the variables. By the definition of Ψ_{sym}^n as a sum, this amounts simply to integrating $\Psi(x_1)$ over a PF-definable subset e_d of another affine curve C'_d , with d a definable function of the parameters. Namely, each $c_i(u)$ is uniformly algebraic over u . Thus, the Zariski closure of $\{(u, c_1(u), \dots, c_k(u)) : u \in C_b\}$ is a curve C''_b . Let

$$C'''_b = \{(x, u, z_1, \dots, z_k) : (u, z_1, \dots, z_k) \in C''_b, x^n + z_1x^{n-1} + \dots + z_n = 0\},$$

and $e_b = \{(x, u, z_1, \dots, z_k) \in C'''_b : z_i = c_i(u)\}$ (a PF-definable subset of C'''_b).

Then,

$$\int_{C_b} \Psi_{\text{sym}}^n(c_1, \dots, c_k) = \int_{e_b} \Psi(x).$$

If necessary, passing from b to a parameter d within $\text{acl}(b)$, break up C'''_b into irreducible components.

The case where ψ_d is a finite set is directly definable via Ψ_{sym}^n , see Remark 3.3 (iv). Using definability of dimension (specifically of dimension zero), we can similarly reduce to integrating $\Psi(x_1)$ over an absolutely irreducible curve C_b . If x_1 is constant on C_b , with value v , the integral is $\Psi(v)\mu(1_{\psi_d})$. Otherwise, by the Weil bound, the answer is 0. ■

We can define, for any formula $\phi(x)$, the Fourier transform

$$F(\phi) := \int_{x \in F^n} \Psi(x \cdot y)\phi(x) d\mu(x).$$

Corollary 4.2. *The Fourier transform of any definable real-valued relation on \mathbb{F}_p^n is also definable, uniformly in p , and uniformly in families of definable functions.*

All definable functions on F^n are bounded, in particular, square-integrable; by the Plancherel theorem the Fourier transform of a definable function is square-integrable with respect to the unnormalized counting measure, hence has countable support; by quantifier-elimination the support of the Fourier transform consists of algebraic numbers, if the original function is 0-definable.

In particular, the Fourier inversion formula does not survive the ‘leading order’ continuous-logic ultraproduct that we are using. It would be interesting to know if it holds using the ‘next to leading order’ measure of Section 6.5, or if a further degree of precision is needed. Fourier inversion does hold for almost periodic functions (uniform limits of finite linear combinations of functions $\Psi(cx)$); I do not know if it holds for any others or if this is the precise domain.

Proposition 4.3. *The images under $\Psi^{(n)}$ of PF-definable subsets D of F^n (in a sufficiently saturated model) are finite unions of cosets of subtori of T^n . Moreover, for fixed D , the pushforward measure $\Psi_*^{(n)}\mu_D$ is a finite, positive \mathbb{Q} -linear combination of Haar measures on such cosets (of dimension equal to $\dim(D)$).*

Proof. Consider the topology τ on F^n , whose closed sets are finite unions of \mathbb{Q} -affine subspaces, cut out by (inhomogeneous) \mathbb{Q} -linear equations. Let $\cup_{i=1}^m H_i$ be the closure of D in this topology. Then

$$D_i = D \cap H_i$$

is τ -dense in H_i ; and it suffices to prove the lemma for each D_i separately (for the measure theoretic statement, only the top-dimensional components D_i play a role, the others having pseudo-finite measure 0). After an $SL_n(\mathbb{Z})$ -change of variable, and translation, we may assume H_i is a coordinate subspace, cut out by $x_{k+1} = \dots = x_n = 0$; and then we may work in F^k instead. By the axioms of PF_+ , Ψ maps D_i surjectively to T^k . Moreover, for any nonzero $m \in \mathbb{Z}^k$, the integral of

$$\prod_{i=1}^k \Psi(x_i)^{m_i} = \Psi(\sum_{i=1}^k m_i x_i)$$

under the pseudo-finite measure evaluates to 0; hence the pushforward to T^m of the pseudo-finite measure on D_i must be a multiple α_i of the Haar measure on T^k . Since the total measure of each D_i is rational, we have $\alpha_i \in \mathbb{Q}^{>0}$. ■

Proposition 4.3 is not directly useful to prove definability, since when $D = D_b$ varies in a definable family, the image tori can jump; notably when $D_b = D \cap H_b$ is a hyperplane section of some fixed definable set D , $\Psi^{(n)}(D \cap H_b)$ has lower dimension when b is rational. The counting measure is definable *despite* this translation.

Corollary 4.4. *Let $D \subset F^n$ be a PF-definable set, and let $h: D \rightarrow \mathbb{C}$ be a PF_+ -definable map. Then $h(D)$ has finitely many connected components. Define $x E_h y$ if $h(x), h(y)$ lie in the same connected component. Then E_h is a PF-definable equivalence relation with finitely many classes.*

Proof. The proof of Proposition 4.3 shows that $\Psi^{(n)}(D) = \cup_{i=1}^k C_i$ is a finite union of cosets C_i of subgroups of T^n ; moreover, we can write

$$D = \cup_{i=1}^k D_i$$

with $D_i = D \cap S_i$ definable with parameters, in fact S_i is a \mathbb{Q} -affine subspace of F^n , so that

$$C_i = \Psi^{(n)}(D_i).$$

Now h can be expressed, after replacing D by a finite cover, as $C \circ \Psi$ with $C: \mathbb{C}^n \rightarrow \mathbb{C}$ continuous. It follows that each D_i is mapped into a single connected component of $h(D)$, and so the classes of E_h are finite unions of the D_i . ■

Corollary 4.5 (PF_+ is conservative over PF for definable sets.). *Let $(F, +, \cdot, \Psi) \models \text{PF}_+$ be sufficiently saturated, let $X \subset F^n$, and assume both X and $F^n \setminus X$ are \wedge -definable (with parameters) in $(F, +, \cdot, \Psi)$. Then X is definable (with parameters) in $(F, +, \cdot)$.*

Proof. The characteristic function of X is definable in $(F, +, \cdot, \Psi)$ in these circumstances, and so E_{1_X} is PF -definable by Lemma 4.4. ■

In view of this conservation result, one may wonder whether PF_+ is at all a proper expansion of PF ; certainly some functions onto the interval $[0, 1]$ become definable. But by no means is this the case for Ψ . In continuous logic, definability would mean that given any two subsets $A \subset B \subset \mathbb{T}$ of the unit circle with A compact and B open, there exists a formula $\phi(x)$ in the language of rings such that $\Psi(a) \in A$ implies $\phi(a)$, while $\phi(a)$ implies $a \in B$. This is ruled out by the undecidability results of Section 2. But let us also see it in a more geometric fashion. Let F be a pseudo-finite field (with nonstandard size p^*), and let $\epsilon > 0$ be standard. A definable set D cannot contain an interval I of length $\geq \epsilon p^*$, unless the complement $D' := F \setminus D$ is finite. To see this recall the stabilizer of [17]. In the special case of rank 1, it can be defined as follows: let q be some rank-one type of elements of D' . Then $St_0(q)$ is the set of elements a with $aq \cap q$ infinite; $St_0(q)$ generates a definable group S in two steps, and $S \setminus St_0(q)$ is finite; if $s_1, \dots, s_k \in S$ are (algebraically) independent then $\cap(s_i + q) \neq \emptyset$. Now \mathbb{F}_p has no proper, nonzero definable subgroups, hence neither does F . So $S = F$. Now one can easily find independent $s_1, \dots, s_k \in S$ such that the union of the translates $s_i + I$ is all of F . But then $\cap(s_i + D') = \emptyset$, a contradiction.

