Chapter 1

Artin's primitive root conjecture for number fields

1.1 Introduction

The unit group of a finite field of n elements is a cyclic group, and it has $\varphi(n-1)$ choices of generator. For example, the finite field $\mathbf{F}_{11}$ of 11 elements has $\varphi(10) = 4$ elements that each generate its unit group. To see if an element x is a generator of $\mathbf{F}_{11}^*$, we could check that x isn't a square or a fifth power, 2 and 5 being the prime divisors of 10, the order of $\mathbf{F}_{11}^*$. We find that the residue classes of 2, 6, 7 and 8 are the four generators of $\mathbf{F}_{11}^*$. We call the integers in these residue classes primitive roots modulo 11. More generally, we call a rational number x a *primitive root* modulo a prime q if q does not divide the denominator of x and x mod q generates $\mathbf{F}_q^*$.

Instead of determining which integers generate the unit group of a given finite field $\mathbf{F}_q$, we can reverse the question and ask modulo which (or how many) primes q a fixed integer (or rational number) is a primitive root. For example, 2 is not only a primitive root modulo 11, but also modulo 3, 5, 13, 19, 29, 37 and numerous other primes.

Question 1.1. *If x is a non-zero rational number, for how many primes q not dividing the numerator and denominator of x is the unit group of $\mathbf{F}_q$ generated by x mod q?*

Only in the simple case when x is -1 or a square, where the answer is finite, can this question be easily answered. In 1927, Emil Artin conjectured that there should be an infinite number of primes q modulo which $x = 2$ is a primitive root, and even that the set of such q should have a natural density. In the 1950s he adapted his conjecture for general $x \in \mathbf{Q}^*$, and this *Artin's primitive root conjecture* was proved by Hooley in 1967 under the assumption of the Generalized Riemann Hypothesis.

In this chapter, we will find a similar answer to the following analogue of the previous Question 1.1 in arbitrary number fields.

Question 1.2. *If K is a number field with $x \in K^*$, for how many primes $\mathfrak{q}$ of the ring of integers $\mathcal{O}_K$ of K with $\operatorname{ord}_{\mathfrak{q}}(x) = 0$ is $(\mathcal{O}_K/\mathfrak{q})^*$ generated by $x \bmod \mathfrak{q}$?*

Let us first turn back to the heuristic argument Artin used to derive his conjectural answer to Question 1.1. Since x is a primitive root modulo q exactly when $\operatorname{ord}_q(x) = 0$ and the index $[\mathbf{F}_q^* : \langle \bar{x} \rangle]$ is 1, one may determine for each prime p the set of primes q for which p divides $[\mathbf{F}_q^* : \langle \bar{x} \rangle]$, and remove these infinitely many sets from the set of all primes to see which primes q are left.

If p is a prime that divides the index $[\mathbf{F}_q^* : \langle \bar{x} \rangle]$, then p divides the group order $\#\mathbf{F}_q^*$, and $x \bmod q$ is an element of the index p subgroup of p-th powers in $\mathbf{F}_q^*$. In this case, $x \bmod q$ is the p-th power of p distinct elements of $\mathbf{F}_q^*$, so the polynomial $X^p - x$ splits completely into distinct linear factors modulo q.

Conversely, if $X^p - x$ splits completely into distinct linear factors modulo q, then x is a p-th power in $\mathbf{F}_q^*$ and has distinct p-th roots. So, $\mathbf{F}_q$ contains a primitive p-th root of unity, and p divides $[\mathbf{F}_q^* : \langle \bar{x} \rangle]$.

We conclude that for prime numbers p and q, we have

$$p \mid [\mathbf{F}_q^* : \langle \bar{x} \rangle] \iff X^p - x \text{ splits completely into distinct linear factors modulo } q.$$

In other words, the set of primes q modulo which x is a primitive root consists of those q (coprime to the numerator and denominator of x) that do not split completely in any of the splitting fields K_p of $X^p - x$, with $p \neq q$ prime.

The Frobenius density theorem (or alternatively the stronger Chebotarëv density theorem; see [32]) tells us that the set of primes q that split completely in K_p has a natural density of $\frac{1}{[K_p:\mathbf{Q}]}$. If x is not a p-th power, this is equal to $\frac{1}{p(p-1)}$.

Artin's heuristic argument now continues: at each prime p, the condition that p does not divide the index $[\mathbf{F}_q^* : \langle \bar{x} \rangle]$ excludes a fraction of $\frac{1}{[K_p:\mathbf{Q}]}$ of all primes q. Since for $x = 2$ the fields K_p are linearly disjoint over $\mathbf{Q}$ (embedding all K_p in a common algebraic closure $\bar{\mathbf{Q}}$), these conditions are independent, and it is reasonable to assume the set of primes modulo which 2 is a primitive root should have density

$$\prod_{p \text{ prime}} \left(1 - \frac{1}{p(p-1)}\right) \approx 0.3739558\ldots, \tag{1.3}$$

a value that is now known as Artin's constant.

When computers became sufficiently powerful in the 1950s to verify this density empirically, Derrick and Emma Lehmer computed for a few small integers x the number of primes $q < 20\,000$ modulo which x is a primitive root. For $x = 2$, the data matched Artin's conjectured density well. For $x = -3$ and $x = 5$ however, the observed densities were notably higher.

When Artin saw this, he realized that for general x, the conditions at the various primes p are not always independent. To see why, consider $x = 5$ and look at the splitting fields K_p of $X^p - 5$ over $\mathbf{Q}$ for primes p. We have $K_2 = \mathbf{Q}(\sqrt{5})$ and

$K_5 = \mathbf{Q}(\zeta_5, \sqrt[5]{5})$ where ζ_5 is a primitive fifth root of unity. In this case we see that K_2 is *contained* in K_5 because we have $\pm\sqrt{5} = \zeta_5 + \zeta_5^{-1} - \zeta_5^2 - \zeta_5^{-2}$ (The sign depends on the choice of ζ_5.) and therefore $\mathbf{Q}(\sqrt{5})$ is a subfield of $\mathbf{Q}(\zeta_5)$. Somewhat informally, we say that $\sqrt{5}$ and ζ_5 are *entangled radicals*.

