Notes on the model theory of finite and pseudo-finite fields

Zoé Chatzidakis - (CNRS (DMA) - ENS)

July 25, 2018

These notes arose from a 12 hour course given during the IMS 2018 Graduate Summer School in Logic, which was held at the National University of Singapore (18 June - 6 July 2018). They are slightly rewritten, and a few details and remarks have been added (in particular in the last section). The order in which results appear is not necessarily the order in which they were presented.

I wish to thank the Institute of Mathematical Sciences and the Department of Mathematics of the National University of Singapore for their generous support and for giving me the opportunity to teach in the summer school. I also thank the audience for their enthusiasm and comments.

Table of contents

1 Finite fields - properties

1.1. Basic properties of finite fields

The characteristic of a unitary commutative ring is the smallest positive integer n such that $1+1+\cdots+1$ (n times) equals 0. If there is no such integer n, one says that the characteristic is 0. If R is a finite ring, and a fortiori a finite field, then its characteristic is finite. Hence, if F is a finite field, the homomorphism $\mathbb{Z} \to F$ which sends $1 \in \mathbb{Z}$ to $1 \in F$ must have kernel a prime ideal, i.e., $p\mathbb{Z}$ for some prime p.

Conversely, if p is a prime number, then $p\mathbb{Z}$ is a maximal ideal of $\mathbb{Z}$ and $\mathbb{Z}/p\mathbb{Z}$ is a field with p elements. This field is denoted by $\mathbb{F}_p$, and it is the *prime field of characteristic* p, i.e., it is contained in every field of characteristic p (by the above). In a field of characteristic 0, the subring generated by 1 is (isomorphic to) $\mathbb{Z}$, and therefore the field also contains the field of fractions of $\mathbb{Z}$, $\mathbb{Q}$. We call $\mathbb{Q}$ the *prime field of characteristic* 0.

Let F be a finite field of characteristic $p > 0$. Since $1 \in F$, it necessarily contains the field $\mathbb{F}_p$, and is therefore a vector space over $\mathbb{F}_p$, whence of cardinality p^n for some $n \in \mathbb{N}$.

Let F be a field of characteristic p having $q = p^n$ elements, let K be an algebraically closed field containing F. Let us consider the multiplicative group $F^\times = F \setminus \{0\}$ of F. It has $q-1$ elements and hence every non-zero element of F satisfies the equation $X^{q-1} - 1 = 0$. [If G is a finite group of size n, then every element g of G satisfies $g^n = 1$]. Thus all elements of F satisfy $X^q - X = 0$. Let $f(X) = X^q - X$, a polynomial over $\mathbb{F}_p$. Then $f'(X) = qX^{q-1} - 1 = -1$ because "$q = 0$" since it is a power of the characteristic. Hence all roots of $f(X) = 0$ are simple roots, and we obtain

$$X^q - X = \prod_{a \in F} (X - a).$$

Indeed, since every element of $\mathbb{F}_q$ satisfies $X^q - X = 0$, we know that each $(X - a)$, $a \in \mathbb{F}_q$, divides $X^q - X$, and therefore so does their product $\prod_{a \in F}(X - a)$. Degree considerations and the fact that the coefficient of X^q is 1 imply that these two polynomials are equal.

Conversely, let us consider the set $S \subset K$ of all solutions of $X^q - X = 0$. As above, its roots are all distinct. S is closed under multiplication, and $S \setminus \{0\}$ by multiplicative inverse. Because we are in characteristic p and q is a power of p, we obtain, using the binomial expansion of $(a+b)^n$ and the fact that "$p = 0$", that $(a+b)^p = a^p + b^p$, and $(a+b)^q = a^q + b^q$. This implies that S is closed under addition, and is therefore a subfield of K. It has cardinality q, and is denoted $\mathbb{F}_q$.

So, we have shown:

Theorem 1.2. *Let F be a finite field. Then for some prime p and $q = p^n$, F has q elements. Its elements are exactly the roots of the equation $X^q - X = 0$.*

1.3. The Frobenius map. We have noticed above that when F is a field of characteristic p, then $(a+b)^p = a^p + b^p$ for $a, b \in F$. The map $x \mapsto x^p$ is therefore a ring morphism (as it obviously is a multiplicative map). Also, as $x^p = 0$ implies $x = 0$, it is injective. The map $x \mapsto x^p$ is called the *Frobenius map*, and I will denote it by Frob_p, or Frob. Similarly, if $q = p^n$, then I will denote Frob^n also by Frob_q.

1.4. The multiplicative group of a finite field. Let $F = \mathbb{F}_q$ be a finite field. We will show that $F^\times$ is cyclic. As it is finite, it can be written as a finite direct sum of cyclic subgroups, and if it is not cyclic, then its exponent[1] m is a proper divisor of $q-1$. But all roots of $X^{q-1} = 1$ are simple roots, whence all roots of $X^m = 1$ are simple as well. This implies that $q - 1 = m$.

1.5. Perfect fields. A field F of characteristic $p > 0$ is *perfect* if every element of F has a p-th root. By convention, every field of characteristic 0 is perfect.

If $F = \mathbb{F}_{p^n}$ is finite, then the order of $F^\times$ is prime to p, which implies that every element is (multiplicatively) divisible by p, i.e., F is perfect. Another way of seeing this is the fact that the map Frob : $x \mapsto x^p$ is injective, because $x^p = 1$ implies $x = 1$: as F is finite, Frob must be onto.

An example of imperfect field is $\mathbb{F}_p(t)$, where t is transcendental over $\mathbb{F}_p$. Then the image by Frob of $\mathbb{F}_p(t)$ is $\mathbb{F}_p(t^p)$.

1.6. The algebraic closure of $\mathbb{F}_p$.

Let m, n be positive integers, p a prime. Then

$$\mathbb{F}_{p^m} \subseteq \mathbb{F}_{p^n} \iff m \text{ divides } n,$$

and in that case we have $[\mathbb{F}_{p^n} : \mathbb{F}_{p^m}] = n/m$.

Indeed, if $\mathbb{F}_{p^m} \subseteq \mathbb{F}_{p^n}$ then $\mathbb{F}_{p^n}$ is in particular an $\mathbb{F}_{p^m}$-vector space, which implies that for some ℓ, $|\mathbb{F}_{p^n}| = |\mathbb{F}_{p^m}|^\ell$, i.e., $p^n = p^{m\ell}$ and $n = m\ell$. We then have $[\mathbb{F}_{p^n} : \mathbb{F}_{p^m}] = \ell$. Conversely, if m divides n, then $p^m - 1$ divides $p^n - 1$, whence all roots of $X^{p^m-1} = 1$ are contained in $\mathbb{F}_{p^n}$, i.e., $\mathbb{F}_{p^m} \subseteq \mathbb{F}_{p^n}$.

It follows easily that for any $m, n \geq 1$,

$$\mathbb{F}_{p^m} \cap \mathbb{F}_{p^n} = \mathbb{F}_{p^d} \text{ and } \mathbb{F}_{p^m}\mathbb{F}_{p^n}$$

where d is the greatest common divisor of m and n, and e is the least common multiple of m and n. Here, $\mathbb{F}_{p^m}\mathbb{F}_{p^n}$ denotes the field composite of $\mathbb{F}_{p^m}$ and $\mathbb{F}_{p^n}$, i.e., the subfield (of the large algebraically closed field K) they generate.

Let α be algebraic over $\mathbb{F}_p$. Then $\mathbb{F}_p(\alpha)$ is a finite-dimensional $\mathbb{F}_p$-vector space, and is therefore also finite. This implies that the algebraic closure $\mathbb{F}_p^{alg}$ of $\mathbb{F}_p$ is $\bigcup_{n\in\mathbb{N}} \mathbb{F}_{p^n}$.

1.7. More on the Frobenius map. Fix a prime p. The Frobenius map is the identity on $\mathbb{F}_p$ (since every element of $\mathbb{F}_p$ satisfies $X^p - X = 0$), and defines an automorphism of each $\mathbb{F}_{p^n}$. Hence it defines an element φ of $\mathrm{Aut}(\mathbb{F}_p^{alg}/\mathbb{F}_p)$. Observe that if $d \in \mathbb{N}$, the elements of $\mathbb{F}_p^{alg}$ which are fixed by φ^d are precisely the elements of $\mathbb{F}_{p^d}$. Furthermore, one checks that the restriction $\varphi|_{\mathbb{F}_{p^d}}$ of φ to $\mathbb{F}_{p^d}$ has order exactly d: φ^ℓ being the identity on $\mathbb{F}_{p^d}$ means exactly that all elements of $\mathbb{F}_{p^d}$ satisfy $X^{p^\ell} = X$, and therefore that d divides ℓ.

[1] The exponent of a group G is the smallest $n > 0$ such that every element $g \in G$ satisfies $g^n = 1$, and ∞ if such an n doesn't exist.

As $[\mathbb{F}_{p^d} : \mathbb{F}_p] = d$, we know that $\mathrm{Aut}(\mathbb{F}_{p^d}/\mathbb{F}_p)$ has size at most d. Since $\varphi \in \mathrm{Aut}(\mathbb{F}_{p^d}/\mathbb{F}_p)$ has order exactly d, this therefore implies that

$$\mathrm{Aut}(\mathbb{F}_{p^d}/\mathbb{F}_p) \simeq \mathbb{Z}/d\mathbb{Z},$$

and that φ generates $\mathrm{Aut}(\mathbb{F}_{p^d}/\mathbb{F}_p)$, and $\mathbb{F}_{p^d}$ is a Galois extension of $\mathbb{F}_p$.

1.8. Description of $\mathrm{Aut}(\mathbb{F}_p^{alg}/\mathbb{F}_p)$. While we will not explicitly use it, we can now describe completely $\mathrm{Aut}(\mathbb{F}_p^{alg}/\mathbb{F}_p)$. As $\mathbb{F}_p^{alg}$ is a direct limit of the finite fields $\mathbb{F}_{p^n}$, it follows, by Galois duality, that

$$\mathrm{Aut}(\mathbb{F}_p^{alg}/\mathbb{F}_p) = \lim_{\leftarrow} \mathbb{Z}/n\mathbb{Z} := \hat{\mathbb{Z}}.$$

The connecting maps are, for n dividing m, the canonical projection $\mathbb{Z}/m\mathbb{Z} \to \mathbb{Z}/n\mathbb{Z}$. That is, the group $\hat{\mathbb{Z}}$ is described as the set of sequences $(a_n)_n \in \prod_{n>1} \mathbb{Z}/n\mathbb{Z}$ such that if n divides m, then $a_n \equiv a_m \mod n$. It is a profinite group, i.e., an inverse limit of finite groups. It is a closed subgroup of $\prod_{n>1} \mathbb{Z}/n\mathbb{Z}$, where each $\mathbb{Z}/n\mathbb{Z}$ is equipped with the discrete topology, and we take the product topology on $\prod_{n>1} \mathbb{Z}/n\mathbb{Z}$. The element $\varphi = \mathrm{Frob}_p$ is a *topological generator* of $\mathrm{Aut}(\mathbb{F}_p^{alg}/\mathbb{F}_p)$: its restriction to any $\mathbb{F}_q$ generates $\mathrm{Aut}(\mathbb{F}_q/\mathbb{F}_p)$.

1.9. An aside: a result of Ax. The fact that $\mathbb{F}_p^{alg}$ is a union of finite fields, has a very nice consequence: Let $\bar{X} = (X_1, \ldots, X_n)$ and $\bar{f}(\bar{X})$ be an n-tuple of polynomials in $\mathbb{F}_p^{alg}[\bar{X}]$. Assume that $\bar{f}(\bar{X})$ defines an injective map $\bar{x} \mapsto \bar{f}(\bar{x})$ from $(\mathbb{F}_p^{alg})^n$ to itself. Then the map $\bar{f}$ is also surjective.

Proof. Indeed, the elements of $\bar{f}$ have their coefficients in some $\mathbb{F}_q$, and therefore the restriction of $\bar{f}$ to $\mathbb{F}_q^n$ is also injective; as $\mathbb{F}_q$ is finite, $\bar{f}|_{\mathbb{F}_q^n}$ is surjective. This being true on all finite fields containing $\mathbb{F}_q$, we obtain the result.

Theorem 1.10. *(Ax, Thm C in [1]) Let $\bar{f}(\bar{X})$ be an n-tuple of polynomials in $\mathbb{C}[\bar{X}]$, $\bar{X} = (X_1, \ldots, X_n)$, and assume that the map $\bar{f}$ it defines $\mathbb{C}^n \to \mathbb{C}^n$ is injective. Then it is also surjective.*

There are two proofs of this result, which I present below. They are essentially equivalent, but one of them uses ultraproducts.

Proof 1 of Ax's result 1.10. One uses the fact that the completions of the theory ACF of algebraically closed fields[2] is obtained by specifying the characteristic. Thus the theory of algebraically closed fields of characteristic 0 is obtained by adding to ACF an infinite set of sentences: for each prime p, an axiom saying that "$p \neq 0$". Any statement true in all (or some) algebraically closed fields of characteristic 0 must therefore be true in all algebraically closed fields of sufficiently large characteristic. Let $m_i(\bar{X})$, $i = 1, \ldots, N(d)$, be an enumeration of the

[2] The theory ACF is axiomatised by adding to the theory of fields for every $n \geq 1$ the axiom expressing that every polynomial of degree exactly n has a solution: $\forall y_0, \ldots, y_{n-1}\ \exists x\ x^n + \sum_{i=0}^{n-1} y_i x^i = 0$.

monomials in $\bar{X} = (X_1, \dots, X_n)$ of degree $\leq d$, and consider the formulas $\varphi(\bar{x})$, $\psi(\bar{x})$, where $\bar{x} = (x_{i,j})_{1\leq i\leq n, 1\leq j\leq N(n)}$, and $\bar{y} = (y_1, \dots, y_n)$, $\bar{z} = (z_1, \dots, z_n)$:

$$\varphi(\bar{x}) : \ \forall \bar{y}, \bar{z} \ \big(\bigwedge_i \sum_j x_{i,j} m_i(\bar{y}) = \sum_i x_{i,j} m_i(\bar{z})\big) \to (\bar{y} = \bar{z})$$

$$\psi(\bar{x}) : \ \forall \bar{z} \exists \bar{y} \ \bigwedge_i \sum_j x_{i,j} m_i(\bar{y}) = z_i.$$

Thus $\varphi(\bar{x})$ says that the map defined by the n-tuple $\bar{f}(\bar{X})$ of polynomials, with $f_i(\bar{X}) = \sum_j x_{i,j} m_j(\bar{X})$, $i = 1, \dots, n$, is injective, while $\psi(\bar{x})$ says that $\bar{f}$ is surjective.

All algebraically closed fields of positive characteristic satisfy $\forall \bar{x} \ \varphi(\bar{x}) \to \psi(\bar{x})$, hence also $\mathbb{C}$ satisfies this sentence. This proves the theorem.

Proof 2. The second proof uses ultraproducts: if $\mathcal{U}$ is a non-principal ultrafilter on the set $\mathbb{P}$ of prime numbers then

$$\mathbb{C} \simeq \prod_{p\in\mathbb{P}} \mathbb{F}_p^{alg}/\mathcal{U}.$$

As every field $\mathbb{F}_p^{alg}$ satisfies the sentence $\forall \bar{x} \ \varphi(\bar{x}) \to \psi(\bar{x})$, and hence so does $\mathbb{C}$.

2 Axiomatisation of a candidate for the theory of finite fields

The model theory of finite fields and pseudo-finite fields was started by J. Ax in [1]. Many of the results below were proved by him, in particular the description of the theory of finite fields, and its decidability. Additional results were shown by M. Jarden and U. Kiehne, and by C. Kiefe. Precise attributions can be found in the book of Fried and Jarden [13]. The purpose of this section is to give the axiomatisations of two theories, the first one coinciding with the infinite models of the second one. Let me temporarily denote them by Psf* and T_f^*. We will show that all finite fields are models of T_f^*. The infinite models of the theory of all finite fields will be called *pseudo-finite* fields, and we will show that they are models of Psf*. We first introduce Psf*.

2.1. Consider the theory Psf* obtained by adding to the theory of fields the following axiom schemes:

– Axiom 1 saying that the field is perfect,

– Axiom $2(\ell)$ saying that the field has exactly one algebraic extension of degree ℓ, for every $\ell > 1$,

– Axiom $3(m, n, d)$ a scheme of axioms expressing that the field is pseudo-algebraically closed (abbreviated by PAC), see definition below 2.9 (for every $m, n, d \in \mathbb{N}$).

2.2. Axiom 1. This one is easy: for each prime p, add the axiom

$$p = 0 \to \forall y \exists x \ y = x^p.$$

Here "$p = 0$" is the sentence $\underbrace{1+1+\cdots+1}_{p \text{ times}} = 0$.

2.3. First half of Axiom 2(ℓ). Fix ℓ, we will first define a formula $\mathrm{Irr}_\ell(\bar{y})$, where $\bar{y} = (y_0, \dots, y_{\ell-1})$, which says that the polynomial $P_\ell(\bar{y})(X) := X^\ell + y_{\ell-1}X^{\ell-1} + y_{\ell-2}X^{\ell-2} + \cdots + y_0$ is irreducible, i.e., is not the product of two polynomials of lower (non-zero) degree. The formula $\mathrm{Irr}_\ell(\bar{y})$ expresses that $\forall z_0, \dots, z_{\ell-1}$, for all $1 \leq d < \ell$ the polynomials $X^\ell + y_{\ell-1}X^{\ell-1} + y_{\ell-2}X^{\ell-2} + \cdots + y_0$ and $(X^d + z_{d-1}X^{d-1} + \cdots + z_0)(X^{\ell-d} + z_{\ell-1}X^{\ell-d-1} + z_{\ell-2}X^{\ell-d-2} + \cdots + z_d)$ are not equal.

I.e., $\mathrm{Irr}_\ell(\bar{y})$ is the disjunction over $j = 0, \dots, \ell-1$, of the formulas

$$y_j \neq \sum_{i=0}^{d-1} z_i z'_{j-i}$$

where $z'_m = z_{\ell-d+m}$ if $0 \leq m < \ell - d$, $z'_m = 1$ if $m = \ell - d$, and $z'_m = 0$ otherwise.

