FIELD THEORY

PETE L. CLARK

Contents

Thanks to Asvin Gothandaraman and David Krumm for pointing out errors in these notes.

About these notes

The purpose of these notes is to give a treatment of the theory of fields. Some aspects of field theory are popular in algebra courses at the undergraduate or graduate levels, especially the theory of finite field extensions and Galois theory. However, a student of algebra (and many other branches of mathematics which use algebra in a nontrivial way, e.g. algebraic topology or complex manifold theory) inevitably finds

that there is more to field theory than one learns in one's standard "survey" algebra courses.[1] When teaching graduate courses in algebra and arithmetic/algebraic geometry, I often find myself "reminding" students of field-theoretic facts that they have not seen before, or at any rate not in the form I wish to use them.

I also wish to fill in some gaps in my own knowledge. Especially, I have long wished to gain a deeper understanding of positive characteristic algebraic geometry, and has become clear that the place to begin study of the "pathologies"[2] of algebraic geometry in characterstic p is the study of finitely generated field extensions in positive characteristic.

These notes are meant to be comprehensible to students who have taken a basic graduate course in algebra. In theory one could get away with less – the exposition is mostly self-contained. As algebraic prerequisites we require a good working knowledge of linear algebra, including tensor products. The reader should also be comfortable with – and fond of – groups and rings. Such a benevolent familiarity is used much more than any specific results of group or ring *theory*. Our approach is sufficiently abstract and streamlined that it is probably inappropriate for most undergraduates. In particular, more often than not our approach proceeds from the general to the specific, and we make no apologies for this.

0.1. **Some Conventions.**

By convention, all of our rings are associative and have a multiplicative unity, called 1. Again by convention, a homomorphism of rings necessarily carries 1 to 1.

These notes contain many exercises, including some which ask for proofs of stated results. In general I am not at all opposed to the idea of a text giving complete details for all of its arguments.[3] However, it is in the nature of this particular subject that there are many more results than proof techniques, to the extent that giving complete proofs of all results would create a lengthy repetitiveness that may discourage the reader to read the proofs that we do give.

As a rule, exercises that ask for proofs of stated results are meant to require no new ideas beyond what was (even recently) exposed in the text. A reader who feels otherwise should contact me: there may be an unintended gap in the exposition. On the other hand, if exercises are given at all, it certainly spruces things up to have some more challenging and interesting exercises. I have also not hesitated to give exercises which can in principle be solved using the material up to that point but become much easier after later techniques are learned.

At some point I fell victim to the disease of not liking the look of a paragraph in which only a few words appear on the last line. Because of this, in the exercises I have sometimes omitted the words "Show that". I hope the meaning remains clear.

---

[1] I make no claim that this phenomenon is unique to g field theory.

[2] The term was used by Mumford, but with evident affection.

[3] In fact I agree with Robert Ash that the prevailing negative reputation of such texts is undeserved: the royal road to a particular destination may or may not exist, but it seems perverse to claim that it *ought not* to exist.

1. Introduction to Fields

A **field** is a commutative ring in which each nonzero element has a multiplicative inverse. Equivalently, a field is a commutative ring R in which the only ideals are (0) and R itself.

So if F is a field, S is a ring, and $\varphi : F \to S$ is a homomorphism of rings, then since the kernel of φ is an ideal of F, φ is either injective (if its kernel is 0) or identically the zero map (if its kernel is F). Moreover, the latter case implies that $1_S = \varphi(1_F) = 0$, which happens iff S is the zero ring. So any homomorphism from a field into a nonzero ring – in particular into any field or integral domain – is injective. Thus if $\varphi : F \to K$ is a homomorphism between fields, we may equally well speak of the **field embedding** φ.

Variations on the definition: In older terminology, a field could be non-commutative, i.e., any ring in which each nonzero element has a two-sided multiplicative inverse. We now call such things "division rings" or "division algebras." One also sometimes encounters non-associative division algebras, e.g. Cayley's octonions.

The two branches of mathematics in which general fields play a principal role are field theory (of course) and linear algebra. Most of linear algebra could be developed over a general division algebra rather than over a general field. In fact for the most part the theory is so similar that it is not really necessary to consider division algebras from the outset: one can just check, if necessary, that a certain result which is true for vectors spaces over a field is also true for left modules over a division algebra. On the other hand, when one studies things like roots of polynomials and lattices of finite degree extensions, one immediately finds that non-commutative division algebras behave in quite different and apparently more complicated ways.

Example 1.1. *There are exactly two complex numbers z such that $z^2 = -1$: $z = i$ and $z = -i$. In general, any nonzero polynomial $P(t)$ with coefficients in a field can have no more solutions than its degree. But in Hamilton's quaternion algebra $\mathbb{H}$ there are clearly at least three solutions: $i^2 = j^2 = k^2 = -1$, and in fact there are uncountably many: a quaternion squares to -1 iff it is of the form $xi + yj + zk$ with $x^2 + y^2 + z^2 = 1$.*

Example 1.2. *Let $K/\mathbb{Q}$ be a quartic field (i.e., a field extension of $\mathbb{Q}$ which has dimension 4 as a $\mathbb{Q}$-vector space). Then there are at most three intermediate subfields $\mathbb{Q} \subsetneq F \subsetneq K$. (More precisely there is either zero, one or three such fields, and the first case happens "most of the time.") However, any noncommutative division algebra $B/\mathbb{Q}$ of degree 4 as a $\mathbb{Q}$-vector space has infinitely many nonisomorphic quadratic subfields.*

The study of division algebras is closely related to field theory – via Brauer groups and Galois cohomology – so that one can put one's understanding of a field F and its finite extensions to excellent use in studying noncommutative division algebras over F. In fact, notwithstanding the above two examples, the finite dimensional, central division algebra over a field F are significantly easier to understand than finite dimensional extension fields of F: e.g. we understand quaternion algebras over $\mathbb{Q}$ far better than quartic number fields.

2. Some examples of fields

2.1. **Examples From Undergraduate Mathematics.**

Example 2.1. *First of all there is the field of real numbers* $\mathbb{R}$*. One also encounters the complex numbers* $\mathbb{C} = \{a + bi \mid a, b \in \mathbb{R},\ i^2 = -1\}$ *and the rational numbers* $\mathbb{Q} = \{\frac{a}{b} \mid a \in \mathbb{Z}, b \in \mathbb{Z} \setminus \{0\}\}$.

Example 2.2. *For a prime* p*, the ring* $\mathbb{F}_p = \mathbb{Z}/p\mathbb{Z}$ *of integers modulo* p *is a field. In fact it is enough to show that it is an integral domain, since any finite integral domain must be a field: if* a *is a nonzero element of a finite integral domain, there must exist* $0 < i < j$ *such that* $a^i = a^j$*, and then by cancellation we get* $1 = a^{j-i} = a^{j-i-1}a$*. To check that* $\mathbb{F}_p$ *is an integral domain, suppose that* x, y *are nonzero elements in* $\mathbb{F}_p$ *such that* $0 = xy$*. Equivalently, we have integers* x, y *not divisible by* p *but such that* $p \mid xy$*. This contradicts the uniqueness of factorization of integers into primes, i.e., the "Fundamental Theorem of Arithmetic."*

Nonexample 2.3. *The ring of integers* $\mathbb{Z}/n\mathbb{Z}$ *is not a field unless* n *is prime: if* $n = n_1 \cdot n_2$ *with* $n_1, n_2 > 1$*, then* $(n_1 \pmod n) \cdot (n_2 \pmod n) = 0 \pmod n$ *exhibits zero divisors.*

Let us reflect on this a bit. Any subring of an integral domain is again an integral domain (if the larger ring has no nonzero divisors of zero, neither does the smaller ring). In particular, any subring of a field must be a domain. Suppose $n \in \mathbb{Z}^+$ is not prime. Then, since $\mathbb{Z}/n\mathbb{Z}$ is not a domain, no ring which contains $\mathbb{Z}/n\mathbb{Z}$ as a subring can be a domain. This leads to the concept of *characteristic*: a ring is said to have characteristic n if it admits $\mathbb{Z}/n\mathbb{Z}$ as a subring, and characteristic zero if it does not have characteristic n for any positive integer n. Equivalently, a ring has characteristic $n > 0$ iff n is the least positive integer such that adding 1 to itself n times in the ring yields 0, and has characteristic zero if there is no such integer. We see therefore that any integral domain – and in particular any field – must have characteristic 0 or characteristic a prime number p.

Exercise 2.1. *Let* R *be a finite ring. Show that* R *has finite characteristic, and that the characteristic divides* $\#R$*.*

Example 2.4. *Suppose there is a field* $\mathbb{F}_4$ *with four elements. Like all fields, it has distinct elements* 0 *and* 1*. Moreover, by the preceding exercise, it must have characteristic* 2*, so* $1 + 1 = 0$*. This leaves two further elements unaccounted for:* x *and* y*. The nonzero elements of any field form a group under multiplication, so in this case the group would have order* 3 *and therefore be cyclic. In particular* x *has order* 3*, hence so does* x^2*, so* x^2 *is equal to neither* 0*,* 1 *or* x*, and therefore* $x^2 = y = x^{-1}$ *and* $y^2 = x = y^{-1}$*. Also* $x + y$ *cannot equal* x *or* y*; if* $x + y = 0$*, then* $x = -y = y$ *since* $-1 = 1$ *in* $\mathbb{F}_4$*. Therefore we must have* $x + y = 1$*, i.e.,* $y = x - 1 = x + 1 = x^2$*. We have thus uniquely worked out the addition and multiplication table for our putative field of order four, and one can check directly that all the field axioms are satisfied: there is, indeed, a field of order four. There is a unique such field up to isomorphism. Finally, as suggested by our analysis above, the map which fixes* 0 *and* 1 *and interchanges* x *and* y *is an automorphism of the field. One can think of it as the map* $a \in \mathbb{F}_4 \mapsto a^2$*.*

Nonexample 2.5. *Suppose* $\mathbb{F}$ *is a field of order* 6*. By Exercise 2.1,* $\mathbb{F}$ *must have characteristic* 2 *or characteristic* 3*. Suppose it has characteristic* 2*. Then, by*

Sylow's Theorem, there exists $x \in (\mathbb{F}, +)$ of order 3: $3x = 0$. But also $2x = 0$, so $x = 3x - 2x = 0$, contradiction.

Exercise 2.2. *Let $\mathbb{F}$ be a finite field. Show that $\#\mathbb{F}$ cannot be divisible by two distinct primes p, q. (Hint: suppose the characteristic is p. Then there exists $a \in \mathbb{Z}^+$ such that $p^a \mid \#\mathbb{F}$, $\frac{\#\mathbb{F}}{p^a}$ is divisible by a second prime $q \neq p$ and $\gcd(p^a, \frac{\#\mathbb{F}}{p^a}) = 1$. By elementary number theory – "Bézout's Lemma" – there exist integers x, y such that $xp^a + y\frac{\#\mathbb{F}}{p^a} = 1$. Now argue as above.)*

Therefore the order of a finite field $\mathbb{F}$ must be a prime power p^f. In particular, $\mathbb{F}$ contains $\mathbb{Z}/p\mathbb{Z}$ as its **prime subring** (i.e., the subring generated by one).

Exercise 2.3. *Give a second proof that a finite field $\mathbb{F}$ must have prime power order: as above, $\mathbb{F}$ contains a unique subfield $\mathbb{F}_p$ of prime order. Argue that $\mathbb{F}$ is a finite-dimensional vector space over $\mathbb{F}_p$ of dimension $f = \log_p \#\mathbb{F}$.*

Exercise 2.4. *The next largest non-prime prime powers are 8 and 9. Try to contruct finite fields of these orders from "first principles", as we did with the case of order 4 above.*

We will see later that for every prime power p^a there is a finite field $\mathbb{F}$ of order p^f, that any two finite fields of order p^f are isomorphic, and that the automorphism group of a finite field of order p^f is cyclic of order f, generated by the "Frobenius map" $x \mapsto x^p$.

2.2. **Fields of Fractions.** If R is an integral domain, then one can define a field F whose elements are viewed as fractions $\frac{a}{b}$ with $a, b \in R, b \neq 0$. Formally speaking one considers ordered pairs $(a, b) \in R^2$, $b \neq 0$ and introduces the equivalence relation $(a, b) \sim (c, d) \iff ad = bc$, i.e., exactly the same construction that one uses to define the rational numbers in terms of the integers. The field F is called the **field of fractions**, (or, sometimes, "quotient field") of the integral domain R.

Exercise 2.5. *(Functoriality of the field of fractions) Let $\varphi : R \to S$ be an injective homomorphism of integral domains. Show that φ extends uniquely to a homomorphism from the fraction field $F(R)$ of R to the fraction field $F(S)$ of S.*

Exercise 2.6. *(Universal property of the field of fractions) Let R be an integral domain with fraction field F and let K be a field. For any injective homomorphism $\varphi : R \to K$, there exists a unique extension to a homomorphism $F \to K$.*

Exercise 2.7. *Let R be an integral domain with field of fractions $F(R)$. Show: $\#R = \#F(R)$.*

Thus any method which produces a supply of integral domains will also produce a supply of fields (of course distinct integral domains may have isomorphic fraction fields, a trivial example being $\mathbb{Z}$ and $\mathbb{Q}$ itself; there are in fact uncountably many isomorphism classes of integral domains with fraction field $\mathbb{Q}$).

Proposition 2.6. *If R is an integral domain, then the univariate polynomial ring $R[t]$ is also an integral domain. Moreover, if F is the fraction field of R, then the fraction field of $R[t]$ is $F(t)$, the field of all quotients of polynomials with F-coefficients.*

Exercise 2.8. *Prove Proposition 2.6.*

Example 2.7. *Applying the Proposition with $R = F$ a field, we get a field $F(t)$ of rational functions in F. E.g., the field $\mathbb{C}(t)$ is the field of meromorphic functions on the Riemann sphere (see the next section). Moreover, for any field F, $F[t]$ is a domain, so $F[t_1, t_2] := F[t_1][t_2]$ is also an integral domain. The fraction field is easily seen to be $F(t_1, t_2)$, i.e., the fraction field of $F[t_1, \ldots, t_n]$ is $F(t_1, \ldots, t_n)$ the field of rational functions in n indeterminates.*

Although successive applications of Proposition 2.6 will yield polynomial rings in only finitely many indeterminates, nothing stops us from considering larger polynomial rings: let $\mathbb{T} = \{t_i\}$ be any set of indeterminates, and R any commutative ring. One can consider the polynomial ring $R[\mathbb{T}]$, defined as the union (or, if you like, direct limit) of polynomial rings $R[S]$ where $S \subset \mathbb{T}$ is a finite subset. In other words, we consider the ring of polynomials in an arbitrary infinite set S of indeterminates, but any given polynomial involves only finitely many indeterminates. One can again show that if R is an integral domain, so is $R[\mathbb{T}]$. The corresponding fraction field $R(\mathbb{T})$ is the field of all quotients of polynomials in all these indeterminates.

Exercise 2.9. *Let F be a field and $\mathbb{T}$ a nonempty set of indeterminates. Show that the cardinality of the rational function field $F(\mathbb{T})$ is* $\max(\aleph_0, \#F, \#\mathbb{T})$.

Another way of manufacturing integral domains is to start with a commutative ring R and take the quotient by a prime ideal $\mathfrak{p}$. Then we can get a field by (if necessary, i.e., if $\mathfrak{p}$ is not maximal) taking the field of fractions of $R/\mathfrak{p}$. For example with $R = \mathbb{Z}$ we get the finite fields $\mathbb{F}_p$.

Example 2.8. *Let $R = F[T]$ and $\mathfrak{p}$ a nonzero prime ideal. Then, since R is a PID, $\mathfrak{p} = (f(t))$, where $f(t)$ is an irreducible polynomial. Moreover, assuming $f(t) \neq 0$, $\mathfrak{p}$ is maximal, so without having to take quotients we get a field*

$$K = F[t]/(f(t)),$$

whose dimension as an F-algebra is the degree of f.

An integral domain R is **finitely generated** (over $\mathbb{Z}$) if there exist $n \in \mathbb{Z}^+$ and elements $\alpha_1, \ldots, \alpha_r \in R$ such that the least subring of R containing all the α_i's is R itself. Another way of saying this is that the natural map

$$\mathbb{Z}[T_1, \ldots, T_n] \to R,\ T_i \mapsto \alpha_i$$

is surjective. In other words, an integral domain is finitely generated iff it is, for some n, the quotient of the ring $\mathbb{Z}[T_1, \ldots, T_n]$ by some prime ideal $\mathfrak{p}$.

Proposition 2.9. *For a field F, the following are equivalent:*
a) There exist $\alpha_1, \ldots, \alpha_n \in F$ so that the only subfield of F containing all the α_i's is F itself.
b) F is the fraction field of $\mathbb{Z}[x_1, \ldots, x_n]/\mathfrak{p}$ for some prime ideal $\mathfrak{p}$.

Exercise 2.10. *Prove Proposition 2.9.*

A field satisfying the equivalent conditions of Proposition 2.9 is said to be **finitely generated**. Applying part b) and Exercise 2.7 we see that any finitely generated field is finite or countably infinite. In particular the fields $\mathbb{R}$, $\mathbb{C}$ are not finitely generated. Conversely, a countable field need not be finitely generated: if $\mathbb{T}$ is a countably infinite set of indeterminates, then by Exercise 2.9 the field $\mathbb{Q}(\mathbb{T})$ is countable. Moreover it is both plausible and true that $\mathbb{Q}(\mathbb{T})$ is not finitely generated, but we lack the tools to prove this at the moment: we will return to this later

on in the context of the concept of **transcendence degree**.

One can also speak of finite generation in a relative sense:

Proposition 2.10. *For a subfield $f \subset F$, the following are equivalent:*
a) There exist elements $\alpha_1, \ldots, \alpha_n \in F$ such that the only subfield of F containing f and the α_i's is F itself.
b) F is isomorphic to the fraction field of $f[x_1, \ldots, x_n]/\mathfrak{p}$ for some prime ideal $\mathfrak{p}$.

Exercise 2.11. *Prove Proposition 2.10.*

If f is a subfield of F and $\alpha_1, \ldots, \alpha_n \in F$, we write $f(\alpha_1, \ldots, \alpha_n)$ for the smallest subfield of F containing f and the α_i's. The notation is sensible because this field can be described concretely as the set of all rational expressions $\frac{P(\alpha_1,\ldots,\alpha_n)}{Q(\alpha_1,\ldots,\alpha_n)}$ for $P, Q \in k[t_1, \ldots, t_n]$. (In particular there is a unique such smallest subfield.)

So for instance one can speak of fields which are finitely generated over the complex numbers $\mathbb{C}$, and such fields are especially important in algebraic geometry.

Proposition 2.11. *Let F be a field.*
a) If F has characteristic 0, there is a unique homomorphism $\iota : \mathbb{Q} \to F$.
b) If F has characteristic p, there is a unique homomorphism $\iota : \mathbb{F}_p \to F$.

Proof. For any ring R, there exists a unique ring homomorphism $\iota : \mathbb{Z} \to R$, which takes the integer n to n times the multiplicative identity in R. For $R = F$ a field, the map ι is an injection iff F has characteristic 0. So if F has characteristic 0, ι is injective, and by Exercise 2.5 it extends uniquely to a homomorphism $\iota : \mathbb{Q} \to F$. Any homomorphism from $\mathbb{Q}$ to F must restrict to the canonical injection on $\mathbb{Z}$ and therefore be ι. If F has characteristic $p > 0$, then ι factors through to give a map $\iota : \mathbb{F}_p \to F$. The uniqueness of ι can be seen in any number of ways: we leave it to the reader to find an explanation that she finds simple and convincing. □

It follows that $\mathbb{Q}$ (resp. $\mathbb{F}_p$) is the unique minimal subfield of any field F of characteristic 0 (resp. $p > 0$). We refer to $\mathbb{Q}$ (resp. $\mathbb{F}_p$) as the **prime subfield** of F. Note that since there are no nontrivial automorphisms of either $\mathbb{Q}$ or $\mathbb{F}_p$ (this follows by applying the proposition with $F = \mathbb{Q}$ or $F = \mathbb{F}_p$), the prime subfield sits inside F in an especially canonical way.

Exercise 2.12. *Let K be a field and k its prime subfield. Show that the concepts "K is finitely generated" and "K is finitely generated over k" coincide.*

Exercise 2.13. *For any field F, there exists a set of indeterminates $\mathbb{T}$ and a prime ideal $\mathfrak{p}$ of $\mathbb{Z}[\mathbb{T}]$ such that F is isomorphic to the fraction field of $\mathbb{Z}[\mathbb{T}]/\mathfrak{p}$.*

If F is infinitely generated (i.e., not finitely generated over its prime subfield) then the set $\mathbb{T}$ in Exercise 2.13 will of course have to be infinite. In such a case this "presentation" of F is not, in truth, so useful: e.g., with certain limited exceptions (to be discussed!) this is not a very insightful way of viewing the complex field $\mathbb{C}$.

Exercise 2.14. *Let R be a commutative ring, $\iota : R \to S$ an injective ring homomorphism, and $\alpha \in S$. Show that there is a unique minimal subring of S containing R and α, namely the set of all polynomials $P(\alpha), P \in R[t]$. This subring is accordingly denoted $R[\alpha]$.*

2.3. **Fields of Functions.**

Let U be a domain – i.e., a nonempty connected open subset – of the complex plane. In complex analysis one studies the set $\mathrm{Hol}(U)$ of all functions holomorphic (a.k.a. analytic) on all of U and also the larger set $\mathrm{Mer}(U)$ of all meromorphic functions on U, i.e., functions which are holomorphic on the complement of a discrete set $X = \{x_i\}$ and such that for each x_i there exists a positive integer n_i such that $z^{n_i} f$ is holomorphic at x_i. Under the usual pointwise addition and multiplication of functions, $\mathrm{Hol}(U)$ is a ring (a subring of the ring of all continuous $\mathbb{C}$-valued functions on U). Similarly, one can view $\mathrm{Mer}(U)$ as a ring in a natural way.

Theorem 2.12. *Let U be a domain in the complex plane.*
a) $\mathrm{Hol}(U)$ *is a domain.*
b) $\mathrm{Mer}(U)$ *is a field.*
c) $\mathrm{Mer}(U)$ *is the field of fractions of* $\mathrm{Hol}(U)$.

Proof. a) A consequence of the principle of analytic continuation is that the zero set of a not-identically-zero holomorphic function is discrete in U. For $0 \neq f, g \in \mathrm{Hol}(U)$, the zero set of fg is the union of the zero sets of f and g so is again discrete and thus certainly a proper subset of U.
b) Because $0 \neq f \in \mathrm{Hol}(U)$ has a discrete zero set $\{x_i\}$ and for each x_i, there exists a positive integer n_i such that $\frac{f}{z^{n_i}}$ extends to a continuous nonzero function at x_i, it follows that $\frac{1}{f_i}$ is meromorphic.
c) This lies deeper: it is a consequence of Weierstrass' factorization theory, in particular of the fact that for any discrete subset $X = \{x_i\}$ of U and any sequence of positive integers $\{n_i\}$ there exists a holomorphic function on U with zero set X and order of vanishing n_i at x_i. □

Exercise 2.15. *Show:* $\mathrm{Mer}(\mathbb{C})$ *is not finitely generated over* $\mathbb{C}$.

More generally, if M is a connected complex manifold, there is a ring $\mathrm{Hol}(M)$ of "global" holomorphic functions on M and a field $\mathrm{Mer}(M)$ of meromorphic functions. It need not be the case that $\mathrm{Mer}(M)$ is the fraction field of $\mathrm{Hol}(M)$.

Example 2.13. *Take $M = \mathbb{C} \cup \{\infty\}$ to be the Riemann sphere. Then the only holomorphic functions on M are the constant functions, whereas* $\mathrm{Mer}(M) = \mathbb{C}(z)$, *the rational functions in z.*

In various branches of geometry one meets many such "fields of functions": a very general example, for the highly trained reader, is that if X is an integral (reduced and irreducible) scheme, then the ring of all functions regular at the generic point η is a field. If X itself is a scheme over a field k, then this field is written $k(X)$ and called the **field of rational functions** on X. For example, the field of rational functions on the complex projective line $\mathbb{P}^1/\mathbb{C}$ is the rational function field $\mathbb{C}(t)$. This is essentially the same example as the Riemann sphere above, but couched in more algebraic language.

In general, one must restrict to functions of a rather special kind in order to get a *field* of functions. Using the ideas of the previous subsection, it seems fruitful to first consider *rings* R of functions on a topological space X. Then we want R to be a domain in order to speak of fraction field $F(R)$ of "meromorphic functions" on X.

Suppose X is a topological space and consider the ring $R = R(X, \mathbb{C})$ of all continuous functions $f : X \to \mathbb{C}$. A moment's thought indicates that for the "reasonable" topological spaces one considers in geometry, R will not be a domain. The question comes down to: do there exist functions $f_1,\ f_2 : X \to \mathbb{C}$ neither of which is zero on all of X but such that the product $f_1 \cdot f_2$ is identically zero?

Here are some easy observations. First, if X is not connected, the answer is certainly yes: write $X = Y_1 \cup Y_2$ where Y_i are disjoint open sets. Take f_1 to be the characteristic function of Y_1 an $f_2 = 1 - f_1$ to be the characteristic function of Y_2.

In fact R is not a domain even if X is the Euclidean plane: let D_1, D_2 be two disjoint closed disks, say with centers z_i and radii equal to 1. Certainly there exist continuous functions $f_i : X \to \mathbb{C}$ such that $f_i(z_i) = 1$ and $f_i(z) = 0$ if z lies outside of D_i. Indeed it is well-known that f_i may be chosen to be infinitely differentiable, and the argument generalizes to all manifolds and indeed to paracompact Hausdorff spaces (the key point being the existence of suitable partitions of unity).

On the other hand, suppose the space X is **irreducible**: that is, if $Y_1,\ Y_2$ are two proper closed subsets of X then $Y_1 \cup Y_2 \neq X$. Then, applying this to $Y_i = f_i^{-1}(0)$, we get that the zero set of $f_1 f_2$ is $Y_1 \cup Y_2 \neq X$, so $R(X, \mathbb{C})$ is a domain, and one can take its fraction field, which consists of functions which are defined on some dense (equivalently, nonempty!) open subset of X. If you have never studied algebraic geometry, you will doubtless be thinking, "What kind of crazy topological space would be irreducible?" However, the Zariski topology on a smooth, connected algebraic variety over (say) the complex field $\mathbb{C}$ is irreducible.

2.4. **Completion.**

None of the constructions of fields we have discussed so far give rise to either $\mathbb{R}$ or $\mathbb{C}$ in a reasonable way. These fields are uncountable, and from a purely algebraic perspective their structure is quite complicated. The right way to think about them is via a mixture of algebra and topology, e.g. one thinks of $\mathbb{R}$ as the completion of the field of rational numbers with respect to the standard absolute value.

An **absolute value** on a field K is a real-valued function $x \to ||x||$ satisfying:

(AV1) $||x|| \geq 0$ for all $x \in K$, with equality iff $x = 0$.
(AV2) $||xy|| = ||x||||y||$ for all $x, y \in K$.
(AV3) $||x + y|| \leq ||x|| + ||y||$ for all $x, y \in K$.

It is immediate that an absolute value gives rise to a metric on K, via $d(x, y) = ||x - y||$. We can therefore complete the metric space to get a metric space $\hat{K}$ with a canonically embedded, dense copy of K. The key point is that $\hat{K}$ also has a canonical field structure.

In brief, we consider the set $\mathcal{C}$ of Cauchy sequences in K. This becomes a ring under the operations of pointwise addition and multiplication. (It is far from being a domain, having many zero divisors and idempotent elements.) Inside this ring we

have $\mathfrak{c}$, the collection of sequences converging to 0. It is not hard to check that $\mathfrak{c}$ is in fact an *ideal* of $\mathcal{C}$, so that we may form the quotient $\mathcal{C}/\mathfrak{c}$. Best of all, this quotient ring is a field: indeed, a nonzero element of the quotient may be represented by a Cauchy sequence $x_\bullet$ in K which does not converge to 0. It follows that there are only finitely many indices n such that $x_n = 0$: otherwise a subsequence of $x_\bullet$ converges to 0 and a Cauchy sequence with a convergent subsequence is itself convergent to the same limit as the subsequence. Consider then the sequence $y_\bullet$ which is defined by $y_n = x_n$ if $x_n = 0$ and $y_n = x_n^{-1}$ otherwise. The product sequence $x_\bullet y_\bullet$ has all sufficiently large terms equal to 1, so differs from the constant sequence 1 (the identity element of $\mathcal{C}$) by a sequence which has only finitely many nonzero terms, so in particular lies in $\mathfrak{c}$. Therefore the class of $y_\bullet$ in $\mathcal{C}/\mathfrak{c}$ is the inverse of $x_\bullet$.

We denote $\mathcal{C}/\mathfrak{c}$ by $\hat{K}$ and call it the **completion** of K with respect to $||\ ||$. There exists a natural embedding $K \hookrightarrow \hat{K}$ – namely we map each element of K to the corresponding constant sequence – and a natural extension of the norm on K to a norm on $\hat{K}$, namely $||x_\bullet|| = \lim_{n\to\infty} ||x_n||$, with respect to which $\iota : K \hookrightarrow \hat{K}$ is an isometry of normed spaces in which the image of K in $\hat{K}$ is dense. For more details, the reader is invited to consult [NTII, Chapter 2].

Example 2.14. *The completion of $\mathbb{Q}$ with the standard Archimedean absolute value $||\frac{p}{q}|| = |\frac{p}{q}|$ is the real field $\mathbb{R}$.*

Remark 2.1. *It is sometimes suggested that there is a circularity in this construction, in that the definition of completion refers to a metric and the definition of a metric refers to the real numbers.*[4] *But one should not worry about this. On the one hand, from our present point of view we can consider the reals as being already constructed and then it is a true, non-tautologous statement that the metric completion of the rationals is the reals. But moreover, a careful look at the construction in terms of equivalence classes of Cauchy sequences shows that one absolutely can construct the real numbers in this way, just by being careful to avoid referring to the real numbers in the course of the completion process. In other words, the real numbers can be defined as the quotient of the ring of Cauchy sequences of rational numbers (where the definition of Cauchy sequence uses only the metric as defined on rational numbers) by the maximal ideal of sequences converging to zero. After one constructs the real numbers in this way, one notes that the $\mathbb{Q}$-valued metric on $\mathbb{Q}$ extends to an $\mathbb{R}$-valued metric on $\mathbb{R}$: no problem.*

Example 2.15. *If k is any field, then defining $||0|| = 0$ and $||x|| = 1$ for all $x \neq 0$ gives an absolute value on k. The induced metric is the discrete metric and therefore k is, in a trivial way, complete and locally compact. This absolute value (and any other absolute value inducing the discrete topology) is called **trivial**; such absolute values are usually either explicitly or implicitly excluded from consideration.*

Example 2.16. $||\frac{a}{b}|| = p^{\mathrm{ord}_p(b)-\mathrm{ord}_p(a)}$, *where for an integer a, $\mathrm{ord}_p(a)$ denotes the largest power of p dividing a. (To get the degenerate cases to work out correctly, we set $\mathrm{ord}_p(0) = \infty$ and $p^{-\infty} = 0$.) The induced metric on $\mathbb{Q}$ is called the p-adic metric: in this metric, a number is close to zero if, after cancelling common factors, its numerator is divisible by a high power of p. Since the induced topology has no*

[4] In particular, Bourbaki's *General Topology* refrains from making any reference to real numbers or metric spaces for many hundreds of pages until the reals can be rigorously constructed.

isolated points, the completeness of the metric would contradict the Baire category theorem, hence the completion is an uncountable field, called $\mathbb{Q}_p$, the field of p-adic numbers.

Example 2.17. *Let k be any field and $K = k(t)$. Any element $r(t) \in K$ can be written as $t^a \frac{P(t)}{Q(t)}$ where $P(0)Q(0) \neq 0$ for a uniquely determined integer a. Define $||r(t)||_\infty := e^{-a}$. (There is no particular reason to use the number $e = 2.718\ldots$; any real number greater than 1 would serve as well.)*

Exercise 2.16. *Show: $|| \; ||_\infty$ gives an absolute value on $K(t)$.*

An element $r(t) \in K(t)$ is close to 0 iff it is divisible by a high power of t.

Exercise 2.17. *Show: the completion of $K(t)$ with respect to $|| \; ||_\infty$ is isomorphic to the **Laurent series field** $K((t))$, whose elements are formal power series $\sum_{n=n_0}^{\infty} a_n t^n$ with $n_0 \in \mathbb{Z}$, $a_n \in f$. (Hint: It is enough to show that the norm $|| \; ||_\infty$ extends to all of $K((t))$ and that $K(t)$ is dense in $K((t))$ in the induced topology.)*

Exercise 2.18. *Show: the fields $\mathbb{Q}_p$ are locally compact in their natural topology. Show: $K((t))$ is locally compact iff K is finite.*

Remark 2.2. *If k is a field complete with respect to an absolute value $| \; |$ and V is a finite-dimensional vector space over k, then viewing $V \cong k^{\dim V}$ gives V the canonical structure of a topological space – i.e., we can endow it with the product topology, and this topology is independent of the choice of basis. In particular, if k is locally compact, so is V. Moreover it has the canonical structure of a uniform space, and if k is complete then so is V. In particular, if $k \hookrightarrow l$ is a field embedding such that l is finite-dimensional as a k-vector space, then l is a complete uniform space and is locally compact iff k is. This implies that any finite extension of the fields $\mathbb{R}$, $\mathbb{Q}_p$ or $\mathbb{F}_p((t))$ have a canonical locally compact topology.*

Theorem 2.18. *(Classification of locally compact valued fields) Let $|| \; ||$ be a nontrivial valuation on a field K. The following are equivalent:*
(i) The metric topology on K is locally compact.
(ii) Either $(K, \; || \; ||) = \mathbb{R}$ or $\mathbb{C}$; or the induced metric is complete and non-Archimedean and the residue field is finite.
(iii) K is a finite extension of $\mathbb{R}$, of $\mathbb{Q}_p$ or of $\mathbb{F}_p((t))$.

Proof. See [NTII, Theorem 5.1]. □

There are more elaborate ways to construct complete fields. For instance, suppose R is a domain and $\mathfrak{p}$ is a prime ideal of R. Then in commutative algebra one learns how to complete R with respect to $\mathfrak{p}$, getting a homomorphism $R \to \hat{R}$ in which $\hat{R}$ is a domain, the image $\mathfrak{p}\hat{R}$ is the unique maximal ideal of $\hat{R}$, and $\hat{R}$ is complete with respect to a canonical uniform structure. We can then take the fraction field to get a complete field $\hat{K}$. Let us just mention one simple example to give the flavor: let f be a field and $R = f[x_1, \ldots, x_n]$ and $\mathfrak{p} = (x_1, \ldots, x_n)$. Then the completion is $\hat{R} = f[[x_1, \ldots, x_n]]$, the ring of formal power series in the indeterminates $x_1, \ldots, x_n$, and its quotient ifeld is $f((x_1, \ldots, x_n))$, the field of formal Laurent series in these indeterminates, i.e., the set of all formal sums $\sum_I a_I x^I$ where $I = (i_1, \ldots, i_n) \in \mathbb{Z}^n$ is a multi-index, $a_I \in k$, $x^I = x^{i_1} \cdots x^{i_n}$, and the set of indices I in which at least one i_j is negative and $a_I \neq 0$ is finite.

Such fields arise in algebraic and analytic geometry: $\mathbb{C}((x_1, \ldots, x_n))$ is the field of germs of meromorphic functions at a nonsingular point P on an n-dimensional analytic or algebraic variety.

Exercise 2.19. *Show: the field $k((x_1, x_2))$ is properly contained in $k((x_1))((x_2))$.*

3. Field Extensions

3.1. **Introduction.**

Let K be a field. If $\iota : K \to L$ is a homomorphism of fields, one says that L is an **extension field** of K. As a matter of psychology, it often seems more convenient to think of L as "lying above K" rather than K as embedding into L. We often write L/K instead of $\iota : K \to L$, notwithstanding the fact that the latter notation hides important information, namely the map ι.[5]

Much of field theory is devoted to an understanding of the various extension fields of a given field K. Since any field K has extensions of all sufficiently large cardinalities – $K(\mathbb{T})$ for any large enough set $\mathbb{T}$ – one obviously cannot literally hope to understand all field extensions of K. However there are two important classes (sets!) of field extensions that one can at least hope to understand: the first is the class of all finitely generated field extensions of K, and the second is the class of all algebraic field extensions of K.

If L/K is a field extension, then L is a K-algebra and in particular a vector space over K. Therefore it has a well-determined (but possibly infinite) dimension, denoted by $[L : K]$. One says that the extension L/K is **finite** if $[L : K] < \infty$, i.e., if L is a finite-dimensional K-vector space. For instance, one has $[\mathbb{C} : \mathbb{R}] = 2 < \infty$, so $\mathbb{C}/\mathbb{R}$ is a finite field extension.

Warning: The term "finite field extension" is ambiguous: it could presumably also refer to an extension of fields L/K in which L and K are both finite fields. In practice, one should expect the term to have the former meaning – i.e., the finiteness refers to the degree of the extension, and not to either field – but be prepared to seek clarification if necessary.

As an immediate application we can rederive the fact that the order of a finite field is necessarily a prime power. Namely, let $\mathbb{F}$ be a finite field, and let $\mathbb{F}_p$ be its prime subfield. Since $\mathbb{F}$ is finite, it is certainly finite-dimensional over $\mathbb{F}_p$ (any infinite dimensional vector space over any field is infinite), say of dimension d. Then $\mathbb{F}$ as an $\mathbb{F}_p$-vector space is isomorphic to $\mathbb{F}_p^d$, so its cardinality is p^d.

Theorem 3.1. *(Degree multiplicativity in towers) Let $F \subset K \subset M$ be field extensions. Then we have*

$$[M : F] = [M : K][K : F].$$

Proof. Let $\{b_i\}_{i \in I}$ be an F-basis for K and $\{a_j\}_{j \in J}$ be a K-basis for M. We claim that $\{a_i b_j\}_{(i,j) \in I \times J}$ is an F-basis for M. This suffices, since then $[K : F] = \#I$, $[M : K] = \#J$, $[M : F] = \#(I \times J) = \#I \times \#J$.

[5]Beware: the notation L/K has nothing to do with cosets or quotients!

Let $c \in M$. Then there exist $\alpha_j \in K$, all but finitely many of which are zero, such that $c = \sum_{j \in J} \alpha_j a_j$. Similarly, for each $j \in J$, there exist $\beta_{ij} \in F$, all but finitely many of which are zero, such that $\alpha_j = \sum_{i,j} \beta_{ij} b_j$, and thus

$$c = \sum_{j \in J} \alpha_j a_j = \sum_{(i,j) \in I \times J} \beta_{ij} a_i b_j,$$

so that $\{a_i b_j\}$ spans K as an F-vector space. Now suppose the set $\{a_i b_j\}$ were linearly dependent. By definition, this means that there is some finite subset $S \subset I \times J$ such that $\{a_i b_j\}_{(i,j) \in S}$ is linearly dependent, and thus there exist $\beta_{ij} \in F$, not all zero, such that

$$\sum_{(i,j) \in S} (\beta_i j b_j) a_i = 0.$$

Since the a_i's are K-linearly independent elements of M, we have that for all i, $\sum \beta_{ij} b_j = 0$, and then similarly, since the b_j's are linearly independent elements of K we have $\beta_{ij} = 0$ for all j. □

Remark 3.1.2: In general the degree $[L : K]$ of a field extension is a cardinal number, and the statement of Theorem 3.1 is to be interpreted as an identity of (possibly infinite) cardinals. On the other hand, when M/K and K/F are finite, the argument shows that M/F is finite and the result reduces to the usual product of positive integers. Moreover the finite case is the one that is most useful.

Let L/K be an extension of fields and $\alpha \in L$. We say that α is **algebraic** over K if there exists some polynomial $P(t) = t^n + a_{n-1} t^n + \ldots + a_1 t + a_0 \in K[t]$ such that $P(\alpha) = 0$. If α is not algebraic over K it is said to be **transcendental** over K. A complex number which is algebraic over $\mathbb{Q}$ is called an **algebraic number**.

Examples 3.1.3: i is algebraic over $\mathbb{R}$ since it satisfies the equation $i^2 + 1 = 0$. It is also algebraic over $\mathbb{Q}$ for the same reason. Indeed for any $a \in \mathbb{Q}$, $a^{\frac{1}{n}}$ is algebraic over $\mathbb{Q}$. This is almost tautological, since by $a^{\frac{1}{n}}$, one generally means any complex number α such that $\alpha^n = a$, so α satisfies $t^n - a = 0$.

The following exercise gives less trivial examples.

Exercise 3.1. *Let $\frac{a}{b} \in \mathbb{Q}$. Show $\cos(\frac{a}{b}\pi)$ and $\sin(\frac{a}{b}\pi)$ are algebraic.*

Exercise 3.2. *a) Show that the set of all algebraic numbers is countably infinite. b) More generally, let K be any infinite field and L/K be any field extension. Show that the cardinality of the set of elements of L which are algebraic over K is equal to the cardinality of K.*

So "most" real or complex numbers are transcendental. This was observed by Cantor and stands as a famous early application of the dichotomy between countable and uncountable sets. Earlier Liouville had constructed particular transcendental numbers, like $\sum_{n=1}^{\infty} 10^{-n!}$: an application of the Mean Value Theorem shows that a number which is "too well approximated" by rational numbers cannot be algebraic. It is of course a different matter entirely to decide whether a particular, not obviously algebraic, number which is given to you is transcendental. Let us say only that both e and π were shown to be transcendental in the 19th century; that there were some interesting results in transcendence theory in the 20th century – e.g.

e^π and $2^{\sqrt{2}}$ are transcendental – and that to this day the transcendence of many reasonable looking constants – e.g. π^e, $\zeta(3) = \sum_{n=1}^\infty \frac{1}{n^3}$ – is much beyond our reach.

The problem of determining whether particular numbers are transcendental, although certainly of interest, has little to do with modern field theory. (Rather it is part of the theory of Diophantine approximation, a branch of number theory.)

Exercise 3.3. *(Universal property of polynomial rings): Let $\iota : R \to S$ be a homomorphism of commutative rings, and let $\alpha_1, \ldots, \alpha_n$ be elements of S. There is a unique R-algebra homomorphism $\Phi : R[t_1, \ldots, t_n] \to S$ which takes $t_i \mapsto \alpha_i$.*

Now let L/K be a field extension and $\alpha \in L$. By Exercise 3.1.6 there is a unique K-algebra homomorphism $\Phi : K[t] \to L$, $t \mapsto \alpha$. Let I be the kernel of Φ. Since $K[t]/I$ embeds in L, it is a domain, so I is a prime ideal. Since $K[t]$ is a principal ideal domain, there are only two choices:

Case 1: $I = 0$, i.e., Φ embeds $K[t]$ into L. This means precisely that α satisfies no polynomial relations with K-coefficients, so occurs iff α is transcendental over K.

Case 2: $I = (P(t))$ is generated by a single irreducible polynomial $P(t)$. Since the units of $K[t]$ are precisely the nonzero elements of K, it follows that there is a unique monic polynomial $P(t)$ (i.e., with leading coefficient 1) that generates I. We call this the **minimal polynomial** of α. Evidently for $Q \in K[t]$ we have $Q(\alpha) = 0 \iff P(t) \mid Q(t)$. In particular $P(\alpha) = 0$, so that α is algebraic, and moreover Φ induces an embedding $K[t]/(P(t)) \hookrightarrow L$. If P has degree d, then we say α is algebraic of degree d; moreover, a K-basis for the left-hand side is $1, t, \ldots, t^{d-1}$, so $[L : K] = d = \deg(P)$.

Let us summarize:

Theorem 3.2. *Let L/K be a field extension and $\alpha \in L$.*
a) TFAE:
(i) α is algebraic of degree d over K.
(ii) The K-vector space $K[\alpha]$ is finite, of degree d.
(iii) The K-vector space $K(\alpha)$ is finite, of degree d.
b) If α is algebraic of degree d, then $K[\alpha] = K(\alpha) \cong K[t]/(P(t))$, where $P(t) \in K[t]$ is the unique monic polynomial of degree d such that $P(\alpha) = 0$.
c) If α is transcendental over K, then $K[t] \cong K[\alpha] \subsetneq K(\alpha) \cong K(t)$.

It follows that the set of all rational expressions $\frac{P(\pi)}{Q(\pi)}$ with $P, Q \in \mathbb{Q}[t]$ is isomorphic to the rational function field $\mathbb{Q}(t)$! In other words, there is no genuinely algebraic distinction to be made between "fields of numbers" and "fields of functions."

A field extension L/K is **algebraic** if every $\alpha \in L$ is algebraic over K.

Corollary 3.3. *A finite extension L/K of fields is algebraic.*

Proof. We go by contraposition: suppose that L/K is transcendental, and let $\alpha \in L$ be transcendental over K. Then by Theorem 3.2c) we have

$$[K(\alpha) : K] \geq [K[\alpha] : K] = [K[t] : K] = \aleph_0,$$

so

$$[L : K] = [L : K(\alpha)][K(\alpha) : K] \geq \aleph_0.$$

□

The converse does not hold: many fields admit infinite algebraic extensions. A detailed analysis of algebraic field extensions is still ahead of us, but it is easy to see that the extension $\mathbb{Q}[\bigcup_{n\geq 2} 2^{\frac{1}{n}}]$ is an infinite algebraic extension, since it contains subextensions of arbitrarily large finite degree.

Exercise 3.4. *(Direct limits) Let $(I, \leq)$ be a directed set: recall that this means that I is partially ordered under $\leq$ and for any $i, j \in I$ there exists $k \in I$ with $i \leq k$ and $j \leq k$. A* ***directed system of sets*** *is a family of sets $\{X_i\}_{i\in I}$ together with maps $\iota(i, j) : X_i \to X_j$ for all $i \leq j$ satisfying the natural compatibility conditions: (i) $\iota_{i,i} = 1_{X_i}$ and (ii) for all $i \leq j \leq k$, $\iota(i, k) = \iota(j, k) \circ \iota(i, j)$. By definition, the* ***direct limit*** *$\lim_I X$ is the quotient of the disjoint union $\coprod_{i\in I} X_i$ by the equivalence relation $(x, X_i) \sim (\iota(i, j)x, X_j)$ for all $i \leq j$.*
a) Show that there are natural maps $\iota_i : X_i \to \lim_I X_i$. State and prove a universal mapping property for the direct limit.
b) Suppose that the maps $\iota(i, j)$ are all injective. Show that the maps $\iota_i : X_i \to \lim_I X_i$ are all injective. Explain why in this case $\lim_I X_i$ is often informally referred to as the "union" of the X_i's.
c) In any concrete category $\mathcal{C}$ – i.e., a category whose objects are sets, for which the set of all morphisms from an object A to an object B is a subset of the set of all functions from A to B, and for which composition and identity of morphisms coincide with the usual notions of functions – one has the notion of a directed system $\{A_i\}$ of objects in $\mathcal{C}$, i.e., we have sets A_i indexed by the directed set $(I, \leq)$ and for all $i \leq j$, the function $\iota(i, j) : A_i \to A_j$ is a morphism in $\mathcal{C}$. Give a definition of the direct limit $\lim_I A_i$ in this more general context. Show that the direct limit exists in the following categories: monoids, groups, commutative groups, rings, commutative rings, fields.
d) Give an example of a concrete category in which directed limits do not necessarily exist.[6]
e) Show that a field extension L/K is algebraic iff it is the direct limit of its finite subextensions.

3.2. **Some Impossible Constructions.**

The results we have derived so far do not look very deep to modern eyes, but they were recognized in the 19th century to imply negative solutions to several of the longest standing open problems in mathematics. Namely, the Greeks were interested in **constructibility** of quantities using a compass and a straightedge. We recall the basic setup: one starts out with two distinct points in the plane, which we may as well view as being a unit distance apart. We have at our disposal an unmarked straightedge, so that given any two points we may construct the line passing through them, and a compass, such that given any previously constructed point P_1 and any previously constructed pair of points P_2, P_3, we may draw a circle whose center is P_1 and whose radius is the distance between P_2 and P_3. Let

[6]Suggestion: impose some finiteness condition on one of the above categories.

us say that a positive real number α is **constructible** if we can after a finite sequence of steps construct points P, P' with distance α (more precisely, α times the unit distance we started with), and let us agree that a negative number α is constructible iff $|\alpha|$ is constructible. Despite the severely constrained toolkit, the supply of constructible numbers is in some sense rather large.

Exercise 3.5. *a) Show: the constructible numbers form a subfield of* $\mathbb{R}$.
b) Show: if $\alpha > 0$ *is constructible, then so is* $\sqrt{\alpha}$.
c) Conclude: the field of constructible numbers has infinite degree over $\mathbb{Q}$.

Now let us look more closely: a constructible number is built up in a sequence of steps: $\alpha_1 = 1$, α_2, $\ldots \alpha_n = \alpha$ corresponding to a tower of fields $F_1 = \mathbb{Q}$, $F_2 = F_1(\alpha_2)$, ..., $F_n = F_{n-1}(\alpha_n)$. To get from F_i to $F_{i+1} = F_i(\alpha_i)$, we are either intersecting two lines – which corresponds to solving a linear equation with coefficients in F_{i-1}, so $F_i = F_{i-1}$ – or intersecting a line defined over F_{n-1} with a circle whose coefficients lie in F_{i-1} which yields solutions in either F_{i-1} or a quadratic extension of F_{i-1} – or we are intersecting two circles with equations defined over F_{i-1}, which leads to solutions over at worst a quadratic extension of a quadratic extension of F_{i-1}. (Note quadratic, not quartic: any two distinct circles intersect in at most two points, and thus the common intersection can also be expressed as the intersection of a line and a circle.)

