J. Eur. Math. Soc. 24, 1189–1283 (2022) DOI 10.4171/JEMS/1112

Peter Koymans · Carlo Pagano

On the distribution of $\mathrm{Cl}(K)[l^\infty]$ for degree l cyclic fields

Received February 27, 2019; revised December 11, 2020

Abstract. Using a recent breakthrough of Smith [22], we prove that l^∞-class groups of cyclic degree l fields have the distribution conjectured by Gerth under GRH.

Keywords. Arithmetic statistics, class groups

Contents

Peter Koymans: Mathematisch Instituut, Leiden University, Niels Bohrweg 1, 2333 CA Leiden, Netherlands; peter.koymans@hotmail.com

Carlo Pagano: Max Planck Institute for Mathematics, Vivatsgasse 7, 53111 Bonn, Germany; carlein90@gmail.com

Mathematics Subject Classification (2020): 11N45, 11R29, 11R34, 11R37, 11R45

1. Introduction

Class groups have a long and rich history going back to Gauss, who studied them in the language of binary quadratic forms. In modern terms, Gauss gave an explicit description of $\mathrm{Cl}(K)[2]$ for K a quadratic number field with narrow class group $\mathrm{Cl}(K)$. This is now known as genus theory. Since then class groups have been extensively studied leading to the development of class field theory and the Langlands conjectures.

Nowadays the class group is typically thought of as a 'random' object. Cohen and Lenstra put forward conjectures on the average behavior of class groups. Their conjecture predicts that for all odd primes p and all finite, abelian p-groups A,

$$\lim_{X\to\infty} \frac{|\{K \text{ quadratic} : 0 < D_K < X \text{ and } \mathrm{Cl}(K)[p^\infty] \cong A\}|}{|\{K \text{ quadratic} : 0 < D_K < X\}|} = \frac{\prod_{i=2}^{\infty}(1-1/p^i)}{|A|\,|\mathrm{Aut}(A)|},$$

where D_K denotes the discriminant of our field K. They also proposed a similar conjecture for imaginary quadratic fields, namely

$$\lim_{X\to\infty} \frac{|\{K \text{ quadratic} : -X < D_K < 0 \text{ and } \mathrm{Cl}(K)[p^\infty] \cong A\}|}{|\{K \text{ quadratic} : -X < D_K < 0\}|} = \frac{\prod_{i=1}^{\infty}(1-1/p^i)}{|\mathrm{Aut}(A)|}.$$

Although the Cohen and Lenstra conjectures have attracted a great deal of attention, there are very few proven instances. Davenport and Heilbronn [2] obtained partial results in the case $p = 3$, while the case $p > 3$ is still wide open. Cohen and Lenstra originally stated their conjectures only for odd p, but the case $p = 2$ is also very interesting. In the case $p = 2$ we have a very explicit description of $\mathrm{Cl}(K)[2]$, and the class group can no longer be thought of as a random object.

Gerth [6] proposed the following modification of the Cohen–Lenstra conjectures; instead of $\mathrm{Cl}(K)[2^\infty]$, it is $(2\mathrm{Cl}(K))[2^\infty]$ that behaves randomly. Fouvry and Klüners [3, 4], building on earlier work of Heath-Brown on 2-Selmer groups [11], proved that $(2\mathrm{Cl}(K))[2]$ has the correct distribution for both imaginary and real quadratic fields. A major breakthrough came when Smith [22], extending his earlier work [17], proved

$$\lim_{X\to\infty} \frac{|\{K \text{ quadratic} : -X < D_K < 0 \text{ and } (2\mathrm{Cl}(K))[2^\infty] \cong A\}|}{|\{K \text{ quadratic} : -X < D_K < 0\}|} = \frac{\prod_{i=1}^{\infty}(1-1/2^i)}{|\mathrm{Aut}(A)|}$$

for all finite, abelian 2-groups A. In the course of the proof Smith develops several powerful and versatile methods. Using the same methods, Smith also deals with the distribution of 2^k-Selmer groups of elliptic curves.

Essential to Smith's method is the explicit description of $\mathrm{Cl}(K)[2]$, i.e. genus theory. This allows us to study complicated sets such as $2^{k-1}\mathrm{Cl}(K)[2^k]$ via its natural inclusion in $\mathrm{Cl}(K)[2]$. Now let l be an odd prime and K be a cyclic degree l field, so that $\mathrm{Cl}(K)$ becomes a $\mathbb{Z}[\zeta_l]$-module in $l-1$ different ways depending on the identification between $\mathrm{Gal}(K/\mathbb{Q})$ and $\langle\zeta_l\rangle$. Fortunately, since K is cyclic, the isomorphism type of $\mathrm{Cl}(K)$ as a $\mathbb{Z}[\zeta_l]$-module does not depend on this identification.

Genus theory gives an explicit description of $\mathrm{Cl}(K)[1-\zeta_l]$. Klys [12] proved conditional on GRH that $((1-\zeta_l)\mathrm{Cl}(K))[1-\zeta_l]$ has the expected distribution [7, p. 312] and also gave an unconditional proof in the case $l=3$. Both these results use the Fouvry–Klüners method [3,4]. Our main theorem proves that $((1-\zeta_l)\mathrm{Cl}(K))[(1-\zeta_l)^\infty]$ has the expected distribution using the breakthrough method of Smith [22].

Theorem 1.1. *Assume GRH. Then for all odd primes l and all finitely generated, torsion $\mathbb{Z}_l[\zeta_l]$-modules A the limit*

$$\lim_{X\to\infty}\frac{|\{K \text{ cyclic of degree } l : \mathrm{rad}(D_K) < X \text{ and } ((1-\zeta_l)\mathrm{Cl}(K))[(1-\zeta_l)^\infty] \cong A\}|}{|\{K \text{ cyclic of degree } l : \mathrm{rad}(D_K) < X\}|}$$

exists and is equal to

$$\frac{\prod_{i=2}^{\infty}(1-1/l^i)}{|A|\,|\mathrm{Aut}_{\mathbb{Z}_l[\zeta_l]}(A)|}.$$

We order our fields by the radical of the discriminant for technical convenience. The interested reader should have no trouble proving Theorem 1.1 when the fields are instead ordered by the absolute value of the discriminant. Let $\mathrm{Field}(N,l)$ be the set of cyclic degree l number fields K over $\mathbb{Q}$ with $\mathrm{rad}(D_K)\le N$. For $0\le j\le n$ let $P(j|n)$ be the probability that a uniformly chosen $n\times(n+1)$ matrix with entries in $\mathbb{F}_l$ has rank $n-j$. Furthermore, for $k\ge 2$ and $n\ge 0$ let $D_{l,k}(n)$ be the set of cyclic degree l fields K satisfying

$$\dim_{\mathbb{F}_l}(1-\zeta_l)^{k-1}\mathrm{Cl}(K)[(1-\zeta_l)^k] = n.$$

Theorem 1.1 will fall as a consequence of the following theorem that we prove in Section 13.

Theorem 1.2. *Assume GRH and let l be an odd prime. There are $c, A, N_0 > 0$ such that for all $N > N_0$, all integers $m\ge 2$ and all sequences $n_2\ge\cdots\ge n_{m+1}\ge 0$ of integers we have*

$$\left|\left|\mathrm{Field}(N,l)\cap\bigcap_{k=2}^{m+1}D_{l,k}(n_k)\right| - P(n_{m+1}|n_m)\cdot\left|\mathrm{Field}(N,l)\cap\bigcap_{k=2}^{m}D_{l,k}(n_k)\right|\right| \le \frac{AN}{(\log\log N)^{\frac{c}{m^{2(l^2+l)^m}}}}.$$

One could also wonder what happens without GRH. One of the first steps in Smith's method is to fix the Rédei matrix of K. This is a rather complicated matter, and Smith proves a weak equidistribution statement with ingenious use of the large sieve. The large

sieve for l-th power residue symbols is currently not as well-developed as the classical quadratic large sieve. It is for this reason only that we need GRH and it may very well be possible to remove this assumption if one obtains a suitable version of the large sieve. An additional benefit of GRH is that it makes several other proofs in this paper substantially easier and shorter.

If K is a quadratic field, $(2\mathrm{Cl}(K))[2^\infty]$ was traditionally studied from the viewpoint of *governing fields*. Cohn and Lagarias [1] conjectured that for each integer $k \geq 1$ and each integer $d \not\equiv 2 \bmod 4$, there exists a normal field extension $M_{d,k}$ over $\mathbb{Q}$ such that the 2^k-rank of $\mathrm{Cl}(\mathbb{Q}(\sqrt{dp}))$ is determined by the splitting of p in $M_{d,k}$. Such a field $M_{d,k}$ is called a governing field. Stevenhagen [23] proved their conjecture for $k \leq 3$. For $k > 3$ the Cohn and Lagarias conjecture is not known to be true or false for any value of d, but widely believed to be false with compelling evidence found by Milovic [19] and later by Koymans and Milovic [13–15].

One of the major insights in Smith's work is the notion of a relative governing field. To explain this notion, let $\{p_{1,0}, p_{1,1}\}, \dots, \{p_{k,0}, p_{k,1}\}$ be primes and let d be a negative squarefree integer. For any function $f : \{1, \dots, k\} \to \{0, 1\}$ define

$$K(f) := \mathbb{Q}\left(\sqrt{d \prod_{i=1}^{k} p_{i,f(i)}}\right).$$

Choose any function $f' : \{1, \dots, k\} \to \{0, 1\}$. Under suitable conditions Smith shows that the 2^k-ranks of $K(f)$ with $f \neq f'$ together with the splitting of $p_{k,0}$ and $p_{k,1}$ in a field depending only on $\{p_{1,0}, p_{1,1}\}, \dots, \{p_{k-1,0}, p_{k-1,1}\}$ determine the 2^k-rank of $K(f')$. This field can be thought of as a relative governing field, and the resulting theorem can be seen as an extremely general 'reflection principle'. Amazingly enough, this is the only algebraic result about class groups used in Smith's paper. The rest of his paper is dedicated to rather ingenious combinatorial and analytical arguments that prove the desired equidistribution.