So the first half of axiom 2(ℓ) will say $\exists \bar{y}\ \mathrm{Irr}_\ell(\bar{y})$. A field F which satisfies this axiom will therefore have an algebraic extension of degree ℓ.

2.4. Second half of Axiom 2(ℓ). To finish the axiomatisation of 2(ℓ), we need to say that this extension is unique. Equivalently, that if $P(X)$ and $Q(X)$ are irreducible polynomials of degree ℓ, then the extension of F generated by a root of $P(X)$ contains all roots of $Q(X)$.

In order to do that, we first need to show that we can interpret, uniformly in the ℓ-tuple $\bar{y}$ (which satisfies Irr_ℓ in the field F) the extension generated over F by a root of the polynomial $P_\ell(\bar{y})(X)$.

2.5. Interpretation of a finite algebraic extension of a field inside the field. Let $\bar{a} = (a_0, \dots, a_{\ell-1})$ be an ℓ-tuple satisfying Irr_ℓ in the field F. Let α be a root of $P_\ell(\bar{a})(X)$, and recall that

$$F(\alpha) \simeq_F F[X]/(P_\ell(\bar{a})(X)).$$

In particular $F(\alpha)$ is an F-vector space of dimension ℓ, with basis $\{1, \alpha, \alpha^2, \dots, \alpha^{\ell-1}\}$. This remark allows us to interpret easily, inside F and uniformly in the ℓ-tuple $\bar{a}$, the structure $(F(\alpha), +, \times, 0, 1, P_F)$, where $+, \times, 0, 1$ are the usual addition, multiplication and constants on the field $F(\alpha)$, and P_F is a unary predicate for the subfield F.

We let $S = F^\ell$ (the direct sum of ℓ copies of F), $+^*$ the usual addition on the vector space S, and $0^* = (0, 0, \dots, 0)$, $1^* = (1, 0, \dots, 0)$, P^*_F the set of elements $\{(b, 0, \dots, 0) \mid b \in F\}$. Clearly these sets, elements and relations are definable in F, with no parameters.

Multiplication by α induces a linear transformation of the vector space $F(\alpha)$, and its matrix is

$$M_\alpha = \begin{pmatrix} 0 & 0 & \cdots & 0 & -a_0 \\ 1 & 0 & \cdots & 0 & -a_1 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & 1 & -a_{\ell-1} \end{pmatrix}$$

since $\alpha^\ell = -\sum_{i=0}^{\ell-1} a_i\alpha^i$. Note that multiplication by α^i is also a linear transformation, and its matrix is simply M_α^i. So, we define $\times^*$ as follows

$$(x_0, \dots, x_{\ell-1}) \times^* (y_0, \dots, y_{\ell-1}) = (x_0 I_\ell + x_1 M_\alpha + \cdots x_{\ell-1} M_\alpha^{\ell-1}) \begin{pmatrix} y_0 \\ y_1 \\ \vdots \\ y_{\ell-1} \end{pmatrix}.$$

Here I_ℓ denotes the identity $(\ell \times \ell)$-matrix. Observe that the definition of $\times^*$ uses the tuple $(a_0, \dots, a_{\ell-1})$, but is totally uniform.

Hence, there is a formula $\theta^*(\bar{x}, \bar{y})$ of the language of fields, such that if $\mathrm{Irr}_\ell(\bar{a})$ and $\mathrm{Irr}_\ell(\bar{b})$ hold for some ℓ-tuples $\bar{a}$ and $\bar{b}$ in F and α is a root of $P_\ell(\bar{a})(X)$, then $F \models \theta^*(\bar{a}, \bar{b})$ if and only if $F(\alpha)$ contains all ℓ roots of $P_\ell(\bar{b})(x)$. So, axiom 2 (ℓ) is:

$$\exists \bar{x}\ \mathrm{Irr}_\ell(\bar{x}) \wedge \forall \bar{y}\, [\mathrm{Irr}_\ell(\bar{y}) \to \theta^*(\bar{x}, \bar{y})].$$

2.6. A comment on this condition. The condition of having at most one extension of each degree is equivalent to the following: whenever L is an algebraic extension of F of degree n, then $\mathrm{Aut}(L/F) \simeq \mathbb{Z}/n\mathbb{Z}$. In particular, $\mathrm{Aut}(L/F)$ is abelian and cyclic. The proof is left as exercise. (It uses some very basic facts on Galois theory).

2.7. Algebraic sets, varieties .... Let F be a **perfect** field, Ω a large algebraically closed field containing it, and F^{alg} the algebraic closure of F (inside Ω). Given an n-tuple $\bar{a}$ in Ω, we look at

$$I(\bar{a}/F) = \{f(\bar{X}) \in F[\bar{X}] \mid f(\bar{a}) = 0\}.$$

We then have the following result:

$$I(\bar{a}/F)\Omega[\bar{X}] \text{ is prime} \iff F(\bar{a}) \cap F^{alg} = F.$$

Here $\bar{X} = (X_1, \dots, X_n)$, $I(\bar{a}/F)\Omega[\bar{X}]$ denotes the ideal generated by $I(\bar{a}/F)$ inside $\Omega[\bar{X}]$. If $I(\bar{a}/F)$ satisfies one of these equivalent conditions, then we say it is *absolutely prime.* This terminology also applies to any prime ideal of $F[\bar{X}]$ which generates a prime ideal in $\Omega[\bar{X}]$[3].

An algebraic subset of Ω^n is the set of solutions of some (finite) set of polynomial equations with coefficients in Ω. The algebraic sets are the closed subsets of a topology on Ω^n, the *Zariski topology.* This topology is Noetherian, and therefore every closed set is the union of finitely many *irreducible closed subsets*[4]. To an algebraic set S (or to any subset of Ω^n) we can associate $I(S)$, the set of polynomials in $\Omega[\bar{X}]$ which vanish at all points of S. An algebraic set S is a *variety* iff the ideal $I(S)$ is prime, iff it is closed and irreducible. It will be *defined over* F if $I(S)$ is generated by its intersection with $F[\bar{X}]$. A point of the algebraic set S is F*-rational* if all its coordinates are in F, and the set of F-rational points is denoted $S(F)$.

[3] Warning – this equivalence uses the fact that we are over a perfect field

[4] A closed set U is irreducible if whenever $U = U_1 \cup U_2$ with U_1, U_2 closed, then $U_1 = U$ or $U_2 = U$

2.8. Coordinate rings, function fields. Let $I \subset F[\bar{X}]$ be a *radical* ideal (if $f^m \in I$ for some $m \in \mathbb{N}$, then $f \in I$), and V the algebraic set defined by the equations $f(\bar{x}) = 0$, $f \in I$ (as I is finitely generated, a finite set of such equations suffices). The *coordinate ring* of the algebraic set V is $F[V] := F[\bar{X}]/I$. If V is F-irreducible, then one defines the *function field* of V to be the field of fractions of $F[V]$. The *dimension* of an F-irredusible set V is the transcendence degree $tr.deg(F(V)/F)$, and the dimension of an algebraic set is the maximal dimension of a component.

2.9. Pseudo-algebraically closed fields. A field F is *pseudo-algebraically closed* (abbreviated by *PAC*) if every variety V defined over F has an F-rational point.

Before showing that being PAC is an elementary property, we will show an easy property of PAC fields:

Lemma 2.10. *Let F be a perfect PAC field, L a field containing F, and assume that $L \cap F^{alg} = F$. Then F is existentially closed in L, denoted $F \prec_1 L$, i.e.: every existential formula with parameters in F which is true in L is true in F.*

Proof. An existential formula of $\mathcal{L}(F)$ is of the form $\exists\ \bar{y}\varphi(\bar{y})$, where $\varphi(\bar{y})$ is a boolean combination of polynomial equations with coefficients in F. Hence a disjunct of conjuncts of polynomial equations and inequations (over F), and we may therefore assume it is a conjunction of equations and inequations. Using the fact that modulo the theory of fields, the formula $x \neq 0$, is equivalent to $\exists y\ xy = 1$, we may assume that $\varphi(\bar{y})$ is a conjunction of polynomial equations with coefficients in F.

Let $\bar{a} \in L$ be a solution of $\varphi(\bar{y})$. Since $L \cap F^{alg} = F$, we know that the ideal $I(\bar{a}/F) = \{f(\bar{X}) \in F[\bar{X}] \mid f(\bar{a}) = 0\}$ is an absolutely prime ideal, see 2.7. I.e., the set V of tuples on which all elements of $I(\bar{a}/F)$ vanish is a variety, which is defined over F. Since F is PAC, it follows that there is some tuple $\bar{b}$ in F on which all polynomials of $I(\bar{a}/F)$ vanish. Hence, $\bar{b}$ satisfies every polynomial equation over F that $\bar{a}$ satisfies, and in particular, will satisfy φ.

2.11. Comments. The condition $L \cap F^{alg} = F$ is clearly necessary: if $\alpha \in L \cap F^{alg}$ and $\alpha \notin F$, and if $p(X)$ is the minimal polynomial of α over F, then $L \models \exists y\, p(y) = 0$, but $F \models \forall y\, p(y) \neq 0$.

It is in general not a sufficient condition. E.g., we will see that one can find a pseudo-finite field F such that $F \cap \mathbb{F}_p^{alg} = \mathbb{F}_p$, and clearly $\mathbb{F}_p \not\prec_1 F$.

Theorem 2.12. *There is a theory (in the language of rings) whose models are exactly the PAC fields.*

Proof. Fix integers m, n, d. We need to express the following:
Let $f_1(\bar{X}), \dots, f_m(\bar{X})$ be polynomials in $\bar{X} = (X_1, \dots, X_n)$ of degree $\leq d$, and assume they generate an absolutely prime ideal. Then they have a common zero.

This follows from results of Hermann, see below 2.15. If d is an integer, then we denote by $F[\bar{X}]_{\leq d}$ the set of polynomials of degree $\leq d$. They form a finite dimensional F-vector space, and are therefore definable in F. The following maps are also definable:
Addition: $F[\bar{X}]_{\leq d} \times F[\bar{X}]_{\leq d} \to F[\bar{X}]_{\leq d}$,
Multiplication: $F[\bar{X}]_{\leq d} \times F[\bar{X}]_{\leq d} \to F[\bar{X}]_{\leq 2d}$.

2.13. Results of Hermann. (For a proof, see [16] or [40].)

(1) There is a constant $A = A(n,d)$ such that for every field F, polynomials $f_1, \dots, f_m, g \in F[\bar{X}]_{\leq d}$, if g belongs to the ideal of $F[\bar{X}]$ generated by $f_1, \dots, f_m$, then there are $h_1, \dots, h_m \in F[\bar{X}]_{\leq A}$ such that $g = \sum_{i=1}^m f_i h_i$.

(2) There is a constant $B = B(n,d)$ such that for every field F, for every ideal I of $F[X]$ generated by elements of $F[X]_{\leq d}$ and for every $g \in F[X]_{\leq d}$, if $g^k \in I$ for some integer k, then $g^B \in I$.

(3) There is a constant $C = C(n,d)$ such that for every field F, ideals I and J generated by elements of $F[X]_{\leq d}$, the ideals $I \cap J$ and $J : I = \{f \in F[X] \mid fI \subseteq J\}$ are generated by elements of $F[X]_{\leq C}$.

(4) There is a constant $D = D(n,d)$ such that for every field F and ideal I of $F[X]$ generated by elements of $F[X]_{\leq d}$, if I is not prime, then there are $g, h \in F[X]_{\leq D}$ such that $gh \in I$ but $g, h \notin I$.

(5) There is a constant $E = E(n,d)$ such that for every field F and ideal I of $F[X]$ generated by elements of $F[X]_{\leq d}$, there are at most E minimal prime ideals containing I, and they are generated by elements of $F[X]_{\leq E}$.

Corollary 2.14. *Let $n, d \geq 1$. There is a formula $\varphi(\bar{y})$, $\bar{y}$ an $mN(d)$-tuple of variables, such that in every field F, for every $mN(d)$-tuple $\bar{a}$ in F, if $f_1, \dots, f_m$ is the m-tuple of elements of $F[\bar{X}]_{\leq d}$ encoded by $\bar{a}$, then*

$$F \models \varphi(\bar{a}) \iff \textit{the ideal of } F[\bar{X}] \textit{ generated by } f_1, \dots, f_m \textit{ is prime.}$$

Proof. Let $D = D(n,d)$, $A = A(n,D)$. Then

$f_1, \dots, f_m$ generate a prime ideal I in $F[\bar{X}]$

if and only if for all $g, h \in F[\bar{X}]_{\leq D}$, either $gh \notin I$ or one of g, h is in I,

if and only if for all $g, h \in F[\bar{X}]_{\leq D}$, either for all $h_1, \dots, h_m \in F[\bar{X}]_{\leq A}$, $gh \neq \sum_{i=1}^m h_i f_i$, or there are $h_1, \dots, h_m \in F[\bar{X}]_{\leq A}$ such that $[g = \sum_{i=1}^m h_i f_i$ or $h = \sum_{i=1}^m h_i f_i]$.

This last statement is clearly an elementary property of the $mN(d)$-tuple $\bar{a}$ of coefficients of $f_1, \dots, f_m$.

Corollary 2.15. *Let $n, d \geq 1$. There is a* **quantifier-free** *formula $\psi(\bar{y})$, $\bar{y}$ an $mN(d)$-tuple of variables such that in every field F, for every $mN(d)$-tuple $\bar{a}$ in F, if $f_1, \dots, f_m$ is the m-tuple of elements of $F[\bar{X}]_{\leq d}$ encoded by $\bar{a}$, then*

$$F \models \psi(\bar{a}) \iff \textit{the ideal of } F^{alg}[\bar{X}] \textit{ generated by } f_1, \dots, f_m \textit{ is prime.}$$

Proof. Take the formula $\varphi(\bar{y})$ given by 2.14. By quantifier-elimination of the theory of algebraically closed fields[5], there is a quantifier-free formula $\psi(\bar{y})$ such that in every algebraically closed field K, for every $mN(d)$-tuple $\bar{a}$ in K we have

$$K \models \varphi(\bar{a}) \iff K \models \psi(\bar{a}).$$

[5] Modulo the theory ACF, every formula is equivalent to a quantifier-free formula.

But if the tuple $\bar{a}$ is in the subfield F of K, we have

$$K \models \psi(\bar{a}) \iff F \models \psi(\bar{a}).$$

Thus $F \models \psi(\bar{a})$ if and only if the m-tuple $(f_1, \ldots, f_m)$ of $F[\bar{X}]_{\leq d}$ encoded by $\bar{a}$ generates a prime ideal in $F^{alg}[\bar{X}]$.

Theorem 2.16. (Lang-Weil) *([30]). For every positive integers n, d, there is positive constant C ($= C(n, d)$) such that for every finite field $\mathbb{F}_q$ and variety V defined by polynomials in $\mathbb{F}_q[X_1, \ldots, X_n]_{\leq d}$,*

$$\left| |V(\mathbb{F}_q)| - q^{\dim(V)} \right| \leq C q^{\dim(V)-1/2}.$$

[Recall that $V(\mathbb{F}_q)$ is the set of points of $V \cap \mathbb{F}_q^n$, and $\dim(V)$ is the dimension of V, i.e., $tr.deg(\mathbb{F}_q(V)/\mathbb{F}_q)$.] In particular, if $q > C^2$, then any variety V as above will have a rational point in $\mathbb{F}_q$. Indeed, we get

$$0 < -Cq^{\dim(V)-1/2} + q^{\dim(V)} \leq |V(\mathbb{F}_q)|.$$

The constant C can be effectively computed.

2.17. Axiom 3(m, n, d), and Axiom 3'(m, n, d). Consider the following two axiom schemes:
Axiom 3(m, n, d): whenever $f_1(\bar{X}), \ldots, f_m(\bar{X})$ are polynomials in $\bar{X} = (X_1, \ldots, X_n)$ of degree $\leq d$ and which generate an absolutely prime ideal, then there is an n-tuple $\bar{a}$ such that $\bigwedge_i f_i(\bar{a}) = 0$.
Axiom 3'(m, n, d): Either the field has less than $C(n, d)^2$ elements, or whenever $f_1(\bar{X}), \ldots, f_m(\bar{X})$ are polynomials in $\bar{X} = (X_1, \ldots, X_n)$ of degree $\leq d$ and which generate an absolutely prime ideal, then there is an n-tuple $\bar{a}$ such that $\bigwedge_i f_i(\bar{a}) = 0$.

We let T_f^* be the theory obtained by taking all the axioms 1, 2(ℓ) and 3'(m, n, d), and Psf* the theory obtained by taking all the axioms 1, 2(ℓ) and 3(m, n, d). Observe that Psf* is $T_f^* \cup \{$ there are infinitely many elements$\}$.

Theorem 2.18. (1) *Finite fields are models of the axioms 1, 2(ℓ) and 3'(m, n, d).*

(2) *Let $\mathcal{Q}$ be the set of all prime powers, and let $\mathcal{U}$ be a non-principal ultrafilter on $\mathcal{Q}$. Then the field $F^* = \prod_{q \in \mathcal{Q}} \mathbb{F}_q / \mathcal{U}$ is a model of Psf*.*

Proof. Clearly any infinite model of the scheme of axioms 3'(m, n, d) is pseudo-algebraically closed, so it suffices to show the first assertion. The result of Lang-Weil 2.16 gives scheme of axioms 3'(m, n, d). We also know that finite fields are perfect, and that they have exactly one algebraic extension of each degree.

3 Showing that $T_f^* \vdash T_f$.

So, we have shown that the theory T_f^* is satisfied by every finite field, and therefore is contained in the theory T_f of all finite fields. In order to show that T_f^* axiomatises the theory of all finite fields, we need to show the converse. I.e., that if a sentence θ is true in all finite fields, then it is true in all models of T_f^*. Since such a sentence is obviously true in all finite models of T_f^*, it remains to show that it is true in all pseudo-finite fields. In other words, we need to show that the pseudo-finite fields are exactly the infinite models of the theory T_f^*.

To do that, it is enough to show that if F is a model of Psf*, then F is elementarily equivalent to an ultraproduct of finite fields. Indeed, this will imply that a formula which is true in all finite fields is also true in this arbitrary infinite model of T_f^* (by Łos' theorem), and therefore that $T_f = T_f^*$ (or rather, $T_f^* \vdash T_f$).