Thus any constructible number α lies in a field which is at the top of a tower of quadratic field extensions, so $[\mathbb{Q}(\alpha) : \mathbb{Q}]$ is a power of 2. The impossibility of three classically sought after constructions follows easily.

First we cannot **double the cube**: given a cube with sides of our unit length, we cannot construct a cube whose volume is twice that of the given cube, because the length of a side would be $\sqrt[3]{2}$, and $[\mathbb{Q}(\sqrt[3]{2}) : \mathbb{Q}] = 3$. Similarly we can construct angles that we cannot trisect; in particular, we can construct an angle of 60 degrees (i.e., we can construct $\cos 60^o = \frac{1}{2}$ and $\sin 60^o = \frac{\sqrt{3}}{2}$), but we cannot construct $\cos 20^o$ since it satisfies an irreducible cubic polynomial over $\mathbb{Q}$. Finally, we cannot **square the circle** i.e., construct a square whose area is that of a unit circle, for that would involve constructing a side length of $\sqrt{\pi}$ and π is not even algebraic!

3.3. **Subfields of Algebraic Numbers.** Let L/K be an arbitrary extension of fields. Consider the set $\mathrm{Cl}_L(K)$ of all elements of L which are algebraic over K. For example, when $K = \mathbb{Q}$, $L = \mathbb{C}$ we are examining the set of all algebraic numbers, which is certainly a proper subset of $\mathbb{C}$.

Proposition 3.4. *The set* $\mathrm{Cl}_L(K)$ *is a subfield of* K.

We often refer to $\mathrm{Cl}_L(K)$ as the **algebraic closure of** K **in** L.

Let us this result in a more general context, that of integral extensions of domains. The generalized proof is not much harder and will be extremely useful for any student of algebra. So: let R be a domain and S a domain which extends R, i.e., there is an injective homomorphism $R \to S$. We say that $\alpha \in S$ is **integral over** R if α satisfies a monic polynomial with R-coefficients:

$$\exists\ a_{n-1}, \ldots, a_0 \in R \mid \alpha^n + a_{n-1}\alpha^{n-1} + \ldots + a_1\alpha + a_0 = 0.$$

We say that the extension S/R is **integral** if every element of S is integral over R.

Note that if R and S are fields, $\alpha \in S$ is integral over R is by definition precisely the same as being algebraic over R. The next result in fact revisits the basic finiteness property of algebraic elements in this more general context.

Theorem 3.5. *Let $R \subset T$ be an inclusion of rings, and $\alpha \in T$. TFAE:*
(i) α is integral over R.
(ii) $R[\alpha]$ is finitely generated as an R-module.
(iii) There exists an intermediate ring $R \subset S \subset T$ such that $\alpha \in S$ and S is finitely generated as an R-module.
(iv) There exists a faithful $R[\alpha]$-submodule M of T which is finitely generated as an R-module.

Proof. (i) $\implies$ (ii): If α is integral over R, there exist $a_0, \ldots, a_{n-1} \in R$ such that

$$\alpha^n + a_{n-1}\alpha^{n-1} + \ldots + a_1\alpha + a_0 = 0,$$

or equivalently

$$\alpha^n = -a_{n-1}\alpha^{n-1} - \ldots - a_1\alpha - a_0.$$

This relation allows us to rewrite any element of $R[\alpha]$ as a polynomial of degree at most $n-1$, so that $1, \alpha, \ldots, \alpha^{n-1}$ generates $R[\alpha]$ as an R-module.
(ii) $\implies$ (iii): Take $T = R[\alpha]$.
(iii) $\implies$ (iv): Take $M = S$.
(iv) $\implies$ (i): Let $m_1, \ldots, m_n$ be a finite set of generators for M over R, and express each of the elements $m_i\alpha$ in terms of these generators:

$$\alpha m_i = \sum_{j=1}^{n} r_{ij} m_j, \ r_{ij} \in R.$$

Let A be the $n \times n$ matrix $\alpha I_n - (r_{ij})$; then recall from linear algebra that

$$AA^* = \det(A) \cdot I_n,$$

where A^* is the "adjugate" matrix (of cofactors). If $m = (m_1, \ldots, m_n)$ (the row vector), then the above equation implies $0 = mA = mAA^* = m \det(A) \cdot I_n$. The latter matrix equation amounts to $m_i \det(A) = 0$ for all i. Thus $\bullet \det(A) = \bullet 0$ on M, and by faithfulness this means $\det(A) = 0$. Since so that α is a root of the monic polynomial $\det(T \cdot I_n - (a_{ij}))$. □

Lemma 3.6. *Let $R \subset S \subset T$ be an inclusion of rings. If $\alpha \in T$ is integral over R, then it is also integral over S.*

Proof. If α is integral over R, there exists a monic polynomial $P \in R[t]$ such that $P(\alpha) = 0$. But P is also a monic polynomial in $S[t]$ such that $P(\alpha) = 0$, so α is also integral over S. □

Lemma 3.7. *Let $R \subset S \subset T$ be rings. If S is finitely generated as an R-module and T is finitely generated as an S-module then T is finitely generated as an R-module.*

Proof. If $\alpha_1, \ldots, \alpha_r$ generates S as an R-module and $\beta_1, \ldots, \beta_s$ generates T as an S-module, then $\{\alpha_i\beta_j\}\{1 \leq i \leq r, 1 \leq j \leq s\}$ generates T as an R-module: for $\alpha \in T$, we have

$$\alpha = \sum_j b_j\beta_j = \sum_i \sum_j (a_{ij}\alpha_i)\beta_j,$$

with $b_j \in S$ and $a_{ij} \in R$. □

Corollary 3.8. *(Transitivity of integrality) If $R \subset S \subset T$ are ring extensions such that S/R and T/S are both integral, then T/R is integral.*

Proof. For $\alpha \in S$, let $\alpha^n + b_{n-1}\alpha^{n-1} + \ldots + b_1\alpha + b_0 = 0$ be an integral dependence relation, with $b_i \in S$. Thus $R[b_1, \ldots, b_{n-1}, \alpha]$ is finitely generated over $R[b_1, \ldots, b_{n-1}]$. Since S/R is integral, $R[b_1, \ldots, b_{n-1}]$ is finite over R. By Lemma 3.7, $R[b_1, \ldots, b_{n-1}, \alpha]$ is a subring of T containing α and finitely generated over R, so by Theorem 3.5, α is integral over R. □

Corollary 3.9. *If S/R is a ring extension, then the set $I_S(R)$ of elements of S which are integral over R is a subring of S, the **integral closure of** R **in** S. Thus $R \subset I_S(R) \subset S$.*

Proof. If $\alpha \in S$ is integral over R, $R[\alpha_1]$ is a finitely generated R-module. If α_2 is integral over R it is also integral over $R[\alpha_1]$, so that $R[\alpha_1][\alpha_2]$ is finitely generated as an $R[\alpha_1]$-module. By Lemmma 3.7, this implies that $R[\alpha_1,\ \alpha_2]$ is a finitely generated R-module containing $\alpha_1 \pm \alpha_2$ and $\alpha_1 \cdot \alpha_2$. By Theorem 3.5, this implies that $\alpha_1 \pm \alpha_2$ and $\alpha_1\alpha_2$ are integral over R. □

If $R \subset S$ such that $I_S(R) = R$, we say R is **integrally closed** in S.

Proposition 3.10. *Let S be a ring. The operator $R \mapsto I_S(R)$ on subrings of R is a closure operator in the abstract sense, namely it satisfies:*
(CL1) $R \subset I_S(R)$,
(CL2) $R_1 \subset R_2 \implies I_S(R_1) \subset I_S(R_2)$.
(CL3) $I_S(I_S(R)) = I_S(R)$.

Proof. (CL1) is the (trivial) Remark 1.1. (CL2) is obvious: evidently if $R_1 \subset R_2$, then every element of S which satisfies a monic polynomial with R_1-coefficients also satisfies a monic polynomial with R_2-coefficients. Finally, suppose that $\alpha \in S$ is such that $\alpha^n + a_{n-1}\alpha^{n-1} + \ldots + a_1\alpha + a_0 = 0$ for $a_i \in I_S(R)$. Then each a_i is integral over R, so $R[a_1, \ldots, a_n]$ is finitely generated as an R-module, and since $R[a_1, \ldots, a_n, \alpha]$ is finitely generated as an $R[a_1, \ldots, a_n]$-module, applying Lemma 3.7 again, we deduce that α lies in the finitely generated R-module $R[a_1, \ldots, a_n, \alpha]$ and hence by Theorem 3.5 is integral over R. □

Proposition 3.11. *Let $R \subset S$ be an integral extension. If R is a field, so is S.*

Proof: Let L be the fraction field of S. If $0 \neq \alpha \in S$ is integral over R, then by Theorem 3.5, $R[\alpha]$ is a finite-dimensional R-submodule of L, so it is a subfield, i.e., is equal to $R(\alpha)$. So $R(\alpha) = R[\alpha] \subset S$, meaning that S contains α^{-1}.

3.4. **Distinguished Classes.**

Here is an organizing principle for classes of field extensions due to S. Lang.

A class $\mathcal{C}$ of field extensions is **distinguished** if it satisfies these two properties:

(DC1) (Tower meta-property) For a tower $M/K/F$, then $M/F \in \mathcal{C}$ iff $M/K \in \mathcal{C}$ and $K/F \in \mathcal{C}$.
(DC2) (Base change meta-property) Let K/F be an element of $\mathcal{C}$, let L/F be any extension such that K and L are contained in a common field. Then $LK/L \in \mathcal{C}$.

We note that (DC1) and (DC2) imply the following

(DC3) (Compositum meta-property) Let K_1/F and K_2/F be elements of $\mathcal{C}$ with K_1, K_2 contained in a common field. Then $K_1K_2/F \in \mathcal{C}$.

Indeed, applying (DC2) we get that $K_1K_2/K_2 \in \mathcal{C}$. Since also $K_2/F \in \mathcal{C}$, applying (DC1) we get that $K_1K_2/F \in \mathcal{C}$.

Exercise 3.6. *a) Show: the class of all finite degree extensions is distinguished.*
b) Show: the class of all algebraic extensions is distinguished.

Some examples of distinguished classes of extensions to come later: finitely generated extensions, separable algebraic extensions, purely inseparable algebraic extensions, solvable extensions, purely transcendental extensions.

Some nonexamples of distinguished classes of extensions to come later: normal extensions, Galois extensions, inseparable extensions, abelian extensions, not-necessarily-algebraic separable extensions.

4. Normal Extensions

4.1. Algebraically closed fields.

Let F be a field. A polynomial $f \in F[t]$ is **split** if every irreducible factor has degree 1. If $f \in F[t]$ is a polynomial and K/F is a field extension, we say f **splits in K** if $f \in K[t]$ is split.

Proposition 4.1. *Let F be a field. The following are equivalent:*
(i) There is no algebraic extension $K \supsetneq F$.
(ii) There is no finite degree extension $K \supsetneq F$.
(iii) There is no finite degree monogenic extension $F(\alpha) \supsetneq F$.
(iv) If $f \in F[t]$ is irreducible, then f has degree 1.
(v) If $f \in F[t]$ is nonconstant, then f has a root in F.
*(vi) Every polynomial $f \in F[t]$ is split. A field satisfying these equivalent conditions is called **algebraically closed**.*

Proof. (i) $\implies$ (ii) $\implies$ (iii) is immediate.
$\neg$ (iv) $\implies$ $\neg$ (iii): if $f \in F[t]$ is an irreducible polynomial of degree $d > 1$ then $K := F[t]/(f)$ is a finite degree monogenic extension of f of degree $d > 1$.
$\neg$ (v) $\implies$ $\neg$ (iv): Suppose f is nonsconstant and admits no root in F. Write $f = f_1 \cdots f_m$ as a product of irreducible polynomials; since linear polynomials have roots in F, no f_i has degree 1.
(iv) $\iff$ (v) $\iff$ (vi) is easy and familiar.
$\neg$ (i) $\implies$ $\neg$ (iv): If $K \supsetneq F$ is a proper algebraic extension, let $\alpha \in K \setminus F$, and let $f \in F[t]$ be the minimal polynomial of α over F, so f is irreducible. By assumption f is also split, so it has degree 1 and is thus of the form $t - \alpha$, contradicting the fact that $\alpha \notin F$. □

Theorem 4.2. *(Fundamental Theorem of Algebra)*
The complex field $\mathbb{C}$ is algebraically closed.

Because the existence of a nonconstant $f \in \mathbb{C}[t]$ without a root in $\mathbb{C}$ leads to absurdities in many areas of mathematics, there are many different proofs, e.g.

using degree theory or complex analysis. It is often held that "fundamental theorem of algebra" is a misnomer, in that the result concerns a structure – the complex numbers – whose definition is in part analytic/topological. We do not dispute this. Nevertheless the true algebraist hankers for an algebraic proof, and indeed this is possible. We may, in fact, view Theorem 4.2 as a special case of the following result, whose proof requires Galois theory so must be deferred until later.

Theorem 4.3. *(Artin-Schreier) Suppose K is a field with the following properties:*
(i) There do not exist *$n \in \mathbb{Z}^+$ and $x_1, \ldots, x_n \in K$ such that $-1 = x_1^2 + \ldots + x_n^2$.*
(ii) Every polynomial $P \in K[t]$ of odd degree has a root in K.
(iii) For any $x \in K^\times$, exactly one of x and $-x$ is a square in K.
Then $K[\sqrt{-1}] = K[t]/(t^2+1)$ is algebraically closed.

Exercise 4.1. *Show: Theorem 4.3 implies Theorem 4.2.*

Proposition 4.4. *Let L/K be a field extension, and let $\mathrm{Cl}_K(L)$ be the algebraic closure of K in L: that is, the set of all elements of L that are algebraic over K. Then $\mathrm{Cl}_K(L)$ is algebraically closed.*

Proof. Put $\overline{K} \coloneqq \mathrm{Cl}_K(L)$. By Proposition 4.1, if $\overline{K}$ is not algebraically closed then there is a monogenic finite degree extension $\overline{K}(\alpha) \supsetneq \overline{K}$. Because α is algebraic over $\overline{K}$ and $\overline{K}$ is algebraic over K, we have by Corollary 3.8 that α is algebraic over K. Let $f \in F[t]$ be the minimal polynomial of α. By Proposition 4.1, as a polynomial over $L[t]$ we have

$$f(t) = (t-\alpha_1)(t-\alpha_2)\cdots(t-\alpha_d)$$

for some $\alpha_1, \ldots, \alpha_d \in L$. Indeed each α_i is algebraic over K so lies in $\overline{K}$. Moreover the $\alpha_1, \ldots, \alpha_d$ are the only roots of f in L, and thus for some i we have $\alpha = \alpha_i \in \overline{K}$, a contradiction. □

Corollary 4.5. *The field $\overline{\mathbb{Q}}$ of all algebraic numbers is algebraically closed.*

Proof. Since $\overline{\mathbb{Q}}$ is the algebraic closure of $\mathbb{Q}$ in $\mathbb{C}$, this follows from Theorem 4.2 and Proposition 4.4. □

Let K be a field. An **algebraic closure** of K is a field extension $\overline{K}/K$ that is both algebraic and algebraically closed. It follows from Proposition 4.1 an algebraic closure of K is precisely a maximal algebraic extension of K, i..e., an algebraic extension that is not properly contained in any other algebraic extension of K.

Exercise 4.2. *Let K/F be an algebraic field extension. Let L/K be a field extension. Show: L is an algebraic closure of K iff L is an algebraic closure of F.*

4.2. **Existence of algebraic closures.**

In this section we will show that every field admits at least one algebraic closure, a basic but nontrivial result.

How might one try to prove this? Probably we can agree to start with the following easy result.

Lemma 4.6. *Let F be a field, and let $f_1, \ldots, f_n \in F[t]$ be nonconstant polynomials, of degrees $d_1, \ldots, d_n$.*
a) There is a finite degree field extension K/F such that each f_i has a root in K.

Moreover, we can choose K so as to get $[K : F] \leq \prod_{i=1}^n d_i$.
b) There is a finite degree field extension K/F such that each f_i splits in K. Moreover, we can choose K so as to get $[K : F] \leq \prod_{i=1}^n d_i!$.

Proof. a) Let M be a field, and let $f \in M[t]$ be a polynomial of degree d. Let g be an irreducible factor of f, say of degree $d' \leq d$. Then $M[t]/(g)$ is a field extension of M of degree $d' \leq d$ in which g (and hence also f) has a root. By applying this procedure successively to $f_1, \ldots, f_n$ we generate a tower of field extensions $F \subset M_1 \subset \ldots \subset M_n$ such that for all $1 \leq i \leq n$, the polynomials $f_1, \ldots, f_i$ all have a root in M_i and $[M_i : F] \leq d_1 \cdots d_i$, so we may take $K := M_n$.
b) Let M be a field, and let $h \in M[t]$ be a polynomal of degree d. Applying part a) to h, there is a field extension M_1/M of degree at most d in which h has a root α_1 and thus we get a factorization $h(t) = (t - \alpha_1)h_2(t) \in M_1[t]$. We apply part a) to h_2 and get a field extension M_2/M_1 of degree at most $d-1$ in which h_2 has a root α_2 and thus we get a factorization $h(t) = (t - \alpha_1)(t - \alpha_2)h_3(t)$. Continuing in this manner, we end up with a field extension M_n of degree at most $d!$ in which h splits. Applying this procedure successively to the polynomials $f_1, \ldots, f_n$ over the field F we get a field extension K of degree at most $\prod_{i=1}^n d_i!$ in which each f_i splits. □

Exercise 4.3. *Let $d_1, \ldots, d_n \in \mathbb{Z}^+$.*
a) Show: there are (necessarily irreducible) polynomials $f_1, \ldots, f_n \in \mathbb{Q}[t]$ such that if $K/\mathbb{Q}$ is a number field (i.e., a finite degree field extension) such that each f_i has a root in K then $\prod_{i=1}^n \deg f_i \mid [K : \mathbb{Q}]$.
b) Show: there are (necessarily irreducible) polynomials $f_1, \ldots, f_n \in \mathbb{Q}[t]$ such that if $K/\mathbb{Q}$ is a number field in which each f_i splits, then $\prod_{i=1}^n d_i! \mid [K : \mathbb{Q}]$.
(Hint/warning: this is most naturally done using basic algebraic number theory.)

Theorem 4.7. *Every field K can be embedded in an algebraically closed field L. Thus every field has at least one algebraic closure, namely $\mathrm{Cl}_L(K)$.*

Proof. Step 1: Let $R = K[\mathbb{T}]$ be a polynomial ring over K indexed by a set of indeterminates t_f that are in bijection with the nonconstant polynomials $f \in K[t]$. Consider the ideal I of R generated by all polynomials of the form $f(t_f)$. We claim that I is proper: if not, there is a finite subset $\{f_1, \ldots, f_n\}$ and elements $g_1, \ldots, g_n \in R$ such that

$$g_1 f_1(t_{f_1}) + \ldots + g_n f_n(t_{f_n}) = 1.$$

By Lemma 4.5, there is a finite degree field extension F/K such that each $f_i(t)$ has a root $\alpha_i \in F$. If we evaluate $t_{f_1} = \alpha_1, \ldots, t_{f_n} = \alpha_n$ in the above equation, we get $0 = 1$: contradiction. So we may choose a maximal ideal $\mathfrak{m} \supset I$. Thus $K_1 := R/\mathfrak{m}$ is a field exetnsion of F in which each t_f is a root of f. Thus K_1/K is a field extension in which each nonconstant polynomial $f \in K[t]$ has a root.
Step 2: The natural question here is whether K_1 is algebraically closed. The remainder of the proof consists of a clever evasion of this question! Namely, we apply the construction of Step 1 to K_1, getting a field extension K_2 in which each polynomial with coefficients in K_1 has a root in K_2, and so forth: we generate a sequence of field extensions

$$K \subset K_1 \subset \ldots \subset K_n \subset \ldots.$$

The union $L = \bigcup_n K_n$ is a field, and any nonconstant polynomial $P \in L[t]$, having only finitely many nonzero coefficients lies in $K_n[t]$ for sufficiently large n, thus has

a root in K_{n+1} and therefore also in L. So L is algebraically closed, and then by Proposition 4.4 the algebraic closure of K in L is an algebraic closure of K. □

Theorem 4.7 lies among the most important results in all of field theory. So we pause to discuss several aspects of it.

First, the proof of Theorem 4.7 used the Axiom of Choice (AC) in a somewhat disguised way: in the assertion that a proper ideal in a ring is contained in a maximal ideal. In fact the statement that every proper ideal in a commutative ring is contained in a maximal ideal implies (AC). So it is natural to wonder whether the existence of an algebraic closure of any field implies (AC). Indeed not: it would be enough to use that every proper ideal is contained in a prime ideal: this gives us a domain, and we can take the fraction field. The assertion that every proper ideal in a commutative ring is contained in a prime ideal is known to be equivalent to the Ultrafilter Lemma (UL), which does *not* imply (AC).

It seems to be an open problem whether the existence of an algebraic closure of every field implies (UL): cf. `http://mathoverflow.net/questions/46566`. However, it is known that (AC) is required for Theorem 4.7 to hold in the sense that there is a model of Zermelo-Fraenkel set theory in which not every field admits an algebraic closure [Je, Thm. 10.13].

The proof of Theorem 4.7 comes from E. Artin by way of Lang [LaFT]. It is unnecessarily (though helpfully) slick in several respects. The use of polynomial rings is a crutch to avoid some mostly set-theoretic unpleasantries: later we wil see that an algebraic closure of F is essentially the direct limit of all finite degree normal field extensions K/F: here the essentially means that we want each K/F to appear exactly once up to F-isomorphism. It just happens that the easiest way to do that is to realize each K inside a fixed algebraically closed field containing F! But by the time the reader has made it to the end of this section, she may consider trying to construct this direct limit directly.

Finally, as we pointed out, the proof constructs an extension K_1/K such that every nonconstant $f \in K[t]$ has a root in K_1 and then nimbly evades the question of whether K_1 contains an algebraic closure of K. It turns out that the answer to this is affirmative. We break this up into two steps. First:

Proposition 4.8. *Let L/K be a field extension. Suppose every nonconstant $f \in K[t]$ splits in L. Then the algebraic closure of K in L is algebraically closed.*

Proof. Let $\overline{K}$ be the algebraic closure of K in L. Suppose $\overline{K}$ is not algebraically closed: then by Proposition 4.1 there is a field extension $M/\overline{K}$ and an element $\alpha \in M \setminus \overline{K}$ that is algebraic over $\overline{K}$. By Corollary 3.8 we have that α is algebraic over K, so has a minimal polynomial $f \in K[t]$. By assumption f splits in $\overline{K}$, and since $f(\alpha) = 0$ one of the factors of f must be $t - \alpha$ and thus $\alpha \in \overline{K}$. □

As for the second step: we will record the answer now, but we will need to know more of the structure theory of algebraic field extensions in order to prove it.

Theorem 4.9. *(Gilmer* [Gi68]*) Let L/K be a field extension. If every nonconstant $f \in K[t]$ has a root in L, then every nonconstant $f \in K[t]$ splits in L.*

Exercise 4.4. *Let K be a field, and let $f \in K[t]$ be a monic polynomial of degree $d \geq 1$. Let $\overline{K}$ be an algebraic closure of K. Over $\overline{K}$, f splits:*

$$f(t) = (t - \alpha_1) \cdots (t - \alpha_d).$$

*We say f is **separable** if the $\alpha_1, \ldots, \alpha_d$ are distinct elements of $\overline{K}$.*
a) Conceivably the above definition depends on the choice of $\overline{K}$. However, let f' be the (formal) derivative of f: the unique K-linear endomorphism of $K[t]$ such that $(t^n)' = nt^{n-1}$. Show: f is separable iff $\gcd(f, f') = 1$.
b) Let K be a field and $n \in \mathbb{Z}^+$. If K has positive characteristic p, assume that $\gcd(n, p) = 1$. Let $a \in K$ be arbitrary. Show: the polynomial $t^n - a$ is separable.
c) Deduce: if K is a field and $n \in \mathbb{Z}^+$ is prime to the characteristic of K if it is positive, then there is a field extension L/K containing n different nth roots of unity: i.e., distinct $z_1, \ldots, z_n$ such that $z_i^n = 1$ for all 1.
d) Deduce: no finite field is algebraically closed.

Exercise 4.5.
a) Show: if K is a field and $\overline{K}$ is an algebraic closure, then $\#\overline{K} = \max(\aleph_0, \#K)$.
b) Show: there are algebraically closed fields of all infinite cardinalities.

4.3. The Magic Mapping Theorem.

Theorem 4.10. *(Magic Mapping Theorem) Let F be a field. Let K/F be an algebraic field extension, and let L/F be a field extension with L algebraically closed. Then there is an F-algebra homomorphism $\varphi : K \hookrightarrow L$.*

Proof. Consider the partially ordered set whose elements are are pairs (M, φ) where M is a subextension of K/F and $\varphi : M \to L$ is an F-algebra homomorphism. We say that $(M_1, \varphi_1) \leq (M_2, \varphi_2)$ if $M_1 \subset M_2$ and the restriction of φ_2 to M_1 is φ_1. In this partially ordered set, any chain has an upper bound given by taking the union of the elements of the chain. So by Zorn's Lemma there is a maximal element (M, φ). We claim that $M = K$. If not, let $\alpha \in K \setminus M$, and consider the field extension $M(\alpha)/M$. Let $f \in M[t]$ be the minimal polynomial of α, so $M(\alpha) \cong M[t]/(f)$. We view L as an M-algebra via φ, and thus we may view $f \in L[t]$. Since L is algebraically closed, there is a root in L, say $\overline{\alpha}$. There is a unique M-algebra homomorphism $M(\alpha) \to L$ that maps α to $\overline{\alpha}$: it is unique because $M(\alpha) = M[\alpha]$ is generated as an M-algebra by α, and it exists because $M(\alpha) \cong M[t]/(f(t))$ so the unique M-algebra map $M[t] \to L$ that carries t to $\overline{\alpha}$ has $f(t)$ in its kernel. It follows that $M = K$. □

Corollary 4.11. *("Uniqueness" of Algebraic Closure) Let $\overline{F_1}$ and $\overline{F_2}$ be two algebraic closures of a field F. Then there is an F-algebra isomorphism $\varphi : \overline{F}_1 \to \overline{F_2}$.*

Proof. We may apply the Magic Mapping Theorem with $K = \overline{F_1}$ and $L = \overline{F_2}$ to get an F-algebra homomorphism $\varphi : \overline{F_1} \hookrightarrow \overline{F_2}$. Then $\overline{F_2}/\varphi(\overline{F_1}))$ is an algebraic extension of an algebraically closed field, so it cannot be proper: we have $\overline{F_2} = \varphi(\overline{F_1})$ and thus φ is an F-algebra isomorphism. □

Note that we speak of "'uniqueness' of the algebraic closure" rather than "uniqueness of the algebraic closure." This is because we have shown that the algebraic closure of F is unique up to F-algebra isomorphism, but given two algebraic closures of F there is in general no *canonical* F-algebra isomorphism between them. If $\varphi, \psi : \overline{F_1} \hookrightarrow \overline{F_2}$ are two F-algebra isomomorphisms, then $\psi^{-1} \circ \varphi$ is an F-algebra

automorphism of $\overline{F_1}$, and conversely: the ambiguity in the choice of isomorphism is precisely measured by the group $G_F := \mathrm{Aut}(\overline{F}_1/F)$. This group is called the **absolute Galois group of F** and is in general a very large, interesting group. In fact, we should not speak of "the" absolute Galois group of F (though we will: it is traditional to do so): it is well-defined up to isomorphism, but switching from one isomorphism $\overline{F_1} \to \overline{F_2}$ to another gives rise to an inner automorphism (i.e., a conjugation) of G. More on this later.

Remark 4.1. *There are models of Zermelo-Fraenkel set theory – i.e., without (AC) – in which a field F can admit non-F-isomorphic algebraic closures.*

Corollary 4.12. *Let K_1/F and K_2/F be two algebraic field extensions. If $\varphi : K_1 \to K_2$ is any F-algebra embedding and $\overline{K_i}$ is any algebraic closure of K_i, then φ extends to an isomorphism $\overline{K_1} \to \overline{K_2}$.*

Exercise 4.6. *Prove Corollary 4.12.*

4.4. **Conjugates.**

Let K/F be an algebraic field extension. We say that elements $\alpha, \beta \in K$ are **conjugate over F** if α and β have the same minimal polynomial over F. If $\alpha \in K \setminus F$ has degree $d \geq 2$ – i.e., $[K(\alpha) : K] = d$ or equivalently the degree of the minimal polynomial of α is d – then the number of conjugates of α is at least 2 and at most d.

If K/F is an algebraic extension and $\overline{F}$ is any algebraic closure of F, then as we know there is an F-algebra homomorphism $\iota : K \hookrightarrow \overline{F}$. If $\alpha \in K$ and $f \in F[t]$ is the minimal polynomial of α, then f splits in $\overline{F}$. We call the roots of f in $\overline{F}$ the **conjugates** of α. Notice that the set of conjugates is defined only in terms of the minimal polynomial, which lies in F, so it is independent of the choice of ι. If $\alpha \in F$ then of course α is the only conjugate of α, whereas if $\alpha \notin F$ then once again if α has degree d then the set of conjugates of α has size at least 2 and at most d.

For the remainder of this section we fix an algebraic closure $\overline{F}$ of F and only consider algberaic extensions K/F that are subextensions of $\overline{F}/F$ (again, every algebraic extensions occurs this way *up to F-algebra isomorphism*). From this perspective, being conjugate over F is an equivalence relation on $\overline{F}$. Moreover, if σ is an F-algebra automorphism of $\overline{F}$, then for all $\alpha \in \overline{F}$, we have that $\sigma(\alpha)$ is a conjugate of α: indeed, for every polynomial $f \in F[t]$, we have

$$f(\alpha) = 0 \iff f(\sigma(\alpha)) = 0$$

and thus α and $\sigma(\alpha)$ have the same minimal polynomial. Conversely, if $\alpha, \beta \in \overline{F}$ are conjugate over F, then there is an F-algebra automorphism σ of $\overline{F}$ such that $\sigma(\alpha) = \beta$. Indeed, let $f \in F[t]$ be the common minimal polynomial of α and β. Then the field extensions $F(\alpha)$ and $F(\beta)$ are both isomorphic to $F[t]/(f(t))$, so there is an isomorphism

$$F(\alpha) \to F(\beta),$$

which by Corollary 4.12 extends to an automorphism of $\overline{F}$.

Remark 4.2. *Recall that if a group G acts on a set X, we say that two elements $x, y \in X$ are **conjugate** if there is $g \in X$ such that $gx = y$. As we just saw, the terminology of conjugate elements of $\overline{F}$ is compatible with this: two elements of $\overline{F}$ are conjugate iff they are conjugate under the action of* $\mathrm{Aut}(\overline{F}/F)$.

4.5. **Splitting Fields.**

It follows from Proposition 4.8 that if K/F is an algebraic field extension such that every nonconstant $f \in F[t]$ splits in K, then K is an algebraic closure of F. This view on algebraic closure opens the door to a natural and important generalization: we go from "all polynomials" to "some polynomials."

Let F be a field, and let $\mathcal{S} \subset F[t]$ be a set of nonconstant polynomials. A **splitting field** for $(F, \mathcal{S})$ is a field extension K/F satisfying the following properties:

(SF1) Every $f_i \in \mathcal{S}$ splits in K.
(SF2) No proper subextension of K satisfies (SF1), i.e., if $F \subset K' \subset K$ and every $f_i \in \mathcal{S}$ splits in K', then $K' = K$.

Exercise 4.7. *Suppose K/F is a splitting field for $(F, \mathcal{S})$, and K' is an F-algebra isomorphic to K. Show: K' is also a splitting field for $(F, \mathcal{S})$.*

Theorem 4.13. *(Existence and "Uniqueness" of Splitting Fields) Let F be a field and $S \subset F[t]$ a set of nonconstant polynomials.*
a) Any algebraic closure $\overline{F}$ contains a unique splitting field for S, namely the subfield of $\overline{F}$ obtained by adjoining to F all roots α_{ij} of all polynomials $P_i \in S$.
b) Splitting fields are unique up to F-algebra isomorphism.

Proof. It is no problem to see that the recipe of part a) does indeed construct a splitting field for F and S: clearly every polynomial in S splits in $F(\alpha_{ij})$ and conversely any subfield of $\overline{F}$ in which all the polynomials in F split must contain all the α_{ij}'s. One way to see the uniqueness up to isomorphism is to reduce to the case of uniqueness up to isomorphism of algebraic closures. Namely, let K_1, K_2 be two splitting fields for F and S. It is easy to see that (SF2) implies that K_i/F is algebraic, so let $\overline{K}_i$ be an algebraic closure of K_i. Since K_i is algebraic over F, $\overline{K}_i$ is equally well an algebraic closure of F, so by Corollary 4.11 there exists an F-algebra isomorphism $\Phi : \overline{K}_1 \to \overline{K}_2$. Then $\Phi(K_1)$ is a subfield of $\overline{K}_2$ which is a splitting field for F and S, and we just saw that each algebraic closure contains a unique splitting field, so $\Phi(K_1) = K_2$ and $\Phi : K_1 \to K_2$ is an F-algebra isomorphism. □

Exercise 5.2.2: Show that the field $K = \mathbb{Q}[\zeta_3, \sqrt[3]{2}]$ discussed in §5.1 is the splitting field of $F, \{t^3 - 2\}$. Conclude that if $L \subset \mathbb{C}$ is such that $L \neq K$, then L is not isomorphic to K.

4.6. **Normal Extensions.**

Lemma 4.14. *Let $\overline{F}$ be an algebraic closure of F, let K be a subextension of $\overline{F}/F$, and let $\sigma : K \hookrightarrow \overline{F}$ be an F-algebra embedding. The following are equivalent:*
(i) $\sigma(K) \subset K$.
(ii) $\sigma(K) \supset K$.
(iii) $\sigma(K) = K$.

Proof. Certainly (iii) implies both (i) and (ii). We will show (i) $\implies$ (iii), and it will be clear how to modify the argument so as to obtain (ii) $\implies$ (iii).
(i) $\implies$ (iii): Let $\alpha \in K$, and let S be the set of F-conjugates of α that lie in K. We observe that S is a finite set containing α. For $\beta \in K$, we have that β is a conjugate of α iff $\sigma(\beta)$ is a conjugate of α, so the set of F-conjugates of α that lie in

$\sigma(K)$ is precisely $\sigma(S)$. By hypohesis we have $\sigma(S) \subset S$; since both are finite sets of the same cardinality we must have $\sigma(S) = S$ and thus $\alpha \in \sigma(S) \subset \sigma(K)$. □

Theorem 4.15. *Let K/F be an algebraic field extension. Let $\overline{F}$ be an algebraic closure of K (hence also of F). The following are equivalent:*
(i) For every F-algebra embedding $\sigma : K \hookrightarrow \overline{F}$ we have $\sigma(K) = K$.
(ii) K/F is the splitting field of a subset $\mathcal{S} \subset F[t]$.
(iii) Every irreducible polynomial $f \in F[t]$ with a root in K splits in K.
(iv) For all $\alpha \in K$, if $\beta \in \overline{F}$ is an F-conjugate of α, then $\beta \in K$.
An extension K/F satisfying these properties is called ***normal****.*

Proof. (i) $\iff$ (iv): We saw above that for $\alpha \in K$ and $\beta \in \overline{F}$, β is a conjugate of α in $\overline{F}$ iff there is an F-algebra homomorphism $\sigma : K \hookrightarrow \overline{F}$ such that $\sigma(\alpha) = \beta$. It follows that as we range over all F-algebra homomorphisms $\sigma : K \hookrightarrow \overline{F}$, we have that $\bigcup_\sigma \sigma(K)$ is the set of all conjugates of all elements of K. Condition (iv) holds iff the set of all conjugates of all elements of K is just K itself iff $\bigcup_\sigma \sigma(K) = K$ iff $\sigma(K) \subset K$ for all σ iff $\sigma(K) = K$ for all $\sigma(K)$: condition (i).
(iii) $\iff$ (iv) is immediate.
(ii) $\iff$ (iv): Condition (ii) can be rephrased by saying that K is generated by adjoining to F a subset $\mathcal{S}$ of $\overline{F}$ that is stable under conjugation. Thus if (iv) holds, then (ii) holds with $\mathcal{S} = K$. Conversely, suppose that K is obtained by adjoining to F a set $\mathcal{S}$ that is stable under conjugation, and let $x \in K$. Then $x = f(\alpha_1, \ldots, \alpha_n)$ is a rational function in elements $\alpha_1, \ldots, \alpha_n \in \mathcal{S}$ with F-coefficients. Every conjugate of x in $\overline{F}$ is of the form $\sigma(x)$ for some F-automorphism σ of $\overline{F}$, and then

$$\sigma(x) = f(\sigma(\alpha_1), \ldots, \sigma(\alpha_n)) \in K,$$

since $\mathcal{S}$ is closed under conjugation. □

Corollary 4.16. *Let K/F be a normal algebraic field extension. Let L/K be any field extension and let $\sigma : L \to L$ be any automorphism. Then $\sigma(K) = K$.*

Proof. Let $\overline{L}$ be an algebraic closure of L. By Corollary 4.11, we may extend σ to an automorphism $\sigma : \overline{L} \to \overline{L}$. Let $\overline{K} = \mathrm{Cl}_{\overline{L}}(K)$, the unique algebraic closure of K contained in L. Since K/F is algebraic, $\overline{K}$ is also an algebraic closure of F. Since $\sigma(\overline{K})$ is also an algebraic closure of K contained in $\overline{L}$, by the aforementioned uniqueness we have $\sigma(\overline{K}) = \overline{K}$. By Theorem 4.15 we have $\sigma(K) = K$. □

Exercise 4.8. *Suppose K/F is finite of degree at most 2. Show: K/F is normal.*

Example 4.17. *For each $n \geq 3$, the extension $K = \mathbb{Q}[\sqrt[n]{2}]/\mathbb{Q}$ is a non-normal extension of degree n. Indeed, let $\zeta_n = e^{2\pi i/n}$; then the other roots of $t^n - 2$ in $\mathbb{C}$ are $\zeta_n^k \cdot \sqrt[n]{2}$ with $0 \leq k < n$, which are not even real numbers unless $k = 0$ or $k = \frac{n}{2}$. So $t^n - 2$ does not split over K. In this case, any extension of K which is normal over $\mathbb{Q}$ must contain all the roots of $t^n - 2$, hence must contain $\sqrt[n]{2}$ and ζ_n. Therefore the smallest normal extension is the splitting field of $t^n - 2$, which is $M = \mathbb{Q}[\sqrt[n]{2}, \zeta_n]$.*

Example 4.18. *Suppose F has characteristic $p > 0$. Suppose $a \in F$ is such that $f(t) = t^{p^n} - a \in F[t]$ is irreducible. Let $K = F[t]/(f(t))$, and write α for the coset of $t + (f(t))$: thus $\alpha^{p^n} = a$. Then as an element of $K[t]$ we have $f(t) = (t - \alpha)^{p^n}$. That is, despite the fact that f has degree p^n, α is conjugate in $\overline{F}$ only to itself. Thus K/F is a normal extension.*

Exercise 4.9. *Show: a direct limit of normal extensions is normal.*

Exercise 4.10. *Show: if K/F is normal and L/K is normal, then L/F need not be normal. (Thus normality does not satisfy the* ***tower meta-property****.)*

Exercise 4.11. *a) Let L/F be an extension and K_1, K_2 be subextensions. Show: $K_1 \cap K_2$ is again an extension field of F.*
b) As above, but with any collection of intermediate field extensions $\{K_i\}_{i \in I}$.

Proposition 4.19. *Let L/F be an extension and $\{K_i\}_{i \in I}/F$ a collection of algebraic subextensions. If each K_i/F is normal, then so is the intersection $K = \bigcap_i K_i$.*

Proof. Using Exercise 5.3.7, we may replace K_i by $\mathrm{Cl}_{K_i}(F)$ and L by $\mathrm{Cl}_L(F)$ and thus assume that all field extensions are algebraic. Let $P \in F[t]$ be an irreducible polynomial. If P has a root in K, P has a root in each K_i, hence K_i contains a splitting field for F. Splitting fields are unique inside any given algebraic extension, so this means that each K_i contains the common splitting field for (F, P), hence K contains it, so P splits in K. □

Let K/F be any field extension. As above, $\overline{K}/F$ is certainly normal. Since the intersection of any family of normal subextensions of $\overline{K}$ is normal, it follows that there is a unique smallest subextension L, $F \subset K \subset L \subset \overline{K}$, such that L/F is normal. If we define a **normal closure** of an extension K/F to be an extension L/K which is normal over F and such that no proper subextension is normal over F, then we just constructed a normal closure, by intersecting all normal subextensions inside an algebraic closure of K. This shows that any normal closure of K/F is algebraic over K, and by the usual tricks with uniqueness up to F-isomorphism of algebraic closure one can show that the normal closure of an extension is also unique up to F-isomorphism.

Proposition 4.20. *Let K/F be finite of degree n. Then the degree of the normal closure M of K/F (inside any algebraic closure $\overline{K}$) is at most $n!$*

Proof. Put $F = F_0$. Write $K = F(\alpha_1, \ldots, \alpha_d)$ and for $1 \leq i \leq d$, put $K_i = F(\alpha_1, \ldots, \alpha_i)$ and $d_i := [K_i : K_{i-1}]$. An argument almost identical to that of Lemma 4.6b) yields a field extension M/K containing all the conjugates of $\alpha_1, \ldots, \alpha_d$ and such that $[M : F] = \prod_{i=1}^d d_i!$. Thus the normal closure of K/F has degree at most $\prod_{i=1}^d d_i!$. Now

$$n = [K : F] = \prod_{i=1}^d [K_i : K_{i-1}] = \prod_{i=1}^d d_i.$$

It follows that $\prod_{i=1}^d d_i! \leq n!$: for intance take sets $S_1, \ldots, S_d$ of cardinalities $d_1, \ldots, d_d$. Then $\prod_{i=1}^d d_i!$ is the number of bijections of $S := \prod_{i=1}^n S_i$ that preserve each coordinate, while $(d_1 \cdots d_n)!$ is the number of bijectiions of S. □

4.7. **Isaacs' Theorem.**

The goal of this section is to prove the following result of Isaacs.

Theorem 4.21. *(Isaacs* [Is80]*) Let F be a field. For an algebraic extension K/F, let $\mathcal{P}(K)$ be the set of polynomials $f \in F[t]$ having a root in K. Then for algebraic extensions K_1/F, K_2/F, the following are equivalent:*

(i) The F-algebras K_1 and K_2 are isomorphic.
(ii) We have $\mathcal{P}(K_1) = \mathcal{P}(K_2)$.

Exercise 4.12. *a) Show: Isaacs' Theorem implies that a field extension L/F that contains a root of every nonconstant $f \in F[t]$ contains an algebraic closure of F.*
b) Deduce Gilmer's Theorem (Theorem 4.9).

5. Separable Algebraic Extensions

Let K/F be an algebraic field extension. We have already explored one desirable property for K/F to have: normality. Normality can be expressed in terms of stability under F-homomorphisms into any extension field, and also in terms of irreducible polynomials: every irreducible polynomial in $F[t]$ with a root in $K[t]$ must split. There is another desirable property of an algebraic extension L/K called **separability**. In some sense it is dual to normality, but this is hard to believe at first because there is a large class of fields F for which all algebraic extensions K/F are separable, including all fields of characteristic 0. (For that matter, there are fields for which every algebraic extension is normal, like $\mathbb{R}$ and $\mathbb{F}_p$.) Like normality, separability can also be expressed in terms of polynomials and also in terms of embedding conditions. We begin with a study of polynomials.

5.1. **Separable Polynomials.**

A nonconstant polynomial $P \in F[t]$ is **separable** if over an algebraic closure $\overline{F}$, $P(t)$ splits into *distinct* linear factors. Equivalently, if P has degree n, then there are n distinct elements $\alpha_1, \ldots, \alpha_n \in \overline{F}$ such that $P(\alpha_i) = 0$ for all i. Note that both of these conditions are easily seen to be independent of the chosen algebraic closure.

Exercise: let F be a field and K/F be any extension. Show that a polynomial $P \in F[t]$ is separable as a polynomial over F iff it is separable when viewed as a polynomial over K.

Lemma 5.1. *Let F be a field of characteristic $p > 0$ and $\alpha \in F^\times \setminus F^{\times p}$. Then for all $n \geq 1$, the polynomial $t^{p^n} - \alpha$ is irreducible.*

Proof. We shall prove the contrapositive: suppose that for some $n \in \mathbb{Z}^+$ the polynomial $t^{p^n} - \alpha$ is reducible; we will show that α is a pth power in F. We may write $t^{p^n} - \alpha = f(t)g(t)$, where $f(t)$ and $g(t)$ are nonconstant monic polynomials. Let K/F be an extension field containing a root β of $t^{p^n} - \alpha$, so that in $K[t]$ we have

$$t^{p^n} - \alpha = t^{p^n} - \beta^{p^n} = (t - \beta)^{p^n}.$$

Since $K[t]$ is a UFD and $f(t)$ and $g(t)$ are monic, we therefore have $f(t) = (t - \beta)^r$ for some $0 < r < p^n$. Write $r = p^m s$ with $\gcd(p, s) = 1$. Note that $m < n$. Then

$$f(t) = (t^{p^m} - \beta^{p^m})^s,$$

so that the coefficient of $t^{p^m(s-1)}$ is $-s\beta^{p^m}$. This lies in F and – since $s \neq 0$ in F – we conclude $\beta^{p^m} \in F$. Thus

$$\alpha = (\beta^{p^m})^{p^{n-m}} \in F^{p^{n-m}} \in F^p$$

since $m < n$. □

Over any field F it is no trouble to come up with a polynomial that is not separable: t^2. What is of more interest is whether there is an inseparable irreducible polynomial in $F[t]$. Note that some authors define a polynomial to be separable if all its irreducible factors are separable and others only discuss in/separability for irreducible polynomials. Although these conventions certainly "work" as well, I find the current definition to be more convenient and more thematic. First, Exercise XX shows that with this definition, separability is faithfully preserved by base extension. Since the way one will check whether an irreducible polynomial is separable is by considering it over the algebraic closure, where of course it is a product of separable (linear!) polynomials, our definition seems simpler. Moreover, in the theory of algebras one does meet reducible polynomials: for any nonconstant $P \in F[x]$, we may consider the finite-dimensional F-algebra $A_P = F[x]/(P(x))$. Then our definition makes it true that P is separable iff A_P is a **separable algebra**, i.e., an algebra which is semisimple and remains semisimple after arbitrary base change.

In general it is far from obvious whether the field extension obtained by adjoining a root of an irreducible polynomial is normal. Fortunately, it is much easier to determine whether a polynomial, especially an irreducible polynomial, is separable.

Exercise 5.1. *Let k be a field.*
a) Show: there is a unique k-linear endomorphism $f \mapsto f'$ of $k[t]$ such that for all $n \in \mathbb{N}$ we have $(t^n)' = nt^{n-1}$.
b) Show: for all $f, g \in k[t]$ we have $(fg)' = f'g + fg'$.
c) Show: for all $f, g \in k[t]$ we have $(f(g(t)))' = f'(g(t))g'(t)$.
d) Suppose k has characteristic 0. Show: if $\deg(f) = n \geq 1$, then $\deg(f') = n - 1$. Deduce that $\{f \in k[t] \mid f' = 0\} = k$.
e) Suppose k has characteristic $p > 0$. Show:

$$\{f \in k[t] \mid f' = 0\} = k[t^p].$$

Proposition 5.2. *(Derivative Criterion)*
Let $f \in F[t]$ be a nonconstant polynomial.
a) The polynomial f is separable iff $\gcd(f, f') = 1$.
b) If f is irreducible, it is separable iff $f' \neq 0$.
c) An irreducible polynomial is always separable in characteristic 0. In characteristic $p > 0$, an irreducible polynomial is inseparable iff there exists $g \in F[t]$ such that $f(t) = g(t^p)$.