If q is a prime that splits completely in K_5, then it also splits completely in the subfield K_2. In other words, the primes q that the condition at $p = 5$ means to exclude, have already all been excluded by the condition at $p = 2$. So, the factor $1 - \frac{1}{5(5-1)}$ in the infinite product must be omitted, and the density of primes modulo which 5 is a primitive root should be $\frac{20}{19}$ times Artin's constant.

It can be shown that in the case of Question 1.1, where the ground field is $\mathbf{Q}$, the only dependency between the splitting fields occurs when the discriminant d of K_2 is odd, in which case $K_2 = \mathbf{Q}(\sqrt{x})$ is contained in $\mathbf{Q}(\zeta_d)$ and thus also in K_d, the compositum of all K_p with $p \mid d$. Lang and Tate gave the necessary correction factor that results from this in 1965 in their preface to Artin's collected works [2]. To prove this corrected conjecture, one needs to show that imposing countably many splitting conditions does indeed give rise to a product density as in (1.3). If one uses the Generalized Riemann Hypothesis to bound the error terms in Frobenius' density theorem, this can be done in the way given by Hooley [16]. To date, there is no unconditional proof.

We now turn to Question 1.2 over a general number field K.

Just as for $K = \mathbf{Q}$, we need to describe the set of primes $\mathfrak{q}$ that do not split completely in any of the splitting fields K_p of $X^p - x$ over K. Imposing the splitting condition in *finitely many* K_p amounts to prescribing the splitting behaviour in the compositum K_n of these K_p (inside a fixed algebraic closure $\bar{K}$), where n is the product of the primes p considered.

More precisely, a prime $\mathfrak{q}$ does not split completely in any of the extensions $K \subset K_p$ with $p \mid n$ if and only if the Frobenius class $\mathrm{Frob}_\mathfrak{q}$ in $G_n = \mathrm{Gal}(K_n/K)$ is non-trivial when restricted to any of the subfields $K_p \subset K_n$. So, by the Chebotarëv density theorem the set of primes $\mathfrak{q}$ that do not split completely in any extension $K \subset K_p$ with $p \mid n$ has a density equal to the ratio $\#S_n/\#G_n$ with

$$S_n = \{\sigma \in G_n : \sigma\,|_{K_p} \neq \mathrm{id} \text{ for all } p \mid n\}.$$

Letting n in the ratio $\#S_n/\#G_n$ tend to the product of all primes then gives a conjectured density of primes $\mathfrak{q}$ modulo which x is a primitive root. Generalizing Hooley's work, Cooke and Weinberger showed in [9] that this conjectured density is the correct density when the Riemann Hypothesis holds for all fields K_n.

In the number field case, there are two complications: the group $\mathrm{Gal}(K_p/K)$ can be significantly harder to compute than for $K = \mathbf{Q}$. Moreover, as we have already seen over $\mathbf{Q}$, the Galois group $G_n = \mathrm{Gal}(K_n/K)$ can be a strict subgroup of the product $\prod_{p|n} \mathrm{Gal}(K_p/K)$, complicating the computation of this density. In the general case, $\mathrm{Gal}(K_n/K)$ can differ from the product of $\mathrm{Gal}(K_p/K)$ in many more ways than over $\mathbf{Q}$.

The fields K_n are "radical extensions" of K generated by all n-th roots of x. For such extensions, the Galois group G_n is a subgroup of the automorphism group of

the *multiplicative group* generated by these roots.

To make this precise, adjoin all n-th roots of x (in $\bar{K}$) to the multiplicative group K^*, resulting in the abelian group $B_n = \langle K^*, \zeta_n, \sqrt[n]{x} \rangle \subset \bar{K}^*$. The group extension $K^* \subset B_n$ is a much simpler structure than the field extension $K \subset K_n$. For example, adjoining a fifth root of unity to the abelian group $\mathbf{Q}^*$ does not also give a square root of 5.

We consider $G_n = \mathrm{Gal}(K_n/K)$ as a subgroup of the group $A_n = \mathrm{Aut}_{K^*}(B_n)$ of group automorphisms of B_n that are the identity on K^*. This larger automorphism group is much easier to compute than G_n itself. For one thing, unlike the Galois group, the group A_n does always factor as $\prod_{p|n} A_p$ (see Lemma 1.12).

Even though we ignored the additive structure of the fields involved, the difference between the Galois groups G_n and the groups A_n is actually quite modest, as reflected by the following theorem, proved in Section 1.4.

Theorem 1.4. *For all n, the Galois group G_n is a normal subgroup of A_n with finite, abelian quotient. There is a group $E = E_{K,x}$ such that for all n divisible by all of a finite set of critical primes, the quotient A_n/G_n equals E.*

This limit group E covers two things. For an individual prime p, the group G_p may be smaller than A_p, although Theorem 1.4 implies this only occurs for a finite number of primes. Additionally, E encodes the interdependencies between the local conditions at all primes p. The entanglement group E admits an explicit description that we derive in Section 1.3, and the set of critical primes in the theorem is given in Section 1.4. For example, when $K = \mathbf{Q}$, this set is empty unless the discriminant d of K_2 is odd, in which case it consists of the primes dividing $2d$.

The correction factor we need for the density statement has a transparent description in terms of the finitely many characters $\chi : E \to \mathbf{C}^*$ that "cut out" the Galois group G_n of K_n/K from the automorphism group A_n.

Theorem 1.5. *If the Generalized Riemann Hypothesis holds, the density of primes $\mathfrak{q}$ of K for which $(\mathcal{O}_K/\mathfrak{q})^*$ is generated by $x \bmod \mathfrak{q}$ exists and it is equal to*

$$C_{K,x} \cdot \prod_{p \text{ prime}} (1 - \frac{1}{\#A_p}), \text{ with } C_{K,x} = \sum_{\chi \in E^\vee} \prod_{\substack{p \text{ prime} \\ \chi(A_p) \neq 1}} \frac{-1}{\#A_p - 1}.$$

The proof of this theorem occupies most of the rest of this chapter. Afterwards, in Section 1.5 we give an explicit method to compute the rational correction factor $C_{K,x}$ with Lemma 1.13, and conclude with several examples.

1.2 Entanglement

In this section we will take a step back from the number theoretic view point of the previous section, and study the Galois group of normal, separable field extensions generated by radicals.

Formally, if $K \subset M$ is any field extension, we call a subgroup $B \subset M^*$ a *radical group* over K if B contains K^* and the quotient group B/K^* is torsion. This last condition means that every element of B has a power that is contained in K^*, or in other words, B consists of radicals over K. The field extension $K(B)$ of K is then called a *radical extension*.