The strategy to do that, is to describe the completions of the theory T_f^*, or rather, of the theory Psf*. Once we have described the completions of Psf*, we will relatively easily obtain the result, as well as some "quantifier-elimination" results.

The main tool in the description of the completions of Psf* is the following

Lemma 3.1. (The embedding Lemma - Simplified version) *Let K, E, K^* be perfect fields, contained in some large algebraically closed field Ω, and such that*

(1) *$K \subset E, K^*$,*
(2) *$K^{alg} \cap K^* = K^{alg} \cap E = K$,*
(3) *K^* is a model of Psf* and $\aleph_1$-saturated,*[6]
(4) *E is countable, has at most one extension of each degree.*

Then there is a field embedding $\varphi : E \to K^*$ such that $\varphi|_K = id$ and $\varphi(E)^{alg} \cap K^* = \varphi(E)$.

I will not give a proof of this result, as it uses the Galois correspondence in an essential way. You can find a proof in the book of Fried and Jarden ([13], Lemma 20.2.2).

Theorem 3.2. *(Kiefe [26]) Let K and L be pseudo-finite fields, containing a common subfield k. Assume that*

$$k^{alg} \cap K \simeq_k k^{alg} \cap L.$$

Then $K \equiv_k L$ (i.e., K and L are elementarily equivalent in the language $\mathcal{L}(k)$ obtained by adding to the language of rings constant symbols for the elements of k).

Proof. If the result is false, then a formula showing it is false will only involve finitely many parameters from k. Hence, we may assume that k is countable. Passing to elementary extensions of K and L, we may also assume that K and L are $\aleph_1$-saturated: if $K \prec K^*$, $L \prec L^*$ and $K^* \equiv_k L^*$, then also $K \equiv_k L$.

We now consider the following family $\mathcal{I}$ of partial isomorphisms: $f : A \to B$, where $A \subset K$ and $B \subset L$, is in $\mathcal{I}$ if and only if it is a field isomorphism, A and B are countable, and $A^{alg} \cap K = A$, $B^{alg} \cap L = B$.

[6] Recall that a model M is $\aleph_1$-saturated if for every countable subset A of M, and set $\Sigma(x)$ of formulas with parameters in A, if Σ is finitely consistent, then it has a realisation in M. Every model has an elementary extension which is $\aleph_1$-saturated.

We will show that the family $\mathcal{I}$ has the *back-and-forth property*, i.e., it is non-empty, and :
– If $f \in \mathcal{I}$ and $a \in K$ there is $g \in \mathcal{I}$ extending f and with a in its domain,
– and if $b \in L$, there is $g \in \mathcal{I}$ extending f and with b in its image.

The first task is to show that $\mathcal{I}$ is non-empty: to do that, we use the assumption: there is an isomorphism $\phi_0 : k^{alg} \cap K \to k^{alg} \cap L$ which is the identity on K. So, $\phi_0 \in \mathcal{I}$.

Suppose we have $f \in \mathcal{I}$ and $a \in K$ as above. Let $E = A(a)^{alg} \cap K$. We first extend f to an automorphism $\tilde{f}$ of the big algebraically closed field Ω in which we are working, and let $E_0 = \tilde{f}(E)$. We wish to use the Embedding lemma 3.1. We already know that $E_0 \cap B^{alg} = B$, since we had $E \cap A^{alg} \subset K \cap A^{alg} = A$. Furthermore, as $E^{alg} \cap K = E$, and K has at most one algebraic extension of each degree, we know that E has at most one algebraic extension of each degree: indeed, if M is an algebraic extension of E of degree n, then MK is an algebraic extension of K of degree n also. This property is preserved by $\tilde{f}$, and we may therefore apply the Embedding lemma to B, E_0, L: there is $\psi : E_0 \to L$ which is the identity on B and such that $\psi(E_0)^{alg} \cap L = \psi(E_0)$. Then $g = \psi\tilde{f}|_E$ is our desired element of $\mathcal{I}$.

The other direction (back) follows by symmetry.

3.3. Back and forth? This is just a saturated version of Ehrenfeucht-Fraïssé games. One shows by induction on the number of quantifiers, that if $f \in \mathcal{I}$, then f preserves all formulas with n quantifiers, i.e., if $\varphi(\bar{x})$ has n quantifiers, and $\bar{a}$ is in the domain of f, then $K \models \varphi(\bar{a})$ if and only if $L \models \varphi(f(\bar{a}))$. Left as exercise.

The Ehrenfeucht-Fraïssé games are defined in a similar way, but they only work for a finite relational language. One defines by induction, for tuples $\bar{a}$ in M and $\bar{b}$ in N of the same length: $\bar{a} \equiv_0 \bar{b}$ iff they satisfy the same atomic formulas; $\bar{a} \equiv_{n+1} \bar{b}$ iff whenever $c \in M$ there is $d \in N$ such that $\bar{a}, c \equiv_n \bar{b}, d$, and conversely whenver $d \in N$ there is $c \in M$ such that $\bar{a}, c \equiv_n \bar{b}, d$. Then $M \equiv N$ iff $\emptyset \equiv_n \emptyset$ for all n. This only works for a **finite relational** language, as can be easily checked for $\mathbb{Q}^{alg}$ and $\mathbb{C}$ in the usual language of rings $\{+, -, \cdot, 0, 1\}$: given $d \in \mathbb{C}$ transcendental over $\mathbb{Q}$, one cannot find $c \in \mathbb{Q}^{alg}$ such that $c \equiv_0 d$.

However, one has the following result: Let M, N be models of a theory, M^*, N^* elementary extensions of M, N respectively, which are κ^+-saturated, where κ is the cardinality of the language. Then, if $M \equiv N$, also $M^* \equiv N^*$; consider the system $\mathcal{I}$ of isomorphisms $A \to B$, where $A \prec M^*$, $B \prec N^*$, and $|A| = |B| \leq \kappa$. Using Löwenheim Skolem and κ^+-saturation of M^* and N^*, one verifies that this set $\mathcal{I}$ is non-empty and has the back-and-forth property.

Definition 3.4. If K is a field and $k_0 \subseteq K$ the prime subfield of K (i.e., the field of fractions of the subring of K generated by 1; it equals $\mathbb{Q}$ or $\mathbb{F}_p$), then the *(field of) absolute numbers of* K is the field $k_0^{alg} \cap K$.

Corollary 3.5. *(Ax [1] Thm 4) The completions of Psf* are obtained by describing the isomorphism type of the field of absolute numbers of a model.*

Proof. Clear from Theorem 3.2: if F_1 and F_2 are models of Psf* and have isomorphic fields of absolute numbers k, then $F_1 \equiv F_2$. On the other hand, a completion of Psf* will determine the characteristic, as well as which polynomials with coefficients in the prime field have a solution or not. (See also the proof of 3.7).

Corollary 3.6. *If $F_1 \subseteq F_2$ are models of Psf* then*

$$F_1 \prec F_2 \iff F_1^{alg} \cap F_2 = F_1.$$

Proof. This follows from Theorem 3.2, with $k = F_2$.

Corollary 3.7. *(Kiefe). Modulo the theory Psf*, any formula $\varphi(\bar{x})$ is equivalent to a Boolean combination of formulas of the form $\exists t\ f(\bar{x}, t) = 0$, where $f(\bar{X}, T) \in \mathbb{Z}[\bar{X}, T]$.*

Proof. By compactness, it suffices to show that if $F_1, F_2 \models$ Psf* have the same characteristic, and $\bar{a}$, $\bar{b}$ are n-tuples in F_1, F_2 respectively, such that for every $f(\bar{X}, T) \in \mathbb{Z}[\bar{X}, T]$,

$$(1) \qquad F_1 \models \exists t\ f(\bar{a}, t) = 0 \iff F_2 \models \exists t\ f(\bar{b}, t) = 0,$$

then for any formula $\varphi(\bar{x})$, we have

$$F_1 \models \varphi(\bar{a}) \iff F_2 \models \varphi(\bar{b}).$$

By Theorem 3.2, this last condition is equivalent to the existence of an isomorphism between the fields $A = k_0(\bar{a})^{alg} \cap F_1$ and $B = k_0(\bar{b})^{alg} \cap F_2$, where k_0 is the prime subfield of F_1 and F_2. We will show that (1) implies that such an isomorphism exists. First of all note that there is an isomorphism $\varphi_0 : k_0(\bar{a}) \to k_0(\bar{b})$ which sends $\bar{a}$ to $\bar{b}$: this is because the tuples $\bar{a}$ and $\bar{b}$ satisfy the same polynomial equations over k_0. Extend φ_0 to $\varphi : k_0(\bar{a})^{alg} \to k_0(\bar{b})^{alg}$. We need to show that such a φ can be chosen with $\varphi(A) = B$. Equivalently, we need to find $\sigma \in \mathrm{Aut}(k_0(\bar{b})^{alg})/k_0(\bar{b}))$ such that $\sigma(\varphi(A)) = B$. We know that for any $f(\bar{X}, T) \in \mathbb{Z}[\bar{X}, T]$, we have

$$\varphi(A) \models \exists t\ f(\bar{b}, t) = 0 \iff B \models \exists t\ f(\bar{b}, t) = 0.$$

In case the characteristic is $p > 0$, consider the perfect closures A_0 and B_0 of $k_0(\bar{a})$ and $k_0(\bar{b})$ respectively. (So they are obtained by closing under taking p-th roots). Note that the extension of φ_0 to A_0 is unique, since any element has a unique p-th root, and therefore $\varphi(A_0) = B_0$. Furthermore, if $f(\bar{X}, T) \in \mathbb{F}_p[\bar{X}, T]$ and $m \in \mathbb{N}$, then $A \models \exists t\, f(\mathrm{Frob}^{-m}(\bar{a}), t) = 0$ if and only if $A \models \exists t\, f(\bar{a}, t) = 0$: this is because A is perfect and Frob is the identity on $\mathbb{F}_p$. As every element of A_0 is in $k_0(\mathrm{Frob}^{-m}(\bar{a}))$ for some m, and the extension of φ_0 to A_0 is unique, this means that whenever a polynomial $f(t) \in A_0[T]$ has a solution in A, then its image under φ has a solution in B, and conversely. The result will now follow from the following lemma:

Lemma 3.8. *Let B be a perfect field, and B_1, B_2 two subfields of B^{alg} which contain B. Assume that for every $f(T) \in B[T]$ we have*

$$B_1 \models \exists t\ f(t) = 0 \iff B_2 \models \exists t f(t) = 0.$$

Then there is $\sigma \in \mathrm{Aut}(B^{alg}/B)$ such that $\sigma(B_1) = B_2$.

Proof. As B is perfect, B^{alg} is the union of finite Galois extensions of B. Hence, it is enough to find $\sigma \in \mathrm{Aut}(B^{alg}/B)$ such that so that for every finite Galois extension L of B, one has $\sigma(B_1 \cap L) = B_2 \cap L$.

For each finite Galois extension L of B consider

$$\mathcal{S}_L = \{\sigma \in \mathrm{Aut}(B^{alg}/B) \mid \sigma(L \cap B_1) = L \cap B_2\}.$$

Claim. $\mathcal{S}_L$ is not empty.

Let $\alpha \in L$ be such that $L \cap B_1 = B(\alpha)$, and let $f(T)$ be its minimal polynomial[7]. Then $B_1 \models f(\alpha) = 0$, and so there is some $\beta \in B_2$ such that $f(\beta) = 0$. Let $\sigma \in \mathrm{Aut}(L/B)$ be such that $\sigma(\alpha) = \beta$. Then certainly $\sigma(B_1) \subseteq B_2$, and therefore $[B_1 : B] \leq [B_2 : B]$. The symmetric argument gives $[B_2 : B] \leq [B_1 : B]$, and this implies that the degrees are equal, and $\sigma(B_1) = B_2$. Lift σ to an element of $\mathrm{Aut}(B^{alg}/B)$.

Thus the family $\mathcal{S}_L$, L ranging over all finite normal extensions of B, has the finite intersection property: If L and M are finite normal extensions of B, then so is their field composite[8] LM and we have $\mathcal{S}_{LM} \subseteq \mathcal{S}_L \cap \mathcal{S}_M$. By compactness of the profinite group $\mathrm{Aut}(B^{alg}/B)$, there is some σ in the intersection of all $\mathcal{S}_L$, and this σ satisfies $\sigma(B_1) = B_2$.

Remarks 3.9. Another way of stating Corollary 3.5 is to say that modulo Psf*, every sentence is equivalent to a Boolean combination of sentences $\exists t\ f(t) = 0$, where $f(T) \in \mathbb{Z}[T]$.

The perfectness condition on B in 3.8 can be omitted, but the proof is a little more involved, since one has to deal with inseparable extensions: not all finte algebraic extensions are generated by one element. Oneshows that if L is a finite algebraic extension of B contained in B_1, and $\hat{L}$ its normal closure over B, $L_1 = \hat{L} \cap B_2$, then there is a B-embedding φ of L into $\bigcup_{\sigma \in \mathrm{Aut}(\hat{L}/B)} L_1^\sigma$. But if B is not perfect, then B is infinite, and one then shows that $\varphi(L)$, being a B-vector space, must be contained in one of the L_1^σ's. I.e., $\hat{L} \cap B_1$ B-embeds into $\hat{L} \cap B_2$.

What are the constraints on fields of absolute numbers of pseudo-finite fields? Actually, none, beside the fact that they must have at most one extension of each degree. [Recall that they must be relatively algebraically closed in a field having exactly one extension of each degree]. Hence $\mathbb{Q}^{alg}$ is allowable, as is any subfield of $\mathbb{F}_p^{alg}$.

To finish the proof that $T_f = T_f^*$, it therefore suffices to prove the following:

Theorem 3.10. *Let $k = \mathbb{F}_p$ or $k = \mathbb{Q}$, and let $E \subseteq k^{alg}$ have at most one extension of each degree. Then there is an ultraproduct K^* of finite fields such that the field of absolute numbers of K^* is isomorphic to E. When the characteristic of E is 0, K^* can be chosen to be an ultraproduct of prime fields.*

Proof. We will start with the easy cases, when the characteristic of E is $p > 0$. The characteristic 0 case will need Chebotarev's theorem, see below 3.11.

Case 1. E is infinite (and of characteristic $p > 0$).

[7] That such an element exists is because L is finite separable over B

[8] the subfield of B^{alg} generated by L and M.

Let n_m be a sequence of integers such that n_m divides n_{m+1}, and $E = \bigcup_m \mathbb{F}_{p^{n_m}}$. For instance, as $\mathbb{F}_p = \bigcup_m \mathbb{F}_{p^{m!}}$, we can define n_m by $\mathbb{F}_{p^{n_m}} = E \cap \mathbb{F}_{p^{m!}}$. Let $\mathcal{U}$ be any non-principal ultrafilter on $\mathbb{N}$, and let $K^* = \prod_m \mathbb{F}_{p^{n_m}}/\mathcal{U}$. Then $K^* \cap \mathbb{F}_p^{alg} \simeq E$.

Indeed, clearly K^* is of characteristic p. Let $d \in \mathbb{N}$. If $\mathbb{F}_p^d \subset E$, then $\mathbb{F}_{p^d}$ will be contained in all fields $\mathbb{F}_{p^{n_m}}$ with $m \geq d$. Hence, by Łos' theorem, K^* will satisfy the sentence "there is an element of multiplicative order exactly $p^d - 1$", and therefore will contain (a copy of) $\mathbb{F}_{p^d}$. On the other hand, if $\mathbb{F}_p^d \not\subset E$, then $\mathbb{F}_{p^d}$ is contained in no $\mathbb{F}_{p^{n_m}}$, therefore K^* will satisfy "there is no element of multiplicative order exactly $q^d - 1$", and K^* will not contain $\mathbb{F}_{q^d}$. Hence we will have $K^* \cap \mathbb{F}_p^{alg} = \mathbb{F}_q$.

Case 2. E is finite.

Let $q = |E|$, so that $E = \mathbb{F}_q$. Consider any non-principal ultrafilter $\mathcal{U}$ on the set $\mathcal{P}$ of prime numbers, and let $K^* = \prod_{\ell \in \mathcal{P}} \mathbb{F}_{q^\ell}/\mathcal{U}$. Then K^* is of characteristic p and contains $\mathbb{F}_q$. But, if $d > 1$, all but at most one field $\mathbb{F}_{q^\ell}$ satisfy "there is no element of multiplicative order exactly $q^d - 1$", and therefore $K^* \cap \mathbb{F}_p^{alg} = \mathbb{F}_q$.

Case 3. E is of characteristic 0.

Write $\mathbb{Q}^{alg}$ as the union of an increasing chain L_n, $n \in \mathbb{N}$, of finite Galois extensions of $\mathbb{Q}$. For each n, let $E_n = L_n \cap F$, and let $I(n)$ be the (finite) set of subfields of L_n which properly contain E_n. We will find a sentence θ_n which describes $L_n \cap F$. Choose a generator α of E_n over $\mathbb{Q}$, and let $f_n(T)$ be its minimal (monic) polynomial over $\mathbb{Q}$. Similarly, for each $M \in I(n)$, choose a generator β_M of M over $\mathbb{Q}$, let $g_M(T)$ be the minimal (monic) polynomial of β_M over $\mathbb{Q}$, and define $g_n(T) = \prod_{M \in I(n)} g_M(T)$. Consider now the sentence $\theta_n : \exists t\ f_n(t) = 0 \wedge \forall t\ g_n(t) \neq 0$. This is a sentence satisfied by E, and if F is any field of characteristic 0, then $F \models \theta_n \iff F \cap L_n \simeq E_n$.

As the L_n's form an increasing chain, so do the E_n's, and we have $\theta_n \to \theta_{n-1}$. In order to find an ultraproduct of prime fields with field of absolute numbers isomorphic to F, it is therefore enough to show that for each n, the set

$$S_n := \{p \in \mathcal{P} \mid \mathbb{F}_p \models \theta_n\}$$

is infinite. As $S_n \supset S_{n+1}$, there will be a non-principal ultrafilter $\mathcal{U}$ containing all S_n's, and if $K^* = \prod_{p \in \mathcal{P}} \mathbb{F}_p/\mathcal{U}$, then $K^* \models \theta_n$ for each n, i.e.: $K^* \cap \mathbb{Q}^{alg} \simeq E$.

