Some Field Theory

Notation: In this handout, $\mathbb{F}$, $\mathbb{K}$, and $\mathbb{L}$ denote fields.

Definition: Let R be a ring with 1. The smallest subring containing 1 is called the *prime subring* of R.

Let $\varphi : \mathbb{Z} \to R$ be defined by $\varphi(n) = n1$ (i.e., the n^{th} additive power of the multiplicative identity of R). For example, $\varphi(3) = 3(1) = 1+1+1$ and $\varphi(-2) = -2(1) = (-1)+(-1)$. Since $(n+m)1 = n1+m1$ and $(nm)1 = n1 \cdot m1$ for any $m, n \in \mathbb{Z}$ (these are laws of additive exponents), we have that φ is a ring homomorphism (and $\varphi(1) = 1$). Notice that $\varphi(\mathbb{Z})$ is the prime subring of R. Moreover, since $\mathbb{Z}$ is a PID, the kernel of φ is a principal ideal: there exists $n \in \mathbb{Z}$ such that $\ker(\varphi) = (n)$. Since $(n) = (-n)$, we can assume $n \geq 0$.

Definition: This non-negative integer generator the kernel is the *characteristic* of R: $\text{char}(R) = n$.

Calling on the first isomorphism theorem, we have $\mathbb{Z}/(\text{char}(R)) = \mathbb{Z}/\ker(\varphi) \cong \varphi(\mathbb{Z})$. In other words, the prime subring of R is isomorphic to $\mathbb{Z}_n$ when $\text{char}(R) = n > 0$ and the prime subring is isomorphic to $\mathbb{Z}$ when $\text{char}(R) = 0$.

Notice that $\text{char}(R) = 1$ means $1 = 0$ and so $R = \{0\}$. Also, if $\text{char}(R) = n = ab$ for $a, b > 1$, then $a1 \cdot b1 = (ab)1 = n1 = 0$ and $a1, b1$ are non-zero since $a, b < n$ and the characteristic is the smallest positive (additive) power of 1 that equals 0. In other words, rings with composite characteristics must have zero divisors.

Corollary: Let R be an integral domain. Then $\text{char}(R)$ is either zero or prime. In particular, fields must be of zero or prime characteristic.

Identification: We identify the prime subring of a ring of characteristic n with $\mathbb{Z}_n$ if $n > 0$ and $\mathbb{Z}$ if $n = 0$.

Thus using this identification, we have that every integral domain contains a copy of $\mathbb{Z}$ or $\mathbb{Z}_p$ for some prime p. When we have a field, each non-zero element must have an inverse, so a field containing a copy of $\mathbb{Z}$ must also contain a copy of $\mathbb{Q}$. Thus if $\text{char}(\mathbb{F}) = p > 0$ (p must be prime), then the prime subring of $\mathbb{F}$ is actually the prime sub*field* of $\mathbb{F}$, namely $\mathbb{Z}_p$. If $\text{char}(\mathbb{F}) = 0$, then the prime subring of $\mathbb{F}$ is $\mathbb{Z}$ and its prime subfield (i.e., its smallest subfield) is $\mathbb{Q}$.

Remark: While having characteristic $n > 0$, brings to mind finite rings like $\mathbb{Z}_n$ or $(\mathbb{Z}_n)^{3\times 3}$, having a positive characteristic $n > 1$ does not preclude being an infinite ring. For example, $\mathbb{Z}_n[x]$ is an infinite ring of characteristic n. [Characteristic 1 always means our ring is trivial, $\{0\}$, since $1 = 0$.] Also, notice that for any prime p, $\mathbb{Z}_p(x)$ (i.e., ratios of polynomials drawn from $\mathbb{Z}_p[x]$) gives us an infinite field of characteristic p.

Definition: Recall from linear algebra that a *vector space* V over a field $\mathbb{F}$ is an Abelian group $(V, +)$ equipped with a scalar multiplication $(s, v) \mapsto sv$ where $s \in \mathbb{F}$ and $v \in V$ such that for all $v, w \in V$ and $s, t \in \mathbb{F}$, we have $1v = v$, $s(tv) = (st)v$, $(s + t)v = sv + tv$, and $s(v + w) = sv + sw$.

While introductory linear algebra courses usually stick to working over the real numbers (i.e., $\mathbb{F} = \mathbb{R}$), most all of the results covered in such a course work over any field. In particular, there is a well-defined notion of dimension. If V is a vector space over $\mathbb{F}$, then $\dim_{\mathbb{F}}(V)$ is the size of any (hence all) bases for V (working over $\mathbb{F}$). For example, $\mathbb{C} = \{a + bi \mid a, b \in \mathbb{R}\}$ has basis $\{1, i\}$ over $\mathbb{R}$ and $\{1\}$ working over itself. In particular, $\dim_{\mathbb{R}}(\mathbb{C}) = 2$ and $\dim_{\mathbb{C}}(\mathbb{C}) = 1$. We could also note that $\dim_{\mathbb{Q}}(\mathbb{C})$ is infinite (specifically $\mathfrak{c} = 2^{\aleph_0}$ continuum in dimension)!

