Chapter 4

Fields and Galois Theory

4.1 Introduction

The rational, real, complex and, much later, the finite fields were the basic inspiration for the study of fields in general. Their ideal theory and the module theory (vector spaces) over them are very simple; so, it was natural to look more deeply inside them. In particular, one can consider solutions of polynomial equations in a field, the automorphisms of a field, the relation of one field to another. We owe to E. Galois the capital idea of applying symmetry in the form of group theory to the study of polynomial equations (coefficients in a field) and their solutions in a (perhaps bigger) field. He was preceded in partial results by such figures as Lagrange, Abel and Gauss and the impetus he provided has sustained the subject until the current day. What concerns one now is not so much the "classical theory" (all of which in smooth modern form is treated below), but questions of basically geometric origin that use an admixture of group theory, ring theory and fields to try to settle vexing questions of apparently "simple" nature. For example, if we adjoin to the rationals all the roots of unity and call the resulting field K, is it true that every homogeneous form of degree $d > 0$ in more than d variables has a non-zero solution in K? This is a conjecture of E. Artin–still open at present.

4.2 Algebraic Extensions

Recall that if A is a commutative ring and B is an over-ring of A (i.e., an A-algebra), an element $\beta \in B$ is *algebraic over* A iff the map $A[X] \longrightarrow A[\beta] \subseteq B$ is *not* injective; the element β is *transcendental over* A iff the map $A[X] \longrightarrow A[\beta]$ is injective. Moreover, $\beta_1, \ldots, \beta_n$ are *independent transcendentals over* A (*algebraically independent over* A) iff $A[X_1, \ldots, X_n] \longrightarrow A[\beta_1, \ldots, \beta_n]$ is injective. The case of interest here is: $A = k$, a field, and B a subring of a field.

Algebraic elements admit of many characterizations:

Proposition 4.1 *Say B is an integral domain containing a field k and $\alpha \in B$. Then, the following are equivalent:*

(1) α is algebraic over k.

(2) $k[\alpha]$ $(\subseteq B)$ is a field.

(3) $k(\alpha) = k[\alpha]$.

(4) $1/\alpha \in k[\alpha]$.

(5) $k[\alpha]$ $(\subseteq B)$ is a finite dimensional k-vector space.

(6) $k[\alpha] \subseteq L$, where $L\,(\subseteq B)$ is a subring of B and L is a finite dimensional k-vector space.

Proof. (1) $\Rightarrow$ (2). By definition there is some polynomial $f \in k[X]$ so that $f(\alpha) = 0$. By unique factorization in $k[X]$, we know that $f = f_1 \cdots f_r$, where each f_j is irreducible. So, $0 = f(\alpha) = \prod_{j=1}^{r} f_j(\alpha)$ and as B is a domain, $f_j(\alpha) = 0$, for some j; so, we may assume that f is irreducible. Look at $k[X]/(f(X))$. Now, as $k[X]$ is a P.I.D and f is irreducible, it follows that $(f(X))$ is a maximal ideal. Thus, $k[X]/(f(X))$ is a field; moreover, $k[\alpha] \cong k[X]/(f(X))$ and (2) holds.

(2) $\Rightarrow$ (3) and (3) $\Rightarrow$ (4) are clear.

(4) $\Rightarrow$ (5). By (4),

$$\frac{1}{\alpha} = \sum_{j=0}^{N} a_j \alpha^j$$

(with $\alpha_N \neq 0$) and this yields $\sum_{j=0}^{N} a_j \alpha^{j+1} = 1$; we deduce

$$\alpha^{N+1} = \frac{1}{a_N} - \sum_{j=0}^{N-1} \frac{a_j}{a_N}\, \alpha^{j+1},$$

i.e., α^{N+1} depends linearly on $1, \alpha, \ldots, \alpha^N$. By an obvious induction, α^{N+i} depends linearly on $1, \alpha, \ldots, \alpha^N$ for all $i \geq 1$ and so, $1, \alpha, \ldots, \alpha^N$ span $k[\alpha]$.

(5) $\Rightarrow$ (6) is a tautology.

(6) $\Rightarrow$ (1). Since $k[\alpha]$ is a subspace of a finite dimensional vector space, $k[\alpha]$ is finite dimensional over k (i.e., (5)). Look at $1, \alpha, \ldots, \alpha^N, \alpha^{N+1}, \ldots$ There must be a linear dependence, so

$$a_N \alpha^N + \cdots + a_1 \alpha + a_0 = 0$$

and α is a root of $f(X) = a_N X^N + \cdots + a_1 X + a_0$. □

Proposition 4.2 *Write $B_{\text{alg}} = \{\alpha \in B \mid \alpha$ is algebraic over $k\}$. Then, B_{alg} is a ring (a domain).*

Proof. Say $\alpha, \beta \in B_{\text{alg}}$. Then, $k[\alpha]$ is finite dimensional over k and $k[\alpha, \beta] = k[\alpha][\beta]$ is finite dimensional over $k[\alpha]$, which implies that $k[\alpha, \beta]$ is finite dimensional over k. As $\alpha \pm \beta$ and $\alpha\beta$ belong to $k[\alpha, \beta]$, by (6), they are algebraic over k. □

Proposition 4.3 *Say $\alpha, \beta \in B_{\text{alg}}$ (with $\beta \neq 0$), then $\alpha/\beta \in B_{\text{alg}}$. Therefore, B_{alg} is actually a field.*

Proof. As before, $k[\alpha, \beta]$ is finite dimensional over $k[\alpha]$. But, $k(\alpha) = k[\alpha]$ and $k[\alpha, \beta] = k[\alpha][\beta]$, so $k[\alpha, \beta] = k(\alpha)[\beta]$. Yet, β is algebraic over $k(\alpha)$; thus, $k(\alpha)[\beta] = k(\alpha)(\beta) = k(\alpha, \beta)$. Consequently, $k[\alpha, \beta] = k(\alpha, \beta)$ and it is finite dimensional over k. As $\alpha/\beta \in k(\alpha, \beta)$, it is algebraic over k. □

Proposition 4.4 *Being algebraic is transitive.*

Given an extension, K/k, the *degree,* $\deg(K/k) = [K\colon k]$, *of* K/k is the dimension of K as a vector space over k. Observe that if $[K\colon k]$ is finite, then K is algebraic over k (for every $\alpha \in K$, there is a linear dependence among $1, \alpha, \ldots, \alpha^n, \ldots$, so, α is the root of some polynomial in $k[X]$). However, an algebraic extension K/k *need not be finite*.

Definition 4.1 Let K/k be a field extension (i.e., $k \subseteq K$ where both are fields and K is a k-algebra). Say $\alpha \in K$ is a root of $f(X) \in k[X]$. Then, α is a *root of multiplicity, m,* iff $f(X) = (X - \alpha)^m g(X)$ in $K[X]$ and $g(\alpha) \neq 0$.

Let A be a commutative ring, B be an A-algebra and C be a B-algebra.

Definition 4.2 An additive map $\delta\colon B \to C$ is an *A-derivation of B with values in C* iff

(1) $\delta(\xi\eta) = \xi\delta(\eta) + \delta(\xi)\eta$ (Leibnitz)

(2) $\delta(\alpha) = 0$ whenever $\alpha \in A$.

Notice that (1) and (2) imply the A-linearity of an A-derivation. The A-derivations of B with values in C form a B-module denoted $\mathrm{Der}_A(B, C)$.

Examples of Derivations.

(1) Let A be a commutative ring, let $B = A[X]$ and let $C = B$.

$$\delta f = \delta\Big(\sum_{j=0}^{N} a_j X^j\Big) = \sum_{j=0}^{N} j a_j X^{j-1} = f'(X)$$

is an A-derivation.

(2) Let A be a commutative ring, $B = A[\{X_\alpha\}_{\alpha\in I}]$, $C = B$ and

$$\delta_\alpha = \frac{\partial}{\partial X_\alpha}.$$

Remark: For Example 1, if h is an independent transcendental from X, we have (DX)

$$f(X + h) = f(X) + f'(X)h + O(h^2).$$

Theorem 4.5 *(Jacobian criterion for multiplicity) Given $f(X) \in k[X]$ and K/k a field extension, for any root α of $f(X)$, we have:*

(1) If the multiplicity of α as a root is $\geq m$, then

$$f(\alpha) = f'(\alpha) = \cdots = f^{(m-1)}(\alpha) = 0.$$

(2) If $\mathrm{char}(k) = 0$ and if $f(\alpha) = f'(\alpha) = \cdots = f^{(m-1)}(\alpha) = 0$ but $f^{(m)}(\alpha) \neq 0$, then α is a root of f of exact multiplicity m.

Proof. We proceed by induction on m. Consider a root, α, of multiplicity 1. This means $f(X) = (X-\alpha)g(X)$ in $K[X]$ and $g(\alpha) \neq 0$. Thus,

$$f'(X) = (X - \alpha)g'(X) + g(X),$$

so, $f'(\alpha) = g(\alpha)$ and $f'(\alpha) \neq 0$. Therefore, (2) holds independently of the characteristic of k in this one case and (1) is trivial.

Now, assume α is a root of multiplicity at least m. As $f(X) = (X - \alpha)^m g(X)$ in $K[X]$, we get

$$f'(X) = (X - \alpha)^{m-1}((X - \alpha)g'(X) + mg(X)),$$

which shows that the multiplicity of α in f' is at least $m - 1$. By the induction hypothesis applied to $f'(X)$, we have $f'(\alpha) = f''(\alpha) = \cdots = f^{(m-1)}(\alpha) = 0$. Also, $f(\alpha) = 0$, so (1) holds.

(2) Again, we proceed by induction. Assume that $f(\alpha) = f'(\alpha) = \cdots = f^{(m-1)}(\alpha) = 0$ but $f^{(m)}(\alpha) \neq 0$, with $\mathrm{char}(k) = 0$. Let q be the exact multiplicity of α. Then, $f(X) = (X-\alpha)^q h(X)$ in $K[X]$, with $h(\alpha) \neq 0$. Now, $f'(\alpha) = (f')'(\alpha) = \cdots = (f')^{(m-2)}(\alpha) = 0$ and the induction hypothesis applied to $f'(X)$ shows that

α is a root of exact multiplicity $m-1$ of f'. So, $f'(X) = (X-\alpha)^{m-1}g(X)$, with $g(\alpha) \neq 0$. We know that α is a root of multiplicity q of f, so by (1), $f(\alpha) = f'(\alpha) = \cdots = f^{(q-1)}(\alpha) = 0$. If $q > m$, then $q-1 \geq m$, so $f^{(m)}(\alpha) = 0$, a contradiction. Thus, $q \leq m$. As

$$f'(X) = (X-\alpha)^{q-1}((X-\alpha)h'(X) + qh(X)),$$

we have

$$(X-\alpha)^{m-1}g(X) = (X-\alpha)^{q-1}((X-\alpha)h'(X) + qh(X)),$$

and since $q \leq m$, we get

$$(X-\alpha)^{m-q}g(X) = (X-\alpha)h'(X) + qh(X).$$

If we let $X = \alpha$, we have $qh(\alpha) \neq 0$, as $h(\alpha) \neq 0$ and $\text{char}(k) = 0$; but then, the left hand side must not be zero, and this implies $m = q$. □

Proposition 4.6 *Say $f \in k[X]$ (k = a field), then there is an extension K/k of finite degree and an element $\theta \in K$ so that $f(\theta) = 0$. If $\widetilde{k}$ is another field and $\mu\colon k \to \widetilde{k}$ is an isomorphism of fields, write $\widetilde{f} \in \widetilde{k}[X]$ for the image of f under μ (i.e., $\mu\Big(\sum g_j X_j\Big) = \sum \mu(g_j)X^j$), then f is irreducible over $k[X]$ iff $\widetilde{f}$ is irreducible over $\widetilde{k}[X]$. Let θ be a root of an irreducible polynomial, $f(X)$, in some extension K/k and let $\widetilde{\theta}$ be a root of $\widetilde{f}$ in some extension $\Omega/\widetilde{k}$. Then, there exists a unique extension of μ to a field isomorphism $k(\theta) \longrightarrow \widetilde{k}(\widetilde{\theta})$, so that $\mu(\theta) = \widetilde{\theta}$.*

Proof. Factor f into irreducible factors in $k[X]$, then a root of an irreducible factor is a root of f, so we may assume that f is irreducible. Now, the ideal $(f(X))$ is maximal in $k[X]$. Therefore, $K = k[X]/(f(X))$ *is* a field and $\overline{X}$ = the image of X in K is θ, a root, and $[K\colon k] = \deg(f) < \infty$.

Next, we have $\mu\colon k \to \widetilde{k}$ and $f \in k[X]$. Of course,

$$k[X] \cong k \otimes_{\mathbb{Z}} \mathbb{Z}[X] \cong \widetilde{k} \otimes_{\mathbb{Z}} \mathbb{Z}[X] \cong \widetilde{k}[X],$$

so f is irreducible iff $\widetilde{f}$ is irreducible. Now, $\theta \in K$ is a root of an irreducible polynomial, f, and $\widetilde{\theta} \in \Omega$ is a root of an irreducible polynomial $\widetilde{f}$. But, $k(\theta) \cong k[X]/(f(X)) \xrightarrow{\mu} \widetilde{k}[X]/(\widetilde{f}(X)) \cong \widetilde{k}(\widetilde{\theta})$. As θ generates $k(\theta)$ over k, the element $\mu(\theta)$ determines the extension of μ to $k(\theta)$. □

Proposition 4.7 *Say k is a field, $f \in k[X]$ and K/k is a field extension. Then, f possesses at most* $\deg(f)$ *roots in K counted with multiplicity and there exists an algebraic extension L/k (in fact, $[L\colon k] < \infty$) where f has exactly* $\deg(f)$ *roots counted with multiplicity.*

Proof. We use induction on $\deg(f)$. If $\alpha \in K/k$ is a root of f, then in $K[X]$, we have

$$f(X) = (X-\alpha)g(X), \quad \text{where } g(X) \in K[X]. \tag{$*$}$$

But, $\deg(g) = \deg(f) - 1$, so there exist at most $\deg(f) - 1$ roots of g in the field, K, containing k. If β is a root of f, either $\beta = \alpha$ or $g(\beta) = 0$ as K is a domain. Then, the first statement is proved. The last statement is again proved by induction. In the above, we can take $K = k(\alpha)$, of finite degree over k. Then, induction and $(*)$ imply our counting statement. □

Corollary 4.8 *(of the proof) The degree $[K\colon k]$ of a minimum field containing all* $\deg(f)$ *roots of f always satisfies $[K\colon k] \leq \deg(f)!$.*

Remarks:

(1) Proposition 4.7 is false if K is a ring but *not* a domain. For example, take

$$K = \underbrace{k \prod k \prod \cdots \prod k}_{n}.$$

Then, if $e_j = (0, \ldots, 0, 1, 0, \ldots, 0)$ with 1 in the j-th place, each e_j solves $X^2 = X$.

(2) Let $K = k[T]/(T^2)$. The elements $\alpha = \lambda\overline{T} \in K$ all satisfy $X^2 = 0$. If k is infinite, there are infinitely many solutions.

(3) Let $k = \mathbb{R}$ and $K = \mathbb{H}$ (the quaternions). We know that $\mathbb{H}$ is a division ring, i.e., every nonzero element has a multiplicative inverse. Consider the equation $X^2 + 1 = 0$. Then, every $\alpha = a\mathbf{i} + b\mathbf{j} + c\mathbf{k}$ with $a^2 + b^2 + c^2 = 1$ satisfies our equation!

(4) Given a field, k, there exists a field extension K/k having two properties:

(a) K/k is algebraic (but in general, $[K\colon k] = \infty$).

(b) For every $f \in K[X]$, there exists $\theta \in K$ so that $f(\theta) = 0$.

We'll prove these facts at the end of the Chapter in Section 4.11

Such a field, K, is called an *algebraic closure of* k and if only (2) holds, K is called *algebraically closed*. The field K is unique up to noncanonical isomorphism. The usual notation for an algebraic closure of k is $\overline{k}$.

4.3 Separable Extensions, Kähler Differentials, Mac Lane's Criterion

Definition 4.3 An algebraic element α over a field k (i.e., $\alpha \in K$ is algebraic over k for some field extension K/k) is *separable over* k iff α is a simple root of its minimal k-polynomial.[1] A polynomial, f, is *separable* iff all its irreducible factors are distinct and separable, and an irreducible polynomial is *separable* if it has one (hence all) separable roots. The field extension K/k is *separable* iff all $\alpha \in K$ are separable over k. We use the adjective *inseparable* to mean not separable.

Proposition 4.9 *Suppose α is inseparable over k. Then,* $\mathrm{char}(k) = p > 0$. *If f is the minimal polynomial for α, then there is some $n \geq 1$ and some irreducible polynomial $g(X) \in k[X]$ so that $f(X) = g(X^{p^n})$. If we choose n maximal then*

(1) $g(X)$ is a separable polynomial and

(2) α^{p^n} is separable over k. Any root β of f has the property that β^{p^n} is separable over k.

Proof. The element α is inseparable iff $f'(\alpha) = 0$ by the $n = 1$ case of the Jacobian criterion. Thus, f divides f', yet $\deg(f') < \deg(f)$. Therefore, $f' \equiv 0$. If $f(X) = \sum_{j=0}^{d} a_j X^j$, then $f'(X) = \sum_{j=0}^{d-1} j a_j X^{j-1}$ and it follows that $j a_j = 0$, for all j. If $\mathrm{char}(k) = 0$, then $a_j = 0$ for all $j \neq 0$ and $f \equiv 0$, as α is a root. Thus, we must have $\mathrm{char}(k) = p > 0$ and if p does not divide j, then $a_j = 0$. We deduce that

$$f(X) = \sum_{r=0}^{e} a_{pr} X^{pr} = h_1(X^p),$$

where $h_1(X) = \sum_{r=0}^{e} a_{pr} X^r$. Note that h_1 must be irreducible and repeat the above procedure if necessary. As $\deg(h_1) < \deg(f)$, this process must stop after finitely many steps. Thus, there is a maximum n with $f(X) = g(X^{p^n})$ and $g(X)$ is irreducible in $k[X]$. Were $g(X)$ inseparable, the first part of the argument would imply that $g(X) = h(X^p)$ and so, $f(X) = h(X^{p^{n+1}})$, contradicting the maximality of n. Therefore, $g(X)$ is separable. Yet, $g(\alpha^{p^n}) = f(\alpha) = 0$, so α^{p^n} is a root of an irreducible separable polynomial and (2) holds. Given β, we have β^{p^n} again a root of g. □

Definition 4.4 A field k of characteristic $p > 0$ is *perfect* iff $k = k^p$, i.e., for every $\lambda \in k$, the element λ has a p-th root in k.

Examples of Perfect and Imperfect Fields.

(1) $\mathbb{F}_p = \mathbb{Z}/p\mathbb{Z}$, where p is prime, is perfect.

(2) Any finite field is perfect.

(3) The field $k(T)$, where $\mathrm{char}(k) = p > 0$ is *always imperfect*.

Proposition 4.10 *If k is a field with characteristic* $\mathrm{char}(k) = p > 0$ *and if $c \notin k^p$ (with $c \in k$), then for every $n \geq 0$, the polynomial $f(X) = X^{p^n} - c$ is irreducible in $k[X]$. Conversely, if for some $n > 0$ the polynomial $X^{p^n} - c$ is irreducible, then $c \notin k^p$.*

Proof. Look at $f(X) = X^{p^n} - c$ and pick a field, K, with a root, $\alpha \in K$, of f. Then, $\alpha^{p^n} - c = 0$, so $f(X) = X^{p^n} - \alpha^{p^n} = (X - \alpha)^{p^n}$, since $\mathrm{char}(k) = p > 0$. Say $\varphi(X) \in k[X]$ is an irreducible factor of $f(X)$, then $\varphi(X) \mid f(X)$ in $k[X]$, and similarly in $K[X]$. By unique factorization in $K[X]$, we have $\varphi(X) = (X - \alpha)^r$, for some $r > 0$, where $\alpha^{p^n} - c = 0$ and $\alpha \in K$.

[1] Recall that the *minimal k-polynomial of α* is the monic polynomial of minimal degree generating the principal ideal consisting of the polynomials in $k[X]$ that vanish on α.

Claim: $X^{p^n} - c$ is a power of $\varphi(X)$.

If not, there is some irreducible polynomial, $\psi(X)$, relatively prime to $\varphi(X)$ and $\psi(X) \mid X^{p^n} - c$ in $k[X]$ (DX). Then, there exist $s(X), t(X)$ with $s(X)\varphi(X) + t(X)\psi(X) = 1$ in $k[X]$. However, $\psi(X)$ divides $X^{p^n} - c$, so $\psi(\alpha) = 0$. If we let $X = \alpha$, we get $1 = s(\alpha)\varphi(\alpha) + t(\alpha)\psi(\alpha) = 0$, a contradiction.

Therefore, $\varphi(X)^l = X^{p^n} - c$. It follows that $rl = p^n$, so $r = p^a$ and $l = p^b$ with $a + b = n$. Then,

$$\varphi(X) = (X - \alpha)^r = (X - \alpha)^{p^a} = X^{p^a} - \alpha^{p^a},$$

which implies $\alpha^{p^a} \in k$. But then, $c = (\alpha^{p^a})^{p^b} \in k^{p^b}$, a contradiction if $b \geq 1$. Thus, $b = 0$ and consequently, $a = n$ and $f(X) = \varphi(X)$ is irreducible.

Conversely, if for some $n > 0$ the polynomial $X^{p^n} - c$ is irreducible and if $c \in k^p$, then $c = b^p$, for some $b \in k$. It follows that

$$X^{p^n} - c = X^{p^n} - b^p = (X^{p^{n-1}} - b)^p$$

contradicting the irreducibility of $X^{p^n} - c$. □

Definition 4.5 An element $\alpha \in K/k$ is *purely inseparable over* k ($\mathrm{char}(k) = p > 0$) iff there is some $n \geq 0$ so that $\alpha^{p^n} \in k$. Equivalently, α is purely inseparable over k iff the minimal k-polynomial for α is of the form $X^{p^n} - c$, for some $c \in k$.

Remark: We have $\alpha \in k$ iff α is separable *and* purely inseparable over k.

Proposition 4.11 *If k is a field, then k is perfect iff every algebraic extension of k is separable.*

Proof. ($\Rightarrow$). Say k is perfect and pick $\alpha \in K/k$, with α algebraic. We know that α has a minimal k-polynomial $f(X)$ and that $f(X) = g(X^{p^n})$, for some irreducible polynomial, $g(X)$, and some $n \geq 0$. We have $g(X) = \sum_{j=0}^{N} b_j X^j$, so $f(X) = \sum_{j=0}^{N} b_j (X^{p^n})^j$. As k is perfect, $k = k^p = k^{p^2} = \cdots = k^{p^n}$. So, $b_j = c_j^{p^n}$, for some $c_j \in k$ and we have

$$f(X) = \sum_{j=0}^{N} c_j^{p^n} (X^{p^n})^j = \Big(\sum_{j=0}^{N} c_j X^j\Big)^{p^n}.$$

This contradicts the irreducibility of $f(X)$ unless $n = 0$, and we know that $\alpha^{p^0} = \alpha$ is separable over k.

($\Leftarrow$). In this case, all algebraic extensions of k are separable and say k is not perfect. Then, there is some $c \in k$, with $c \notin k^p$. Hence, by Proposition 4.10, the polynomial $X^p - c$ is irreducible over k. Let $K = k(\alpha)$ where α is some root of $X^p - c$. Then, $\alpha^p = c \in k$ and it follows that α is purely inseparable over k. But, α is separable over k, a contradiction, as $\alpha \notin k$. □

Corollary 4.12 *For a field, k, the following are equivalent:*

(1) k is imperfect.

(2) k possesses nontrivial inseparable extensions.

(3) k possesses purely inseparable extensions.

Say K/k is algebraic and inseparable. It can happen that there does not exist $\alpha \in K$ with α purely inseparable over k.

To go further, we need derivations and Kähler differentials. Consider the situation where A, B are commutative rings and B is an A-algebra. On B-modules, we have an endofunctor:

$$M \rightsquigarrow \mathrm{Der}_A(B, M).$$

Is the above functor representable? This means, does there exist a B-module, $\Omega_{B/A}$, and an element, $d \in \mathrm{Der}_A(B, \Omega_{B/A})$, so that functorially in M

$$\theta_M \colon \mathrm{Hom}_B(\Omega_{B/A}, M) \xrightarrow{\sim} \mathrm{Der}_A(B, M)?$$

(Note: For every $\varphi \in \mathrm{Hom}_B(\Omega_{B/A}, M)$, we have $\theta_M(\varphi) = \varphi \circ d$, see below).

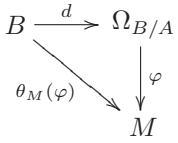

Theorem 4.13 *The functor $M \rightsquigarrow \mathrm{Der}_A(B, M)$ is representable by a pair $(\Omega_{B/A}, d)$, as above.*

Proof. Consider $B \otimes_A B$ and the algebra map $B \otimes_A B \xrightarrow{\mu} B$, where μ is multiplication, i.e., $\mu(b \otimes \widetilde{b}) = b\widetilde{b}$. Let $I = \mathrm{Ker}\ \mu$ and write $I/I^2 = \Omega_{B/A}$. We let B act on $B \otimes_A B$ *via* the left action $b(\xi \otimes \eta) = b\xi \otimes \eta$. Then, $\Omega_{B/A}$ is a B-module. Given $b \in B$, set

$$db = d(b) = (1 \otimes b - b \otimes 1) \bmod I^2.$$

Now, for $b, \widetilde{b} \in B$, we have

$$(1 \otimes b - b \otimes 1)(1 \otimes \widetilde{b} - \widetilde{b} \otimes 1) \in I^2,$$

and we get

$$1 \otimes b\widetilde{b} + b\widetilde{b} \otimes 1 - (b \otimes \widetilde{b} + \widetilde{b} \otimes b) \in I^2.$$

So, modulo I^2, the above is zero and

$$1 \otimes b\widetilde{b} - \widetilde{b} \otimes b = b \otimes \widetilde{b} - b\widetilde{b} \otimes 1 \quad \text{in } \Omega_{B/A}.$$

Obviously, d is additive and zero on A, so we only need to check the Leibnitz rule. We have

$$\begin{aligned} bd(\widetilde{b}) &= b(1 \otimes \widetilde{b} - \widetilde{b} \otimes 1) \quad \bmod I^2 \\ &= b \otimes \widetilde{b} - b\widetilde{b} \otimes 1 \quad \text{in } \Omega_{B/A} \\ &= 1 \otimes b\widetilde{b} - \widetilde{b} \otimes b \quad \text{in } \Omega_{B/A} \\ &= 1 \otimes b\widetilde{b} - b\widetilde{b} \otimes 1 + b\widetilde{b} \otimes 1 - \widetilde{b} \otimes b \quad \text{in } \Omega_{B/A} \\ &= d(b\widetilde{b}) - \widetilde{b}(1 \otimes b - b \otimes 1) \quad \text{in } \Omega_{B/A}. \end{aligned}$$

So, $bd(\widetilde{b}) = d(b\widetilde{b}) - \widetilde{b}db$ in $\Omega_{B/A}$, namely $d(b\widetilde{b}) = \widetilde{b}db + bd(\widetilde{b})$. The rest of the proof is routine. □

Definition 4.6 The B-module $\Omega_{B/A}$ (together with the derivation d) is called the *module of relative Kähler differentials of B over A*.

Examples of Relative Kähler Differentials.

(1) Let $B = A[T_1, \ldots, T_n]$. Say D is a derivation of $B \longrightarrow M$ trivial on A. So, we know $D(T_1), \ldots, D(T_n)$; these are some elements in M. Say we are given $T_l^r \in B$. Then,

$$D(\underbrace{T_l \cdots T_l}_{r}) = rT_l^{r-1}D(T_l) = \frac{\partial}{\partial T_l}(T_l^r)D(T_l).$$

Now,

$$D(T_k^r T_l^s) = T_k^r D(T_l^s) + T_l^s D(T_k^r) = T_k^r \frac{\partial}{\partial T_l}(T_l^s) D(T_l) + T_l^s \frac{\partial}{\partial T_k}(T_k^r) D(T_k).$$

In general

$$D(T_1^{a_1} \cdots T_n^{a_n}) = \sum_{j=1}^{n} T_1^{a_1} \cdots \widehat{T_l^{a_l}} \cdots T_n^{a_n} \frac{\partial}{\partial T_l}(T_l^{a_l}) D(T_l), \tag{$\dagger$}$$

and

$$D\Big(\sum \alpha_{(a)} T^{(a)}\Big) = \sum_{(a)} \alpha_{(a)} D(T^{(a)}) D(T_l), \tag{$\dagger\dagger$}$$

as $D \restriction A \equiv 0$. Conversely, ($\dagger$) on the linear base of monomials in the $T_1, \ldots, T_n$ of B gives a derivation. Therefore,

$$\mathrm{Der}_A(B, M) \longrightarrow \coprod_{i=1}^{n} M,$$

via $D \mapsto (D(T_1), \ldots, D(T_n))$ is a functorial isomorphism. Consequently,

$$\Omega_{B/A} \cong \coprod_{j=1}^{n} B dT_j,$$

where the dT_j are A-linearly independent elements of $\Omega_{B/A}$ (case $M = \Omega_{B/A}$).

(2) Let B be a f.g. algebra over A, i.e., $B = A[T_1, \ldots, T_n]/(f_1, \ldots, f_p)$. We have

$$\mathrm{Der}_A(B, M) = \{\varphi \in \mathrm{Der}_A(A[T_1, \ldots, T_n], M) \mid \varphi(f_i) = 0,\ i = 1, \ldots, p\}.$$

But,

$$\varphi(f_i) = \sum_{j=1}^{n} \frac{\partial f_i}{\partial T_j} \varphi(T_j) = \sum_{j=1}^{n} \frac{\partial f_i}{\partial T_j} \overline{\varphi}(dT_j),$$

where $\overline{\varphi}\colon \Omega_{B/A} \longrightarrow M$ (and $\varphi = \overline{\varphi} \circ d$). We let $M = \Omega_{B/A}$ to determine it, and we see that

$$\overline{\varphi} \text{ must kill } df_i.$$

It follows that

$$\Omega_{B/A} = \left(\coprod_{j=1}^{n} B dT_j\right) \Big/ (\text{submodule } df_1 = \cdots = df_n = 0).$$

(3) Let $B = \mathbb{C}[X, Y]/(Y^2 - X^3)$ and $A = \mathbb{C}$. From (2) we get

$$\Omega_{B/A} = (B dX \amalg B dY)/(2Y dY - 3X^2 dX).$$

The module $\Omega_{B/A}$ is *not* a free B-module (due to the singularity at the origin of the curve $Y^2 = X^3$).

(4) Let $A = \mathbb{R}$ or $\mathbb{C}$ and $B =$ the ring of functions on a small neighborhood of a smooth r-dimensional manifold (over A). Derivations on B over A have values in B. Let $\xi_1, \ldots, \xi_r$ be coordinates on this neighborhood. Then, $\partial/\partial\xi_j$ is a derivation defined so that

$$\frac{\partial f}{\partial \xi_j} = \lim_{h \longrightarrow 0} \frac{f(\ldots, \xi_j + h, \ldots) - f(\ldots, \xi_j, \ldots)}{h}.$$

Look near a point, we may assume $\xi_1 = \cdots = \xi_r = 0$, there. By Taylor,

$$f(\xi_1 + h_1, \ldots, \xi_r + h_r) = f(\xi_1, \ldots, \xi_r) + \sum_{j=1}^{r} \frac{\partial f}{\partial \xi_j} h_j + O(\|h\|^2).$$

Hence, $\Omega_{B/A}$ is generated by $d\xi_1, \ldots, d\xi_r$ and they are linearly independent over B because the implicit function theorem would otherwise imply that some ξ_j is a function of the other ξ_i's near our point, a contradiction.

Definition 4.7 Given an A-algebra, B, the algebra B is *étale over* A iff

(1) The algebra B is flat over A.

(2) The algebra B is f.p. as an A-algebra.

(3) $\Omega_{B/A} = (0)$.

The algebra B is *smooth over* A iff (1), (2) and (3′): $\Omega_{B/A}$ is a locally-free B-module, hold.

Remark: Putting aside (2), we see that checking that an algebra is *étale* or smooth is local on A, i.e., it is enough to check it for $B_{\mathfrak{p}}$ over $A_{\mathfrak{p}}$ for every $\mathfrak{p} \in \operatorname{Spec} A$. This is because (DX)

$$(\Omega_{B/A})_{\mathfrak{p}} = \Omega_{B_{\mathfrak{p}}/A_{\mathfrak{p}}}.$$

It turns out that smooth means: Locally on A, the algebra B looks like

$$A \hookrightarrow A[T_1, \ldots, T_r] \longrightarrow B$$

where $B/A[T_1, \ldots, T_r]$ is *étale*.

We can apply the concepts of relative Kähler differentials and étale homomorphisms to field theory. For this, given a field, write $p = \operatorname{char}(k)$ and if $p > 0$, let $k^{1/p}$ be the field

$$k^{1/p} = \{x \in \overline{k} \mid x^p \in k\}.$$

Theorem 4.14 *(Main theorem on separability (alg. case).) Let K/k be an algebraic extension. Then, in the following statements: (1) implies any of the others; (2), (2a) and (3) are equivalent; (1) and (4) are equivalent; all are equivalent if K/k is finite.*

(1) The extension K/k is separable.

