LECTURE NOTES FOR COURSE 88-909
COHOMOLOGY OF GROUPS
BAR-ILAN UNIVERSITY, SPRING 5782

MICHAEL M. SCHEIN

Before starting, we'll say a few words of motivation. We wish to understand the infinite Galois group $G = \mathrm{Gal}(\overline{\mathbb{Q}}/\mathbb{Q})$, which is very important in number theory. In particular, we wish to understand G-modules, namely modules over some ring that are endowed with an action of G. Such objects tend to be complicated. Recall that in algebraic topology, we define the homology and cohomology of simplicial complexes. These are collections of invariants that can be used, for instance, to prove that two simplicial complexes are not homeomorphic. One loses information when passing from a complex to its cohomology – two different complexes can have the same cohomology groups in every dimension – but retains enough to do some useful things. On the other hand, the cohomology groups have the advantage of sometimes being computable; for simple complexes they can be computed directly, and for more complicated ones they can often be deduced with tools such as the long exact cohomology sequence, the Mayer-Vietoris sequence, etc. The aim of this course is to develop an analogous theory in the setting of G-modules.

1. Profinite groups

Definition 1.1. Let I be a directed partially ordered set. This means that every pair of elements has a common upper bound: if $i, j \in I$ then there exists $k \in I$ such that $i \leq k$ and $j \leq k$. A *projective system* of groups indexed by I consists of a collection of groups $\{G_i : i \in I\}$, and, for each pair $i \geq j$ of elements of I, a group homomorphism $\varphi_{ij} : G_i \to G_j$ such that φ_{ii} is always the identity. Moreover, we demand that these maps be compatible, in the sense that if $i \geq j \geq k$ then $\varphi_{ik} = \varphi_{jk} \circ \varphi_{ij}$.

Definition 1.2. Let $\{G_i\}_{i \in I}$ be a projective system. Its projective limit $\varprojlim_I G_i$ is the subgroup of $\prod_{i \in I} G_i$ consisting of "compatible" tuples $(a_i)_{i \in I}$, namely those satisfying $\varphi_{ij}(a_i) = a_j$ for all $i > j$.

If the G_i are topological groups and the homomorphisms φ_{ij} are continuous maps, then we can define the *projective limit topology* on $\varprojlim_I G_i$ as the weakest topology such that the maps

$$\begin{aligned} \pi_i : \varprojlim_I G_i &\to G_i \\ (a_j)_j &\mapsto a_i \end{aligned}$$

are continuous for all $i \in I$.

Proposition 1.3. *The projective limit topology on $\varprojlim_I G_i$ coincides with the subspace topology induced from the product topology on $\prod_{i \in I} G_i$.*

Date: June 20, 2022.

Proof. A base of open subsets of the projective limit topology is given by

$$\{\pi_i^{-1}(U_i) : i \in I, U_i \subset G_i \text{ open}\}.$$

Observe that $\pi_i^{-1}(U_i) = (\varprojlim_I G_i) \cap (U_i \times \prod_{j \in I \setminus \{i\}} G_j)$ is open in the subspace topology.

Conversely, a base of the subspace topology is given, by definition, by sets of the form $U = (\varprojlim_I G_i) \cap \prod_{i \in S} U_i \times \prod_{i \in I \setminus S} G_i$, where $S \subset I$ is finite and $U_i \subset G_i$ is open. Clearly $U = \bigcap_{i \in S} \pi_i^{-1}(U_i)$, which is open in the projective limit topology. □

For our purposes, the G_i will usually be finite groups endowed with the discrete topology.

As an example, let I be the set of all finite Galois extensions of $\mathbb{Q}$, ordered by inclusion. Then $\{\mathrm{Gal}(K/\mathbb{Q})\}_{K \in I}$ is a projective system, with the obvious homomorphisms. We observe that $\mathrm{Gal}(\overline{\mathbb{Q}}/\mathbb{Q}) \simeq \varprojlim_I \mathrm{Gal}(K/\mathbb{Q})$. Indeed, define a homomorphism $f : \mathrm{Gal}(\overline{\mathbb{Q}}/\mathbb{Q}) \to \varprojlim_I \mathrm{Gal}(K/\mathbb{Q})$ by $f(\sigma) = (\sigma_{|K})_{K \in I}$. Since every element of $\overline{\mathbb{Q}}$ lies in some finite Galois extension of $\mathbb{Q}$, we see that f is injective. Moreover, given an element $(\sigma_K)_K$ of the projective limit, define $\sigma \in \mathrm{Gal}(\overline{\mathbb{Q}}/\mathbb{Q})$ by $\sigma(\alpha) = \sigma_K(\alpha)$, where $\alpha \in \overline{\mathbb{Q}}$ and K is a finite Galois extension of $\mathbb{Q}$ containing α. The compatibility condition of the projective limit exactly ensures that the σ_K glue and σ is well-defined.

Since we are interested in representations of $\mathrm{Gal}(\overline{\mathbb{Q}}/\mathbb{Q})$, it is natural to study projective limits of finite groups. It will be convenient to have an alternative characterization of such groups.

Definition 1.4. A topological space X is called *totally disconnected* if its only non-empty connected subsets are single points. It is called *totally separated* if for any pair of distinct points $x, y \in X$ there exist open sets $U, V \subset X$ such that $x \in U$, $y \in V$, $U \cap V = \varnothing$, and $U \cup V = X$.

Remark 1.5. Some books define "totally disconnected" to be what we have called "totally separated," so one must exercise caution when using the literature. These two notions are not, in general, equivalent. Indeed, a totally separated space is obviously both totally disconnected and Hausdorff. By contrast, a totally disconnected space need not be totally separated or even Hausdorff; see the exercises for examples.

However, a compact Hausdorff space is totally disconnected if and only if it is totally separated; see the exercises for a guided proof.

Definition 1.6. A topological group is called *profinite* if it is compact and totally separated.

Remark 1.7. Profinite groups are most commonly defined in the literature to be compact, Hausdorff, and totally disconnected. In view of Remark 1.5, this definition is equivalent to the one above.

The following basic results will be used frequently throughout the course.

Proposition 1.8. *Let G be a compact group. Any subgroup $H \subset G$ is open if and only if it is closed and has finite index in G.*

Proof. Suppose that H is open. The cosets of H are open, disjoint, and cover G. Since G is compact, there are only finitely many of them. Moreover, the complement of H is a union of cosets and hence open; thus H is closed.

Conversely, if H is closed and of finite index, then its complement is a finite union of closed cosets, thus closed. Hence H is open. □

Proposition 1.9. *Let G be a topological group and $H \leq G$ a subgroup. Then H is open if and only if it contains an open neighborhood of the identity.*

Proof. One direction is obvious. If H contains an open set $e \in U$, then $H = \bigcup_{h\in H} hU$ is open. $\square$

Theorem 1.10. *A topological group G is profinite if and only if $G \simeq \varprojlim_I G_i$, where $\{G_i\}_{i\in I}$ is a projective system of (discrete) finite groups.*

Proof. Suppose that $G = \varprojlim_I G_i$, and consider the natural inclusion $f : G \to \prod_{i\in I} G_i$. By Proposition 1.3, the map f is a topological isomorphism onto its image. We claim that $f(G)$ is closed; since each G_i is finite and thus compact, and hence $\prod_{i\in I} G_i$ is compact by Tychonoff's theorem, this will imply the compactness of G. Now, for each pair $i > j$, let

$$X_{ij} = \{(a_k)_k \in \prod_k G_k : \varphi_{ij}(a_i) = a_j\}.$$

Observe that X_{ij} is the preimage of the diagonal under the continuous map (recall that the φ_{ij} are continuous)

$$\begin{aligned} \prod_k G_k &\to G_j \times G_j \\ (a_k)_k &\mapsto (a_j, \varphi_{ij}(a_i)). \end{aligned}$$

Therefore X_{ij} is closed; here we have finally used the discreteness of G_j. Thus $G = \bigcap_{i>j} X_{ij}$ is closed and hence compact.

The G_i are discrete and hence totally separated, and a product of totally separated spaces is clearly totally separated. (Indeed, let $a = (a_i)_i$ and $b = (b_i)_i$ be distinct elements of $\prod_{i\in I} X_i$. Let j be such that $a_j \neq b_j$, and let $U_j, V_j \subset X_j$ be open sets such that $a_j \in U_j$, $b_j \in V_j$, $U_j \cup V_j = X_j$, $U_j \cap V_j = \varnothing$. Set $U = U_j \times \prod_{i\neq j} X_i$ and $V = V_j \times \prod_{i\neq j} X_i$.) But a subspace of a totally separated space is clearly totally separated. Thus G is profinite.

Conversely, suppose that G is a profinite group, and let $\{H_q\}_{q\in Q}$ be the family of its open normal subgroups, ordered by inclusion: $q > q'$ if $H_q \subset H_{q'}$. Note that the quotient subgroups G/H_q are all finite (by Proposition 1.8) and form a projective system of groups. We have a natural homomorphism of groups

$$\begin{aligned} f : G &\to \varprojlim_Q G/H_q \\ g &\mapsto (gH_q)_q \end{aligned}$$

We claim that f is a topological isomorphism. The pre-image under the quotient map $G \to G/H_q$ of any subset of G/H_q is a union of cosets of the open subgroup H_q and hence is open. From this it follows that f is continuous. Any closed subset of the compact group G is compact, hence its image under f is compact and hence closed, since $\varprojlim_Q G/H_q$ is Hausdorff by the first half of this proof. Thus, it suffices to show that f is an isomorphism of abstract groups.

To show that f is surjective, let $(g_qH_q)_q \in \varprojlim_Q G/H_q$. We claim that $\bigcap_{q\in Q} g_qH_q \neq \varnothing$. If this is true, then $(g_qH_q)_q = f(g)$ for any g in the above intersection. Assume, by way of contradiction, that the intersection is empty. This means that $\bigcup_{q\in Q}(G \setminus g_qH_q) = G$. Since G is compact and each g_qH_q is closed, there is a finite subcover $G = \bigcup_{i=1}^r (G \setminus g_{q_i}H_{q_i})$. Hence

$\bigcap_{i=1}^r g_{q_i} H_{q_i} = \varnothing$. The subgroup $H_{q_1} \cap \cdots \cap H_{q_r}$ is open and normal, so it is H_q for some $q \in Q$. Clearly $g_q \in \bigcap_{i=1}^r g_{q_i} H_{q_i}$, contradicting the emptiness of this intersection.

It remains to show that f is injective, and this will follow immediately from the following claim. □

Proposition 1.11. *Let G be a profinite group, and let $\{H_q\}_{q \in Q}$ be the family of all open normal subgroups of G (note that G is in the family, so it is non-empty). Then $\bigcap_{q \in Q} H_q = \{e\}$.*

Proof. Let $x \in G$ be a non-trivial element. We need to construct an open normal subgroup $H \trianglelefteq G$ such that $x \notin H$. By total separatedness of G, there exist open and closed sets U and V that separate e and x. Suppose that $e \in U$. Then $x \notin U$. We will construct an open normal subgroup H that is contained in U and hence does not contain x.

It suffices to show that there exists an open subgroup K of G such that $K \subset U$. Indeed, K has finite index by Proposition 1.8. Then $H = \cap_{g \in G} g K g^{-1} = \cap_{g_i K \in G/K} g_i K g_i^{-1}$ is an open (since the second intersection is finite) normal subgroup of G and $H \subset U$.

For any subset $A \subset G$, we denote by A^n the set of all products of n elements of A and by A^{-1} the set of all inverses of elements of A. Thus A is a subgroup if and only if $A^2 \subset A$ and $A^{-1} \subset A$. Note that U^2 is the image of the compact set $U \times U$ (closed subset of compact $G \times G$) under the continuous multiplication map $\mu : G \times G \to G$. Hence U^2 is compact, and thus closed, since G is Hausdorff. By an easy induction, U^n is compact for all $n \in \mathbb{N}$.

Set $W = (G \setminus U) \cap U^2$; then W consists of all elements of U^2 that do not lie in U. This is a closed subset of G, hence compact. Note that $U \subset G \setminus W$. We claim that for any $u \in U$, there exist open neighborhoods $u \in X_u$ and $e \in Y_u$, contained in U, such that $X_u Y_u \subset (G \setminus W) \cap U^2 \subset U$ (the last inclusion holds since the middle set consists of all elements of U^2 that do lie in U). Indeed, $u \in G \setminus W$. Thus $\mu^{-1}(G \setminus W) \subset G \times G$ is open and contains (u, e). In particular, it contains an open box of the form $X'_u \times Y'_u$, where $u \in X'_u$ and $e \in Y'_u$ are open. Setting $X_u = X'_u \cap U$ and $Y_u = Y'_u \cap U$ (recall that $e \in U$), we get the desired neighborhoods.

Now $\{X_u\}_{u \in U}$ is an open cover of U, and since U is compact there exists a finite subcover $\{X_{u_1}, \dots, X_{u_r}\}$. Set $Y = Y_{u_1} \cap \cdots \cap Y_{u_r}$. This is an open set containing e. Set $Z = Y \cap Y^{-1}$. This is still an open set containing e, since the inversion map is continuous. Moreover, we have $Z \subset Y \subset U$.

Observe that $UZ = \cup_{i=1}^r X_{u_i} Z \subset U$, since $Z \subset Y_{u_i}$ for each i. By induction, $UZ^n \subset U$ for all $n \geq 1$. Since $e \in U$, this implies $Z^n \subset U$ for all $n \geq 1$. Moreover, Z is closed under inversion by construction, and it follows easily that so is each Z^n. Set $K = \langle Y \rangle = \bigcup_{n=1}^\infty Z^n$. Then K is a subgroup of G that is contained in U. Since K contains the open set Z, it is open. We have now completed the proof of our claim and consequently that of Theorem 1.10. □

Corollary 1.12. *Let G be a profinite group. Any open neighborhood of e contains an open normal subgroup.*

Proof. Let U be an open neighborhood of e. Observe that if U is open and closed, then the claim follows from the proof of the previous proposition.

Let $\{H_q\}_{q \in Q}$ be the family of open normal subgroups of G. From the previous proposition it follows that $\bigcup_{q \in Q}(G \setminus H_q) = G \setminus \{e\}$, hence $U \cup \bigcup_{q \in Q}(G \setminus H_q) = G$. By compactness of G, there is a finite subcover $G = U \cup (G \setminus H_{q_1}) \cup \cdots \cup (G \setminus H_{q_r})$. Hence $(G \setminus U) \cap H_{q_1} \cap \cdots \cap H_{q_r} = \varnothing$.

Set $H = H_{q_1} \cap \cdots \cap H_{q_r}$. This is an open normal subgroup of G, and $H \subset U$. □

Corollary 1.13. *Let G be a profinite group and $H \subset G$ a subgroup. Then H is closed (normal) if and only if it is an intersection of open (normal) subgroups.*

Proof. Since open subgroups are closed, one direction is trivial. Conversely, let H be a closed subgroup. Let $\{K_q\}_{q\in Q}$ be the family of open normal subgroups of G. Observe that for every $q \in Q$,

$$HK_q = \{hk : h \in H,\, k \in K_q\},$$

since the K_q are normal, and that HK_q is normal if H is normal as well. We claim that

$$\bigcap_{q\in Q} HK_q = H\left(\bigcap_{q\in Q} K_q\right)$$

for any closed H. Since the right-hand side is equal to H by Proposition 1.11, this implies our claim. It is obvious that the right-hand side is contained in the left-hand side for any subgroup H. To prove the opposite containment, let $g \in \bigcap_{q\in Q} HK_q$. We want to show that g is contained in the right-hand side, or, equivalently, that $Hg \cap \left(\bigcap_{q\in Q} K_q\right) \neq \varnothing$.

Indeed, if this intersection were empty, then, by the argument that is by now standard for us, $(G\setminus Hg)\cup\bigcup_{q\in Q}(G\setminus K_q)$ would be an open cover of G, hence would have a finite subcover, hence $Hg\cap(K_{q_1}\cap\cdots\cap K_{q_r}) = \varnothing$. But $K = K_{q_1}\cap\cdots\cap K_{q_r}$ is itself an open normal subgroup, hence by assumption $g \in HK$. Thus we have arrived at a contradiction. □

Proposition 1.14. *Let G be a profinite group and $H \subset G$ a closed subgroup. Then H is profinite. Moreover, if H is normal, then G/H (with the quotient topology) is profinite.*

Proof. Easy to see that H and G/H are compact and totally separated. Alternatively, if $G \simeq \varprojlim_I G_i$ and $H_i = \pi_i(H) \subset G_i$ for each $i \in I$, then one can show that $H \simeq \varprojlim_I H_i$ and $G/H \simeq \varprojlim_I G_i/H_i$. □

Proposition 1.15. *Suppose that G is profinite and $K \subset H \subset G$ are closed subgroups. There exists a continuous section $s : G/H \to G/K$ of the natural surjection $G/K \to G/H$.*

Proof. First consider the case $[H : K] < \infty$. Then K is open in H, so there exists an open set $V \subset G$ such that $K = V \cap H$. Let $U \trianglelefteq G$ be an open normal subgroup contained in V (this exists by Corollary 1.12), hence that $H \cap U \subset K$.

Decompose G into double cosets $G = \sqcup_{i=1}^{r} Ug_iH$; there are only finitely many double cosets because U has finite index in G. On each piece, define $s : Ug_iH/H \to Ug_iH/K$ by $s(ug_ihH) = ug_iK$. This is well-defined, since if $ug_ih = u'g_ih'$, then $h(h')^{-1} = g_i^{-1}u^{-1}u'g_i \in U \cap H \subset K$, since U is normal in G. Thus $ug_iK = u'g_iK$, so s is indeed well-defined. It is obvious that s is continuous.

We treat the general case by a Zorn's Lemma argument. Consider the set $\mathcal{S}$ of pairs (T, s), where $H \supseteq T \supseteq K$ is a closed subgroup of G and s is a continuous section of the surjection $G/T \to G/H$. We define an order on $\mathcal{S}$ by setting $(T, s) \geq (T', s')$ if $T \subseteq T'$ and the obvious diagram commutes:

$$\begin{tikzcd}
 & G/H \arrow[dl, "s"'] \arrow[dr, "s'"] & \\
G/T \arrow[rr] & & G/T'
\end{tikzcd}$$

If $(T_1, s_1) \leq (T_2, s_2) \leq \cdots$ is a chain, we can set $T_\infty = \bigcap_{i=1}^\infty T_i$ (this is a closed subgroup containing K). It is easy to see that $G/T_\infty = \varprojlim G/T_i$, and by the universal property of the projective limit (which we haven't stated – can also just say that the s_i glue) we get a section $s_\infty : G/H \to G/T_\infty$ that is compatible with all the other data. Thus the chain has an upper bound, and by Zorn's Lemma the set $\mathcal{S}$ contains a maximal element (T, s).

We need to prove that $T = K$. Since K is the intersection of all open subgroups containing K by Corollary 1.13, it suffices to show that $T \subset U$ for any open subgroup $U \supset K$. However, for any such U, we have that $V = T \cap U$ is a closed subgroup such that $T \supseteq V \supseteq K$ and $[T : V] < \infty$.

By the first case we treated in this proof, there is a continuous section $G/T \to G/V$. Composing it with s gives a continuous section $G/H \to G/V$. By maximality of (T, s), this implies that $V = T$ and hence $T \subset U$. □

1.1. **Profinite completion.** Let G be an "abstract group." By this we mean that G is a set with a binary operation satisfying the group axioms, as defined in the first lecture of your group theory course, without any topology or any other further structure. We associate a canonical profinite group to G.

Definition 1.16. Let G be a group. Let $\{N_i\}_{i \in I}$ be the collection of normal subgroups of G of finite index, partially ordered by reverse inclusion. If $j \leq i$, i.e. $N_i \subseteq N_j$, then there is a natural surjection of finite groups $\varphi_{ij} : G/N_i \to G/N_j$. The *profinite completion* of G is the profinite group $\widehat{G} = \varprojlim_I G/N_i$.

There is a natural homomorphism $\eta : G \to \widehat{G}$ given by $\eta(g) = (gN_i)_i$. Observe that η is not necessarily injective: its kernel is $\bigcap_{i \in I} N_i$. A group is called residually finite if the intersection of all its normal subgroups of finite index is trivial. Thus η is an embedding precisely when G is residually finite. We conclude this section with two essential properties of profinite completion.

Proposition 1.17. *Let G be an abstract group, and let $\widehat{G}$ be its profinite completion. The image $\eta(G)$ is dense in $\widehat{G}$.*

Proof. Let $U \subset \widehat{G}$ be an open subset. We need to show that $U \cap \eta(G) \neq \varnothing$. For every normal subgroup $N \trianglelefteq G$ of finite index, let $\widehat{N} \subset \widehat{G}$ be the kernel of the map $\pi_N : \widehat{G} \to G/N$. By the definition of the profinite topology, we have $g\widehat{N} \subset U$ for some $g \in \widehat{G}$ and $N \trianglelefteq G$ of finite index. Let $x \in G$ be such that $xN = \pi_N(g)$. Clearly $\pi_N(\eta(x)) = xN = \pi_N(g)$, so g and $\eta(x)$ lie in the same coset of $\widehat{N} = \ker f_N$. Hence $\eta(x) \in g\widehat{N} \subset U$. □

Proposition 1.18. *Let G be an abstract group and let H be a profinite group. Let $f : G \to H$ be a homomorphism of abstract groups. Then there exists a unique continuous homomorphism $\widehat{f} : \widehat{G} \to H$ such that $\widehat{f} \circ \eta = f$.*

Proof. Let $H = \varprojlim_K H_k$ for a projective system $\{H_k\}_{k \in K}$ of finite groups. For every $k \in K$, let f_k be the composition $G \xrightarrow{f} H \to H_k$. Then $\ker f_k \trianglelefteq G$ is a normal subgroup of finite index, so $\ker f_k = N_{i(k)}$ for some $i(k) \in I$. This induces a homomorphism $G/N_{i(k)} \to H_k$, which we continue to call f_k, and hence a continuous homomorphism $\widehat{f_k} : \widehat{G} \to H_k$, where H_k of course has the discrete topology, given by $\widehat{f_k}((g_i N_i)_i) = f_k(g_{i(k)} N_{i(k)})$. The $\widehat{f_k}$ are compatible with the transition maps of the projective system $\{H_k\}_k$, so we obtain a continuous homomorphism

$\widehat{f} : \widehat{G} \to H$ by the universal property of projective limits. It is clear from the construction that $\widehat{f} \circ \eta = f$. Since $\widehat{f}$ is continuous and is determined by f on the dense (by the previous proposition) subset $\eta(G) \subset \widehat{G}$, it is unique. □

Exercises

(1) Let $X = \mathbb{N} \cup \{x, y\}$, where x and y are distinct elements not contained in $\mathbb{N}$. Declare a subset $U \subset X$ to be open if and only if either $U \subset \mathbb{N}$ or U contains all but finitely many elements of $\mathbb{N}$. Prove that this gives a topology on X, and that X is compact and totally disconnected but not Hausdorff.

(2) Let X be a compact Hausdorff topological space. Let $x \in X$, let C be the connected component of x, and let Q be the intersection of all neighborhoods of x that are both closed and open ("clopen"); this Q is called the quasi-component of x. The aim of this exercise is to prove that $C = Q$.
   (a) Show that $C \subset Q$. Thus it remains to show that Q is connected.
   (b) Suppose that Q is not connected. Show that there exist disjoint open subsets U and V of X such that $Q = (Q \cap U) \cup (Q \cap V)$ and that $Q \cap U$ and $Q \cap V$ are both non-empty.
   (c) Prove that there exist finitely many clopen neighborhoods $U_1, \ldots, U_r$ of x such that $(X \setminus (U \cup V)) \cap (U_1 \cap \cdots \cap U_r) = \varnothing$.
   (d) Let $U' = U_1 \cap \cdots \cap U_r$. Then U' is a clopen neighborhood of x, but also $U' = (U' \cap U) \cup (U' \cap V)$. Derive a contradiction.

(3) Let X be a compact Hausdorff topological space. Prove that X is totally disconnected if and only if it is totally separated.

2. G-modules

Definition 2.1. Let G be a topological group. A *G-module* is a topological abelian group M endowed with a continuous action $G \times M \to M$ such that every $g \in G$ acts by group homomorphisms. In other words, we require:

- $g_1(g_2m) = (g_1g_2)m$ for all $g_1, g_2 \in G$ and $m \in M$.
- $em = m$ for all $m \in M$.
- $g(m_1 + m_2) = gm_1 + gm_2$ for all $g \in G$ and $m_1, m_2 \in M$.

Unless otherwise stated, we will assume from now on that G is profinite and M has the discrete topology. The continuity of the action of G is then equivalent to the following condition: for every $m \in M$ the stabilizer $\mathrm{stab}_G(m) = \{g \in G : gm = m\}$ is an open subgroup of G.

Remark 2.2. If G is profinite and M is discrete (as is now our running assumption), then clearly $M = \bigcup_U M^U$, where U runs over open subgroups of G, and $M^U = \{m \in M : \forall u \in U, um = m\}$ is the space of U-invariants.

Example 2.3. Let K be a number field and $G = \mathrm{Gal}(\overline{K}/K)$. Then G acts on $\overline{K}$ in the obvious way. Moreover, for any $\alpha \in \overline{K}$, the field $K(\alpha)$ has finite degree over K, so that $\mathrm{stab}_G(\alpha) = \mathrm{Gal}(\overline{K}/K(\alpha))$ is open in G. (If $K(\alpha)/K$ is Galois, then $\mathrm{Gal}(\overline{K}/K(\alpha))$ is the kernel of the continuous map $G \to \mathrm{Gal}(K(\alpha)/K)$ and hence open. Otherwise, it contains the absolute Galois group of the Galois closure of $K(\alpha)$, which is open.) Thus $\overline{K}$ is a G-module.

Example 2.4. Let $G = \{e\}$. (Note that finite groups, which we always assume to have the discrete topology, are profinite). A G-module is just an abelian group, with no extra structure.

We write Mod_G for the category of (discrete) G-modules. The morphisms are, of course, group homomorphisms $f : M \to N$ that respect the G-action, i.e. such that $f(gm) = gf(m)$ for all $g \in G$ and $m \in M$. Such maps are often called G-*equivariant*. One checks that Mod_G is an abelian category.

Definition 2.5. Let G be a profinite group, $H \subset G$ a closed subgroup, and M an H-module. The induced module $\mathrm{Ind}_H^G M$ is the space of functions $f : G \to M$ such that

(1) For all $h \in H$ and $g \in G$, we have $f(hg) = hf(g)$.
(2) f is continuous, i.e. locally constant (recall that M has the discrete topology).

The group G acts on this space by right translation: $(gf)(x) = f(xg)$ for any $x, g \in G$ and any $f \in \mathrm{Ind}_H^G M$.

Proposition 2.6. *The G-action defined above gives $\mathrm{Ind}_H^G M$ the structure of a G-module.*

Proof. Let $f \in \mathrm{Ind}_H^G M$. Then for every $g \in G$ there exists an open set $g \in V_g$ such that f is constant on V_g. Since a base of the topology on G is given by the cosets of open subgroups, there exists an open subgroup U_g such that $g \in gU_g \subset V_g$. Then $G = \bigcup_{g \in G} gU_g$, and by compactness of G there is a finite subcover $G = g_1 U_{g_1} \cup \cdots \cup g_r U_{g_r}$ such that f is constant on each $g_i U_{g_i}$. Let $U = U_{g_1} \cap \cdots \cap U_{g_r}$. This is an open subgroup of G that is contained in $\mathrm{stab}_G(f)$. Indeed, for any $u \in U$ and $x \in G$, if $x \in g_i U_{g_i}$, then $xu \in g_i U_{g_i}$ as well, and thus $(uf)(x) = f(xu) = f(x)$. □

Proposition 2.7. *Let G be a profinite group and $H \subset G$ a closed subgroup.*

(1) $\mathrm{Ind}_H^G : \mathrm{Mod}_H \to \mathrm{Mod}_G$ *is a functor.*
(2) *If $K \subset H \subset G$ are closed subgroups, then $\mathrm{Ind}_K^G \simeq \mathrm{Ind}_H^G \mathrm{Ind}_K^H$.*
(3) *The functor Ind_H^G is exact. In other words, if $0 \to M \to N \to P \to 0$ is an exact sequence of H-modules, then $0 \to \mathrm{Ind}_H^G M \to \mathrm{Ind}_H^G N \to \mathrm{Ind}_H^G P \to 0$ is also exact.*

Proposition 2.8 (Frobenius Reciprocity)**.** *Let G be a profinite group and $H \subset G$ a closed subgroup. Let $M \in \mathrm{Mod}_G$ and $N \in \mathrm{Mod}_H$. Then*

$$\mathrm{Hom}_G(M, \mathrm{Ind}_H^G N) \simeq \mathrm{Hom}_H(M, N).$$

Proof. We can write down an explicit bijection. Indeed, if $\varphi \in \mathrm{Hom}_G(M, \mathrm{Ind}_H^G N)$, define $A(\varphi) \in \mathrm{Hom}_H(M, N)$ by $A(\varphi)(m) = (\varphi(m))(e) \in N$ for any $m \in M$.

In the other direction, if $\psi \in \mathrm{Hom}_H(M, N)$, we define $B(\psi) \in \mathrm{Hom}_G(M, \mathrm{Ind}_H^G N)$ by $(B(\psi)(m))(g) = \psi(gm)$. To check that $B(\psi)$ is G-equivariant, observe for any $g, x \in G$ and $m \in M$ that

$$(B(\psi)(xm))(g) = \psi(gxm) = (B(\psi)(m))(gx) = (x(B(\psi)(m)))(g).$$

and hence $(B(\psi))(xm) = x(B(\psi)(m))$. It is simple to check that $B(A(\varphi)) = \varphi$ and $A(B(\psi)) = \psi$. For instance, for all $m \in M$ and $g \in G$ and $\varphi \in \mathrm{Hom}_G(M, \mathrm{Ind}_H^G N)$, we have

$$B(A(\varphi))(m)(g) = A(\varphi)(gm) = \varphi(gm)(e) = g(\varphi(m))(e) = \varphi(m)(g).$$

Similarly, $A(B(\psi))(m) = (B(\psi)(m))(e) = \psi(em) = \psi(m)$. □

Definition 2.9. A G-module I is called *injective* if for any G-module morphism $f : M \to I$ and any injective $g : M \hookrightarrow N$, there exists $\tilde{f} : N \to I$ that completes the diagram:

$$\begin{tikzcd} M \arrow[r, "f"] \arrow[d, "g"'] & I \\ N \arrow[ur, "\tilde{f}"] & \end{tikzcd}$$

Proposition 2.10. *An abelian group is an injective $\{e\}$-module if and only if it is divisible. In particular, $\mathbb{Q}$ and $\mathbb{Q}/\mathbb{Z}$ are injective $\{e\}$-modules.*

Proof. Recall that a group G, written additively, is called divisible if for any $x \in G$ and any $n \in N$ there exists $y \in G$ such that $ny = x$.

Let G be a divisible abelian group, and f and g as in Definition 2.9 above. View M as a submodule of N via g, and consider the set of all pairs (M', h), where $M \subset M' \subset N$ and h extends f. Define a partial order on this set by extension: $(M', h') \leq (M'', h'')$ if $M' \subset M''$ and $h''_{|M'} = h'$. Any ascending chain glues, so has an upper bound, so by Zorn's Lemma there exists a maximal pair (M', h). We need to show that $M' = N$. Suppose not. Then there exists $n \in N \setminus M'$. Let s be the order of n in N/M'. Then the map $\tilde{h} : M' + \langle n \rangle \to G$ given by

$$\tilde{h}(m' + an) = \begin{cases} h(m') & : s = \infty \\ h(m') + at & : s < \infty \end{cases}$$

extends h, contradicting the maximality of (M', h). Here $a \in \mathbb{Z}$ is arbitrary and $t \in G$ is an element satisfying $st = h(sn)$.

Now we prove the converse direction, which will not be used later. Suppose that G is an injective $\{e\}$-module and let $x \in G$. Consider the map $f : \mathbb{Z} \to G$ given by $f(m) = mx$ for all $m \in \mathbb{Z}$. Given $n \in \mathbb{N}$, consider the injection $g : \mathbb{Z} \to \mathbb{Z}$ given by $g(m) = mn$. By injectivity of G there exists $\tilde{f} : \mathbb{Z} \to G$ such that $f = \tilde{f} \circ g$. Let $y = \tilde{f}(1)$. Then $x = f(1) = \tilde{f}(n) = ny$, so G is divisible. $\square$

Corollary 2.11. *Let G be a profinite group and $H \subset G$ a closed subgroup. If I is an injective H-module, then $\mathrm{Ind}_H^G I$ is an injective G-module.*

Proof. Suppose we have $f : M \to \mathrm{Ind}_H^G I$ and $\varepsilon : M \hookrightarrow N$ is an embedding of G-modules. By Frobenius reciprocity, $f \in \mathrm{Hom}_G(M, \mathrm{Ind}_H^G I)$ corresponds to an H-module morphism $A(f) \in \mathrm{Hom}_H(M, I)$. By injectivity of I, we get an H-module map $\tilde{f} : N \to I$ satisfying $\tilde{f} \circ \varepsilon = A(f)$, hence, by a second application of Frobenius reciprocity, a map $B(\tilde{f}) : N \to \mathrm{Ind}_H^G I$ of G-modules.

We claim that $B(\tilde{f})$ is the map we are looking for, namely that it satisfies $B(\tilde{f}) \circ \varepsilon = f$. Indeed, for any $m \in M$ and any $g \in G$, since ε is a map of G-modules we have

$$\begin{aligned} (B(\tilde{f}) \circ \varepsilon)(m)(g) &= \tilde{f}(g\varepsilon(m)) = \tilde{f}(\varepsilon(gm)) = A(f)(gm) = f(gm)(e) = \\ & \quad g \cdot (f(m))(e) = f(m)(eg) = f(m)(g). \end{aligned}$$

Thus $(B(\tilde{f}) \circ \varepsilon)(m) = f(m) \in \mathrm{Ind}_H^G I$ for each $m \in M$, which proves our claim. $\square$

Corollary 2.12. *If G is a profinite group and M is any G-module, there exists a G-module I into which M embeds (i.e. the category* Mod_G has enough injectives*).*

Proof. The claim is true for $\{e\}$-modules, i.e. abelian groups. Indeed, let M be any abelian group and $m \in M$ a non-trivial element. If $\langle m \rangle$ is a torsion group, then it embeds in $I_m = \mathbb{Q}/\mathbb{Z}$, and otherwise it embeds in $I_m = \mathbb{Q}$. Since $\langle m \rangle \subset M$, we can extend this embedding to a map $f_m : M \to I_m$ satisfying $f_m(m) \neq 0$. We can collect all of these into a map $M \to \prod_{m \in M} I_m$ which clearly has trivial kernel. (Observe that a direct product of injective modules is injective.)

Now let M be any G-module, and let I be an injective group such that there exists an embedding $\psi : M \to I$ of abelian groups. By Frobenius reciprocity, it corresponds to a G-module map $B(\psi) : M \to \mathrm{Ind}_{\{e\}}^G I$. Note that $\mathrm{Ind}_{\{e\}}^G I$ is an injective G-module by Corollary 2.11, so it remains only to show that $B(\psi)$ is injective. But for any $m \in M$, we have that $B(\psi)(m)$ is the map $g \mapsto \psi(gm)$. Since ψ is an embedding, this can be the zero map only when $m = 0$. □

Exercises

(1) Let G be a torsion abelian group with the discrete topology and let $H \subset G$ be a subgroup. Let $g \in G$ be such that $g \notin H$. Prove that there exists a homomorphism $f : G \to \mathbb{Q}/\mathbb{Z}$ such that $H \subset \ker f$ but $g \notin \ker f$.

(2) This exercise is a variant of the previous one. Suppose now that G is a profinite abelian group and $H \subset G$ is a closed subgroup. Let $g \in G \setminus H$ as above, and show that there is a continuous homomorphism $f : G \to \mathbb{Q}/\mathbb{Z}$ such that $H \subset \ker f$ and $g \notin \ker f$.

(3) Let I be a directed partially ordered set. For every $i \in I$, let A_i be a topological (unital) ring. One can define a projective system analogously to Definition 1.1, where now $\varphi_{ij} : A_i \to A_j$ is a continuous ring homomorphism whenever $i \geq j$. One obtains a topological ring $\varprojlim_I A_i$ analogously to Definition 1.2.

   (a) We define a topological ring A to be profinite if it is compact and totally separated. Prove that A is profinite if and only if $A \simeq \varprojlim_I A_i$ for some projective system $\{A_i\}_{i \in I}$ of finite rings with the discrete topology.

   (b) Let A be a topological ring. Consider the group $A^\vee = \mathrm{Hom}(A, \mathbb{Q}/\mathbb{Z})$ of continuous group homomorphisms, where we consider A as an abelian group under addition. Endow $A^\vee$ with the compact-open topology; this is the coarsest topology such that the sets $B(K, U) = \{f \in A^\vee : f(K) \subset U\}$ are open for all compact $K \subset A$ and all open $U \subset \mathbb{Q}/\mathbb{Z}$. The topology on $\mathbb{Q}/\mathbb{Z}$ is taken to be the subspace topology arising from the embedding $\iota : \mathbb{Q}/\mathbb{Z} \hookrightarrow \mathbb{C}$ given by $\iota(x) = e^{2\pi i x}$. Prove that $A^\vee$ is discrete if A is compact and that $A^\vee$ is a profinite group if A is a discrete torsion group.

   (c) A topological abelian group M is said to be an A-module if it is endowed with an A-module structure such that the scalar multiplication map $A \times M \to M$ is continuous. Show that $A^\vee$ is an A-module if we set $(af)(b) = f(ba)$ for all $f \in A^\vee$ and $a, b \in A$.

   (d) Let A be a topological ring. Suppose that A is compact and Hausdorff, and let $C \subset A$ be the connected component of 0. Prove that $af = 0$ for any $a \in C$ and $f \in A^\vee$. Deduce that $C = \{0\}$.

   (e) Show that a topological ring A is profinite if and only if it is compact and Hausdorff.

(4) Give an example showing that a compact Hausdorff group is not necessarily profinite. Why does your solution to the previous exercise fail for groups?

(5) Prove Proposition 2.7.

(6) Let G be a profinite group, and let $H \subset G$ be an open subgroup. Let $M \in \mathrm{Mod}_G$ and $N \in \mathrm{Mod}_H$. Prove the following version of Frobenius reciprocity:
$$\mathrm{Hom}_G(\mathrm{Ind}_H^G N, M) \simeq \mathrm{Hom}_H(N, M).$$
The maps, in notation corresponding to that of Proposition 2.8, are $A(\varphi)(n) = \varphi(f_n)$ and $B(\psi)(f) = \sum_{Hg \in H\backslash G} g^{-1}\psi(f(g))$, where $f_n \in \mathrm{Ind}_H^G N$ is the function $f_n : G \to N$ satisfying $f_n(h) = hn$ for $h \in H$ and $f_n(g)$ for $g \notin H$.

(7) Let G be a profinite group and let M be a G-module.
   (a) Let $m \in M$, and let $U = \mathrm{stab}_G(m)$. Let $\mathbf{1}_U$ be the trivial U-module; more precisely, $\mathbf{1}_U$ is an infinite cyclic group with the action $ua = a$ for all $u \in U$ and $a \in \mathbf{1}_U$. Show that there is a G-module homomorphism $f : \mathrm{Ind}_U^G \mathbf{1}_U \to M$ such that the image of f is the submodule of M generated by m.
   (b) Let $N \subset M$ be a *monogenic* submodule, i.e. N is generated by a single element. Prove that N is isomorphic to a quotient of $\mathrm{Ind}_U^G \mathbf{1}_U$ for some open subgroup $U \subset G$.

(8) Show that $I \in \mathrm{Mod}_G$ is injective if and only if $\mathrm{Hom}_G(-, I)$ is an exact contravariant functor.

(9) Let I be an injective G module, and let $H \subset G$ be an open subgroup. Prove that I, with the natural H-action obtained by restriction, is an injective H-module.

(10) Let $H \trianglelefteq G$ be a closed normal subgroup, and let I be an injective G-module. Prove that $I^H = \{i \in I : hi = i \,\text{for all}\, h \in H\}$ is an injective G/H-module.

3. Cohomology of G-modules

3.1. **Definition of cohomology.** For any profinite group G, there is a left exact functor $\mathrm{Mod}_G \to \mathrm{Ab}$ sending a G-module M to the space of invariants $M^G = \{m \in M : \forall g \in G, gm = m\}$. The cohomology of M will be defined as the right derived functors of this functor. To make the exposition more elegant, we will define some notions (first introduced by Grothendieck in the famous Tohoku paper) that axiomatize the properties of derived functors.

Definition 3.1. A *(cohomological) δ-functor* $\mathrm{Mod}_G \to \mathrm{Ab}$ (the same definition can be made for any two abelian categories) consists of the following data:

- A series of covariant functors $H^i : \mathrm{Mod}_G \to \mathrm{Ab}$, for $i \in \mathbb{N}_0$.
- For each exact sequence of G-modules $0 \to M \to N \to P \to 0$ and each $i \geq 0$, a G-module homomorphism $\delta^i : H^i(P) \to H^{i+1}(M)$ with the following properties:
   (1) For $0 \to M \to N \to P \to 0$ exact as above, there is a long exact sequence of cohomology groups
$$0 \to H^0(M) \to H^0(N) \to H^0(P) \xrightarrow{\delta^0} H^1(M) \to H^1(N) \to H^1(P) \xrightarrow{\delta^1} H^2(M) \to \cdots$$
   (2) Given a morphism of short exact sequences
$$\begin{tikzcd} 0 \arrow[r] & M \arrow[r] \arrow[d] & N \arrow[r] \arrow[d] & P \arrow[r] \arrow[d] & 0 \\ 0 \arrow[r] & M' \arrow[r] & N' \arrow[r] & P' \arrow[r] & 0 \end{tikzcd}$$

(where the diagram is commutative with exact rows), for each $i \geq 0$ the corresponding diagram commutes:

$$\begin{tikzcd} H^i(P) \arrow[r, "\delta^i"] \arrow[d] & H^{i+1}(M) \arrow[d] \\ H^i(P') \arrow[r, "\delta^i"] & H^{i+1}(M') \end{tikzcd}$$

Proposition 3.2. *Let G be a profinite group. There exists a unique δ-functor $(H^i(G,-),\delta^i)$ from Mod_G to Ab satisfying the following two properties:*

(1) *For every $M \in \mathrm{Mod}_G$, we have $H^0(G,M) = M^G$.*
(2) *If $I \in \mathrm{Mod}_G$ is injective, then $H^i(G,I) = 0$ for every $i > 0$.*

Proof. Suppose that such a δ-functor exists. What can we say about it? Any $M \in \mathrm{Mod}_G$ embeds into an injective G-module I^0; let N be the cokernel, so that we have a short exact sequence $0 \to M \xrightarrow{\iota} I^0 \xrightarrow{\pi} N \to 0$. By (2) and the long exact sequence, we would have an exact sequence

$$0 \to M^G \xrightarrow{\iota^G} (I^0)^G \xrightarrow{\pi^G} N^G \to H^1(G,M) \to H^1(G,I^0) \to \cdots,$$

whence $H^1(G,M) = \operatorname{coker} \pi^G$ and we must have $H^{i+1}(M) \simeq H^i(N)$ for all $i \geq 1$. Since we know that $H^0(G,M) = M^G$ for any $M \in \mathrm{Mod}_G$, the observations above provide a recursive algorithm for computing $H^i(G,M)$.

We can streamline this process by fixing an injective resolution of M, namely an exact sequence

$$0 \to M \to I^0 \xrightarrow{\alpha_0} I^1 \xrightarrow{\alpha_1} I^2 \xrightarrow{\alpha_2} \cdots.$$

This is possible because the category Mod_G has enough injectives by Corollary 2.12. Indeed, M embeds in an injective module I^0 as above. The quotient $I^0/M = N$ embeds into an injective module I^1, and this lifts to a map $\alpha_0 : I^0 \to I^1$. Then $I^1/\alpha_0(I^0)$ embeds into I^2, and this lifts to a map $\alpha_1 : I^1 \to I^2$. Continue forever.

Now apply the G-invariants functor to the injective resolution above, to obtain a complex

$$0 \to M^G \to (I^0)^G \xrightarrow{d_0} (I^1)^G \xrightarrow{d_1} (I^2)^G \xrightarrow{d_2} \cdots \tag{1}$$

This complex need not be exact, since the G-invariants functor is only left exact. However, it obviously inherits the property that $d_{i+1} \circ d_i = 0$ for all $i \geq 0$, and hence $\operatorname{im} d_{i-1} \subseteq \ker d_i$ for all $i \geq 1$.

We claim that $H^1(G,M) \simeq \operatorname{coker} \pi^G \simeq (\ker d_1)/(\operatorname{im} d_0)$. Indeed, since $\alpha_0 : I^0 \to I^1$ is the composition of π with an embedding $\varepsilon : N \hookrightarrow I^1$, we see that $\operatorname{im} d_0 = \varepsilon(\operatorname{im} \pi^G)$. On the other hand, $\alpha_1 : I^1 \to I^2$ is the composition $I^1 \twoheadrightarrow I^1/\alpha_0(I^0) = I^1/\varepsilon(N) \hookrightarrow I^2$, hence we see that $\ker \alpha_1 = \varepsilon(N)$ and $\ker d_0 = \varepsilon(N^G)$. Thus, the embedding ε induces an isomorphism $(\ker d_1)/(\operatorname{im} d_0) \simeq N^G/(\operatorname{im} \pi^G) = \operatorname{coker} \pi^G$.

Observe that $0 \to N \to I^1 \to I^2 \to \cdots$ is an injective resolution of N. Thus, by the argument of the previous paragraph, $H^2(G,M) \simeq H^1(G,N) \simeq (\ker d_2)/(\operatorname{im} d_1)$. By induction, for all $i \geq 1$, we get

$$H^i(G,M) = (\ker d_i)/(\operatorname{im} d_{i-1}).$$

So far we have shown that if a δ-functor with the desired properties exists, then it is unique. In fact, this construction works, so the δ-functor does exist. However, it is not at all obvious that the $H^i(G,-)$ are well-defined functors and that there exist maps δ^i that turn all this data into a δ-functor. We will soon sketch a proof of these facts, but for the time being we will assume them. $\square$

The construction of the cohomology groups $H^i(G,M)$ that we just gave does not, in practice, provide a useful way of computing them. Indeed, the injective modules constructed in Corollaries 2.11 and 2.12 are huge and unwieldy. Moreover, the definition of the maps α_i involves extending maps into injective modules, and this ultimately relies on the argument with Zorn's Lemma in the proof of Proposition 2.10 and thus is completely non-constructive. Hence an injective resolution as in the proof of Proposition 3.2 cannot usually be constructed explicitly. Thus, our first priority is to find a more explicit way to compute cohomology.

Definition 3.3. A G-module J is called *acyclic* if $H^i(G,J)=0$ for all $i>0$.

Clearly, all injective G-modules are acyclic. Note that in the construction of the previous proof, all we used about injective modules was their acyclicity and the existence of injective resolutions. The injectivity will be used more intensively in the proof that the above construction actually gives a δ-functor. In the meantime, we can conclude the following, which is useful for computing cohomology in practice.