That S_n is infinite follows from Chebotarev's density theorem. Here is the consequence of Chebotarev's theorem that we will use:

3.11. *Let $f_1(T), \dots, f_m(T), g(T) \in \mathbb{Z}[T]$, T a single variable. Let L be the Galois extension of $\mathbb{Q}$ obtained by adjoining all roots of the polynomials $f_i(T)$, $i = 1, \dots, m$. Assume that there is a subfield E of L such that* $\mathrm{Aut}(L/E)$ *is cyclic and*

$$E \models \bigwedge_{i=1}^{m} \exists t\ f_i(t) = 0 \wedge \forall t\ g(t) \neq 0.$$

Then the set of prime numbers p such that $\mathbb{F}_p \models \bigwedge_{i=1}^m \exists t\ f_i(t) = 0 \wedge \forall t\ g(t) \neq 0$ is infinite.

For a precise statement, see e.g. [13], Thm 6.3.1.

3.12. So this shows that Psf* is the theory of all pseudo-finite fields, and therefore that T_f^* is the theory of all finite fields. From now on, I will drop the *.

3.13. Decidability issues. Observe first that by Theorem 3.10 we have

$$\mathrm{Psf} \subset \mathrm{Psf}_0 \subset T_{\text{prime}} \text{ and } \mathrm{Psf} \subset T_f \subset T_{\text{prime}}.$$

(Here Psf_0 denotes the theory of pseudo-finite fields of characteristic 0 and T_{prime} the theory of all prime fields.) We will first show that the theory Psf is decidable, that is, that there is an algorithm which decides, given a sentence θ, whether it is true in all pseudo-finite fields or not. From this we will be able to derive the decidability of the other theories.

We have an enumeration of a set Γ consisting of axioms for the theory Psf (this assumes that the bounds given in 2.13 on degrees of polynomials can be computed effectively, but they can). Hence, we can produce an enumeration of the set of all proofs made using axioms of Γ, and therefore of the theory Psf (by the completeness theorem, if a sentence is true in all pseudo-finite fields, then it is provable from Γ). Similarly, we have an enumeration of a set Γ_0 of axioms for the theory Psf_0 of all pseudo-finite fields of characteristic 0, and of the theory Psf_0. Note that $\Gamma_0 = \Gamma \cup \{p \neq 0 \mid p \text{ a prime}\}$.

This tells us that if θ is in Psf, then going through the enumeration of Psf we will find it. However, we need another procedure to decide if $\theta \notin \mathrm{Psf}$. This is what we will do below. Let us fix a sentence θ.

Let ψ_n, $n \in \mathbb{N}$, be an enumeration of all sentences which are Boolean combinations of sentences of the form $\exists t\ f(t) = 0$, where $f(T) \in \mathbb{Z}[T]$. By 3.9, we know that $\Gamma \vdash \theta \leftrightarrow \psi_n$ for some n, i.e., $\theta \leftrightarrow \psi_n \in \mathrm{Psf}$, and therefore we can effectively find this ψ_n. Note that the proof of $\theta \leftrightarrow \psi_n$ uses only a finite number of axioms expressing the PAC property, and we can therefore find a constant C_1 (given by Lang-Weil (2.16)) such that in all finite fields $\mathbb{F}_q$ with $q > C_1$ we have

$$\mathbb{F}_q \models \theta \leftrightarrow \psi_n.$$

It now remains to decide whether ψ_n is true in all pseudo-finite fields. I.e., we need to show that if k is a prime field, and $E \subseteq k^{alg}$ has at most one algebraic extension of each degree, then $E \models \psi_n$.

Step 1. Decide whether $\psi_n \in \mathrm{Psf}_0$ or not.

We know that ψ_n is (equivalent to) a disjunction of sentences of the form $\bigwedge_i \exists t\ f_i(t) = 0 \wedge \forall t\ g(t) \neq 0$. Let L be the extension of $\mathbb{Q}$ generated by all roots of all polynomials appearing in ψ_n. Then one can compute effectively $\mathrm{Aut}(L/\mathbb{Q})$, as well as those subfields E of L such that $\mathrm{Aut}(L/E)$ is cyclic. Hence we can decide whether or not ψ_n is true in all subfields E of L such that $\mathrm{Aut}(L/E)$ is cyclic. If it is not, then $\psi_n \notin \mathrm{Psf}_0$ and therefore $\psi_n \notin \mathrm{Psf}$, i.e., $\theta \notin \mathrm{Psf}_0$, $\theta \notin \mathrm{Psf}$.

Step 2. Decide whether $\psi_n \in \mathrm{Psf}$.

Assume that $\psi_n \in \mathrm{Psf}_0$. Then it is provable from Γ_0, and its proof only uses finitely many axioms expressing that the characteristic is $\neq p$; therefore there is a constant C_2 such that ψ_n holds in all pseudo-finite fields of characteristic $p > C_2$. It therefore remains to check whether

ψ_n holds in all pseudo-finite fields of characteristic $p \leq C_2$. Fix one such p. Then, as in step 1 we let $\mathbb{F}_{p^m}$ be the extension of $\mathbb{F}_p$ generated by all roots of polynomials appearing in ψ_n. It then suffices to check whether $\mathbb{F}_{p^d} \models \psi_n$ or not for all d dividing m. This is certainly decidable, and finishes the proof that Psf is decidable.

Step 3. Decidability of T_f and of T_{prime}.

We assume now that all pseudo-finite fields satisfy θ. Hence there is a proof of θ from Γ, and this proof will involve only finitely many axioms saying that varieties have points. Hence, there is a constant C_3 such that $T_f \cup \{$there are at least C_3 elements$\}$ proves θ. It now remains to check whether θ holds in the finitely many finite fields of size $< C_3$. But this is decidable.

Similarly, assume that all pseudo-finite fields of characteristic 0 satisfy θ. Then the proof of θ from Psf_0 uses only finitely many axioms saying that varieties have points and that "$p \neq 0$", and there is a constant C_4 such that $T_f \cup \{$"$p \neq 0$", $p < C_4\}$ proves θ. It now remains to check whether θ holds in the finitely many prime fields of size $< C_4$.

[So we didn't need C_1 after all].

4 More results on pseudo-finite fields

If M is a structure, and $\varphi(\bar{x})$ is a formula in the language of M, we denote by $\varphi(M)$ the set of tuples in M satisfying φ.

4.1. Examples of pseudo-finite fields. If F is an infinite subfield of $\mathbb{F}_p^{alg}$, then F is PAC by the theorem of Lang-Weil (2.16), and is perfect. Hence, any infinite subfield F of $\mathbb{F}_p^{alg}$ is pseudo-finite as soon as it satisfies axiom $2(\ell)$ for all ℓ. (By group theory results, it actually suffices to have it for all primes). Hence, if f is any function from the set of prime numbers to the positive integers, and F is the field composite of all $\mathbb{F}_{p^{f(\ell)}}$, ℓ a prime, then F is pseudo-finite.

This gives us many pseudo-finite fields of positive characteristic. In characteristic 0, there are no such explicit examples. However a result of Jarden (for a proof see [13], 18.5.6 with $K = \mathbb{Q}$ and $e = 1$, and 18.6.1) shows that there are many such fields. The profinite group $\text{Aut}(\mathbb{Q}^{alg}/\mathbb{Q})$ is compact, and has a unique Haar probability measure. In the sense of this measure, for almost all $\sigma \in \text{Aut}(\mathbb{Q}^{alg}/\mathbb{Q})$, the subfield of $\mathbb{Q}^{alg}$ fixed by σ is pseudo-finite. Other examples are of course non-principal ultraproducts of prime fields.

Quantifier-elimination results for pseudo-finite fields. An easy consequence of Kiefe's result 3.7 is:

Theorem 4.2. *Let $\mathcal{L}'$ be the language obtained by adding to the language of rings an $(n+1)$-ary predicate Sol_n for every $n > 1$, and add to the theory Psf the axioms*

$$Sol_n(x_0, \dots, x_n) \iff \exists y \sum_{i=0}^{n} x_i y^i = 0,$$

to obtain a theory Psf$'$. *Then* Psf$'$ *eliminates quantifiers.*

Proof. Let F_1 and F_2 be pseudo-finite fields, containing a common $\mathcal{L}'$-substructure A. Then A is a subring. We need to show that $F_1 \equiv_A F_2$. By Lemma 3.8 (or its proof), there is an isomorphism $f : A^{alg} \cap F_1 \to A^{alg} \cap F_2$ which is the identity on A. Now apply Theorem 3.2.

4.3. Quantifier-elimination for finite fields. The predicates Sol_n need to be slightly modified to take into account the models which are finite. The predicates Sol'_n are defined as follows:

$$Sol'_n(x_0, \dots, x_n) := \big(|F| = n \land \forall y\, y = 0 \lor \bigvee_{i=1}^{n-1} y = x_0^i\big) \lor \big(|F| \neq n \land Sol_n(x_0, \dots, x_n)\big).$$

(I leave you to find the correct translation of $|F| = n$ in first-order logic). I gave a slight variation in class and defined Sol''_n: when $|F| = n$, I imposed $x_0 = 1$, $x_1 = \cdots = x_n = 0$. Note that $Sol_n(1, 0, \dots, 0)$ is never satisfied.
One then expands the theory T_f with the defining axioms of the predicates Sol_n and Sol'_n.
Exercise: show that this expansion of T_f has quantifier elimination.

4.4. A language in which Psf is model complete. We form the language $\mathcal{L}_c$ by adjoining to the language $\mathcal{L}$ of rings new constant symbols $c_{i,n}$, where $2 \leq n \in \mathbb{N}$ and $0 \leq i \leq n-1$. The theory Psf_c is obtained by adding to the theory Psf for each n an axiom stating that the polynomial $X^n + \sum_{i=0}^{n-1} c_{i,n} X^i$ is irreducible.

Note that every pseudo-finite field expands to a model of Psf_c: if F is pseudo-finite, for each n choose the $c_{i,n}$ to be the coefficients of some (monic) irreducible polynomial of degree n.

Recall that a theory T is *model complete* if whenever $M \subseteq N$ are models of T, then $M \prec N$. If T is model complete, then every formula is equivalent modulo T to an existential formula (and to a universal formula).

Theorem 4.5. *The theory* Psf_c *is model complete.*

Proof. Let $F_1 \subseteq F_2$ be models of Psf_c. If L is an algebraic extension of F_1 of degree n, then L is generated over F_1 by a solution of the equation $X^n + \sum_{i=0}^{n-1} c_{i,n} X^i$. Since $F_i \models \mathrm{Psf}_c$, this polynomial stays irreducible over F_2, i.e., $F_2 \cap L = F_1$. By 3.6, we obtain $F_1 \prec F_2$.

Corollary 4.6. *Psf_c is model complete. Moreover, given a formula $\varphi(\bar{x})$, there is an existential formula $\exists \bar{y}\, \psi(\bar{x}, \bar{y})$ which is equivalent to $\varphi(\bar{x})$ modulo Psf_c, such that $\psi(\bar{x}, \bar{y})$ defines an algebraic set, and that in any model F of Psf_c, the image by the projection $F^{|\bar{x}|+|\bar{y}|} \to F^{|\bar{x}|}$ of the algebraic set defined by $\psi(\bar{x}, \bar{y})$ has finite fibers.*

For a proof, see Proposition 2.7 in [6].

4.7. Other quantifier-elimination results. Fried and Sacerdote introduce a more geometric language, in which one has quantifier elimination. They are using "Galois formulas", and the process is called "elimination through Galois stratification". The elimination procedure is primitive recursive. For details see [14], [11] or chapter 30 in [13]. One should note that this is the language that Denef and Loeser found more convenient to set up motivic integration in [7].

4.8. Results of Kiefe on Zeta and Poincaré series. Recall that if R is a ring, then $R[[t]]$, the *ring of formal power series over* R, is the set of formal sums $\sum_{i=0}^{\infty} a_i t^i$. Addition and multiplication are defined by

$$\sum a_i t^i + \sum b_j t^j = \sum (a_i + b_i) t^i, \qquad \sum_i a_i t^i \sum b_j t^j = \sum_n (\sum_{i+j=n} a_i b_j) t^n.$$

[Note that there are only finitely many non-negative integers such that $i + j = n$, so that $\sum_{i+j=n} a_i b_j$ is a finite sum and is well defined].

Let $\varphi(\bar{x})$ be an $\mathcal{L}$-formula, with parameters in $\mathbb{F}_q$, some $q = p^n$, p a prime. For each $s \geq 1$, we define

$$N_s(\varphi) = |\varphi(\mathbb{F}_{q^s})|.$$

We then define two formal series over $\mathbb{Q}$, the Poincaré series π_φ and the zeta series ζ_φ by

$$\pi_\varphi(t) = \sum_{s=1}^{\infty} N_s(\varphi) t^s, \qquad \zeta_\varphi(t) = \exp(\sum_{s=1}^{\infty} \frac{N_s(\varphi)}{s} t^s).$$

Theorem 4.9. *(Kiefe [26]).* $\pi_\varphi(t)$ *is rational in* t *(i.e., is of the form* $p(t)/q(t)$*, with* $p(t), q(t) \in \mathbb{Q}[t]$*, and* $q(0) \neq 0$*).*

I will give indications for the proof of the first result. (For a complete proof, see the book of Fried and Jarden (31.3.7 in [13]), for instance. The proof given there uses Galois formulas. It also gives a result for the zeta series.)

Proof. Dwork proved that if U is an algebraic set, then $\pi_U(t)$ and $\zeta_U(t)$ are rational functions. We will use his result. Let us first do quantifier-free formulas of the language of rings. Using normal form, we may write such a formula as a disjunction of conjunctions of atomic and negations of atomic formulas. I.e., the set it defines in a field F will be of the form $\bigcup_{i=1}^{n} V_i(F) \setminus W_i(F)$ for some algebraic sets V_i, W_i, with $W_i \subset V_i$. Then

$$\pi_{V_i \setminus W_i}(t) = \pi_{V_i}(t) - \pi_{W_i}(t).$$

We now use the following formula, proved by induction on the integer m, for finite sets $P_1, \ldots, P_m$:

$$|\bigcup_i P_i| = \sum_{B \subseteq \{1,\ldots,m\}} (-1)^{|B|+1} |\bigcap_{i \in B} P_i|.$$

This gives the result, since a finite intersection of sets of the form $V \setminus W$ has the same form.

It then remains to do the case of formulas with quantifiers. We use Kiefe's elimination of quantifiers result, and the following observation: let $\varphi(\bar{x})$ be a formula of the language of Kiefe, and $\varphi_0(\bar{x})$ be the formula obtained by replacing each occurence of a predicate Sol'_n by Sol_n. Then the set of integers s where $\varphi(\mathbb{F}_{q^s})$ and $\varphi_0(\mathbb{F}_{q^s})$ differ is finite. Therefore the Poincaré series associated to the two formulas will differ by a polynomial. This allows us to reduce to formulas only containing the predicates Sol_n. Using the same trick as before to get rid of negations, we may assume that our formula is positive. In class, I explained (or tried to) explain Kiefe's proof, with a certain amount of hand-waving. I now realise that another, more direct proof is easier and more correct, and I will present it. As in the previous case, we may reduce to the case where our definable set is defined by a conjunction of atomic formulas, i.e., by $\bigwedge_i f_i(\bar{x}) = 0$ and $\bigwedge_j \exists t_j \, g_j(\bar{x}, t_j) = 0$. Refining further and distinguishing more cases, we may assume that for any field F and tuple $\bar{a}$ in F, if $\bigwedge_i f_i(\bar{a}) = 0$, then none of the $g_j(\bar{a}, t_j)$ is identically 0. (This is where I did some handwaving in class: I was using the other normal form, where my formulas were disjunctions of atomic, and there distinguishing cases is a little more delicate since it involves conjunctions). Let N be the product of the degrees of the $g_j(\bar{x}, t)$ in t, consider the algebraic set V defined by

$$\bigwedge_i f_i(\bar{x}) = 0, \bigwedge_j g_j(\bar{x}, t_j) = 0.$$

We are interested in its projection U, and we know that the projection $f : V \to U$ has finite fibers, of size $\leq N$. Let F be any field. Then U is the disjoint union of the set U_ℓ's with $1 \leq \ell \leq N$, where each U_ℓ denotes the set of points $\bar{x}$ in U with $|f^{-1}(\bar{x})| = \ell$. Then

$$\pi_U(t) = \sum_{\ell=1}^{N} \pi_{U_\ell}(t)/\ell.$$

Let m be the number of polynomials g_j. We now consider, for each $\ell \leq N$, the algebraic set W_ℓ defined by the system of equations expressing that:
$\bigwedge_i f_i(\bar{x}) = 0$; the points $(\bar{x}, \bar{t}_k)$, $k = 1, \ldots, \ell$, are distinct points of V.
(This latter condition, which doesn't look completely equational, can be made so by adding, for each pair $\bar{t}_i, \bar{t}_j$ with $i \neq j$, a variable $z_{i,j}$ and the equation saying that $z_{i,j}$ is the inverse of one of the coordinates of $\bar{t}_i - \bar{t}_j$.) Then the projection of W_ℓ on $\bar{x}$-space will be $\bigcup_{i \geq \ell} U_i$. Its cardinality will equal

$$\sum_{i=\ell}^{N} \frac{(\ell + i)!}{i!} |U_i|.$$

But we know how to compute $\pi_{W_\ell}(t)$, hence we also know how to compute the $\pi_{U_\ell}(t)$ (we have a triangular system of N linear equations in N unknowns), and finally $\pi_U(t)$ by the above.

5 Measure, definability, and other applications

5.1. Counting points. We saw in Theorem 3.10 that every pseudo-finite field is elementarily equivalent to an ultraproduct of finite fields. This implies in fact that every pseudo-finite field elementarily embeds into an ultraproduct of finite fields (an ultrapower of ultraproducts is an ultraproduct). Every finite field can be equipped with a measure (the counting measure), and one would think that the ultraproduct of these measures might define something interesting on F. It turns out that this is the case, and we will see below how it works. The main tool is the following

Theorem 5.2. *([6]). Let $\varphi(\bar{x}, \bar{y})$ be a formula, $\bar{x}$ an n-tuple of variables ($\bar{y}$ an m-tuple of variables). Then there is a finite set $D \subset \{0, 1, \ldots, n\} \times \mathbb{Q}^{>0} \cup \{(0,0)\}$ of pairs (d, μ), and a constant $C > 0$, formulas $\varphi_{d,\mu}(\bar{y})$ for $(d, \mu) \in D$ such that:*

(1) *If $\mathbb{F}_q$ is a finite field and $\bar{a}$ an m-tuple in $\mathbb{F}_q$, then there is some $(d, \mu) \in D$ such that*

$$\big| |\varphi(\mathbb{F}_q, \bar{a})| - \mu q^d \big| < C q^{d-1/2}. \qquad (*)$$

(2) *The formula $\varphi_{d,\mu}(\bar{y})$ defines in each $\mathbb{F}_q$ the set of tuples $\bar{a}$ such that $(*)$ holds.*

I am not going to give a complete proof of this result, although I will later sketch a strategy for the proof. With some work one can show that the constant C can be found effectively, see [12], and also [14], [11]. First a few remarks.