The extensions $K \subset K_n$ from the previous section are examples of radical extensions, with $B_n = \langle K^*, \zeta_n, \sqrt[n]{x} \rangle$ as their generating radical groups. Here ζ_n denotes a primitive n-th root of unity in a fixed algebraic closure $\bar{K}$. We write $\mu_n \subset \bar{K}^*$ for the group of all n-th roots of unity of $\bar{K}$.

We will only consider radical groups B satisfying

$$\forall x \in B : \exists n \in \mathbf{Z}_{>0} : \operatorname{char} K \nmid n, x^n \in K^* \text{ and } \mu_n \subset B,$$

and call such groups *Galois* radical groups. The field extension $K(B)/K$ generated by such a group of radicals is separable due to the condition $\operatorname{char} K \nmid n$, and normal since we require B to contain sufficiently many roots of unity, so $K(B)/K$ is a Galois field extension.

Since any field automorphism of $K(B)$ is defined by its action on B, we can consider $\operatorname{Gal}(K(B)/K)$ as a subgroup of the group $\operatorname{Aut}_{K^*}(B)$ of group automorphisms of B that are the identity on K^*, also known as K^*-*automorphisms*.

The K^*-automorphism group genuinely depends on the generating radical group, and not just on the radical field extension it generates. For example, the radical group $B = \langle \mathbf{Q}^*, \zeta_5 \rangle$ over $\mathbf{Q}$ has $\mathbf{Q}^*$-automorphism group equal to the Galois group $\operatorname{Gal}(\mathbf{Q}(\zeta_5)/\mathbf{Q}) \cong (\mathbf{Z}/5\mathbf{Z})^*$. However, $B' = \langle \mathbf{Q}^*, \zeta_5, \sqrt{5} \rangle$, which generates the same field as B, has a $\mathbf{Q}^*$-automorphism group isomorphic to $\operatorname{Aut}_{\mathbf{Q}^*}(B) \times \operatorname{Aut}_{\mathbf{Q}^*}(\langle \mathbf{Q}^*, \sqrt{5} \rangle) \cong (\mathbf{Z}/5\mathbf{Z})^* \times C_2$. In this case, the Galois group $\operatorname{Gal}(\mathbf{Q}(B')/\mathbf{Q})$ is a normal subgroup of index 2 in $\operatorname{Aut}_{\mathbf{Q}^*}(B')$.

In two important cases, the K^*-automorphism group is equal to the Galois group of the generated field extension.

The first case is that of cyclotomic extensions of $\mathbf{Q}$. If μ is a multiplicative group of roots of unity of $\bar{\mathbf{Q}}$, the radical group $B = \langle \mathbf{Q}^*, \mu \rangle$ has $\mathbf{Q}^*$-automorphism group naturally isomorphic to $\operatorname{Aut}(\mu)$, since any automorphism of μ induces a $\mathbf{Q}^*$-automorphism of B. Any automorphism of μ also induces a field automorphism of $\mathbf{Q}(\mu)$, so here we find that $\operatorname{Gal}(\mathbf{Q}(B)/\mathbf{Q})$ is equal to $\operatorname{Aut}_{\mathbf{Q}^*}(B)$.

The second important case is the case of Kummer extensions. We call $x \in \bar{K}^*$ a *Kummer radical* over K if x^w is an element of K^* for some w with $\mu_w \subset K$. Radical extensions generated by Kummer radicals are called *Kummer extensions*. For instance, all Kummer extensions of $\mathbf{Q}$ are of the form $\mathbf{Q}(\sqrt{W})$ where we adjoin all square roots of elements of some set $W \subset \mathbf{Q}^*$.

For a group $B \subset \bar{K}^*$ of Kummer radicals, any K^*-automorphism σ of B multiplies each radical in B with a root of unity of K, and this fully determines σ. Writing μ_K for the set of roots of unity of K, the following therefore defines an

injective homomorphism:

$$\begin{aligned} \omega : \mathrm{Aut}_{K^*}(B) &\longrightarrow \mathrm{Hom}(B/K^*, \mu_K) \\ \sigma &\longmapsto \left(x \mapsto \frac{\sigma(x)}{x} \right). \end{aligned}$$

Kummer theory (see, e.g., [17], §VI.8) tells us the following composed map is an isomorphism:

$$\mathrm{Gal}(K(B)/K) \xrightarrow{\mathrm{res}} \mathrm{Aut}_{K^*}(B) \xrightarrow{\omega} \mathrm{Hom}(B/K^*, \mu_K).$$

We find that the natural restriction $\mathrm{Gal}(K(B)/K) \to \mathrm{Aut}_{K^*}(B)$ is an isomorphism.

The remainder of this section is devoted to proving the following main structure theorem.

Theorem 1.6. *If B is a Galois radical group over a field K, then the Galois group $\mathrm{Gal}(K(B)/K)$ is a normal subgroup of $\mathrm{Aut}_{K^*}(B)$ with an abelian quotient.*

We write $E(B)$ for this quotient group, and call it the *entanglement group* of B over K.

In both the cyclotomic case and the Kummer case, the automorphism group is abelian. Restricting to the maximal abelian subextension B_{ab} is the main ingredient of the proof of this theorem, and we can characterise B_{ab} for an arbitrary radical extension using the following theorem of Schinzel. We state this theorem and prove two lemmas before proceeding with the proof of Theorem 1.6.

Theorem 1.7. *Let F be a field, $a \in F$, and n a positive integer not divisible by $\mathrm{char}\, K$. Let w be the number of n-th roots of unity in F. Then, a splitting field of $X^n - a$ is abelian over F if and only if there exists $b \in F$ with $a^w = b^n$.*

Proof. See [28], or alternatively, Corollary 2.21 in the next chapter. □

Lemma 1.8. *If $C \subset D$ are two radical groups over K that are both Galois, any K^*-automorphism of C can be extended to an automorphism of D.*

Proof. Let $\varphi \in \mathrm{Aut}_{K^*}(C)$ be a K^*-automorphism of C. It follows from Zorn's lemma that the set of subgroups of D with an injective homomorphism to D that extends φ has a maximal element M with an injection $\psi : M \to D$.

To show that M is in fact equal to D, assume it is not, and take $x \in D \setminus M$. We will extend ψ to an injection $\langle M, x \rangle \to D$.