Proof. a) Let $d \in F[t]$ be a greatest common divisor of f and f'. This means that

$$\{\alpha f + \beta f' \mod \alpha, \beta \in F[t]\} = \{\gamma d \mid \gamma \in F[t]\}.$$

If K/F is any field extension, it follows that

$$\{\alpha f + \beta f' \mod \alpha, \beta \in K[t]\} = \{\gamma d \mid \gamma \in K[t]\},$$

so d is again a greatest common divisor of f and f' in $K[t]$. Moreover, if $\overline{F}$ is an algebraic closure of F and $\overline{K}$ is an algebraic closure of K, then the Magic Mapping Theorem gives an F-algebra homomorphism $\overline{F} \hookrightarrow \overline{K}$, so $f \in F[t]$ is separable iff $f \in K[t]$ is separable. Thus both of the conditions of part a) are stable under replacing F by an extension field, so we may assume that F is algebraically closed and thus f is split. If f is not separable, then for some $\alpha \in F$ we have

$$f = (t - \alpha)^2 g$$

and thus

$$f' = (t-\alpha)^2 g' + 2(t-\alpha)g = (t-\alpha)h$$

so $(t-\alpha) \mid \gcd(f, f')$. Conversely, if f is separable, then for every root α of f we have

$$f = (t-\alpha)g \text{ with } g(\alpha) \neq 0,$$

so

$$f' = (t-\alpha)g' + g,$$

so $f'(\alpha) = g(\alpha) \neq 0$ and thus $(t-\alpha) \nmid f'$. Thus $\gcd(f, f') = 1$.
b) If f is irreducible, then since $\gcd(f, f') \mid f$, if $\gcd(f, f') \neq 1$ then $\gcd(f, f') = f$ so $f \mid f'$. Since $\deg(f') < \deg(f)$, this occurs if and only if $f' = 0$.
c) This follows from part b) and the previous exercise. □

Lemma 5.3. *Let F be a field of characteristic $p > 0$, $a \in \mathbb{Z}^+$, and $\alpha \in F^\times$. TFAE:*
(i) There exists $\beta \in F$ such that $\beta^p = \alpha$.
(ii) The polynomial $P(t) = t^{p^r} - \alpha$ is reducible over F.

Proof. Because the polynomial $t^p - \alpha$ is inseparable, it has a unique root in an algebraic closure $\overline{F}$, namely an element β such that $\beta^p = \alpha$. We must show that the reducibility of $P(t)$ is equivalent to this β lying in F. Moreover, let γ be an element of $\overline{F}$ such that $\gamma^{p^a} = \alpha$. Then $P(t) = (t-\gamma)^{p^a}$, so that the element γ is unique; moreover, since $(\gamma^{p^{a-1}})^p = \alpha$ and α has a unique pth root in $\overline{F}$, we must have $\gamma^{p^{a-1}} = \beta$. That (i) implies (ii) is now easy: if $\beta \in F$, then we may write $P(t) = (t-\beta)^p$ so P is irreducible over F.

Conversely, assume that $P(t)$ is reducible over F, i.e., there exist $0 < i < p^r$ such that $(t-\gamma)^i \in F[t]$. The coefficient of t^{i-1} in this polynomial is $-i\gamma$, so if i is prime to p this implies that $\gamma \in F$, hence $\beta \in F$ which gives (i). So we may therefore assume that $i = p^b \cdot j$ where $1 \leq b \leq a-1$ and $\gcd(p, j) = 1$. Then $(t^{p^b} - \gamma^{p^b})^j \in F[t]$, and arguing as before we get that $\gamma^{p^b} \in F$, and therefore $\beta = (\gamma^{p^b})^{p^{a-b}} \in F$. □

A field F is **perfect** if every irreducible polynomial over F is separable. It follows immediately from Prop XXc) that every field of characteristic 0 is perfect. In other words, the entire discussion of separability is nonvacuous only in positive characteristic, so for the remainder of this section we assume that all fields are of positive characteristic. Unless otherwise specified, p shall always denote a prime number which is the characteristic of the field(s) in question.

If F has characteristic $p > 0$, we consider the Frobenius homomorphism

$$\mathfrak{f} : F \to F,\ x \mapsto x^p.$$

Let $F^p = \mathfrak{f}(F)$ be the image, a subfield of F.

Proposition 5.4. *a) A field of characteristic $p > 0$ is perfect iff the Frobenius homomorphism is surjective: $F^p = F$.*
b) Therefore finite fields and algebraically closed fields are perfect.

Proof. Assume $F^p = F$, and let $P(t) = \sum_i a_i t^{pi}$ be an irreducible inseparable polynomial. We can then write $a_i = b_i^p$ and then

$$P(t) = \sum_i (b_i)^p (t^i)^p = (\sum_i b_i t^i)^p = Q(t)^p,$$

hence $P(t)$ is not irreducible after all. Therefore F is perfect. Inversely, if the Frobenius homomorphism is not surjective, then there exists some $\alpha \in F$ which is not a pth power, and then by Lemma XX the inseparable polynomial $t^p - \alpha$ is irreducible, so F is not perfect. This gives part a). As for part a), like any field homomorphism, the Frobenius map is injective, and an injective map from a finite set to itself is necessarily surjective. If F is algebraically closed, then for any $\alpha \in F$ the polynomial $t^p - \alpha$ has a root in F, i.e., $\alpha \in F^p$. □

For any positive integer a, we may consider the $\mathfrak{p}^a$, the map which takes $\alpha \mapsto \alpha^{p^a}$, which can also be described as the ath power of the Frobenius map. We write $F^{p^a} = \mathfrak{p}^a(F)$. If F is not perfect then we get an infinite descending chain of proper subfields

$$F \supsetneq F^p \supsetneq F^{p^2} \supsetneq \dots$$

Indeed, if $\alpha \in F \setminus F^p$, then $\alpha^{p^{a-1}} \in F^{p^{a-1}} \setminus F^{p^a}$. This gives another proof that an imperfect field is infinite.

Exercise X.X.X: Let F be a field of characteristic p, with an algebraic closure $\overline{F}$. Define $F^{1/p} = \{\beta \in \overline{F} \mid \beta^p \in F\}$.
a) Show that $F^{1/p}$ is a subextension of $\overline{F}/F$.
b) Similarly define a tower of subextensions

$$F \subset F^{1/p} \subset F^{1/p^2} \subset \dots F^{1/p^a} \subset \dots \subset \overline{F},$$

and show that if F is imperfect, all these inclusions are strict.
c) Define $F^{1/p^\infty} = \bigcup_{a=1}^\infty F^{1/p^a}$. Show that F^{1/p^∞} is perfect and is the intersection of all perfect subextensions of $\overline{F}$. It is called the **perfect closure** of F.

Purely inseparable polynomials: Say that a polynomial $P(t) \in F[t]$ is **purely inseparable** if there exists exactly one $\alpha \in \overline{F}$ such that $P(\alpha) = 0$. As above, there are certainly purely inseparable polynomials over F – $(t - \alpha)^n$ for any $\alpha \in F$ and $n \in \mathbb{Z}^+$ – and what is of interest is the purely inseparable irreducible polynomials, which by the discussion thus far clearly can only exist in characteristic $p > 0$.

Proposition 5.5. *Let F be a field of characteristic $p > 0$. The irreducible, purely inseparable monic polynomials $P(t) \in F[t]$ are precisely those of the form $t^{p^a} - \alpha$ for some $a \in \mathbb{Z}^+$ and some $\alpha \in F \setminus F^p$.*

Proof. By Lemma XX, any polynomial of the form $t^{p^a} - \alpha$ for $\alpha \in F \setminus F^p$ is irreducible. Conversely, let $P(t)$ be a purely inseparable polynomial. By XXXXX, there exists a polynomial $P_2(t)$ such that $P(t) = P_2(t^p)$. Since P is irreducible, so is P_2. If there exist distinct $\alpha, \beta \in \overline{F}$ such that $P_2(\alpha) = P_2(\beta)$ then there are unique and distinct elements $\alpha^{\frac{1}{p}}, \beta^{\frac{1}{p}}$ in $\overline{F}$ such that $P(\alpha^{\frac{1}{p}}) = P(\beta^{\frac{1}{p}}) = 0$, contradicting the pure inseparability of α. Therefore P_2 must itself be irreducible purely inseparable, and an evident inductive argument finishes the proof. □

Exercise XX: Show that the polynomial $t^6 - x$ over the field $\mathbb{F}_3[x]$ is irreducible and inseparable but not purely inseparable.

5.2. **Separable Algebraic Field Extensions.**

Let F be a field and $P(t)$ an irreducible, inseparable polynomial over F of degree $d > 1$. Consider the finite field extension $K = F[t]/(P(t))$ of F. It exhibits

some strange behavior. First, the only F-algebra embedding $\sigma : K \to \overline{K}$ is the inclusion map. Indeed, such embeddings correspond bijectively to the assignments of $t \in K$ to a root α of P in $\overline{K}$, and by assumption there are less than d such elements. It follows that the group $\mathrm{Aut}(K/F)$ of F-algebra automorphisms of F has cardinality smaller than d.

For an extension K/F, the **separable degree** $[K : F]_s$ is the cardinality of the set of F-algebra embeddings $\sigma : K \to \overline{F}$.

Exercise: Show that the separable degree may be computed with respect to embeddings into any algebraically closed field containing F.

Theorem 5.6. *The separable degree is multiplicative in towers: if $L/K/F$ is a tower of finite field extensions, then $[L : F]_s = [L : K]_s[K : F]_s$.*

Proof. Let $\sigma : F \hookrightarrow C$ be an embedding of F into an algebraically closed field. Let $\{\sigma_i\}_{i \in I}$ be the family of extensions of σ to K, and for each $i \in I$ let $\{\tau_{ij}\}_{j \in J_i}$ be the family of extensions of σ_i to L. Each σ_i admits precisely $[L : K]_s$ extensions to embeddings of L into C: in particular, the cardinality of J_i is independent of i and there are thus precisely $[L : K]_s[K : F]_s$ F-algebra embeddings τ_{ij} overall. These give all the F-algebra embeddings $L \hookrightarrow C$, so $[L : F]_s = [L : K]_s[K : F]_s$. □

Corollary 5.7. *Let K/F be a finite degree field extension. Then*

$$[K : F]_s \leq [K : F].$$

In particular, the separable degree is finite.

Proof. We employ dévissage: break up K/F into a finite tower of simple extensions. Each simple extension has finite degree and by Theorem 5.6 the degree is multiplicative in towers. We are therefore reduced to the case $K = F(\alpha) \cong F[t]/(P(t))$, where $P(t)$ is the minimal polynomial for α. In this case the result is clear, since an F-algebra homomorphism of $F[t]/(P(t))$ into any field M is given by sending the image of t to a root of $P(t)$ in M, and the degree $[K : F]$ polynomial has at most $[K : F]$ roots in any field. □

In the situation of the proof of Corollary 5.7 we can say more: the separable degree $[F(\alpha) : F]_s$ is equal to the number of distinct roots of the minimal polynomial $P(t)$ of α. In particular it is equal to the degreee of the field extension iff $P(t)$ is a separable polynomial. Let us record this result.

Proposition 5.8. *For K/F a field extension and $\alpha \in K$ algebraic over F, TFAE:*
(i) The minimal polynomial of α is a separable polynomial.
(ii) $[F(\alpha) : F]_s = [F(\alpha) : F]$.

More generally:

Theorem 5.9. *For a finite degree field extension K/F, TFAE:*
(i) Every element of K is separable over F.
(ii) $[K : F]_s = [K : F]$.
*A field extension satisfying these equivalent conditions is said to be **separable**.*

Proof. (i) $\implies$ (ii): We may write K/F as a finite tower of simple extensions:

$$F = F_0 \subset \ldots \subset F_n = K$$

such that for all i we have $F_{i+1} = F_i(\alpha_{i+1})$. Since α_{i+1} is separable over F, it is separable over F_i: indeed, the minimal polynomial for α_{i+1} over the extension field divides the minimal polynomial over the ground field. Therefore Proposition 5.8 applies and $[F_{i+1} : F_i]_s = [F_{i+1} : F_i]$ for all i. Since both the separable degree and the degree are multiplicative in towers, we conclude $[K : F]_s = [K : F]$.
(ii) $\implies$ (i): Seeking a contradiction, we suppose that there exists $\alpha \in K$ which is not separable over F. By Proposition 5.8, it follows that $[F(\alpha) : F]_s < [F(\alpha) : F]$. Now applying Theorem 5.6 and Corollary 5.7 we get

$$[K : F]_s = [K : F(\alpha)]_s[F(\alpha) : F]_s < [K : F(\alpha)][F(\alpha) : F] = [K : F].$$

□

Corollary 5.10. *Finite degree separable extensions are a distinguished class of field extensions: that is, they satisfy (DC1) and (DC2) of* §3.4 *and thus also (DC3).*

Exercise: Prove Corollary 5.10.

Theorem 5.11. *Let L/F be an algebraic field extension. TFAE:*
(i) Every finite subextension of L/F is separable.
(ii) Every irreducible polynomial $P \in F[t]$ which has a root in L is separable.
(iii) L is obtained by adjoining to F a set of roots of separable polynomials.
An extension satisfying these equivalent properties is called a ***separable algebraic extension****.*

Exercise: Prove Theorem 5.11.

Corollary 5.12. *Algebraic separable extensions are a distinguished class of field extensions.*

Exercise: Prove Corollary 5.12.

Corollary 5.13. *For a family $\{K_i/F\}_{i \in I}$ of algebraic field extensions inside a common algebraically closed field M, TFAE: (i) For all $i \in I$, K_i/F is a separable algebraic field extension.*
(ii) The compositum $\prod_i K_i$ is a separable algebraic field extension.

Exercise: Prove Corollary 5.13.

Corollary 5.13 has the following important consequence: for any field extension K/F, there exists a unique maximal separable algebraic subextension $\mathrm{SepCl}_K(F)$, the **separable closure of F in K**.

5.3. **Purely Inseparable Extensions.**

Theorem 5.14. *For an algebraic field extension K/F, TFAE:*
(i) There is only one F-algebra embedding $K \hookrightarrow \overline{K}$.
(ii) Every irreducible polynomial $P \in F[t]$ with a root in K is purely inseparable.
(iii) K is obtained by adjoining to F roots of purely inseparable polynomials.
(iv) The separable closure of K in F is F.

Exercise: Prove Theorem 5.14. An extension satisfying the conditions of Theorem 5.14 is **purely inseparable**.

Exercise:

a) Show that finite degree purely inseparable extensions form a distinguished class.
b) Show that the purely inseparable algebraic extensions form a distinguished class which is closed under composita.

In light of Exercise X.Xb), for any algebraic field extension K/F we may define the **purely inseparable closure** of F in K to be the largest subextension of K which is purely inseparable over F.

Exercise X.X.X: Show that the purely inseparable closure of F in an algebraic closure $\overline{F}$ is the perfect closure F^{1/p^∞}.

Corollary 5.15. *Let K/F be a purely inseparable extension of finite degree. Then $[K:F]$ is a power of p.*

Proof. One may reduce to the case of a simple extension $K = F[\alpha]$, and then α is purely inseparable over F so has minimal polynomial of the form $t^{p^a} - \alpha$ for some $a \in \mathbb{Z}^+$. □

Corollary 5.16. *A purely inseparable extension is normal.*

Proof. This follows immediately from condition (i) of Theorem 5.14. □

The flavor of these results is that many formal properties are common to both separable and purely inseparable extensions. The exceptions to this rule are the following: first, purely inseparable extensions are always normal, whereas this is most certainly not the case for separable extensions. A more subtle difference is expressed in Theorem XX: if K/F is **not** purely inseparable, then it must have a nontrivial separable subextension. However, if K/F is **not** separable, that does not mean that it has a nontrivial purely inseparable subextension.

Example [Mo96, p. 48]: Let k be a field of characteristic 2, $F = k(x,y)$ (rational function field), u a root in $\overline{F}$ of the separable irreducible quadratic polynomial t^2+t+x, $S = F(u)$ and $K = S(\sqrt{uy})$. Clearly K/S is purely inseparable and S/F is separable. But there is no nontrivial purely inseparable subextension of K/F. Equivalently, we will show that if $a \in K$, $a^2 \in F$, then already $a \in F$. An F-basis for K is $1,\ u,\ \sqrt{uy},\ u\sqrt{uy}$. If $a^2 \in F$, write

$$a = \alpha + \beta u + \gamma\sqrt{uy} + \delta u\sqrt{uy},\ \alpha,\beta,\gamma,\delta \in F.$$

Since $a^2 \in F$, the coefficient of $u = 0$, i.e.,

$$\beta^2 + (\gamma+\delta)^2 y + \delta^2 xy = 0.$$

If $\delta = 0$ then $\beta^2 + \gamma^2 y = 0$, so $\gamma = 0$ since y is not a square in F. But then $\beta = 0$ and $a \in F$. If $\delta \neq 0$, then

$$x = \frac{\beta^2 + (\gamma+\delta)^2 y}{\delta^2 y} = (\frac{\gamma}{\delta} + 1)^2 + (\frac{\beta}{\delta})^2 y,$$

so that $x \in F^2(y)$, which is not the case. So $\delta = 0$ and $a \in F$.

5.4. **Structural Results on Algebraic Extensions.**

Proposition 5.17. *Suppose an algebraic extension K/F is both separable and purely inseparable. Then $K = F$.*

Proof. For such an extension, let $\alpha \in K$. Then the minimal polynomial of α over F is both separable and purely inseparable. The only such polynomials have degree one, i.e., $\alpha \in F$. □

Proposition 5.18. *For any algebraic field extension K/F, the extension $K/\operatorname{SepCl}_K(F)$ is purely inseparable.*

Proof. Since $\operatorname{SepCl}_K(F)$ is the maximal separable subextension of K/F, there cannot be a proper nontrivial separable extension of $K/\operatorname{SepCl}_K(F)$, so it is purely inseparable. □

In general this result is not valid the other way around: an algebraic field extension K/F *need not* be separable over its purely inseparable closure. Indeed, in the example of the previous section the purely inseparable closure F_i was F and K/F was not separable. The following two results give more information on when K is separable over F_i.

Theorem 5.19. *For an algebraic extension K/F, let F_s and F_i be, respectively, the separable and purely inseparable closures of F in K. TFAE:*
(i) $K = F_sF_i$.
(ii) K is separable over F_i.

Proof. (i) $\implies$ (ii): K is obtained by adjoining to F_i roots of separable polynomials with coefficients in F, hence by polynomials with coefficients in F_s.
(ii) $\implies$ (i): If K/F_i is separable, then K/F_iF_s is separable. Similarly, since K/F_s is inseparable, K/F_iF_s is inseparable. By Proposition 5.17, $K = F_iF_s$. □

Corollary 5.20. *The equivalent conditions of Theorem 5.19 hold when K/F is normal. In particular they hold for $\overline{F}/F$, giving $\overline{F} = F^{\mathrm{sep}}F^{1/p^\infty}$.*

Proof. Let $\alpha \in K \setminus F^i$. Then α is not purely inseparable over F, i.e., the minimal polynomial P of α has at least one other distinct root, say β, in an algebraic closure. But since K/F is normal, $\beta \in F$. By the Extension Theorem, there exists an F-algebra automorphism s of L such that $s(\alpha) = \beta$. This shows that the set of elements in F which are fixed by every automorphism of L/F is precisely F^i. Let Q be the minimal polynomial of α over F^i, and let $\alpha_1, \ldots, \alpha_r$ be the distinct roots of Q in $\overline{F}$. Since the group G of automorphisms of K/F^i acts on the α_i's by permutations, the separable polynomial $R(t) = \prod_{i=1}^r (t - \alpha_i)$ is invariant under G, i.e., it lies in $F^i[x]$. This shows that K/F_i is obtained by adjoining roots of separable polynomials and is therefore separable. The second sentence of the Corollary follows immediately from the first. □

Corollary 5.21. *For a finite extension K/F, $[K : F]_s = [\operatorname{SepCl}_K(F) : F]$. In particular $[K : F]_s \mid [K : F]$.*

Proof. We have $[K : F]_s = [K : \operatorname{SepCl}_K(F) : F]_s[\operatorname{SepCl}_K(F) : F]_s$. But the separable degree of a purely inseparable extension is 1, so the conclusion follows. □

For a finite extension K/F one may therefore define the **inseparable degree** $[K : F]_i$ of a finite extension to be $[K : F]/[K : F]_s = [K : \operatorname{SepCl}_K(F)]$.

A field is **separably closed** if it admits no proper separable algebraic field extension.

Proposition 5.22. *The separable closure of a field in any algebraically closed field is separably closed.*

Exercise: Prove Proposition 5.22.

One often writes F^{sep} for a separable closure of F. Like the algebraic and normal closures, this extension is unique up to non-canonical F-algebra isomorphism.

Corollary 5.23. *Let K/F be a normal algebraic extension. Then the separable closure F^s of F in K is also normal.*

Proof. For any embedding σ of K into $\overline{F}$, the image $\sigma(F^s)$ lies in K (by normality of K) and is evidently also a separable subextension of K/F. Therefore we must have $\sigma(F^s) = F^s$. □

Corollary 5.24. *A field F is perfect iff its separable closure is algebraically closed.*

Proof. If F is perfect then all algebraic extensions are separable, so the result is clear. Inversely, suppose that F is not perfect, so there exists $\alpha \in F \setminus F^p$ and a corresponding purely inseparable field extension $F[\alpha^{1/p}]/F$ defined by the irreducible inseparable polynomial $P = t^p - \alpha$. By Theorem 5.11, only a separable irreducible polynomial can acquire a root in a separable field extension, so the polynomial P remains irreducible over the separable closure of F. □

6. Norms, traces and discriminants

6.1. Dedekind's Lemma on Linear Independence of Characters.

Theorem 6.1. *(Dedekind's Lemma) Let M be a monoid and K a field. The set $X(M, K)$ of all monoid homomorphisms $M \to K^\times$ is linearly independent as a subset of the K-vector space K^M of all functions from M to K.*

Proof. By definition, a subset of a vector space is linearly independent iff every nonempty finite subset is linearly independent. So it's enough to show that for all $N \in \mathbb{Z}^+$, every N-element subset of $X(M, K)$ is linearly independent in K^M. We show this by induction on N. The base case, $N = 1$, is immediate: the only one element linearly dependent subset of K^M is the zero function, and elements of $X(M, K)$ are nonzero at all values of M. So suppose $N \geq 2$, that every $N - 1$ element subset of $X(M, K)$ is linearly independent, and let $\chi_1, \ldots, \chi_N$ be distinct elements of $X(M, K)$. Let $\alpha_1, \ldots, \alpha_N \in K$ be such that for all $x \in M$, we have

$$\alpha_1\chi_1(x) + \ldots + \alpha_N\chi_N(x) = 0. \tag{1}$$

Our goal is to show that $\alpha_1 = \ldots = \alpha_N = 0$. Since $\chi_1 \neq \chi_N$, there is $m \in M$ such that $\chi_1(m) \neq \chi_N(m)$. Substituting mx for x in (1), we get that for all $x \in M$,

$$\alpha_1\chi_1(m)\chi_1(x) + \alpha_2\chi_2(m)\chi_2(x) + \ldots + \alpha_N\chi_N(m)\chi_N(x) = 0. \tag{2}$$

Multiplying (2) by $\chi_1(m)^{-1}$ and subtracting this from (1), we get

$$\forall x \in M,\ \alpha_2\left(\frac{\chi_2(m)}{\chi_1(m)} - 1\right)\chi_2(x) + \ldots + \alpha_N\left(\frac{\chi_N(m)}{\chi_1(m)} - 1\right)\chi_N(x) = 0. \tag{3}$$

By induction, $\chi_2, \ldots, \chi_N$ are linearly independent, so $\alpha_N\left(\frac{\chi_N(m)}{\chi_1(m)} - 1\right) = 0$ and thus $\alpha_N = 0$. Thus (1) gives a linear dependence relation among the $N - 1$ characters $\chi_1, \ldots, \chi_{N-1}$, so by induction $\alpha_1 = \cdots = \alpha_{N-1} = 0$. □

6.2. **The Characteristic Polynomial, the Trace and the Norm.**

Let L/K be a field extension of degree $n < \infty$. For $x \in L$, the map $x\bullet : L \to L$ given by $y \in L \mapsto xy$ is an endomorphism of L as a K-vector space. That is, for all $\alpha \in K$ and $y_1, y_2 \in$L, we have $x(\alpha y_1 + y_2) = x(\alpha y_1 + y_2) = \alpha x y_1 + x y_2 = \alpha(x y_1) + (x y_2)$. We may therefore analyze the element $x \in L$ using tools of linear algebra.

Choose a K-basis $b_1, \ldots, b_n$ for L. With respect to such a basis, the linear transformation $x\bullet$ is represented by an $n \times n$ matrix, say $M(x)$.

Example: Take $K = \mathbb{R}$, $L = \mathbb{C}$, and the basis $(1, i)$. Let $x = a + bi$. Then $x \bullet 1 = a \cdot 1 + b \cdot i$ and $x \bullet i = -b \cdot 1 + a \cdot i$. Therefore

$$M(x) = \left[\begin{array}{cc} a & -b \\ b & a \end{array} \right].$$

Example: if x lies in K, then $M(x) = m_{i,j}$ is simply the scalar matrix $\mathrm{diag}(x, \ldots, x)$. Note that Proposition 6.2 below gives a generalization of this simple observation.

We define the characteristic polynomial of x:

$$P_x(t) = \det(tI_n - M(x)) = \prod_{i=1}^{n} (t - \lambda_i).$$

Similarly we define the **trace**

$$\mathrm{Tr}_{L/K}(x) = \mathrm{tr}(M(x)) = \sum_{i=1}^{n} m_{i,i} = \sum_{i=1}^{n} \lambda_i$$

and the **norm**

$$N_{L/K}(x) = \det(M(x)) = \prod_{i=1}^{n} \lambda_i.$$

Proposition 6.2. *Let $L/K/F$ be a tower of field extensions with $m = [K : F]$ and $n = [L : K]$. Let $x_1, \ldots, x_m$ be a basis for K/F and $y_1, \ldots, y_n$ a basis for L/K.*
a) For any element $\alpha \in K$, if M is the matrix representing $x\bullet \in \mathrm{End}_F(K)$ with respect to $\{x_1, \ldots, x_m\}$, the matrix representation of $x\bullet \in \mathrm{End}_F(L)$ with respect to the basis $\{x_i y_j\}_{1 \le i \le m, 1 \le j \le n}$, reverse lexicographically ordered, is the block diagonal matrix $\mathrm{diag}(M, \ldots, M)$, i.e., n blocks, each equal to M. It follows that:
b) Let $f(t)$ be the characteristic polynomial of $\alpha\bullet \in \mathrm{End}_F(K)$ and $g(t)$ be the characteristic polynomial of $x\bullet \in \mathrm{End}_F(L)$. Then $g(t) = f(t)^{[L:K]}$.
c) $N_{L/F}(x) = N_{K/F}(x)^{[L:K]}$.
d) $\mathrm{Tr}_{L/F}(x) = [L : K] \mathrm{Tr}_{K/F}(x)$.

Proof. We have $\alpha x_i = \sum_{k=1}^{m} m_{ki} x_k$ and hence $\alpha x_i y_j = \sum_{k=1}^{m} m_{ki}(x_k y_j)$. This establishes part a). The remaining parts follow easily by standard linear algebraic considerations. □

Corollary 6.3. *Let L/F be a finite degree field extension. Let α be an element of L, let $f(t)$ be the minimal polynomial of α over F, and let $g(t)$ be the characteristic polynomial of $\alpha\bullet \in \mathrm{End}_F(L)$. Then $g(t) = f(t)^{[L:F(\alpha)]}$.*

Proof. Put $K = F(\alpha)$. The minimal polynomial f of α over F is the characteristic polynomial of $x\bullet \in \mathrm{End}_F(K)$. So the result follows from Proposition 6.2. □

Proposition 6.4. *Let $L/K/F$ be a tower of finite degree field extensions. Then:*
a) $\mathrm{Tr}_{K/F} : K \to F$ is an F-linear map.
b) For all $x, y \in K$, $N_{K/F}(xy) = N_{K/F}(x)N_{K/F}(y)$.
c) For all $c \in F$ and $x \in K$, $N_{K/F}(cx) = c^{[K:F]}N_{K/F}(x)$.

Proof. Parts a) and b) are standard properties of the trace and determinant of any F-linear map. Part c) follows by applying part b) and observing that for $c \in F$, $N_{K/F}(c)$ is the determinant of the scalar matrix $\mathrm{diag}(c, \ldots, c)$, i.e., $c^{[K:F]}$. □

The following key result identifies the eigenvalues of $\alpha\bullet$ in field-theoretic terms.

Theorem 6.5. *Let K/F be a field extension of degree $n < \infty$ and separable degree n_s. Put $p^e = \frac{n}{n_s} = [K : F]_i$. Let $\overline{K}$ be an algebraic closure of K. Let $\alpha \in K$ and let $f(t)$ be the characteristic polynomial of $\alpha\bullet \in \mathrm{End}_F(K)$. Let $\tau_1, \ldots, \tau_{n_s}$ be the distinct F-algebra embeddings of K into $\overline{K}$. Then*

$$f(t) = \prod_{i=1}^{n_s} (t - \tau_i(\alpha))^{p^e}.$$

It follows that

$$N_{K/F}(\alpha) = (\prod_{i=1}^{n_s} \tau_i(\alpha))^{p^e} \tag{4}$$

and

$$\mathrm{Tr}_{K/F}(\alpha) = p^e \sum_{i=1}^{m} \tau_i(\alpha). \tag{5}$$

Proof. Put $L = F[\alpha]$. Let $d = [L : F]$, $d_s = [L : F]$ and $d_i = [L : F]_i$. Let $\sigma_1, \ldots, \sigma_{d_s}$ be the distinct F-algebra homomorphisms from L into $\overline{F}$. For each $1 \le i \le d_s$, σ_i extends to $\frac{n_s}{d_s}$ F-algebra homomorphisms from K into $\overline{F}$. Let

$$f(t) = (\prod_{i=1}^{d_s} (t - \sigma_i(\alpha)))^{d_i}$$

be the minimal polynomial of α over F, and let $g(t)$ be the characteristic polynomial of $\alpha\bullet$ on K, so by Corollary 6.3 we have

$$g(t) = f(t)^{[K:L]} = (\prod_{i=1}^{d_s} (t - \sigma_i(\alpha))^{d_i \frac{n}{d}} = \left((\prod_{i=1}^{d_s} (t - \sigma_i(\alpha))^{\frac{n_s}{d_s}} \right)^{n_i}$$

$$= \left(\prod_{i=1}^{n_s} (t - \tau_i(\alpha)) \right)^{p^i}.$$

Equations (4) and (5) follow immediately. □

Corollary 6.6. *Let $\mathbb{F}_{q^d}/\mathbb{F}_q$ be an extension of finite fields. Then the norm map $N : \mathbb{F}_{q^d}^\times \to \mathbb{F}_q^\times$ is surjective.*

Proof. Let $\sigma : x \mapsto x^q$, so that $\mathrm{Aut}(\mathbb{F}_{q^d}/\mathbb{F}_q) = \langle 1, \sigma, \ldots, \sigma^{d-1} \}$. Thus for $x \in \mathbb{F}_{q^d}$,

$$N(x) = \prod_{i=0}^{d-1} \sigma^i(x) = \prod_{i=0}^{d-1} x^{q^i} = x^{\sum_{i=0}^{d-1} q^i} = x^{\frac{q^d-1}{q-1}}.$$

Therefore $\operatorname{Ker} N$ consists of all elements of the finite cyclic group $\mathbb{F}_{q^d}^\times$ of order dividing $\frac{q^d-1}{q-1}$, so $\# \operatorname{Ker} N = \frac{q^d-1}{q-1}$. Since $\mathbb{F}_{q^d}^\times / \operatorname{Ker} N \cong N(\mathbb{F}_{q^d}^\times)$, we deduce that $\# N(\mathbb{F}_{q^d}^\times) = q-1$: N is surjective. □

6.3. **The Trace Form and the Discriminant.**

Let F be a field and V a finite-dimensional F-vector space equipped with a bilinear form, i.e., a function $\langle , \rangle : V \times V \to F$ such that for all $v, 1, v_2 \in V$ and $\alpha \in F$,

$$\langle \alpha v_1 + v_2, v_3 \rangle = \alpha \langle v_1, v_3 \rangle + \langle v_2, v_3 \rangle$$

and

$$\langle v_1, \alpha v_2 + v_3 \rangle = \alpha \langle v_1, v_2 \rangle + \langle v_1, v_3 \rangle.$$

Let $V^\vee = \mathrm{Hom}(V, K)$ be the dual space of V. A bilinear form on V induces a linear map $\Phi : V \to V^\vee$, namely

$$\Phi(v) = \langle v, \ \rangle.$$

(Note that a more careful notation would be something like $\Phi_L : v \mapsto \langle v, \ \rangle$, to distinguish it from the *other* obvious map $\Phi_R : v \mapsto \langle \ , v \rangle$. We have $\Phi_L = \Phi_R$ iff the bilinear form is **symmetric**, an assumption which we have not (yet) made. But in the general case the two maps are equally good, so let us work with $\Phi = \Phi_L$ for simplicity.) We say that the bilinear form $\langle , \rangle$ is **nondegenerate** if $\Phi : V \to V^\vee$ is an isomorphism. Since Φ is a linear map between two finite-dimensional vector spaces of the same dimension, Φ is an isomorphism iff it is injective, i.e., for each $v \in V$, if $\langle v, w \rangle = 0$ for all $w \in V$, then $v = 0$.

Let $\langle , \rangle$ be a bilinear form on V, and fix a K-basis $e_1, \ldots, e_n$ of V. We define the **Gram matrix** M of the bilinear form as $M(i,j) = \langle e_i, e_j \rangle$. Then for all $v, w \in V$, we have

$$\langle v, w \rangle = v^T M w.$$

We claim that the nondegeneracy of the form is equivalent to the nonsingularity of the Gram matrix M. If M is singular, so is M^T, so there exists $0 \neq v$ such that $v^T M = (Mv)^T = 0$, and thus $\langle v, w \rangle = 0$ for all $w \in V$. Conversely, if M is nonsingular, then for all $0 \neq v \in V$, $v^T M$ is nonzero, so it has at least one nonzero component i, so $v^T M e_i = \langle v, e_i \rangle \neq 0$. (Note that this argument also makes clear that Φ_L is an isomorphism iff Φ_R is an isomorphism.

Moreover, our fixed basis $(e_1, \ldots, e_n)$ induces a **dual basis** $(e_1^\vee, \ldots, e_n^\vee)$, characterized by $e_i^\vee(e_j) = \delta_{i,j}$ (Kronecker delta) for all $1 \leq i, j \leq n$. Thus, given a nondegenerate bilinear form $\langle , \rangle$ on V, we may pull back the dual basis $(e_1^\vee, \ldots, e_n^\vee)$ under Φ^{-1} to get a basis $(e^1, \ldots, e^n)$ of V with the characteristic property $\langle e_i, e^j \rangle = \delta_{i,j}$. Conversely, if a basis $(e^1, \ldots, e^n)$ of V exists which is dual to the given basis $(e_1, \ldots, e_n)$ in the above sense, then the bilinear form is easily seen to be nondegenerate. In summary:

Proposition 6.7. *Let V be an n-dimensional vector space over a field K, let $\langle , \rangle$ be a bilinear form on V, and let $(e_1, \ldots, e_n)$ be any K-basis of V. Then the following are equivalent:*
(i) The induced map $\Phi = \Phi_L : V \to V^\vee$ given by $v \mapsto \langle v, \ \rangle$ is an isomorphism.
(ii) The induced map $\Phi_R : V \to V^\vee$ given by $v \mapsto \langle \ , v\rangle$ is an isomorphism.
(iii) The Gram matrix $M(i,j) = \langle e_i, e_j \rangle$ is nonsingular.
(iv) There exists a basis $(e^1, \ldots, e^n)$ of V such that $\langle e_i, e^j \rangle = \delta_{i,j}$.

And now, back to field theory: let K/F be a finite-dimensional field extension. Define the **trace form** $T : K \times K \to F$, $T(x,y) := \mathrm{Tr}(x \bullet y\bullet)$. The bilinearity of T follows immediately from the linearity of the trace map. Note that T is also **symmetric** in the sense that $T(x,y) = T(y,x)$ for all $x, y \in K$. A natural question is when the trace form is nondegenerate.

Theorem 6.8. *Let K/F be a field extension of finite degree n. TFAE:*
(i) The trace form $T : K \times K \to F$ is nondegenerate.
(ii) There exists some $x \in K$ such that $\mathrm{Tr}(x) \neq 0$.
(iii) The trace function $\mathrm{Tr} : K \to F$ is surjective.
(iv) The extension K/F is separable.

Proof. The implications (i) $\implies$ (ii) $\implies$ (iii) may safely be left to the reader.
(iii) $\implies$ (iv): we argue by contraposition. If K/F is not separable, then $\mathrm{char}(F) = p > 0$, $[K:F]_i = p^e$ is divisible by p, and thus (5) shows that the trace function is identically zero.
(iv) $\implies$ (i): By the Primitive Element Corollary, we have $K = F[\alpha]$ for some $\alpha \in K$. Then $(1, \alpha, \ldots, \alpha^{n-1})$ is an F-basis of K. Let $x \in K$. By Proposition 6.7, it is enough to show that the Gram matrix $M(i,j) = \mathrm{Tr}(\alpha^{i-1}\alpha^{j-1}) = \mathrm{Tr}(\alpha^{i+j-2})$ is nonsingular. To see this, let $\alpha_1, \ldots, \alpha_n$ be the distinct F-conjugates of α in $\overline{K}$. Then $\mathrm{Tr}(\alpha) = \sum_{i=1}^n \alpha_i$, so that for any $N \in \mathbb{N}$, $\mathrm{Tr}(\alpha^N) = \sum_{i=1}^n \alpha_i^N$. Now we introduce the Vandermonde matrix $V = V(\alpha_1, \ldots, \alpha_n)$: $V(i,j) = \alpha_j^{i-1}$. Why? Well, we compute that the (i,j) entry of VV^T is $\sum_{k=1}^n \alpha_k^{i-1}\alpha_k^{j-1} = M(i,j)$. Therefore

$$\det M = \det VV^T = (\det V)^2 = (\prod_{i>j} \alpha_i - \alpha_j)^2 \neq 0.$$

□

Example (Trace form of a quadratic extension): Let F be a field of characteristic different from 2, and let $K = F(\sqrt{D})$ be a quadratic field extension. We wish to explicitly compute the trace form. A natural choice of F-basis for K is $(1, \sqrt{D})$. The Gram matrix is then

$$M = \begin{bmatrix} \mathrm{Tr}(1) & \mathrm{Tr}(\sqrt{D}) \\ \mathrm{Tr}(\sqrt{D}) & \mathrm{Tr}(D) \end{bmatrix} = \begin{bmatrix} 2 & 0 \\ 0 & 2D \end{bmatrix}.$$

Thus the corresponding quadratic form is $(2, 2D)$, of discriminant $D \in K^\times/K^{\times 2}$.

7. The primitive element theorem

Theorem 7.1. *Let K/F be a finite field extension. TFAE:*
(i) The set of subextensions L of K/F is finite.
(ii) K/F is simple: there exists $\alpha \in K$ such that $K = F[\alpha]$.

Proof. [LaFT, pp. 243-244]: Suppose first that $K = \mathbb{F}_q$ is finite. Then (i) is clear, and (ii) holds because $K^\times$ is cyclic of order $q-1$: if α is a generator of the multiplicative group $K^\times$, then $K = F[\alpha]$. Henceforth we suppose that K is infinite. (i) $\implies$ (ii): observe that for any subextension E of K/F, since (i) holds for K/F, it also holds for E/F. Writing $K = F[\alpha_1, \ldots, \alpha_n]$, we see that it is enough to prove the result in the case of extensions which are generated by two elements: a simple dévissage/induction argument then recovers the general case.

So suppose that $K = F[\alpha, \beta]$. As c ranges over the infinitely many elements of F, there are only finitely many distinct subfields of K of the form $F[\alpha + cb]$, so there exist distinct elements c_1, c_2 of F such that

$$E = F[\alpha + c_1\beta] = F[\alpha + c_2\beta].$$

It then follows, successively, that $(c_1 - c_2)\beta \in E$, $\beta \in E$, $\alpha \in E$, so

$$F[\alpha + c_1\beta] = E = F[\alpha, \beta] = K.$$

(ii) $\implies$ (i): Suppose $K = F[\alpha]$, and let $f(t) \in F[t]$ be the minimal polynomial for α over F. For each subextension E of K/F, let $g_E(t) \in E[t]$ be the minimal polynomial for α over E. Let E' be the subextension of K/F generated by the coefficients of g_E. So $F \subset E' \subset E \subset K$; since g_E is irreducible over E, it is also irreducible over E', and thus $[K : E'] = [E'[\alpha] : E'] = [E[\alpha] : E] = [K : E]$. It follows that $E = E'$. In other words, E can be recovered from g_E and thus the map $E \mapsto g_E$ is bijective. However, we also have that g_E divides f for all E, so g_E is a monic polynomial whose multiset of roots in any algebraic closure is a subset of the multiset of roots of f. So there are only finitely many possibilities for E. □

Corollary 7.2. *("Primitive Element Corollary") The equivalent conditions of Theorem 7.1 hold when K/F is finite and separable. In particular, every such extension is of the form $K = F[\alpha]$.*

Proof. Once again we may assume that F is infinite, and once again by dévissage/induction, it is enough to treat the case of a degree n separable extension of the form $K = F[\alpha, \beta]$. We may, and shall assume, that neither of α and β lie in F. Let $\sigma_1, \ldots, \sigma_n$ be the distinct F-algebra embeddings of K into an algebraic closure $\overline{F}$. Put

$$P(t) = \prod_{i \neq j} \left(\sigma_i\alpha + t\sigma_i\beta - \sigma_j\alpha - t\sigma_j\beta\right).$$

Then, e.g. by Theorem 6.5, $P(t)$ is a nonzero polynomial. Since F is infinite, there exists $c \in F$ such that $P(c) \neq 0$. Then for $1 \leq i \leq n$, the elements $\sigma_i(\alpha + c\beta)$ are distinct, so that $[F[\alpha + c\beta] : F] \geq n = [F[\alpha, \beta] = [K : F]$. Thus $K = F[\alpha + c\beta]$. □

Remark: What we are calling the Primitive Element Corollary is often itself referred to as the Primitive Element Theorem.

Corollary 7.3. *(Lang) Let K/F be a separable algebraic extension such that: there is $n \in \mathbb{Z}^+$ such that for all $\alpha \in K$, $[F(\alpha) : F] \leq n$. Then $[K : F] \leq n$.*

Proof. Let $\alpha \in K$ be such that $[F(\alpha) : F]$ has maximal degree – it is no loss of generality to assume that this degree is n. We claim that $K = F(\alpha)$, which will establish the result.

Suppose that $K \supsetneq F(\alpha)$, and let $\beta \in K \setminus F(\alpha)$. Since $F(\alpha, \beta)/F$ is finite separable, by the Primitive Element Corollary (Corollary 7.2) there exists $\gamma \in K$

such that $F(\alpha, \beta) = F(\gamma)$. But then we must have $[F(\gamma) : F] > [F(\alpha) : F]$, contradiction. □

Exercise: Give an example to show that the conclusion of Corollary 7.3 does not hold without the separability hypothesis.

Remark: A more natural proof of Corollary 7.2 would be obtained by taking the normal closure M of K/F and using the *Galois correspondence*: the lattice of subextensions of M/F is anti-isomorphic to the lattice of subgroups of $\mathrm{Aut}(M/F)$, hence there are certainly only finitely many of the former, which of course implies that there are only finitely many subextensions of K/F. This brings us to our next topic, Galois Theory.

8. Galois Extensions

8.1. **Introduction.**

For any field extension K/F we define $\mathrm{Aut}(K/F)$ to be the group of F-algebra automorphisms of K, i.e., the set of all field isomorphisms $\sigma : K \to K$ such that $\sigma(x) = x$ for all $x \in F$. This is a group under composition.

Let G be a subgroup of $\mathrm{Aut}(K/F)$, i.e., a group of F-algebra automorphisms of K. We define the **fixed field**

$$K^G = \{x \in K \mid \sigma(x) = x \ \forall \sigma \in G\}.$$

Note that the notation comes from representation theory: if R is a commutative ring, M an R-module and G is a group, then one has the notion of an R-linear representation of G on M, i.e., a homomorphism from G to the group of R-module automorphisms of M. In such a situation one can "take invariants", i.e., consider the subset of M on which G acts trivially: this is denoted M^G. The present definition is an instance of this with $R = F$, $M = K$.

It is immediate to check that K^G is a subextension of K/F. (In fact in the more general setting detailed above, one checks that M^G is an R-submodule of M.)

A field extension K/F is **weakly Galois** if $K^{\mathrm{Aut}(K/F)} = F$. Equivalently, for any element $x \in K \setminus F$, there exists $\sigma \in \mathrm{Aut}(K/F)$ such that $\sigma(x) \neq x$.

A field extension K/F is **Galois** if for all subextensions L of K/F, $K^{\mathrm{Aut}(K/L)} = L$.

Remark: The terminology "weakly Galois" is not standard. In fact, it is usual to consider Galois theory only for *algebraic* extensions and in this case it will turn out to be the case that the notions of weakly Galois and Galois coincide.

This "top down" definition of a weakly Galois extension is the generalization to arbitrary extensions of a definition of E. Artin for finite extensions. It has the merit of making it easy to exhibit a large class of weakly Galois extensions: if K is any field and G is any group of automorphisms of K, then K/K^G is, tautologically, a Galois extension.

Example: Let G be the 2-element subgroup of the complex numbers generated by complex conjugation. Then $\mathbb{C}^G = \mathbb{R}$, so $\mathbb{C}/\mathbb{R}$ is a Galois extension.

Example: Let L/K be a separable quadratic extension, so that $L = K[t]/(P(t))$, where $P(t)$ is a separable polynomial. Then $P(t)$ splits over L into $(t-\alpha)(t-\overline{\alpha})$, so that the automorphism group of L/K has order 2, the nontrivial element being the unique K-automorphism σ of L which sends $\alpha \mapsto \overline{\alpha}$. Since $L^{\mathrm{Aut}(L/K)}$ is a subextension of the degree 2 extension L/K, it could only be L or K, and since $\sigma(\alpha) = \overline{\alpha} \neq \alpha$, we conclude that the fixed field is K and the extension is Galois. In contrast the automorphism group of an inseparable quadratic extension is trivial, so this extension is not Galois.

Example: Let $K = \mathbb{Q}[t]/(t^3-2) = \mathbb{Q}[\sqrt[3]{2}]$. Since K contains exactly one of the three roots of t^3-2 in $\overline{\mathbb{Q}}$, $\mathrm{Aut}(K/\mathbb{Q})$ is the trivial group and $K/\mathbb{Q}$ is not Galois. On the other hand, the automorphism group of the normal closure $M = \mathbb{Q}[\zeta_3, \sqrt[3]{2}]$ of $K/\mathbb{Q}$ has order 6: since everything is separable, there are three embeddings of $\mathbb{Q}[\sqrt[3]{2}]$ into M, and each of these extends in two ways to an automorphism of M. Any automorphism s of M is determined by an $i \in \{0,1,2\}$ and $j \in \{0,1\}$ such that

$$s : \sqrt[3]{2} \mapsto \zeta_3^i \sqrt[3]{2},\ \zeta_3 \mapsto (\zeta_3)^{(-1)^j}.$$

Since there are six possibilities and six automorphisms, all of these maps must indeed give automorphisms. In particular, there is an order 3 automorphism σ which takes $\sqrt[3]{2} \mapsto \zeta_3 \sqrt[3]{2}$ and fixes ζ_3 and an order 2 automorphism τ which fixes $\sqrt[3]{2}$ and maps $\zeta_3 \mapsto zeta_3^{-1}$. One checks that $\tau\sigma\tau = \tau\sigma\tau^{-1} = \sigma^{-1}$, i.e., $\mathrm{Aut}(L/\mathbb{Q}) \cong S_3$, the symmetric group on three elements. Indeed, these three elements can be viewed as the three roots of t^3-2 in M. Finally, the subgroup fixed by $\{1,\sigma\}$ is precisely K, whereas the generator $\sqrt[3]{2}$ of $K/\mathbb{Q}$ is not fixed by σ, so that we conclude that $M^{\mathrm{Aut}(M/\mathbb{Q})} = \mathbb{Q}$ and $M/\mathbb{Q}$ is Galois.

These examples already suggest that a finite extension K/F is Galois iff it is normal and separable, and in this case $\#\,\mathrm{Aut}(K/F) = [K:F]$. We will show in the next section that these conditions are all equivalent.

Example: The extension $\overline{\mathbb{Q}}/\mathbb{Q}$ is Galois. We cannot show this by some sort of direct computation of $G_{\mathbb{Q}} := \mathrm{Aut}(\overline{\mathbb{Q}}/\mathbb{Q})$: this group is uncountably infinite and has a very complicated structure. Indeed, as an algebraic number theorist I am more or less honorbound to inform you that the group $G_{\mathbb{Q}}$ is the single most interesting group in all of mathematics! We will see that the Galois theory of infinite algebraic extensions cannot be developed in exactly the same way as in the finite case, but is, in theory, easily understood by a reduction to the finite case.

Example: The extension $\mathbb{C}/\mathbb{Q}$ is Galois, as is $\mathbb{C}/\overline{\mathbb{Q}}$. In particular the automorphism group of the complex field is (much) larger than just $\{1,c\}$. In fact we will show that if F has characteristic zero and K is algebraically closed, then K/F is Galois. These results are not part of "Galois theory" as it is usually understood, but rather are facts about automorphism groups of transcendental extensions. These results will be shown in §10.1.

Example: For any field F, $\mathrm{Aut}(F(t)/F)$ is the group of linear fractional transformations: the group $GL_2(F)$ of 2×2 matrices $[[ab]][cd]]$ with $ad \neq bc$ acts by automorphisms on $F(t)$, via $t \mapsto \frac{at+b}{ct+d}$. Scalar matrices – those with $b = d = 0$, $a = c$ – act trivially, so the action factor through to the quotient $PGL_2(F)$ of $GL_2(F)$ by the subgroup $F^\times$ of scalar matrices. It is a standard fact (more in the vein of algebraic geometry than pure field theory) that this is the entire automorphism group of $F(t)$.