Our paper borrows heavily from the ideas introduced by Smith; and in particular his proof strategy. We start by generalizing his reflection principle. To do so, we introduce a generalized notion of Smith's relative governing fields. One has to be slightly careful, since Smith relies on the fact that $\zeta_2 \in \mathbb{Q}$. Furthermore, Smith uses that $\mathrm{Gal}(K/\mathbb{Q}) \cong \mathbb{F}_2$ has trivial automorphism group to make several important identifications. However, if K is cyclic of degree l with $l > 2$, we have $\mathrm{Gal}(K/\mathbb{Q}) \cong \mathbb{F}_l$, which does not have trivial automorphism group. To work around this, we need to work with characters $\chi : G_{\mathbb{Q}} \to \langle \zeta_l \rangle$ instead of fields. During our proofs, it will be very important to carefully keep track of the characters $\chi : G_{\mathbb{Q}} \to \langle \zeta_l \rangle$ that we have chosen, since we use these characters to make the necessary identifications in a canonical way.

Once we have generalized Smith's notion of relative governing fields, the proof is mostly a straightforward adaptation of Smith's work with the exception of three major changes. Suppose that we have chosen characters $\chi_p : G_{\mathbb{Q}} \to \langle \zeta_l \rangle$ of conductor p for all

$p \equiv 1 \bmod l$. We need to deal with sums of the type

$$\sum_{X<p<Y} \chi_p(\mathrm{Frob}(q))$$

for fixed q. If χ runs over quadratic characters, one may prove cancellation of such sums by an application of Chebotarev or the large sieve. However, if χ runs over more general characters, such a sum may be biased for a bad choice of the characters χ_p. To work around this issue, we average over all choices of characters χ_p, and use a mixture of Chebotarev and combinatorial arguments to show that there is cancellation for most choices of characters χ_p.

The second issue is the earlier mentioned lack of an appropriate large sieve, and we work around this by assuming GRH. Finally, there is one important point where the analogy between $\mathrm{Cl}(K)[2^\infty]$ for K imaginary quadratic and $\mathrm{Cl}(K)[(1-\zeta_l)^\infty]$ for K degree l cyclic breaks down. Indeed, the relation between the ramified prime ideals in $\mathrm{Cl}(K)[2]$ is explicitly given by Gauss genus theory. On the other hand, the relation between the ramified prime ideals in $\mathrm{Cl}(K)[1-\zeta_l]$ should be thought of as being random. It is for this reason that $\mathrm{Cl}(K)[(1-\zeta_l)^\infty]$ is more similar to $\mathrm{Cl}(K)[2^\infty]$ for K real quadratic.

It is not hard to adapt our arguments to also deal with real quadratic fields. The main obstacle is that one needs to make the Markov chain analysis of Gerth [6] effective, and also generalize his Markov chain analysis to include even discriminants. One further needs to be careful with ramification at 2, because $\zeta_2 \in \mathbb{Q}_2$ while $\zeta_l \notin \mathbb{Q}_l$ for odd l, but this is a mere technicality. The authors plan to conduct a deeper study of the arithmetic of real quadratic fields in future work including the Markov chain analysis necessary to extend our results here to real quadratics.

2. Set-up

In this section we introduce the most important objects and notation. Our first subsection defines the central objects in this paper. Once this is done, we devote the next subsection to the necessary notation and conventions.

2.1. The Artin pairing

Let l be an odd prime number, which is treated as fixed throughout the paper. Whenever we use $O(\cdot)$ or $\ll$, the implicit constant may depend on l and we shall not record this dependence. Fix once and for all $\overline{\mathbb{Q}}$, an algebraic closure of $\mathbb{Q}$. If $K \subseteq \overline{\mathbb{Q}}$ is a number field, we denote by $G_K := \mathrm{Gal}(\overline{\mathbb{Q}}/K)$. Also fix an element ζ_l of $\overline{\mathbb{Q}}^*$ with multiplicative order equal to l; this is a generator of the group $\mu_l(\overline{\mathbb{Q}}) := \{\alpha \in \overline{\mathbb{Q}} : \alpha^l = 1\}$. We define

$$\Gamma_{\mu_l}(\mathbb{Q}) := \mathrm{Hom}_{\mathrm{top.gr.}}(G_{\mathbb{Q}}, \mu_l(\overline{\mathbb{Q}})).$$

Here $G_{\mathbb{Q}}$ has the Krull topology and $\mu_l(\overline{\mathbb{Q}})$ the discrete topology. For a character $\chi \in \Gamma_{\mu_l}(\mathbb{Q})$, we denote by

$$K_\chi := \overline{\mathbb{Q}}^{\ker(\chi)}$$

the corresponding extension of $\mathbb{Q}$. This is a cyclic extension with degree dividing l, and equal to 1 if and only if χ is the trivial character. We denote by $\mathrm{Cl}(K_\chi)$ the class group of K_χ. Observe that $\mathrm{Cl}(K_\chi)[l^\infty]$ is a $\mathbb{Z}_l[\mathrm{Gal}(K_\chi/\mathbb{Q})]$-module. Since $\mathbb{Z}$ is a PID, the norm element

$$N_{\mathrm{Gal}(K_\chi/\mathbb{Q})} := \sum_{g \in \mathrm{Gal}(K_\chi/\mathbb{Q})} g$$

acts trivially on $\mathrm{Cl}(K_\chi)$. From this, we deduce that $\mathrm{Cl}(K_\chi)[l^\infty]$ has naturally the structure of a $\frac{\mathbb{Z}_l[\mathrm{Gal}(K_\chi/\mathbb{Q})]}{N_{\mathrm{Gal}(K_\chi/\mathbb{Q})}}$-module. Moreover, χ gives a natural isomorphism of $\mathbb{Z}_l$-algebras

$$\chi : \frac{\mathbb{Z}_l[\mathrm{Gal}(K_\chi/\mathbb{Q})]}{N_{\mathrm{Gal}(K_\chi/\mathbb{Q})}} \to \mathbb{Z}_l[\zeta_l] := \mathbb{Z}[\zeta_l] \otimes_{\mathbb{Z}} \mathbb{Z}_l.$$

In this manner $\mathrm{Cl}(K_\chi)[l^\infty]$ is naturally equipped with the structure of a $\mathbb{Z}_l[\zeta_l]$-module. In what follows, it is *always* with respect to this structure that we will talk about $\mathrm{Cl}(K_\chi)[l^\infty]$ as a $\mathbb{Z}_l[\zeta_l]$-module.

The ring $\mathbb{Z}_l[\zeta_l]$ is a local PID with the unique maximal ideal generated by $1-\zeta_l$. Therefore for every finite $\mathbb{Z}_l[\zeta_l]$-module A, there is a unique function $f_A : \mathbb{Z}_{\geq 1} \to \mathbb{Z}_{\geq 0}$ such that

$$A \simeq_{\mathbb{Z}_l[\zeta_l]} \bigoplus_{i \in \mathbb{Z}_{\geq 1}} \left(\frac{\mathbb{Z}_l[\zeta_l]}{(1-\zeta_l)^i \mathbb{Z}_l[\zeta_l]} \right)^{f_A(i)}.$$

Since A is finite, the map f_A has finite support and it can be reconstructed from the decreasing sequence of numbers

$$k \mapsto \mathrm{rk}_{(1-\zeta_l)^k} A := \dim_{\mathbb{F}_l} (1-\zeta_l)^{k-1} A[(1-\zeta_l)^k],$$

defined for every positive integer k. Therefore the sequence

$$\{\mathrm{rk}_{(1-\zeta_l)^k} \mathrm{Cl}(K_\chi)\}_{k \in \mathbb{Z}_{\geq 1}}$$

determines completely the structure of the $\mathbb{Z}_l[\zeta_l]$-module $\mathrm{Cl}(K_\chi)[l^\infty]$. Here, for brevity, $\mathrm{rk}_{(1-\zeta_l)^k} \mathrm{Cl}(K_\chi)$ stands for $\mathrm{rk}_{(1-\zeta_l)^k} \mathrm{Cl}(K_\chi)[l^\infty]$, which has been defined above. The following $\mathbb{Z}_l[\zeta_l]$-module will have a big role for us:

$$N := \frac{\mathbb{Q}_l(\zeta_l)}{\mathbb{Z}_l[\zeta_l]}.$$

For a finitely generated, torsion $\mathbb{Z}_l[\zeta_l]$-module A, we define

$$A^\vee := \mathrm{Hom}_{\mathbb{Z}_l[\zeta_l]}(A, N).$$

The following is not hard to see.

Proposition 2.1. *The $\mathbb{Z}_l[\zeta_l]$-modules A and $A^\vee$ are isomorphic.*

For every integer $k \geq 1$ we next define a pairing of $\mathbb{Z}_l[\zeta_l]$-modules

$$\mathrm{Art}_k(A) : (1-\zeta_l)^{k-1}A[(1-\zeta_l)^k] \times (1-\zeta_l)^{k-1}A^\vee[(1-\zeta_l)^k] \to N[1-\zeta_l].$$

Let $a \in (1-\zeta_l)^{k-1}A[(1-\zeta_l)^k]$ and $\chi \in (1-\zeta_l)^{k-1}A^\vee[(1-\zeta_l)^k]$. Let $\psi \in A^\vee$ be an element such that $(1-\zeta_l)^{k-1}\psi = \chi$. We put

$$\mathrm{Art}_k(A)(a,\chi) := \psi(a).$$

Observe that since $a \in (1-\zeta_l)^{k-1}A[(1-\zeta_l)^k]$, the definition does not depend on the choice of ψ. The following fact is straightforward.

Proposition 2.2. *The left kernel of* $\mathrm{Art}_k(A)$ *is* $(1-\zeta_l)^k A[(1-\zeta_l)^{k+1}]$ *and the right kernel is* $(1-\zeta_l)^k A^\vee[(1-\zeta_l)^{k+1}]$.