(2) For all K-modules, M, we have $\operatorname{Der}_k(K, M) = (0)$.

(2a) $\Omega_{K/k} = (0)$*, i.e., when K/k is finite, it is étale.*

(3) Every derivation of k to M (where M is a K-vector space) which admits an extension to K (i.e., becomes a derivation $K \longrightarrow M$) admits a unique extension.

When $\operatorname{char}(k) = p > 0$*,*

(4) Mac Lane I*: The natural map $k^{1/p} \otimes_k K \longrightarrow K^{1/p}$ is injective.*

(5) Mac Lane II*: $kK^p = K$.*

In order to prove Theorem 4.14, we first need the following subsidiary statement:

Proposition 4.15 *If K/k is separable and if M is a K-vector space, then every derivation $D\colon k \to M$ admits an extension to a derivation of K with values in M.*

Proof. We use Zorn's lemma. Let $\mathcal{S}$ be the set of all pairs, (L, D_L), where

(1) L is a subextension of K/k (i.e., $k \subseteq L \subseteq K$).

(2) D_L is an extension of D to L with values in M.

As $(k, D) \in \mathcal{S}$, the set $\mathcal{S}$ is nonempty. Define a partial order on $\mathcal{S}$ by: $(L, D_L) \leq (L', D_{L'})$ iff $L \subseteq L'$ and $D_{L'} \restriction L = D_L$. The set $\mathcal{S}$ is inductive. (If $\{L_\alpha\}_\alpha$ is a chain, then $L = \bigcup_\alpha L_\alpha$ *is* a field, and define $D_L(\xi) = D_{L_\alpha}(\xi)$, where $\xi \in L_\alpha$; this is well-defined (DX).) By Zorn's lemma, there exists a maximal extension, say (L, D_L).

If $K \neq L$, then there is some $\beta \in K$ with $\beta \notin L$. Let $g(X) \in L[X]$ be the minimum L-polynomial for β. We try to extend D_L to $L(\beta)$. For this, we must define $D_{L(\beta)}(\beta)$ and the only requirement it needs to satisfy is

$$0 = D_{L(\beta)}(g(\beta)) = g'(\beta)D_{L(\beta)}(\beta) + D_L(g)(\beta).$$

Here, if $g(X) = \sum_{j=1}^r a_j X^j$, then $D_L(g)(\alpha)$ is $\sum_{j=1}^r \alpha^j D_L(a_j) \in M$. Since β is separable, $g'(\beta) \neq 0$, so we can find the value of $D_{L(\beta)}(\beta)$, contradicting the maximality of our extension. Therefore, $L = K$. □

Proof of Theorem 4.14. (1) ⇒ (2). Pick $D \in \mathrm{Der}_k(K, M)$ and $\alpha \in K$; by (1), the element α is separable over k, i.e., α has a minimal k-polynomial, $g(X)$, so that $g(\alpha) = 0$ and $g'(\alpha) \neq 0$. As D is a derivation, the argument of Proposition 4.15 implies that

$$0 = D(g(\alpha)) = g'(\alpha)D(\alpha) + D(g)(\alpha).$$

But, $D(g) = 0$, because the coefficients of g are in k and $D \restriction k \equiv 0$. Since $g'(\alpha) \neq 0$, we get $D(\alpha) = 0$, i.e., (2) holds.

(2) ⇒ (2a). We have the functor $M \rightsquigarrow \mathrm{Der}_k(K, M)$ and $\mathrm{Der}_k(K, M) = (0)$. By Yoneda's lemma, the representing object, $\Omega_{K/k}$, must vanish.

(2a) ⇒ (2). We have $\mathrm{Hom}_k(\Omega_{K/k}, M) \cong \mathrm{Der}_k(K, M)$ and $\Omega_{K/k} = (0)$, so (2) holds.

(2) ⇒ (3). Say D and $\widetilde{D}$ are two extensions of the same derivation on k. Then, $D - \widetilde{D}$ is a derivation and $(D - \widetilde{D}) \restriction k \equiv 0$. By (2), $(D - \widetilde{D}) \in \mathrm{Der}_k(K, M) = (0)$, so $D - \widetilde{D} = 0$.

(3) ⇒ (2). Choose $D \in \mathrm{Der}_k(K, M)$, so $D \restriction k \equiv 0$. But then, D extends 0 and 0 extends 0; by (3), $D \equiv 0$.

(1) ⇒ (5). If $\alpha \in K$, then α is separable over k, so α is separable over kK^p (as $kK^p \supseteq k$). Yet, $\alpha^p \in K^p$, so $\alpha^p \in kK^p$; thus, α is purely inseparable over kK^p. As α is both separable and purely inseparable over kK^p, by a previous remark, $\alpha \in kK^p$. This shows $K \subseteq kK^p$. On the other hand, $kK^p \subseteq K$, always. Therefore, $K = kK^p$, i.e., (5) holds.

Before discussing the equivalence of (4) with (1), we need to elucidate the meaning of the Mac Lane conditions.

For (5), say $\{\xi_\lambda\}_\lambda$ spans K as a k-vector space. Then, $\{\xi_\lambda^p\}_\lambda$ spans K^p as a k^p-space. As $k^p \subseteq k$, $\{\xi_\lambda^p\}_\lambda$ spans kK^p as a *k-space*. Hence, *Mac Lane II means: If $\{\xi_\lambda\}_\lambda$ spans K as a k-space, so does $\{\xi_\lambda^p\}_\lambda$.*

For (4), say $\{\xi_\lambda\}_\lambda$ is a linearly independent family (for short, an *l.i.* family) over k in K. Then, we know that the elements $1 \otimes \xi_\lambda$ are linearly independent in $k^{1/p} \otimes_k K$ as $k^{1/p}$-vectors ($k^{1/p}$ acts on the left on $k^{1/p} \otimes_k K$). The map $k^{1/p} \otimes_k K \longrightarrow K^{1/p}$ is just

$$\sum_\lambda a_\lambda \otimes \xi_\lambda \mapsto \sum_\lambda a_\lambda \xi_\lambda.$$

If the map is injective and if there is a linear dependence of the ξ_λ (in $K^{1/p}$) *over* $k^{1/p}$, we get $\sum_\lambda a_\lambda \xi_\lambda = 0$, for some $a_\lambda \in k^{1/p}$. But then, $\sum_\lambda a_\lambda \otimes \xi_\lambda$ would go to zero and by injectivity

$$\sum_\lambda a_\lambda \otimes \xi_\lambda = \sum_\lambda (a_\lambda \otimes 1)(1 \otimes \xi_\lambda) = 0$$

in $k^{1/p} \otimes_k K$. But, $\{1 \otimes \xi_\lambda\}_\lambda$ is linearly independent in $k^{1/p} \otimes_k K$, so $a_\lambda = 0$, for all λ. Consequently, the family $\{\xi_\lambda\}_\lambda$ is still linearly independent over $k^{1/p}$. Conversely (DX), if any l.i. family $\{\xi_\lambda\}_\lambda$ (with $\xi_\lambda \in K$) over k remains l.i. over $k^{1/p}$, then our map $k^{1/p} \otimes_k K \longrightarrow K^{1/p}$ is injective. By using the isomorphism $x \mapsto x^p$, we get: *Mac Lane I says that any l.i. family* $\{\xi_\lambda\}_\lambda$ *over* k*, has the property that* $\{\xi_\lambda^p\}_\lambda$ *is still l.i. over* k*.*

Now, say K/k is finite, with $[K\colon k] = n$. Then, $\xi_1, \ldots, \xi_n$ is l.i. over k iff $\xi_1, \ldots, \xi_n$ span K. Condition (4) implies $\xi_1^p, \ldots, \xi_n^p$ are l.i. and since there are n of them, they span K, i.e. (5) holds. Conversely, if (5) holds then $\xi_1^p, \ldots, \xi_n^p$ span K and there are n of them, so they are l.i., i.e., (4) holds. *Therefore, (4) and (5) are equivalent if* K/k *is finite*. We can show that (1) and (4) are equivalent (when $\mathrm{char}(k) = p > 0$).

(4) $\Rightarrow$ (1). Pick $\alpha \in K$. We know that α^{p^n} is separable over k for some $n \geq 0$. Further, the minimal polynomial for $\beta = \alpha^{p^n}$ is $h(X)$, where $f(X) = h(X^{p^n})$ and f is the minimal k-polynomial for α. Say, $\deg(f) = d$. So, $d = p^n d_0$, with $d_0 = \deg(h)$. Now, $1, \alpha, \ldots, \alpha^{d-1}$ are l.i. over k. By (4), repeatedly, $1, \alpha^{p^n}, (\alpha^2)^{p^n}, \ldots, (\alpha^{d-1})^{p^n}$ are l.i., i.e., $1, \beta, \ldots, \beta^{d-1}$ are l.i. Yet, $1, \beta, \ldots, \beta^{d_0}$ is the maximum l.i. family for the powers of β, so $d \leq d_0$. This can only happen if $n = 0$ and α is separable over k. □

(1) $\Rightarrow$ (4). Say $\{\xi_\lambda\}_\lambda$ is l.i. in K/k. As linear independence is checked by examining finite subfamilies, we may assume that our family is $\xi_1, \ldots, \xi_t$. We must prove, $\xi_1^p, \ldots, \xi_t^p$ are still l.i. over k. Let $L = k(\xi_1, \ldots, \xi_t)$, then L/k is a *finite* extension. For such an extension, (4) and (5) are equivalent. But, we just proved that (1) implies (5), so (1) implies (4).

Finally, in the case K/k is finite there remains the proof of (2) $\Rightarrow$ (1). For this, it is simplest to prove a statement we'll record as Corollary 4.16 below. This is:

Corollary 4.16 *If* $\alpha_1, \ldots, \alpha_t$ *are each separable over* k*, then the field* $k(\alpha_1, \ldots, \alpha_t)$ *is separable over* k*. In particular, if* K/k *is algebraic and* K_{sep} *denotes the set of all elements of* K *that are separable over* k*, then* K_{sep} *is a field.*

To prove these statements, we will apply Mac Lane II; this will suffice as $L = k(\alpha_1, \ldots, \alpha_t)$ is finite over k. Now $kL^p = k(\alpha_1^p, \ldots, \alpha_t^p)$ and each α_j is therefore purely inseparable over kL^p. However, each α_j is separable over k and therefore over kL^p. It follows that each $\alpha_j \in kL^p$ so that $L = kL^p$ and Mac Lane II applies. For the proof, proper, that (2) $\Rightarrow$ (1), assume (2) and that (1) is false. Then $K_{\mathrm{sep}} \neq K$, so we can find $\alpha_1, \ldots, \alpha_s \in K$, each purely inseparable over K_{sep}, and so that

$$K = K_{\mathrm{sep}}(\alpha_1, \ldots, \alpha_s) > K_{\mathrm{sep}}(\alpha_1, \ldots, \alpha_{s-1}) > \cdots > K_{\mathrm{sep}}(\alpha_1) > K_{\mathrm{sep}}.$$

Consider the zero derivation on $K_{\mathrm{sep}}(\alpha_1, \ldots, \alpha_{s-1})$. Now, $\beta = \alpha_s^{p^r} \in K_{\mathrm{sep}}(\alpha_1, \ldots, \alpha_{s-1})$ for some minimal $r > 0$, thus to extend the zero derivation to K we need only assign a value to $D(\alpha_s)$ so that $D(\alpha_s^{p^r}) = p^r \alpha_s^{p^{r-1}} D(\alpha_s) = 0$. Any nonzero element of M will do, contradicting (2). □

Corollary 4.17 *Every algebraic extension of a perfect field is perfect. In particular, every finite field is perfect and every absolutely algebraic field (i.e., algebraic over a prime field) is perfect.*

Proof. If K/k is algebraic and k is perfect, then K/k is separable. By Mac Lane II, we have $K = kK^p$. But, $k = k^p$ (k perfect), so $K = k^p K^p = (kK)^p = K^p$. A finite field is algebraic over $\mathbb{F}_p$ and by little Fermat, $\mathbb{F}_p^p = \mathbb{F}_p$, i.e., perfect. (Second proof by counting: The map $\xi \mapsto \xi^p$ is injective, taking $\mathbb{F}_q$ to itself. But, the image has cardinality q; by finiteness, the image is all of $\mathbb{F}_q$.) By the first part of the proof, an absolutely algebraic field is perfect.

Corollary 4.18 *Say $\alpha_1, \dots, \alpha_t$ are each separable over k. Then, the field $k(\alpha_1, \dots, \alpha_t)$ is a separable extension of k. In particular, if K/k is algebraic and we set*

$$K_{\text{sep}} = \{\alpha \in K \mid \alpha \text{ is separable over } k\}$$

then K_{sep} is a subfield of K/k called the separable closure of k in K.

Corollary 4.19 *Say K/k is an algebraic extension and $\alpha_1, \dots, \alpha_t \in K$. If each α_j is separable over $k(\alpha_1, \dots, \alpha_{j-1})$, then $k(\alpha_1, \dots, \alpha_t)$ is separable over k. In particular, separability is transitive.*

Proof. We use induction on t. When $t = 1$, this is Corollary 4.18. Assume that the induction hypothesis holds for $t - 1$. So, $L = k(\alpha_1, \dots, \alpha_{t-1})$ is separable over k and it is a finite extension, therefore Mac Lane II yields $kL^p = L$. Let $M = k(\alpha_1, \dots, \alpha_t)$, then $M = L(\alpha_t)$. So, M is separable over L, by the case $t = 1$. Therefore, $M = LM^p$, by Mac Lane II. Now,

$$M = LM^p = kL^pM^p = k(LM)^p = kM^p.$$

By Mac Lane II, again, M is separable over k. □

Corollary 4.20 *If K/k is an algebraic extension, then K is purely inseparable over K_{sep}.*

Corollary 4.21 *Pure inseparability is transitive.*

The implication (2) ⇒ (1) does not hold if K/k is not finite. Here is an example: Set $k = \mathbb{F}_p(T)$, where T is an indeterminate. Define, inductively, the chain of fields

$$k = k_0 < k_1 < \cdots < k_n < \cdots$$

via the rule

$$\alpha_0 = T; \quad \alpha_j = \alpha_{j-1}^{1/p}; \quad k_j = k_{j-1}(\alpha_j).$$

Let $K = k_\infty = \bigcup_{j=0}^{\infty} k_j$. Then a derivation on K, trivial on k is determined by its values on the α_j. Yet, we have $\alpha_{j+1}^p = \alpha_j$, therefore $D(\alpha_j) = 0$ for every j; hence, $\text{Der}_k(K, -) = 0$. But, K/k is not separable; indeed it is purely inseparable.

Notation: For a field, k, of characteristic $p > 0$, set $[K\colon k]_s \stackrel{\text{def}}{=} [K_{\text{sep}}\colon k]$, the *separable degree of K/k* and $[K\colon k]_i \stackrel{\text{def}}{=} [K\colon K_{\text{sep}}]$, the *purely inseparable degree of K/k* (if K/k is finite, $[K\colon k]_i$ is a power of p). Clearly,

$$[K\colon k] = [K\colon k]_i[K\colon k]_s.$$

4.4 The Extension Lemma and Splitting Fields

We begin with a seemingly "funny" notion: Two fields K, L are *related*, denoted $K \underset{\text{rel}}{\sim} L$, iff there is some larger field, W, so that $K \subseteq W$ and $L \subseteq W$ (as sets, *not* isomorphic copies). This notion is reflexive and symmetric, but *not* transitive.

Theorem 4.22 *(Extension Lemma) Let K/k be a finite extension and say $\widetilde{k}$ is another field isomorphic to k via $\theta\colon k \to \widetilde{k}$. Suppose Γ is another field related to $\widetilde{k}$, but otherwise arbitrary. Then, there exists a finite extension, $\widetilde{K}/\widetilde{k}$, with $\widetilde{K} \underset{\text{rel}}{\sim} \Gamma$ and an extension of θ to an isomorphism $\widetilde{\theta}\colon K \to \widetilde{K}$.*

$$\begin{array}{ccccc} K & \xrightarrow{\widetilde{\theta}} & \widetilde{K} & \underset{\text{rel}}{\sim} & \Gamma \\ {\scriptstyle\text{finite}}\uparrow & & \uparrow{\scriptstyle\text{finite}} & & \| \\ k & \xrightarrow[\theta]{} & \widetilde{k} & \underset{\text{rel}}{\sim} & \Gamma \end{array}$$

Proof. We proceed by induction on the number, n, of adjunctions needed to obtain K from k.

Case $n = 1$: $K = k(\alpha)$. Let $g(X) \in k[X]$ be the minimum k-polynomial for α. Write $\widetilde{g}(X) \in \widetilde{k}[X]$ for the image, $\theta(g)(X)$, of $g(X)$. Of course, $\widetilde{g}(X)$ is $\widetilde{k}$-irreducible. Now, there exists a field, W, with $W \supseteq \widetilde{k}$ and $W \supseteq \Gamma$. Thus, $\widetilde{g}(X) \in W[X]$; moreover, there exists an extension W'/W of W and some $\widetilde{\alpha} \in W'$, so that $\widetilde{g}(\widetilde{\alpha}) = 0$. It follows that $\widetilde{k}(\widetilde{\alpha}) \subseteq W'$ and $\Gamma \subseteq W \subseteq W'$, so $\widetilde{k}(\widetilde{\alpha}) \underset{\text{rel}}{\sim} \Gamma$. But we know by Proposition 4.6 that θ extends to an isomorphism $\widetilde{\theta}\colon k(\alpha) \to \widetilde{k}(\widetilde{\alpha})$. This proves case 1.

Induction step. Assume that the induction hypothesis holds for all $t \leq n-1$. We have $K = k(\alpha_1, \ldots, \alpha_n)$ and let $L = k(\alpha_1, \ldots, \alpha_{n-1})$. By the induction hypothesis, there is a finite extension, $\widetilde{L}$, and an isomorphism, $\theta'\colon L \to \widetilde{L}$, extending θ;

$$\begin{array}{ccccc} L(\alpha_n) = K & \xrightarrow{\widetilde{\theta}} & \widetilde{K} & \underset{\text{rel}}{\sim} & \Gamma \\ \uparrow & & \uparrow & & \| \\ L & \xrightarrow{\theta'} & \widetilde{L} & \underset{\text{rel}}{\sim} & \Gamma \\ \uparrow & & \uparrow & & \| \\ k & \xrightarrow{\theta} & \widetilde{k} & \underset{\text{rel}}{\sim} & \Gamma \end{array}$$

We complete the proof using the argument in case 1 (a single generator), as illustrated in the above diagram. □

Corollary 4.23 *If K/k is a finite extension and $k \underset{\text{rel}}{\sim} \Gamma$, then there is a k-isomorphism $K/k \longrightarrow \widetilde{K}/\widetilde{k}$ and $\widetilde{K} \underset{\text{rel}}{\sim} \Gamma$.*

Proof. This is the case $k = \widetilde{k}$; $\theta = \text{id}$. □

Definition 4.8 A field extension L/k is a *splitting field for the polynomial* $f(X) \in k[X]$ iff $L = k(\alpha_1, \ldots, \alpha_n)$ and $\alpha_1, \ldots, \alpha_n$ are all the roots of $f(X)$ in some larger field ($n = \deg(f)$).

Remarks:

(1) When we view $f(X) \in L[X]$, then $f(X)$ splits into linear factors

$$f(X) = c(X - \alpha_1) \cdots (X - \alpha_n)$$

in $L[X]$, hence the name. Conversely, if M/k is a field extension and in $M[X]$, the polynomial $f(X)$ splits into linear factors, then M contains some splitting field for f. (Here, $f(X) \in k[X]$.)

(2) Suppose L/k and L'/k are two splitting fields for the same polynomial $f(X) \in k[X]$. Then $L = L'$ iff $L \underset{\text{rel}}{\sim} L'$ (L and L' are *identical*, not just isomorphic).

Proof. The implication ($\Rightarrow$) is obvious. Conversely, assume $L \underset{\text{rel}}{\sim} L'$. Say Ω is a common extension of L and L' in which $f(X)$ splits. In Ω, the polynomial f has just n roots, say $\beta_1, \ldots, \beta_n$. Yet, $L = k(\beta_1, \ldots, \beta_n)$ and $L' = k(\beta_1, \ldots, \beta_n)$, too. Therefore, $L = L'$.

(3) Suppose L/k is a splitting field for $f(X) \in k[X]$ and $k \cong \widetilde{k}$ *via* some isomorphism, θ. If $\widetilde{f}(X)$ is the image of $f(X)$ by θ, and if θ extends to an isomorphism $L \cong \widetilde{L}$ for some extension $\widetilde{L}/\widetilde{k}$, then $\widetilde{L}$ is a splitting field for $\widetilde{f}(X)$.

Proposition 4.24 *Say $f(X) \in k[X]$ and $\theta\colon k \to \widetilde{k}$ is an isomorphism. Write $\widetilde{f}(X)$ for the image of $f(X)$ by θ. Then, θ extends to an isomorphism from any splitting field of f to any splitting field of $\widetilde{f}$. In particular, any two splitting fields of $f(X)$ are k-isomorphic (case $k = \widetilde{k}$; $f = \widetilde{f}$).*

Proof. Apply the extension lemma to the case where K is any chosen splitting field for f and Γ is any chosen splitting field for $\widetilde{f}$. The extension lemma yields an extension $\widetilde{K}/\widetilde{k}$ and an extension $\widetilde{\theta}\colon K \to \widetilde{K}$ with $\widetilde{K} \underset{\text{rel}}{\sim} \Gamma$. By Remark (3), the field $\widetilde{K}$ is a splitting field for $\widetilde{f}$. By Remark (2), as $\widetilde{K}$ and Γ are both splitting fields and $\widetilde{K} \underset{\text{rel}}{\sim} \Gamma$, they are equal. □

Definition 4.9 An algebraic field extension, M/k, is *normal* iff for all irreducible k-polynomials, $g(X)$, whenever some root of g is in M, all the roots of g are in M.

Proposition 4.25 *Say M/k is a finite extension and write $M = k(\beta_1, \ldots, \beta_t)$. Then, the following are equivalent:*

(1) M/k is normal.

(2) M is the splitting field of a family, $\{g_\alpha\}_\alpha$, of k-polynomials (the family might be infinite).

(3) M is the splitting field of a single k-polynomial (not necessarily irreducible).

(4) M is identical to all its k-conjugates; here two fields are k-conjugate iff they are both related and k-isomorphic.

Proof. (1) $\Rightarrow$ (2). For each β_i, there is an irreducible k-polynomial, say g_i with $g_i(\beta_i) = 0$. By (1), all the other roots of g_i are in M. Therefore, M contains the splitting fields of each g_i. But, clearly, M is contained in the field generated by all these splitting fields. It follows that M is equal to the splitting field of the (finite) family of k-polynomials $g_1, \ldots, g_t$.

(2) $\Rightarrow$ (3). Say $\{g_\alpha\}$ is the family of k-polynomials for which M is the splitting field. (Note that we may assume that $\deg(g_\alpha) > 1$ for all g_α.) Pick a countable (at most) subset $\{g_1, g_2, \ldots, \}$ of our family. Then, M contains the splitting field of g_1, call it M_1. We have $M \supseteq M_1 \supseteq k$ and $[M\colon M_1] < [M\colon k]$. If $M \neq M_1$, then M contains the splitting field, M_2, of g_1 and g_2, where we may assume that the splitting field of g_2 is distinct from M_1. Thus, we have $M \supseteq M_2 \supseteq M_1 \supseteq k$. Since M is finite over k, the above process stops and we deduce that M is the splitting field of a finite subfamily $\{g_1, \ldots, g_t\}$. Then, take $g = \prod_{i=1}^t g_i$, and (3) holds.

(3) $\Rightarrow$ (4). If $\widetilde{M}$ is a k-conjugate of M, then

(a) $\widetilde{M}$ is a splitting field (k-isomorphic to M)

(b) $\widetilde{M} \underset{\text{rel}}{\sim} M$.

But, we know that (a) and (b) imply that $\widetilde{M} = M$.

(4) $\Rightarrow$ (1). Pick an irreducible k-polynomial, g, and $\alpha \in M$ with $g(\alpha) = 0$. Consider the extension lemma in the situation where $k = \widetilde{k}$ and $\Gamma = M$. Pick in an algebraic closure, $\overline{M}$, of M, any root β of g. We get the diagram

$$\begin{array}{ccccc}
M & \xrightarrow{\widetilde{\theta}} & \widetilde{M} & \underset{\text{rel}}{\sim} & M \\
\uparrow & & \uparrow & & \| \\
k(\alpha) & \xrightarrow{\theta} & \widetilde{k}(\beta) & \underset{\text{rel}}{\sim} & M \\
\uparrow & & \uparrow & & \| \\
k & = & k & \underset{\text{rel}}{\sim} & M.
\end{array}$$

By the extension lemma applied to the upper portion of the above diagram, there exists $\widetilde{M}$ with $\widetilde{M} \underset{\text{rel}}{\sim} M$ and an extension $\widetilde{\theta}\colon M \to \widetilde{M}$. But, $\widetilde{\theta} \restriction k = \theta \restriction k = \text{id}$, so $\widetilde{\theta}$ is a k-isomorphism and $\widetilde{M} \underset{\text{rel}}{\sim} M$. By (4), we get $\widetilde{M} = M$. Since $\beta \in \widetilde{M}$, we have $\beta \in M$. $\square$

Corollary 4.26 *Say $M \supseteq K \supseteq k$ and M is normal over k. Then, M is normal over K.*

Proof. Use (3), i.e., M is the splitting field of some $g \in k[X]$. Yet, $g \in K[X]$, and use (3) again. $\square$

M normal over K and K normal over k does *not* imply M normal over k.

Here is a counter-example to the transitivity of normality. Let $k = \mathbb{Q}$; $K = \mathbb{Q}(\sqrt{2})$; the extension K/k is normal. Let $\alpha = \sqrt{2}$ and $L = K(\sqrt{\alpha})$; again, L/K is normal of degree 2. Observe that L is the splitting field over K of $X^2 - \alpha \in K[X]$. But, $L/\mathbb{Q}$ is *not* normal. This is because the polynomial $X^4 - 2$ has a root, $\sqrt{\alpha}$, in L, yet $i\sqrt{\alpha}$ is *not* in L because $L \subseteq \mathbb{R}$.

M normal over k and $M \supseteq K \supseteq k$ does *not* imply K normal over k.

Corollary 4.27 *(SMA, I[2]) Say M is normal over k and g is any irreducible k-polynomial with a root $\alpha \in M$. Then, a n.a.s.c. that an element $\beta \in M$ be a root of g is that there exists σ, a k-automorphism of M (i.e., $\sigma \restriction k = \text{id}$) so that $\sigma(\alpha) = \beta$.*

Proof. ($\Leftarrow$). If α is a root and $g \in k[X]$, then

$$0 = \sigma(0) = \sigma(g(\alpha)) = g(\sigma(\alpha)) = g(\beta).$$

So, β is a root.

[2]SMA = **s**ufficiently **m**any **a**utomorphisms.

($\Rightarrow$). Say $\beta \in M$ is a root, then there is a k-isomorphism $k(\alpha) \longrightarrow k(\beta)$. Now, $k(\beta) \underset{\text{rel}}{\sim} M$; so, in the extension lemma, take $\Gamma = M$:

$$\begin{array}{ccccc}
M & \xrightarrow{\widetilde{\theta}} & \widetilde{M} & \underset{\text{rel}}{\sim} & M \\
\uparrow & & \uparrow & & \| \\
k(\alpha) & \xrightarrow{\theta} & \widetilde{k}(\beta) & \underset{\text{rel}}{\sim} & M \\
\uparrow & & \uparrow & & \| \\
k & = & k & \underset{\text{rel}}{\sim} & M.
\end{array}$$

We get $\widetilde{\theta}\colon M \to \widetilde{M}$, a k-isomorphism and $M \underset{\text{rel}}{\sim} \widetilde{M}$. By (4), $M = \widetilde{M}$. So, $\widetilde{\theta} = \sigma$ is our required automorphism (it takes α to β). $\square$

Corollary 4.28 *(SMA, II) Let M be normal over k and say K, K' are subextensions of the layer M/k (i.e, $M \supseteq K \supseteq k$ and $M \supseteq K' \supseteq k$). If $\theta\colon K \to K'$ is a k-isomorphism, then there is a k-automorphism, σ, of M so that $\sigma \restriction K = \theta$.*

Proof. Apply the extension lemma with $\Gamma = M$ to the situation

$$\begin{array}{ccccc}
M & \xrightarrow{\widetilde{\theta}} & \widetilde{M} & \underset{\text{rel}}{\sim} & M \\
\uparrow & & \uparrow & & \| \\
K & \xrightarrow{\theta} & K' & \underset{\text{rel}}{\sim} & M.
\end{array}$$

There exist $\widetilde{\theta}$ and $\widetilde{M}$ with $\widetilde{\theta}$ a k-isomorphism and $M \underset{\text{rel}}{\sim} \widetilde{M}$. By (4), $M = \widetilde{M}$. Therefore, $\sigma = \widetilde{\theta}$ is our automorphism.

Corollary 4.29 *Say K/k is a finite extension of degree $[K\colon k] = n$, then there exists $M \supseteq K$ with*

(1) M is normal over k and

(2) Whenever W is normal over k, $W \supseteq K$ and $W \underset{\text{rel}}{\sim} M$, then automatically $W \supseteq M$.

(3) $[M\colon k] \leq n!$.

The field, M, is called a normal closure of K/k.

Proof. (DX).

4.5 The Theorems of Dedekind and Artin; Galois Groups & the Fundamental Theorem

Recollect that a K-representation of a group, G, is just a $K[G]$-module. So, a K-representation of a group, G, is just a K-vector space plus a (linear) G-action on it (by K-automorphisms); that is, a homomorphism $G \longrightarrow \mathrm{Aut}(V)$. If $\dim_K V < \infty$, we have a finite dimensional representation. In this case, $\mathrm{Aut}(V) = \mathrm{GL}(V) \cong \mathrm{GL}_n(K)$, where $n = \dim_K(V)$ is the degree of the representation. Say $\rho\colon G \to \mathrm{GL}_n(K)$ is our representation. Then, $\chi_\rho(\sigma) = \mathrm{tr}(\rho(\sigma))$, the trace of $\rho(\sigma)$, is a function $G \longrightarrow K$ independent of the basis chosen, called the *character* of our representation. The case $n = 1$ is very important. In this case, the characters *are* the representations, $\chi_\rho = \rho$. Therefore, we have functions $\chi\colon G \to K^*$, with $\chi(\sigma\tau) = \chi(\sigma)\chi(\tau)$. From now on, we use only one-dimensional characters.

Definition 4.10 Suppose $\{\chi_\alpha\}_\alpha$ is a given family of characters, $\chi_\alpha\colon G \to K^*$, of the group G. Call the family *independent* iff the relation

$$\sum_{j=1}^{n} a_j\chi_j(\sigma) = 0, \quad \textit{for all } \sigma \in G$$

implies $a_j = 0$, for $j = 1, \ldots, n$ (all applicable n).

Theorem 4.30 *(R. Dedekind, about 1890) If G is a group and $\{\chi_\alpha\}_\alpha$ is a family of mutually distinct characters of G (with values in K^*), then they are independent.*

Proof. We may assume our family is finite and we use induction on the number of elements, n, in this family. The case $n = 1$ holds trivially. Assume that the result holds for all $t \leq n - 1$ characters. Say $\chi_1, \ldots, \chi_n$ are distinct characters of G and suppose

$$\sum_{j=1}^{n} a_j\chi_j(\sigma) = 0, \quad \text{for all } \sigma \in G. \tag{$*$}$$

The induction hypothesis implies that if the conclusion of the theorem is false, then $a_j \neq 0$, for all $j = 1, \ldots, n$. Since the χ_j are distinct, there is some $\alpha \neq 1$ with $\chi_1(\alpha) \neq \chi_n(\alpha)$. Divide $(*)$ by $a_n \neq 0$, to obtain

$$\sum_{j=1}^{n-1} b_j\chi_j(\sigma) + \chi_n(\sigma) = 0, \quad \text{for all } \sigma \in G. \tag{$**$}$$

Consider the group element $\alpha\sigma$, then $(**)$ is true for it and we have

$$\sum_{j=1}^{n-1} b_j\chi_j(\alpha)\chi_j(\sigma) + \chi_n(\alpha)\chi_n(\sigma) = 0, \quad \text{for all } \sigma \in G.$$

If we multiply by $\chi_n(\alpha)^{-1}$, we get

$$\sum_{j=1}^{n-1} (b_j\chi_n(\alpha)^{-1}\chi_j(\alpha))\chi_j(\sigma) + \chi_n(\sigma) = 0, \quad \text{for all } \sigma \in G. \tag{$\dagger$}$$

If we subtract $(\dagger)$ from $(**)$, we get

$$\sum_{j=1}^{n-1} b_j(1 - \chi_n(\alpha)^{-1}\chi_j(\alpha))\chi_j(\sigma) = 0, \quad \text{for all } \sigma \in G.$$

By the induction hypothesis, $b_j(1 - \chi_n(\alpha)^{-1}\chi_j(\alpha)) = 0$, for $j = 1, \ldots, n-1$. If we take $j = 1$ and we remember that $b_1 = a_1/a_n \neq 0$, we get

$$1 - \chi_n(\alpha)^{-1}\chi_1(\alpha) = 0,$$

i.e., $\chi_n(\alpha) = \chi_1(\alpha)$, a contradiction. □

Corollary 4.31 *Say $\{\chi_\alpha\}_\alpha$ is a family of mutually distinct isomorphisms of a field L with another, $\widetilde{L}$. Then, the χ_α are independent.*

Proof. Take $G = L^*$ and $K = \widetilde{L}$ in Dedekind's theorem. □

Definition 4.11 If $\{\chi_\alpha\}_\alpha$ is a family of isomorphisms $K \longrightarrow \widetilde{K}$, then set

$$\mathrm{Fix}(\{\chi_\alpha\}) = \{\xi \in K \mid (\forall \alpha, \beta)(\chi_\alpha(\xi) = \chi_\beta(\xi))\}.$$

Observe that $\mathrm{Fix}(\{\chi_\alpha\})$ is always a subfield of K, so we call it the *fixed field* of $\{\chi_\alpha\}_\alpha$.