A slightly different argument that PF, PF_+ differ: it is clear that the additive group of an ultraproduct of finite fields with additive character admits a \wedge -definable subgroup such that the quotient is connected in the logic topology, namely the kernel of Ψ . But according to results in [17] we have $G^0 = G^{00}$ for any definable group over any base set, i.e. no nontrivial connected quotients exist.

5. Prime fields with standard character

In this section we will consider the class $P+$ of prime fields with their standard character

$$n \bmod p \mapsto \exp(2\pi i p/n)$$

embodying the natural map

$$\mathbb{F}_p = \mathbb{Z}/p\mathbb{Z} \subset \mathbb{R}/p\mathbb{Z} \cong \mathbb{T}.$$

(Note that the last isomorphism is unique if $\mathbb{R}/p\mathbb{Z}, \mathbb{T}$ are viewed as compact groups, up to inversion.)

We have found an explicit theory admitting quantifier-elimination, and true asymptotically in the family of finite fields with any nontrivial additive character. It continues to hold, of course, in the smaller class $P+$. So every sentence is equivalent to a quantifier-free one (in the language with Ψ_{sym}); to axiomatize $P+$ amounts to finding the valid quantifier-free sentences, describing, for each number field L , the possibilities for $F \cap L$ and for $\Psi|(F \cap L)$, where (F, Ψ) is an ultraproduct of elements of $P+$.

One group of additional axioms, the axioms SP below relating Ψ on \mathbb{Q} to roots of unity in F , is certainly needed. We conjecture that this gives a full axiomatization of the class, and show that this follows from a number theoretic conjecture, closely related to work of Duke, Friedlander, Iwaniec, and Toth. This is the counterpart for fields with an additive character of the use made by Ax of the Chebotarev density theorem.

We start with a description of the axiom scheme SP; it consists of axioms SP[1], SP[2], . . . , where the n th axiom asserts: SP[n]. For some $k < n$ prime to n we have:

- $\Psi\left(\frac{1}{n}\right) = \exp(2\pi i k/n)$; and
- $[F \cap \mathbb{Q}[\sqrt[n]{1}]] = \text{Fix}(\alpha_k)$,

α_k being the automorphism of $\mathbb{Q}[\sqrt[n]{1}]$ acting on the n th roots of 1 by k th power.

For instance, for $n \leq 4$ we have

$$\Psi(1) = 1, \quad \Psi\left(\frac{1}{2}\right) = -1, \quad \Psi\left(\frac{1}{3}\right) = e^{2\pi i/3}$$

iff F contains a primitive cube root of 1 and $4\pi i/3$ otherwise,

$$\begin{aligned} \Psi\left(\frac{1}{4}\right) &= i && \text{if } i \in F, \text{ and} \\ \Psi\left(\frac{1}{4}\right) &= -i && \text{otherwise.} \end{aligned}$$

Note that it follows from any SP[n] that $\Psi(1) = 1$, and hence $\Psi(\mathbb{Z}) = \{1\}$, i.e. Ψ induces a character of \mathbb{Q}/\mathbb{Z} .

Remark 5.1. The prime field axioms are nicer to formulate in the language of difference fields, restricted to the fixed field; this coincides with the language of pseudo-finite fields with the addition, for each n , of an algebraic imaginary constant naming (coherently) a generator σ_n of the Galois group of the order n extension. Now both the character group and the Galois group have a distinguished generator. In the standard models, the σ_n is interpreted as the Frobenius automorphism Fr ; in particular, on the n th roots of 1 (with n prime to p) the distinguished automorphism determines an integer $k(n)$, namely

the unique integer modulo n such that $\text{Fr}(\omega^k) = \omega$ for each n th root ω of 1. In this language the n th axiom asserts:

$$\text{If } \sigma_n(\omega) = \omega^k \text{ then } \Psi\left(\frac{-k}{n}\right) = \exp(2\pi i/n),$$

relating the choice of generator of the orientation module of n th roots of unity to the choice of generator of the order n Galois group, via the additive character ‘to resolution $1/n$ ’.

The truth of SP in the prime fields \mathbb{F}_p^+ with standard character, asymptotically as $p \rightarrow \infty$, is an agreeable exercise. For instance \mathbb{F}_p contains all n th roots of 1 iff $p = 1 \pmod n$ iff $-1/n$ is represented by $(p - 1)/n$ in \mathbb{F}_p . In this case we have

$$\psi(-1/n) = \exp(2\pi i/n \cdot (1 - 1/p)),$$

so the limit with large p is indeed $\exp(2\pi i/n)$. More generally, if $p = k \pmod n$ then $\Psi_p(-k/n)$ approaches $\exp(2\pi i/n)$ as $p \rightarrow \infty$, and the value of $\Psi(1/n)$ follows.

Conjecture 5.2. The characteristic zero asymptotic theory of prime fields with their standard additive character is axiomatized by $\text{PP}_+ := \text{PF}_+ + \text{SP}$.

Since we already have quantifier elimination, the conjecture must be fundamentally number-theoretic; we will now attempt to bring this out. Let \mathbb{Q}^{ab} denote the field generated by all roots of unity, over \mathbb{Q} . For any field F containing \mathbb{Q} , call a homomorphism $h: F \rightarrow \mathbb{T}$ *acceptable* if for each $n \in \mathbb{N}, n \geq 1, h(1/n)$ is a primitive n th root of unity. The condition is only on $h|_{\mathbb{Q}}$; it implies that h vanishes on \mathbb{Z} , and indeed that the kernel of h is precisely \mathbb{Z} . If h is acceptable, there exists a unique automorphism θ_h of \mathbb{Q}^{ab} such that for any n and any primitive n th root ω_n of 1, for some k prime to n , $\theta(\omega_n) = \omega_n^k$ where $h(-k/n) = \exp(2\pi i/n)$. When $F = \text{Fix}(\sigma)$, where $\sigma \in \text{Aut}(L), L$ a finite Galois extension of \mathbb{Q} , we say h, σ are *compatible* if h is acceptable, and there exists an automorphism of $L\mathbb{Q}^{ab}$ extending both σ and θ_h .

In case $L \cap \mathbb{Q}^{ab} = \mathbb{Q}[\omega]$ for some primitive m th root of unity ω , compatibility amounts to the equality $\theta_h(\omega) = \sigma(\omega)$.