Hence, instead of directly dealing with

$$\{\mathrm{rk}_{(1-\zeta_l)^k}\,\mathrm{Cl}(K_\chi)\}_{k\in\mathbb{Z}_{\geq 1}},$$

our goal is to control the sequence of pairings

$$\{\mathrm{Art}_k(\mathrm{Cl}(K_\chi))\}_{k\in\mathbb{Z}_{\geq 1}}.$$

Here $\mathrm{Art}_k(\mathrm{Cl}(K_\chi))$ is an abbreviation for $\mathrm{Art}_k(\mathrm{Cl}(K_\chi)[l^\infty])$, which has been defined above. Proving equidistribution for this sequence of pairings is the main goal of this paper. We start with some algebraic tools, which culminate in an extremely general reflection principle. In the next section we fix identifications between some cyclic groups of order l that occur in this paper, as well as some other important conventions regarding notation.

2.2. Identifications and conventions

Throughout the paper we will encounter the groups $\mathbb{F}_l$, $\langle\zeta_l\rangle$ and $N[1-\zeta_l]$. These three groups are isomorphic, but not in a canonical way. Working with each group has its own advantages. Kummer theory is most naturally stated using $\langle\zeta_l\rangle$, while $\mathbb{F}_l$ has a natural product structure that we will take advantage of. Finally, $N[1-\zeta_l]$ is a subgroup of N, which is the image of the various Artin pairings. We need to identify these three groups at several points in the paper, and it is of utmost importance this is done in a consistent manner. We refer to the following diagram whenever such an identification is made.

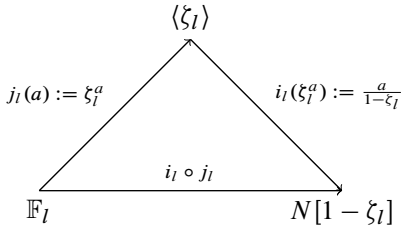

Any other identification is made by inverting the arrows and maps. The symbol $\mathbb{C}$ will denote, as usual, the complex numbers. The symbol i denotes a fixed element of $\mathbb{C}^*$ of multiplicative order equal to 4. The function $\exp : \mathbb{C} \to \mathbb{C}^*$ denotes the exponential map. The group $\mu_l(\mathbb{C})$ is generated by the element $\exp(2\pi \mathrm{i}/l)$. We also fix the identification $h_l : \mu_l(\mathbb{C}) \to \langle \zeta_l \rangle$ given by

$$h_l(\exp(2\pi \mathrm{i}/l)) := \zeta_l.$$

We denote $\Gamma_{\mathbb{F}_l}(\mathbb{Q}) := \mathrm{Hom}_{\text{top.gr.}}(G_{\mathbb{Q}}, \mathbb{F}_l)$. The map j_l induces an isomorphism $\Gamma_{\mathbb{F}_l}(\mathbb{Q}) \to \Gamma_{\mu_l}(\mathbb{Q})$. We also define $\Gamma_{\mu_l(\mathbb{C})}(\mathbb{Q}) := \mathrm{Hom}_{\text{top.gr.}}(G_{\mathbb{Q}}, \mu_l(\mathbb{C}))$, so h_l induces an isomorphism $\Gamma_{\mu_l(\mathbb{C})}(\mathbb{Q}) \to \Gamma_{\mu_l}(\mathbb{Q})$.

If q is either equal to l or to a prime number that is congruent to 1 modulo l, then there exists a unique cyclic degree l extension of $\mathbb{Q}$ that is totally ramified at q and unramified elsewhere. By class field theory, if $q \neq l$, this is the unique cyclic degree l extension contained in $\mathbb{Q}(\zeta_q)/\mathbb{Q}$. If $q = l$ this is the unique cyclic degree l extension contained in $\mathbb{Q}(\zeta_{l^2})/\mathbb{Q}$. Here ζ_q and ζ_{l^2} are elements of $\overline{\mathbb{Q}}$ of multiplicative order equal to q and l^2 respectively. We denote these extensions by L_q, for $q \neq l$, and by L_{l^2} in the case $q = l$. For each q congruent 1 modulo l we fix a character

$$\chi_q \in \Gamma_{\mu_l}(\mathbb{Q})$$

such that $\ker(\chi_q) = G_{L_q}$. There is no way to make such a choice in a canonical manner, and we fix one simply for notational purposes. Similarly, we fix a character

$$\chi_l \in \Gamma_{\mu_l}(\mathbb{Q})$$

such that $\ker(\chi_l) = G_{L_{l^2}}$. All our algebraic results work for a fixed choice of characters, but later on we will have to vary the choice of characters to make our analytic results work.

The set $\{\chi_q\}_{q \equiv 1 \bmod l} \cup \{\chi_l\}$ is a *basis* for $\Gamma_{\mu_l}(\mathbb{Q})$. In particular any cyclic degree l extension ramifies only at primes congruent 1 modulo l or at l, see Proposition 4.5 for a generalization of this fact. By the conductor-discriminant formula we see that a positive integer D equals $\Delta_{K_\chi/\mathbb{Q}}$ for some $\chi \in \Gamma_{\mu_l}(\mathbb{Q})$ if and only if

$$D = (q_1 \cdot \ldots \cdot q_r)^{l-1}$$

with $\{q_i\}_{1 \leq i \leq r}$ a set of r distinct elements each either a prime equal to 1 modulo l, or equal to l^2. In case D admits such a factorization then $D = \Delta_{K_\chi/\mathbb{Q}}$ for $(l-1)^r$ different choices of χ, which amounts to a total of $(l-1)^{r-1}$ different fields. From now on we say that a positive integer D is *l-admissible* if it is the discriminant of a cyclic degree l extension of $\mathbb{Q}$. For each l-admissible integer D we define an *amalgama for D* to be a map $\epsilon : \{q \mid D\}_{q \text{ prime}} \to \{1, \ldots, l-1\}$. For an l-admissible integer D, the set of characters $\chi \in \Gamma_{\mu_l}(\mathbb{Q})$ such that $\Delta_{K_\chi/\mathbb{Q}} = D$ corresponds bijectively to the set of amalgamas for D, via the assignment

$$\epsilon \mapsto \chi_\epsilon(D) := \prod_{q \mid D} \chi_q^{\epsilon(q)}.$$

We denote by $\chi \mapsto \epsilon_\chi$ the inverse assignment. Let $\chi \in \Gamma_{\mu_l}(\mathbb{Q})$ and $q \mid \Delta_{K_\chi/\mathbb{Q}}$. Then there is a unique prime ideal in O_{K_χ} lying above q. We denote such a prime ideal by $\mathrm{Up}_{K_\chi}(q)$. For a positive integer b and for a prime number q dividing b, we write $\epsilon_b(q)$ for the unique integer in $\{0, \ldots, l-1\}$ with $\epsilon_b(q) \equiv v_{\mathbb{Q}_q}(b) \bmod l$. In particular, $\prod_{q|b} q^{\epsilon_b(q)}$ equals b in $\mathbb{Q}^*/\mathbb{Q}^{*l}$. We also define $[d] := \{1, \ldots, d\}$ for any integer d.

We shall frequently encounter maps from some profinite group G to some finite set X. Whenever we encounter such a map, it will be continuous with respect to the discrete topology on X.

Lemma 2.3. *Let G be a profinite group, let X be a discrete topological space and $\phi : G \to X$ a continuous map. There exists a largest (by inclusion) open normal subgroup N_ϕ of G such that the map ϕ factors through the canonical projection $G \to G/N_\phi$.*

Proof. This is straightforward. ∎

Thanks to Lemma 2.3 we can make the following definition.

Definition 2.4. Let $\phi : G_\mathbb{Q} \to X$ be a continuous map, where X is a discrete topological space. The *group of definition* of ϕ is the open normal subgroup N_ϕ in Lemma 2.3. Furthermore, we define $L(\phi)$ to be the fixed field of N_ϕ, which will be called the *field of definition* of ϕ.

3. Ambiguous ideals and genus theory

In this section we study $\mathrm{Cl}(K_\chi)[1-\zeta_l]$ and $\mathrm{Cl}(K_\chi)^\vee[1-\zeta_l]$. The material collected here is well-known to experts and can be found in various forms in the literature, but we have decided to include it for the sake of completeness.

3.1. Ambiguous ideals

The material in this subsection is known as the theory of ambiguous ideals. Since in $\mathbb{Z}_l[\zeta_l]$ we have the equality of ideals $(1-\zeta_l)^{l-1} = (l)$, in particular $\mathrm{Cl}(K_\chi)[1-\zeta_l]$ is an $\mathbb{F}_l$-vector space. From the definition of the structure of $\mathrm{Cl}(K_\chi)[l^\infty]$ as a $\mathbb{Z}_l[\zeta_l]$-module, it is clear that $\mathrm{Cl}(K_\chi)[1-\zeta_l] = \mathrm{Cl}(K_\chi)^{\mathrm{Gal}(K_\chi/\mathbb{Q})}$. Thus we can obtain a description of $\mathrm{Cl}(K_\chi)[1-\zeta_l]$ by taking Galois invariants of the sequence

$$1 \to \mathrm{Pr}(K_\chi) \to \mathcal{I}_{K_\chi} \to \mathrm{Cl}(K_\chi) \to 1,$$

where $\mathcal{I}_{K_\chi}$ denotes the group of fractional ideals of O_{K_χ} and $\mathrm{Pr}(K_\chi)$ denotes the group of principal fractional ideals of O_{K_χ}. To take advantage of this sequence we shall begin with the following simple fact.

Proposition 3.1. *We have*

$$H^1(\mathrm{Gal}(K_\chi/\mathbb{Q}), \mathrm{Pr}(K_\chi)) = 0.$$

Proof. We take the exact sequence

$$1 \to O_{K_\chi}^* \to K_\chi^* \to \mathrm{Pr}(K_\chi) \to 1.$$

Thanks to Hilbert 90, we can canonically identify the H^1 in the statement with

$$\ker\big(H^2(\mathrm{Gal}(K_\chi/\mathbb{Q}), O_{K_\chi}^*) \to H^2(\mathrm{Gal}(K_\chi/\mathbb{Q}), K_\chi^*)\big).$$

Hence it is *sufficient* to show that $H^2(\mathrm{Gal}(K_\chi/\mathbb{Q}), O_{K_\chi}^*) = 0$. Since $\mathrm{Gal}(K_\chi/\mathbb{Q})$ is cyclic, it follows from [25, Section 6.2] that this last H^2 is isomorphic to

$$\frac{\mathbb{Z}^*}{N_{K_\chi/\mathbb{Q}}(O_{K_\chi}^*)} = \frac{\langle -1\rangle}{\langle -1\rangle} = \{1\},$$

where in the first equality we have used that l is odd. This concludes the proof. ∎

Therefore we have the following corollary.