Note that $\mathrm{Fix}(\{\chi_\alpha\})$ contains the prime field of K.

Theorem 4.32 *(E. Artin, 1940) If $\{\chi_\alpha\}_\alpha$ is a family of pairwise distinct isomorphisms $K \longrightarrow \widetilde{K}$ and if $k = \mathrm{Fix}(\{\chi_\alpha\})$, then*

(1) $[K\colon k] \geq \min(\aleph_0, \#(\{\chi_\alpha\}))$.

(2) Say $\{\chi_\alpha\}$ forms a group under composition (so, $K = \widetilde{K}$ and all χ_α's are automorphisms of K), then if $\#(\{\chi_\alpha\}) = n < \infty$, we have $[K\colon k] = n$ and if $n = \infty$ then $[K\colon k] = \infty$.

Proof. (1) First, we consider the case where we have a finite set, $\{\chi_1, \ldots, \chi_n\}$, of isomorphisms $K \longrightarrow \widetilde{K}$. Let $k = \mathrm{Fix}(\{\chi_j\}_{j=1}^n)$ and assume that $[K\colon k] < n$. Then, there exists a basis, $\omega_1, \ldots, \omega_r$, of K/k and $r < n$. Consider the r equations in n unknowns (y_j's)

$$\sum_{j=1}^{n} y_j \chi_j(\omega_i) = 0, \quad 1 \leq i \leq r.$$

As $r < n$, this system has a nontrivial solution, call it $(\alpha_1, \ldots, \alpha_n)$ (with $\alpha_i \in \widetilde{K}$). So, we have

$$\sum_{j=1}^{n} \alpha_j \chi_j(\omega_i) = 0, \quad 1 \leq i \leq r.$$

Pick any $\xi \in K$, as the ω_i's form a basis, we can write $\xi = \sum_{i=1}^r a_i\omega_i$, for some (unique) $a_i \in k$. We have

$$\sum_{j=1}^{n} \alpha_j \chi_j(\xi) = \sum_{j=1}^{n} \alpha_j \chi_j\Big(\sum_{i=1}^{r} a_i\omega_i\Big) = \sum_{j=1}^{n}\sum_{i=1}^{r} \alpha_j \chi_j(a_i)\chi_j(\omega_i).$$

But, $\chi_j(a_i) = \chi_l(a_i)$, for all j, l, as $a_i \in k$ and $k = \mathrm{Fix}(\{\chi_j\})$. Write $b_i = \chi_j(a_i)$ (independent of j). So, we have

$$\sum_{j=1}^{n} \alpha_j \chi_j(\xi) = \sum_{i=1}^{r} b_i\Big(\sum_{j=1}^{n} \alpha_j \chi_j(\omega_i)\Big).$$

But, $\sum_{j=1}^n \alpha_j\chi_j(\omega_i) = 0$, by the choice of $\alpha_1, \ldots, \alpha_n$, so

$$\sum_{j=1}^{n} \alpha_j \chi_j(\xi) = 0, \quad \text{for all } \xi.$$

This contradicts Dedekind's theorem and thus, $[K\colon k] \geq n$.

Now, consider the case where $\#(\{\chi_\alpha\})$ is infinite. If $[K\colon k]$ were finite, then pick any $n > [K\colon k]$ and repeat the above argument with the subset $\{\chi_1, \ldots, \chi_n\}$. We deduce that $[K\colon k]$ must be infinite.

(2) Now, suppose $\{\chi_1, \ldots, \chi_n\}$ forms a group under composition (i.e., they are a group of automorphisms of K). Then, one of the χ_j's is the identity, say $\chi_1 = \mathrm{id}$. It follows that for every $a \in k$, we have $\chi_j(a) = \chi_1(a) = a$, so

$$k = \mathrm{Fix}(\{\chi_j\}) = \{a \in K \mid \chi_j(a) = a, \quad j = 1, \ldots, n\}.$$

By part (1), we know $[K\colon k] \geq n$; so, assume $[K\colon k] > n$. In this case, there exist $r > n$ elements, $\omega_1, \ldots, \omega_r \in K$, linearly independent over k. Consider the n equations in r unknowns (x_i's)

$$\sum_{i=1}^{r} x_i \chi_j(\omega_i) = 0, \quad j = 1, \ldots, n.$$

Again, there is a nontrivial solution, say $a_1, \ldots, a_r$, with $a_j \in K$. So, we have

$$\sum_{i=1}^{r} a_i \chi_j(\omega_i) = 0, \quad j = 1, \ldots, n. \tag{$\dagger$}$$

Note that for any nontrivial solution, the a_i's can't all be in k. If they were, then ($\dagger$) with $j = 1$ gives $\sum_{i=1}^{r} a_i \omega_i = 0$, contradicting the linear independence of the ω_i's.

Pick a solution containing a minimal number of nonzero a_i's, say $a_1, \ldots, a_s \neq 0$ and $a_{s+1} = \cdots = a_r = 0$. If we divide ($\dagger$) by a_s, we get

$$\sum_{i=1}^{s-1} b_i \chi_j(\omega_i) + \chi_j(\omega_s) = 0, \quad j = 1, \ldots, n. \tag{$\dagger\dagger$}$$

By the remark above, there is some i, with $1 \leq i \leq s-1$, so that $b_i \notin k$. By relabelling, we may assume that $b_1 \notin k$. As $b_1 \notin k$, there is some ρ ($1 \leq \rho \leq n$) with $\chi_\rho(b_1) \neq b_1$. Apply χ_ρ to ($\dagger\dagger$); we get

$$\sum_{i=1}^{s-1} \chi_\rho(b_i)(\chi_\rho \circ \chi_j)(\omega_i) + (\chi_\rho \circ \chi_j)(\omega_s) = 0, \quad j = 1, \ldots, n.$$

As χ_j ranges over $\{\chi_1, \ldots, \chi_n\}$, so does $\chi_\rho \circ \chi_j$; consequently, we have

$$\sum_{i=1}^{s-1} \chi_\rho(b_i) \chi_\xi(\omega_i) + \chi_\xi(\omega_s) = 0, \quad \xi = 1, \ldots, n. \tag{$*$}$$

If we subtract ($*$) from ($\dagger\dagger$), we obtain

$$\sum_{i=1}^{s-1} (b_i - \chi_\rho(b_i)) \chi_\xi(\omega_i) = 0, \quad \xi = 1, \ldots, n.$$

But, we know that $b_1 \neq \chi_\rho(b_1)$. For this ρ, not all the coefficients are zero, so we get a solution with strictly fewer nonzero components, a contradiction to the minimality of $(a_1, \ldots, a_s)$. $\square$

Definition 4.12 If Ω is a finite, normal extension of k, the *Galois group* of Ω/k, denoted $\mathcal{G}(\Omega/k)$, is the group of all k-automorphisms of Ω (i.e., the automorphisms, σ, of Ω so that $\sigma \restriction k = \mathrm{id}$). Say $f \in k[X]$ and let Ω be a splitting field for $f(X)$ over k. The *Galois group of the polynomial, $f(X)$,* over k, denoted $\mathcal{G}_k(f)$, is just $\mathcal{G}(\Omega/k)$.

Lemma 4.33 *Suppose Ω is finite, normal over k and $\mathcal{G} = \mathcal{G}(\Omega/k)$ is its Galois group. Then, a n.a.s.c. that $\xi \in \Omega$ lie in $\mathrm{Fix}(\mathcal{G})$ is that ξ be purely inseparable over k.*

Proof. If ξ is purely inseparable over k, then there is some $s \geq 0$ so that $\xi^{p^s} \in k$. Then, for every $\sigma \in \mathcal{G}$, we have $\sigma\big(\xi^{p^s}\big) = \xi^{p^s}$. But, $\sigma\big(\xi^{p^s}\big) = (\sigma(\xi))^{p^s}$, so $(\sigma(\xi))^{p^s} = \xi^{p^s}$; since $\mathrm{char}(k) = p$, it follows that $(\sigma(\xi) - \xi)^{p^s} = 0$. Therefore, $\sigma(\xi) - \xi = 0$, i.e., ξ is fixed by σ and $\xi \in \mathrm{Fix}(\mathcal{G})$. Conversely, assume that $\xi \in \mathrm{Fix}(\mathcal{G})$. First, pick an element $\alpha \in \Omega$, with α separable over k and $\alpha \notin k$, if such an element exists. Then, α is a simple root of some irreducible k-polynomial g. But, Ω is normal, so all the roots of g lie in Ω and as $\alpha \notin k$, we have $\deg(g) > 1$. Consequently, there is another root, $\beta \in \Omega$, of g with $\beta \neq \alpha$ and by SMA, I, there exists $\sigma \in \mathcal{G}$ so that $\sigma(\alpha) = \beta$. Now, consider our $\xi \in \mathrm{Fix}(\mathcal{G})$. As $\xi \in \Omega$, there is some power, ξ^{p^r}, of ξ that is separable over k. Since ξ is fixed by all $\sigma \in \mathcal{G}$, so is ξ^{p^r}. If ξ^{p^r} were not in k, then ξ^{p^r} could play the role of α above, so it could be moved to some $\beta \neq \alpha$, a contradiction. This implies that $\xi^{p^r} \in k$, i.e., ξ is purely inseparable over k. □

Nomenclature & Notation.

Say Ω/k is a normal (not necessarily finite) extension. Pick an extension, K, in the layer Ω/k, i.e, $k \subseteq K \subseteq \Omega$. Define

$$K^{(*)} = \left\{\xi \in \Omega \mid \xi^{p^r} \in K, \quad \text{for some } r \geq 0\right\}.$$

(Obviously, $p = \mathrm{char}(k)$.) Note that $K^{(*)} = \Omega \cap K^{p^{-\infty}}$ in some algebraic closure (where $K^{p^{-\infty}}$ is defined as $\{\xi \in \overline{K} \mid (\exists r \geq 0)(\xi^{p^r} \in K)\}$). Also define

$$K_{(*)} = \{\xi \in K \mid \xi \quad \text{is separable over } k\}.$$

Note: $K^{(*)}$ and $K_{(*)}$ are subfields of Ω/k and we have $K_{(*)} \subseteq K \subseteq K^{(*)} \subseteq \Omega$.

We say that K *is Galois equivalent to* K' (where $k \subseteq K \subseteq \Omega$ and $k \subseteq K' \subseteq \Omega$) iff $K^{(*)} = K'^{(*)}$; write K gal K'. This equivalence relation fibers the subextensions of Ω/k into Galois equivalence classes.

Corollary 4.34 *If Ω/k is finite, normal, then $\mathrm{Fix}(\mathcal{G}(\Omega/k)) = k^{(*)}$. In particular, if $k \subseteq L \subseteq \Omega$, then $\mathrm{Fix}(\mathcal{G}(\Omega/L)) = L^{(*)}$.*

Corollary 4.35 *If Ω/k is finite, normal, then $\#(\mathcal{G}(\Omega/k))$ divides $[\Omega\colon k]$; in particular, $\#(\mathcal{G}(\Omega/k)) \leq [\Omega\colon k] < \infty$.*

Proof. By Artin's theorem (Theorem 4.32) $\#(\mathcal{G}(\Omega/k)) = [\Omega\colon \mathrm{Fix}(\mathcal{G}(\Omega/k))]$. By Lemma 4.33, we have $\mathrm{Fix}(\mathcal{G}(\Omega/k)) = k^{(*)}$. Therefore, $\#(\mathcal{G}(\Omega/k)) = [\Omega\colon k^{(*)}]$ which divides $[\Omega\colon k]$. □

Corollary 4.36 *If Ω/k is finite, normal and k is perfect, e.g.* $\mathrm{char}(k) = 0$, *then $\#(\mathcal{G}(\Omega/k)) = [\Omega\colon k]$.*

Corollary 4.37 *Say f is a separable, irreducible k-polynomial with degree* $\deg(f) = n$. *Then, there is an injection $\mathcal{G}_k(f) \hookrightarrow \mathfrak{S}_n$ (where $\mathfrak{S}_n$ denotes the symmetric group on n elements) and this injection is unique up to inner automorphisms in $\mathfrak{S}_n$. In particular, $\#(\mathcal{G}_k(f)) \mid n!$.*

Proof. Write $\alpha_1, \ldots, \alpha_n$ for all the roots of f (they are all distinct) in some order. Given $\sigma \in \mathcal{G}_k(f)$, the element $\sigma(\alpha_i)$ is some other root of f, call it $\alpha_{p_\sigma(i)}$. Then, p_σ is a permutation of the n roots, i.e., $p_\sigma \in \mathfrak{S}_n$. Clearly, the map $\sigma \mapsto p_\sigma$ is a homomorphism $\mathcal{G}_k(f) \longrightarrow \mathfrak{S}_n$. If $p_\sigma = \mathrm{id}$, then $\sigma(\alpha_i) = \alpha_i$ for all i, so $\sigma \restriction \Omega = \mathrm{id}$, as Ω, the splitting field of f, is generated over k by the α_i's. So, $\sigma = id$ in $\mathcal{G}_k(f) = \mathcal{G}(\Omega/k)$, and the our map $\mathcal{G}_k(f) \longrightarrow \mathfrak{S}_n$ is an injection. We can reorder (relabel) the $\alpha_1, \ldots, \alpha_n$; to do so introduces an inner automorphism of $\mathfrak{S}_n$. □

Remarks: (On Galois equivalence)

(1) *If $K \subseteq K'$, then $K^{(*)} \subseteq K'^{(*)}$.* Indeed, if $\xi \in K^{(*)}$, then $\xi^{p^r} \in K \subseteq K'$ (for some $r \geq 0$), so $\xi \in K'^{(*)}$.

(2) *For all K in the layer Ω/k (of course, Ω/k is a finite normal extension), we have K gal $K^{(*)}$. Hence, the Galois equivalence class of any field possesses a unique least upper bound, namely $K^{(*)}$ for any K in the class.* For, $K \subseteq K^{(*)}$, so $K^{(*)} \subseteq (K^{(*)})^{(*)}$. Also, if $\xi \in (K^{(*)})^{(*)}$, then $\xi^{p^r} \in K^{(*)}$, for some r; but then, $(\xi^{p^r})^{p^s} \in K$, for some s, i.e., $\xi^{p^{r+s}} \in K$, which means that $\xi \in K^{(*)}$. Consequently, $(K^{(*)})^{(*)} \subseteq K^{(*)}$ and so $(K^{(*)})^{(*)} = K^{(*)}$, i.e. K gal $K^{(*)}$. If K gal L then $K^{(*)} = L^{(*)}$; $K \subseteq K^{(*)}$ and $L \subseteq L^{(*)}$, so $K^{(*)} = L^{(*)}$ is indeed the least upper bound of the equivalence class of K and L.

(3) *If K belongs to the layer Ω/k (where Ω/k is normal), then $K_{(*)}$ gal K and $K_{(*)}$ is the unique greatest lower bound for the Galois equivalence class of K.*

Proof. If we prove that $(K_{(*)})^{(*)} = K^{(*)}$ and $(K^{(*)})_{(*)} = K_{(*)}$, we are done. The first equation will prove that $K_{(*)}$ gal K. As $K_{(*)} \subseteq K$, we get $(K_{(*)})^{(*)} \subseteq K^{(*)}$. Pick $\xi \in K^{(*)}$, then $\xi^{p^r} \in K$, for some r and $(\xi^{p^r})^{p^s} = \xi^{p^{r+s}} \in K_{(*)}$, for some s, so $\xi \in (K_{(*)})^{(*)}$; hence, $(K_{(*)})^{(*)} = K^{(*)}$. Now, pick $\xi \in K_{(*)}$, then $\xi \in K^{(*)}$ (as $K_{(*)} \subseteq K \subseteq K^{(*)}$) and since ξ is separable over k, we have $\xi \in (K^{(*)})_{(*)}$. Conversely, if $\xi \in (K^{(*)})_{(*)}$, then $\xi \in K^{(*)}$, which means that ξ is inpurely separable over K. Yet, ξ is separable over k, so ξ is separable over K. As ξ is purely inseparable over K and separable over K, we get $\xi \in K$; moreover, as ξ is separable over k, we get $\xi \in K_{(*)}$.

(4) *We have K gal L iff $K_{(*)} = L_{(*)}$, hence in each Galois equivalence class, there is a unique greatest lower bound, it is the common $K_{(*)}$.* If K gal L, then $K^{(*)} = L^{(*)}$, so

$$K_{(*)} = (K^{(*)})_{(*)} = (L^{(*)})_{(*)} = L_{(*)},$$

by (3). Conversely, if $K_{(*)} = L_{(*)}$, then

$$K^{(*)} = (K_{(*)})^{(*)} = (L_{(*)})^{(*)} = L^{(*)},$$

again, by (3), i.e., K gal L.

(5) *Suppose K gal L and $K, L \subseteq \Omega/k$, Then, $\mathcal{G}(\Omega/K) = \mathcal{G}(\Omega/L)$, hence the maps*

$$\mathcal{G}(\Omega/L^{(*)}) \hookrightarrow \mathcal{G}(\Omega/L) \hookrightarrow \mathcal{G}(\Omega/L_{(*)})$$

are equalities. All we need show is $\mathcal{G}(\Omega/L) = \mathcal{G}(\Omega/L^{(*)})$. We already know $\mathcal{G}(\Omega/L^{(*)}) \subseteq \mathcal{G}(\Omega/L)$, as $L \subseteq L^{(*)}$. Say $\sigma \in \mathcal{G}(\Omega/L)$ and pick any $\xi \in L^{(*)}$. Then, $\xi^{p^r} \in L$, for some $r \geq 0$. Consequently, $\sigma\big(\xi^{p^r}\big) = \xi^{p^r}$, as $\sigma \restriction L = \mathrm{id}$. As σ is an automorphism, we get $(\sigma(\xi))^{p^r} = \xi^{p^r}$, i.e., $(\sigma(\xi) - \xi)^{p^r} = 0$, and so, $\sigma(\xi) = \xi$. As ξ is arbitrary in $L^{(*)}$, we have $\sigma \restriction L^{(*)} = \mathrm{id}$; since σ is arbitrary, the proof is complete.

Theorem 4.38 *(Fundamental Theorem of Galois Theory) Suppose Ω/k is a finite, normal extension. Write $\mathcal{G}$ for $\mathcal{G}(\Omega/k)$ and write $[K]$ for the Galois class of $K \subseteq \Omega/k$. Then, the maps*

$$\mathcal{H} \mapsto [\mathrm{Fix}(\mathcal{H})] \quad \textit{and} \quad [L] \mapsto \mathcal{G}(\Omega/L)$$

establish a one-to-one order-reversing correspondence between all subgroups of $\mathcal{G}$ and all the Galois classes of subextensions $L/k \subseteq \Omega/k$. Here, $[K] \subseteq [L]$ means $K^{(*)} \subseteq L^{(*)}$ as fields. In this correspondence, $\mathcal{G}(\Omega/L) \lhd \mathcal{G}$ iff $L^{(*)}$ is a normal extension of k iff $L_{(*)}$ is a normal extension of k. When the latter is the case, then there is a canonical exact sequence

$$0 \longrightarrow \mathcal{G}(\Omega/L) \longrightarrow \mathcal{G}(\Omega/k) \longrightarrow \mathcal{G}(L^{(*)}/k) \longrightarrow 0.$$

Claim 1. If $L = \mathrm{Fix}(\mathcal{H})$, then $L = L^{(*)}$.

Pick $\xi \in L^{(*)}$, so $\xi^{p^r} \in L$, for some $r \geq 0$. Then, for all $\sigma \in \mathcal{H}$, we have $\sigma\big(\xi^{p^r}\big) = \xi^{p^r}$, and by a standard argument, $\xi \in \mathrm{Fix}(\mathcal{H}) = L$. Consequently, $L^{(*)} \subseteq L$, yet $L \subseteq L^{(*)}$, so $L = L^{(*)}$.

Proof of Theorem 4.38. Now say $\mathcal{H} \subseteq \widetilde{\mathcal{H}}$ and look at $\mathrm{Fix}(\widetilde{\mathcal{H}})$. If $\xi \in \mathrm{Fix}(\widetilde{\mathcal{H}})$, then for every $\tau \in \widetilde{\mathcal{H}}$, we have $\tau(\xi) = \xi$ and so, for every $\sigma \in \mathcal{H}$, we have $\sigma(\xi) = \xi$, i.e., $\xi \in \mathrm{Fix}(\mathcal{H})$. Consequently, $\mathrm{Fix}(\widetilde{\mathcal{H}}) \subseteq \mathrm{Fix}(\mathcal{H})$ and so, $[\mathrm{Fix}(\widetilde{\mathcal{H}})] \subseteq [\mathrm{Fix}(\mathcal{H})]$, by Claim (1). Now, if $[L] \subseteq [\widetilde{L}]$, then $L^{(*)} \subseteq \widetilde{L}^{(*)}$. If $\sigma \in \mathcal{G}(\Omega/\widetilde{L})$, then $\sigma \in \mathcal{G}(\Omega/\widetilde{L}^{(*)})$ (by Remark (5), above); so, $\sigma \in \mathcal{G}(\Omega/L^{(*)}) = \mathcal{G}(\Omega/L)$ (again, by Remark (5)). Thus, $\mathcal{G}(\Omega/\widetilde{L}) \subseteq \mathcal{G}(\Omega/L)$.

Given $\mathcal{H} \subseteq \widetilde{\mathcal{H}}$, say we know $\mathrm{Fix}(\mathcal{H}) = \mathrm{Fix}(\widetilde{\mathcal{H}})$. By Artin's theorem, we have

$$\#(\mathcal{H}) = [\Omega\colon \mathrm{Fix}(\mathcal{H})] = [\Omega\colon \mathrm{Fix}(\widetilde{\mathcal{H}})] = \#(\widetilde{\mathcal{H}}).$$

As $\mathcal{H} \subseteq \widetilde{\mathcal{H}}$ and $\#(\mathcal{H}) = \#(\widetilde{\mathcal{H}})$, we get $\mathcal{H} = \widetilde{\mathcal{H}}$.

Choose a subgroup, $\mathcal{H}$, of $\mathcal{G}$ and let $L = \mathrm{Fix}(\mathcal{H})$; write $\widetilde{\mathcal{H}}$ for $\mathcal{G}(\Omega/L) = \mathcal{G}(\Omega/\mathrm{Fix}(\mathcal{H}))$. If $\sigma \in \mathcal{H}$, then σ fixes L, so $\sigma \in \widetilde{\mathcal{H}}$ and $\mathcal{H} \subseteq \widetilde{\mathcal{H}}$. But, $\mathrm{Fix}(\widetilde{\mathcal{H}}) = \mathrm{Fix}(\mathcal{G}(\Omega/L)) = L^{(*)}$, by Corollary 4.34. Thus, $\mathrm{Fix}(\widetilde{\mathcal{H}}) = (\mathrm{Fix}(\mathcal{H}))^{(*)}$. Claim 1 implies that $(\mathrm{Fix}(\mathcal{H}))^{(*)} = \mathrm{Fix}(\mathcal{H})$, so $\mathrm{Fix}(\widetilde{\mathcal{H}}) = \mathrm{Fix}(\mathcal{H})$ and, by the above, we get $\mathcal{H} = \widetilde{\mathcal{H}}$. Therefore, $\mathcal{H} = \mathcal{G}(\Omega/\mathrm{Fix}(\mathcal{H}))$.

Consider L, make $\mathcal{G}(\Omega/L)$ and form $\mathrm{Fix}(\mathcal{G}(\Omega/L))$. By Corollary 4.34, we have $\mathrm{Fix}(\mathcal{G}(\Omega/L)) = L^{(*)}$ and L gal $L^{(*)}$, so $[L] = [\mathrm{Fix}(\mathcal{G}(\Omega/L))]$.

Having proved all the statements about the order inverting correspondence, we see that only normality statements remain.

Claim 2. If $L \subseteq \Omega/k$, then L is normal over k iff for every $\sigma \in \mathcal{G}(\Omega/k)$, we have $\sigma(L) = L$.

($\Rightarrow$). For every $\sigma \in \mathcal{G}(\Omega/k)$, the field $\sigma(L)$ is k-conjugate to L. As L is normal over k, we find $\sigma(L) = L$.

($\Leftarrow$). Assume $\sigma(L) = L$, for every $\sigma \in \mathcal{G}(\Omega/k)$. Let g be any irreducible k-polynomial and assume that $\alpha \in L$ is a root of g. But, $\alpha \in \Omega$ and Ω is normal; consequently, *all* the roots of g lie in Ω. Say $\beta \in \Omega$ is any other root of g. By SMA, I, there is some $\sigma \in \mathcal{G}$ so that $\sigma(\alpha) = \beta$. So, $\beta \in \sigma(L)$, and as $\sigma(L) = L$, we get $\beta \in L$. Thus, L contains all the roots of g which means that L is normal over k.

Assume $\mathcal{G}(\Omega/L) \lhd \mathcal{G}$. Look at $L^{(*)}$ and choose any $\sigma \in \mathcal{G}$ and any $\eta \in \sigma(L^{(*)})$. Then, $\sigma^{-1}(\eta) \in L^{(*)}$ and for all $\tau \in \mathcal{G}(\Omega/L) = \mathcal{G}(\Omega/L^{(*)})$, we have

$$(\sigma\tau\sigma^{-1})(\eta) = \sigma(\tau(\sigma^{-1}(\eta))) = (\sigma\sigma^{-1})(\eta) = \eta,$$

because $\sigma^{-1}(\eta) \in L^{(*)}$. Thus, $(\sigma\mathcal{G}(\Omega/L)\sigma^{-1})(\eta) = \eta$, and as $\mathcal{G}(\Omega/L) \lhd \mathcal{G}$, we get $\mathcal{G}(\Omega/L)(\eta) = \eta$, so $\eta \in \mathrm{Fix}(\mathcal{G}(\Omega/L)) = L^{(*)}$, as we know. In summary, if $\eta \in \sigma(L^{(*)})$, then $\eta \in L^{(*)}$, i.e., $\sigma(L^{(*)}) \subseteq L^{(*)}$. If we apply this to σ^{-1}, we get $\sigma^{-1}(L^{(*)}) \subseteq L^{(*)}$, i.e. $L^{(*)} \subseteq \sigma(L^{(*)})$. Therefore, $L^{(*)} = \sigma(L^{(*)})$ and by Claim 2, the extension $L^{(*)}/k$ is normal.

Now, say $L^{(*)}$ is normal over k. Then, we know $\sigma(L^{(*)}) = L^{(*)}$, for all $\sigma \in \mathcal{G}(\Omega/k)$. For any $\xi \in L^{(*)}$ and any $\tau \in \mathcal{G}(\Omega/L)$, we have

$$(\sigma\tau\sigma^{-1})(\xi) = \sigma(\tau(\sigma^{-1}(\xi))) = (\sigma\sigma^{-1})(\xi) = \xi,$$

because $\sigma^{-1}(\xi) \in \sigma^{-1}(L^{(*)}) = L^{(*)}$, by hypothesis. So, $\sigma\tau\sigma^{-1} \in \mathcal{G}(\Omega/L^{(*)}) = \mathcal{G}(\Omega/L)$ and thus, $\mathcal{G}(\Omega/L) \lhd \mathcal{G}$.

Suppose $L^{(*)}$ is normal. We have a map $\mathcal{G}(\Omega/k) \longrightarrow \mathcal{G}(L^{(*)}/k)$ *via* $\sigma \mapsto \sigma \restriction L^{(*)}$ ($\sigma \restriction L^{(*)} \in \mathcal{G}(L^{(*)}/k)$, by normality). This map is onto because, given any $\sigma \in \mathcal{G}(L^{(*)}/k)$, we have the diagram

$$\begin{array}{ccc} \Omega & \xrightarrow{\widetilde{\sigma}} & \Omega \\ \uparrow & & \uparrow \\ L^{(*)} & \xrightarrow{\sigma} & L^{(*)} \\ \uparrow & & \uparrow \\ k & = & k, \end{array}$$

and by SMA, II, the automorphism σ lifts to an automorphism, $\widetilde{\sigma}$, of Ω. The kernel of our map is clearly $\mathcal{G}(\Omega/L)$.

Lastly, we need to show that $L^{(*)}$ is normal iff $L_{(*)}$ is normal. Say $L^{(*)}$ is normal and $\sigma \in \mathcal{G}$. If $\xi \in L_{(*)}$, then $\xi \in L^{(*)}$ and $\sigma(\xi) \in L^{(*)}$ (as $L^{(*)}$ is normal). But, $\sigma(\xi)$ is separable over k as ξ is. It follows that $\sigma(\xi) \in (L^{(*)})_{(*)} = L_{(*)}$ and so, $\sigma(L_{(*)}) \subseteq L_{(*)}$. By the usual argument, $\sigma(L_{(*)}) = L_{(*)}$ and $L_{(*)}$ is normal. If $L_{(*)}$ is normal and $\xi \in L^{(*)}$, then $\xi^{p^r} \in L_{(*)}$, for some $r \geq 0$. It follows that $\sigma\big(\xi^{p^r}\big) \in \sigma(L_{(*)}) = L_{(*)}$, so $(\sigma(\xi))^{p^r} \in L_{(*)}$, i.e., $\sigma(\xi) \in (L_{(*)})^{(*)} = L^{(*)}$; thus, $\sigma(L^{(*)}) \subseteq L^{(*)}$ and, by the usual argument, we conclude that $L^{(*)}$ is normal. $\square$

Proposition 4.39 *Suppose Ω is normal over k and $L/k \subseteq \Omega/k$. Then $L = L^{(*)}$ iff Ω is separable over L.*

Proof. ($\Rightarrow$). Say Ω is separable over L, then as $L^{(*)} \subseteq \Omega$, we find $L^{(*)}$ is separable over L. Yet, $L^{(*)}$ is purely inseparable over L. It follows that $L = L^{(*)}$.

($\Leftarrow$). We must prove that Ω is separable over $L^{(*)}$. Pick $\alpha \in \Omega$ and consider $\mathcal{G}(\Omega/L^{(*)})$. Choose $\sigma_1, \ldots, \sigma_n \in \mathcal{G}(\Omega/L^{(*)})$ so that

(1) $\sigma_1 = \mathrm{id}$ and $\alpha = \sigma_1(\alpha), \sigma_2(\alpha), \ldots, \sigma_n(\alpha)$ are mutually distinct,

(2) n is maximal, i.e., no further $\sigma \in \mathcal{G}(\Omega/L^{(*)})$ can be added while preserving (1).

Consider $g(X) = \prod_{i=1}^{n}(X - \sigma_j(\alpha))$. If $\sigma \in \mathcal{G}(\Omega/L^{(*)})$, the elements $\sigma\sigma_1(\alpha), \ldots, \sigma\sigma_n(\alpha)$ are a permutation of $\alpha, \sigma_2(\alpha), \ldots, \sigma_n(\alpha)$, so $\sigma g(X) = g(X)$. This implies that the coefficients of $g(X)$ belong to $\mathrm{Fix}(\mathcal{G}(\Omega/L^{(*)})) = L^{(*)}$. Thus, $g(X) \in L^{(*)}[X]$, but the roots of $g(X)$ are distinct and α is among them. Therefore, α is separable over $L^{(*)}$. $\square$

Corollary 4.40 *Assume Ω/k is a finite normal extension. Then, the following are equivalent:*

(1) Ω is separable over k.

(2) $k^{(*)} = k$.

(3) For all subextensions, L, of Ω/k, we have $L^{(*)} = L_{(*)}$.

(3a) For all subextensions, L, of Ω/k, the equivalence class $[L]$ has but one element.

(4) Same as (3) but for some extension $L/k \subseteq \Omega/k$.

(4a) Same as (3a) but for some extension $L/k \subseteq \Omega/k$.

(5) $\Omega = \Omega_{(*)}$.

Proof. First, observe that the equivalences (3) $\iff$ (3a) and (4) $\iff$ (4a) are obvious.

(1) $\Rightarrow$ (2). This is Proposition 4.39 when $L = k$.

(2) $\Rightarrow$ (3). Given $L \subseteq \Omega/k$, then $L^{(*)} \subseteq \Omega$. By Proposition 4.39, Ω is separable over k. Thus, $L^{(*)}$ is separable over k and so $L^{(*)}$ is separable over $L_{(*)}$; yet, $L^{(*)}$ is purely inseparable over $L_{(*)}$, so $L^{(*)} = L_{(*)}$.

(3) $\Rightarrow$ (4) is a tautology.

(4) $\Rightarrow$ (5). We have $L^{(*)} = L_{(*)}$, for some $L \subseteq \Omega/k$. Proposition 4.39 implies that Ω is separable over $L_{(*)}$. But, $L_{(*)}$ is always separable over k and separability is transitive, so Ω is separable over k, i.e., $\Omega = \Omega_{(*)}$.