Proposition 3.4. *Let M be a G-module, and consider an acyclic resolution*

$$0 \to M \to J^0 \to J^1 \to J^2 \to \cdots,$$

namely an exact sequence where the J^i are acyclic for all $i \geq 0$. Taking G-invariants gives rise to a complex

$$0 \to M^G \to (J^0)^G \xrightarrow{d_0} (J^1)^G \xrightarrow{d_1} (J^2)^G \cdots$$

Then for all $i>0$ we have $H^i(G,M) \simeq (\ker d_i)/(\operatorname{im} d_{i-1})$.

Definition 3.5. Let G be a profinite group and M a G-module. For each $i \geq 0$, define $\mathcal{C}^i(G,M)$ to be the space of continuous (i.e. locally constant) functions $\varphi : G^{i+1} \to M$. Let G act on this space by

$$(g\varphi)(g_0,\dots,g_i) = g \cdot (\varphi(g_0 g,\dots,g_i g))$$

for $\varphi \in \mathcal{C}^i(G,M)$ and $g, g_0, \dots, g_i \in G$.

We have an obvious injective map $f : M \to \mathcal{C}^0(G,M)$ that sends each $m \in M$ to the constant function $\varphi_m : G \to M$ satisfying $\varphi_m(g) = m$ for all $g \in G$. Observe that f is G-equivariant, since $(x\varphi_m)(g) = x\varphi_m(gx) = xm = \varphi_{xm}(g)$ for all $x, g \in G$.

Moreover, for all $i \geq 0$ we have a map $f_i : \mathcal{C}^i(G,M) \to \mathcal{C}^{i+1}(G,M)$ given by

$$(f_i\varphi)(g_0,\dots,g_{i+1}) = \sum_{j=0}^{i+1} (-1)^j \varphi(g_0,\dots,\hat{g}_j,\dots,g_{i+1}), \tag{2}$$

where $\hat{g}_j$ denotes omission of the variable g_j, so that each summand on the right-hand side provides $i+1$ arguments for the function φ. It is clear that f_i is G-equivariant.

Lemma 3.6. *The complex*

$$0 \to M \xrightarrow{f} \mathcal{C}^0(G,M) \xrightarrow{f_0} \mathcal{C}^1(G,M) \xrightarrow{f_1} \mathcal{C}^2(G,M) \xrightarrow{f_2} \cdots \tag{3}$$

is exact.

Proof. It is easy to check that this is indeed a complex, namely that the composition of any two consecutive arrows is the zero map. Observe that f is obviously injective. For any $\varphi \in \mathcal{C}^0(G, M)$ we have $(f_0\varphi)(g_0, g_1) = \varphi(g_1) - \varphi(g_0)$. Hence, if $\varphi \in \ker f_0$ then φ is a constant function and thus contained in the image of f.

Now let $i > 0$. Suppose that $\varphi \in \mathcal{C}^i(G, M)$ lies in the kernel of f_i, and define a locally constant map $\psi : G^i \to M$ by $\psi(g_1, \ldots, g_i) = -\varphi(e, g_1, \ldots, g_i)$. Then one checks that

$$\begin{aligned}(f_{i-1}(\psi))(g_1, \ldots, g_{i+1}) = \quad & \sum_{j=1}^{i+1} (-1)^{j-1} \psi(g_1, \ldots, \hat{g}_j, \ldots, g_{i+1}) = \\ & \sum_{j=1}^{i+1} (-1)^{j} \varphi(e, g_1, \ldots, \hat{g}_j, \ldots, g_{i+1}) = \\ (f_i(\varphi))(e, g_1, \ldots, g_{i+1}) + \varphi(g_1, \ldots, g_{i+1}) = \quad & \varphi(g_1, \ldots, g_{i+1}),\end{aligned}$$

where the last equality follows from $\varphi \in \ker f_i$. Thus $\varphi = f_{i-1}(\psi)$. This gives exactness at $\mathcal{C}^i(G, M)$. □

3.2. **Universal δ-functors and Shapiro's Lemma.** If we knew that the complex (3) were an acyclic resolution of M, we would be able to apply Proposition 3.4 to compute cohomology. To prove the acyclicity of $\mathcal{C}^i(G, M)$ we will need a few more tools at our disposal.

Definition 3.7. A functor $\mathcal{F} : \mathrm{Mod}_G \to \mathrm{Ab}$ is called *effaceable*[1] if for any $M \in \mathrm{Mod}_G$ there exists an injection $\iota : M \hookrightarrow N$ such that $\mathcal{F}(\iota) : \mathcal{F}(M) \to \mathcal{F}(N)$ is the zero map.

Definition 3.8. A δ-functor (H^i, δ^i) is called *universal* if, for any δ-functor (G^i, ε_i) and any natural transformation $f_0 : H^0 \to G^0$, there exists a unique sequence of natural transformations $f_i : H^i \to G^i$ that commutes with the δ^i maps.

(Recall that a natural transformation of functors $f : \mathcal{F} \to \mathcal{G}$ provides, for each $M \in \mathrm{Ob}(\mathcal{F})$, a map $f_M : \mathcal{F}(M) \to \mathcal{G}(M)$, such that for any arrow $M \to N$ the obvious square below commutes.)

$$\begin{tikzcd} \mathcal{F}(M) \arrow[r, "f_M"] \arrow[d] & \mathcal{G}(M) \arrow[d] \\ \mathcal{F}(N) \arrow[r, "f_N"] & \mathcal{G}(N). \end{tikzcd}$$

Lemma 3.9 (Grothendieck)**.** *If (H^i, δ^i) is a δ-functor such that the functors H^i are effaceable for all $i > 0$, then (H^i, δ^i) is a universal δ-functor.*

Proof. Let (G^i, ε^i) be another δ-functor, and let $f_0 : H^0 \to G^0$ be a natural transformation. The first step is to construct the unique natural transformations $f_i : H^i \to G^i$ for all $i > 0$. Assume, by induction, that f_j has been constructed and shown to be unique for all $j < i$. Let $M \in \mathrm{Mod}_G$ and let $\iota : M \to J$ be an embedding, where J is a G-module such that

[1] "Effaceable" is a French word taken directly from Grothendieck's Tohoku paper. S. Lang argues forcefully that authors writing in English should employ its English translation "erasable." Nevertheless, here we follow the standard usage.

$H^i(\iota) : H^i(M) \to H^i(J)$ is the zero map; this exists because (H^i, δ^i) is effaceable. Let C be the cokernel. Then the short exact sequence $0 \to M \xrightarrow{\iota} J \to C \to 0$ gives rise to the diagram

$$\begin{tikzcd}
H^{i-1}(M) \arrow[r] \arrow[d] & H^{i-1}(J) \arrow[r] \arrow[d] & H^{i-1}(C) \arrow[r, "\delta^{i-1}"] \arrow[d, "f_{i-1,C}"] & H^i(M) \arrow[r, "H^i(\iota)"] \arrow[d, dotted] & H^i(J) \\
G^{i-1}(M) \arrow[r] & G^{i-1}(J) \arrow[r] & G^{i-1}(C) \arrow[r, "\varepsilon^{i-1}"] & G^i(M) \arrow[r] & G^i(J).
\end{tikzcd}$$

The dotted arrow $f_{i,M}$ is constructed by a simple diagram chase. Indeed, let $c \in H^i(M)$. Since δ^{i-1} is surjective, there exists $b \in H^{i-1}(C)$ such that $\delta^{i-1}(c)$. Then we define $f_{i,M}(c) = \varepsilon^{i-1} f_{i-1,C}(b)$. It is easy to check that this is independent of the choice of b. It remains to show that $f_{i,M} : H^i(M) \to G^i(M)$ is independent of the choice of ι and that f_i is a natural transformation. We will omit these details here. □

Corollary 3.10. *For any profinite group G, the δ-functor $(H^i(G, -), \delta^i)$ constructed in Proposition 3.2 is universal.*

Proof. Since every G-module M injects into an injective module, the functors $H^i(G, -)$ are effaceable for all $i > 0$. □

Proposition 3.11 (Shapiro's Lemma)**.** *Let G be a profinite group, $H \subset G$ a closed subgroup, and M an H-module. Then there is a natural isomorphism $H^i(H, M) \simeq H^i(G, \mathrm{Ind}_H^G M)$ for all $i \geq 0$.*

By "natural" we mean that these isomorphisms arise from a natural transformation of functors $H^i(H, -) \to H^i(G, -) \circ \mathrm{Ind}_H^G$.

Proof. The strategy of the proof is to show that both functors are universal δ-functors extending the same H^0. Indeed, observe that $H^0(G, \mathrm{Ind}_H^G M) = (\mathrm{Ind}_H^G M)^G$ consists of (continuous) functions $f : G \to M$ such that $f(hg) = h \cdot f(g)$ for all $h \in H, g \in G$ and $f(gx) = f(g)$ for all $g, x \in G$. The second condition forces f to be a constant function $f_m : g \mapsto m$ for some $m \in M$, while the first condition imposes $hm = m$ for all $h \in H$, in other words that $m \in M^H$. Thus we have $M^H \xrightarrow{\sim} (\mathrm{Ind}_H^G M)^G$ for all $M \in \mathrm{Mod}_H$, and this is clearly a natural isomorphism of functors.

We already know by Corollary 3.10 that $(H^i(H, -), \delta^i)$ is a universal δ-functor. Since $(H^i(G, -), \delta^i)$ is a δ-functor and Ind_H^G is exact, it is not hard to see that $(H^i(G, \mathrm{Ind}_H^G -))$ is a δ-functor. For any H-module M, let $\iota : M \to I$ be an embedding of M into an injective H-module. The induced map $H^i(G, \mathrm{Ind}_H^G M) \to H^i(G, \mathrm{Ind}_H^G I) = 0$ is the zero map when $i > 0$, since $\mathrm{Ind}_H^G I$ is an injective G-module by Corollary 2.11. Thus the functors $H^i(G, \mathrm{Ind}_H^G -)$ are effaceable for $i > 0$, and universality follows from Lemma 3.9. □

Lemma 3.12. *All $\{e\}$-modules are acyclic. If G is any profinite group and M any $\{e\}$-module, then $\mathrm{Ind}_{\{e\}}^G M$ is an acyclic G-module.*

Proof. Since the functor taking $\{e\}$-invariants is just the identity functor, the complex (1) is exact for any $\{e\}$-module M. Thus M is acyclic. The second part of the claim now follows from Shapiro's Lemma. □

Before returning to our explicit computation of the cohomology $H^i(G, M)$, we point out a simple consequence of Lemma 3.12.

Corollary 3.13. *Let L/K be a finite Galois extension of fields. Then $H^i(\mathrm{Gal}(L/K), L) = 0$ for all $i > 0$.*

Proof. Here $G = \mathrm{Gal}(L/K)$ acts on L in the obvious way. Since any finite separable extension is simple, there exists an element $\alpha \in L$ such that $L = K(\alpha)$. By the Normal Basis Theorem, α may be chosen so that $\{\sigma(\alpha) : \sigma \in G\}$ is a K-basis of L. Thus any $\beta \in L$ can be written in the form $\beta = \sum_{\sigma\in G} \beta_\sigma \sigma(\alpha)$, where $\beta_\sigma \in K$.

We can obtain a G-module isomorphism $L \to \mathrm{Ind}_{\{e\}}^G K$ by sending $\beta \in L$ to the function $f_\beta : G \to K$ given by $f_\beta(\sigma) = \beta_{\sigma^{-1}}$. Indeed, for any $\tau \in G$ we have $\tau(\beta) = \sum_{\sigma\in G} \beta_\sigma \tau\sigma(\alpha) = \sum_{\sigma\in G} \beta_{\tau^{-1}\sigma}\sigma(\alpha)$. On the other hand, we have

$$(\tau f_\beta)(\sigma) = f_\beta(\sigma\tau) = \beta_{\tau^{-1}\sigma^{-1}} = f_{\tau(\beta)}(\sigma)$$

for all $\sigma, \tau \in G$. The claim now follows from Lemma 3.12. □

3.3. **Explicit computation of cohomology.** We are now ready to return to the G-modules $\mathcal{C}^i(G, M)$ that were defined earlier.

Lemma 3.14. *For every $i \geq 0$, let $\mathcal{LC}^i(G, M)$ be the abelian group of locally constant functions $f : G^i \to M$, viewed as a $\{e\}$-module. There is a G-module isomorphism $\theta : \mathcal{C}^i(G, M) \xrightarrow{\sim} \mathrm{Ind}_{\{e\}}^G \mathcal{LC}^i(G, M)$.*

Proof. Given $\varphi \in \mathcal{C}^i(G, M)$, define $\theta_\varphi \in \mathrm{Ind}_{\{e\}}^G \mathcal{LC}^i(G, M)$ by

$$\theta_\varphi(g)(g_1, \ldots, g_i) = g\varphi(g, g_1 g, \ldots, g_i g).$$

Clearly θ is a group homomorphism. We note that θ is G-equivariant. Indeed, if $x \in G$, then

$$\theta_{x\varphi}(g)(g_1, \ldots, g_i) = g(x\varphi)(g, g_1 g, \ldots, g_i g) = gx\varphi(gx, g_1 gx, \ldots, g_i gx) = \theta_\varphi(gx)(g_1, \ldots, g_i).$$

Thus $\theta_{x\varphi}(g) = \theta_\varphi(gx) = (x\theta_\varphi)(g)$ for all $g \in G$.

Note that if $\theta_\varphi(g)$ is the zero element of $\mathcal{LC}^i(G, M)$ for all $g \in G$, then $\varphi(g, g_1 g, \ldots, g_i g) = 0$ for all $g, g_1, \ldots, g_i \in G$, whence $\varphi = 0$. Thus θ is injective. Moreover, given $\psi \in \mathcal{LC}^i(G, M)$ we can define $\eta \in \mathcal{C}^i(G, M)$ by $\eta(g_0, \ldots, g_i) = g_0^{-1}\psi(g_0)(g_1 g_0^{-1}, \ldots, g_i g_0^{-1})$. It is easy to see that $\psi = \theta_\eta$, so that θ is also surjective. □

Corollary 3.15. *Let G be a profinite group and M a G-module. The G-modules $\mathcal{C}^i(G, M)$ are acyclic for all $i \geq 0$.*

Proof. This is immediate from Lemma 3.14 and Lemma 3.12. □

Therefore, the complex (3) is an acyclic resolution of M, and by Proposition 3.4 it may be used to compute the cohomology of M. We will now do this and observe some consequences.

First of all, it will be useful to find a convenient parametrization of the spaces $\mathcal{C}^i(G, M)^G$ of G-invariants. Recall that these spaces consist of locally constant functions $\varphi : G^{i+1} \to M$ such that $g \cdot \varphi(g_0 g, \ldots, g_i g) = \varphi(g_0, \ldots, g_i)$ for all $g, g_0, \ldots, g_i \in G$.

If $i > 0$, let $C^i(G, M)$ be the space of locally constant functions $\psi : G^i \to M$, with no G-module structure. There is a group homomorphism $C^i(G, M) \to \mathcal{C}^i(G, M)^G$ sending $\psi \in C^i(G, M)$ to the map

$$\varphi_\psi(g_0, \ldots, g_i) = g_0^{-1}\psi(g_0 g_1^{-1}, g_1 g_2^{-1}, \ldots, g_{i-1} g_i^{-1}).$$

We wish to show that this is an isomorphism of groups by constructing an inverse map. Given $\varphi \in \mathcal{C}^i(G,M)^G$, we want to find $\psi_\varphi \in C^i(G,M)$ such that $g_0^{-1}\psi_\varphi(g_0g_1^{-1}, g_1g_2^{-1}, \ldots, g_{i-1}g_i^{-1}) = \varphi(g_0,\ldots,g_i) = g_0^{-1}\varphi(e, g_1g_0^{-1},\ldots,g_ig_0^{-1})$. The map

$$\psi_\varphi(x_1,\ldots,x_i) = \varphi(e, x_1^{-1}, (x_1x_2)^{-1}, \ldots, (x_1x_2\cdots x_i)^{-1}) \tag{4}$$

works. Note that all of this holds for $i=0$ as well, where $C^0(G,M)$ is the space of constant functions. The next step is to determine the map $d_i : C^i(G,M) \to C^{i+1}(G,M)$ induced by f_i. Given $\psi \in C^i(G,M)$, a straightforward computation shows that

$$\begin{aligned}(f_i\varphi_\psi)(g_0,\ldots,g_{i+1}) \;=\;& g_1^{-1}\psi(g_1g_2^{-1},\ldots,g_ig_{i+1}^{-1}) + \\ & \sum_{j=1}^{i}(-1)^j g_0^{-1}\psi(g_0g_1^{-1},\ldots,g_{j-1}g_{j+1}^{-1},\ldots,g_ig_{i+1}^{-1}) + \\ & (-1)^{i+1}g_0^{-1}\psi(g_0g_1^{-1},\ldots,g_{i-1}g_i^{-1}).\end{aligned}$$

It follows that $d_i\psi = \psi_{f_i(\varphi_\psi)}$ is given by the following formula:

$$\begin{aligned}d_i(\psi)(x_1,\ldots,x_{i+1}) \;=\;& x_1\psi(x_2,\ldots,x_{i+1}) - \psi(x_1x_2,x_3,\ldots,x_{i+1}) + \\ & \psi(x_1,x_2x_3,x_4,\ldots,x_{i+1}) + \cdots + (-1)^i\psi(x_1,x_2,\ldots,x_ix_{i+1}) + \\ & (-1)^{i+1}\psi(x_1,\ldots,x_i).\end{aligned} \tag{5}$$

Finally, note that if $\varphi \in \mathcal{C}^0(G,M)^G$, then $\varphi(g_1) = g\varphi(g_1g)$ for all $g, g_1 \in G$. Taking $g_1 = e$, we see that $m = g\varphi(g)$ is independent of g. Moreover, $\psi_{f_0(\varphi)}(g) = f_0\varphi(e,g^{-1}) = \varphi(g^{-1}) - \varphi(e)$. Thus $\operatorname{im} d_0$ consists of functions $\psi : G \to M$ of the form $\psi(g) = gm - m$ for a fixed $m \in M$.

For each $i > 0$, let $Z^i(G,M) = \ker d_i$ and $B^i(G,M) = \operatorname{im} d_{i-1}$ denote the cocycles and coboundaries, respectively. The Proposition 3.4 tells us that:

Proposition 3.16. *Let G be a profinite group and M a G-module. Then for all $i > 0$, we have $H^i(G,M) \simeq Z^i(G,M)/B^i(G,M)$.*

Remark 3.17. In particular, if $\psi \in C^1(G,M)$, then $d_1(\psi)(g_1,g_2) = g_1\psi(g_2) - \psi(g_1g_2) + \psi(g_1)$. Thus the 1-cocycles are maps $\psi : G \to M$ satisfying $\psi(g_1g_2) = \psi(g_1) + g_1\psi(g_2)$. These are called *crossed homomorphisms*. Similar, $B^1(G,M)$ consists of maps $\psi : G \to M$ of the form $\psi(g) = gm - m$ for a fixed $m \in M$.

Example 3.18. If G is a finite cyclic group and M is a G-module, then $Z^1(G,M)$ is isomorphic to the abelian group $N(M) = \{m \in M : \sum_{g\in G} gm = 0\}$.

Proof. Let σ be a generator of G and let $\varphi \in Z^1(G,M)$. Observe that $\varphi(e) = \varphi(e) + e\varphi(e)$, whence $\varphi(e) = 0$. Similarly, for all $j \geq 1$ we have $\varphi(\sigma^j) = \varphi(\sigma) + \sigma(\varphi(\sigma^{j-1}))$. It follows by induction that $\varphi(\sigma^j) = \sum_{k=0}^{j-1}\sigma^k(\varphi(\sigma))$ for all $j \geq 1$. Taking $j = |G|$, we get

$$0 = \varphi(e) = \varphi(\sigma^j) = \sum_{k=0}^{j-1}\sigma^k(\varphi(\sigma)) = \sum_{\tau\in G}\tau(\varphi(\sigma)).$$

Thus $\varphi(\sigma) \in N(M)$, and $\varphi \mapsto \varphi(\sigma)$ is our candidate for an isomorphism between $Z^1(G,M)$ and $N(M)$. Clearly it is injective, since φ is determined by $\varphi(\sigma)$ by the recursive formula above. To show surjectivity, let $m \in N(M)$ and define $\varphi_m : G \to M$ by $\varphi_m(\sigma^j) = \sum_{k=0}^{j-1}\sigma^k(m)$. It is easy to check that this is indeed a 1-cocycle, and clearly $\varphi_m(\sigma) = m$. □

The following important result is a generalization, by Emmy Noether, of Theorem 90 in Hilbert's *Zahlbericht*, which itself originated with Kummer.

Theorem 3.19 (Hilbert 90)**.** *If L/K is a finite Galois extension, then $H^1(\mathrm{Gal}(L/K), L^\times) = 0$.*

Proof. We write the abelian group $L^\times$ multiplicatively and denote $G = \mathrm{Gal}(L/K)$. Let $\varphi \in Z^1(G, L^\times)$. For every $\alpha \in L^\times$ consider the "Poincaré series"

$$b(\alpha) = \sum_{\sigma \in G} \varphi(\sigma)\sigma(\alpha) \in L.$$

By Dedekind's Lemma the automorphisms σ are linearly independent over L, so there must exist some $\alpha \in L^\times$ such that $b = b(\alpha) \neq 0$. This means that for all $\sigma, \tau \in G$ we have $\varphi(\sigma\tau) = \varphi(\sigma)\sigma(\varphi(\tau))$. Now for every $\sigma \in G$ we have

$$\sigma(b) = \sum_{\tau \in G} \sigma(\varphi(\tau))\sigma\tau(\alpha) = \sum_{\tau \in G} \varphi(\sigma\tau)\varphi(\sigma)^{-1}\sigma\tau(\alpha) = (\varphi(\sigma))^{-1} \sum_{\tau \in G} \varphi(\sigma\tau)\sigma\tau(\alpha) = (\varphi(\sigma))^{-1}b.$$

This means that $\varphi(\sigma) = \sigma(b^{-1}) \cdot b = \sigma(m)m^{-1}$, where $m = b^{-1}$. Thus $\varphi \in B^1(G, L^\times)$. □

We now derive the original Hilbert 90 as a corollary.

Corollary 3.20. *Let L/K be a cyclic Galois extension and let σ be a generator of $G = \mathrm{Gal}(L/K)$. Let $\alpha \in L$ be an element with $\mathrm{N}_{L/K}(\alpha) = 1$. Then there exists $\beta \in L^\times$ such that $\alpha = \sigma(\beta)\beta^{-1}$.*

Proof. In the notation of Example 3.18, we have that $N(L^\times) = \{\gamma \in L^\times : \mathrm{N}_{L/K}(\gamma) = 1\}$. By that example, there exists $\varphi \in Z^1(G, L^\times)$ such that $\varphi(\sigma) = \alpha$. By Hilbert 90, we see that $Z^1(G, L^\times) = B^1(G, L^\times)$, hence there exists $\beta \in L^\times$ such that $\varphi(\tau) = \tau(\beta)\beta^{-1}$ for all $\tau \in G$. □

3.4. **Some homological algebra.** At this point we finally go back to the construction in Proposition 3.2 and show why it indeed gives a well-defined δ-functor. Recall that a complex $\mathbf{C}$ consists of the data

$$C^0 \stackrel{\partial^0}{\to} C^1 \stackrel{\partial^1}{\to} C^2 \stackrel{\partial^2}{\to} C^3 \cdots$$

where $\partial^{i+1} \circ \partial^i = 0$ for all $i \geq 0$. The cohomology of the complex is given by $H^i(\mathbf{C}) = (\ker \partial^{i+1})/(\mathrm{im}\, \partial^i)$ for all $i \geq 0$.

Definition 3.21. Let $\mathbf{C}$ and $\mathbf{D}$ be two complexes. A *cochain map* $\varphi : \mathbf{C} \to \mathbf{D}$ is a collection of morphisms $\varphi^i : C^i \to D^i$ that commute with the coboundary maps of the complexes. In other words, the diagram

$$\begin{tikzcd}
\cdots \arrow[r, "\partial^{i-1}"] & C^i \arrow[r, "\partial^i"] \arrow[d, "\varphi^i"'] & C^{i+1} \arrow[r, "\partial^{i+1}"] \arrow[d, "\varphi^{i+1}"'] & \cdots \\
\cdots \arrow[r, "\partial^{i-1}"] & D^i \arrow[r, "\partial^i"] & D^{i+1} \arrow[r, "\partial^{i+1}"] & \cdots
\end{tikzcd}$$

commutes for each i.

Observe that a cochain map $\varphi : \mathbf{C} \to \mathbf{D}$ induces maps $H^i(\varphi) : H^i(\mathbf{C}) \to H^i(\mathbf{D})$ in a natural way. We will sometimes write ∂_C^i and ∂_D^i to distinguish the coboundary maps coming from our two complexes. The main tool for proving that two complexes have the same cohomology is the notion of homotopy.

Definition 3.22. Let $\mathbf{C}$ and $\mathbf{D}$ be two complexes, and let $\varphi, \psi : \mathbf{C} \to \mathbf{D}$ be two cochain maps. A *homotopy* $\Sigma : \varphi \to \psi$ is a family of morphisms $\Sigma^i : C^i \to D^{i-1}$ such that $\psi - \varphi = \partial\Sigma + \Sigma\partial$, i.e. $\psi^i - \varphi^i = \partial^{i-1} \circ \Sigma^i + \Sigma^{i+1} \circ \partial^i$ for all i.

Lemma 3.23. *Let $\mathbf{C}$ and $\mathbf{D}$ be two complexes, and let $\varphi, \psi : \mathbf{C} \to \mathbf{D}$ be two cochain maps. Suppose that there exists a homotopy $\Sigma : \varphi \to \psi$. Then $H^i(\varphi) = H^i(\psi)$ for all i.*

Proof. Suppose that $x \in C^i$ lies in the kernel of ∂^i. We need to show that the two elements $\varphi^i(x)$ and $\psi^i(x)$ of D^i, which clearly lie in the kernel of ∂^i, differ by an element of the image of ∂^{i-1}. However,

$$\psi^i(x) - \varphi^i(x) = \partial^{i-1}(\Sigma^i(x)) + \Sigma^{i+1}(\partial^i(x)) = \partial^{i-1}(\Sigma^i(x))$$

by the definition of homotopy, and this is exactly what we need. □

We say that φ and ψ are *homotopic* if there exists a homotopy $\Sigma : \varphi \to \psi$ and leave it to the reader to verify that this is an equivalence relation on the set of cochain maps.

Proposition 3.24. *Let $\mathbf{C}$ and $\mathbf{D}$ be two complexes. Suppose that $\mathbf{C}$ is acyclic, i.e. that $H^i(\mathbf{C}) = 0$ for all $i \geq 1$, whereas $\mathbf{D}$ is injective, namely that all the D^i are injective objects. Let $\eta : H^0(\mathbf{C}) \to H^0(\mathbf{D})$ be a homomorphism. Then there exists a cochain map $\varphi : \mathbf{C} \to \mathbf{D}$ inducing η. Moreover, any two cochain maps inducing η are homotopic.*

Proof. Observe that $H^0(\mathbf{C}) = \ker \partial_C^0 \subset C^0$, and similarly $H^0(\mathbf{D}) \subset D^0$. Thus η gives a homomorphism $H^0(\mathbf{C}) \to D^0$, which extends to a homomorphism $\varphi^0 : C^0 \to D^0$ by injectivity of D^0. Now assume by induction that φ^{i-1} has been constructed. Observe that φ^{i-1} maps $\ker \partial_C^{i-1}$ to $\ker \partial_D^{i-1}$. Indeed, if $i = 1$ and $z \in \ker \partial_C^0 = H^0(\mathbf{C})$, then $\varphi^0(z) = \eta(z) \in H^0(\mathbf{D}) = \ker \partial_D^0$. If $i > 1$ and $z \in \ker \partial_C^{i-1}$, then $z \in \operatorname{im} \partial_C^{i-2}$ by the acyclicity of $\mathbf{C}$. Since $\varphi^{i-1} \circ \partial_C^{i-2} = \partial_D^{i-2} \circ \varphi^{i-2}$, this implies that $\varphi^{i-1}(z) \in \operatorname{im} \partial_D^{i-2} \subset \ker \partial_D^{i-1}$ as claimed.

If $y \in \operatorname{im} \partial_C^{i-1}$, then choose $x \in C^{i-1}$ such that $\partial_C^{i-1}(x) = y$ and define $\varphi^i(y) = \partial_D^{i-1}(\varphi^{i-1}(x))$. This is well-defined; indeed, if x and x' are two preimages of x, then $z = x' - x \in \ker \partial_C^{i-1}$ and thus $\varphi^{i-1}(z) \in \ker \partial_D^{i-1}$. Hence we get a map $\varphi^i : \operatorname{im} \partial_C^{i-1} \to D^i$. It extends to a homomorphism $\varphi^i : C^i \to D^i$ by the injectivity of D^i, and our definition ensures that the following square commutes:

$$\begin{tikzcd} C^{i-1} \arrow[r, "\partial_C^{i-1}"] \arrow[d, "\varphi^{i-1}"] & C^i \arrow[d, "\varphi^i"'] \\ D^{i-1} \arrow[r, "\partial_D^{i-1}"] & D^i \end{tikzcd}$$

Now let φ and ψ be two cochain maps inducing $H^0(\varphi)$; we will construct a homotopy Σ between them. If $y \in \operatorname{im} \partial_C^0$, then define $\Sigma^1(y) = \psi^0(x) - \varphi^0(x)$, where $\partial^0(x) = y$; observe that this is well-defined. We can extend this to a homomorphism $\Sigma^1 : C^1 \to D^0$ by injectivity of D^0.

Suppose now that Σ^i has been constructed. If $y \in \operatorname{im} \partial_C^i$, then define $\Sigma^{i+1}(y) = \psi^i(x) - \varphi^i(x) - \partial_D^{i-1}(\Sigma^i(x))$, where $x \in C^i$ satisfies $\partial_C^i(x) = y$. If this is well-defined, it will extend to a homomorphism $\Sigma^{i+1} : C^{i+1} \to D^i$ by the injectivity of D^i. It is indeed well-defined, because if x' is another pre-image of y, then $z = x' - x \in \ker \partial_C^i = \operatorname{im} \partial_C^{i-1}$ by acyclicity of $\mathbf{C}$. Let $w \in C^{i-1}$ satisfy $\partial_C^{i-1}(w) = z$. Then

$$\begin{aligned}(\psi^i - \varphi^i)(z) = \ & (\psi^i - \varphi^i)(\partial_C^{i-1}(w)) = \\ & \partial_D^{i-1}(\psi^{i-1} - \varphi^{i-1})(w) = \partial_D^{i-1}(\partial_D^{i-2}\Sigma^{i-1}w + \Sigma^i \partial_C^{i-1} w) = \partial_D^{i-1}\Sigma^i(z). \quad \square\end{aligned}$$

Corollary 3.25. *Let G be a profinite group and M a G-module. The cohomology groups $H^i(G, M)$, as constructed in the proof of Proposition 3.2, do not depend on the choice of injective resolution. Moreover, $H^i(G, -) : \mathrm{Mod}_G \to \mathrm{Ab}$ is a functor for all $i \geq 0$.*

Proof. Let $0 \to M \to I^0 \xrightarrow{\partial_I^0} I^1 \xrightarrow{\partial_I^1} \cdots$ and $0 \to M \to J^0 \xrightarrow{\partial_J^0} J^1 \xrightarrow{\partial_J^1} \cdots$ be two injective resolutions of M. Note that $H^0(\mathbf{I}) = H^0(\mathbf{J}) = M$. By the previous proposition, the identity maps in either direction extend to cochain maps $\varphi : \mathbf{I} \to \mathbf{J}$ and $\psi : \mathbf{J} \to \mathbf{I}$. Moreover, the cochain map $\psi \circ \varphi : \mathbf{I} \to \mathbf{I}$ induces the identity on $H^0(\mathbf{I})$ and is thus homologous to the identity cochain map $\mathbf{1_I} : \mathbf{I} \to \mathbf{I}$. Let $\Sigma : \mathbf{1_I} \to \psi \circ \varphi$ be a homotopy.

Now apply the G-invariants functor to everything in sight. The resulting complexes $\mathbf{I}^G$ and $\mathbf{J}^G$ are no longer necessarily acyclic or injective, so we can't apply the previous proposition directly. However, the relation $\psi^i \circ \varphi^i - 1 = \partial_I^{i-1} \circ \Sigma^i + \Sigma^{i+1} \circ \partial_I^i$ clearly survives the application of the G-invariants functor. Since $\mathbf{1}_{\mathbf{I}^G}$ obviously induces the identity on cohomology, we find that the composition $H^i(\mathbf{I}^G) \xrightarrow{\mathrm{H}^i(\varphi)} H^i(\mathbf{J}^G) \xrightarrow{H^i(\psi)} H^i(\mathbf{I}^G)$ is the identity by Lemma 3.23. Hence we get isomorphic results when we compute $H^i(G, M)$ via the injective resolutions $\mathbf{I}$ and $\mathbf{J}$.

Finally if $f : M \to N$ is any map of G-modules, then we can choose injective resolutions $0 \to M \to \mathbf{I}$ and $0 \to N \to \mathbf{J}$. Then $H^0(\mathbf{I}) \simeq M$ and $H^0(\mathbf{J}) \simeq N$, so the homomorphism $f : H^0(\mathbf{I}) \to H^0(\mathbf{J})$ is induced by a cochain map φ. By the previous proposition, all possible choices of φ are homologous and thus induce the same maps on cohomology. Applying the G-invariants functor as above, we find that $H^i(f) : H^i(G, M) \to H^i(G, N)$ is well-defined and functorial. $\square$

It remains to show that the functors $H^i(G, -)$ may be supplemented with boundary maps δ^n to obtain a δ-functor. This is a consequence of the Snake Lemma.

Lemma 3.26 (Snake Lemma)**.** *Suppose we have a commutative diagram of G-modules*

$$\begin{tikzcd} & A \arrow[r, "f"] \arrow[d, "\alpha"'] & B \arrow[r, "g"] \arrow[d, "\beta"'] & C \arrow[r] \arrow[d, "\gamma"'] & 0 \\ 0 \arrow[r] & A' \arrow[r, "f'"] & B' \arrow[r, "g'"] & C' & \end{tikzcd}$$

in which the horizontal rows are exact. There exists a canonical exact sequence

$$\ker \alpha \xrightarrow{\tilde{f}} \ker \beta \xrightarrow{\tilde{g}} \ker \gamma \xrightarrow{\delta} \operatorname{coker} \alpha \xrightarrow{\tilde{f}'} \operatorname{coker} \beta \xrightarrow{\tilde{g}'} \operatorname{coker} \gamma,$$

in which $\tilde{f}$ and $\tilde{g}$ are the restrictions of f and g, respectively, whereas $\tilde{f}'$ and $\tilde{g}'$ are induced from f' and g'. Moreover, $\tilde{f}$ is injective if f is, and $\tilde{g}'$ is surjective if g' is.

Proof. There is only one natural way to obtain a map $\delta : \ker\gamma \to \operatorname{coker}\alpha$. Let $x \in \ker\gamma$. Since g is surjective, there exists $y \in B$ such that $g(y) = x$. Moreover, $\beta(y) \in \ker g' = \operatorname{im} f'$ by commutativity of the right square and exactness at B'. Since f' is injective, there is a unique $z \in A'$ satisfying $f'(z) = \beta(y)$, and we set $\delta(x) = z + \operatorname{im}\alpha$.

If $y' \in B$ satisfies $g(y') = x$ and $z' \in A'$ satisfies $f'(z') = \beta(y')$, then $y' - y \in \ker g = \operatorname{im} f$, so $y' - y = f(w)$ for some $w \in A$. But then $\alpha(w) = z' - z$, whence δ is well-defined. This argument is presented, among other sources, in the opening scene of the 1980 film *It's My Turn*. It remains to prove the exactness of the claimed sequence. This diagram chase is left as an exercise for the reader. ☐

To apply the Snake Lemma, we will need to find injective resolutions $\mathbf{I}'$, $\mathbf{I}$, and $\mathbf{I}''$ of M, N, and P, respectively, and chain maps φ and ψ extending ι and π, respectively, such that $0 \to I'^{,i} \xrightarrow{\varphi^i} I^i \xrightarrow{\psi^i} I''^{,i} \to 0$ is a short exact sequence for every $i \geq 0$. This is a stronger claim than we can get from Proposition 3.24 directly.

Lemma 3.27. *Suppose that $0 \to M \to N \to P \to 0$ is a short exact sequence of G-modules. Suppose that we are given embeddings $\varepsilon' : M \hookrightarrow I'$ and $\varepsilon'' : P \hookrightarrow I''$, where I' and I'' are injective objects. Then we can complete this picture to a commutative diagram*

$$\begin{tikzcd}
0 \arrow[r] & M \arrow[r, "\iota"] \arrow[d, "\varepsilon'"] & N \arrow[r, "\pi"] \arrow[d, dotted, "\varepsilon"] & P \arrow[r] \arrow[d, "\varepsilon''"] & 0 \\
0 \arrow[r, dotted] & I' \arrow[r, dotted] & I \arrow[r, dotted] & I'' \arrow[r, dotted] & 0,
\end{tikzcd}$$

where I is injective, $\varepsilon : N \to I$ is an embedding, and the bottom row is a split short exact sequence.

Proof. Let $I = I' \oplus I''$; clearly this is an injective object. The maps on the bottom row will be the standard embedding and projection. It remains to define ε. By injectivity of I' there exists a map $\varphi : N \to I'$ such that $\varphi \circ \iota = \varepsilon'$. The diagram above already includes a map from N to I'', and we take these to be the components of $\varepsilon = (\varphi, \varepsilon'' \circ \pi)$. This works. ☐

Corollary 3.28. *Suppose that $0 \to M \to N \to P \to 0$ is a short exact sequence and $0 \to M \to \mathbf{I}'$ and $0 \to P \to \mathbf{I}''$ are injective resolutions. These data can be extended to a* *commutative diagram*

$$\begin{tikzcd}
0 \arrow[r] & M \arrow[r, "\iota"] \arrow[d, dotted, "\varepsilon'"'] & N \arrow[r, "\pi"] \arrow[d, dotted, "\varepsilon"'] & P \arrow[r] \arrow[d, dotted, "\varepsilon''"'] & 0 \\
0 \arrow[r, dotted] & (I^0)' \arrow[r, dotted] \arrow[d, dotted] & I^0 \arrow[r, dotted] \arrow[d, dotted] & (I^0)'' \arrow[r, dotted] \arrow[d, dotted] & 0 \\
0 \arrow[r, dotted] & (I^1)' \arrow[r, dotted] \arrow[d, dotted] & I^1 \arrow[r, dotted] \arrow[d, dotted] & (I^1)'' \arrow[r, dotted] \arrow[d, dotted] & 0 \\
 & \vdots & \vdots & \vdots &
\end{tikzcd}$$

where the columns are injective resolutions and the rows are short exact sequences. Moreover, all the rows, except for possibly the top one, are split.

Proof. We are given embeddings $\varepsilon' : M \hookrightarrow (I^0)'$ and $\varepsilon'' : P \hookrightarrow (I^0)''$ of M and P into injective objects. By the previous lemma we can fill in the second row of the diagram. Recalling that the map ε factors through the inclusion $(I^0)' \to I^0$, we check that the cokernels naturally lie in a short exact sequence

$$0 \to \operatorname{coker} \varepsilon' \to \operatorname{coker} \varepsilon \to \operatorname{coker} \varepsilon'' \to 0.$$

The injective resolutions we have been given provide for embeddings of $\operatorname{coker} \varepsilon'$ and $\operatorname{coker} \varepsilon''$ into injective objects $(I^1)'$ and $(I^1)''$, respectively. We apply the previous subclaim again and continue forever. $\square$

Remark 3.29. The proofs of Lemma 3.27 and Corollary 3.28 work in any abelian category with enough injectives.

Proposition 3.30. *Let* $0 \to M \xrightarrow{\iota} N \xrightarrow{\pi} P \to 0$ *be a short exact sequence of* G*-modules. For every* $i \geq 0$ *there exists a map* $\delta^i : H^i(G, P) \to H^{i+1}(G, M)$ *satisfying the properties of a cohomological* δ*-functor.*

Proof. Fix injective resolutions $0 \to M \to \mathbf{I}$ and $0 \to P \to \mathbf{K}$. Note that $M = H^0(\mathbf{I})$ and $P = H^0(\mathbf{K})$. By Corollary 3.28 there is an injective resolution $0 \to N \to \mathbf{J}$ and cochain maps $\varphi : \mathbf{I} \to \mathbf{J}$ and $\psi : \mathbf{J} \to \mathbf{K}$ extending ι and π, respectively, such that $0 \to I^i \xrightarrow{\varphi^i} J^i \xrightarrow{\psi^i} K^i \to 0$ is a short exact sequence for every $i \geq 0$. Moreover, by the construction in the proof of Lemma 3.27 we can take this sequence to be split. Since left exact functors preserve split short exact sequences (prove this!), the sequence $0 \to (I^i)^G \xrightarrow{\varphi^i} (J^i)^G \xrightarrow{\psi^i} (K^i)^G \to 0$ is still split for all $i \geq 0$.

For each $X \in \{I, J, K\}$, let $d^i_X : (X^i)^G \to (X^{i+1})^G$ be as in (1). Now we apply the Snake Lemma to the diagram

$$\begin{tikzcd} & \operatorname{coker} d_I^{i-1} \arrow[r, "\varphi^i"] \arrow[d, "\partial_I^i"'] & \operatorname{coker} d_J^{i-1} \arrow[r, "\psi^i"] \arrow[d, "\partial_J^i"'] & \operatorname{coker} d_K^{i-1} \arrow[r] \arrow[d, "\partial_K^i"'] & 0 \\ 0 \arrow[r] & \ker d_I^{i+1} \arrow[r, "\varphi^{i+1}"] & \ker d_J^{i+1} \arrow[r, "\psi^{i+1}"] & \ker d_K^{i+1} & \end{tikzcd} \tag{6}$$

where ∂^i_X is induced from d^i_X. The exactness of the rows can be verified as an easy exercise, or, alternatively, by applying the Snake Lemma to the diagram

$$\begin{tikzcd} 0 \arrow[r] & (I^j)^G \arrow[r, "\varphi^j"] \arrow[d, "d_I^j"'] & (J^j)^G \arrow[r, "\psi^j"] \arrow[d, "d_J^j"'] & (K^j)^G \arrow[r] \arrow[d, "d_K^j"'] & 0 \\ 0 \arrow[r] & (I^{j+1})^G \arrow[r, "\varphi^{j+1}"] & (J^{j+1})^G \arrow[r, "\psi^{j+1}"] & (K^{j+1})^G \arrow[r] & 0 \end{tikzcd}$$

for $j \in \{i-1, i+1\}$. Note that $\ker \partial^i_I = H^i(G, M)$ and $\operatorname{coker} \partial^i_I = H^{i+1}(G, M)$, and similarly for the other two columns of (6). Applying the Snake Lemma to (6) thus gives an exact sequence

$$H^i(G, M) \to H^i(G, N) \to H^i(G, P) \xrightarrow{\delta^i} H^{i+1}(G, M) \to H^{i+1}(G, N) \to H^{i+1}(G, P),$$

where the unlabeled arrows arise from ι and π by the functoriality of $H^i(G, -)$ and $H^{i+1}(G, -)$. It remains to show that δ^i is itself functorial in the sense of the second part of Definition 3.1. Given a morphism of short exact sequences

$$\begin{tikzcd} 0 \arrow[r] & M \arrow[r] \arrow[d] & N \arrow[r] \arrow[d] & P \arrow[r] \arrow[d] & 0 \\ 0 \arrow[r] & M' \arrow[r] & N' \arrow[r] & P' \arrow[r] & 0 \end{tikzcd}$$

we obtain the diagram

$$\begin{tikzcd} \cdots \arrow[r] & H^i(G, N) \arrow[r] \arrow[d] & H^i(G, P) \arrow[r, "\delta^i"] \arrow[d] & H^{i+1}(G, M) \arrow[r] \arrow[d] & H^{i+1}(G, N) \arrow[r] \arrow[d] & \cdots \\ \cdots \arrow[r] & H^i(G, N') \arrow[r] & H^i(G, P') \arrow[r, "(\delta')^i"] & H^{i+1}(G, M') \arrow[r] & H^{i+1}(G, N') \arrow[r] & \cdots \end{tikzcd}$$

in which all the unlabeled maps arise from the functors $H^i(G, -)$ and $H^{i+1}(G, -)$. In particular,, the squares on the left and right commute. We need to show that the central square commutes. This is proved by following the definition of the connecting map δ in the proof of the Snake Lemma and is left as an exercise. $\square$

Remark 3.31. A careful inspection of our arguments shows that the only property of the G-invariants functor that figures in the proof that $H^i(G,-)$ indeed gives a δ-functor is that $M \mapsto M^G$ is left exact. Thus we can make the same construction for any left exact functor $\mathcal{F} : \mathcal{A} \to \mathcal{B}$, where $\mathcal{A}$ and $\mathcal{B}$ are abelian categories and $\mathcal{A}$ has enough injectives. Indeed, if M is an object of $\mathcal{A}$, we can find a resolution by injective objects:

$$0 \to M \to I^0 \xrightarrow{\partial^0} I^1 \xrightarrow{\partial^1} I^2 \cdots$$

Applying $\mathcal{F}$, we get a complex in $\mathcal{B}$ that need no longer be exact:

$$\mathcal{F}(I^0) \xrightarrow{\mathcal{F}(\partial^0)} \mathcal{F}(I^1) \xrightarrow{\mathcal{F}(\partial^1)} \mathcal{F}(I^2) \cdots$$

Note that $\ker \mathcal{F}(\partial^0) = \mathcal{F}(M)$. By the arguments above, the cohomology of this complex is independent of the choice of injective resolution. We obtain the series of *right derived functors* of $\mathcal{F}$, given by

$$R^i\mathcal{F}(M) = (\ker \mathcal{F}(\partial^i))/(\operatorname{im} \mathcal{F}(\partial^{i-1})).$$

Observe that $R^0\mathcal{F} = \mathcal{F}$. The connecting maps to make this a δ-functor are obtained by a Snake Lemma argument as above.

Exercises

(1) Prove that the G-invariants functor is left exact. In other words, given a short exact sequence $0 \to M \to N \to P \to 0$ of G-modules, show that $0 \to M^G \to N^G \to P^G$ is an exact sequence.

Prove that the functor is not exact, i.e. that $0 \to M^G \to N^G \to P^G \to 0$ need not be an exact sequence.

(2) Let G be a profinite group, and let k be a field. Let $\mathrm{Rep}_k(G)$ be the category of all k-vector spaces V such that V, viewed as an abelian group, is endowed with a G-module structure that respects the k-linear structure: in other words, for every $g \in G$ the map $v \mapsto gv$ is a k-linear automorphism of V. The morphisms in this category are, of course, k-linear G-equivariant maps. Let Vec_k be the category of k-vector spaces.

(a) Verify that the left exact functor $(-)^G : \mathrm{Rep}_k(G) \to \mathrm{Vec}_k$ gives rise to right derived functors $H^i_k(G,-)$.