Proposition 5.3. *The following statements are equivalent.*

- : Conjecture 5.2.
- ♠ : *Let L be a finite Galois extension of \mathbb{Q} . Let $\sigma \in \text{Aut}(L)$, and let $h: \text{Fix}(\sigma) \rightarrow \mathbb{T}$ be an acceptable homomorphism, compatible with σ . Then $(\text{Fix}(\sigma), +, \cdot, h)$ embeds into an ultraproduct of enriched prime fields from $P+$.*
- ♣ : *Let L be a finite Galois extension of \mathbb{Q} . Assume $L \cap \mathbb{Q}^{ab} = L \cap \mathbb{Q}[\omega]$, where ω is a primitive m th root of 1. Let $\sigma \in \text{Aut}(L), \sigma(\omega) = \omega^k$, and let f be an irreducible polynomial over \mathbb{Q} of degree d , with a root in $a \in L$ satisfying $\sigma(a) = a$. Then for any $\epsilon > 0$ and any $u_1, \dots, u_{d-1} \in \mathbb{T}$, there exist infinitely many primes $p \in \mathbb{N}$ with $p = k \pmod m$, and $b \in \mathbb{F}_p, f(b) = 0$, with $|\Psi_p(b^i) - u_i| < \epsilon, i = 1, \dots, d - 1$.*

Note that the statement $p = k \pmod m$ in the conclusion of ♣ could equally well be replaced by $|\Psi_p(-k/m) - u_0| < \epsilon$, where $u_0 = \exp(2\pi i/m)$.

Proof. We saw above that PP_+ holds asymptotically in the class P_+ .

Let us show ♣ follows from Conjecture 5.2. Let $L, \omega, \sigma, f, a, k$ be as in ♣. Let

$$h: \text{Fix}(\sigma) \rightarrow \mathbb{T}$$

be a homomorphism vanishing on \mathbb{Z} , with $h(a^i) = u_i$, and $h(-k/m) = \exp(2\pi i/m)$. Define an automorphism τ of \mathbb{Q}^{ab} by setting $\tau(\mu) = \mu^j$, where μ is a primitive n th root of unity, and $h(-j/n) = \exp(2\pi i/n)$; note τ agrees with σ on $\mathbb{Q}[\omega]$. Since L, \mathbb{Q}^{ab} are linearly disjoint over $\mathbb{Q}[\omega]$, one can find an automorphism $\bar{\sigma}$ of \mathbb{Q}^{alg} extending both σ and τ . Then,

$$(\text{Fix}(\bar{\sigma}), h) \models \text{SP}.$$

Let M be a model of PF_+ extending $(\text{Fix}(\bar{\sigma}), +, \cdot, h)$. By Conjecture 5.2, M embeds elementarily into an ultraproduct of standard enriched prime fields $(\mathbb{F}_p, +, \cdot, \Psi_p)$. Since in these fields $\Psi_p(-k/m)$ approaches $\exp(2\pi i/m)$, it follows that $p = k \pmod m$ for almost all of these p , proving ♣.

Now assume ♠. We will show that any model M of PP_+ is elementarily equivalent to an ultraproduct K of prime fields with their standard character. By Proposition 3.9, it will suffice if

$$(K \cap \mathbb{Q}^a, +, \cdot, \Psi) \cong (M \cap \mathbb{Q}^a, +, \cdot, \Psi).$$

Thus, we have to prove that for any $\sigma \in \text{Aut}(\mathbb{Q}^a)$ and any compatible homomorphism $h: \text{Fix}(\sigma) \rightarrow \mathbb{T}$, the structure $(\text{Fix}(\sigma), +, \cdot, h)$ is isomorphic to $K \cap \mathbb{Q}^a$ for some ultraproduct $K = (K, +, \cdot, \Psi)$ of prime fields with standard character. By compactness (of the class of ultraproducts of P_+), it suffices to find such an isomorphism on a prescribed number field L , Galois over \mathbb{Q} ; or even a finitely generated subring $R = \mathcal{O}_L[1/N]$ of L ; i.e. a field isomorphism

$$\alpha: L \cap \text{Fix}(\sigma) \rightarrow L \cap K,$$

with $\Psi \circ \alpha = h$ holding on R . Enlarging L if necessary, we may assume it contains a primitive N th root ω of 1.

We have reduced to a statement that is nearly provided by ♠; except that ♠ only speaks of an embedding, whereas we require the image of α to be all of $K \cap L$. To bridge this gap we employ the Chebotarev field crossing argument; see [13] for a full account, including the existence of the auxiliary primes l that appear below.

Let d' be the order of σ in $\text{Aut}(L)$. Let l be a prime such that $d' \mid (l - 1)$ and $\mathbb{Q}[\omega']$ is linearly disjoint from L over \mathbb{Q} , where ω' is a primitive l th root of 1. Let τ be an automorphism of $\mathbb{Q}[\omega']$ of order d' . Combine σ and τ to an automorphism $\bar{\sigma}$ of $L[\omega']$; then $\bar{\sigma}$ still has order d' . Let $h': \mathbb{Q} \rightarrow \mathbb{T}$ be compatible with $\bar{\sigma}$ (i.e. with some automorphism of \mathbb{Q}^a extending $\bar{\sigma}$). Since $\sigma(\omega) = \bar{\sigma}(\omega)$, and the pairs (h, σ) and $(h', \bar{\sigma})$ are both compatible, we have $h'|\mathbb{Z}[1/N] = h|\mathbb{Z}[1/N]$. Now,

$$\mathcal{O}_L[1/N] \cap \mathbb{Q} = \mathbb{Z}[1/N].$$

Thus, there exists a homomorphism $\mathbb{Q} + \mathcal{O}_L[1/N] \rightarrow \mathbb{T}$ extending both $h'|\mathbb{Z}[1/N]$ and $h|\mathcal{O}_L[1/N]$; extend it further to $\bar{h}: \text{Fix}(\rho) \rightarrow \mathbb{T}$. The pair $(\bar{h}, \bar{\sigma})$ remains compatible since this is tested on \mathbb{Q} .

By ♠ applied to $(\text{Fix}(\bar{\sigma}), \bar{h})$, there exists an ultraproduct (K, Ψ) of elements of $P+$ and an embedding

$$\alpha: (\text{Fix}(\bar{\sigma}), \bar{h}) \rightarrow (K, \Psi).$$

Restricting the ring R , we have $\bar{h}|R = h|R$, and so

$$\Psi \circ \alpha|R = \bar{h}|R = h|R.$$

It remains only to show that $K \cap L$ is no bigger than $\text{Fix}(\sigma)$. If $K \cap L > \text{Fix}(\sigma)$, since K contains $\text{Fix}(\bar{\sigma})$, it would follow that $K \cap \mathbb{Q}[\omega'] > \text{Fix}(\tau)$. However, as both K and $\text{Fix}(\bar{\sigma})$ satisfy SP, the fields $\mathbb{Q}[\omega'] \cap L', \mathbb{Q}[\omega'] \cap K$ are determined by $\bar{h}(1/l) = \Psi(1/l)$ and so must be equal. This contradiction shows that

$$K \cap L = \text{Fix}(\sigma),$$

as required.