(5) $\Rightarrow$ (1). By definition, $\Omega_{(*)}$ is separable over k and $\Omega = \Omega_{(*)}$, so Ω is separable over k. $\square$

Proposition 4.41 *Say Ω/k is a finite normal extension. Then, $\Omega = \Omega_{(*)}k^{(*)}$ (= the smallest field containing $\Omega_{(*)}$ and $k^{(*)}$). The natural map*

$$\Omega_{(*)} \otimes_k k^{(*)} \longrightarrow \Omega$$

is an isomorphism. Indeed, for all $L/k \subseteq \Omega/k$, we have

(1) $L^{(*)} = Lk^{(*)} = L_{(*)}k^{(*)}$.

(2) $L_{(*)} = L \cap \Omega_{(*)}$.

(3) The natural map

$$L_{(*)} \otimes_k k^{(*)} \longrightarrow L^{(*)}$$

is an isomorphism.

Proof. We just have to prove (1)–(3) for $L/k \subseteq \Omega/k$ and then set $L = \Omega$ to get the rest.

(1) Since $L^{(*)} \supseteq k^{(*)}$ and $L^{(*)} \supseteq L \supseteq L_{(*)}$, we deduce that $L^{(*)} \supseteq L_{(*)}k^{(*)}$ and $L^{(*)} \supseteq Lk^{(*)}$. If $\xi \in L^{(*)}$, then ξ is purely inseparable over $L_{(*)}$, so ξ is purely inseparable over $L_{(*)}k^{(*)}$. If $\xi \in L^{(*)}$, then ξ is separable over $k^{(*)}$ (by Proposition 4.39), so ξ is separable over $L_{(*)}k^{(*)}$. Thus, $L^{(*)}$ is both separable and purely inseparable over $L_{(*)}k^{(*)}$, which means that $L^{(*)} = L_{(*)}k^{(*)}$

(2) This is the definition of $L_{(*)}$, as $L \subseteq \Omega$.

(3) The (illegal definition of the) map is $\alpha \otimes \beta \mapsto \alpha\beta$. The image is $L_{(*)}k^{(*)} = L^{(*)}$. So, we need to prove our map is injective. Now, $k^{(*)} \subseteq k^{p^{-\infty}}$ (where $k^{p^{-\infty}} = \{\xi \in \overline{k} \mid \xi^{p^r} \in k, \text{ for some } r \geq 0\}$). By Mac Lane I and right limits, we get

$$L_{(*)} \otimes_k k^{p^{-\infty}} \longrightarrow L_{(*)}k^{p^{-\infty}}$$

is injective (because $L_{(*)}$ is separable over k). But, $0 \longrightarrow k^{(*)} \longrightarrow k^{p^{-\infty}}$ is exact and vector spaces over a field are flat, so

$$0 \longrightarrow L_{(*)} \otimes_k k^{(*)} \longrightarrow L_{(*)} \otimes_k k^{p^{-\infty}}$$

is still exact. Then, the diagram

$$\begin{array}{ccccc} 0 & \longrightarrow & L_{(*)} \otimes_k k^{(*)} & \longrightarrow & L_{(*)} \otimes_k k^{p^{-\infty}} \\ & & \downarrow & & \downarrow \\ & & L_{(*)}k^{(*)} & \hookrightarrow & L_{(*)}k^{p^{-\infty}} \end{array}$$

commutes, and this shows that $L_{(*)} \otimes_k k^{(*)} \longrightarrow L_{(*)}k^{(*)} = L^{(*)}$ is injective. $\square$

Proposition 4.42 *Suppose Ω/k is a finite normal extension and $\mathcal{G} = \mathcal{G}(\Omega/k)$. Let $L/k \subseteq \Omega/k$ and $\mathcal{H} = \mathcal{G}(\Omega/L)$. Then,*

(1) $[\Omega\colon L^{(*)}] = \#(\mathcal{H})$.

(2) $[L_{(*)}\colon k] = (\mathcal{G}\colon \mathcal{H})$.

Moreover, we have $[\Omega\colon \Omega_{(*)}] = [L^{(*)}\colon L_{(*)}] = [k^{(*)}\colon k] =$ *a p-power (the degree of inseparability of* Ω/k*).*

Proof. We know $\mathrm{Fix}(\mathcal{H}) = L^{(*)}$. So, (1) is just Artin's theorem (Theorem 4.32).

Claim: The map $\sigma \mapsto \sigma \restriction \Omega_{(*)}$ is an isomorphism $\mathcal{G} \xrightarrow{\sim} \mathcal{G}(\Omega_{(*)}/k)$.

We know $\Omega_{(*)}$ is normal over k, so $\sigma \restriction \Omega_{(*)}$ takes $\Omega_{(*)}$ to itself. Therefore, the map $\mathcal{G} \longrightarrow \mathcal{G}(\Omega_{(*)}/k)$ given by $\sigma \mapsto \sigma \restriction \Omega_{(*)}$ is well defined. If $\sigma \mapsto \mathrm{id} \in \mathcal{G}(\Omega_{(*)}/k)$, then $\sigma \restriction \Omega_{(*)}$ leaves $\Omega_{(*)}$ element-wise fixed. If $\xi \in \Omega$, then $\xi^{p^r} \in \Omega_{(*)}$, for some r. Therefore, $\sigma\big(\xi^{p^r}\big) = \xi^{p^r}$. By the usual argument, we conclude that $\sigma(\xi) = \xi$. Therefore, $\sigma = \mathrm{id}$ on Ω and our map is injective. Pick $\widetilde{\sigma} \in \mathcal{G}(\Omega_{(*)}/k)$. We have the diagram

$$\begin{array}{ccc} \Omega & \xrightarrow{\sigma} & \Omega \\ \uparrow & & \uparrow \\ \Omega_{(*)} & \xrightarrow{\widetilde{\sigma}} & \Omega_{(*)} \\ \uparrow & & \uparrow \\ k & = & k. \end{array}$$

By SMA, II, our automorphism $\widetilde{\sigma}$ comes from a $\sigma\colon \Omega \to \Omega$; so, our map is onto.

We have $\mathrm{Fix}(\mathcal{G}(\Omega_{(*)}/k)) = k$ (as $k^{(*)} = k$ in $\Omega_{(*)}$). By Artin's theorem, $[\Omega_{(*)}\colon k] = \#(\mathcal{G})$. Now,

$$[\Omega\colon L_{(*)}] = [\Omega\colon L^{(*)}][L^{(*)}\colon L_{(*)}] = [\Omega\colon \Omega_{(*)}][\Omega_{(*)}\colon L_{(*)}],$$

and

$$\mathcal{H} = \mathcal{G}(\Omega/L) = \mathcal{G}(\Omega/L_{(*)}) = \mathcal{G}(\Omega/L^{(*)}) = \mathcal{G}(\Omega_{(*)}/L_{(*)}),$$

by what's just been proved. By Artin's theorem, $[\Omega\colon L^{(*)}] = \#(\mathcal{H})$, so

$$[\Omega\colon L_{(*)}] = \#(\mathcal{H})[L^{(*)}\colon L_{(*)}] = [\Omega\colon \Omega_{(*)}]\#(\mathcal{H});$$

it follows that $[L^{(*)}\colon L_{(*)}] = [\Omega\colon \Omega_{(*)}]$, for all L. As remarked above,

$$\#(\mathcal{G}) = [\Omega_{(*)}\colon k] = [\Omega_{(*)}\colon L_{(*)}][L_{(*)}\colon k] = \#(\mathcal{H})[L_{(*)}\colon k].$$

Consequently, $[L_{(*)}\colon k] = (\mathcal{G}\colon \mathcal{H})$. □

A picture of the situation is shown in Figure 4.1.

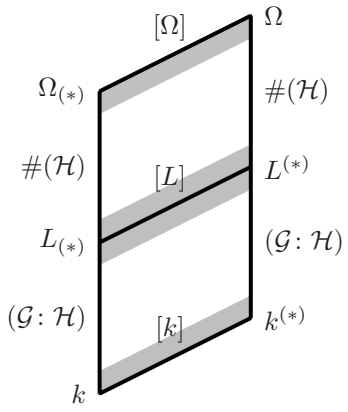

Figure 4.1: Structure of Normal Extensions

4.6 Primitive Elements, Natural Irrationalities, Normal Bases

Proposition 4.43 *If G is a finite subgroup of $K^* = \mathbb{G}_m(K)$, where K is a field, then G is cyclic.*

Proof. An abelian finite group is cyclic iff its p-Sylow subgroups are cyclic (DX). So, we may assume that $\#(G) = p^r$, for some $r > 0$ and some prime p. Let $x \in G$ be an element of maximal order, $q = p^t$, with $t \leq r$. Pick any $y \in G$; the order of y is equal to p^s for some s. But order$(y) \leq$ order(x), so $s \leq t$. As order$(y) \mid$ order(x), we must have $y^q = 1$. So, for every $y \in G$, the element y is a root of $T^q - 1$. As K is a field, this polynomial has at most q roots. But, there exist q roots in G: $1, x, \ldots, x^{q-1}$. Therefore, G is generated by x. □

Corollary 4.44 *In any field, the n-th roots of unity in the field form a cyclic group. It is a finite subgroup of $\mathbb{G}_m(K)$.*

Corollary 4.45 *The multiplicative group of a finite field is always cyclic. Every nonzero element of a finite field is a root of unity.*

Theorem 4.46 *(Artin's Theorem of the Primitive Element) Suppose K/k is a finite extension of fields, then there is some $\alpha \in K$ so that $K = k(\alpha)$ iff there are only finitely many fields, L, with $k \subseteq L \subseteq K$. (Such an α is called a primitive element).*

Proof. ($\Rightarrow$). Assume $K = k(\alpha)$. Let L be any subfield of K, write $f(X)$ for the minimal k-polynomial of α. We know that $f(X)$ is irreducible in $k[X]$. Let $g(X)$ be the minimum L-polynomial for α. As $k(\alpha) = L(\alpha)$, we have $[k(\alpha)\colon L] = [L(\alpha)\colon L] = \deg(g)$. Take L' to be the field obtained by adjoining the coefficients of g to k; we have $L' \subseteq L$. Thus, $g(X) \in L'[X]$ and $g(X)$ is irreducible over L'. Consequently, $[L'(\alpha)\colon L'] = \deg(g)$. But, $L'(\alpha) = k(\alpha)$, so

$$\deg(g) = [k(\alpha)\colon L'] = [k(\alpha)\colon L][L\colon L'] = \deg(g)[L\colon L'],$$

and we deduce that $L = L'$. This means that L is uniquely determined by g. However, every $g(X)$ is a factor of $f(X) \in K[X]$ and since there are only finitely many factors of $f(X)$, there are only finitely many subfields L.

($\Leftarrow$). Say K/k possesses just finitely many subfields.

Claim: Given $\alpha, \beta \in K$, there is some $\gamma \in K$ with $k(\alpha, \beta) \subseteq k(\gamma)$.

If the claim holds, we can finish the proof by induction on the number of generators, n, for K/k. The cases $n = 1, 2$, are clear. Assume that the induction hypothesis holds for $n - 1 \geq 1$, and let $K = k(\alpha_1, \ldots, \alpha_n) = k(\alpha_1, \ldots, \alpha_{n-2})(\alpha_{n-1}, \alpha_n)$. The claim implies that $K = k(\alpha_1, \ldots, \alpha_{n-2})(\gamma)$, and the induction hypothesis finishes the proof. So, we just have to prove the claim.

If k is finite, so is K. Consequently, $K^* = \mathbb{G}_m(K)$ is cyclic, which means that $K^* = \mathrm{Gp}\{\alpha\}$ and $K = k(\alpha)$. Thus, we may assume k is infinite. Make a map from k to the subfields of $k(\alpha, \beta)$ *via*

$$\lambda \mapsto k(\alpha + \lambda\beta)\ (\subseteq k(\alpha, \beta)).$$

Since k is infinite and since there are only finitely many subfields, there is a pair $(\lambda, \widetilde{\lambda})$, with $\lambda \neq \widetilde{\lambda}$, and

$$k(\alpha + \lambda\beta) = k(\alpha + \widetilde{\lambda}\beta) = L.$$

Thus, both $\alpha + \lambda\beta, \alpha + \widetilde{\lambda}\beta \in L$, so $(\lambda - \widetilde{\lambda})\beta \in L$. But $\lambda - \widetilde{\lambda} \neq 0$, so $\beta \in L$, and then, $\alpha \in L$. It follows that $k(\alpha, \beta) \subseteq L = k(\alpha + \lambda\beta)$, and $\gamma = \alpha + \lambda\beta$ does the job. □

Corollary 4.47 *(Kronecker's Theorem of the Primitive Element) Suppose K/k is a finite separable field extension, then there is some $\alpha \in K$ so that $K = k(\alpha)$.*

Proof. If Ω is the normal closure of K, then it is normal and separable over k. By the main theorem of Galois theory, there is a one-to-one correspondence between subfields of Ω/k and subgroups of $\mathcal{G}(\Omega/k)$. As $\mathcal{G}(\Omega/k)$ is finite, there are only finitely many subfields of Ω/k. But, any subfield of K/k is a subfield of Ω/k, which means that there are only finitely many subfields of K/k. Then, Theorem 4.46 (Artin) implies that α exists. □

Corollary 4.48 *Say K/k is a finite degree field extension and Ω is some field with $k \subseteq \Omega$. Then, the number of k-monomorphisms $K \longrightarrow \Omega$ is at most $[K\colon k]_s$. If $\widetilde{K}$ is a field k-isomorphic to K and $\widetilde{K} \underset{\text{rel}}{\sim} \Omega$, then the number of k-monomorphisms $K \longrightarrow \Omega$ is equal to $[K\colon k]_s$ iff Ω contains the normal closure of $\widetilde{K}$. In particular, if $K \subseteq \Omega$, then the number of k-monomorphisms $K \longrightarrow \Omega$ is equal to $[K\colon k]_s$ iff Ω contains the normal closure of K.*

Proof. Look at $K_{(*)}$, then we know that $[K_{(*)}\colon k] = [K\colon k]_s$. By Kronecker's theorem of the primitive element, there is some $\alpha \in K_{(*)}$ so that $K_{(*)} = k(\alpha)$. To give a k-monomorphism $K \longrightarrow \Omega$ implies that we have a k-monomorphism $K_{(*)} \longrightarrow \Omega$ and the latter is determined by its value on α. Furthermore, two k-monomorphisms of K to Ω which agree on $K_{(*)}$ necessarily agree on K. Hence, the choice of an image of α in Ω determines a k-monomorphism of $K \longrightarrow \Omega$. The image of α, say β, satisfies the minimal k-polynomial, $g(X)$, for α. Consequently, the number of k-monomorphisms $K \longrightarrow \Omega$ is at most equal to the number of roots of $g(X)$ in Ω, which is at most $\deg(g) = [K\colon k]_s$.

Take $\widetilde{K}$ with $\widetilde{K} \underset{\text{rel}}{\sim} \Omega$ and say $\widetilde{K}$ is k-isomorphic to K. Since $\widetilde{K} \cong K$, we are reduced to the case $K = \widetilde{K}$, i.e., $\Omega \underset{\text{rel}}{\sim} K$. We obtain the maximum number of k-monomorphisms iff Ω contains all the roots of any irreducible k-polynomial one root of which lies in k. For then all the conjugates of α are there and their p^rth roots for suitable r. □

Theorem 4.49 *(Natural Irrationalities) Say Ω/k is finite normal and $\widetilde{k} \supseteq k$ is some field with $\widetilde{k} \underset{\text{rel}}{\sim} \Omega$. Write $\widetilde{\Omega}$ for the compositum of $\widetilde{k}$ and Ω, denoted $\Omega\widetilde{k}$ (the smallest field containing Ω and $\widetilde{k}$). Then,*

(1) $\widetilde{\Omega}/\widetilde{k}$ is a normal extension (finite degree).

(2) The map $\sigma \mapsto \sigma \restriction \Omega$ gives a canonical injection $\mathcal{G}(\widetilde{\Omega}/\widetilde{k}) \hookrightarrow \mathcal{G}(\Omega/k)$. The image of this injection is $\mathcal{G}(\Omega/D)$, where $D = \Omega \cap \widetilde{k}$.

Proof. (1) We know $\Omega = k(\alpha_1, \ldots, \alpha_t)$, where $\alpha_1, \ldots, \alpha_t$ are all the roots of a k-polynomial, f. Now, $\widetilde{\Omega} = \widetilde{k}(\alpha_1, \ldots, \alpha_t) =$ a splitting field of the same f, but now viewed as a $\widetilde{k}$-polynomial. So (1) holds.

(2) Given $\sigma \in \mathcal{G}(\widetilde{G}/\widetilde{k})$, look at $\sigma \restriction \Omega$. We know $\sigma(\Omega)$ is a k-conjugate to Ω (inside $\widetilde{\Omega}$). As Ω is normal, $\sigma(\Omega) = \Omega$, and so, $\sigma \restriction \Omega$ is an automomorphism of Ω. As σ fixes $\widetilde{k}$, it fixes $k \subseteq \widetilde{k}$. Thus, $\sigma \restriction \Omega \in \mathcal{G}(\Omega/k)$. If $\sigma \restriction \Omega$ were the identity, we would have $\sigma(\alpha_j) = \alpha_j$, for all j. Also, $\sigma \restriction \widetilde{k} = \mathrm{id}$ and thus, σ fixes all of $\widetilde{k}(\alpha_1, \ldots, \alpha_t) = \widetilde{\Omega}$. Therefore, $\sigma = \mathrm{id}$ in $\mathcal{G}(\widetilde{G}/\widetilde{k})$, i.e., our map is injective.

Let $D = \Omega \cap \widetilde{k}$ and let $\mathcal{H}$ be the image of $\mathcal{G}(\widetilde{\Omega}/\widetilde{k})$ in $\mathcal{G}(\Omega/k)$. We have $\mathcal{H} \cong \mathcal{G}(\widetilde{\Omega}/\widetilde{k})$. As $D \subseteq \widetilde{k}$, we see that $\mathcal{H}$ fixes D, so $\mathcal{H} \subseteq \mathcal{G}(\Omega/D)$. Let $L = \mathrm{Fix}(\mathcal{H})$. We know that $L = L^{(*)}$. As D is fixed, $D \subseteq L = L^{(*)} \subseteq \Omega$. Now, all elements of $\mathcal{H}$ come from $\mathcal{G}(\widetilde{\Omega}/\widetilde{k})$, which implies that $\mathrm{Fix}(\mathcal{H}) \subseteq \mathrm{Fix}(\mathcal{G}(\widetilde{\Omega}/\widetilde{k})) = \widetilde{k}^{(*)}$, by Corollary 4.34. So, $D \subseteq L = L^{(*)} \subseteq \widetilde{k}^{(*)}$ and $D \subseteq L = L^{(*)} \subseteq \Omega$. Pick $\xi \in L$. Then, $\xi \in \widetilde{k}^{(*)}$, so $\xi^{p^r} \in \widetilde{k}$, for some r. But, $\xi \in L \subseteq \Omega$, so $\xi^{p^r} \in \Omega$, and thus, $\xi^{p^r} \in \widetilde{k} \cap \Omega = D$. It follows that $L \subseteq D^{(*)}$. As $L = L^{(*)}$, we have $L^{(*)} \subseteq D^{(*)}$. Yet, $D \subseteq L$, so $D^{(*)} \subseteq L^{(*)}$ and therefore $L^{(*)} = D^{(*)}$. It follows that

$$\mathcal{G}(\Omega/D) = \mathcal{G}(\Omega/L) = \mathcal{G}(\Omega/\mathrm{Fix}(\mathcal{H})) = \mathcal{H},$$

by the fundamental theorem of Galois theory. □

Corollary 4.50 *(Original Form of Natural Irrationalities) Say f is a k-polynomial and $k \subseteq \widetilde{k}$. Then, $\mathcal{G}_{\widetilde{k}}(f)$ is a subgroup of $\mathcal{G}_k(f)$ in a natural way and in fact, $\mathcal{G}_{\widetilde{k}}(f) = \mathcal{G}_D(f)$, where $D = \Omega \cap \widetilde{k}$ and $\Omega \underset{\text{rel}}{\sim} \widetilde{k}$ is a splitting field of f.*

Explanation: Let Ω be a given splitting field of f. The elements of Ω were termed the *natural irrationalities* of f. The reduction in $\mathcal{G}_k(f)$ effected by considering f over $\widetilde{k}$ is the same as that achieved by considering f over the field of those natural irrationalities of f contained in $\widetilde{k}$.

Theorem 4.51 *(Normal Basis Theorem) Suppose K/k is a finite normal and separable extension and let $\mathcal{G}(K/k)$ be its Galois group. Then, there is some $\theta \in K$ so that $\{\sigma\theta \mid \sigma \in \mathcal{G}(K/k)\}$ is a k-basis for K. (This is called a normal basis for K/k).*

Proof. By Kronecker's theorem, $K = k(\alpha)$, for some $\alpha \in K$; let $f(X)$ be the minimum k-polynomial for α. We know $K = k[X]/(f(X))$. Examine two rings: $K[X]$ and $A = K[X]/(f(X))$. Note,

$$K \otimes_k K = K \otimes_k (k[X]/(f(X))) \cong K[X]/(f(X)) = A.$$

For $\sigma \in \mathcal{G} = \mathcal{G}(K/k)$, write α_σ for $\sigma(\alpha)$. Consider the K-polynomials

$$g_\sigma(X) = \frac{f(X)}{f'(\alpha_\sigma)(X - \alpha_\sigma)}.$$

Note that $g_1(X) = f(X)/(f'(\alpha)(X - \alpha))$, so $\sigma g_1(X) = g_\sigma(X)$. The g_σ's satisfy the following properties:

(1) Each $g_\sigma(X)$ has degree $\deg(f) - 1$.

(2) If $\sigma \neq \tau$, then $g_\sigma(\alpha_\tau) = 0$.
Also, by Taylor's theorem,

$$f(X) = f(\alpha_\sigma + (X - \alpha_\sigma)) = f(\alpha_\sigma) + f'(\alpha_\sigma)(X - \alpha_\sigma) + O((X - \alpha_\sigma)^2),$$

so, $g_\sigma(X) = 1 + O(X - \alpha_\sigma)$ and therefore,

(3) $g_\sigma(\alpha_\sigma) = 1$.

Consider the polynomial $\sum_{\sigma \in \mathcal{G}} g_\sigma(X) - 1 \, (\in K[X])$. By (2) and (3), we see that this polynomial vanishes on the n elements $\alpha, \alpha_{\sigma_2}, \ldots, \alpha_{\sigma_n}$, where $\mathcal{G} = \{1, \sigma_2, \ldots, \sigma_n\}$. By (1), this polynomial has degree $n - 1$. Hence, the polynomial is identically zero and we have

$$\sum_{\sigma \in \mathcal{G}} g_\sigma(X) = 1. \qquad \text{(partition of unity)} \quad (*)$$

In A, we get

$$\sum_{\sigma \in \mathcal{G}} \overline{g_\sigma(X)} = 1. \qquad (\overline{*})$$

Pick σ, τ, with $\sigma \neq \tau$, and look at $g_\sigma(X)g_\tau(X)$. For all $\rho \in \mathcal{G}$, we have $g_\sigma(\alpha_\rho)g_\tau(\alpha_\rho) = 0$. But, $f(X) = \prod_{\rho \in \mathcal{G}}(X - \alpha_\rho)$, so $f(X) \mid g_\sigma(X)g_\tau(X)$ if $\sigma \neq \tau$. If we read this in A, we get

$$\overline{g_\sigma(X)}\,\overline{g_\tau(X)} = 0 \quad \text{in } A, \text{ if } \sigma \neq \tau. \qquad \text{(orthogonality)} \quad (**)$$

If we multiply $(*)$ by $g_\tau(X)$, we get

$$\sum_{\sigma \in \mathcal{G}} g_\tau(X)g_\sigma(X) = g_\tau(X),$$

and if we read this in A and use $(**)$, we get

$$(\overline{g_\sigma(X)})^2 = \overline{g_\sigma(X)} \quad \text{in } A. \qquad \text{(idempotence)} \quad (***)$$

Write $e_\sigma = \overline{g_\sigma(X)}$, so $e_\sigma \in A = K \otimes_k K$. Then, $(*)$, $(**)$ and $(***)$ say:

$$\sum_{\sigma\in\mathcal{G}} e_\sigma = 1; \quad e_\sigma e_\tau = \delta_{\sigma\tau} e_\sigma.$$

Therefore, the e_σ's are an orthogonal decomposition of 1 by idempotents, and so,

$$K \otimes_k K \cong \prod_{\sigma\in\mathcal{G}} K e_\sigma \cong \prod_{\sigma\in\mathcal{G}} K.^3$$

Order the elements of $\mathcal{G}$ in some fashion as we did above: $1, \sigma_2, \ldots, \sigma_n$, and consider the matrix

$$(g_{\sigma\tau}(X)) \in \mathrm{M}_n(K[X]).$$

Let $D(X) = \det(g_{\sigma\tau}(X))$. In order to compute $D(X)$ in A, consider $D(X)^2$. Since $\det(g_{\sigma\tau}(X)) = \det(g_{\sigma\tau}(X))^\top$, we can compute $D(X)^2$ by multiplying columns by columns and summing. We get

$$\sum_{\sigma\in\mathcal{G}} g_{\sigma\tau}(X) g_{\sigma\rho}(X) = \sum_{\sigma\in\mathcal{G}} \sigma(g_\tau(X))\sigma(g_\rho(X)) = \sum_{\sigma\in\mathcal{G}} \sigma(g_\tau(X) g_\rho(X)).$$

If we read this in A, we get

$$\begin{aligned}
\sum_{\sigma\in\mathcal{G}} \overline{g_{\sigma\tau}(X)}\,\overline{g_{\sigma\rho}(X)} &= \sum_{\sigma\in\mathcal{G}} \sigma(\overline{g_\tau(X)}\,\overline{g_\rho(X)}) = 0, \quad \text{if} \quad \tau \neq \rho; \quad \text{and} \\
&= \sum_{\sigma\in\mathcal{G}} \sigma(\overline{g_\rho(X)}), \quad \text{if} \quad \tau = \rho \\
&= \sum_{\sigma\in\mathcal{G}} \overline{g_{\sigma\rho}(X)} \\
&= \sum_{\tau\in\mathcal{G}} \overline{g_\tau(X)} = 1, \quad \text{if} \quad \tau = \rho.
\end{aligned}$$

Therefore, we find that in A, the matrix $(g_{\sigma\tau}(X))(g_{\sigma\tau}(X))^\top$ is the identity matrix and so, $\overline{D(X)^2} = 1$. Consequently, $D(X)^2 \equiv 1 \pmod{f(X)}$, which shows that $D(X) \neq 0$.

If k is infinite, then there is some $\xi \in k$ with $D(\xi) \neq 0$. Let $\theta = g_1(\xi)$. Then, $\sigma\tau\theta = \sigma\tau g_1(\xi) = g_{\sigma\tau}(\xi)$. Consequently, $\det(\sigma\tau(\theta)) = \det(g_{\sigma\tau}(\xi)) = D(\xi) \neq 0$. If $\{\sigma\theta\}_{\sigma\in\mathcal{G}}$ were linearly dependent, we would have

$$\sum_{\tau\in\mathcal{G}} a_\tau \tau\theta = 0,$$

for some $a_\tau \in k$, not all zero. If we apply σ, we get

$$\sum_{\tau\in\mathcal{G}} a_\tau \sigma\tau\theta = 0.$$

So, (a_τ) would be a nontrivial simultaneous solution to the linear system of equations

$$\sum_{\tau\in\mathcal{G}} X_\tau \sigma\tau\theta = 0, \quad \text{for} \quad \sigma \in \mathcal{G},$$

a contradiction to the fact that $\det(\sigma\tau(\theta)) \neq 0$. Therefore, $\{\sigma\theta\}_{\sigma\in\mathcal{G}}$ is linearly independent and the case where k is infinite is proved.

If k is finite, we don't need the $g_\sigma(X)$ and $D(X)$. We do need the following facts to be proved below:

[3] At this stage, we are essentially done. However, we've not kept track of the $\mathcal{G}$ action; so, a little more argument is needed.

(1) The Galois group $\mathcal{G}(K/k)$ is cyclic.

(2) The Galois group $\mathcal{G}(K/k)$ has a canonical generator, $\mathbf{F}$, where $\mathbf{F}(\xi) = \xi^{\#(k)}$, for all $\xi \in K$.

Recall that for a linear transformation, T, on a finite dimensional vector space, V, if $m(X)$ is the minimal polynomial for T then there exists a vector, $v \in V$, so that $m(T)v = 0$ but *no polynomial of smaller degree than $m(X)$ kills v*. Now, our K plays the role of V and the automorphism $\mathbf{F}$ plays the role of T. If we can show that the minimum polynomial of $\mathbf{F}$ is exactly $X^n - 1$, where $n = [K\colon k]$, then we take a Θ in K so that no polynomial of smaller degree than $\mathbf{F}^n - 1$ kills Θ. This means that

$$\Theta, \mathbf{F}(\Theta), \cdots, \mathbf{F}^{n-1}(\Theta)$$

are linearly independent; so by (1) and (2) we have our normal basis.

Of course, by (1) and (2), $\mathbf{F}^{n-1} - 1 \equiv 0$ on K; therefore, whatever is the minimal polynomial for $\mathbf{F}$, it divides $X^n - 1$ and its degree is at most n. Were $m(X) = a_0X^d + a_1X^{d-1} + \cdots + a_d$ the minimal polynomial for $\mathbf{F}$ and $d < n$, then

$$0 = a_0\mathbf{F}^d(\xi) + a_1\mathbf{F}^{d-1}(\xi) + \cdots + a_{d-1}\mathbf{F}(\xi) + a_0\mathbf{F}^0(\xi) \qquad (\dagger)$$

for all $\xi \in K$. But this is a contradiction of Dedekind's Theorem as (†) is a linear dependence among I, $\mathbf{F}, \cdots, \mathbf{F}^d$, and we are done. □

Remark: The argumeent actually proves (independently of previous arguments) that *every cyclic extension possesses a normal basis*.

The facts concerning finite fields were proved by E.H. Moore. Here is his theorem:

Theorem 4.52 *(E.H. Moore, 1892) If k is a finite field then* $\mathrm{char}(k) = p > 0$ *and $\#(k) = p^l$, for some prime p and some $l \geq 1$. If $\mathbb{F}_p$ is the prime field of characteristic p, then for each integer $l \geq 1$, there exists one and only one finite field of cardinality p^l, up to $\mathbb{F}_p$-isomorphism. If K/k is a finite extension of degree n and k is a finite field, then K/k is always normal and separable; the Galois group $\mathcal{G}(K/k)$ is cyclic of order n and has a canonical generator, $\mathbf{F}$. This $\mathbf{F}$ is the Frobenius automorphism, and it is given by $\xi \mapsto \mathbf{F}(\xi) = \xi^{\#(k)}$, for all $\xi \in K$. Each finite field has exactly one extension of degree n for each $n \geq 1$.*

Proof. The statement in the first sentence is well-known. Pick $l \geq 1$ and look at the splitting field of the polynomial $X^{p^l} - X \in \mathbb{F}_p[X]$. Note, if ξ and η are roots of this polynomial, then $\xi \pm \eta$, $\xi\eta$, ξ/η are also roots of the polynomial. Thus, the set of roots is a field and it contains $\mathbb{F}_p$, because for all $\xi \in \mathbb{F}_p$, we have $\xi^p = \xi$. It follows that the splitting field is exactly the entire set of roots and as the derivative of $X^{p^l} - X$ is -1, the roots are distinct. Therefore, we get a field with p^l elements. Conversely, any field with p^l elements has multiplicative group of order $p^l - 1$. So, this group has a generator of order $p^l - 1$ and for this generator, θ, we get $\theta^{p^l} = \theta$. Consequently, any power of θ satisfies $X^{p^l} - X = 0$ and so, our field is a splitting field of $X^{p^l} - X$; such fields are unique up to $\mathbb{F}_p$-isomorphism.

Suppose K/k has degree n, then K is a splitting field, so K/k is normal. Moreover, finite fields are perfect, so K/k is separable.