(b) Let $V \in \mathrm{Rep}_k(G)$ be a two-dimensional representation of G. For any $g \in G$, let $\rho(g) \in \mathrm{Aut}_k(V)$ denote the map $v \mapsto gv$. Suppose that V is reducible, i.e. it has a G-invariant non-trivial proper subspace. Show that there exists a k-basis of V with respect to which

$$\rho(g) = \begin{pmatrix} \chi_1(g) & c(g) \\ 0 & \chi_2(g) \end{pmatrix}$$

for all $g \in G$, where $\chi_1, \chi_2 : G \to k^\times$ are group homomorphisms and $c : G \to k$ is a function. Show that $(g \mapsto c(g)\chi_2^{-1}(g)) \in Z^1_k(G, \chi_1\chi_2^{-1})$. Here $\chi_1\chi_2^{-1}$ denotes a one-dimensional k-vector space with the G-action $gv = \chi_1(g)\chi_2^{-1}(g)v$. Let $c_V \in H^1_k(G, \chi_1\chi_2^{-1})$ be the corresponding cohomology class.

(c) Show that V determines c_V up to multiplication by a non-zero scalar in k. Conversely, show that every one-dimensional subspace of $H^1_k(G, \chi_1\chi_2^{-1})$ gives rise to a two-dimensional representation $V \in \mathrm{Rep}_k(G)$ admitting a short exact sequence $0 \to \chi_1 \to V \to \chi_2 \to 0$.

(3) Prove that $Z^2(G,M)$ consists of locally constant functions $\psi : G^2 \to M$ satisfying

$$\psi(g_1,g_2) + \psi(g_1g_2,g_3) = \psi(g_1,g_2g_3) + g_1\psi(g_2,g_3)$$

for all $g_1,g_2,g_3 \in G$. Show that $\psi \in B^2(G,M)$ if and only if there exists a function $\varphi : G \to M$ such that $\psi(g_1,g_2) = \varphi(g_1) - \varphi(g_1g_2) + g_1\varphi(g_2)$ for all $g_1,g_2 \in G$.

(4) Let G be a profinite group and let A be a finite abelian group. An extension of G by A is a short exact sequence $0 \to A \xrightarrow{\iota} E \xrightarrow{\pi} G \to 0$ of profinite groups, where the maps are continuous group homomorphisms. (Beware that some authors call this an extension of A by G). By abuse of notation, we will call the extension E. Let $u : G \to E$ be a continuous section of π, namely a continuous map (not necessarily a homomorphism) such that $\pi \circ u$ is the identity on G; this exists by Proposition 1.15.

(a) For $g \in G$ and $a \in A$, define $ga = u(g)au(g)^{-1}$. Show that this action is independent of the choice of u and endows A with a G-module structure.

(b) Let E be an extension as above. For all $g_1,g_2 \in G$, show that there exists $\psi(g_1,g_2) \in A$ such that $\psi_u(g_1,g_2) = u(g_1)u(g_2)u(g_1g_2)^{-1}$. Show that the map $\psi_u : G \times G \to A$ is a 2-cocycle, namely that it lies in $Z^2(G,A)$.

(c) Show that if we choose a different section $u' : G \to E$, then the 2-cocycles ψ_u and $\psi_{u'}$ give rise to the same class in $H^2(G,A)$.

(d) Two extensions E_1 and E_2 are called congruent if there exists a continuous homomorphism $f : E_1 \to E_2$ such that the diagram

$$\begin{tikzcd}
0 \arrow[r] & A \arrow[r] \arrow[d, "="] & E_1 \arrow[r] \arrow[d, "f"] & G \arrow[r] \arrow[d, "="] & 0 \\
0 \arrow[r] & A \arrow[r] & E_2 \arrow[r] & G \arrow[r] & 0
\end{tikzcd}$$

commutes. Prove that f is necessarily an isomorphism and that congruence of extensions is an equivalence relation.

(e) Prove that congruent extensions E_1 and E_2 give rise to the same class in $H^2(G,A)$.

(f) Now, let A be a finite abelian group with a G-module structure, and let $\psi \in Z^2(G,A)$. Consider the set $E = A \times G$ with the product topology and the multiplication

$$(a_1,g_1)(a_2,g_2) = (a_1 + g_1a_2 + \psi(g_1,g_2), g_1g_2).$$

Show that this construction naturally produces an extension of G by A that depends, up to congruence, only on the class $[\psi] \in H^2(G,A)$.

(g) Prove that the above gives a bijection between $H^2(G,A)$ and the set $\mathcal{E}(G,A)$ of congruence classes of extensions of G by A. In particular, this endows $\mathcal{E}(G,A)$ with the structure of an abelian group. Show that its identity element corresponds to the *trivial extension* $0 \to A \to\xrightarrow{\iota} A \times G \xrightarrow{\pi} G \to 0$ with $\iota(a) = (a, e_G)$ and $\pi(a,g) = g$.

4. Restriction, Corestriction, and Inflation

Let G be a profinite group and $H \subset G$ a closed subgroup. The "restriction of scalars" functor $\mathrm{Res}^G_H : \mathrm{Mod}_G \to \mathrm{Mod}_H$ is clearly an exact functor, and hence $(H^i(H, \mathrm{Res}^G_H -), \delta^i)$ is a δ-functor.

Definition 4.1. For any G-module M we have an embedding $M^G \hookrightarrow M^H$. This gives a natural transformation of functors $H^0(G,-) \to H^0(H, \mathrm{Res}_H^G -)$. By the universality of $(H^i(G,-))$, we get natural transformations $\mathrm{res} : H^i(G,-) \to H^i(H, \mathrm{Res}_H^G -)$ for all $i > 0$ as well. This is called *restriction*.

Lemma 4.2. *Let $H \subset G$ be an open subgroup. For each $i > 0$, the functor $H^i(H, \mathrm{Res}_H^G -)$ is effaceable.*

Proof. Let M be a G-module. As we saw in the proof of Corollary 2.12, as an abelian group M injects into a divisible group I. We concluded, using Frobenius reciprocity, that M injects, as a G-module, into the injective G-module $\mathrm{Ind}_{\{e\}}^G I$. To establish our claim, it suffices to show that $\mathrm{Res}_H^G \mathrm{Ind}_{\{e\}}^G I$ is an injective H-module. In the exercises for Section 2 we showed that any restriction of an injective module is injective. Here we give a more direct proof.

Observe that $\mathrm{Res}_H^G \mathrm{Ind}_{\{e\}}^G I = \bigoplus_{gH \in G/H} \Phi_{gH}$, where Φ_{gH} is the abelian group of locally constant functions $gH \to I$; note that each Φ_{gH} is stable under the H-action on $\mathrm{Res}_H^G \mathrm{Ind}_{\{e\}}^G I$. For each $g \in G$, there is an H-module isomorphism

$$\begin{array}{rcl} \Phi_{gH} & \to & \mathrm{Ind}_{\{e\}}^H I \\ f & \mapsto & (h \mapsto f(gh)). \end{array}$$

Since $\mathrm{Ind}_{\{e\}}^H I$ is an injective H-module by Corollary 2.11, and a direct sum of injective H-modules is injective, we are done. $\square$

Definition 4.3. Let $H \subset G$ be an open subgroup and M a G-module. If $m \in M^H$ and $g \in G$, then gm depends only on the left coset gH. Therefore, since H has finite index in G, we have a map

$$\begin{array}{rcl} M^H & \to & M^G \\ m & \mapsto & \displaystyle\sum_{gH \in G/H} gm. \end{array}$$

It is simple to check that the right-hand side is indeed G-invariant. This is a natural transformation $\mathrm{tr}_{G/H} : H^0(H, \mathrm{Res}_H^G -) \to H^0(G,-)$. Since $(H^i(H, \mathrm{Res}_H^G -), \delta^i)$ is a universal δ-functor by Lemma 3.9 and the previous lemma, we obtain a natural transformation $\mathrm{cor} : H^i(H, \mathrm{Res}_H^G -) \to H^i(G,-)$ for all $i \geq 0$. This is called *corestriction*.

Lemma 4.4. *Let $H \subset G$ be an open subgroup and M a G-module. For any $i \geq 0$, the composition* $\mathrm{cor} \circ \mathrm{res} : H^i(G,M) \to H^i(G,M)$ *is multiplication by the index $[G:H]$.*

Proof. If $m \in M^G$, then $\mathrm{tr}_{G/H}(m) = \sum_{gH \in G/H} gm = \sum_{gH \in G/H} m = [G:H]m$, so the claim is true if $i = 0$. It follows for $i > 0$ by universality. $\square$

For any open subgroup $H \subset G$ and any G-module M, consider the G-module map $i : M \to \mathrm{Ind}_H^G M$ given by $i(m) = (g \mapsto gm)$; note that $i(m)$ is locally constant because M is a discrete G-module. Similarly, consider the trace map $\mathrm{tr} : \mathrm{Ind}_H^G M \to M$ given by $\mathrm{tr}(f) = \sum_{Hg \in H \backslash G} g^{-1} f(g)$; one quickly checks that this is well-defined.

Lemma 4.5. *Let $H \subset G$ be an open subgroup and M a G-module. Then for each $i \geq 0$ there is a commutative diagram*

$$\begin{tikzcd} H^i(G,M) \arrow[r, "i"] \arrow[d, "\mathrm{res}"] & H^i(G, \mathrm{Ind}_H^G M) \arrow[d, "\mathrm{tr}"] \\ H^i(H,M) \arrow[r, "\mathrm{cor}"'] \arrow[ur, "\sim"] & H^i(G,M) \end{tikzcd}$$

where the diagonal isomorphism comes from Shapiro's Lemma.

Proof. Since all the cohomology functors here are universal, it suffices to show that in the case $i = 0$ our maps arise from a commutative diagram of natural transformations. This is easily checked using the definitions above and the proof of Shapiro's Lemma. The only point that is not completely obvious is the bottom triangle: here one uses the observation that if $\{g_1, \dots, g_r\}$ is a system of coset representatives of G/H, then $\{g_1^{-1}, \dots, g_r^{-1}\}$ are coset representatives of $H\backslash G$. □

The construction of the restriction map can be generalized. Let $f : G' \to G$ be a continuous group homomorphism, and let M be a G-module. Then M can be given the structure of a G'-module by $g' \cdot m = f(g') \cdot m$ for all $g' \in G'$ and $m \in M$. We denote this G'-module (why is it discrete?) by f^*M. Observe that if $H \subset G$ is a subgroup and $f : H \to G$ is the inclusion map, then $f^*M = \mathrm{Res}_H^G M$. Note that $M \mapsto f^*M$ is an exact functor, and clearly $M^G \subset (f^*M)^{G'}$, so by universality this inclusion gives rise to functors $H^i(G,M) \to H^i(G', f^*M)$ for all $i \geq 0$. Moreover, if $h : f^*M \to M'$ is a G'-module homomorphism, then we can consider the composition

$$(f,h)^* : H^i(G,M) \to H^i(G', f^*M) \to H^i(G', M').$$

Definition 4.6. Let G be a profinite group and $H \trianglelefteq G$ a closed normal subgroup. Consider the natural projection $f : G \to G/H$, which is continuous by definition of the quotient topology. Let M be a G-module. Then M^H is a G-submodule (since for any $g \in G$, $h \in H$, and $m \in M^H$ we have $h(gm) = g(g^{-1}hg)m = gm$). It has an obvious G/H-module structure. The inclusion $h : f^*M^H \to M$ is G-equivariant, and the map $(f,h)^*$ is called *inflation*:

$$\mathrm{inf} : H^i(G/H, M^H) \to H^i(G,M).$$

It is often very helpful to be able to compute explicitly with cocyles. Given a cocyle $\psi \in Z^i(G,M)$, we denote its cohomology class by $[\psi]$. The following explicit description of the action of the boundary map is useful.

Lemma 4.7. *Let $0 \to M \xrightarrow{\iota} N \xrightarrow{\pi} P \to 0$ be a short exact sequence of G-modules. Let $\psi \in Z^i(G,P)$. The connecting map $\delta^i : H^i(G,P) \to H^{i+1}(G,M)$ sends $[\psi]$ to the cohomology class of $\eta \in Z^{i+1}(G,M)$, where $\eta : G^{i+1} \to M$ is given by*

$$\begin{aligned} \eta(g_1, \dots, g_{i+1}) \quad = \quad & g_1\psi(\widetilde{g_2, \dots, g_{i+1}}) - \psi(g_1g_2, \widetilde{g_3, \dots}, g_{i+1}) + \psi(g_1, g_2\widetilde{g_3, g_4}, \dots, g_{i+1}) + \cdots + \\ & (-1)^i\psi(g_1, \widetilde{g_2, \dots,} g_ig_{i+1}) + (-1)^{i+1}\psi(\widetilde{g_1, \dots,} g_i). \end{aligned}$$

Here, for any $p \in P$, we denote by $\tilde{p}$ an arbitrary lift of p to N.

Proof. This essentially follows from (5) and the proof of Proposition 3.30. However, some care must be taken, since the resolutions

$$0 \to M \to \mathcal{C}^0(G,M) \to \mathcal{C}^1(G,M) \to \mathcal{C}^2(G,M) \to \cdots$$

that we are using to compute cohomology need not be injective. Thus the machinery of Section 3.4 does not immediately apply. We observe directly that the sequences

$$0 \to \mathcal{C}^i(G,M) \to \mathcal{C}^i(G,N) \to \mathcal{C}^i(G,P) \to 0,$$

where the maps are given by pre-composition with ι and π, are exact, although not necessarily split. Moreover, after applying the G-invariants functor we are left with an exact sequence

$$0 \to C^i(G,M) \to C^i(G,N) \to C^i(G,P) \to 0.$$

Now we may continue as in the proof of Proposition 3.30 and compute the connecting maps δ^i by applying the Snake Lemma to the diagrams

$$\begin{tikzcd}
0 \arrow[r] & C^i(G,M) \arrow[r] \arrow[d, "d_i"] & C^i(G,N) \arrow[r] \arrow[d, "d_i"] & C^i(G,P) \arrow[r] \arrow[d, "d_i"] & 0 \\
0 \arrow[r] & C^{i+1}(G,M) \arrow[r] & C^{i+1}(G,N) \arrow[r] & C^{i+1}(G,P) \arrow[r] & 0.
\end{tikzcd}$$

□

It will also be useful to know how restriction and inflation act on cochains.

Lemma 4.8. *Let G be a profinite group, and let $H \subset G$ be a closed subgroup. Let M be a G-module.*

(1) *Let $\psi \in Z^i(G,M)$. Then $\mathrm{res}^G_H([\psi]) = [\psi_{|H^i}]$.*
(2) *Suppose that H is normal. Let $\psi \in Z^i(G/H, M^H)$. Then $\inf([\psi]) = [\tilde{\psi}]$, where $\tilde{\psi} : G^i \to M$ is the composition $G^i \twoheadrightarrow (G/H)^i \xrightarrow{\psi} M$.*

Proof. The claim holds trivially for $i = 0$. Since the formulas in our claim are compatible with the boundary maps by Lemma 4.7, we obtain the claim in general by the universal property used to define the restriction and inflation maps. □

Definition 4.9. Let G be a profinite group and let $\{G_i\}_{i \in I}$ be a projective system of finite groups, with connecting homomorphisms $\varphi_{ij} : G_i \to G_j$ for $i \geq j$, such that $G = \varprojlim G_i$. Suppose that for each $i \in I$ we have a G_i-module M_i. Moreover, suppose that whenever $i \geq j$, we have a G_i-module homomorphism $h_{ij} : \varphi_{ij}^* M_j \to M_i$. Then we can define a G-module structure on the direct limit $M = \varinjlim M_i$: if $g = (g_i) \in G$ and $m = (m_i) \in M$, we define $gm = (g_i m_i) \in M$. (Why is this a discrete G-module?)

Proposition 4.10. *Let $G = \varprojlim G_i$ be a profinite group, and let $\{M_i\}$ be a system of G_i-modules as above. Then*

$$H^k(G,M) \simeq \varinjlim H^k(G_i, M_i)$$

for all $k \geq 0$. The connecting homomorphisms on the right-hand side are the maps $(\varphi_{ij}, h_{ij})^*$.

Proof. If $\pi_i : G \to G_i$ are the natural projections, then the maps $(\pi_i : M_i \to M)^* : H^k(G_i, M_i) \to H^k(G,M)$ are clearly compatible with the $(\varphi_{ij}, h_{ij})^*$, so by the universal property of direct limits we get a map $\varinjlim H^k(G_i, M_i) \to H^k(G,M)$. We claim that it is an isomorphism. This can be established by checking on cocycles: it is easy to show that $C^k(G,M) \simeq \varinjlim C^k(G_i, M_i)$. □

Remark 4.11. The previous proposition does not assume that the G_i are finite groups. Indeed, if $G_i = G$ for all i, and all the homomorphisms φ_{ij} are identities, then we clearly have $G = \varprojlim G_i$ and the proposition states that cohomology commutes with direct limits:

$$H^k(G, \varinjlim M_i) = \varinjlim H^k(G, M_i).$$

Corollary 4.12. *Let K be a field, and let its absolute Galois group $G_K = \mathrm{Gal}(\overline{K}/K)$ act on $\overline{K}^\times$ in the natural way. Then $H^1(G_K, \overline{K}^\times) = 0$.*

Proof. This is immediate from Proposition 4.10 and Hilbert 90. □

Corollary 4.13. *Let G be a profinite group, and let M be a $\mathbb{Q}$-vector space with a G-module structure. Then $H^i(G, M) = 0$ for all $i > 0$.*

Proof. Let $\{H_j\}_{j\in J}$ be the family of open normal subgroups of G, ordered by reverse inclusion. Since $M = \bigcup_{j\in J} M^{H_j}$, it is easy to see that $M = \varinjlim M^{H_j}$, where the connecting homomorphisms are the inclusions $M^{H_j} \subset M^{H_k}$ for $j \le k$, i.e. $H_k \subseteq H_j$. Thus $H^i(G, M) = \varinjlim H^i(G/H_j, M^{H_j})$ by Proposition 4.10. Since the M^{H_j} are all $\mathbb{Q}$-vector spaces, it suffices to prove our claim in the case where G is a finite group.

So let G be finite. In this case, $\{e\} \subset G$ is an open subgroup, so by Lemma 4.4 the composition

$$H^i(G, M) \overset{\text{res}}{\to} H^i(\{e\}, \mathrm{Res}^G_{\{e\}} M) \overset{\text{cor}}{\to} H^i(G, M)$$

is multiplication by $|G|$. On the other hand, this composition is the zero map, since we have $H^i(\{e\}, \mathrm{Res}^G_{\{e\}} M) = 0$ for $i > 0$ by Lemma 3.12. Since $H^i(G, M)$ is naturally a $\mathbb{Q}$-vector space (the spaces of cochains are, and the boundary maps commute with the $\mathbb{Q}$-vector space structure), it follows that $H^i(G, M) = 0$. □

Definition 4.14. A G-module M is said to be a *torsion module* if every element of m is annihilated by some integer.

Lemma 4.15. *Let G be a profinite group. Suppose there exists $i \ge 1$ such that $H^i(G, M) = 0$ for all G-modules M of finite cardinality. Then $H^j(G, M) = 0$ for all $j \ge i$ and for all torsion G-modules M.*

Proof. First we prove the claim in the case $j = i$. Observe that every element m of a torsion G-module M is contained in a finite G-module. Indeed, since $\mathrm{stab}_G(m)$ has finite index in G, the G-orbit of m has only finitely many elements, say $m_1, \dots, m_r$. The subgroup $\langle m_1, \dots, m_r\rangle$ of the abelian group M is thus finite, and it is clearly stable under the G-action. Now let $\{M_k\}_{k\in K}$ be the family of finite G-submodules of M, ordered by inclusion and with the natural inclusions as the connecting homomorphisms. By the considerations above we have $M = \varinjlim M_k$. By Proposition 4.10 (note the remark following it) we have $H^i(G, M) = \varinjlim H^i(G, M_k) = 0$.

Now suppose that the claim is known for $j-1$. We will prove it for j by a "dimension shifting" argument. For brevity, we will write $\mathrm{Ind}^G_{\{e\}} M$ for $\mathrm{Ind}^G_{\{e\}} \mathrm{Res}^G_{\{e\}} M$. Observe that this is a torsion G-module. Indeed, every element is a locally constant function $f : G \to M$. Since G is compact, the function f only takes on finitely many values, so there is an integer that annihilates all of them. There is a natural injection $\varepsilon : M \to \mathrm{Ind}^G_{\{e\}} M$, where $\varepsilon(m)$ is the function $g \mapsto gm$. Let Q be the cokernel; it is a quotient of a torsion module and thus is a torsion module itself. The short exact sequence $0 \to M \to \mathrm{Ind}_{\{e\}}^G M \to Q \to 0$ gives rise to the following bit of the long exact sequence:

$$\cdots \to H^{j-1}(G, \mathrm{Ind}_{\{e\}}^G M) \to H^{j-1}(G, Q) \to H^j(G, M) \to H^j(G, \mathrm{Ind}_{\{e\}}^G M) \to \cdots$$

Since $H^{j-1}(G, Q) = 0$ by the inductive hypothesis and $H^j(G, \mathrm{Ind}_{\{e\}}^G M) = 0$ by Shapiro's Lemma, we conclude that $H^j(G, M) = 0$. □

Our next goal is to prove the "inflation-restriction exact sequence": if G is a profinite group, $H \trianglelefteq G$ is a normal subgroup, and M is a G-module, then the following sequence is exact:

$$0 \to H^1(G/H, M^H) \stackrel{\mathrm{inf}}{\to} H^1(G, M) \stackrel{\mathrm{res}}{\to} H^1(H, \mathrm{Res}_H^G M)^{G/H} \to H^2(G/H, M^H) \stackrel{\mathrm{inf}}{\to} H^2(G, M).$$

Here the action of G on $H^1(H, \mathrm{Res}_H^G M)$ is the one arising naturally from the G-action on cocycles, and the unlabelled map has yet to be defined. The desired result could be proven by direct computations on cocycles; while long and unpleasant, such a proof would still be shorter than the one we will give. However, we will take the opportunity to develop the machinery of spectral sequences in the next section. This is a tool for computing cohomology that turns out to be ubiquitous in number theory and algebraic geometry. The inflation-restriction sequence will turn out to be a special case of a very general phenomenon.

Exercises

(1) A triple $(x, y, z) \in \mathbb{Z}^3$ is called a *Pythagorean triple* if $x^2 + y^2 = z^2$. It has been known since ancient times that (x, y, z) is a Pythagorean triple if and only if it is proportional to $(m^2 - n^2, 2mn, m^2 + n^2)$ for some integers m, n. It was observed by N. Elkies that this fact can be deduced straightforwardly from Hilbert 90. Do it.

Hint: Let (x, y, z) be a Pythagorean triple such that $z \neq 0$. Consider the element $\frac{x+y\sqrt{-1}}{z} \in \mathbb{Q}(\sqrt{-1})$.

(2) Generalize your solution to the preceding exercise to obtain a parametrization of all solutions (x, y, z) to the Diophantine equation $x^2 + axy + by^2 = z^2$, where $a, b \in \mathbb{Z}$ are such that $a^2 - 4b$ is not a perfect square.

(3) Let K be a field of characteristic prime to $n \geq 1$. Let $\mu_n \subset \overline{K}^\times$ be the subgroup of n-th roots of unity; clearly it is preserved by the action of the absolute Galois group G_K. Prove that $H^1(G_K, \mu_n) \simeq K^\times/(K^\times)^n$.

(4) Let G be a finite group and let M be a G-module. Show that, for any $i \geq 0$ and any $c \in H^i(G, M)$, the equality $|G|c = 0$ holds.

(5) Let G be a finite group and let M be a G-module which is finitely generated as an abelian group. Prove that $H^i(G, M)$ is finite.

5. Spectral sequences

5.1. **Basic definitions.** The goal of this section is to understand the Hochschild-Serre spectral sequence, which is an important tool for computing cohomology. We start off with a general treatment of spectral sequences. The underlying theory is not complicated, once one sees past all the cluttered diagrams. Let R be a ring.

Definition 5.1. (1) A *bigraded R-module* E is a family of R-modules $\{E^{p,q} : p, q \in \mathbb{Z}\}$.

(2) A *differential* $d : E \to E$ of degree (r, s) is a family of R-module homomorphisms $d : E^{p,q} \to E^{p+r,q+s}$ such that $d \circ d = 0$.

Definition 5.2. A *spectral sequence* is a sequence $\{E_r\}_{r\geq t}$, for some integer t, of bigraded R-modules $E_r = (E_r^{p,q})$ equipped with differentials $d_r : E_r \to E_r$ of degree $(r, 1-r)$ such that, for every $p, q \in \mathbb{Z}$ and every integer $r \geq r_0$, the following holds:

$$E_{r+1}^{p,q} \simeq \ker(d_r : E_r^{p,q} \to E_r^{p+r,q-r+1})/\mathrm{im}(d_r : E_r^{p-r,q+r-1} \to E_r^{p,q}). \tag{7}$$

The spectral sequence is called *positive* if $E_r^{p,q} = 0$ whenever $p < 0$ or $q < 0$. Most of the interesting applications of spectral sequences happen in the case $t = 2$, and we will assume this from now on, except in Example 5.6 below.

From now on, all our spectral sequences will be assumed to be positive. The bigraded modules E_r are called *sheets* of the spectral sequence. For a fixed r, think of the modules $E_r^{p,q}$ as lattice points on the plane. Then the differential maps are slanted arrows that form complexes which run in slanted lines across the plane. The (co)homology of these complexes computes the modules $E_{r+1}^{p,q}$ of the following sheet.

The terms of the form $E_r^{p,0}$ and $E_r^{0,q}$ are called base terms and fiber terms, respectively.

Lemma 5.3. *Let $\{E_r\}$ be a positive spectral sequence. Fix $p, q \in \mathbb{N}$, and let $r_0 = \max\{p, q+1\}+1$. Then $E_r^{p,q} \simeq E_{r_0}^{p,q}$ for all $r \geq r_0$.*

Proof. If $r \geq r_0$, then $E_r^{p+r,q-r+1}$ and $E_r^{p-r,q+r-1}$ both vanish. The claim is then immediate from (7). □

Definition 5.4. A *filtered R-module* is an R-module A together with a family of submodules

$$A = F^0 A \supseteq F^1 A \supseteq F^2 A \supseteq \cdots$$

We will assume throughout that $\bigcap_{i=0}^{\infty} F^i A = 0$. The graded pieces are the quotients $\mathrm{gr}^i A = F^i A/F^{i+1} A$.

Definition 5.5. Let $\{E_r^{p,q}\}$ be a spectral sequence.

(1) For any $p, q \in \mathbb{N}$, we define $E_\infty^{p,q} = E_{r_0}^{p,q}$, in the notation of Lemma 5.3.
(2) We say that the spectral sequence $\{E_r^{p,q}\}$ *converges to* or *abuts to* the family $\{A^n\}_{n\in\mathbb{N}}$ of filtered modules if $E_\infty^{p,q} \simeq \mathrm{gr}^p A^{p+q}$ for all $p, q \in \mathbb{N}$. In this case, one writes $E_r^{p,q} \Rightarrow A^n$ or $E_2^{p,q} \Rightarrow A^n$.

Example 5.6. As our simplest example of a spectral sequence, let $C^0 \xrightarrow{\partial^0} C^1 \xrightarrow{\partial^1} \cdots$ be a cochain complex. Define $E_1^{p,q} = C^p$ whenever p and q are both non-negative, and set $d_1 : E_1^{p,q} \to E_1^{p+1,q}$ to be ∂^p. Then (7) forces $E_2^{p,q} = H^p(\mathbf{C})$. The only natural choice for the differential is the zero map. Thus we get $E_\infty^{p,q} = E_2^{p,q} = H^p(\mathbf{C})$ for all p, q. If for all n we define $H^n = \bigoplus_{i=0}^{\infty} H^i(\mathbf{C})$, with the grading $F^j H^n = \bigoplus_{i=j}^{\infty} H^i(\mathbf{C})$, then clearly $E_1^{p,q} \Rightarrow H^n$. While this example is silly, it already suggests that spectral sequences could be useful for computing cohomology.

5.2. **The five-term exact sequence.** In this section we derive a general five-term exact sequence associated to any positive spectral sequence. The inflation-restriction sequence will be the special case of this result for the Hochschild-Serre spectral sequence.

Lemma 5.7. *Let $\{E_r^{p,q}\} \Rightarrow A^n$ be a spectral sequence. For each $n \geq 1$ there is a natural injection $e_B : E_\infty^{n,0} \hookrightarrow A^n$ and a natural surjection $e_F : A^n \twoheadrightarrow E_\infty^{0,n}$. Moreover, in the case $n = 1$ the sequence*

$$E_\infty^{1,0} \xrightarrow{e_B} A^1 \xrightarrow{e_F} E_\infty^{0,1}$$

is exact.

Proof. Observe by the definition of abutment that if $m > n$, then $\mathrm{gr}^m A^n \simeq E_\infty^{m,n-m} = 0$, since our spectral sequences are positive. Since $\bigcap_i F^i A^n = 0$, this implies that $F^{n+1}A^n = 0$, and hence that $\mathrm{gr}^n A^n = F^n A^n$ is a submodule of A^n. Since $E_\infty^{n,0} \simeq \mathrm{gr}^n A^n$, this provides our injection.

Similarly, $E_\infty^{0,n} \simeq \mathrm{gr}^0 A^n = A^n/F^1 A^n$. Since this is a quotient of A^n, we obtain the desired surjection. It follows from the above that $\ker e_F = F^1 A^n$ for all n, whereas $\mathrm{im}\, e_B = F^n A^n$, and hence we get exactness when $n = 1$. □

Lemma 5.8. *Let $\{E_r^{p,q}\}$ be a spectral sequence. Then for all $n \geq 1$ there is a natural surjection $E_2^{n,0} \twoheadrightarrow E_\infty^{n,0}$ and injection $E_\infty^{0,n} \hookrightarrow E_2^{0,n}$.*

Proof. This is immediate from (7). Indeed, for every $r \geq 2$ we have $E_r^{n+r,-r+1} = 0$ by positivity, whence $E_{r+1}^{n,0} = E_r^{n,0}/\mathrm{im}(d_r : E_r^{n-r,r-1} \to E_r^{n,0})$ is a quotient of $E_r^{n,0}$. By induction, every $E_r^{n,0}$ is thus naturally a quotient of $E_2^{n,0}$.

Similarly, $E_{r+1}^{0,n} = \ker(d_r : E_r^{0,n} \to E_r^{r,n-r+1}) \subseteq E_r^{0,n}$ for every $r \geq 1$, and hence $E_r^{0,n} \subseteq E_2^{0,n}$ for every $r \geq 2$. □

Definition 5.9. The compositions of the maps from Lemmas 5.7 and 5.8 provide maps $E_2^{n,0} \to A^n$ and $A^n \to E_2^{0,n}$, which we abusively call e_B and e_F, respectively. Note that these maps are not necessarily injective or surjective.

Definition 5.10. Let $n \geq 1$. The spectral sequence $\{E_r^{p,q}\}$ is said to *satisfy condition* $(*)_n$ if $E_2^{p,q} = 0$ for all pairs p, q such that $1 \leq q \leq n-1$ and $p + q \in \{n-1, n, n+1\}$. Observe that the condition $(*)_1$ is vacuous and thus is satisfied by all spectral sequences.

Remark 5.11. Observe that if condition $(*)_n$ holds, then for all $2 \leq r \leq n$ we have $E_2^{r,n-r+1} = 0$, and hence $E_r^{r,n-r+1} = 0$ by (7). This implies that the inclusions $E_{n+1}^{0,n} \subseteq E_n^{0,n} \subseteq \cdots \subseteq E_2^{0,n}$ of Lemma 5.8 are isomorphisms. Similarly, we have $E_2^{(n+1)-r,r-1} = 0$ for all $2 \leq r \leq n$, and thus the projections $E_2^{n+1,0} \twoheadrightarrow E_3^{n+1,0} \twoheadrightarrow \cdots \twoheadrightarrow E_{n+1}^{n+1,0}$ are also all isomorphisms. This means that, for all $n \geq 1$, we may consider the composition

$$E_2^{0,n} \simeq E_{n+1}^{0,n} \stackrel{d_{n+1}}{\to} E_{n+1}^{n+1,0} \simeq E_2^{n+1,0},$$

which is called the transgression map and will be denoted d_{n+1}.

We finally have all the necessary ingredients to establish the "five-term exact sequences" associated to the spectral sequence $\{E_r^{p,q}\}$.

Proposition 5.12. *Let $\{E_r^{p,q} \Rightarrow A^n\}$ be a spectral sequence satisfying condition $(*)_n$ for some $n \geq 1$. Then there is an exact sequence*

$$0 \to E_2^{n,0} \stackrel{e_B}{\to} A^n \stackrel{e_F}{\to} E_2^{0,n} \stackrel{d_{n+1}}{\to} E_2^{n+1,0} \stackrel{e_B}{\to} A^{n+1}.$$

Proof. We successively verify exactness at each node.

Exactness at $E_2^{n,0}$: We need to show that $e_B : E_2^{n,0} \to A^n$ is injective. Since $E_{n+1}^{n,0} = E_\infty^{n,0} \to A^n$ is injective by Lemma 5.8, it suffices to show that the maps $E_r^{n,0} \twoheadrightarrow E_{r+1}^{n,0}$ are injective, hence isomorphisms, for all $2 \leq r \leq n$. Now by (7) we have $\ker(E_r^{n,0} \to E_{r+1}^{n,0}) = \mathrm{im}(d_r : E_r^{n-r,r-1} \to E_r^{n,0}) = 0$ for each $2 \leq r \leq n$, since $E_r^{n-r,r-1} = 0$ by condition $(*)_n$.

Exactness at A^n: In view of the injectivity of e_B and the end of the proof of Lemma 5.7, we find that $\mathrm{im}\, e_B = F^n A^n$, whereas $\ker\, e_F = F^1 A^n$. However, these are equal, since condition $(*)^n$ implies that for all $1 \leq i \leq n-1$ we have $\mathrm{gr}^i A^n = E_\infty^{i,n-i} = 0$.

Exactness at $E_2^{0,n}$: Observe by Lemma 5.3 that $E_\infty^{0,n} = E_{n+2}^{0,n}$ and $A^n \to E_\infty^{0,n}$ is surjective. Thus $\operatorname{im} e_F = \operatorname{im}(E_{n+2}^{0,n} \hookrightarrow E_{n+1}^{0,n}) = \ker(d_{n+1} : E_{n+1}^{0,n} \to E_{n+1}^{n+1,0})$, which implies exactness by the definition of the transgression map.

Exactness at $E_2^{n+1,0}$: Again by Lemma 5.3, we see that $E_{n+2}^{n+1,0} = E_\infty^{n+1,0}$ injects into A^{n+1}, and hence that $\ker e_B = \ker(E_{n+1}^{n+1,0} \twoheadrightarrow E_{n+2}^{n+1,0}) = \operatorname{im}(d_{n+1} : E_{n+1}^{0,n} \to E_{n+1}^{n+1,0})$, which is what we need by the definition of the transgression map. □

Corollary 5.13. *Let $E_r^{p,q} \Rightarrow A^n$ be any positive spectral sequence. Then there is an exact sequence*

$$0 \to E_2^{1,0} \xrightarrow{e_B} A^1 \xrightarrow{e_F} E_2^{0,1} \xrightarrow{d_2} E_2^{2,0} \xrightarrow{e_B} A^2.$$

Proof. As noted above, the condition $(*)_1$ is vacuous. □

5.3. **Double complexes.** So far the only example we have seen of a spectral sequence is the trivial one in Example 5.6. In this section we will study a general construction that produces many useful examples of spectral sequences.

Definition 5.14. A filtered complex is a complex $C^0 \xrightarrow{\partial^0} C^1 \xrightarrow{\partial^1} \cdots$ of filtered modules whose filtrations are compatible with the boundary maps. In other words, for every i, j, we have $\partial^i(F^j C^i) \subseteq F^j C^{i+1}$.

A filtration of a complex induces a filtration on its cohomology. Indeed, for every $i, j \geq 0$ we may define $F^j H^i(\mathbf{C}) = \operatorname{im}(H^i(F^j\mathbf{C}) \to H^i(\mathbf{C}))$, where the maps whose image we are considering is the one induced by the natural inclusion of complexes $F^j\mathbf{C} \hookrightarrow \mathbf{C}$. The family of filtered modules $\{H^i(\mathbf{C})\}$ will be denoted $H(\mathbf{C})$.

For every $r \in \mathbb{Z}$, we set $Z_r^{p,q} = \{x \in F^pC^{p+q} : \partial^{p+q}(x) \in F^{p+r}C^{p+q+1}\}$. Similarly, set $B_r^{p,q} = \partial^{p+q-1} Z_{r-1}^{p-r+1,q+r-2} = \partial^{p+q-1}F^{p-r+1}C^{p+q-1} \cap F^pC^{p+q}$, the second equality following immediately from the definition. Observe that $B_r^{p,q} \subseteq Z_r^{p,q}$, since the composition of two boundary maps is zero. Similarly, note that $Z_{r-1}^{p+1,q-1} \subseteq Z_r^{p,q}$. Now set

$$E_r^{p,q} = Z_r^{p,q}/(B_r^{p,q} + Z_{r-1}^{p+1,q-1}). \tag{8}$$

Proposition 5.15. *Let $\mathbf{C}$ be a filtered complex as above.*

(1) *There exists a spectral sequence, with $E_r^{p,q}$ defined as in* (8) *and differentials induced by the boundary maps of $\mathbf{C}$.*
(2) *If the filtration of each C^i is bounded, so that there exists some j such that $F^jC^i = 0$ (we allow j to depend on i), then $E_r^{p,q} \Rightarrow H(\mathbf{C})$.*

Proof. The definitions of $Z_r^{p,q}$ and $E_r^{p,q}$ are exactly the ones needed to make this claim work. Since $\partial^{p+q} Z_r^{p,q} \subseteq Z_r^{p+r,q-r+1}$ by definition, the boundary map ∂^{p+q} clearly induces a map $d_r : E_r^{p,q} \to E_r^{p+r,q-r+1}$, and the composition of two such maps is zero.

Again by definition, $\operatorname{im}(\partial^{p+q-1} : Z_r^{p-r,q+r-1} \to Z_r^{p,q}) = B_{r+1}^{p,q}$. Observing that $B_r^{p,q} \subseteq B_{r+1}^{p,q}$, we conclude that

$$\operatorname{im}(d_r : E_r^{p-r,q+r-1} \to E_r^{p,q}) = (B_{r+1}^{p,q} + Z_{r-1}^{p+1,q-1})/(B_r^{p,q} + Z_{r-1}^{p+1,q-1}).$$

Furthermore, for $x \in Z_r^{p,q}$, we have $\partial^{p+q}(x) \in Z_r^{p+r+1,q-r}$ if and only if $\partial^{p+q}(x) \in F^{p+r+1}C^{p+q+1}$, which is equivalent to $x \in Z_{r+1}^{p,q}$. Also, $B_r^{p+r,q-r+1} = \partial^{p+q}(Z_{r-1}^{p+1,q-1})$. Thus, $\partial^{p+q}(x) \in B_r^{p+r,q-r+1}$ implies $x \in Z_{r-1}^{p+1,q-1} + \ker \partial^{p+q}$.

Hence, $\ker(d_r : E_r^{p,q} \to E_r^{p+r,q-r+1}) = (Z_{r+1}^{p,q} + Z_{r-1}^{p+1,q-1})/(B_r^{p,q} + Z_{r-1}^{p+1,q-1})$.

Therefore, $\ker(d_r : E_r^{p,q} \to E_r^{p+r,q-r+1})/\mathrm{im}\,(d_r : E_r^{p-r,q+r-1} \to E_r^{p,q}) \simeq E_{r+1}^{p,q}$, as claimed.

To prove the second part of the proposition, let p,q be fixed. For sufficiently large r, we have that $F^{p+r}C^{p+q+1} = 0$, and hence $Z_r^{p,q} = \{x \in F^pC^{p+q} : \partial^{p+q}(x) = 0\}$. But for sufficiently large r we have $F^{p-r+1}C^{p+q-1} = C^{p+q-1}$ and hence $B_r^{p,q} = \partial^{p+q-1}C^{p+q-1} \cap F^pC^{p+q}$. It follows that $E_\infty^{p,q}$ is exactly $\mathrm{im}\,(H^{p+q}(F^p\mathbf{C}) \to H^{p+q}(\mathbf{C}))/\mathrm{im}\,(H^{p+q}(F^{p+1}\mathbf{C}) \to H^{p+q}(\mathbf{C}))$. $\square$

Definition 5.16. A double complex is a family $\mathbf{K} = \{K^{p,q}\}$ of R-modules, where the indices p and q run over the natural numbers, equipped with horizontal and vertical differential maps

$$\begin{aligned} \partial' : K^{p,q} &\to K^{p+1,q} \\ \partial'' : K^{p,q} &\to K^{p,q+1} \end{aligned}$$

such that $\partial' \circ \partial' = 0$, $\partial'' \circ \partial'' = 0$, and $\partial' \circ \partial'' + \partial'' \circ \partial' = 0$.

Observe that a double complex is not a commutative diagram, since the squares

$$\begin{tikzcd} K^{p,q} \arrow[r, "\partial'"] \arrow[d, "\partial''"'] & K^{p+1,q} \arrow[d, "\partial''"'] \\ K^{p,q+1} \arrow[r, "\partial'"] & K^{p+1,q+1} \end{tikzcd}$$

anti-commute. Moreover, it is conventional to refer to the maps ∂' and ∂'' as horizontal and vertical, respectively, so one visualizes $K^{p,q}$ as lying in the q-th row and p-th column of the double complex. Given a double complex as in Definition 5.16, we define a (usual) cochain complex $\mathbf{C}$ as follows: for each $i \geq 0$ set

$$C^i = \bigoplus_{p+q=i} K^{p,q} \tag{9}$$

and define the boundary maps $\partial^i : C^i \to C^{i+1}$ by $\partial^i = \partial' + \partial''$. It is simple to check that $\partial^{i+1} \circ \partial^i = 0$ for all $i \geq 0$; this is the reason for the condition $\partial' \circ \partial'' + \partial'' \circ \partial' = 0$ in the definition of double complexes. The complex $\mathbf{C}$ is called the *total complex* of $\mathbf{K}$ and is often denoted $\mathrm{Tot}\,\mathbf{K}$ in the literature.

We give $\mathbf{C}$ the structure of a filtered complex in two different ways. Define two filtrations

$$\begin{aligned} {}'F^jC^i &= \bigoplus_{\substack{p+q=i \\ p\geq j}} K^{p,q} \\ {}''F^jC^i &= \bigoplus_{\substack{p+q=i \\ q\geq j}} K^{p,q} \end{aligned}$$

Both filtrations are compatible with the boundary maps and are obviously bounded. By Proposition 5.15 we get two spectral sequences, $\{{}'E_r^{p,q}\}$ and $\{{}''E_r^{p,q}\}$, that both abut to the family $H(\mathbf{C})$, although with different filtrations on $H(\mathbf{C})$.

We now investigate the terms of these sequences. For any element $x \in A$ of a filtered module A with bounded filtration, let $\deg(x) = \max\{p : x \in F^pA\}$. Considering the filtrations ${}'F^jC^i$, observe that if $x \in C^i$, then $\deg(\partial'(x)) = \deg(x) + 1$, whereas $\deg(\partial''(x)) = \deg(x)$. It follows that

$${}'Z_1^{p,q} = \ker(\partial'' : K^{p,q} \to K^{p,q+1}) \oplus {}'F^{p+1}C^{p+q},$$

whereas ${}'B_1^{p,q} = \partial' F^p C^{p+q-1} \cap {}'F^p C^{p+q} = \partial' F^p C^{p+q-1}$. Now ${}'Z_0^{p+1,q-1} = {}'F^{p+1}C^{p+q}$ by definition, and $\partial'({}'F^p C^{p+q-1}) \subset {}'F^{p+1}C^{p+q}$, so

$${}'B_1^{p,q} + {}'Z_0^{p+1,q-1} = \mathrm{im}\,(\partial'' : K^{p,q-1} \to K^{p,q}) \oplus {}'F^{p+1}C^{p+q}.$$

We conclude that

$${}'E_1^{p,q} = \ker(\partial'' : K^{p,q} \to K^{p,q+1})/\mathrm{im}\,(\partial'' : K^{p,q-1} \to K^{p,q}),$$

Thus the spectral sequence $\{{}'E_1^{p,q}\}$ computes the cohomology of the columns of the double complex $\{K^{p,q}\}$. Precisely, ${}'E_1^{p,q} = H^q(\mathbf{K}^{p,\bullet})$, where $K^{p,\bullet}$ is the cochain complex $K^{p,0} \xrightarrow{\partial''} K^{p,1} \xrightarrow{\partial''} K^{p,2} \to \cdots$.

Similarly, by definition we have

$${}''Z_1^{p,q} = \ker(\partial' : K^{q,p} \to K^{q+1,p}) \oplus {}''F^{p+1}C^{p+q}.$$

By an analogous argument to the one just above, we conclude that

$${}''E_1^{p,q} = \ker(\partial' : K^{q,p} \to K^{q+1,p})/\mathrm{im}\,(\partial' : K^{q-1,p} \to K^{q,p}) = H^q(\mathbf{K}^{\bullet,p}) \tag{10}$$

computes the cohomology of the rows of the original double complex.

Such a setup is very useful, since often one can arrange a double complex for which one of these two spectral sequences is of independent interest, whereas the other one is less interesting but has a readily computable limit. The prime example of such a situation is described in the following section.

5.4. **Grothendieck's theorem.** Recall, from Remark 3.31, the right derived functors $R^i\mathcal{F}$ of a left exact functor $\mathcal{F} : \mathcal{A} \to \mathcal{B}$, where $\mathcal{A}$ and $\mathcal{B}$ are abelian categories and $\mathcal{A}$ has enough injectives.

Definition 5.17. Let $\mathcal{G} : \mathcal{B} \to \mathcal{C}$ be a left exact functor between two abelian categories, where $\mathcal{B}$ has enough injectives. An object $B \in \mathrm{Ob}(\mathbf{B})$ is said to be $\mathcal{G}$-*acyclic* if $R^i\mathcal{G}(B) = 0$ for all $i > 0$.