Remarks 5.3. (1) Observe that the pair $(0,0)$ has been put in D to take care of the case when $\varphi(\mathbb{F}_q, \bar{a})$ is empty.

(2) If $\varphi(\bar{x}, \bar{a})$ defines a variety V, then this is simply the Theorem of Lang-Weil, with $d = \dim(V)$ and $\mu = 1$.

(3) Thus, if $\varphi(\bar{x}, \bar{a})$ defines an algebraic set W, all of whose irreducible components are defined over $\mathbb{F}_q$, then d will be the maximal dimension of the irreducible components of W, and μ the number of these components of maximal dimension. Note that therefore, if $\varphi(\bar{x}, \bar{y})$ is quantifier-free, then the associated set of pairs will be contained in $\{0, \ldots, n\} \times \mathbb{N}^{>0} \cup \{(0,0)\}$.

(4) If q is sufficiently large, the formulas $\varphi_{d,\mu}(\bar{y})$ will define a partition of the parameter set $\mathbb{F}_q^m$.

(5) If $n = 1$, then there are positive numbers $A \in \mathbb{N}$ and $r \in \mathbb{Q}$ such that for every $\mathbb{F}_q$ and tuple $\bar{a}$ in $\mathbb{F}_q$,

$$\text{either } |\varphi(\mathbb{F}_q, \bar{a})| < A \text{ or } |\varphi(\mathbb{F}_q, \bar{a})| \geq rq.$$

Indeed, let D be the set of pairs (d, μ) associated to $\varphi(x, \bar{y})$; define $A_0 = \sup\{\mu \mid (0, \mu) \in D\}$, $r_0 = \inf\{\mu \mid (1, \mu) \in D\}$. Let $r = r_0/2$ and $A = \sup\{A_0 + C, 4C^2/r_0^2\}$. Using $(*)$, this gives the assertion.

(6) Observe that if q is sufficiently large, $(0, \mu) \in D$ and $\mathbb{F}_q \models \varphi_{0,\mu}(\bar{a})$, then, because $q^{-1/2}$ becomes very small, and in particular $< \frac{1}{2C}$, the number μ must give the exact size of the set $\varphi(\mathbb{F}_q, \bar{a})$ defined by $\varphi(\bar{x}, \bar{a})$.

5.4. Some simple applications of this result.

(1) There is no formula of the language of rings which defines in each field $\mathbb{F}_q^2$ the subfield $\mathbb{F}_q$.

(2) We know that the multiplicative group of $\mathbb{F}_q$ is cyclic, of order $q - 1$. There is no formula which defines in all fields $\mathbb{F}_q$ the set of generators of the multiplicative group $\mathbb{F}_q^\times$.

(3) Let G, H be groups definable in the pseudo-finite field F, and assume that $f : G \to H$ is definable, $\mathrm{Ker}(f)$ is finite, and $\dim(G) = \dim(H) = d$. Then

$$\mu(G)[H : f(G)] = \mu(H)|\mathrm{Ker}(f)|.$$

(I proved a weaker version in class, with G, H connected algebraic groups of the same dimension, and f an algebraic morphism with finite kernel. It follows from this one).

(4) (not done in class)With a little more work than in (1), one can show that there is no formula allowing to interpret uniformly $\mathbb{F}_q$ in $\mathbb{F}_{q^2}$, by showing that if S is definable, and E is a definable equivalence relation on E, then the interpretation of S/E in $\mathbb{F}_{q^2}$ has size $O(q^{2d})$ for some integer d.

Proof. (1) If $\varphi(x)$ is a formula, there are $A > 0$ and $r \in \mathbb{Q}^{>0}$ such that for every finite field $\mathbb{F}_q$, the size of the set defined by φ is either $\leq A$ or greater than rq. hence, we cannot have a formula which defines in all $\mathbb{F}_{q^2}$ a set of size $q = \sqrt{q^2}$.

(2) (Not done in class, but an amusing example) The function ϕ (called the Euler function) giving the number of generating elements of a cyclic group can be computed. Note that if m, n are relatively prime integers then $\phi(nm) = \phi(n)\phi(m)$ (since $\mathbb{Z}/mn\mathbb{Z} \simeq \mathbb{Z}/n\mathbb{Z} \times \mathbb{Z}/m\mathbb{Z}$). Also, $\phi(p^n) = (p-1)p^{n-1}$, since any lifting of a generator of $\mathbb{Z}/p\mathbb{Z}$ to $\mathbb{Z}/p^n/\mathbb{Z}$ is a generator of $\mathbb{Z}/p^n\mathbb{Z}$.

First observe that if $p^n > 2$, then $\phi(p^n) \geq \sqrt{n}$. Hence, for every $A \in \mathbb{N}$, the set of integers n such that $\phi(n) < A$ is finite.

We will now show that for every $\epsilon > 0$, there is some prime power q such that $\phi(q-1) < \epsilon(q-1)$. Observe that

$$\phi(n)/n = \prod_{\ell \text{ a prime divisor of } n} (1 - 1/\ell).$$

Fix some prime p, and let $\ell_1, \dots, \ell_m$ be distinct prime numbers, $M = \prod_{i=1}^m (\ell_i - 1)$. Then for every i, we have $p^M \equiv 1 \mod (\ell_i)$ and therefore $\phi(p^M - 1) \leq (p^M - 1) \prod_{i=1}^m (1 - 1/\ell_i)$. Hence we can find arbitrarily small values of $\frac{\phi(p^M-1)}{p^M-1}$, which shows our assertion.

The existence of a formula defining the set of generators in all $\mathbb{F}_q$ would then, as in (1), contradict 5.3(5).

(3) Let $F^* = \prod_{i \in I} \mathbb{F}_{q_i}/\mathcal{U}$ be an elementary extension of F, let a be a tuple of elements of F needed to define f, G and H (and their group law, and $(a(i))_i$ a sequence such that $[a(i)_i]_{\mathcal{U}} = a$.

Let $\varphi_1(\bar{x}, \bar{a})$ be the formula defining G, $\psi_1(\bar{x}, \bar{y}, \bar{z}, \bar{a})$ the one defining its group law, $\varphi_2(\bar{x}, \bar{a})$ the formula defining H, $\psi_2(\bar{x}, \bar{y}, \bar{z}, \bar{a})$ the one defining its group law, and $\theta(\bar{x}, \bar{y}, \bar{a})$ the formula defining the graph of f. The following property is then a first order property of the parameter $\bar{a}$:

$\psi_i(\bar{x}, \bar{y}, \bar{z}, \bar{a})$ is the graph of a group operation on the set defined by $\varphi_i(\bar{x}, \bar{a})$ $(i = 1, 2)$, and $\theta(\bar{x}, \bar{y}, \bar{a})$ is the graph of a group morphism between the set defined by $\varphi_1(\bar{x}, \bar{a})$ and the set defined by $\varphi_2(\bar{x}, \bar{a})$, whose kernel is of size m.

Hence, by Łos' theorem, for a set $J \in \mathcal{U}$, we have, for all $j \in J$, that the following statement holds in $\mathbb{F}_{q_j}$:

$\psi_1(\bar{x}, \bar{y}, \bar{z}, \bar{a}(j))$ is the graph of a group operation on the set G_j defined by $\varphi_1(\bar{x}, \bar{a}(j))$, $\psi_2(\bar{x}, \bar{y}, \bar{z}, \bar{a}(j))$ is the graph of a group operation on the set H_j defined by $\varphi_2(\bar{x}, \bar{a}(j))$, and $\theta(\bar{x}, \bar{y}, \bar{a}(j))$ is the graph of a group morphism $f_j : G_j \to H_j$, whose kernel has size m.

But G_j and H_j are finite!! Hence we have $|G_j|[H_j : f_j(G_j)] = |H_j||\mathrm{Ker}(f_j)|$. For q_j sufficiently large, dividing by q_j^d, we get $\mu(G_j)[H_j : f_j(G_j)] = \mu(H_j)|\mathrm{Ker}(f)|$.

There is a first-order formula which expresses that fact, is satisfied in all $\mathbb{F}_{q_j}$ for $j \in J$, and therefore is satisfied by $\bar{a}$ in F^*, whence also in F. This gives the result.

5.5. Very rough sketch of the proof of Theorem 5.1. The result is proved by induction on the complexity of formulas.

Let us first assume that $\varphi(\bar{x}, \bar{y})$ is positive quantifier-free, that is, a disjunction of conjunction of equations (over $\mathbb{Z}$).

Let $\mathbb{F}_q$ be a finite field, and $\bar{a}$ a tuple in $\mathbb{F}_q$. Consider the set S defined by $\varphi(\bar{x}, \bar{a})$. Then $S = W(\mathbb{F}_q)$, where W is the algebraic set given by the equations of $\varphi(\bar{x}, \bar{a})$. However, we do not know that the Theorem of Lang-Weil can tell us the estimate of how many points there are: we will be able to apply this theorem only if *all irreducible components of W are defined over $\mathbb{F}_q$*. In order to be able to use Lang-Weil, we must therefore find an algebraic set W' such that $W'(\mathbb{F}_q) = W(\mathbb{F}_q)$ and all irreducible components of W' are defined over $\mathbb{F}_q$. This is done in the following fashion, which works over any field F, and is called the intersection decomposition procedure:

5.6. Intersection decomposition procedure. Let F be a (perfect) field, V an algebraic set defined over F, defined by m polynomials of degree d in n variables. Write $V = W_1 \cup \cdots \cup W_\ell$, where the W_i's are varieties (defined over the algebraic closure of F). The number ℓ is bounded in terms of the data (m, n, d), by the results of Herman 2.13. If all W_i are defined over F, we are done. Otherwise, we may assume that V is irreducible over F, and then $W_2, \dots, W_\ell$ are the conjugates of W_1 under the action of $\mathcal{G}al(F^{alg}/F)$. Then $V(F) = \bigcap_{i=1}^{\ell} W_i(F)$. So we replace V by the intersection W_1' of the W_i's. The degree of the defining polynomials of W_1' is bounded in terms of (m, n, d), and $\dim(W_1') < \dim(V)$. Repeat the procedure.

The point of all this is that given polynomials $f_i(\bar{X}, \bar{Y}) \in \mathbb{Z}[\bar{X}, \bar{Y}]$, $i = 1, \dots, m$, one can find finitely many formulas $\theta_j(\bar{y})$ and polynomials $g_{i,j,k}(\bar{X}, \bar{Y}, \bar{Z})$, $j \leq \ell_i$, $k \leq m_{i,\ell}$ such that for any field F and tuple $\bar{a}$ in F, there is an index i such that $F \models \theta_i(\bar{a})$, and the intersection decomposition procedure described above gives rise to ℓ_i varieties W_j which are defined over F by the equations $g_{i,j,k}(\bar{x}, \bar{a}, \bar{\alpha}) = 0$, where $\bar{\alpha}$ is a tuple of elements in F. (In fact the $\bar{\alpha}$ live in the algebraic closure of the field generated by $\bar{a}$, and define an extension of bounded degree). This procedure is effective, but of very high complexity: a tower of exponential of height $\dim(V)$.

5.7. Continuation of the sketch of the proof. We apply this decomposition-intersection procedure to the algebraic set W and the field $\mathbb{F}_q$, and obtain the varieties W_j's, which are defined over $\mathbb{F}_q$. Since $|W_j(\mathbb{F}_q)| \sim q^{\dim(W_i)}$ by Lang-Weil, we obtain the result, with $d = \max\{\dim(W_j)\}$, and μ the number of indices i such that $\dim(W_j) = d$. The particular shape of the W_j's only depends on the index i such that $\mathbb{F}_q \models \theta_i(\bar{a})$; moreover, one can refine θ_i so that the dimensions of the W_j's occurring in the decomposition are always the same, and this gives the formulas $\varphi_{d,\mu}$. The constant C is manufactured using the various constants C_j associated to the W_j's.

The case of a quantifier-free formula $\varphi(\bar{x}, \bar{y})$ follows, observing that modulo the theory of fields, an inequality $z \neq 0$ is equivalent to $\exists y\ yz = 1$. Thus, every quantifier-free definable set is in bijection, via a projection, with an algebraic set. We then use the first case.

Let us now assume that $\varphi(\bar{x}, \bar{y})$ is arbitrary. Then, Theorem 4.4 tells us, using compactness, that there are positive quantifier-free $\mathcal{L}_c$-formulas $\psi_1(\bar{x}, \bar{y}, \bar{z}), \ldots, \psi_m(\bar{x}, \bar{y}, \bar{z})$ such that

$$\mathrm{Psf}_c \vdash \forall \bar{x}, \bar{y}\ (\varphi(\bar{x}, \bar{y}) \leftrightarrow \exists \bar{z} \bigvee_j \psi_j(\bar{x}, \bar{y}, \bar{z})),$$

and furthermore such that for some integer N, in any field F one has

$$F \models \forall \bar{x}, \bar{y}\ (\exists \bar{z}\ \psi_j(\bar{x}, \bar{y}, \bar{z}) \to \exists^{\leq N} \bar{z}\ \psi_j(\bar{x}, \bar{y}, \bar{z})).$$

The same equivalence holds in sufficiently large finite fields, say of size $\geq C'$ for some C' (only depending on $\varphi(\bar{x}, \bar{y})$). Given some sufficiently large finite field $\mathbb{F}$ and tuple $\bar{a}$ in $\mathbb{F}$, we know by the previous steps how to estimate the size of the sets defined by the formulas $\psi_i(\bar{x}, \bar{a}, \bar{z})$. The problem is that the set defined by $\bigvee_j \psi_j(\bar{x}, \bar{a}, \bar{z})$ is not in bijection with the set defined by $\varphi(\bar{x}, \bar{a})$: given some $\bar{x}$ in that set, there may be several $\bar{z}$ such that $\psi_j(\bar{x}, \bar{a}, \bar{z})$ holds. As with Kiefe's result on the Poincaré series, one uses a trick to transform the algebraic sets defined by the ψ_j, in such a way that we are able to count how many $\bar{z}$ are sitting above an $\bar{x}$. Then we use some counting arguments and induction to conclude.

5.8. Definition of the measure on pseudo-finite fields. Let $\varphi(\bar{x}, \bar{y})$ be a formula ($\bar{x}$ an n-tuple of variables), and D, $\varphi_{d,\mu}(\bar{y})$ the set and formulas given by Theorem 5.1. It follows from Remark 5.3(6) that if F is a pseudo-finite field and $\bar{a}$ a tuple in F, then there will be a unique pair $(d, \mu) \in D$ such that $F \models \varphi_{d,\mu}(\bar{a})$. We then define $\dim(\varphi(\bar{x}, \bar{a})) = d$ and $\mu(\varphi(\bar{x}, \bar{a})) = \mu$. If S is the set defined by $\varphi(\bar{x}, \bar{a})$, then we also write $\dim(S)$ and $\mu(S)$ respectively.

Proposition. Let F be a pseudo-finite field, S, T two definable sets.

(1) If V is a variety defined over F, then $\dim(V(F)) = \dim(V)$ and $\mu(V(F)) = 1$.

(2) Assume that $T \cap S = \emptyset$. Then

$$\mu(S \cup T) = \begin{cases} \mu(S) + \mu(T) & \text{if } \dim(S) = \dim(T), \\ \mu(S) & \text{if } \dim(S) > \dim(T), \\ \mu(T) & \text{if } \dim(S) < \dim(T). \end{cases}$$

(3) Assume that $f : S \to T$ is a definable function, which is onto. If for all $\bar{y} \in T$, $\dim(f^{-1}(\bar{y})) = d$, then $\dim(S) = \dim(T) + d$. If moreover for every $\bar{y} \in T$, $\mu(f^{-1}(\bar{y})) = m$, then $\mu(S) = m\mu(T)$.

(4) Let us define a function m_S on definable subsets of S as follows. Assume that $T \subseteq S$ is definable, and let $(d, \mu) = (\dim(S), \mu(S))$, $(e, \nu) = (\dim(T), \mu(T))$. Then

$$m_S(T) = \begin{cases} 0 & \text{if } e < d, \\ \nu/\mu & \text{if } d = e. \end{cases}$$

Then m_S is a finitely additive measure on the set of definable subsets of S.

(5) Let $\bar{S}$ be the Zariski closure of S (in F^{alg}. I.e., the smallest Zariski closed set containing S. It is defined over F). Then $\dim(S) = \dim(\bar{S})$. [That is, we are saying that the algebraic dimension of the algebraic set $\bar{S}$ coincides with the model-theoretic dimension of the set S]

Proof. (1) is clear.

Recall that F embeds elementarily in some ultraproduct $\prod_{q\in\mathcal{Q}} \mathbb{F}_q/\mathcal{U}$ of finite fields. Assume that S is defined by $\varphi(\bar{x}, \bar{a})$, write $\bar{a} = [\bar{a}_q]_{\mathcal{U}}$, and S_q for the subset of $\mathbb{F}_q^n$ defined by $\varphi(\bar{x}, \bar{a}_q)$. Note that for some set $A \in \mathcal{U}$, we will then have $\mathbb{F}_q \models \varphi_{d,\mu}(\bar{a}_q)$ for all $q \in A$, and therefore $|S_q| \sim \mu q^d$. A moment's thought shows that this gives items (2) - (4).

(5) By 4.4, there is an algebraic set $W(F) \subset F^{n+\ell}$ such that $S = \pi(W(F))$ and the restriction of the projection π to W is finite-to-one. Without loss of generality, $W(F)$ is Zariski dense in W, and by (3) we obtain that $\dim(W) = \dim(S)$. Working now in F^{alg}, we have that π is also finite-to-one on a Zariski-dense open subset of W, and therefore $\dim(W) = \dim(V)$ (algebraic dimensions). Since $V \supseteq \bar{S}$, we get that $\dim(V) = \dim(\bar{S})$.