Definition: Let $\mathbb{F} \subseteq \mathbb{K}$. Then we say $\mathbb{F}$ is a *subfield* of $\mathbb{K}$ and $\mathbb{K}$ is an *extension field* of $\mathbb{F}$. Oddly, this is often denoted $\mathbb{K}/\mathbb{F}$ even though this has nothing to do with quotients or cosets. Notice that $\mathbb{K}$ (or any ring containing $\mathbb{F}$) can be thought of as a vector space over $\mathbb{F}$ (it's an Abelian group and we can "scale" (i.e., multiply) elements of $\mathbb{K}$ by elements of $\mathbb{F}$). We call $[\mathbb{K} : \mathbb{F}] = \dim_{\mathbb{F}}(\mathbb{K})$ the *degree* of the extension.

Theorem: Finite Fields has Prime Power Orders: Let $\mathbb{F}$ be a finite field. Then $|\mathbb{F}| = p^k$ for the prime $p = \text{char}(\mathbb{F})$ and some positive integer k.

Proof: Let $\mathbb{F}$ be a finite field. Thus $\mathbb{F}$ cannot contain an infinite set, so it cannot contain a copy of $\mathbb{Q}$. Thus its prime subfield must be $\mathbb{Z}_p$ where $\text{char}(\mathbb{F}) = p$. Now we have that $\mathbb{F}$ is an extension field of its prime subfield. Again, $\mathbb{F}$ is finite so $[\mathbb{F} : \mathbb{Z}_p] = \dim_{\mathbb{Z}_p}(\mathbb{F}) = k < \infty$. Recall that any vector space V of dimension $n < \infty$ over some field $\mathbb{K}$ is isomorphic (as a vector space) to $\mathbb{K}^n$. Thus $\mathbb{F} \cong (\mathbb{Z}_p)^k$ (as vector spaces). Therefore, $|\mathbb{F}| = |\mathbb{Z}_p|^k = p^k$. ■

Theorem: (The Degree Formula) Let $\mathbb{L}$ extend $\mathbb{K}$ which extends $\mathbb{F}$. Then $[\mathbb{L} : \mathbb{F}] = [\mathbb{L} : \mathbb{K}] \cdot [\mathbb{K} : \mathbb{F}]$.

Proof: Let α be a basis for $\mathbb{L}$ working over $\mathbb{K}$ and β be a basis for $\mathbb{K}$ working over $\mathbb{F}$. Consider $\gamma = \alpha \cdot \beta = \{uv \mid u \in \alpha \text{ and } v \in \beta\}$. We show that γ is a basis for $\mathbb{L}$ working over $\mathbb{F}$.

Suppose that $\sum_{i,j} c_{ij}u_iv_j = 0$ for some finite linear combination of u_iv_j drawn from γ (i.e., $u_i \in \alpha$ and $v_j \in \beta$) and scalars $c_{ij} \in \mathbb{F}$. Then $\sum_i(\sum_j c_{ij}v_j)u_i = 0$ where $\sum_j c_{ij}v_j \in \mathbb{K}$. We have that the u_i's are linearly independent over $\mathbb{K}$. Thus for each i, we have $\sum_j c_{ij}v_j = 0$. But the v_j's are linearly independent over $\mathbb{F}$. Thus for each i, $c_{ij} = 0$ for all j. Therefore, we have that γ is linearly independent.

Next, suppose $x \in \mathbb{L}$. Then $x = \sum_i y_iu_i$ for some u_i's in α and $y_i \in \mathbb{K}$ (since α is a basis for $\mathbb{L}$ over $\mathbb{K}$). For each i, $y_i = \sum_j c_{ij}v_j$ for some v_j's in β and $c_{ij} \in \mathbb{F}$ (since $y_i \in \mathbb{K}$ and β is a basis for $\mathbb{K}$ over $\mathbb{F}$). Therefore, $x = \sum_i y_iu_i = \sum_i(\sum_j c_{ij}v_j)u_i = \sum_{i,j} c_{ij}u_iv_j$ and thus x lies in the span of γ.

Therefore, γ is a linearly independent (over $\mathbb{F}$) and $\text{span}_{\mathbb{F}}\gamma = \mathbb{L}$. We have that it is a basis for $\mathbb{L}$ over $\mathbb{F}$.

Finally, notice that $|\gamma| = |\alpha| \cdot |\beta|$: If $u_iv_j = u_kv_\ell$, we must have $v_j = v_\ell$ and $u_i = u_k$ since v's belong to $\mathbb{K}$ and the u's are a basis over $\mathbb{K}$ and working over a basis coordinates (i.e., coefficients of u's) must match. ■

Proposition: Let $0 \neq f(x) \in \mathbb{F}[x]$ and let $n = \deg(f(x))$.
Then, working over $\mathbb{F}$, $\{1 + (f(x)), x + (f(x)), \ldots, x^{n-1} + (f(x))\}$ is a basis for $\mathbb{F}[x]/(f(x))$.