Corollary 3.2. *The natural map* $\mathcal{I}_{K_\chi}^{\mathrm{Gal}(K_\chi/\mathbb{Q})} \to \mathrm{Cl}(K_\chi)^{\mathrm{Gal}(K_\chi/\mathbb{Q})}$ *induces an isomorphism*

$$\frac{\mathcal{I}_{K_\chi}^{\mathrm{Gal}(K_\chi/\mathbb{Q})}}{\mathrm{Pr}(K_\chi)^{\mathrm{Gal}(K_\chi/\mathbb{Q})}} \simeq \mathrm{Cl}(K_\chi)^{\mathrm{Gal}(K_\chi/\mathbb{Q})}.$$

We next focus on the group $\mathcal{I}_{K_\chi}^{\mathrm{Gal}(K_\chi/\mathbb{Q})}$. Recall from Section 2.2 that for a prime q with $q \mid \Delta_{K_\chi/\mathbb{Q}}$ the symbol $\mathrm{Up}_{K_\chi}(q)$ denotes the unique prime ideal of O_{K_χ} lying above q. We have the following fact.

Proposition 3.3. *For any element* $I \in \mathcal{I}_{K_\chi}^{\mathrm{Gal}(K_\chi/\mathbb{Q})}$ *there is a unique pair* (ϵ, n) *where* ϵ *is a map* $\epsilon : \{q \mid \Delta_{K_\chi/\mathbb{Q}}\}_{q \text{ prime}} \to \{0, \dots, l-1\}$ *and* n *is a positive rational number, with the property*

$$I = (n) \prod_{q \mid \Delta_{K_\chi/\mathbb{Q}}} \mathrm{Up}_{K_\chi}(q)^{\epsilon(q)}.$$

Proof. Let I be in $\mathcal{I}_{K_\chi}^{\mathrm{Gal}(K_\chi/\mathbb{Q})}$ and factor I as a product of prime ideals. Then inert primes can clearly be bunched together into a rational fractional ideal. For split primes, the Galois invariance and unique factorization of ideals imply that every exponent is invariant in each Galois orbit of split primes. Hence also the split primes can be bunched together to give a total contribution that is a rational fractional ideal.

The remaining primes are exactly the ramified primes. For each ramified prime, we can always pick the largest multiple of l smaller than the exponent, and throw this contribution into a rational fractional ideal. This shows the existence part of the proposition.

For the uniqueness, suppose that the pairs (ϵ_1, n_1) and (ϵ_2, n_2) give the same ideal. Observe that when we norm down to $\mathbb{Q}$, we obtain for q a prime not dividing the discriminant that $l \cdot v_{\mathbb{Q}_q}(n_1) = l \cdot v_{\mathbb{Q}_q}(n_2)$ and hence $v_{\mathbb{Q}_q}(n_1) = v_{\mathbb{Q}_q}(n_2)$. Finally, if q is ramified, we obtain

$$l \cdot v_{\mathbb{Q}_q}(n_1) + \epsilon_1(q) = l \cdot v_{\mathbb{Q}_q}(n_2) + \epsilon_2(q).$$

Since $\epsilon_1(q)$ and $\epsilon_2(q)$ are in $\{0,\dots,l-1\}$, it must be that $\epsilon_1(q)=\epsilon_2(q)$ and $v_{\mathbb{Q}_q}(n_1)=v_{\mathbb{Q}_q}(n_2)$. So $\epsilon_1=\epsilon_2$ and the two rational numbers n_1 and n_2 have the same valuation at all finite places and they are both positive, hence they coincide. ∎

Since the surjection $\mathcal{I}_{K_\chi}^{\mathrm{Gal}(K_\chi/\mathbb{Q})} \twoheadrightarrow \mathrm{Cl}(K_\chi)[1-\zeta_l]$ factors through $\mathcal{I}_{\mathbb{Q}}$, it induces a surjective map of $\mathbb{F}_l$-vector spaces

$$\overline{\mathrm{Cl}}(K_\chi) \twoheadrightarrow \mathrm{Cl}(K_\chi)[1-\zeta_l],$$

where $\overline{\mathrm{Cl}}(K_\chi)$ denotes the subgroup consisting of those α in $\frac{\mathbb{Q}^*}{\mathbb{Q}^{*l}}$ such that $v(\alpha)$ is divisible by l for all places $v\in\Omega_{\mathbb{Q}}$ not dividing $\Delta_{K_\chi/\mathbb{Q}}$. It follows from Proposition 3.3 that we have an identification

$$\overline{\mathrm{Cl}}(K_\chi) \simeq \frac{\mathcal{I}_{K_\chi}^{\mathrm{Gal}(K_\chi/\mathbb{Q})}}{\mathcal{I}_{\mathbb{Q}}}$$

via the norm map, which sends every invariant non-zero ideal to its norm in $\mathbb{Q}^*/\mathbb{Q}^{*l}$. By construction, the kernel of the map $\overline{\mathrm{Cl}}(K_\chi) \twoheadrightarrow \mathrm{Cl}(K_\chi)[1-\zeta_l]$ equals

$$\frac{\mathrm{Pr}(K_\chi)^{\mathrm{Gal}(K_\chi/\mathbb{Q})}}{\mathcal{I}_{\mathbb{Q}}}.$$

Due to Hilbert 90 this group is canonically isomorphic to $H^1(\mathrm{Gal}(K_\chi/\mathbb{Q}), O_{K_\chi}^*)$.

Proposition 3.4. *The group $H^1(\mathrm{Gal}(K_\chi/\mathbb{Q}), O_{K_\chi}^*)$ is a 1-dimensional $\mathbb{F}_l$-vector space.*

Proof. Fix a non-trivial element σ of $\mathrm{Gal}(K_\chi/\mathbb{Q})$. Then σ generates $\mathrm{Gal}(K_\chi/\mathbb{Q})$. Since $\mathrm{Gal}(K_\chi/\mathbb{Q})$ is a cyclic group, an elementary calculation with 1-cocycles shows that the H^1 in the statement is isomorphic to

$$\frac{\{\alpha\in O_{K_\chi}^* : N_{K_\chi/\mathbb{Q}}(\alpha)=1\}}{\{\sigma(\beta)/\beta : \beta\in O_{K_\chi}^*\}}.$$

Observe that this is an $\mathbb{F}_l$-vector space, which also follows from the size of the Galois group being l. Thus we can compute it also by first completing at l, i.e. considering

$$\{\alpha\in O_{K_\chi}^* : N_{K_\chi/\mathbb{Q}}(\alpha)=1\}\otimes_{\mathbb{Z}}\mathbb{Z}_l = O_{K_\chi}^*\otimes_{\mathbb{Z}}\mathbb{Z}_l.$$

Through χ, this can be naturally viewed as a $\mathbb{Z}_l[\zeta_l]$-module. But the $\mathbb{Z}_l$-torsion of $O_{K_\chi}^*\otimes_{\mathbb{Z}}\mathbb{Z}_l$ is trivial, since l is odd and the $\mathbb{Z}$-torsion of $O_{K_\chi}^*$ is equal to $\langle -1\rangle$. Hence the $\mathbb{Z}_l[\zeta_l]$-torsion is also trivial. Therefore, using Dirichlet's Unit Theorem and the fact that l is odd, it must be that

$$O_{K_\chi}^*\otimes_{\mathbb{Z}}\mathbb{Z}_l \simeq_{\mathbb{Z}_l[\zeta_l]} \mathbb{Z}_l[\zeta_l].$$

Therefore the H^1 we are after is isomorphic to $\frac{\mathbb{Z}_l[\zeta_l]}{(1-\zeta_l)}\simeq\mathbb{F}_l$. ∎

Combining Corollary 3.2 with Proposition 3.4, we deduce the following.

Corollary 3.5. *The map*

$$\overline{\mathrm{Cl}}(K_\chi) \twoheadrightarrow \mathrm{Cl}(K_\chi)[1-\zeta_l]$$

has 1*-dimensional kernel. Moreover, a generator of the kernel can be obtained by taking* $N_{K_\chi/\mathbb{Q}}(\gamma) \in \frac{\mathbb{Q}^*}{\mathbb{Q}^{*l}}$ *for any* $\gamma \in K_\chi^*$ *such that*

$$\sigma(\gamma)/\gamma \in O_{K_\chi}^* - \{\sigma(\beta)/\beta : \beta \in O_{K_\chi}^*\},$$

where σ *is any non-trivial element of* $\mathrm{Gal}(K_\chi/\mathbb{Q})$.

The second part of Corollary 3.5 suggests that we should not expect a simple description for the relation among the ramified prime ideals as is the case for the 2-torsion of imaginary quadratic number fields. Instead we should expect it to be a genuine 'random' piece of data. It is for this reason that cyclic degree l fields are analogous to real quadratic number fields.

3.2. Genus theory

The material of the previous subsection is sufficient to determine the structure of $\mathrm{Cl}(K_\chi)^\vee[1-\zeta_l]$ right away. We remind the reader that we have fixed an identification i_l between $\langle\zeta_l\rangle$ and $N[1-\zeta_l]$ in Section 2.2. Also recall that we have a canonical identification $\mathrm{Cl}(K_\chi) = \mathrm{Gal}(H_{K_\chi}/K_\chi)$ via the Artin map, where H_{K_χ} is the Hilbert class field of K_χ. Finally, χ_p denotes a fixed choice of an element in $\Gamma_{\mu_l}(\mathbb{Q})$ of conductor dividing p^∞.