First of all, if x is a p-th root of unity, then $\langle M, x \rangle$ equals $M \oplus \mu_p$ and we can extend ψ with the identity on μ_p. This contradicts the fact that M is maximal, so M contains all torsion of D of prime order.

Otherwise, take the minimal $k \in \mathbf{Z}_{>1}$ such that $x^k \in M$ and the minimal $n \in \mathbf{Z}_{>1}$ such that $x^n \in K^*$. The injection ψ maps x^k to ζx^k for some $\zeta \in D$ with $\zeta^{n/k} = 1$. Since D is Galois, there exists $\xi \in D$ with $\xi^n = 1$ and $\xi^k = \zeta$. We can now define the injection $\psi' : \langle x \rangle \to D$ by $x \mapsto \xi x$. Since $\psi'(x^k)$ then equals $\psi(x^k)$, the injections ψ and ψ' are compatible on the intersection of M and $\langle x \rangle$.

The group $\langle M, x\rangle \subset D$ can be written as a fibered sum (or push-out):

$$\langle M, x\rangle \cong M \oplus_{\langle x^k\rangle} \langle x\rangle.$$

The pair of injections $\psi : M \to D$ and $\psi' : \langle x\rangle \to D$ together with the universal property of this push-out now defines a homomorphism $\chi : \langle M, x\rangle \to D$ that extends φ.

We claim that χ is injective. Since χ multiplies all elements by torsion elements, the kernel of χ is torsion. However, all elements ζ of D of prime order are contained in M, so we have $\chi(\zeta) = \psi(\zeta) \neq 1$. The kernel of χ is therefore trivial, so χ is an injective homomorphism $\langle M, x\rangle \to D$. This contradicts the maximality of M, so M is equal to D.

The injection $\psi : M \to D$ is now necessarily an automorphism, since for any $n > 0$ and $x \in D$ it permutes the finitely many n-th roots of x. □

Let B be a Galois radical extension over K. We will write B_{tors} for the subgroup of torsion elements of B.

Lemma 1.9. *Let x be an element of B and n the minimal positive integer such that $x^n \in K^*$. Then the following are equivalent:*

1. $\exists w \in \mathbf{Z}_{>0} : x^w \in B_{\text{tors}}K^*$ and $\mu_w \subset K^*$;
2. $\text{Aut}_{K^*}(\langle K^*, \zeta_n, x\rangle)$ *is abelian;*
3. $\text{Gal}(K(\zeta_n, x)/K)$ *is abelian.*

Proof.

(1) ⇒ (2). If $x^w \in B_{\text{tors}}K^*$ for a positive integer w with $\mu_w \subset K^*$, then $\langle K^*, \zeta_n, x\rangle$ is a subset of $B' = \mu_{\bar{K}}\sqrt[w]{K^*}$. Since any K^*-automorphism of B' sends roots of unity to roots of unity and w-th roots of elements of K^* to w-th roots of elements of K^*, there are restriction maps $\text{Aut}_{K^*}(B') \to \text{Aut}(\mu_{\bar{K}})$ and $\text{Aut}_{K^*}(B') \to \text{Aut}_{K^*}(\sqrt[w]{K^*})$. This implies there is an injective homomorphism from $\text{Aut}_{K^*}(B')$ to $\text{Aut}(\mu_{\bar{K}}) \times \text{Aut}_{K^*}(\sqrt[w]{K^*})$, which, as we saw, is abelian. Finally, by Lemma 1.8, the restriction map from $\text{Aut}_{K^*}(B')$ to $\text{Aut}_{K^*}(\langle K^*, \zeta_n, x\rangle)$ is surjective, and therefore the latter is also abelian.

(2) ⇒ (3). This is trivial since $\text{Gal}(K(\zeta_n, x)/K)$ can be considered a subgroup of $\text{Aut}_{K^*}(\langle K^*, \zeta_n, x\rangle)$.

(3) ⇒ (1). Since $K(\zeta_n, x)$ is a splitting field of $X^n - x^n$ over K, by Schinzel's Theorem 1.7 there is an element $b \in K$ with $x^{nw} = b^n$ if we take w to be the number of n-th roots of unity in K. Then we have $x^w \in \mu_n b$, proving the lemma. □

We are now ready to prove Theorem 1.6. The group of radicals

$$B_{\text{ab}} = \{x \in B : \exists w : x^w \in B_{\text{tors}}K^* \text{ and } \mu_w \subset K^*\},$$

consisting of the elements of B satisfying the conditions from Lemma 1.9, has an abelian group of K^*-automorphisms $\text{Aut}_{K^*}(B_{\text{ab}})$. Since B_{ab} contains all roots of unity in B, any K^*-automorphism of B maps B_{ab} into itself. So, there is a well-defined restriction map $\text{Aut}_{K^*}(B) \to \text{Aut}_{K^*}(B_{\text{ab}})$ with kernel $\text{Aut}_{B_{\text{ab}}}(B)$, which is surjective by Lemma 1.8.

Thus, we get the following exact sequence.

$$0 \to \mathrm{Aut}_{B_{\mathrm{ab}}}(B) \to \mathrm{Aut}_{K^*}(B) \xrightarrow{\mathrm{res}} \mathrm{Aut}_{K^*}(B_{\mathrm{ab}}) \to 0.$$

This sequence is the K^*-automorphism equivalent of the exact sequence of Galois groups of the tower of extensions $K \subset K(B_{\mathrm{ab}}) \subset K(B)$. Combining the two gives the following diagram, where the rows are exact and the vertical arrows are injective. Since the only maps involved are natural injections and restrictions, the squares are both commutative.

$$\begin{array}{ccccccccc}
0 & \longrightarrow & \mathrm{Gal}(K(B)/K(B_{\mathrm{ab}})) & \longrightarrow & \mathrm{Gal}(K(B)/K) & \xrightarrow{\pi} & \mathrm{Gal}(K(B_{\mathrm{ab}})/K) & \longrightarrow & 0 \\
 & & \downarrow f & & \downarrow g & & \downarrow h & & \\
0 & \longrightarrow & \mathrm{Aut}_{B_{\mathrm{ab}}}(B) & \longrightarrow & \mathrm{Aut}_{K^*}(B) & \xrightarrow{\pi'} & \mathrm{Aut}_{K^*}(B_{\mathrm{ab}}) & \longrightarrow & 0
\end{array}$$