Proof: Let $g(x) \in \mathbb{F}[x]$. There are unique polynomials $q(x), r(x) \in \mathbb{F}[x]$ where $r(x) = 0$ or $\deg(r(x)) < \deg(f(x)) = n$ such that $g(x) = q(x)f(x) + r(x)$. Therefore, there is a unique polynomial $r(x)$ where $r(x) = 0$ or $\deg(r(x)) < n$ such that $g(x) + (f(x)) = r(x) + (f(x))$. In other words, every coset in $\mathbb{F}[x]/(f(x))$ has a unique representation as $a_0 + a_1x + \cdots + a_{n-1}x^{n-1} + (f(x))$ for some $a_0, \ldots, a_{n-1} \in \mathbb{F}$. Existence of these scalars $a_0, \ldots, a_{n-1}$ implies our proposed basis is spans the quotient, and the uniqueness of these scalars implies the proposed basis is linearly independent. ■

For example, $\mathbb{Q}[x]/(x^2-1)$ has a basis $\{1 + (x^2-1), x + (x^2-1)\}$ working over $\mathbb{Q}$. Notice that $\mathbb{F}[x]/(0) \cong \mathbb{F}[x]$ has the standard basis $\{1, x, x^2, \ldots\}$ and thus is a countable dimensional vector space over $\mathbb{F}$.

Another example, $\mathbb{Q}[x]/(x^4-1) = \{a_3x^3 + a_2x^2 + a_1x + a_0 + (x^4-1) \mid a_3, a_2, a_1, a_0 \in \mathbb{Q}\}$ and $a_3x^3 + a_2x^2 + a_1x + a_0 + (x^4-1) = b_3x^3 + b_2x^2 + b_1x + b_0 + (x^4-1)$ if and only if $a_3 = b_3$, $a_2 = b_2$, $a_1 = b_1$, and $a_0 = b_0$. Similarly, this is why $\mathbb{Q}[x]/(x^3-2) \cong \mathbb{Q}[\sqrt[3]{2}] = \{a + b\,2^{1/3} + c\,2^{2/3} \mid a, b, c \in \mathbb{Q}\}$ and $[\mathbb{Q}[\sqrt[3]{2}] : \mathbb{Q}] = 3$.

Definition: Let $\mathbb{K}$ extend $\mathbb{F}$ and suppose $\alpha \in \mathbb{K}$. If there exists some $0 \neq f(x) \in \mathbb{F}[x]$ such that $f(\alpha) = 0$, we say α is *algebraic* over $\mathbb{F}$. If no such polynomial exists, α is *transcendental* over $\mathbb{F}$.

If no base field is mentioned, we usually assume one means that α is algebraic or transcendental over $\mathbb{Q}$. For example, $\sqrt{2}$ is algebraic (over $\mathbb{Q}$) since it is the root of $x^2 - 2$. It can be shown (through fairly technical calculations) that both π and Euler's number e are transcendental (over $\mathbb{Q}$). Notice that π and e are algebraic over $\mathbb{R}$ – they're roots of $x - \pi$ and $x - e$ in $\mathbb{R}[x]$ respectively. Base fields matter!

One can also speak over algebraic or transcendental functions. Here the base field is usually assumed to be $\mathbb{C}(x)$ (i.e., ratios of complex polynomials). Notice that $\sqrt{x}$ is a root of $Y^2 - x \in (\mathbb{C}(x))[Y]$. Thus $\sqrt{x}$ is an algebraic function. On the other hand, it is possible to show that $\sin(x), \cos(x)$, and e^x are transcendental functions (working over $\mathbb{C}(x)$).

Suppose that $\alpha \in \mathbb{K}$ is algebraic over $\mathbb{F}$. Then $I = \{f(x) \in \mathbb{F}[x] \mid f(\alpha) = 0\}$ is a non-zero ideal of $\mathbb{F}[x]$ (Why? If $f(\alpha) = 0$ and $g(\alpha) = 0$, then $(f-g)(\alpha) = f(\alpha) - g(\alpha) = 0 - 0 = 0$ so $f - g \in I$ and if $f(\alpha) = 0$ then $hf(\alpha) = h(\alpha)f(\alpha) = h(\alpha) \cdot 0 = 0$ for any $h(x) \in \mathbb{F}[x]$). But $\mathbb{F}[x]$ is a PID, so I must have a non-zero generator, say $I = (m(x))$. We know that a generator for an ideal in a Euclidean domain is an element of lowest degree in that ideal. Moreover, associates generate the same ideal, so one can assume $m(x)$ is monic. This leads to the following definition:

Definition: Let $\alpha \in \mathbb{K}$ be algebraic over $\mathbb{F}$. The monic generator of $I = \{f(x) \in \mathbb{F}[x] \mid f(\alpha) = 0\}$ (i.e., the lowest degree, monic polynomial in $\mathbb{F}[x]$ with root α) is called the *minimal polynomial* for α (over $\mathbb{F}$).