Proposition 3.6. *Let* χ *be in* $\Gamma_{\mu_l}(\mathbb{Q})$. *The set* $\{i_l \circ \chi_p|_{\ker(\chi)}\}_{p|\Delta_{K_\chi/\mathbb{Q}}}$ *is a generating set for* $\mathrm{Cl}(K_\chi)^\vee[1-\zeta_l]$. *Moreover, any relation among these characters is a multiple of the trivial relation* $\chi|_{G_{K_\chi}} = 1$.

Proof. It is clear that the set $\{i_l \circ \chi_p|_{\ker(\chi)}\}_{p|\Delta_{K_\chi/\mathbb{Q}}}$ belongs to $\mathrm{Cl}(K_\chi)^\vee[1-\zeta_l]$, with any relation a multiple of the trivial one. This combined with Proposition 2.1 and Corollary 3.5 gives the conclusion by counting. ∎

It is possible to give a more direct proof of Proposition 3.6 by completely different considerations. This relies on the following fundamental fact that will anyway play a crucial role for us.

Proposition 3.7. *Let* F/E *be a cyclic degree* l *extension of number fields. Let* $v \in \Omega_E$ *be a place of* E *ramifying in* F/E. *Suppose moreover that* L/F *is an abelian extension of* F, *Galois over* E *and unramified at the unique place in* Ω_F *lying above* v. *Then the exact sequence*

$$1 \to \mathrm{Gal}(L/F) \to \mathrm{Gal}(L/E) \to \mathrm{Gal}(F/E) \to 1$$

splits, i.e. the surjection $\mathrm{Gal}(L/E) \twoheadrightarrow \mathrm{Gal}(F/E)$ *admits a section.*

Proof. For a number field F, we let Ω_F be the set of places of F. Since L/F is unramified at the unique place in Ω_F lying above v, and since v is totally ramified in F/E, we deduce that for every place $\tilde{v} \in \Omega_L$ lying above v in L/E, the inertia group $I_{\tilde{v}/v}$ has size exactly l. Therefore either it is fully contained in $\mathrm{Gal}(L/F)$ or it intersects $\mathrm{Gal}(L/F)$ trivially providing the claimed section.

We assume that $I_{\tilde{v}/v}$ is fully contained in $\mathrm{Gal}(L/F)$ and derive a contradiction. Since $\mathrm{Gal}(L/F)$ is a normal subgroup, it follows that all conjugates of $I_{\tilde{v}/v}$ are also contained in $\mathrm{Gal}(L/F)$. Then we conclude that

$$\mathrm{Gal}(L/F) \supseteq \prod_{w \in \Omega_L : w|v} I_{w/v},$$

which is equivalent to

$$F \subseteq \bigcap_{w \in \Omega_L : w|v} L^{I_{w/v}}.$$

This implies that v is unramified in F/E, which is the desired contradiction. ∎

Alternative proof of Proposition 3.6. Let L/K_χ be a degree l extension coming from a character in $\mathrm{Cl}(K_\chi)^\vee[1-\zeta_l]$. Note that L is a Galois extension of $\mathbb{Q}$ with degree l^2. By Proposition 3.7 we conclude that

$$\mathrm{Gal}(L/\mathbb{Q}) \simeq \mathbb{F}_l \times \mathbb{F}_l.$$

Hence the extension is of the shape KK_χ/K_χ, where K is a cyclic degree l extension of $\mathbb{Q}$. Let $\chi' \in \Gamma_{\mu_l}(\mathbb{Q})$ be a character with $K_{\chi'} = K$. If χ' ramifies at any prime q not dividing $\Delta_{K_\chi/\mathbb{Q}}$, then the resulting extension of K_χ will ramify at the primes of K_χ lying above q. Hence χ' must be one of the characters listed in Proposition 3.6, which clearly satisfy only the trivial relation stated there.

4. Central extensions

In this section we prove several important facts about central $\mathbb{F}_l$-extensions that we will extensively use in the coming sections to deal with the first and higher Artin pairings. The main tool established here is Theorem 4.10, which provides sufficient conditions to realize central embedding problems with as little ramification as possible.

Let E be a field and fix a separable closure E^{sep} of E. We extend the notation from Subsection 2.2 in the natural way to our more general setting; in particular we will use G_E and $\Gamma_{\mathbb{F}_l}(E)$ without further introduction. Let $F \subseteq E^{\mathrm{sep}}$ be a finite Galois extension of E. The most important object in this section is the set

$$\mathrm{Cent}_{\mathbb{F}_l}(F/E)$$

consisting of degree dividing l extensions $\tilde{F}/F$ in E^{sep} that are Galois over E and such that $\mathrm{Gal}(\tilde{F}/F)$ is a central subgroup of $\mathrm{Gal}(\tilde{F}/E)$. For $\tilde{F}_1, \tilde{F}_2$ in $\mathrm{Cent}_{\mathbb{F}_l}(F/E)$ we say that $\tilde{F}_1$ is equivalent to $\tilde{F}_2$ if at least one of the following two statements is true:

- $\tilde{F}_1 = \tilde{F}_2 = F$;
- the compositum $\tilde{F}_1\tilde{F}_2$ contains a non-trivial cyclic degree l extension of F that is obtained as $\tilde{E}F$, where $\tilde{E}$ is a cyclic degree l extension of E.

One can easily see that this is an equivalence relation, and we use the symbol $\tilde{F}_1 \sim \tilde{F}_2$ to express the fact that $\tilde{F}_1$ is equivalent to $\tilde{F}_2$.

Every cyclic extension of F with degree dividing l can be obtained as $(E^{\mathrm{sep}})^{\ker(\chi)}$ for some $\chi \in \Gamma_{\mathbb{F}_l}(F)$. This can be done only with the trivial character if the extension is trivial, and in all the other cases with precisely the $l-1$ non-zero multiples of a non-trivial character. It can be easily seen that for this extension to be in $\mathrm{Cent}_{\mathbb{F}_l}(F/E)$ it is necessary and sufficient that χ is fixed by the natural action of $\mathrm{Gal}(F/E)$ on $\Gamma_{\mathbb{F}_l}(F)$. Therefore we have a natural surjective map

$$\Gamma_{\mathbb{F}_l}(F)^{\mathrm{Gal}(F/E)} \twoheadrightarrow \mathrm{Cent}_{\mathbb{F}_l}(F/E),$$

which descends to a map

$$\frac{\Gamma_{\mathbb{F}_l}(F)^{\mathrm{Gal}(F/E)}}{\Gamma_{\mathbb{F}_l}(E)} \twoheadrightarrow \mathrm{Cent}_{\mathbb{F}_l}(F/E)/\sim.$$

The map attains the class of F/E precisely on the trivial element of $\frac{\Gamma_{\mathbb{F}_l}(F)^{\mathrm{Gal}(F/E)}}{\Gamma_{\mathbb{F}_l}(E)}$ and on the remaining points is a $(l-1):1$ assignment.

Given a class $\chi \in \frac{\Gamma_{\mathbb{F}_l}(F)^{\mathrm{Gal}(F/E)}}{\Gamma_{\mathbb{F}_l}(E)}$ we can naturally attach a class

$$r_1(\chi) \in H^2(\mathrm{Gal}(F/E), \mathbb{F}_l),$$

where the implicit action is declared to be trivial. This uses the group-theoretic interpretation of $H^2(\mathrm{Gal}(F/E), \mathbb{F}_l)$, and goes as follows. Using χ we have an identification

$$\mathrm{Gal}((E^{\mathrm{sep}})^{\ker(\chi)}/F) \simeq \mathbb{F}_l.$$

Therefore this transforms the sequence

$$1 \to \mathrm{Gal}((E^{\mathrm{sep}})^{\ker(\chi)}/F) \to \mathrm{Gal}((E^{\mathrm{sep}})^{\ker(\chi)}/E) \to \mathrm{Gal}(F/E) \to 1$$

into a sequence

$$0 \to \mathbb{F}_l \to \mathrm{Gal}((E^{\mathrm{sep}})^{\ker(\chi)}/E) \to \mathrm{Gal}(F/E) \to 1,$$

which naturally provides us with a class $r_1(\chi) \in H^2(\mathrm{Gal}(F/E), \mathbb{F}_l)$. It is not hard to show that the resulting map

$$r_1 : \frac{\Gamma_{\mathbb{F}_l}(F)^{\mathrm{Gal}(F/E)}}{\Gamma_{\mathbb{F}_l}(E)} \to H^2(\mathrm{Gal}(F/E), \mathbb{F}_l)$$

is an injective group homomorphism. Hence we can use r_1 to identify the group

$\frac{\Gamma_{\mathbb{F}_l}(F)^{\mathrm{Gal}(F/E)}}{\Gamma_{\mathbb{F}_l}(E)}$ with its image in $H^2(\mathrm{Gal}(F/E), \mathbb{F}_l)$. We define

$$\widetilde{\mathrm{Cent}}_{\mathbb{F}_l}(F/E) := \mathrm{Im}(r_1).$$

Our next goal is to characterize the elements of $H^2(\mathrm{Gal}(F/E), \mathbb{F}_l)$ that belong to $\widetilde{\mathrm{Cent}}_{\mathbb{F}_l}(F/E)$. To do so we write another natural map

$$r_2 : \frac{\Gamma_{\mathbb{F}_l}(F)^{\mathrm{Gal}(F/E)}}{\Gamma_{\mathbb{F}_l}(E)} \to H^2(\mathrm{Gal}(F/E), \mathbb{F}_l),$$

this time using Galois cohomology in the following manner. Firstly observe that $\Gamma_{\mathbb{F}_l}(F) = H^1(G_F, \mathbb{F}_l)$ and $\Gamma_{\mathbb{F}_l}(E) = H^1(G_E, \mathbb{F}_l)$, since the action of G_E on $\mathbb{F}_l$ has been declared to be trivial. In this manner the natural map coming from restriction $\Gamma_{\mathbb{F}_l}(E) \to \Gamma_{\mathbb{F}_l}(F)^{\mathrm{Gal}(F/E)}$ becomes the natural restriction homomorphism

$$\mathrm{Res} : H^1(G_E, \mathbb{F}_l) \to H^1(G_F, \mathbb{F}_l)^{\mathrm{Gal}(F/E)}.$$

Therefore the generalized inflation-restriction long exact sequence gives us a connecting homomorphism r_2, which provides a canonical isomorphism

$$r_2 : \frac{H^1(G_F, \mathbb{F}_l)^{\mathrm{Gal}(F/E)}}{H^1(G_E, \mathbb{F}_l)} \simeq \ker\big(\mathrm{Inf} : H^2(\mathrm{Gal}(F/E), \mathbb{F}_l) \to H^2(G_E, \mathbb{F}_l)\big).$$

The map Inf maps a 2-cocycle for $\mathrm{Gal}(F/E)$ to a 2-cocycle for G_E simply by precomposing the 2-cocycle with the projection of G_E onto $\mathrm{Gal}(F/E)$. So the kernel consists of the 2-cocycles $\theta \in H^2(\mathrm{Gal}(F/E), \mathbb{F}_l)$ for which there exists a continuous 1-cochain $\phi : G_E \to \mathbb{F}_l$ with $d(\phi) = \theta$.