We can now finish the proof of Theorem 1.6. On the left side of the diagram, the radical extension $K(B)/K(B_{\mathrm{ab}})$ is a Kummer extension since B_{ab} contains all roots of unity of B. Therefore, the image of the Galois group $\mathrm{Gal}(K(B)/K(B_{\mathrm{ab}}))$ under f is the image of the restriction map

$$\mathrm{Aut}_{K(B_{\mathrm{ab}})^*}(K(B_{\mathrm{ab}})^*B) \to \mathrm{Aut}_{B \cap K(B_{\mathrm{ab}})^*}(B).$$

We claim that this restriction is a surjection. To see this, choose any automorphism $\sigma \in \mathrm{Aut}_{B \cap K(B_{\mathrm{ab}})^*}(B)$. We have that $K(B_{\mathrm{ab}})^*B$ is the following fibered sum:

$$K(B_{\mathrm{ab}})^*B = K(B_{\mathrm{ab}})^* \oplus_{(B \cap K(B_{\mathrm{ab}})^*)} B.$$

The automorphism σ induces an injective homomorphism $\varphi_\sigma : B \to K(B_{\mathrm{ab}})^*B$. By the universal property of the fibered sum, the injection φ_σ together with the inclusion $K(B_{\mathrm{ab}})^* \subset K(B_{\mathrm{ab}})^*B$ induces an automorphism of $K(B_{\mathrm{ab}})^*B$ that is the identity on $K(B_{\mathrm{ab}})^*$ and that extends σ. This proves the claim.

Because $K(B_{\mathrm{ab}})$ is abelian over K, the intersection $B \cap K(B_{\mathrm{ab}})^*$ is contained in B_{ab} by Lemma 1.9. Of course, B_{ab} is also contained in $B \cap K(B_{\mathrm{ab}})^*$, so we have $B_{\mathrm{ab}} = B \cap K(B_{\mathrm{ab}})^*$ and we conclude that f is an isomorphism.

On the right side of the diagram, $\mathrm{Aut}_{K^*}(B_{\mathrm{ab}})$ is abelian, so π' maps the commutator subgroup H of $\mathrm{Aut}_{K^*}(B)$ to 1, so H is contained in $\mathrm{Aut}_{B_{\mathrm{ab}}}(B)$. Since f is a surjection, this implies H is in fact contained in the image of $\mathrm{Gal}(K(B)/K)$. From this we can directly conclude that the image of $\mathrm{Gal}(K(B)/K)$ is a normal subgroup of $\mathrm{Aut}_{K^*}(B)$ with an abelian cokernel, which concludes the proof of Theorem 1.6.

1.3 Computing the entanglement groups

Let B again be a Galois radical group over a field K. In the previous section, we have seen that the entanglement group of B over K is equal to the entanglement group $E(B_{\mathrm{ab}})$ of the subgroup B_{ab} of B. As mentioned in the proof of Lemma 1.9, this group B_{ab} is a subgroup of the group generated by all Kummer radicals and all roots of unity in $\bar{K}$. If B_{ab} is *itself* generated by Kummer radicals and roots of unity and if the characteristic of K is 0, there is an explicit way to describe $E(B_{\mathrm{ab}})$ as a Galois group, which we give in this section.

All the Galois radical groups that play a role for the results of this chapter are of this form, but this is not the case in general. As an example, let α be a fourth root of -4 in $\bar{\mathbf{Q}}$ and consider the radical extension $\langle \mathbf{Q}^*, \alpha \rangle$ over $\mathbf{Q}$. This is a Galois radical extension since $\frac{1}{2}\alpha^2$ is a primitive 4th root of unity. Its automorphism group is of order 4 and abelian, but $\langle \mathbf{Q}^*, \alpha \rangle$ is not generated over $\mathbf{Q}^*$ by a Kummer radical or a root of unity. We will see in Chapter 5 how to handle this case.

For radical groups over fields of non-zero characteristic we refer to the more general treatment in Chapter 2 and Chapter 3, and in particular Section 3.4.

Now suppose that K is of characteristic 0 and that $B_{\mathrm{ab}} = \mu W$ is a group generated by a group of roots of unity μ and a group of Kummer radicals $W \supset K^*$. Since the Galois group $\mathrm{Gal}(K(\mu W)/K)$ is the kernel of the homomorphism $\mathrm{Aut}(\mu W) \to E(\mu W)$, the image of a group automorphism σ in $E(\mu W)$ determines whether or not σ is the restriction of a field automorphism.

In the examples in the previous section, we saw that any K^*-automorphism of W can be uniquely extended to a field automorphism of $K(W)$, and any automorphism of μ can be uniquely extended to a field automorphism of $\mathbf{Q}(\mu)$. To determine if a given element $\sigma \in \mathrm{Aut}(\mu W)$ is the restriction of an element of $\mathrm{Gal}(K(\mu W)/K)$, it therefore makes sense to compare the obtained field automorphisms of $K(W)$ and $\mathbf{Q}(\mu)$, and see if they are compatible.

To this end, we define the homomorphisms φ_1 and φ_2 as follows.

$$\begin{array}{rcl} \varphi_1 : \mathrm{Aut}_{K^*}(\mu W) & \xrightarrow{\text{res}} & \mathrm{Aut}_{K^*}(W) \xrightarrow{\sim} \mathrm{Gal}(K(W)/K) \\ \varphi_2 : \mathrm{Aut}_{K^*}(\mu W) & \xrightarrow{\text{res}} & \mathrm{Aut}_{\mu \cap K^*}(\mu) \xrightarrow{\sim} \mathrm{Gal}(\mathbf{Q}(\mu)/\mathbf{Q}(\mu \cap K^*)) \end{array}$$

Finally, we write $K_0 = K(W) \cap \mathbf{Q}(\mu)$. This is an abelian Galois field extension of $\mathbf{Q}$ since $\mathbf{Q}(\mu)/\mathbf{Q}$ is abelian. We then define φ as the difference of φ_1 and φ_2:

$$\begin{array}{rcl} \varphi : \mathrm{Aut}_{K^*}(\mu W) & \longrightarrow & \mathrm{Gal}(K_0/\mathbf{Q}) \\ \sigma & \longmapsto & \varphi_1(\sigma)|_{K_0} \cdot \varphi_2(\sigma)|_{K_0}^{-1}. \end{array}$$

Because $\mathbf{Q}(\mu)$ is abelian over $\mathbf{Q}$, the subextension $K_0/\mathbf{Q}$ is also abelian, so φ is a group homomorphism. Furthermore, for $\sigma \in \mathrm{Aut}_{K^*}(\mu W)$, both $\varphi_1(\sigma)|_{(\mu \cap W)}$ and $\varphi_2(\sigma)|_{(\mu \cap W)}$ are equal to $\sigma|_{(\mu \cap W)}$. Because φ is the difference of the two, $\varphi(\sigma)$ restricts to the identity on $\mu \cap W$. We see that the image of φ is in fact contained in $\mathrm{Gal}(K_0/\mathbf{Q}(W \cap \mu))$.