Proposition 8.1. *The extension $F(t)/F$ is weakly Galois iff F is infinite.*

Proof. We will need to use a fact from the next section: if G is a finite group of automorphisms acting on a field K, then $[K : K^G] = \#G < \infty$. Therefore if F is finite, $F(t)^{\mathrm{Aut}(F(t)/F)}$ has finite index in $F(t)$, so is certainly not equal to F. Conversely assume F is infinite... □

Remark Aside: I am not aware of a simple necessary and sufficient condition for an extension K/F which is finitely generated, but of infinite degree, to be Galois. When K/F is regular of transcendence degree 1 (two terms which we have not yet defined), one can give such a criterion in terms of the Jacobian $J(C)$ of the corresponding algebraic curve $C_{/F}$, namely K/F is Galois iff $\dim J(C) = 0$ or ($\dim J(C) = 1$ and $J(C)(F)$ is infinite). In particular no such field of *genus* $g \geq 2$ is Galois. One can give some examples of Galois extensions of higher transcendence degree – e.g. the proof of Proposition XX easily adapts to show that $F(t_1, \ldots, t_n)/F$ is Galois if F is infinite – but the general problem seems to be a quite subtle one in birational arithmetic geometry.

8.2. **Finite Galois Extensions.**

Theorem 8.2. *If K/F is a finite field extension,* $\mathrm{Aut}(K/F)$ *is a finite group of cardinality at most* $[K : F]$.

Proof. First recall that the set of F-algebra embeddings σ of K into an algebraic closure $\overline{F}$ is finite, so in particular the subset of such with $\sigma(K) = K$ is finite. This holds because $K = F(\alpha_1, \ldots, \alpha_n)$, and an embedding σ is determined by sending each α_i to one of the at most $d_i = [F[\alpha_i] : F]$ roots of the minimal polynomial of α_i over F in $\overline{F}$. Therefore the set of such embeddings has cardinality at most $d_1 \cdots d_n$. Note that when $K = F[\alpha]$ is simple this is exactly the bound we want, so that e.g. if K/F is separable we are already done.

Now for the general case. Let $\mathrm{Aut}(K/F) = \{\sigma_1, \ldots, \sigma_N\}$ and suppose, for a contradiction, that $N > m = [K : F]$. Let $\alpha_1, \ldots, \alpha_m$ be an F-basis for K, and consider the $N \times m$ matrix A whose (i, j) entry is $\sigma_i(\alpha_j)$. This matrix has rank at most $m < N$, so that its rows are K-linearly dependent: there exist $c_1, \ldots, c_N \in K$, not all 0, such that for all $1 \leq j \leq m$ we have

$$\sum_i c_i \sigma_i(\alpha_j) = 0.$$

For each $x \in K^\times$, there exist $a_1, \ldots, a_m$ in F such that $x = \sum_j a_j \alpha_j$. Then

$$\sum_i c_i \sigma_i(x) = \sum_i c_i \sigma_i(\sum_j a_j \alpha_j) = \sum_i c_i (a_j \sum_j \sigma_j(\alpha_j))$$

$$= \sum_j a_j (\sum_i c_i \sigma_i(\alpha_j)) = 0.$$

But taking $M = K^\times$ all the automorphisms σ_i give characters $M \to K^\times$ hence are K-linearly independent. Therefore in the last equation we must have $c_i = 0$ for all i, a contradiction. □

Proposition 8.3. *(Artin) Let K be a field and G a finite group of automorphisms of K, of cardinality n. Then $[K : K^G] = n$.*

Proof. Step 1: We show that K/K^G has finite degree.[7]

Let $\alpha \in K$, and let $S = \sigma_1, \ldots, \sigma_r$ be a maximal subset of G such that the elements $\sigma_i(\alpha)$ are distinct in K. It follows that for all $\tau \in G$, the r-tuple $v = (\tau\sigma_1\alpha, \ldots, \tau\sigma_r\alpha)$ differs from $w = (\sigma_1\alpha, \ldots, \sigma_r\alpha)$ by a permutation: indeed, since τ is injective, the components of w are all distinct, and if they were not simply a reordering of the components of v, this would contradict the maximality of S. Therefore α is a root of the polynomial

$$f(t) = \prod_{i=1}^{t} (t - \sigma_i\alpha),$$

a polynomial with coefficients in K^G. Moreover, $f(t)$ is separable, and thus K/K^G is separable. Corollary 7.3 applies to show that K/K^G has finite degree, indeed degree equal to the maximal degree $[K^G(\alpha) : K^G]$ of an element $\alpha \in K$.
Step 2: Above, for each α we constructed a polynomial satisfied by α of degree $r \leq n$, it follows that $[K^G : K] \leq n$. On the other hand, by Theorem 8.2 we have $n = \#G \leq \#\operatorname{Aut}(K/K^G) \leq [K^G : K]$. We conclude $[K : K^G] = n$ and $G = \operatorname{Aut}(K/K^G)$. □

Theorem 8.4. *(Omnibus theorem for finite Galois extensions) Let K/F be a finite extension. TFAE:*
(i) $K^{\operatorname{Aut}(K/F)} = F$ ("K/F is Galois.")
(ii) $\#Aut(K/F) = [K : F]$.
(iii) K/F is normal and separable.
(iv) K/F is the splitting field of a separable polynomial.

Proof. Let $G = \operatorname{Aut}(K/F)$. (i) implies (ii) by Proposition XX. (ii) implies (i): we have $F \subset K^G \subset K$, and $[K : K^G] = \#G = [K : F]$, so $K^G = F$.

(iii) implies (iv): if K/F is separable then by the Primitive Element Theorem $K = F[t]/(P(t))$ for some irreducible, separable polynomial P. Since it is normal, P splits in K and therefore K/F is the splitting field of the separable polynomial P. (iv) implies (iii) is essentially the same: since K/F is a splitting field, it is normal; since it is obtained by adjoining roots of separable polynomials, it is separable.

(iv) $\iff$ (ii): We know that the number of embeddings of K into $\overline{F}$ is equal to the separable degree of K/F and that this equals $[K : F]$ iff K/F is separable; moreover, every F-algebra embedding $s : K \to \overline{F}$ has $s(K) = K$ – i.e., gives an automorphism of K iff K/F is normal. □

[7] In many standard treatments of finite Galois theory, the finiteness of K/K^G is an additional assumption. Our source for this stronger version is Lang's *Algebra*.

Corollary 8.5. *A finite extension is a subextension of a finite Galois extension iff it is separable. Any algebraic closure $\overline{F}$ of K contains a unique minimal extension M of K such that M/F is Galois, namely the normal closure of K/F in $\overline{F}$.*

Proof. Since Galois extensions are separable and subextensions of separable extensions are separable, for K/F to be contained in a finite Galois extension it is clearly necessary for it to be separable. If so, then the normal closure M of K/F, being a compositum of the separable extensions $s(K)$ as s ranges over the finite set of distinct F-algebra embeddings of K into $\overline{F}$ is separable and normal, hence Galois. M/K is even the minimal extension of K which is normal over F, so certainly it is the minimal such Galois extension. □

Remark: In view of Corollary XX, it is reasonable to call the normal closure of a finite separable field extension the **Galois closure**.

Theorem 8.6. *(Natural Irrationalities) Let K/F be a finite Galois extension, and let L/F be an arbitrary (not necessarily algebraic) field extension. Then:*
a) The field extension KL/L is Galois.
b) The restriction map $r : \mathrm{Aut}(KL/L) \to \mathrm{Aut}(K/K \cap L)$ is an isomorphism.
c) We have $[KL : L] = [K : K \cap L]$.

Proof. a) This is the assertion that finite Galois extensions have the *base change meta-property*. But all of the following properties have the base-change meta property: being of finite degree, normality and separability. Alternately, since K/F is finite Galois, it is the splitting field of the separable polynomial $f \in F[x]$. Then KL/L is the splitting field of the polynomial $f \in L[x]$, which is still separable because of the Derivative Criterion.
b) Let $\sigma \in \mathrm{Aut}(KL/L)$, and let $r(\sigma)$ denote the restriction of σ to K. Since σ fixes L pointwise and $F \subset L$, also σ fixes F pointwise. So for all $x \in K$, $r(\sigma)(x)$ is an F-conjugate of x; since K/F is normal, this implies $r(\sigma)(x) \in K$ and thus $r(\sigma) \in \mathrm{Aut}(K/F)$. Indeed, because σ pointwise fixes L, $r(\sigma)$ pointwise fixes $K \cap L$ and $r(\sigma) \in \mathrm{Aut}(K/K \cap L)$. This defines a map

$$r : \mathrm{Aut}(KL/L) \to \mathrm{Aut}(K/K \cap L).$$

That r is a group homomorphism is immediate. Moreover, the kernel of r consists of the set of automorphisms α of KL that pointwise fix both K and L and thus also pointwise fix KL: r is injective. Finally we must show that α is surjective. Its image is a subgroup of $\mathrm{Aut}(K/K \cap L)$, which by the Galois correspondence is therefore of the form $\mathrm{Aut}(K/E)$ for some $K \cap L \subset E \subset K$. Now observe that E is pointwise fixed by every $\alpha \in \mathrm{Aut}(KL/L)$, so hence $E \subset (KL)^{\mathrm{Aut}(KL/L)} = L$. It follows that $E \subset K \cap L$ and thus $E = K \cap L$ and α is surjective.
c) By part b) we have

$$[KL : L] = \#Aut(KL/L) = \#\,\mathrm{Aut}(K/K \cap L) = [K : K \cap L].$$ □

8.3. **An Abstract Galois Correspondence.**

Let X be a set and G a group of automorphisms of X, i.e., a subgroup of the group $\mathrm{Sym}(S)$ of all bijections $s : X \to X$. Let $\Lambda(X)$ be the collection of all subsets of X and $\Lambda(G)$ be the collection of all subgroups of G. Both $\Lambda(X)$ and $\Lambda(G)$ are partially ordered sets under inclusion.

For a subset $Y \subset X$, we define

$$G_Y = \{g \in G \mid gy = y \forall y \in Y\},$$

which is a subgroup of G. Dually, for a subgroup H of G, we define

$$X^H = \{x \in X \mid gx = x \forall g \in H\},$$

which is a subgroup of H. (We could define in the same way X^S for any subset $S \subset G$, but one checks immediately that if H is the subgroup generated by S, $X^S = X^H$, so this extra generality leads nowhere.) To be very formal about it, we have thus defined a map

$$\Phi : \Lambda(X) \to \Lambda(G),\ Y \mapsto G_Y$$

and a map

$$\Psi : \Lambda(G) \to \Lambda(X),\ H \mapsto X^H.$$

Let us explore what can be said about these two maps in this extreme level of generality. Statements that we do not prove are exercises in unwinding the definitions and left to the reader. (We do recommend that the reader peform these exercises!)

First, both Φ and Ψ are **anti**-homomorphisms of the partially ordered sets, i.e., if $Y_1 \subset Y_2$, then $\Phi(Y_2) \subset \Phi(Y_1)$, and similarly if $H_1 \subset H_2$ then $\Psi(H_2) \subset H_1$. This implies that $\Psi \circ \Phi : \Lambda(X) \to \Lambda(X)$ and $\Phi \circ \Psi : \Lambda(G) \to \Lambda(G)$ are homomorphisms of partially ordered sets:

$$Y_1 \subset Y_2 \implies X^{G_{Y_1}} \subset X^{G_{Y_2}},$$

$$H_1 \subset H_2 \implies G_{X^{H_1}} \subset G_{X^{H_2}}.$$

Moreover, for all $Y \subset X$ and $H \subset G$ we have

(GC) $\quad Y \subset X^H \iff H \subset G_Y.$

Indeed, both containments assert precisely that every element of H acts trivially on every element of Y. If $H = G_Y$ we certainly have the second containment, therefore by (GC) we have

$$Y \subset X^{G_Y}. \tag{6}$$

Dually with $Y = X^H$ we certainly have the first containment hence (GC) gives

$$H \subset G_{X^H}. \tag{7}$$

Proposition 8.7. *Let H be a subgroup of G, Y a subset of X and $\sigma \in G$. We have:*
a) $\sigma G_Y \sigma^{-1} = G_{\sigma Y}$.
b) $\sigma X^H = X^{\sigma H \sigma^{-1}}$.

Proof. We have $g \in G_{\sigma Y} \iff \forall y \in Y,\ g\sigma y = \sigma y \iff \forall y \in Y,\ \sigma^{-1} g \sigma y = y \iff \sigma^{-1} g \sigma \in G_Y \iff g \in \sigma G_Y \sigma^{-1}$. Similarly, $y \in \sigma X^H \iff \sigma^{-1} y \in X^H \iff \forall h \in h,\ h\sigma^{-1} y = \sigma^{-1} y \iff \forall h \in H,\ (\sigma h \sigma^{-1}) y = y \iff y \in \sigma H \sigma^{-1}$. □

Let us now introduce the following simplified (and symmetric) notation: for $Y \subset X$, we write Y' for G_Y; for $H \subset G$, we write H' for X^H. Equations (6) and (7) now read as $Y \subset Y''$ and $H \subset H''$. Let us call a subset Y of X (resp. a subgroup H of G) **closed** if $Y'' = Y$ (resp. if $H'' = H$).

Proposition 8.8. *For any $Y \in \Lambda(X)$ and $H \in \Lambda(G)$, we have $Y' = Y'''$ and $H' = H'''$. Hence Y' is a closed subgroup of G and X' is a closed subset of X.*

Proof. By (6) we have $Y' \subset (Y')''$ and $Y \subset Y''$. Applying a prime to the latter containment reverses it and hence gives $Y' \supset (Y'')'$. Therefore $Y' = Y'''$. The argument for H is identical. □

Remark: This shows that the operators $''$ on the posets $\Lambda(X)$ and $\Lambda(G)$ are what are called **closure operators**. In general, if $(S, \leq)$ is a partially ordered set, then a map $c : S \to S$ is a closure operator if for all $s \in S$, $s \leq c(s)$, $s \leq t \implies c(s) \leq c(t)$ and $c(c(s)) = c(s)$ for all $s \in S$.

Corollary 8.9. *Let $\Lambda_c(X)$ be the closed subsets of X and $\Lambda_c(G)$ be the closed subgroups of G. Let Φ_c be Φ restricted to $\Lambda_c(X)$ and Ψ_c be Ψ restricted to $\Lambda_c(G)$. Then*

$$\Phi_c : \Lambda_c(X) \to \Lambda_c(G), \ \Psi_c : \Lambda_c(G) \to \Lambda_c(X)$$

give mutually inverse anti-automorphisms of posets.

In fact the proof is immediate from the previous result; again, it is a good exercise for the reader to chase through the definitions and notation to see this.

Corollary 8.10. *a) A closed subgroup H of G is normal iff its corresponding closed subset $Y = H' = H'''$ is stable under all automorphisms of G: for all $\sigma \in G$, $\sigma Y = Y$.*
b) A closed subset Y of X is stable under all automorphisms of G iff the corresponding closed subgroup $H = Y' = Y'''$ is normal in G.

Again, this follows immediately from Proposition XX.

Exercise X.X: Show that if H is a normal subgroup of G, so is its closure H''.

Example: Suppose $\#X > 1$; let $x \in X$ and take $Y = X \setminus x$. Then Y is not a closed subset of X, since any group of automorphisms of X which fixes every element of Y must also fix x.

Example: If $X = \{1, 2\}$ and $G = S_2$ is the full symmetry group of X, then the closed subsets are $\emptyset$ and X; the corresponding closed subgroups are G and the trivial subgroup e. In particular all subgroups are closed. If $X = \{1, 2, 3\}$ and G is the full symmetry group S_3. The closed subsets are $\emptyset$, $\{1\}$, $\{2\}$, $\{3\}$ and X. The corresponding closed subgroups are S_3, $\langle(23)\rangle$, $\langle(13)\rangle$, $\langle(12)\rangle$ and the trivial subgoup e. In particular the (unique) subgroup $H = \langle(123)\rangle$ of order 3 is not closed: $H' = \emptyset$ and $H'' = S_3$. If X is any set, the only subsets invariant under $G = \mathrm{Sym}(X)$ are $\emptyset$ and X itself, so $\mathrm{Sym}(S)$ does not have any nontrivial, proper closed normal subgroups. On the other hand, if $\#X \geq 3$ then $\mathrm{Sym}(S)$ always has a nontrivial, proper normal subgroup (i.e., it is not a simple group): if S is finite, so $\mathrm{Sym}(S) \cong S_n$ take the alternating group A_n (the only possible choice if $n \geq 5$); if S is infinite, take the the subgroup H of elements $g \in \mathrm{Sym}(S)$ such that $X \setminus X^{\langle g \rangle}$ is finite. Then $\#H = \#S$ while $\#\mathrm{Sym}(S) = 2^{\#\mathrm{Sym}(S)}$.

Key example: Let K/F be a field extension, $X = K$ and $G = \mathrm{Aut}(K/F)$. Then every closed subset of X is a subextension K^H of K/F. Corollary XX shows that

there is a bijective correspondence between the closed subextensions of K/F and the closed subgroups of $\mathrm{Aut}(K/F)$. Of course the key word in the previous sentence is "closed": if e.g. $\mathrm{Aut}(K/F)$ is the trivial group (e.g. $K = \mathbb{R}$) then the statement is completely vacuous. In the next section we will show that if K/F is a finite Galois extension, the best possible behavior occurs.

Exercise X.X: Let Λ and Λ' be two partially ordered sets. A **Galois connection** between Λ and Λ' is a pair of order-reversing maps $\Phi : \Lambda \to \Lambda'$, $\Psi : \Lambda' \to \Lambda$ satisfying the analogoue of identity (GC)[8] above: for $x \in \Lambda$, $y \in \Lambda'$, $\Phi(x) \leq y \iff x \leq \Psi(y)$.
a) Check that the entire discussion (except for the bit about conjugation and normality) goes through in this level of generality: we get closure operators on Λ and Λ' such that Φ and Ψ give mutually inverse anti-automorphisms on the subsets of closed elements: $\Phi : \Lambda_c \xrightarrow{\sim} \Lambda'_c$, $\Psi : \Lambda'_c \xrightarrow{\sim} \Lambda_c$.
b) Look for Galois connection. in your everyday (mathematical) life, paying special attention to the closure process. For example, consider the polynomial ring $R = k[t_1, \ldots, t_n]$ over an algebraically closed field k. Let Λ be the set of ideals I of R, and let Λ' be the set of **algebraic** subsets of affine n-space $\mathbb{A}^n$ over k: that is, the subsets of k^n of the form $\bigcap_{i \in I}^r P_i^{-1}(0)$, where $\{P_i\}_{i \in I}$ is a set of elements of R. Define $\Phi : \Lambda \to \Lambda'$ by $I \mapsto V(I)$, the set of points of $(a_1, \ldots, a_n)$ such that $P(a_1, \ldots, a_n)$ for all $P \in I$. Define $\Psi : \Lambda' \to \Lambda$ by $S \mapsto I(S)$, the ideal of all elements of R which vanish at every $(x_1, \ldots, x_n) \in S$. It is no problem to see that this gives a Galois connection. What are the closed ideals? What are the closed algebraic subsets?

8.4. **The Finite Galois Correspondence.**

Let K/F be a finite Galois extension, so that by the general nonsense of the previous section, we get a bijective correspondence between closed subextensions L of K and closed subgroups of $G = \mathrm{Aut}(K/F)$.

Theorem 8.11. *(Fundamental theorem of Galois theory) If K/F is finite Galois, then every subgroup H of $G = \mathrm{Aut}(K/F)$ is closed, i.e., of the form $H = \mathrm{Gal}(K/L)$ for a unique subextension L/K. Conversely, every subextension L is closed, i.e., of the form K^H for a unique subgroup H of G. Therefore the maps $L \mapsto \mathrm{Gal}(K/L)$ and $H \mapsto K^H$ give mutually inverse inclusion-reversing bijections between the set of subextensions of L/K and the set of subgroups of G. Moreover, a subextension L is Galois over F iff the corresponding subgroup $\mathrm{Gal}(K/L)$ is normal in G, and in this case $\mathrm{Aut}(L/F)$ is canonically isomorphic to the quotient $\mathrm{Aut}(K/F)/\mathrm{Aut}(K/L)$.*

Proof: Let L be a subextension of K/F. It is clear that $L \subset K^{\mathrm{Aut}(K/L)}$. But by XXX we know that $[K : K^{\mathrm{Aut}(K/L)}] = \#\,\mathrm{Aut}(K/L)$. Since K/F is Galois, so is K/L, hence $\#\,\mathrm{Aut}(K/L) = [K : L]$. Therefore we must have $K^{\mathrm{Aut}(K/L)} = L$. Moreover, if H is a subgroup of G, we again clearly have $H \subset G_{K^H}$; but we also have $[G : G_{K^H}] = \frac{\#G}{\#\,\mathrm{Aut}(K/K^H)} = [G : H]$, so $H = G_{K^H}$. This shows the Galois correspondence is perfect. Now applying Corollary XX we get that $H = \mathrm{Gal}(K/L)$ is normal in G iff L is stable under all F-algebra automorphisms σ of K. Since K/F is itself normal, this holds iff L is stable under all F-algebra embeddings into an algebraic closure $\overline{F}$, i.e., iff L/F is normal. Finally, suppose that L/F is normal.

[8] In particular "GC" stands for **G**alois **C**onnection.

Then every F-automorphism of K restricts to an F-automorphism of L, giving a natural map $\mathrm{Aut}(K/F) \to \mathrm{Aut}(L/F)$ which is easily checked to be a homomorphism of groups. The map is surjective by the Extension Theorem XX. Its kernel is the subgroup of F-algebra automorphisms of K which fix every element of L, i.e., $\mathrm{Aut}(K/L)$.

This theorem is probably the single most important result in field theory. It reduces the study of the lattice of subextensions of a finite Galois extension K/F to the corresponding lattice of subgroups of the finite group $\mathrm{Aut}(K/F)$, which is much easier to study, e.g. is *a priori* finite. Indeed, if K/F is any finite separable extension, then one may – and should! – apply the Galois correspondence to the Galois closure M/F.

Exercise X.X: Use the Galois Correspondence to give a more natural proof of the Primitive Element Corollary (7.2).

When K/F is Galois, we write $\mathrm{Gal}(K/F)$ for $\mathrm{Aut}(K/F)$ and speak of $\mathrm{Gal}(K/F)$ as the **Galois group** of K/F. We note that some authors (e.g. Shifrin, Kaplansky) use the notation $\mathrm{Gal}(K/F)$ for the automorphism group of an arbitrary field extension, but from the perspective of infinite Galois theory (coming up!) and modern number theory this seems dangerously misleading. Namely, it would then be tempting to call any automorphism group of a finite extension "a Galois group" and this is most certainly at odds with contemporary terminology. Indeed, perhaps the single outstanding problem in field theory is to decide whether, for any finite group G, there is a Galois extension $K/\mathbb{Q}$ such that $\mathrm{Gal}(K/\mathbb{Q}) \cong G$. However, the corresponding statement that any finite group is the automorphism group of some finite extension $K/\mathbb{Q}$ – possibly with $[K : \mathbb{Q}] > \#G$ – is a much weaker one, and indeed this is a known theorem of E. Fried and J. Kollar [FK78]

Composita of Galois extensions: let F be a field and K_1, K_2 two Galois extensions of F. After choosing an algebraic closure $\overline{F}$ of F, since K_1 and K_2 are splitting fields, there is a unique F-algebra embedding of K_i into $\overline{F}$. Since composita of normal (resp. separable) extensions are normal (resp. separable), the compositum $K = K_1 \vee K_2$ is a finite Galois extension. What is the relationship of $\mathrm{Gal}(K/F)$ to $\mathrm{Gal}(K_1/F)$ and $\mathrm{Gal}(K_2/F)$? As above we get surjective restriction maps $\iota_i : \mathrm{Gal}(K/F) \to \mathrm{Gal}(K_i/F)$, and hence a diagonal map $\iota = (\iota_1, \iota_2) : \mathrm{Gal}(K/F) \to \mathrm{Gal}(K_1/F) \times \mathrm{Gal}(K_2/F)$. This composite homomorphism ι need not (of course?) be surjective: e.g. it will not be if $K_1 = K_2$ are nontrivial extensions of F. Rather ι is always injective: since K is generated as a field by K_1 and K_2, a pair of automorphisms σ_i of K_i can extend in at most one way to an automorphism of K. Therefore $\mathrm{Gal}(K/F)$ can naturally be viewed as a subgroup of the product $\mathrm{Gal}(K_1/F) \times \mathrm{Gal}(K_2/F)$.

This is in fact rather useful: let $\mathcal{C}$ be any class of finite groups which is closed under formation of direct products and passage to subgroups, and suppose that K_i/F are two $\mathcal{C}$-Galois extensions, i.e., finite Galois extensions whose Galois groups lie in $\mathcal{C}$. Then the compositum $K_1 \vee K_2$ is a $\mathcal{C}$-Galois extension. E.g. we may profitably take $\mathcal{C}$ to be the class of all finite abelian groups, or the class of all finite solvable

groups. When we turn to infinite Galois theory we will see that we are allowed to take infinite composita as well, and this observation will show that any field admits a maximal $\mathcal{C}$-Galois extension.

Exercise: Let $K_1,\ K_2/F$ be two finite Galois extensions, and $K = K_1K_2$ their compositum. Let H be the image of the map $\iota : \mathrm{Gal}(K/F) \to \mathrm{Gal}(K_1/F) \times \mathrm{Gal}(K_2/F)$. Show that H is normal in $\mathrm{Gal}(K_1/F) \times \mathrm{Gal}(K_2/F)$, and that the quotient $(\mathrm{Gal}(K_1/F) \times \mathrm{Gal}(K_2/F))/H \cong \mathrm{Gal}(K_1 \cap K_2/F)$. In particular, ι is an isomorphism iff $K_1 \cap K_2 = F$.

8.5. **The Normal Basis Theorem.**

Let K/F be a finite degree field extension. Then a basis $\{\alpha_1, \ldots, \alpha_n\}$ of K as an F-vector space is a **normal basis** if all of its elements lie in the same $\mathrm{Aut}(K/F)$-orbit, i.e., if for all $1 \leq i \leq n$ there exists $\sigma \in \mathrm{Aut}(K/F)$ such that $\alpha_i = \sigma\alpha_1$.

Exercise: If a finite extension K/F admits a normal basis, it is Galois.

The main result of this section is the converse: every finite Galois extension admits a normal basis. A lot of literature has been written on this result. Our treatment follows [CW50] and [We09, §3.6]. It is certainly not the shortest treatment available, but it proceeds by establishing several preliminary results which are of some interest in their own right.

Every known proof of the existence of normal bases must negotiate a fundamental dichotomy between finite fields and infinite fields. This dichotomy comes up several times in field theory, algebra and algebraic geometry (another good example of a theorem for which the finite field case must be taken separately is the **Noether Normalization Theorem**), but often without much fanfare our explanation. To our mind at least, the source of the trouble is the different behavior of the evaluation map on polynomials over finite domains versus infinite integal domains. (A geometer might point to the fact that for any $n \in \mathbb{Z}^+$, a field K is infinite iff the K-rational points of affine n-space over K are Zariski dense, but in fact this comes down to the same algebraic observation.)

Lemma 8.12. *Let $R \subset S$ be an extension of domains and $n \in \mathbb{Z}^+$. TFAE:*
(i) For all $f \in S[t_1, \ldots, t_n]$, $f(a_1, \ldots, a_n) = 0$ for all $(a_1, \ldots, a_n) \in R^n \implies f = 0$.
(ii) R is infinite.

Proof. (i) $\implies$ (ii): We prove the contrapositive. Note that any finite domain is a field, so suppose $R = \mathbb{F}_q$. Let $f(t) = t_1^q - t_1$. Then for all $a = (a_1, \ldots, a_n) \in \mathbb{F}_q^n$, $f(a) = a_1^q - a_1 = 0$.
(ii) $\implies$ (i): We go by induction on n.
Base Case ($n = 1$): suppose $f \in S[t]$ is a polynomial which is not the zero polynomial. Then it has degree $d \geq 0$ and by the Root-Factor Theorem has at most d roots in the fraction field of R, hence *a fortori* at most d roots in R. But $\#R \geq \aleph_0 > d$, so there exists $a_1 \in R$ with $f(a_1) = 0$.
Induction Step: Suppose $n > 1$ and that every polynomial in $n-1$ variables with S-coefficients which is not the zero polynomial has a R-rational root. Let $f(t_1, \ldots, t_{n-1}, z) \in S[t_1, \ldots, t_{n-1}, z]$. Put $S' = S[t_1, \ldots, t_{n-1}]$, so f may

be identified with a nonzero polynomial $g(z) \in S'[z]$. Applying the Base Case, there exists $A \in R'$ such that $0 \neq g(A) \in R'$. Now $g(A)$ is a nonzero element of $S' = S[t_1, \ldots, t_{n-1}]$, so by induction there exist $a_1, \ldots, a_{n-1} \in R$ such that $g(A(a_1, \ldots, a_{n-1})) \neq 0$. Putting $a_n = A(a_1, \ldots, a_{n-1})$ we have

$$f(a_1, \ldots, a_{n-1}, a_n) = g(A(a_1, \ldots, a_{n-1})) \neq 0.$$

□

Proposition 8.13. *Any finite cyclic extension K/F admits a normal basis.*

Proof. Let K/F be cyclic of degree n with $\mathrm{Gal}(K/F) = \langle \alpha \rangle$. We may endow K with the structure of an $F[t]$-module extending its F-module structure by putting $t \cdot x = \sigma(x)$ for all $x \in K$. Then $t^n - 1$ annihilates K; moreover, by linear independence of characters, no smaller degree polynomial does so. It follows that as an $F[t]$-module, K is isomorphic to $F[t]/(t^n - 1)$. Thus there exists $\alpha \in K$ such that $\mathrm{ann}(\alpha) = (t^n - 1)$ – take, e.g., the preimage of 1 (mod $t^n - 1$) under an isomorphism – so the elements $\alpha, \sigma\alpha, \sigma^2\alpha, \ldots, \sigma^{n-1}\alpha$ are F-linearly independent and thus give a normal basis. □

Lemma 8.14. *Let K/F be a degree n Galois extension, and write $\mathrm{Aut}(K/F) = \{\sigma_i\}_{i=1}^n$. For $\alpha_1, \ldots, \alpha_n \in K$, TFAE:*
(i) $\alpha_1, \ldots, \alpha_n$ is an F-basis of K.
(ii) The matrix $A \in M_n(K)$ with $A_{ij} = \sigma_i\alpha_j$ is nonsingular.

Proof. (i) $\implies$ (ii) follows almost immediately from the (K-)linear independence of the characters $\sigma_1, \ldots, \sigma_n$: details are left to the reader.
(ii) $\implies$ (i): We argue by contraposition: suppose $\alpha_1, \ldots, \alpha_n$ is *not* an F-basis for K, so there exist $a_1, \ldots, a_n \in F$, not all zero, with $a_1\alpha_1 + \ldots + a_n\alpha_n = 0$. Then for all i we have

$$\sum_{j=1}^n a_j A_{ij} = \sum_{j=1}^n a_j \sigma_i \alpha_j = \sigma_i(\sum_{j=1}^n a_i \alpha_j) = 0,$$

which shows that the columns of the matrix A are linearly dependent. □

By linear independence of characters, for any field extension K/F, any finite set of automorphisms $\sigma_1, \ldots, \sigma_n \in \mathrm{Aut}(K/F)$ is K-linearly independent. If K/F is a Galois extension **and F is infinite**, we have the following significantly stronger independence result.

Theorem 8.15. *Let K/F be a finite degree Galois extension of infinite fields. Then the elements $\sigma_1, \ldots, \sigma_n$ of $\mathrm{Aut}(K/F)$ are **algebraically independent** – if $0 \neq f(t_1, \ldots, t_n) \in K[t_1, \ldots, t_n]$, there exists $\alpha \in K$ such that $f(\sigma_1(\alpha), \ldots, \sigma_n(\alpha)) \neq 0$.*

Proof. As a matter of notation, for an n-tuple $(x_1, \ldots, x_n) \in K^n$, we will denote by $(x_1, \ldots, x_n)^\bullet$ the corresponding column vector, i.e., element of $M_{n,1}(K)$. If it brings no confusion, we will suppress indices by writing $x^\bullet$ for $(x_1, \ldots, x_n)^\bullet$.

Let $\alpha_1, \ldots, \alpha_n$ be a basis for K/F. Define $A \in M_n(K)$ by $A_{ij} = \sigma_i\alpha_j$. By Lemma 8.14, A is nonsingular. Now let $c = (c_1, \ldots, c_n) \in F^n$ and put

$$\alpha = \sum_{j=1}^n c_j \alpha_j.$$

Then for all $1 \leq i \leq n$,

$$\sigma_i(\alpha) = \sum_{j=1}^{n} A_{ij} c_j,$$

so

$$\sigma(\alpha)^\bullet = Ac^\bullet.$$

Seeking a contradiction, we suppose that for all $\sigma \in K$, $f(\sigma_1(\alpha), \ldots, \sigma_n(\alpha)) = 0$. By the above, this can be reexpressed as

$$0 = f(\sigma(\alpha)^\bullet) = f(Ac^\bullet)$$

for all $c \in F^n$. Thus the polynomial

$$g(t) = g(t_1, \ldots, t_n) = f(At^\bullet) \in K[t_1, \ldots, t_n]$$

vanishes at every $c \in F^n$, so by Lemma 8.12, $g = 0$. So $f(t) = g(A^{-1}t^\bullet) = 0$. □

Exercise: Show that Theorem 8.15 fails for every finite extension of finite fields.

Theorem 8.16. *(Normal Basis Theorem) Let K/F be a finite Galois extension of degree n. Then there exists $\alpha \in K$ such that the set $\{\sigma\alpha\}_{\sigma \in \mathrm{Gal}(K/F)}$ is a basis of K as an F-vector space.*

Proof. By Proposition 8.12 we may assume that F, and hence also K, is infinite. Write out the elements of $\mathrm{Aut}(K/F)$ as $1 = \sigma_1, \sigma_2, \ldots, \sigma_n$. Let $t_1, \ldots, t_n$ be independent indeterminates, and consider the matrix B with $B_{ij} = t_k$, where $\sigma_i\sigma_j = \sigma_k$. In this matrix each t_i appears exactly once in each row and column, so the specialization $t_1 = 1$, $t_i = 0$ for all $i > 1$ gives rise to a permutation matrix with determinant ± 1. It follows that $d(t_1, \ldots, t_n) = \det B$ is a nonzero element of the polynomial ring $K[t_1, \ldots, t_n]$. Applying Theorem 8.15, there exists $\alpha \in K$ such that $d(\sigma_1(\alpha), \ldots, \sigma_n(\alpha)) \neq 0$.

For $1 \leq j \leq n$, put $\alpha_j = \sigma_j(\alpha)$. Then the matrix A with $A_{ij} = \sigma_i\alpha_j = \sigma_i\sigma_j\alpha = \sigma_k\alpha$ nonsingular, so by Lemma 8.14 $\sigma_1\alpha, \ldots, \sigma_j\alpha$ is an F-basis of K. □

Exercise: Explain how the Normal Basis Theorem gives a stronger result than the Primitive Element Corollary in the case of a Galois extension.

8.6. **Hilbert's Theorem 90.**

Let G be a group, and let M be a G-module, i.e., commutative group on which G acts $\mathbb{Z}$-linearly: that is, we are given a homomorphism $G \to \mathrm{Aut}_{\mathbb{Z}}(M)$. Let $Z^1(G, M)$ be the set of all maps $f : G \to M$ which satisfy the **cocycle condition**:

$$\forall \sigma, \tau \in G,\ f(\sigma\tau) = f(\sigma) + \sigma(f(\tau)).$$

Let $B^1(G, M)$ be the set of maps $f : G \to M$ such that there is $a \in M$ with $f(\sigma) = \sigma(a) - a$ for all $\sigma \in G$.

Exercise: a) Show that $Z^1(G, M)$ and $B^1(G, M)$ are commutative groups under pointwise addition.
b) Show that $B^1(G, M) \subset Z^1(G, M)$.

We may therefore define

$$H^1(G, M) = Z^1(G, M)/B^1(G, M),$$

the **first cohomology group of G with coefficients in M**.

Exercise: Suppose that G acts trivially on M. Show that $H^1(G,M) = \mathrm{Hom}(G,M)$, the group of all homomorphisms from G to M.

Now observe that if K/F is a field extension and $G = \mathrm{Aut}(K/F)$, then both K (as an additive group) and $K^\times$ (as a multiplicative group) are G-modules.

Theorem 8.17. *Let K/F be a finite Galois extension, with Galois group $G = \mathrm{Aut}(K/F)$.*
a) $H^1(G,K) = 0$.
b) $H^1(G,K^\times) = 0$.

Proof. a) Let $f : G \to K$ be a 1-cocycle. Since K/F is finite separable, by Theorem X.X there is $c \in K$ with $\mathrm{Tr}_{K/F}(c) = 1$. Put

$$b = \sum_{\sigma \in G} f(\sigma)\sigma(c),$$

so

$$\tau(b) = \sum_{\sigma \in G} \tau(f(\sigma))(\tau\sigma)(c)$$

$$= \sum_{\sigma \in G} \left(f(\tau\sigma) - f(\tau)\right)(\tau\sigma)(c) = \sum_{\sigma \in G} f(\tau\sigma)(\tau\sigma)(c) - \sum_{\sigma \in G} f(\tau)(\tau\sigma)(c)$$

$$= b - f(\tau) \cdot \tau\left(\sum_{\sigma \in G} \sigma(c)\right) = b - f(\tau).$$

Thus $f(\tau) = b - \tau(b)$ for all $\tau \in G$, so $f \in B^1(G,K)$.
b) Let $f : G \to K^\times$ be a 1-cocycle. By independence of characters, there is $c \in K$ such that $\sum_{\sigma \in G} f(\sigma)\sigma(c) \neq 0$; fix such a c and put $b = \sum_{\sigma \in G} f(\sigma)\sigma(c)$. Then

$$\tau(b) = \sum_{\sigma \in G} \tau(f(\sigma))(\tau\sigma)(c),$$

so

$$f(\tau)\tau(b) = \sum_{\sigma \in G} f(\tau)\tau(f(\sigma)) \cdot (\tau\sigma)(c) = \sum_{\sigma \in G} f(\tau\sigma) \cdot (\tau\sigma)(c) = b,$$

i.e., $f(\tau) = b/\tau(b)$. So $f \in B^1(G,K^\times)$. □

The following is a basic result from group cohomology.

Theorem 8.18. *Let $n \in \mathbb{Z}^+$, and let $G = \langle \sigma \mid \sigma^n = 1 \rangle$ be a finite cyclic group. For any G-module M, we have*

$$H^1(G,M) \cong \{x \in M \mid (1 + \sigma + \ldots + \sigma^{n-1})(x) = 0\}/\{\sigma x - x \mid x \in M\}.$$

Combining Theorems 8.17 and 8.18 we immediately deduce the following famous result of D. Hilbert, the 90th theorem in his *Zahlbericht*. However, because our focus here is on field-theoretic methods, we will not give a proof of Theorem 8.18 but rather a purely field-theoretic proof of Hilbert's Satz 90.

Theorem 8.19. *(Hilbert's Satz 90) Let K/F be a finite Galois extension with cyclic Galois group $G = \langle \sigma \mid \sigma^n = 1 \rangle$.*
a) For $c \in K$, the following are equivalent:
(i) $\mathrm{Tr}_{K/F}(c) = 0$.

(ii) There is $a \in K$ such that $c = a - \sigma(a)$.
b) For $c \in K$, the following are equivalent:
(i) $N_{K/F}(c) = 1$.
(ii) There is $a \in K^\times$ such that $c = \frac{a}{\sigma(a)}$.

Proof. Step 1: Because Galois conjugate elements have the same norm and trace, in both parts a) and b) the implications (ii) $\implies$ (i) are immediate.
Step 2: Let $c \in K$ be such that $\mathrm{Tr}_{K/F}(c) = 0$. Since K/F is separable, by Theorem X.X there is $b \in K$ with $\mathrm{Tr}_{K/F}(b) = 1$.[9]

Put

$$a = cb + (c + \sigma(c))\sigma(b) + \ldots + (c + \sigma(c) + \ldots + \sigma^{n-2}(c))\sigma^{n-2}(b).$$

Then

$$\sigma(a) = \sigma(c)\sigma(b) + (\sigma(c) + \sigma^2(c))\sigma^2(b) + \ldots + (\sigma(c) + \ldots + \sigma^{n-1}(c))\sigma^{n-1}(b).$$

Since $\mathrm{Tr}_{K/F}(c) = c + \sigma(c) + \ldots + \sigma^{n-1}(c) = 0$, we have

$$a - \sigma(a) = cb + c\sigma(b) + \ldots + c\sigma^{n-1}b = c\,\mathrm{Tr}_{K/F}(b) = c.$$

Step 3: Let $c \in K$ be such that $N_{K/F}(c) = 1$. By Dedekind's linear independence of characters, there is $b \in K$ with

$$a = b + c\sigma(b) + c\sigma(c)\sigma^2(b) + \ldots + c\sigma(c)\cdots\sigma^{n-2}(c)\sigma^{n-1}(b) \neq 0.$$

Then

$$c\sigma(a) = c\sigma(b) + c\sigma(c)\sigma^2(b) + \ldots + c\sigma(c)\cdots\sigma^{n-1}(c)b = a,$$

so

$$c = \frac{a}{\sigma(a)}.$$

□

We will use Theorem 8.19 later on in our study of cyclic extensions.

include application to Pythagorean triples

8.7. **Infinite Algebraic Galois Theory.**

Theorem 8.20. *For an algebraic field extension K/F, TFAE:*
(i) $K^{\mathrm{Aut}(K/F)} = F$. ("$K/F$ is Galois.")
(ii) K is normal and separable.
(iii) K is the splitting field of a set (possibly infinite) of separable polynomials.

Proof. The equivalence of (ii) and (iii) follows from our characterization of normal and separable algebraic extensions.
(i) $\implies$ (ii): (Morandi, p. 40something) FIXME!!!
(ii) $\implies$ (i): Let $\alpha \in K \setminus F$. Then the minimal polynomial P for α over K splits in K and has at least one other distinct root β. There is a unique F-algebra embedding $\sigma : F[\alpha] \to K$ that sends α to β; as usual, we can extend σ to an automorphism of $\overline{F}$ and then the restriction of σ to K is an automorphism of K (since K is normal) for which $\sigma(\alpha) \neq \alpha$. Therefore $K^{\mathrm{Aut}(K/F)} = F$. □

[9] Alternately, since K/F is Galois, $\mathrm{Tr}_{K/F}(x) = x + \sigma(x) + \ldots + \sigma^{n-1}(x)$. It follows from Dedekind's linear independence of characters that $\mathrm{Tr}_{K/F}$ is not identically zero, and since it is an F-linear functional it must then be surjective.

Let us now revisit the abstract setting of section XX in the somewhat less trivial present framework: $X = K = F^{\mathrm{sep}}$, $G = \mathrm{Aut}(K/F)$. Then the maps $L \mapsto \mathrm{Gal}(K/L)$ and $H \mapsto K^H$ give a bijective correspondence between **closed** subextensions L of K/F and **closed** subgroups H of G. The key fact is the following

Lemma 8.21. *Every subextension L of K/F is closed, i.e. $K^{\mathrm{Gal}(K/L)} = L$.*

Proof. FIX ME!!! □

8.8. **A Characterization of Normal Extensions.**

Lemma 8.22. *a) Let K be a field with algebraic closure $\overline{K}$. Let L/K be a purely inseparable extension, and let $\sigma : L \hookrightarrow \overline{K}$ be a K-algebra embedding. Then σ is an L-algebra embedding, i.e., for all $x \in L$, $\sigma(x) = x$.*
b) Let K/F be an algebraic field extension and F_i the purely inseparable closure of F in K. Then $\mathrm{Aut}(K/F_i) = \mathrm{Aut}(K/F)$.

Proof. a) Any element $y \in L$ satisfies a purely inseparable polynomial $P(t) = t^{p^n} - x$ for some $x \in K$. The map σ must send y to some root of $P(t)$, of which there is only one.
b) Choose an algebraic closure $\overline{K}$ of K. Let $\sigma \in \mathrm{Aut}(K/F)$; by X.X σ extends to an automorphism of $\overline{K}$, which we continue to denote by σ. Applying part a) to σ with $K_i = L$, we get that σ fixes K_i pointwise, qed. □

Theorem 8.23. *For an algebraic extension K/F, TFAE:*
(i) The extension $K^{\mathrm{Aut}(K/F)}/F$ is purely inseparable.
(ii) K/F is normal.

Proof. (i) $\implies$ (ii): Put $L = K^{\mathrm{Aut}(K/F)}$. Let $\overline{F}$ be an algebraic closure of K and let $\sigma : K \to \overline{F}$ be an F-algebra embedding, which we may extend to an automorphism of $\overline{F}$. Since L/F is purely inseparable, by Lemma 8.22b) we have $\sigma \in \mathrm{Aut}(\overline{K}/L)$. In other words, σ fixes L pointwise. But K/L is Galois, hence normal, so for any embedding $\sigma : K \hookrightarrow \overline{F}$ which fixes L pointwise we have $\sigma(K) = K$.
(ii) $\implies$ (i): Let F_i be the purely inseparable closure of F in K. Since K/F is normal, so is K/F_i. Moreover, by Theorem X.X and Corollary X.X, K/F_i is separable. Thus K/F_i is Galois, so (applying Lemma) 8.22) we get

$$K^{\mathrm{Aut}(K/F)} = K^{\mathrm{Aut}(K/F_i)} = F_i.$$

□

9. Solvable Extensions

9.1. **Cyclotomic Extensions.**

9.1.1. *Basics.*

Let K be a field. An element $x \in K^\times$ is a **root of unity** if there is $n \in \mathbb{Z}^+$ such that $x^n = 1$; equivalently, x lies in the torsion subgroup of $K^\times$. We put

$$\mu_n(K) = \{x \in K \mid x^n = 1\},$$

the nth roots of unity in K. We put $\mu(K) = \bigcup_{n \geq 1} \mu_n(K)$. Thus $\mu_n(K)$ and $\mu(K)$ are subgroups of $K^\times$ and $\mu(K) = K^\times[\mathrm{tors}]$.

Lemma 9.1. *For any field K and $n \in \mathbb{Z}^+$, we have $\#\mu_n(K) \leq n$.*

Proof. The elements of $\mu_n(K)$ are the roots of the polynomial $t^n - 1$ over K, and a nonzero polynomial over a field cannot have more roots than its degree. □

Lemma 9.2. *For any field K and $n \in \mathbb{Z}^+$, $\mu_n(K)$ is a finite cyclic group.*

Proof. By Lemma 9.2, $\mu_n(K)$ is finite. We use the **Cyclicity Criterion**: a finite group G is cyclic iff for all $d \in \mathbb{Z}^+$ there are at most d element of order n in G. This holds in $\mu_n(K)$ since the polynomial $t^d - 1$ can have no more than d roots. □

Example: Fix $n \in \mathbb{Z}$. For $0 \leq k < n$, the elements $e^{\frac{2\pi k i}{n}}$ are distinct nth roots of unity in $\mathbb{C}$. So $\#\mu_n(\mathbb{C}) = n$.

Exercise: Let K be an ordered field. Show that $\mu(K) = \{\pm 1\}$.

An element of $K^\times$ of exact order n is called a **primitive nth root of unity**.

Proposition 9.3. *Let K be an algebraically closed field. For $n \in \mathbb{Z}^+$, TFAE:*
(i) $\operatorname{char} K \nmid n$.
(ii) $\#\mu_n(K) = n$.
(iii) K admits a primitive nth root of unity.
(iv) K admits precisely $\varphi(n)$ primitive nth roots of unity.

Proof. (i) $\iff$ (ii): Let $f(t) = t^n - 1$. Then $f'(t) = nt^{n-1}$. Thus $\operatorname{char} K \nmid n \iff \gcd(f, f') = 1 \iff t^n - 1$ has n distinct roots $\iff \#\mu_n(K) = n$.
(ii) $\iff$ (iii): By Lemma 9.2, $\mu_n(K)$ is a finite, cyclic n-torsion abelian group. Thus it has order n iff it has an element of order n.
(ii) $\implies$ (iv): (ii) holds $\iff$ $\mu_n(K)$ is cyclic of order n, in which case it has precisely $\varphi(n)$ generators.
(iv) $\implies$ (iii): Since for all $n \in \mathbb{Z}^+$, $\varphi(n) \geq 1$, this is clear. □

Exercise: a) Let K be an algebraically closed field of characteristic zero. Show that $\mu(K) \cong \varinjlim_{n \in \mathbb{Z}^+} \mathbb{Z}/n\mathbb{Z}$.
b) Let K be an algebraically closed field of characteristic $p \geq 0$. Show that $\mu(K) \cong \varinjlim_{n \in \mathbb{Z}^+,\ p \nmid n} \mathbb{Z}/n\mathbb{Z}$.

Exercise: Show that for any field K, $\mu(K^{\mathrm{sep}}) = \mu(\overline{K})$.

Henceforth we only consider μ_n for $\operatorname{char} K \nmid n$.