We stress that this does not imply that θ is trivial as an element of $H^2(\mathrm{Gal}(F/E), \mathbb{F}_l)$, since the field of definition of ϕ need not be a subfield of F. We claim that r_1 and r_2 are actually the same map. Indeed, by the general formula for the connecting homomorphism for the inflation-restriction exact sequence, one finds that the map r_2 can be written as follows. For each element σ of $\mathrm{Gal}(F/E)$ fix a lift $\tilde{\sigma} \in G_E$ and take $\chi \in \Gamma_{\mathbb{F}_l}(F)^{\mathrm{Gal}(F/E)}$. Then $r_2(\chi)$ is represented in $H^2(\mathrm{Gal}(F/E), \mathbb{F}_l)$ by the cocycle

$$(\sigma_1, \sigma_2) \mapsto \chi(\widetilde{\sigma_1\sigma_2}\tilde{\sigma_2}^{-1}\tilde{\sigma_1}^{-1}).$$

It is a pleasant exercise to show directly, using that $\chi \in \Gamma_{\mathbb{F}_l}(F)^{\mathrm{Gal}(F/E)}$, that the choice of the lift only changes the expression by a coboundary. On the other hand, one readily sees that this is the same class as $r_1(\chi)$. Therefore we get the following fundamental fact.

Proposition 4.1. *We have*

$$\widetilde{\mathrm{Cent}}_{\mathbb{F}_l}(F/E) = \ker\big(H^2(\mathrm{Gal}(F/E), \mathbb{F}_l) \to H^2(G_E, \mathbb{F}_l)\big).$$

In other words, a class $\theta \in H^2(\mathrm{Gal}(F/E), \mathbb{F}_l)$ is equal to $r_1(\chi)$ for some $\chi \in \Gamma_{\mathbb{F}_l}(F)^{\mathrm{Gal}(F/E)}$ if and only if θ is trivial when viewed as a class in $H^2(G_E, \mathbb{F}_l)$. In this case, the set ${r_1}^{-1}(\theta)$ consists precisely of a single coset for the group $\Gamma_{\mathbb{F}_l}(E)$.

Loosely speaking, the above criterion tells us that the group-theoretic $\mathbb{F}_l$-extensions of $\mathrm{Gal}(F/E)$ that can be realized with a field extension are precisely equal to the classes of $H^2(\mathrm{Gal}(F/E), \mathbb{F}_l)$ that become trivial in $H^2(G_E, \mathbb{F}_l)$. This criterion takes an even simpler form if we assume that E contains an element $\zeta_l \in E^{\mathrm{sep}}$ of multiplicative order exactly l.

Recall from Section 2.2 that we identified $\mathbb{F}_l$ and $\langle \zeta_l \rangle$ with the isomorphism $j_l(a) = \zeta_l^a$ for each $a \in \mathbb{F}_l$. Observe that if $\zeta_l \in E$, then j_l is an identification of G_E-modules. In particular j_l induces an isomorphism

$$j_l : H^2(G_E, \mathbb{F}_l) \to \mathrm{Br}_E[l].$$

Also the group $\frac{\Gamma_{\mathbb{F}_l}(F)^{\mathrm{Gal}(F/E)}}{\Gamma_{\mathbb{F}_l}(E)}$ can be identified via j_l with

$$\frac{\left(\frac{F^*}{F^{*l}}\right)^{\mathrm{Gal}(F/E)}}{E^*},$$

and one can quickly re-obtain in this special case a proof of Proposition 4.1 by using Kummer sequences, which provide a natural identification

$$\frac{\left(\frac{F^*}{F^{*l}}\right)^{\mathrm{Gal}(F/E)}}{E^*} = \ker\left(H^2(\mathrm{Gal}(F/E), \mathbb{F}_l) \to \mathrm{Br}_E[l]\right).$$

It turns out that as long as $\mathrm{char}(E) \neq l$ we can always verify the realizability of a cohomology class θ in terms of the vanishing of a class attached to θ in a Brauer group.

Proposition 4.2. *Suppose F/E is a finite Galois extension with $\mathrm{char}(E) \neq l$. Then the group $\widetilde{\mathrm{Cent}}_{\mathbb{F}_l}(F/E)$ consists of those classes $\theta \in H^2(\mathrm{Gal}(F/E), \mathbb{F}_l)$ such that $j_l \circ \mathrm{Res}_{G_{E(\zeta_l)}}(\theta)$ is trivial in $\mathrm{Br}_{E(\zeta_l)}$.*

Proof. By Proposition 4.1 we know that θ is realizable if and only if θ becomes trivial in $H^2(G_E, \mathbb{F}_l)$. There is a map Co-Res : $H^2(G_{E(\zeta_l)}, \mathbb{F}_l) \to H^2(G_E, \mathbb{F}_l)$ with the property

$$\text{Co-Res} \circ \mathrm{Res}(\theta) = [E(\zeta_l) : E] \cdot \theta.$$

But $l \cdot \theta = 0$ and $[E(\zeta_l) : E]$ divides $|\mathrm{Aut}_{\mathrm{gr}}(\langle \zeta_l \rangle)| = l - 1$. Hence θ becomes trivial in $H^2(G_E, \mathbb{F}_l)$ if and only if $\mathrm{Res}_{G_{E(\zeta_l)}}(\theta)$ is trivial in $H^2(G_{E(\zeta_l)}, \mathbb{F}_l)$. Finally, since j_l is an isomorphism, the triviality of θ in $H^2(G_{E(\zeta_l)}, \mathbb{F}_l)$ is equivalent to the triviality of $j_l \circ \theta$ in $\mathrm{Br}_{E(\zeta_l)}$. ∎

In particular for $E = \mathbb{Q}$ we derive the following criterion. The group $\mathbb{F}_l$ will be implicitly considered as a trivial G-module whenever a symbol suggests an action of a group G on $\mathbb{F}_l$.

Proposition 4.3. *Let $F/\mathbb{Q}$ be a finite Galois extension. Then the following are equivalent for $\theta \in H^2(\mathrm{Gal}(F/\mathbb{Q}), \mathbb{F}_l)$:*

(1) *$\theta \in \widetilde{\mathrm{Cent}}_{\mathbb{F}_l}(F/\mathbb{Q})$, i.e. $\theta = r_1(\chi)$ for some $\chi \in \Gamma_{\mathbb{F}_l}(F)^{\mathrm{Gal}(F/\mathbb{Q})}$;*

(2) *there exists a continuous* 1*-cochain* $\phi : G_{\mathbb{Q}} \to \mathbb{F}_l$ *such that* $d(\phi) = \theta$*;*

(3) *for every* $v \in \Omega_{\mathbb{Q}(\zeta_l)}$ *we have* $\mathrm{inv}_v(j_l \circ \theta) = 0$.

Proof. This is an immediate consequence of Proposition 4.1, Proposition 4.2 and the fact that $\mathrm{Br}_{\mathbb{Q}(\zeta_l)}$ embeds in the direct sum $\bigoplus_{v \in \Omega_{\mathbb{Q}(\zeta_l)}} \mathrm{Br}_{\mathbb{Q}(\zeta_l)_v}$. ∎

The following fact will help us to cut down in practice the set of places one needs to check in part (3) of Proposition 4.3. The proposition could easily be derived from local class field theory. Instead we give an elementary proof here based on Proposition 4.2.

Proposition 4.4. *Let* E *be a local field and* F/E *a finite unramified extension. Then* $\widetilde{\mathrm{Cent}}_{\mathbb{F}_l}(F/E) = H^2(\mathrm{Gal}(F/E), \mathbb{F}_l)$*, i.e. each* $\theta \in H^2(\mathrm{Gal}(F/E), \mathbb{F}_l)$ *is trivial in* $H^2(G_E, \mathbb{F}_l)$*. In particular, if* $\mu_l(E) \neq \{1\}$*, then* $\mathrm{inv}_E(j_l \circ \theta) = 0$ *for each* $\theta \in H^2(\mathrm{Gal}(F/E), \mathbb{F}_l)$.

Proof. Since $\mathrm{Gal}(F/E)$ is cyclic, $H^2(\mathrm{Gal}(F/E), \mathbb{F}_l)$ equals $\mathrm{Ext}(\mathrm{Gal}(F/E), \mathbb{F}_l)$. Indeed, if G is a group with a central subgroup H and cyclic quotient G/H, then G is abelian. We can assume that l divides $[F : E]$ otherwise the H^2 collapses and the statement becomes a triviality. In this case $\mathrm{Ext}(\mathrm{Gal}(F/E), \mathbb{F}_l)$ is cyclic of order l. Therefore it is enough to provide one non-trivial element of $\mathrm{Cent}_{\mathbb{F}_l}(F/E)$ for which we take the unramified degree l extension of F. The other statements now follow from Propositions 4.1 and 4.2. ∎

Thus, in practice, when we use Proposition 4.3, it is enough to check at the places $v \in \Omega_{\mathbb{Q}(\zeta_l)}$ that ramify in $F(\zeta_l)/\mathbb{Q}(\zeta_l)$. The following general fact explains why we need only deal with elements of $\Omega_{\mathbb{Q}(\zeta_l)}$ with residue field degree 1.