Proposition: Minimal polynomials are irreducible.

Proof: Suppose $m(x)$ is the minimal polynomial for α working over $\mathbb{F}$. Since $m(x)$ is non-zero and has a root α, it is a non-zero constant polynomial (i.e., not zero or a unit in $\mathbb{F}[x]$). Thus if it is not irreducible, it must have a proper factorization, say $m(x) = p(x)q(x)$. But $0 = m(\alpha) = p(\alpha)q(\alpha)$. Thus either $p(\alpha) = 0$ or $q(\alpha) = 0$. However, $m(x)$ is the lowest degree polynomial with root α. Therefore, no such proper factorization can exist! ■

Example: The minimal polynomial of $\sqrt[3]{5}$ is $x^3 - 5$ if we are working over $\mathbb{Q}$. On the other hand, it's $x - \sqrt[3]{5}$ if we're working over $\mathbb{R}$.

Theorem: Let α be algebraic over $\mathbb{F}$ with minimal polynomial $m(x)$. Then $\mathbb{F}[x]/(m(x)) \cong \mathbb{F}[\alpha]$ and $[\mathbb{F}[\alpha] : \mathbb{F}]$. In other words, the degree of the extension is the same as the degree of the element's minimal polynomial.

Proof: Consider $\varphi : \mathbb{F}[x] \to \mathbb{K}$ where $\varphi(f(x)) = f(\alpha)$. Then φ is a homomorphism whose kernel is the ideal generated by α's minimal polynomial and the image of φ is nothing more than the subfield of $\mathbb{K}$ generated by $\mathbb{F}$ and α (i.e., $\mathbb{F}[\alpha]$). The result follows from the first isomorphism theorem. ■

If one tries this with a transcendental element α, one finds that $\mathbb{F}[x] \cong \mathbb{F}[x]/(0) \cong \mathbb{F}[\alpha]$. Thus adjoining a transcendental element does not yield a field (just an integral domain). In fact, working over $\mathbb{F}$, α is algebraically indistinguishable from an intedeterminate like x. In particular, $\mathbb{Q}[\pi] \neq \mathbb{Q}(\pi)$ (i.e., the ring adjoining π to $\mathbb{Q}$ is not the same as the field generated by $\mathbb{Q}$ and π). Also, $\mathbb{Q}[\pi] \cong \mathbb{Q}[x]$. In fact, one could define algebraic: any element α such that $\mathbb{F}[\alpha] = \mathbb{F}(\alpha)$ whereas transcendental means $\mathbb{F}[\alpha] \neq \mathbb{F}(\alpha)$.

Example: We have that $\sqrt{2}$ is algebraic over $\mathbb{Q}$ since it is a root of $x^2 - 2$ (this is $\sqrt{2}$'s minimal polynomial over $\mathbb{Q}$). We have that $\{1, \sqrt{2}\}$ is a basis for $\mathbb{Q}[\sqrt{2}]$ working over $\mathbb{Q}$ and that $[\mathbb{Q}[\sqrt{2}] : \mathbb{Q}] = 2$. Likewise, $\sqrt[3]{5}$ is algebraic over $\mathbb{Q}[\sqrt{2}]$ with minimal polynomial $x^3 - 5$ (notice that $x^3 - 5$ has no roots of the form $a + b\sqrt{2}$ where $a, b \in \mathbb{Q}$). Thus $\{1, \sqrt[3]{5}, (\sqrt[3]{5})^2\}$ is a basis for $\mathbb{Q}[\sqrt{2}, \sqrt[3]{5}]$ working over $\mathbb{Q}[\sqrt{2}]$ and $[\mathbb{Q}[\sqrt{2}, \sqrt[3]{5}] : \mathbb{Q}[\sqrt{2}]] = 3$. The degree formula (and its proof) then tell us that $\{1, \sqrt{2}, \sqrt[3]{5}, \sqrt[3]{5} \cdot \sqrt{2}, (\sqrt[3]{5})^2, (\sqrt[3]{5})^2 \cdot \sqrt{2}\}$ is a basis for $\mathbb{Q}[\sqrt{2}, \sqrt[3]{5}]$ working over $\mathbb{Q}$ and that $[\mathbb{Q}[\sqrt{2}, \sqrt[3]{5}] : \mathbb{Q}] = [\mathbb{Q}[\sqrt{2}, \sqrt[3]{5}] : \mathbb{Q}[\sqrt{2}]] \cdot [\mathbb{Q}[\sqrt{2}] : \mathbb{Q}] = 3 \cdot 2 = 6$.