Consider $\mathbf{F}_k \in \mathcal{G}(K/k)$ where $\mathbf{F} = \mathbf{F}_k$ is defined by $\mathbf{F}(\xi) = \xi^{\#(k)}$. Look at $1 = \mathbf{F}^0, \mathbf{F}^1, \mathbf{F}^2, \ldots, \mathbf{F}^{n-1}$. These are distinct, as $\mathbf{F}^r(\theta) = \mathbf{F}^s(\theta)$ implies $\mathbf{F}^{r-s}(\theta) = \theta$; that is, $\theta^{q^{r-s}-1} = 1$. Yet, $q^{r-s} < \#(K)$, a contradiction. Now, $\mathbf{F}^n(\xi) = \xi^{q^n}$. It follows from linear algebra that $q^n = \#(K)$ and by the above, $\xi^{q^n} = \xi$ implies $\mathbf{F}^n = 1$. Observe, $\mathbf{F}(\xi) = \xi$ when $\xi \in k$, which implies that $\mathbf{F}$ is a k-automorphism and $\mathbf{F}$ has the proper order. □

Interpretations of the Normal Basis Theorem

(1) **Algebraic Interpretation**

Assume K/k is normal and separable, let $\mathcal{G} = \mathcal{G}(K/k)$ with $\#(\mathcal{G}) = n$. We claim that there is a natural ring homomorphism

$$K \otimes_k K \longrightarrow \prod_{\sigma \in \mathcal{G}} K.$$

(Here $\prod_{\sigma\in\mathcal{G}} K$ consists of n factors of K under coordinatewise multiplication.) Take $\alpha, \beta \in K$, and send (α, β) to the n-tuple

$$\langle \alpha\beta, \alpha\sigma_2\beta, \ldots, \alpha\sigma_n\beta \rangle,$$

where $\mathcal{G} = \{\sigma_1 = 1, \sigma_2, \ldots, \sigma_n\}$. This is a bilinear map, so we get a map

$$K \otimes_k K \longrightarrow \prod_{\sigma \in \mathcal{G}} K.$$

On the left hand side, we have a K-vector space *via* $\alpha \in K$ acts as $\alpha \otimes 1$. The righthand side is a K-vector space *via* the action of K on each factor; thus, the above map is K-linear. We also have

$$\begin{aligned} (\alpha \otimes \beta)(\gamma \otimes \delta) &= (\alpha\gamma \otimes \beta\delta) \\ (\alpha\sigma\beta)_\sigma(\gamma\sigma\delta)_\sigma &= (\alpha\gamma\sigma(\beta\delta))_\sigma. \end{aligned}$$

The normal basis theorem says that this ring map is an isomorphism. Say θ is our normal basis element, then

$$1 \otimes \theta, 1 \otimes \sigma_2\theta, \ldots, 1 \otimes \sigma_n\theta$$

is a basis for $K \otimes_k K$ over K. Now, as

$$1 \otimes \tau\theta \mapsto \langle \sigma\tau\theta \rangle_{\sigma\in\mathcal{G}},$$

a basis on the left hand side goes to a basis on the right hand side; so, the map is an isomorphism. Check the converse.

(2) **Geometric Interpretation**

Say X is a space; G is a group, and suppose G acts on X: There is a map $G \times X \longrightarrow X$ denoted $(\sigma, x) \mapsto \sigma x$.

Definition 4.13 A space X is a *principal homogeneous space* for G (*PHS* for G) if

(1) X is a *homogeneous space*, i.e., for all $x, y \in G$, there is some $\sigma \in G$ with $\sigma x = y$ (G acts transitively), i.e., X is equal to an orbit of G under the action.

(2) The group element $\sigma \in G$ in (1) is *uniquely* determined by x and y.

Proposition 4.53 *The following statements are equivalent:*

(A) X is a PHS for G.

(B) The map $G \prod X \longrightarrow X \prod X$ via $(\sigma, x) \mapsto (\sigma x, x)$ is an isomorphism.

Proof. $(A) \Rightarrow (B)$. Given $(\xi, \eta) \in X \prod X$, there is a $\sigma \in G$ with $\sigma\xi = \eta$. Thus, $(\sigma, \xi) \mapsto (\eta, \xi)$ under our map, which shows its surjectivity. The map is injective by property (2) of the definition.

$(B) \Rightarrow (A)$. This is a tautology. □

Let G be a group and let k be a field. Write $A(G)$ for the k-algebra of all functions $f\colon G \to k$ under pointwise operations (e.g., $(fg)(\sigma) = f(\sigma)g(\sigma)$, etc.). The k-algebra $A(G)$ has a basis,
$\{e_\sigma\}$, where $e_\sigma(\tau) = \delta_{\sigma\tau}$.

Suppose now G is a finite group, then there is a k-algebra map $\Delta\colon A(G) \to A(G) \otimes_k A(G)$ given by (convolution)

$$\Delta(e_\tau) = \sum_{\sigma \in G} e_\sigma \otimes e_{\sigma^{-1}\tau}.$$

I claim: For all k-algebras, R,

$$A(G)(R) = \mathrm{Hom}_k(A(G), R)$$

is a group. Given $\varphi, \psi \in A(G)(R)$, we define $\varphi\psi$ as the composition

$$A(G) \xrightarrow{\Delta} A(G) \otimes_k A(G) \xrightarrow{\varphi\otimes\psi} R \otimes_k R \xrightarrow{\text{mult}} R.$$

Let us see what $(\varphi\psi)(e_\rho)$ is. We have

$$\Delta(e_\rho) = \sum_{\sigma \in \mathcal{G}} e_\sigma \otimes e_{\sigma^{-1}\rho} \quad \text{and} \quad (\varphi \otimes \psi)(\Delta(e_\rho)) = \sum_{\sigma \in \mathcal{G}} \varphi(e_\sigma) \otimes \psi(e_{\sigma^{-1}\rho}),$$

so

$$(\varphi\psi)(e_\rho) = \sum_{\sigma \in \mathcal{G}} \varphi(e_\sigma)\psi(e_{\sigma^{-1}\rho}).$$

(Note: We can form $k[G]$ = the group algebra and the reader should check that:

(1) As linear spaces, $A(G)$ and $k[G]$ are naturally dual.

(2) Multiplication in $A(G)$ goes over to Δ for $k[G]$ and Δ for $A(G)$ goes over to ordinary multiplication in $k[G]$.)

The space $\mathrm{Spec}\, A(G) = \underline{G}$ is a geometric object (at least it's a topological space). Indeed, it is described by the equations $X_\sigma X_\tau = \delta_{\sigma\tau} X_\sigma$ and $\sum_{\sigma \in G} X_\sigma = 1$ (the e_σ have been replaced by the X_σ for convenience of more usual notation). To find solutions in a ring R is to give a homomorphism $A(G) \longrightarrow R$, as above. If $\mathrm{Spec}\, R$ is connected (i.e., $e^2 = e$ implies $e = 0$ or $e = 1$) then solutions correspond just to the set G and we recover the multiplication in G from our funny multiplication using Δ.

We know that

$$\mathrm{Spec}(B \otimes_A C) = \mathrm{Spec}\, B \prod_{\mathrm{Spec}\, A} \mathrm{Spec}\, C.$$

The meaning of this is exactly that

$$\mathrm{Hom}_{A\text{-alg}}(B \otimes_A C, R) = \mathrm{Hom}_{A\text{-alg}}(B, R) \prod \mathrm{Hom}_{A\text{-alg}}(C, R),$$

where on the right we have the ordinary cartesian product of sets.

Look at $A(G) \otimes_k K$, where $G = \mathcal{G}(K/k)$. Remember, $A(G)$ has the e_σ's and $K \otimes_k K$ has the $g_\sigma(X) = e_\sigma$'s, too. So, there is an isomorphism of rings

$$A(G) \otimes_k K \cong K \otimes_k K.$$

Upon taking Spec's we see that

$$\underline{G} \prod \mathrm{Spec}\, K \cong \mathrm{Spec}\, K \prod \mathrm{Spec}\, K.$$

Therefore, the fact $\mathrm{Spec}\, K$ *is a PHS for* $\underline{G}$ *is exactly the normal basis theorem.*

4.7 Galois Cohomology, Norms and Traces

Recall that in Chapter 1, Section 1.4, we introduced the notion of cohomology of a group, G, with coefficients in a G-module, M. I urge you to review the appropriate parts of Section 1.4 now.

If $0 \longrightarrow M' \longrightarrow M \longrightarrow M'' \longrightarrow 0$ is an exact sequence of G-modules, then, for each $r \geq 0$, the sequence

$$0 \longrightarrow C^r(G, M') \longrightarrow C^r(G, M) \longrightarrow C^r(G, M'') \longrightarrow 0$$

is again exact and a commutative diagram of G-modules

$$\begin{array}{ccccccccc} 0 & \longrightarrow & M' & \longrightarrow & M & \longrightarrow & M'' & \longrightarrow & 0 \\ & & \downarrow & & \downarrow & & \downarrow & & \\ 0 & \longrightarrow & N' & \longrightarrow & N & \longrightarrow & N'' & \longrightarrow & 0 \end{array}$$

yields a similar commmutative diagram

$$\begin{array}{ccccccccc} 0 & \longrightarrow & C^r(G, M') & \longrightarrow & C^r(G, M) & \longrightarrow & C^r(G, M'') & \longrightarrow & 0 \\ & & \downarrow & & \downarrow & & \downarrow & & \\ 0 & \longrightarrow & C^r(G, N) & \longrightarrow & C^r(G, N) & \longrightarrow & C^r(G, N'') & \longrightarrow & 0 \end{array}$$

for all $r \geq 0$. We'll see in the next chapter (Chapter 5, Lemma 5.7 and Corollary 5.8) that these statements imply the following facts:

Fact I. If $0 \longrightarrow M' \longrightarrow M \longrightarrow M'' \longrightarrow 0$ *is an exact sequence of* G*-modules, then we have the long exact sequence of cohomology*

$$\begin{array}{l} 0 \longrightarrow H^0(G, M') \longrightarrow H^0(G, M) \longrightarrow H^0(G, M'') \xrightarrow{\delta^{(0)}} \\ \quad H^1(G, M') \longrightarrow H^1(G, M) \longrightarrow H^1(G, M'') \xrightarrow{\delta^{(1)}} \\ \quad H^2(G, M') \longrightarrow \cdots \qquad\qquad \cdots \xrightarrow{\delta^{(r-1)}} \\ \quad H^r(G, M') \longrightarrow H^r(G, M) \longrightarrow H^r(G, M'') \xrightarrow{\delta^{(r)}} \\ \quad H^{r+1}(G, M') \longrightarrow \cdots \end{array}$$

(The maps $\delta^{(r)}$ are the *connecting homomorphisms* of the long exact sequence.)

Fact II. A small commutative diagram of G*-modules*

$$\begin{array}{ccccccccc} 0 & \longrightarrow & M' & \longrightarrow & M & \longrightarrow & M'' & \longrightarrow & 0 \\ & & \downarrow & & \downarrow & & \downarrow & & \\ 0 & \longrightarrow & N' & \longrightarrow & N & \longrightarrow & N'' & \longrightarrow & 0 \end{array}$$

yields a large (long) commutative diagram of cohomology:

$$\begin{array}{ccccccccccc}
0 & \longrightarrow & H^0(G,M') & \longrightarrow & \cdots & \longrightarrow & H^r(G,M) & \longrightarrow & H^r(G,M'') & \longrightarrow & H^{r+1}(G,M') & \longrightarrow & \cdots \\
 & & \downarrow & & & & \downarrow & & \downarrow & & \downarrow & & \\
0 & \longrightarrow & H^0(G,N') & \longrightarrow & \cdots & \longrightarrow & H^r(G,N) & \longrightarrow & H^r(G,N'') & \longrightarrow & H^{r+1}(G,N') & \longrightarrow & \cdots
\end{array}$$

The proofs of these facts do not use any of the material below, so we will assume them now without circularity in our reasoning.

Suppose B is an abelian group. We can make, from B, a G-module, $\mathrm{Map}(G,B)$, as follows:

$$\mathrm{Map}(G,B) = \{f \mid f\colon G \to B, \quad \text{i.e., } f \text{ is a function from } G \text{ to } B\}.$$

The module structure is

$$(\sigma f)(\tau) = f(\tau\sigma)$$

and one checks that if B is actually a G-module, there is a G-module *injection*

$$\epsilon_B\colon B \to \mathrm{Map}(G,B)$$

given by

$$\epsilon_B(b)(\sigma) = \sigma b. \quad \text{(DX)}$$

The module $\mathrm{Map}(G,B)$ is special in that it is "cohomologically trivial." This is

Proposition 4.54 *For every abelian group, B and every $n > 0$, we have*

$$H^n(G, \mathrm{Map}(G,B)) = (0).$$

Proof. Choose $f \in Z^n(G, \mathrm{Map}(G,B))$ and assume $n > 0$. Then f is a function of n variables chosen from G and has values in $\mathrm{Map}(G,B)$. We define a function, g, of $n-1$ variables chosen from G with values in $\mathrm{Map}(G,B)$ as follows:

$$g(\sigma_1,\ldots,\sigma_{n-1})(\tau) = f(\tau,\sigma_1,\ldots,\sigma_{n-1})(1).$$

Let us prove that $\delta g = f$, which will finish the argument.

$$(\delta g)(\sigma_1,\ldots,\sigma_n) = \sigma_1 g(\sigma_2,\ldots,\sigma_n) + \sum_{r=1}^{n-1}(-1)^r g(\sigma_1,\ldots,\sigma_r\sigma_{r+1},\ldots,\sigma_n) + (-1)^n g(\sigma_1,\ldots,\sigma_{n-1}).$$

So, upon evaluating δg on an arbitrary element, τ, we get

$$\begin{aligned}(\delta g)(\sigma_1,\ldots,\sigma_n)(\tau) &= g(\sigma_2,\ldots,\sigma_n)(\tau\sigma_1) + \sum_{r=1}^{n-1}(-1)^r g(\sigma_1,\ldots,\sigma_r\sigma_{r+1},\ldots,\sigma_n)(\tau) + (-1)^n g(\sigma_1,\ldots,\sigma_{n-1})(\tau) \\ &= f(\tau\sigma_1,\sigma_2,\ldots,\sigma_n)(1) + \sum_{r=1}^{n-1}(-1)^r f(\tau,\sigma_1,\ldots,\sigma_r\sigma_{r+1},\ldots,\sigma_n)(1) + (-1)^n f(\tau,\sigma_1,\ldots,\sigma_{n-1})(1).\end{aligned}$$

Now, $f(\sigma_1,\ldots,\sigma_n)(\tau) = (\tau f)(\sigma_1,\ldots,\sigma_n)(1)$, and

$$\begin{aligned}0 = \delta f(\tau,\sigma_1,\ldots,\sigma_n) = (\tau f)(\sigma_1,\ldots,\sigma_n) - f(\tau\sigma_1,\sigma_2,\ldots,\sigma_n) + \sum_{s=2}^{n}(-1)^s f(\tau,\sigma_1,\ldots,\sigma_{s-1}\sigma_s,\ldots,\sigma_n) \\ +(-1)^{n+1} f(\tau,\sigma_1,\ldots,\sigma_{n-1}).\end{aligned}$$

Therefore,

$$(\tau f)(\sigma_1,\ldots,\sigma_n) = f(\tau\sigma_1,\sigma_2,\ldots,\sigma_n) + \sum_{s=2}^{n}(-1)^{s-1}f(\tau,\sigma_1,\ldots,\sigma_{s-1}\sigma_s,\ldots,\sigma_n) + (-1)^n f(\tau,\sigma_1,\ldots,\sigma_{n-1}).$$

Let $n = s-1$ in the sum above and evaluate both sides at 1. We get immediately

$$f(\sigma_1,\ldots,\sigma_n)(\tau) = \delta g(\sigma_1,\ldots,\sigma_n)(\tau). \qquad \square$$

Proposition 4.54 is extremely useful and very powerful. Rather than explain this in abstract terms, let's begin to use Proposition 4.54 and, in so doing, show *how* to use it and *why* it is powerful. One of the facts left unproved in Chapter 1 was the fact that $H^r(G,M)$ is $\#(G)$-torsion if $r>0$ (any module, M). Based on Proposition 4.54, we can now prove this and, while our proof is not the most elegant known, it certainly requires the least machinery:

Proposition 4.55 *If G is a finite group and M is any G-module, then $H^r(G,M)$ is $\#(G)$-torsion if $r>0$.*

Proof. Take the case $r=1$, first. If $f\in Z^1(G,M)$, we know

$$f(\sigma\rho) = \sigma f(\rho) + f(\sigma).$$

Write α for the element $-\sum_{\rho\in G} f(\rho)$ of M. We compute $\sigma\alpha$:

$$\begin{aligned}\sigma\alpha = -\sum_{\rho\in G}\sigma f(\rho) &= -\sum_{\rho\in G}(f(\sigma\rho)-f(\sigma)) \\ &= -\sum_{\rho\in G} f(\sigma\rho) + \#(G)f(\sigma) \\ &= \alpha + \#(G)f(\sigma).\end{aligned}$$

Therefore, $(\#(G)f)(\sigma) = (\delta\alpha)(\sigma)$, and the case $r=1$ is done.

Now, use induction on r—here is where Proposition 4.54 enters. Assume as induction hypothesis that given r ($r\geq 1$), for all modules, N, we have $H^r(G,N)$ is $\#(G)$-torsion. The step from r to $r+1$ goes like this:

Choose M, embed M in $\mathrm{Map}(G,M)$, to get

$$0 \longrightarrow M \xrightarrow{\epsilon_M} \mathrm{Map}(G,M) \longrightarrow \mathrm{coker} \longrightarrow 0.$$

Apply cohomology (i.e., use the long exact sequence of Fact I), we get

$$\ldots \longrightarrow H^r(G,\mathrm{Map}(G,M)) \longrightarrow H^r(G,\mathrm{coker}) \longrightarrow H^{r+1}(G,M) \longrightarrow H^{r+1}(G,\mathrm{Map}(G,M)) \longrightarrow \ldots \qquad (*)$$

The ends of $(*)$ vanish by Proposition 4.54 and we obtain the isomorphism

$$H^r(G,\mathrm{coker}) \xrightarrow{\sim} H^{r+1}(G,M), \quad \text{for all } r\geq 1. \qquad (**)$$

But, the left side of $(**)$ is $\#(G)$-torsion by our induction hypothesis, therefore $H^{r+1}(G,M)$ is also $\#(G)$-torsion. $\square$

The special case when G is cyclic is both instructive and important for some material to follow. For arbitrary (finite) G, and any G-module, M, we define the *norm map*, $\mathcal{N}_G$, taking M to itself by

$$\mathcal{N}_G(m) = \sum_{\sigma\in G}\sigma m.$$

Note that the image of $\mathcal{N}_G$ lies in M^G. Further, $\mathcal{N}_G$ is actually a G-module map, for

$$\mathcal{N}_G(\tau m) = \sum_{\sigma\in G} \sigma\tau m = \mathcal{N}_G(m) = \tau\mathcal{N}_G(m).$$

(In cases of interest below, the map $\mathcal{N}_G$ is usually called *trace* and when M is written multiplicatively then $\mathcal{N}_G$ is called the *norm*.) Now, the equation $\mathcal{N}_G(\tau m) = \tau\mathcal{N}_G(m)$ shows that the elements $\tau m - m$ all lie in $\operatorname{Ker}\mathcal{N}_G$. The submodule generated by all $\tau m - m$, as τ runs over G and m over M, is denoted IM; so, $IM \subseteq \operatorname{Ker}\mathcal{N}_G(M)$.

Proposition 4.56 *If G is a (finite) cyclic group and σ is one of its generators, then for any module, M:*

(a) The map $f \mapsto f(\sigma) \in M$ is a G-isomorphism of $Z^1(G, M)$ with $\operatorname{Ker}\mathcal{N}_G(M)$,

(b) The submodule IM is generated by $\sigma m - m$ for this fixed σ and m varying over M,

(c) There is an isomorphism $H^1(G, M) \xrightarrow{\sim} \operatorname{Ker}\mathcal{N}_G/IM$.

Proof. The elements of G are $1, \sigma, \ldots, \sigma^{n-1}$. Let $f \in Z^1(G, M)$, so $f(\rho\tau) = \rho f(\tau) + f(\rho)$ for all ρ and τ of G. Apply this successively to the powers of σ:

$$f(\sigma^2) = f(\sigma\sigma) = \sigma f(\sigma) + f(\sigma); \quad f(\sigma^3) = f(\sigma\sigma^2) = \sigma f(\sigma^2) + f(\sigma) = \sigma^2 f(\sigma) + \sigma f(\sigma) + f(\sigma), \quad \text{etc.} \quad (*)$$

We find that

$$f(1) = f(\sigma^n) = \sigma^{n-1} f(\sigma) + \sigma^{n-2} f(\sigma) + \cdots + f(\sigma) = \mathcal{N}_G(f(\sigma)).$$

But, $f(1) = f(1 \cdot 1) = f(1) + f(1)$; so, $f(1) = 0$. Thus, when $f \in Z^1(G, M)$, we get $f(\sigma) \in \operatorname{Ker}\mathcal{N}_G(M)$.

From $(*)$ above, we see that $f(\sigma)$ determines f when f is a cocycle, conversely an easy argument using the inductive definition of $f(\sigma^i)$ given by $(*)$ (namely, $\sigma f(\sigma^{i-1}) + f(\sigma)$) shows that if $f(\sigma) \in \operatorname{Ker}\mathcal{N}_G$ our definition makes f a 1-cocycle (DX). This gives an abelian group isomorphism $Z^1(G, M) \xrightarrow{\sim} \operatorname{Ker}\mathcal{N}_G$. Since $Z^1(G, M)$ is a G-module *via* M, the map is a G-module isomorphism, and (a) is proved.

To prove (b), all we need to show is that $\tau m - m$ is in the submodule generated by $\sigma\widetilde{m} - \widetilde{m}$ as $\widetilde{m}$ ranges over M, where τ is a fixed arbitrary element of G. But, $\tau = \sigma^i$; so,

$$\tau m - m = \sigma^i m - m = \sigma^i m - \sigma^{i-1} m + \sigma^{i-1} m - m = \sigma^{i-1}(\sigma m - m) + \sigma^{i-1} m - m.$$

A clear induction finishes the argument.

(c) The group $B^1(G, M)$ consists exactly of those f for which $f(\tau) = \tau m - m$ for some $m \in M$. Hence, $f(\sigma) = \sigma m - m \in IM$ and part (a) now shows that in the isomorphism $Z^1(G, M) \xrightarrow{\sim} \operatorname{Ker}\mathcal{N}_G$ the subgroup $B^1(G, M)$ corresponds to IM; (c) is thereby proved. □

Given a finite normal (field) extension K/k, we can consider the cohomology groups of the Galois group $\mathcal{G} = \mathcal{G}(K/k)$. These cohomology groups give a sequence of very interesting invariants of the layer K/k. As nomenclature, the groups $H^r(\mathcal{G}(K/k), M)$ are called the *Galois cohomology groups of K/k with values in M*, and as notation we write $H^r(K/k, M)$ for $H^r(\mathcal{G}(K/k), M)$. Probably, the most useful facts about Galois cohomology are the two forming the statement of the next proposition.

Proposition 4.57 *(Hilbert Theorem 90[4].) If K/k is a finite normal extension, then*

(1) $H^r(K/k, K^+) = (0)$, all $r > 0$ and

(2) $H^1(K/k, K^*) = (0)$.

[4]When K/k is a cyclic extension, statement (2) is the essential content of Theorem 90 (§54) of Hilbert's magnificient paper [23]. The general case of a normal extension is due to E. Noether.

Proof. For (1), we examine the layer $K/k^{(*)}$ and apply the normal basis theorem to it. I claim that, as $\mathcal{G} = \mathcal{G}(K/k)$-modules, $\mathrm{Map}(\mathcal{G}, k^{(*)})$ and K are isomorphic. If we show this, then Proposition 4.54 and our isomorphism establish (1).

If $f \in \mathrm{Map}(\mathcal{G}, k^{(*)})$, we send f to $\sum_{\sigma\in\mathcal{G}} f(\sigma)\sigma^{-1}\theta$, where θ is a normal basis element for $K/k^{(*)}$. The linear independence of the elements $\{\sigma\theta\}_{\sigma\in\mathcal{G}}$ shows our map is injective; that it is surjective is obvious. As for the $\mathcal{G}$-action, call our map Θ then,

$$\begin{aligned}\Theta(\tau f) = \sum_{\sigma\in\mathcal{G}} (\tau f)(\sigma)\sigma^{-1}\theta &= \sum_{\sigma\in\mathcal{G}} f(\sigma\tau)\sigma^{-1}\theta \\ &= \sum_{\rho\in\mathcal{G}} f(\rho)\tau\rho^{-1}\theta \\ &= \tau\cdot\sum_{\rho\in\mathcal{G}} f(\rho)\rho^{-1}\theta \\ &= \tau\Theta(f),\end{aligned}$$

as contended.

The proof of (2) has a similar flavor but depends on Dedekind's theorem (our Theorem 4.30). We take as family of characters of K^* the elements of $\mathcal{G} = \mathcal{G}(K/k)$. By Dedekind's theorem, they are independent; that is, any relation (with $x_\sigma \in K^*$)

$$\sum_{\sigma\in\mathcal{G}} x_\sigma\sigma(\lambda) = 0, \quad \text{all } \lambda \in K^*$$

necessarily implies that all the $x_\sigma = 0$. Given $f \in Z^1(K/k, K^*)$, take as the x_σ the elements $f(\sigma) \in K^*$. None of the x_σ are zero, so there must be a $\lambda \in K^*$ with

$$\beta = \sum_{\sigma\in\mathcal{G}} f(\sigma)\sigma(\lambda) \neq 0.$$

Now, $\tau\beta = \sum_{\sigma\in\mathcal{G}} \tau f(\sigma)\tau\sigma(\lambda)$, and as f is a 1-cocycle, we have $\tau f(\sigma)\cdot f(\tau) = f(\tau\sigma)$. Thus,

$$\begin{aligned}\beta = \sum_{\sigma\in\mathcal{G}} f(\tau\sigma)(\tau\sigma)(\lambda) &= \sum_{\sigma\in\mathcal{G}} (\tau f(\sigma)\cdot f(\tau))(\tau\sigma)(\lambda) \\ &= f(\tau)\sum_{\sigma\in\mathcal{G}} \tau f(\sigma)(\tau\sigma)(\lambda) \\ &= f(\tau)\cdot(\tau\beta).\end{aligned}$$

Let $\alpha = 1/\beta$, then $(\tau\alpha)/\alpha = f(\tau)$, as required. □

Remark: Proposition 4.57 gives yet another interpretation of the normal basis theorem. It shows that *for K normal over k, the $\mathcal{G}(K/k)$-module K is of the form* $\mathrm{Map}(\mathcal{G}, -)$*; namely, it is* $\mathrm{Map}(\mathcal{G}, k^{(*)})$.

Norms and Traces.

If K/k is a field extension and $\alpha \in K$, then the k-vector space map

$$T_\alpha\colon K \to K \quad \textit{via} \quad T_\alpha(\lambda) = \alpha\lambda$$

has a trace and a determinant.

Definition 4.14 The *trace,* $\mathrm{tr}_{K/k}(\alpha)$*, of α from K to k* is the trace of T_α; the *norm,* $\mathcal{N}_{K/k}(\alpha)$*, of α from K to k* is $\det T_\alpha$.

The following three facts are extremely simple to prove and are left as (DX):

Fact I. $\mathrm{Tr}_{K/k}(\alpha)$ *is additive;* $\mathcal{N}_{K/k}(\alpha)$ *is multiplicative.*

Fact II. If $\alpha \in k$*, then*

$$\mathrm{Tr}_{K/k}(\alpha) = [K\colon k]\alpha \quad \textit{and} \quad \mathcal{N}_{K/k}(\alpha) = \alpha^{[K\colon k]}.$$

Fact III. If $L \supseteq K \supseteq k$*, then*

$$\mathrm{Tr}_{L/k}(\alpha) = \mathrm{Tr}_{K/k}(\mathrm{Tr}_{L/K}(\alpha)) \quad \textit{and} \quad \mathcal{N}_{L/k}(\alpha) = \mathcal{N}_{K/k}(\mathcal{N}_{L/K}(\alpha)).$$

Of course, from Facts II and III , we find

$$\mathrm{Tr}_{K/k}(\alpha) = [K\colon k(\alpha)]\mathrm{Tr}_{k(\alpha)/k}(\alpha) \quad \text{and} \quad \mathcal{N}_{K/k}(\alpha) = (\mathcal{N}_{k(\alpha)/k}(\alpha))^{[K\colon k(\alpha)]}.$$

When K/k is normal, more can be said. First, assume K/k is both normal and separable, then

$$\mathrm{Tr}_{K/k}(\alpha) = \sum_{\alpha\in\mathcal{G}(K/k)} (\sigma\alpha) \quad \text{and} \quad \mathcal{N}_{K/k}(\alpha) = \prod_{\alpha\in\mathcal{G}(K/k)} (\sigma\alpha).$$

Both of these statements are very easy corresponding to the fact that the roots of the characteristic polynomial of T_α are exactly the various $\sigma\alpha$ as σ ranges over $\mathcal{G}(K/k)$.

Now allow inseparability. We have

$$\mathcal{N}_{K/k}(\alpha) = \mathcal{N}_{K_{(*)}/k}(\mathcal{N}_{K/K_{(*)}}(\alpha)) = \mathcal{N}_{k^{(*)}/k}(\mathcal{N}_{K/k^{(*)}}(\alpha)).$$

Hence, we must first investigate $\mathcal{N}_{K/k}(\alpha)$ when K/k is purely inseparable. I claim the value of this norm is $\alpha^{[K\colon k]}$. To see this, observe that

$$\mathcal{N}_{K/k}(\alpha) = \mathcal{N}_{k(\alpha)/k}(\mathcal{N}_{K/k(\alpha)}(\alpha)) = \mathcal{N}_{k(\alpha)/k}(\alpha)^{[K\colon k(\alpha)]}. \tag{$\dagger$}$$

Now, the minimal and characteristic polynomials for T_α on the vector space $k(\alpha)$ are $X^q - c$, where $q = [k(\alpha)\colon k] = p^r$, and $c = \alpha^q$. Here, $p = \mathrm{char}(k)$. Therefore, the norm of α is $\det(T_\alpha) = c$ if p is odd and $-c = c$ if p is 2. Hence, $\mathcal{N}_{k(\alpha)/k}(\alpha) = \alpha^q = \alpha^{[k(\alpha)\colon k]}$. Put this together with ($\dagger$) above and obtain our claim. The general case now is

$$\mathcal{N}_{K/k}(\alpha) = \mathcal{N}_{K_{(*)}/k}(\alpha^{[K\colon k]_i}) = (\mathcal{N}_{K/k^{(*)}}(\alpha))^{[K\colon k]_i}.$$

Proposition 4.58 *(Original Form of Hilbert Theorem 90[5].) Suppose that* K/k *is normal and that* $K_{(*)}/k$ *is a cyclic extension. Then, a necessary and sufficient condition that* $\mathcal{N}_{K/k}(\alpha) = 1$ *is that there exists a* $\beta \in K_{(*)}$ *so that*

$$\alpha^{[K_{(*)}(\alpha)\colon K_{(*)}]} = \frac{\sigma\beta}{\beta}.$$

Here, σ *is an a priori chosen generator of* $\mathcal{G}(K_{(*)}/k)$*.*

Proof. This is merely the confluence of Propositions 4.56 and 4.57. If $\alpha \in K_{(*)}$, statements (b) and (c) of Proposition 4.56 and (2) of Proposition 4.57 give the statement that $\mathcal{N}_{K_{(*)}/k}(\alpha) = 1$ iff $\alpha = \sigma\beta/\beta$ for some $\beta \in K_{(*)}$. But, $\mathcal{N}_{K/k}(\alpha) = (\mathcal{N}_{K_{(*)}/k}(\alpha))^{[K\colon K_{(*)}]}$ in this case, and $[K\colon K_{(*)}]$ is a p-power. Therefore, $\mathcal{N}_{K/k}(\alpha) = 1$ iff $\mathcal{N}_{K_{(*)}/k}(\alpha) = 1$.

[5]Of course, Hilbert dealt only with the separable case.

Suppose now that $\alpha \in K$ yet $\alpha \notin K_{(*)}$. Then, $\alpha^{[K_{(*)}(\alpha)\colon K_{(*)}]}$ is in $K_{(*)}$. But,

$$\mathcal{N}_{K/k}(\alpha) = \mathcal{N}_{K_{(*)}/k}(\mathcal{N}_{K_{(*)}(\alpha)/K_{(*)}}(\alpha))^{[K\colon K_{(*)}]}.$$

As $[K\colon K_{(*)}]$ is a p-power, the left hand side is 1 iff $\mathcal{N}_{K_{(*)}/k}(\mathcal{N}_{K_{(*)}(\alpha)/K_{(*)}}(\alpha))$ is 1. By our remarks above, this last quantity is exactly $\mathcal{N}_{K_{(*)}/k}(\alpha^{[K_{(*)}(\alpha)\colon K_{(*)}]})$; so, we can apply the first part of the proof to the element $\alpha^{[K_{(*)}(\alpha)\colon K_{(*)}]}$, and we are done. □

It is not clear that β in Proposition 4.58 is of the form γ^q (where $q = [K_{(*)}(\alpha)\colon K_{(*)}]$), because in the proof of Proposition 4.57 part (2), the element λ may not be a qth power. If it proves to be so, then α would be $\sigma\gamma/\gamma$.

4.8 Krull's Galois Theory

In our treatment of Galois theory, the extensions were assumed finite. W. Krull discovered a natural way to treat (possibly) infinite algebraic extensions, His method leads to a non-trivial topology on the Galois group. We begin with the generalization of the extension lemma.