Proposition 4.5. *Let* $F/\mathbb{Q}$ *a finite Galois extension of degree a power of* l*. If a prime* q *divides* $\Delta_{F/\mathbb{Q}}$ *then* $q = l$ *or* $q \equiv 1 \bmod l$.

Proof. Let q be a prime different from l that ramifies in $F/\mathbb{Q}$ and choose some inertia subgroup $I_q \leq \mathrm{Gal}(F/\mathbb{Q})$ at q. Observe that since $q \neq l$ it must be that $\gcd(|I_q|, q) = 1$. Hence $|I_q|$ divides $q^{f_q(F/\mathbb{Q})} - 1$. On the other hand, $f_q(F/\mathbb{Q})$ divides $[F : \mathbb{Q}]$, which is a power of l. We conclude that $q^{f_q(F/\mathbb{Q})} - 1 \equiv q - 1 \bmod l$. Because $|I_q|$ is a non-trivial power of l, this implies that q is 1 modulo l. ∎

Recall that if we have two conjugate subgroups H_1, H_2 of $G_{\mathbb{Q}}$ and a class θ in $H^2(G_{\mathbb{Q}}, \mathbb{F}_l)$ then $\mathrm{Res}_{H_1}(\theta) = 0$ if and only if $\mathrm{Res}_{H_2}(\theta) = 0$. Thus the following definition makes sense.

Definition 4.6. Let $v \in \Omega_{\mathbb{Q}}$ be a place and θ be a class in $H^2(G_{\mathbb{Q}}, \mathbb{F}_l)$. We say that θ is *locally trivial* at v if $\mathrm{Res}_{i^*(G_{\mathbb{Q}_v})}(\theta) = 0$ for some (equivalently any) choice of an embedding $i : \overline{\mathbb{Q}} \to \overline{\mathbb{Q}_v}$.

Proposition 4.7. *Let* $F/\mathbb{Q}$ *be a finite Galois extension. Let* θ *be a class in* $H^2(\mathrm{Gal}(F/\mathbb{Q}), \mathbb{F}_l)$*. Then the following are equivalent:*

(1) $\theta \in \widetilde{\mathrm{Cent}}_{\mathbb{F}_l}(F/\mathbb{Q})$;

(2) *the inflation of* θ *to* $G_{\mathbb{Q}}$ *is locally trivial at all places* $v \in \Omega_{\mathbb{Q}}$;

(3) *the inflation of* θ *to* $G_{\mathbb{Q}}$ *is locally trivial at all places* $v \in \Omega_{\mathbb{Q}}$ *that ramify in* $F/\mathbb{Q}$;

(4) *for any* $v \in \Omega_{\mathbb{Q}}$ *and* $\tilde{v} \in \Omega_F$ *lying above* v, *we have* $\theta \in \widetilde{\mathrm{Cent}}_{\mathbb{F}_l}(F_{\tilde{v}}/\mathbb{Q}_v)$.

Proof. Parts (2) and (4) are equivalent in view of Proposition 4.1. Thanks to the same proposition, (1) certainly implies (2). Furthermore, (2) trivially implies (3). On the other hand, thanks to Proposition 4.4, (3) implies (2) as well. It remains to show that (2) implies (1). But (2) implies that part (3) of Proposition 4.3 holds. Now use Proposition 4.3. ∎

The equivalence between (1) and (4) in Proposition 4.7 tells us that a class θ is realizable if and only if it is realizable locally everywhere, and moreover it is sufficient to check that it is realizable locally at the places of $\mathbb{Q}$ that are ramified in F.

In our applications, we will not merely be interested in writing the relevant θ as $r_1(\chi)$ for some $\chi \in \Gamma_{\mathbb{F}_l}(F)^{\mathrm{Gal}(F/\mathbb{Q})}$, but it will also be important for us to find a representative χ in the $\Gamma_{\mathbb{F}_l}(\mathbb{Q})$-coset $r_1^{-1}(\theta)$ such that $\overline{\mathbb{Q}}^{\ker(\chi)}/\mathbb{Q}$ has as little ramification as possible.

Proposition 4.8. *Let* $F/\mathbb{Q}$ *a finite Galois extension and let* $\theta \in \widetilde{\mathrm{Cent}}_{\mathbb{F}_l}(F/\mathbb{Q})$. *Then there exists* $\chi \in \Gamma_{\mathbb{F}_l}(F)^{\mathrm{Gal}(F/\mathbb{Q})}$ *such that* $r_1(\chi) = \theta$ *and* $\overline{\mathbb{Q}}^{\ker(\chi)}/\mathbb{Q}$ *is unramified at all primes* q *not dividing* $\Delta_{F/\mathbb{Q}}$.

Proof. Take any $\chi \in \Gamma_{\mathbb{F}_l}(F)^{\mathrm{Gal}(F/\mathbb{Q})}$ with $r_1(\chi) = \theta$. If $\overline{\mathbb{Q}}^{\ker(\chi)}/\mathbb{Q}$ is unramified at all primes q not dividing $\Delta_{F/\mathbb{Q}}$, we are done. So suppose that there is a rational prime q that does not ramify in F but does ramify in $\overline{\mathbb{Q}}^{\ker(\chi)}$. Then there is a prime $\mathfrak{q}$ of O_F above q that ramifies in $\overline{\mathbb{Q}}^{\ker(\chi)}$.

As observed in the proof of Proposition 4.4, the group $H^2(\mathrm{Gal}(F_{\mathfrak{q}}/\mathbb{Q}_q), \mathbb{F}_l)$ is cyclic of order l, generated by an unramified character of $G_{F_{\mathfrak{q}}}$ of order l. This means that we can always find $\chi' \in G_{\mathbb{Q}_q}$ of order l such that $\chi + \chi'$ is an unramified character for $G_{F_{\mathfrak{q}}}$. Moreover, we can take χ' to be a multiple of χ_q (see Subsection 2.2 for the notation). So we can find $\chi' \in \Gamma_{\mathbb{F}_l}(\mathbb{Q})$ such that $\mathfrak{q}$ does not ramify in $\overline{\mathbb{Q}}^{\ker(\chi+\chi')}$.

We claim that this implies that $\chi + \chi'$ does not ramify at *any* prime above q. Indeed, for each $\sigma \in G_{\mathbb{Q}}$, the character $\sigma \cdot (\chi + \chi')$ is certainly unramified at $\sigma(\mathfrak{q})$. On the other hand, $\sigma \cdot (\chi + \chi') = \chi + \chi'$, since by assumption $\chi \in \Gamma_{\mathbb{F}_l}(F)^{\mathrm{Gal}(F/\mathbb{Q})}$. This proves our claim. Finally, observe that $\overline{\mathbb{Q}}^{\ker(\chi+\chi')}/F$ does not ramify at any new prime, since χ_q ramifies only at q. Hence continuing in this manner we get rid of all such q and we have proved the proposition. ∎

A stronger control on the ramification can be achieved at the cost of having a stronger notion of local triviality, which will be given in the next definition. Recall again that if H_1, H_2 are conjugate subgroups of a finite group G and if $\theta \in H^2(G, \mathbb{F}_l)$, then $\mathrm{Res}_{H_1}(\theta) = 0$ if and only if $\mathrm{Res}_{H_2}(\theta) = 0$. This shows that the following definition makes sense.

Definition 4.9. Let $F/\mathbb{Q}$ be a finite Galois extension, $\theta \in H^2(\mathrm{Gal}(F/\mathbb{Q}), \mathbb{F}_l)$ and q a prime number. We say that θ is *locally split at* q if the restriction of θ to one (equivalently any) subgroup $D_{\mathfrak{q}/q}$ is trivial, where $\mathfrak{q}$ is a prime above q in F and $D_{\mathfrak{q}/q}$ is the corresponding decomposition group. Moreover we say that θ is *locally split at inertia at* q if the restriction of θ to some (equivalently any) subgroup $I_{\mathfrak{q}/q}$ is trivial, where $I_{\mathfrak{q}/q}$ denotes the inertia subgroup relative to $\mathfrak{q}$.

Theorem 4.10. (1) *Let $F/\mathbb{Q}$ be a finite Galois extension. If $\theta \in H^2(\mathrm{Gal}(F/\mathbb{Q}), \mathbb{F}_l)$ is locally split at all primes dividing $\Delta_{F/\mathbb{Q}}$, then $\theta \in \widetilde{\mathrm{Cent}}_{\mathbb{F}_l}(F/\mathbb{Q})$. Moreover, there exists $\chi \in \Gamma_{\mathbb{F}_l}(F)^{\mathrm{Gal}(F/\mathbb{Q})}$ such that $r_1(\chi) = \theta$ and $\overline{\mathbb{Q}}^{\ker(\chi)}/F$ is unramified.*

(2) *Suppose that all the primes dividing $\Delta_{F/\mathbb{Q}}$ are 1 modulo l. Then the same conclusion as in* (1) *can be reached assuming only that θ is locally trivial and locally split at inertia at all primes dividing $\Delta_{F/\mathbb{Q}}$.*

(3) *Suppose that $[F : \mathbb{Q}]$ is of degree a power of l and does not ramify at l. Then the same conclusion as in* (1) *can be reached assuming only that θ is locally trivial and locally split at inertia at all primes dividing $\Delta_{F/\mathbb{Q}}$.*

Proof of (1). Observe that if θ is locally split at a prime q, then it is certainly also locally trivial at q. Hence by Proposition 4.7 we deduce that $\theta \in \widetilde{\mathrm{Cent}}_{\mathbb{F}_l}(F/\mathbb{Q})$. Due to Proposition 4.8 there is $\chi \in \Gamma_{\mathbb{F}_l}(F)^{\mathrm{Gal}(F/\mathbb{Q})}$ with $\theta = r_1(\chi)$ such that $\overline{\mathbb{Q}}^{\ker(\chi)}/F$ is unramified at all primes $\mathfrak{q}$ in O_F that lie above primes q in $\mathbb{Z}$ not dividing $\Delta_{F/\mathbb{Q}}$.