Definition: Let $\mathbb{K}$ extend $\mathbb{F}$. We say that $\mathbb{K}$ is a *finite extension* of $\mathbb{F}$ if $[\mathbb{K} : \mathbb{F}] < \infty$. If every element of $\mathbb{K}$ is algebraic over $\mathbb{F}$, we say that $\mathbb{K}$ is an *algebraic extension* of $\mathbb{F}$.

Theorem: Finite extensions are algebraic

Proof: Let $[\mathbb{K} : \mathbb{F}] = n < \infty$ and let $\alpha \in \mathbb{K}$. Then $\{1, \alpha, \alpha^2, \ldots, \alpha^n\}$ must be linearly dependent (if not, we would have a linearly independent set with more than $n = \dim_{\mathbb{F}}(\mathbb{K})$ elements). Thus there exists scalars $a_0, \ldots, a_n \mathbb{F}$ (not all zero) such that $a_0 + a_1\alpha + \cdots + a_n\alpha^n = 0$ (i.e., α is the root of some non-zero polynomial). This means α is algebraic over $\mathbb{F}$. ■

The converse of this theorem is not true. For example, $\overline{\mathbb{Q}} = \{\alpha \in \mathbb{C} \mid \alpha \text{ is algebraic over } \mathbb{Q}\}$, called the field of algebraic numbers, is an infinite extension of $\mathbb{Q}$. But it is algebraic over $\mathbb{Q}$.

Let $f(x) \in \mathbb{F}[x]$ be a non-constant polynomial. Let $g(x)$ be an irreducible factor of $f(x)$. Then $\mathbb{K} = \mathbb{F}[x]/(g(x))$ is a field containing a copy of $\mathbb{F}$. Moreover, $\alpha = x + (g(x))$ is a root of $g(x)$ in $\mathbb{K}$ since $g(x + (g(x))) = g(x) + (g(x)) = 0 + (g(x))$. Thus we write: $\mathbb{K} = \mathbb{F}[\alpha]$ where α is some root of $g(x)$. By iterating this process one can extend $\mathbb{F}$ to a field $\mathbb{L}$ such that $\mathbb{L}$ contains all of the roots of $f(x)$. Moreover, constructing $\mathbb{L}$ this way we have $\mathbb{L} = \mathbb{F}[\alpha_1, \ldots, \alpha_n]$ where $f(x) = (x - \alpha_1) \cdots (x - \alpha_n)$. We have just constructed the splitting field of $f(x)$ (over $\mathbb{F}$).

Definition: Let $f(x) \in \mathbb{F}[x]$. Suppose $\mathbb{L}$ is an extension of $\mathbb{F}$ such that $f(x) = (x - \alpha_1) \cdots (x - \alpha_n)$ and $\mathbb{L} = \mathbb{F}[\alpha_1, \ldots, \alpha_n]$ (i.e., $f(x)$ *splits* into linear factors in $\mathbb{L}[x]$ and $\mathbb{L}$ is generated by the roots of $f(x)$ and $\mathbb{F}$). Then $\mathbb{L}$ is called the *splitting field* for $f(x)$ (working over $\mathbb{F}$).

The discussion above this definition shows that splitting fields always exist. We state without proof:

Theorem: Splitting fields are unique up to isomorphism. In particular, given $f(x) \in \mathbb{F}$. Suppose $\mathbb{L} = \mathbb{F}[\alpha_1, \ldots, \alpha_n]$ and $\mathbb{K} = \mathbb{F}[\beta_1, \ldots, \beta_n]$ are splitting fields. Then there is some permutation $\sigma \in S_n$ such and an isomorphism $\varphi : \mathbb{L} \to \mathbb{K}$ such that $\varphi(x) = x$ for all $x \in \mathbb{F}$ and $\varphi(\alpha_j) = \beta_{\sigma(j)}$ (i.e., $\mathbb{F}$ is fixed pointwise and roots of $f(x)$ map to roots of $f(x)$).

For example, $\mathbb{Q}[i, \sqrt{5}]$ is a splitting field for $f(x) = (x^2 + 1)(x^2 - 5)$ (working over $\mathbb{Q}$). The splitting field for $x^2 - 5$ over $\mathbb{Q}$ is $\mathbb{Q}[\sqrt{5}]$ whereas the splitting field for $x^2 - 5$ over $\mathbb{R}$ is just $\mathbb{R}$ ($x^2 - 5$ already splits in $\mathbb{R}[x]$). The splitting field for $x^2 + 1$ over $\mathbb{R}$ is $\mathbb{R}[i] = \mathbb{C}$.

The proof of the above theorem is not that difficult, but it would require us setting up a fair amount of machinery. This kind of result is really the first step toward developing Galois theory.

Definition: We say that $\overline{\mathbb{F}}$ is *algebraically closed* if every non-constant polynomial $f(x) \in \overline{\mathbb{F}}[x]$ has a root in $\overline{\mathbb{F}}$. Equivalently, in $\overline{\mathbb{F}}[x]$, any irreducible polynomial must be linear.