To close the circle we prove ♠ from ♣. Let L, σ, h be as in ♠. Let a be a primitive generator for $\text{Fix}(\sigma)/\mathbb{Q}$, i.e.

$$\text{Fix}(\sigma) = \mathbb{Q}[a], \quad d = [\text{Fix}(\sigma) : \mathbb{Q}].$$

Then h is determined by $h(a), \dots, h(a^{d-1})$ and $h|\mathbb{Q}$. In turn, $h|\mathbb{Q}$ is determined by $h(1/N)$ for the various N . Of course, $h(1/N), h(1/N')$ are determined by the single value of $h(1/(NN'))$. Using compactness, it suffices to prove that for any $m \in \mathbb{N}$, the field $\text{Fix}(\sigma)$ embeds into an ultraproduct of enriched prime fields from $P+$, in such a way that $h(a^i) = \Psi(a^i)$ for $i = 1, \dots, d$, and $h(1/m) = \Psi(1/m)$.

We have $L \cap \mathbb{Q}^{ab} \subset \mathbb{Q}[\nu]$ for some root of unity ν ; we may assume m divides the order of ν , so that $L[\nu]$ contains a primitive m th root ω of 1. Now,

$$L[\nu] \cap \mathbb{Q}^{ab} = \mathbb{Q}[\nu].$$

Extend σ to $L[\nu]$, compatibly with h .

Let k be such that

$$h(-k/m) = \exp(2\pi i/m).$$

By SP (i.e. acceptability), upon possibly replacing σ by another automorphism of L with the same fixed field, we have $\sigma(\omega) = \omega^k$. Let f be the minimal polynomial of a over \mathbb{Z} . By ♣ for $L[\nu], \sigma, a$, there exist primes $p_i = k \pmod m$ and $b_i \in \mathbb{F}_{p_i}, f(b_i) = 0$, such that

$$|\Psi_p(b^i) - h(a^i)| < 1/k, \quad i = 1, \dots, d - 1.$$

Since $p_i = k \pmod m$, we also have

$$\Psi(-k/m) = \exp(2\pi i/m).$$

Let $(K, +, \cdot, \Psi, b)$ be a nonprincipal ultraproduct of the $(\mathbb{F}_{p_i}, +, \cdot, \Psi_{p_i})$, and embed $\text{Fix}(\sigma) = \mathbb{Q}[a]$ in K via $a \mapsto b$. Then h, Ψ agree on $-k/m$ and hence on $1/m$, and on each power a^i . This finishes the proof. ■

Remark 5.4. In case ♣ has counterexamples, the characteristic zero asymptotic theory of prime fields with their standard additive character is axiomatized by PP_+ along with the list of assertions

$$(\forall x) \left(f(x) = 0 \rightarrow \bigvee_{i=1}^{d-1} |\Psi(x^i) - u_i| \geq \epsilon \vee \left| \Psi\left(\frac{-k}{m}\right) - e^{2\pi i/m} \right| \geq \epsilon \right)$$

for each such counterexample.

5.5. Equidistribution

It is natural to expect not only existence, but also equidistribution of the primes in ♣. We will see that a considerably simpler (but still deeply conjectural!) one-variable equidistribution statement implies it. The required statement, Conjecture 5.7, appears to be quite close to a hypothesis raised in Duke, Friedlander and Iwaniec [11], that we recall below as a conjecture.

Conjecture 5.6 (DFI). Let $f \in \mathbb{Z}[X]$ be an irreducible polynomial of degree $d \geq 2$. Let $0 < \alpha < \beta \leq 1$. Let P be the set of primes. Consider,

$$S(x) := \{(p, v) : p \in P, p \leq x, v \in \mathbb{Z}, 0 \leq v < p, f(v) \equiv 0 \pmod{p}\}.$$

View $S(x)$ as a probability space, with the normalized counting measure. Then the probability that $\alpha \leq \frac{v}{p} < \beta$ approaches $\beta - \alpha$ as $|S(x)| \rightarrow \infty$.

The quadratic case was proved in [11] and [28]; the general case was qualified as ‘far away’ in [11], a view that to my outsider knowledge has not required much revision despite a quarter-century of progress.

We will need what seems to be a slight extension. Let L be a Galois extension of \mathbb{Q} , of degree $\delta = [L : \mathbb{Q}]$, with ring of integers \mathcal{O}_L . Let $P'(L, x)$ denote the set of maximal ideals \mathfrak{p} of \mathcal{O}_L , whose residue field $k(\mathfrak{p})$ has cardinality at most x . According to the prime ideal theorem, $P'(L, x)$ has size asymptotic to $x / \log(x)$. There are at most δ such ideals above a given rational prime p . If $k(\mathfrak{p})$ is not a prime field, then

$$p^2 \leq |k(\mathfrak{p})| \leq x;$$

so $p \leq \sqrt{x}$; there are thus at most $\delta\sqrt{x}$ such primes \mathfrak{p} , a negligible number *vis-a-vis* $|P'(L, x)|$; we will disregard them. We will thus be interested in the set $P(L, x)$ of primes $\mathfrak{p} \in P'(L, x)$ with residue field \mathbb{F}_p for some prime p . We will always let p denote the cardinality of the residue field of \mathfrak{p} .

Let a be an element of L . For all but finitely many primes $\mathfrak{p} \in P(L, \infty)$, we have $a \in \mathcal{O}_{\mathfrak{p}}$, and so $\text{res}_{\mathfrak{p}}(a) \in k(\mathfrak{p}) = \mathbb{F}_p$. Let Ψ_p be the standard additive character

$$n \bmod p \mapsto \exp(2\pi i n/p).$$

Let $\mu_{L,x}$ be the normalized counting measure on $P(L, x)$, and let $\mu[L, a; x]$ be the push-forward under the map

$$\mathfrak{p} \mapsto \Psi_p(\text{res}_{\mathfrak{p}}(a)).$$

Conjecture 5.7 (Extended DFI). Let L be a Galois extension of \mathbb{Q} , and $a \in L \setminus \mathbb{Q}$. Then the probability measures $\mu[L, a; x]$ converge weakly to the Haar measure on \mathbb{T} as $x \rightarrow \infty$.

One can also state the conjecture in the style of Conjecture 5.6: the fraction of $\mathfrak{p} \in P'(L, x)$ such that $\alpha \leq \text{res}_{\mathfrak{p}}(a) < \beta$ approaches $\beta - \alpha$ as $x \rightarrow \infty$.

Using Weyl’s criterion for equidistribution, as in [11] (3,4), a third equivalent form is the truth for each $h \in \mathbb{Z}, h \neq 0$ of the estimate:

$$\lim_{x \rightarrow \infty} \int_{\mu[L, a, x]} \Psi_p(\text{res}_{\mathfrak{p}}(a))^h dp = 0. \tag{1}$$

Note that Conjecture 5.7 considers the value of Ψ_p at a single field element a , whereas ♣ requires the values at a, a^2, \dots, a^{d-1} . Despite this we will see in Proposition 5.11 that Conjecture 5.7 does imply Conjecture 5.2.

Remark 5.8. The case $L = \mathbb{Q}(a)$ of Conjecture 5.7 is equivalent to Conjecture 5.6.