Theorem 4.59 *(General Extension Lemma) Suppose K/k is an algebraic extension and $\widetilde{k}$ is another field isomorphic to k via $\theta\colon k \to \widetilde{k}$. Let Γ be a field related to $\widetilde{k}$, but otherwise arbitrary. Then, there exists an algebraic extension, $\widetilde{K}/\widetilde{k}$, with $\widetilde{K} \underset{\text{rel}}{\sim} \Gamma$ and an extension of θ to an isomorphism $\widetilde{\theta}\colon K \to \widetilde{K}$.*

$$\begin{array}{ccccc} K & \xrightarrow{\widetilde{\theta}} & \widetilde{K} & \underset{\text{rel}}{\sim} & \Gamma \\ {\scriptstyle\text{alg}}\uparrow & & {\scriptstyle\text{alg}}\uparrow & & \| \\ k & \xrightarrow[\theta]{} & \widetilde{k} & \underset{\text{rel}}{\sim} & \Gamma \end{array}$$

Proof. This is a standard use of Zorn's lemma. We let

$$\mathcal{S} = \{(L, \varphi, \widetilde{L}) \mid L/k \text{ is algebraic, } \varphi \text{ extends } \theta \text{ and is an isomorphism } L \longrightarrow \widetilde{L} \text{ and } \widetilde{L} \underset{\text{rel}}{\sim} \Gamma\}.$$

Notice that the $\widetilde{L}$ in $\mathcal{S}$ are automatically algebraic over $\widetilde{k}$. We partially order $\mathcal{S}$ *via* the usual:

$$(L, \varphi, \widetilde{L}) \leq (M, \psi, \widetilde{M}) \quad \text{iff} \quad L \subseteq M;\ \widetilde{L} \subseteq \widetilde{M};\ \psi \restriction L = \varphi.$$

Of course, $\mathcal{S}$ is inductive; so, let $(L_0, \varphi_0, \widetilde{L}_0)$ be a maximal element of $\mathcal{S}$. Were $L_0 \neq K$, there would be some $\alpha \in K$ with $\alpha \notin L_0$. Then, the extension lemma for the finite extension $L_0(\alpha)/L_0$ would yield $(L_0(\alpha), \widetilde{\varphi}_0, \widetilde{L}_0)$ an element in $\mathcal{S}$ bigger that our maximal element—a contradiction. Therefore, $L_0 = K$. □

The material on splitting fields, etc. of Section 4.4 carries over provided no statement involving finiteness is used (e.g., statement (3) of Proposition 4.25 would be omitted in the general case that M/k was algebraic, not necessarily finite). The corollaries SMA, I and SMA II (Corollary 4.27 and Corollary 4.28) go over as does the existence of a normal closure.

Proposition 4.60 *Suppose K/k is an algebraic extension and write $\{K_\alpha/k \mid \alpha \in \Lambda\}$ for the family of sub-extensions of K/k of finite degree. Then, our family is a right mapping family in a natural way and we have*

$$K = \varinjlim_{\alpha} K_\alpha.$$

If K/k is normal, we may restrict the K_α/k to the finite normal extensions. Conversely, if $K = \varinjlim_{\alpha} K_\alpha$ and each K_α is normal over k, then so is K.

Proof. Of course, we define $\alpha \leq \beta$ (in Λ) when and only when $K_\alpha \subseteq K_\beta$ (everything takes place inside K). The map $K_\alpha \longrightarrow K_\beta$ is the inclusion. Since we have the inclusions $K_\alpha \hookrightarrow K$, consistent with the $K_\alpha \longrightarrow K_\beta$, we get the canonical homomorphism

$$\varinjlim_{\alpha} K_\alpha \longrightarrow K.$$

Choose $\xi \in K$. Then $k(\xi)$ is some K_α, and it is clear that $\xi \mapsto \text{can}_\alpha(\xi) \in \varinjlim_{\alpha} K_\alpha$ is well-defined and provides an inverse map to that above.

Of course, the family of finite normal extensions M_α/k is final in the family of all finite extensions provided K/k is itself normal. So, all we need prove is the last statement. We have $K = \varinjlim_{\alpha} K_\alpha$ and each

K_α is normal over k. If $\xi \in K$, there is an α so that $\xi \in K_\alpha$. Then all the k-conjugates of ξ lie in K_α; hence, they are in K. □

If Ω is an algebraic normal extension of K, then we consider the group $\mathrm{Aut}_k(\Omega)$. We topologize $\mathrm{Aut}_k(\Omega)$ by taking as a fundamental set of neighborhoods about 1 the subgroups of finite index in $\mathrm{Aut}_k(\Omega)$. Of course, it is the same to take the *normal* subgroups of finite index as our basic neighborhoods of $\{1\}$ in $\mathrm{Aut}_k(\Omega)$. (Remember: To get the neighborhoods about $\sigma \in \mathrm{Aut}_k(\Omega)$, we take the cosets σH, where the H are our neighborhoods about 1.) This renders $\mathrm{Aut}_k(\Omega)$ a Hausdorff topological group (use ordinary Galois Theory to see this) and it is this group *together with its topology* that we call the *Galois group of* Ω *over* k and denote by $\mathcal{G}(\Omega/k)$. The topology itself is the *Krull topology*.

Theorem 4.61 *The group $\mathcal{G} = \mathcal{G}(\Omega/k)$ is compact and totally disconnected in its Krull topology. In fact, we have $\mathcal{G}(\Omega/k) = \varprojlim_{\mathfrak{H}} \mathcal{G}/\mathfrak{H}$, where the left limit is taken over all open subgroups, $\mathfrak{H}$, of $\mathcal{G}$. Thus, $\mathcal{G}(\Omega/k)$ is a profinite group. Moreover, if we write $\Omega = \varinjlim_{\alpha} \Omega_\alpha$, where each Ω_α is a finite normal extension of k, then $\mathcal{G}(\Omega/k) = \varprojlim_{\alpha} \mathcal{G}(\Omega_\alpha/k)$.*

Proof. If $\sigma \in \mathcal{G}(\Omega/k)$ and Ω_α is one of the finite normal subextensions of Ω/k, then $\sigma \restriction \Omega_\alpha$ is in $\mathcal{G}(\Omega_\alpha/k)$. The maps $\pi_\alpha \colon \mathcal{G}(\Omega/k) \to \mathcal{G}(\Omega_\alpha/k)$ are consistent and hence we obtain the commutative diagram

$$\begin{array}{ccc} \mathcal{G}(\Omega/k) & \xrightarrow{\quad\varphi\quad} & \varprojlim_{\beta} \mathcal{G}(\Omega_\beta/k) \\ & {\scriptstyle \pi_\alpha} \searrow \quad \swarrow {\scriptstyle \mathrm{can}_\alpha} & \\ & \mathcal{G}(\Omega_\alpha/k) & \end{array}$$

If $\xi \in \varprojlim_{\beta} \mathcal{G}(\Omega_\beta/k)$, then ξ consists in a collection (ξ_β) where $\xi_\beta \in \mathcal{G}(\Omega_\beta/k)$ and when $k \subseteq \Omega_\beta \subseteq \Omega_\gamma$, we have $\xi_\gamma \restriction \Omega_\beta = \xi_\beta$. Since each $x \in \Omega$ lies in some finite normal extension of k, we have $x \in \Omega_\beta$ for various β. Then $\xi_\beta(x)$ is well-defined and our collection $(\xi_\beta) = \xi$ gives rise to an element of $\mathcal{G}(\Omega/k)$. Therefore we have a map

$$\varprojlim_{\beta} \mathcal{G}(\Omega_\beta/k) \xrightarrow{\psi} \mathcal{G}(\Omega/k)$$

plainly inverse to φ. Now, a neighborhood of 1 in $\varprojlim_{\beta} \mathcal{G}(\Omega_\beta/k)$ consists of those tuples (ξ_β) for which finitely many β are the identity and otherwise arbitrary (though consistent). Such tuples when restricted to the compositum of the Ω_β for which $\xi_\beta = 1$, are the identity on the compositum which is a field of finite degree over k, call it L. I claim $\mathcal{G}(\Omega/L)$ has finite index in $\mathcal{G}(\Omega/k)$. For L is normal and the usual argument shows $\mathcal{G}(\Omega/L) \lhd \mathcal{G}(\Omega/k)$. Moreover, SMA II (in its extended form) implies that $\mathcal{G}(\Omega/k) \longrightarrow \mathcal{G}(\Omega/L)$ is surjective. Hence, the exact sequence

$$0 \longrightarrow \mathcal{G}(\Omega/L) \longrightarrow \mathcal{G}(\Omega/k) \longrightarrow \mathcal{G}(L/k) \longrightarrow 0$$

gives the finite index assertion immediately. But then, we see that open neighborhoods of 1 in the Krull topology on $\mathcal{G}(\Omega/k)$ correspond to open neighnorhoods of 1 in the natural (product) topology on $\varprojlim_{\beta} \mathcal{G}(\Omega_\beta/k)$. Consequently, our maps φ and ψ are homeomorphisms.

Since $\mathcal{G}(\Omega_\beta/k)$ is compact, so is $\mathcal{G}(\Omega/k)$ in the Krull topology, and of course $\mathcal{G}(\Omega/k)$ is a profinite group. Every profinite group is totally disconnected (DX); so, $\mathcal{G}(\Omega/k)$ is totally disconnected.

That $\mathcal{G}$ is $\varprojlim_{\mathfrak{H}} \mathcal{G}/\mathfrak{H}$ as $\mathfrak{H}$ ranges over all open normal subgroups of $\mathcal{G}$ is the same kind of argument (remember that $\mathfrak{H}$ will be closed and of finite index). Or, it follows immediately from the next lemma whose proof is easy. □

Proposition 4.62 *Suppose $\mathcal{G}$ is a compact (Hausdorff) group and $\mathcal{G}_\alpha, \mathfrak{H}_\alpha$ are two families of closed subgroups with $\mathfrak{H}_\alpha \lhd \mathcal{G}_\alpha$ for every α. Assume that the indices $\alpha, \beta, \ldots$ form a directed set and that for every $\beta \geq \alpha$ we have $\mathcal{G}_\beta \subseteq \mathcal{G}_\alpha$ and $\mathfrak{H}_\beta \subseteq \mathfrak{H}_\alpha$. Then, the groups $\mathcal{G}_\alpha/\mathfrak{H}_\alpha$ form an inverse mapping family in a natural way and*

$$\varprojlim_\alpha \mathcal{G}_\alpha/\mathfrak{H}_\alpha = \bigcap_\alpha \mathcal{G}_\alpha \Big/ \bigcap_\alpha \mathfrak{H}_\alpha.$$

Using the same notations as in our treatment of standard Galois Theory, we can now extend the fundamental theorem to the general case. First of all, Lemma 4.33 and the material on Galois equivalence (between Lemma 4.33 and Theorem 4.38) go over word for word (together with no change in their proofs). So, here is the theorem.

Theorem 4.63 *(Fundamental Theorem of Galois Theory, General Case) If Ω/k is a normal (not necessarily finite) algebraic extension, then the mappings*

$$\begin{array}{rcl} [L] & \mapsto & \mathcal{G}(\Omega/L) \\ \mathcal{H} & \mapsto & [\mathrm{Fix}(\mathcal{H})] \end{array}$$

establish a one-to-one order-inverting correspondence between Galois classes of extension fields of k and **closed** *subgroups of $\mathcal{G}(\Omega/k)$. In this correspondence:*

(a) $L^{(*)}$ is normal over k iff $L_{(*)}$ is normal over k iff $\mathcal{G}(\Omega/L)$ is a normal subgroup of $\mathcal{G}(\Omega/k)$.

(b) Under the conditions of (a), we have a natural exact sequence

$$0 \longrightarrow \mathcal{G}(\Omega/L) \longrightarrow \mathcal{G}(\Omega/k) \longrightarrow \mathcal{G}(L_{(*)}/k) \longrightarrow 0$$

of compact topological groups.

(c) A necessary and sufficient condition that $L_{(*)}$ be a finite extension of k is that $\mathcal{G}(\Omega/L)$ be an open subgroup of $\mathcal{G}(\Omega/k)$. In this case,

$$(\mathcal{G}(\Omega/k)\colon \mathcal{G}(\Omega/L)) = [L_{(*)}\colon k].$$

Proof. If $\alpha \in \Omega$, I claim $\{\sigma \mid \sigma(\alpha) = \alpha\}$ is an open (hence closed) subgroup of $\mathcal{G}(\Omega/k)$. Notice that if this claim is proved, then

$$\mathcal{G}(\Omega/L) = \{\sigma \mid (\forall \alpha \in L)(\sigma(\alpha) = \alpha)\} = \bigcap_{\alpha \in L} \{\sigma \mid \sigma(\alpha) = \alpha\}$$

is a closed subgroup. Now, $k(\alpha)$ has finite degree over k, so its normal closure, L, also has finite degree. In the proof of Theorem 4.61, we showed $\mathcal{G}(\Omega/L)$ has finite index in $\mathcal{G}(\Omega/k)$. But,

$$\mathcal{G}(\Omega/L) \subseteq \mathcal{G}(\Omega/k(\alpha)) \subseteq \mathcal{G}(\Omega/k)$$

and therefore $(\mathcal{G}(\Omega/k)\colon \mathcal{G}(\Omega/k(\alpha))) < \infty$. By definition of the Krull topology, the subgroup $\mathcal{G}(\Omega/k(\alpha))$ is open, as contended.

Next, just as in the usual (finite) case we see that $\mathcal{G}(\Omega/L^{(*)}) = \mathcal{G}(\Omega/L)$ and if $L = \mathrm{Fix}(\mathcal{G})$, then $L = L^{(*)}$. So, if we start with $[L]$, then we get $\mathcal{G}(\Omega/L)$ which is $\mathcal{G}(\Omega/L^{(*)})$. However, as mentioned, $\mathrm{Fix}(\mathcal{G}(\Omega/L^{(*)}))$ is $(L^{(*)})^{(*)} = L^{(*)}$ and so the correspondence inverts if we start from the field side.

Now take a *closed* subgroup $\mathfrak{H}$ and form $[\mathrm{Fix}(\mathfrak{H})]$. If $L = \mathrm{Fix}(\mathfrak{H})$, consider $\mathcal{G}(\Omega/L)$. Now $L = L^{(*)}$, so in what follows we consider only those subfields, M, of Ω/k with $M = M^{(*)}$. The Galois group don't change and all fields are now separable over the base field, $k^{(*)}$. For notation, drop all mention of "upper stars."'

We must show $\mathfrak{H} = \mathcal{G}(\Omega/L)$. We know $\mathfrak{H} \subseteq \mathcal{G}(\Omega/L)$ by definition of L. Observe that $\Omega = \varinjlim_{\alpha} K_\alpha$ for fields, K_α, finite and normal over k. We find as well that $L = \varinjlim_{\alpha} K_\alpha \cap L$. The Galois group $\mathcal{G}(\Omega/K_\alpha)$ is then an open, normal subgroup of $\mathcal{G}(\Omega/k)$ by definition of the Krull topology. Consider the subgroups $\mathfrak{H}\mathcal{G}(\Omega/K_\alpha)$, which contains $\mathfrak{H}$. I claim: $\mathrm{Fix}(\mathfrak{H}\mathcal{G}(\Omega/K_\alpha))$ is just $L\mathcal{K}_\alpha$. For, the elements of $\mathfrak{H}\mathcal{G}(\Omega/K_\alpha)$ are products $\sigma\tau$, where $\sigma \in \mathfrak{H}$ and $\tau \in \mathcal{G}(\Omega/K_\alpha)$. A ξ in $\mathrm{Fix}(\mathfrak{H}\mathcal{G}(\Omega/K_\alpha))$ satisfies $\sigma\tau(\xi) = \xi$, for all such σ and τ. In particular, when $\sigma = 1$, we find $\xi \in \mathrm{Fix}(\mathcal{G}(\Omega/K_\alpha) = K_\alpha$ (remember: $K_\alpha = K_\alpha^{(*)}$), and when $\tau = 1$, we find $\xi \in \mathrm{Fix}(\mathfrak{H}) = L$; hence, $\xi \in K_\alpha \cap L$. Conversely, if $\xi \in K_\alpha \cap L$ it is fixed by both $\mathfrak{H}$ and $\mathcal{G}(\Omega/K_\alpha)$; therefore, our cliam is proved. Then, we have the commutative diagram

$$\begin{array}{ccccccccc}
0 & \longrightarrow & \mathcal{G}(\Omega/K_\alpha) & \longrightarrow & \mathcal{G}(\Omega/k) & \longrightarrow & \mathcal{G}(K_\alpha/k) & \longrightarrow & 0 \\
 & & \| & & \uparrow & & \uparrow & & \\
0 & \longrightarrow & \mathcal{G}(\Omega/K_\alpha) & \longrightarrow & \mathcal{G}(\Omega/K_\alpha \cap L) & \longrightarrow & \mathcal{G}(K_\alpha/K_\alpha \cap L) & \longrightarrow & 0,
\end{array}$$

and from it we see that $\mathfrak{H}\mathcal{G}(\Omega/K_\alpha)$ corresponds to a subgroup of $\mathcal{G}(K_\alpha/k)$, *via* $\mathfrak{H}\mathcal{G}(\Omega/K_\alpha) \mapsto \mathfrak{H}\mathcal{G}(\Omega/K_\alpha)/\mathcal{G}(\Omega/K_\alpha)$. But then, the lower line of our diagram and the finite case of ordinary Galois theory show that

$$\mathfrak{H}\mathcal{G}(\Omega/K_\alpha)/\mathcal{G}(\Omega/K_\alpha) \cong \mathcal{G}(K_\alpha/K_\alpha \cap L).$$

We pass these isomorphisms to the projective limit over α; on the left hand side, Lemma 4.62 implies that we get

$$\bigcap_\alpha \mathfrak{H}\mathcal{G}(\Omega/K_\alpha)/\bigcap_\alpha \mathcal{G}(\Omega/K_\alpha) = \bigcap_\alpha \mathfrak{H}\mathcal{G}(\Omega/K_\alpha)$$

while on the right hand side we get $\mathcal{G}(\Omega/L)$. But, $\bigcap_\alpha \mathfrak{H}\mathcal{G}(\Omega/K_\alpha)$ is the closure of $\mathfrak{H}$ and $\mathfrak{H}$ is already closed. Therefore, $\mathfrak{H} = \mathcal{G}(\Omega/L)$.

The proofs of assertions (a) and (b) are now just as they were in the finite case. As for (c), we know that $\mathcal{G}(\Omega/L_{(*)}) = \mathcal{G}(\Omega/L)$ and that this subgroup is of finite index in $\mathcal{G}(\Omega_{(*)}/k)$. We take, as in the proof above, a family of fields, K_α, of finite degree and normal over k so that $\varinjlim_{\alpha} K_\alpha = \Omega$. Then

$$(\mathcal{G}(K_\alpha/k)\colon \mathcal{G}(K_\alpha/K_\alpha \cap L_{(*)})) = [K_\alpha \cap L_{(*)}\colon k]$$

by usual Galois theory. Pass to the limit over α, observe that the left side tends to $(\mathcal{G}(\Omega/k)\colon \mathcal{G}(\Omega/L))$ and in fact is constant as soon as $K_\alpha \supseteq L_{(*)}$, and we get (c). □

We can now extend the notions and results of the previous section on Galois cohomology to the general (not necessarily finite) case. All that is necessary is to sprinkle the word "continuous" in the appropriate places and use some case. The group G will be a profinite group, for example $G = \mathcal{G}(\Omega/k)$. All modules will be given the discrete topology unless otherwise noted and our action $G \times M \longrightarrow M$ will be assumed continuous. This means that for $m \in M$, there is an *open* subgroup, U, of G so that $Um = m$. We define

$$C^n(G, M) = \{f\colon G^r \to M \mid f \quad \text{is continuous}\},$$

and use the usual formula for δ, thus δ is a continuous function ($C^r(G, M)$ inherits the discrete topology from M as G is compact). The continuity of the cochains shows up as follows: If $N \lhd G$ (N is, of course, closed), then M^N is a G/N-module. If $\widetilde{N} \subseteq N$ and $\widetilde{N} \lhd N$, then there are maps $G/\widetilde{N} \longrightarrow G/N$ and $M^N \longrightarrow N^{\widetilde{N}}$. The latter two combine to give a map

$$C^r(G/N, M^N) \longrightarrow C^r(G/\widetilde{N}, M^{\widetilde{N}})$$

called *inflation from G/N to $G/\widetilde{N}$*; it is injective. We have

Proposition 4.64 *For a continuous G-module, M, for the profinite group, G, the modules $C^r(G/U, M^U)$ form a right mapping system as U runs over the open normal subgroups of G—the map being inflation. We have*

$$C^r(G, M) = \varinjlim_U C^r(G/U, M^U),$$

and passing to cohomology, we also have

$$H^r(G, M) = \varinjlim_U H^r(G/U, M^U).$$

The proof of this is now routine and may be safely left to the reader (DX). The cohomological triviality of $\mathrm{Map}(G, B)$ (continuous functions, of course) expressed by Proposition 4.54 carries over and so does Proposition 4.57 (Hilbert Theorem 90).

4.9 Kummer Theory

In this section we consider base fields containing prescribed roots of unity. At first, we assume our base field contains a primitive mth root of unity. Notice that something mildly subtle is happening. We are not merely assuming all mth roots of 1 lie in k, for that would be true is $m = \text{char}(k) > 0$, yet there is *no* primitive mth root of unity in this case because 1 is the only mth root of unity when $m = \text{char}(k) > 0$. When a primitive mth root of 1 lies in k, then necessarily $(\text{char}(k), m) = 1$. (Else, $p = \text{char}(k) \mid m$ and $X^m - 1 = (X^q - 1)^p$, when $m = pq$. So each mth root of 1 is already a qth root of 1 with $q < m$, contradicting primitivity.)

Proposition 4.65 *Suppose the field k contains a primitive mth root of 1. A necessary and sufficient condition that K/k be a normal, separable extension whose Galois group is cyclic of order m is that $K = k(\beta)$ where the minimal k-polynomial for β is $X^m - b$.*

Proof. ($\Leftarrow$). We assume $K = k(\beta)$ and β is a root of the irreducible k-polynomial $X^m - b$. Write μ for a primitive mth root of 1 in k, then $1, \mu, \ldots, \mu^{m-1}$ are all the mth roots of 1, they are all distinct and lie in k. Of course, $(m, \text{char}(k)) = 1$ shows K/k is separable and $[K\colon k] = m$. Now, the elements

$$\beta, \mu\beta, \mu^2\beta, \ldots, \mu^{m-1}\beta$$

are all distinct and all are roots of $X^m - b$, therefore K is a splitting field of $X^m - b$; so, K/k is indeed normal. If we consider the k-isomorphism $k(\beta) \longrightarrow k(\mu\beta)$, we see (even without SMA, I) that it gives an element, σ, of $\mathcal{G}(K/k)$. The powers of σ operate on β *via*

$$\sigma^r(\beta) = \mu^r\beta$$

and so $1, \sigma, \ldots, \sigma^{m-1}$ are m distinct elements of $\mathcal{G}(K/k)$. They thereby exhaust $\mathcal{G}(K/k)$ and ($\Leftarrow$) is proved.

($\Rightarrow$). We suppose here that K/k is normal, separable and $\mathcal{G}(K/k)$ is cyclic of order $m = [K\colon k]$. Now we know $\mathcal{N}_{K/k}(\mu) = \mu^m = 1$, so we can apply the original form of Hilbert Theorem 90. We find there exists a $\beta \in K$ so that $\sigma\beta = \mu\beta$, where σ is a generator of $\mathcal{G}(K/k)$. Then, of course, $\sigma^r\beta = \mu^r\beta$, and so β is left fixed only by the trivial subgroup of $\mathcal{G}(K/k)$. By the fundamental theorem of Galois Theory, $k(\beta) = K$. The minimal k-polynomial of β is then

$$\prod_{r=0}^{m-1}(X - \sigma^r\beta) = \prod_{r=0}^{m-1}(X - \mu^r\beta).$$

On the other hand, if $b = \beta^m$, then each σ^r fixes b and each $\mu^r\beta$ is a root of $X^m - b$. The minimal polynomial for β and $X^m - b$ both have degree m; so it is clear the latter polynomial *is* the minimal polynomial for β. □

Corollary 4.66 *Suppose the field k contains a primitive mth root of 1. If n is any divisor of m, then*

(1) A n.a.s.c. that an extension K/k of degree n be normal with cyclic Galois group is that $K = k(\alpha)$ where the minimal polynomial of α is $X^n - a$.

(2) The k-polynomial $X^m - a$ is irreducible in $k[X]$ if and only if for all divisors, d, of m with $d > 1$, we have $a \notin k^{*d}$.

Proof. First, as $n \mid m$, we can write $m = nd$. Then for our primitive mth root of 1 in k, μ, the element μ^d is a primitive nth root of 1 in k and (1) is simply a restatement of Proposition 4.65 with n replacing m.

For statement (2), first consider $X^m - a$ and let α be a root in an overfield, Ω, of k. Write d for the smallest power of α lying in k. Then, the usual division algorithm argument shows that if $\alpha^q \in k$, we have $d \mid q$; in particular, $d \mid m$. I claim the polynomial $X^d - b$ is the minimal k-polynomial for α (here, $b = \alpha^d \in k$), in particular it is irreducible. To see this, let $f(X)$ be the minimal k-polynomial for α and have degree t.

Thus, $f(X) \mid (X^d - b)$ and $t \leq d$. Yet, the roots of $X^d - b$ are $\alpha, \zeta\alpha, \ldots, \zeta^{d-1}\alpha$, where $\zeta = \mu^n$ is a primitive dth root of 1. Thus,

$$f(X) = (X - \zeta^{i_1}\alpha) \cdots (X - \zeta^{i_t}\alpha)$$

and its constant term is therefore $\pm\Big(\prod_{l=1}^{t} \eta^{i_l}\Big)\alpha^t$. But then, $\alpha^t \in k$; so, $d \mid t$ and $t \leq d$. We find $t = d$ and $f(X) = X^d - b$.

Now if $a \notin k^{*q}$ for any $q \mid m$ with $q > 1$, then the smallest power of α to lies in k is the mth. Else, $\alpha^d \in k$ implies $dq = m$ (as above) and $a = \alpha^m = (\alpha^d)^q \in k^{*q}$ and $q > 1$ if $d < m$. By our claim, $X^m - a$ is irreducible in $k[X]$.

Finally, assume $X^m - a$ is irreducible. Were $a \in k^{*d}$ where $d \mid m$ and $d > 1$, then as $\alpha^m = a$, we have $(\alpha^q)^d = \beta^d$ for some $\beta \in k^*$. Therefore, $\alpha^q = z\beta$ for some z a dth root of 1 (hence, in k). It follows that the smallest power of α in k is, say, δ where $\delta \leq q < m$. Just as before, $\delta \mid m$ and $X^\delta - b$ is k-irreducible, where $b = \alpha^\delta$. Write $\delta r = m$ and look at $\alpha, \mu\alpha, \ldots, \mu^{r-1}\alpha$ (as usual μ is our primitive mth root of 1). Each of these elements has δth power in k and δ is minimal. Set $\zeta = \mu^\delta$, then $X^\delta - \zeta^i b$ is the minimal k-polynomial for $\mu^i\alpha$ by our claim above. But,

$$X^m - a = (X^\delta - b)(X^\delta - \zeta b) \cdots (X^\delta - \zeta^{r-1} b),$$

contradicting the irreducibility of $X^m - a$. □

An important part of the proof above should be isolated and recorded:

Corollary 4.67 *If k contains a primitive mth root of 1 and K is an overfield of k, then given $\alpha \in K$ with $\alpha^m \in k^*$, the minimal k-polynomial for α is $X^d - \alpha^d$, where d is the smallest positive integer so that $\alpha^d \in k$. In fact, $d \mid m$.*

Now, we can make an obvious attempt to "classify" the cyclic overfields of degree n ($n \mid m$) of k when k possesses a primitive mth root of 1. Namely, such a K is $k(\alpha)$ and we could send α to $\overline{\alpha^n}$ where $\overline{\alpha^n}$ is the image of α^n in k^*/k^{*n}. But, α is not unique and its choice depends on μ and σ (a generator of $\mathcal{G}(K/k)$). There is a better way:

Theorem 4.68 *(Kummer) Suppose k is a field possessing a primitive mth root of 1. Write Ω for the maximal, abelian, m-torsion extension of k and denote by $\mathcal{G}$ its (Krull topologized) Galois group. Then there is a natural continuous pairing*

$$\mathcal{G} \prod k^*/k^{*m} \longrightarrow \mu_m \,(=\ m\text{th root of } 1),$$

and it makes $\mathcal{G}$ the Pontrjagin dual of k^*/k^{*m}.

Proof. Choose $\sigma \in \mathcal{G}$ and $\overline{a} \in k^*/k^{*m}$. Lift $\overline{a}$ to some $a \in k^*$ and take an mth root of a in an overfield, call it α. We know that $K = k(\alpha)$ is cyclic of degree d and $d \mid m$ by our above propositions. So, $K \subseteq \Omega$ (we fix an algebraic closure of k and work inside it) and $\sigma\alpha$ makes sense. We set

$$(\sigma, \overline{a}) = \frac{\sigma\alpha}{\alpha}.$$

Note that

$$\Big(\frac{\sigma\alpha}{\alpha}\Big)^m = \frac{\sigma(\alpha^m)}{\alpha^m} = \frac{\alpha^m}{\alpha^m} = 1,$$

therefore $(\sigma, \overline{a}) \in \mu_m$. Let's check that $(\sigma, \overline{a})$ is well-defined. First, if we change the mth root of a we get $\zeta\alpha$ where ζ is some mth root of 1 (hence, $\zeta \in k^*$). Then

$$\frac{\sigma(\zeta\alpha)}{\zeta\alpha} = \frac{\sigma\alpha}{\alpha},$$

so there is no problem with the choice of α. If we lift $\overline{a}$ to some $b \in k^*$, then $b = \lambda^m a$ for some $\lambda \in k^*$. Thus, β, an mth root of b is $\zeta\lambda\alpha$ for some ζ as above. Once again,

$$\frac{\sigma\beta}{\beta} = \frac{\sigma(\zeta\lambda\alpha)}{\zeta\lambda\alpha} = \frac{\sigma\alpha}{\alpha},$$

and so $(\sigma, \overline{a})$ is a well -defined mth root of 1.

It is easy to see that $(\sigma, \overline{a})$ is bi-multiplicative (DX), so assume $(\sigma, \overline{a}) = 1$ for all $\sigma \in \mathcal{G}$. If a lifts $\overline{a}$ and $a \notin k^{*m}$ (i.e., $\overline{a} \neq 1$) then $K = k(\alpha)$ is a non-trivial cyclic degree d extension of k and $d \mid m$. But then, a generator, τ, of $\mathcal{G}(K/k)$ comes from some $\sigma \in \mathcal{G}$ and $\sigma\alpha = \tau\alpha = \zeta\alpha$ (some dth root of 1, say $\zeta \neq 1$). Hence, $(\sigma, \overline{a}) = \zeta \neq 1$, a contradiction. Therefore, $(\sigma, \overline{a})$ is non-degenerate on the right.

If $(\sigma, \overline{a}) = 1$ for all $\overline{a} \in k^*/k^{*m}$, then I claim σ must be 1. For, notice that when K/k is finite normal ($K \subseteq \Omega$), then $\mathcal{G}(K/k)$ is an abelian m-torsion group. Hence, $\mathcal{G}(K/k)$ is a product of various $\mathbb{Z}/d\mathbb{Z}$, where each $d \mid m$. This means that K is generated as a field by elements, α, for which $k(\alpha)$ is a cyclic extension of k. As K is arbitrary, it follows immediately that Ω is a field generated by such elements α. However, Proposition 4.65 and our assumption $(\sigma, \overline{a}) = 1$ (all $\overline{a}$), now yield $\sigma\alpha = \alpha$ for all the α's generating Ω. Thus, $\sigma = 1$, as claimed.

Lastly continuity of $\langle \sigma, \overline{a} \rangle \mapsto (\sigma, \overline{a})$ follows because on the entire open subgroup $\mathcal{G}(\Omega/k(\alpha))$ the pairing $\langle -, \overline{a} \rangle \mapsto (-, \overline{a})$ is identically 1. Here, α is, of course, an mth root of a. The product $\mathcal{G}(\Omega/k(\sigma)) \prod \{\overline{a}\}$ is an open neighborhood of 1 in $\mathcal{G}(\Omega/k) \prod k^*/k^{*m}$. □

Corollary 4.69 *Under the assumptions and notations of Theorem 4.68, there is a one-to-one correspondence between subgroups, S, of k^*/k^{*m} and sub-extensions K/k of Ω/k. It is given by*

$$S \longleftrightarrow K = k(S^{1/m}).$$

In all the foregoing, m was relatively prime to $\mathrm{char}(k) = p > 0$. What happens if $p \mid m$? Of course, we can factor m as $p^r\widetilde{m}$ with $(\widetilde{m}, p) = 1$. It's not hard to see that the case for this breaks up into the p^r case and the previous case. So, we'll assume $m = p^r$. Here, we will use the additive part of Hilbert 90 and the isomorphism $(\mathrm{Ker}\ \mathrm{Tr}_{K/k})/I_{K/k}K^+ \cong H^1(K/k, K^+)$ in case $\mathcal{G}(K/k)$ is cyclic.