5.9. The Zariski topology revisited. Recall that if K is a field, we define on K^n the *Zariski topology*, by taking as a basis of closed sets the subsets of K^n defined by polynomial equations. This topology is Noetherian, and finitely many equations suffice to define a closed set.
Suppose now that L is a field extension of K. Then L^n has its own Zariski topology, and one can wonder whether the topology it induces on the subset K^n coincides with the Zariski topology of K^n. It does! Let $S \subset L^n$ be defined by the equations $f_1(\bar{x}) = \cdots = f_m(\bar{x}) = 0$. Fix a basis B of the K-vector space L, and write

$$f_i(\bar{x}) = \sum_{b\in B} g_{i,b}(\bar{x})b$$

where the $g_{i,b}$ have their coefficients in K. (Note that all but finitely many of them are 0). Then for any point $\bar{a} \in K^n$, we have the following, for every i:

$$f_i(\bar{a}) = 0 \iff \sum_b g_{i,b}(\bar{a})b = 0 \iff g_{i,b}(\bar{a}) = 0\,\forall b \in B.$$

This proves our assertion.

5.10. Zariski closure. Let K be a field, and $S \subset K^n$. The Zariski closure of S, $\bar{S}$ is defined as the smallest Zariski closed set of K^n which contains S. In other words, if I is the ideal of polynomials with coefficients in K which vanish on all elements of S, then $\bar{S}$ is defined as the zero-set of a set of generators of I. I claim that if $S \subset K_0^n$ for some (perfect) subfield K_0 of K, then $\bar{S}$ can be defined over K_0. Indeed, by the preceding paragraph 5.9, we may assume that $K = \Omega$ is algebraically closed. But if ρ is any automorphism of Ω which is the identity on K_0, then $\rho(S) = S$, and therefore $\rho(\bar{S}) = \bar{S}$, i.e., $\bar{S}$ is definable over K_0 (by elimination of imaginaries in ACF). So, the concept of Zariski closure of a set does not depend on the field in which we work, as soon as it contains the coordinates of the points of our set. One says that a set S is Zariski dense inside an algebraic set V if $\bar{S} = V$.

5.11. Lemma. Let F be a PAC field, and V a variety defined over F. Them $V(F)$ is Zariski dense in V.

Proof. Let $g(\bar{x}) \in K[\bar{x}]$, and suppose that g does not vanish identically on $V(\Omega)$ (Ω a large algebraically closed field containing F). Then the algebraic set W defined by $\bar{x} \in V, g(\bar{x})y-1=0$ is also a variety, defined over K, and therefore $W(F) \neq \emptyset$. Hence, $V(F)$ is not contained in the hypersurface defined by $g(\bar{x})=0$, i.e., its Zariski closure is all of V.

5.12. Existence of certain bounds. Let $\varphi(\bar{x}, \bar{y})$ be a formula.

(1) (Not the Strict Order Property) There is a number M such that in any finite or pseudo-finite field F, the length of a chain of definable subsets of F^n defined by formulas $\varphi(\bar{x}, \bar{a})$ for some tuples $\bar{a}$ in F, is bounded by M.

(2) (Finite Shelah rank) There is a number M such that in any finite field or pseudo-finite field F, if S is a definable set and $(\bar{a}_i)_{i\in I}$ is a set of tuples such that each $\varphi(\bar{x}, \bar{a}_i)$ defines a subset of S of the same dimension d as S, and for $i \neq j$, $\dim(\varphi(\bar{x}, \bar{a}_i) \wedge \varphi(\bar{x}, \bar{a}_j)) < d$, then $|I| \leq M$.

Proof. These two facts follow from general properties of measures. It suffices to show them for all pseudo-finite fields, since then they will be true in all sufficiently large finite fields, whence, taking into account the finitely many small finite fields, we will get the bound M.

(1) Assume that this is not the case, i.e., that there are such chains of arbitrarily large length. Then, going to a sufficiently saturated pseudo-finite field F, we can find a sequence $(\bar{a}_i)_{i\in\mathbb{N}}$ of tuples in F such that if $i<j$ then the set S_j defined by $\varphi(\bar{x}, \bar{a}_j)$ is strictly contained in the set S_i defined by $\varphi(\bar{x}, \bar{a}_i)$. Let D be the finite set of pairs associated to φ. Because D is finite, we may, going to a subsequence, assume that for every $i \in \mathbb{N}$, $\dim(S_i)=d$ and $\mu(S_i)=\mu$. The proof is by induction on d.

If $d=0$, then we know that μ is the size of the set S_i, and therefore $|I|=1$. Assume $d>0$ and that the result holds for all definable sets of smaller dimension. For $i>0$ let $T_i = S_0 \setminus S_i$. Then the sets T_i, $i \in \mathbb{N}$, form a strictly increasing chain of subsets of S_0, and we have $\dim(T_i) < d$ (since $(\dim(S_i), \mu(S_i)) = (\dim(S_0), \mu(S_0))$). This contradicts the induction hypothesis and proves the result.

(2) Let D be the set of pairs associated to the formula $\varphi(\bar{x}, \bar{y})$, and let ν be the inf of all μ such that $(d,\mu) \in D$. If $\varphi(\bar{x}, \bar{a}_i)$, $i \in I$, define subsets S_i of S such that $\dim(S_i)=d$ ($=\dim(S)$) and $\dim(S_i \cap S_j) < d$, then we get $m_S(S_i) \geq \nu/\mu(S)$ and $m_S(S_i \cap S_j)=0$. This gives $|I| \leq \mu(S)/\nu$.

Definition 5.13. Recall that a formula $\varphi(\bar{x}, \bar{y})$ has the *independence property* (in the model M) iff for every n, there are tuples $\bar{a}_i$, $1 \leq i \leq n$, and $\bar{b}_s$, $s \in \mathcal{P}(\{1,\dots,n\})$, in M such that

$$M \models \varphi(\bar{a}_i, \bar{b}_s) \iff i \in s$$

for every i and s. A complete theory has the independence property if there is a formula which has the independence property.

Theorem 5.14. *(Duret [9]) The theory of any pseudo-algebraically closed field which is not separably closed, has the independence property.*

I will give here a simple case of his proof, for a pseudo-finite field of characteristic $\neq 2$. (In characteristic 2, the example can be modified). Let F be a pseudo-finite field and consider the formula $\varphi(x, y)$ which says that $x+y$ is a square and $x \neq y$. Let $a_1, \ldots, a_n$ be distinct elements of F, s a subset of $\{1, \ldots, n\}$, we want to find an element b such that $b + a_i$ is a square if and only if $i \in s$. Renumbering the a_i's, we may assume that $i \in s \iff i \leq r$.

Because F is pseudo-finite, it contains an element c which is not a square. Then, as F has a unique extension of degree 2, we have $F^\times = F^{\times 2} \cup cF^{\times 2}$, and therefore

$$F \models \forall x\, [(\forall y\; y^2 \neq x) \leftrightarrow (\exists y\; y^2 = cx)].$$

Let t be transcendental over F, and consider the extension

$$L = F(t, \sqrt{t + a_1}, \ldots, \sqrt{t + a_r}, \sqrt{c(t + a_{r+1})}, \ldots, \sqrt{c(t + a_n)}).$$

Then $L \cap F^{alg} = F$ (This needs a proof which I will not give). Hence, by Lemma 2.10, in F there is an element d such that $d + a_1, \ldots, d + a_r, c(d + a_{r+1}), \ldots, c(d + a_n)$ are squares. I.e., $d + a_i$ is a square if and only if $i \leq r$.

5.15. Graphs interpretable in pseudo-finite fields. The above proof shows that the random graph is interpretable in any pseudo-finite field (of characteristic $\neq 2$), by the formula expressing that $x + y$ is a square and $x \neq y$.

Observe that if -1 is a square in F, then the formula $\psi(x, y)$ saying that $x - y$ is a square and is non-zero would work as well. If -1 is not a square in F, then the formula $\psi(x, y)$ defines the random tournament. (A *tournament* is a binary relation not intersecting the diagonal and such that given two distinct elements, one exactly of (a, b), (b, a) is in the relation. The *random tournament* is a tournament in which given any two disjoint finite sets A and B there is an element c such that $\bigwedge_{a \in A} R(c, a) \wedge \bigwedge_{b \in B} R(b, c)$.

Hrushovski proves in [19] that one cannot interpret the *random triangle-free* graph in any countable pseudo-finite field. Beyarslan proves in [2] that one can interpret in any pseudo-finite field the random n-hypergraph. Recall that if $n \geq 2$, an *n-hypergraph* is an n-ary relation R satisfying:

– $R(x_1, \ldots, x_n) \to \bigwedge_{i \neq j} x_i \neq x_j$,

– $R(x_1, \ldots, x_n) \to \bigwedge_{\sigma \in \mathrm{Sym}(\{1, \ldots, n\})} R(x_{\sigma(1)}, \ldots, x_{\sigma(n)})$.

The *random n-hypergraph* is the existentially closed countable n-hypergraph. I.e., it is countable, and satisfies, for all m, ℓ,

if $a_1, \ldots, a_m, b_1, \ldots, b_\ell$ are distinct $(n-1)$-element subsets, then there is an element x such that $\bigwedge_i R(x, a_i) \wedge \bigwedge_j \neg R(x, b_j)$.

5.16. Another interesting result, in the vein of 5.4. Say that $I \subset \mathbb{F}_p$ is an interval (p a prime) if there is some interval J in $\mathbb{Z}$ such that I is the image of J under the natural reduction modulo $p : \mathbb{Z} \to \mathbb{Z}/p\mathbb{Z} = \mathbb{F}_p$.

Proposition (Kowalski [27]). If $\varphi(x)$ is a formula of the language of rings which defines in all prime field $\mathbb{F}_p$ an interval, then there is a number N such that for every prime p, one of $|\varphi(\mathbb{F}_p)|$, $|\neg\varphi(\mathbb{F}_p)|$ has size $\leq N$.

Note that an argument using measures is not sufficient, since the interval $[0, \frac{p-1}{2}]$ has size approximately $\frac{p}{2}$.

6 Generalizations, and various applications

Definition 6.1. We say that a class $\mathcal{C}$ of finite structures (in a language $\mathcal{L}$) is *N-dimensional asymptotic* if for every formula $\varphi(\bar{x}, \bar{y})$, there is a finite set D of pairs $(d, \mu) \in \{0, 1, \dots, N|\bar{x}|\} \times \mathbb{R}^{>0} \cup \{(0, 0\}$ and a constant $C > 0$ such that for every structure M in $\mathcal{C}$ and tuple $\bar{b}$ in M, there is a pair $(d, \mu) \in D$ such that

$$\big||\varphi(M, \bar{b})| - \mu|M|^{d/N}\big| < C|M|^{d/N-1/2}. \tag{$*$}$$

Furthermore, for each pair (d, μ) in D there is a formula which defines in any structures of $\mathcal{C}$ the set of tuples for which $(*)$ holds.

This definition was introduced by Macpherson and Steinhorn ([32]) for $N = 1$, and extended by Elwes to arbitrary N, although I didn't quite take his definition: he writes $< o(|M|)^{d/M}$. For more details one can consult the survey paper by Elwes and Macpherson [10]. I will concentrate on the case $N = 1$. As with finite fields, this allows to define a relative measure on ultraproducts of structures of $\mathcal{C}$, which satisfies the usual properties of a measure. This can be in fact generalised as follows, by picking out the important properties of a measure ([32]):

6.2. Measurable structures. A structure M is *measurable* if there is a function $h = (\dim, \mu)$ (dimension and measure) from the set $\mathrm{Def}(M)$ of definable (with parameters in M) subsets of cartesian powers of M, taking values in $\mathbb{N} \times \mathbb{R}^{>0} \cup \{(0, 0)\}$, and satisfying the following conditions:

(1) For every formula $\varphi(\bar{x}, \bar{y})$ there is a finite set $D = D_\varphi$ such that for any $\bar{a}$ in M $h(\varphi(M, \bar{a})) \in D$.

(2) If S is finite (or empty), then $h(S) = (0, |S|)$,

(3) For every formula φ, and $(d, \mu) \in D_\varphi$, the set $\{\bar{a} \mid h(\varphi(M, \bar{a})) = (d, \mu)\}$ is definable in M (without parameters).

(4) (Additivity) Let S, T be disjoint definable subsets of M^n. Then $\dim(S \cup T) = \sup\{\dim(S), \dim(T)\}$, and

$$\mu(S \cup T) = \begin{cases} \mu(S) + \mu(T) & \text{if } \dim(S) = \dim(T), \\ \mu(S) & \text{if } \dim(S) > \dim(T), \\ \mu(T) & \text{if } \dim(S) < \dim(T). \end{cases}$$

(5) (Fubini) Assume that $f : S \to T$ is a definable function, which is onto. If for all $\bar{y} \in T$, $\dim(f^{-1}(\bar{y})) = d$, then $\dim(S) = \dim(T) + d$. If moreover for every $\bar{y} \in T$, $\mu(f^{-1}(\bar{y})) = m$, then $\mu(S) = m\mu(T)$.

6.3. Comments. The motivating example is pseudo-finite fields. Because of the definability condition and the Fubini condition, note that measurability is preserved by elementary equivalence. Also, one can restrict one's attention to definable subsets of M when checking whether a structure is measurable or not: if a structure satisfies the conditions for formulas $\varphi(x, \bar{y})$ with $|x| = 1$, then a simple induction argument shows it also satisfies it for arbitrary formulas. As with pseudo-finite fields, it follows that the theory of a measurable structure is supersimple.

If $\mathcal{C}$ is an N-dimensional asymptotic class of finite structures, then any structure in the elementary class generated by $\mathcal{C}$ will be measurable, and we also see that the definitions and measures will be uniform through the structures in this elementary class.

The converse is however not true: there exists measurable structures which are not elementarily equivalent to ultraproducts of finite structures. Below (6.6) we will give such an example.

6.4. Definitions.

(1) Recall that a first order theory T is *strongly minimal* if in any model M of T, any definable (with parameters) subset of M is finite or cofinite.

(2) In a strongly minimal theory one can define a rank, called the *Morley rank* as well as a multiplicity, the *Morley degree*, of definable sets. I will not give precise definitions, let me say that the rank satisfies the obvious axioms of a dimension (see properties (4) and (5) of dim in 6.2), and that the Morley degree of a definable set S is the maximal number n such that S can be definably partitioned into sets of the same rank as S. Furthermore, the Morley rank of the universe is 1, and of a finite set is 0. A strongly minimal theory T has the DMP (*definable multiplicity property*) if for any model M, formula $\varphi(\bar{x}, \bar{y})$ and integer $n > 0$, the set of tuples $\bar{a}$ in M such that $\varphi(M, \bar{a})$ has Morley degree n, is definable.

Important examples of strongly minimal theories with the DMP are the completions of the theory ACF of algebraically closed fields. Moreover, if T_1 and T_2 are strongly minimal with the DMP, then so is the Hrushovski fusion T_3 constructed from T_1 and T_2.

Theorem 6.5. *(Ryten-Tomasic [39]) Let T be a strongly minimal theory with the DMP and which eliminates imaginaries. Consider the theory T_σ of models of T with an automorphism σ (so, structures in the language $\mathcal{L}$ to which one has added a unary function symbol σ), let N be an existentially closed model of T_σ, and let F be the $\mathcal{L}$-structure $\{a \in N \mid \sigma(a) = a\}$. Then F is measurable of dimension* 1.

6.6. Example of a measurable structure not arising from an asymptotic class of finite structures. So, let T_1 be the theory of algebraically closed fields of characteristic 2 (in a language $\mathcal{L}_1$), and T_2 the theory of algebraically closed fields of characteristic 3 (in the language $\mathcal{L}_2$). Let T_3 be the Hrushovski fusion of T_1 and T_2, and let F be as above. Then the reduct of F to $\mathcal{L}_1$ is a pseudo-finite field of characteristic 2, the reduct of F to $\mathcal{L}_2$ is a pseudo-finite field of characteristic 3, and F cannot be elementarily equivalent to an ultraproduct of finite $\mathcal{L}_1 \cup \mathcal{L}_2$-structures: a power of 2 can never equal a power of 3 unless they both equal 1. But, by the result of Ryten-Tomasic, the $\mathcal{L}_1 \cup \mathcal{L}_2$-structure F is measurable of dimension 1.

6.7. An aside - Zilber's conjecture. Hrusvhoski's construction gives a counterexample to *Zilber's conjecture*, which I will recall below. In a strongly minimal structure, the algebraic closure (acl) allows to define a pregeometry: say that a subset A of M is independent if whenever $a \in A$ then $a \notin \mathrm{acl}(A \setminus \{a\})$. The pregeometry is *trivial* iff whenever $A \subseteq M$, then $\mathrm{acl}(A) = \bigcup_{a \in A} \mathrm{acl}(a)$; it is *locally modular* or *modular* if whenever A and B are algebraically closed with non-empty intersection, then they are independent over their intersection. (Where independance is defined in terms of the pregeometry).
To a pregeometry one associates a geometry as follows: we consider the quotient of $M \setminus \mathrm{acl}(\emptyset)$ by the equivalence relation $a \sim b \iff a \in \mathrm{acl}(b)$ ($\iff b \in \mathrm{acl}(a)$). Two pregeometries are equivalent if they have the same associated geometry. Here is a statement of Zilber's conjecture, more precise than the one I gave in class:

(Zilber's conjecture) *Let M be strongly minimal. Then M is geometrically equivalent to one of the following structures:*

(1) *(Trivial pregeometry) A set with no structure.*
(2) *(The pregeometry on M is locally modular, and non-trivial) A vector space $(V, +, \alpha\cdot -)_{\alpha \in \Delta}$, where Δ is a division ring, and $\alpha\cdot$ denotes scalar multiplication by $\alpha \in \Delta$.*
(3) *(The pregeometry is* **not** *locally modular). An algebraically closed field of characteristic p or 0, with maybe additional constants.*

Hrushovski disproved the conjecture in two papers ([17] and [18]) by exhibiting a strongly minimal set living exactly between (1) and (2), and by exhibiting a strongly minimal non-locally modular strict expansion of an algebraically closed field, namely, a structure with two distinct structures of algebraically closed field. Later, together with Zilber, he exhibited a set of axioms (*Zariski geometries*) which, when satisfied by the geometry, guarantee that the conclusion of Zilber's conjecture holds ([25]). This result on Zariski geometries has been applied outside the strongly minimal context, for sets of "dimension 1", yielding a dichotomy locally modular/field.