Theorem 1.10. *The homomorphism φ induces an isomorphism*

$$\psi : E(\mu W) \xrightarrow{\sim} \mathrm{Gal}(K_0/\mathbf{Q}(W \cap \mu)).$$

Proof. For ψ to be well-defined and injective, we show that the kernel of φ equals $\mathrm{Gal}(K(\mu W)/K)$.

We can write $\mathrm{Aut}_{K^*}(\mu W)$ as a fibered product:

$$\mathrm{Aut}_{K^*}(\mu W) \cong \mathrm{Aut}_{\mu \cap K^*}(\mu) \times_{\mathrm{Aut}_{\mu \cap K^*}(\mu \cap W)} \mathrm{Aut}_{K^*}(W).$$

The two factors are naturally isomorphic to $\mathrm{Gal}(\mathbf{Q}(\mu)/\mathbf{Q}(\mu \cap K^*))$ respectively $\mathrm{Gal}(K(W)/K)$. Using this structure, an element of $\mathrm{Aut}_{K^*}(\mu W)$ can be uniquely represented by a pair (σ, τ) with $\sigma \in \mathrm{Gal}(\mathbf{Q}(\mu)/\mathbf{Q}(\mu \cap K^*))$ and $\tau \in \mathrm{Gal}(K(W)/K)$. By construction of φ, the pair (σ, τ) is in the kernel of φ if and only if $\sigma|_{K_0}$ equals $\tau|_{K_0}$.

We now observe that $\mathrm{Gal}(K(\mu W)/K)$ admits the following fibered product structure:

$$\mathrm{Gal}(K(\mu W)/K) \cong \mathrm{Gal}(K(W)/K) \times_{\mathrm{Gal}(K(W) \cap K(\mu)/K)} \mathrm{Gal}(K(\mu)/K)$$

Composing this with the restriction $\mathrm{Gal}(K(\mu)/K) \xrightarrow{\sim} \mathrm{Gal}(\mathbf{Q}(\mu)/K \cap \mathbf{Q}(\mu))$ results in an isomorphism

$$\mathrm{Gal}(K(\mu W)/K) \xrightarrow{\sim} \mathrm{Gal}(K(W)/K) \times_{\mathrm{Gal}(K_0/\mathbf{Q}(\mu) \cap K)} \mathrm{Gal}(\mathbf{Q}(\mu)/\mathbf{Q}(\mu) \cap K)$$

Therefore, we see that the pair (σ, τ) extends to an automorphism of the field $\mathrm{Gal}(K(\mu W)/K)$ if and only if $\sigma|_{K_0}$ equals $\tau|_{K_0}$. This implies that the Galois group $\mathrm{Gal}(K(\mu W)/K)$ is precisely the kernel of φ and therefore that ψ is well-defined and injective.

For the surjectivity of ψ, we show that $\psi(\mathrm{Aut}_W(\mu W))$ already gives the full image $\mathrm{Gal}(K_0/\mathbf{Q}(W \cap \mu))$. Note that $\varphi_1(\mathrm{Aut}_W(\mu W))$ is trivial, so we only need to follow $\mathrm{Aut}_W(\mu W)$ through the composite map φ_2. The restriction map $\mathrm{Aut}_W(\mu W) \to \mathrm{Aut}_{(W \cap \mu)}(\mu)$ is surjective because μW is the fibered sum of μ and W over $\mu \cap W$, using the same argument as in the proof of Theorem 1.6. Its image in $\mathrm{Gal}(\mathbf{Q}(\mu)/\mathbf{Q})$ then equals $\mathrm{Gal}(\mathbf{Q}(\mu)/\mathbf{Q}(W \cap \mu))$. So, $\varphi_2(\mathrm{Aut}_W(\mu W))$ is equal to the Galois group $\mathrm{Gal}(\mathbf{Q}(\mu)/\mathbf{Q}(\mu \cap W))$.

Finally, we see that ψ is surjective since the restriction $\varphi_2(\mathrm{Aut}_W(\mu W))|_{K_0}$ equals $\mathrm{Gal}(K_0/\mathbf{Q}(W \cap \mu))$. □

1.4 Proof of main results

In this section we tie together the results from the previous two sections to prove Theorem 1.4 and the main theorem of this chapter, Theorem 1.5.

Let K be a number field, and $x \in K$ non-zero. As described in the introduction, a main ingredient of our computation of the density of primes $\mathfrak{q}$ of K for which $(\mathcal{O}_K/\mathfrak{q})^*$ is generated by x mod $\mathfrak{q}$, is determining for squarefree positive integers n the entanglement groups of the Galois radical groups $B_n = \langle K^*, \zeta_n, \sqrt[n]{x} \rangle$ over K.

We will prove Theorem 1.4 by directly proving the following more explicit variant.

Theorem 1.11. *Let w be the product of all primes p with $\zeta_p \in K$, and W the group of radicals $\langle K^*, \sqrt[w]{x}\rangle$. Let K_0 be the maximal subfield of $K(W)$ that is abelian over $\mathbf{Q}$, tamely ramified, and unramified at all primes $p \mid w$. Let n be a positive squarefree integer divisible by w and by all primes ramified in K_0, and write $r = n/w$. Then with $\mu = \mu_r$, the map defined in Theorem 1.10 induces an isomorphism from $E(B_n)$ to the finite abelian group $E = \mathrm{Gal}(K_0/\mathbf{Q})$.*

Proof. First, we show that we can apply Theorem 1.10 to compute the entanglement group $E(B_n) = E(B_{n,\mathrm{ab}})$. We will show that $B_{n,\mathrm{ab}}$ equals $\mu_r W$.