So, assume K/k is a cyclic extension of degree p^r; choose a generator, σ, of $\mathcal{G}(K/k)$. For the element $1 \in k^+$, we have $\mathrm{Tr}_{K/k}(1) = [K\colon k] \cdot 1 = 0$, so there exists $\theta \in K^+$ with

$$\sigma(\theta) - \theta = 1, \quad \text{i.e.,} \quad \sigma(\theta) = \theta + 1.$$

The action of the Galois group on θ is given by

$$\sigma^i(\theta) = \theta + 1 \qquad 0 \leq i \leq p^r - 1.$$

Observe that only $\theta, \theta + 1, \ldots, \theta + (p-1)$ are distinct, after that we repeat these in order. Thus, the polynomial

$$g(X) = (X - \theta)(X - (\theta + 1)) \cdots (X - (\theta + (p-1)))$$

is the minimal k-polynomial for θ and $L = k(\theta)$ is a cyclic p-extension. The only case when $K = L$ is when $r = 1$; so, *from now on we'll assume K/k is cyclic of degree p.* Hence, $K = k(\theta)$ and $\sigma(\theta) = \theta + 1$. We can compute the minimal polynomial $g(X)$ as follows: Write $Y = X - \theta$, then $g(X) = Y(Y-1)(Y-2)\cdots(Y-(p-1))$. But, the elements $1, 2, \ldots, p-1$ are the $(p-1)$st roots of unity (and lie in $\mathbb{F}_p$, the prime field), therefore

$$g(X) = Y(Y^{p-1} - 1) = Y^p - Y = X^p - X - (\theta^p - \theta).$$

If we write $\wp(\theta) = \theta^p - \theta$, then we've proved the first part of

Theorem 4.70 *(E. Artin & O. Schreier, 1929) If k is a field of characteristic $p > 0$, then every cyclic p-extension, K/k, has the form $K = k(\theta)$ where $\wp(\theta) = \theta^p - \theta$ lies in k and the Galois group, $\mathcal{G}$, acts by a (prechosen) generator, σ, taking θ to $\theta + 1$. The minimal k-polynomial for θ is $X^p - X - \wp(\theta)$. Conversely, the polynomial $X^p - X - a$ is k-irreducible when and only when $a \notin \wp(k^+)$. If it is irreducible and θ is a root, then $k(\theta)$ is a normal, separable, cyclic p-extension of k.*

Proof. If $a \in \wp(k)$ so that $a = b^p - b$ for some $b \in k$, then

$$(X - b)(X - (b + 1)) \cdots (X - (b + (p - 1))$$

is exactly $X^p - X - \wp(b) = X^p - X - a$; so, our polynomial splits in $k[X]$.

If $a \notin \wp(k^+)$, the polynomial $X^p - X - a$ has no root in k. Adjoin a root to k, we get an extension $K = k(\theta)$. Now, $\theta^p - \theta = a$, so $(\theta + i)^p - (\theta + i) = a$, too, where $0 \leq i \leq p - 1$. Therefore, all the roots of $X^p - X - a$ lie in K and K is a normal extension. But the roots of $X^p - X - a$ are all distinct, therefore $X^p - X - a$ is separable and we find that K/k is a normal, separable extension.

If d is the degree of θ over k, then $\theta, \theta + i_2, \ldots, \theta + i_d$ are the roots of its minimal k-polynomial. Were $d \neq p$, there would be an integer, j, so that $\theta + j$ is not a root of the minimal k-polynomial for θ. Yet, $k(\theta + j) = k(\theta)$, so $\theta + j$ also has degree d over k and $\theta + j, \theta + j_2, \ldots, \theta + j_d$ are all the conjugates of $\theta + j$ and all distinct from each $\theta + i_l$. Continue in this way, we find the p roots $\theta, \theta + 1, \ldots, \theta + (p - 1)$ partition themselves into t blocks of d elements each. But then, $dt = p$ and p is prime. As $\theta \notin k$, we have $d > 1$ therefore $d = p$ and so $X^p - X - a$ is indeed irreducible. □

The analog of Kummer's theorem is

Theorem 4.71 *(E. Artin & O. Schreier) Suppose k is a field of characteristic $p > 0$ and write Ω for the maximal, abelian, p-torsion extension of k. If $\mathcal{G} = \mathcal{G}(\Omega/k)$ is the Galois group of Ω/k (with Krull topology), then there is a natural continuous pairing*

$$\mathcal{G} \prod k^+/\wp(k^+) \longrightarrow \mathbb{Z}/p\mathbb{Z} \ (\subseteq \mathbb{R}/\mathbb{Z})$$

and it makes $\mathcal{G}$ the Pontrjagin dual of $k^+/\wp(k^+)$.

Proof. Pick $\sigma \in \mathcal{G}$ and $\overline{a} \in k^+/\wp(k^+)$. Lift $\overline{a}$ to some $a \in k^+$ and let $\theta \in \overline{k}$ be a root of $X^p - X - a$. Define

$$(\sigma, \overline{a}) = \sigma\theta - \theta.$$

Note that unless $\overline{a} = 0$, the field $k(\theta)$ has degree p over k and is normal, separable cyclic. If $\overline{a} = 0$, then $\theta \in k$. Therfore, $k(\theta)$ is contained in Ω and $\sigma\theta$ makes sense. Now $\sigma\theta$ is a root of $X^p - X - a$ and so $\sigma\theta = \theta + j$ for some $j \in \mathbb{Z}/p\mathbb{Z}$; therefore, $(\sigma, \overline{a})$ is indeed in $\mathbb{Z}/p\mathbb{Z}$.

As in the proof of Kummer's theorem, $\langle \sigma, \overline{a} \rangle \mapsto (\sigma, \overline{a})$ is a pairing of the groups $\mathcal{G}(\Omega/k)$ and $k^+/\wp(k^+)$. Just as in the proof of that theorem, the field Ω is generated by the various θ's as above; so, if $(\sigma, \overline{a}) = 0$ for all $\overline{a}$, we find σ fixes all the θ's and thereby $\sigma = 1$. If $(\sigma, \overline{a}) = 0$ for all $\sigma \in \mathcal{G}(\Omega/k)$, then $\overline{a}$ must be 0 else the polynomial $X^p - X - a$ would be irreducible (Theorem 4.70) and $k(\theta)$, where θ is one of its roots, would be a cyclic p-extension. Then, $\sigma\theta = \theta + 1$ for some $\sigma \in \mathcal{G}(k(\theta)/k)$ and upon lifting σ to $\mathcal{G}(\Omega/k)$ we'd get $(\sigma, \overline{a}) \neq 0$, a contradiction.

Continuity is proved exactly as in Kummer's theorem, the open neighborhood on which $(\sigma, \overline{a})$ vanishes being $\mathcal{G}(\Omega/k(\theta)) \prod \{\overline{a}\}$. □

Corollary 4.72 *If* $\mathrm{char}(k) = p > 0$ *there is a one-to-one correspondence between subgroups, T, of $k^+/\wp(k^+)$ and p-torsion, abelian overfields of k. It is given by*

$$T \longleftrightarrow K = k(\wp^{-1}(T)).$$

What happens for p^r, $r > 1$? Here, the situation is sufficiently complicated that the solution had to wait until 1937. Then E. Witt introduced a ring, $W(k)$, called the ring of *Witt vectors over* k and he proved that even if $\mathrm{char}(k) = p > 0$, the ring $W(k)$ is an integral domain of characteristic 0. Now, it turns out that

$$W(k) = \varprojlim_n W_n(k),$$

where the $W_n(k)$ are "truncated" Witt vector rings. There is a map $F\colon W_n^n(k) \to W_n(k)$ playing the role of $\xi \mapsto \xi^p$ and one gets $\wp = F - \mathrm{id}$. When $n = 1$, the ring $W_1(k)$ is just k, and the exact sequence

$$0 \longrightarrow \mathbb{Z}/p\mathbb{Z} \longrightarrow k^+ \xrightarrow{\wp} \wp(k^+) \longrightarrow 0$$

becomes an exact sequence

$$0 \longrightarrow \mathbb{Z}/p^r\mathbb{Z} \longrightarrow W_r(k)^+ \xrightarrow{\wp} \wp(W_r(k)^+) \longrightarrow 0,$$

in the general case. It then turns out that if Ω is the maximal, abelian p^r-torsion extension of k, the Galois group, $\mathcal{G}(\Omega/k)$, is naturally Pontrjagin dual to $W_n(k)/\wp(W_n(k))$ by a pairing similar to the Artin-Schreier pairing. See Witt [49] for the details.

4.10 An Amazing Theorem of Galois Theory

Question: If k is a field with $k \neq \overline{k}$, when is $[\overline{k}\colon k]$ finite?

An example: $k = \mathbb{R}$; $K = \mathbb{C} = \mathbb{R}(i)$.

The answer of our question depends on an irreducibility criterion:

Theorem 4.73 *(Artin's Irreducibility Criterion) Given a field, k, consider the polynomial $X^n - a$, where $a \in k$. If p is prime and p divides n, assume $a \in (k^*)^p$. If $4 \mid n$, then assume as well that $a \notin (-4(k^*)^4)$. Under these conditions, $X^n - a$ is irreducible in $k[X]$.*

We will assume this theorem for the moment, and based on it we can prove

Theorem 4.74 *(Artin) Say k is a field, $\overline{k}$ is an algebraic closure of k and $1 < [\overline{k}\colon k] < \infty$. Then we have:*

(1) $\overline{k} = k(i) \quad (i^2 = -1)$.

(2) $\mathrm{char}(k) = 0$.

Proof. We claim that $\overline{k}/k$ is separable.

If not, let $\overline{k}_{(*)} = L$, then $\overline{L}/L$ is purely inseparable and $\overline{k} \neq L$. So, $L \neq L^p$ implies that there is $a \notin L^p$ (where $p = \mathrm{char}(k)$). We know $X^{p^n} - a$ is irreducible in $L[X]$ implies that L has extensions of degree p^n, for all n; yet, all these extensions are contained in $\overline{k}$, a contradiction.

Look at $k(i) \subseteq \overline{k}$; as $\overline{k}/k$ is separable, $\overline{k}/k(i)$ is normal, separable. Let $\mathcal{G} = \mathcal{G}(\overline{k}/k(i))$. We need to show that $\#(\mathcal{G}) = 1$.

Pick a prime, p, with $p \mid \#(\mathcal{G})$; let $\mathcal{H}$ be the subgroup of $\mathcal{G}$ of order p and write $L = \mathrm{Fix}(\mathcal{H})$.

Step 1. $p \neq \mathrm{char}(k)$.

If $p = \mathrm{char}(k)$, then, by separability, there is some $\beta \in \overline{k}$ so that $\mathrm{tr}_{k/L}(\beta) = 1$. We know that M/L is separable iff the bilinear form

$$(u, v) \mapsto \mathrm{tr}_{M/L}(uv)$$

is non-degenerate on the vector space M over the field L. Now, there is $\widetilde{\beta}$ so that $\mathrm{tr}_{\overline{k}/L}(1 \cdot \widetilde{\beta}) \neq 0$. Let $\lambda = \mathrm{tr}(\widetilde{\beta}) \in L$ and form $\beta = (1/\lambda)\widetilde{\beta}$. Then, we have

$$\mathrm{tr}_{\overline{k}/L}(\beta) = (1/\lambda)\mathrm{tr}_{\overline{k}/L}(\widetilde{\beta}) = \lambda/\lambda = 1.$$

As the trace is a sum and

$$\mathrm{tr}(\beta^p) = \mathrm{tr}(\beta)^p \quad (p = \mathrm{char}(k)),$$

we get $\mathrm{tr}_{\overline{k}/L}(\beta^p - \beta) = 0$. Our extension $\overline{k}/L$ is cyclic of degree p, say σ is a generator of $\mathcal{H}$. Note that for every $\xi \in \overline{k}$, we have

$$\mathrm{tr}_{\overline{k}/L}(\xi) = \mathrm{tr}_{\overline{k}/L}(\sigma\xi),$$

so $\mathrm{tr}_{\overline{k}/L}(\sigma\xi - \xi) = 0$. By "Additive Hilbert 90", every element of zero trace in $\overline{k}$ has the form $(\sigma\gamma - \gamma)$, for some $\gamma \in \overline{k}$. As $\mathrm{tr}_{\overline{k}/L}(\sigma\xi - \xi) = 0$, there is some $\gamma \in \overline{k}$ so that $\beta^p - \beta = \sigma\gamma - \gamma$. Now, the polynomial $X^p - X - \gamma \in \overline{k}[X]$ has a root in $\overline{k}$, as $\overline{k}$ is algebraically closed. Say $\alpha \in \overline{k}$ is such a root, then $\gamma = \alpha^p - \alpha$. We have

$$\beta^p - \beta = \sigma(\alpha^p - \alpha) - (\alpha^p - \alpha) = \sigma(\alpha)^p - \alpha^p - (\sigma(\alpha) - \alpha),$$

and so

$$\sigma(\alpha) - \alpha\beta = \sigma(\alpha)^p - \alpha^p - \beta^p = (\sigma(\alpha) - \alpha - \beta)^p.$$

Consequently, $\sigma(\alpha) - \alpha - \beta \in \mathbb{F}_p$, call it ν. It follows that $\sigma(\alpha) - \alpha = \beta + \nu$. Taking $\mathrm{tr}_{\overline{k}/L}$ on both sides, we get

$$0 = \mathrm{tr}_{\overline{k}/L}(\beta) + \mathrm{tr}_{\overline{k}/L}(\nu) = 1 + \mathrm{tr}_{\overline{k}/L}(\nu).$$

As $\nu \in \mathbb{F}_p \subseteq L$, we have

$$\mathrm{tr}_{\overline{k}/L}(\nu) = [\overline{k}\colon L]\nu = p\nu = 0,$$

which implies that $0 = 1 + 0 = 1$, a contradiction. Therefore, $\mathrm{char}(k) \neq p$, as claimed.

Step 2. L does not exist, i.e., as no p divides $\#(\mathcal{G})$, we have $\#(\mathcal{G}) = 1$; thus, $\overline{k} = k(i)$.

Adjoin to L a p-th root of unity, say ζ is such a primitive root. Then, $[\overline{k}\colon L(\zeta)] \mid p$; so, $[L(\zeta)\colon L]$ also divides p. But, $L(\zeta)/L$ has degree at most $p-1$. Indeed,

$$X^p - 1 = (X-1)(X^{p-1} + X^{p-2} + \cdots + X + 1),$$

so $L(\zeta) = L$ already, i.e., $\zeta \in L$. So, $[\overline{k}\colon L] = p$. As L has the p-th roots of unity, by Kummer's theorem (Theorem 4.65), we know $\overline{k} = l(\alpha)$, where α is a root of $X^p - a$, with $a \notin (L^*)^p$. But, if p is odd, the Artin irreducibility criterion implies that $X^p - a$ is also irreducible, so $[\overline{k}\colon L] \geq p^l$, for all $l \geq 0$, a contradiction. Therefore, we must have $p = 2$. Now, our situation is

(a) $\overline{k} = L(\alpha)$, where α is a root of $X^2 - a$, with $a \notin (L^*)^2$.

(b) $i \in k \subseteq L$.

Since $X^{2^r} - a$ cannot be irreducible for all $r \geq 0$, since otherwise we would have $[\overline{k}\colon K] \geq 2^r$ for all $r \geq 0$, it must be that $a \in (-4(L^*)^4)$ (by Artin's irreducibility). Thus, $a = -4b^4$, for some $b \in L^*$; it follows that $\alpha = \sqrt{a} = \pm 2ib^2$. As $2, i, b \in L$, we deduce that $\alpha \in L$, a contradiction. Therefore, $\#(\mathcal{G}) = 1$.

Step 3. $\mathrm{char}(k) = 0$.

If not, then say $q = \mathrm{char}(k)$ and write $\mathbb{F}_q$ for the prime field of k. Pick $r >> 0$, adjoin to $\mathbb{F}_q$ a primitive 2^r-th root of unity, call it ζ. Apply natural irrationalities to the picture show in Figure 4.2:

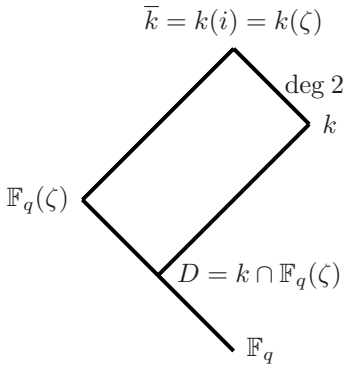

Figure 4.2: The Extension $\overline{k}/\mathbb{F}_q$

If σ is the generator of $\mathcal{G}(\overline{k}/k) = \mathbb{Z}/2\mathbb{Z}$, then $\sigma \restriction \mathbb{F}_q(\zeta)$ yields an automorphism and we know

$$\mathcal{G}(\overline{k}/k) \cong \mathcal{G}(\mathbb{F}_q(\zeta)/D) \hookrightarrow \mathcal{G}(\mathbb{F}_q(\zeta)/\mathbb{F}_q).$$

The extension $\mathbb{F}_q(\zeta)/\mathbb{F}_q$ is *cyclic* of degree 2^s. As a cyclic group has a unique subgroup of every possible order, there is a unique subfield of degree 2 in the extension $\mathbb{F}_q(\zeta)/\mathbb{F}_q$. If $\mathbb{F}_q < D$, then D contains this

unique extension. But, $\mathbb{F}_q(i)$ has degree 1 or 2 over $\mathbb{F}_q$, so $\mathbb{F}_q(i) \subseteq D$, which yields $i \in D$, and finally, $i \in k$ ($D = k \cap \mathbb{F}_q(\zeta)$). Thus, $\overline{k} = k$, a contradiction. (Note: i is a fourth root of unity; so, as $\overline{k} \neq k$, we have $\text{char}(k) = q \neq 2$). Therefore, $D = \mathbb{F}_q$.

Now,

$$\mathbb{Z}/2\mathbb{Z} = \mathcal{G}(\overline{k}/k) = \mathcal{G}(\mathbb{F}_q(\zeta)/\mathbb{F}_q),$$

a group of order 2^s. If we let r tend to ∞, then s tends to ∞, a contradiction. Therefore, $\text{char}(k) = 0$. □

Finally, here the proof of Theorem 4.73:

Proof of Artin's Irreducibility Criterion. Assume at first that we know the result for n a prime power–here is how to prove the general case: Use induction on the number of primes dividing n. If $n = p^r m$ with $(p, m) = 1$, we may assume p is odd. Now, $X^m - a$ is irreducible by our induction hypothesis; let $\alpha_1, \ldots, \alpha_m$ be its roots. Then,

$$X^m - a = \prod_{j=1}^{m} (X - \alpha_j)$$

and

$$X^n - a = (X^{p^r})^m - a = \prod_{j=1}^{m} (X^{p^r} - \alpha_j).$$

Suppose for some j that α_j is a pth power in $k(\alpha_j)$. Now $X^m - a$ is irreducible so its Galois group acts transitively on its roots. Therefore, each α_i is $\sigma(\alpha_j)$ for some σ and so each of the α_i is a pth power in $k(\alpha_i)$. There exist $\beta_i \in k(\alpha_i)$ with $\beta_i^p = \alpha_i$ for $i = 1, \ldots, m$. We find that

$$\mathcal{N}_{k(\alpha_i)/k}(\alpha_i) = \mathcal{N}_{k(\alpha_i)/k}(\beta_i)^p.$$

But,

$$\prod_{j=1}^{m} \alpha_j = (-1)^{m+1} a = \mathcal{N}_{k(\alpha_i)/k}(\alpha_i) = \mathcal{N}_{k(\alpha_i)/k}(\beta_i)^p.$$

If m is odd, this gives $a \in k^{*p}$, contrary our assumptions. If m is even, then $a = -(\mathcal{N}_{k(\alpha_i)/k}(\beta_i)^p)$. But, p is odd so

$$a = (-\mathcal{N}_{k(\alpha_i)/k}(\beta_i))^p,$$

again contrary to hypothesis. We conclude that none of the α_i are pth powers in the field $k(\alpha_i)$. By the one prime case, the polynomials $X^{p^r} - \alpha_i$ are irreducible for $i = 1, \ldots, m$ and all r.

Let ξ be a root of $X^n - a$. Then, ξ satisfies $X^{p^r} - \alpha_j = 0$ for at least one j. According to the irreducibility of $X^{p^r} - \alpha_j$, we find that $[k(\xi)\colon k(\alpha_j)] = p^r$ (of course $k(\alpha_j) \subseteq k(\xi)$). However, $[k(\alpha_j)\colon k] = m$ by the induction hypothesis and so $[k(\xi)\colon k] = n$. But this means the minimal polynomial for ξ has degree n and ξ *is* a root of $X^n - a$; so, $X^n - a$ is irreducible.

We've achieved a reduction to the heart of the matter, the one prime case. Here, $n = p^r$ and when $p = \text{char}(k)$ we already know the result. Therefore, we may and do assume $p \neq \text{char}(k)$. Now use induction on r. Say $r = 1$, adjoin the pth roots of 1 to k—call ζ a primitive pth root of 1:

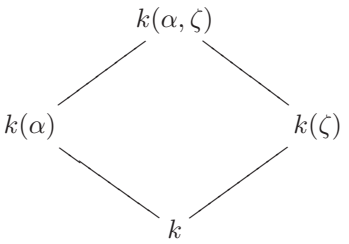

Here, α is a root of $X^p - a$. Were $[k(\alpha)\colon k] \neq p$, i.e., were $X^p - a$ reducible, we would have $[k(\alpha,\zeta)\colon k(\zeta)] < p$. Now over $k(\zeta)$, the Galois group of $X^p - a$ is cyclic of order p or trivial according as $a \notin k(\zeta)^p$ or $a \in k(\zeta)^p$. We then would have $a \in k(\zeta)^p$; hence $\alpha \in k(\zeta)$. We know that $[k(\zeta)\colon k] = r \leq p-1$ and so, $(r,p) = 1$. Write $1 = sr + tp$, for some s, t. Now,

$$a^r = \mathcal{N}_{k(\zeta)/k}(a) = \mathcal{N}_{k(\zeta)/k}(\alpha)^p.$$

But,

$$a = a^{sr+tq} = \mathcal{N}_{k(\zeta)/k}(\alpha)^{ps} a^{pt} = (\mathcal{N}_{k(\zeta)/k}(\alpha)^p \cdot a^t)^p \in k^{*p},$$

a contradiction. We conclude $X^p - a$ is irreducible.

Induction Step. Consider $X^{p^r} - a$ and assume p is *odd*. Write α for a root of $X^p - a$, and further write $\prod_{j=1}^{p}(X - \alpha_j) = X^p - a$, with $\alpha = \alpha_1$. Now α is not a pth power in $k(\alpha)$. For, if it were β^p, then

$$a = (-1)^{p+1}a = \mathcal{N}_{k(\zeta)/k}(\alpha) = \mathcal{N}_{k(\zeta)/k}(\beta)^p \qquad (p \text{ is odd})$$

contrary to the hypothesis that a is not a pth power. Again by transitivity of the Galois group on the α_j, no α_j is a pth power in $k(\alpha_j)$, and therefore, by induction all the polynomials

$$X^{p^{r-1}} - \alpha_j, \quad j = 1, 2, \ldots, p$$

are irreducible (over $k(\alpha_j)$). If ξ is a root of $X^{p^r} - a$, then ξ is a root of $X^{p^{r-1}} - \alpha_j$ for some j, and as before

$$[k(\xi)\colon k(\alpha_j)] = p^{r-1} \quad \text{and} \quad [k(\alpha_j)\colon k] = p;$$

so, $[k(\xi)\colon k] = p^r$. We conclude $X^{p^r} - a$ is indeed irreducible.

Finally, we have the case $p = 2$. We know $X^2 - a$ is irreducible, we must prove $X^{2^r} - a$ is irreducible. As usual and with familiar notation, we have

$$X^{2^r} - a = \prod_{j=1}^{2}(X^{2^{r-1}} - \alpha_j); \quad \alpha = \alpha_1; \quad \alpha^2 = a.$$

So, if $X^{2^{r-1}} - \alpha_j$ is irreducible for $j = 1, 2$, the usual degree argument will show $X^{2^r} - a$ is irreducible. The only way $X^{2^{r-1}} - \alpha_j$ will be reducible, by the induction hypothesis, is if $\alpha_j \in k(\alpha_j)^{*2}$ or $\alpha_j \in -4k(\alpha_j)^{*4}$. We will show each of these is untenable.

(1) Say $\alpha \in k(\alpha)^{*2}$; so, the same is true of α_2. Now,

$$-a = \mathcal{N}_{k(\alpha)/k}(\alpha) = \mathcal{N}_{k(\alpha)/k}(\beta)^2 = b^2$$

yet $a \notin k^{*2}$, so $(-1) \notin k^{*2}$. Hence, $i \notin k$ and we factor $X^{2^r} - a$ over $k(i)$:

$$X^{2^r} - a = X^{2^r} + b^2 = (X^{2^{r-1}} + ib)(X^{2^{r-1}} - ib).$$

If, on the right hand side, one of the factors is reducible, the induction hypothesis shows ib (or $-ib$) $\in k(i)^{*2}$ or ib (or $-ib$) $\in -4k(i)^{*4}$. Since -4 is square in $k(i)$, the cases $ib \in -4k(i)^{*4}$ or $-ib \in -4k(i)^{*4}$ reduce respectively to $ib \in k(i)^{*2}$ or -$ib \in k(i)^{*2}$. But -1 is also a square, so these two cases are just the one case: $ib \in k(i)^{*2}$.

Write $ib = (\gamma + i\delta)^2$, with $\gamma, \delta \in k$. Then $\gamma^2 = \delta^2$ and $b = 2\gamma\delta$. However, $\gamma\delta = \pm\gamma^2$, and so $b^2 = 4\gamma^4$. But then, $a = -b^2 \in -4k^{*4}$, a contradiction. We are left with

(2) $\alpha \in -4k(\alpha)^{*4}$. Again,

$$-a = \mathcal{N}_{k(\alpha)/k}(\alpha) = \mathcal{N}_{k(\alpha)/k}(-4)\mathcal{N}_{k(\alpha)/k}(\beta)^4.$$

Since $\mathcal{N}_{k(\alpha)/k}(-4) = (-4)^2 = 16$, we deduce $-a$ is a square and we've assumed a is not a square. As above, $i \notin k$ and we now repeat the argument of (1) to finish the proof. □

4.11 Algebraic Closures; Steinitz's Theory of Fields

In the twentieth century, E. Steinitz examined the theory of fields, especially transcendental extensions. He had at his disposal the then new technique of transfinite induction which he used in the form of Zermelo's well-ordering principle. Of course, the latter is equivalent to Zorn's Lemma or the Axiom of Choice. Here, we'll examine Steinitz's results both in the purely algebraic case (existence of an algebraic closure) and in the general case (transcendence bases).

Recall that in Remark (4) at the close of Section 4.2 we made the following definition (but informally):

Definition 4.15 A field, K, is *algebraically closed* (AC) iff for every $f \in K[X]$, there exists a $\theta \in K$, so that $f(\theta) = 0$.

We also defined an *algebraic closure* of the field k as a field, K, which was itself AC and moreover was algebraic over k. Here, we'll prove the existence of an algebraic closure for each field, k, and its (essential) uniqueness. First, for technical agility we'll need equivalent forms of the condition (AC):

Proposition 4.75 *For a field, K, the following conditions are equivalent:*

(1) K has AC

(2) For every $f \in K[X]$, all the roots of f (in any extension of K) are already in K

(3) Every polynomial $f \in K[X]$ factors into linear factors in $K[X]$

(4) The only irreducible K-polynomials are the linear ones

(5) If L/K is algebraic, then $L = K$ (so, K is algebraically closed in any of its overfields)

(6) If k is a subfield of K for which K/k is algebraic, then for any algebraic extension, L, of k, there exists a k-monomorphism $L \longrightarrow K$.

(7) If k is a subfield of K for which K/k is algebraic and if $\widetilde{k}$ is a field isomorphic to k, via an isomorphism, φ, then for any algebraic extension, $\widetilde{L}$ or $\widetilde{k}$, there exists a monomorphism $\Phi\colon \widetilde{L} \to K$ extending φ.

The proofs of the equivalences (1)–(7) are trivial (DX); in (6) and (7) one makes use of the extension lemma.

Remark: An algebraicaly closed field is always infinite. For, were it finite and $\theta_1, \ldots, \theta_n$ a listing of its elements, then $f(T) = 1 + \prod_{j=1}^{n}(T - \theta_j)$ would be a polynomial with no root in our field.

Now for the proof of the existence of algebraic closures, we need a very basic existence theorem.

Theorem 4.76 *(Basic Existence Theorem) Suppose k is a field and K_λ ($\lambda \in \Lambda$) is a family of overfields of k. Then, there exists a field extension K/k so that for every $\lambda \in \Lambda$ we have a k-monomorphism $\varphi_\lambda\colon K_\lambda \to K$. That is, K contains a k-isomorphic copy of each field K_λ. Moreover, we may even choose K so that it is generated by all the subfields $\varphi_\lambda(K_\lambda)$.*

Proof. The proof is very simple using our techniques. We form the commutative ring $A = \bigotimes_{\lambda\in\Lambda} K_\lambda$. Of course,

$$A = \varinjlim_{S\in\mathcal{L}(\Lambda)} \Big(\bigotimes_{\lambda\in\Lambda} K_\lambda\Big),$$

where $\mathcal{L}(\Lambda)$ is the family of *finite* subsets of Λ. The ring A is a k-algebra (the tensor products are taken over k) and we embed k in A as usual *via* $\alpha \mapsto \alpha \cdot (1 \otimes 1 \otimes \cdots \otimes 1 \otimes \cdots)$. Choose any maximal ideal, $\mathfrak{M}$, of A and write $K = A/\mathfrak{M}$. Of course, K is a field extension of k and as each K_λ has a k-algebra homomorphism to K ($K_\alpha \longrightarrow \bigotimes_\mu K_\mu = A \longrightarrow A/\mathfrak{M} = K$) taking 1 to 1, we see that each K_λ is embedded in K *via* this homomorphism. Now the images in A of the K_λ generate A; so, their images in K generate K. $\square$

Theorem 4.77 *(Steinitz) If k is a field, then k possesses an algebraic closure, Ω. If Ω and $\widetilde{\Omega}$ are two algebraic closures of k, then there exists a (non-canonical) k-isomorphism $\Omega \xrightarrow{\sim} \widetilde{\Omega}$. The set $\mathrm{Isom}_k(\Omega, \widetilde{\Omega})$ is in one-to-one correspondence with $\mathcal{G}(\Omega/k)$.*

Proof. We wish to use the Basic Existence Theorem, so the only problem is to find a good way of parametrizing all finite extensions of k. Here, a better idea is to parametrize all the *finitely generated* extensions of k. Take $A = k[X_j]_{j=0}^{\infty}$, the polynomial algebra on $\aleph_0$ independent transcendentals over k. View, for each $n \geq 0$, the finitely generated polynomial rings $k[X_0, \ldots, X_n]$ as a subring of A. In each ring $k[X_0, \ldots, X_n]$ we have the family of its maximal ideals, $\mathfrak{M}$. Write $K(n, \mathfrak{M})$ for the field $k[X_0, \ldots, X_n]/\mathfrak{M}$ and consider the collection of all these $K(n, \mathfrak{M})$.[6] By the Basic Existence Theorem, there is field, L, over k containing a k-isomorphic copy of each $K(n, \mathfrak{M})$. But each finite degree extension, M, of k is k-isomorphic to at least one $K(n, \mathfrak{M})$; and so, each finite degree extension is "contained" in L. Write

$$\begin{aligned} \Omega = L_{\mathrm{alg}} &= \{\xi \in L \mid \xi \text{ is algebraic over } k\} \\ &= \text{ algebraic closure of } k \text{ in } L. \end{aligned}$$

By construction, Ω is algebraic over k; by choice of L, each finite degree extension, M, of k is k-isomorphic to a $K(n, \mathfrak{M})$ so the latter is algebraic over k; hence, in Ω. And now, (the obvious modification of) Proposition 4.75 #(6) shows Ω is algebraically closed.

Having proved existence, we now investigate uniqueness. Say $\widetilde{\Omega}$ is another algebraic closure of k. Now for Ω and $\widetilde{\Omega}$ we have

$$\Omega = \varinjlim_{K/k \text{ finite}, K \subseteq \Omega} K \tag{$\dagger$}$$

$$\widetilde{\Omega} = \varinjlim_{\widetilde{K}/k \text{ finite}, \widetilde{K} \subseteq \widetilde{\Omega}} \widetilde{K}. \tag{$\ddagger$}$$

Since Ω is algebraically closed, for each such $\widetilde{K}/k$ we get a k-injection $\widetilde{K} \longrightarrow \Omega$. We may assume each such $\widetilde{K}$ is normal over k and choose a maximal chain of such $\widetilde{K}$. Then, twisting if necessary by the $\mathcal{G}(\widetilde{K}/k)$, we obtain a consistent family of k-injections of these $\widetilde{K}$ into Ω. By ($\ddagger$), there results the k-injection $\widetilde{\Omega} \longrightarrow \Omega$. But the image of $\widetilde{\Omega}$ is algebraically closed and Ω is algebraic over it. We deduce from Proposition 4.75 (5) that $\Omega =$ image of $\widetilde{\Omega}$.

Lastly, if φ and ψ are two k-isomorphisms from Ω to $\widetilde{\Omega}$, the map $\psi^{-1} \circ \varphi$ is in $\mathcal{G}(\Omega/k)$. Hence, the k-isomorphisms $\varphi \circ \sigma$ run over all k-isomorphisms $\Omega \longrightarrow \widetilde{\Omega}$ whenever φ is one such and σ runs over $\mathcal{G}(\Omega/k)$. □

There remains the general case of a field extension k/k. The important concept here is the notion of *transcendence basis*.

Definition 4.16 A subset, S, of a field extension, K/k, is a *transcendence basis* for K/k iff

(1) S is algebraically independent over k and

(2) K is algebraic over $k(S)$.

We need some technique in handling algebraically independent elements. The most useful technical observation is the following:

[6] For readers with a foundational mind, note: In the first place, the pairs $(n, \mathfrak{M})$ are elements of the set $\mathbb{N} \coprod \mathcal{P}(A)$, where $\mathcal{P}(A)$ is the power set of A; so, our indexing is done by a set. Next, each field, $K(n, \mathfrak{M})$, is itself in $\mathcal{P}(A)$; so, the whole collection is perfectly valid from the point of view of set theory.

Proposition 4.78 *Suppose that K/k is a field extension and A and B are subsets of K. Then the three conditions below are mutually equivalent:*

(1) $A \cap B = \emptyset$ and $A \cup B$ is algebraically independent over k

(2) A is algebraically independent over k and B is algebraically independent over $k(A)$

(3) Same statement as (2) with A and B interchanged.