6.8. Examples of finite dimensional asymptotic classes. (Not done in class) By Theorem 5.1, the collection of all finite fields forms a 1-dimensional asymptotic class, as does any subclass. Also, for a fixed $n > 1$, the class of all $\mathrm{GL}_n(\mathbb{F}_q)$ is an (n^2)-dimensional asymptotic class. The definability assumption comes from the fact that there is a uniform interpretation of the field $\mathbb{F}_q$ in these groups ([36]). Here is another example

(Ryten [38]) Fix a prime p, and relatively prime integers m, n with $m \geq 1$ and $n > 1$. Let $\mathcal{C}_{(m,n,p)}$ be the class of all fields $\mathbb{F}_{p^{kn+m}}$ with a distinguished automorphism Frob^k, for $k \in \mathbb{N}^{>0}$. One can show that there is no formula of the field language which defines in each field $\mathbb{F}_{p^{kn+m}}$ the graph of Frob^k. These structures appear in a significant way in the study of certain finite simple groups: for instance $\mathcal{C}_{(1,2,2)}$ is uniform parameter biinterpretable with the classes of

Suzuki groups ${}^2B_2(2^{2k+1})$ and the Ree groups ${}^2F_4(2^{2k+1})$, and $\mathcal{C}_{(1,2,3)}$ with the class of Ree groups ${}^2G_2(3^{2k+1})$.

6.9. Consequences of measurability. As with pseudo-finite fields, one gets the following two results:

(1) (Not the strict order property): if $\varphi(\bar{x},\bar{y})$ is a formula, then there is a bound on the length of a (strictly) decreasing chain of subsets of M^n defined by formulas $\varphi(\bar{x},\bar{a}_i)$.

(2) (the S1 property): if $\varphi(\bar{x},\bar{y})$ and $\psi(\bar{x},\bar{z})$ are formulas, there is a number N such that if $\bar{a}\in M$ and $\dim(\varphi(M,\bar{a}))=n$, then there are at most N tuples $\bar{b}_j$ such that for every $i\neq j$ one has

$$\dim(\varphi(M,\bar{a})\cap\psi(M,\bar{b}_i))=n>\dim(\varphi(M,\bar{a})\cap\psi(M,\bar{b}_i)\cap\psi(M,\bar{b}_j)).$$

Proof. Exercise.

Definition 6.10. Let M be a structure, $\varphi(\bar{x},\bar{y})$ a formula. Then the formula $\varphi(\bar{x},\bar{y})$ is *stable* (with respect to the complete theory $\mathrm{Th}(M)$) if there is a number N such that whenever $(\bar{a}_i,\bar{b}_i)$ are tuples for $1\leq i\leq k$ which satisfy

$$M\models\varphi(\bar{a}_i,\bar{b}_j)\iff i<j,$$

then $k\leq N$.

Remarks 6.11. The stability of the formula depends on the partition of the set of variables. Stability is preserved under Boolean combinations. Furthermore, if $\varphi(\bar{x},\bar{y})$ is stable, then so is its "transpose" $\varphi^t(\bar{y},\bar{x})$.

Proposition 6.12. *Let M be a measurable structure, and $\varphi(\bar{x},\bar{y})$, $\psi(\bar{x},\bar{z})$ formulas and n an integer such that for any tuples $\bar{b}$ and $\bar{c}$ in M, we have* $\dim(\varphi(M,\bar{b}))\leq n$, $\dim(\psi(M,\bar{c}))\leq n$. *Then the formula $\delta(\bar{y},\bar{z})$ defined by $M\models\delta(\bar{a},\bar{b})$ if and only if* $\dim(\varphi(M,\bar{a})\wedge\psi(M,\bar{b}))<n$, *is stable.*

Proof. Suppose this is not the case. We may assume that M is sufficiently saturated. Then we can find an infinite sequence $(\bar{a}_i,\bar{b}_i)_{i\in\mathbb{N}}$ such that

$$M\models\delta(\bar{a}_i,\bar{b}_j)\ \text{ if and only if } i\leq j.$$

The usual Ramsey type argument and compactness allow us to find an indiscernible sequence $(\bar{a}_i,\bar{b}_j)$ indexed by the elements of $\mathbb{Z}$ and satisfying the same condition. There are two cases to consider.
Case 1: $\dim(\varphi(M,\bar{a}_1)\wedge\varphi(M,\bar{a}_2)\wedge\psi(M,\bar{b}_3))=n$.
Then $\dim(\varphi(M,\bar{a}_1)\wedge\varphi(M,\bar{a}_i)\wedge\psi(M,\bar{b}_{i+1}))=n$ whenever $i>1$. However, we have $\dim(\varphi(M,\bar{a}_i)\wedge\psi(M,\bar{b}_{i+1})\wedge\varphi(M,\bar{a}_j)\wedge\psi(M,\bar{b}_{j+1}))<n$ whenever $1<i+1<j$. Hence each member of the family $\varphi(M,\bar{a}_1)\wedge\varphi(M,\bar{a}_{2i})\wedge\psi(M,\bar{b}_{2i+1})$ $(i>0)$ has dimension n, but the intersection of any two of them has dimension $<n$. This contradicts 5.12(2).
Case 2: $\dim(\varphi(M,\bar{a}_1)\wedge\varphi(M,\bar{a}_2)\wedge\psi(M,\bar{b}_3))<n$.
Then $\dim(\varphi(M,\bar{a}_i)\wedge\varphi(M,\bar{a}_j)\wedge\psi(M,\bar{b}_3))<n$ for all $i<j<3$. But $\dim(\varphi(M,\bar{a}_i)\wedge\psi(M,\bar{b}_3))=n$ for all $i<3$, and this gives us again a contradiction.

6.13. Comments. This result (which appears in [23] for pseudo-finite fields and some other PAC fields) is the key to the study of definable groups using stability-theoretic techniques. Indeed, if Δ is a (finite or infinite) set of stable formulas (with fixed distinguished variables $\bar{x}$, and arbitrary parameter variables $\bar{y}$), we know that any Boolean combination of formulas from Δ is stable, and we call these formulas Δ-formulas. A Δ-type (in the variables $\bar{x}$) over a set A is a maximal consistent set of Δ-formulas $\varphi(\bar{x}, \bar{a})$ where $\bar{a}$ in A. By stability, it is definable (over models). Here are a few results by Hrushovski and Pillay on definable groups and definable subgroups of algebraic groups:

Applications to groups

Theorem 6.14. *(Hrushovski-Pillay, [22], Thm C) Let G be a group definable in a pseudo-finite field F. Then there is a definable subgroup G_1 of G, an algebraic group H defined over F, and a definable group homomorphism $G_1 \to H(F)$ with finite kernel.*

6.15. Algebraic groups. An algebraic group is an algebraic set G equipped with two morphisms (maps defined everwhere by rational functions) $G \times G \to G$ and $G \to G$ defining respectively the group multiplication and the inverse map. These maps are continuous for the Zariski topology. If it is not a variety, then the irreducible component of G which contains the identity is a subgroup of finite index, denboted by G^0, and will always be defined over the field of definition of G. If $G = G^0$, one says that G is *connected.*

Typically we will be working with affine groups, and a result of Chevalley says that they are algebraic subgroups of $\mathrm{GL}_n(\Omega)$ for some n (GL_n are the invertible n by n matrices). One has the following easy results:

Remarks 6.16. Let F be a pseudo-finite field, G an algebraic group defined over F.

(1) If G is a variety, then $G(F)$ is Zariski dense in G. If $U \subset G$ is Zariski open, then every element of $G(F)$ is in $U \cdot U^{-1}$.

(2) If G is a variety and $g \in Z(G(F))$, then $g \in Z(G)$. (If H is a group, $Z(H)$ is the *center* of H, i.e., $\{g \in H \mid \forall h \in H\, g^{-1}h^{-1}gh = 1\}$.)

(3) H a definable subgroup of $G(F)$. Then $\dim(H) = \dim(G)$ if and only if $[G(F) : H] < \infty$.

(4) Let H be another algebraic group defined over F, assume that G and H are connected, have the same dimension, and that $f : H \to G$ is a morphism (of algebraic groups) with finite kernel. Then
$$|\mathrm{Ker}(f) \cap H(F)| = [G(F) : f(H(F))].$$

Proof. The first part of (1) follows from 5.11; for the second, let $g \in G(F)$, and consider the Zariski open set $U \cap gU$: if $h \in U(F) \cap gU(F)$, and $u = g^{-1}h$, then $u \in U$, $h \in U$, and $g = hu^{-1}$. For (2), note that commuting with g is a Zariski closed condition, which defines a subgroup of G, hence the result follows from (1). Items (3) and (4) are easy consequences of the properties of dimension and measure. Note the following consequence: if $\mathrm{char}(F) \neq 2$, then the map $x \mapsto x^2$ is never surjective in a pseudo-finite field.

Theorem 6.17. *(Prop. 2.1 in [23]) Let F be a pseudo-finite field, G an algebraic group defined over F, and S_i, $i \in I$, a family of definable subsets of $G(F)$ which contain 1, and whose Zariski closure $\bar{S}_i$ is a variety. Let H be the subgroup of G generated by the S_i's. Then its Zariski closure $\bar{H}$ coincides with the subgroup generated by the $\bar{S}_i$'s, is connected, and H has finite index in $\bar{H}(F)$. Moreover, there are indices $i_1, \dots, i_m$ (maybe with repetitions) and $\varepsilon_1, \dots, \varepsilon_m \in \{+1, -1\}$ such that*

$$H = S_{i_1}^{\varepsilon_1} \cdots S_{i_m}^{\varepsilon_m},$$

and in particular H is definable. (Here I mean that every element of H is a product $g_1 g_2 \cdots g_m$ with $g_j \in S_{i_j}^{\varepsilon_j}$.

Sketch of proof. Consider the set of all products of the form $\bar{S}_{i_1}^{\varepsilon_1} \cdots \bar{S}_{i_m}^{\varepsilon_m}$. These sets are varieties, and therefore attain a maximal dimension ($\leq \dim(G)$). We choose indices $i_1, \dots, i_m$ and exponents $\varepsilon_1, \dots, \varepsilon_m$ at which this maximum is attained, to give the variety $\bar{S}$. Then $\bar{S}$ is a group - since it is stable under multiplication by elements of $\bar{S}_i^{\pm 1}$, $i \in I$, and therefore is the group generated by all $\bar{S}_i$'s. (This is a classical result in algebraic geometry, due to Chevalley). We now consider the corresponding definable set $U = S_{i_1}^{\varepsilon_1} \cdots S_{i_m}^{\varepsilon_m}$. An easy argument (for instance passing to a saturated elementary extension and taking "generic" elements of the S_{i_j}'s) shows that its Zariski closure $\bar{U}$ has the same dimension as $\bar{S}$, and therefore equals $\bar{S}$. One then shows (I will not do it, but this is where the techniques of stability come into play - see 1.13 in [23] for a proof) that a finite number of products of $U \cup U^{-1}$ define a subgroup H_0 of $\bar{H}(F)$, and therefore that H_0 is definable. As $[H(F) : H_0] < \infty$ (they have the same dimension), and $H_0 \subset \langle S_i \mid i \in I \rangle$, it follows that the group $\langle S_i \mid i \in I \rangle$ is definable, as it is a finite union of translates of the definable group H_0.

Corollary 6.18. *([23]) Let F be a pseudo-finite field, G an algebraic group defined over F, and S_i, $i \in I$, definable subsets of G. Then there is a definable group H contained in $\langle S_i \mid i \in I \rangle$ such that for each $i \in I$, HS_i is the union of finitely many cosets of H.*

Proof. Without loss of generality, we may assume that the Zariski closures of the S_i's are varieties. If $1 \notin S_i$, define $T_i = g_i^{-1} S_i$ for some $g_i \in S_i$, if $1 \in S_i$ let $T_i = S_i$, and let H be the group generated by the T_i's. By Theorem 6.17, we know that H is definable, and it is contained in $\langle S_i \mid i \in I \rangle$. Moreover, each set HS_i has the same dimension as H, is definable, and therefore is covered by a finite number of cosets of H (since distinct cosets are disjoint).

Corollary 6.19. *([23]) Let G be an algebraic group defined over the pseudo-finite field F. If $G(F)$ is definably simple (i.e., every definable normal subgroup of $G(F)$ is either* (1) *or $G(F)$), and non-abelian, then $G(F)$ is simple (as an abstract group).*

Proof. First of all, note that necessarily $G = G^0$, since we saw that G^0 is also definable. Also, $Z(G) = (1)$: it is definable, normal, and $\neq G(F)$ since G is non-abelian. If $1 \neq g \in G(F)$, consider the centraliser $C(g) = \{h \in G(F) \mid hg = gh\}$; if it has finite index in $G(F)$, then so does the intersection H of its conjugates; but H is normal and definable, hence $H = G(F)$, and $g = 1$ because G has no center. Hence, for any element $1 \neq g \in G(F)$, $[G(F) : C(g)] = \infty$.

Equivalently, if $1 \neq g \in G(F)$, the set $g^{G(F)} = \{h^{-1}gh \mid h \in G(F)\}$ is infinite.
Assume by way of contradiction that H is a normal infinite subgroup of $G(F)$, $H \neq G(F)$. Choose $1 \neq g \in H$, consider the set $Y = g^{-1}g^{G(F)}$. By the discussion above, Y is infinite. Also, observe that the Zariski closure $\bar{X}$ of $g^{G(F)}$ is a variety: this is because $g^{G(F)}$ is the image of $G(F)$ under a morphism, is "isomorphic" to $G(F)/C(g)$, and because G is connected (Exercise). Hence, so is the Zariski closure $\bar{Y}$ of Y; as it contains 1, we may apply Theorem 6.17, and obtain that the group H_0 generated by Y is definable. Observe that H_0 is normal, contained in H, and infinite, and this gives the desired contradiction.

6.20. Another amusing application. We work in $\mathrm{GL}_n(K) \subset \mathrm{Mat}_n(K)$, K any field. Recall that a matric A is *nilpotent* if $A^n = 0$, and is *unipotent* if $(A - I_n)^n = 0$. If $\mathrm{char}(K) = p > n$ and A is unipotent, then $A^p = I_n$.
If A is unipotent, define

$$\log(A) = \sum_{i=1}^{n-1} (-1)^{i+1} \frac{(A - I_n)^i}{i},$$

and if B is nilpotent, then define

$$\exp(B) = \sum_{i=0}^{n-1} \frac{B^i}{i!}.$$

The maps log and exp define bijections between the set of unipotent matrices of $\mathrm{GL}_n(K)$ and the set of nilpotent matrices of $\mathrm{Mat}_n(K)$. They are inverse of each other.
Note that when A and B are unipotent and commute, then $\log(AB) = \log(A) + \log(B)$. Also, the algebraic set

$$X(A, K) = \{\exp(t \log(A)) \mid t \in K\}$$

is a subgroup of $\mathrm{GL}_n(K)$, which is isomorphic to the additive group K. Moreover, in characteristic $p > 0$, one has $X(A, \mathbb{F}_p^{alg}) \cap \mathrm{GL}_n(\mathbb{F}_p) = \langle A \rangle$.

Theorem 6.21. *(4.3 in [23]) Let $n > 1$. There is an integer k such that whenever G is a subgroup of $GL_n(\mathbb{F}_p)$ generated by elements of order p, then $G = \langle g_1 \rangle \cdots \langle g_k \rangle$, with the g_i's of order p. Moreover, there is an integer d (depending only on n) such that if G^* is the algebraic subgroup of $GL_n(\mathbb{F}_p)$ generated by the subgroups $X(u, \mathbb{F}_p^{alg})$ where $u \in G$ has order p, then $[G^*(\mathbb{F}_p) : G] \leq d$, and G contains all elements of $G^*(\mathbb{F}_p)$ of order p. If $p > d$, then $G^*(\mathbb{F}_p) = G$.*

Proof. If F is a field and u a unipotent matrix, then the group $X(u, F)$ is definable in F (and contained in $\mathrm{GL}_n(F) \cap X(u, F^{alg})$). Suppose first that there is no such integer k. We can then find an increasing sequence of prime numbers $p(i)$, $i \in \mathbb{N}$, and for each i, a subset A_i of $\mathrm{GL}_n(\mathbb{F}_{p(i)})$ consisiting of elements of order $p(i)$, and such that the subgroup G_i of $\mathrm{GL}_n(\mathbb{F}_{p(i)})$ generated by A_i cannot be written as $\langle g_1 \rangle \langle g_2 \rangle \cdots \langle g_i \rangle$ with the g_i's of order $p(i)$.
Consider the structure $M_i = (\mathbb{F}_{p(i)}, +, \cdot, 0, 1, A_i)$, and let $M = (F, +, \cdot, 0, 1, A)$ a non-principal ultraproduct of the M_i's. Then F is pseudo-finite, has characteristic 0, and A is a set of unipotent elements of $\mathrm{GL}_n(F)$. Then, for any i and $g_1, \ldots, g_i \in A$, if $X(g_1, F) \cdots X(g_i, F)$ is a group, then there is some element of A which does not belong to it. However, this contradicts

Theorem 6.17: when $g \in A$, then the Zariski closure of $X(g, F)$ is an algebraic connected subgroup of $\mathrm{GL}_n(F^{alg})$, and therefore the subgroup generated by all $X(g, F)$, $g \in A$, is definable, and of the form $X(g_1, F) \cdots X(g_r, F)$ for some r. This gives the desired contradiction.

By Theorem 6.17 again, we know that $[G^*(F) : G] < \infty$, where $G = \langle X(g, F) \mid g \in A\rangle$. This gives us the bound d on the index of the $\langle A_i\rangle$'s in $G^*(\mathbb{F}_{p(i)})$.

If $p = p(i) > d$, let $g \in G^*(\mathbb{F}_p)$ be an element of order p. If $g \notin \langle A_i\rangle$, then all cosets $g^j\langle A_i\rangle$'s are distinct, which contradicts the bound.

Remark 6.22. The second assertion was originally proved by Nori [35].