Since n is squarefree, $B_{\mathrm{tors}}K^*$ equals $\mu_r K^*$. If y is an element of $B_{n,\mathrm{ab}}$, then we have $y^w \in B_{\mathrm{tors}}K^* = \mu_r K^*$. Since w and r are coprime, this makes y the product of an element of μ_r and a w-th root of an element of K^*.

Theorem 1.10 then gives us an explicit isomorphism from $E(\mu_r W)$ to the Galois group $\mathrm{Gal}(K(W)\cap\mathbf{Q}(\mu_r)/\mathbf{Q}(\mu_r \cap W))$. Note that since r and w are coprime, $W \cap \mu_r$ is trivial, so we find an isomorphism

$$E(B_n) \xrightarrow{\sim} \mathrm{Gal}(K(W) \cap \mathbf{Q}(\mu_r)/\mathbf{Q}).$$

Since for a prime p, the field $\mathbf{Q}(\mu_p)$ is only (tamely) ramified at p, the field K_0 defined in this theorem is equal to $K(W) \cap \mathbf{Q}(\hat{\mu})$ where $\hat{\mu}$ is the group generated by primitive p-th roots of unity for all $p \nmid w$. We see that because n is divisible by w and by all primes ramified in $K_0/\mathbf{Q}$, we have $K(W) \cap \mathbf{Q}(\hat{\mu}) = K(W) \cap \mathbf{Q}(\mu_r)$, so $E(B_n)$ is isomorphic to $\mathrm{Gal}(K_0/\mathbf{Q})$.

Since $K(W)$ has finite degree over $\mathbf{Q}$, so does the subfield K_0, and we conclude that the limit entanglement group $E = \mathrm{Gal}(K_0/\mathbf{Q})$ is finite. □

Theorem 1.4 is now a direct corollary of Theorem 1.6 and Theorem 1.11. To derive the explicit formula for the density, we need one last ingredient.

Lemma 1.12. *For every squarefree positive integer n, there is a natural isomorphism*

$$\mathrm{Aut}_{K^*}(B_n) \cong \prod_{p \mid n \text{ prime}} \mathrm{Aut}_{K^*}(B_p).$$

Proof. Let n be a squarefree positive integer. Since B_p is a Galois radical group, there is a natural restriction map from $A_n = \mathrm{Aut}_{K^*}(B_n)$ to $A_p = \mathrm{Aut}_{K^*}(B_p)$ for every prime $p \mid n$. Since B_n is generated by all B_p with $p \mid n$, the combined map $\varphi : A_n \to \prod A_p$ is an injection.

The restriction maps $A_n \to A_p$ are surjective by Lemma 1.8. To see that the map to $\prod A_p$ is also surjective, let $(\sigma_p)_p$ be an element of $\prod A_p$. We construct $\sigma \in A_n$ with $\varphi(\sigma) = (\sigma_p)_p$ as follows: define $\sigma : B_n \to B_n$ by $\prod b_p \mapsto \prod \sigma_p(b_p)$ (with $b_p \in B_p$). Since every element of B_n can be uniquely written as $\prod b_p$ (up to multiplication with elements of K^*), the map σ is a well-defined homomorphism. It is invertible because its inverse is given by applying the same procedure to $(\sigma_p^{-1})_p$. We see that σ is contained in A_n and $\varphi(\sigma)$ is indeed $(\sigma_p)_p$. □

This factorization of A_n into a product of A_p for $p \mid n$ now allows us to prove the main density theorem.

Proof of Theorem 1.5. Recall that the density (under GRH, as described in the introduction) is given by the limit of $\#S_n/\#G_n$ when we let n tend to all primes. Let n therefore be a squarefree positive integer that is large enough for $E(B_n)$ to equal the (finite) entanglement group $E = \mathrm{Gal}(K_0/\mathbf{Q})$, as defined by Theorem 1.11.

Also recall the definition of S_n:

$$S_n = \{\sigma \in G_n : \text{ for all } p \mid n : \sigma\,|_{K_p} \neq \mathrm{id}\}.$$

As an analogue of S_n inside the K^*-automorphism group A_n, define

$$T_n = \{\sigma \in A_n : \text{ for all } p \mid n : \sigma\,|_{B_p} \neq \mathrm{id}\}.$$

We then have $S_n = T_n \cap G_n$ inside A_n. Also, under the natural isomorphism of A_n with $\prod A_p$, the subset T_n is mapped to $\prod A_p \setminus \{1\}$. Now we rewrite $\#S_n/\#G_n$ as follows, using the characteristic function 1_{G_n} of G_n inside A_n.

$$\delta_n = \frac{\#S_n}{\#G_n} = \frac{\#(T_n \cap G_n)}{\#G_n} = \frac{\sum_{s \in T_n} 1_{G_n}(s)}{\#G_n}$$

Exploiting the fact that E is abelian, we can rewrite 1_{G_n}.

$$\delta_n = \frac{1}{\#G_n \#E} \sum_{s \in T_n} \sum_{\chi \in E^\vee} \chi(s)$$

Under the natural isomorphism of A_n with $\prod A_p$, the subset T_n is mapped to the product $\prod A_p \setminus \{1\} = \prod T_p$.

$$\begin{aligned}\delta_n &= \frac{1}{\#A_n} \sum_{\chi \in E^\vee} \prod_{p \mid n} \sum_{s_p \in T_p} \chi(s_p) = \frac{1}{\#A_n} \sum_{\chi \in E^\vee} \prod_{p \mid n} \left(-1 + \sum_{s_p \in A_p} \chi(s_p) \right) \\ &= \frac{\#T_n}{\#A_n} \sum_{\chi \in E^\vee} \prod_{p \mid n} \frac{1}{\#T_p} \left(-1 + \sum_{s_p \in A_p} \chi(s_p) \right)\end{aligned}$$

Because $\sum_{s_p \in A_p} \chi(s_p)$ equals $\#A_p$ if $\chi(A_p)$ is trivial, and 0 otherwise, we get:

$$\begin{aligned}\delta_n &= \frac{\#T_n}{\#A_n} \sum_{\chi \in E^\vee} \prod_{\substack{p \text{ prime} \\ \chi(A_p) \neq 1}} \frac{-1}{\#A_p - 1} \\ &= \prod_{p \mid n} \left(1 - \frac{1}{\#A_p}\right) \sum_{\chi \in E^\vee} \prod_{\substack{p \text{ prime} \\ \chi(A_p) \neq 1}} \frac{-1}{\#A_p - 1} \\ &= C_{K,x} \prod_{p \mid n} \left(1 - \frac{1}{\#A_p}\right)\end{aligned}$$

Taking the limit of n to infinity now gives the desired formula. □

1.5 Explicit densities

Computing the correction factor $C_{K,x}$ explicitly requires determining for each character $\chi \in E^\vee$ for which primes p the image $\chi(A_p)$ is non-trivial. We call a character χ *bad* at p in this case. We start this section with an explicit criterion for determining at which primes characters are bad.