Proof. By symmetry, (2) $\iff$ (3); all that remains is to prove (1) $\iff$ (2).

(1) $\implies$ (2). As $A \subseteq A \cup B$ and the latter is algebraically independent over k, we find that A is algebraically independent over k. If B is algebraically dependent over $k(A)$, there are elements $b_1, \ldots, b_t$ and a nonzero polynomial $f(T_1, \ldots, T_t) \in k(A)[T_1, \ldots, T_t]$ with $f(b_1, \ldots, b_t) = 0$. But, the coefficients may be chosen from $k[A]$ and involve only finitely many elements $a_1, \ldots, a_s$ from A. Then, $f(T_1, \ldots, T_t)$ is actually a nonzero polynomial of the form $\widetilde{f}(a_1, \ldots, a_s, T_1, \ldots, T_t)$, and $\widetilde{f}$ is a polynomial over k in variables $U_1, \ldots, U_s, T_1, \ldots, T_t$. It is satisfied by $\{a_1, \ldots, a_s, b_1, \ldots, b_t\} \subseteq A \cup B$ contradicting (1).

(2) $\implies$ (1). No element, ξ, can be in $A \cap B$, else the polynomial $T - \xi \in k(A)[T]$ is satisfied by $\xi \in B$ contradicting (2). We need only show each finite subset of $A \cup B$ is algebraically independent and, of course, this is immediate if that finite subset is in A or B. So, we may assume that our subset is $a_1, \ldots, a_s, b_1, \ldots, b_t$. Any polynomial $f(U_1, \ldots, U_s, T_1, \ldots, T_t) \in k[U_1, \ldots, U_s, T_1, \ldots, T_t]$ which vanishes on $a_1, \ldots, a_s, b_1, \ldots, b_t$ gives a polynomial

$$f(a_1, \ldots, a_s, T_1, \ldots, T_t) \in k(A)[T_1, \ldots, T_t]$$

which vanishes on $b_1, \ldots, b_t$. By (2), all the coefficients of $f(a_1, \ldots, a_s, T_1, \ldots, T_t)$ have to vanish. By (2), again, these coefficients which are just different polynomials $g_j(U_1, \ldots, U_s)$ (coeffs in k) must be zero *as polynomials*. Therefore, our original polynomial $f(U_1, \ldots, U_s, T_1, \ldots, T_t)$ is identically zero. This proves (1). □

We derive many corollaries from Proposition 4.78.

Corollary 4.79 *Let K/k be a field extension and A be a subset of K. Then, A is algebraically independent over k iff for all $\xi \in A$, the element ξ is transcendental over $k(A - \{\xi\})$.*

Proof. By taking $A - \{\xi\}$ and $\{\xi\}$ as the two subsets of Proposition 4.78, we see that ($\Longrightarrow$) is proved. To prove ($\Longleftarrow$), take a finite subset of A, say $a_1, \ldots, a_t$, and suppose it is algebraically dependent over k. We may assume no smaller subset of $a_1, \ldots, a_t$ is dependent by passing to that smaller subset. Apply Proposition 4.78 to the sets $\{a_1\}$ and $\{a_2, \ldots, a_t\}$. Since $a_1, \ldots, a_t$ is **not** independent, it follows that a_1 is not independent over $k(a_2, \ldots, a_t)$. Hence, a_1 is not independent over the bigger field $k(A - \{a_1\})$. This contradicts our hypothesis when $\xi = a_1$. □

Corollary 4.80 *If K/k be a field extension and A is an algebraically independent subset (over k) of K, and if $\xi \in K$ has the property that ξ is transcendental over $k(A)$, then $A \cup \{\xi\}$ is again algebraically independent over k.*

Proof. This is immediate either from Proposition 4.78 or Corollary 4.79.

Corollary 4.81 *Suppose K/k is a field extension and A is an algebraically independent subset of K. A necessary and sufficient condition that A be a transcendence basis for K/k is that A be a maximal element (under partial ordering by set inclusion) among the algebraically independent subsets of K.*

Proof. If A is a transcendence basis for L/k yet is not maximal, there is an independent set, B, of K and $B > A$. In Proposition 4.78, let $B - a$ and A be the two sets, the $K \subseteq k(A)(B - A)$ and $K/K(A)$ is algebraic. So, $B - A$ is not algebraically independent over $k(A)$ contradicting Proposition 4.78.

Conversely, if A is maximal among algebraically independent sets and $\xi \in K$ but not in $K(A)$, then ξ cannot be transcendental over $k(A)$ by Proposition 4.78 (set $B = \{\xi\}$, $A = A$). So, ξ is algebraic over $K(A)$; that is, $K/k(A)$ is algebraic. □

Theorem 4.82 *(Steinitz) If K/k is a field extension and if $S \subseteq T$ are two subsets of K so that*

(a) K is algebraic over $k(S)$ and

(b) T is algebraically independent over k, then there exists a transcendence basis, B, for K/k with $T \subseteq B \subseteq S$. In particular, every field extension possesses a transcendence basis.

Proof. We let $\mathcal{S}$ denote the collection of subsets of S which both contain T and are algebraically independent over k. Of course, as $T \in \mathcal{S}$, we have $\mathcal{S} \neq \emptyset$. Partially order $\mathcal{S}$ by set-theoretic inclusion and note that $\mathcal{S}$ is inductive. Let B be a maximal element of $\mathcal{S}$, it exists by Zorn's Lemma. Consider the extension $k(S)/k(B)$. We know if each element of S is algebraic over $k(S)$, then $k(S)$ will be algebraic over $k(B)$. But by the maximality of B and Proposition 4.78 (or Corollary 4.80), we see that every element of S is indeed algebraic over $k(B)$. Thus, from the facts that K is algebraic over $k(S)$ and $k(S)$ is algebraic over $k(B)$, we find K is algebraic over $k(B)$.

Upon taking $S = K$ and $T = \emptyset$, we deduce each field extension has a transcendence basis. □

Doubtless you will have noticed an analogy between the familiar theory of linear dependence and independence for vector spaces and our theory of algebraic independence and indepnedence for field extensions. For example, Proposition 4.78 can be translated into the linear case. Steinitz noticed this explicitly and transformed the analogy into an axiomatic treatment of both cases simulataneously. In the linear case, the notion of Span is a crucial ingredient and Steinitz generalized this by setting

$$\Sigma(A) = \{\xi \in K \mid \xi \text{ is algebraic over } k(A) \qquad (*)$$

for A a subset of K and K/k a field extension. Of course we can then write: A is a transcendence basis for K/k iff A is algebraically independent over k and $K = \Sigma(A)$. The axioms for the Σ operation are the dictated by the linear case:

(1) $A \subseteq \Sigma(A)$,

(2) If $A \subseteq B$, then $\Sigma(A) \subseteq \Sigma(B)$.

(3) $\Sigma(\Sigma(A)) = \Sigma(A)$.

(4) If $\xi \in \Sigma(A)$, then there is a finite subset, $\widetilde{A}$, of A so that $\xi \in \Sigma(\widetilde{A})$.

(5) If $\eta \in \Sigma(A \cup \{\xi\})$ but $\eta \notin \Sigma(A)$, then $\xi \in \Sigma(A \cup \{\eta\})$.

Conditions (1)–(3) are obvious both in the linear case (when $\Sigma(A) = \text{Span}(A)$) and in the algebraic case (when $\Sigma(A)$ is as above). However, (4) and (5) deserve some comment. Property (4) makes the formation of $\Sigma(A)$ a property "of finite character", and allows the application of Zorn's Lemma in proofs of statements about $\Sigma(A)$ or independence. Property (5) is called the *Steinitz Exchange Lemma*—it is well-known in the linear case. Here it is in the algebraic case:

Proposition 4.83 *(Steinitz Exchange Lemma) For a field extension, K/k, a subset $A \subseteq K$ and element ξ, η of K we have*

If $\eta \in \Sigma(A \cup \{\xi\})$ but $\eta \notin \Sigma(A)$, then $\xi \in \Sigma(A \cup \{\eta\})$.

Here, $\Sigma(A)$ is as above in $(*)$.

Proof. In $k(A)$ we can choose a transcendence basis (over k) contained in A by Theorem 4.83. As $k(A)$ is algebraic over $k(B)$, it is algebraic over $k(B \cup \{\xi\})$. Now, ξ is algebraic over $k(B \cup \{\xi\})$; so, $k(A \cup \{\xi\})$ is algebraic over $k(B \cup \{\xi\})$ and therefore $\eta \in k(B \cup \{\xi\})$. If the exchange lemma were valid when A was independent, we would deduce $\xi \in \Sigma(B \cup \{\eta\}) \subseteq \Sigma(A \cup \{\eta\})$.

This achieves a reduction to the case where A is algbraically independent. The silly case $\xi = \eta$ is a tautology and so we have $\xi \neq \eta$ and $A \cup \{\xi, \eta\}$ is algebraically dependent. But then, Proposition 4.78 applied to the sets $A \cup \{\eta\}$, $\{\xi\}|$ shows that $\xi \in \Sigma(A \cup \{\eta\})$, as desired. □

Clearly, the Exchange Lemma is susceptible of generalizations. But one must be careful. "Obvious" generalizations may be false. For example, the statement: If A, B, C are subsets of K (an extension of k) and if $C \subseteq \Sigma(A \cup B)$ but $C \not\subseteq \Sigma(A)$, then $B \subseteq \Sigma(A \cup C)$ is *false*. Indeed, even the weaker statememnt (because the hypotheses are stronger): If $C \subseteq \Sigma(A \cup B)$ but *no element of* C *is in* $\Sigma(A)$, then $B \subseteq \Sigma(A \cup C)$ is *false*. To see why the latter is false, just take $A = \emptyset$ and $K = k(X, Y, \sqrt{X})$, where X and Y are algebraically independent over k. Set $B = \{X, Y\}$ and $C = \{\sqrt{X}\}$.

Proposition 4.84 *(General Steinitz Exchange Lemma) Suppose K/k is a field extension and $A, B, C \subseteq K$. Assume that $C \subseteq \Sigma(A \cup B)$ but $C \not\subseteq \Sigma(A)$. Then, there exists a subset, B', of B with properties*

(1) $B \subseteq \Sigma(A \cup C \cup B')$.

(2) $B \neq B'$.

(3) $B' \cap C = \emptyset$.

Before proving this form of the exchange lemma, we should remark that:

(a) The hypotheses are those of the previous strong (but false) statement—the conclusion is weaker: we need B'. In the example where $A = \emptyset$, $C = \{\sqrt{X}\}$, $B = \{X, Y\}$, it is clear that $B' = \{Y\}$.

(b) The name come from the fact that B'' $(= B - B')$ and C have been exchanged. That is, we conclude $B'' \subseteq \Sigma(A \cup B' \cup C)$ from the hypotheses $C \subseteq \Sigma(A \cup B' \cup B'')$ and $C \not\subseteq \Sigma(A)$.

Proof. Here, the notation Σ refers to algebraic dependence *over* k. Let $\widetilde{A}$ be a maximal algebraically independent subset of A, so that $\Sigma(\widetilde{A}) = \Sigma(A)$. Write $\widetilde{C}$ for a subset of C maximal with respect to algebraic independence *over* $k(\widetilde{A})$. Because $C \not\subseteq \Sigma(\widetilde{A})$, we see that $\widetilde{C} \neq \emptyset$ and that $\widetilde{A} \cup \widetilde{C}$ is algebraically independent over k by Proposition 4.78. Now $C \subseteq \Sigma(\widetilde{A} \cup \widetilde{C})$ and $A \subseteq \Sigma(\widetilde{A}) \subseteq \Sigma(\widetilde{A} \cup \widetilde{C})$. We find that

$$\Sigma(A \cup B) = \Sigma(\widetilde{A} \cup \widetilde{C}).$$

Write $T = \widetilde{A} \cup \widetilde{C} \cup B$. Now, $\Sigma(A \cup B) = \Sigma(\widetilde{A} \cup B)$ and by hypothesis we find that $C \subseteq \Sigma(\widetilde{A} \cup B)$. Therefore, $\Sigma(T) = \Sigma(A \cup B)$; call this field $\widetilde{K}$. In it, we have $T \supseteq \widetilde{A} \cup \widetilde{C}$, the former set generates and the latter is algebraically independent. By the existence of transcendence bases (Theorem 4.83), there is a transcendence basis for $\widetilde{K}/k$, call it S, so that

$$T \supseteq S \supseteq \widetilde{A} \cup \widetilde{C}.$$

We set $B' = S - (\widetilde{A} \cup \widetilde{C}) \subseteq B$. Of course, $B' \cap \widetilde{C} = \emptyset$. We know

$$\Sigma(T) = \Sigma(S) = \Sigma(\widetilde{A} \cup \widetilde{C} \cup B')$$

and $B \subseteq \Sigma(T)$; so, conclusion (1) is proved. Were $B' = B$, we'd have $S = \widetilde{A} \cup \widetilde{C} \cup B$. Now $\widetilde{C} \subseteq C$ and by hypothesis $C \subseteq \Sigma(A \cup B) = \Sigma(\widetilde{A} \cup B)$. As S is algebraically independent, we have a contradiction of Proposition 4.78; this proves (2). Finally, if $\xi \in B' \cap C$, then $\xi \in C \subseteq \Sigma(\widetilde{A} \cup \widetilde{C})$ implies that the subset of $B' \cup \widetilde{A} \cup \widetilde{C} = S$ consisting of $\{\xi\} \cup \widetilde{A} \cup \widetilde{C}$ is dependent; contradiction on how we chose S. □

The main use of the standard exchange lemma is to prove that transcendence bases have the same cardinality. Here's how the finite case goes.

Theorem 4.85 *Suppose K/k is a field and S is a finite subset of K while T is any subset of K. Assume that $\#(T) > \#(S)$ and $T \subseteq \Sigma(S)$. Then T is algebraically dependent. In particular, if K/k has a finite transcendence basis, then all transcendence bases of K/k are finite with the same cardinality.*

Proof. First, replace K by $\Sigma(S)$, second replace S by a transcendence basis (for $K = \Sigma(S)$) which is a subset of S. Therefore, we may assume S is a transcendence basis for K/k. If $\#(T) = \infty$, then then replace T by any finite subset with $\#(T) > \#(S)$; so we may assume T is finite, too. Now suppose the result is false and choose a counter-example pair S, T so that $\#(S)+\#(T)$ is minimal. Of course, in this case $\#(T) = \#(T)+1$, else we could reduce the sum by choosing a subset of T with $\#(T) = \#(S) + 1$.

Our situation is now that

(a) S and T are finite algebraically independent sets

(b) $T \subseteq \Sigma(S)$

(c) $\#(S) = n$, $\#(T) = n + 1$

(d) n is minimal among integers having (a), (b), (c).

Label the elements of S as $s_1, \ldots, s_n$ but refrain from labelling T as yet. Consider $S - \{s_1\}$. There must be some $t \in T$ so that $t \notin \Sigma(S - \{s_1\})$, else $T \subseteq \Sigma(S - \{s_1\})$ and (d) would show T dependent contradicting (a). Call this element t_1. Note that $\{s_2, \ldots, s_n, t_1\}$ is an independent set at $t_1 \notin \Sigma(s_2, \ldots, s_n)$. Since $t_1 \in \Sigma((S-\{s_1\})\cup\{s_1\})$, the standard exchange lemma (Proposition 4.83) shows that $s_1 \in \Sigma(s_2, \ldots, s_n, t_1)$. All the other elements of S lie in $\Sigma(s_2, \ldots, s_n, t_1)$ so $T - \{t_1\}$ is certainly in $\Sigma(s_2, \ldots, s_n, t_1)$. However, $T - \{t_1\}$ cannot be contained in $\Sigma(s_3, \ldots, s_n, t_1)$. For if it were, the sets $\{s_3, \ldots, s_n, t_1\}$, $T - \{t_1\}$ would satisfy (a) and (b), their cardinalities would be $n - 1$ and n respectively and (d) would be contradicted.

Since $T - \{t_1\} \not\subseteq \Sigma(s_3, \ldots, s_n, t_1)$, there is an element $t_2 \in T - \{t_1\}$ with $t_2 \notin \Sigma(s_3, \ldots, s_n, t_1)$. This means $\{s_3, s_4, \ldots, s_n, t_1, t_2\}$ is an independent set and allows the exchange lemma to be applied once more to $\eta = t_2$, $\xi = s_2$, and $\{s_3, \ldots, s_n, t_1\}$. We conclude that $s_2 \in \Sigma(s_3, \ldots, s_n, t_1, t_2)$ and so all of T (thus also $T - \{t_1, t_2\}$) is in $\Sigma(s_3, \ldots, s_n, t_1, t_2)$. It is clear how to continue the process and equally clear what is happening: We are systematically replacing the elements $s_1, s_2, \ldots$ of S by elements $t_1, t_2, \ldots$ from T. In the end, we find $T - \{t_1, \ldots, t_n\} \subseteq \Sigma(t_1, \ldots, t_n)$; but, $\#(T) = n + 1$, so $t_{n+1} \in \Sigma(t_1, \ldots, t_n)$–our final contradiction (on (a)). As (a)–(d) are untenable, no counter-example exists.

To prove if K/k has finite transcendence basis, all transcendence bases have the same cardinality, we choose a transcendence basis, S, of minimal (so, finite) cardinality and any other transcendence basis, T. Of course, $\#(T) \geq \#(S)$. If $\#(T) > \#(S)$, then $T \subseteq \Sigma(S)$ immediately implies from the above that T is dependent, which is not true. Thus, $\#(T) \leq \#(S)$, and we are done. □

If the reader will go through the argument, he will see we have used only Steinitz's rules (1)—(5) on Σ. Thus, the argument works in the linear case—though a direct argument is simpler. In carrying this out, one sees that Corollary 4.79 gives a way of defining algebraic independence solely in terms of the Σ operation. Namely, A is algebraically independent iff for every $\xi \in A$, the element ξ is not in $\Sigma(A - \{\xi\})$.

We can now handle the infinite case.

Theorem 4.86 *For every field extension K/k, any two transcendence bases have the same cardinality.*

Proof. Is S and T are given transcendence bases for K/k, then, by Theorem 4.85, the sets S and T are simultaneously finite or infinite. Of course, the only case of concern is when S and T are infinite.

I claim two statements, which taken together will quickly finish the proof.

(I) *For each $\xi \in K$*, there exists a *unique finite* subset of S, call it $S(\xi)$, characterized by

  (a) $\xi \in \Sigma(S(\xi))$ and

  (b) If $\widetilde{S} \subseteq S$ and $\xi \in \Sigma(\widetilde{S})$, then $S(\xi) \subseteq \widetilde{S}$.

(II) For ξ and η in T, if $\xi \neq \eta$, then $S(\xi) \neq S(\eta)$.

Suppose we assume (I) and (II) and write $\mathcal{FP}(S)$ for the collection of all *finite* subsets of S. Then the map $\xi \mapsto S(\xi)$ is, by (I) and (II), a well defined injection of T to $\mathcal{FP}(S)$. We find that $\#(T) \leq \#(\mathcal{FP}(S))$ and we know $\#(S) = \#(\mathcal{FP}(S))$ because $\#(S)$ is infinite. Thus, $\#(T) \leq \#(S)$; by symmetry, $\#(S) \leq \#(T)$. Then, the Cantor-Schröder-Bernstein Theorem yields $\#(S) = \#(T)$.

Both (I) and (II) are consequences of the exchange lemma. For (I), choose $\xi \in K$. If ξ is algebraic over k, the set $S(\xi) = \emptyset$ satisfies (a) and (b); we may assume ξ is transcendental over k. There is a finite subset, $\{s_1, \ldots, s_n\}$, of S so that $\xi \in \Sigma(s_1, \ldots, s_n)$. We may assume no smaller subset of $\{s_1, \ldots, s_n\}$ has ξ in the Σ formed from it. Suppose $\{\sigma_1, \ldots, \sigma_q\}$ is another subset of S and $\xi \in \Sigma(\sigma_1, \ldots, \sigma_q)$. Choose any s_j, apply the exchange lemma to ξ, s_j and $S - \{s_j\}$. We find that $s_j \in \Sigma(s_1, \ldots \widehat{s}_j, \ldots, s_n, \xi)$. Now, $\xi \in \Sigma(\sigma_1, \ldots, \sigma_q)$, therefore

$$s_j \in \Sigma(s_1, \ldots \widehat{s}_j, \ldots, s_n, \sigma_1, \ldots, \sigma_q).$$

The elements s_j and σ_l are in the independent set S therefore s_j must be one of the $\sigma_1, \ldots, \sigma_q$. Since s_j is arbitrary in $\{s_1, \ldots, s_n\}$ we get $\{s_1, \ldots, s_n\} \subseteq \{\sigma_1, \ldots, \sigma_q\}$ and so $S(\xi) = \{s_1, \ldots, s_n\}$ has (a) and (b).

To prove (II), first note that if $\xi \in S$, then $S = \{\xi\}$. Pick $\xi, \eta \in T$ and assume $S(\xi) = S(\eta)$. Write $\{s_1, \ldots, s_n\}$ for the listing of the elements of $S(\xi)$. A standard application of the exchange lemma shows $s_1 \in \Sigma(s_2, \ldots, s_n, \xi)$. Therefore, $S(\xi) \subseteq \Sigma(s_2, \ldots, s_n, \xi)$. It follows, as $S(\xi) = S(\eta)$, that $\eta \in \Sigma(s_2, \ldots, s_n, \xi)$. By Claim (I) property (b), we find

$$\{s_1, \ldots, s_n\} = S(\eta) \subseteq \{s_2, \ldots, s_n, \xi\}.$$

Hence, $\xi = s_1$. By symmetry, $\eta = s_1$, too; and so, $\xi = \eta$ (or $\xi \in S$ and hence so is η; therefore $\{\xi\} = S(\xi) = S(\eta) = \{\eta\}$). We are done. □

Definition 4.17 The common cardinal number of all the transcendence bases for K/k is the *transcendence degree of* K/k. It is denoted $\text{tr.d.}_k(K)$. The field K is *purely transcendental over* k iff $K = k(S)$ where S is a transcendence base for K/k.

To finish up this section, we have only to discuss the notion of separability for general field extensions (i.e., not necessarily algebraic). For this, we essentially make Mac Lane I into a definition:

Definition 4.18 A field extension, K/k, is *separable* iff either $\text{char}(k) = 0$ or when $\text{char}(k) = p > 0$, then the natural map

$$k^{1/p} \otimes_k K \longrightarrow K^{1/p}$$

is injective.

There is a related (but stronger) concept, namely the notion of *separable generation*:

Definition 4.19 A field extension, K/k, is *separably generated* iff there exists a transcendence base, B, for K/k so that K is separable (algebraic) over $k(B)$. Such a transcendence bases is called a *separating transcendence base for* K/k.

Separable non-algebraic field extensions exist:

Proposition 4.87 *If $K = k(B)$ and B is an algebraically independent set, then K/k is a separable extension.*

Proof. By the argument of Section 4.3, the definition of separability is that when $u_1, u_2, \ldots$ are element of K linearly independent over k, then $u_1^p, \ldots, u_n^p, \ldots$ are again linearly independent over k. If we apply this to all the monomials formed from the elements of B, we see that we must prove: The elements u^p, where u ranges over B, are algebraically independent over k. But, any non-trivial polynomial relation

$$f(u_{i_1}^p, \ldots, u_{i_t}^p) = 0$$

is, *a fortiori*, a polynomial relation for $u_{i_1}, \ldots, u_{i_t}$; hence, cannot be non-trivial. □

We can make two remarks that will be helpful for what follows:

Remarks:

(I) *Separability is transitive*. To see this, say L is separable over K and K is separable over k. Then, the two maps

$$k^{1/p} \otimes_k K \longrightarrow K^{1/p}; \quad K^{1/p} \otimes_K L \longrightarrow L^{1/p} \tag{$*$}$$

are injective. But then, we get the injection

$$(k^{1/p} \otimes_k K) \otimes_K L \longrightarrow K^{1/p} \otimes_K L \tag{$**$}$$

(as L is flat over K). The left hand side of $(**)$ is $k^{1/p} \otimes_k L$ and the right hand side injects into $L^{1/p}$ by $(*)$; so, we are done.

(II) *Any field extension of a perfect field is separable*. For, if k is perfect, then $k = k^p$; that is, $k^{1/p} = k$. But then, $k^{1/p} \otimes_k K \cong K$ and $K \subseteq K^{1/p}$, as required.

(III) *If $K \supseteq L \supseteq k$ and K/k is separable, then L/k is separable*. For consider the map

$$k^{1/p} \otimes_k L \longrightarrow L^{1/p}.$$

Let its kernel be $\mathfrak{A}$. By the flatness of K over L, we see that

$$0 \longrightarrow \mathfrak{A} \otimes_L K \longrightarrow (k^{1/p} \otimes_k L) \otimes_L K = k^{1/p} \otimes_k K \longrightarrow L^{1/p} \otimes_L K$$

is exact. Now, the composed map $k^{1/p} \otimes_k K \longrightarrow L^{1/p} \otimes_L K \longrightarrow K^{1/p}$ is injective by hypothesis; so, $\mathfrak{A} \otimes_L K = (0)$. But, K is faithfully flat over L, therefore $\mathfrak{A} = (0)$.

Corollary 4.88 *If K/k is separably generated, then K/k is separable.*

Proof. Write B for a separating transcendence base for K/k. Then, K is separable over $k(B)$ and the latter is separable over k by Proposition 4.87. Now Remark (I) applies. □

Separable generation is, in general, a strictly stronger concept than separability. Here is a standard example: Let k be a perfect field (i.e., $k = \mathbb{F}_p$) and write $k = k(T, T^{1/p}, T^{1/p^2}, \ldots)$. Thus, $K = \varinjlim_n K_n$, where $K_n = k(T^{1/p^n})$ and each K_n, being pure transcendental over k, is separable over k. Of course, $K^{1/p} = K$ and $k^{1/p} = k$ by choice of k; so K/k is separable. We will now see it is **not** separably generated. Let's write STB for the phrase separating transcendence basis. We know $\mathrm{tr.d}_k K = 1$ as each K_n is algebraic over K_1. Were an element $z \in K$ an STB, we'd have $z \in K_n$ for some n. Now, we may ignore $K_1, \ldots, K_{n-1}$ and still have $K = \varinjlim_m K_m$ $(m \geq n)$, so we may assume $z \in K_1$. i.e., $z \in k(T)$. but then, the diagram of *algebraic* extensions

$$\begin{array}{c} K \\ | \\ k(T) = K_1 \\ | \\ k(Z) \end{array}$$

and the fact that K is separable over $k(Z)$ would show that K/K_1 is a separable algebraic extension and this is nonsense.

Remark. In the general case when K/k is separable and L is a subextension of the layer K/k it does **not** follow that K/L is separable. For example, $L = K_1$ in the above example shows that K/K_1 is **not** separable even though K/k is separable.

This remark indicates that separability is not a good notion in the general case; separable generation is a much better notion. We are going to show now that the two concepts coalesce when the big field is a finitely generated extension; so, the cause of most difficulties is infinite generation in the general case (as should be clear from the counter-example above). Still, even in the finitely generated case, there are problems: K_2/k is separably generated yet it is **not** separably generated over K_1 (or separable–notations as above). The moral is: *be careful with separability (or separable generation) in the non-algebraic case, especially with infinitely generated extensions*.

Theorem 4.89 *If K/k is a finitely generated field extension, then K/k is separable if and only if it is separably generated.*

Proof. One direction is Corollary 4.88; so, assume K/k is separable and finitely generated, say $K = k(T_1, \ldots, T_n)$. We let $r = \mathrm{tr.d}_k K$ and use induction on $n - r$. If the latter is zero, $T_1, \ldots, T_r$ are already an STB; so, assume $n = r + 1$ (this turns out to be the essential case). Now $T_1, \ldots, T_{r+1}$ are algebraically dependent and, by rearranging their order, we may assume $T_1, \ldots, T_r$ are a transcendence base. Then there is a polynomial of smallest degree in X_{n+1} coefficients in $k[X_1, \ldots, X_r]$ having content 1, say $f(X_1, \ldots, X_{r+1})$, so that $f(T_1, \ldots, T_r, T_{r+1}) = 0$. The degree of this polynomial in X^{r+1} must be positive and if its leading coefficient is $a_0(X_1, \ldots, X_r)$, we can localize $k[X_1, \ldots, X_r]$ with respect to a_0 and make f monic in $k[X_1, \ldots, X_r]_{a_0}[X_{r+1}]$. The division algorithm for monic polynomials shows then that if $g \in k[X_1, \ldots, X_{r+1}]$ vanishes on $T_1, \ldots, T_{r+1}$, we have

$$a_0^s g = f \cdot g \quad \text{in} \quad k[X_1, \ldots, X_{r+1}]$$

for some $s \geq 0$. by unique factorization in $k[X_1, \ldots, X_{r+1}]$ it shows further that f is *irreducible*.

Suppose we could show that $f(X_1, \ldots, X_{r+1}) \notin k[X_1^p, \ldots, X_{r+1}^p]$. If so, at least one variable occurs in f with exponent indivisible by p, call this variable X_i. Then T_i is dependent on $T_1, \ldots, T_{i-1}, T_{i+1}, \ldots, T_{r+1}$ and the latter must be algebraically independent by Theorem 4.82. Moreover, as the exponent of T_i is not divisible by p, the element T_i is separable over $k(T_1, \ldots, T_{i-1}, T_{i+1}, \ldots, T_{r+1})$ and so, $T_1, \ldots, T_{i-1}, T_{i+1}, \ldots$, T_{r+1} form a separating transcendence basis, as required.

We use the separability of K/k to prove that $f(X_1, \ldots, X_{r+1}) \notin k[X_1^p, \ldots, X_{r+1}^p]$. Were the contrary true, there would be a polynomial

$$f(X_1, \ldots, X_{r+1}) = g(X_1^p, \ldots, X_{r+1}^p).$$

The monomials $m_1, \ldots, m_t$ comprising g all have degree less than of f, so the elements $m_1(T_1, \ldots, T_{r+1}), \ldots$, $m_t(T_1, \ldots, T_{r+1})$ are linearly independent over k. By separability, the elements

$$m_1(T_1, \ldots, T_{r+1})^p = m_1(T_1^p, \ldots, T_{r+1}^p), \ldots, m_t(T_1, \ldots, T_{r+1})^p = m_t(T_1^p, \ldots, T_{r+1}^p)$$

are still linearly independent over k. Yet the relation

$$g(T_1^p, \ldots, T_{r+1}^p) = f(T_1, \ldots, T_{r+1}) = 0$$

is a non-trivial linear relation among $m_1^p, \ldots, m_t^p$, a contradiction.

For use below, we record what we have just proved:

If K/k is separable, of transcendence degree r, then any set of $r+1$ elements of K, say $T_1, \ldots, T_{r+1}$, which contains a transcendence basis for K/k, already contains a separating transcendence basis for $k(T_1, \ldots, T_{r+1})$ over k. (All wee need note is that, by Remark III above, the field $k(T_1, \ldots, T_{r+1})$ is separable over k.)

Now, let's continue with our induction and finish the proof. We have $n-r>1$ and we assume separable generation for all separable field extensions, K/k, generated by less than $n-r$ elements ($\mathrm{tr.d}_k K = r$). By Remark III, $k(T_1,\ldots,T_{n-1})$ is separably generated; so we can take an STB $U_1,\ldots,U_t$ for it. There are only two possibilities for t: either $t=r-1$ or $t=r$. Since K is then separable (algebraic) over $k(U_1,\ldots,U_t,T_n)$, we need only show the latter field is separably generated over k. But, the transcendence degree of $k(U_1,\ldots,U_t,T_n)$ over k is r and $t+1\leq r+1$ by the above. Again, Remark III shows $k(U_1,\ldots,U_t,T_n)$ is separable over k, and so our argument above (summarized in italics above), implies the required separable generation of $k(U_1,\ldots,U_t,T_n)$ over k. □

We can augment the reasoning in the proof of Theorem 4.89 to obtain a useful theorem of Mac Lane:

Theorem 4.90 *(Mac Lane) Suppose K/k is a finitely generated, separable field extension. Then, any set of generators for K/k already contains a separating transcendence basis for K/k.*

Proof. Write $r=\mathrm{tr.d}\,K$ and say $K=k(T_1,\ldots,T_n)$. We use, as usual, induction on $n-r$, the case $n-r=0$ is trivial and the case $n-r=1$ is covered by the italicized statement in the middle of the proof of Theorem 4.89. For the induction step, use the notation of the last part of Theorem 4.89 and note that, by the induction hypothesis, STB $U_1,\ldots,U_t$ may be chosen from among $T_1,\ldots,T_{n-1}$. Then the $r+1=t+1$ generators $U_1,\ldots,U_t,T_n$ for $k(U_1,\ldots,U_t,T_n)$ are among $T_1,\ldots,T_{n-1},T_n$ and so the case $n-r=1$ now applies and finishes the proof. □

An important corollary of our theorems is this result:

Corollary 4.91 *(F.K. Schmidt) If k is a perfect field, every finitely generated field extension of k is separably generated over k.*

Proof. We apply Remark II and Theorem 4.89 (or 4.90) to our finitely generated extension of k. □

4.12 Further Readings

Some basics of Galois theory is covered in most algebra texts (see Section 2.9). Emil Artin's classic [1] is a must. Other references include Kaplanski [31], Zariski and Samuel [50], Bourbaki (Algebra, Chapter IV) [6], Lafon [33], Morandi [41], Escofier [14] and Van Der Waerden [47].