For a field K, we denote by K^{cyc} the field obtained by adjoining to K all roots of unity in a fixed algebraic closure $\overline{K}$. Then K^{cyc} is the splitting field of the set $\{t^n - 1\}_{\operatorname{char} K \nmid n}$ of separable polynomials, so is an algebraic Galois extension, the **maximal cyclotomic extension of K**. For $n \in \mathbb{Z}^+$ with $\operatorname{char} K \nmid n$, let $K(\mu_n)$ be the splitting field of the separable polynomial $t^n - 1$, the **nth cyclotomic extension**. Thus $K^{\mathrm{cyc}} = \varinjlim K(\mu_n)$.

For a field K, it is traditional to denote by ζ_n a primitive nth root of unity in K^{sep}. When $K = \mathbb{C}$, the standard choice is $\zeta_n = e^{\frac{2\pi i}{n}}$. There is an advantage to this choice: for all $m \mid n$, we have

$$\zeta_n^{\frac{n}{m}} = \zeta_m. \tag{8}$$

Exercise: Let K be any algebraically closed field.
a) Show that one may choose, for all $n \in \mathbb{Z}^+$ with $\operatorname{char} K \nmid n$, a primitive nth root of unity ζ_n such that the compatibility relation (8) holds.
b) In how many ways is it possible to do this?
(Suggestion: express your answer as an inverse limit of finite sets.)

Proposition 9.4. *Let K be a field and $n \in \mathbb{Z}^+$ with $\operatorname{char} K \nmid n$.*
a) We have $K(\mu_n) = K(\zeta_n)$.
b) There is a canonical injection $a_n : \operatorname{Aut}(K(\zeta_n)/K) \hookrightarrow (\mathbb{Z}/n\mathbb{Z})^\times$.

Proof. a) In other words, the assertion is that by adjoining any one primitive root of unity, we get the splitting field of the polynomial $t^n - 1$. Since every nth root of unity is a power of ζ_n, this is clear.
b) For $\sigma \in \operatorname{Aut}(K(\zeta_n)/K)$, $\sigma(\zeta_n)$ is a primitive nth root of unity: any automorphism of a field preserves the order of elements of the multiplicative group of that field. Thus $\sigma(\zeta_n) = \zeta_n^{a_n(\sigma)}$ for a unique $a_n(\sigma) \in (\mathbb{Z}/n\mathbb{Z})^\times$. It is immediate that $\sigma \mapsto a_n(\sigma)$ is a group homomorphism. Finally, if $a_n(\sigma) = 1$, then $\sigma(\zeta_n) = \zeta_n$, so σ fixes $K(\zeta_n)$ and is thus trivial. □

Exercise: a) In order to define a_n we chose a primitive nth root of unity $\zeta_n \in K^{\mathrm{sep}}$. Show that the homomorphism a_n is in fact independent of this choice.
b) Suppose that $m \mid n$. Show that we have a commutative diagram

$$\operatorname{Aut}(K(\zeta_n)/K) \stackrel{a_n}{\to} (\mathbb{Z}/n\mathbb{Z})^\times$$

$$\operatorname{Aut}(K(\zeta_m)/K) \stackrel{a_m}{\to} (\mathbb{Z}/m\mathbb{Z})^\times,$$

where the map $(\mathbb{Z}/n\mathbb{Z})^\times \to (\mathbb{Z}/m\mathbb{Z})^\times$ is the induced map on units of the quotient map $\mathbb{Z}/n\mathbb{Z} \to \mathbb{Z}/m\mathbb{Z}$.
c) Deduce that there is an injection

$$a : \operatorname{Aut}(K^{\mathrm{cyc}}/K) \hookrightarrow \varprojlim_{n \in \mathbb{Z}^+, \operatorname{char} K \nmid n} (\mathbb{Z}/n\mathbb{Z})^\times.$$

d) In particular, for any prime $\ell \neq \operatorname{char} K$, there is an injection

$$\chi_\ell : \operatorname{Aut}(\varinjlim K(\mu_{\ell^n})/K) \to \mathbb{Z}_\ell^\times,$$

called the **ℓ-adic cyclotomic character**. When $\operatorname{char} K = 0$, there is an injection

$$\chi : \operatorname{Aut}(K^{\mathrm{cyc}}/K) \hookrightarrow \hat{\mathbb{Z}}^\times,$$

the **adelic cyclotomic character**.

9.1.2. *Cyclotomic Polynomials.*

For $n \in \mathbb{Z}^+$, let $\Phi_n(t)$ be the unique monic polynomial with roots the primitive nth roots of unity in $\mathbb{C}$.[10]

Proposition 9.5. *a) For all $n \in \mathbb{Z}^+$, we have*

$$\prod_{d \mid n} \Phi_d(t) = t^n - 1. \tag{9}$$

[10]The use of $\mathbb{C}$ here is somewhere between tradition and psychology: any algebraically closed field of characteristic zero – e.g. $\overline{\mathbb{Q}}$ – would serve as well.

b) For all $n \in \mathbb{Z}^+$, $\Phi_n(t) \in \mathbb{Z}[t]$.
c) For all $n \in \mathbb{Z}^+$, we have

$$\Phi_n(t) = \prod_{d|n} (t^d - 1)^{\mu(\frac{n}{d})}. \tag{10}$$

Proof. a) Both sides of (9) are monic polynomials with $\mathbb{C}$ coefficients whose roots are precisely the nth roots of unity in $\mathbb{C}$. So they are equal.
b) By strong induction on n. The base case is clear: $\Phi_1(t) = t - 1$. Now suppose $n > 1$ and that $\Phi_d(t) \in \mathbb{Z}[t]$ for all proper divisors d of n. Then $Q(t) = \prod_{d|n,\ d \neq n} \Phi_d(t) \in \mathbb{Z}[t]$ is a monic polynomial and $Q(t)\Phi_n(t) = t^n - 1$. Now imagine actually performing polynomial long division of $t^n - 1$ by $Q(t)$ to get $\Phi_n(t)$: since $t^n - 1, \Phi_n(t) \in \mathbb{Z}[t]$ are monic, the quotient $\Phi_n(t)$ has $\mathbb{Z}$-coefficients.
c) This follows from part a) by the Möbius Inversion Formula applied in the commutative group $\mathbb{Q}(t)^\times$.[11] □

Theorem 9.6. *Let $n \in \mathbb{Z}^+$ and let K be a field of characteristic p. Regard $\Phi_n(t) \in \mathbb{F}_p[t] \subset K[t]$. Then $\Phi_n(t)$ is a separable polynomial whose roots in $\overline{K}[t]$ are precisely the primitive nth roots of unity.*

Proof. By the Derivative Criterion $t^n - 1 \in K[t]$ is separable; by (8) so is $\Phi_n(t)$. It is clear that the $\varphi(n)$ roots of $\Phi_n(t)$ in $\overline{K}$ are nth roots of unity; that they are the $\varphi(n)$ *primitive* nth roots of unity follows by an easy induction argument. □

Exercise: Let p be a prime number and $a \in \mathbb{Z}^+$.
a) Show that $\Phi_p(t) = 1 + t + \ldots + t^{p-1}$.
b) Show that $\Phi_{2p}(t) = 1 - t + \ldots + (-t)^{p-1}$.
c) Show that $\Phi_{p^a}(t) = \Phi_p(t^{p^{a-1}})$.

Exercise: For $n \in \mathbb{Z}^+$, let $r(n) = \prod_{p|n} p$. Show:

$$\Phi_n(t) = \Phi_{r(n)}(t^{\frac{n}{r(n)}}).$$

Exercise: Let $n \in \mathbb{Z}^+$.
a) Show: for all $n \geq 2$, the constant coefficient of $\Phi_n(t)$ is 1.
b) Show: for all $n \neq 2$, the product of the primitive nth roots of unity in $\mathbb{C}$ is 1.

Theorem 9.7. *(Gauss-Kronecker) For all $n \in \mathbb{Z}^+$, $\Phi_n(t) \in \mathbb{Q}[t]$ is irreducible.*

Proof. Since $\Phi_n(t) \in \mathbb{Z}[t]$ is monic and $\mathbb{Z}$ is a UFD, by Gauss's Lemma it is equivalent to show that Φ_n is irreducible in $\mathbb{Z}[t]$. We may write $\Phi_n(t) = f(t)g(t)$ with $f, g \in \mathbb{Z}[t]$ monic and f irreducible, and the goal is to show $g = 1$.
Step 1: Let α be a root of $f(t) \in \overline{\mathbb{Q}}$ (hence a primitive nth root of unity) and let p be a prime number not dividing n. We CLAIM that α^p is also a root of $f(t)$.
PROOF OF CLAIM: Suppose not; then, since $p \nmid n$, α^p is a primitive nth root of unity, so α^p is a root of g. Thus α is a root of $h(t^p)$. Since f is monic irreducible and $f(\alpha) = 0$, f is the minimal polynomial for α, so there is $h \in \mathbb{Z}[t]$ with $f(t)h(t) = g(t^p)$. Now apply the homomorphism $\mathbb{Z}[t] \to \mathbb{Z}/p\mathbb{Z}[t]$, $f \mapsto \overline{f}$: we get

$$\overline{\Phi_n} = \overline{fh}.$$

[11] In fact we don't need this in what follows – it is just a pretty formula.

For any polynomial $a(t) \in \mathbb{F}/p\mathbb{Z}[t]$ we have $a(t^p) = a(t)^p$, and thus

$$\overline{g}^p = \overline{fh}.$$

Let $\overline{q}$ be an irreducible factor of $\overline{f}$. Then $\overline{q} \mid \overline{f} \mid \overline{g^p}$, so $\overline{q} \mid \overline{g}$. It follows that $\overline{g}^2 \mid \overline{f}\overline{g} = \overline{\Phi_n}$. But since $p \nmid n$, the Derivative Criterion still holds to show that $\overline{\Phi_n} \in \mathbb{Z}/p\mathbb{Z}[t]$ is separable: contradiction.
Step 2: Let β be any root of $\Phi_n(t)$ in $\overline{\mathbb{Q}}$. Then β and α are both primitive nth roots of unity, so that there is a sequence of (not necessarily distinct) prime numbers $p_1, \ldots, p_r$ with $\gcd(p_1, \ldots, p_r, n) = 1$ and $\alpha^{p_1 \cdots p_r} = \beta$. Applying Step 1 successively to $\alpha, \alpha^{p_1}, \ldots, \alpha^{p_1 \cdots p_{n-1}}$ we find that β is also a root of $f(t)$. Thus f has as its roots *all* primitive nth roots of unity, i.e., $f = \Phi_n$, and Φ_n is irreducible. □

9.1.3. *Some Applications.*

Corollary 9.8. *For any $n \in \mathbb{Z}^+$, the extension $\mathbb{Q}(\mu_n)/\mathbb{Q}$ is Galois, with $\mathrm{Aut}(\mathbb{Q}(\mu_n)/\mathbb{Q})$ canonically isomorphic to $(\mathbb{Z}/n\mathbb{Z})^\times$.*

Exercise: Prove it.

Exercise: Let $m, n \in \mathbb{Z}^+$ with $m \mid n$.
a) Show: $\mathbb{Q}(\mu_m) \subseteq \mathbb{Q}(\mu_n)$.
b) Show: $\mathbb{Q}(\mu_m) = \mathbb{Q}(\mu_n)$ iff $m = n$ or (m is odd and $n = 2m$). c) Show:

$$\mathbb{Q}(\mu_m, \mu_n) = \mathbb{Q}(\mu_{\mathrm{lcm}(m,n)}). \tag{11}$$

$$\mathbb{Q}(\mu_m) \cap \mathbb{Q}(\mu_n) = \mathbb{Q}(\mu_{\gcd(m,n)}). \tag{12}$$

Theorem 9.9. *Let $n \in \mathbb{Z}^+$. There are infinitely many primes p with $p \equiv 1 \pmod{n}$.*

Proof. We may assume $n \geq 2$. Let S be a finite set (possibly empty) of primes $p \equiv 1 \pmod{n}$, and let $q = \prod_{p \in S} p$. For sufficiently large $k \in \mathbb{Z}$, we have

$$N = \Phi_n(knq) > 1.$$

Since the constant term of Φ_n is 1, for any prime $p \mid knq$, $N \equiv 1 \pmod{p}$. Since $N > 1$, there is a prime p with $\Phi_n(knq) = N \equiv 0 \pmod{p}$, so $p \nmid knq$: in particular $p \notin S$. By Theorem 9.6, $knq \in \mathbb{F}_p$ is a primitive nth root of unity. By Lagrange's Theorem, $n \mid p - 1$. We've produced a prime $p \notin S$ with $p \equiv 1 \pmod{n}$. □

Lemma 9.10. *Let G be a finite abelian group. Then there are $k, n \in \mathbb{Z}^+$ and a surjective homomorphism of groups $(\mathbb{Z}/n\mathbb{Z})^k \to G$.*

Exercise: Prove it.

Corollary 9.11. *For any finite abelian group G, there is a Galois extension $L/\mathbb{Q}$ with $\mathrm{Aut}(L/\mathbb{Q}) \cong G$.*

Proof. Step 1: By Lemma 9.10, G is a quotient of $(\mathbb{Z}/n\mathbb{Z})^k$ for some $k, n \in \mathbb{Z}^+$. Since any group which is a quotient of a finite Galois group over a field K is also a finite Galois group over that field, it suffices to treat the case $G = (\mathbb{Z}/n\mathbb{Z})^k$.
Step 2: By Theorem 9.9, there are prime numbers $p_1, \ldots, p_k$ such that $n \mid (p_i - 1)$

for $1 \leq i \leq k$. The group $(\mathbb{Z}/p_i\mathbb{Z})^\times$ is cyclic of order $\varphi(p_i) = p_i - 1$, so there is a surjection $q_i : (\mathbb{Z}/p_i\mathbb{Z})^\times \to \mathbb{Z}/n\mathbb{Z}$. Let

$$q = (q_1, \ldots, q_k) : \prod_{i=1}^k (\mathbb{Z}/p_i\mathbb{Z})^\times \to (\mathbb{Z}/n\mathbb{Z})^k,$$

a surjective group homomorphism. Put $N = p_1 \cdots p_k$. Since the p_i's are distinct, by the Chinese Remainder Theorem there is an isomorphism

$$\mathbb{Z}/N\mathbb{Z} \xrightarrow{\sim} \prod_{i=1}^k \mathbb{Z}/p_i\mathbb{Z}$$

and thus, passing to unit groups, an isomorphism

$$\Phi : (\mathbb{Z}/N\mathbb{Z})^\times \xrightarrow{\sim} \prod_{i=1}^k (\mathbb{Z}/p_i\mathbb{Z})^\times.$$

Thus we get a surjective map

$$\mathrm{Aut}(\mathbb{Q}(\mu_N)/\mathbb{Q}) \xrightarrow{\sim} (\mathbb{Z}/N\mathbb{Z})^\times \xrightarrow{\Phi} \prod_{i=1}^k (\mathbb{Z}/p_i\mathbb{Z})^\times \xrightarrow{q} (\mathbb{Z}/n\mathbb{Z})^k.$$

By Galois Theory, there is a subextension L of $\mathbb{Q}(\mu_N)/\mathbb{Q}$ with $\mathrm{Aut}(L/\mathbb{Q}) \cong (\mathbb{Z}/n\mathbb{Z})^k$. □

Exercise: Show: for any number field K and any finite abelian group G, there is a Galois extension L/K with $\mathrm{Aut}(L/K) \cong G$.

Exercise: Let $n \in \mathbb{Z}^+$.
a) (Parker: [Pa74]) Show: there is a number field $K \subset \mathbb{R}$ such that $K/\mathbb{Q}$ is Galois and $\mathrm{Aut}(K/\mathbb{Q}) \cong \mathbb{Z}/n\mathbb{Z}$.
b) Prove or disprove: for every finite abelian group G, there is a number field $K \subset \mathbb{R}$ such that $K/\mathbb{Q}$ is Galois and $\mathrm{Aut}(K/\mathbb{Q}) \cong G$.

9.2. **Cyclic Extensions I: Kummer Theory.**

A field extension K/F is **cyclic** if it is of finite degree and $\mathrm{Aut}(K/F)$ is a cyclic group of order $[K : F]$. In particular a cyclic extension is necessarily Galois. By a **generator** of a cyclic extension L/K, we mean an element σ which generates $\mathrm{Aut}(L/K)$. (Of course σ is not unique if $n > 2$.)

Example 9.12. *A quadratic extension K/F is cyclic iff it is separable. Thus if F does not have characteristic 2 then every quadratic extension K/F is cyclic, and moreover – as the quadratic formula holds here – is of the form $F(\sqrt{a})$ for some $a \in F \setminus F^2$.*

Let F be a field of characteristic 0. Since adjunction of square roots of elements of F yields cyclic extensions, it is natural to try to construct cyclic extensions of degree n by adjunction of nth roots. This is a good idea, but it works only under certain restrictions.

Example 9.13. *We revisit Example 4.17. For $n \geq 3$, let $p_n(t) = t^n - 2$, and let $F_n = \mathbb{Q}[t]/(p_n(t))$. We may embed $\mathbb{Q} \hookrightarrow \mathbb{R} \hookrightarrow \mathbb{C}$, and then, since p_n has a unique root $\sqrt[n]{2}$ in $\mathbb{R}$, and in such a way we view $F_n \hookrightarrow \mathbb{R}$. If ζ_n is a primitive nth root*

of unity, then the conjugates of $\sqrt[n]{2}$ over $\mathbb{Q}$ are $\zeta_n^i \sqrt[n]{2}$ for $0 \leq i < n$. The only conjugate that lies in $\mathbb{R}$, let alone F_n, is $\sqrt[n]{2}$, so $F_n/\mathbb{Q}$ is not normal (so certainly not cyclic). The splitting field of $F_n/\mathbb{Q}$ is

$$K_n := \mathbb{Q}(\zeta_n, \sqrt[n]{2}).$$

Because the subgroup $\mathrm{Aut}(K_n/F_n)$ of $\mathrm{Aut}(K_n/\mathbb{Q})$ is not normal, the group $\mathrm{Aut}(K_n/\mathbb{Q})$ is not commutative, hence certainly not cyclic.

Now let $n = 3$. Then the polynomial $p_3(t)$ remains irreducible over $\mathbb{Q}(\zeta_3)$: indeed, every irreducible cubic polynomial remains irreducible over a quadratic field extension, so $K_3/\mathbb{Q}(\zeta_3)$ is Galois of degree 3, hence cyclic. A generator for its automorphism group is the automorphism that sends $\sqrt[3]{2}$ to $\zeta_3 \sqrt[3]{2}$. We also compute in this way that the automorphism group $\mathrm{Aut}(K_3/\mathbb{Q})$ is noncommutative of order 6 and thus, as a permutation group on the conjugates $\sqrt[3]{2}, \zeta_3 \sqrt[3]{2}, \zeta_3^2 \sqrt[3]{2}$, is the full symmetric group S_3. Indeed, it has order $[K_3 : \mathbb{Q}(\zeta_3)][\mathbb{Q}(\zeta_3) : \mathbb{Q}] = 6$ and is noncommutative since the order 2 subgroup $\mathrm{Aut}(K_3/F_3)$ is not normal.

Proposition 9.14. *Let K be a field of characteristic $p \geq 0$, let $n \in \mathbb{Z}^+$, and let $a \in K$ be such that the polynomial $f(t) = t^n - a$ is irreducible in $K[t]$. Let $L := K[t]/(f(t)) = K(\sqrt[n]{a})$. The following are equivalent:*
(i) The extension L/K is cyclic.
(ii) The field K contains a primitive nth root of unity. (In particular, $p \nmid n$).

Proof. Note first that $f'(t) = nt^{n-1}$, so by the Derivative Criterion, L/K is separable iff $p \nmid n$. It follows that if $p \mid n$ then *neither* (i) nor (ii) holds, so we may assume henceforth that $p \nmid n$. In this case the roots of $f(t)$ in a splitting field are of the form $\zeta_n^i \sqrt[n]{a}$, where ζ_n is a primitive nth root of unity.
(i) $\implies$ (ii): In particular, $\frac{\zeta_n \sqrt[n]{a}}{\sqrt[n]{a}} = \zeta_n$ lies in any splitting field for f, so if L/K is normal then ζ_n lies in L.
(ii) $\implies$ (i): The above discussion shows that if K contains a primitive nth root of unity – say ζ_n – then L/K is normal and separable, thus Galois.

It remains to show that the group $\mathrm{Aut}(L/K)$ is cyclic. For this, observe that there is a unique $\sigma \in \mathrm{Aut}(L/K)$ such that $\sigma(\sqrt[n]{a}) = \zeta_n \sqrt[n]{a}$: such an automorphism exists because the automorphism group of a Galois extension $K[t]/(f)/K$ acts transitively on the roots of f, and it is unique because $L = K(\sqrt[n]{a})$. For any $i \in \mathbb{Z}^+$, $\sigma^i : \sqrt[n]{a} \mapsto \zeta_n^i \sqrt[n]{a}$, and thus the order of σ is

$$\langle \sigma \rangle = n = [L : K] = \#\,\mathrm{Aut}(L/K). \qquad \square$$

There is an important converse to Proposition 9.14. To prove it, we need first the following result, which despite its innocuous appearance is actually quite famous.

Lemma 9.15. *Let K be a field, $\zeta_n \in K$ a primitive nth root of unity. Let L/K be a cyclic extension of degree n, with generator σ. There is $\alpha \in L$ such that $\zeta_n = \frac{\sigma(\alpha)}{\alpha}$.*

Proof. Equivalently, we need to show that ζ_n is an eigenvalue for the K-linear endomorphism $\sigma : L \to L$. Since σ has order n, by Dedekind's Theorem the transformations $1, \sigma, \ldots, \sigma^{n-1}$ are all K-linearly independent, and therefore the minimal polynomial of σ is indeed $p(t) = t^n - 1$. Thus ζ_n is a root of the minimal polynomial for σ and therefore also a root of its characteristic polynomial. $\square$

Theorem 9.16. *(Kummer) Let $n \in \mathbb{Z}^+$, and let K be a field containing a primitive nth root of unity ζ_n. Let L/K be a degree n cyclic extension with generator σ.*

a) There exists $a \in K$ such that $\sigma(\sqrt[n]{a}) = \zeta_n \sqrt[n]{a}$ and $L = K(\sqrt[n]{a})$.
b) If $b \in K$ is such that $\sigma(\sqrt[n]{b}) = \zeta_n \sqrt[n]{b}$ and $L = K(\sqrt[n]{b})$, then $\frac{a}{b} \in K^n$.

Proof. a) By Lemma 9.15, there is $\alpha \in L$ such that $\sigma(\alpha) = \zeta_n \alpha$. Thus for all $i \in \mathbb{Z}^+$ $\sigma(\alpha^i) = \zeta_n^i \alpha$. In particular $a = \alpha^n \in K$, and the subgroup of $\langle \sigma \rangle = \mathrm{Aut}(L/K)$ fixing $K(\alpha)$ pointwise is the identity. It follows that $L = K(\alpha) = K(\sqrt[n]{a})$.
b) We have $\sigma \sqrt[n]{\frac{a}{b}} = \sqrt[n]{\frac{a}{b}}$, so $\frac{\sqrt[n]{a}}{\sqrt[n]{b}} = u \in K$. Take nth powers: $\frac{a}{b} = u^n \in K^n$. □

We continue with our study of cyclic extensions under the existence of sufficiently many roots of unity. Note that an important feature of the next result is that we analyze extensions of the form $K(\sqrt[n]{a})$ without the hypothesis that $t^n - a$ is irreducible in $K[t]$.

Proposition 9.17. *Let K be a field containing a primitive nth root of unity ζ_n, and let L/K be a field extension such that $L = K(\alpha)$ and $\alpha^n = a \in K$.*
a) L/K is a cyclic extension.
b) The degree $m = [L : K]$ is equal to the order of the image of a in $K^\times / K^{\times n}$.
c) There exists $b \in K$ such that the minimal polynomial of α over K is $t^m - b$.

Proof. □

Proposition 9.18. *Let K be a field containing a primitive nth root of unity ζ_n, and let $L = K(\sqrt[n]{a})$ for $a \in K$. Then any subextension M of L/K is of the form $K(\sqrt[m]{a})$ for some divisor m of n.*

If K is a field of characteristic not dividing n but not containing a primitive nth root of unity, there is in general no simple description of the degree n cyclic extensions of K. A lot of work has been done on special cases: for instance **global class field theory** gives a kind of description of all abelian extensions of a number field or function field in one variable over a finite field. Cyclic extensions have a distinguished role to play in this theory (e.g. via the **Hasse Norm Theorem**), but restricting class field theory to the cyclic case does not make it easier.

Perhaps surprisingly, the positive characteristic case is much more auspicious. If K is a field of characteristic $p > 0$, then none of the results of this section describe cyclic extensions of of K of order a power of p. But in fact there is a very satisfactory description of these extensions, due to Artin-Schreier and Witt.

9.3. The equation $t^n - a = 0$.

In this section we analyze the structure of the splitting field of a polynomial $t^n - a = 0$ *without* assuming that the ground field contains a primitive nth root of unity. We closely follow [LaFT, §VI.9].

Lemma 9.19. *Let F be a field of characteristic $p > 0$ and $a \in F^\times \setminus F^{\times p}$. Then for all $n \geq 1$, the polynomial $t^{p^n} - a$ is irreducible.*

Proof. We shall prove the contrapositive: suppose that for some $n \in \mathbb{Z}^+$ the polynomial $t^{p^n} - \alpha$ is reducible; we will show that α is a pth power in F. We may write $t^{p^n} - \alpha = f(t)g(t)$, where $f(t)$ and $g(t)$ are nonconstant monic polynomials. Let K/F be an extension field containing a root β of $t^{p^n} - \alpha$, so that in $K[t]$ we have

$$t^{p^n} - \alpha = t^{p^n} - \beta^{p^n} = (t - \beta)^{p^n}.$$

Since $K[t]$ is a UFD and $f(t)$ and $g(t)$ are monic, we therefore have $f(t) = (t-\beta)^r$ for some $0 < r < p^n$. Write $r = p^m s$ with $\gcd(p,s) = 1$. Note that $m < n$. Then

$$f(t) = (t^{p^m} - \beta^{p^m})^s,$$

so that the coefficient of $t^{p^m(s-1)}$ is $-s\beta^{p^m}$. This lies in F and – since $s \neq 0$ in F – we conclude $\beta^{p^m} \in F$. Thus

$$\alpha = (\beta^{p^m})^{p^{n-m}} \in F^{p^{n-m}} \in F^p$$

since $m < n$. □

Theorem 9.20. *Let $n \geq 2$, let F be a field, and let $a \in F^\times$. We suppose:*
- *For all prime numbers $p \mid n$, we have $a \notin F^p$, and*
- *If $4 \mid n$, then $a \notin -4F^4$.*

Then $f(t) \coloneqq t^n - a$ is irreducible in $F[t]$.

Proof. We begin by establishing several special cases.
Step 1: Suppose $n = p^e$ is a prime power, $a \in F \setminus F^p$ and p is the characteristic of F. This case is covered by Lemma 9.19.
Step 2: Suppose $n = p^e$ is a prime power, $a \in F \setminus F^p$ and p is *not* the characteristic of F. First we claim that $t^p - a$ is irreducible. Otherwise, there is some root $\alpha \in \overline{F}$ of $t^p - a$ such that $[F(\alpha) : F] = d < p$. Let N denote the norm map from $F(\alpha)$ to F: since $\alpha^p = a$, we have

$$N(\alpha)^p = N(a) = a^d.$$

Since $\gcd(d,p) = 1$, there are $x, y \in \mathbb{Z}$ such that $xd + yp = 1$, and thus

$$a = a^{xd}a^{up} = (N(\alpha)^x a^u)^p \in F,$$

contradiction. Now write

$$t^p - a = \prod_{i=1}^{p}(t - \alpha_i),$$

with $\alpha_1, \ldots, \alpha_p \in \overline{F}$ and $\alpha_1 = \alpha$. We may thus also write

$$t^{p^e} - a = \prod_{i=1}^{p}(t^{p^{e-1}} - \alpha_i).$$

Suppose first that $\alpha \notin F(\alpha)^p$. Let A be root of $t^{p^{e-1}} - \alpha$. If p is odd, then by induction A has degree p^{e-1} over $F(\alpha)$ and thus degree p^e over F and it follows that $t^{p^e} - a$ is irreducible. If $p = 2$, suppose $\alpha = -4\beta^4$ for some $\beta \in F(\alpha)$. Again let N be the norm from $F(\alpha)$ to F. Then

$$-a = N(\alpha) = 16N(\beta)^4,$$

so $-a \in F^2$. Since $p = 2$ it follows that $\sqrt{-1} \in F(\alpha)$ but $\alpha = (\sqrt{-1}2\beta^2)^2$, a contradiction. By induction, A has degree p^e over F. So we may assume that there is $\beta \in F(\alpha)$ such that $\beta^p = \alpha$. ...

□

The following is an immediate consequence.

Corollary 9.21. *Let p be a prime number, F a field, and $a \in F \setminus F^p$. If p is either odd or equal to the characteristic of F, then for all $n \in \mathbb{Z}^+$ the polynomial $t^{p^n} - a$ is irreducible in $F[t]$.*

Let F be a field. Let n be a positive integer that is not divisible by the characteristic of F, let $a \in F^\times$, and let K be the splitting field of the separable polynomial $p(t) = t^n - a$. We address the following question: what is the Galois group $G \coloneqq \mathrm{Aut}(K/F)$? Let α be a root of $p(t)$ in K, so $K = (\alpha, \zeta_n)$. Then an element $\sigma \in G$ is determined by its action on α and ζ_n, and we have

$$\sigma(\alpha) = \zeta^{b(\sigma)}\alpha,\ b(\sigma) \in \mathbb{Z}/n\mathbb{Z},$$

$$\sigma(\zeta_n) = \zeta_n^{d(\sigma)},\ d(\sigma) \in (\mathbb{Z}/n\mathbb{Z})^\times.$$

Consider the group

$$G(n) \coloneqq \left\{ \left[\begin{array}{cc} 1 & 0 \\ b & d \end{array} \right] \in \mathrm{GL}_2(\mathbb{Z}/n\mathbb{Z}) \right\}.$$

The identity

$$\left[\begin{array}{cc} 1 & 0 \\ 0 & d \end{array} \right] \left[\begin{array}{cc} 1 & 0 \\ b & 1 \end{array} \right] \left[\begin{array}{cc} 1 & 0 \\ 0 & d^{-1} \end{array} \right] = \left[\begin{array}{cc} 1 & 0 \\ bd & 1 \end{array} \right].$$

shows that the subgroup

$$N = \left\{ \left[\begin{array}{cc} 1 & 0 \\ b & 1 \end{array} \right] \in \mathrm{GL}_2(\mathbb{Z}/n\mathbb{Z}) \right\}$$

is normal. It also cyclic of order n, and it follows easily that

$$G(n) \cong \mathbb{Z}/n\mathbb{Z} \rtimes (\mathbb{Z}/n\mathbb{Z})^\times,$$

with the homomorphism given by the canonical isomorphism

$$\varphi : (\mathbb{Z}/n\mathbb{Z})^\times \to \mathrm{Aut}\, \mathbb{Z}/n\mathbb{Z}.$$

A straightforward computation shows that the commutator subgroup of $G(n)$ is contained in N; since $G(n)/N \cong (\mathbb{Z}/n\mathbb{Z})^\times$ is commutative, N must be the commutator subgroup of $G(n)$. The map $\sigma \mapsto d(\sigma)$ is precisely the mod n cyclotomic character, so $\zeta_n \in F \iff G \subset N$. In general, let $C_n \subset (\mathbb{Z}/n\mathbb{Z})^\times$ be the image of the cyclotomic character, viewed as a subgroup of diagonal matrices $\left[\begin{array}{cc} 1 & 0 \\ 0 & d \end{array} \right]$ as above. Then

$$G \subset \mathbb{Z}/n\mathbb{Z} \rtimes C_n.$$

On the other hand, if $p(t) = t^n - a$ is irreducible then K contains $F[t]/(p(t))$ hence $n \mid \#G$. So this gives us the answer in some cases.

Proposition 9.22. *Suppose $t^n - a$ is irreducible and $\gcd(n, \varphi(n)) = 1$. Then*

$$G \cong \mathbb{Z}/n\mathbb{Z} \rtimes C_n.$$

Proof. We know that G is a subgroup of $\mathbb{Z}/n\mathbb{Z} \rtimes C_n$, of corder $n\#C_n$. As above, irreducibility implies $n \mid \#G$. We also have $C_n \subset G$, so $\#C_n \mid \#G$. Since $\gcd(n, \varphi(n)) = 1$ and $\#C_n \mid \varphi(n) = 1$, also $\gcd(n, \#C_n) = 1$ and thus $n\#C_n \mid \#G$. It follows that $G = \mathbb{Z}/n\mathbb{Z} \rtimes C_n$. □

Theorem 9.23. *Let n be an odd positive integer prime to the characteristic of F, and suppose that $[F(\zeta_n) : F] = \varphi(n)$: equivalently, the mod n cyclotomic character is surjective. Let $a \in F$ be such that $a \in F \setminus F^p$ for all primes $p \nmid n$. Let K be the splitting field of $t^n - a$ over F, and let $G \coloneqq \mathrm{Aut}(K/F)$ be its Galois group. Then $G = G(n)$, and the commutator subgroup of G is $\mathrm{Aut}(K/F(\zeta_n))$.*

Proof. Noe first that since n is odd, by Theorem 9.20 the polynomial $t^n - a$ is irreducible in F. Let $\alpha \in K$ be a root, so $[F(\alpha) : F] = n$.
Step 1: Suppose $n = p$ is prime. Since $\gcd(p, \varphi(p)) = \gcd(p, p-1) = 1$, Proposition 9.22 applies to give $G = G(n)$. The commutator subgroup is N, which is precisely the set of automorphisms that pointwise fix ζ_n, so the commutator subgroup is $\mathrm{Aut}(K/F(\zeta_n))$. (This latter argument holds in the general case.)
Step 2: Now suppose that n is composite; we may write $n = pm$ with p prime. Since the mod n cyclotomic character is surjective and $m \mid n$, also the mod m cyclotomic character is surjective. Put $\beta \coloneqq \alpha^p$, so of course β is a root of $t^m - a$, and by induction the result applies to $t^m - a$. In particular we have

$$n = pm = [F(\alpha) : F] = [F(\alpha) : F(\beta)][F(\beta) : F],$$

so $[F(\alpha) : F(\beta)] = p$. This implies that $t^p - \beta$ is irreducible over $F(\beta)$: otherwise, the minimal polynomial of α over $F(\beta)$ would have degree less than p, contradiction. Consider the subfield

$$L \coloneqq F(\alpha) \cap F(\beta, \zeta_n) \subset K.$$

Certainly $F(\beta) \subset L$. On the other hand, $L/F(\beta)$ is an abelian extension. On the other hand, L is also the splitting field of $t^p - \beta$ over $F(\beta)$, so by Step 1, the maximal abelian subextension of $K/F(\beta)$ is $F(\beta, \zeta_p)$, and thus

$$L \subset F(\alpha) \cap F(\beta, \zeta_p) = F(\beta) :$$

if it were any larger, then $F(\alpha)$ would contain a nontrivial subextension of $F(\zeta_p)/F$, contradicting $[F(\zeta_n) : F] = \varphi(n)$. Thus

$$[F(\alpha, \zeta_n) : F(\beta, \zeta_n)] = p :$$

if not, then these fields would be equal and thus

$$F(\beta) \subset F(\alpha) \subset F(\beta, \zeta_n),$$

so $F(\alpha)/F(\beta)$ would be abelian, again contradicting Step 1. An argument identical to the above but using induction instead of Step 1 shows that

$$F(\zeta_n) \cap F(\beta) = F$$

and then using Natural Irrationalities we get

$$[F(\beta, \zeta_n) : F(\beta)] = [F(\zeta_n) : F] = \varphi_n.$$

It follows that

$$[K : F] = [K : F(\beta, \zeta_n)][F(\beta, \zeta_n) : F(\zeta_n)][F(\zeta_n) : F] = n\varphi(n) = \#G(n),$$

so $\mathrm{Aut}(K/F) = G_n$. The conclusion on commutator subgroups follows. □

Exercise 9.1. *a) Let $f(t) = x^8 - 2 \in \mathbb{Q}[t]$. Show: the splitting field of f is $\mathbb{Q}(\sqrt[8]{2}, \zeta_4)$.*
b) Observe that f satisfies all of the hypotheses of Theorem 9.23 except that 8 is not odd, and that the conclusion does not hold: $[K : F] = \frac{n\varphi(n)}{2}$, not $n\varphi(n)$.

We remark that the essential content content of Theorem 9.23 lies in the assertion that (under the hypotheses and notation used therein) $F(\zeta_n) \cap F(\alpha) = F$, since by Natural Irrationalities this implies that $[F(\zeta_n, \alpha) : F(\zeta)] = n$. It is also natural to think in terms of *linear disjointness* (cf. §12): because $F(\zeta_n)/F$ is Galois, the

identity $F(\zeta_n) \cap F(\alpha) = F$ holds iff $F(\zeta_n)$ and $F(\alpha)$ are linearly disjoint over F. Since this holds iff

$$F(\zeta_n) \otimes_F F(\alpha) = F(\zeta_n)[t]/(t^n - a)$$

is a field, another equivalent condition is that the polynomial $t^n - a$ remains irreducible over $F(\zeta_n)$. In the situation of the above exercise we have

$$\mathbb{Q}(\zeta_8) \cap \mathbb{Q}(\sqrt[8]{2}) = \mathbb{Q}(\sqrt{2})$$

and thus the polynomial $t^8 - 2$, which is irreducible over $\mathbb{Q}$, becomes reducible over $\mathbb{Q}(\zeta_8)$: indeed we have

$$t^8 - 2 = (t^4 - \sqrt{2})(t^4 + \sqrt{2}).$$

9.4. **Cyclic Extensions II: Artin-Schreier Theory.**

9.5. **Cyclic Extensions III: Witt's Theory.**

9.6. **Abelian Extensions of Exponent n: More Kummer Theory.**

9.7. **Solvable Extensions I: Simple Solvable Extensions.**

9.8. **Solvable Extensions II: Solvability by Radicals.**

10. Computing Galois Groups

11. Structure of Transcendental Extensions

11.1. **Transcendence Bases and Transcendence Degree.**

Let K/F be an extension. A finite set $S = \{x_1, \ldots, x_n\} \subset K$ is **algebraically independent** over F if for the only polynomial $P(t_1, \ldots, t_n) \in F[t_1, \ldots, t_n]$ such that $P(x_1, \ldots, x_n) = 0$ is $P = 0$. An arbitrary set $S \subset K$ is **algebraically independent** if all of its finite subsets are algebraically independent. (To be precise, we must impose some ordering on the elements of S in order to substitute them in as values of an n-variable polynomial, but the definition is obviously independent of the chosen ordering.) We say that K/F is **purely transcendental** if it is of the form $F(S)$ for some algebraically independent subset S of K.

Proposition 11.1. *Let K/F be an extension and $S = \{x_i\}$ be an ordered set of elements of K. TFAE:*
(i) The natural map $\Phi : F[\{t_i\}] \to K$ given by $t_i \mapsto x_i$ is an injection.
(ii) The map Φ extends uniquely to an isomorphism $F(\{t_i\}) \to F(S)$.
(iii) S is algebraically independent over F.

A subset S of K/F is a **transcendence basis** if it is algebraically independent and $K/F(S)$ is algebraic. In other words, a transcendence basis for K/F effects a decomposition of K/F into a tower $K/F(S)/F$ of a purely transcendental extension followed by an algebraic extension.

Example: The empty set is – perhaps by definition – always algebraically independent. If K/F is algebraic, then the only algebraically independent subset is the empty set, which is a transcendence basis.

Lemma 11.2. *Let K/F be an extension, $S \subset K$ be algebraically independent, and $x \in K$. Then $S \cup \{x\}$ is algebraically independent iff x is transcendental over $F(S)$.*

Proof. f S is an algebraically independent subset and $x \in K$ is transcendental over $F(S)$, then suppose for a contradiction that $S \cup \{x\}$ were dependent: i.e., there exists finite ordered subset $S_n = (x_1, \ldots, x_n)$ of S and a nonzero polynomial $P \in F[t_1, \ldots, t_n, t_{n+1}]$ such that $P(x_1, \ldots, x_n, x) = 0$. But the transcendence of x over $F(S)$ implies that the polynomial $P(x_1, \ldots, x_n, t_{n+1})$ is identically zero, so that the polynomial $Q(t_1, \ldots, t_n) := P(t_1, \ldots, t_n, 0)$ is not identically zero and $Q(x_1, \ldots, x_n) = 0$, contradicting the independence of $(x_1, \ldots, x_n)$. The other direction is even easier. □

Corollary 11.3. *a) An algebraically independent subset S of K is a transcendence basis iff it is not properly contained in any other algebraically independent set.*
b) Every algebraically independent subset of K is contained in a transcendence basis.

Proof. Part a) follows immediately from Lemma 11.2: a maximal algebraically independent set S is precisely one for which $K/F(S)$ is algebraic, i.e., a transcendence basis. Moreover the union of a chain of algebraically independent sets is algebraically independent, so part b) follows from part a) by Zorn's Lemma. □

Applying Corollary 11.3 to $S = \emptyset$, we deduce that every field extension K/F admits a transcendence basis.

Exercise X.X.X: Let $\{x_i\}_{i \in S}$ be a transcendence basis for the (nonalgebraic) field extension K/F. Let $n_\bullet : S \to \mathbb{Z}^+$ be any function. Show that $\{x_i^{n_i}\}$ is also a transcendence basis.

Definition: The **transcendence degree** of a field extension K/F is the minimum cardinality of a transcendence basis. We defer the obvious question – can there exist two transcendence bases of different cardinalities? – until §$X.X$.

The transcendence degree of an extension is related to $\#K$ and $\#F$ as follows:

Proposition 11.4. *Let K/F be a transcendental field extension, with transcendence degree κ. Then*

$$\#K = \max(\#F, \kappa, \aleph_0).$$

Proof. Since K/F is transcendental, K is infinite. Moreover, κ and $\#F$ are cardinalities of subsets of K, so clearly $\#K \geq \max(\#F, \kappa, \aleph_0)$. Conversely, let S be a transcendence basis; then $F(S)$ has cardinality $\max(\#, \kappa)$ and $K/F(S)$ is algebraic and $F(S)$ is infinite, so $\#K = \#F(S)$. □

11.2. **Applications to Algebraically Closed Fields.**

Theorem 11.5. *(Automorphism extension theorem) Let K be an extension of F, with K algebraically closed. Then every automorphism of F can be extended to at least one automorphism of K.*

Proof. Let $\{x_i\}_{i \in S}$ be a transcendence basis for K/F. There is a unique automorphism of $F(S)$ which extends ι and maps each x_i to itself. Since K is the algebraic closure of $F(S)$, by Corollary XX we can further extend to an automorphism of K. □

For any field K, let $\mathbb{F}$ be its prime subfield. An **absolute transcendence basis** for K is a transcendence basis for K/F.

Corollary 11.6. *a) Two algebraically closed fields K_1 and K_2 are isomorphic iff they have the same characteristic and the same absolute transcendence degree.*
b) Suppose K_1, K_2 are two algebraically closed fields of the same characteristic and $\#K_1 = \#K_2$ is uncountable. Then $K_1 \cong K_2$.

Proof. Evidently any pair of isomorphic fields $K_1 \cong K_2$ have the same characteristic and absolute transcendence degree. If K_1 is algebraically closed with prime subfield $\mathbb{F}$ and transcendence degree κ, then for a set S of indeterminates of cardinality κ, then K_1 is isomorphic to the algebraic closure of $\mathbb{F}(S)$, which shows that the characteristic and the absolute transcendence degree determine the isomorphism class of an algebraically closed field. Proposition **??** implies that the absolute transcendence degree of any uncountable field is equal to its cardinality, and part b) then follows immediately from part a). □

Remark: The fact that any two algebraically closed fields of given cardinality and, say, continuum cardinality, are isomorphic has important applications in model theory: via the Tarski-Vaught test, it shows that the first order theory of algebraically closed fields of a given characteristic is **complete**.

Theorem 11.7. *Let K/F be an extension of fields, of transcendence degree κ. TFAE:*
(i) For any extension field K' of F with transcendence degree $\kappa' \leq \kappa$, there exists an F-algebra embedding $K' \hookrightarrow K$.
(ii) K is algebraically closed.

Exercise X.X.X: Prove Theorem X.X.

Theorem 11.8. *Let K be an algebraically closed field. The group $\mathrm{Aut}(K)$ of all automorphisms of K has cardinality $2^{\#K}$.*

Proof. Step 0: Note that $2^{\#K}$ is also the cardinality of the set of all functions from K to K, so is the largest conceivable value of $\#\mathrm{Aut}(K)$.

Step 1: We must check the result for $\overline{\mathbb{F}_p}$ and $\overline{\mathbb{Q}}$. In the former case we have identified the automorphism group as $\hat{\mathbb{Z}}$, which indeed has cardinality $c = 2^{\aleph_0} = 2^{\#\overline{\mathbb{F}_p}}$. In the latter case we can by no means "identify" $\mathrm{Aut}(\overline{\mathbb{Q}})$, but to see that it has continuum cardinality it suffices, by the automorphism extension theorem, to exhibit a simpler Galois extension $K/\mathbb{Q}$ which has continuum cardinality. Indeed one can take K to be quadratic closure of $\mathbb{Q}$, i.e., the compositum of all quadratic field extensions of $\mathbb{Q}$. The automorphism group here is $(\mathbb{Z}/2\mathbb{Z})^{\aleph_0} = c$.

Step 2: By the automorphism extension theorem, the cardinality of the automorphism group of any algebraically closed field is at least that of the continuum, which by Step 0 gives the answer for all countable fields, i.e., for all fields of countable absolute transcendence degree.

Step 3: Otherwise K is uncountable so there exists an absolute transcendence basis S with $\#S = \#K$. Now the natural action of $\mathrm{Sym}(S)$ on S gives rise to an injection $\mathrm{Sym}(S) \hookrightarrow \mathrm{Aut}(\mathbb{F}(S))$, i.e., by permutation of indeterminates. By the automorphism extension theorem, this shows that $\#\mathrm{Aut}(K) \geq \#\mathrm{Sym}(S) = 2^{\#S}$. □

Corollary 11.9. *Suppose K/F is an extension with K algebraically closed. Then $K^{\mathrm{Aut}(K/F)}$ is the purely inseparable closure of F in K. In particular, $K^{\mathrm{Aut}(K/F)} = F$ iff F is perfect.*

Proof. If x lies in the purely inseparable closure of F in K, then for some $e \in \mathbb{Z}^+$, $x^{p^e} \in F$. Since x has no Galois conjugates, we must have $\sigma(x) = x$ for every $\sigma \in \mathrm{Aut}(K/F)$. Let $\overline{F}$ be the algebraic closure of F in K. By the usual Galois theory we have $\overline{F}^{\mathrm{Aut}(\overline{F}/F)}$ is the purely inseparable closure of F in $\overline{F}$, and by the automorphism extension theorem we conclude that $K^{\mathrm{Aut}(K/F)} \cap \overline{F}$ is the purely inseparable closure of F in K. If $x \in K$ is transcendental over F, then by Theorem X.X.X there exists an ordered transcendence basis $S = (x, \{x_\alpha\})$ containing x. By Exercise X.X.X, $S' = (x^2, \{x_\alpha\})$ is also a transcendence basis hence there exists an automorphism $F(S) \to F(S')$ sending $x \mapsto x^2$, which, as usual, extends to an F-algebra automorphism σ of K with $\sigma(x) = x^2 \neq x$. □

Another fact which is true about automorphism groups of algebraically closed field extensions K/F is that any bijection φ between algebraically independent subsets I and I' of K extends to an F-automorphism of F. For this it is necessary and sufficient that φ extend to a bijection on transcendence bases $S \supset I$, $S' \supset I'$. A moment's thought shows that this holds provided that all transcendence bases of K/F have the same cardinality and need not hold otherwise. This brings us to the next section.

11.3. **An Axiomatic Approach to Independence.**

We wish to prove the following result.

Theorem 11.10. *Let K/F be a field extension. Then any two transcendence bases for K/F have the same cardinality, so that the transcendence degree of K/F is the cardinality of* any *transcendence basis.*

Of course this is strikingly similar to the situation in ordinary linear algebra. We could therefore go back to our linear algebra texts, consult the proof of the cardinality independence of bases in vector spaces, and attempt to mimic it in the present context. This approach will succeed. Of course in order to do this we will have to find some sort of precise analogy between linear independence and algebraic independence. In mathematics, once we determine that situations A and B are analogous (to the extent that certain proofs can be carried over from one context to the other), do we just dutifully copy down the similar proofs and keep the analogy in the back of our mind in case we need it later? Depending on taste, this is a reasonable approach to take, perhaps more reasonable for the mind which is able to quickly remember what it once knew. As for myself, I would at the same time worry that it would take me some time and energy to recreate the analogy if I hadn't written it down, and I would also be curious whether A and B might be common instances of a more general construction that it might be interesting or useful to know explicitly. So we shall follow the second course here, with apologies to those with different tastes.

Let us begin by placing alongside the analogies between linear independence of a subset S of an F-vector space V and algebraic independence of a subset S of an F-algebra K.