Theorem 5.18 (Grothendieck)**.** *Let $\mathcal{A}, \mathcal{B}$, and $\mathcal{C}$ be abelian categories, and suppose that $\mathcal{A}$ and $\mathcal{B}$ have enough injectives. Let $\mathcal{F} : \mathcal{A} \to \mathcal{B}$ and $\mathcal{G} : \mathcal{B} \to \mathcal{C}$ be additive left exact functors, and suppose that $\mathcal{F}$ takes injective objects of $\mathcal{A}$ to $\mathcal{G}$-acyclic objects of $\mathcal{B}$. Then for each object $A \in \mathrm{Ob}(\mathcal{A})$ there exists a spectral sequence*

$$E_2^{p,q} = (R^p\mathcal{G} \circ R^q\mathcal{F})(A) \Rightarrow R^{p+q}(\mathcal{G} \circ \mathcal{F})(A).$$

Proof. Let A be an object of $\mathcal{A}$, and let $A \to I^0 \to I^1 \to I^2 \to \cdots$ be a resolution of A by injective objects. This exists because $\mathcal{A}$ has enough injectives. We denote it by $\mathbf{I}$. Suppose we can construct a "resolution of the resolution $\mathcal{F}(\mathbf{I})$," namely a commutative diagram of objects of $\mathcal{B}$ as follows, in which each row is a cochain complex (i.e. the composition of two consecutive maps is zero) and for each $p \geq 0$, the p-th column is an injective resolution of

$\mathcal{F}(I^p)$:

$$\begin{tikzcd}
\mathcal{F}(I^0) \arrow[r] \arrow[d, hook] & \mathcal{F}(I^1) \arrow[r] \arrow[d, hook] & \mathcal{F}(I^2) \arrow[r] \arrow[d, hook] & \cdots \\
J^{0,0} \arrow[r] \arrow[d] & J^{1,0} \arrow[r] \arrow[d] & J^{2,0} \arrow[r] \arrow[d] & \cdots \\
J^{0,1} \arrow[r] \arrow[d] & J^{1,1} \arrow[r] \arrow[d] & J^{2,1} \arrow[r] \arrow[d] & \cdots \\
J^{0,2} \arrow[r] \arrow[d] & J^{1,2} \arrow[r] \arrow[d] & J^{2,2} \arrow[r] \arrow[d] & \cdots \\
\vdots & \vdots & \vdots &
\end{tikzcd}$$

A minor annoyance is that this diagram is not a double complex, since the squares are commutative rather than anti-commutative. We remedy this by changing the sign of the vertical arrows in every other column, i.e. by replacing $d' : J^{p,q} \to J^{p,q+1}$ with $(-1)^p d'$. Clearly the columns remain injective resolutions. Now apply $\mathcal{G}$ to this diagram, to produce a double complex of objects of $\mathcal{C}$. We write $K^{p,q}$ for $\mathcal{G}(J^{p,q})$. Since the $\mathcal{F}(I^p)$ are $\mathcal{G}$-acyclic by assumption, the columns of the double complex $\{K^{p,q}\}$ remain exact. In particular, considering the two spectral sequences associated to this double complex, we find that ${}'E_1^{p,q} = H^q(\mathbf{K}^{p,\bullet}) = 0$ whenever $q > 0$, whereas ${}'E_1^{p,0} = \mathcal{G}(\mathcal{F}(I^p))$. Thus the differential maps $d_1 : {}'E_1^{p,q} \to {}'E_1^{p+1,q}$ induce

$${}'E_2^{p,q} = \begin{cases} R^p(\mathcal{G} \circ \mathcal{F})(A) & : q = 0 \\ 0 & : q > 0. \end{cases}$$

Since the differentials of ${}'E_2$ connect objects in different rows of the sheet, it is clear that they are all zero maps. Hence ${}'E_\infty^{p,q} = {}'E_2^{p,q}$ for all p, q. Thus, for every $p \geq 0$ we see that $H^p(\mathbf{C})$ has only one non-zero graded piece, where $\mathbf{C}$ is the filtered complex defined as in (9) from the double complex $\{K^{p,q}\}$. We conclude for all p that

$$H^p(\mathbf{C}) = R^p(\mathcal{G} \circ \mathcal{F})(A). \tag{11}$$

To compute the left-hand side of this equation in a different way, we will use the second spectral sequence associated to the double complex $\mathbf{K}$. It is determined by the cohomology of the rows of $\mathbf{K}$, over which we don't have much control in the generality in which we have worked so far in this proof. We will need to construct the diagram of resolutions $\{J^{p,q}\}$ in a rather specific way. Thereby we will also prove that such diagrams exist; recall that their existence was only assumed above.

We started with an injective resolution $0 \to A \to \mathbf{I}$ of $A \in \mathrm{Ob}(\mathcal{A})$ and applied the functor $\mathcal{F}$ to it to obtain a complex

$$0 \to \mathcal{F}(A) \to \mathcal{F}(I^0) \xrightarrow{d^0} \mathcal{F}(I^1) \xrightarrow{d^1} \cdots$$

Defining $B^i = \operatorname{im} d^{i-1}$ and $Z^i = \ker d^i$ for each $i \geq 0$, we refine this sequence to a sequence

$$Z^0 \hookrightarrow \mathcal{F}(I^0) \twoheadrightarrow B^1 \hookrightarrow Z^1 \hookrightarrow \mathcal{F}(I^1) \twoheadrightarrow B^2 \hookrightarrow Z^2 \hookrightarrow \mathcal{F}(I^2) \cdots$$

which, by construction, is exact at each $\mathcal{F}(I^i)$. We wish to construct a commutative diagram of the form

$$\begin{tikzcd}
Z^0 \arrow[r, hook] \arrow[d, hook] & \mathcal{F}(I^0) \arrow[r, two heads] \arrow[d, hook] & B^1 \arrow[r, hook] \arrow[d, hook] & Z^1 \arrow[r, hook] \arrow[d, hook] & \mathcal{F}(I^1) \arrow[r, two heads] \arrow[d, hook] & \cdots \\
L^{0,0} \arrow[r, hook] \arrow[d] & J^{0,0} \arrow[r, two heads] \arrow[d] & N^{1,0} \arrow[r, hook] \arrow[d] & L^{1,0} \arrow[r, hook] \arrow[d] & J^{1,0} \arrow[r, two heads] \arrow[d] & \cdots \\
L^{0,1} \arrow[r, hook] \arrow[d] & J^{0,1} \arrow[r, two heads] \arrow[d] & N^{1,1} \arrow[r, hook] \arrow[d] & L^{1,1} \arrow[r, hook] \arrow[d] & J^{1,1} \arrow[r, two heads] \arrow[d] & \cdots \\
\vdots & \vdots & \vdots & \vdots & \vdots &
\end{tikzcd}$$

whose columns are injective resolutions and whose rows are exact at each $J^{p,q}$. Note that the rows will not in general be complexes, i.e. the composition of two consecutive horizontal arrows need not be zero. Choosing arbitrary injective resolutions of $Z^0 = \mathcal{F}(A)$ and of B^1, which exist since $\mathcal{B}$ has enough injectives, we obtain the first three columns by applying Corollary 3.28 to the short exact sequence $0 \to Z^0 \to \mathcal{F}(I^0) \to B^1 \to 0$. Choose an arbitrary injective resolution of Z^1/B^1 and apply Corollary 3.28 to the sequence $0 \to B^1 \to Z^1 \to Z^1/B^1 \to 0$ to get a resolution $\mathbf{L}^{1,\bullet}$ of Z^1; observe that the maps $N^{1,q} \to L^{1,q}$ are injective as desired. Choose an arbitrary injective resolution of B^2 and apply Corollary 3.28 to the short exact sequence $0 \to Z^1 \to \mathcal{F}(I^1) \to B^2 \to 0$ to finish the next three columns of our diagram, with exactness at each $J^{1,q}$. Continue forever. By construction, the horizontal segments $0 \to L^{p,q} \to J^{p,q} \to N^{p,q} \to 0$ are split short exact sequences for every pair (p, q).

Remark 5.19. A consequence of our construction is the following. For each $p > 0$, we applied Corollary 3.28 to $0 \to B^p \to Z^p \to Z^p/B^p \to 0$. At that point, the injective resolution $0 \to B^p \to \mathbf{N}^{p,\bullet}$ had already been chosen, and we chose an arbitrary injective resolution $0 \to Z^p/B^p \to \mathbf{M}^{p,\bullet}$ and obtained an injective resolution $0 \to Z^p \to \mathbf{L}^{p,\bullet}$ such that the rows $0 \to N^{p,q} \to L^{p,q} \to M^{p,q} \to 0$ are split short exact sequences. In particular, $M^{p,q} \simeq L^{p,q}/N^{p,q}$ and thus

$$Z^p/B^p \hookrightarrow L^{p,0}/N^{p,0} \to L^{p,1}/N^{p,1} \to L^{p,2}/N^{p,2} \to \cdots$$

is an injective resolution. Observe also that $Z^p/B^p = R^p\mathcal{F}(A)$ by the construction of right derived functors. Thus the q-th cohomology of the complex $\mathcal{G}(\mathbf{L}^{p,\bullet}/\mathbf{N}^{p,\bullet})$ computes $R^q\mathcal{G}(R^p\mathcal{F}(A))$.

Consider the commutative diagram $\{J^{p,q}\}$ obtained from the previous construction by deleting two out of every three columns. We apply $\mathcal{G}$ to it and flip the sign of the vertical maps in the odd-numbered columns as above to obtain a double complex

$$\begin{tikzcd}
\mathcal{G}(\mathcal{F}(I^0)) \arrow[r] \arrow[d, hook] & \mathcal{G}(\mathcal{F}(I^1)) \arrow[r] \arrow[d, hook] & \mathcal{G}(\mathcal{F}(I^2)) \arrow[r] \arrow[d, hook] & \cdots \\
K^{0,0} \arrow[r] \arrow[d] & K^{1,0} \arrow[r] \arrow[d] & K^{2,0} \arrow[r] \arrow[d] & \cdots \\
K^{0,1} \arrow[r] \arrow[d] & K^{1,1} \arrow[r] \arrow[d] & K^{2,1} \arrow[r] \arrow[d] & \cdots \\
\vdots & \vdots & \vdots &
\end{tikzcd}$$

where $K^{p,q} = \mathcal{G}(J^{p,q})$. The analysis of $'E_2^{p,q}$ given above applies to this double complex. Now we consider the second associated spectral sequence $''E_r^{p,q}$.

Since all the $N^{p,q}$ are injective, it is possible to complete each triangle

$$\begin{tikzcd}
N^{p,q} \arrow[r, hook] \arrow[d, "\mathrm{id}"'] & L^{p,q} \arrow[dl, dotted] \\
N^{p,q} &
\end{tikzcd}$$

The triangles are preserved after applying $\mathcal{G}$. It follows that the maps $\mathcal{G}(N^{p,q}) \to \mathcal{G}(L^{p,q})$ are all injections, and similarly for the maps $\mathcal{G}(L^{p,q}) \to \mathcal{G}(J^{p,q}) = K^{p,q}$. By construction, we have $J^{p,q} = L^{p,q} \oplus N^{p,q}$, so the maps $J^{p,q} \twoheadrightarrow N^{p+1,q}$ have sections, and thus the maps $K^{p,q} \to \mathcal{G}(N^{p+1,q})$ are surjective. Finally, applying the left exact functor $\mathcal{G}$ to the short exact sequence $0 \to L^{p,q} \to J^{p,q} \to N^{p+1,q} \to 0$, we find that $\ker(K^{p,q} \to \mathcal{G}(N^{p+1,q})) = \mathrm{im}\,(\mathcal{G}(L^{p,q}) \to K^{p,q})$.

In particular, the kernel of the horizontal map $\partial' : K^{q,p} \to K^{q+1,p}$ is $\mathbf{G}(L^{q,p})$, and its image is $\mathbf{G}(N^{q+1,p})$. The second spectral sequence of the double complex $\mathbf{K}$ computes the cohomology of the rows. It follows from (10) that

$$''E_1^{p,q} = \mathcal{G}(L^{q,p})/\mathcal{G}(N^{q,p}) = \mathcal{G}(L^{q,p}/N^{q,p}),$$

where the second equality follows from $\mathcal{G}$ preserving split short exact sequences and $N^{0,p} = 0$. The differential $d_1 : {''E}_1^{p,q} \to {''E}_1^{p+1,q}$ is induced by the vertical map $L^{q,p} \to L^{q,p+1}$. Thus it follows from Remark 5.19 that

$$''E_2^{p,q} = R^p\mathcal{G}(Z^q/B^q) = R^p\mathcal{G}(R^q\mathcal{F}(A)).$$

We know that $''E_\infty^{p,q} = F^pH^{p+q}(\mathrm{Tot}\,\mathbf{K})$, for a suitable filtration of $H^{p+q}(\mathrm{Tot}\,\mathbf{K})$. We determined in (11) that $H^{p+q}(\mathrm{Tot}\,\mathbf{K}) = R^{p+q}(\mathcal{G} \circ \mathcal{F})(A)$. Hence $R^p\mathcal{G}(R^q\mathcal{F}(A)) \Rightarrow R^{p+q}(\mathcal{G} \circ \mathcal{F})(A)$, as claimed. $\square$

5.5. **The Hochschild-Serre spectral sequence.** Having developed the tools we need from the theory of spectral sequences, we can finally specialize them to the study of G-modules.

Let G be a profinite group and $H \trianglelefteq G$ a normal subgroup. We apply Grothendieck's Theorem 5.18 to the following situation. Let $\mathcal{F} : \mathrm{Mod}_G \to \mathrm{Mod}_{G/H}$ be the functor sending a G-module M to the submodule M^H, with the obvious G/H-module structure. Let $\mathcal{G}$: $\mathrm{Mod}_{G/H} \to \mathrm{Ab}$ be the functor $M \mapsto M^{G/H}$. Both of these functors are left exact, and the corresponding right derived functors are $R^p\mathcal{F} = H^p(H, \mathrm{Res}^G_H -)$ and $R^q\mathcal{G} = H^q(G/H, -)$. Both Mod_G and $\mathrm{Mod}_{G/H}$ have enough injectives.

Thus, to verify the hypotheses of Theorem 5.18 it remains to show that, for every injective G-module I, the G/H-module I^H is $\mathcal{G}$-acyclic. We claim that if I^H is, in fact, an injective G/H-module. Indeed, if we are given a diagram

$$\begin{tikzcd} M \arrow[r, hook] \arrow[d] & N \\ I^H & \end{tikzcd}$$

of G/H-modules, we can consider it as a diagram of G-modules via the natural projection $\pi : G \to G/H$:

$$\begin{tikzcd} \pi^\times M \arrow[r, hook] \arrow[d] & \pi^\times N \arrow[dl, dotted] \\ I^H \subset I & \end{tikzcd}$$

By injectivity of I, we can fill in the dotted G-module map $N \to I$. Since H acts trivially on N, the image of this map is contained in I^H.

Definition 5.20. Let G be a profinite group and $H \trianglelefteq G$ a normal subgroup. In view of the preceding discussion, for every G-module M Theorem 5.18 provides us with an explicit spectral sequence

$$H^p(G/H, H^q(H, \mathrm{Res}^G_H M)) \Rightarrow H^{p+q}(G, M).$$

Thus is called the *(Lyndon)-Hochschild-Serre spectral sequence*.

Theorem 5.21. *Let G be a profinite group, let $H \trianglelefteq G$ be a normal subgroup, and let M be a G-module. The following sequence is exact:*

$$0 \to H^1(G/H, M^H) \overset{\inf}{\to} H^1(G, M) \overset{\mathrm{res}}{\to} H^1(H, \mathrm{Res}^G_H M)^{G/H} \to H^2(G/H, M^H) \overset{\inf}{\to} H^2(G, M).$$

Proof. We apply Corollary 5.13 to the Hochschild-Serre spectral sequence. It remains to check that the maps induced by the spectral sequence are indeed inflation and restriction. □

Corollary 5.22. *Let G be a profinite group, let $H \trianglelefteq G$ be a normal subgroup, and let M be a G-module. Suppose that $H^i(H, \mathrm{Res}^G_H M) = 0$ for all $i > 0$. Then $H^i(G/H, M^H) \simeq H^i(G, M)$ for all $i > 0$.*

Proof. If $i = 1$ this is immediate from the inflation-restriction sequence. Otherwise, our hypothesis implies that the Hochschild-Serre spectral sequence for M satisfies condition $(*)_n$ for all $n \geq 1$. □

Exercises

(1) Prove Theorem 5.21 (the inflation-restriction sequence) directly, without using spectral sequences.

(2) Let X be a topological space. Recall that a presheaf $\mathcal{F}$ on X consists of the following data:
- For every open set $U \subset X$, a set $\mathcal{F}(U)$, called the set of sections;
- For every inclusion $V \subset U$ of open sets, a restriction map $\mathrm{res}^{\mathcal{F}}_{U,V} : \mathcal{F}(U) \to \mathcal{F}(V)$ such that $\mathrm{res}^{\mathcal{F}}_{U,U} : \mathcal{F}(U) \to \mathcal{F}(U)$ is the identity and $\mathrm{res}^{\mathcal{F}}_{V,W} \circ \mathrm{res}^{\mathcal{F}}_{U,V} = \mathrm{res}^{\mathcal{F}}_{U,W}$ for $W \subset V \subset U$.

Often $\mathcal{F}(U)$ is the collection of functions with suitable properties defined on U. If $s \in \mathcal{F}(U)$, we write $s_{|V}$ for $\mathrm{res}^{\mathcal{F}}_{U,V}(s)$ by analogy with restriction of functions. A presheaf $\mathcal{F}$ is called a sheaf if it satisfies two additional properties:
- (Locality) If $U = \bigcup_{i\in I} U_i$ is an open cover and $s, t \in \mathcal{F}(U)$ satisfy $s_{|U_i} = t_{|U_i}$ for all $i \in I$, then $s = t$;
- (Gluing) Let $U = \bigcup_{i\in I} U_i$ be an open cover as above. Suppose we are given $s_i \in \mathcal{F}(U_i)$ for all $i \in I$ such that $(s_i)_{|U_i\cap U_j} = (s_j)_{|U_i\cap U_j}$ for all $i, j \in I$. Then there exists a section $s \in \mathcal{F}(U)$ such that $s_{|U_i} = s_i$ for all $i \in I$.

We will assume that the sets of sections $\mathcal{F}(U)$ are abelian groups and that the restriction maps are group homomorphisms. A morphism of presheaves $f : \mathcal{F} \to \mathcal{G}$ is a family of group homomorphism $f(U) : \mathcal{F}(U) \to \mathcal{G}(U)$ for every open $U \subset X$ that are compatible with the restriction maps, i.e. the square

$$\begin{tikzcd}
\mathcal{F}(U) \arrow[r, "f(U)"] \arrow[d, "\mathrm{res}_{U,V}"'] & \mathcal{G}(U) \arrow[d, "\mathrm{res}_{U,V}"] \\
\mathcal{F}(V) \arrow[r, "f(V)"] & \mathcal{G}(V)
\end{tikzcd}$$

commutes for every $V \subset U$. Thus we may speak of the category Sh_X of sheaves on X. Define the kernel of f to be the presheaf $(\ker f)(U) = \ker f(U) \subset \mathcal{F}(U)$; the restriction maps are restrictions of those of $\mathcal{F}$.

(a) Prove that $\ker f$ is indeed a presheaf, and that it is a sheaf if $\mathcal{F}$ and $\mathcal{G}$ are sheaves.

(b) Now we want to define the image of f. This is a bit trickier. First define a presheaf $\mathrm{preim}\, f$ by $(\mathrm{preim}\, f)(U) = \mathrm{im}\, f(U) \subset \mathcal{G}(U)$, where the restriction maps are induced by those of $\mathcal{G}$. Prove that this is indeed a presheaf. Thus, in the category Presh_X of presheaves on X, we may take $\mathrm{preim}\, f$ as the image of f. In particular, we can define exact sequences of objects of Presh_X.

(c) Show that $\mathrm{preim}\, f$ need not be a sheaf even if $\mathcal{F}$ and $\mathcal{G}$ are both sheaves. Observe that the locality property always holds if $\mathcal{G}$ is a sheaf, but that gluing may fail. Hint: let $X = S^1$ be the unit circle in $\mathbb{C}$, with the usual topology, and for every $U \subset X$ let $\mathcal{C}(U)$ be the (additive) abelian group of continuous functions $U \to \mathbb{C}$. Show that this is indeed a sheaf. Let $\mathcal{C}^\times(U)$ be the (multiplicative) abelian group of non-vanishing continuous functions on U. Show that $\mathcal{C}^\times$ is also a sheaf and that

$$\begin{aligned} \exp : \mathcal{C} &\to \mathcal{C}^\times \\ \varphi(x) &\mapsto e^{\varphi(x)} \end{aligned}$$

is a morphism. Prove that preim exp is not a sheaf.

(d) Let $f : \mathcal{F} \to \mathcal{G}$ be a morphism of sheaves. We are led to define $\operatorname{im} f$ as the smallest subsheaf of $\mathcal{G}$ that contains $\operatorname{preim} f$. Prove that this is indeed a well-defined object.

(e) Let $f : \mathcal{F} \to \mathcal{G}$ be a morphism of sheaves. Show that it is an epimorphism (i.e. that $\operatorname{im} f = \mathcal{G}$) if and only if for every open $U \subset X$ and every section $s \in \mathcal{G}(U)$ there exists an open cover $U = \bigcup_{i\in I} U_i$ such that $s_{|U_i} \in \operatorname{im} f(U_i)$ for all $i \in I$.

(3) Let X be a topological space. A sheaf $\mathcal{F}$ on X is said to be a sheaf of rings if the abelian group $\mathcal{F}(U)$, for every open $U \subset X$, is endowed with a multiplication that gives it the structure of a unital ring, and if all the restriction maps $\operatorname{res}_{U,V}$ are ring homomorphisms. A *ringed space* $(X, \mathcal{O}_X)$ consists of a topological space X and a sheaf of rings $\mathcal{O}_X$ on X.

Given a ringed space $(X, \mathcal{O}_X)$, an $\mathcal{O}_X$-module is a sheaf $\mathcal{M}$ on X such that $\mathcal{M}(U)$ is endowed with an $\mathcal{O}_X(U)$-module structure for every open $U \subset X$, and the restriction maps of $\mathcal{M}$ are compatible with those of $\mathcal{O}_X$. More precisely, if $V \subset U$, then for every $a \in \mathcal{O}_X(U)$ and $s \in \mathcal{M}(U)$ we require that

$$\operatorname{res}^{\mathcal{M}}_{U,V}(as) = \operatorname{res}^{\mathcal{O}_X}_{U,V}(a) \cdot \operatorname{res}^{\mathcal{M}}_{U,V}(s).$$

(a) Let R be a unital ring. Show that there exists a *constant sheaf* of rings $\underline{R}$ on X such that $\underline{R}(U) = R$ for all open $U \subset X$, and all the restriction maps are identity maps.

(b) Show that a $\underline{\mathbb{Z}}$-module on X is the same thing as a sheaf.

(c) Generalize all the notions developed in the previous exercise to $\mathcal{O}_X$-modules on a ringed space. In particular, define the category $\mathcal{O}_X - \mathrm{Mod}$ of $\mathcal{O}_X$-modules and the notion of a short exact sequence $0 \to \mathcal{M} \to \mathcal{N} \to \mathcal{P} \to 0$ of $\mathcal{O}_X$-modules.

(d) Show that $\mathcal{M} \mapsto \mathcal{M}(X)$ gives a left exact functor from $\mathcal{O}_X - \mathrm{Mod}$ to the category $\mathrm{Mod}_{\mathcal{O}_X(X)}$ of left modules over the ring $\mathcal{O}_X(X)$. We denote this functor by Γ. Since the category $\mathcal{O}_X - \mathrm{Mod}$ has enough injectives[2] we may define the right derived functors $H^i(X, -) : \mathcal{O}_X - \mathrm{Mod} \to \mathrm{Mod}_{\mathcal{O}_X(X)}$ of Γ.

(4) The definition of sheaf cohomology in the previous exercise, as right derived functors of the global sections functor, doesn't give us a way to compute it. In this exercise we will provide a way and will finally obtain another application of spectral sequences.

6. The Brauer group

If L/K is a finite Galois extension with Galois group $G = \mathrm{Gal}(L/K)$, then $H^0(G, L^\times) = K^\times$, whereas $H^1(G, L^\times) = 0$ by Hilbert 90. Our next aim is to understand $H^2(G, L^\times)$, which turns out to be most conveniently done in terms of the theory of central simple algebras.

[2] Insert exercise proving this

6.1. **Non-abelian cohomology.** Let G be a profinite group, and let M be a group, not necessarily abelian. We write the operation of M multiplicatively. For the purposes of this section, M will be called a *(discrete) G-module* if it is equipped with a G-action such that $g(m_1m_2) = (gm_1)(gm_2)$ for all $g \in G$ and $m_1, m_2 \in M$, and if $\mathrm{stab}_G(m)$ is an open subgroup of G for every $m \in M$.

Example 6.1. Let L/K be a finite Galois extension. For every $n \geq 1$, the Galois group $G = \mathrm{Gal}(L/K)$ acts on $\mathrm{GL}_n(L)$ by acting on the matrix elements. Since matrix multiplication is given by polynomials in the matrix elements, it is respected by the action of G. Thus, the group $M = \mathrm{GL}_n(L)$, which is of course non-abelian if $n > 1$, is naturally a G-module.

We would like to have a cohomology theory in this situation. We still have a G-invariants functor $M \mapsto M^G$ from the category of G-modules to the category of groups. Unfortunately, the category of groups is not abelian, so most of the standard constructions of homological algebra don't work. For instance, the image of an injective group homomorphism $\iota : H \to G$ need not be a *normal* subgroup of G, and thus it is not possible in general to extend the exact sequence $0 \to H \to G$. However, for small i we can try to mimic the definition of $H^i(G, M)$ "by hand" and see what we get.

Definition 6.2. Let G be a profinite group and M a G-module.

(1) Set $H^0(G, M) = M^G$.
(2) Let $Z^1(G, M)$ be the set of functions $\psi : G \to M$ satisfying $\psi(g_1g_2) = \psi(g_1) \cdot (g_1(\psi(g_2)))$. Note that there is no natural group structure on this set, but it does have a distinguished element, namely the trivial function given by $\psi(g) = e_M$ for all $g \in G$.
(3) Given $\psi, \eta \in Z^1(G, M)$, say that $\psi \sim \eta$ if there exists an element $m \in M$ such that $\psi(g) = m^{-1}\eta(g) \cdot gm$ for all $g \in G$. This is clearly an equivalence relation, and we define $H^1(G, M)$ to be the set of equivalence classes $Z^1(G, M)/\sim$. Again, $H^1(G, M)$ is a pointed set, namely a set with a distinguished element.

Observe that if M is abelian, then these definitions coincide (as pointed sets) with the cohomology groups we have defined already. Moreover, if PtSet denotes the category of pointed sets, where the morphisms are set maps $A \to B$ sending the distinguished element of A to that of B, then $H^1(G, -)$ is a naturally a functor from the category of G-modules to PtSet; the proof is left to the reader.

If $f : A \to B$ is a morphism of pointed sets, then its image is a pointed subset of B. If we set the kernel of f to be the preimage of the distinguished element of B, then clearly $\ker f$ is a pointed subset of A. Thus, we have a notion of exact sequences of pointed sets.

Proposition 6.3. *Let G be a profinite group, and let $0 \to M \xrightarrow{\varepsilon} N \xrightarrow{\pi} P \to 0$ be an exact sequence of G-modules. Then there is an exact sequence of pointed sets*

$$0 \to H^0(G, M) \to H^0(G, N) \to H^0(G, P) \xrightarrow{\delta^0} H^1(G, M) \to H^1(G, N) \to H^1(G, P),$$

where, for $p \in P^G$, we set $\delta^0(p)$ to be the equivalence class of the 1-cocycle $m \mapsto \tilde{p}^{-1} \cdot g\tilde{p}$. Here $\tilde{p}$ is an arbitrary lift of p to N.

Furthermore, suppose that $\varepsilon(M)$ lies in the center of N. In this case M is abelian, so $H^2(G, M)$ is defined. Then the exact sequence above can be extended by the map $H^1(G, P) \xrightarrow{\delta^1} H^2(G, M)$, where, for $\psi \in Z^1(G, P)$, we set

$$\delta^1([\psi]) = [(g_1, g_2) \mapsto \widetilde{\psi(g_1)} \cdot g_1\widetilde{\psi(g_2)} \cdot \psi\widetilde{(g_1g_2)}^{-1}].$$

Proof. Explicit computation. □

Proposition 6.4. *Let L/K be a finite Galois extension with Galois group $G = \mathrm{Gal}(L/K)$, and let V be a finite-dimensional L-vector space endowed with a semi-linear G-action: for every $\sigma \in G$, $a \in L$, and $v \in V$ we have $\sigma(av) = \sigma(a) \cdot \sigma(v)$. Then $\dim_K V^G = \dim_L V$ and the natural map*

$$\begin{array}{ccc} V^G \otimes_K L & \xrightarrow{\gamma} & V \\ v \otimes a & \mapsto & av \end{array}$$

is an isomorphism of L-vector spaces.

Proof. First we will show that γ is surjective. Suppose not, and let $u \in V$ be an element that is not contained in $\operatorname{im}\gamma$. Then there exists a subspace $W \subset V$ such that $\operatorname{im}\gamma \subset W$ and $V \simeq W \oplus Lu$, and projection onto the second component provides a non-zero linear functional $\lambda : V \to L$ such that $\lambda(\operatorname{im}\gamma) = 0$. For any $a \in L$ and $v \in V$, the element $w = \sum_{\sigma \in G} \sigma(av)$ lies in V^G. Thus $w = \gamma(w \otimes 1)$, and hence

$$0 = \lambda(w) = \sum_{\sigma \in G} \lambda(\sigma(av)) = \sigma(a) \sum_{\sigma \in G} \lambda(\sigma(v)).$$

In particular, for all $a \in L$ we have $\sum_{\sigma \in G} \lambda(\sigma(v)) \cdot \sigma(a) = 0$. By Dedekind's Lemma this means that $\lambda(\sigma(v)) = 0$ for all $\sigma \in G$ and all $v \in V$, contradicting $\lambda \neq 0$.

Hence γ is surjective. In view of this, to establish the remaining parts of our claim it suffices to show that $\dim_K V^G \leq \dim_L V$. Note that V^G need not be an L-vector subspace of V, but it is a K-vector subspace. Let $\{e_1, \ldots, e_r\}$ be a K-basis of V^G. We claim that these vectors are linearly independent over L. Set

$$B = \left\{ (b_1, \ldots, b_r) \in L^r : \sum_{i=1}^{r} b_i e_i = 0 \right\}.$$

Clearly B is an L-linear subspace of L^r. We know that $B \cap K^r = \{0\}$, since the e_i are linearly independent over K. Also, since the e_i are G-invariants, we observe that if $b = (b_1, \ldots, b_r) \in B$ and $\sigma \in G$, then $(\sigma(b_1), \ldots, \sigma(b_r)) \in B$. Thus, for all $a \in L$ and all $b \in B$, we have

$$(\mathrm{Tr}_{L/K}(ab_1), \ldots, \mathrm{Tr}_{L/K}(ab_r)) = \sum_{\sigma \in G} (\sigma(ab_1), \ldots, \sigma(ab_r)) \in B \cap K^r = \{0\}. \tag{12}$$

We claim that the trace map $\mathrm{Tr}_{L/K}$ is not identically zero. (In fact, it is true that a finite extension L/K is separable if and only if the trace map $\mathrm{Tr}_{L/K} : L \to K$ is not identically zero). Since we are assuming that L/K is Galois, we can give a quick proof. Indeed, if $\operatorname{char} K = 0$, then $\mathrm{Tr}_{L/K}(1) = [L : K] \neq 0$. In general, since L/K is finite and separable, it is primitive, so that $L = K(\alpha)$ for some $\alpha \in L$. We may assume that $\alpha \neq 0$, since otherwise $L = K$ and the claim is trivial. Let $\alpha = \alpha_1, \ldots, \alpha_n$ be the roots of the minimal polynomial $f_\alpha(x) \in K[x]$ of α; since L/K is Galois, these are all distinct and contained in K. Moreover, $\alpha_1^m, \ldots, \alpha_n^m$ are the conjugates of α^m for any $m \in \mathbb{Z}$. Let $f'_\alpha \in K[x]$ be the formal derivative of f_α; observe that $f'_\alpha(\alpha) \neq 0$. Thus, in the power series ring $L[[x]]$, we have

$$\frac{f'_\alpha}{f_\alpha} = \sum_{i=1}^{n} \frac{1}{x - \alpha_i} = -\sum_{m=0}^{\infty} \sum_{i=1}^{n} \alpha_i^{-(m+1)} x_n = -\sum_{m=0}^{\infty} \mathrm{Tr}_{L/K}(\alpha^{-(m+1)}) x^m,$$

where the first equality is immediate from the Leibniz rule. Since the left-hand side is non-zero, the right-hand side is also non-zero. Thus $\mathrm{Tr}_{L/K}(\alpha^{-m+1}) \neq 0$ for some $m \geq 0$.

Now, by (12) we have that $\mathrm{Tr}_{L/K}(ab_i) = 0$ for all $a \in L$ and all $1 \leq i \leq r$. Thus we must have $b_i = 0$ for all $1 \leq i \leq r$. Hence $B = \{0\}$, and so $\{e_1, \dots, e_r\}$ is linearly independent over L. We conclude that $\dim_K V^G \leq \dim_L V$, and the claim follows. □

We now deduce a corollary, which may be viewed as a non-abelian version of Hilbert 90.

Corollary 6.5. *Let L/K be a finite Galois extension with Galois group $G = \mathrm{Gal}(L/K)$. Then $H^1(G, \mathrm{GL}_n(L))$ has one element for all $n \geq 1$.*

Proof. Let $\psi \in Z^1(G, \mathrm{GL}_n(L))$. Observe that $V = L^n$ carries a semilinear G-action by $v_\sigma = (\sigma(a_1), \dots, \sigma(a_n))$, where $\sigma \in G$ and $v = (a_1, \dots, a_n) \in V$. We consider a modified G-action, setting $\sigma(v) = \psi(\sigma) \cdot v_\sigma$. This is clearly semilinear, and it is an action because

$$(\tau\sigma)(v) = \psi(\tau\sigma) \cdot v_{\tau\sigma} = \psi(\tau) \cdot \tau(\psi(\sigma)) \cdot (v_\sigma)_\tau = \psi(\tau) \cdot (\psi(\sigma) \cdot v_\sigma)_\tau = \tau(\sigma(v)),$$

where the second equality holds because ψ is a 1-cocycle.

By Proposition 6.4, V has an L-basis $\{e_1, \dots, e_n\}$ consisting of G-invariants. Let $A \in \mathrm{GL}_n(L)$ be the matrix whose columns are $e_1, \dots, e_n$. Since $\psi(\sigma)(e_i)_\sigma = e_i$ for all $\sigma \in G$ and $1 \leq i \leq n$, and since the columns of the matrix $\sigma(A)$ are the $(e_i)_\sigma$, we find that $A^{-1}\psi(\sigma)\sigma(A) = A^{-1}A = I_n$ for all $\sigma \in G$, where I_n is the identity matrix. Thus ψ is equivalent to the constant function $\sigma \mapsto I_n$ and lies in the distinguished class of $H^1(G, \mathrm{GL}_n(L))$. □

For every $a \in L^\times$, let $\varepsilon(a) \in \mathrm{GL}_n(L)$ denote the scalar matrix aI_n. Then the image of $\varepsilon : L^\times \hookrightarrow \mathrm{GL}_n(L)$ is the center of $\mathrm{GL}_n(L)$, and the cokernel is, by definition, the projective linear group $\mathrm{PGL}_n(L)$.

Lemma 6.6. *Let L/K be a finite Galois extension with Galois group G, let $n = [L : K]$, and consider the short exact sequence $0 \to L^\times \xrightarrow{\varepsilon} \mathrm{GL}_n(L) \to \mathrm{PGL}_n(L) \to 0$ of G-modules. The map*

$$\delta^1 : H^1(G, \mathrm{PGL}_n(L)) \to H^2(G, L^\times)$$

defined in Proposition 6.3 is a surjection of pointed sets.

Proof. Let $V = L[G]$ be an n-dimensional L-vector space spanned by $\{e_\sigma : \sigma \in G\}$. Fixing an enumeration of the elements of G, identify $\mathrm{Aut}_L V$ with $\mathrm{GL}_n(L)$. If $A \in \mathrm{GL}_n(L)$, denote by $\overline{A}$ its image in $\mathrm{PGL}_n(L)$.

Let $\psi \in Z^2(G, L^\times)$. For each $\sigma \in G$, define $\varphi(\sigma) \in \mathrm{Aut}_L V \simeq \mathrm{GL}_n(L)$ by $\varphi(\sigma)e_\tau = \psi(\sigma, \tau)e_{\sigma\tau}$ for all $\tau \in G$. For every $\eta \in G$ we have by definition that $\varphi(\sigma\tau)e_\eta = \psi(\sigma\tau, \eta)e_{\sigma\tau\eta}$, whereas

$$(\varphi(\sigma) \cdot \sigma\varphi(\tau))e_\eta = \varphi(\sigma)(\sigma(\psi(\tau, \eta))e_{\tau\eta}) = \sigma(\psi(\tau, \eta))\psi(\sigma, \tau\eta)e_{\sigma\tau\eta} = \psi(\sigma, \tau)\psi(\sigma\tau, \eta)e_{\sigma\tau\eta},$$

where the last equality holds because it follows from (5) that

$$\sigma(\psi(\tau, \eta)) \cdot \psi(\sigma\tau, \eta)^{-1}\psi(\sigma, \tau\eta)\psi(\sigma, \tau)^{-1} = 0 \tag{13}$$

for all $\sigma, \tau, \eta \in G$. Hence, although the map $\varphi : G \to \mathrm{GL}_n(L)$ need not be a 1-cocycle, the map $\overline{\varphi} : G \to \mathrm{PGL}_n(L)$ given by $\overline{\varphi}(\sigma) = \overline{\varphi(\sigma)}$ is a 1-cocycle.

Finally, it follows from (13) that $\varphi(\sigma) \cdot \sigma\varphi(\tau) \cdot \varphi(\sigma\tau)^{-1}e_\eta = \psi(\sigma, \sigma^{-1}\eta) \cdot \sigma\psi(\tau, \tau^{-1}\sigma^{-1}\eta) \cdot \psi(\sigma\tau, \tau^{-1}\sigma^{-1}\eta)e_\eta = \psi(\sigma, \tau)e_\eta$. Hence

$$\delta^1([\overline{\varphi}]) = [(\sigma, \tau) \mapsto \varphi(\sigma) \cdot \sigma\varphi(\tau) \cdot \varphi(\sigma\tau)^{-1}] = [\psi]$$

and δ^1 is surjective. □

Remark 6.7. It follows from Corollary 6.5 and the exact sequence of Proposition 6.3 that δ^1 has trivial kernel for all $n \geq 1$. However, since δ^1 is only a morphism of pointed sets and not of groups, this does not imply that it is injective. Nevertheless, we will manage in Theorem 6.36 below to get some substantial mileage out of this observation.

We state the following result, which follows immediately from the Skolem-Noether Theorem. The precise statement and proof of Skolem-Noether (Theorem 6.28 below) will have to wait until we have developed the theory of central simple algebras.

Lemma 6.8 (Skolem-Noether)**.** *Let L be a field, let $m \geq 1$, and let $f, g : M_m(L) \to M_m(L)$ be two L-algebra homomorphisms. Then there exists an invertible matrix $b \in M_m(L)$ such that $g(a) = b \cdot f(a) \cdot b^{-1}$ for all $a \in A$.*

Using the previous lemma, we may describe the pointed sets $H^1(G, \mathrm{PGL}_m(L))$ in terms of central simple algebras; see Definition 6.10 below.

Lemma 6.9. *Let L/K be a finite Galois extension with Galois group $G = \mathrm{Gal}(L/K)$, and for every $m \geq 1$ let $A(L/K, m)$ denote the set of K-isomorphism classes of central simple K-algebras A such that $A \otimes_K L \simeq M_m(L)$. Then $A(L/K, m) \simeq H^1(G, \mathrm{PGL}_m(L))$ as pointed sets, where the distinguished element of $A(L/K, m)$ is (the isomorphism class of) $M_m(K)$.*

Proof. Let A be a K-algebra such that $A \otimes_K L \simeq M_m(L)$. Fix such an isomorphism of L-algebras, and let $\sigma \in G$. On the one hand, σ acts on matrices $c \in M_m(L)$ by acting on each matrix element; we denote the image by $\sigma(c)$. On the other hand, σ acts on $A \otimes_K L$ by sending $c = \sum_{j=1}^{r} a_j \otimes b_j$ to $(1 \otimes \sigma)c = \sum_{j=1}^{r} a_j \otimes \sigma(b_j)$, where $a_j \in A$ and $b_j \in L$. We view this as an action on $M_m(L)$ by transport of structure via the isomorphism we fixed.

The maps $c \mapsto \sigma(c)$ and $c \mapsto (1 \otimes \sigma)c$ are semilinear, not linear, over L. However, the map $f : M_m(L) \to M_m(L)$ given by $f(c) = \sigma^{-1}((1 \otimes \sigma)(c))$ is L-linear. We can apply Lemma 6.8, taking g to be the identity map, to obtain an invertible matrix $b \in \mathrm{GL}_m(L)$ such that $c = b \cdot \sigma^{-1}(1 \otimes \sigma)c \cdot b^{-1}$ for all $c \in M_m(L)$. Hence $(1 \otimes \sigma)c = \Psi(\sigma)\sigma(c)\Psi(\sigma)^{-1}$, where $\Psi(\sigma) = \sigma(b^{-1})$. Observe that such b is well-defined up to left multiplication by an element of the center of $M_m(L)$. Thus $\Psi(\sigma)$ depends on the choice of b, but $\overline{\Psi(\sigma)} \in \mathrm{PGL}_m(L)$ depends only on σ.

We claim that $(\sigma \mapsto \overline{\Psi(\sigma)}) \in Z^1(G, \mathrm{PGL}_m(L))$. Indeed,

$$(1 \otimes \sigma\tau)c = (1 \otimes \sigma)((1 \otimes \tau)c) = \Psi(\sigma)\sigma(\Psi(\tau)\tau(c)\Psi(\tau)^{-1})\Psi(\sigma)^{-1}.$$

However, the previous construction depended on the choice of an isomorphism $\iota : A \otimes_K L \xrightarrow{\sim} M_m(L)$. Any automorphism of $M_m(L)$ is inner (i.e. conjugation by some $\beta \in M_m(L)$) by Lemma 6.8, and any isomorphism $A \otimes_K L \to M_m(L)$ is the composition of ι with such an automorphism of $M_m(L)$. This amounts to replacing the map $1 \otimes \sigma$ defined as above using ι with the map $c \mapsto \beta \cdot (1 \otimes \sigma)(\beta^{-1} c \beta) \cdot \beta^{-1}$. We know that $\beta^{-1} c \beta = b \cdot \sigma^{-1}(1 \otimes \sigma)(\beta^{-1} c \beta) \cdot b^{-1}$ by the defining property of b. Hence

$$c = \beta b \sigma^{-1}(\beta^{-1}) \cdot \sigma^{-1}(\beta \cdot (1 \otimes \sigma)(\beta^{-1} c \beta) \cdot \beta^{-1}) \cdot \sigma^{-1}(\beta) b^{-1} \beta^{-1}.$$

Thus b is replaced with $\beta b \sigma^{-1}(\beta^{-1})$, up to scalar multiple, and the cocycle $(\sigma \mapsto \overline{\Psi(\sigma)})$ is replaced with the cocycle $(\sigma \mapsto \overline{\beta} \cdot \overline{\Psi(\sigma)} \cdot \sigma(\overline{\beta}^{-1}))$, which clearly is equivalent to it under the equivalence relation of Definition 6.2. In all, we have obtained a well-defined map $A(L/K, m) \to H^1(G, \mathrm{PGL}_m(L))$.

Conversely, given a cocycle $(\sigma \mapsto \overline{\Psi(\sigma)})$, we can define an action of G on $M_m(L)$ by $(1 \otimes \sigma)c = \Psi(\sigma)\sigma(c)\Psi(\sigma)^{-1}$. The set of elements $c \in M_m(L)$ that are invariant under every

$1 \otimes \sigma$ is a central simple K-algebra A, which satisfies $A \otimes_K L \simeq M_m(L)$ by Proposition 6.4. One checks that this gives a well-defined map $H^1(G, \mathrm{PGL}_m(L)) \to A(L/K, m)$: indeed, any 1-cocycle that is cohomologous to $(\sigma \mapsto \overline{\Psi(\sigma)})$ will give rise to a K-algebra of G-invariants that is conjugate to, and hence isomorphic to, A. It is obvious from the construction that this map is inverse to the map $A(L/K, m) \to H^1(G, \mathrm{PGL}_m(L))$ that was defined above.

Finally, if $A = M_m(K)$, then the two actions $c \mapsto \sigma(c)$ and $c \mapsto (1 \otimes \sigma)(c)$ coincide, and thus we may take $\Psi(\sigma)$ to be the identity matrix for every $\sigma \in G$. Thus our maps preserve distinguished elements. □

6.2. **Central simple algebras.** Before we can proceed, we'll need a crash course in the theory of central simple algebras. We'll now do a bit of beautiful pure ring theory, although for our purposes in this course it is a tool for computing cohomology. As Rowen writes about central simple algebras (*Ring Theory*, Volume II, p.187): "There is some question among experts as to whether this theory belongs more properly to ring theory, field theory, cohomology theory, or algebraic K-theory."

Let F be a field.

Definition 6.10.
(1) An *F-algebra* is a ring A, not necessarily commutative, equipped with a ring homomorphism $F \to Z(A)$. Here $Z(A)$ denotes the center of A. This homomorphism is necessarily injective, so we view F as a subring of A.
(2) An F-algebra A is called *simple* if has no non-zero proper two-sided ideals. It is called *central simple* if $F \simeq Z(A)$.
(3) An F-algebra A is called a *division algebra* if $A \setminus \{0\}$ is a group under multiplication.
(4) Given an F-algebra A, we let A^{op} be the algebra with the same underlying abelian group as A, but with a multiplication operation $*$ given by $a * b = ba$. Observe that A^{op} is simple if and only if A is.

Lemma 6.11 (Schur)**.** *Let A be an F-algebra, and let M and N be simple A-modules (i.e. they have no non-trivial A-submodules). If $f \in \mathrm{Hom}_A(M, N)$, then either $f = 0$ or f is an isomorphism. In particular, if M is a simple A-module, then $\mathrm{End}_A(M)$ is a division algebra.*

Proof. If $f \in \mathrm{Hom}_A(M, N)$, then $\ker f$ is a submodule of M and $\mathrm{im}\, f$ is a submodule of N. This implies the first claim. Taking $M = N$, we find that any non-zero element of $\mathrm{End}_A(M)$ is an isomorphism, so it has a multiplicative inverse. □

Proposition 6.12. *Let D be a division ring, and let M be a left D-module. Then M is free. Moreover, any D-linearly independent subset of M may be extended to a basis.*

Proof. This claim is proved in every linear algebra course in the case where D is a commutative division ring, i.e. a field. It is usually not pointed out that the commutativity isn't necessary. By a standard Zorn's Lemma argument, there exists a D-linearly independent subset $S \subset M$ that is maximal under inclusion. We claim that S spans M. Indeed, suppose that $v \notin \mathrm{span}_D(S)$. Then $S \cup \{v\}$ is linearly dependent by the maximality of S, so there exists a non-trivial linear relation $dv + \sum_{s \in S} d_s s = 0$, where only finitely many of the d_s are non-zero. Since $d \neq 0$ by the linear independence of S, we find that $-v = \sum_{s \in S} d_s s$, contradicting $v \notin \mathrm{span}_D(S)$. Thus M is free as a D-module. In fact, we have shown that any maximal linearly independent subset of M is a basis; this implies the second claim. □

In the sequel we shall only have cause to work with algebras that have finite dimension over F. In order to avoid writing this condition over and over, we shall make it a running hypothesis: *from now on, all F-algebras are assumed to be finite-dimensional.*

Theorem 6.13 (Jacobson Density Theorem)**.** *Let A be a finite-dimensional F-algebra, and let M be a simple A-module such that $\dim_F(M) < \infty$. Let $D = \mathrm{End}_A(M)$. Suppose that $m_1, \ldots, m_r \in M$ are linearly independent over D, and let $n_1, \ldots, n_r \in M$. Then there exists an element $a \in A$ such that $am_i = n_i$ for all $1 \leq i \leq r$.*

Proof. By Proposition 6.12 the set $\{m_1, \ldots, m_r\}$ extends to a D-basis of M. Thus there exists a D-submodule $N \subset M$ such that $M = Dm_1 \oplus Dm_2 \oplus \cdots \oplus Dm_r \oplus N$. In particular, the map $\varphi \in \mathrm{End}_D(M)$ sending $(d_1m_1, \ldots, d_rm_r, n)$ to $(d_1n_1, \ldots, d_rn_r, 0)$ with respect to this decomposition satisfies $\varphi(m_i) = n_i$ for all $1 \leq i \leq r$.