Proof. Assume $L = \mathbb{Q}(a)$. Let f be the minimal polynomial of a over \mathbb{Z} . Then, for some n ,

$$n\mathbb{Z}[a] \subset \mathcal{O}_L \subset \frac{1}{n}\mathbb{Z}[a].$$

Thus, putting asides the finitely many primes dividing n , a prime \mathfrak{p} of \mathcal{O}_L in $P'(L, x)$ corresponds to a prime $p \leq x$ of \mathbb{Z} along with a homomorphism $\mathbb{Z}[a] \rightarrow \mathbb{F}_p$; this homomorphism is determined by the image of a , i.e. by choice in \mathbb{F}_p of a root of f . Thus, $\mu[x]$ is also induced by the pairs (p, b) with $p \leq x$ and $b \in \mathbb{F}_p, f(b) = 0$. This is what is counted in Conjecture 5.6. ■

Going a little further, evidence that removing the condition $L = \mathbb{Q}(a)$ is not entirely unreasonable is given by the following remark, treating cyclotomic extensions of $\mathbb{Q}(a)$. We state it for quadratic a , but in the same way we can show that if Conjecture 5.6 holds true, then so does Conjecture 5.7 for the case $L = \mathbb{Q}[a, v]$ with v a root of unity.

Remark 5.9. Conjecture 5.7 is true for cyclotomic L and quadratic $a \in L$ (using [11,28]).

Proof. We may assume $L = \mathbb{Q}[\mu], \mu^N = 1$; we have $\sigma \in \text{Aut}(L), \sigma(a) = a, \sigma(\mu) = \mu^k$. Let $H = \text{Aut}(L/\mathbb{Q}[a])$; so H is an index-two subgroup of $\text{Aut}(L) = (\mathbb{Z}/N\mathbb{Z})^*$. So $k \in H$. We consider the map

$$\mathfrak{p} \mapsto (\Psi_p(a \bmod \mathfrak{p}), p \bmod N)$$

on primes \mathfrak{p} of $\mathbb{Q}[a]$ with prime residue field \mathbb{F}_p . Let $\mu[x]$ be the normalized counting measure on the set of primes \mathfrak{p} of $\mathbb{Q}[a]$ with norm $p \leq x$. We have to show that the image of this map becomes dense as $x \rightarrow \infty$; we will even prove equidistribution in $\mathbb{T} \times H$, using Weyl’s method. Thus, given a character z^n of \mathbb{T} and a character $\delta: H \rightarrow \mathbb{T}$, not both trivial, we must show that

$$\int \Psi_p(a)^n \delta(p) d\mu[x](\mathfrak{p}) = o(1).$$

In case $n = 0$, this does not involve Ψ_p , and follows from the Chebotarev density theorem. If $n \neq 0$, we have

$$\Psi_p(a)^n \delta(p) = \Psi_p(na + r)$$

for an appropriate rational number r , with $\Psi_p(r) = \delta(p)$. The results now follows from the quadratic case of Conjecture 5.6, a theorem of [11] and [28]. ■

Remark 5.10. In Proposition 5.3, we obtained Conjecture 5.2 from ♣ qualitatively. A similar deduction could be made for the equidistribution versions; as is the case with the usual use of the Chebotarev field crossing argument. The point is that the behavior of σ on l th roots of 1 for various $l = k \pmod m$ is statistically independent among the various l , and hence also statistically independent from a fixed event that we would like to measure, at least at the limit. From a stability point of view, this is the standard fact that a Morley sequence is asymptotically independent from any given element, specialized to measure algebras; see [4]. This allows relativizing to such an l without losing track of probabilities.

Proposition 5.11. *Assume Conjecture 5.7. Let L be a Galois extension of \mathbb{Q} and a an element of L with $d = [\mathbb{Q}[a] : \mathbb{Q}] > 1$. Let $\mu^d[L, a; x]$ be the pushforward of $\mu_{L,x}$ under the map*

$$p \mapsto \Psi_p(\text{res}_p(a), \text{res}_p(a^2), \dots, \text{res}_p(a^{d-1})).$$

Then $\mu^d[L, a; x]$ weakly approaches the Haar measure on \mathbb{T}^{d-1} as $x \rightarrow \infty$.

For any $\sigma \in \text{Aut}(L)$ with $\sigma(a) = a$ and $\sigma(\omega) = \omega^k$, ω an m th root of 1 in L , ♣ holds.

Proof. Again we use the Weyl equidistribution criterion, this time on \mathbb{T}^{d-1} . A character of \mathbb{T}^{d-1} has the form

$$(z_1, \dots, z_{d-1}) \mapsto z_1^{n_1} \cdot \dots \cdot z_{d-1}^{n_{d-1}}.$$

Thus, we have to prove that for any $(n_1, \dots, n_{d-1}) \in \mathbb{Z}^{d-1} \setminus (0)$, as in (1), we have

$$\lim_{x \rightarrow \infty} \int_{\mu[L,a,x]} \prod_{i=1}^{d-1} \Psi_p(\text{res}_p(a^i))^{n_i} = 0,$$

but

$$\prod_{i=1}^{d-1} \Psi_p(\text{res}_p(a^i))^{n_i} = \Psi_p \text{res}_p \sum_{i=1}^{d-1} n_i a^i;$$

as a has degree d , we have

$$\sum_{i=1}^{d-1} n_i a^i \notin \mathbb{Q};$$

hence the equation follows from (1), applied to the element $\sum_{i=1}^{d-1} n_i a^i$ in place of a .

If we are also given a primitive m th root of unity $\omega \in L$, and $\sigma \in \text{Aut}(L)$ with $\sigma(\omega) = \omega^k$ and $\sigma(a) = a$, let H be the subgroup of $\text{Aut}(\mathbb{Q}[\omega]/\mathbb{Q}) \cong (\mathbb{Z}/m\mathbb{Z})^*$ generated by σ . So the image of k in $(\mathbb{Z}/m\mathbb{Z})^*$ lies in H . Exactly as in Lemma 5.9, we obtain equidistribution in $\mathbb{T}^{d-1} \times H$. But this clearly implies ♣. ■

By way of illustration, consider the case where $f \in \mathbb{Z}[X]$ is monic, irreducible, and $\mathbb{Q}[X]/f$ is Galois. Then, for (all but finitely many) primes p , the number of roots of f in \mathbb{F}_p is 0 or d . On the face of it, Conjecture 5.6 is compatible with a picture where for each prime p where f has roots at all, either all d roots are represented by numbers in $\{0, \dots, (p-1)/2\}$, or all d roots are represented by numbers in $(p+1)/2, \dots, p-1$; as long as the number of primes choosing the first option is about equal to those choosing the second. By the Proposition, Conjecture 5.7 (at least) implies that this is not the case. Applying Ψ to the finite set of roots gives an element of $\mathbb{T}^d / \text{Sym}(d)$; by varying p among primes $\leq x$ and taking the limit $x \rightarrow \infty$, these elements equidistribute on the subset

$$\left\{ (z_1, \dots, z_d) \in \mathbb{T}^d : \prod_{i=1}^d z_i = 1 \right\} / \text{Sym}(d).$$

In particular, for many (most) p there will be roots in both lower and upper half-intervals.

We note also that [12] proves ♣ for the case $L = \mathbb{Q}[a]$, $k = 0$, $m = 1$, assuming the Bunyakovsky conjecture.