In both contexts we have a set, say X, and a collection of subsets S of X that we are calling **independent**, subject to:

(LI1) The empty set is independent.
(LI2) A set is independent iff all its finite subsets are independent.
(LI3) Any subset of an independent set is independent.

Notice that it follows from (LI2) and (LI3) that the union $S = \bigcup_i S_i$ of any chain of independent subsets is independent: if not, there would exist a finite dependent subset S' of S, but S' would have to be a subset of some S_i, contradicting the independence of S_i. Combining this with (LI1) and aplying Zorn's Lemma, we get

(A) Maximal independent sets exist, and every independent set is contained in some maximal independent set.

Could it be that (LI1) through (LI3) imply the following desirable property?

(B) All maximal independent sets have the same cardinality.

Unfortunately this is not the case. Suppose we have a set X which is partitioned into disjoint subsets:

$$X = \coprod_i X_i.$$

Call a subset $S \subset X$ independent iff it is contained in X_i for some i. Then (LI1) through (LI3) are satisfied and the maximal independent sets are simply the X_i's, which we are evidently not entitled to conclude have the same cardinality.

So we need another axiom. Consider the following:

(LI4) If S_1 and S_2 are independent subsets of X with $\#S_1 < \#S_2$, then there exists $x \in X \setminus S_1$ such that $S_1 \bigcup \{x\}$ is independent.

A set X equipped with a family of subsets $\{S_i\}$ satisfying axioms (LI1) through (LI4) is called an **independence space**.

In an independence space, if S_1 and S_2 are independent sets with $\#S_1 < \#S_2$, then S_1 is non-maximal. Therefore a maximal independent set has cardinality at least as large as any other independent set, so by symmetry all maximal independent sets have the same cardinality: independence spaces satisfy (B). Conversely, (LI1) through (LI3) and (B) clearly imply (LI4).

In this new language, Theorem 11.10 takes the form

Theorem 11.11. *If K/F is a field extension, then the collection of algebraically independent subsets of K is an independence space.*

Unfortunately it is not so obvious how to show that the collection of algebraically independent subsets of K satisfies (LI4). So let us try a different approach, in terms

of something called spanning sets. We notice that to each subset S of a vector space its linear span $\overline{S}$ gives an abstract closure operator: namely we have

(CL1) $S \subset \overline{S}$
(CL2) $S \subset S' \implies \overline{S} \subset \overline{S'}$
(CL3) $\overline{\overline{S}} = \overline{S}$.

But the linear span satisfies two other properties, the first of which is not surprising in view of what has come before:

(SO4) if $x \in \overline{S}$, there exists a finite subset $S' \subset S$ such that $x \in \overline{S'}$.

Famously, linear span also satisfies the following **Exchange Lemma**:[12]

(SO5) If $y \in \overline{S \cup x}$ and y is *not* in $\overline{S}$, then $x \in \overline{S \cup y}$.

(Proof: If $y \in \overline{S \cup x}$, there exist $s_1, \ldots, s_n \in S$ and scalars $a_1, \ldots, a_n, a$ such that $y = a_1 s_1 + \ldots + a_n s_n + ax$. If y is not in the span of S, then $a \neq 0$, so $x = y - \frac{-a_1}{a} s_1 + \ldots + \frac{-a_n}{a} s_n \in \overline{S \cup y}$.)

Now, suppose K/F is a field extension and S is a subset of K. We will define $\overline{S}$ to be the algebraic closure of $F(S)$ in K. It is immediate that this "algebraic closure" operator satisfies (SO1) through (SO4). Let us check that it also satisfies (SO5): suppose $y \in \overline{S \cup x}$ and y is not in the algebraic closure of S. Then there exists a finite subset $x_1, \ldots, x_n$ of S such that y is algebraic over $F(x_1, \ldots, x_n, x)$: i.e., there exists a polynomial $f(t_1, \ldots, t_n, t_{n+1}, t_{n+2})$ with F-coefficients such that $f(x_1, \ldots, x_n, x, t_{n+2}) \neq 0$ and $f(x_1, \ldots, x_n, x, y) = 0$. Writing

$$f(x_1, \ldots, x_n, t_{n+1}, t_{n+1}) = \sum_{i=0}^{g} A_i(x_1, \ldots, x_n, t_{n+2}) t_{n+1}^i,$$

observe that not all the polynomials $A_i(x_1, \ldots, x_n, t_{n+2})$ can be zero. Since y is not algebraic over $F(S)$, it follows that not all of the elements $A(x_1, \ldots, x_n, y)$ are zero, and therefore $f(x_1, \ldots, x_n, t_{n+1}, t_{n+1}, y) \neq 0$. Since $f(x_1, \ldots, x_n, x, y) = 0$, it follows that x is algebraic over $F(S, y)$ as asserted.

Suppose again that X is any set equipped with a **spanning operator** $S \mapsto \overline{S}$, i.e., an operator satisfying the three closure axioms (CL1) through (CL3) and also (CL4) and (CL5). A subset S of X is a **spanning set** if $\overline{S} = X$. A subset S of X is **independent** if for all $s \in S$, s is not in $\overline{S \setminus s}$. A **basis** is an independent spanning set.

Note that it is immediate to show that the independent sets for a spanning operator satisfy (LI1) through (LI3). In particular, we have (A), that bases exist and any independent set is contained in a basis. Again it is not obvious that (LI4) is satisfied. Rather we will show (B) directly – which is what we really want anyway

[12]This is an absolutely prototypical example of a *lemma*: the exchange lemma is the essential kernel of content in the theory of linearly independence, and yet it is itself not very memorable or appealing, so is doomed to be overshadowed by the figurehead theorems that it easily implies.

– and by the above remarks that implies (LI4).

In the following results X is always a set equipped with a spanning operator $S \mapsto \overline{S}$.

Proposition 11.12. *For a subset $S \subset X$, TFAE:*
(i) S is a minimal spanning set of X.
(ii) S is a maximal independent set of X.
(iii) S is a basis.

Proof. (This is the usual thing.) (i) $\implies$ (iii): Suppose S is minimal spanning but not dependent; then by definition there exists $s \in S$ such that $x \in \overline{S \setminus s}$), so that $\overline{S \setminus s}$, being a closed set containing S, also contains the closure of S, i.e., X, and we found a smaller spanning set. (iii) $\implies$ (ii): if S is a basis and $S \cup \{x\}$ is independent then x does not lie in $\overline{S}$ which is absurd since S is a spanning set. (ii) $\implies$ (i) is similar: if S were a maximal independent set but not a spanning set, then there exists $x \in X \setminus \overline{S}$ and then $S \cup \{x\}$ is independent. □

Theorem 11.13. *Let S be an independent subset of X and T a spanning set. There exists a subset $T' \subset T$ such that $S \cup T'$ is a basis and $S \cap T' = \emptyset$.*

Proof. Let $\mathcal{I}$ be the collection of all subsets T' of T such that $S \cap T' = \emptyset$ and $S \cup T'$ is indepdendent. Observe that $\emptyset \in \mathcal{I}$, so $\mathcal{I}$ is not itself empty. As usual, $\mathcal{I}$ is closed under unions of increasing chains so by Zorn's Lemma has a maximal element T'. Let $x \in T$, and suppose that x is not in $\overline{S \cup T'}$. Then $T'' := T' \cup \{x\}$ is a strictly larger subset of T such that $S \cup T''$ is still independent, contradicting the maximality of T'. Therefore

$$X = \overline{T} \supset \overline{\overline{S \cup T'}} = \overline{S \cup T'},$$

so $S \cup T'$ is a basis. □

Corollary 11.14. *If X admits a finite spanning set, it admits a finite basis.*

Proof. Apply Theorem 11.13 with $S = \emptyset$. □

Theorem 11.15. *Any two bases B, B' of X have the same cardinality.*

Proof. Case 1: Suppose $B = \{x_1, \ldots, x_n\}$ is a finite basis, and let B' be any other basis. Let $m = \#B \cap B'$. If $m = n$ then $B \subset B'$ and by Proposition 11.12 distinct bases are at least incomparable, so $B = B'$. So suppose (WLOG) that $B \cap B' = \{x_1, \ldots, x_m\}$ with $m < n$. The set $B \setminus x_{m+1}$ cannot be a spanning set, whereas B' is, so there exists $y \in B' \overline{B \setminus x_{m+1}}$. The set $B_1 := (B \setminus x_{m+1}) \cup y$ is independent. By the Exchange Lemma (SO5), $x_{m+1} \in \overline{(B_1)}$. Hence $B \subset \overline{B_1}$, and since B is a spanning set, so is B_1. Thus B_1 is a basis. Notice that B_1 has n elements and also $\{x_1, \ldots, x_m, y\} \subset B_1 \cap B'$, so that we have replaced B by another basis of the same cardinality and sharing at least one more element with B'. Repeating this procedure will produce a finite sequence of bases B_2, B_3, each of cardinality n, such that the last basis B_k is contained in, and thus equal to, B'.

Case 2: We may now suppose that B and B' are both infinite. For every $x \in X$, we claim the existence of a subset E_x with the property that $x \in \overline{E_x}$ and for any subset E of B such that $x \in \overline{E}$, $E_x \subset E$. Assuming the claim for the moment, we complete the proof. Consider the subset $S = \bigcup_{x \in B'} E_x$ of B. Since each E_x

is finite, $\#S \leq \#B'$. On the other hand, for all $x \in B'$, $x \in \overline{E_x} \subset \overline{S}$, so $B' \subset \overline{S}$ and therefore $\overline{S} \supset \overline{B'} = X$. Therefore S is a spanning subset of the basis B, so $S = B$ and thus $\#B \leq \#B'$. By reversing the roles of B and B' in the argument we conclude $\#B = \#B'$.

It remains to prove the claim on the existence of E_x. In turn we claim that if E' and E'' are two subsets of B such that $x \in \overline{E'} \cap \overline{E''}$ and x is not in the span of any proper subset of E', then $E' \subset E''$; this certainly suffices. Assuming to the contrary that there exists $y \in E' \setminus E''$. Then x is not in the span of $E' \setminus y$ and is in the span of $(E' \setminus y) \cup y$, so by (SO5) y is in the span of $(E' \setminus y) \cup x$. Since x is in the span of E'', we get that y is in the span of $(E' \setminus y) \cup E''$. But this contradicts the fact that the $(E' \setminus y)cupE'' \cup \{y\}$, being a subset of B, is independent. □

Remark: A set X endowed with a spanning operator as above is often called a **finitary matroid**. (The word "finitary" refers to (SO4).) Combinatoricists are especially interested in finite matroids, which includes the class of finite-dimensional vector spaces over finite fields but not that of independent subsets of a field extension (except in the trivial case of an algebraic field extension).

For future reference, for a field extension L/K, we will refer to the matroid with sets the subsets of L, spanning operator $S \mapsto \overline{S}$ the algebraic closure of $K(S)$ in L and (it follows) with independent sets the algebraically independent subsets the **transcendence matroid of L/K**.

We saw above how to go from a finitary matroid to an independence space, namely by decreeing a subset $S \subset X$ to be dependent if there exists $x \in S$ such that $x \in \overline{S \setminus x}$. Conversely, to every independence space we can associate a finitary matroid: define the span $\overline{Y}$ of a subset Y to be the set of $x \in X$ such that $S \cup x$ is dependent. This complete equivalence between concepts of linear independence and spanning seems a bit unexpected, even in the context of vector spaces.

For finite matroids, combinatorialists know at least half a dozen other equivalent axiomatic systems: e.g. in terms of graphs, circuits, "flat" subspaces and projective geometry. As above, demonstrating the equivalence of any two of these systems is not as easy as one might expect. This phenomenon of multiple nonobviously equivalent axiomatizations has been referred to, especially by G. Rota, as **cryptomorphism**. Of course every twenty-first century student of mathematics has encountered crytomorphism (although it seems that the multiplicity is especially large for finite matroids!). In several essays, Rota saw cryptomorphism as a warning not to take any particular axiomatization of a theory or structure too seriously. This seems fair, but since the different axiomatizations can lead to different and possibly easier proofs, perhaps it should also be viewed as an instance of the inherent richness of mathematical concepts.

11.4. **More on Transcendence Degrees.**

Proposition 11.16. *Let L/K be a field extension and T a subset of L such that $L = K(T)$. Then* $\mathrm{trdeg}(L/K) \leq \#T$.

Proof. In the transcendence matroid of L/K, T is a spanning set. According to Theorem 11.13 with $S = \varnothing$, some subset T' of T is a basis for the matroid, i.e., a

transcendence basis for L/K. Thus

$$\mathrm{trdeg}(L/K) = \#T' \leq \#T.$$

□

Theorem 11.17. *Let $F \subset K \subset L$ be a tower of field extensions.*
a) If $\{x_i\}_{i \in I}$ is a transcendence basis for K/F and $\{y_j\}_{j \in J}$ is a transcendence basis for L/K, then $\{x_i, y_j\}$ is a transcendence basis for L/F.
b) We have $\mathrm{trdeg}(L/F) = \mathrm{trdeg}(L/K) + \mathrm{trdeg}(K/F)$.

Proof. a) We first show that $\{x_i, y_j\}$ is an algebraically independent set. Choose any finite subsets of $\{x_i\}$ and $\{y_j\}$: for ease of notation, we rename the elements $x_1, \ldots, x_m, y_1, \ldots, y_n$. Suppose there exists a polynomial $P \in F[t_1, \ldots, t_{m+n}]$ such that $P(x_1, \ldots, x_m, y_1, \ldots, y_n) = 0$. Put $Q(t_1, \ldots, t_n) = P(x_1, \ldots, x_m, t_1, \ldots, t_n) \in K[t]$. Then $Q(y_1, \ldots, y_n) = 0$ implies $Q(t_1, \ldots, t_n) = 0$. Each coefficient of this polynomial is a polynomial expression in $x_1, \ldots, x_m$ with F-coefficients, and the algebraic independence of the x_i's implies that each of these coefficients is equal to 0. Thus $P = 0$. Let $K_0 = F(\{x_i\})$, so K/K_0 is algebraic. Let $L_0 = K(\{y_j\}$, so L/L_0 is algebraic. Let $z \in L$. Then z satisfies a polynomial equation with coefficients in L_0. Since K/K_0 is algebraic, z also satisfies a polynomial equation with coefficients in $K_0(\{y_j\}) = F(\{x_i, y_j\})$.
b) By part a), $\{x_i, y_j\}$ is a transcendence basis for L/F, of cardinality $\#I + \#J = \mathrm{trdeg}(K/F) + \mathrm{trdeg}(L/K)$. □

Exercise: Let M/F be a field extension, and let K, L be subextensions of M/F. Suppose K/F is finite and L/F is purely transcendental. Show $[LK : L] = [K : F]$. (Suggestion: reduce to the case $K = F[t]/(p(t))$ and $L = F(t)$. An idea for this case is that if the polynomial $p(t)$ factors over $F(t)$, then by taking $t = a$ for $a \in F$ we get a factorization over F. One has to be a little careful here in order to avoid values a which make the denominator of one of the rational functions equal to 0.)[13]

Theorem 11.18. *For $F \subset K \subset L$ be a tower of field extensions, TFAE:*
(i) K/F and L/K are both finitely generated.
(ii) L/F is finitely generated.

Proof. (i) $\implies$ (ii): If $K = F(x_1, \ldots, x_m)$ and $L = K(y_1, \ldots, y_n)$, then $L = F(x_1, \ldots, x_m, y_1, \ldots, y_n)$.
(ii) $\implies$ (i): It is immediate that if L/F is finitely generated then so is L/K for any subextension K of L/F: any finite generating set for L/F is also a finite generating set for L/K. Let $z_1, \ldots, z_e$ be a transcendence basis for K/F. Then $F(z_1, \ldots, z_e)/F$ is finitely generated, so it suffices to show that the algebraic extension $K/F(z_1, \ldots, z_e)$ is finitely generated. Moreover, $L/F(z_1, \ldots, z_e)$ is finitely generated, so it is enough to prove the result with $F(z_1, \ldots, z_e)$ in place of F and thus we may assume that K/F is algebraic.

We are thus reduced to showing: if $L/K(t_1, \ldots, t_n)$ is a finite extension of a rational function field and K/F is an algebraic extension, then L/F finitely generated implies K/F finitely generated – or, equivalently since K/F is algebraic – that K/F is finite. But suppose not: then for all $d \in \mathbb{Z}^+$ there exists a subextension K_d of K/F such that $[K_d : F] \geq d$. By the preceding exercise we have

[13]This result will become much more clear following our later discussion of **linear disjointness**. The reader may prefer to defer the exercise until then.

$[K_d(t_1,\dots,t_n):F(t_1,\dots,t_n)] = [K_d:F] \geq n$. Thus $L/F(t_1,\dots,t_n)$ is an algebraic extension but

$$[L:F(t_1,\dots,t_n)] \geq [K(t_1,\dots,t_n):F(t_1,\dots,t_n) \geq \aleph_0,$$

so it is algebraic of infinite degree, hence not finitely generated: contradiction! □

Exercise: Let k be any field. Consider the polynomial ring $R = k[x,y]$: note that it is finitely generated as a k-algebra. Show that there is a k-subalgebra of R which is *not* finitely generated. (Thus Theorem 11.18 exhibits a property of field extensions without analogue in the study of commutative rings.)

12. Linear Disjointness

12.1. **Definition and First Properties.**

Let E/F be a field extension, and let R, S be F-subalgebras of E. We say that R and S are **F-linearly disjoint in E** if the canonical map $R \otimes_F S \to E$ is injective. (If the *ambient field* E is understood, we will just say that R, S are F-linearly disjoint, or that they are **linearly disjoint over F**. In fact the depenence on E is often suppressed, for reasons that will be explored soon enough.)

Lemma 12.1. *Let E/F be a field extension, and let K, L be subextensions of finite degree over F. Then K and L are linearly disjoint over F iff*
$[KL:F] = [K:F][L:F]$.

Proof. The canonical map $\tau : K \otimes_F L \to KL$ is always surjective. Since its source and target are both finite-dimensional F-vector spaces, τ is injective iff

$$[K:F][L:F] = \dim_F K \otimes_F L = \dim_F LK. \qquad \square$$

Exercise 12.1. *Let K, L be finite degree extensions of a field F of coprime degrees. Show: K, L are F-linearly disjoint.*

Lemma 12.2. *If R, S are F-linearly disjoint in E, then $R \cap S = F$.*

Proof. By contraposition: suppose there exists $u \in (R \cap S) \setminus F$. We may then choose F-bases A of R and B of S such that $\{1, u\} \subset A \cap B$. The elements $1 \otimes u$ and $u \otimes 1$ are then F-linearly independent in $R \otimes_F S$ but under $\iota : R \otimes_F S \to E$ they both get mapped to u, so ι is not injective. □

Exercise 12.2. *a) Let $F = \mathbb{Q}$ and $E = \mathbb{C}$. Show that $K = \mathbb{Q}(\sqrt[3]{2})$ and $L = \mathbb{Q}(e^{\frac{2\pi i}{3}}\sqrt[3]{2})$ are not linearly disjoint over F, even though $K \cap L = F$.*
b) Try to generalize the result of part a), for instance as follows: if K/F is algebraic and not *normal, then inside any algebraic closure E of K there exists a field extension L/F such that $K \cap L = F$ but K, L are not F-linearly disjoint in E.*

Exercise 12.3. *Let R, S be F-subalgebras of E/F. Show: the following are equivalent:*
(i) R and S are linearly disjoint over F.
(ii) For all F-linearly independent subsets $\{a_i\}_{i\in I}$ of R and $\{b_j\}_{j\in J}$ of S, $\{a_ib_j\}_{(i,j)\in I\times J}$ is F-linearly independent in E.
(iii) For all positive integers m and n, if $a_1,\dots,a_m$ are F-linearly independent in R and $b_1,\dots,b_n$ are F-linearly independent in S, then $a_1b_1,\dots,a_m,b_1,a_2b_1,\dots,a_mb_n$ are F-linearly indpendent in E.

Exercise 12.4. *(Linear disjointness is preserved by direct limits) Let R be an F-subalgebra of E/F. Suppose $R = \varinjlim R_i$ is a direct limit of a family $\{R_i\}_{i \in I}$ of F-subalgebras. Show: for any F-subalgebra S of E/F, R and S are linearly disjoint iff for all $i \in I$, R_i and S are linearly disjoint.*

Exercise 12.5. *Suppose R, S are linearly disjoint subalgebras of E/F. Let $R' \subset R$ and $S' \subset S$ be F-subalgebras. Show: R' and S' are linearly disjoint over F.*

Lemma 12.3. *Two subalgebras R and S of E/F are linearly disjoint over F iff the subfields they generate, say K and L, are linearly disjoint over F.*

Proof. Suppose that R and S are linearly disjoint over F. It is enough to show that if $k_1, \ldots, k_m$ are F-linearly independent elements of K and $l_1, \ldots, l_n$ are F-linearly independent elements of L, then $\{k_i l_j\}_{1 \leq i \leq m, 1 \leq j \leq n}$ are F-linearly independent in E. There exist $a, a_1, \ldots, a_m \in R$ such that $k_i = \frac{a_i}{a}$ for all i, and simiarly there exist $b, b_1, \ldots, b_n \in S$ such that $l_j = \frac{b_j}{b}$ for all j. Then if $\alpha_{ij} \in F$ is such that $\sum_{i,j} \alpha_{ij} \frac{a_i b_j}{ab} = 0$, then multiplying by ab gives $\sum_{i,j} \alpha_{ij} a_i b_j = 0$, and by assumption $\alpha_{ij} = 0$ for all i and j.
The converse is immediate from Exercise 12.5. □

Thus it is no loss of generality to speak of linear disjointness of sub*fields* of E/F, but it is often convenient to phrase things in terms of subdomains of these fields.

Proposition 12.4. *Let K, L be subextensions of a field extension L/F. TFAE:*
(i) K and L are linearly disjoint over F.
(ii) Every F-linearly independent subset S of K is L-linearly independent in E.
(ii′) Every F-linearly independent subset T of L is K-linearly independent in E.
(iii) There is an F-basis A of K which is L-linearly independent as a subset of E.
(iii′) There is an F-basis B of L which is K-linearly independent as a subset of E.

Proof. (i) $\implies$ (ii): Let A be F-linearly independent in K. Consider any finite subset of elements of A, say $k_1, \ldots, k_n$, and let $\beta_1, \ldots, \beta_n \in L$ be such that

$$\beta_1 k_1 + \ldots + \beta_n k_n = 0. \tag{13}$$

Choose an F-basis $\{l_j\}_{j \in J}$ for L, so that there are unique $\alpha_{ij} \in F$ such that for all i, $\beta_i = \sum_j \alpha_{ij} l_j$. Substituting this into (13) gives

$$\sum_{i,j} \alpha_{ij} k_i l_j = 0.$$

By Exercise 12.3 this forces $\alpha_{ij} = 0$ for all i, j and thus $\beta_j = 0$ for all k, so the k_i's are L-linearly independent.
(i) $\implies$ (ii′): The above proof works with the roles of K and L reversed.
(ii) $\implies$ (i): By Exercise 12.3, it is enough to fix $m, n \in \mathbb{Z}^+$ let $k_1, \ldots, k_m$ be F-linearly independent elements of K and $l_1, \ldots, l_n$ be F-linearly independent elements of L and show that $\{k_i l_j\}$ are F-linearly independent elements of E. Suppose that $\alpha_{ij} \in F$ are such that $\sum_{i,j} \alpha_{ij} k_i l_j = 0$. But we may rewrite this as

$$(\alpha_{11} l_1 + \ldots + \alpha_{1n} l_n) k_1 + \ldots + (\alpha_{m1} l_1 + \ldots + \alpha_{mn} l_n) k_m = 0.$$

By hypothesis the k_i's are L-linearly independent, so this forces all the coefficients of the above equation to be equal to zero, which in turn, since the l_j's are F-linearly independent, forces all the α_{ij}'s to be zero.
(ii′) $\implies$ (i) in the same way.

(ii) $\implies$ (iii) and (ii$'$) $\implies$ (iii$'$) are immediate.
(iii) $\implies$ (ii): Let S be an F-linearly independent subset of K, and complete it to a basis A' of K. Let $\varphi : A' \to A$ be a bijection and Φ the induced F-linear automorphism of K. Suppose that A' is not L-linearly independent, i.e., there exists a finite subset $a'_1, \ldots, a'_n$ of A' and $\beta_1, \ldots, \beta_n \in L$, not all zero, such that $\sum_i \beta_i a'_i = 0$. Applying Φ to this relation gives $\sum_i \beta_i a_i = 0$, so that A is not L-linearly independent, contradiction. Thus A' is L-linearly independent independent and *a fortiori* so is its subset S.
(iii$'$) $\implies$ (ii$'$) in the same way. □

Remark: Some source take condition (ii) of Proposition 12.4 to be the definition of linear disjointness. This has the advantage of not requiring any knowledge of tensor products on the part of the reader. All the other advantages, however, seem to lie with the tensor product definition. For instance, it is clearly symmetric with respect to K and L.

Exercise 12.6. *Let K, L be subfields of E/F, and let R be an F-subalgebra of K with fraction field K. Suppose that there exists a K-basis of R which is L-linearly independent in E. Show that K, L are F-linearly disjoint in E.*

Proposition 12.5. *Let E_1, E_2 be subextensions of E/F and K_1 a subextension of E_1/F. Then E_1, E_2 are linearly disjoint over F if and only if:*
(i) K_1, E_2 are linearly disjoint over F, and
(ii) K_1E_2, E_1 are linearly disjoint over K_1.

Proof. JacobsonII, p. 526. □

12.2. **Intrinsic Nature of Linear Disjointness.**

The definiton of linear disjointness is initially hard to process because it involves four different algebras. In fact the dependence of the definition on the "ambient" field E is in many cases rather weak. One easy of instance of this is given in the following exercise.

Exercise 12.7. *Let K, L be subextensions of a field extension E/F, and let E'/E be any field extension. Show: K, L are F-linearly disjoint as subfields of E iff they are F-linearly disjoint as subfields of E'.*

We now look more deeply into the dependence on the ambient field E, following a MathOverflow discussion led by Andrew Critch. Let F be a field, and let K, L be field extensions of F. We say that K, L are **somewhere linearly disjoint over F** if there exists a field extension E/F and F-algebra embeddings of K and L into E such that K, L are F-linearly disjoint in E. Further, we say that K, L are **everywhere linearly disjoint over F** if for all field extensions E/F and all F-algebra embeddings of K, L into E, K, L are F-linearly disjoint in E.

Certainly we want everywhere linearly disjoint over F to imply somewhere linearly disjoint over F. To see this there is a minor technicality to be disposed of, which is treated in the next exercise.

Exercise 12.8. *a) Let F be a field and K, L be field extensions of F. Show: there exists a field extension E and F-algebra embeddings of K and L into E. Show that for instance one may take E to be any algebraically closed field such that* $\mathrm{trdeg}(E/F) \geq \max \mathrm{trdeg}(K/F), \mathrm{trdeg}(L/F)$.

b) Deduce: if K, L are everywhere linearly disjoint over F then they are somewhere linearly disjoint over F.

Exercise 12.9. *Let F be any field, and put $K = L = F(t)$.*
a) Take $E = F(t)$ to show that K, L are not everywhere linearly disjoint.
b) Take $E = F(a, b)$ (rational function field in two variables) to show that K and L are somewhere linearly disjoint.

Proposition 12.6. *Let F be a field, and let K and L be field extensions of F. TFAE:*
(i) K, L are somewhere F-linearly disjoint.
(ii) The tensor product $K \otimes_F L$ is a domain.

Proof. If $K \otimes_F L$ can be embedded into a field, then it is a domain. Conversely, if $K \otimes_F L$ is a domain, it can be embedded into its fraction field. □

Corollary 12.7. *Let F be a field, K, L be field extensions of F and R an F-subalgebra of K which is a domain. Then K, L are somewhere F-linearly disjoint iff $R \otimes_K L$ is a domain.*

Proof. By Proposition 12.6, K, L are somewhere F-linearly disjoint iff $K \otimes_F L$ is a domain. If it is, then the subring $R \otimes_F L$ is also a domain. Conversely, if $R \otimes_F L$ is a domain, then taking E to be its fraction field shows that K, L are somewhere F-linearly disjoint. □

Proposition 12.8. *Let F be a field, and let K and L be field extensions of F. TFAE:*
(i) K, L are everywhere F-linearly disjoint.
(ii) $K \otimes_F L$ is a field.

Proof. (i) $\implies$ (ii): In order to show that the (evidently nonzero, since it contains F) ring $R = K \otimes_F L$ is a field, it suffices to show that the only maximal ideal is (0). So let $\mathfrak{m}$ be a maximal ideal of R. Then $E = R/\mathfrak{m}$ is a field extension of K, L and the induced map $K \otimes_F L \to E$ is precisely the quotient map $R \to R/\mathfrak{m}$. Since this map is injective, $\mathfrak{m} = (0)$.
(ii) $\implies$ (i): If $R = K \otimes_F L$ is a field, then every homomorphism into a nonzero ring – and in particular, any F-algebra homomorphism – is injective. □

Theorem 12.9. *Let K, L be field extensions of F.*
a) Suppose that K, L are everywhere F-linearly disjoint. Then at least one of K, L is algebraic over F.
b) Conversely, suppose that at least one of K, L is algebraic over F. Then K, L are somewhere F-linearly disjoint iff they are everywhere F-linearly disjoint.

Proof. a) If K and L are transcendental over F, then they admit subextensions $K' = F(a)$, $L' = F(b)$. By Exercise 12.5, it suffices to show that $F(a)$ and $F(b)$ are not everywhere F-linearly disjoint over F. To see this take $E = F(t)$ and map $K' \to E$ by $a \mapsto t$ and $L' \to E$ by $b \mapsto t$ and apply Lemma 12.2.
b) Because every algebraic extension is a direct limit of finite extensions, by Exercise 12.4 it is no loss of generality to assume that K/F is finite, and in light of Propositions 12.6 and 12.8, we must show that if $K \otimes_F L$ is a domain then it is a field. But if $\{k_1, \ldots, k_n\}$ is a basis for K/F, then $k_1 \otimes 1, \ldots, k_n \otimes 1$ is a basis for $K \otimes_F L$ over L, so $K \otimes_F L$ is a domain and a finitely generated L-module. Therefore

it is a field, by an elementary argument which we have seen before (and which is a special case of the preservation of Krull dimension in an integral extension). □

In conclusion: the notion of F-linear disjointness of two field extensions K, L is *intrinsic* – independent of the embeddings into E – iff at least one of K, L is algebraic over F. In most of our applications of linear disjointness this hypothesis will be satisfied, and when it is we may safely omit mention of the ambient field E.

Here is a first result with our new convention in force.

Theorem 12.10. *Let K/F be purely transcendental and L/F be algebraic. Then K, L are F-linearly disjoint.*

Proof. By Exercise 12.4 and Lemma 12.3 it is enough to show that for all $n \in \mathbb{Z}^+$, $F[x_1, \ldots, x_n]$, L are F-linearly disjoint. By Corollary 4.5, this holds iff $F[x_1, \ldots, x_n] \otimes_F L$ is a domain. It is clear that the F-basis of $F[x_1, \ldots, x_n]$ consisting of monomials remains L-linearly independent in $L[x_1, \ldots, x_n]$ and by Proposition 12.4 this implies that $F[x_1, \ldots, x_n]$ and L are F-linearly disjoint. In particular, the natural map $F[x_1, \ldots, x_n] \otimes_K L \to L[x_1, \ldots, x_n]$ is an isomorphism of L-algebras. □

Theorem 12.11. *Let K, L be two field extensions of F with K/F purely transcendental. Then $K \otimes_F L$ is a domain.*

Proof. The F-algebra $K \otimes_F L$ is the direct limit of the F-algebras $K_i \otimes_F L_i$ as K_i ranges over finitely generated subextensions of K/F and L_i ranges over finitely generated subextensions of L/F. Since the direct limit of domains is a domain, we have reduced to the case in which K and L are finitely generated over F, say $E_2 = F(s_1, \ldots, s_m)$, and $E_1 = F(t_1, \ldots, t_n, x_1, \ldots, x_p)$, where the t_i's are independent indeterminates over F and for all $1 \leq k \leq p$, $F(t_1, \ldots, t_n, x_1, \ldots, x_k)/F(t_1, \ldots, t_n, x_1, \ldots, x_{k-1})$ has finite degree. Put $K_1 = F(t_1, \ldots, t_n)$. Let E be the algebraic closure of the fraction field of $F[s_1, \ldots, s_m] \otimes_F F[t_1, \ldots, t_n]$. We may embed E_2 and L in E, and then E_2 and K_1 are linearly disjoint over F. Since K_1E_2/K_1 is purely transcendental and E_1/K_1 is algebraic, by Theorem 12.14 K_1E_2 and E_1 are linearly disjoint over K_1. By Proposition 12.5, E_1 and E_2 are linearly disjoint over F, hence $E_1 \otimes_F E_2$ is a domain. □

A field extension K/F is **regular** if it satisfies the conclusion of Theorem 12.11: that is $K \otimes_F L$ is a domain for all field extensions L/F. Thus purely transcendental extensions are regular, whereas nontrivial algebraic extensions are not regular. Later we will give a characterization of regular extensions.

12.3. **Linear Disjointness and Normality.**

Proposition 12.12. *Let E/F be a field extension, and let K, L be two finite degree subextensions, with K/F and L/F both Galois extensions. Then K, L are F-linearly disjoint [in E, but by Theorem 12.9 this does not matter] iff $K \cap L = F$.*

Proof. The forward direction holds for any pair of subextensions by Lemma 12.2. Conversely, assume $K \cap L = F$. The image of $K \otimes_F L$ in E is the compositum KL, which is finite Galois over F since normality, separability and finiteness of degree are all preserved by finite composita. Let $d = [KL : F]$, $d_K = [K : F]$ and $d_L = [L : F]$. We have a surjective F-linear map $\iota : K \otimes_F L \to KL$ between two finite-dimensional F-vector spaces, so ι is injective iff $[KL : F] = [K \otimes_F L : F] = d_K d_L$. Let

$H_K = \mathrm{Aut}(KL/K)$ and $H_L = \mathrm{Aut}(KL/L)$. Since K/F and L/F are Galois, H_K and H_L are normal in $G = \mathrm{Aut}(L/K)$. Moreover, since KL is the compositum of $(KL)^{H_K}$ and $(KL)^{H_L}$, $KL = KL^{H_K \cap H_L}$, i.e., $H_K \cap H_L = \{e\}$. Therefore $H_K H_L$ is a subgroup of G, the internal direct product of H_K and H_L. Moreover,

$$F = K \cap L = (KL)^{H_K} \cap (KL)^{H_L} = KL^{\langle H_K, H_L \rangle} = KL^{H_K H_L},$$

so $H_K H_L = G$. It follows that $G = H_K \times H_L$ and therefore

$$d = \#G = \#H_K \#H_L = \frac{d}{d_K}\frac{d}{d_L}$$

and thus $[K \otimes L : F] = d_K d_L = d = [KL : F]$. □

Theorem 12.13. *Let E/F be a field extension, and let K, L be two algebraic subextensions such that K/F is Galois. Then K, L are F-linearly disjoint [in E, but...] iff $K \cap L = F$.*

Proof. By a now familiar argument involving Exercise 12.5 and Proposition 12.8, we reduce to the case in which K/F is finite Galois. Now by the theorem of Natural Irrationalities, we have

$$[KL : F] = [KL : L][L : F] = [K : K \cap L][L : F],$$

so

$$[KL : F] = [K : F][L : F] \iff K \cap L = F.$$ □

12.4. **Linear Disjointness and Separability.**

Lemma 12.14. *Let K/F be a separable field extension of characteristic $p > 0$, and let $a_1, \ldots, a_n$ be F-linearly independent elements of K. Then for all $e \in \mathbb{Z}^+$, $a_1^{p^e}, \ldots, a_n^{p^e}$ are F-linearly independent.*

Proof. By replacing K by $F(a_1, \ldots, a_n)$, we may assume that $a_1, \ldots, a_n$ is a basis for K and thus $[K : F] = n$.
Step 1: Let $K' = F(a_1^{p^e}, \ldots, a_n^{p^e})$, so that K' is a subextension of K/F. Observe that for all i, the element a_i is both separable and purely inseparable over K', so $a_i \in K'$ for all i and thus $K' = K$ and $[K' : F] = n$.
Step 2: Let V be the F-subspace spanned by $a_1^{p^e}, \ldots, a_n^{p^e}$. It is enough to show that V is closed under multiplication: then it is a subring of a field which is finite dimensional as an F-algebra and therefore a field and therefore the F-subalgebra generated by $a_1^{p^1}, \ldots, a_n^{p^e}$. By Step 1, this means $V = K' = K$ and thus $[V : F] = n$. Therefore the n-element spanning set $a_1^{p^e}, \ldots, a_n^{p^e}$ is linearly independent.
Step 3: To show that V is a subalgebra, it is enough to show that the product of two basis elements is an F-linear combination of the basis elements. To see this, fix any $1 \leq i, j \leq n$. Since $a_1, \ldots, a_n$ span K over F, there exist $\alpha_1, \ldots, \alpha_n$ such that

$$a_i a_j = \sum_i \alpha_i a_i$$

Raising both sides to the p^eth power gives

$$a_i^{p^e} a_j^{p^e} = \sum_i \alpha_i^{p^e} a_i^{p^e},$$

which shows that $a_i^{p^e} a_j^{p^e}$ lies in V. □

Proposition 12.15. *Let K/F be a separable algebraic extension. Then K and $F^{p^{-\infty}}$ are F-linearly disjoint.*

Proof. Since $F^{p^{-\infty}} = \varinjlim F^{p^{-e}}$, by Exercises 12.3, 12.4 and Proposition 12.4 it is enough to show that for all $e, m \in \mathbb{Z}^+$, if $a_1, \ldots, a_n$ are F-linearly independent elements of K, they are also $F^{p^{-e}}$-linearly independent. But this last statement holds iff $a_1^{p^e}, \ldots, a_n^{p^e}$ are F-linearly independent, which they are by Lemma 12.14. □

The natural question to ask at this point is: can an inseparable extension K/F be linearly disjoint from $F^{p^{-\infty}}$? It follows immediately from what we already know about separable extensions that the answer is *no* if K/F is inseparable and normal, for then by XXX it contains a nontrivial purely inseparable subextension and thus $F \subsetneq K \cap F^{p^{-1}} \subset K \cap F^{p^{-\infty}}$. In fact, as we are about to see, among algebraic field extensions K/F, being linearly disjoint from $F^{p^{-\infty}}$ characterizes separable extensions. But actually we can go further, with the following definitions.

A **separating transcendence basis** for a field extension K/F is an algebraically independent subset S of K such that $K/F(S)$ is separable algebraic.

It is clear that separating transcendence bases need not exist, e.g. an inseparable algebraic extension will not admit a separating transcendence basis. On the other hand, it is clear that separable algebraic extensions and purely transcendental extensions both admit separating transcendence bases: as with being linearly disjoint from the perfect closure, this is something that these apparently very different classes of extensions have in common.

We say that a field extension K/F is **separably generated** if it admits a separating transcendence basis.

Exercise 12.10. *Give an example of a separably generated field extension admitting a transcendence basis that* is not *a separating transcendence basis.*

An arbitrary field extension K/F is **separable** if every finitely generated subextension admits a separating transcendence basis.

And now the main theorem on separable extensions.

Theorem 12.16. *(Mac Lane) Let F be a field of characteristic $p > 0$, and let E/F be a field extension. The following are equivalent:*
(i) E/F is separable: every finitely generated subextension is separably generated.
(ii) E and $F^{p^{-\infty}}$ are F-linearly disjoint.
(iii) E and $F^{p^{-1}}$ are F-linearly disjoint.

Proof. (i) $\implies$ (ii): Since every field extension is the direct limit of its finite generated subextensions, by Exercise LD2 we may assume that E/F is finitely generated and thus separably generated, so let B be a transcendence basis for E/F such that $E/F(B)$ is separable algebraic. By Proposition 12.15, E and $F(B)^{p^{-\infty}}$ are $F(B)$-linearly disjoint. Since $F(B)^{p^{-\infty}} \supset F^{p^{-\infty}}(B)$, it follows that $F^{p^{-\infty}}(B)$ are $F(B)$-linearly disjoint. By Proposition 12.5, E and $F^{p^{-\infty}}$ are F-linearly disjoint.

(ii) $\implies$ (iii) is immediate.
(iii) $\implies$ (i): Suppose that E and $F^{p^{-1}}$ are F-linearly disjoint. We will prove by induction on n that for all $n \in \mathbb{N}$, if $K = F(a_1, \ldots, a_n)$ is a finitely generated subextension of E/F then there exists a subset $S \subset \{a_1, \ldots, a_n\}$ which is a separating transcendence basis for K/F. When $n = 0$, $K = F$ and the result is trivial. The result is also clear if $a_1, \ldots, a_n$ are algebraically independent. Hence we may assume (after relabelling) that there exists $r < n$ such that $a_1, \ldots, a_r$ are a transcendence basis for K/F. Let $f \in K[t_1, \ldots, t_{r+1}]$ be a polynomial of minimal total degree such that $f(a_1, \ldots, a_{r+1}) = 0$; necessarily f is irreducible.

We CLAIM f is *not* of the form $g(t_1^p, \ldots, t_r^p)$. If it were, there would exist $h \in F^{p^{-1}}[t_1, \ldots, t_{r+1}]$ such that $g(t_1^p, \ldots, t_r^p) = h(t_1, \ldots, t_r)^p$ with $h(a_1, \ldots, a_{r+1}) = 0$. Let $\{m_i\}$ be the monomials occurring in h. Then the elements $m_i(a_1, \ldots, a_{r+1})$ are $F^{p^{-1}}$-linearly dependent, so *by hypothesis* they are F-linearly dependent. This gives a nontrivial polynomial relation in the a_i of degree less than the degree of h, contradiction.

It follows that there is at least one i, $1 \leq i \leq r+1$, such that $f(t_1, \ldots, t_{r+1})$ is not a polynomial in t_i^p. Then a_i is algebraic over $F(a_1, \ldots, a_{i-1}, a_{i+1}, \ldots, a_{r+1})$ and thus $\{a_1, \ldots, a_{i-1}, a_{i+1}, \ldots, a_{r+1}\}$ is a transcendence basis for K/F. So

$$F[a_1, \ldots, a_{i-1}, t, a_{i+1}, \ldots, a_{r+1}] \cong F[t_1, \ldots, t_{r+1}],$$

so $f(a_1, \ldots, a_{i-1}, t, a_{i+1}, \ldots, a_{r+1})$ is ireducible in $F[a_1, \ldots, a_{i-1}, t, a_{i+1}, \ldots, a_{r+1}]$, so by Gauss's Lemma it is irreducible in $F(a_1, \ldots, a_{i-1}, a_{i+1}, \ldots, a_{r+1})[t]$. Since a_i is a root of $f(a_1, \ldots, a_{i-1}, t, a_{i+1}, \ldots, a_{r+1})$ and this is not a polynomial in t^p, a_i is separable algebraic over $F(a_1, \ldots, a_{i-1}, a_{i+1}, \ldots, a_{r+1})$ and hence over $L := F(a_1, \ldots, a_{i-1}, a_{i+1}, \ldots, a_n)$. The induction hypothesis applies to L to give a subset $\{a_{i_1}, \ldots, a_{i_r}\}$ of $\{a_1, \ldots, a_{i-1}, a_{i+1}, \ldots, a_n\}$ that is a separating transcendence base for L/F. Since a_i is separable algebraic over L, it is separable algebraic over $F(a_{i_1}, \ldots, a_{i_r})$. So $\{a_{i_1}, \ldots, a_{i_r}\}$ is a separating transcendence basis for K/F. □

Example 12.17. *(Mac Lane): Let F be any field of characteristic $p > 0$, $F(t)$ a rational function field. Let $E = F(t, t^{p^{-1}}, t^{p^{-2}}, \ldots)$. Then any finitely generated subextension of E/F is isomorphic to $F(t)$ and thus separably generated. But E itself does not admit a separating transcendence basis. Thus E/F is separable but not separably generated.*

Exercise 12.11. *Show: a separably generated extension is separable.*

Exercise 12.12. *Show: a subextension of a separable extension is separable.*

Exercise 12.13. *a) Show: separably generated extensions do not satisfy the base change property (DC2).*
(Suggestion: let F be any field of characteristic $p > 0$, let $K = F(t)$, and let $L = F(t^{\frac{1}{p}})$.)
b) Conclude that separably generated extensions – and thus also separable extensions – do not form a distinguished class in the sense of Lang.
c) Prove or disprove: the compositum of two separably generated extensions is separably generated.
d) Prove or disprove: the compositum of two separable extensions is separable.

13. Derivations and Differentials

13.1. **Derivations.**

Let R be a commutative ring, and let M be an R-module. A **derivation** of R into M is a map $D : R \to M$ satisfying both of the following:

(D1) For all $x, y \in R$, $D(x + y) = D(x) + D(y)$ (i.e., D is a homomorphism of additive groups),
(D2) For all $x, y \in R$, $D(xy) = xD(y) + D(x)y$ ("Leibniz rule").

Exercise DER0: Let $D : R \to M$ be a derivation, let $x \in R$ and let $n \in \mathbb{Z}^+$. Show that $D(x^n) = nx^{n-1}D(x)$.

Suppose we are given a subring k of R. Then a **k-derivation** is a derivation $D : R \to M$ satisfying the additional property

(D3) For all $x \in k$, $D(x) = 0$.

We often have $M = R$ and then we speak of derivations and k-derivations *on* R. We denote the set of all k-derivations on R by $\mathrm{Der}_k(R)$.

Exercise DER0.5:
a) Show that any k-derivation $D : R \to M$ is a k-linear map: for all $c \in k$ and $x \in R$, $D(cx) = cD(x)$.
b) Show that $\mathrm{Der}_k(R)$ is a k-submodule of $\mathrm{Hom}_k(R, R)$.
c) Show that $\mathrm{Der}_k(R)$ in fact has the structure of an R-module: if $D \in \mathrm{Der}_k(R)$ and $\alpha \in R$, then $\alpha D \in \mathrm{Der}_k(R)$.
d) Show that if $D_1, D_2 \in \mathrm{Der}_k(R)$, $D_1 \circ D_2$ need not be a derivation of R.
e) Show that if $D_1, D_2 \in \mathrm{Der}_k(R)$, then the map $[D_1, D_2] : R \to R$ defined by $[D_1, D_2] : x \mapsto D_1(D_2(x)) - D_2(D_1((x))$ is a k-derivation of R.
f) Suppose that k is a field of characteristic $p > 0$. Show that for any $D \in \mathrm{Der}_k(R)$, the p-fold composition $D^{\circ p}$ is a k-derivation on R.

Exercise DER1: Let $D : R \to M$ be a derivation, and let $C = \{x \in R \mid D(x) = 0\}$ be its kernel. Show that C is a subring of R and is in fact the unique maximal subring of k of R such that D is a k-derivation. (It is sometimes called the **constant subring** of R.)

Example: Let k be a field and $R = k[t]$. The usual polynomial derivative $f \mapsto f'$ is a k-derivation on R; we will denote it by ∂. The derivation ∂ is the unique k-derivation D such that $D(t) = 0$.

Exercise DER2: Compute the constant subring of $\partial : k[t] \to k[t]$. Note that the answer in positive characteristic is very different from characteristic zero!

Exercise DER3: Let k be a domain, $n \in \mathbb{Z}^+$, and let $R = k[t_1, \ldots, t_n]$ be the polynomial ring in n variables over k. Show that for each $1 \leq i \leq n$ there is a unique k-derivation ∂_i on R such that $\partial_i(t_j) = \delta_{ij}$.

At least when $k = \mathbb{R}$, it is well known that we may differentiate not only polynomials but also rational functions. This generalizes nicely to our abstract algebraic context.

Theorem 13.1. *Let R be a domain with fraction field K and D a derivation on R.*
a) There is a unique extension of D to a derivation on K, given by

$$D_K\left(\frac{x}{y}\right) = \frac{yD(x) - xD(y)}{y^2}. \tag{14}$$

b) If D is a k-derivation for some subring k of R with fraction field $f(k)$, then D_K is an $f(k)$-derivation.