While we will not pursue this here, it can be shown that every field $\mathbb{F}$ is contained in an algebraically closed field (essentially one just takes each irreducible polynomial in $\mathbb{F}[x]$ and tacks on all of its roots). It's essentially the union of all of the splitting fields (but "union" and "all" must be dealt with carefully). The smallest (in terms of containment) algebraically closed field containing $\mathbb{F}$ is called its *algebraic closure*. Much like splitting fields, it can be shown that algebraic closures are unique up to isomorphism (fixing $\mathbb{F}$ pointwise). *The* algebraic closure of $\mathbb{F}$ is denoted by $\overline{\mathbb{F}}$.

The fundamental theorem of algebra states that $\mathbb{C}$ is algebraically closed, so $\overline{\mathbb{R}} = \mathbb{C}$. The algebraic closure of the rational numbers, denoted $\overline{\mathbb{Q}}$, is called the field of *algebraic numbers* (mentioned above). Note that $\overline{\mathbb{Q}} \subsetneq \mathbb{C}$. For example, $\pi \notin \overline{\mathbb{Q}}$. In fact, $\mathbb{C}$ is continuum in size while $\overline{\mathbb{Q}}$ is only countable.

Theorem: Finite subgroups of $\mathbb{F}^\times$ must be cyclic. Moreover, for every positive integer n, there is at most one subgroup of order n.

Proof: Let G be a finite subgroup of $\mathbb{F}^\times$, say $|G| = n < \infty$.

Lemma: Let $a \in G$ be an element of maximal order $|a| = m$ in a finite Abelian group. If $b \in G$, then $|b|$ divides m.

Let $d = \gcd(|b|, m)$. Then let $c = b^d$ so $|c| = |b|/d$ and thus the orders of c and a are relatively prime. Now suppose $(ac)^\ell = 1$. Since G is Abelian, $1 = (ac)^\ell = a^\ell c^\ell$ so that $a^\ell = c^{-\ell}$. Since $x = a^\ell = c^{-\ell}$ is a power of a and a power of c, its order divides both m and $|c|$. But m and $|c|$ are relatively prime. Therefore, the order of $x = a^\ell = c^{-\ell}$ is 1 and so $a^\ell = 1 = c^\ell$.

Thus the smallest positive power such that $(ac)^\ell = 1$ is the smallest power such that $a^\ell = c^\ell = 1$. This must be $\operatorname{lcm}(m, |c|)$ ($\geq m$). But the order of a (i.e., $|a| = m$) is maximal. Therefore, $\operatorname{lcm}(m, |c|) = m$ and since m and $|c|$ are relatively prime, we have $|c| = 1$ (i.e., $b^d = 1$) and so the order of b divides m.■

Back to proving our theorem. We now know that the order of every element in G divides some maximal order m (which by Lagranges theorem divides $|G| = n$). Therefore, for all $g \in G$, we have $g^m = 1$ so $g^m - 1 = 0$. Thus every element of G is a root of $x^m - 1$ (i.e., $x^m - 1$ has at least n roots). On the other hand, $x^m - 1$ cannot have more than $\deg(x^m - 1) = m$ roots. Therefore, we have $m \geq n$. However, $m \leq n$ since m divides n. Thus $m = n$ (i.e., the maximal element order is $|G| = n$). Hence G is cyclic.

Now suppose H and G both have order n. Thus all of there elements are roots of $x^n - 1$. If $H \neq G$, then $x^n - 1$ would have too many roots! Thus $H = G$. ■

Corollary: The group of units of a finite field must be cyclic.

In $\mathbb{R}$, the only elements of finite (multiplicative) order are ± 1. Thus the only possible elements of finite subgroups of $\mathbb{R}^\times$ are ± 1. Note that the only finite subgroups of $\mathbb{R}^\times$ are thus $\{1\}$ and $\{1, -1\}$ (both cyclic). In $\mathbb{C}$, there is a finite subgroup of every positive integer order n. In fact, the unique subgroup of order n is exactly the n^{th}-roots of unity. A generator of such a subgroup is called a primitive n^{th}-root.

If $\mathbb{F}$ is a finite field, one calls $\zeta \in \mathbb{F}$ a primitive element if $\langle \zeta \rangle = \mathbb{F}^\times$ (i.e., it generates the group of units of our finite field). While the above corollary guarantees these elements exist, there is no nice way of finding them (i.e., we are stuck just blindly searching).

Theorem: There is a unique finite field of order p^k for every prime p and positive integer k. In particular, *the* field of order p^k, denoted $\mathbb{F}_{p^k}$ or $\mathrm{GF}(p^k)$ (GF = Galois Field), is the splitting field for $x^{p^k} - x \in \mathbb{Z}_p[x]$.