6. Some open-ended remarks

6.1. p-adic additive character

The map $\mathbb{Z}[1/p]/\mathbb{Z} \rightarrow \mathbb{T}$, $a \mapsto \exp((2\pi i)a)$, induces a homomorphism $\Psi_p: \mathbb{Q}_p \rightarrow \mathbb{T}$. Construed in discrete first-order logic, the theory of (\mathbb{Q}_p, Ψ_p) is undecidable for reasons similar to Proposition 2.2; pulling back appropriate arcs in \mathbb{T} , and rescaling, one can interpret long intervals $[1, \dots, a]$ in $\mathbb{Z}/p^n\mathbb{Z}$. Likewise, the asymptotic theory (over all \mathbb{Q}_p) is undecidable in discrete logic. However it is natural to expect that the continuous-logic presentation will be decidable, and with definable integration with respect to the p-adic measure, both for a single \mathbb{Q}_p and asymptotically. It is relevant that integration with an additive character reduces motivically to exponential sums over the finite residue field (Cluckers–Loeser, Hrushovski–Kazhdan). Developing this theory would be interesting.

6.2. Transformal geometry

The theory of fields with an automorphism σ has a model companion ACFA; see e.g. [7]. The fixed field of σ is denoted k . The theory of pseudo-finite fields is precisely the theory of k ; in this sense PF is contained in ACFA. The Frobenius difference field $K_p = (\mathbb{F}_p^a, \sigma_p)$ with $\sigma_p(x) = x^p$ is not a model of ACFA, but any nonprincipal ultraproduct K of such difference fields is ([16, 30]). Over any difference field, there is a notion of transformal dimension; it can be defined using transformal transcendence degree. Over K one can also characterize the transformal dimension of a definable set $D = \{x : \phi(x)\}$ as the unique n such that $D_p := \phi(K_p)$ has dimension n for almost all p . Note that D_p is a constructable set, since σ_p is algebraic.

Any variety over the fixed field k has transformal dimension 0; thus, the world of varieties over finite fields is represented within the difference varieties of transformal dimension 0. However, the latter category is considerably bigger. For one example, ultra-products of Suzuki groups, defined using a square root of Frobenius, live therein.

When D has transformal dimension 0, a finer dimension called the total dimension can be defined, and one has $|D_p| = O(p^m)$ iff D has total dimension $\leq m$. One can thus assign a measure μ_0 to any definable set of transformal dimension zero, namely the limit along the ultrafilter of $|D_p|/p^m$. This measure depends only on the theory of K (along with any parameters used in ϕ), and not on the specific presentation as an ultraproduct; it has the same definability properties as shown in [9] for pseudo-finite fields.

An example of a difference variety of transformal dimension 1 is given by the transformal curve $D_f: \sigma(y) - y = f(x)$, where x varies over an algebraic curve C , and f is a regular (algebraic or transformal) function on C . The specializations $D_{f,p}$ are algebraic curves; their smooth completions, were used by Weil ([31]) to bridge the gap between varieties and exponential sums. In a different way, not requiring smooth completions, the $D_{f,p}$ were used by Grothendieck and Deligne to express additive characters in terms of l -adic monodromy. Katz raised the question of a uniform treatment of such situations; Kazhdan and Kowalski suggested specifically that model theory may be useful; the model theory of difference geometry seems to be a very natural framework.

A survey of the model theory of ACFA, with relevant open questions, is planned. It is difficult to include a summary here of reasonable length, but some minimal observations seem to be called for.

6.3. ACFA with additive character on the fixed field

Consider a characteristic 0 model $K \models ACFA$, with constant field k , and expand k to a model of PF_+ , in the language considered above.¹

This simplest expansion requires no additional preparatory work: by the stable embeddedness of k in K , or directly from the nature of the quantifier-elimination of K , the new theory – let us temporarily name it $ACFA_+$ – admits quantifier-elimination. Nevertheless, it presents already some aspects worth noting.

(i) In any continuous-logic theory, and any sort S of that theory, one can define a compact analogue of the absolute Galois group along the lines of [22]. Namely, consider all \bigwedge -definable equivalence relations E on S with a bounded number of classes, independently of the model. There exists a smallest such E , called E_{KP} (recall we do not employ the convention of allowing arbitrary parameters). If M is a saturated model, the set of E -classes carries a natural compact topology, the logic topology, and the permutations induced from $\text{Aut}(M)$ become a compact topological group. This group coincides with Lascar's G_c , the compact part of the general Lascar group [26]; in simple theories, it is the full Lascar group. We will call it the Kim–Pillay group.

¹In the same language, we can extend the additive character to k^a . Namely, on a finite extension L of k , define $\psi(x) = \Psi\left(\frac{1}{[L:k]} \text{tr}_{L/k} x\right)$. We will not use this extension at present.

Let \mathcal{G} be the Kim–Pillay $ACFA_+$. By contrast with $ACFA$ (see Remark 3.21) for $ACFA_+$, this compact group has a nontrivial connected component of the identity, \mathcal{G}^0 . If B is any 0-definable torsor of the additive group $(k, +)$, one can define an equivalence relation E by:

$$xEy \text{ iff } \Psi(x - y) = 0.$$

Then B/E is a T -torsor. In case B determines a complete type over \emptyset , the automorphism group of B/E is T . For instance, if a is any 0-definable element of K , then

$$B_a = \{y : y - y^\sigma = a\}$$

is such a torsor T_a . In fact, \mathcal{G}^0 admits a homomorphism into T^m for each algebraic element of K with m conjugates. Of course, the image of \mathcal{G}^0 in $\prod_{a \in \text{acl}(0)} \text{Aut}(T_a)$ is not all of $T^{\text{acl}(0)}$, but reflects the additive relations among the elements a .

(ii) Conversely, I believe it can be shown that \mathcal{G}^0 is commutative, and may be precisely the above image. Let F be a difference field, relatively algebraically closed in K , and let F' be the base structure consisting of F and B/E as above, for all $\text{acl}(F)$ -definable k -torsors B . The key is to prove 3-amalgamation over F ; we are given, symbolically, the 2-types of each pair from a, b, c over F' , compatible on the 1-types and with each 2-type independent over F , and must find an independent 3-type extending them. As in Proposition 3.20, the purely algebraic amalgam is known to exist; and to define Ψ , it suffices to show that the known values of Ψ on $k(Fab), k(Fac), k(Fbc)$ where Fab denotes the relative algebraic closure of $F(a, b)$, are compatible with any additive relations among them. If X is an F^a -definable k -torsor, and $x \in X(Fa), y \in X(Fb), z \in X(Fc)$, then $x - y, x - z, y - z$ are such a triple of elements whose sum is zero. By methods similar to [8], it should be possible to show that all additive relations on elements of $k(Fab), k(Fac), k(Fbc)$ are generated in this way. These are taken into account in advance in the 1-types, since Ψ can be defined on the torsor X , with values in a torsor of T . (We may replace a by the k -internal part of a over F' . In case a, b, c lie in the same principal homogeneous space X for a k -definable group G , we have

$$F(k(Fab)) = F(d)$$

for d a generic of G , and likewise e for bc and f for ac ; if any unexpected additive relation holds, one can find $d' \in F(d), e' \in F(e), f' \in F(f)$ with $d' + e' + f' = 0$; then a standard argument shows yields a homomorphism from G to the additive group G_a , say with kernel N ; and the relation is already accounted for by the images of a, b, c in X/N . In general, a can be taken to lie in a translation variety in the sense of [8], and the situation requires some further analysis.)