Consider the element $m = (m_1, \ldots, m_r) \in M^r$. Since M^r is a semisimple A-module, being a direct sum of simple modules, any submodule is a direct summand. Thus there exists an A-submodule $P \subset M^r$ such that $M^r = Am \oplus P$. Let $\pi \in \mathrm{End}_A(M^r) = M_r(D)$ be projection onto the component Am.

Since φ is D-linear, the map $\varphi^r : M^r \to M^r$ given by $\varphi^r(x_1, \ldots, x_r) = (\varphi(x_1), \ldots, \varphi(x_r))$ is a map of $M_r(D)$-modules, and thus

$$(n_1, \ldots, n_r) = \varphi^r(m_1, \ldots, m_r) = \varphi^r\pi(m_1, \ldots, m_r) = \pi\varphi^r(m_1, \ldots, m_r) = \pi(n_1, \ldots, n_r).$$

Hence $(n_1, \ldots, n_r) \in Am$, which is exactly our claim. $\square$

Corollary 6.14. *Suppose that A is a central simple F-algebra. Then $A \otimes_F A^{\mathrm{op}} \simeq M_n(F)$, where $n = \dim_F(A)$.*

Proof. We wish to prove that $A \otimes_F A^{\mathrm{op}} \simeq \mathrm{End}_F(A)$. Let $\{a_1, \ldots, a_n\}$ be an F-basis of A. Since $A \otimes_F A^{\mathrm{op}}$ acts on A by $(x \otimes y)a = xay$, for $a, x \in A$ and $y \in A^{\mathrm{op}}$, we get an F-linear map $f : A \otimes_F A^{\mathrm{op}} \to \mathrm{End}_F(A)$. Since the two algebras are both of dimension n^2 over F, it suffices to show that f is surjective.

Observe that $\mathrm{End}_A(A) \simeq A^{\mathrm{op}}$ via the map $\theta \mapsto \theta(1)$. Thus any element $\psi \in \mathrm{End}_{A\otimes_F A^{\mathrm{op}}}(A)$ is of the form $\psi(a) = ab$ for some $b \in A$. In order for this to be compatible with the A^{op}-action as well, by analogous considerations we must have $\psi(a) = ca$ for all $a \in A$ and some fixed $c \in A$. It follows that $b = c \in Z(A) \simeq F$. Thus $D = \mathrm{End}_{A\otimes_F A^{\mathrm{op}}}(A) \simeq F$, so that $a_1, \ldots, a_n$ are linearly independent over D.

Note that A is a simple $A \otimes_F A^{\mathrm{op}}$-module, since any submodule would be a two-sided ideal. Let $\varphi \in \mathrm{End}_F(A)$. By the Jacobson density theorem, there exists an element $c \in A \otimes_F A^{\mathrm{op}}$ such that $ca_i = \varphi(a_i)$ for all $1 \leq i \leq n$. Then $f(c) = \varphi$, so f is indeed surjective. $\square$

Corollary 6.15. *Let A be a central simple F-algebra and B any simple F-algebra. Then $A \otimes_F B$ is a simple F-algebra.*

Proof. As in the proof of Corollary 6.14, let $\{a_1, \ldots, a_n\}$ be an F-basis of A. By that corollary, for each $1 \leq i \leq n$ there exists an element $c_i \in A \otimes_F A^{\mathrm{op}}$ such that

$$c_i(a_j) = \delta_{ij} = \begin{cases} 1 & : i = j \\ 0 & : i \neq j. \end{cases}$$

Suppose that $I \subset A \otimes_F B$ is a two-sided ideal and thus is preserved by left and right multiplication by elements of $A \otimes_F F$. Let $\sum_{j=1}^n a_j \otimes b_j \in I$. Then $\sum_{j=1}^n c_i(a_j) \otimes b_j = 1 \otimes b_i \in I$, for all $1 \leq i \leq n$, since I is a two-sided ideal. However, $1 \otimes b_i \in F \otimes_F B$.

Since $J = I \cap (F \otimes_F B)$ is a two-sided ideal of $F \otimes_F B \simeq B$, we have either $J = 0$ or $J = B$. In the first case, $b_i = 0$ for all i, and hence $I = 0$. In the second case, $1 \otimes b \in I$ for every $b \in B$, so clearly $I = A \otimes_F B$. $\square$

Corollary 6.16. *If A and B are central simple F-algebras, then so is $A \otimes_F B$.*

Proof. By the previous corollary, it suffices to verify that $Z(A \otimes_F B) \simeq F$. As before, let $\{a_1, \ldots, a_n\}$ be an F-basis of A, and suppose that $\sum_{j=1}^n a_j \otimes b_j \in Z(A \otimes_F B)$. Then $1 \otimes b$ commutes with this element for all $b \in B$, so that $\sum_{j=1}^n a_j \otimes (b_j b - b b_j) = 0$ for every $b \in B$. Since the a_j are F-linearly independent, it follows that $b_j b - b b_j = 0$ for all $b \in B$, hence that $b_j \in Z(B) = F$. Hence $c = \alpha \otimes 1$ for some $\alpha \in A$. Since c commutes with $a \otimes 1$ for every $a \in A$, we find that $\alpha \in Z(A) = F$, hence $c \in F$. □

Let R be a unital ring. Recall that an element $e \in R$ such that $e^2 = e$ is called an idempotent. If $e \in R$ is an idempotent, then $eRe = \{eae : a \in R\}$ is easily seen to be a ring. A left (respectively, right) ideal $I \subset R$ is called minimal if $I \neq (0)$ and there does not exist any left (respectively, right) ideal $J \subset R$ such that J is strictly contained in I; in other words, I is a minimal element of the set of non-zero left (respectively, right) ideals of R, partially ordered by inclusion.

Lemma 6.17 (Brauer)**.** *Let R be a ring, and let $I \subset R$ be a minimal left ideal such that $I^2 \neq 0$. Then there exists an idempotent $e \in I$ such that $I = Re$ and eRe is a division ring.*

Proof. Since $I^2 \neq 0$, there exist $x, y \in I$ such that $yx \neq 0$. Let $Ix = \{ax : a \in I\}$. Clearly Ix is a left ideal of R, and $(0) \neq Ix \subseteq Rx \subseteq I$, where the last inclusion follows from $x \in I$. By minimality of I we conclude $Ix = I$. In particular, there exists an element $e \in I$ such that $ex = x$. Then $e^2 x = e(ex) = ex = x$, so $(e^2 - e)x = 0$. Therefore $e^2 - e \in I \cap \mathrm{Ann}_R(x)$, where $\mathrm{Ann}_R(x) = \{a \in R : ax = 0\}$ is the annihilator obtained by viewing R as a left module over itself.

Now $I \cap \mathrm{Ann}_R(x)$ is a left ideal of R contained in I. Moreover, the containment is strict, since $ex = x$ and hence $e \notin \mathrm{Ann}_R(x)$. Thus $I \cap \mathrm{Ann}_R(x) = (0)$ by minimality of I. Hence $e^2 - e = 0$ and e is an idempotent. Furthermore, $(0) \neq Re \subseteq I$, so $I = Re$, again by minimality of I.

It remains to show that eRe is a division ring. Observe that the multiplicative identity of eRe is e. Let $0 \neq a \in eRe$. Then $a = ebe$ for some $b \in R$. Since $eRe = eI \subset I$, we have $(0) \neq Ra \subseteq I$ and hence $Ra = I$ by minimality of I. Since $e \in I$, there exists $r \in R$ such that $ra = e$. Now r need not be contained in eRe, but ere certainly is. Moreover,

$$(ere)a = (ere)(ebe) = ere^2be = erebe = era = e^2 = e.$$

It remains to show that $a(ere) = e$; then $ere \in eRe$ is a two-sided inverse of a, and we will conclude that eRe is a division ring. But we have just shown that every non-zero element of eRe has a left inverse. Since $ere \neq 0$, there exists $x \in eRe$ such that $x(ere) = e$. Then $a = ea = x(ere)a = xe = x$. Hence $a = x$ and $a(ere) = e$ as desired. □

Now we are ready to prove Wedderburn's structure theorem; see the exercises for a different proof using the Jacobson density theorem. Recall that a simple ring is one with no non-zero proper two-sided ideals.

Theorem 6.18. *Let R be a simple ring. Suppose that R has a minimal left ideal. Then there exists a division ring D and a natural number $n \in \mathbb{N}$ such that $R \simeq M_n(D)$.*

Proof. Let $I \subset R$ be a minimal left ideal. For any two subsets $S, T \subset R$, we write ST for the collection of finite sums of products st, with $s \in S$ and $t \in T$. This "multiplication" of sets is clearly associative. Moreover, $RI = I$ since I is a left ideal. On the other hand, IR is easily seen to be a two-sided ideal of R. Since $(0) \neq I \subseteq IR$, we have $IR = R$ by the simplicity of R. Therefore, $R = RR = IRIR = I^2R$. In particular, $I^2 \neq (0)$ and we may apply Brauer's Lemma 6.17. Thus there exists an idempotent $e \in I$ such that $I = Re$ and $D = eRe$ is a division ring. Clearly $D \subseteq Re = I$. Since I is closed under right multiplication by elements of I, it naturally has the structure of a right D-module. Let $\mathrm{End}_D(I)$ be the ring of right D-module homomorphisms $\alpha : I \to I$ whose multiplication is given by right-to-left composition, so that $(\alpha\beta)(a) = \alpha(\beta(a))$ for $\alpha, \beta \in \mathrm{End}_D(I)$ and $a \in I$.

We claim that the rings R and $\mathrm{End}_D(I)$ are isomorphic. Indeed, for every $r \in R$ define $\alpha_r : I \to I$ by $\alpha_r(a) = ra$ for all $a \in I$. Clearly α_r respects the D-module structure of I, so we obtain a map $f : R \to \mathrm{End}_D(I)$ by $f(r) = \alpha_r$. It is easy to check that f is a ring homomorphism. Recall that $IR = R$. If $r \in \ker f$, then $rI = (0)$ and hence $rR = rIR = (0)$, so necessarily $r = 0$. Thus f is injective.

Finally, since $1 \in R = IR = ReR$, we may express $1 = \sum_{i=1}^m r_i e s_i$ for some $m \in \mathbb{N}$ and $r_i, s_i \in R$. Let $\alpha \in \mathrm{End}_D(I)$. Then for any $r \in R$ we have

$$\alpha(re) = \alpha(1re) = \alpha\left(\sum_{i=1}^m r_i e s_i re\right) = \alpha\left(\sum_{i=1}^m r_i e \cdot e s_i re\right) = \sum_{i=1}^m \alpha(r_i e) \cdot e s_i re,$$

where the last equality holds since $es_ire \in eRe = D$ and α is a homomorphism of right D-modules. We conclude that

$$\alpha(re) = \left(\sum_{i=1}^m \alpha(r_i e) e s_i\right) re = \alpha_x(re),$$

where $x = \sum_{i=1}^m \alpha(r_i e) e s_i \in R$. Thus $\alpha = f(x)$, so f is surjective and we have shown that $R \simeq \mathrm{End}_D(I)$.

In particular, we now know that $\mathrm{End}_D(I)$ is a simple ring. It is easy to check that

$$J = \{\alpha \in \mathrm{End}_D(I) : \alpha(I) \text{ is a finitely generated } D\text{-module}\}$$

is a two-sided ideal. We claim $J \neq (0)$. Indeed, since $I \neq (0)$ and I is a free D-module by Proposition 6.12, we may express I as a direct sum $I = I' \oplus I''$, where I' is a D-module of rank one. The projection onto the first component is a non-zero element of J. Hence $J = \mathrm{End}_D(I)$ and $I = \mathrm{id}(I)$ is a finitely-generated D-module. Thus I is a free D-module of finite rank $n \in \mathbb{N}$ and $R \simeq \mathrm{End}_D(I) \simeq M_n(D)$, completing the proof. □

The following special case is the form in which the theorem was originally proved by Wedderburn.

Corollary 6.19. *Let F be a field, and let A be a simple finite-dimensional F-algebra. Then there exist a division algebra D and a natural number $r \in \mathbb{N}$ such that $A \simeq M_r(D)$.*

Proof. Since any left ideal $I \subset A$ is closed under multiplication with $b \cdot 1$ for any $b \in F$, we see that I is an F-subspace of A. Since A is finite-dimensional, there must exist minimal left ideals. Hence the previous theorem applies. Moreover, its proof shows that $D = eRe$, which is naturally an F-algebra. □

Lemma 6.20. *Let D be a division ring and $r \in \mathbb{N}$. Any non-zero simple $M_r(D)$-module is isomorphic to D^r, with the natural action of $M_r(D)$. In particular, if A is a simple finite-dimensional F-algebra, then there is only one isomorphism class of non-zero simple A-modules.*

Proof. The second claim is immediate from the first one by Corollary 6.19. To prove the first claim, let M be a simple $M_r(D)$-module, and let $0 \neq m \in M$. There exists a matrix $c \in M_r(D)$ with only one non-zero column such that $cm \neq 0$, since every element of $M_r(D)$ is the sum of sum matrices. Suppose that the j-th column of c is not all zeroes, and define a map $\varphi : D^r \to M$ by

$$\varphi((a_1, \dots, a_r)^T) = \begin{pmatrix} 0 & \cdots & a_1 & \cdots & 0 \\ 0 & \cdots & a_2 & \cdots & 0 \\ \vdots & & \vdots & & \vdots \\ 0 & \cdots & a_r & \cdots & 0 \end{pmatrix} m,$$

where all the elements of the matrix are zero apart from those in the j-th column. We claim that φ is $M_r(D)$-linear. Indeed, this follows from the definition of φ and the observation that if $c \in M_r(D)$ and $(a_1, \dots, a_r)^T \in D^r$ and we set $(b_1, \dots, b_r)^T = c(a_1, \dots, a_r)^T$, then

$$\begin{pmatrix} 0 & \cdots & b_1 & \cdots & 0 \\ 0 & \cdots & b_2 & \cdots & 0 \\ \vdots & & \vdots & & \vdots \\ 0 & \cdots & b_r & \cdots & 0 \end{pmatrix} = c \begin{pmatrix} 0 & \cdots & a_1 & \cdots & 0 \\ 0 & \cdots & a_2 & \cdots & 0 \\ \vdots & & \vdots & & \vdots \\ 0 & \cdots & a_r & \cdots & 0 \end{pmatrix}.$$

Since M is simple and φ has non-zero image, it must be surjective. Since D^r is clearly a simple $M_r(D)$-module and φ is non-zero, it must be injective. Thus φ is an isomorphism of $M_r(D)$-modules. □

Lemma 6.21. *Let D be a division algebra, and let $r \geq 1$. Consider D^r as a left $M_r(D)$-module in the natural way. Then* $\mathrm{End}_{M_r(D)}(D^r) \simeq D^{\mathrm{op}}$.

Proof. We write elements of D^r as columns. Let $f \in \mathrm{End}_{M_r(D)}(D^r)$, and let $\varepsilon = (1, 0, \dots, 0)^T \in D^r$. Since the $M_r(D)$-orbit of ε is all of D^r, we see that f is determined by $f(\varepsilon)$. If $C \in M_r(D)$ is any matrix whose first column consists entirely of zeroes, then $C\varepsilon = (0, \dots, 0)^T$ and hence $Cf(\varepsilon) = (0, \dots, 0)^T$. It follows that $f(\varepsilon) = (d, 0, \dots, 0)^T$ for some $d \in D$. Thus, for any $d_1, \dots, d_r \in D$, we must have

$$f \begin{pmatrix} d_1 \\ d_2 \\ \vdots \\ d_r \end{pmatrix} = f \left(\begin{pmatrix} d_1 & 0 & \cdots & 0 \\ d_2 & 0 & \cdots & 0 \\ \vdots & \vdots & & \vdots \\ d_r & 0 & \cdots & 0 \end{pmatrix} \begin{pmatrix} 1 \\ 0 \\ \vdots \\ 0 \end{pmatrix} \right) = \begin{pmatrix} d_1 & 0 & \cdots & 0 \\ d_2 & 0 & \cdots & 0 \\ \vdots & \vdots & & \vdots \\ d_r & 0 & \cdots & 0 \end{pmatrix} \begin{pmatrix} d \\ 0 \\ \vdots \\ 0 \end{pmatrix} = \begin{pmatrix} d_1 d \\ d_2 d \\ \vdots \\ d_r d \end{pmatrix}.$$

On the other hand, any f as above is indeed a $M_r(D)$-endomorphism of D^r, since $M_r(D)$ acts by left multiplication on the components of elements of D^r, and this commutes with right multiplication by d. Identifying the element f as above with $d \in D$, we clearly obtain an isomorphism $\mathrm{End}_{M_r(D)}(D^r) \simeq D^{\mathrm{op}}$. □

A consequence of the previous lemma is that the division ring D and the natural number n in the statement of Wedderburn's theorem are unique.

Corollary 6.22. *Let D and D' be division algebras, and let $m, n \in \mathbb{N}$. If $M_n(D) \simeq M_m(D')$ as rings, then $D \simeq D'$ and $n = m$.*

Proof. Let $R = M_n(D)$. Combining Lemmas 6.20 and 6.21, we may recover the division algebra D from $D^{\mathrm{op}} \simeq \mathrm{End}_R(M)$, where M is any non-zero simple R-module. It can be proved just as for vector spaces that any two bases of a D-module have the same cardinality; this is called its rank. Viewing $D \subset R$ as a the subring of scalar matrices, we obtain a D-module structure on M and recover n as the rank of M. $\square$

Definition 6.23. Let A and B be two central simple F-algebras. We say that they are *equivalent* if there exists a division algebra D such that $A \simeq M_r(D)$ and $B \simeq M_s(D)$. Let $\mathrm{Br}(F)$ be the set of equivalence classes of central simple F-algebras.

Lemma 6.24. *The set* $\mathrm{Br}(F)$ *is an abelian group under the operation* $[A][B] = [A \otimes_F B]$.

Proof. The only bit of the proof that is non-trivial at this point is to verify that the operation is well-defined. This follows from the fact that $M_r(D) \otimes_F M_s(D') \simeq M_{rs}(D \otimes D')$. If we view F as an algebra over itself, then the class $[F] \in \mathrm{Br}(F)$ is an identity element. Since $[A][A^{\mathrm{op}}] = [A \otimes_F A^{\mathrm{op}}] = [M_{\dim_F(A)}(F)] = [F]$, every element of $\mathrm{Br}(F)$ has an inverse. $\square$

Definition 6.25. Let L/K be an extension of fields. There is a natural homomorphism of groups

$$\begin{aligned} \mathrm{Br}(K) &\to \mathrm{Br}(L) \\ [A] &\mapsto [A \otimes_K L]. \end{aligned}$$

We write $\mathrm{Br}(L/K)$ for the kernel.

Proposition 6.26 (Double Centralizer Theorem)**.** *Suppose that A is a central simple F-algebra, and let $B \subset A$ be a simple subalgebra. Write $C_A(B)$ for the centralizer of B, namely the set of all $a \in A$ such that $ab = ba$ for every $b \in B$. Then $C_A(B)$ is a simple F-algebra. Moreover, $\dim_F C_A(B) = \frac{\dim_F(A)}{\dim_F(B)}$ and $C_A(C_A(B)) = B$.*

Proof. Since $B \otimes_F A^{\mathrm{op}}$ is simple by Corollary 6.15, we have $B \otimes_F A^{\mathrm{op}} \simeq M_r(D)$ for some division algebra D and some $r \geq 1$ by Wedderburn's theorem. The algebra A is a $B \otimes_F A^{\mathrm{op}}$-module in the obvious way; thus it is semisimple. Since any simple $M_r(D)$-module is isomorphic to D^r by Wedderburn's theorem, we have $A \simeq (D^r)^s$ as a $M_r(D)$-module, for some integer s.

Now observe that any $f \in \mathrm{End}_{B \otimes_F A^{\mathrm{op}}}(A)$ is determined by $f(1)$. Moreover, if $f(1) = c$, then for any $a \in A$ and $b \in B$ we have the equality $bca = (b \otimes a)f(1) = f((b \otimes a)1) = f(ba) = f((1 \otimes ba)1) = (1 \otimes ba)f(1) = cba$, which forces $c \in C_B(A)$. This gives an isomorphism $C_A(B) \simeq \mathrm{End}_{B \otimes_F A^{\mathrm{op}}}(A)$. But $\mathrm{End}_{B \otimes_F A^{\mathrm{op}}}(A) \simeq \mathrm{End}_{M_r(D)}((D^r)^s) \simeq M_s(\mathrm{End}_{M_r(D)}(D^r)) \simeq M_s(D^{\mathrm{op}})$, where the last isomorphism comes from Lemma 6.21. Hence $C_A(B)$ is simple.

It follows from the previous paragraph that

$$r^2 \dim_F D = \dim_F M_r(D) = \dim_F(B \otimes_F A^{\mathrm{op}}) = (\dim_F A)(\dim_F B).$$

On the other hand, $A \simeq (D^r)^s$, so $\dim_F A = rs \dim_F D$. This in turn implies that $\dim_F B = \frac{r}{s}$. We have also shown that $C_A(B) \simeq M_s(D^{\mathrm{op}})$, hence $\dim_F C_A(B) = s^2 \dim_F D$, and the claimed relation of dimensions follows immediately.

Finally, since $C_A(B)$ is simple, the second part of the claim implies that $\dim_F C_A(C_A(B)) = \frac{\dim_F A}{\dim_F C_A(B)} = \dim_F(B)$. Clearly $B \subseteq C_A(C_A(B))$. Since these two F-algebras have the same dimension, they are equal. $\square$

The following statement was essentially established in the proof of the Double Centralizer Theorem, but it is not traditionally stated as part of that theorem.

Corollary 6.27. *Let A be a central simple F-algebra and $B \subset A$ a simple subalgebra. Let $n = \dim_F B$. Then $A \otimes_F B^{\mathrm{op}} \simeq M_n(C_A(B))$.*

Proof. We saw in the proof of Proposition 6.26 that $B \otimes_F A^{\mathrm{op}} \simeq M_r(D)$ for a suitable division algebra D. Hence $A \otimes_F B^{\mathrm{op}} \simeq (B \otimes_F A^{\mathrm{op}})^{\mathrm{op}} \simeq (M_r(D))^{\mathrm{op}} \simeq M_r(D^{\mathrm{op}})$. It was shown later in the same proof that $C_A(B) \simeq M_s(D^{\mathrm{op}})$. Thus $A \otimes_F B^{\mathrm{op}} \simeq M_{r/s}(C_A(B))$. But $\dim_F B = \frac{r}{s}$, which establishes our claim. □

Theorem 6.28 (Skolem-Noether)**.** *Let F be a field, and let A and B be two simple F-algebras. Suppose that $Z(B) = F$ and that B has finite dimension over F. Let $f, g : A \to B$ be two F-algebra homomorphisms. Then there exists a unit $b \in B$ such that $g(a) = b \cdot f(a) \cdot b^{-1}$ for all $a \in A$.*

Proof. First we treat the special case $B = M_n(F) = \mathrm{End}_F(F^n)$. In this case, we can view the maps f and g as specifying two A-actions on the vector space F^n. Let V_f and V_g be the corresponding A-modules. Since there is only one simple A-module up to isomorphism, every finite-dimensional A-module is a direct sum of simple A-modules, and $\dim_F V_f = \dim_F V_g = n$, we see that V_f and V_g must be isomorphic as A-modules. Let $b : V_f \to V_g$ be an A-module isomorphism. Forgetting the A-module structure, we have that $b : F^n \to F^n$ is a linear transformation and hence $b \in M_n(F) = B$. Since B respects the A-module structure, for all $a \in A$ we have $bf(a) = g(a)b$ as claimed.

Now consider the general case. Since B is central simple, we know that $B \otimes_F B^{\mathrm{op}} \simeq M_n(F)$ by Corollary 6.14, where $n = \dim_F(B)$. Moreover, $A \otimes_F B^{\mathrm{op}}$ is a simple F-algebra by Corollary 6.15. We obtain a map

$$\begin{aligned} f \otimes 1 : A \otimes_F B^{\mathrm{op}} &\to B \otimes_F B^{\mathrm{op}} \simeq M_n(F) \\ a \otimes c &\mapsto f(a) \otimes c, \end{aligned}$$

for all $a \in A$ and $c \in B^{\mathrm{op}}$. We define $g \otimes 1$ similarly. By the case that we have proved already, there exists an element $\beta \in B \otimes_F B^{\mathrm{op}}$ such that

$$(g \otimes 1)(a \otimes c) = \beta \cdot (f \otimes 1)(a \otimes c) \cdot \beta^{-1} \tag{14}$$

for all $a \in A$ and $c \in B^{\mathrm{op}}$. Taking $a = 1$, we find that $\beta(1 \otimes c)\beta^{-1} = 1 \otimes c$, namely that $\beta \in C_{B \otimes_F B^{\mathrm{op}}}(F \otimes_F B^{\mathrm{op}}) = B \otimes_F F$. To establish the last equality, note that $B \otimes_F F$ is obviously contained in the centralizer, and it has the same dimension as the centralizer by the Double Centralizer Theorem. Hence $\beta = b \otimes 1$ for some $b \in B$, and taking $c = 1$ in (14), we see that b has the property we want. □

Note that the matrix algebras $M_m(F)$ are central simple F-algebras, and Lemma 6.8, which was already used above, is just the previous theorem in the case $A = B = M_m(F)$. Having repaid our Skolem-Noether debt, we deduce some further corollaries of the Double Centralizer Theorem.

Corollary 6.29. *Let D be a central division F-algebra. Then $\dim_F D$ is a square, and any maximal subfield L of D has degree $[L : F] = \sqrt{\dim_F D}$ over F. Moreover, $D \otimes_F L \simeq M_{[L:F]}(L)$.*

Proof. Observe that D does have subfields, since F is one. Let L be a maximal subfield of D. Since L is commutative, we have $L \subseteq C_D(L)$. This inclusion is in fact an equality; otherwise, we could take $x \in C_D(L) \setminus L$, and the subalgebra $L(x)$ would be a commutative division algebra, hence a field, contradicting the maximality of L. Hence $L = C_D(L)$. Then the Double Centralizer Theorem tells us that $\dim_F L = \frac{\dim_F D}{\dim_F L}$, whence $\dim_F D = [L : F]^2$. Since $L = L^{\mathrm{op}}$, the final claim follows from Corollary 6.27. $\square$

Definition 6.30. Let A be a simple F-algebra. Let L/F be a field extension. If $A \otimes_F L \simeq M_n(L)$ for some $n \in \mathbb{N}$, we say that L is a splitting field of A, or that A splits over L.

It is an exercise to show that if L is a splitting field of A, then so is any field containing L.

Lemma 6.31. *Let A be a central simple F-algebra, and let $A \simeq M_r(D)$ for a division algebra D and $r \in \mathbb{N}$. Then A splits over an extension L/F if and only if D splits over L.*

Proof. If $D \otimes_F L \simeq M_n(L)$ for some $n \in \mathbb{N}$, then clearly $A \otimes_F L \simeq M_r(D) \otimes_F L \simeq M_r(D \otimes_F L) \simeq M_r(M_n(L)) \simeq M_{rn}(L)$.

Conversely, suppose that $A \otimes_F L \simeq M_n(L)$. We know that $D \otimes_F L$ is a simple F-algebra by Corollary 6.15, hence $D \otimes_F L \simeq M_m(D')$ by Wedderburn, for some division F-algebra D' and some $m \in \mathbb{N}$. Thus $M_n(L) \simeq A \otimes_F L \simeq M_r(D \otimes_F L) \simeq M_{rm}(D')$. By the uniqueness of Corollary 6.22 we must have $D' \simeq L$, and thus D splits over L. $\square$

Corollary 6.32. *Let A be a central simple F-algebra, and let $A \simeq M_r(D)$. If L is a maximal subfield of D, then $A \otimes_F L \simeq M_n(L)$ for some $n \in \mathbb{N}$. In particular, $A \in \mathrm{Br}(L/F)$.*

Proof. Embedding $D \subset A$ as a the subalgebra of scalar matrices, we see that $Z(D) = Z(A) \simeq F$. Thus D is a central division F-algebra. The claim now follows from Corollary 6.29 and Lemma 6.31. $\square$

Thus we have shown that for every central simple F-algebra A, there is a finite extension L/F splitting A. We will need a bit more, namely that the extension L/F may be taken to be separable. If F has characteristic zero or is a perfect field of positive characteristic, then any finite extension of F is separable and this is automatic, but a bit of work is needed to obtain this claim in general.

Proposition 6.33. *Let D be a central division F-algebra, and let $L \subset D$ be a subfield such that L/F is a separable extension. Then there exists a maximal subfield of D that contains L and is separable over F.*

Proof. We start with two reduction steps. First, it clearly suffices to prove our claim in the case where L is maximal among subfields of D that are separable over F. In this case, the claim is that L is itself a maximal subfield of D. We will now assume that we are in this case. Secondly, any subfield $K \subset D$ that contains L must be contained in the centralizer $C_D(L)$. If K/L is separable, then K/F is also separable since L/F is separable. By the Double Centralizer theorem, we know that $Z(C_D(L)) \subset C_D(C_D(L)) = L$ and hence $C_D(L)$ is a central division L-algebra.[3] Thus, replacing D by $C_L(D)$, we may assume without loss of generality that $L = F$ and that there is no non-trivial separable extension of F contained in D.

Let $p = \operatorname{char} F$; as noted above, we may assume that $p > 0$. Let $n = \sqrt{\dim_F D}$. We will now prove that $D = F$; by Corollary 6.29 this is equivalent to the claim that F is a maximal subfield of D. Indeed, suppose that $n > 1$ and let $a \in D \setminus F$. Then $F(a)/F$ is a purely inseparable field extension, so the minimal polynomial of a has the form $x^{p^t} - c$ for some $c \in F$ and $t \in \mathbb{N}$. Moreover, the degree p^t is bounded by n, which is the degree of a maximal subfield of D. Thus there exists $t \in \mathbb{N}$ such that $a^{p^t} \in F = Z(D)$ for *all* $a \in D$; for

[3]Explain why division algebra.

instance, we may take $t = \lfloor \log_p n \rfloor$. Hence all $a, b \in D$ satisfy the identity $a^{p^t} b - b a^{p^t} = 0$. This identity continues to hold after extending scalars, in the ring $D \otimes_F K \simeq M_n(K)$, where K is a maximal subfield of D; indeed, the identity clearly holds for pure tensors, and it is easily checked to hold for mixed tensors since $\operatorname{char} F = p$. This gives rise to a contradiction, however, since $n > 1$ and the element $a = \operatorname{diag}(1, 0, \ldots, 0) \in M_n(K)$ is an idempotent that does not lie in the center; thus a suitable $b \in M_n(K)$ will violate the identity. □

Corollary 6.34. *Let A be a central simple F-algebra. There exists a finite separable extension L/F that splits A.*

Proof. This is immediate from Lemma 6.31 and Proposition 6.33. □

Using the tools we have developed, we can finally understand the structure of the cohomology group $H^2(\operatorname{Gal}(L/K), L^\times)$. Observe that if $m|q$ are two natural numbers, then the diagonal embedding $\Delta : \mathrm{GL}_m(L) \hookrightarrow \mathrm{GL}_q(L)$ sending $A \in \mathrm{GL}_m(L)$ to the block-diagonal matrix $\operatorname{diag}(A, A, \ldots, A)$ maps scalar matrices, and only scalar matrices, to scalar matrices in $\mathrm{GL}_q(L)$. Thus it induces an embedding $\mathrm{PGL}_m(L) \hookrightarrow \mathrm{PGL}_q(L)$ and hence a map $\Delta : H^1(G, \mathrm{PGL}_m(L)) \to H^1(G, \mathrm{PGL}_q(L))$. Recall that $A(L/K, m)$ denotes the set of K-isomorphism classes of central simple K-algebras A such that $A \otimes_K L \simeq M_m(L)$. We want to understand the map $A(L/K, m) \to A(L/K, q)$ that corresponds to Δ under the bijection of Lemma 6.9.

Let $(\sigma \mapsto \overline{\Psi(\sigma)}) \in Z^1(G, \mathrm{PGL}_m(L))$ be a 1-cocycle. Recall that in Lemma 6.9 we defined an action of G on $M_m(L)$ as follows: an element $\sigma \in G$ sends $c \in M_m(L)$ to the matrix $(1 \otimes \sigma)(c) = \Psi(\sigma)\sigma(c)\Psi(\sigma)^{-1}$, where $\sigma(c)$ denotes the action of G on matrix elements as in Example 6.1. Then the element of $A(L/K, m)$ associated to our cocycle is (the isomorphism class of) the algebra A of G-invariants under this new action.

It is easy to check that the algebra associated to the cocyle $(\sigma \mapsto \overline{\operatorname{diag}(\Psi(\sigma), \ldots, \Psi(\sigma))}) \in Z^1(G, \mathrm{PGL}_q(L))$ consists of all matrices of the form

$$c = \begin{pmatrix} c_{11} & c_{12} & \cdots & c_{1d} \\ c_{21} & c_{22} & \cdots & c_{2d} \\ \vdots & \vdots & \ddots & \vdots \\ c_{d1} & c_{d2} & \ldots & c_{dd} \end{pmatrix},$$

where $d = \frac{q}{m}$ and each $c_{ij} \in M_m(L)$ is contained in A. Thus, if a class $\psi \in H^1(G, \mathrm{PGL}_m(L))$ corresponds to an algebra isomorphic to $M_r(D)$ for some division algebra D and some $r \geq 1$, then $\Delta(\psi) \in H^1(G, \mathrm{PGL}_q(L))$ corresponds to $M_{rd}(D)$. We have thus proved the following.

Lemma 6.35. *The bijections of Lemma 6.9 induce a bijection of pointed sets*

$$\varinjlim H^1(G, \mathrm{PGL}_m(L)) \to \operatorname{Br}(L/K). \tag{15}$$

Proof. The content of the claim is that the direct limit $\varinjlim A(L/K, m)$ induced by the maps of Lemma 6.9 is exactly $\operatorname{Br}(L/K)$, as a set. This is immediate from the calculation we just did. □

Since $\operatorname{Br}(L/K)$ is a group under the operation $[A][B] = [A \otimes_K L]$, it follows that the injective limit on the left-hand side of (15) has a natural group structure even though the sets $H^1(G, \mathrm{PGL}_m(L))$ do not. Moreover, recalling the maps $\delta^1 : H^1(G, \mathrm{PGL}_m(L)) \to H^2(G, L^\times)$ from Lemma 6.6, the universal property of the injective limit and the previous lemma produce a map $\operatorname{Br}(L/K) \to H^2(G, L^\times)$, which we continue to label δ^1.

Theorem 6.36. *The map $\delta^1 : \mathrm{Br}(L/K) \to H^2(G, L^\times)$ is an isomorphism of abelian groups.*

Proof. It is enough to show that δ^1 is a group homomorphism. Indeed, δ^1 is surjective by Lemma 6.6, and by Corollary 6.5 it has trivial kernel, which for group homomorphisms implies injectivity.

So let A and B be two central simple K-algebras such that $[A], [B] \in \mathrm{Br}(L/K)$. Let $A \in A(L/K, m)$ and $B \in A(L/K, m')$. If A and B correspond to the cocycles $(\sigma \mapsto \overline{\Psi(\sigma)})$ and $(\sigma \mapsto \overline{\Phi(\sigma)})$ in $H^1(G, \mathrm{PGL}_m(L))$ and $H^1(G, \mathrm{PGL}_{m'}(L))$, respectively, then it is easy to see that $A \otimes_K B$ corresponds to $(\sigma \mapsto \overline{\Xi(\sigma)})$, where $\Xi(\sigma) \in M_{mm'}(L)$ corresponds to the endomorphism of $L^{mm'} = L^m \otimes_L L^{m'}$ satisfying

$$(\Xi(\sigma))(v \otimes v') = (\Psi(\sigma))v \otimes \Phi(\sigma)v'$$

for all $v \in L^m$ and $v' \in L'$. By the definition of δ^1, we have

$$\Xi(g_1) \circ g_1\Xi(g_2) \circ \Xi(g_1g_2)^{-1} = \delta^1([A \otimes_K B])I_{mm'}$$

as automorphisms of $L^m \otimes_L L^{m'}$, for any $g_1, g_2 \in G$. Applying both sides to a vector of the form $v \otimes v'$, we find that $\delta^1([A \otimes_K L]) = \delta^1([A])\delta^1([B])$, as claimed. $\square$

Corollary 6.37. *Let K be a field, let $\overline{K}$ denote a separable closure, and let $G_K = \mathrm{Gal}(\overline{K}/K)$ be the absolute Galois group of K. Then $\mathrm{Br}(K) \simeq H^2(G_K, \overline{K}^\times)$.*

Proof. On one hand, $H^2(G_K, \overline{K}^\times) = \varinjlim H^2(\mathrm{Gal}(L/K), L^\times)$ by Proposition 4.10, where L/K runs over finite Galois extensions. On the other hand, by Corollary 6.34, every central simple K-algebra is split by some finite Galois extension L/K, so that $\mathrm{Br}(K)$ is the union of the $\mathrm{Br}(L/K)$. Thus $\mathrm{Br}(K) = \varinjlim \mathrm{Br}(L/K)$, where the connecting maps are inclusions. Finally, one checks that the maps $\delta^1 : \mathrm{Br}(L/K) \xrightarrow{\sim} H^2(\mathrm{Gal}(L/K), L^\times)$ are compatible with the injective systems. $\square$

Corollary 6.38. *Let L/K be a finite Galois extension of degree $n = [L : K]$. Then $\mathrm{Br}(L/K)$ is a torsion group with exponent dividing n.*

Proof. Let $G = \mathrm{Gal}(L/K)$. Then $\{e\}$ is an open subgroup of G, and as in the proof of Corollary 4.13, we observe that the composition

$$H^2(G, L^\times) \xrightarrow{\mathrm{res}} H^2(\{e\}, \mathrm{Res}^G_{\{e\}} L^\times) \xrightarrow{\mathrm{cor}} H^2(G, L^\times)$$

is zero since the middle group is trivial. On the other hand, this composition is multiplication by $[G : \{e\}] = n$ by Lemma 4.4. Thus all elements of $H^2(G, L^\times) \simeq \mathrm{Br}(L/K)$ are killed by multiplication by n. $\square$

Exercises

(1) Let G be a profinite group, and let $\mathcal{M}_G$ be the category of groups M, not necessarily abelian, with a G-module structure as in Section 6.1. The morphisms are, of course, G-equivariant group homomorphisms. Show that $H^1(G, -) : \mathcal{M}_G \to \mathrm{PtSet}$ is a covariant functor.

(2) Let A be a central simple F-algebra and let $B \subset A$ be a subfield containing F. Prove that the following statements are equivalent:
  (a) B is maximal as a commutative subring of A;
  (b) $C_A(B) = B$;
  (c) $\dim_F A = (\dim_F B)^2$.

7. Examples of Brauer groups and the invariant map

7.1. **Trivial Brauer groups.** So far we have not explicitly computed a single Brauer group. We will now remedy this situation, developing useful techniques along the way. Some of the arguments we give are much more complicated than is necessary to compute the relevant Brauer groups, but the idea is to introduce tools that will serve us well later in the course, and in later life.

Proposition 7.1. *If K is an algebraically closed field, then* $\mathrm{Br}(K)$ *is trivial.*

Proof. Let A be a central simple K-algebra. By Corollary 6.32 we see that $[A] \in \mathrm{Br}(L/K)$ for some finite extension L/K. Since K has no nontrivial finite extensions, it follows that $[A] \in \mathrm{Br}(K/K) = 0$. □

Lemma 7.2 (Wedderburn's little theorem)**.** *Let D be a finite division ring. Then D is a field.*

Proof. Let $K = Z(D)$ be the center of D. Clearly this is a field, and D is a central division K-algebra. Suppose, by way of contradiction, that D is not a field, so that $n = \sqrt{\dim_K D} > 1$. Any element $x \in D$ is contained in the subfield $K(x) \subset D$ and hence in some maximal subfield of D. By Corollary 6.29, all maximal subfields of D have degree n over K and thus have cardinality $|K|^n$.

Let L and L' be two maximal subfields of D. Since they have the same finite cardinality, there exists an isomorphism $\iota : L \to L'$, which may be taken to be K-linear. Now we apply Skolem-Noether to the two maps $f, g : L \to D$, where f is the natural inclusion of L in D, and g is the composition of ι with the natural inclusion of L' in D. We find that L and L' are conjugate in D, and hence that the group $D^\times$ is the union of the conjugates of the subgroup $L^\times$. Moreover, since D is not a field and hence $D \neq L$, we have that $L^\times$ is a proper subgroup of $D^\times$. We have arrived at a contradiction, since a finite group cannot be a union of conjugates of a proper subgroup.

Indeed, if G is a finite group and H is a subgroup, then the number of conjugates of H is $[G : N_G(H)]$. Since all the conjugates contain the identity element of G, we find that the cardinality of the union of the subgroups conjugate to H is at most $[G : N_G(H)](|H|-1)+1 \leq [G : H](|H| - 1) + 1$, and this is strictly smaller than $|G|$ if H is proper. □

Proposition 7.3. *If K is a finite field, then* $\mathrm{Br}(K)$ *is trivial.*

Proof. Let A be a central simple K-algebra. Then, by Wedderburn's big theorem (Corollary 6.19), we have $A \simeq M_r(D)$, where D is a division ring that is central as a K-algebra. Since D is finite-dimensional over K, it is finite. Thus D is a field by Wedderburn's little theorem (Lemma 7.2), so that $D = Z(D) = K$. □

7.2. **Brauer groups of local fields.** We now begin to consider one of the most interesting cases, namely that of local fields. The study of their Brauer groups will lead to a deeper understanding of their Galois cohomology. From this point onwards in the course, we will rely on standard theorems from algebraic number theory.

Proposition 7.4. *Let $K/\mathbb{Q}_p$ be a finite extension and let D be a division K-algebra. There exists an* unramified *finite extension L/K such that D splits over L.*

Proof. Let k be the residue field of K, and let $n = \sqrt{\dim_K(D)}$. Let $|\cdot|_K$ denote the (multiplicative) valuation of K. We claim that one can construct a non-Archimedean valuation $|\cdot|_D : D \to \mathbb{R}_{\geq 0}$ such that

- For any $x \in D$, we have $|x|_D = 0$ if and only if $x = 0$.
- For any $x, y \in D$, we have $|xy|_D = |x|_D|y|_D$.
- For any $x, y \in D$, we have $|x + y|_D \leq \max\{|x|_D, |y|_D\}$.

Indeed, for every $x \in D$, we define $|x|_D = |\det A_x|_K$, where A_x is a matrix, with entries in K, representing the K-linear map

$$\begin{array}{ccc} D & \to & D \\ y & \mapsto & yx \end{array}$$

with respect to some K-basis of D. It is clear that $|\cdot|_D$, thus defined, satisfies the first two of the three desired properties, and that $|x|_D = |x|_K^{\dim_K(D)}$ whenever $x \in K$. Note that the third property is equivalent to the claim that $|1 + z|_D \leq 1$ whenever $z \in D$ is such that $|z|_D \leq 1$. As in the previous proof, $K(z)$ is a subfield of D. Judiciously choosing a K-basis of D of the form $\alpha_i\beta_j$, where $\{\alpha_i\}$ is a $K(z)$-basis of D and $\{\beta_j\}$ is a K-basis of $K(z)$, we see that, if $t \in K(z)$, then A_z is a block-diagonal matrix and that

$$|t|_D = |\det A_t|_K = |\mathrm{N}_{K(z)/K}(t)|_K^{\dim_{K(z)}(D)}. \tag{16}$$

We know from algebraic number theory that the right-hand side of (16) defines a multiplicative valuation of $K(z)$. Considering $t = z$ and $t = 1 + z$, we obtain the third property as well.

As in the commutative case, we define $\mathcal{O}_D = \{x \in D : |x|_D \leq 1\}$. By our three properties, this is a ring, and $\mathfrak{m}_D = \{x \in D : |x|_D < 1\}$ is a two-sided ideal. It is a maximal ideal since all elements of $\mathcal{O}_D \setminus \mathfrak{m}_D$ have inverses in $\mathcal{O}_D$, and hence the quotient $\Delta = \mathcal{O}_D/\mathfrak{m}_D$ is a division ring. For every subfield $K \subset L \subset D$ we have $\mathcal{O}_D \cap L = \mathcal{O}_L$ by (16), and thus $\mathcal{O}_D$ consists precisely of the elements of D that are integral over K. The same argument as in the commutative case shows that (the underlying abelian group of) $\mathcal{O}_D$ is a finitely generated $\mathcal{O}_K$-module. Since $\mathcal{O}_D$ is torsion-free and $\mathcal{O}_K$ is a principal ideal domain, $\mathcal{O}_D$ is a free $\mathcal{O}_K$-module, and its rank must be $\dim_K D = n^2$, since $\mathcal{O}_D \otimes_{\mathcal{O}_K} K = D$. Since $\mathfrak{m}_K \subset \mathfrak{m}_D$, we find that Δ is naturally a finite-dimensional vector space over k. Hence Δ is finite, and thus it is a field by Lemma 7.2. Set $f = \dim_k \Delta$.

Since all extensions of finite fields are finite and separable and hence simple, there exists an element $\overline{\delta} \in \Delta$ such that $\Delta = k(\overline{\delta})$. Let $\delta \in \mathcal{O}_D$ be a lift of $\overline{\delta}$. Then $K(\delta)$ is a subfield of D with residue field Δ. Thus $f \leq [K(\delta) : K] \leq n$, where the second inequality is Corollary 6.29.