Proof: Let $\mathbb{F}$ be the splitting field for $f(x) = x^{p^k} - x \in \mathbb{Z}_p[x]$. Notice that the (formal) derivative of this polynomial is $f'(x) = p^k x^{p^k} - 1 = -1$ since $p^k = 0$ in $\mathbb{Z}_p$. Therefore, $f(x)$ and $f'(x)$ are relatively prime. This implies that $f(x)$ has no repeated roots. Since $\mathbb{F}$ contains all of the roots of $x^{p^k} - x$ and the roots are distinct, we have $|\mathbb{F}| \geq p^k$.

Now let $q = p^k$ (for convenience) and let $\mathbb{K} = \{x \in \mathbb{F} \mid x^q - x\}$ (our set of q roots of $x^q - x$). We have $1 \in \mathbb{K}$ since $1^q - 1 = 1 - 1 = 0$. Let $a, b \in \mathbb{K}$. Then $a^q - a = 0$ and $b^q - b = 0$ so $a^q = a$ and $b^q = b$. Next, $(a+b)^q = a^q + qa^{q-1}b + \cdots + qab^{q-1} + b^q = a^q + b^q = a + b$ since the binomial coefficients (except the first and last) are all divisible by p and thus are 0 in $\mathbb{Z}_p$. Therefore, $(a+b)^q - (a+b) = 0$ so $a + b$ is a root of $x^q - x$ and thus $a + b \in \mathbb{K}$. Next, $(-a)^q = (-1)^q a^q = (-1)^q a = -a$ if q is odd. On the other hand, if q is even, we are working in characteristic 2 and so $a = -a$. Either way, $(-a)^q - (-a) = 0$ and thus $-a$ is a root of $x^q - x$ so that $-a \in \mathbb{K}$. Also, $(ab)^q = a^q b^q = ab$. Thus $(ab)^q - (ab) = 0$ so ab is a root of $x^q - x$. Thus $ab \in \mathbb{K}$. Finally, if $a \neq 0$, $(a^{-1})^q = a^{-q} = (a^q)^{-1} = a^{-1}$. Therefore, a^{-1} is a root of $x^q - x$ and so $a^{-1} \in \mathbb{K}$. We have shown that the set of roots $\mathbb{K}$ is in fact a subfield of $\mathbb{F}$. Since the splitting field is the smallest field in which $x^q - x$ splits and such a field must at least contain the roots, we have $\mathbb{K}$ must be the splitting field. Therefore, $\mathbb{F} = \mathbb{K}$ and so $\mathbb{F}$ is a desired field of order $q = p^k$.

From this we learn that a field of order p^k is nothing more than *the* splitting field of $x^{p^k} - x$ over $\mathbb{Z}_p$. By uniqueness of splitting fields, this field is unique up to isomorphism. ■

Using essentially the argument in the proof above, one has that $\varphi : \mathbb{F}_{p^k} \to \mathbb{F}_{p^k}$ defined by $\varphi(x) = x^p$ is an automorphism of $\mathbb{F}_{p^k}$. This is called the *Frobenius* automorphism of $\mathbb{F}_{p^k}$. Everything is $\mathbb{F}_{p^k}$ is fixed by φ^k. It turns out that this automorphism generates the group of automorphisms of $\mathbb{F}_{p^k}$.

Suppose $\mathbb{F}_{p^m} \subseteq \mathbb{F}_{p^n}$. Then $\mathbb{F}_{p^n}$ is a vector space over $\mathbb{F}_{p^m}$. Thus $\mathbb{F}_{p^n} \cong (\mathbb{F}_{p^m})^k$ (as vector spaces) where $\dim_{\mathbb{F}_{p^m}}(\mathbb{F}_{p^n}) = k$. Comparing sizes, we have $p^n = (p^m)^k = p^{mk}$. Therefore, m must be a divisor of n. Conversely,

if m divides n, one can show that $x^{p^m} - x$ divides $x^{p^n} - x$ (in $\mathbb{Z}_p[x]$). Thus the splitting field of $x^{p^m} - x$ can be embedded in the splitting field of $x^{p^n} - x$. In particular, we can identify $\mathbb{F}_{p^m}$ with a subfield of $\mathbb{F}_{p^n}$. Assuming this identification:

Theorem: $\mathbb{F}_{p^m}$ is a subfield of $\mathbb{F}_{p^n}$ if and only if m divides n.

Therefore, $\mathbb{F}_4$ is a subfield of $\mathbb{F}_{16}$ (since $4 = 2^2$ and $16 = 2^4$ and 2 divides 4) but $\mathbb{F}_4$ is not a subfield of $\mathbb{F}_8$ (since $4 = 2^2$ and $8 = 2^3$ and 2 does not divide 3).

It is possible to use these identifications to build a tower: $\mathbb{F}_p \subseteq \mathbb{F}_{p^{2!}} \subseteq \mathbb{F}_{p^{3!}} \subseteq \cdots$ since $n!$ divides $(n+1)!$. Then one can define $\mathbb{F}_{p^\infty} = \bigcup_{n=1}^{\infty} \mathbb{F}_{p^{n!}}$. This turns out to be the algebraic closure every finite field of characteristic p.