(iii) Let f be a regular function on a curve C . We can also view C as a transformal curve (as such, it is denoted $C[\sigma]$ in [16]). The transformal curve D defined by

$$\sigma(y) - y = f(x)$$

is a transformal étale cover of C (smooth and zero-dimensional). Define D_Ψ as the quotient of $D \times \mathbb{T}$ by the identification of (d, t) with $(d + a, t\Psi(a))$; we obtain a circle bundle; if instead we take the quotient of $D \times \mathbb{C}$ by the same relation, we find an archimedean analogue of an l -adic local system over C . The structural automorphism σ lifts, given an element a of $C(k)$, to translation by $\Psi(f(a))$ on D_Ψ .

l -adic local systems are a central object of Grothendieck's theory leading towards the Weil conjectures; part of the reason for taking l -adic coefficients is their approximability by finite coefficients, where Grothendieck proved constructability theorems for cohomology. Deligne then had to devise methods to move over to archimedean fields, first in his solution of the Weil conjectures, but again in the case of his equidistribution theorem. As explained in [24], Deligne obtains equidistribution in terms of the natural measure on conjugacy classes in compact Lie groups by what feels like a Robinson-style model-theoretic transfer; the compact groups themselves do not really make a geometric appearance.

Could one dream of a direct archimedean analogue of Grothendieck's l -adic local systems, with his constructability theorems replaced by quantifier-elimination results? The fundamental group of a variety V is a quotient of the Shelah–Galois group of $ACF_{F(V)}$: it is the limit of $\text{Aut}(F(V')/F(V))$, where V'/V is finite and unramified. When one moves to simple theories, such as $ACFA_+$, the Kim–Pillay Galois group has quotients that may be viewed as a generalized fundamental group, no longer totally disconnected. In view of the probable abelian nature of the connected parts for $ACFA_+$ noted in (ii), and other reasons, it seems clear that $ACFA_+$ is not the full answer, and perhaps one needs to go beyond the Kim–Pillay part of Galois. In any case the possibility is tantalizing.

6.4. Difference geometry in transformal dimension one

Let us also take a quick look at how Weil's ideas in [31] may generalize. A notion of smooth transformal varieties exists; it may be defined in elementary terms using the usual Jacobian criterion applied to difference polynomials in place of polynomials, where differentiation treats any monomial X^σ as a constant. There is also a notion of transformal blowing-up, [16]. It is plausible that the projective completion of the curves $D_{f,\sigma}$ can be made smooth upon blowing up, at least for ordinary polynomials f . (The existence of a smooth difference curve with a given function field has not been investigated, but carries its own interest.) Assuming a smooth projective model $E = E_{f,\sigma}$ exists, a moving lemma on E^2 can be formulated, but has only been proved for the transformalization of smooth projective algebraic varieties; it plausibly follows from a transposition to difference geometry of the classical 'synthetic geometry' proof in [14]. One can now define an intersection product on $E \times E$. The coefficient ring is a 'motivic' ring constructed out of zero-dimensional difference varieties, with a dynamic 'preservation of number' principle built in. Whereas the Grothendieck ring of zero-dimensional schemes has roughly the complexity of the natural numbers (the number of points with multiplicities), zero-dimensional difference schemes have at least the complexity of zeta functions; in particular, one can apply Frobenius specialization for almost all p , to obtain a sequence of ordinary numbers. Using [16], this coefficient ring admits a natural homomorphism

into the field $\mathbb{R}(\epsilon)$, and so one can work with these coefficients when interested in ‘next to highest order’ estimates (see below). A Hodge index inequality is valid, though a purely transformal proof is not at present known to me; it follows using the main theorem of [16] from the case of ordinary algebraic varieties; presumably a direct proof is also possible. A Weil-style trace can be defined for correspondences on E using the intersection product with the diagonal; and exponential sums can be related to the trace of the structural automorphism σ .

6.5. Next-to-leading-order measure

Consider a definable family $(D_a : a \in P)$ of definable sets over a pseudo-finite field F . We take F with the induced structure as a fixed field of a model of ACFA. (It is very likely that the statements below are true generally for definable sets of transformal dimension zero over $K \models ACFA$.)

Using Weil’s Riemann hypothesis for curves, Will Johnson has shown [21] that counting modulo a prime l is definable. Equivalently, the function $a \mapsto |D_a|$ from P into \mathbb{Z} , if viewed as a function into the compact set \mathbb{Z}_l , the l -adic completion of \mathbb{Z} , is definable in the sense of continuous logic.²

For an archimedean analogue, it is necessary to renormalize since the image of \mathbb{Z} in \mathbb{R} is not relatively compact. We defined μ_0 by renormalizing by $p^{-\dim}$. The definability of μ_0 implies, in particular, that the relation

$$\mu_0(X) = \mu_0(X')$$

is definable on definable families.

If we consider the formal expression $[X] - [X']$, it can be viewed as a function whose μ_0 -integral is 0. (If X, X' are subsets of some ambient definable set D , we can represent $[X] - [X']$ by a function on D , the difference of characteristic functions $1_X - 1_{X'}$.) We can now go one step further, and consider the counting measure μ_1 normalized so that

$$\int (1_X - 1_{X'}) d\mu_1$$

is finite and (in general) nonzero; namely the ultralimit of

$$q^{\frac{1}{2}-\dim(X)} (|X| - |X'|).$$

Of course, it can be efficient to use a single invariant combining μ_0 and μ_1 ; simply let

$$\mu(D) = \mu_0(D) + \epsilon \mu_1(D).$$

²For $l \neq p$ also follows, though more elaborately, from Grothendieck’s l -adic theory, likewise dependent on Weil’s results; see [10, Theorems 4.4.10, and 7.1.1]; see also p. 31 for the explanation of the notion of orientation, corresponding to the richer-than-pure field language we take here. (See also §6.2 of [21].)

This takes values in the ordered field $\mathbb{R}(\epsilon)$, with $0 < \epsilon < 1/n$ for each n . (So, ϵ^2 'is' $\frac{1}{p}$.) But here we will consider them separately.

Now μ_0 is automatically definable with discreteness properties. μ_1 is not discrete, even in pseudo-finite fields, and cannot be expected to be definable in the pseudo-finite field itself. But in continuous logic, the measure μ_1 can be added to the structure; note that μ_1 is bounded in bounded families. One would like to know if μ_1 is definable in a tame geometric expansion of the theory in continuous logic; to begin with, as an expansion of PF . Definability of μ on families of curves (one-dimensional integration) would imply definability over all varieties.

Here is a precise question: let (F, μ_1) be an ultraproduct of finite fields with the $p^{-1/2}$ -normalized counting measure. Is it true that $\text{Th}(F, \mu_1)$ is simple as a continuous logic structure, and every definable subset of F^n is definable over the pseudo-finite field F alone?