Since $\mathcal{O}_D$ is discretely valued, we can prove exactly as for commutative discretely valued rings that any two-sided ideal of $\mathcal{O}_D$ is principal, generated by an element of maximal valuation, and hence is either zero or of the form $\mathfrak{m}_D^r$ for some $r \geq 0$. Thus we can define the ramification index of D/k, as in the commutative case, to be the natural number e satisfying $\mathfrak{m}_K\mathcal{O}_D = \mathfrak{m}_D^e$. Finally, if π_D and π_K are uniformizers (i.e. elements of maximal valuation) in $\mathcal{O}_D$ and $\mathcal{O}_K$, respectively, then (16) implies that

$$|\pi_D|_D = |\mathrm{N}_{K(\pi_D)/K}(\pi_D)|_K^{\dim_{K(\pi_D)}(D)} = |\pi_K^{\frac{[K(\pi_D):K]}{e(K(\pi_D)/K)}}|_K^{\dim_{K(\pi_D)}(D)} = |\pi_K|_D^{\frac{1}{e(K(\pi_D)/K)}},$$

where $e(K(\pi_D)/K)$ is the ramification index of the field extension $K(\pi_D)/K$. Hence $e = e(K(\pi_D)/K) \leq [K(\pi_D) : K] \leq n$. Now $ef = \mathrm{rank}_{\mathcal{O}_K} \mathcal{O}_D = n^2$; one can check that the proof for the commutative case given in, say, Proposition II.6.8 of Neukirch's *Algebraic Number Theory* transfers verbatim to our case. However, we have already shown that $f \leq n$ and $e \leq n$. Hence $e = f = n$. In particular, we must have $[K(\delta) : K] = f = [\Delta : k]$, so that $K(\delta)/K$ is an unramified extension. However, $[K(\delta) : K] = n$, so $D \otimes_K K(\delta)$ splits over $K(\delta)$ by Corollaries 6.29 and 6.32. $\square$

For the rest of this section, K will always denote a local field. Let k denote the residue field $\mathcal{O}_K/\mathfrak{m}_K$, and let K^{nr} be the maximal unramified extension of K; this is the compositum of all the unramified extensions of K inside some fixed algebraic closure. Recall that K^{nr}/K is an infinite Galois extension, and that $\mathrm{Gal}(K^{\mathrm{nr}}/K)$ is isomorphic in a standard way to the absolute Galois group $G_k = \mathrm{Gal}(\overline{k}/k)$ of the residue field k. The kernel of the projection $G_K \twoheadrightarrow G_k$ is $I_K = \mathrm{Gal}(\overline{K}/K^{\mathrm{nr}})$, which is called the *inertia subgroup* of G_K. Furthermore, G_k has a dense cyclic subgroup generated by the *arithmetic Frobenius* element $\Phi_k : \overline{k} \to \overline{k}$, where $\Phi_k(x) = x^q$ for all $x \in \overline{k}$ and q is the cardinality of k. We write Frob_k for the *geometric* Frobenius, namely the inverse of Φ_k. Of course, Frob_k and Φ_k generate the same cyclic subgroup of G_k. Be aware that some books use the notation Frob_k for the arithmetic Frobenius.

Corollary 7.5. *Let $K/\mathbb{Q}_p$ be a finite extension. Then $\mathrm{Br}(K) = \mathrm{Br}(K^{\mathrm{nr}}/K) \simeq H^2(G_k, (K^{\mathrm{nr}})^\times)$.*

Proof. The first equality is the content of Proposition 7.4. The isomorphism of $\mathrm{Br}(K^{\mathrm{nr}}/K)$ with $H^2(G_k, (K^{\mathrm{nr}})^\times)$ is proved in exactly the same way as Corollary 6.37. □

Before proceeding, we make a brief study of the cohomology of the group $\widehat{\mathbb{Z}} = \varprojlim \mathbb{Z}/n\mathbb{Z}$. The transitions maps are the surjections $\mathbb{Z}/m\mathbb{Z} \to \mathbb{Z}/n\mathbb{Z}$ for $n|m$. It is easy to see that $\widehat{\mathbb{Z}}$ is indeed the profinite completion of $\mathbb{Z}$.

Proposition 7.6. *If k is any finite field, then $G_k = \mathrm{Gal}(\overline{k}/k) \simeq \hat{\mathbb{Z}}$.*

Proof. This follows from the fact that k has a unique extension of any degree $n \geq 1$ inside a fixed algebraic closure $\overline{k}$. Moreover, this extension is Galois with a cyclic Galois group. □

Lemma 7.7. *Let M be a finite $\widehat{\mathbb{Z}}$-module. Then $H^2(\widehat{\mathbb{Z}}, M) = 0$.*

Proof. Let E be an extension of $\widehat{\mathbb{Z}}$ by M, namely a short exact sequence $0 \to M \xrightarrow{\iota} E \xrightarrow{\pi} \widehat{\mathbb{Z}} \to 0$ of profinite groups, where the maps are continuous group homomorphisms. Let $F \in \widehat{\mathbb{Z}}$ be a topological generator, and let $x \in E$ be a preimage of F. The map $\mathbb{Z} \to E$ given by $n \mapsto x^n$ extends to a continuous homomorphism $\varepsilon : \widehat{\mathbb{Z}} \to E$ by Proposition 1.18. Since $\pi \circ \varepsilon$ is continuous and is equal to the identity on the dense subgroup $\langle F \rangle \subset \widehat{\mathbb{Z}}$, it must be the identity. Hence π has a section, and we conclude that $E \simeq \widehat{\mathbb{Z}} \times M$ is a split extension. Since the group $H^2(\widehat{\mathbb{Z}}, M)$ classifies extensions up to congruence by an exercise in Section 3, it must be trivial.

An alternative proof of this lemma, relying on Tate cohomology rather than the properties of profinite completion, will be given at the end of this section. □

Corollary 7.8. *Let M be a torsion $\widehat{\mathbb{Z}}$-module. Then $H^i(\widehat{\mathbb{Z}}, M) = 0$ for all $i \geq 2$.*

Proof. This is immediate from the previous lemma by Lemma 4.15. □

The next result often provides a handy way to verify that cohomology groups of a G-module M vanish, if we can find a filtration of M with tractable graded pieces. In principle we could have proved it much earlier, but we had no need for it until now.

Lemma 7.9. *Let G be a finite group and let M be a G-module. Suppose that we have a descending filtration of M by open submodules:*

$$M = M_0 \supseteq M_1 \supseteq M_2 \supseteq M_3 \supseteq \cdots$$

such that $\bigcap_{j=0}^{\infty} M_j = 0$. *Suppose that* M *is complete for the topology defined by the filtration* $\{M_j\}$ *(cosets of the submodules* M_j *form a base of open sets). Let* $i \geq 0$ *and suppose that* $H^i(G, M_j/M_{j+1}) = 0$ *for all* $j \geq 0$. *Then* $H^i(G, M) = 0$.

Proof. Let $\psi_0 : G^i \to M$ be an i-cocycle; our aim is to show that it is also a coboundary. Composing ψ_0 with the natural projection $M = M_0 \twoheadrightarrow M_0/M_1$, we get a cocycle $\overline{\psi_0} \in Z^i(G, M_0/M_1)$. Since $H^i(G, M_0/M_1) = 0$ by assumption, there exists a cochain $\overline{\varphi_1} \in C^{i-1}(G, M_0/M_1)$ such that $d_{i-1}(\overline{\varphi_1}) = \overline{\psi_0}$. Let $\varphi_1 : G^{i-1} \to M_0$ be any lift of $\overline{\varphi_1}$; since G is finite, φ_1 is automatically continuous and thus a cochain. Now define $\psi_1 = \psi_0 - d_{i-1}(\varphi_1)$. Then $d_i\psi_1 = 0$ and the image of ψ_1 lies in M_1, so $\psi_1 \in Z^i(G, M_1)$.

We can now play the same game with ψ_1. Continuing this process, we obtain a sequence of cocycles $\psi_j \in Z^i(G, M_j)$ and of cochains $\varphi_j \in Z^i(G, M_{j-1})$ such that for each $j \geq 0$ we have $\psi_j = d_{i-1}(\varphi_{j+1}) + \psi_{j+1}$. Now define $\varphi = \sum_{j=1}^{\infty} \varphi_j$. By our assumptions on M the series converges. The difference $\psi_0 - d_{i-1}\varphi$ is congruent to zero modulo every M_j, hence it is zero. We have proved that ψ_0 is a coboundary, as claimed. □

Our next goal is to obtain an explicit description of the Brauer group $\mathrm{Br}(K)$ of a finite extension $K/\mathbb{Q}_p$. To this end, we prepare some final tools. First of all, let $v_K : K^\times \twoheadrightarrow \mathbb{Z}$ be the normalized additive valuation of K: if $x \in K^\times$ and $m \in \mathbb{Z}$, then $v_K(x) = m$ if and only if $x = u\pi_K^m$, where $u \in \mathcal{O}_K^\times$. We continue to denote the unique extension of v_K to $(K^{\mathrm{nr}})^\times$ by v_K. Since π_K is still a uniformizer in K^{nr}, we have that $v_K((K^{\mathrm{nr}})^\times) = \mathbb{Z}$. Since the action of G_k on (K^{nr}) preserves the ideals of $\mathcal{O}_{K^{\mathrm{nr}}}$ and hence preserves valuation, we see that $v_K : (K^{\mathrm{nr}})^\times \to \mathbb{Z}$ is a G_k-module map, where G_k acts on $\mathbb{Z}$ trivially. This induces a natural map on cohomology:

$$v_K : \mathrm{Br}(K) \simeq H^2(G_k, (K^{\mathrm{nr}})^\times) \to H^2(G_k, \mathbb{Z}).$$

Next consider the short exact sequence $0 \to \mathbb{Z} \to \mathbb{Q} \to \mathbb{Q}/\mathbb{Z} \to 0$ of abelian groups. We view all three groups as G_k-modules, with a trivial action of G_k. Looking at the following bit of the long exact cohomology sequence:

$$\cdots \to H^1(G_k, \mathbb{Q}) \to H^1(G_k, \mathbb{Q}/\mathbb{Z}) \to H^2(G_k, \mathbb{Z}) \to H^2(G_k, \mathbb{Q}) \to \cdots$$

and observing that the leftmost and rightmost groups vanish by Corollary 4.13, we conclude that $H^2(G_k, \mathbb{Z}) \simeq H^1(G_k, \mathbb{Q}/\mathbb{Z})$.

Finally we have a map $\gamma : H^1(G_k, \mathbb{Q}/\mathbb{Z}) \to \mathbb{Q}/\mathbb{Z}$ given by $\gamma([\psi]) = \psi(\mathrm{Frob}_k^{-1})$ for every $\psi \in Z^1(G_k, \mathbb{Q}/\mathbb{Z})$. This is well-defined: since G_k acts trivially on $\mathbb{Q}/\mathbb{Z}$, there are no non-zero 1-coboundaries, and so every cohomology class in $H^1(G_k, \mathbb{Q}/\mathbb{Z})$ contains only one 1-cocycle. Furthermore, again because G_k acts trivially on $\mathbb{Q}/\mathbb{Z}$, the 1-cocycles are nothing more than continuous group homomorphisms $G_k \to \mathbb{Q}/\mathbb{Z}$. This implies that γ is an isomorphism. Indeed, γ is injective because a *continuous* homomorphism $G_k \to \mathbb{Q}/\mathbb{Z}$ is determined by its restriction to the dense cyclic subgroup $\langle \mathrm{Frob}_k \rangle$. It is surjective because any homomorphism $\langle \mathrm{Frob}_k \rangle \to \mathbb{Q}/\mathbb{Z}$ can be extended to G_k by the injectivity of $\mathbb{Q}/\mathbb{Z}$. (Why can we always extend to a *continuous* homomorphism?)

Putting all this together, we define the *Hasse invariant* $\mathrm{Inv}_K : \mathrm{Br}(K) \to \mathbb{Q}/\mathbb{Z}$ as a composition of four maps:

$$\mathrm{Br}(K) \simeq H^2(G_k, (K^{\mathrm{nr}})^\times) \xrightarrow{v_K} H^2(G_k, \mathbb{Z}) \simeq H^1(G_k, \mathbb{Q}/\mathbb{Z}) \xrightarrow{\gamma} \mathbb{Q}/\mathbb{Z}.$$

Remark 7.10. To visualize the Hasse invariant map, observe that any non-trivial element of $\mathbb{Q}/\mathbb{Z}$ contains a unique representative $\frac{a}{b} \in \mathbb{Q}$ such that $0 < a < b$ and $(a, b) = 1$. Let K_b/K

be the unique unramified extension of K of degree b. Define a vector space

$$D = K_b \cdot 1 \oplus K_b\pi \oplus K_b\pi^2 \oplus \cdots \oplus K_b\pi^{b-1}$$

and endow it with a K-algebra structure as follows. Of course, π^i is the i-th power of $\pi \in D$ for all $0 \leq i \leq b-1$. Let $\beta_0, \ldots, \beta_{b-1} \in K$ be such that $v_K(\beta_0) = 1$ and $v_K(\beta_i) \geq 1$ for all $1 \leq i \leq b-1$. Set $\pi^b = \beta_0 + \beta_1\pi + \cdots + \beta_{b-1}\pi^{b-1}$. This defines all powers of π. Finally, if $\alpha \in K_b$, then set $\pi\alpha = \mathrm{Frob}_k^{-a}(\alpha)\pi$. Recall that $\mathrm{Gal}(K_b/K) \simeq \mathrm{Gal}(k_b/k)$, where the residue field k_b of K_b is the extension of degree b of the finite field k, and that $\mathrm{Frob}_k \in \mathrm{Gal}(K_b/K)$ corresponds to the inverse of the map $(x \mapsto x^p) \in \mathrm{Gal}(k_b/k)$.

We leave it as an exercise for the reader to show that D is a central simple K-algebra and that $\mathrm{Inv}_K(D) = \frac{a}{b}$.

Proposition 7.11. *Let $K/\mathbb{Q}_p$ be a finite extension. The map $\mathrm{Inv}_K : \mathrm{Br}(K) \to \mathbb{Q}/\mathbb{Z}$ is an isomorphism of abelian groups.*

Proof. We defined Inv_K as a composition of four maps, and we already know that three of them are isomorphisms. So it suffices to show that $v_K : H^2(G_k, (K^{\mathrm{nr}})^\times) \to H^2(G_k, \mathbb{Z})$ is an isomorphism. The short exact sequence $0 \to \mathcal{O}_{K^{\mathrm{nr}}}^\times \to (K^{\mathrm{nr}})^\times \xrightarrow{v_K} \mathbb{Z} \to 0$ of G_k-modules produces the fragment

$$\cdots H^2(G_k, \mathcal{O}_{K^{\mathrm{nr}}}^\times) \to H^2(G_k, (K^{\mathrm{nr}})^\times) \xrightarrow{v_K} H^2(G_k, \mathbb{Z}) \to H^3(G_k, \mathcal{O}_{K^{\mathrm{nr}}}^\times) \to \cdots$$

of the long exact cohomology sequence, and thus it suffices to prove that $H^i(G_k, \mathcal{O}_{K^{\mathrm{nr}}}^\times) = 0$ for $i \in \{2,3\}$. Proposition 4.10 tells us that

$$H^i(G_k, \mathcal{O}_{K^{\mathrm{nr}}}^\times) = \varinjlim H^i(\mathrm{Gal}(\ell/k), \mathcal{O}_L^\times),$$

where L/K runs over all finite unramified Galois extensions and ℓ is the residue field of L; since L/K is unramified, note that $\mathrm{Gal}(\ell/k) \simeq \mathrm{Gal}(L/K)$. Therefore it suffices to show that $H^i(\mathrm{Gal}(\ell/k), \mathcal{O}_L^\times) = 0$ for any finite unramified Galois L/K and $i \in \{2,3\}$.

So let L/K be a finite unramified extension and consider the following filtration of $\mathcal{O}_L^\times$ by $\mathrm{Gal}(\ell/k)$-submodules:

$$\mathcal{O}_L^\times \supset 1 + \mathfrak{m}_L \supset 1 + \mathfrak{m}_L^2 \supset 1 + \mathfrak{m}_L^3 \supset \cdots$$

where $\mathfrak{m}_L$ is the maximal ideal of $\mathcal{O}_L$. Since $\mathcal{O}_L^\times/(1+\mathfrak{m}_L) \simeq \ell^\times$ and $(1+\mathfrak{m}_L^j)/(1+\mathfrak{m}_L^{j+1}) \simeq \ell$ for all $j \geq 1$, by Lemma 7.9 it suffices to show that $H^i(\mathrm{Gal}(\ell/k), \ell) = 0$ and $H^i(\mathrm{Gal}(\ell/k), \ell^\times) = 0$ for $i \in \{2,3\}$.

We know that $H^i(\mathrm{Gal}(\ell/k), \ell) = 0$ for all $i \geq 1$ by Corollary 3.13, and furthermore we know that $H^2(\mathrm{Gal}(\ell/k), \ell^\times) = \mathrm{Br}(\ell/k) = 0$, where the first equality is Theorem 6.36 and the second is immediate from Proposition 7.3. It remains only to prove that $H^3(\mathrm{Gal}(\ell/k), \ell^\times) = 0$; however, our proof of this will also establish the other three vanishings.

Indeed, consider the Hochschild-Serre spectral sequence of Definition 5.20, with $G = \mathrm{Gal}(\overline{k}/k)$, $H = \mathrm{Gal}(\overline{k}/\ell)$, and $M = \ell^\times$. We obtain a spectral sequence

$$E_2^{p,q} = H^p(\mathrm{Gal}(\ell/k), H^q(G_\ell, \overline{k}^\times)) \Rightarrow H^{p+q}(G_k, \overline{k}^\times).$$

Since $G_\ell \simeq \widehat{\mathbb{Z}}$ and $\overline{k}^\times = \overline{\ell}^\times$ is a union of unit groups of finite fields and thus torsion, we have $H^q(G_\ell, \ell^\times) = 0$ for all $q \geq 2$ by Corollary 7.8. In addition, $H^1(G_\ell, \overline{\ell}^\times) = 0$ by Hilbert 90. Thus the E_2-sheet of our spectral sequence has only one non-zero row, and we conclude that $H^p(\mathrm{Gal}(\ell/k), \ell^\times) = E_2^{p,0} = E_\infty^{p,0} = F^pH^p(G_k, \overline{k}^\times) = 0$ for all $p \geq 1$, where the last equality holds since $H^p(G_k, \overline{k}^\times) = 0$ for all $p \geq 1$, again by Corollary 7.8 for $p \geq 2$ and Hilbert 90 for $p = 1$. $\square$

Remark 7.12. It is possible to give shorter and more direct proofs of the isomorphism $\mathrm{Br}(K) \simeq \mathbb{Q}/\mathbb{Z}$. In particular, the isomorphism may be viewed purely as a statement about central simple K-algebras, and it may be proved within that theory. We have chosen a proof that allows us to practice cohomological techniques that are useful in a wide variety of situations.

7.3. **Tate cohomology.** Let G be a *finite* group and let M be a G-module. In this case, M is equipped with a norm map $\mathrm{N} : M \to M$ given by

$$\mathrm{N}(m) = \sum_{g \in G} gm \tag{17}$$

for all $m \in M$. We now define a slightly modified version of cohomology as follows:

Definition 7.13. Let G be a finite group and M a G-module. Then we set

$$\hat{H}^i(G, M) = \begin{cases} M^G/N(M) & : i = 0 \\ H^i(G, M) & : i > 0. \end{cases}$$

Proposition 7.14. *Let G be a finite group, and let $0 \to A \to B \to C \to 0$ be a short exact sequence of G modules. There is a long exact sequence*

$$\hat{H}^0(G, A) \to \hat{H}^0(G, B) \to \hat{H}^0(G, C) \xrightarrow{\delta^0} \hat{H}^1(G, A) \to \hat{H}^1(G, B) \to \hat{H}^1(G, C) \to \hat{H}^2(G, A) \to \cdots$$

where, if $c \in C^G$ belongs to the class $[c] \in \hat{H}^0(G, C)$, then $\delta^0([c]) = [g \mapsto g\tilde{c} - \tilde{c}]$, where $\tilde{c}$ is an arbitrary lift of c to B.

Proof. We only have to check exactness at the first four groups, since from then onwards this is the usual long exact cohomology sequence. The checking is straightforward. $\square$

Remark 7.15. Tate also defined $\hat{H}^i(G, M)$ for $i < 0$. These groups see the homology of the G-module M and allow the exact sequence of the previous proposition to be continued infinitely to the left, accounting for the failure of the map $\hat{H}^0(G, A) \to \hat{H}^0(G, B)$ to be injective. Since we have not introduced homology in this course, we do not consider Tate cohomology in negative degree – but see the exercises!

Now let G be a finite cyclic group of order n, and let $\sigma \in G$ be a generator. Let M be a G-module and define a map $D : M \to M$ by $\mathrm{D}(m) = \sigma(m) - m$ for all $m \in M$. We define a complex $\mathcal{K}(M)$ as follows:

$$\cdots \xrightarrow{\partial^{-2}} M^{-1} \xrightarrow{\partial^{-1}} M^0 \xrightarrow{\partial^0} M^1 \xrightarrow{\partial^1} M^2 \xrightarrow{\partial^2} \cdots$$

where M^i is a copy of M for each $i \in \mathbb{Z}$, and the maps $\partial^i : M^i \to M^{i+1}$ are given by

$$\partial^i(m) = \begin{cases} \mathrm{D}(m) & : m \,\text{even} \\ \mathrm{N}(m) & : m \,\text{odd}. \end{cases}$$

for all $m \in M$. Note that this is indeed a complex, since for any $m \in M$ we have

$$\mathrm{N}(\mathrm{D}(m)) = \mathrm{D}(\mathrm{N}(m)) = (\sigma - 1)(1 + \sigma + \sigma^2 + \cdots + \sigma^{n-1})m = (\sigma^n - 1)m = 0.$$

Here we implicitly treat M as a module over the group ring $\mathbb{Z}[G]$. Let $H^i(\mathcal{K}(M)) = (\ker \partial^i)/(\operatorname{im} \partial^{i-1})$ denote the cohomology of this complex. It is manifestly clear that the abelian groups $H^i(\mathcal{K}(M))$ depend only on the parity of i.

Proposition 7.16. *Let G be a finite cyclic group and M a G-module. Given a generator $\sigma \in G$, let $\mathcal{K}(M)$ be the complex defined as above. Then $H^i(\mathcal{K}(M)) \simeq \hat{H}^i(G, M)$ for all $i \geq 0$. In particular, $H^i(G, M) \simeq H^{i+2}(G, M)$ for all $i \geq 1$.*

Proof. Since G is cyclic, we have $M^G = \{m \in M : \sigma(m) = m\}$. It follows immediately that $H^0(\mathcal{K}(M)) = M^G/N(M) = \hat{H}^0(G, M)$. Similarly, we have already observed in Example 3.18 that there is an isomorphism

$$\begin{array}{rcl} Z^1(G, M) & \xrightarrow{\sim} & \ker \mathrm{N} \\ \psi & \mapsto & \psi(\sigma). \end{array}$$

Under this isomorphism, $B^1(G, M)$ corresponds to the set of elements of M of the form $\sigma(m) - m$, namely to $\operatorname{im} \mathrm{D}$. Thus $\hat{H}^1(G, M) = H^1(G, M) \simeq H^1(\mathcal{K}(M))$.

It is easy to see that $H^i(\mathcal{K}(-)) : \mathrm{Mod}_G \to \mathrm{Ab}$ is a functor for every $i \geq 0$. Moreover, if $0 \to M \to N \to P \to 0$ is a short exact sequence of G-modules, then we clearly get a short exact sequence of complexes $0 \to \mathcal{K}(M) \to \mathcal{K}(N) \to \mathcal{K}(P)$, and by the Zigzag Lemma this produces a long exact cohomology sequence. Thus $\{H^i(\mathcal{K}(-))\}$ is a δ-functor. We claim that it is universal; this implies our proposition.

As usual, we prove universality by means of Lemma 3.9, so we need to show that the functors $H^i(\mathcal{K}(-))$ are effaceable for all $i \geq 1$. In the course of proving Corollary 2.12, we showed that any G-module M embeds in a module of the form $\mathrm{Ind}_{\{e\}}^G I$, where I is an injective abelian group. Let $f \in (\mathrm{Ind}_{\{e\}}^G I)^G$. This means that f is a constant function sending every $g \in G$ to some fixed element $\iota \in I$. Define $h_\iota \in \mathrm{Ind}_{\{e\}}^G I$ by

$$h_\iota(\sigma) = \begin{cases} \iota & : \sigma = e \\ 0 & : \sigma \neq e. \end{cases}$$

Then it is clear that $\mathrm{N}(f_\iota) = f$. Hence $\hat{H}^0(G, \mathrm{Ind}_{\{e\}}^G I) = H^0(\mathcal{K}(\mathrm{Ind}_{\{e\}}^G I)) = 0$, which in turn implies that $H^i(\mathcal{K}(\mathrm{Ind}_{\{e\}}^G I)) = 0$ for any even i.

Similarly, suppose that $f \in \mathrm{Ind}_{\{e\}}^G I$ lies in the kernel of N. This means that $\sum_{\tau \in G} f(\tau) = 0$. Thus the element $h \in \mathrm{Ind}_{\{e\}}^G I$ given by $h(\sigma^j) = \sigma_{k=0}^{j-1} f(\sigma^k)$ for all $j \geq 0$ is well-defined. One checks that $f(\tau) = h(\tau\sigma) - h(\tau)$ for all $\tau \in G$ and hence that $f \in \operatorname{im} \mathrm{D}$. It follows that $H^i(\mathcal{K}(\mathrm{Ind}_{\{e\}}^G I)) = 0$ for all odd i, and we have obtained the required effaceability. □

We now deduce two corollaries of the previous proposition. The first gives an alternative argument, avoiding the Hochschild-Serre spectral sequence, for the last step of the proof of Proposition 7.11 (the isomorphism $\mathrm{Br}(K) \simeq \mathbb{Q}/\mathbb{Z}$ for any finite extension $K/\mathbb{Q}_p$).

Corollary 7.17. *Let ℓ/k be an extension of finite fields. Then $H^i(\mathrm{Gal}(\ell/k), \ell^\times) = 0$ for any odd $i \geq 1$.*

Proof. If $i \geq 1$ is odd, then $H^i(\mathrm{Gal}(\ell/k), \ell^\times) \simeq H^1(\mathrm{Gal}(\ell/k), \ell^\times)$ by Proposition 7.16. Now $H^1(\mathrm{Gal}(\ell/k), \ell^\times) = 0$ by Hilbert 90. □

We can also now give the promised alternative proof of Lemma 7.7:

Corollary 7.18. *Let M be a finite $\widehat{\mathbb{Z}}$-module. Then $H^i(\widehat{\mathbb{Z}}, M) = 0$ for any even $i \geq 2$.*

Proof. By Proposition 7.16 it suffices to show that $\hat{H}^0(\widehat{\mathbb{Z}}, M) = 0$. It is an exercise to show that every open subgroup of $\widehat{\mathbb{Z}}$ is of the form $n\widehat{\mathbb{Z}}$ for $n \in \mathbb{N}$. Thus $\hat{H}^0(\widehat{\mathbb{Z}}, M) =$

$M^{\widehat{\mathbb{Z}}}/N(M) = \varinjlim M^{n\widehat{\mathbb{Z}}}/N(M^{n\widehat{\mathbb{Z}}})$. The modules $M^{n\widehat{\mathbb{Z}}}/N(M^{n\widehat{\mathbb{Z}}})$ need not be trivial. However, unraveling the definitions of Section 4 shows that the transition maps, when $n|m$, are the maps $\varinjlim M^{m\widehat{\mathbb{Z}}}/N(M^{m\widehat{\mathbb{Z}}}) \to M^{n\widehat{\mathbb{Z}}}/N(M^{n\widehat{\mathbb{Z}}})$ induced by multiplication by $\frac{m}{n}$ in M. In particular, this is the zero map whenever $\frac{m}{n} = |M|$, which implies our claim. □

Exercises

(1) Let ℓ/k be an extension of finite fields. Show that the norm map $\mathrm{N}_{\ell/k} : \ell \to k$ is surjective and use this to obtain another proof of Proposition 7.3.

(2) Prove, using direct computations with cocycles, that $H^2(\mathrm{Gal}(\mathbb{C}/\mathbb{R}), \mathbb{C}^\times) \simeq \mathbb{Z}/2\mathbb{Z}$.

(3) Prove that any central simple $\mathbb{R}$-algebra is isomorphic either to $\mathbb{R}$ itself or to Hamilton's quaternion algebra $\mathbb{H}$. Prove directly that $\mathrm{Br}(\mathbb{R}) \simeq \mathbb{Z}/2\mathbb{Z}$.

(4) Prove the claims of Remark 7.10.

(5) Let G be a profinite group, and let M be a G-module. Let $M' \subseteq M$ be the submodule generated by all elements of the form $gm - m$, for $g \in G$ and $m \in M$. Define the *G-coinvariants* of M to be $M_G = M/M'$. Prove that the functor $M \mapsto M_G$ is a right exact functor $\mathrm{Mod}_G \to \mathrm{Ab}$.

(6) A G-module P is called *projective* if, given a surjection $g : M \twoheadrightarrow N$ of G-modules and a map $f : P \to N$, there exists $\tilde{f} : P \to M$ completing the triangle:

$$\begin{tikzcd} & P \arrow[dl, dotted, "\tilde{f}"'] \arrow[d, "f"] \\ M \arrow[r, "g"] & N. \end{tikzcd}$$

This is dual to the notion of an injective G-module from Definition 2.9. Prove that every G-module M has a projective resolution, namely an exact sequence

$$\cdots \to P_2 \to P_1 \to P_0 \to M \to 0.$$

(7) Let M be a G-module. Choose a projective resolution as above and apply the G-coinvariants functor to get a complex

$$\cdots \to (P_2)_G \to (P_1)_G \to (P_0)_G.$$

The homology of this complex is called the homology of M and denoted $H_i(G, M)$. Prove that it is independent of the choice of projective resolution.

(8) Let G be a finite group and let M be a G-module. Observe that the map $N : M \to M$ of (17) above vanishes on M' and thus defines a map $N : M_G \to M^G$. Define Tate cohomology in negative degrees by

$$\hat{H}^i(G, M) = \begin{cases} \ker N \subseteq M_G & : i = -1 \\ H_{-(i+1)}(G, M) & : i \leq -2. \end{cases}$$

Given an exact sequence $0 \to M \to N \to C \to 0$ of G-modules, show that there is long exact sequence that extends infinitely in both directions:

$$\cdots \to \hat{H}^{-2}(G, C) \to \hat{H}^{-1}(G, M) \to \hat{H}^{-1}(G, N) \to \hat{H}^{-1}(G, C) \to \hat{H}^0(G, M) \to$$
$$\hat{H}^0(G, N) \to \hat{H}^0(G, C) \to \hat{H}^1(G, M) \to \cdots$$

(9) If G is finite cyclic and M is a G-module, prove that $\hat{H}^i(G, M) \simeq \hat{H}^{i+2}(G, M)$ for any $i \in \mathbb{Z}$.

8. Cohomological dimension

Our next big aim in this course is to understand the cohomology of local fields, namely to obtain an explicit description of $H^i(G_K, M)$, where $K/\mathbb{Q}_p$ is a finite extension and G_K, as usual, is the absolute Galois group of K. It turns out that essentially nothing of interest happens when $i > 2$. In this section we will prove this result and fit it into the more general framework of the theory of cohomological dimension.

8.1. **Basic properties.** If G is any abelian group and $n \geq 1$, then we denote the n-torsion of G by $G[n] = \{g \in G : ng = 0\}$. Of course, $G[n]$ is a subgroup of G. The following fact is basic but essential.

Lemma 8.1. *Let G be a profinite group and M a torsion G-module. Then $H^i(G, M)$ is a torsion group for any $i \geq 0$. If $M[n] = M$ for some $n \in \mathbb{N}$, then $H^i(G, M)[n] = H^i(G, M)$.*

Proof. This is obvious for $H^0(G, M) = M^G$, so we assume $i \geq 1$. Clearly it suffices to show that $Z^i(G, M)$ is a torsion group. Since G^i is compact, any continuous map $\varphi : G^i \to M$ has finite image. Let $\{m_1, \dots, m_r\}$ be the image of φ. Since M is torsion, there exist integers $n_1, \dots, n_r$ such that $n_i m_i = 0$ for all $1 \leq i \leq r$. Let $N = \mathrm{lcm}(n_1, \dots, n_r)$. Then $N\varphi = 0$.

In particular, if all the n_i divide some n, then so does N. This implies the final statement of the claim. □

Definition 8.2. Let G be a profinite group and p a prime. The *cohomological dimension of G at p*, denoted $\mathrm{cd}_p(G)$, is defined to be the largest i such that there exists a torsion G-module M satisfying $H^i(G, M)[p] \neq 0$. We say that $\mathrm{cd}_p(G) = \infty$ if the set of such i is unbounded.

The *cohomological dimension of G* is then defined to be

$$\mathrm{cd}(G) = \sup_p \mathrm{cd}_p(G).$$

Since $H^i(G, M)$ is a torsion group whenever M is torsion, by Lemma 8.1, we observe that $\mathrm{cd}(G)$ is, equivalently, the largest degree i for which there exists a torsion G-module M satisfying $H^i(G, M) \neq 0$.

Example 8.3. It is immediate from Lemma 3.12 that $\mathrm{cd}(\{e\}) = 0$.

Example 8.4. We have $\mathrm{cd}(G_k) = 1$ for any finite field k. Indeed, in the language of cohomological dimension Corollary 7.8 states precisely that $\mathrm{cd}(G_k) \leq 1$. In the discussion preceding the definition of the Hasse invariant map Inv_K, we saw that $H^1(G_k, \mathbb{Q}/\mathbb{Z}) \simeq \mathbb{Q}/\mathbb{Z}$, which implies $\mathrm{cd}(G_k) \geq 1$.

Definition 8.5. Let $K/\mathbb{Q}_p$ be a finite extension. For any $n \geq 1$, we set $\mu_n \subset \overline{K}^\times$ to be the subgroup of n-th roots of unity. This is naturally a discrete G-module, where G is any closed subgroup of G_K.

Example 8.6. Let $K/\mathbb{Q}_p$ be a finite extension. Then $\mathrm{cd}_\ell(G_K) \geq 2$ for all primes ℓ. Indeed, for any $n \geq 1$ we may consider the short exact sequence

$$0 \to \mu_n \to \overline{K}^\times \xrightarrow{x \mapsto x^n} \overline{K}^\times \to 0 \tag{18}$$

of G_K-modules. This gives rise to the following short exact sequence of cohomology groups:

$$H^1(G_K, \overline{K}^\times) \to H^2(G_K, \mu_n) \to H^2(G_K, \overline{K}^\times) \to H^2(G_K, \overline{K}^\times).$$

The leftmost group is trivial by Hilbert 90. Thus $H^2(G_K, \mu_n)$ is the kernel of $H^2(G_K, \overline{K}^\times) \to H^2(G_K, \overline{K}^\times)$. This map is just multiplication by n. Hence it follows from Corollary 6.37 and Proposition 7.11 that $H^2(G_K, \mu_n) \simeq \frac{1}{n}\mathbb{Z}/\mathbb{Z}$. Clearly μ_n is a torsion module, and if $\ell | n$ then $H^2(G_K, \mu_n)[\ell] \neq 0$.

The reader may well wonder whether the cohomological dimension of a profinite group G tells us anything at all about the cohomology of non-torsion modules. It turns out that it tells us quite a lot.

Definition 8.7. Let G be a profinite group and p a prime. The *strict cohomological dimension of G at p*, denoted $\operatorname{scd}_p(G)$, is defined to be the largest i such that there exists a G-module M satisfying $H^i(G, M)[p] \neq 0$. As usual, we say that $\operatorname{cd}_p(G) = \infty$ if the set of such i is unbounded. Similarly, we define $\operatorname{scd}(G) = \sup_p \operatorname{scd}_p(G)$.

Proposition 8.8. *Let G be a profinite group. Then* $\operatorname{cd}_p(G) \leq \operatorname{scd}_p(G) \leq \operatorname{cd}_p(G) + 1$ *for any prime p.*

Proof. The first inequality is obvious, so we only prove the second. Moreover, there is nothing to prove if $\operatorname{cd}_p(G) = \infty$, so we assume that $n = \operatorname{cd}_p(G)$ is finite. We need to show that $H^q(G, M)[p] = 0$ for all G-modules M whenever $q > n+1$. Consider the short exact sequence $0 \to M[p] \to M \xrightarrow{\pi} pM \to 0$ arising from multiplication by p. Since $M[p]$ is a torsion module, we have $H^q(G, M[p]) = H^q(G, M[p])[p] = 0$, and thus $H^q(G, M) \xrightarrow{\pi} H^q(G, pM)$ is injective.

On the other hand, consider the short exact sequence $0 \to pM \xrightarrow{\varepsilon} M \to M/pM \to 0$. Again M/pM is a p-torsion module, hence $H^{q-1}(G, M/pM) = H^{q-1}(G, M/pM)[p] = 0$, so that $H^q(G, pM) \xrightarrow{\varepsilon} H^q(G, M)$ is injective. Clearly the composition $\varepsilon \circ \pi : M \to M$ is just multiplication by p, and the same is true of the induced maps on cohomology. But if multiplication by p is injective on $H^q(G, M)$, then $H^q(G, M)[p] = 0$. ☐

Proposition 8.9. *Let G be a profinite group, and let $H \subseteq G$ be a closed normal subgroup. Then* $\operatorname{cd}(G) \leq \operatorname{cd}(H) + \operatorname{cd}(G/H)$.

Proof. Let M be a torsion G-module and let $i > \operatorname{cd}(H) + \operatorname{cd}(G/H)$. If $p, q \in \mathbb{N}$ are integers such that $p + q = i$, then either $p > \operatorname{cd}(H)$, in which case $H^p(H, \operatorname{Res}^G_H M) = 0$ and hence

$$H^q(G/H, H^p(H, \operatorname{Res}^G_H M)) = 0, \tag{19}$$

or else $q > \operatorname{cd}(G/H)$, in which case (19) still holds because $H^p(H, \operatorname{Res}^G_H M)$ is a torsion module by Lemma 8.1. Thus we have shown that the Hochschild-Serre spectral sequence of Definition 5.20 associated to the triple (G, H, M) satisfies $E_2^{p,q} = 0$ for all p, q such that $p + q = i$. Since $E_{r+1}^{p,q}$ is a subquotient of $E_r^{p,q}$ for any $r \geq 2$, it follows that $E_\infty^{p,q} = 0$ for all p, q such that $p, q = i$. By the abutment of the Hochschild-Serre spectral sequence, we have $F^pH^i(G, M) = 0$ for all $0 \leq p \leq i$, and hence $H^i(G, M) = 0$. ☐

8.2. **The cohomological dimension of inertia.** Let $K/\mathbb{Q}_p$ be a finite extension. We can now state precisely the claim to which we alluded at the beginning of this chapter. We would like to prove that $\operatorname{cd}(G_K) = 2$. We already know by Example 8.6 that $\operatorname{cd}(G_K) \geq 2$. Moveover, we know by Example 8.4 that $\operatorname{cd}(G_k) = 1$. Since $G_K/I_K \simeq G_k$, by Proposition 8.9 it suffices to show that $\mathrm{cd}(I_K) = 1$. The beginning of this section introduces notions and lemmas that could have been proved much earlier in the course, but we did not then have need for them.

Definition 8.10. Let p be a prime. A *pro-p group* is a profinite group whose finite quotients are all p-groups.

It follows from the proof of Theorem 1.10 that a profinite group G is pro-p if and only if $G \simeq \varprojlim_I G_i$, where the G_i are all finite p-groups. The following fun fact is foundational in mod p representation theory and accounts for much of its difference from representation theory over fields of characteristic zero.

Lemma 8.11. *Let G be a pro-p group and let M be a G-module. Suppose that M has the structure of an $\mathbb{F}_p$-vector space, or that M is finite and that its cardinality is a power of p. Then $M^G \neq \{0\}$.*

Proof. If M is an $\mathbb{F}_p$-vector space and $m \in M$ is a non-zero element, then, since $\mathrm{stab}_G(m)$ has finite index in G, the G-orbit of m spans a finite-dimensional subspace. Thus we may assume without loss of generality that M is finite of p-power cardinality. In this case, $U = \bigcap_{m \in M} \mathrm{stab}_G(m)$ is an open subgroup of G, hence of finite index. Moreover U is normal, since all the conjugates of $\mathrm{stab}_G(m)$ are stabilizers of translates of m. Thus the action of G on M factors through G/U, and we may assume without loss of generality that G is a finite p-group.

The finite group M decomposes into a disjoint union of G-orbits, whose cardinalities all divide $|G|$ and so are powers of p. Since $|M|$ is divisible by p, it follows that $|M^G|$, which is just the number of G-orbits of cardinality 1, must be divisible by p. In particular, $|M^G| > 1$. $\square$

Corollary 8.12. *Let G be a profinite group, let $H \trianglelefteq G$ be a pro-p normal subgroup, and let M be a G-module which has either an $\mathbb{F}_p$-vector space structure or finite p-power cardinality. Suppose that M is a simple G-module, in the sense that it has no proper non-trivial G-submodules. Then H acts trivially on M.*

Proof. Observe that M^H is a G-submodule of M. Indeed, for any $g \in G$, $h \in H$, and $m \in M^H$ we have $hgm = g(g^{-1}hg)m = gm$ and hence $gm \in M^H$. By the previous lemma, $M^H \neq \{0\}$. Therefore $M^H = M$. $\square$

Definition 8.13. Let G be a profinite group and let p be prime. A *pro-p-Sylow subgroup* of G is a maximal closed pro-p subgroup.

Remark 8.14. Suppose that $G \simeq \varprojlim_I G_i$, where $\{G_i\}_{i \in I}$ is a projective system of finite groups. Let $H \subseteq G$ be a pro-p-Sylow subgroup. It is a pleasant exercise to show that there exist p-Sylow subgroups $H_i \subseteq G_i$ for every $i \in I$ that are compatible, in the sense that $\varphi_{ij}(H_i) \subseteq H_j$ whenever $i, j \in I$ satisfy $i \geq j$, and that $H \simeq \varprojlim_I H_i$. It then follows from the usual Sylow theorems for finite groups that any two pro-p-Sylow subgroups of G are conjugate.

Lemma 8.15. *Let G be a profinite group, let p be prime, and let M be a G-module of finite p-power cardinality. Let $H \subseteq G$ be a pro-p-Sylow subgroup. Then for any $i \geq 0$ the restriction map* $\mathrm{res} : H^i(G, M) \to H^i(H, M)$ *is injective.*

Proof. As in the proof of Corollary 4.13 we can reduce to the case where G is finite. Indeed, we saw in that proof that $H^i(G, M) \simeq \varinjlim H^i(G/U_j, M^{U_j})$, where $\{U_j\}$ is the family of open normal subgroups of G. If H_j is the image of H under the natural projection $G \twoheadrightarrow G/U_j$, then $H_j \subseteq G/U_j$ is a p-Sylow subgroup. Moreover, $H \simeq \varprojlim H_j$, and the restriction map $\mathrm{res} :$ $H^i(G, M) \to H^i(H, M)$ arises from the maps $H^i(G/U_j, M^{U_j}) \overset{\mathrm{res}}{\to} H^i(H_j, M^{U_j}) \to H^i(H, M)$.

So let G be a finite group. Then any p-Sylow subgroup $H \subseteq G$ is open. Since $Z^i(G, M)$ is a finite group of exponent $|M|$, it has p-power order, and hence so does $H^i(G, M)$. By Lemma 4.4, the composition $H^i(G, M) \overset{\text{res}}{\to} H^i(H, M) \overset{\text{cor}}{\to} H^i(G, M)$ is multiplication by the index $[G : H]$. This is an isomorphism of $H^i(G, M)$, since $[G : H]$ is prime to p and hence to the order of $H^i(G, M)$. Thus the restriction map is injective, and the corestriction is surjective. □

Lemma 8.16. *Let K be a finite extension of $\mathbb{Q}_p$ and $n \geq 1$. Then $H^2(I_K, \mu_n) = 0$.*

Proof. Since $I_K = \mathrm{Gal}(\overline{K}/K^{\mathrm{nr}})$, we have $H^2(I_K, \overline{K}^\times) \simeq \mathrm{Br}(K^{\mathrm{nr}})$. Now, let A be a central simple K^{nr}-algebra. Choose a K^{nr}-basis $b_1, \ldots, b_r$ of A, and let $c_{ij}^\ell \in K^{\mathrm{nr}}$ be the structure constants satisfying $b_i b_j = \sum_{\ell=1}^r c_{ij}^\ell b_\ell$. Let L/K be the finite extension generated by the constants c_{ij}^ℓ, and let A' be the L-algebra spanned by the basis $b_1, \cdots, b_\ell$ and with multiplication defined as above. Clearly $A' \otimes_L K^{\mathrm{nr}} \simeq A$, and A' is simple: if $I \subset A'$ were a two-sided ideal, then $I \otimes_L K^{\mathrm{nr}} \subset A$ would also be one. Finally, the center of A' is $Z(A) \cap L = L$. Then A' is split by an unramified finite extension L'/L by Proposition 7.4. Thus $A \simeq A' \otimes_L K^{\mathrm{nr}} \simeq M_{\sqrt{r}}(K^{\mathrm{nr}})$, and so $\mathrm{Br}(K^{\mathrm{nr}}) = 0$.

Furthermore, $H^1(I_K, \overline{K}^\times) = 0$ by Hilbert 90 and Proposition 4.10. Considering the long exact cohomology sequence arising from the short exact sequence (18), whose terms are now viewed as I_K-modules, we see that $H^2(I_K, \mu_n)$ embeds in $H^2(I_K, \overline{K}^\times) = 0$. This establishes the claim. □

Corollary 8.17. *Let $K/\mathbb{Q}_p$ be a finite extension, and let M be a torsion I_K-module. Then $H^i(I_K, M) = 0$ for any $i \geq 2$.*

Proof. By Lemma 4.15 it suffices to show that $H^2(I_K, M)$ for any I_K-module M of finite cardinality. Since the abelian group M decomposes into a direct sum of groups of prime power order, since each of these direct summands is clearly preserved by the action of I_K, and since cohomology commutes with direct sums, we may assume without loss of generality that M has prime power order. So let $|M| = \ell^r$, where ℓ is a prime number.

If $H \subset I_K$ is a pro-ℓ-Sylow subgroup, then by Lemma 8.15 it suffices to prove that $H^2(H, M) = 0$. We show this by induction on r. If $r = 1$, then M has prime order and is thus a simple H-module. It follows from Corollary 8.12 that there is only one possible H-action on M, namely the trivial one. Thus $H^2(H, M) = H^2(H, \mu_\ell)$.

Let $F \subset \overline{K}$ be the fixed field of H. If L/K is a finite extension, then $[I_K : I_L] = [\mathrm{Gal}(\overline{K}/K^{\mathrm{nr}}) : \mathrm{Gal}(\overline{K} : L^{\mathrm{nr}})] = [L^{\mathrm{nr}} : K^{\mathrm{nr}}] = e_{L/K}$, where $e_{L/K}$ is the ramification index of L/K. In particular, if $\ell \nmid e_{L/K}$ then we must have $H \subseteq I_L$, since otherwise the image of H under the natural projection $I_K \twoheadrightarrow I_K/I_L$ would be a non-trivial subgroup of ℓ-power order. Thus

$$H \subseteq \bigcap_{L/K \text{ finite}, \ell \nmid e_{L/K}} I_L. \tag{20}$$

On the other hand, suppose that $x \in F$. Then $\ell \nmid e_{K(x)/K}$, hence F is contained in the compositum of the fields L^{nr}, where L runs over all finite extensions of K satisfying $\ell \nmid e_{L/K}$. This implies the inverse inclusion to (20). If we order these fields by inclusion, then the groups I_L naturally from a projective system whose connecting homomorphisms are inclusions. Then $H = \varprojlim I_L$, so that $H^2(H, \mu_\ell) = \varinjlim H^2(I_L, \mu_\ell)$. Hence $H^2(H, \mu_\ell) = 0$ by Lemma 8.16.