Lemma 1.13. *A character $\chi \in E^\vee$ is bad at a prime p if and only if one of the following two conditions holds:*

1. *χ is ramified at p, or*
2. *$\zeta_p \in K$ and $x \notin K^{*p}$ and the restriction to K_0 of the K-automorphism of $K(\sqrt[w]{x})$ given by $\sqrt[w]{x} \mapsto \zeta_p \sqrt[w]{x}$ is not in the kernel of χ.*

Proof. Let n be large enough for $E(B_n)$ to be isomorphic to E. Recall that the isomorphism $\varphi : A_n/G_n \xrightarrow{\sim} \mathrm{Gal}(K_0/\mathbf{Q})$ is given by $\varphi(\sigma) = \varphi_1(\sigma)\varphi_2(\sigma)^{-1}$ as defined in Section 1.3.

First, let p be a prime. We claim that at least one of $\varphi_1(A_p)$ and $\varphi_2(A_p)$ is trivial. If this is not the case, then both $\mathrm{Aut}_{K^*}(B_p \cap W)$ and $\mathrm{Aut}_{\mu \cap K^*}(B_p \cap \mu_n)$ are non-trivial. This first condition implies that B_p/K^* has a non-trivial element of order dividing w, and the second that B_p/K^* has an element of order not dividing w. This contradicts the fact that B_p/K^* has prime exponent p.

Now let χ be an element of $E^\vee$.

Assume that $\varphi_1(A_p)$ is non-trivial. Then $\varphi_2(A_p)$ is trivial, and χ is bad precisely if $\varphi_1(A_p)\,|_{K_0} = \mathrm{Gal}(K(W) \cap \mathbf{Q}(\mu_p)/\mathbf{Q})$ is not contained in the kernel of χ. This in turn is equivalent to the first condition from this lemma.

Alternatively, assume that $\varphi_2(A_p)$ is non-trivial. In that case, p divides w and $\varphi_1(A_p)$ is trivial. Since K contains ζ_p, the image $\varphi_2(A_p)\,|_{K_0}$ is trivial if x is a p-th power in K. Otherwise, it is generated by the automorphism $\sqrt[w]{x} \mapsto \zeta_p \sqrt[w]{x}$. We see that in this case, χ is bad precisely if the second condition from the lemma holds. □

We start by comparing our results with the known (under GRH) densities in the classical case over $\mathbf{Q}$.

Consider $x = 2$. Since $\mathbf{Q}$ only contains 2 roots of unity, we have $w = 2$. Using Theorem 1.11, we see that the limit entanglement group E is $\mathrm{Gal}(K_0/\mathbf{Q})$ for K_0 the maximal subfield of $\mathbf{Q}(\sqrt{2})$ that is abelian over $\mathbf{Q}$, tamely ramified, and unramified at 2. Since $\mathbf{Q}(\sqrt{2})$ is ramified at 2, K_0 equals $\mathbf{Q}$, and E is trivial, so the correction factor $C_{\mathbf{Q},2}$ is 1. This results in a density of

$$\prod_{p \text{ prime}} \left(1 - \frac{1}{p(p-1)}\right) \approx 0.373955\ldots,$$

as expected.

Next, we consider $K = \mathbf{Q}$ and $x = 5$. In this case $K_0 = \mathbf{Q}(\sqrt{5})$ and $E = \mathrm{Gal}(K_0/\mathbf{Q})$ is a group of order two. The non-trivial character in $E^\vee$ is only ramified at 5, so it is bad at 5 and potentially at primes dividing $w = 2$. Since $K_0 = K(\sqrt{5})$,

the automorphism sending $\sqrt{5}$ to $-\sqrt{5}$ is clearly not in the kernel of χ, so χ is indeed bad at 2.

We find $a_2 = 2$ and $a_5 = 20$, so we have $C_{\mathbf{Q},5} = 1 + \frac{1}{19} = \frac{20}{19}$. This leads to a conjectured density of $\frac{20}{19}$ times Artin's constant, as was also observed by the Lehmers.

As an example where K is larger than the rationals, consider the case $K = \mathbf{Q}(\sqrt{-7})$ and $x = 21$. Since K doesn't contain any new roots of unity, w equals 2 as before. The field $K(\sqrt[w]{x}) = \mathbf{Q}(\sqrt{-3}, \sqrt{-7})$ is now abelian and only (tamely) ramified above 3 and 7, so $K_0 = \mathbf{Q}(\sqrt{-3}, \sqrt{-7})$. We now need to determine at which primes the four characters in $\mathrm{Gal}(K_0/\mathbf{Q})^\vee$ are bad. We write $\mathrm{Gal}(K_0/\mathbf{Q})^\vee = \{\mathrm{id}, \chi_{-3}, \chi_{-7}, \chi_{21}\}$ where χ_t is such that $\mathbf{Q}(\sqrt{t})$ is the invariant field of $\ker \chi_t$. This yields the following table, where the minus signs indicate at which primes the characters are bad.

	2	3	7
1	+	+	+
χ_{-3}	−	−	+
χ_{-7}	+	+	−
χ_{21}	−	−	−

Since $a_2 = 2$, $a_3 = 6$ and $a_7 = 42$, we find $C_{K,x} = 1 + \frac{1}{5} - \frac{1}{41} - \frac{1}{5 \cdot 41} = \frac{48}{41}$.

To verify this empirically, the following table lists approximations to the density computed with Sage [30]. For each N in the table, it lists the fraction of primes $\mathfrak{q}$ with norm smaller than N for which $\bar{x} = \overline{21}$ is a primitive root modulo $\mathfrak{q}$.

N	density
10^5	0.443679...
10^6	0.436286...
10^7	0.437864...
10^8	0.437940...
10^9	0.437870...
10^{10}	0.437818...
conjectured	0.437801...

The conjectured density matches the observed approximations well.