Acknowledgments. Many thanks to Chieu-Minh Tran and to Jamshid Derakhshan for early discussions, and to Zoé Chatzidakis, Alexis Chevalier, Anand Pillay, Michał Szachniewicz, and two anonymous referees, for their comments.

References

- [1] J. Ax, The elementary theory of finite fields. *Ann. of Math. (2)* **88** (1968), 239–271
Zbl [0195.05701](#) MR [229613](#)
- [2] M. Bays and B. Hart, Pseudofinite-dimensional Schrödinger representations, preliminary notes, August 24, 2016
- [3] I. Ben-Yaacov, Positive model theory and compact abstract theories. *J. Math. Log.* **3** (2003), no. 1, 85–118 Zbl [1028.03034](#) MR [1978944](#)
- [4] I. Ben Yaacov, A. Berenstein, C. W. Henson, and A. Usvyatsov, Model theory for metric structures. In *Model theory with applications to algebra and analysis. Vol. 2*, pp. 315–427, London Math. Soc. Lecture Note Ser. 350, Cambridge Univ. Press, Cambridge, 2008
Zbl [1233.03045](#) MR [2436146](#)
- [5] C.-C. Chang and H. J. Keisler, *Continuous model theory*. Ann. of Math. Stud. 58, Princeton Univ. Press, Princeton, N.J., 1966 Zbl [0149.00402](#) MR [0231708](#)
- [6] C.-C. Chang and H. J. Keisler, *Model theory*. Third edn., Stud. Logic Found. Math. 73, North-Holland Publishing Co., Amsterdam, 1990 Zbl [0697.03022](#) MR [1059055](#)
- [7] Z. Chatzidakis and E. Hrushovski, Model theory of difference fields. *Trans. Amer. Math. Soc.* **351** (1999), no. 8, 2997–3071 Zbl [0922.03054](#) MR [1652269](#)
- [8] Z. Chatzidakis and E. Hrushovski, Difference fields and descent in algebraic dynamics. I. *J. Inst. Math. Jussieu* **7** (2008), no. 4, 653–686 Zbl [1165.03014](#) MR [2469450](#)
- [9] Z. Chatzidakis, L. van den Dries, and A. Macintyre, Definable sets over finite fields. *J. Reine Angew. Math.* **427** (1992), 107–135 Zbl [0759.11045](#) MR [1162433](#)
- [10] P. Deligne, *Cohomologie étale*. Lecture Notes in Math. 569, Springer, Berlin, 1977
Zbl [0345.00010](#) MR [463174](#)
- [11] W. Duke, J. B. Friedlander, and H. Iwaniec, Equidistribution of roots of a quadratic congruence to prime moduli. *Ann. of Math. (2)* **141** (1995), no. 2, 423–441 Zbl [0840.11003](#)
MR [1324141](#)

- [12] T. Foo, The Bouniakowsky conjecture and the density of polynomial roots to prime moduli. *Acta Arith.* **144** (2010), no. 1, 1–4 Zbl [1248.11069](#) MR [2660552](#)
- [13] M. D. Fried and M. Jarden, *Field arithmetic*. *Ergeb. Math. Grenzgeb.* (3) **11**, Springer, Berlin, 1986 Zbl [0625.12001](#) MR [868860](#)
- [14] W. L. Hoyt, On the moving lemma for rational equivalence. *J. Indian Math. Soc. (N.S.)* **35** (1971), 47–66 (1972) Zbl [0241.14004](#) MR [318146](#)
- [15] E. Hrushovski, Approximate equivalence relations, preprint
- [16] E. Hrushovski, The first-order theory of the Frobenius automorphisms, 2004 [arXiv:math/0406514](#)
- [17] E. Hrushovski, Pseudo-finite fields and related structures. In *Model theory and applications*, pp. 151–212, *Quad. Mat.* **11**, Aracne, Rome, 2002 Zbl [1082.03035](#) MR [2159717](#)
- [18] E. Hrushovski, Groupoids, imaginaries and internal covers. *Turkish J. Math.* **36** (2012), no. 2, 173–198 Zbl [1260.03065](#) MR [2912035](#)
- [19] E. Hrushovski, On pseudo-finite dimensions. *Notre Dame J. Form. Log.* **54** (2013), no. 3–4, 463–495 Zbl [1345.03059](#) MR [3091666](#)
- [20] A. E. Ingham, On the difference between consecutive primes. *Quart. J. Math. (Oxford Ser.)* **8** (1937), 255–266 Zbl [63.0903.04](#)
- [21] W. Johnson, Pseudo-finite counting functions mod n , preprint
- [22] B. Kim and A. Pillay, Simple theories. Joint AILA-KGS Model Theory Meeting (Florence, 1995). *Ann. Pure Appl. Logic* **88** (1997), no. 2–3, 149–164 Zbl [0897.03036](#) MR [1600895](#)
- [23] E. Kowalski, Exponential sums over definable subsets of finite fields. *Israel J. Math.* **160** (2007), 219–251 Zbl [1146.03020](#) MR [2342497](#)
- [24] E. Kowalski, Some aspects and applications of the Riemann hypothesis over finite fields. *Milan J. Math.* **78** (2010), no. 1, 179–220 Zbl [1271.11113](#) MR [2684778](#)
- [25] E. Kowalski, A survey of algebraic exponential sums and some applications. In *Motivic integration and its interactions with model theory and non-Archimedean geometry. Volume II*, pp. 178–201, *London Math. Soc. Lecture Note Ser.* **384**, Cambridge Univ. Press, Cambridge, 2011 Zbl [1276.11136](#) MR [2885344](#)
- [26] D. Lascar, On the category of models of a complete theory. *J. Symbolic Logic* **47** (1982), no. 2, 249–266 Zbl [0498.03019](#) MR [654786](#)
- [27] I. Tomašić, Exponential sums in pseudofinite fields and applications. *Illinois J. Math.* **48** (2004), no. 4, 1235–1257 Zbl [1069.03028](#) MR [2114155](#)
- [28] Á. Tóth, Roots of quadratic congruences. *Internat. Math. Res. Notices* (2000), no. 14, 719–739 Zbl [1134.11339](#) MR [1776618](#)
- [29] C.-M. Tran, Tame structures via character sums over finite fields, 2019 [arXiv:1704.03853](#)
- [30] Y. Varshavsky, Intersection of a correspondence with a graph of Frobenius. *J. Algebraic Geom.* **27** (2018), no. 1, 1–20 Zbl [1388.14034](#) MR [3722688](#)
- [31] A. Weil, On some exponential sums. *Proc. Nat. Acad. Sci. U.S.A.* **34** (1948), 204–207 Zbl [0032.26102](#) MR [27006](#)
- [32] B. Zilber, The semantics of the canonical commutation relation, 2016 [arXiv:1604.07745](#)

Received 22 April 2021.

Ehud Hrushovski

Mathematical Institute, University of Oxford, Woodstock Road, Oxford OX2 6GG, UK;
hrushovski@maths.ox.ac.uk