Now suppose the claim is known for $|M| \leq \ell^{r-1}$. Let $|M| = \ell^r$. Let $N \subseteq M$ be a simple H-submodule; this exists by the finiteness of M. By Lemma 8.11, the action of H on N is trivial, so that any subgroup of N is an H-submodule. Since any finite ℓ-group has a subgroup of order ℓ, it follows that $|N| = \ell$. Now the short exact sequence $0 \to N \to M \to M/N \to 0$ of H-modules gives rise to the exact sequence

$$H^2(H, N) \to H^2(H, M) \to H^2(H, M/N),$$

where the two outer groups are trivial by induction. Therefore $H^2(H, M) = 0$ and we are done.

Note that this induction would fail if we tried to work with I_K directly rather than with a pro-ℓ-Sylow subgroup, since then there could be simple G-modules of order ℓ^r for $r > 1$. □

8.3. **Finite G_K-modules.** We are now in a position to compute the cohomological dimension of G_K and to make an observation about the cohomology of G_K-modules that will lay the groundwork for local duality, which will be developed in the next section.

Proposition 8.18. *Let $K/\mathbb{Q}_p$ be a finite extension. Then $\mathrm{cd}(G_K) = 2$.*

Proof. As noted above, Example 8.6 shows that $\mathrm{cd}(G_K) \geq 2$, so it suffices to prove that $\mathrm{cd}(G_K) \leq 2$. Consider the closed normal subgroup $I_K \subset G_K$. We have $\mathrm{cd}(I_K) \leq 1$ by Corollary 8.17. Since $G_K/I_K \simeq G_k$ and $\mathrm{cd}(G_k) = 1$ by Example 8.4, we have $\mathrm{cd}(G_K) \leq 2$ by Proposition 8.9. □

Corollary 8.19. *Let $K/\mathbb{Q}_p$ be a finite extension, and let M be a G_K-module of finite cardinality. Then $H^i(G_K, M)$ is a finite group for all $i \geq 0$.*

Proof. Observe that $\bigcap_{m \in M} \mathrm{stab}_G(m)$ is a finite intersection of open subgroups of G_K, hence open. Similarly, there exists an open subgroup of G_K that acts trivially on $\mu_{|M|}$. Recall that $\{G_L\}$, where L/K runs over all finite Galois extensions, is a base of open neighborhoods of the identity of G_K. Therefore we can find a finite Galois extension L/K such that G_L is contained in the intersection of these two open subgroups and so acts trivially both on M and on $\mu_{|M|}$, and thus also on μ_n for any $n||M|$. By the structure theorem of abelian groups, we have $M \simeq \mu_{n_1} \oplus \cdots \oplus \mu_{n_r}$ as abelian groups for some integers n_j such that $n_1 n_2 \cdots n_r = |M|$. This is also an isomorphism of G_L-modules, since G_L acts trivially on both sides.

We claim that $H^i(G_L, M) \simeq \oplus_{j=1}^r H^i(G_L, \mu_{n_j})$ is finite for all $i \geq 0$. Of course it is enough to show that $H^i(G_L, \mu_n)$ is finite for all $i \geq 0$ and all $n \geq 1$. If $i \geq 3$, then $H^i(G_L, \mu_n) = 0$ by Proposition 8.18. Note also that $H^0(G_L, \mu_n) = \mu_n^{G_L} = \mu_n$ is finite, and that $H^2(G_L, \mu_n) \simeq \frac{1}{n}\mathbb{Z}/\mathbb{Z}$ by Example 8.6 (recall that this used all the machinery of the Brauer group), so in particular $H^2(G_L, \mu_n)$ is finite. Finally, viewing the short exact sequence (18) as a sequence of G_L-modules, we obtain the following bit of the long exact cohomology sequence:

$$H^0(G_L, \overline{K}^\times) \to H^0(G_L, \overline{K}^\times) \to H^1(G_L, \mu_n) \to H^1(G_L, \overline{K}^\times).$$

The rightmost group vanishes by Hilbert 90, and the leftmost map is raising to the power n. Hence $H^1(G_L, \mu_n) \simeq L^\times/(L^\times)^n$, and this will be shown to be finite in Proposition 8.25 below.

Thus $H^i(G_L, M)$ is indeed finite for all $i \geq 0$. Consider the Hochschild-Serre spectral sequence

$$E_2^{p,q} = H^p(\mathrm{Gal}(L/K), H^q(G_L, \mathrm{Res}_{G_L}^{G_K} M)) \Rightarrow H^{p+q}(G_K, M).$$

Observe that each $E_2^{p,q}$ is finite: since $\mathrm{Gal}(L/K)$ and $H^p(G_L, \mathrm{Res}^{G_K}_{G_L} M)$ are both finite, there are only finite many possible p-cocycles. Of course, this p clashes with the p of $K/\mathbb{Q}_p$, but it would be bizarre to index the modules $E_2^{p,q}$ by other parameters. Thus each $E_\infty^{p,q}$ is also finite, since it is isomorphic to a subquotient of $E_2^{p,q}$. Since the filtration of each $H^i(G_K, M)$ produces only finitely many graded pieces, each of which is isomorphic to $E_\infty^{i-q,q}$ for some $0 \leq q \leq i$, we see that $H^i(G_K, M)$ is finite for any $i \geq 0$, as claimed. □

One of the statements reached in the course of the proof of Corollary 8.19 is important in its own right; it encapsulates the results of Kummer theory. We state it as a proposition of its own.

Proposition 8.20. *Let $n \geq 1$, and let K be a field whose characteristic is prime to n. Then $H^1(G_K, \mu_n) \simeq K^\times/(K^\times)^n$.*

Proof. This was established in the proof of Corollary 8.19, where K was called L. While the setup there assumes that K is a p-adic field and that G_K acts trivially on μ_n, i.e. that $\mu_n \subset K^\times$, none of this figures in the proof of our statement. We only require that $K(\mu_n)$ be separable over K, in order to obtain a G_K-action on μ_n, and this is ensured by our condition on the characteristic of K. The diligent reader will notice that this statement appeared as an exercise a few chapters ago. □

Remark 8.21. While Proposition 8.20 is a very general statement, the proof of Corollary 8.19 made use of all of our work on the Brauer group of a finite extension $K/\mathbb{Q}_p$ and is thus specific to the p-adic setting. Indeed, Corollary 8.19 may fail for other fields. For instance, let $K = \mathbb{Q}$, and let $M = \mathbb{Z}/2\mathbb{Z}$, where the $G_\mathbb{Q}$-action on M is trivial. Then the non-trivial elements of $H^1(G_K, M) = \mathrm{Hom}(\mathrm{Gal}(\overline{\mathbb{Q}}/\mathbb{Q}), \mathbb{Z}/2\mathbb{Z})$ correspond to quadratic extensions of $\mathbb{Q}$, and there are infinitely many of these.

8.4. **Finiteness of $K^\times/(K^\times)^n$.** Before moving on to the next section, we prove an important statement about local fields that has already been used in the proof of Corollary 8.19 and will appear several more times in our development of local class field theory. If $K/\mathbb{Q}_p$ is a finite extension, write $\mathcal{O}_K$ for its valuation ring, $\mathfrak{m} \triangleleft \mathcal{O}_K$ for the maximal ideal, and $\pi \in \mathfrak{m}$ for a choice of uniformizer. As usual, k denotes the residue field; let q be its cardinality, and let $|\cdot|$ denote the normalized multiplicative valuation on K. For $\ell \in \mathbb{N}$, denote $U_\ell = 1 + \mathfrak{m}^\ell$.

Lemma 8.22. *Let $n \in \mathbb{N}$ be coprime to p, and let $\alpha \in U_1 = 1 + \mathfrak{m}$. There exists $\beta \in \mathcal{O}_K$ such that $\beta^n = \alpha$.*

Proof. This is immediate from (the usual formulation of) Hensel's Lemma. Consider the polynomial $f(x) = x^n - \alpha \in \mathcal{O}_K[x]$. Its reduction $\overline{f}(x) = x^n - 1 \in k[x]$ factors as $\overline{f}(x) = (x-1)\overline{g}(x)$, where $\overline{g}(x) \in k[x]$ is coprime to $x - 1$; indeed, 1 is not a root of $\overline{f}' = nx^{n-1}$, so it is not a multiple root of $\overline{f}$. Lifting this factorization of $\overline{f}$ to $\mathcal{O}_K$ produces an n-th root of α. □

We need an analogue of the previous lemma when n is divisible by p. One way to achieve this is by using the following variant of Hensel's Lemma.

Proposition 8.23 (The Hensel-Rychlik lemma)**.** *Let $\mathcal{O}$ be a complete discrete valuation ring, and let $f(x) \in \mathcal{O}[x]$. Suppose that $\gamma \in \mathcal{O}$ satisfies $|f(\gamma)| < |f'(\gamma)|^2$. Then there exists a unique $\beta \in \mathcal{O}$ such that $f(\beta) = 0$ and $|\beta - \gamma| < |f'(\gamma)|$. In fact, one has $|\beta - \gamma| = \frac{|f(\gamma)|}{|f'(\gamma)|}$.*

Proof. This can be proved using a variant of the iterative argument in the usual proof of Hensel's Lemma, that will be very reminiscent of Newton's method for finding roots of polynomials over the reals. Indeed, the claim is sometimes called the Hensel-Newton method in the literature. The idea is to define a recursive sequence by setting $a_0 = \gamma$ and $a_{n+1} = a_n - \frac{f(a_n)}{f'(a_n)}$. One proves by induction that each a_n is well-defined (i.e. that $f'(a_{n-1}) \neq 0$) and that $a_n \in \mathcal{O}$ for every $n \geq 0$. Moreover, the sequence $\{a_n\}$ converges (quite rapidly, in fact) to the β we want. □

Corollary 8.24. *Let $K/\mathbb{Q}_p$ be a finite extension, and let $n \in \mathbb{N}$. Suppose that $n = p^s n'$, where n' is coprime to p, and let $r \geq 0$. If $\alpha \in U_{2s+r+1}$, then there exists $\beta \in \mathcal{O}_K$ satisfying $\beta^n = \alpha$.*

Proof. Consider the polynomial $f(x) = x^n - \alpha \in \mathcal{O}_K[x]$, and take $\gamma = 1$. Then $|f(1)| = |1-\alpha| \leq q^{-(2s+r+1)}$, whereas $|f'(1)| = |n| = q^{-s}$. Thus the hypotheses of Proposition 8.23 are satisfied, and there exists $\beta \in \mathcal{O}_K$ such that $\beta^n = \alpha$ and $|\beta - 1| < q^{-(r+s+1)}$, so that $\beta \in U_{r+s+2}$. □

Proposition 8.25. *Let $K/\mathbb{Q}_p$ be a finite extension, and let $n \in \mathbb{N}$. Then $K^\times/(K^\times)^n$ and $\mathcal{O}_K^\times/(\mathcal{O}_K^\times)^n$ are finite groups.*

Proof. Observe that $K^\times$ decomposes as an internal direct product $K^\times \simeq \langle \pi \rangle \times \mathcal{O}_K^\times$, where $\mathcal{O}_K^\times = \varprojlim (\mathcal{O}_K/\mathfrak{m}^\ell)^\times$ is profinite and hence compact. We have shown that the endomorphism $\varphi : K^\times \to K^\times$ given by $\varphi(x) = x^n$ is an open mapping. Indeed, any open set $U \subset K^\times$ contains a coset of U_{r+s+2} for a sufficiently large r. Thus $\varphi(U)$ contains a coset of $\varphi(U_{r+s+2})$, but we have just shown that the open subgroup U_{2s+r+1} is contained in $\varphi(U_{r+s+2})$.

Now $(K^\times)^n = \varphi(K^\times) = \langle \pi^n \rangle \times \varphi(\mathcal{O}_K^\times)$, so that $K^\times/(K^\times)^n \simeq \mathbb{Z}/n\mathbb{Z} \times \mathcal{O}_K^\times/\varphi(\mathcal{O}_K^\times)$. Since $\varphi(\mathcal{O}_K^\times) = (\mathcal{O}_K^\times)^n$ is open, it has finite index in the compact group $\mathcal{O}_K^\times$, and we are done. □

Exercises

(1) Let G be a profinite group. Prove that $\mathrm{cd}(G) = 0$ if and only if $G = \{e\}$.

(2) This exercise outlines an alternative approach to Corollary 8.24. Recall the classical Taylor series

$$(1+x)^a = \sum_{i=0}^{\infty} \binom{a}{i} x^i = \sum_{i=0}^{\infty} \frac{a(a-1)\cdots(a-i+1)}{i!} x^i.$$

Let $K/\mathbb{Q}_p$ be a finite extension with uniformizer $\pi \in \mathcal{O}_K$, and let $\alpha = 1 + \pi^N y \in U_N$. Let $n \in \mathbb{N}$, and substitute into the Taylor series with $a = 1/n$ to obtain a series

$$\sum_{i=0}^{\infty} \frac{1 \cdot (1-n)(1-2n)\cdots(1-(i-1)n)}{n^i i!} \pi^{Ni} y^i \in K[[y]].$$

(a) Show that if N is sufficiently large, then the series converges for all $y \in \mathcal{O}_K$.
(b) For a fixed y, let β be the sum. Prove that $\beta^n = \alpha = 1 + \pi^N y$.
(c) What can you say about the valuation of $\beta - 1$?

9. The cup product and local duality

One can define a cup product in the context of group cohomology, analogous to that appearing in algebraic topology. This will turn out to be a very useful tool.

9.1. **The cup product.** Let G be a profinite group, and let M and N be two G-modules. Viewing their underlying abelian groups as $\mathbb{Z}$-modules, we form the tensor product $M \otimes_{\mathbb{Z}} N$; hereafter we will omit $\mathbb{Z}$ from the notation. Then $M \otimes N$ carries a G-action determined by $g(m \otimes n) = gm \otimes gn$ for any $g \in G$, $m \in M$, and $n \in N$. It is easy to see that the action is discrete, so that $M \otimes N$ is a G-module. There is a natural map $M^G \times N^G \to (M \otimes N)^G$ given by $(m, n) \mapsto m \otimes n$. The *cup product* will generalize this to a map

$$\begin{array}{rcl} H^i(G, M) \times H^j(G, N) & \to & H^{i+j}(G, M \otimes N) \\ (\varphi, \psi) & \mapsto & \varphi \cup \psi. \end{array}$$

Recall the spaces $\mathcal{C}^i(G, M)$ of Definition 3.5. These are spaces of continuous functions $G^{i+1} \to M$, with a G-action as defined there. Consider the map

$$\mathcal{C}^i(G, M) \times \mathcal{C}^j(G, N) \xrightarrow{\cup} \mathcal{C}^{i+j}(G, M \otimes N) \tag{21}$$

sending the pair (φ, ψ) to the function $\varphi \cup \psi \in \mathcal{C}^{i+j}(G, M \otimes N)$ given by

$$(\varphi \cup \psi)(g_0, \ldots, g_{i+j}) = \varphi(g_0, \ldots, g_i) \otimes \psi(g_i, \ldots, g_{i+j}).$$

Our first task is to check that the map of (21) does induce induce a map on cohomology. One verifies immediately that this map is G-equivariant, where G acts diagonally on the left-hand side. Thus G-invariants are preserved, and by the correspondence (4) we get an induced map $C^i(G, M) \times C^j(G, N) \xrightarrow{\cup} C^{i+j}(G, M \otimes N)$ of cochains.

Lemma 9.1. *Let M and N be G-modules, and suppose $i, j \geq 0$. Let $\varphi \in C^i(G, M)$ and $\psi \in C^j(G, N)$ be cochains. Then*

$$d_{i+j}(\varphi \cup \psi) = d_i\varphi \cup \psi + (-1)^i \varphi \cup d_j\psi.$$

Proof. Recall the maps $f_i : \mathcal{C}^i(G, M) \to \mathcal{C}^{i+1}(G, M)$ defined by (2); these correspond to the boundary maps d_i. It is straightforward to check that, for any $\varphi \in \mathcal{C}^i(G, M)$ and $\psi \in \mathcal{C}^j(G, N)$ and any $g_0, \ldots, g_{i+j+1} \in G$, we have

$$\begin{array}{rcl} (f_{i+j}(\varphi \cup \psi))(g_0, \ldots, g_{i+j+1}) & = & (f_i\varphi)(g_0, \ldots, g_{i+1}) \otimes \psi(g_{i+1}, \ldots, g_{i+j+1}) + \\ & & (-1)^i \varphi(g_0, \ldots, g_i) \otimes (f_j\psi)(g_i, \ldots, g_{i+j+1}). \end{array}$$

In other words, $f_{i+j}(\varphi \cup \psi) = f_i\varphi \cup \psi + (-1)^i \varphi \cup f_j\psi$. Since the correspondence (4) is linear, the claim follows immediately. □

Corollary 9.2. *Let M and N be G-modules. For any $i, j \geq 0$, the map of* (21) *induces a cup product map on cohomology:*

$$H^i(G, M) \times H^j(G, N) \xrightarrow{\cup} H^{i+j}(G, M \otimes N).$$

Proof. It is immediate from Lemma 9.1 that if $\varphi \in Z^i(G, M)$ and $\psi \in Z^j(G, N)$, then $d_{i+j}(\varphi \cup \psi) = 0$, so that $\varphi \cup \psi \in Z^{i+j}(G, M \otimes N)$. Moreover, if $\varphi \in B^i(G, M)$, then there exists $\eta \in C^{i-1}(G, M)$ such that $\varphi = d_{i-1}\eta$. The lemma then implies that

$$\varphi \cup \psi = d_{i-1}\eta \cup \psi = d_{i-1}\eta \cup \psi + (-1)^{i-1} \eta \cup d_j\psi = d_{i+j-1}(\eta \cup \psi) \in B^{i+j}(G, M \otimes N).$$

An analogous observation holds if ψ is a j-coboundary and φ is any i-cocyle. Hence the cup product induces a map on cohomology. □

We now observe several basic properties of the cup product. Most of them can be established by simple manipulation of cocycles.

Proposition 9.3. *Let M and N be G-modules.*

1. *The cup product $H^0(G,M) \times H^0(G,N) \xrightarrow{\cup} H^0(G, M \otimes N)$ is just the natural map $M^G \times N^G \to (M \otimes N)^G$ given by $(m,n) \mapsto m \otimes n$.*
2. *If $\varphi \in C^i(G,M)$ and $\psi \in C^j(G,N)$, then the inhomogeneous cochain $\varphi \cup \psi \in C^{i+j}(G, M \otimes N)$ is given explicitly as*
$$(\varphi \cup \psi)(g_1, \ldots, g_{i+j}) = \varphi(g_1, \ldots, g_i) \otimes g_1 g_2 \cdots g_i \psi(g_{i+1}, \ldots, g_{i+j}).$$
3. *If $m \in M^G = H^0(G,M)$ and $\psi \in Z^j(G,N)$, then $m \cup [\psi] = [m \otimes \psi] \in H^j(G, M \otimes N)$, where $m \otimes \psi \in Z^j(G, M \otimes N)$ is the cocycle $(g_1, \ldots, g_j) \mapsto m \otimes \psi(g_1, \ldots, g_j)$.*
4. *If $\varphi \in Z^1(G,M)$ and $\psi \in Z^1(G,N)$, then $[\varphi] \cup [\psi] = [\eta]$, where $\eta \in Z^2(G, M \otimes N)$ is the cocycle $\eta(g_1, g_2) = \varphi(g_1) \otimes g_1 \psi(g_2)$.*
5. *The cup product is associative: if P is a third G-module, then for any $\varphi \in H^i(G,M)$, $\psi \in H^j(G,N)$, and $\eta \in H^k(G,P)$ we have $(\varphi \cup \psi) \cup \eta = \varphi \cup (\psi \cup \eta)$ as elements of $H^{i+j+k}(G, M \otimes N \otimes P)$.*
6. *Let $H \subset G$ be a closed subgroup. If $\varphi \in H^i(G,M)$ and $\psi \in H^j(G,N)$, then*
$$\mathrm{res}^G_H \varphi \cup \mathrm{res}^G_H \psi = \mathrm{res}^G_H (\varphi \cup \psi).$$
If, in addition, H is normal and $\varphi \in H^i(G/H, M^H)$ and $\psi \in H^j(G/H, N^H)$, then
$$\inf \varphi \cup \inf \psi = \inf (\varphi \cup \psi).$$

Proof. The first claim is obvious from the definition of the cup product. The second follows from an easy computation using (4). The third and fourth claims are special cases of the second. The fifth and sixth claims follow from the second and the explicit descriptions of the restriction and inflation maps from Lemma 4.8. □

Lemma 9.4. *Let $0 \to M' \to M \to M'' \to 0$ be a short exact sequence of G-modules. Observe that the corresponding long exact sequence of cohomology groups produces a map $\delta^i : H^i(G, M'') \to H^{i+1}(G, M')$ for each $i \geq 0$. Let N be a G-module such that the sequence $0 \to M' \otimes N \to M \otimes N \to M'' \otimes N \to 0$ is also short exact. As above, we obtain maps $\delta^i : H^i(G, M'' \otimes N) \to H^{i+1}(G, M')$. Let $\varphi'' \in H^i(G, M'')$ and $\psi \in H^j(G,N)$ for some $i, j \geq 0$. Then*
$$(\delta^i \varphi'') \cup \psi = \delta^{i+j}(\varphi'' \cup \psi).$$

Proof. This is obtained from a computation on cocycles using Proposition 9.3(2) and the explicit description of the connecting maps in Lemma 4.7. □

Analogously we obtain the following statement.

Lemma 9.5. *Let $0 \to N' \to N \to N'' \to 0$ be a short exact sequence of G-modules. Let M be a G-module, and suppose that the sequence $0 \to M \otimes N' \to M \otimes N \to M \otimes N'' \to 0$ is also short exact. Let $\varphi \in H^i(G,M)$ and $\psi'' \in H^j(G, N'')$ for some $i, j \geq 0$. Then*
$$\varphi \cup (\delta^j \psi'') = (-1)^i \delta^{i+j}(\varphi \cup \psi''),$$
where the δ^i are connecting maps arising from the appropriate long exact sequences.

Lemma 9.6. *Let M and N be G-modules, and let $H \subset G$ be an open subgroup. Suppose that $\varphi \in H^i(H,M)$ and $\psi \in H^j(G,N)$. Then $(\mathrm{cor}\, \varphi) \cup \psi = \mathrm{cor}(\varphi \cup \mathrm{res}\, \psi) \in H^{i+j}(G, M \otimes N)$.*

Proof. If $(i,j) = (0,0)$, we can prove this directly using the explicit description of the corestriction map in Definition 4.3. Indeed,

$$(\operatorname{cor}\varphi) \cup \psi = \left(\sum_{gH \in G/H} g\varphi\right) \cup \psi = \sum_{gH\in G/H} (g\varphi \cup \psi) = \sum_{gH \in G/H} (g\varphi \cup g\psi) = \sum_{gH\in G/H} g(\varphi \cup \psi) = \operatorname{cor}(\varphi \cup \operatorname{res}\psi),$$

by the bilinearity of the cup product and since the restriction map $\operatorname{res} : H^0(G,N) \to H^0(H,N)$ is just the natural inclusion $N^G \hookrightarrow N^H$.

Now we can use a dimension shifting argument. Suppose that $i > 1$ and the claim is known for the pair $(i-1, j)$. Consider the short exact sequence

$$0 \to M \xrightarrow{\varepsilon} \operatorname{Ind}^G_{\{e\}} \operatorname{Res}^G_{\{e\}} M \to Q \to 0 \tag{22}$$

that appeared in the proof of Lemma 4.15. Since $\operatorname{Ind}^G_{\{e\}} \operatorname{Res}^G_{\{e\}} M$ is acyclic by Lemma 3.12, the boundary map $\delta^{i-1} : H^{i-1}(G,Q) \to H^i(G,M)$ is surjective. Let $\eta \in H^{i-1}(G,Q)$ be such that $\delta^{i-1}(\eta) = \varphi$. Since corestriction commutes with boundary maps by construction, we would like to conclude

$$\operatorname{cor}\varphi \cup \psi = \delta^{i-1}(\operatorname{cor}\eta) \cup \psi = \delta^{i+j-1}(\operatorname{cor}\eta \cup \psi) = \delta^{i+j-1}(\operatorname{cor}(\eta \cup \operatorname{res}\psi)) = \operatorname{cor}(\delta^{i+j-1}(\eta \cup \operatorname{res}\psi)) = \operatorname{cor}(\delta^{i-1}\eta \cup \operatorname{res}\psi) = \operatorname{cor}(\varphi \cup \operatorname{res}\psi),$$

where the third equality holds by our claim for $(i-1,j)$ and we would like to apply Lemma 9.4 to establish the second and fifth equalities. For this to be legitimate, we need to verify that the hypotheses of Lemma 9.4 hold, namely that $0 \to M \otimes N \to (\operatorname{Ind}^G_{\{e\}} \operatorname{Res}^G_{\{e\}} M) \otimes N \to Q \otimes N \to 0$ is a short exact sequence for an arbitrary G-module N. Since the functor $- \otimes N$ is always right exact (see, for instance, Proposition XVI.2.6 in Lang's *Algebra*, 3rd ed. or, better yet, prove it yourself), it suffices to show that the map $\varepsilon \otimes 1 : M \otimes N \to (\operatorname{Ind}^G_{\{e\}} \operatorname{Res}^G_{\{e\}} M) \otimes N$ is injective. It is an exercise to check that the following map is an isomorphism of G-modules:

$$\begin{array}{rcl} (\operatorname{Ind}^G_{\{e\}} \operatorname{Res}^G_{\{e\}} M) \otimes N & \xrightarrow{\sim} & \operatorname{Ind}^G_{\{e\}} \operatorname{Res}^G_{\{e\}} (M \otimes N) \\ f \otimes n & \mapsto & (g \mapsto f(g) \otimes gn). \end{array}$$

Recalling the definition of ε, it is easy to see that $\varepsilon \otimes 1$ sends the pure tensor $m \otimes n$ to the map $(g \mapsto gm \otimes gn)$, and thus $(\varepsilon \otimes 1)(v)(e) = v$ for any $v \in M \otimes N$. Hence $\varepsilon \otimes 1$ is injective.

Thus the claim holds for (i,j) if it holds for $(i-1,j)$. Similarly, suppose that the claim holds for $(i, j-1)$. By an argument analogous to the one above, using a short exact sequence $0 \to N \xrightarrow{\varepsilon} \operatorname{Ind}^G_{\{e\}} \operatorname{Res}^G_{\{e\}} N \to Q' \to 0$ and applying Lemma 9.5, we prove it for (i,j). This two-dimensional induction allows us to prove the claim for all pairs (i,j). $\square$

Dimension-shifting arguments as in the previous proof and, in a somewhat simpler version, in the proof of Lemma 4.15, are a very useful tool for working with cohomology. We now use a similar idea to show that the cup product is commutative up to sign. This property is called *anti-commutativity*.

Proposition 9.7. *Let M and N be G-modules. If $\varphi \in H^i(G,M)$ and $\psi \in H^j(G,N)$, then $\varphi \cup \psi = (-1)^{ij}(\psi \cup \varphi)$ as elements of $H^{i+j}(G, M \otimes N)$, where we identify $M \otimes N$ with $N \otimes M$*

by means of the isomorphism

$$\begin{array}{rcl} M \otimes N & \xrightarrow{\sim} & N \otimes M \\ m \otimes n & \mapsto & n \otimes m. \end{array}$$

Proof. If $(i,j) = (0,0)$, then $\varphi \cup \psi = \varphi \otimes \psi \in (M \otimes N)^G$, so the claim is obvious. We now use a dimension-shifting argument as in the proof of Lemma 9.6 to prove the proposition by induction on the pair (i,j). As in that proof, suppose that the claim is known for the pair $(i-1,j)$ and arbitrary G-modules, consider the short exact sequence (22), and choose $\eta \in H^{i-1}(G,Q)$ such that $\varphi = \delta^{i-1}(\eta)$. Then

$$\begin{array}{rcl} \varphi \cup \psi = \delta^{i-1}(\gamma) \cup \psi = \delta^{i+j-1}(\gamma \cup \psi) & = & (-1)^{(i-1)j} \delta^{i+j-1}(\psi \cup \gamma) = \\ & & (-1)^{(i-1)j}(-1)^j (\psi \cup \delta^{i-1}(\gamma)) = (-1)^{ij}(\psi \cup \varphi). \end{array}$$

Here the second equality is Lemma 9.4, the third comes from the inductive hypothesis, and the fourth is Lemma 9.5; we checked in the course of the proof of Lemma 9.6 that the hypotheses of Lemmas 9.4 and 9.5 hold. Hence we obtain the claim for (i,j). Similarly, if the claim is known for $(i,j-1)$, then it is obtained for (i,j) by an analogous argument that exchanges the roles of M and N. Since we already know the claim for $(0,0)$, these two results imply it for arbitrary pairs (i,j). □

9.2. **Statement of local duality.** The cup product enables us to state local duality, which is an analogue of the Poincaré duality arising in algebraic geometry, in the cohomology of a complex curve. We will see later that this theorem has wide-ranging implications.

Definition 9.8. Let $K/\mathbb{Q}_p$ be a finite extension.

1. Define $\mu_\infty = \bigcup_{n \geq 1} \mu_n \subset \overline{K}^\times$ to be the G_K-submodule consisting of all roots of unity. Note that $\mu_\infty \simeq \mathbb{Q}/\mathbb{Z}$ as a group.
2. Let M be a G_K-module. Its *dual module* is $M^* = \mathrm{Hom}(M, \mu_\infty)$ with the G_K-action defined by $(\sigma f)(m) = \sigma(f(\sigma^{-1} m))$ for all $\sigma \in G_K$, $f \in M^*$, and $m \in M$.
3. Let A be a finite abelian group. Its *Pontryagin dual* is $A^\vee = \mathrm{Hom}(A, \mathbb{Q}/\mathbb{Z})$.

Example 9.9. Let $M = \mu_n$. Since we view μ_n as embedded in μ_∞, any element $f \in M^*$ has the form $f(\zeta) = \zeta^i$ for all $\zeta \in \mu_n$ and some $i \in \mathbb{Z}/n\mathbb{Z}$. Moreover, since any $\sigma \in G_K$ acts on μ_n by $\sigma(\zeta) = \zeta^a$ for some $a \in (\mathbb{Z}/n\mathbb{Z})^\times$, we find that $(\sigma f)(\zeta) = \sigma(f(\sigma^{-1}(\zeta))) = \sigma(f(\zeta^b)) = \sigma(\zeta^{ib}) = \zeta^{aib} = \zeta^i = f(\zeta)$ for all $\zeta \in \mu_n$, where $ab \equiv 1 \bmod n$. Thus $M^* = \mathbb{Z}/n\mathbb{Z}$, with trivial G_K-action.

Lemma 9.10. *If A is a finite abelian group, then $A^\vee \simeq A$. Moreover, the evaluation map gives a* canonical *isomorphism $(A^\vee)^\vee \simeq A$.*

If M is a finite G_K-module, then $M \simeq (M^*)^*$ canonically as G_K-modules.

Proof. The first claim is immediate from the fact that any finite abelian group A is a direct product of finite cyclic groups. Now let M be a finite G_K-module. The evaluation map provides a homomorphism

$$\begin{array}{rcl} M & \to & (M^*)^* \\ m & \mapsto & (f \mapsto f(m)), \end{array}$$

where $m \in M$ and $f \in M^*$. It follows from the definition of the G_K-action on dual modules that this map is G_K-equivariant. Since M is a direct product of finite cyclic groups, for any distinct elements $m, m' \in M$ it is easy to construct $f \in M^*$ such that $f(m) \neq f(m')$.

Thus the map above is injective. But $M \simeq M^* \simeq (M^*)^*$ as abelian groups, so it must also be surjective. The same argument, ignoring the G_K-module structure, shows that the evaluation map provides an isomorphism of A with $(A^\vee)^\vee$ if A is a finite abelian group. $\square$

Definition 9.11. Let A be any topological abelian group. Write $A^\vee$ for the group of continuous homomorphisms $\mathrm{Hom}(A, \mathbb{Q}/\mathbb{Z})$, where $\mathbb{Q}/\mathbb{Z}$ has the discrete topology. We endow $A^\vee$ with the compact-open topology; this is the coarsest topology such that $\{f \in A^\vee : f(K) \subset U\}$ is open for all compact $K \subset A$ and all open (i.e. all) $U \subset \mathbb{Q}/\mathbb{Z}$. Observe that this is consistent with our previous definition of $A^\vee$ in the case of finite A.

Proposition 9.12. *Let A be a topological abelian group such that the quotients A/nA are finite for all $n \in \mathbb{N}$ and any continuous homomorphism $f : A \to \mathbb{Q}/\mathbb{Z}$ factors through A/nA for some $n \in \mathbb{N}$. Then $(A^\vee)^\vee \simeq \varprojlim A/nA$ as topological abelian groups.*

Proof. We have $A^\vee = \mathrm{Hom}(A, \mathbb{Q}/\mathbb{Z}) = \varinjlim \mathrm{Hom}(A/nA, \mathbb{Q}/\mathbb{Z})$ by assumption. Thus

$$(A^\vee)^\vee = \mathrm{Hom}(\varinjlim \mathrm{Hom}(A/nA, \mathbb{Q}/\mathbb{Z}), \mathbb{Q}/\mathbb{Z}) = \\ \varprojlim \mathrm{Hom}(\mathrm{Hom}(A/nA, \mathbb{Q}/\mathbb{Z}), \mathbb{Q}/\mathbb{Z}) = \varprojlim((A/nA)^\vee)^\vee = \varprojlim A/nA,$$

where the last equality uses the finiteness of the A/nA and Lemma 9.10. We leave it as an exercise to check that this is a homeomorphism. $\square$

Proposition 9.13. *If A is a profinite abelian group, then the evaluation map gives an isomorphism $A \xrightarrow{\sim} (A^\vee)^\vee$.*

Proof. If the quotients A/nA are finite, then this follows easily from the previous proposition. This is the case, for instance, when $A = \mathbb{Z}_p$ or, more generally, when A is the additive group $\mathcal{O}_K$ or the multiplicative group $\mathcal{O}_K^\times$ for any finite extension $K/\mathbb{Q}_p$; recall Proposition 8.25.

More generally, it is an exercise to show that if $A = \varprojlim A_i$, then $A^\vee = \varinjlim A_i^\vee$ and, conversely, if $A = \varinjlim A_i$, then $A = \varprojlim A_i^\vee$, with the obvious transition maps. Applying both of these observations, we see that if A is a profinite group realized as $A = \varprojlim A_i$ with finite A_i, then $(A^\vee)^\vee = \varprojlim((A_i)^\vee)^\vee$. The evaluation map induces an isomorphism on each A_i by Lemma 9.10, and the claim follows. $\square$

Proposition 9.14. *Let TorMod_{G_K} be the category of torsion G_K-modules. Then $M \mapsto M^*$ is an exact contravariant functor from TorMod_G to itself.*

Proof. The only non-obvious part of the statement is the exactness. This is an exercise; note that it uses the fact that $\mu_\infty \simeq \mathbb{Q}/\mathbb{Z}$ is injective as an abelian group. $\square$

For any G_K-module M, the evaluation map $M \times M^* \to \mu_\infty$ given by $(m, f) \mapsto f(m)$ is clearly bilinear over $\mathbb{Z}$, and thus it factors through a map $\alpha : M \otimes M^* \to \mu_\infty$. Moreover, α is G_K-equivariant. Indeed, for all $\sigma \in G_K$ we have

$$\alpha(\sigma(m \otimes f)) = \alpha(\sigma m \otimes \sigma f) = (\sigma f)(\sigma m) = \sigma(f(\sigma^{-1}\sigma m)) = \sigma(f(m)) = \sigma(\alpha(m \otimes f)).$$

Thus α induces a map $\alpha^* : H^i(G_K, M \otimes M^*) \to H^i(G_K, \mu_\infty)$ on cohomology for all $i \geq 0$.

Definition 9.15. Let $K/\mathbb{Q}_p$ be a finite extension and let M be a G_K-module. For each $i \in \{0, 1, 2\}$ we define the pairing $\langle , \rangle_K : H^i(G_K, M) \times H^{2-i}(G_K, M^*) \to \mathbb{Q}/\mathbb{Z}$ as the composition

$$H^i(G_K, M) \times H^{2-i}(G_K, M^*) \xrightarrow{\cup} H^2(G_K, M \otimes M^*) \xrightarrow{\alpha^*} H^2(G_K, \mu_\infty) \xrightarrow{\mathrm{Inv}_K} \mathbb{Q}/\mathbb{Z}.$$

The image under this pairing of a pair $(\varphi, \psi) \in H^i(G_K, M) \times H^{2-i}(G_K, M^*)$ is denoted $\langle \varphi, \psi \rangle_K$.

Now suppose that M is a finite G_K-module. Clearly M^* is also finite, and the groups $H^i(G_K, M)$ and $H^{2-i}(G_K, M^*)$ are finite by Corollary 8.19. The pairing defined above gives rise to maps

$$A : H^{2-i}(G_K, M^*) \to H^i(G_K, M)^\vee \qquad B : H^i(G_K, M) \to H^{2-i}(G_K, M^*)^\vee$$
$$\psi \mapsto (\varphi \mapsto \langle \varphi, \psi \rangle_K) \qquad \varphi \mapsto (\psi \mapsto \langle \varphi, \psi \rangle_K).$$

The pairing $\langle , \rangle_K$ is called a *perfect pairing* if the maps A and B are isomorphisms for each $i \in \{0, 1, 2\}$.

The rest of this section is devoted to proving the following theorem, which is the local duality advertised above.

Theorem 9.16. *Let $K/\mathbb{Q}_p$ be a finite extension and let M be a finite G_K-module. Then $\langle , \rangle_K$ is a perfect pairing for every $i \in \{0, 1, 2\}$.*

9.3. **Reduction to $i = 2$.** Although Theorem 9.16 consists a priori of three entirely separate statements, for the three possible values of i, we shall see that the case $i = 2$ implies the other two without much difficulty. Throughout this section, $K/\mathbb{Q}_p$ is a finite extension, and M is a finite G_K-module.

The following statement makes sense by Lemma 9.10.

Lemma 9.17. *The pairing $\langle , \rangle_K : H^i(G_K, M) \times H^{2-i}(G_K, M^*) \to \mathbb{Q}/\mathbb{Z}$ of Definition 9.15 satisfies*

$$\langle b, a \rangle_K = \begin{cases} \langle a, b \rangle_K & : i \in \{0, 2\} \\ -\langle a, b \rangle_K & : i = 1, \end{cases}$$

for all $a \in H^i(G_K, M)$ and $b \in H^{2-i}(G_K, M^*)$.

Proof. This follows from Proposition 9.7, which gives the analogous commutativity properties for the cup product. Note that $(-1)^{i(2-i)}$ is equal to 1 if $i \in \{0, 2\}$ and to -1 if $i = 1$. □

Corollary 9.18. *If Theorem 9.16 holds for $i = 2$, then it holds for $i = 0$.*

Proof. This follows from the first part of Lemma 9.17, with M^* in the role of M. □

Proposition 9.19. *If Theorem 9.16 holds for $i = 2$, then it holds for $i = 1$.*

Proof. Let M be a finite G_K-module. As an abelian group, M is a direct sum of finitely many, say r, finite cyclic groups, so it embeds into the divisible abelian group $I = (\mathbb{Q}/\mathbb{Z})^r$. As in the proof of Corollary 2.12, by Frobenius reciprocity we obtain an embedding of G_K-modules $M \hookrightarrow \mathrm{Ind}_{\{e\}}^{G_K} I$. We identify M with a submodule of $\mathrm{Ind}_{\{e\}}^{G_K} I$ by means of this embedding. Since I is a torsion group, and any $f \in \mathrm{Ind}_{\{e\}}^{G_K} I$ is locally constant and thus takes only finitely many values since G_K is compact, we conclude that $\mathrm{Ind}_{\{e\}}^{G_K} I$ is a torsion G_K-module; see the proofs of Lemma 4.15 and Lemma 8.1 for similar arguments. Moreover, since every element of a torsion module is contained in a finite submodule (again, more details in the proof of Lemma 4.15), we see that $\mathrm{Ind}_{\{e\}}^{G_K} I \simeq \varinjlim N_\alpha$, where the N_α are finite G_K-modules such that $M \subset N_\alpha$.

By Proposition 4.10 we have $\varinjlim H^i(G_K, N_\alpha) = H^i(G_K, \mathrm{Ind}_{\{e\}}^{G_K} I) = 0$ for every $i > 0$, where the last equality is because $\mathrm{Ind}_{\{e\}}^{G_K} I$ is an injective G_K-module by Corollary 2.11. For every $i > 0$, since $H^i(G_K, M)$ is finite by Corollary 8.19, it follows that there exists a finite submodule $M \subset N_i$ such that $H^i(G_K, M) \to H^i(G_K, N_i)$ is the zero map; indeed, every element $\varphi \in H^i(G_K, M)$ has zero image in some $H^i(G_K, N_{\alpha_\varphi})$, and since there are finitely many elements we can take N_i to be an upper bound of all the N_{α_φ}. Take a finite submodule $M \subset N \subset \mathrm{Ind}_{\{e\}}^{G_K} I$ that is an upper bound of N_1 and N_2. Since $H^i(G_K, N) = 0$ for all $i \geq 3$ by Proposition 8.18, we conclude that $H^i(G_K, M) \to H^i(G_K, N)$ is the zero map for all $i \geq 1$.

The short exact sequence $0 \to M \to N \to Q \to 0$ of G_K-modules, where $Q = N/M$, gives rise to a long exact sequence of cohomology groups containing the segment $H^0(G_K, N) \to H^0(G_K, Q) \to H^1(G_K, M)$.

The sequence $0 \to Q^* \to N^* \to M^* \to 0$ of duals is also an exact sequence of G_K-modules by Proposition 9.14. Hence we get an exact sequence $H^1(G_K, M^*) \to H^2(G_K, Q^*) \to H^2(G_K, N^*)$. Taking Pontryagin duals, we get a diagram

$$\begin{tikzcd}
H^0(G_K,N) \arrow[r] \arrow[d, "B"', "\sim"] & H^0(G_K,Q) \arrow[r] \arrow[d, "B"', "\sim"] & H^1(G_K,M) \arrow[r, "0"] \arrow[d, "B"'] & H^1(G_K,N) \\
H^2(G_K,N^*)^\vee \arrow[r] & H^2(G_K,Q^*)^\vee \arrow[r] & H^1(G_K,M^*)^\vee,
\end{tikzcd}$$

where the rows are exact. This diagram is commutative, as one sees by studying the effect of the maps on cochains. Since Theorem 9.16 holds for $i = 2$ by hypothesis and hence the leftmost and central vertical maps are isomorphisms by Corollary 9.18, we verify by a diagram chase that the rightmost vertical map is an injection.

By the same argument with M^* instead of M, we find that $A : H^1(G_K, M^*) \to H^1(G_K, M)^\vee$ is injective. If $f : \Gamma \to \Delta$ is an injective homomorphism of abelian groups, the corresponding map $f^\vee : \Delta^\vee \to \Gamma^\vee$ is surjective by the injectivity of $\mathbb{Q}/\mathbb{Z}$. Hence $B = A^\vee : H^1(G_K, M) \to H^1(G_K, M^*)$ is surjective, since $H^1(G_K, M)$ is finite by Corollary 8.19 and hence $(H^1(G_K, M)^\vee)^\vee \simeq H^1(G_K, M)$ by Lemma 9.10. We conclude that B is an isomorphism. Similarly, A is an isomorphism, and thus the local pairing is perfect in the case $i = 1$. $\square$

9.4. **Reduction to a finite Galois extension.** By the results of the previous section, it suffices to prove Theorem 9.16 in the case $i = 2$. In this section, we will show that it suffices to prove the theorem after replacing K by any finite Galois extension L/K. We will then choose L wisely in the next section to complete the proof.

We will need to understand the behavior of the Hasse invariant map, which figures in the definition of local duality, with regards to field extensions.