Proof. a) Our first order of business is to show that D_K is well-defined, i.e., if $x_1, x_2, y_1, y_2 \in R$ are such that $y_1 y_2 \neq 0$ and $x_1 y_2 = x_2 y_1$, then

$$\frac{y_1 D(x_1) - x_1 D(y_1)}{y_1^2} = \frac{y_2 D(x_2) - x_2 D(y_2)}{y_2^2}.$$

We check this by a straightforward if somewhat unenlightening calculation:

$$\begin{aligned} &y_2^2\left(x_1 D(y_1) - y_1 D(x_1)\right) - \left(y_1^2(x_2 D(y_2) - y_2 D(x_2)\right) \\ &= y_2^2 x_1 D(y_1) - y_2^2 y_1 D(x_1) - y_1^2 x_2 D(y_2) - y_1^2 y_2 D(x_2) \\ &= (y_2 x_1 d(y_1 y_2) - y_1 y_2 D(x_1 y_2)) - (y_1 x_2 D(y_1 y_2) + y_1 y_2 D(y_1 x_2)) \\ &= (x_1 y_2 - x_2 y_1) D(y_1 y_2) - y_1 y_2 D(x_1 y_2 - x_2 y_1) = 0. \end{aligned}$$

Next we check that D_K is a derivation:

$$D_K\left(\frac{x_1}{y_1} + \frac{x_2}{y_2}\right) = D_K\left(\frac{x_1 y_2 + x_2 y_1}{y_1 y_2}\right) = \ldots.$$

$$D_K\left(\frac{x_1}{y_1}\frac{x_2}{y_2}\right) = D_K\left(\frac{x_1 x_2}{y_1 y_2}\right) = \ldots.$$

Let $\mathcal{D}$ be any derivation on K extending D. For $x, y \in K$ with $y \neq 0$, we have

$$\mathcal{D}(x) = \mathcal{D}\left(\frac{x}{y} \cdot y\right) = \frac{x}{y}\mathcal{D}(y) + y\mathcal{D}\left(\frac{x}{y}\right),$$

so

$$\mathcal{D}\left(\frac{x}{y}\right) = \frac{\mathcal{D}(x)}{y} - \frac{x\mathcal{D}(y)}{y^2} = \frac{y\mathcal{D}(x) - x\mathcal{D}(y)}{y^2} = D_K\left(\frac{x}{y}\right),$$

completing the proof of part a).
b) Since D_K extends D and $D(x) = 0$ for all $x \in k$, certainly $D_K(x) = 0$ for all $x \in k$. Using (14) it follows that for all $x, y \in k$ with $y \neq 0$, $D\left(\frac{x}{y}\right) = 0$. □

We now concentrate our studies on $\mathrm{Der}_K(L)$ for a field extension L/K.

Proposition 13.2. *Let L/K be a field extension, and let $D \in \mathrm{Der}_K(L)$. Let $f \in K[t_1, \ldots, t_n]$ and $a = (a_1, \ldots, a_n) \in L^n$. Then*

$$D(f(a)) = \sum_{i=1}^{n} \partial_i f(a_1, \ldots, a_n) D(a_i).$$

Exercise DER4: Prove Proposition 13.2.

Let L/K be a field extension and $S \subset L$. Any derivation D on L restricts to a function $D_S : S \to L$. We say that D is **S-finite** if $\{x \in S \mid D(x) \neq 0\}$ is finite. Of course S-finiteness is automatic if S itself is a finite set. The S-finite derivations form an L-subspace of $\mathrm{Der}_K(L)$ which we will denote by $\mathrm{Der}_K^S(L)$.

Proposition 13.3. *Let L/K be a field extension. Let $S \subset L$ be such that $L = K(S)$.*
a) We have

$$\dim_L \mathrm{Der}_K^S(L) \leq \#S.$$

b) In particular if L can be generated as a field extension by $n < \infty$ elements, then $\dim_L \mathrm{Der}_K(L) \leq n$.

Proof. Let $L^{(S)}$ be the set of all finitely nonzero functions from S to L. This is an L-vector space with basis canonically in bijection with S: indeed, for $s \in S$, let δ_s be the function which takes the value 1 at s and zero elsewhere. Then $\{\delta_s\}_{s \in S}$ is an L-basis for $L^{(S)}$.
The natural restriction map $\mathrm{Der}_K^S(L) \to L^{(S)}$ is L-linear and injective. The L-linearity is a triviality: the injectivity follows from the fact that every element of L is a rational function in the elements of S with coefficients in K. Since $\dim L^{(S)} = \#S$, part a) follows immediately! Part b) is also immediate from the observation that S-finiteness is a vacuous condition when S itself is a finite set. □

We will see later on that equality holds in Proposition 13.3 when S is a separating transcendence basis for L/K. At the other extreme, Proposition 13.3 together with the Primitive Element Corollary shows that for any finite separable extension L/K, $\dim_L \mathrm{Der}_K(L) \leq 1$. In fact we can do better.

Theorem 13.4. *(Derivation Extension Theorem) See Lang's algebra, pages 369-370, for the statement and proof.*

Corollary 13.5. *Let L/K be a separable algebraic extension.*
a) Every derivation on K extends uniquely to a derivation on L.
b) We have $\mathrm{Der}_K(L) = 0$.

Proof. a) Step 1: suppose that L/K is finite.
Step 2: Suppose that L/K is an infinite degree separable extension. It is therefore the direct limit of its finite separable subextensions L_α. By Step 1, there exists a unique $D_\alpha \in \mathrm{Der}_F(L)$ extending D. Because of the uniqueness, it is automatic that these derivations fit together to give a derivation D_L on L: that is, for any $x \in L$, we choose α such that $x \in L_\alpha$ and put $D_L(x) = D_{L_\alpha}(x)$. If $x \in L_\alpha \cap L_\beta$ then the uniqueness forces $D_{L_\alpha}(x) = D_{L_\alpha L_\beta}(x) = D_{L_\beta}(x)$.
b) Let $D \in \mathrm{Der}_K(L)$. Then D extends $0 \in \mathrm{Der}(K)$, as does $0 \in \mathrm{Der}\, L$. By part a), we must have $D = 0$. □

Exercise DER5: Let $F \subset K \subset L$ be a tower of fields with L/K separable algebraic, and let M be a subextension of L/K. Show that for any F-derivation D of L such that $D(K) \subset K$, we have also $D(M) \subset M$.

Corollary 13.6. *Let K be a field of chracteristic $p > 0$, and let $L = K(x)$ be a nontrivial purely inseparable field extension of K.*

a) For each $D \in \operatorname{Der} K$ and $\alpha \in L$, there exists a unique $D_L \in \operatorname{Der} L$ extending D_K and such that $D_L(x) = \alpha$.
b) In particular $\dim_L \operatorname{Der}_K(L) = 1$; a basis is given by $\{D_x\}$, where D_x is the unique K-derivation with $D_x(x) = 1$.

Proof. a) ...
b) By part a) there is a unique derivation D_x on L extending the zero derivation on K and such that $D_x(x) = 1$. Thus $\dim_L \operatorname{Der}_K(L) \geq 1$. On the other hand, by Proposition 13.3, $\dim_L \operatorname{Der}_K(L) \leq 1$. We deduce that $\operatorname{Der}_K(L)$ is a one-dimensional L-vector space and thus a basis is given by any nonzero vector in that space, e.g. D_x. □

Corollary 13.7. *Let $L = K(t)$ be a univariate rational function field.*
a) For each $D \in \operatorname{Der} K$ and $\alpha \in L$, there exists a unique $D_L \in \operatorname{Der} L$ extending D_K and such that $D_L(t) = \alpha$.
b) In particular $\dim_L \operatorname{Der}_K(L) = 1$; a basis is given by ∂_t, where ∂_t is the unique K-derivation with $\partial_t(t) = 1$.

Proof. ... □

Corollary 13.8. *Let L/K be any field extension, and let D be a derivation of K. Then there is at least one extension of D to a derivation on L.*

Proof. Consider the set of pairs (M, D_M) where M is a subextension of L/K and $D_M \in \operatorname{Der} M$ extends D. This set is partially ordered as follows: $(M_1, D_{M_1}) \leq (M_2, D_{M_2})$ if $M_1 \subset M_2$ and D_{M_2} extends D_{M_1}. It is easy to see that the hypothesis of Zorn's Lemma is satisfied, so that we get a maximal element (M, D_M).

Suppose first that L/M is not algebraic. Then there exists an element $t \in L$ which is transcendental over M, so that $M(t) \subset L$ is a rational function field. By Corollary 13.7, D_M extends to $M(t)$, contradicting the maximality of M.
Next suppose that L/M is algebraic, and let M^s be the separable closure of M in M. If $M^s \supsetneq L$, then by Corollary 13.5 D_M extends to a derivation on M_s, contradicting the maximality of M. So it must be the case that L/M is purely inseparable. Thus if $L \supsetneq M$, there exists $x \in L \setminus M$ and then $M(x)/M$ is a proper purely inseparable extesion. By Corollary 13.6, D_M extends to a derivation on $M(x)$, contraditing the maximality of M.

It follows that $M = L$, i.e., the derivation D can be extended to L. □

Proposition 13.9. *Let K be a field, S a set and $L = K(\{t_s\}_{s \in S}$ be the rational function field over K.*
a) For $s \in S$, there is a unique K-derivation δ_s of L such that $\delta_s(t_s) = 1$ and $\delta_s(t_{s'}) = 0$ for all $s' \neq s$.
b) The set $\{\delta_s\}_{s \in S}$ is an L-basis for $\operatorname{Der}_K^S(L)$.

Theorem 13.10. *Let L/K be a finitely generated separable field extension.*
a) We have $\operatorname{trdeg}(L/K) = \dim_L \operatorname{Der}_K(L)$.
b) If $\{x_1, \ldots, x_n\}$ is a separating transcendence basis for L/K, then there is a basis $\{D_i\}_{1 \leq i \leq n}$ for $\operatorname{Der}_K(L)$ such that for all $1 \leq i \leq n$, the restriction of D_i to $K(x_1, \ldots, x_n)$ is ∂_i.

Proof. Let $\{x_1, \ldots, x_n\}$ be a separating transcendence basis for L/K and put $M = K(x_1, \ldots, x_n)$. By Theorem **??**, for each $1 \leq i \leq n$ there exists a unique K-derivation of L extending ∂_i on M: call it D_i. We claim that $\{D_1, \ldots, D_n\}$ is an

L-basis for $\mathrm{Der}_K(L)$: this will establish both parts of the theorem.

□

13.2. **Differentials.**

14. Applications to Algebraic Geometry

15. Ordered Fields

15.1. **Ordered Abelian Groups.**

An **ordered abelian group** $(G, +, <)$ is an abelian group $(G, +)$ equipped with a total ordering $<$ which is compatible with the group law in the sense that

(OAG) For all $x, y, z \in G,\ x \leq y \implies x + z \leq y + z$.

A homomorphism of ordered abelian groups $f : (G, <) \to (H, <)$ is a group homomorphism which is **isotone**: for all $x_1 \leq x_2$, $f(x_1) \leq f(x_2)$.

Lemma 15.1. *For x, y, z in an ordered abelian group G, if $x < y$ then $x+z < y+z$.*

Proof. Since $x < y$, certainly $x \leq y$, so by (OAG) $x + z \leq y + z$. If $x + z = y + z$ then adding $-z$ to both sides gives $x = y$, a contradiction. □

Lemma 15.2. *Let x_1, x_2, y_1, y_2 be elements of an ordered abelian group G with $x_1 \leq x_2$ and $y_1 \leq y_2$. Then $x_1 + y_1 \leq x_2 + y_2$.*

Proof. Applying (OAG) with x_1, x_2, y_1 gives $x_1 + y_1 \leq x_2 + y_1$. Applying (OAG) with y_1, y_2, x_2 gives $x_2+y_1 = y_1+x_2 \leq y_2+x_2$. By transitivity $x_1+y_1 \leq x_2+y_2$. □

To an ordering on a commutative group we associate its **positive cone**:

$$G^+ = \{x \in G \mid x > 0\}.$$

Elements of G^+ are called **positive**. We also define

$$G^- = \{x \in G \mid x < 0\}.$$

Elements of G^- are called **negative**.

Lemma 15.3. *Let x be a nonzero element of the ordered abelian group G. Then exactly one of x, $-x$ is positive. Thus $G = \{0\} \coprod G^+ \coprod G^-$.*

Proof. If $x > 0$ and $-x > 0$ then adding gives $0 > 0$, a contradiction.
If x is not positive then $x < 0$. By Lemma 15.1 we may add $-x$ to both sides, getting $0 = x + (-x) < 0 + x = -x$. □

Lemma 15.4. *Let x_1, x_2 be elements of an ordered abelian group.*
a) If $x_1, x_2 \in G^+$, then $x_1 + x_2 \in G^+$.
b) If $x_1, x_2 \in G^-$, then $x_1 + x_2 \in G^-$.

Proof. a) Since $x_1 > 0$ and $x_2 > 0$, by Lemma 15.1 $x_1 + x_2 > 0$.
b) If $x_1 < 0$ and $x_2 < 0$, then by Lemma 15.3 $-x_1, -x_2 > 0$, so by part a) $-x_1 - x_2 = -(x_1 + x_2) > 0$, so by Lemma 15.1 again $x_1 + x_2 < 0$. □

In an ordered abelian group we define $|x|$ to be x if $x \geq 0$ and $-x$ otherwise.

Exercise: Let x, y be elements of an ordered abelian group G.
a) Suppose $x \leq y$ and $n \in \mathbb{N}$. Show that $nx \leq ny$.
b) Suppose $x \leq y$ and n is a negative integer. Show that $nx \geq ny$.

Example: Let $(G, <)$ be an ordered abelian group and H a subgroup of G. Restricting $<$ to H endows H with the structure of an ordered abelian group.

Example (Lexicographic ordering): Let $\{G_i\}_{i \in I}$ be a nonempty indexed family of ordered abelian groups. Suppose that we are given a well-ordering on the index set I. We may then endow the direct product $G = \prod_{i \in I} G_i$ with the structure of an ordered abelian group, as follows: for $(g_i), (h_i) \in G$, we decree $(g_i) < (h_i)$ if for the least index i such that $g_i \neq h_i$, $g_i < h_i$.

Theorem 15.5. *(Levi* [Lev43]*) For an abelian group G, TFAE:*
(i) G admits at least one ordering.
(ii) G is torsionfree.

Proof. (i) $\implies$ (ii) Let $<$ be an ordering on G, and let $x \in G^\bullet$. By Lemma 15.4 we have $nx \neq 0$ for all $n \in \mathbb{Z}^+$.
(ii) $\implies$ (i): Let G be a torsionfree abelian group. Then G is a flat $\mathbb{Z}$-module. Tensoring the injection $\mathbb{Z} \hookrightarrow \mathbb{Q}$ gives us an injection $G \hookrightarrow G \otimes \mathbb{Q}$. Since $\mathbb{Q}$ is a field, the $\mathbb{Q}$-module $G \otimes \mathbb{Q}$ is free, i.e., it is isomorphic to $\bigoplus_{i \in I} \mathbb{Q}$. Choose a total ordering on I. Give each copy of $\mathbb{Q}$ its standard ordering as a subfield of $\mathbb{R}$ and put the lexicographic ordering on $\bigoplus_{i \in \mathbb{Q}} \mathbb{Q} \cong G \otimes \mathbb{Q}$. Via the injection $G \hookrightarrow G \otimes \mathbb{Q}$ this induces an ordering on G. □

An **anti-isomorphism** of abelian groups is an order-reversing group isomophism. For every ordered abelian group $(G, <)$, the inversion map $x \in G \mapsto -x$ is an anti-isomorphism of G.

Exercise X.X: a) Show that the abelian group $\mathbb{Z}$ admits exactly two orderings $<_1$ and $<_2$, such that inversion gives an isomorphism $(\mathbb{Z}, <_1) \xrightarrow{\sim} (\mathbb{Z}, <_2)$.
b) Give an example of an abelian group G admitting orderings $<_1$ and $<_2$ such that $(G, <_1)$ is not isomorphic or anti-isomorphic to $(G, <_2)$.

The comparability quasi-ordering: For $x, y \in G$, we write $x \prec y$ if there exists $n \in \mathbb{Z}^+$ such that $|x| \leq n|y|$. We claim that $\prec$ is a **quasi-ordering** on G, i.e., a reflexive, transitive but not necessarily anti-symmetric binary relation. Indeed the reflexivity is immediate; if $x \prec y$ and $y \prec z$ then there exist $n_1, n_2 \in \mathbb{Z}^+$ such that $|x| \leq n_1|y|$ and $|y| \leq n_2|z|$, and thus $|x| \leq n_1 n_2 |z|$.

As is the case for any quasi-ordering, the relation $x \prec y$ and $y \prec x$ is an equivalence relation, and the quasi-ordering descends to a partial ordering on equivalence classes. Write $x \approx y$ for the resulting equivalence relation on the ordered group G: explicitly, there exist $n_1, n_2 \in \mathbb{Z}^+$ such that $|x| \leq n_1|y|$ and $|y| \leq n_2|x|$.

Exercise: Show that the resulting partial ordering on $G/\approx$ is a total ordering.

In any ordered abelian group G, $\{0\}$ is its own $\approx$-equivalence class, hence any nontrivial ordered abelian group has at least two $\approx$-equivalence classes. We refer to nonzero $\approx$-equivalence classes as **Archimedean equivalence classes** and denote the set of all such equivalence classes as $\Omega(G)$.

An ordered abelian group with $\#\Omega(G) \leq 1$ is called **Archimedean**. Equivalently, for all $x, y \in G^\bullet$, there are $n_1, n_2 \in \mathbb{Z}^+$ such that $|x| \leq n_1|y|$ and $|y| \leq |x|$.

Example: The group $(\mathbb{R}, +)$ is Archimedean. That is, for any $x \in \mathbb{R}^{>0}$ there are positive integers n_1 and n_2 such that $\frac{1}{n_1} \leq x \leq n_2$. Indeed the second inequality follows from the least upper bound axiom: if this were not the case then the set $\mathbb{Z}^+$ of positive integers would be bounded above in $\mathbb{R}$, and this set cannot have a least upper bound. The first inequality follows from the second upon taking reciprocals.

Example: A subgroup of an Archimedean ordered abelian group is Archimedean. In particular, any subgroup of $(\mathbb{R}, +)$ is Archimedean in the induced ordering.

Rather remarkably, the converse is also true.

Theorem 15.6. *(Hölder* [Hö01]*) Let $(G, +)$ be an ordered abelian group. If G is Archimedean, there exists an embedding $(G, +) \hookrightarrow (\mathbb{R}, +)$.*

Proof. We may assume G is nontrivial. Fix any positive element x of G. We will construct an order embedding of G into $\mathbb{R}$ mapping x to 1.

Namely, let $y \in G$. Then the set of integers n such that $nx \leq y$ has a maximal element n_0. Put $y_1 = y - n_0 x$. Now let n_1 be the largest integer n such that $nx \leq 10y_1$: observe that $0 \leq n_1 < 10$. Continuing in this way we get a set of integers $n_1, n_2, \ldots \in \{0, \ldots, 9\}$. We define $\varphi(y)$ to be the real number $n_0 + \sum_{k=1}^\infty \frac{n_k}{10^k}$. It is not hard to show that φ is isotone – $y \leq y' \implies \varphi(y) \leq \varphi(y')$ – and also that φ is injective: we leave these tasks to the reader.

But let us check that φ is a homomorphism of groups. For $y \in G$, and $r \in \mathbb{Z}^+$, let $\frac{n}{10^r}$ be the rational number obtained by truncating $\varphi(y)$ at r decimal places. The numerator n is characterized by $nx \leq 10^r y < (n+1)x$. For $y' \in G$, if $n'x \leq 10^r y' \leq (n'+1)x$, then

$$(n + n')x \leq 10^r(y + y') < (n + n' + 2)x,$$

so

$$\varphi(y + y') - (n + n')10^{-r} < \frac{2}{10^r}$$

and thus

$$|\varphi(y + y') - \varphi(y) - \varphi(y')| < \frac{4}{10^r}.$$

Since r is arbitrary, we conclude $\varphi(y + y') = \varphi(y) + \varphi(y')$. □

Proposition 15.7. *Let G be an Archmidean ordered abelian group. Then exactly one of the following holds:*
(i) G is trivial.
(ii) G is order-isomorphic to $\mathbb{Z}$.
(iii) The ordering on G is dense.

Proof. We may suppose that G is nontrivial.
Step 1: Suppose G^+ has a least element x. Let $y \in G^+$. Since the ordering is

Archimedean there is a largest $n \in \mathbb{Z}^+$ such that $nx \leq y$. Then $y - nx \geq 0$; if $y > 0$ then $y - nx \geq x$ so $y \geq (n+1)x$, contradicting the maximality of n. Thus $y = nx$, i.e., every positive element of G^+ is a multiple of x. It follows that there is a unique order isomorphism from G to $(\mathbb{Z}, <)$ carrying x to 1.
Step 2: Suppose G is not isomorphic to $(\mathbb{Z}, <)$, so there is no least positive element. In other words, given any positive element x there exists 0 with $0 < y < x$. Now let $a, b \in G$ with $a < b$. If $0 < y < b - a$ then $a < y < b$. So the ordering is dense. □

Theorem 15.8. *(Pierce) Let $(G, <)$ be an ordered abelian group. Let G^D be the Dedekind completion of the linearly ordered set $(G, <)$.*
a) There is a unique commutative monoid structure $+$ on G^D such that $(G^D, +)$ is an ordered commutative monoid and the natural map $(G, +, <) \hookrightarrow (G^D, +, <)$ is a homomorphism of ordered commutative monoids.
b) The following are equivalent:
(i) G^D is a group.
(ii) $(G, <)$ is Archimedean.

15.1.1. *Hahn Embedding Theorem.*

The following generalization of Hölder's Theorem was proven by H. Hahn.

Theorem 15.9. *(Hahn Embedding Theorem) Let G be an ordered abelian group. Then there is an embedding of ordered abelian groups*

$$\mathfrak{h} : G \hookrightarrow \mathbb{R}^{\Omega(G)}.$$

Here $\Omega(G)$ denotes the Archimedean equivalence classes of $G^\bullet$ and $\mathbb{R}^{\Omega(G)}$ is lexicographically ordered.

15.2. **Introducing Ordered Fields.**

An **ordered ring** is a ring $(R, +, \cdot)$ together with a total ordering $\leq$ on R compatible with the commutative group $(R, +)$ and satisfying the additional property

(OR) $\forall x, y \geq 0$, $xy \geq 0$.

In these notes the ordered rings we will study are ordered fields.

Example: The real numbers $\mathbb{R}$ with the standard $<$ form an ordered field.

Example: Let $F = \mathbb{Q}(\sqrt{2})$. There are two embeddings $F \hookrightarrow \mathbb{R}$ which differ from each other by the nontrivial automorphism of F, which carries $\sqrt{2} \mapsto -\sqrt{2}$. In one of these embeddings, $\sqrt{2}$ goes to the positive real number whose square is 2, and in the other one it goes to the negative real number whose square is 2. Thus the two embeddings give different orderings, and it is easy to check that these are the only two orderings of F.

A homomorphism $f : (F, <) \to (F', <')$ is a monotone field homomorphism: i.e., a field homomorphism such that $x < y \implies f(x) < f(y)$.

Exercise: Let $(K, <)$ be an ordered field and let F be a subfield of K. Denote by $<_F$ the restriction to F of $<$. Show that $(F, <_F)$ is an ordered field and the

inclusion of F into K is an homomorphism of ordered fields.

We denote by $X(K)$ the set of all field orderings on K.

Exercise: Show that there is a natural action of $\mathrm{Aut}(K)$ on $X(K)$. Give an example where the orbit space $\mathrm{Aut}(K)\backslash X(K)$ consists of more than one element.

Proposition 15.10. *Every ordered field $(K, \leq)$ has characteristic 0.*

Proof. Apply Theorem 15.5 to $(K, +)$. □

For a subset $S \subset K$, put $S^\bullet = S \setminus \{0\}$.

We consider the following conditions on a subset P of a field K:

(PO1) $P + P \subset P$, and $PP \subset P$.
(PO2) $\Sigma_\square(K) = \{x_1^2 + \ldots + x_n^2 \mid x_i \in K\} \subset P$.
(PO3) $-1 \notin P$.
(PO3′) $P \cap (-P) = \{0\}$.
(PO3″) $P^\bullet + P^\bullet \subset P^\bullet$.
(PO3‴) $P \neq K$.
(PO4) $P \cup (-P) = K$.

Exercise: Let $P \subset K$ satisfy (PO1) and (PO2).
a) Show: (PO3), (PO3′), and (PO3″) are equivalent conditions on P.
b) Suppose $\mathrm{char}\, K \neq 2$. Show: (PO3‴) and (PO3) are equivalent conditions on P. (Hint: $x = \left(\frac{x+1}{2}\right)^2 - \left(\frac{x-1}{2}\right)^2$.)
c) Suppose $\mathrm{char}\, K = 2$. Show: P satisfies (PO1) and (PO2) iff P is a subfield of K containing K^2.

Exercise: Let $P \subset K$ satisfy (PO1) and (PO4). Show: P satisfies (PO2).

Lemma 15.11. *Let K be a field.*
a) If $\leq$ is a field ordering on K, put $P = K^{\geq 0}$. Then K satisfies (PO1) through (PO4) above, and also $1 \in P$.
b) Let $P \subset K$ satisfy (PO1) through (PO4). Define a relation $\leq$ on K by $x \leq y \iff y - x \in P$. Then $\leq$ is a field ordering on K.

Proof. a) By Proposition X.X, K has characteristic 0. Lemma 146 implies P satisfies (PO1), and Lemma 147 implies P satisfies (PO3′) and (PO4). By Exercise X.X, P satisfies (PO2). Finally, by (PO4), exactly one of 1, -1 lies in P. But if $-1 \in P$, then $(-1)^2 = 1 \in P$, so $1 \in P$.
b) By (PO3) and (PO4), $\leq$ is a total ordering on K. Given $x, y, z \in K$ with $x \leq y$, then $(y + z) - (x + z) = y - x \in P$, so $x + z \leq y + z$: $(K, \leq)$ is an ordered abelian group. Finally, if $x, y \geq 0$ then $x, y \in P$, so by (PO1) $xy \in P$, i.e., $xy \geq 0$. □

In view of this result, we refer to a subset $P \subset K$ satisying (PO1), (PO2) and (PO3) as being an ordering on K, and we often refer to the ordered field (K, P).

Exercise: Let P_1, P_2 be two orderings on a field K. Show: $P_1 \subset P_2 \implies P_1 = P_2$.

The alert reader may now be wondering why we have introduced (PO2) at all since it is implies by the other axioms for an ordering.[14] The reason is that it is a key idea to entertain a more general structure.

A subset $P \subset K$ satisying (PO1), (PO2) and (PO3) is called a **preordering** of K. (Note that our choice of (PO3) instead of (PO3$'''$) shows that a field of characteristic 2 admits no preorderings.)

Exercise X.X: Let T be a preordering on F and $x, y \in T$.
Show that $x, y \in T,\ x + y = 0 \implies x = y = 0$.

A field K is **formally real** if $-1 \notin \Sigma_\square(K)$.

Lemma 15.12. *Suppose* $\operatorname{char} K \neq 2$. *Then the following are equivalent:*
(i) K *is formally real.*
(ii) $\Sigma_\square(K) \subsetneq K$. *(iii)* $\Sigma_\square(K)$ *is a preordering on* K.
(iv) For all $n \geq 1$ *and all* $x_1, \ldots, x_n \in K$, $x_1^2 + \ldots + x_n^2 = 0 \implies x_1 = \ldots = x_n = 0$.

Exercise: Prove it.

Remark: Condition (iv) above makes a connection with quadratic form theory. A quadratic form $q(x) = a_1 x_1^2 + \ldots + a_n x_n^2$ over a field K is **isotropic** if there exists some nonzero $x \in K^n$ with $q(x) = 0$ and otherwise **anisotropic**. Thus by Lemma X.X, a field is formally real iff for each n the sum of n squares form is anisotropic.

Lemma 15.13. *Let* F *be a field such that* $\Sigma_\square(F) \cap (-\Sigma_\square(F)) = \emptyset$ *and* $\Sigma_\square(F) \cup (-\Sigma_\square(F)) = F^\times$. *Then* $P = \Sigma_\square(F)$ *is the unique ordering on* F.

Exercise: Prove Lemma 15.13.

Exercise: Use Lemma 15.13 to show that each of the following fields admits a *unique* ordering: $\mathbb{R}$, $\mathbb{Q}$, the field of constructible numbers.

Proposition 15.14. *If* (F, P) *is an ordered field,* F *is formally real.*

Proof. The contrapositive is clear: if F is not formally real, then -1 is a sum of squares, so it would be – along with 1 – in the positive cone of any ordering. □

It follows that any ordered field has characteristic 0.

Much more interestingly, the converse of Proposition 15.14 is also true. In order to prove this celebrated result we need the following innocuous one.

Lemma 15.15. *Let* F *be a field,* $T \subseteq F$ *a preordering on* F, *and* $a \in F^\times$. *TFAE:*
(i) The set $T[a] := \{x + ya \mid x, y \in T\}$ *is a preordering.*
(ii) a *is* not *an element of* $-T$.

Proof. Since by (PC4) no preordering can contain both a and $-a$, (i) $\implies$ (ii) is clear. Conversely, assume (ii). It is immediate to verify that $T[a]$ satisfies (PC1),

[14]The less than alert reader may now be asleep, and we owe him our apologies: things will liven up shortly!

(PC2) and (PC5), so it suffices to show that there is no $x \in F^\times$ such that x and $-x$ both lie in $T[a]$. If so, we deduce

$$-1 = -x \cdot x \cdot (\frac{1}{x})^2 \in T[a].$$

But now suppose $-1 = x + ya$ for $x, y \in T$. Then $-ya = 1 + x$ is a nonzero element of T, so $a = (-y)^{-2}(-y)(1+x) \in -T$, a contradiction. □

Theorem 15.16. *Let $\mathfrak{t}$ be a preordering on a field K. Then:*
a) $\mathfrak{t}$ is the intersection of all orderings $P \supset \mathfrak{t}$.
b) (Artin-Schreier) If K is formally real, then it admits an ordering.

Proof. a) Step 1: Let $\mathcal{S}$ be the set of all preorderings on K containing $\mathfrak{t}$. The union of a chain of preorderings is again a preordering. Applying Zorn's Lemma, we get a maximal element $T \supset \mathfrak{t}$. By Lemma 15.15 we have that for all $a \in F$, if $-a \notin T$ then $a \in T$, so T satisfies (PO4) and is therefore an order.
Step 2: Let $b \in K \setminus \mathfrak{t}$. We must construct an ordering $P \supset \mathfrak{t}$ with $b \notin P$. But by Lemma 15.15, $\mathfrak{t}[-b]$ is a preordering, which by Step 1 extends to an ordering P, and since $-b \in P$, $b \notin P$.
b) If K is formally real then $\Sigma_\square(K)$ is a preordering on K. In particular, by (PO3) a preordering is a proper subset of K, whereas by part a) if there were no orderings on K then the intersection over all orderings containing $\Sigma_\square(T)$ would be the empty intersection, and thus would equal K. □

The following special case of Theorem X.Xa) is important in its own right.

Corollary 15.17. *(Artin) For $x \in K^\times$, $\mathrm{char}(K) \neq 2$, the following are equivalent:*
(i) For every ordering P on K, $x \in K$.
(ii) The element x is a sum of squares.

Remark: Corollary 15.17 is an important step towards the solution of Hilbert's 17th problem: show that any positive semidefinite polynomial $f \in \mathbb{R}[t_1, \ldots, t_n]$ is a sum of squares of rational functions.

Remark: Corollary 15.17 does *not* extend to all fields of characteristic 2. Indeed, for a field F of characteristic 2, we simply have $\Sigma_\square(F) = F^2$, so every element of F is a sum of squares iff F is perfect. (In no case are there any orderings on F.)

15.3. **Extensions of Formally Real Fields.**

Let L/K be a field extension. If L is formally real, then by Artin-Schreier it admits an ordering P, which restricts to an ordering $\mathfrak{p}$ on K.[15] However, there is a related but much more subtle question: suppose $\mathfrak{p}$ is an ordering on a field K and L/K is an extension field. Can the ordering $\mathfrak{p}$ be extended to L?

An obvious necessary condition is that L be formally real: if not it admits no orderings at all, let alone an extension of $\mathfrak{p}$. But this condition is not sufficient: let $K = \mathbb{R}(t)$. By Example X.X above there is a unique ordering $\mathfrak{p}$ on K extending the unique ordering on $\mathbb{R}$ and such that $x \leq t$ for all $x \in \mathbb{R}$. Take $L = K(\sqrt{-t})$. Clearly $\mathfrak{p}$ does not extend to L, since if so the negative element $-t$ would be a

[15] particular, a subfield of a formally real field is formally real. But that was clear anyway from the definition.

square. However, the element $\sqrt{-t}$ is transcendental over K, so there is a K-algebra automorphism $K(\sqrt{-t}) \to K(t)$, and thus L is certainly formally real.

In general the extension problem for orderings is a rich one with a large literature. But we will give one fundamental and useful result, an extension of the Artin-Schreier Theorem. First:

Lemma 15.18. *For an ordered field $(K, \mathfrak{p})$, an extension L/K, and $c \in L$, TFAE:*
(i) There are $a_1, \ldots, a_n \in \mathfrak{p}^\bullet$ and $x_1, \ldots, x_n \in L$ such that

$$c = a_1 x_1^2 + \ldots + a_n x_n^2.$$

(ii) $c \in \bigcap_{P \supset \mathfrak{p}} \mathbb{P}$, the intersection being over all orderings of L extending $\mathfrak{p}$.

Proof. Let

$$\mathfrak{t} = \{a_1 x_1^2 + \ldots + a_n x_n^2 \mid a_i \in \mathfrak{p},\ x_i \in L\},$$

and note that the desired equivalence can be rephrased as $\mathfrak{t} = \bigcap_{P \supset \mathfrak{p}} P$. Moreover $\mathfrak{t}$ satisfies (PO1) and (PO2), and an ordering P of L contains $\mathfrak{t}$ iff it contains $\mathfrak{p}$.
Case 1: Suppose $-1 \notin \mathfrak{t}$. Then $\mathfrak{t}$ is a preordering, and by Theorem X.X, $\mathfrak{t} = \bigcap_{P \supset \mathfrak{p}} P$.
Case 2: If $-1 \in \mathfrak{t}$, there is no ordering on L extending $\mathfrak{p}$. Then – since K has ordered and thus not of characteristic 2! – by Exercise X.X, we have $\mathfrak{t} = K = \bigcap_{P \supset \mathfrak{p}} P$. □

We are now ready for one of our main results.

Theorem 15.19. *For an ordered field $(K, \mathfrak{p})$ and an extension field L/K, TFAE:*
(i) There is an ordering on L extending $\mathfrak{p}$.
(ii) For all $a = (a_1, \ldots, a_n) \in \mathfrak{p}^n$, the quadratic form

$$q_a(x) = a_1 x_1^2 + \ldots + a_n x_n^2$$

is anisotropic over L: if $x = (x_1, \ldots, x_n) \in L^n$ is such that $q(x) = 0$, then $x = 0$.

Proof. (i) $\implies$ (ii) is immediate.
(ii) $\implies$ (i): If for any $a \in \mathfrak{p}^n$ the quadratic form $q_a(x)$ represents -1, then the form $q_{a,1}(x) = a_1 x_1^2 + \ldots + a_n x_n^2 + x_{n+1}^2$ would be isotropic, contrary to our hypothesis. It follows that

$$-1 \notin \mathfrak{t} = \{a_1 x_1^2 + \ldots + a_n x_n^2 \mid a_i \in \mathfrak{p},\ x_i \in L\},$$

so – as in the proof of the previous result – $\mathfrak{t}$ is a preordering of L containing $\mathfrak{p}$. By Theorem X.X, $\mathfrak{t}$ must extend to at least one ordering of L. □

Exercise: Deduce from Theorem 15.19 that every formally real field L admits an ordering. (Hint: we wrote L, not K!)

We will now deduce several sufficient conditions for extending orderings.

Theorem 15.20. *Let $(K, \mathfrak{p})$ be an ordered field, and let $L = K(\{\sqrt{x}\}_{x \in \mathfrak{p}})$ be the extension obtained by adjoining all square roots of positive elements. Then the ordering $\mathfrak{p}$ extends to L.*

Proof. By Theorem 15.19, it suffices to show that for any $n, r \in \mathbb{Z}^+$ and any $b_1, \ldots, b_r, c_1, \ldots, c_n \in \mathfrak{p}$, if $x_1, \ldots, x_n \in F(\sqrt{b_1}, \ldots, \sqrt{b_r})$ are such that

$$c_1 x_1^2 + \ldots + c_n x_n^2 = 0, \tag{15}$$

then $x_1 = \ldots = x_n = 0$. For any fixed n, we prove this by induction on r. Suppose by induction that the equation $c_1 x_1^2 + \ldots + c_n x_n^2 = 0$ has no nontrivial solutions

over K_{r-1}, and let $(z_1, \ldots, z_n) \in K_r^n$ be a solution to (15). Write $z_i = x_i + \sqrt{b_r} y_i$, with $x_i, y_i \in K_{r-1}$. Then equating "rational parts" in the equation

$$0 = \sum c_i z_i^2 = \sum c_i x_i^2 + \sum b_r c_i y_i^2 + 2 \sum c_i x_i y_i \sqrt{b_r}$$

shows that $(x_1, \ldots, x_n, y_1, \ldots, y_n) \in K_{r-1}^{2n}$ is a solution of

$$c_1 t_1^2 + \ldots + c_n t_n^2 + b_r c_1 t_{n+1}^2 + \ldots + b_r c_n t_{2n}^2 = 0.$$

By induction, $x_1 = \ldots = x_n = y_1 = \ldots = y_n = 0$, i.e., $z_1 = \ldots = z_n = 0$. □

To obtain further results we take a perspective arising from quadratic form theory. Let us say a field extension L/K is **anistropic** if every anisotropic quadratic form $q(x_1, \ldots, x_n) \in K[x_1, \ldots, x_n]$ remains aniostropic when extended to L. (In the algebraic theory of quadratic forms one studies the **Witt kernel** of a field extension: the kernel of the natural ring homomorphism $W(K) \to W(L)$. An anisotropic extension is precisely one in which the Witt kernel is trivial.) From Theorem 15.19 we immediately deduce the following result.

Corollary 15.21. *If $(K, \mathfrak{p})$ is an ordered field and L/K is an anisotropic extension, then the ordering $\mathfrak{p}$ extends to L.*

Exercise: a) Let K be a field and let $\{L_i\}_{i \in I}$ be a directed system of anisotropic extensions of K. Show that $\varinjlim L_i/K$ is an anisotropic extension.
b) Let $(K, \mathfrak{p})$ be an ordered field and L/K a field extension. Suppose that $\mathfrak{p}$ extends to an ordering on any finitely generated subextension of L/K. Show that $\mathfrak{p}$ extends to an ordering on L.

The next results give the two basic examples of anisotropic extensions.

Theorem 15.22. *A purely transcendental extension L/K is anisotropic.*

Proof. Step 0: It suffices to prove that $K(t)/K$ is anisotropic. Indeed, if so then an immediate induction gives that $K(t_1, \ldots, t_n)/K$ is anisotropic, and we finish by applying Exercise X.X.
Step 1: Let K be any field, and let $(f_1, \ldots, f_n) \in K(t)^n$ be an n-tuple of rational functions, not all zero. Then there exists a nonzero rational function f such that $(ff_1, \ldots, ff_n)$ is a **primitive vector** in $K[t]$, i.e., each $ff_i \in K[t]$ and $\gcd(ff_1, \ldots, ff_n) = 1$. Indeed this holds with $K[t]$ and $K(t)$ replaced by any UFD and its fraction field.

Step 2: Let $q = a_1 x_1^2 + \ldots + a_n x_n^2$ be a nonsingular quadratic form over K such that $q_{K(t)}$ is isotropic: that is, there exist rational functions $f_1, \ldots, f_n$, not all zero, such that

$$a_1 f_1^2 + \ldots + a_n f_n^2 = 0.$$

Let $f \in K(t)^\times$ be the rational function as in Step 1; then multiplying through by f^2 we get a primitive polynomial solution, i.e., there exist polynomials $p_1(t), \ldots, p_n(t) \in K[t]$ with $\gcd(p_1(t), \ldots, p_n(t)) = 1$ and

$$a_1 p_1(t)^2 + \ldots + a_n p_n(t)^2 = 0.$$

Now we substitute $t = 0$ (or any value of K): we cannot have $p_1(0) = \ldots = p_n(0) = 0$, because then all of the p_i's would be divisible by t, contradicting primitivity. Therefore $q(p_1(0), \ldots, p_n(0)) = 0$ shows that q is isotropic over K. □

Remark: The proof of Proposition X.X used only that q was a form – i.e., a homogeneous polynomial – not that it was a *quadratic* form. Indeed any system of homogeneous polynomials would work as well, so the argument really shows: if $V_{/K}$ is a projective variety which has a $K(t)$-rational point, then it has a K-rational point.[16]

The following was conjectured by Witt in 1937 and proven by Springer in 1952.[17]

Theorem 15.23. *(Springer* [Sp52]*) Let L/K be a field extension of finite* odd *degree d. Then L/K is anisotropic.*

Proof. We go by induction on the degree, the case $d = 1$ being trivial. Suppose the result holds for all field extensions of odd degree less than d, and L/K be an extension of odd degree d. If L/K had any proper subextension, then we would be done by a dévissage argument. So we may assume in particular that L is monogenic over K: $L = K[x]$. Let $p(t) \in K[t]$ be the minimal polynomial of x. Let q be an anisotropic quadratic form over K which becomes isotropic over L: i.e., there exists an equation

$$q(g_1(t), \ldots, g_n(t)) = h(t)p(t) \tag{16}$$

with polynomials $g_i, h \in K[t]$, not all $g_i = 0$, and $M := \max \deg g_i \leq d - 1$. As in the proof of Proposition 15.22, we may also assume that $(g_1, \ldots, g_n)$ is a primitive vector in $K[t]$. Since q is anisotropic, the left hand side of (16) has degree $2M \leq 2d - 2$, so $\deg h$ is *odd* and at most $d - 2$. In particular, h has an irreducible factor $\tilde{h}$ of odd degree at most $d-2$; let y be a root of $\tilde{h}$ in $\overline{K}$. Taking $t = y$ in (16), we see that $q(g_1(y), \ldots, g_n(y)) = 0$. Note that since $K[t]$ is a PID, the condition $\gcd(g_1, \ldots, g_n) = 1$ is equivalent to the fact that $1 \in \langle g_1, \ldots, g_n \rangle$, which implies that the polynomials $g_1, \ldots, g_n$ remain setwise coprime as elements of $K[y][t]$. In particular, not all $g_i(y)$ are equal to 0, so that $q_{K[y]}$ is isotropic. By induction, this implies that q was isotropic, contradiction! □

Exercise: Let $(K, \mathfrak{p})$ be an ordered field. Show that the formal power series field $K((t))$ admits a unique ordering extending $\mathfrak{p}$ in which $0 < t < x$ for all $x \in K$.

Exercise: In the algebraic theory of quadratic forms it is shown that the Witt kernel of a quadratic extension $L = K(\sqrt{p})/K$ is the principal ideal generated by $\alpha = \langle 1, -p \rangle$: [QF, Thm. II.20]. In other words, it consists of quadratic forms $a_1x_1^2 + \ldots + a_nx_n^2 - pa_1x_1^2 - \ldots - pa_nx_n^2$ for $a_1, \ldots, a_n \in K^\times$. Use this (and induction) to give another proof of Theorem 15.20.

15.4. **The Grand Artin-Schreier Theorem.**

A field F is **real-closed** if it is formally real and admits no proper formally real algebraic extensions. For instance, $\mathbb{R}$ is evidently real-closed since its unique nontrivial algebraic extension is $\mathbb{C} = \mathbb{R}(\sqrt{-1})$, which is not formally real.

Example (Puiseux series): The **Puiseux series** field $\bigcup_{n \in \mathbb{Z}^+} \mathbb{R}((t^{\frac{1}{n}}))$ is real-closed.

[16]The same conclusion holds for arbitrary varieties over any infinite field, or for complete varieties over a finite field. But taking the projective line over $\mathbb{F}_q$ and removing its $\mathbb{F}_q$-rational points shows that *some* hypothesis is necessary!.

[17]According to D. Hoffmann, Artin orally conveyed a proof of Witt's conjecture to Witt in 1939: he calls the result the Artin-Springer Theorem.

The previous examples of real-closed fields F were obtained by showing that F is formally real and $F(\sqrt{-1})$ is algebraically closed. In fact this is a characterization of real-closed fields. In particular the absolute Galois group of a real-closed field is finite and nontrivial. Remarkably, this too is a characterization of real-closed fields! These assertions are part of the following result, one of the most striking and celebrated theorems in all of field theory.

Theorem 15.24. *(Grand Artin-Schreier Theorem) For a field F, TFAE:*
(i) F is formally real and admits no proper formally real algebraic extension.
(ii) F is formally real, every odd degree polynomial over F has a root, and for each $x \in F^\times$, one of x, $-x$ is a square.
(iii) F is formally real and $F(\sqrt{-1})$ is algebraically closed.
(iv) The absolute Galois group of F is finite and nontrivial.

The proofs of $(i) \implies (ii) \implies (iii) \implies (iv)$ follow relatively easily from what we have already done. We give these first and then tackle $(iv) \implies (i)$, the hardest implication.[18]

Proof. (i) $\implies$ (ii): Since an odd degree polynomial has an odd degree irreducible factor, an odd degree polynomial without a root would yield a proper odd degree extension K/F. By Proposition X.X, K would be formally real, contradicting the definition of real closure. Suppose that neither x nor $-x$ is a square. One of them is positive; WLOG say it is x. By Proposition X.X, $F(\sqrt{x})$ is a proper formally real extension field, contradiction.
(ii) $\implies$ (iii) Since F is formally real, certainly $[F(\sqrt{-1}) : F] = 2$. Let $\overline{F}$ be an algebraic clsoure of $F(\sqrt{-1})$: we wish to show that $\overline{F} = F(\sqrt{-1})$. By hypothesis on odd degree polynomials having a root, the absolute Galois group of F is a pro-2-group, and thus so is the absolute Galois group of $F(\sqrt{-1})$. If $F(\sqrt{-1}) \neq \overline{F}$ then, we are entitled to a proper finite extension M of $F(\sqrt{-1})$, which is Galois over $F(\sqrt{-1})$ and has degree a power of 2. By the basic theory of 2-groups together with the Galois correspondence, there must exist a subextension G of $M/F(\sqrt{-1})$ with $[G : F(\sqrt{-1})] = 2$. But we claim that the hypotheses on F imply that $F(\sqrt{-1})$ is quadratically closed. Indeed, let a, b be arbitrary elements of F. We claim that there are $c, d \in F$ such that

$$a + b\sqrt{-1} = (c + d\sqrt{-1})^2.$$

This amounts to the system $a = c^2 - d^2$, $b = 2cd$. Substituting $d = \frac{b}{2c}$, we get the equation $c^2 = a + \frac{b^2}{4c^2}$, or $c^4 - ac^2 - \frac{b^2}{4} = 0$. The quadratic formula gives

$$c^2 = \frac{a \pm \sqrt{a^2 + b^2}}{2}.$$

Since inside the radical we have a sum of squares, the squareroot does exist in F. If we choose the plus sign in the squareroot, it is easy to see that the expression is again non-negative, so we can solve for c in our field F.
(iii) $\implies$ (iv) is immediate. □

Now we begin the proof of (iv) $\implies$ (i), so suppose that F is a field with algebraic closure $\overline{F}$ such that $1 < [\overline{F} : F] < \infty$.

[18]Our proof of $(iv) \implies (i)$ closely follows lecture notes of Keith Conrad.

Step 1: We claim that $\overline{F}/F$ is Galois.

Proof: Certainly $\overline{F}/F$ is normal, so it suffices to show that it is separable. If F has characteristic 0 (which it cannot, in fact, but we haven't shown that yet), then there is nothing to say, so suppose F has characteristic $p > 0$. We claim that the hypotheses imply that F is perfect, and thus that every algebraic extension of F is separable. Indeed, if F is not perfect, then there exists $\alpha \in F \setminus F^p$ and then by Lemma 9.19, the polynomials $t^{p^n} - \alpha$ are irreducible for all $n \in \mathbb{Z}^+$ so $[\overline{F} : F] = \infty$.

Step 2: Let $G = \mathrm{Gal}(\overline{F}/F)$. We wish to show that $\#G = 2$. If not, then by Sylow theory there exists a subgroup H of order either 4 or an odd prime ℓ. We wish to derive a contradiction.

We will consider the cases $\#G = \ell$ a prime number and $\#G = 4$ in turn. First we suppose $\#G = \ell$ and let σ be a generator of the cyclic group G.

Step 3: We claim that the characteristic of F is not equal to ℓ. If it were, then Artin-Schreier theory would apply, so that $\overline{F} = F(\alpha)$, where α is a root of an Artin-Schreier polynomial $t^p - t - a \in F[t]$. We may write any element $b \in \overline{F}$ as

$$b = b_0 + b_1\alpha + \ldots + b_{\ell-1}\alpha^{\ell-1}$$

for unique $b_0, \ldots, b_{\ell-1} \in F$. Thus

$$b^\ell - b = \sum_{i=0}^{\ell-1} b_i^\ell \alpha^{\ell i} - b_i\alpha^i = \sum_{i=0}^{\ell-1} b_i^\ell(\alpha + a)^i - b_i\alpha^i = \left(b_{p-1}^p - b_{p-1}\right)\alpha^{p-1} + O(\alpha^{p-2}),$$

where by $O(\alpha^{p-2})$ we mean a polynomial in α of degree at most $p-2$. Choose $b \in \overline{F}$ such that $b^p - b = a\alpha^{p-1}$, and then equating coefficients of α^{p-1} gives $b_{p-1}^p - b_{p-1} - a = 0$. Since $b_{p-1} \in F$, this contradicts the irreducibility of $t^p - t - a$.