Theorem 6.23. *(Jordan) Given $n \geq 1$, there is $J(N)$ such that if $\Gamma \leq \mathrm{GL}_n(\mathbb{C})$ is finite, then Γ has an abelian normal subgroup of index $\leq J(n)$.*

This result is false in positive characteristic: as we saw above, when $n > 2$, $\mathrm{GL}_n(\mathbb{F}_p^{alg})$ contains arbitrarily large finite nilpotent subgroups, and no such bound exists. The problem comes from the elements of order p. M. Larsen and R. Pink formulated the correct analogue in positive characteristic:

Theorem 6.24. *(Larsen-Pink, [31]). Fix n. There is $J'(n)$ such that if $\Gamma \leq \mathrm{GL}_n(k)$ is finite, then there are normal subgroups $\Gamma_3 \leq \Gamma_2 \leq \Gamma_1$ of Γ such that*

(a) $[\Gamma : \Gamma_1] \leq J'(n)$

(b) Either $\Gamma_1 = \Gamma_2$, or char$(k) = p > 0$ and Γ_1/Γ_2 is a direct product of finite simple groups of Lie type in characteristic p.

(c) Γ_2/Γ_3 is abelian of order not divisible by p.

(d) Either $\Gamma_3 = (1)$, or char$(k) = p > 0$ and Γ_3 is a p-group.

They also obtain:

Theorem 6.25. *Any finite subgroup Γ of $GL_n(k)$ (char$(k) = p > 0$) has an abelian normal subgroup of order prime to p, and of index $\leq J'(n)|\Gamma_{(p)}|^3$, where $\Gamma_{(p)}$ is a p-Sylow of Γ.*

A key technical result in the proof was also obtained by Hrushovski and Wagner ([24]), and is the following:

Theorem 6.26. *(3.1 in [24]) Let G be a simple algebraic group, $G_i = G(k_i)$, k_i a field, $i \in I$, and $\Gamma_i \leq G_i$ a finite subgroup of G_i. Assume that in a non-principal ultraproduct $G^* := G(\prod_i k_i/\mathcal{U}) = \prod_i G_i/\mathcal{U}$, the subgroup $\Gamma^* = \prod_i \Gamma_i/\mathcal{U}$ is Zariski dense in G. Then for any subvariety $V^* := \prod_i V_i/\mathcal{U}$ of G^*, we have*

$$|V_i \cap \Gamma_i| = O(|\Gamma_i|^{\dim(V)/\dim(G)}).$$

The proof is set in a very abstract context, with several notions of dimension intervening. One checks that our setting verifies the hypotheses of the main result (Thm 1.6 in [24]) to obtain the following:

Let G^*, Γ^* as above. Let $d = \dim$ be the usual algebraic dimension (of the Zariski closure of a set), and δ be the quasi-dimension on G^* defined as follows: If $X \subset G^{*m}$ is definable, then

$$\delta(X) = \pi \log(\mu(X)).$$

Here, μ is the ultraproduct of the counting measures on the finite sets $X_i \cap \Gamma_i$, $\log : \mathbb{R}^{*>0} \to \mathbb{R}^*$ is the ultrapower of the log map, and $\pi : \mathbb{R}^* \to \mathbb{R}^*/I$ is the projection, with I the convex hull of $\mathbb{Z}$ inside $\mathbb{R}^*$. Then

$$d(G^*)\delta(X) \leq d(X)\delta(G^*).$$

Applications to regularity lemmas

Theorem 6.27. *(Szemerédi's regularity Lemma, graph theoretic version) For every $\varepsilon > 0$, there is an integer M such that if $G = (V, E)$ is a finite graph with at least M vertices, there exists a partition $V = V_1 \cup \cdots \cup V_k$ (ewith $k \leq M$) and a set $\Sigma \subset [k]^2$ such that*

(a) (Exceptional set) $\sum_{(i,j)\in\Sigma} |V_i||V_j| \leq \varepsilon |V|^2$

(b) (ε-regularity) For every (i, j) not in Σ, for every $A \subseteq V_i$, $B \subseteq V_j$

$$\big| |E \cap (A \times B)| - \delta_{i,j}|A||B| \big| < \varepsilon,$$

where $\delta_{i,j} = d_E(V_i, V_j) = \frac{E\cap(V_i\times V_j)}{|V_i||V_j|}$.

Remarks 6.28. (1) One can rephrase the regularity condition as:

$$|d_E(V_i, V_j) - d_E(A, B)| < \varepsilon \text{ whenever } A \subseteq V_i,\ B \subseteq V_j \text{ and } |A| \geq \varepsilon|V_i|,\ |B| \geq \varepsilon|V_j|.$$

Moreover, one can choose the V_i's so that $\big||V_i| - |V_j|\big| \leq 1$ (equipartition).

(2) (Gowers) M is bounded by a tower of exponentials of length $O(\varepsilon^{-5})$.

(3) The exceptional set is necessary: suppose for instance that $V = W$ and E is the relation corresponding to a linear order (the so-called *half-graph*). Then the members of the partition which intersect the diagonal cannot be ε-regular, unless they are very small. If the graph is stable, however this problem does not occur. The result first appears in a paper by M. Malliaris and S. Shelah [34]. I will give here the version proved by Malliaris and Pillay [33], in the context of Keisler measures. Recall that a Keisler measure (on a structure M) is a finitely additive probability measure μ on the set Def(M) of definable subsets of cartesian powers of M.

Theorem 6.29. *([33], Thm 1.1) Let (V, W, R) be a definable bipartite graph (inside a saturated M), where R is stable, and given Keisler measures μ on V and ν on W, we have the followoing: Given $\varepsilon > 0$, we can partition V into finitely many definable sets $V_1, \ldots, V_m$, each defined by a Δ-formula, and partition W into finitely many $W_1, \ldots, W_m$ each defined by a Δ^*-formula, such that for each V_i, W_j exactly one of the following holds:*

(1) *For all $a \in V_i$ outside a set of μ-measure $\leq \varepsilon\mu(V_i)$, for all $b \in W_j$ ouside a set of ν-measure $\leq \varepsilon\nu(W_j)$, we have $R(a,b)$ and dually, for all $b \in W_j$ outside a set of ν-measure $\leq \varepsilon\nu(W_j)$ for all $a \in V_i$ outside a set of μ-measure $\leq \varepsilon\mu(V_i)$, $R(a,b)$ holds,* **or**

(2) *For all $a \in V_i$ outside a set of μ-measure $\leq \varepsilon\mu(V_i)$, for all $b \in W_j$ ouside a set of ν-measure $\leq \varepsilon\nu(W_j)$, we have $\neg R(a,b)$ and dually, for all $b \in W_j$ outside a set of ν-measure $\leq \varepsilon\nu(W_j)$ for all $a \in V_i$ outside a set of μ-measure $\leq \varepsilon\mu(V_i)$, $\neg R(a,b)$ holds.*

A Δ-formula is a finite Boolean combination of formulas of the form $R(x,b)$ or $x = a$ where $a \in V$, $b \in W$, and a Δ^*-formula is a finite Boolean combination of formulas $R(a,y)$ or $y = b$, $a \in V$, $b \in W$. The main ingredient of the proof is the following

Lemma 6.30. *Given $\varepsilon > 0$, we can write V as a disjoint union of Δ-formulas over M, $V = V_1 \cup \cdots \cup V_m$, such that for each i there is a complete Δ-type p_i over M such that $\mu(p_i) > 0$, $V_i \in p$ and $\mu(V_i \setminus p_i) \leq \varepsilon\mu(V_i)$.*

The proof is by induction on the Cantor-Bendixon rank of the space of Δ-types over M. This rank is finite. Recall that if X is a Boolean space, then one defines by induction: $X_0 = X$; X_{i+1} is the closed subset of X_i consisting of non-isolated points. Then the rank is the least ordinal at which the procedure stops, if there is one. In our case it is finite. One then works with the (finitely many) Δ-types p_i of maximal CB-rank. The model theory of pseudo-finite fields, and in particular the stability-theoretic techniques which can be used, allow Pillay and Starchenko to give a new proof of a regularity lemma used by T. Tao ([41]) in the proof of a result on finite fields. This in turn has an application and gives a new proof of an algebraic regularity lemma used by T. Tao in the proof of a result on finite fields. I will first state Tao's result, and then the two versions of the algebraic regularity lemma.

Theorem 6.31. *(Tao, [41] Thm 1) For any degree d, there is a constant C such that the following holds. Let $\mathbb{F}$ be a finite field of characteristic at least C, and let $P \in \mathbb{F}[X,Y]$ be a polynomial of degree at most d. Then at least one of the following statement holds:*

(a) (Additive structure) One has

$$P(X,Y) = Q(F(X) + G(Y))$$

for some polynomials Q, F, G over $\mathbb{F}$.

(b) (Multiplicative structure) One has

$$P(X,Y) = Q(F(X) \cdot G(Y))$$

for some polynomials Q, F, G over $\mathbb{F}$.

(c) (Moderate asymmetric expansion) One has

$$|P(A,B)| \geq C^{-1}|\mathbb{F}|$$

whenever $A, B \subseteq \mathbb{F}$ with $|A||B| \geq C|\mathbb{F}|^{2-1/8}$.

There are other versions with variations on the conditions.

Lemma 6.32. *(Pillay-Starchenko, Cor. 1.2 in [37]) Let F be a pseudo-finite field, and let V, W and $E \subseteq V \times W$ be definable sets. Assume* $\dim(V) = n$ *and* $\dim(W) = k$. *Then we can dpartition W into* $\mathrm{acl}(A)$*-definable sets $Q_1, \dots, W_m$ such that for each $1 \leq i, j \leq m$ there is $c_{i,j} \in \mathbb{Q}^{>0}$ and* $\mathrm{acl}(A)$*-definable subset $D_{i,j}$ of $W_i \times W_j$ with* $\dim(D_{i,j}) < 2k$*, such that either* $\dim((E(x, a) \cap E(x, b)) < n$ *for all $(a, b) \in (W_i \times W_j) \setminus D_{i,j}$, or the pair* $(\dim, \mu)$ *associated to* $(E(x, a) \cap E(x, b))$ *equals $(n, c_{i,j})$ for all $(a, b) \in W_i \times W_j \setminus D_{i,j}$.*

Corollary 6.33. *(Pillay-Starchenko, Cor. 1.3 in [37]) If $M > 0$ there exists $C = C_M > 0$ such that whenever $\mathbb{F}$ is a finite field of cardinality $> C$, and V, W are non-empty definable seets (in $\mathbb{F}$) of complexity at most M and $E \subseteq V \times W$ is another definable set of complexity at most M, then there exists partitions $V = V_1 \cup \cdots \cup V_a$ and $W = W_1 \cup \cdots \cup W_b$ of V and W such that*

- *for all $1 \leq i \leq a$ and $1 \leq j \leq b$, $|V_i| \geq |V|/C$ and $|W_j| \geq |W|/C$.*
- *The V_i's and W_j's are definable with complexity at most C.*
- *For all $1 \leq i \leq a$ and $1 \leq j \leq b$, $A \subseteq V_i$ and $B \subseteq W_j$,*

$$\big||E \cap (A \times B)| - d_E(V_i, W_j)|A||B|\big| \leq C|\mathbb{F}|^{-1/4}|V_i||W_j|.$$

6.34. Concluding remarks. Other applications were made by E. Hrushovski on approximate subgroups: let G be a group, $k \in \mathbb{N}$, and assume that A is a large finite subset of G such that $|A \cdot A| \leq k|A|$ (*small doubling*). What can one say about A? Same question if $|A \cdot A^{-1} \cdot A| \leq k|A|$ (*small tripling*). See [20], and also results by Breuillard-Green-Tao ([3], [4], ...).
The use of various *dimensions*, as in the proof of 6.26, is currently a hot topic in model theory. A thorough study is started in Hrushovski's paper [21], and is full of open questions.

References

[1] J. Ax, The elementary theory of finite fields, Annals of Math. 88 (1968), 239 – 271.

[2] O. Beyarslan, Interpreting Random Hypergraphs in Pseudofinite Fields, J. of Inst. Math. Jussieu 9 No 1 (2010, 29 – 47.

[3] E. Breuillard, B. Green, Ben; T. Tao, Linear approximate groups. Electron. Res. Announc. Math. Sci. 17 (2010), 57 – 67.

[4] E. Breuillard, B. Green, T. Tao, Approximate groups - I: The torsion-free nilpotent case. J. Inst. Math. Jussieu 10 (2011), no. 1, 37 – 57.

[5] C.C. Chang, H.J. Keisler, *Model theory*, North-Holland, Amsterdam 1977.

[6] Z. Chatzidakis, L. van den Dries, A. Macintyre, Definable sets over finite fields, J. reine u. ang. Math. 427 (1992), 107 – 135.

[7] J. Denef, F. Loeser, Definable sets, motives and p-adic integrals, J. Amer. Math. Soc. 14 (2001), no. 2, 429–469.

[8] L. van den Dries, K. Schmidt, Bounds in the theory of polynomials rings over fields. A non-standard approach. Invent. Math. 76 (1984), 77 – 91.

[9] J. -L. Duret, Les corps faiblement algébriquement clos non séparablement clos ont la propriété d'indépendance, in: *Model theory of Algebra and Arithmetic, Pacholski et al. ed.*, Springer Lecture Notes 834 (1980), 135 –157.

[10] R. Elwes, H.D. Macpherson, A survey of asymptotic classes and measurable structures, in *Model theory and applications to algebra and analysis* (Eds. Z. Chatzidakis, H.D. Macpherson, A. Pillay, A.J. Wilkie), Cambridge University Press, 2008, 125 – 159.

[11] M. Fried, D. Haran, M. Jarden, Galois stratification over Frobenius fields, Adv. in Math. 51 (1984), 1 – 35.

[12] M. Fried, D. Haran, M. Jarden, Effective counting of the points of definable sets over finite fields, Israel J. Math. 85 (1994), 103 – 133.

[13] M. Fried, M. Jarden, Field Arithmetic, Ergebnisse 11, 3rd edition, Springer Berlin-Heidelberg 2008.

[14] M. Fried, G. Sacerdote, Solving diophantine problems over all residue class fields of a number field and all finite fields, Annals of Math. 104 (1976), 203 – 233.

[15] I. Halupczok, A measure for perfect PAC fields with pro-cyclic Galois group. Journal of Algebra 310 (2007), 371-395.

[16] G. Hermann, Die Frage der endlich vielen Schritte in der Theorie der Polynomideale, Math. Ann. 95 (1926), no. 1, 736 – 788.

[17] E. Hrushovski, A new strongly minimal set, Annals of Pure and Applied Logic 62 (1993), 147 – 166.

[18] E. Hrushovski, Strongly minimal expansions of algebraically closed fields. Israel Journal of Mathematics, vol. 79 (1992), 129 – 151.

[19] E. Hrushovski, Pseudo-finite fields and related structures, in: Model Theory and Applications, Bélair et al. ed., Quaderni di Matematica Vol. 11, Aracne, Rome 2005, 151 – 212.

[20] E. Hrushovski, Stable group theory and approximate subgroups. J. Amer. Math. Soc. 25 (2012), no. 1, 189 – 243.

[21] E. Hrushovski, On pseudo-finite dimensions, Notre Dame J. Form. Log. 54 (2013), no. 3 - 4, 463 – 495.

[22] E. Hrushovski, A. Pillay, Groups definable in local fields and pseudo-finite fields, Israel J. of Math. 85 (1994), 203 – 262.

[23] E. Hrushovski, A. Pillay, Definable subgroups of algebraic groups over finite fields, J. reine angew. Math. 462 (1995), 69 – 91.

[24] E. Hrushovski, F.O. Wagner, Counting and dimensions, London Math. Soc. Lecture Note Ser., 350, Cambridge Univ. Press, Cambridge, 2008, 161 – 176.

[25] E. Hrushovski, B. Zilber, Zariski geometries, J. of the AMS, Vol 9 Nr 1 (1996), 1 – 56.

[26] C. Kiefe, Sets definable over finite fields: Their Zeta function, Trans. of Amer. Math. Soc. 223 (1976), 45 – 59.

[27] E. Kowalski, Exponential sums over definable subsets of finite fields, Israel J. Math. 160 (2007), 219–251.

[28] S. Lang, *Introduction to algebraic geometry*, Addison-Wesley Pub. Co., Menlo Park 1973.

[29] S. Lang, *Algebra*, Addison-Wesley Pub. Co., Menlo Park 1984.

[30] S. Lang, A. Weil, Number of points of varieties in finite fields, Am. J. of Math. 76 (1954), 819 –827.

[31] M. Larsen, R. Pink, Finite subgroups of algebraic groups, J. Amer. Math. Soc. 24 (2011), no. 4, 1105 – 1158.

[32] H.D. Macpherson, C. Steinhorn, One-dimensional asymptotic classes of finite structures, Trans. Amer. Math. Soc. 360(2008), 411–448.

[33] M. Malliaris, A. Pillay, The stable regularity lemma revisited, Proc. of Amer. Math. Soc. 144 (2016), 1761 – 1765.

[34] M. Malliaris, S. Shelah, Regularity lemmas for stable graphs, Trans. of Amer. Math. Soc. 366 (2014) No 3, 1551 – 1585.

[35] M.V. Nori, On subgroups of $GL_n(\mathbb{F}_p)$, Inventiones Math. 88 (1987), 257 – 275.

[36] F. Point, Ultraproducts and Chevalley groups, Arch. Math. Logic 38 (1999), 355–372.

[37] A. Pillay, S. Starchenko, Remarks on Tao's algebraic regularity lemma. Preprint 2013, arXiv 1310.7538.

[38] M. Ryten, *Results around asymptotic and measurable groups*, PhD thesis, University of Leeds, 2008.

[39] M. Ryten, I. Tomašić, ACFA and measurability, Selecta Mathematica New series, 11 (2005), 523-537.

[40] A. Seidenberg, Constructions in algebra, Trans. Amer. Math. Soc. 197 (1974), 273 – 313.

[41] T. Tao, Expanding polynomials over finite fields of large characteristic, and a regularity lemma for definable sets. Contrib. to discrete math., 10 (2015), no1, 22 – 98.

DMA - Ecole Normale Supérieure
45, rue d'Ulm
75230 Paris Cedex 05
France
e-mail: zoe.chatzidakis@ens.fr