Suppose $f(x) \in \mathbb{Z}_p[x]$ is irreducible with root α. Then $\mathbb{Z}_p[\alpha] = \mathbb{Z}_p[x]/(f(x))$ is a field of degree $[\mathbb{Z}_p[\alpha] : \mathbb{Z}_p] = \deg(f(x))$. This means $|\mathbb{Z}_p[\alpha]| = p^{\deg(f(x)}$. So by the uniqueness of finite fields, we must have $\mathbb{Z}_p[\alpha] \cong \mathbb{F}_{p^{\deg(f(x))}}$.

In particular, if we want to construct a field of order p^k, we should find an irreducible polynomial $f(x) \in \mathbb{Z}_p[x]$ with $\deg(f(x)) = k$. Then $\mathbb{F}_{p^k} = \mathbb{Z}_p[x]/(f(x))$. It can be shown that $\mathbb{Z}_p[x]$ does in fact have irreducible polynomials of all positive degrees, so this is always possible.

Example: Let's construct the field of order 4. We need an irreducible quadratic in $\mathbb{Z}_2[x]$ (since $2^2 = 4$). Notice that $x^2 + x + 1$ has no roots in $\mathbb{Z}_2$, so $x^2 + x + 1$ is irreducible in $\mathbb{Z}_2[x]$. Therefore, $\mathbb{Z}_2[x]/(x^2 + x + 1) \cong \mathbb{F}_{2^2}$.

Let's make this more concrete. We have $\mathbb{Z}_2[x]/(x^2 + x + 1) \cong \mathbb{Z}_2[\alpha] = \{0, 1, \alpha, \alpha + 1\}$ where $\alpha^2 + \alpha + 1 = 0$ (we only need up to linear stuff in α since α is a root of a polynomial of degree 2). Notice $\alpha^2 = -\alpha - 1 = \alpha + 1$ since we're working mod 2. We can thus fill out the following addition and multiplication tables:

$+$	0	1	α	$\alpha+1$
0	0	1	α	$\alpha+1$
1	1	0	$\alpha+1$	α
α	α	$\alpha+1$	0	1
$\alpha+1$	$\alpha+1$	α	1	0

$\cdot$	0	1	α	$\alpha+1$
0	0	0	0	0
1	0	1	α	$\alpha+1$
α	0	α	$\alpha+1$	1
$\alpha+1$	0	$\alpha+1$	1	α

For example, $\alpha^2 = \alpha + 1$ (since $\alpha^2 + \alpha + 1 = 0$). Likewise, $(\alpha + 1)^2 = \alpha^2 + 2\alpha + 1 = \alpha^2 + 1 = (\alpha + 1) + 1 = \alpha$ and $\alpha(\alpha + 1) = \alpha^2 + \alpha = 1$ both using the relation $\alpha^2 + \alpha + 1 = 0$ to simplify.

Also, let $f(x)$ be an irreducible (over $\mathbb{Z}_p[x]$) of degree k. Then its roots belong to the field $\mathbb{F}_{p^k}$. But this field is precisely the roots of $x^{p^k} - x$. It follows that that the factors of $f(x)$ are factors of $x^{p^k} - x$. Thus $f(x)$ divides $x^{p^k} - x$. Thus...

Theorem: Working in $\mathbb{Z}_p[x]$, monic irreducible polynomials of degree k are precisely the irreducible (degree k) factors of $x^{p^k} - x$. Moreover, $x^{p^k} - x$ has no repeated factors.

For example, working in $\mathbb{Z}_2[x]$, we have that $x^2 - x = x^2 + x = x(x + 1)$. Thus the linear irreducibles are x and $x + 1$. Next, $x^4 - x = x^4 + x = x(x + 1)(x^2 + x + 1)$. Therefore, the irreducible quadratic (there is only one in $\mathbb{Z}_2[x]$) is $x^2 + x + 1$. We could have figured this out by listing all quadratics: $x^2, x^2 + 1, x^2 + x, x^2 + x + 1$ and then plugging in $x = 0$ and $x = 1$ to see which ones have roots. If we wanted to find all irreducible (monic) cubics in $\mathbb{Z}_2[x]$, we could factor $x^{2^3} - x$. This polynomial would also contain all of the irreducibles of degree k for any k dividing 3. Notice that x and $x + 1$ are the only irreducible linears, so the other $2^3 - 1 - 1 = 6$ degrees of factors of $x^{2^3} - x$ must consist of the irreducible cubics. From this (and the fact that no factors are repeated in $x^{2^3} - x$) we can conclude that $\mathbb{Z}_2[x]$ must have exactly 2 irreducible cubics! This kind of reasoning can be generalized and thus gives us an iterative technique for finding all irreducibles in $\mathbb{Z}_p[x]$ of degree k.