Step 4: Since the characteristic of F is not $\ell = \#G$, $\overline{F}$ contains a primitive ℓth root of unity ζ. Indeed, since $[F(\zeta) : F] \leq \ell - 1$ and $(\ell - 1, \ell) = 1$, we must have $\zeta \in F$. Therefore Kummer Theory applies to give $\overline{F} = F(\gamma)$, where $\gamma^\ell = c \in F$. Choose $\beta \in \overline{F}$ such that $\beta^\ell = \gamma$, so $\beta^{\ell^2} = c$. Thus $\beta^{\ell^2} = \sigma(\beta^{\ell^2}) = (\sigma\beta)^{\ell^2}$, so $\sigma(\beta) = \omega\beta$ with $\omega^{\ell^2} = 1$. Then ω^ℓ, being an ℓth root of unity, lies in F. If $\omega^\ell = 1$, then $(\sigma(\beta))^\ell = \beta^\ell$, so $\sigma(\beta^\ell) = \beta^\ell$ and then $\beta^\ell = \gamma \in F$, contradiction. So ω is a primitive (ℓ^2)th root of unity. It follows easily that there exists $k \in \mathbb{Z}$ such that

$$\sigma\omega = \omega^{1+\ell k}.$$

From $\sigma\beta = \omega\beta$, we get

$$\beta = \sigma^p\beta = \sigma^{\ell-1}\omega\beta = \omega\sigma(\omega)\cdots\sigma^{\ell-1}(\omega)\beta = \omega^{1+(1+\ell k)+\ldots+(1+\ell k)^{\ell-1}}\beta.$$

From this we deduce

$$\sum_{i=0}^{\ell-1} 1 + (1+\ell k) + \ldots + (1+\ell k)^{\ell-1} \equiv 0 \pmod{\ell^2}.$$

Expanding out the binomial and reducing modulo ℓ^2, we get

$$0 \equiv \sum_{i=0}^{\ell-1}(1+i\ell k) \equiv \ell + \frac{(\ell-1)(\ell)}{2}(\ell k) \pmod{\ell^2}.$$

If ℓ is odd, this gives $0 \equiv \ell \pmod{\ell^2}$, a contradiction. When $\ell = 2$, we get

$$2 + 2k \equiv 0 \pmod 4,$$

so that k is odd. In this case ω has order 4 and $\sigma\omega = \omega^{1+2k} = \omega^3$, so $\sigma\omega \neq \omega$ and $\omega \notin F$. Let us write ω as i. In summary: if $\#G$ is prime, then it equals 2, $i \notin F$ and F does not have characteristic 2.

Step 5: Now suppose that $\#G = 4$. Then there exists at least one subextension K of $\overline{F}/F$ with $[\overline{F} : K]$. Then the above reasoning shows that $i \notin K$, hence not in F, but then $F(i)$ is a subfield of $\overline{F}$ with $[\overline{F} : F(i)] = 2$ and containing a 4th root of unity, contradicting the above analysis.

In summary, we have shown so far that if $1 < [\overline{F} : F] < \infty$, then F does not have characteristic 2 and $\overline{F} = F(i)$. It remains to be shown that F is formally real, and this is handled by the following result.

Lemma 15.25. *Let F be a field in which -1 is not a square and such that every element of $F(\sqrt{-1})$ is a square in $F(\sqrt{-1})$. Then:*
a) $\Sigma_\square(F) = F^2$,
b) $\mathrm{char}(F) = 0$, *and*
c) F is formally real.

Proof. Put $i = \sqrt{-1}$. To show part a), it is enough to see that the sum of two squares in $F(i)$ is again a square in $F(i)$. Let $a, b \in F$. By hypothesis, there are $c, d \in F$ such that $(a + bi) = (c + di)^2$, so $a = c^2 - d^2$ and $b = 2cd$ and thus $a^2 + b^2 = (c^2 + d^2)^2$.
b) If F had positive characteristic p, then -1 is a sum of $p - 1$ squares but not itself a square, contradicting part a).
c) Since -1 is not a square, F does not have characteristic 2, and thus by part a) -1 is not a sum of squares and F is formally real. □

The following exercises give strengthenings and variations on the Artin-Schreier theorem.

Exercise X.X: (E. Fried): Let F be a field. Suppose that there exists a positive integer d such that for every irreducible polynomial $P \in K[t]$, $\deg(P) \leq d$. Show that F is real-closed or algebraically closed.

Exercise X.X (Knopfmacher-Sinclair) Let F be a field. Suppose that the set of isomorphism classes of finite-dimensional field extensions of F is finite. Show that F is real-closed or algebraically closed.

Exercise X.X (K. Conrad): A field K is real-closed iff $1 < [K^{\mathrm{sep}} : K] < \infty$.

Exercise X.X (E. Fried): Let C be an algebraically closed field and K a subfield of

C with $K \neq C$. Suppose that C is finitely generated over K. Then K is real-closed and $C = K(\sqrt{-1})$.

Corollary 15.26. *Let R be a real-closed field and K be a subfield of R. Let K' be the algebraic closure of K in R. Then K' is real-closed.*

Proof. Certainly K' is formally real. If $P(t) \in K'[t]$ is an irreducible polynomial of odd degree, then $K'[t]/(P)$ is formally real, so P has a root in R and therefore also in K'. Moreover, if $0 \neq \alpha \in K'$, then exactly one of $\alpha, -\alpha$ is a square in R, so that $t^2 \pm \alpha$ has a root in R and thus in K'. By Theorem 15.24, K' is real-closed. □

15.5. **Sign Changing in Ordered Fields.**

Let $(K, \mathfrak{p})$ be an ordered field, and let $f \in K[t]$ be a polynomial. If for $a, b \in K$ we have $f(a)f(b) < 0$, then we say f **changes sign between a and b**. If such a, b exist we say f **changes sign**.

Lemma 15.27. *Let $(K, \mathfrak{p})$ be an ordered field.*
a) Every odd degree $f \in K[t]$ changes sign.
b) For all $a > 0$, the polynomial $t^2 - a$ changes sign.

Exercise: Prove Lemma 15.27.

Proposition 15.28. *For an ordered field $(F, \mathfrak{p})$, the following are equivalent:*
*(i) (**Polynomial Intermediate Value Theorem**) Let $f \in F[t]$ and let $a < b \in F$ be such that $f(a)f(b) < 0$. Then there is $c \in F$ such that $a < c < b$ and $f(c) = 0$.*
(ii) F is real-closed.

Proof. (i) $\implies$ (ii): Suppose the Polynomial Intermediate Value Theorem holds in F. By Lemma 15.27, every odd degree polynomial $f \in K[t]$ change sign hence has a root. Similarly, if $a \in F^\times$, then either a or $-a$ is positive; without loss of generality $a > 0$, and by Lemma 15.27, $t^2 - a$ changes sign so has a root. Thus there is $b \in F$ with $b^2 = a$. By Theorem 15.24 F is real-closed.
(ii) $\implies$ (i): Without loss of generality we may assume that $f(a) < 0$,, $f(b) > 0$ and that f is monic irreducible. By Theorem 15.24 f has degree 1 or 2. At this point the proof is an amusing callback to high school algebra. If f has degree 1 then it is $f(a) + \left(\frac{f(b)-f(a)}{b-a}\right) x$, so it has a unique root and is moreover increasing, so its unique root must occur in (a, b). Otherwise $f(t) = t^2 + ct + d$, so by the quadratic formula if it does not have a root then $c^2 - 4d < 0$, but then for all $x \in K$, $f(x) = \left(x + \frac{c}{2}\right)^2 + \left(d - \frac{c^2}{4}\right) > 0$, contradiction! □

Proposition 15.29. *Let $(K, \mathfrak{p})$ be an ordered field, and let $f \in K[t]$ be an irreducible polynomial which changes sign. Then the field $L = K[t]/(f)$ admits an ordering extending $\mathfrak{p}$.*

Proof. We go by induction on $n = \deg f$, the base case $n = 1$ being trivial. So suppose $n \geq 2$, that the result holds for all smaller degrees and – seeking a contradiction – that it fails for some irreducible f of degree n. By Theorem 15.19 then there are $a_i \geq 0$ and $f_i \in K[t]$, each of degree at most $n - 1$, such that

$$1 + \sum_i a_i f_i(t)^2 \equiv 0 \pmod{f}$$

and thus there is $0 \neq h \in K[t]$ with $\deg h \leq n-2$ such that

$$1 + \sum_i a_i f_i(t)^2 = f(t)h(t).$$

Plugging in $t = a$ and $t = b$ we find $f(a)h(a) > 0$ and $f(b)h(b) > 0$ and thus $h(a)h(b) < 0$. There must then be at least one irreducible factor $g(t)$ of $h(t)$ such that $g(a)g(b) < 0$. Since

$$\deg g \leq \deg h \leq n-2 < n = \deg f$$

and

$$1 + \sum_i a_i f_i(t)^2 \equiv 0 \pmod{g},$$

this contradicts our induction hypothesis. □

Exercise: Use Proposition 15.29 to deduce new proofs of many (as many as possible!) of the results of § 16.3.

15.6. **Real Closures.**

Proposition 15.30. *For every formally real field K, there exists an algebraic extension K^{rc} which is real-closed.*

Proof. Let $\overline{K}$ be an algebraic closure of K, and consider the partially ordered set of formally real subextensions of $\overline{K}/K$. Since the union of a chain of formally real fields is formally real, Zorn's Lemma applies to give a maximal formally real subextension, which is by definition real-closed. □

Definition: A **real closure** of a formally real field K is a real-closed algebraic extension of K.

Lemma 15.31. *Let K be a field, let R/K be a real-closed extension field of K, and let R_0 be the algebraic closure of K in R. Then R_0 is a real closure of K.*

Exercise: Prove Lemma 15.31.

Thus we have shown the *existence* of real closures for formally real fields. What about uniqueness? By comparison with the case of algebraically closed fields, one might guess that any two real closures of a given formally real field K are isomorphic as K-algebras. However, this is in general very far from being the case!

Example: Let $K = \mathbb{Q}(t)$. There is a unique embedding $\iota : K \to \mathbb{R}$ in which t gets sent to π. Let K_1 be the algebraic closure of $\iota(K)$ in $\mathbb{R}$. On the other hand, let $\iota_2 : K \to \bigcup_n \mathbb{R}(t^{\frac{1}{n}})$ be the natural embedding of K into the Puiseux series field, and let K_2 be the algebraic closure of $\iota_2(K)$ in $\bigcup_n \mathbb{R}(t^{\frac{1}{n}}$. By Corollary 15.26, K_1 and K_2 are both real-closed fields. In particular, they each admit a unique ordering, in which the positive elements are precisely the nonzero squares. However, the ordering on K_1 is Archimedean and the ordering on K_2 is not, since t is an infinitesimal element. Therefore K_1 and K_2 are not isomorphic as fields, let alone as K-algebras.

Theorem 15.32. *Let $(F, \mathfrak{p})$ be an ordered field. Then there is an algebraic extension R/F which is real-closed and such that the unique ordering on R extends $\mathfrak{p}$.*

Proof. Let $K = F(\{\sqrt{x}\}_{x \in \mathfrak{p}})$. By Theorem 15.20, K is formally real, and now by Proposition 15.30, there exists a real-closed algebraic extension R of K. Let $P = \{x^2 \mid x \in R^\times\}$ be the unique ordering on R. Every $x \in \mathfrak{p}$ is a square in K and hence also in R: that is, $\mathfrak{p} \subset P \cap F$. Conversely, if $x \in F^\times \setminus \mathfrak{p}$, then $-x \in \mathfrak{p} \subset P \cap F \subset P$, so that $x \notin P$, hence $x \notin P \cap F$. Thus $P \cap F = \mathfrak{p}$. □

In the above situation, we say R is a real closure of the ordered field $(F, \mathfrak{p})$.

Exercise: Use Theorem 15.32 to give a third proof of X.X and X.X

Theorem 15.33. *(Sylvester) Let $(K, \mathfrak{p})$ be an ordered field, and let (R, P) be a real-closed extension. Let $f \in K[t]$ be a nonzero monic separable polynomial, and put $A = K[t]/(f)$. Let B_f be the* ***trace form*** *on the K-algebra A, i.e., the bilinear form $\langle x, y \rangle = \mathrm{Tr}_{A/K}(x, y)$. Let $C = R(\sqrt{-1})$. Then:*
a) The number of roots of f in R is equal to the signature of B_f.
b) Half the number of roots of f in $C \setminus R$ is equal to the number of hyperbolic planes appearing in the Witt decomposition of B_f.

Proof. Let $f(t) = f_1(t) \cdots f_r(t)$ be the factorization of f over $R[t]$. Since f is separable, the polynomials f_i are distinct, and since $R(\sqrt{-1})$ is algebraically closed, each f_i has degree 1 or 2. Since $A \otimes_K R \cong R[t]/(f)$, the trace form of $A \otimes_K R$ is simply the scalar extension to R of the trace form B_f. Further, by the Chinese Remainder Theorem

$$R[t]/(f) \cong \prod_{i=1}^{r} R[t]/(f_i),$$

so

$$(B_f)/R \cong \bigoplus_{i=1}^{r} B_{f_i}.$$

It is easy to see that if $\deg f_i = 1$ then the trace form is just $\langle 1 \rangle$, whereas the computation at the end of Section 7 shows that when $\deg f_i = 2$ – so that $R[t]/(f_i) \cong C$ – the trace form is congruent to $\langle 2, -2 \rangle \cong \langle 1, -1 \rangle = \mathbb{H}$, the hyperbolic plane. Both parts of the theorem follow immediately. □

Sylvester's Theorem may look rather specific and technical at first glance. Let us explicitly extract from it the following key consequence: let $f \in K[t]$ by a polynomial defined over an ordered field (K, P). Then if f has a root in one real-closed field extending (K, P), it has a root in every real-closed field extending (K, P). This is a very special case of **Tarski's transfer principle**, which a logician would express in the form "The theory of real-closed fields is model complete." Although it is a very special case, it has enough teeth to be the driving force behind the powerful theorems we will now establish.

Theorem 15.34. *Let $(E, P)/(K, \mathfrak{p})$ be an algebraic extension of ordered fields. Let R be a real-closed field, and let $\sigma : K \to R$ be an ordered field embedding. Then there is a unique order embedding $\rho : E \hookrightarrow R$ extending σ.*

Corollary 15.35. *Let (K, P) be an ordered field, and for $i = 1, 2$ let $\sigma_i : (K, P) \to R_i$ be real closures. There is a unique K-algebra isomorphism $\rho : R_1 \to R_2$.*

Proof. Applying Theorem 15.34 with $R_1 = E$ and $R_2 = R$, $\sigma_2 = \sigma$, there is a unique order embedding $\rho : R_1 \to R_2$ extending σ_2. Since $R_2/\rho(R_1)$ is an algebraic

extension of real-closed fields we must have $\rho(R_1) = R_2$. Finally, if $\tau : R_1 \to R_2$ is any K-algebra homomorphism, then for all $\alpha > 0$ in R_1, we have

$$\tau(\alpha) = \tau(\sqrt{\alpha})^2 > 0.$$

Thus τ is order-preserving, so $\tau = \rho$. □

15.7. **Artin-Lang and Hilbert.**

Lemma 15.36. *Let K be real-closed, and let $h_1, \ldots, h_n \in K[t]^\bullet$. Let P be an ordering on $K(t)$. Then there are infinitely many $a \in K$ such that*

$$\forall 1 \leq i \leq n,\ \operatorname{sgn}(h_i) = \operatorname{sgn}(h_i(a)).$$

Proof. Let $h \in K[t]^\bullet$. Then we may write

$$h = u(t - c_1) \cdots (t - c_r) q_1(t) \cdots q_s(t)$$

with $u, c_1, \ldots, c_r \in K^\times$ and $q_j(t)$ a monic irreducible quadratic for all $1 \leq j \leq s$. For any j,

$$q_j(t) = q(t) = t^2 + bt + c = (t + \frac{b}{2})^2 + (c - \frac{b^2}{4}),$$

and since q_j is irreducible over the real-closed field K, $c - \frac{b^2}{4} > 0$. It follows that $q > 0$ and that for all $a \in K$, $q(a) > 0$. Thus

$$\operatorname{sgn} h = \operatorname{sgn} u \prod_{i=1}^{r} \operatorname{sgn}(t - c_i),$$

$$\forall a \in K,\ \operatorname{sgn} h(a) = \operatorname{sgn} u \prod_{i=1}^{r} \operatorname{sgn}(a - c_i).$$

We may thus assume that each h_i is monic and $\prod_{i=1}^n h_i$ has distinct roots in K. Let c be the smallest root which is strictly greater than t, or ∞ if there is no such root. Then for all $a \in (t, c)$, $\operatorname{sgn} h_i(a) = \operatorname{sgn} h_i(t)$ for all i: this is an infinite set. □

Theorem 15.37 (Artin-Lang Homomorphism Theorem)**.** *Let R be a real-closed field, and let $E = R(x_1, \ldots, x_m)$ be a finitely generated field extension. If E is formally real, then there is an R-algebra map $R[x_1, \ldots, x_m] \to R$.*

Proof. Let d be the transcendence degree of E/R. The case $d = 0$ is trivial: then $E = R[x_1, \ldots, x_m] = R$.
Step 1: We reduce to the $d = 1$ case. Indeed, let E' be a subextension of E/R of transcendence degree 1. Let $\mathcal{R}$ be a real-closure of E, and let $\mathcal{R}'$ be the algebraic closure of E' in $\mathcal{R}$, so by Lemma 15.31 $\mathcal{R}'$ is real-closed. Assuming the result in transcendence degree 1, there is a homomorphism of $\mathcal{R}'$-algebras

$$\varphi : \mathcal{R}'[x_1, \ldots, x_m] \to \mathcal{R}'.$$

Then

$$\operatorname{trdeg}(K(\varphi(x_1), \ldots, \varphi(x_m))/K) \leq \operatorname{trdeg}(\mathcal{R}'/K) = \operatorname{trdeg}(E'/K) = \operatorname{trdeg}(E/K) - 1,$$

so by induction on d we may assume there is a K-algebra map $K[\varphi(x_1), \ldots, \varphi(x_m)] \to K$. Composing with the restriction of φ to $K[x_1, \ldots, x_m]$ we get a K-map to K.
Step 2: Suppose $E = K(x, y_1, \ldots, y_r)$, with x transcendental over K and $y_1, \ldots, y_r$

algebraic over K. We want a K-algebra map $K[x, y_1, \ldots, y_r] \to K$. By the Primitive Element Corollary, there is $y \in E$ such that $E = K(x)[y]$; further, we may take y to be integral over $K[x]$. Then:

$$\exists g_1, \ldots, g_r \in K[x, y],\ h \in K[x]^\bullet \text{ such that } \forall 1 \leq i \leq r,\ y_i = \frac{g_i(x, y)}{h(x)}.$$

If $\varphi : K[x, y] \to K$ is such that $\varphi(h) \neq 0$, then φ induces a K-algebra map $K[x, y_1, \ldots, y_r] \to K$. Thus it is enough to show: there are infinitely many K-algebra maps $\varphi : K[x, y] \to K$. Indeed, if $0 = \varphi(h) = h(\varphi(x))$, then φ maps x to one of the finitely many roots of h in K; since y is algebraic over $K(x)$, having fixed $\varphi(x)$ there are only finitely many choices for $\varphi(y)$.

Step 3: Let

$$f = (x, Y) = Y^n + c_{n-1}(x)Y^{n-1} + \ldots + c_0(x)$$

be the minimal polynomial for y over $K(x)$. Since y is integral over $K[x]$, we have $c_i(x) \in K[x]$ for all i. For $a \in K$, put $f_a(Y) = f(a, Y) \in K[Y]$. We look for roots of f_a in K. For if $b \in K$ is such that $f_a(b) = f(a, b) = 0$, there is a unique K-algebra map $\varphi : K[x, y] \to K$ with $\varphi(x) = a$, $\varphi(y) = b$. So it is enough to show: there are infinitely many $a \in K$ such that there is $b \in K$ with $f_a(b) = 0$.

Step 4: Finally we use that E is formally real! Let P be an ordering on E and let R be a real-closure of (E, P). Then $f(x, Y) \in K[x][Y]$ has a root in R, namely $y \in E \subset R$. By Sylvester's Theorem, $\operatorname{sgn}(B_f)/K(x) > 0$. If we can show that there are infinitely many $a \in K$ such that $\operatorname{sgn}((B_f)_a)/K) > 0$, then applying Sylvester's Theorem again we will get infinitely many a such that $f_a(Y)$ has a root in K and be done. We may diagonalize the quadratic form corresponding to B_f as $\langle h_1(x), \ldots, h_n(x)\rangle$, say. Staying away from the finitely many a such that $h_i(a)$ is zero or undefined for some i, we have that $B_{f_a} \cong \langle h_1(a), \ldots, h_n(a)\rangle$. By Lemma 15.36 there are infinitely many a such that $\operatorname{sgn} B_{f_a} = \operatorname{sgn} B_f > 0$, and we're done. □

Actually Lang proved a stronger result, giving in particular a necessary and sufficient condition for E to be formally real. His result uses the language of arithmetic geometry, so unfortunately will probably not be accessible to all readers of these notes, but here it is anyway.

Theorem 15.38 (Lang [La53])**.** *Let $V_{/R}$ be a geometrically integral algebraic variety over a real-closed field R, with function field $E = R(V)$. Then E is formally real iff V has a* nonsingular *R-point.*

The Artin-Lang homomorphism theorem is powerful enough to yield a quick proof of the following result, which when one takes $K = R = \mathbb{R}$, was the 17th of Hilbert's Problems proposed to the worldwide mathematical community in 1900.

Theorem 15.39 (Artin)**.** *Let K be a formally real field admitting a unique ordering, and let R be a real closure of K. If $f \in K[t_1, \ldots, t_m]$ is such that*

$$f(a_1, \ldots, a_n) \geq 0\ \forall (a_1, \ldots, a_n) \in R^n,$$

then f is a sum of squares in $K(t_1, \ldots, t_m)$.

Proof. We argue by contraposition: suppose $f \in K[t_1, \ldots, t_m]$ is not a sum of squares in $K(t_1, \ldots, t_m)$. By Corollary 15.17, there is an ordering P on $E = K(t_1, \ldots, t_m)$ such that $f <_P 0$. Let $\mathcal{R}$ be a real closure of (E, P). Then $f < 0$

in $\mathcal{R}$, so there is $w \in \mathcal{R}$ with $w^2 = -f$. By Lemma 15.31, the algebraic closure R_0 of K in $\mathcal{R}$ is real-closed, hence is a real-closure of the *ordered* field K since K admits exactly one ordering. By uniqueness of real closures $R_0 = R$. The field $R(t_1, \ldots, t_m, w)$ is a subfield of the real-closed field $\mathcal{R}$, hence by Artin-Lang there is an R-algebra map

$$\varphi : R[t_1, \ldots, t_n, w, \frac{1}{w}] \to R.$$

Note that the effect of including $\frac{1}{w}$ is that $\varphi(w)\varphi(\frac{1}{w}) = 1$, hence $\varphi(w) \neq 0$. For $1 \leq i \leq n$, put $a_i = \varphi(t_i)$; then $(a_1, \ldots, a_n) \in R^n$ and

$$f(a_1, \ldots, a_n) = \varphi(f) = -\varphi(w)^2 < 0.$$

□

Exercise: Let (K, P) be an ordered field with real-closure R. Suppose $f \in K[t_1, \ldots, t_n]$ has the property that $f(a) \geq 0$ for all $a \in R^n$. Show that there is a positive definite quadratic form $q_{/K}$ such that q represents f over $K(t_1, \ldots, t_n)$: there are $x_1, \ldots, x_n \in K(t_1, \ldots, t_n)$ such that $q(x_1, \ldots, x_n) = f$.

15.8. **Archimedean and Complete Fields.**

As usual, a subset S of an ordered field F is called **bounded above** if there exists a single element $x \in F$ such that $s \leq x$ for all $s \in S$; **bounded below** is defined similarly.

An ordered field F is **Archimedean** if the subfield $\mathbb{Q}$ is *not* bounded above.

Example: If x is any rational number, then $x + 1$ is a larger rational number. Thus the field $\mathbb{Q}$ is Archimedean.

Exercise X.X: Show that all of the following conditions on an ordered field are equivalent to the Archimedean property.
(i) For all $x \in F$, there exists $n \in \mathbb{Z}^+$ with $n > x$.
(ii) If $x, y \in F$ with $x > 0$, then there eixsts $n \in \mathbb{Z}^+$ with $nx > y$.
(iii) If $x \in F$ is non-negative and such that $x < \frac{1}{n}$ for all $n \in \mathbb{Z}^+$, then $x = 0$.

An ordered field is **non-Archimedean** if it is not Archimedean.

Exercise: Show that any subfield of an Archimedean ordered field is Archimedean, but a subfield of a non-Archimedean ordered field may be Archimedean.

An element x of an ordered field is **infinitely large** if $x > n$ for all $n \in \mathbb{Z}^+$ and **infinitesimal** if $0 < x < \frac{1}{n}$ for all $n \in \mathbb{Z}^+$. Thus x is infinitely large iff $\frac{1}{x}$ is infinitesimal, and by Exercise X.X, the ordering is non-Archimedean iff infinitely large elements exist iff infinitsimal elements exist.

Exercise X.X: Suppose x is an infinitely large element of an ordered field. Show that for all $y \in \mathbb{Q}$, $x - y$ is infinitely large.

Exercise X.X: Let K be an ordered field; consider the rational function field $K(t)$.
a) Observe that Proposition **??** shows that $K(t)$ admits at least one non-Archimedean

ordering. Show that in fact $K(t)$ admits at least four non-Archimedean orderings. Can you improve upon 4?
b) Use part a) to show that for every infinite cardinal κ, there exists a non-Archimedean ordered field of cardinality κ.

A partially ordered *set* $(S, \leq)$ is **Dedekind complete** if every nonempty subset which is bounded above has a least upper bound.

Exercise X.X: Show that a partially ordered set is Dedekind complete iff every subset which is bounded below has a greatest lower bound.

Proposition 15.40. *Let F be a Dedekind complete ordered field. Then the ordering is Archimedean.*

Proof. We go by contraposition: if F is non-Archimedean, then the subset $\mathbb{Z}^+$ is bounded above, and the set of upper bounds is precisely the set of infinitely large elements. However, Exercise X.X shows in particular that the set of infinitely large elements has no least element: if x is infinitely large, so is $x - 1$. □

Famously, $\mathbb{R}$ satisfies the least upper bound axiom, i.e., its ordering is Dedekind complete. So by Proposition 15.40 the ordering on $\mathbb{R}$ is Archimedean. (Probably the reader was not in doubt of this, but this is an especially clean approach.) Thus every subfield of $\mathbb{R}$ is Archimedean.

The order topology: let $(S, \leq)$ be any linearly ordered space. Recall that we can use the ordering to endow S with a topology, the **order topology**, in which a base of open sets consists of all open intervals.[19] Order topologies have several pleasant properties: for instance, any order topology is a hereditarily normal space (i.e., every subspace is normal: for us, this includes Hausdorff).

Proposition 15.41. *Let K be an ordered field. Then the order topology endows K with the structure of a topological field. That is, the addition and multiplication operations are continuous as functions from $K \times K$ to K.*

Exercise: Prove Proposition 15.41. (Suggestion: use the characterization of continuous functions as those which preserve limits of nets.)

Proposition 15.42. *For any Archimedean ordered field F, $\mathbb{Q}$ is dense in the order topology on F.*

Proof. It is sufficient to show that for $a, b \in F$ with $0 < a < b$, there exists $x \in \mathbb{Q}$ with $a < x < b$. Because of the nonexistence of infinitesimals, there exist $x_1, x_2 \in \mathbb{Q}$ with $0 < x_1 < a$ and $0 < x_2 < b - a$. Thus $0 < x_1 + x_2 < b$. Therefore the set $S = \{n \in \mathbb{Z}^+ \mid x_1 + nx_2 < b\}$ is nonempty. By the Archimedean property S is finite, so let N be the largest element of S. Thus $x_1 + Nx_2 < b$. Moreover we must have $a < x_1 + Nx_2$, for if $x_1 + Nx_2 \leq a$, then $x_1 + (N+1)x_2 = (x_1 + Nx_2) + x_2 < a + (b - a) = b$, contradicting the definition of N. □

Exercise: Deduce from Proposition 15.42 that the order topology on any Archimedean ordered field is second countable. (Hint: show in particular that open intervals with

[19] If there is a bottom element $\mathfrak{b}$ of S, then the intervals $[\mathfrak{b}, b)$ are deemed open. If there is a top element $\mathfrak{t}$ of S, then the intervals $(a, \mathfrak{t}]$ are deemed open. Of course, no ordering on a field has either top or bottom elements, so this is not a relevant concern at present.

rational endpoints form a base for the topology.) From the normality of all order topologies cited above and Urysohn's Metrization Theorem, it follows that the order topology on an Archimedean ordered field is metrizable.[20]

In particular the order topology on K endows $(K, +)$ with the structure of a commutative topological group. In such a situation we can define Cauchy nets, as follows: a net $x_\bullet : I \to G$ in a commutative topological group G is **Cauchy** if for each neighborhood U of the identity $0 \in G$ there exists $i \in I$ such that for all $j, k \geq i$, $x_j - x_k \in U$. A topological group is **complete** if every Cauchy net converges.

Let F be an ordered field. We define the absolute value function from F to $F^{\geq 0}$, of course taking $|x|$ to be x if $x \geq 0$ and $-x$ otherwise.
Exercise: Let F be an ordered field. Show that the triangle inequality holds: for all $x, y \in F$, $|x + y| \leq |x| + |y|$.

Thus for any ordered field F, one can define the function $\rho : F \times F \to F^{\geq 0}$ by $\rho(x, y) = |x - y|$ and this has all the formal properties of a metric except that it is F-valued. In particular, for any net $x_\bullet$ in F we have $x_\bullet \to x$ iff $|x_\bullet - x| \to 0$. In general it can be of some use to consider "F-valued metrics" where F is a non-Archimedean ordered field. But here is the key point: if the ordering on F is Archimedean, then the convergence can be expressed by inequalities involving rational numbers (rather than the infinitesimal elements that would be required in the non-Archimedean case): namely, for an Archimedean ordered field F, a net $x_\bullet : I \to F$ converges to $x \in F$ iff for all $n \in \mathbb{Z}^+$, there exists $i_n \in I$ such that $j \geq i \implies |x_j - x| < \frac{1}{n}$. Topologically speaking, we are exploiting the fact that the topology of an Archimedean ordered field has a countable neighborhood base at each point. Thus it is sufficient to replace nets by sequences. In particular we have the following simple but important result.

Lemma 15.43. *Let K be an Archimedean ordered field. Then TFAE:*
(i) Every Cauchy net in K is convergent.
(ii) Every Cauchy sequence in K is convegent.

Proof. Of course (i) $\implies$ (ii). Now suppose that every Cauchy sequence in K converges, and let $x_\bullet : I \to K$ be a Cauchy net. We may assume that I has no maximal element, for otherwise the net is certainly convergent. Choose $i_1 \in I$ such that $j, k \geq i_1$ implies $|x_j - x_k| < 1$. Now pick $i_2 \in I$ such that $i_2 > I_1$ and $j, k \geq i_2$ implies $|x_j - x_k| < \frac{1}{2}$. Continuing in this manner we get an increasing sequence $\{i_n\}$ in I such that for all n, if $j, j \geq i_n$, $|x_j - x_k| < \frac{1}{n}$. Thus from the net we have extracted a Cauchy subsequence, which by hypothesis converges, say to x. From this it follows immediately that the net $x_\bullet$ converges to x. □

Remark: The proof here is based on [Wi, Thm. 39.4], which asserts that the uniform structure associated to a complete metric is a complete uniform structure iff the metric is a complete metric.

Theorem 15.44. *For an Archimedean ordered field K, TFAE:*
(i) The ordering on K is Dedekind complete: every nonempty subset which is

[20]However, we are not going to use this fact in our discussion. Rather, as will become clear, an ordered field K comes with a canonical "K-valued metric", which will be just as useful to us as an "$\mathbb{R}$-valued metric" – a special case!

bounded below has a greatest lower bound.
(ii) $(K,+)$ *is a Cauchy-complete topological group: every Cauchy net converges.*

Proof. (i) $\implies$ (ii): Dedekind complete implies Archimedean implies second countable implies first countable implies it is enough to look at Cauchy sequences. The argument is then the usual one from elementary real analysis: suppose K is Dedekind complete, and let x_n be a Cauchy sequence in K. Then the sequence is bounded, so there exists a least upper bound x. We can construct a subsequence converging to x in the usual way: for all $k \in \mathbb{Z}^+$, let x_{n_k} be such that $|x_{n_k} - x| < \frac{1}{n}$. (That this implies that the subseuence converges is using the the Archimedean property that for all $x > 0$, there exists $n \in \mathbb{Z}^+$ with $\frac{1}{n} < x$.) Then, as usual, a Cauchy sequence with a convergent subsequence must itself be convergent.
(ii) $\implies$ (i): let $S \subset K$ be nonempty and bounded below. Let $\mathcal{B}$ be the set of all lower bounds of S, with the ordering induced from K. What we want to show is that $\mathcal{B}$ has a greatest element: we will prove this by Zorn's Lemma. Let $\mathcal{C}$ be a nonempty chain in $\mathcal{B}$. We may view this as a net $x : \mathcal{C} \to K$. We claim that it is Cauchy: i.e., for every open neighborhood U of 0, there exists an index i such that for all $j, k \geq i$, $x_i - x_j \in U$. Because the ordering is Archimedean, this is equivalent to $|x_i - x_j| < \epsilon$ for some positive rational number ϵ. But since $\mathcal{C}$ is a set of lower bounds for the nonempty set S, it is certainly bounded above, and if the desired conclusion were false there would exist infinitely many pairs of indices (i,j) with $j > i$ and $x_j - x_i \geq \epsilon$, and by the Archimedean nature of the ordering this would imply that $\mathcal{C}$ is unbounded above, contradiction! Therefore the net $x_\bullet$ is Cauchy and converges by assumption to $x \in K$. This element x is an upper bound for $\mathcal{C}$ and a lower bound for S. Thus by Zorn's Lemma $\mathcal{B}$ has a maximal element, i.e., S has a greatest lower bound. □

An Archimedean ordered field satisfying the equivalent conditions of Theorem 15.44 will simply be said to be **complete**.

Proposition 15.45. *(Strong Rigidity for Archimedean ordered fields) Let K be an Archimedean ordered field and let $f : K \to K$ be an endomorphism, i.e., an order-preserving field homomorphism from K to itself. Then $f = 1_K$ is the identity map.*

Proof. Suppose not, and let $x \in K$ be such that $f(x) \neq x$. Without loss of generality we may suppose that $x < f(x)$, and then by Proposition 15.42 there exists $q \in \mathbb{Q}$ with $x < q < f(x)$. Applying the isotone map f gives $f(x) < f(q) = q$, a contradiction! □

Lemma 15.46. *Let R and S be topological rings and D a dense subring of R. Suppose that $f : R \to S$ is a continuous* set map *from R to S which upon restriction to D is a homomorphism of rings. Then f is itself a homomorphism of rings.*

Exercise: Prove Lemma 15.46. (Hint: use the net-theoretic characterization of dense subspaces: for any $x \in R$, there exists a net $x_\bullet : I \to D$ which converges to x.)

Theorem 15.47. *(Main Theorem on Archimedean Ordered Fields)*
A complete Archimedean field R is a final object in the category of fields. That is:
(i) For any Archimedean field K and Dedekind complete field R, there exists a unique embedding of ordered fields $K \hookrightarrow R$.
(ii) Any two Dedekind complete fields are canonically – even uniquely! – isomorphic.

Proof. (i) The idea here is that we have copies of $\mathbb{Q}$ inside both K and L and that in an Archimedean ordered field an element is uniquely specified by all of its order relations with elements of $\mathbb{Q}$. Formally, we define a map $\varphi : K \to L$ as follows: we map x to $\sup\{q \in \mathbb{Q} \mid q < x\}$. As above, it is clear that φ is order-preserving. When restricted to the dense subring $\mathbb{Q}$ it is certainly a homomorphism, so in order to apply Lemma 15.46 we need only check that φ is continuous. But again, a base for the topology of any Archimedean field is given by open intervals (a, b) with $a, b \in \mathbb{Q}$. Evidently φ maps the interval (a, b) of K to the interval (a, b) of L, so it is therefore continuous: done.
(ii) Let R_1 and R_2 be complete Archimedean fields. By (i), there exist embeddings of ordered fields $\varphi : R_1 \to R_2$ and $\varpi : R_2 \to R_1$. Applying Proposition 15.45 to the endomorphisms $\varpi \circ \varphi$ and $\varpi \circ \varphi$, we get $\varpi \circ \varphi = 1_{R_1}$ and $\varphi \circ \varpi = 1_{R_2}$, thus ϖ and φ are mutually inverse isomorphisms: so $R_1 \cong R_2$ as ordered fields. Moreover the same argument applies to show that any two isomorphisms φ_1, φ_2 from R_1 to R_2 are inverses of the isomorphism ϖ, so $\varphi_1 = \varphi_2$: there is only one isomorphism from R_1 to R_2. □

We have already identified the real numbers $\mathbb{R}$ as a complete Archimedean field, so we know that the final object referred to in Theorem 15.47 indeed exists. Let us restate things in a more concrete fashion using $\mathbb{R}$.

Corollary 15.48. *For any Archimedean ordered field K, there is a unique embedding of ordered fields $K \hookrightarrow \mathbb{R}$. Thus we may identify the Archimedean ordered fields – up to unique isomorphism – as precisely the subfields of $\mathbb{R}$ with the inherited ordering.*

It may be worth asking at this point: exactly how do we know that this field "of real numbers" we've heard so much about actually exists? We've proven some fairly remarkable facts about it: maybe rumors of its existence are greatly exaggerated!

The previous paragraph is silly. A rigorous construction of $\mathbb{R}$ was first given by R. Dedekind in the late 19th century. Accounts of his method (using what are now called) "Dedekind cuts" may be found in many texts. However, our Cauchy-theoretic perspective also gives an easy answer to this question. Namely, one has the notion of **Cauchy completion** of any commutative topological group G: namely, given G there exists a complete topological group $\hat{G}$ and a homomorphism of topological groups $G \to \hat{G}$ which is *universal* for homomorphisms from G into a complete topological group (If G is Hausdorff the map to the completion is an embedding.) The construction can be given in terms of an equivalence relation on the class of Cauchy nets on G, for instance. Moreover, when G is the additive group of an ordered field F, it is not hard to show that $\hat{F}$ is also an ordered field. Note well that we can therefore construct many **Cauchy complete** non-Archimedean ordered fields. However what we want is a Dedekind complete ordered field, and for this, according to Theorem 15.44 it is sufficient – and clearly also necessary – to complete an *Archimedean* ordered field, like $\mathbb{Q}$.

The construction of the Cauchy completion of a commutative topological group is more abstruse than is necessary for this application, though. As in Lemma 15.43 above, we can get away with Cauchy sequences rather than Cauchy nets. Thus we may construct $\mathbb{R}$ from $\mathbb{Q}$ in the following appealingly algebraic way: take the

ring $\mathcal{C}(\mathbb{Q})$ of all Cauchy sequences in $\mathbb{Q}$ and mod out by the maximal ideal $\mathfrak{c}_0$ of sequences converging to 0. Therefore the quotient is a Cauchy complete field, say $\mathcal{R}$. It is easy to check that the ordering on $\mathcal{Q}$ extends to $\mathcal{R}$ and that $\mathcal{Q}$ is dense in $\mathcal{R}$ in the order topology, which implies that the ordering on $\mathcal{R}$ is Archimedean. Thus $\mathcal{R}$ is a Cauchy complete, Archimedean ordered field, so it is Dedekind complete.

15.9. **The Real Spectrum.**

For a field F, let $X(F)$ be the set of all orderings on F. There is a natural topology on $X(F)$: namely the open sets are given by finite intersections of (subbasic) open sets of the form

$$H(a) = \{P \in X(F) \mid a \in P\}$$

as a ranges through nonzero elements of F: that is, $H(a)$ is the set of orderings which regard a as positive. Note that $H(-a) = X(F) \setminus H(a)$, so that the $H(a)$ and (and hence also all the basis elements) are closed as well open: this implies that $X(F)$ is totally disconnected and Hausdorff. It is also compact. To see this, note that an ordering P of F gives rise to an element of $Y = \{\pm 1\}^{F^\times}$, namely for each nonzero element a, we assign $+1$ if $a \in P$ and -1 if $-a \in P$. Giving $\{\pm 1\}$ the discrete topology and Y the product topology, it is a compact Hausdorff totally disconnected space by Tychonoff's theorem. It remains to be shown first that the topology on $X(F)$ defined above is the same as the topology it gets as a subspace of Y[21], and second that $X(F)$ is closed as a subspace of Y. Neither of these is very difficult and we leave them to the reader.

If $F_1 \hookrightarrow F_2$ is a field embedding, then the aforementioned process of restricting orders on F_2 to orders on F_1 gives a map $X(F_2) \to X(F_1)$ which is easily seen to be continuous.

A topological space which is compact Hausdorff and totally disconnected is often called **Boolean**, since these are precisely the spaces which arise as spectra of maximal ideals of Boolean algebras.

Theorem 15.49. *(Craven* [Cr75]*) Any Boolean space X is homeomorphic to $X(F)$ for some field F.*

The following exercises develop a proof in the special case in which X is second countable. Exercise: Let $F = \varinjlim_\alpha F_\alpha$ be a direct limit (i.e., directed union) of fields. Show $X(F) = \varprojlim_\alpha X(F_\alpha)$ as topological spaces.

Exercise: Let $F/\mathbb{Q}$ be a (possibly infinite) formally real Galois extension. Show that $\mathrm{Aut}(F) = \mathrm{Gal}(F/\mathbb{Q})$ acts continuously and simply transitively on $X(F)$, and conclude that in this case $X(F)$ is homeomorphic to the underlying topological space of a profinite group. In particular, if $F/\mathbb{Q}$ is infinite, $X(F)$ is an infinite profinite space without isolated points and with a countable basis, so is homeomorphic to the Cantor set. (A good example is $F = \mathbb{Q}(\{\sqrt{p}\})$ as p ranges over all the prime numbers: here $\mathrm{Aut}(F) \cong (\mathbb{Z}/2\mathbb{Z})^{\aleph_0}$ really looks like the Cantor set.)

[21]One might wonder why we didn't save ourselves the trouble and define the topology on $X(F)$ in this latter way. It turns out that the sets $H(a)$, called the *Harrison subbasis*, are important in their own right.

Exercise: Use weak approximation of valuations to show that any inverse system

$$\ldots \to S_{n+1} \to S_n \to \ldots \to S_1$$

of finite sets can be realized as the system of $X(F_n)$'s where

$$F_1 \ldots \hookrightarrow F_n \hookrightarrow F_{n+1} \hookrightarrow \ldots$$

is a tower of number fields. Conclude that any profinite space with a countable basis arises as the space of orderings of an algebraic field extension of $\mathbb{Q}$.

References

[Ar44] E. Artin, *Galois Theory*. Second edition. Notre Dame Mathematical Lectures, no. 2. University of Notre Dame, Notre Dame, Ind., 1944.

[BAI] N. Jacobson, *Basic algebra. I.* Second edition. W. H. Freeman and Company, New York, 1985.

[BAII] N. Jacobson, *Basic algebra. II.* Second edition. W. H. Freeman and Company, New York, 1989.

[Bar51] D. Barbilian, *Solution exhaustive du problème de Steinitz*. Acad. Repub. Pop. Romîne. Stud. Cerc. Mat. 2, (1951). 195–259 (misprinted 189–253).

[BJ01] F. Borceux and G. Janelidze, *Galois theories*. Cambridge Studies in Advanced Mathematics, 72. Cambridge University Press, Cambridge, 2001.

[CW50] J.W.S. Cassels and G.E. Wall, *The normal basis theorem*. J. London Math. Soc. 25 (1950), 259–264.

[Ch70] A. Charnow, *The automorphisms of an algebraically closed field*. Canad. Math. Bull. 13 (1970), 95–97.

[Cr75] T.C. Craven, *The Boolean space of orderings of a field.* Trans. Amer. Math. Soc. 209 (1975), 225–235.

[DG94] M. Dugas and R. Göbel, *Automorphism groups of fields*. Manuscripta Math. 85 (1994), no. 3-4, 227–242.

[DG97] M. Dugas and R. Göbel, *Automorphism groups of fields. II.* Comm. Algebra 25 (1997), 3777–3785.

[Di74] J. Dieudonné, *Sur les automorphismes des corps algébriquement clos*. Bol. Soc. Brasil. Mat. 5 (1974), 123–126.

[FK78] E. Fried and J. Kollar, *Automorphism groups of algebraic number fields*. Math. Z. 163 (1978), 121–123.

[Gi68] R. Gilmer, *Classroom Notes: A Note on the Algebraic Closure of a Field.* Amer. Math. Monthly 75 (1968), 1101-1102.

[Hö01] O. Hölder, *Die Axiome der Quantitat und die Lehre vom Mass*. Ber. Verh. Sachs. Ges. Wiss. Leipzig, Math.-Phys. Cl. 53 (1901), 1–64.

[Is80] I.M. Isaacs, *Roots of polynomials in algebraic extensions of fields*. Amer. Math. Monthly 87 (1980), 543–544.

[Ja1] N. Jacobson, *Basic algebra. I.* Second edition. W. H. Freeman and Company, New York, 1985.

[Ja2] N. Jacobson, *Basic algebra. II.* Second edition. W. H. Freeman and Company, New York, 1989.

[Je] T.J. Jech, *The axiom of choice*. Studies in Logic and the Foundations of Mathematics, Vol. 75. North-Holland Publishing Co., New York, 1973.

[Kap95] I. Kaplansky, *Fields and rings.* Reprint of the second (1972) edition. Chicago Lectures in Mathematics. University of Chicago Press, Chicago, IL, 1995.

[Ka89] G. Karpilovsky, *Topics in field theory*. North-Holland Mathematics Studies, 155. Notas de Matematica [Mathematical Notes], 124. North-Holland Publishing Co., Amsterdam, 1989.

[Kr53] W. Krull, *Über eine Verallgemeinerung des Normalkörperbegriffs*. J. Reine Angew. Math. 191 (1953), 54–63.

[Ku65] W. Kuyk, *The construction of fields with infinite cyclic automorphism group*. Canad. J. Math. 17 (1965), 665–668.

[La53] S. Lang, *The theory of real places.* Ann. of Math. (2) 57 (1953), 378–391.
[LaFT] S. Lang, *Algebra.* Revised third edition. Graduate Texts in Mathematics, 211. Springer-Verlag, New York, 2002.
[Las97] D. Lascar, *The group of automorphisms of the field of complex numbers leaving fixed the algebraic numbers is simple. Model theory of groups and automorphism groups (Balubeuren, 1995)*, 110-114, London Math. Soc. Lecture Note Ser., 244, Cambridge Univ. Press, Cambridge, 1997.
[Le55] H. Leptin, *Ein Darstellungssatz für kompakte, total unzusammenhngende Gruppen.* Arch. Math. (Basel) 6 (1955), 371–373.
[Lev43] F.W. Levi, *Contributions to the theory of ordered groups.* Proc. Indian Acad. Sci., Sect. A. 17 (1943), 199–201.
[LoI] F. Lorenz, *Algebra. Vol. I. Fields and Galois theory.* Translated from the 1987 German edition by Silvio Levy. With the collaboration of Levy. Universitext. Springer, New York, 2006.
[LoII] F. Lorenz, *Algebra. Vol. II. Fields with structure, algebras and advanced topics.* Translated from the German by Silvio Levy. With the collaboration of Levy. Universitext. Springer, New York, 2008.
[Mac39] S. Mac Lane, *Steinitz field towers for modular fields.* Trans. Amer. Math. Soc. 46, (1939). 23–45.
[Mo96] P. Morandi, *Field and Galois theory.* Graduate Texts in Mathematics, 167. Springer-Verlag, New York, 1996.
[NTII] P.L. Clark, *Algebraic Number Theory II: Valuations, Local Fields and Adeles.* `http://math.uga.edu/~pete/8410FULL.pdf`
[Pa74] T. Parker, *Some applications of Galois theory to normal polynomials.* Amer. Math. Monthly 81 (1974), 1009–1011.
[QF] P.L. Clark, *Lecture notes on quadratic forms.* Available online.
[Rom06] S. Roman, *Field theory.* Second edition. Graduate Texts in Mathematics, 158. Springer, New York, 2006.
[Rot98] J. Rotman, *Galois theory.* Second edition. Universitext. Springer-Verlag, New York, 1998.
[Sc92] B. Schnor, *Involutions in the group of automorphisms of an algebraically closed field.* J. Algebra 152 (1992), 520–524.
[Sp52] T.A. Springer, *Sur les formes quadratiques d'indice zéro.* C. R. Acad. Sci. Paris 234 (1952), 1517–1519.
[St10] E. Steinitz, *Algebrasiche Theorie der Körper.* Journal für die reine und angewandte Mathematik, 1910.
[Su99] B. Sury, *On an example of Jacobson.* Amer. Math. Monthly 106 (1999), 675-676.
[Wa74] W.C. Waterhouse, *Profinite groups are Galois groups.* Proc. Amer. Math. Soc. 42 (1974), 639–640.
[We09] S.H. Weintraub, *Galois theory.* Second edition. Universitext. Springer, New York, 2009.
[Wi] S. Willard, *General topology.* Reprint of the 1970 original. Dover Publications, Inc., Mineola, NY, 2004.
[Ya66] P.B. Yale, *Automorphisms of the Complex Numbers.* Math. Magazine 39 (1966), 135–141.