Proposition 9.20. *Let K be a finite extension of $\mathbb{Q}_p$, and let L/K be a finite extension. Both squares in the diagram below commute:*

$$\begin{tikzcd}
\mathrm{Br}(K) \simeq H^2(G_K, \overline{K}^\times) \arrow[r, "\mathrm{Inv}_K"] \arrow[d, shift left, "\mathrm{res}"] & \mathbb{Q}/\mathbb{Z} \arrow[d, shift left, "{[L:K]}"] \\
\mathrm{Br}(L) \simeq H^2(G_L, \overline{K}^\times) \arrow[r, "\mathrm{Inv}_L"] \arrow[u, shift left, "\mathrm{cor}"] & \mathbb{Q}/\mathbb{Z}. \arrow[u, shift left, "\mathrm{id}"]
\end{tikzcd}$$

Proof. The proof is just a matter of unpacking the definition of the Hasse invariant map. First, denoting the residue fields of L and K by ℓ and k, respectively, we will show that the following diagrams commute:

$$\begin{tikzcd} H^2(G_K, \overline{K}^\times) \arrow[d, shift left, "\mathrm{res}"] & H^2(G_k, (K^{\mathrm{nr}})^\times) \arrow[l, "\sim"', "\mathrm{inf}"] \arrow[d, shift left, "\mathrm{res}"] \\ H^2(G_L, \overline{K}^\times) \arrow[u, shift left, "\mathrm{cor}"] & H^2(G_\ell, (L^{\mathrm{nr}})^\times). \arrow[l, "\sim"', "\mathrm{inf}"] \arrow[u, shift left, "e_{L/K}\mathrm{cor}"] \end{tikzcd} \tag{23}$$

Recall that the horizontal inflation maps are known to be isomorphisms by Corollary 7.5. Our claim can be checked by a direct computation on cocycles using the formulas of Section 4. Alternatively, since all the maps in this diagram can be defined using the universality of group cohomology as a δ-functor, it suffices to verify the commutativity of the corresponding diagrams for H^0:

$$\begin{tikzcd} K^\times = (\overline{K}^\times)^{G_K} \arrow[d, shift left, "\mathrm{id}"] & ((K^{\mathrm{nr}})^\times)^{G_k} = K^\times \arrow[l, "\mathrm{id}"'] \arrow[d, shift left, "\mathrm{id}"] \\ L^\times = (\overline{K}^\times)^{G_L} \arrow[u, shift left, "\mathrm{tr}_{L/K}"] & ((L^{\mathrm{nr}})^\times)^{G_\ell} = L^\times. \arrow[l, "\mathrm{id}"'] \arrow[u, shift left, "e_{L/K}\mathrm{tr}_{\ell/k}"] \end{tikzcd}$$

Here $\mathrm{tr}_{\ell/k} : L^\times \to K^\times$ denotes the map $\mathrm{tr}_{\ell/k}(x) = \sum_{\sigma G_\ell \in G_k/G_\ell} \sigma(x)$. and this is obvious. By a similar universality argument, and the compatibility of δ-functors with the connecting maps arising in long exact cohomology sequences, we get the commutation of the diagrams

$$\begin{tikzcd} H^2(G_k, (K^{\mathrm{nr}})^\times) \arrow[r, "v_K"] \arrow[d, shift left, "\mathrm{res}"] & H^2(G_k, \mathbb{Z}) \arrow[r, "\delta"] \arrow[d, shift left, "e_{L/K}\mathrm{res}"] & H^1(G_k, \mathbb{Q}/\mathbb{Z}) \arrow[d, shift left, "e_{L/K}\mathrm{res}"] \\ H^2(G_\ell, (L^{\mathrm{nr}})^\times) \arrow[r, "v_L"] \arrow[u, shift left, "e_{L/K}\mathrm{cor}"] & H^2(G_\ell, \mathbb{Z}) \arrow[r, "\delta"] \arrow[u, shift left, "\mathrm{cor}"] & H^1(G_\ell, \mathbb{Q}/\mathbb{Z}). \arrow[u, shift left, "\mathrm{cor}"] \end{tikzcd}$$

Finally, noting that $\mathrm{Frob}_L = \mathrm{Frob}_K^{f_{L/K}}$, we find that the following commutes:

$$\begin{tikzcd} H^1(G_k, \mathbb{Q}/\mathbb{Z}) \arrow[r] \arrow[d, shift left, "e_{L/K}\mathrm{res}"] & \mathbb{Q}/\mathbb{Z} \arrow[d, shift left, "e_{L/K}f_{L/K} = {[L:K]}"] \\ H^1(G_\ell, \mathbb{Q}/\mathbb{Z}) \arrow[r] \arrow[u, shift left, "\mathrm{cor}"] & \mathbb{Q}/\mathbb{Z}. \arrow[u, shift left, "\beta"] \end{tikzcd}$$

Since the composition of the two vertical maps on the right must be multiplication by $e_{L/K}[G_k : G_\ell] = [L : K]$ by Lemma 4.4, the map β is forced to be the identity. Concatenating all these diagrams and taking the perimeter, we obtain the commutative diagrams of the claim. $\square$

Lemma 9.21. *Let M be a finite G_K-module, and let L/K be a finite extension. The corestriction map* $\mathrm{cor} : H^2(G_L, M) \to H^2(G_K, M)$ *is surjective.*

Proof. Consider the map $\mathrm{tr} : \mathrm{Ind}_{G_L}^{G_K}\mathrm{Res}_{G_L}^{G_K}M \to M$ given by $\mathrm{tr}(f) = \sum_{gG_L \in G_K/G_L} gf(g^{-1})$. It is easy to see that this is well-defined, and it is G_K-equivariant since for all $x \in G_K$ and all $f \in \mathrm{Ind}_{G_L}^{G_K}\mathrm{Res}_{G_L}^{G_K}M$ we have

$$\mathrm{tr}(xf) = \sum_{gG_L \in G_K/G_L} g \cdot xf(g^{-1}) =$$
$$\sum_{gG_L \in G_K/G_L} gf(g^{-1}x) = x\left(\sum_{gG_L \in G_K/G_L} x^{-1}gf((x^{-1}g)^{-1})\right) = x(\mathrm{tr}(f)).$$

Finally, tr is surjective: if $m \in M$, let $f_m \in \mathrm{Ind}_{G_L}^{G_K}\mathrm{Res}_{G_L}^{G_K}M$ be the function defined by $f_m(g) = gm$ if $g \in G_L$ and $f_m(g) = 0$ otherwise; note that f_m is locally constant. Then $\mathrm{tr}(f_m) = m$. Now consider the exact sequence

$$0 \to N \to \mathrm{Ind}_{G_L}^{G_K}\mathrm{Res}_{G_L}^{G_K}M \xrightarrow{\mathrm{tr}} M \to 0, \tag{24}$$

where N is the appropriate kernel.

The induced map $\mathrm{tr}^* : H^2(G_K, \mathrm{Ind}_{G_L}^{G_K}\mathrm{Res}_{G_L}^{G_K}M) \to H^2(G_K, M)$ is surjective, as we see by observing the bit

$$H^2(G_K, \mathrm{Ind}_{G_L}^{G_K}\mathrm{Res}_{G_L}^{G_K}M) \xrightarrow{\mathrm{tr}^*} H^2(G_K, M) \to H^3(G_K, N)$$

of the long exact sequence and noting that $H^3(G_K, N) = 0$ because $\mathrm{cd}(G_K) = 2$ and N is a torsion module since $\mathrm{Ind}_{G_L}^{G_K}\mathrm{Res}_{G_L}^{G_K}M$ is.

The corestriction map in the statement of the lemma is just the composition

$$H^2(G_L, \mathrm{Res}_{G_L}^{G_K}M) \simeq H^2(G_K, \mathrm{Ind}_{G_L}^{G_K}\mathrm{Res}_{G_L}^{G_K}M) \xrightarrow{\mathrm{tr}^*} H^2(G_K, M),$$

where the isomorphism comes from Shapiro's Lemma. Thus the corestriction is a composition of two surjective maps, so it is surjective. □

Lemma 9.22. *Let M be a finite G_K-module, and let L/K be a finite Galois extension. If Theorem 9.16 holds for the G_L-module $\mathrm{Res}_{G_L}^{G_K}M$ when $i = 2$, then it holds for the G_K-module M when $i = 2$.*

Proof. The diagram

$$\begin{tikzcd}
H^0(G_K, M^*) = \mathrm{Hom}_{G_K}(M, \mu_\infty) \arrow[r, "A"] \arrow[d] & H^2(G_K, M)^\vee \arrow[d, "\mathrm{cor}^\vee"] \\
H^0(G_L, (\mathrm{Res}_{G_L}^{G_K}M)^*) = \mathrm{Hom}_{G_L}(M, \mu_\infty) \arrow[r, "A"] & H^2(G_L, \mathrm{Res}_{G_L}^{G_K}M)^\vee
\end{tikzcd}$$

where the left vertical map is the natural inclusion, commutes as a consequence of Proposition 9.20. Since the right vertical map is injective by the previous lemma, we see that the top horizontal map is also injective. Since M carries an action of G_K, the two groups in the bottom row carry a natural action of $G_K/G_L = \mathrm{Gal}(L/K)$. The map A in the bottom row commutes with this action (this needs to be justified). Thus we get a diagram

$$\begin{tikzcd}
\mathrm{Hom}_{G_K}(M, \mu_\infty) \arrow[r, "A"] \arrow[d] & H^2(G_K, M)^\vee \arrow[d, "\mathrm{cor}^\vee"'] \\
\mathrm{Hom}_{G_K}(M, \mu_\infty)^{\mathrm{Gal}(L/K)} \arrow[r, "A"] & (H^2(G_K, M)^\vee)^{\mathrm{Gal}(L/K)}.
\end{tikzcd}$$

But now the left vertical map is an isomorphism (in fact, an equality). The bottom horizontal map is assumed to be an isomorphism, so the composition $\mathrm{cor}^\vee \circ A$ is an isomorphism, and this forces the top horizontal map to be an isomorphism. $\square$

Now we finally complete the proof of Theorem 9.16.

9.5. **Properties of the pairing.** We now prove several properties of the perfect pairing. These are all consequences of the definitions and of the functoriality of the Hasse invariant map.

Lemma 9.23. *Let L/K be a finite extension and $i \in \{0, 1, 2\}$. Let M be a finite G_K-module. Let $\varphi \in H^i(G_K, M)$ and $\psi \in H^{2-i}(G_L, M^*)$. Then $\langle \varphi, \mathrm{cor}\,\psi\rangle_K = \langle \mathrm{res}\,\varphi, \psi\rangle_L$.*

Proof. Recall the definition of the local duality pairing as the composition

$$\langle\, , \,\rangle_K : H^i(G_K, M) \times H^{2-i}(G_K, M^*) \xrightarrow{\cup} H^2(G_K, M \otimes M^*) \to H^2(G_K, \mu_\infty) \xrightarrow{\mathrm{Inv}_K} \mathbb{Q}/\mathbb{Z}.$$

We know that $\mathrm{Inv}_L = \mathrm{Inv}_K \circ \mathrm{cor}$ by Proposition 9.20. Thus $\langle \varphi, \mathrm{cor}\,\psi\rangle_K = \mathrm{Inv}_K(\varphi \cup \mathrm{cor}\,\psi) = \mathrm{Inv}_K(\mathrm{cor}(\mathrm{res}\,\varphi \cup \psi)) = \mathrm{Inv}_L(\mathrm{res}\,\varphi, \psi) = \langle \mathrm{res}\,\varphi, \psi\rangle_L$, where we have used Lemma 9.6 and Proposition 9.7. $\square$

Lemma 9.24. *Let L/K be a finite extension and $i \in \{0, 1, 2\}$. Let M be a finite G_K-module. Let $\varphi \in H^i(G_K, M)$ and $\psi \in H^{2-i}(G_K, M^*)$. Then $\langle \mathrm{res}\,\varphi, \mathrm{res}\,\psi\rangle_L = [L : K]\langle \varphi, \psi\rangle_K$.*

Proof. We showed that $\mathrm{Inv}_L \circ \mathrm{res} = [L : K]\,\mathrm{Inv}_K$ in Proposition 9.20. Thus $\langle \mathrm{res}\,\varphi, \mathrm{res}\,\psi\rangle_L = \mathrm{Inv}_L(\mathrm{res}\,\varphi \cup \mathrm{res}\,\psi) = \mathrm{Inv}_L(\mathrm{res}(\varphi \cup \psi)) = [L : K]\,\mathrm{Inv}_K(\varphi \cup \psi) = [L : K]\,\langle \varphi, \psi\rangle_K$, where the second equality is Proposition 9.3(5). $\square$

Exercises

(1) Let G be a topological group. The usual definition of the Pontryagin dual of G is $G^\vee = \mathrm{Hom}(G, \mathbb{R}/\mathbb{Z})$, where $\mathbb{R}/\mathbb{Z}$ is the circle group with its usual topology, and $G^\vee$ is endowed with the compact-open topology.
  (a) Show that a subgroup $H \subset \mathbb{R}/\mathbb{Z}$ is closed if and only if $H = \mathbb{R}/\mathbb{Z}$ or H is finite.
  (b) Prove that $G^\vee$ is discrete if G is compact, and that $G^\vee$ is compact if G is discrete.
  (c) Show that if G is profinite, then $\mathrm{Hom}(G, \mathbb{R}/\mathbb{Z}) = \mathrm{Hom}(G, \mathbb{Q}/\mathbb{Z})$.

10. The Artin reciprocity map

We are now in a position to prove the main results of local class field theory. Let $K/\mathbb{Q}_p$ be a finite extension as in the previous sections. The unit group $K^\times$ has a profinite, hence compact, subgroup $\mathcal{O}_K^\times = \varprojlim_n (\mathcal{O}_K/\mathfrak{m}_K^n)^\times$, where $\mathfrak{m}_K \lhd \mathcal{O}_K$ is the maximal ideal. Consider the topology on $K^\times$, where a base of open neighborhoods of the identity is that from the topology of $\mathcal{O}_K^\times$. This makes $K^\times$ into a locally compact group.

Lemma 10.1. *Let $f : K^\times \to \mathbb{Q}/\mathbb{Z}$ be a continuous group homomorphism, where $\mathbb{Q}/\mathbb{Z}$ is endowed with the discrete topology. Then f has finite image.*

Proof. The restriction of f to the compact subgroup $\mathcal{O}_K^\times$ has finite image; since $\mathbb{Q}/\mathbb{Z}$ is discrete, the fibers of f are open, and finitely many of them must cover $\mathcal{O}_K^\times$. If $\pi \in \mathcal{O}_K$ is a uniformizer, then $K^\times$ is generated by $\mathcal{O}_K^\times$ and $\langle \pi \rangle$. This $f(K^\times)$ is generated by the finite subgroup $f(\mathcal{O}_K^\times)$ and the finite-order element $f(\pi) \in \mathbb{Q}/\mathbb{Z}$, so it is finite. □

If G is a profinite group, its commutator subgroup $[G, G]$ need not in general be closed. We define the abelianization of G to be $G^{\mathrm{ab}} = G/\overline{[G, G]}$, where $\overline{[G, G]}$ is the closure of $[G, G]$. Then G^{ab} is profinite by Proposition 1.14. In fact, G^{ab} is the maximal abelian profinite homomorphic image of G. If $K/\mathbb{Q}_p$ is a finite extension, then the commutator subgroup $[G_K, G_K]$ is actually a closed subgroup of G_K (see exercises), but we will not need this fact below. If K^{ab} is the maximal abelian subextension of $\overline{K}/K$, then $G_K^{\mathrm{ab}} = \mathrm{Gal}(K^{\mathrm{ab}}/K)$.

Proposition 10.2. *There is an isomorphism $\mathrm{Hom}(G_K^{\mathrm{ab}}, \mathbb{Q}/\mathbb{Z}) \xrightarrow{\sim} \mathrm{Hom}(K^\times, \mathbb{Q}/\mathbb{Z})$ of abelian groups, where in both cases we consider continuous homomorphisms.*

Proof. By the previous lemma, every $f \in \mathrm{Hom}(K^\times, \mathbb{Q}/\mathbb{Z})$ has finite image and thus $\ker f$ has finite index in $K^\times$. In particular, $(K^\times)^n \subseteq \ker f$, where $n = [K^\times : \ker f]$, so that f factors through $K^\times/(K^\times)^n$. Thus we have

$$\mathrm{Hom}(K^\times, \mathbb{Q}/\mathbb{Z}) = \varinjlim \mathrm{Hom}(K^\times/(K^\times)^n, \mathbb{Q}/\mathbb{Z}) \simeq \varinjlim \mathrm{Hom}(H^1(G_K, \mu_n), \mathbb{Q}/\mathbb{Z}), \tag{25}$$

where the final isomorphism was obtained in Proposition 8.20 ("Kummer theory") and is induced by the boundary map $\delta^0 : H^0(G_K, \overline{K}^\times) \to H^1(G_K, \mu_n)$ arising from the short exact sequence $0 \to \mu_n \to \overline{K}^\times \xrightarrow{n} \overline{K}^\times \to 0$ of G_K-modules.

Recall from Example 9.9 that $(\mu_n)^* = \frac{1}{n}\mathbb{Z}/\mathbb{Z}$, with trivial G_K-action. Thus the perfect pairing of Theorem 9.16, in the case $i = 1$, tells us exactly that

$$\mathrm{Hom}(H^1(G_K, \mu_n), \mathbb{Q}/\mathbb{Z}) = H^1(G_K, \mu_n)^\vee = H^1(G_K, \frac{1}{n}\mathbb{Z}/\mathbb{Z}). \tag{26}$$

But $H^1(G_K, \frac{1}{n}\mathbb{Z}/\mathbb{Z}) = \mathrm{Hom}(G_K, \frac{1}{n}\mathbb{Z}/\mathbb{Z})$ since the G_K-action is trivial, so, putting together (25) and (26), we have

$$\mathrm{Hom}(K^\times, \mathbb{Q}/\mathbb{Z}) \simeq \varinjlim \mathrm{Hom}(G_K, \frac{1}{n}\mathbb{Z}/\mathbb{Z}) = \mathrm{Hom}(G_K, \mathbb{Q}/\mathbb{Z}) = \mathrm{Hom}(G_K^{\mathrm{ab}}, \mathbb{Q}/\mathbb{Z}),$$

where the final equality holds since any homomorphism from G_K to the abelian group $\mathbb{Q}/\mathbb{Z}$ must factor through G_K^{ab}, as its kernel is closed and contains the commutator subgroup $[G_K, G_K]$. □

Consider the group

$$\widehat{K^\times} = \varprojlim_n K^\times/(K^\times)^n.$$

Theorem 10.3. *There is an isomorphism $r_K : \widehat{K^\times} \to G_K^{\mathrm{ab}}$.*

Proof. We apply the functor $\mathrm{Hom}(-, \mathbb{Q}/\mathbb{Z})$ to the isomorphism of Proposition 10.2. By Proposition 9.12, whose hypotheses are satisfied in the case of $A = K^\times$ by Proposition 8.25 and Lemma 10.1, we have that $\mathrm{Hom}(K^\times, \mathbb{Q}/\mathbb{Z})^\vee = ((K^\times)^\vee)^\vee \simeq \widehat{K^\times}$. Since G_K^{ab} is profinite, we know that $\mathrm{Hom}(G_K^{\mathrm{ab}}, \mathbb{Q}/\mathbb{Z})^\vee \simeq G_K^{\mathrm{ab}}$ canonically by Proposition 9.13. □

The isomorphism r_K is called the Artin reciprocity map. Its definition may be restated pictorially as follows.

Corollary 10.4. *There exists a unique injection $r_K : K^* \to G_K^{\mathrm{ab}}$ with dense image such that the following diagram commutes for all $n \in \mathbb{N}$:*

$$\begin{tikzcd}
\mathrm{Hom}(K^\times, \mathbb{Z}/n\mathbb{Z}) \times & K^\times/(K^\times)^n \arrow[r, "\mathrm{eval}"] & \mathbb{Z}/n\mathbb{Z} \arrow[d, "\times\frac{1}{n}"] \\
H^1(G_K, \mathbb{Z}/n\mathbb{Z}) \arrow[u] \times & H^1(G_K, \mu_n) \arrow[u, "\sim"'] \arrow[r, "\langle\, ,\,\rangle_K"] & \mathbb{Q}/\mathbb{Z}.
\end{tikzcd}$$

Here the top horizontal map is the evaluation map, and the bottom horizontal map is local duality, which makes sense by Example 9.9. The middle vertical map is the inverse of the isomorphism $d : K^\times/(K^\times)^n \xrightarrow{\sim} H^1(G_K, \mu_n)$ of Kummer theory (Proposition 8.20), and the leftmost vertical map is the homomorphism

$$H^1(G_K, \mathbb{Z}/n\mathbb{Z}) = \mathrm{Hom}(G_K, \mathbb{Z}/n\mathbb{Z}) = \mathrm{Hom}(G_K^{\mathrm{ab}}, \mathbb{Z}/n\mathbb{Z}) \xrightarrow{\varphi \mapsto \varphi \circ r_K} \mathrm{Hom}(K^\times, \mathbb{Z}/n\mathbb{Z}).$$

Proposition 10.5. *Let L/K be a finite extension. Then the following diagram commutes,*

$$\begin{tikzcd}
L^\times \arrow[r, "r_L"] \arrow[d, "N_{L/K}"'] & G_L^{\mathrm{ab}} \arrow[d] \\
K^\times \arrow[r, "r_K"] & G_K^{\mathrm{ab}}
\end{tikzcd}$$

where $N_{L/K} : L^\times \to K^\times$ is the norm map and $G_L^{\mathrm{ab}} \hookrightarrow G_K^{\mathrm{ab}}$ is the natural embedding.

Proof. It suffices to check for all $b \in L^*$ and all continuous homomorphisms $\chi : G_K^{\mathrm{ab}} \to \mathbb{Q}/\mathbb{Z}$ that $\chi(r_K(N_{L/K}(b))) = \chi(r_L(b))$ holds. Indeed, if $\sigma, \tau \in G_K^{\mathrm{ab}}$ are distinct elements, then there exists a finite abelian extension K'/K such that the projections of σ and τ in the finite abelian group $\mathrm{Gal}(K'/K)$ remain distinct. But we have already seen that there exists a homomorphism $\chi' : \mathrm{Gal}(K'/K) \to \mathbb{Q}/\mathbb{Z}$ sending σ and τ to distinct images, as will the map $\chi : G_K^{\mathrm{ab}} \to \mathbb{Q}/\mathbb{Z}$ obtained from χ' by inflation.

Since G_K^{ab} is profinite and hence compact, the image of any χ as above is finite and thus contained in $\frac{1}{n}\mathbb{Z}/\mathbb{Z}$ for some $n \in \mathbb{N}$, which we identify with $\mathbb{Z}/n\mathbb{Z}$ by means of the "multiplication by n" map.

If $a \in K^*$ and $\chi \in \mathrm{Hom}(G_K^{\mathrm{ab}}, \mathbb{Z}/n\mathbb{Z}) = H^1(G_K, \mathbb{Z}/n\mathbb{Z})$, then the diagram of Corollary 10.4 says precisely that $\frac{1}{n}\chi(r_K(a)) = \langle \chi, da \rangle_K$. Now $\langle \chi, da \rangle_K = \mathrm{Inv}_K(\chi \cup da)$ by the definition of the local duality pairing,

Consider the corestriction map $\mathrm{cor} : L^\times = H^0(G_L, \mathrm{res}^{G_K}_{G_L} \overline{K}^\times) \to H^0(G_K, \overline{K}^\times) = K^\times$. Given $b \in L^\times$, by Definition 4.3 we have $\mathrm{cor}\, b = \sum_{\sigma G_L \in G_K/G_L} \sigma(b) = N_{L/K}(b)$. The Kummer theory isomorphism $d : K^\times/(K^\times)^n \xrightarrow{\sim} H^1(G_K, \mu_n)$ arises from a boundary map $H^0(G_K, \overline{K}^\times) \to H^1(G_K, \mu_n)$ and thus commutes with corestriction. We will also write $d : L^\times/(L^\times)^n \xrightarrow{\sim} H^1(G_L, \mu_n)$.

Putting all this together, and taking $a = N_{L/K}(b) \in K^\times$ two paragraphs above, we have

$$\frac{1}{n}\chi(r_K(N_{L/K}(b))) = \mathrm{Inv}_K(\chi \cup d(\mathrm{cor}\, b)) = \mathrm{Inv}_K(\chi \cup \mathrm{cor}\, db) = \mathrm{Inv}_K(\mathrm{cor}(\mathrm{res}\,\chi \cup db)) =$$
$$\mathrm{Inv}_L(\mathrm{res}\,\chi \cup db) = \frac{1}{n}\chi(r_L(b)),$$

where the third equality comes from Lemma 9.6 (and the anti-commutativity of Proposition 9.7), the fourth one holds as $\mathrm{Inv}_L = \mathrm{Inv}_K \circ \mathrm{cor}$ by Proposition 9.20, and the last equality uses Corollary 10.4 again. ☐

Proposition 10.6. *Let L/K be a finite abelian extension. Then r_K induces an isomorphism between $K^\times/N_{L/K}(L^\times)$ and $\mathrm{Gal}(L/K)$.*

Proof. Consider the following diagram:

$$\begin{tikzcd}
0 \arrow[r] & L^\times \arrow[r, "N_{L/K}"] \arrow[d, "r_L"] & K^\times \arrow[r] \arrow[d, "r_K"] & K^\times/N_{L/K}(L^\times) \arrow[r] \arrow[d, "r"] & 0 \\
0 \arrow[r] & G_L^{\mathrm{ab}} \arrow[r] & G_K^{\mathrm{ab}} \arrow[r] & \mathrm{Gal}(L/K) \arrow[r] & 0
\end{tikzcd}$$

The top row is obviously exact, and the bottom one is exact since L/K is abelian and thus a Galois subextension of K^{ab}/K. The maps r_L and r_K are injective and have dense images, and the square on the left commutes by the previous proposition, thus inducing the vertical map r on the right. We need to show that r is an isomorphism.

Observe that, since the r_K here is a restriction to $K^\times$ of the map $r_K : \widehat{K^\times} \to G_K^{\mathrm{ab}}$ obtained in Theorem 10.3, the relevant topology on $K^\times$ is the subspace topology induced from the profinite group $\widehat{K^\times} = \varprojlim K^\times/(K^\times)^n$.

To show that r is injective, suppose that $x \in \ker r \subset K^\times/N_{L/K}(L^\times)$ and let $y \in K^\times$ be a lift of x. A very simple diagram chase shows that $r_K(y)$ must lie in $G_L^{\mathrm{ab}} \subset G_K^{\mathrm{ab}}$. Since the image of r_L is dense in G_L^{ab}, there exists a sequence $\{y_n\}$ of elements of $L^\times$ such that $\{r_L(y_n)\}$ converges to $r_K(y)$.[4] By the injectivity of r_L and r_K we find that $\{N_{L/K}(y_n)\}$ converges to y. If we knew that $N_{L/K}(L^\times)$ were closed in $K^\times$, we could conclude that $y \in N_{L/K}(L^\times)$ and hence $x = 0$.

It remains to show that $N_{L/K}(L^\times) \subset K^\times$ is closed in the topology induced from $\widehat{K^\times}$. So let $\{b_n\}$ be a sequence of elements of $L^\times$ such that $\{N_{L/K}(b_n)\}$ converges to $a \in K^\times$ in the relevant topology. We need to show that $a \in N_{L/K}(L^\times)$. For any $b \in L^\times$, recall

[4] This is true because G_L^{ab} is a first countable topological space, i.e. every point has a countable base of open neighborhoods. It suffices to show that the identity has a countable base of open neighborhoods, since translates of these will give a base of open neighborhoods for any other point. A base of open neighborhoods of the identity is given by the subgroups $\mathrm{Gal}(L^{\mathrm{ab}}/L')$, where L'/L is a finite abelian extension. Since L has finitely many extensions of any fixed finite degree, there are countably many of these.

Here is a sketch of a proof that L has finitely many extensions of any fixed degree. Since L has a unique unramified extension of any degree, it suffices to show that there are finitely many totally ramified extensions of L of degree n. The minimal polynomial of a uniformizer in any such extension L' is an Eisenstein polynomial $x^n + a_{n-1}x^{n-1} + \cdots + a_1 x + a_0 \in \mathcal{O}_L[x]$, where $a_i \in \mathfrak{m}_L$ for all i and $a_0 \notin \mathfrak{m}_L^2$. By Krasner's Lemma, perturbing each coefficient a_i within some open set will give rise to the same extension L'. By compactness of $\mathcal{O}_L$ we conclude that there are only finitely many possibilities for L'.

that $v_K(N_{L/K}(b)) = f_{L/K}v_L(b)$, where $f_{L/K} = [\ell : k]$ is the inertia degree. Indeed, all the conjugates of b have the same valuation, so that $v_L(N_{L/K}(b)) = [L : K]v_L(b)$. But $(v_L)_{|K^\times} = e_{L/K}v_K$. Since the sequence $\{v_K(N_{L/K}(b_n))\}$ converges to $v_K(a)$, we have that $v_K(a) \in f_{L/K}\mathbb{Z}$. Now choose $c \in L^\times$ such that $v_L(c) = v_K(a)/f_{L/K}$. Then $N_{L/K}(b_n/c) \to a/N_{L/K}(c) \in \mathcal{O}_K^\times$. Clearly a lies in the image of $N_{L/K}$ if and only if $a/N_{L/K}(c)$ does, so we may replace a with $a/N_{L/K}(c)$ and assume without loss of generality that $a \in \mathcal{O}_K^\times$.

Now, note that $K^\times \simeq \langle \pi_K \rangle \times \mathcal{O}_K^\times$ naturally, where $\pi_K \in \mathcal{O}_K$ is a uniformizer. Thus $K^\times/(K^\times)^n \simeq \langle \pi_K \rangle / \langle \pi_K^n \rangle \times \mathcal{O}_K^\times/(\mathcal{O}_K^\times)^n$. Taking projective limits, we get $\widehat{K^\times} \simeq \widehat{\langle \pi_K \rangle} \times \mathcal{O}_K^\times \simeq \widehat{\mathbb{Z}} \times \mathcal{O}_K^\times$. Indeed, we showed at the end of the proof of Proposition 8.25 that the groups $\mathcal{O}_K^\times/(\mathcal{O}_K^\times)^n$ are all finite, and hence that $\mathcal{O}_K^\times$ is their projective limit. The point of this is that a base of open neighborhoods of the identity of $K^\times$ in the $\widehat{K^\times}$-topology is provided by sets of the form $\langle \pi^N \rangle \times U$, where $N \geq 0$ and $U \subset \mathcal{O}_K^\times$ is open in the usual profinite topology of $\mathcal{O}_K^\times$.

Since the sequence of integers $\{v_L(b_n)\}$ must stabilize at 0, it follows that for all open neighborhoods $e \in U \subset \mathcal{O}_K^\times$ we have $N_{L/K}(b_n) \in Ua$ for all sufficiently large n. In other words, $N_{L/K}(b_n) \to a$ in the usual topology of $\mathcal{O}_K^\times$. But $\mathcal{O}_L^\times$, in its usual topology, is compact, and thus the sequence $\{b_n\}$ has a limit point b; again, throwing away finitely many b_n's we may assume that all the b_n's lie in $\mathcal{O}_L^\times$. We must have $N_{L/K}(b) = a$, so that $a \in N_{L/K}(L^\times)$. We may finally conclude that the map r is injective.

To prove surjectivity of r, observe that $K^\times/N_{L/K}(L^\times)$ is compact. Indeed, it is clear that $\pi_K^{[L:K]} = N_{L/K}(\pi_K) \in N_{L/K}(L^\times)$. Thus $K^\times/N_{L/K}(L^\times)$ is a quotient of $K^\times/\langle \pi_K^{[L:K]} \rangle$, and this in turn is a disjoint union $K^\times/\langle \pi_K^{[L:K]} \rangle = \coprod_{i=0}^{[L:K]-1} \pi_K^i \mathcal{O}_K^\times$ of finitely many compact sets and hence is compact.

Now let $x \in \mathrm{Gal}(L/K)$ and lift it to $z \in G_K^{\mathrm{ab}}$. By the density of the image of r_K, there is a sequence $\{y_n\}$ of elements of $K^\times$ such that $\{r_K(y_n)\}$ converges to z. By the compactness we just proved, the sequence of images of the y_n in $K^\times/N_{L/K}(L^\times)$ has a convergent subsequence. Let $y \in K^\times/N_{L/K}(L^\times)$ be the limit. By the commutativity of the square on the right, we must have $r(y) = x$. Hence r is surjective. $\square$

Lemma 10.7. *Let L/K be a finite extension. There exists a continuous homomorphism $V : G_K \to G_L$ with the following property: for any $n \in \mathbb{N}$ and any $\chi \in \mathrm{Hom}(G_L, \mathbb{Z}/n\mathbb{Z}) = H^1(G_L, \mathbb{Z}/n\mathbb{Z})$, we have $\operatorname{cor} \chi = \chi \circ V \in H^1(G_K, \mathbb{Z}/n\mathbb{Z})$.*

Proof. This is a special case of a general construction; see the exercises. $\square$

Corollary 10.8. *Let L/K be a finite extension, Consider the natural inclusion $\iota : K^\times \hookrightarrow L^\times$. Then the diagram*

$$\begin{tikzcd}
K^\times \arrow[r, "r_K"] \arrow[d, "\iota"'] & G_K^{\mathrm{ab}} \arrow[d, "V"] \\
L^\times \arrow[r, "r_L"] & G_L^{\mathrm{ab}}
\end{tikzcd}$$

commutes, where $V : G_K^{\mathrm{ab}} \to G_L^{\mathrm{ab}}$ is the map from Lemma 10.7.

Proof. Since r_K has dense image and $r_L \circ \iota$ is injective, observe that a map $G_K^{\mathrm{ab}} \to G_L^{\mathrm{ab}}$ completing the diagram is unique if it exists.

As in the proof of Proposition 10.5, it suffices to show that $\chi(V(r_K(a))) = \chi(r_L(\iota(a)))$ for all $a \in K^\times$ and $\chi \in \mathrm{Hom}(G_L^{\mathrm{ab}}, \mathbb{Z}/n\mathbb{Z}) = \mathrm{Hom}(G_L, \mathbb{Z}/n\mathbb{Z}) = H^1(G_L, \mathbb{Z}/n\mathbb{Z})$ for any $n \in \mathbb{N}$. To simplify the notation, we drop ι and view $K^\times$ as a subgroup of $L^\times$. Then, as in the proof of Proposition 10.5, we have

$$\chi(r_L(a)) = n\,\mathrm{Inv}_L(\chi \cup da) = n\,\mathrm{Inv}_K(\mathrm{cor}(\chi \cup da)) = n\,\mathrm{Inv}_K((\mathrm{cor}\,\chi) \cup da) = \\ n\,\mathrm{Inv}_K((\chi \circ V) \cup da) = \chi(V(r_K(a))),$$

where the first and last equalities come from Corollary 10.4, the second is from Proposition 9.20, the third is from Lemma 9.6, and the fourth holds by Lemma 10.7, which we just proved. □

Lemma 10.9. *Consider the short exact sequence* $0 \to \mathbb{Z} \to \mathbb{Q} \to \mathbb{Q}/\mathbb{Z} \to 0$ *of* G_K*-modules, all with the trivial action, where* $K/\mathbb{Q}_p$ *is a finite extension as usual. This gives rise to a boundary map* $\delta^1 : H^1(G_K, \mathbb{Q}/\mathbb{Z}) \to H^2(G_K, \mathbb{Z})$.

Now let $\chi \in \mathrm{Hom}(G_K, \mathbb{Q}/\mathbb{Z}) = H^1(G_K, \mathbb{Q}/\mathbb{Z})$ *and* $a \in K^\times = H^0(G_K, \overline{K}^\times)$. *Consider the element* $a \cup \delta^1(\chi) \in H^2(G_K, \overline{K}^\times \otimes \mathbb{Z}) = H^2(G_K, \overline{K}^\times)$. *Then we have*

$$\chi(r_K(a)) = -\mathrm{Inv}_K(a \cup \delta^1(\chi)) \in \mathbb{Q}/\mathbb{Z}.$$

Proof. Let n be such that the image of χ is contained in $\frac{1}{n}\mathbb{Z}/\mathbb{Z}$. We have

$$\chi(r_K(a)) = \mathrm{Inv}_K(\chi \cup da) = -\mathrm{Inv}_K(da \cup \chi) = -\mathrm{Inv}_K(\partial^1(a \cup \chi)),$$

where the first equality is Corollary 10.4 and the second is Proposition 9.7. In the rightmost term, ∂^1 denotes the connecting map $\partial^1 : H^1(G_K, \overline{K}^\times \otimes \mathbb{Z}) \to H^1(G_K, \mu_n \otimes \mathbb{Z})$. For the third equality, we would like to use Lemma 9.4, and we may do so. Indeed, the short exact sequence $0 \to \mu_n \to \overline{K}^\times \to \overline{K}^\times \to 0$, as well as any other short exact sequence, clearly remains exact after tensoring with $\mathbb{Z}$. This is because $\mathbb{Z}$ is torsion-free, and so flat, as a $\mathbb{Z}$-module. Finally, a simple computation involving the formulas of Lemma 4.7 and Proposition 9.3(2), which specify what connecting maps and the cup product do to cocycles, concludes that $\partial^1(a \cup \chi) = a \cup \delta^1(\chi) \in H^2(G_K, \overline{K}^\times \otimes \mathbb{Z})$. The claim follows. □

Let k be a finite field of order q, and let $G_k = \mathrm{Gal}(\overline{k}/k)$ be its absolute Galois group. We denote the *inverse* of the element $(x \mapsto x^q) \in G_k$ by Frob_k; this is the so-called *geometric* Frobenius. Recall that the cyclic subgroup of $G_k \simeq \widehat{\mathbb{Z}}$ generated by Frob_k is dense.

The reader should beware that about half of the literature uses the notation Frob_k for the *arithmetic* Frobenius $(x \mapsto x^q) \in G_k$.

Proposition 10.10. *Let* $K/\mathbb{Q}_p$ *be a finite extension, let* $v_K : K^\times \twoheadrightarrow \mathbb{Z}$ *be the normalized valuation, and let* k *be the residue field of* K. *Then the diagram*

$$\begin{tikzcd} K^\times \arrow[r, "r_K"] \arrow[d, "v_K"'] & G_K^{\mathrm{ab}} \arrow[d, "\omega"] \\ \mathbb{Z} \arrow[r] & G_k \simeq \mathrm{Gal}(K^{\mathrm{nr}}/K) \end{tikzcd}$$

of continuous group homomorphisms commutes. Here the vertical map on the right is induced from the projection $G_K \twoheadrightarrow G_k$, which factors through G_K^{ab} since G_k is abelian. The horizontal map at the bottom is $n \mapsto (\mathrm{Frob}_k^{-1})^n$.

Proof. The claim follows from the previous lemma and the definition of the Hasse invariant map $\mathrm{Inv}_K : H^2(G_K, \overline{K}^\times)$, so we should recall the latter. Indeed, Inv_K was obtained, in the discussion preceding Proposition 7.11, as a composition of four isomorphisms:

$$H^2(G_K, \overline{K}^\times) \simeq H^2(G_k, (K^{\mathrm{nr}})^\times) \stackrel{v_K^*}{\to} H^2(G_k, \mathbb{Z}) \stackrel{(\delta^1)^{-1}}{\to} H^1(G_k, \mathbb{Q}/\mathbb{Z}) \stackrel{\psi \mapsto \psi(\mathrm{Frob}_k^{-1})}{\to} \mathbb{Q}/\mathbb{Z}.$$

Here $H^2(G_K, \overline{K}^\times) \simeq H^2(G_k, (K^{\mathrm{nr}})^\times)$ arises from realizing $H^2(G_K, \overline{K}^\times)$ as the Brauer group $\mathrm{Br}(K)$ and observing that every central simple K-algebra splits over an unramified extension of K, whereas $\delta^1 : H^1(G_k, \mathbb{Q}/\mathbb{Z}) \to H^2(G_k, \mathbb{Z})$ is the connecting map arising from the short exact sequence $0 \to \mathbb{Z} \to \mathbb{Q} \to \mathbb{Q}/\mathbb{Z} \to 0$ of G_k-modules with trivial action.

Now let $a \in K^\times$. We need to show that $\mathrm{Frob}_k^{v_K(a)} = \omega(r_K(a))$. As in several previous arguments, it suffices to show that $\chi(\mathrm{Frob}_k^{v_K(a)}) = \chi(\omega(r_K(a)))$ for all $\chi \in \mathrm{Hom}(G_k, \mathbb{Q}/\mathbb{Z}) = H^1(G_k, \mathbb{Q}/\mathbb{Z})$. By Lemma 10.9 we have

$$\chi(\omega(r_K(a))) = -\mathrm{Inv}_K(a \cup \delta^1(\chi \circ \omega)) = -\mathrm{Inv}_K(a \cup \delta^1(\chi),$$

where we use the same notation for the connecting map $\delta^1 : H^1(G_K, \mathbb{Q}/\mathbb{Z}) \to H^2(G_K, \mathbb{Z})$ in the middle of the previous formula and for $\delta^1 : H^1(G_k, \mathbb{Q}/\mathbb{Z}) \to H^2(G_k, \mathbb{Z})$ on the right-hand side. From now on in this proof, only $\delta^1 : H^1(G_k, \mathbb{Q}/\mathbb{Z}) \to H^2(G_k, \mathbb{Z})$ will appear. Also in the above formula, $a \in K^\times$ is viewed as an element of $H^0(G_K, \overline{K}^\times)$ in the middle and of $H^0(G_k, (K^{\mathrm{nr}})^\times)$ on the right.

By the (de)construction of Inv_K, we see that

$$-\mathrm{Inv}_K(a \cup \delta^1\chi) = -\delta_1^{-1}(v_K^*(a \cup \delta^1\chi))(\mathrm{Frob}_k^{-1}).$$

The map $v_K^* : H^2(G_k, (K^{\mathrm{nr}})^\times \otimes \mathbb{Z}) = H^2(G_k, (K^{\mathrm{nr}})^\times) \to H^2(G_k, \mathbb{Z})$ sends $a \cup \delta^1\chi$ to $v_K(a)\delta^1\chi$. Thus we end up with

$$\chi(\omega(r_K(a))) = -\delta_1^{-1}(v_K(a)\delta^1\chi)(\mathrm{Frob}_k^{-1}) = -v_K(a)\delta_1^{-1}(\delta^1\chi)(\mathrm{Frob}_k^{-1}) = \chi(\mathrm{Frob}_k^{v_K(a)})$$

as claimed. □

Corollary 10.11. *Let $\pi \in K$ be a uniformizer. Then $r_K(\pi) \in G_K^{\mathrm{ab}}$ is a lift of the geometric Frobenius element* $\mathrm{Frob}_k \in G_k$.

Proof. This is immediate from the previous proposition, since $v_K(\pi) = 1$. □

Remark 10.12. We could just as well have taken the isomorphism $H^1(G_k, \mathbb{Q}/\mathbb{Z}) \simeq \mathbb{Q}/\mathbb{Z}$ given by $\psi \mapsto \psi(\mathrm{Frob}_k)$ in the definition of the Hasse invariant map Inv_K. Had we done that, we would end up with the opposite normalization of the Artin reciprocity map that would send uniformizers to lifts of the arithmetic Frobenius.

Definition 10.13. Let $K/\mathbb{Q}_p$ be a finite extension. Recall the short exact sequence

$$0 \to I_K \to G_K = \mathrm{Gal}(\overline{K}/K) \stackrel{\varpi}{\to} G_k = \mathrm{Gal}(\overline{k}/k) \to 0.$$

The *Weil group* W_K is the pre-image $W_K = \varpi^{-1}(\langle \mathrm{Frob}_k \rangle)$.

We endow W_K with a topology by insisting that $I_K \subset W_K$ be an open embedding. In other words, a base of open neighborhoods of the identity in W_K is given by a base of open neighborhoods of the identity in usual topology of I_K as a subset of G_K. Observe that this topology of W_K is not the subspace topology, since I_K has infinite index in G_K and thus is not open in G_K.

Theorem 10.14. *The Artin reciprocity map of Theorem 10.3 induces a topological isomorphism $r_K : K^\times \to W_K^{\mathrm{ab}}$.*

Proof. There is an obvious injection $W_K^{\mathrm{ab}} \to G_K^{\mathrm{ab}}$, and it is clear from Proposition 10.10 that the image of $r_K : K^\times \hookrightarrow G_K^{\mathrm{ab}}$ is contained in W_K^{ab}. We thus obtain a diagram

$$\begin{tikzcd} 0 \arrow[r] & \mathcal{O}_K^\times \arrow[r] \arrow[d] & K^\times \arrow[r] \arrow[d, "r_K"] & \mathbb{Z} \arrow[r] \arrow[d, "1 \mapsto \mathrm{Frob}_k"] & 0 \\ 0 \arrow[r] & I_{K^{\mathrm{ab}}/K} \arrow[r] & W_K^{\mathrm{ab}} \arrow[r] & \langle \mathrm{Frob}_k \rangle \arrow[r] & 0, \end{tikzcd}$$

where the rows are exact (in particular, this defines $I_{K^{\mathrm{ab}}/K}$), all three vertical maps are injective, and the rightmost one is obviously an isomorphism. It remains only to show that the map $\mathcal{O}_K^\times \to I_{K^{\mathrm{ab}}/K}$ is surjective, since then we will conclude from the Five Lemma that the middle vertical map $r_K : K^\times \to W_K^{\mathrm{ab}}$ is an isomorphism.

For every finite abelian extension L/K, we know by Proposition 10.6 that r_K induces an isomorphism $K^\times/N_{L/K}(L^\times) \to \mathrm{Gal}(L/K)$. Let $I_{L/K}$ denote the kernel of the surjection $\mathrm{Gal}(L/K) \to \mathrm{Gal}(\ell/k)$. From the proof of Proposition 10.6, it follows that r_K induces a surjection of $\mathcal{O}_K^\times$ onto $I_{L/K}$. But $I_{K^{\mathrm{ab}}/K} = \varprojlim I_{L/K}$, where L/K runs over finite abelian extensions. It follows that the image of $\mathcal{O}_K^\times \to I_{K^{\mathrm{ab}}/K}$ is dense. But it is also compact and hence closed, since $\mathcal{O}_K^\times$ is. Thus the image is all of $I_{K^{\mathrm{ab}}/K}$. $\square$

Exercises

(1) Suppose that we are given an isomorphism $s_K : K^\times \xrightarrow{\sim} W_K^{\mathrm{ab}}$ for every finite extension $K/\mathbb{Q}_p$. Suppose that this collection of maps has the following properties:
- If $a \in K^\times$ and $v_K(a) = m$, then the image of $s_K(a)$ in G_k is Frob_K^m.
- If L/K is a finite extension, then $s_K(N_{L/K}(b)) = s_L(b)$ for all $b \in L^\times$.
- If L/K is a finite abelian extension, then s_K induces an isomorphism $s_K : K^\times/N_{L/K}(L^\times) \xrightarrow{\sim} \mathrm{Gal}(L/K)$.

We have proved in this section that the collection of Artin reciprocity maps r_K satisfies these three properties. In this exercise we will prove that $s_K = r_K$ for all K; in other words, these three properties determine the Artin reciprocity maps.

(a) Prove that it suffices to show that the induced maps $s_K : \mathcal{O}_K^\times \to \mathrm{Gal}(L/K)$ and $r_K : \mathcal{O}_K^\times \to \mathrm{Gal}(L/K)$ coincide whenever L/K is a finite abelian totally ramified extension.
Hint: $\mathrm{Gal}(L^{\mathrm{nr}}/K) \simeq \mathrm{Gal}(L/K) \times \mathrm{Gal}(K^{\mathrm{nr}}/K)$.

(b) Let L/K be finite, abelian, and totally ramified. Given $\sigma \in \mathrm{Gal}(L/K)$, let $L_\sigma \subset L^{\mathrm{nr}}$ be the fixed field of $(\sigma, \mathrm{Frob}_K) \in \mathrm{Gal}(L/K) \times \mathrm{Gal}(K^{\mathrm{nr}}/K) \simeq \mathrm{Gal}(L^{\mathrm{nr}}/K)$. Show that the fields L_σ and L_τ are distinct for distinct elements $\sigma, \tau \in \mathrm{Gal}(L/K)$, and show that $[L_\sigma : K] = [L : K]$ for all $\sigma \in \mathrm{Gal}(L/K)$.

(c) Let $\pi_L \in \mathcal{O}_L$ be a uniformizer, and let $a \in \mathcal{O}_K^\times$. Show that the subgroup $N_{L_{s_K(a)}/K}(L_{s_K(a)}^\times) \subset K^\times$ is generated by $N_{L/K}(\mathcal{O}_L^\times)$ and $aN_{L/K}(\pi_L)$. Conclude from this that the maps $s_K, r_K : \mathcal{O}_K^\times \to \mathrm{Gal}(L/K)$ coincide, which is enough by the first part of this question.

(2) Let G be a profinite group and let $H \subset G$ be an open subgroup. Let $s : H\backslash G \to G$ be a continuous section of the natural projection $G \twoheadrightarrow H\backslash G$; recall that this exists by Proposition 1.15.

(a) Let M be a G-module, with the discrete topology. Fix $\varphi \in Z^1(H, M)$, and define a map $\Phi : G \to M$ by

$$\Phi(g) = \sum_{Hx \in H\backslash G} s(x)^{-1} \cdot \varphi(s(x)gs(xg)^{-1}).$$

Prove that $\Phi \in Z^1(G, M)$ and that $[\Phi] = \mathrm{cor}_H^G[\varphi] \in H^1(G, M)$.

(b) Prove that

$$V(g) = \prod_{xH \in H\backslash G} s(x)gs(xg)^{-1}$$

induces a homomorphism $V : G^{\mathrm{ab}} \to H^{\mathrm{ab}}$ satisfying $\varphi \circ V = \mathrm{cor}\,\varphi$ for all $\varphi \in \mathrm{Hom}(H, \mathbb{Q}/\mathbb{Z}) = H^1(H, \mathbb{Q}/\mathbb{Z})$. Here, of course, the action of G on $\mathbb{Q}/\mathbb{Z}$ is trivial. The map V is called the transfer map (“Verschiebung” in German, hence the notation).