Now take a prime q dividing $\Delta_{F/\mathbb{Q}}$ and let $\mathfrak{q}$ in O_F be a prime above q that ramifies in $\overline{\mathbb{Q}}^{\ker(\chi)}$. By assumption θ is locally split at each such prime q. Hence locally at each such q the character χ is a character from $G_{\mathbb{Q}_q}$. Following the logic of the proof of Proposition 4.8, we may employ multiples of χ_q to get rid of this additional ramification whenever that is required. ∎

Proof of (2). The assumption that θ is locally trivial at all primes dividing $\Delta_{F/\mathbb{Q}}$ guarantees that $\theta \in \widetilde{\mathrm{Cent}}_{\mathbb{F}_l}(F/\mathbb{Q})$ due to Proposition 4.7. Again, by Proposition 4.8, write $\theta = r_1(\chi)$ with $\overline{\mathbb{Q}}^{\ker(\chi)}/F$ unramified at all primes of O_F above a rational prime not dividing the discriminant.

Let q be a prime divisor of $\Delta_{F/\mathbb{Q}}$ and let $\mathfrak{q}$ be a prime above it in O_F. Let $F_{\mathfrak{q}}^{\mathrm{unr}}$ be the largest unramified extension inside $F_{\mathfrak{q}}/\mathbb{Q}_q$. This is precisely the field fixed by the inertia subgroup $I_{\mathfrak{q}/q}$. The assumption that θ is locally split at inertia at q guarantees precisely that χ restricted to any copy of $G_{F_{\mathfrak{q}}}$ in G_F equals the restriction of a character coming from $G_{F_{\mathfrak{q}}^{\mathrm{unr}}}$. Since q is 1 modulo l, any such character equals the product of a multiple of χ_q and an unramified character. Hence we can use the same logic as in part (1). ∎

Proof of (3). This follows from part (2) and Proposition 4.5. ∎

We state the following simple fact, which is a consequence of elementary properties of $H^2(\mathbb{F}_l^2, \mathbb{F}_l)$. Fix a field K and a separable closure K^{sep}. Denote by G_K the group of K-algebra automorphisms of K^{sep}. For a continuous character $\chi : G_K \to \mathbb{F}_l$, let $K(\chi)$ be the corresponding field extension of K.

Proposition 4.11. *Let χ_1, χ_2 be two independent continuous characters from G_K to $\mathbb{F}_l$. Then $\chi_1 \cup \chi_2$ is in $\widetilde{\mathrm{Cent}}_{\mathbb{F}_l}(K(\chi_1)K(\chi_2)/K)$ if and only if there exists a Galois extension L/K containing $K(\chi_1)K(\chi_2)$ such that*

$$\mathrm{Gal}(L/K) \simeq_{\mathrm{gr.}} \frac{\mathbb{Z}_l[\zeta_l]}{(1-\zeta_l)^2} \rtimes \langle \zeta_l \rangle.$$

We end this section by mentioning the following fact concerning extensions having Galois group $\frac{\mathbb{Z}_l[\zeta_l]}{(1-\zeta_l)^2} \rtimes \langle \zeta_l \rangle$ that can be used to prove Corollary A.4 for fields of characteristic different from l. One way to prove this fact is to use the material in this section. The interested reader can also look at [18, Theorem 3.1].

Proposition 4.12. *Suppose K has a primitive l-th root of unity. If $b \in K^*$, we define $\chi_b : G_K \to \mathbb{F}_l$ to be the unique character such that for each $\beta \in K^{\mathrm{sep}}$ with $\beta^l = b$ we have*

$$\sigma(\beta) = (j_l \circ \chi_b(\sigma))\beta.$$

Let $b_1, b_2 \in K^*$ be independent in K^*/K^{*l}. Then

$$\chi_{b_1} \cup \chi_{b_2} \in \widetilde{\mathrm{Cent}}_{\mathbb{F}_l}(K(\chi_{b_1})K(\chi_{b_2})/K) \iff \exists \omega \in K(\chi_{b_1})^* : b_2 = N_{K_{\chi_{b_1}}/K}(\omega).$$

In that case the image of $\chi_{b_1} \cup \chi_{b_2}$ in $\mathrm{Cent}_{\mathbb{F}_l}(K(\chi_{b_1})K(\chi_{b_2})/K)$ is obtained by taking the extension $K(\chi_{b_1})K(\chi_{b_2})(\sqrt[l]{\alpha})$ with

$$\alpha := \prod_{i=0}^{l-2} \sigma^i(\omega^{l-1-i}),$$

where σ is a generator of $\mathrm{Gal}(K(\chi_{b_1})/K)$.

5. The first Artin pairing

In this section we study the first Artin pairing. This culminates in a description of the spaces $(1-\zeta_l)\mathrm{Cl}(K_\chi)[(1-\zeta_l)^2]$ and $(1-\zeta_l)\mathrm{Cl}(K_\chi)^\vee[(1-\zeta_l)^2]$ given respectively in Corollary 5.5 and Corollary 5.6.

Let χ be in $\Gamma_{\mu_l}(\mathbb{Q})$ and $b \in \overline{\mathrm{Cl}}(K_\chi)$. We extend the notation introduced in Section 2.2 by defining $\mathrm{Up}_{K_\chi}(b)$ to be the unique product of ramified prime ideals of K_χ whose norm is precisely b. We will sometimes attribute properties of $\mathrm{Up}_{K_\chi}(b)$ to $b \in \overline{\mathrm{Cl}}(K_\chi)$. For example, we shall often say b is in $(1-\zeta_l)^k\mathrm{Cl}(K_\chi)$ for some positive integer k, which means that $\mathrm{Up}_{K_\chi}(b)$ is in $(1-\zeta_l)^k\mathrm{Cl}(K_\chi)$.

From the description of $\mathrm{Cl}(K_\chi)[1-\zeta_l]$ and $\mathrm{Cl}(K_\chi)^\vee[1-\zeta_l]$ combined with Proposition 2.2, we readily obtain the following description of $(1-\zeta_l)\mathrm{Cl}(K_\chi)[(1-\zeta_l)^2]$ and of $(1-\zeta_l)\mathrm{Cl}(K_\chi)^\vee[(1-\zeta_l)^2]$.

Proposition 5.1. (1) *An element $b \in \overline{\mathrm{Cl}}(K_\chi)$ is in $(1-\zeta_l)\mathrm{Cl}(K_\chi)[(1-\zeta_l)^2]$ if and only if for every prime q dividing $\Delta_{K_\chi/\mathbb{Q}}$, the Artin symbol*

$$\left[\frac{K_{\chi_q}K_\chi/K_\chi}{\mathrm{Up}_{K_\chi}(b)}\right]$$

is the identity.

(2) *A character $\chi' \in \mathrm{Cl}(K_\chi)^\vee[1-\zeta_l]$ is in $(1-\zeta_l)\mathrm{Cl}(K_\chi)^\vee[(1-\zeta_l)^2]$ if and only if for every prime q dividing $\Delta_{K_\chi/\mathbb{Q}}$,*

$$\chi'(\mathrm{Frob}_{\mathrm{Up}_{K_\chi}(q)}) = 1.$$

Proof. This follows immediately from the material in Subsection 3.1, Subsection 3.2 and Proposition 2.2. ∎

Proposition 5.2. *Let χ be in $\Gamma_{\mu_l}(\mathbb{Q})$. Suppose that q and q' are two distinct primes dividing $\Delta_{K_\chi/\mathbb{Q}}$. Then*

$$\chi_q(\mathrm{Frob}_{\mathrm{Up}_{K_\chi}(q')}) = \chi_q(\mathrm{Frob}_{q'}).$$

Proof. Observe that q' splits in K_{χ_q} if and only if $\mathrm{Up}_{K_\chi}(q')$ splits in $K_\chi K_{\chi_q}$. Hence we can safely assume that they are both non-split. In that case, since $q \neq q'$, they must be both inert. It follows from the defining property of Frobenius that restricting $\mathrm{Frob}_{\mathrm{Up}_{K_\chi}(q')}$ to $K_{\chi_q}/\mathbb{Q}$ gives $\mathrm{Frob}_{q'}$. Hence we obtain the desired conclusion. ∎

Recall that if χ is in $\Gamma_{\mu_l}(\mathbb{Q})$, then ϵ_χ denotes the unique amalgama for $\Delta_{K_\chi/\mathbb{Q}}$ with the property

$$\chi = \prod_{q|\Delta_{K_\chi/\mathbb{Q}}} \chi_q^{\epsilon_\chi(q)}.$$

We will now define an $\omega(\Delta_{K_\chi/\mathbb{Q}}) \times \omega(\Delta_{K_\chi/\mathbb{Q}})$ matrix with coefficients in $\mathbb{F}_l$. We index the matrix by the product set $\{q \text{ prime} : q \mid \Delta_{K_\chi/\mathbb{Q}}\} \times \{q \text{ prime} : q \mid \Delta_{K_\chi/\mathbb{Q}}\}$, where in the row $\{q\} \times \{q' \text{ prime} : q' \mid \Delta_{K_\chi/\mathbb{Q}}\}$ we put for each $q \neq q'$ the element

$$j_l^{-1} \circ \chi_{q'}^{\epsilon_\chi(q')}(\mathrm{Frob}_q),$$

and we impose that the sum on each row is 0. This uniquely determines the so-called Rédei matrix that we denote as

$$\text{Rédei}(K_\chi) \in \mathbb{F}_l^{\{q \text{ prime} : q|\Delta_{K_\chi/\mathbb{Q}}\} \times \{q \text{ prime} : q|\Delta_{K_\chi/\mathbb{Q}}\}}.$$

In what follows, exponentiation by an element v of $\mathbb{F}_l$ has to be read as the conventional powering with the only integer between $\{0, \ldots, l-1\}$ that is congruent to v modulo l. Then we have the following important conclusion.

Corollary 5.3. (1) *The elements b of $\overline{\mathrm{Cl}}(K_\chi)$ that are in $(1-\zeta_l)\mathrm{Cl}(K_\chi)[(1-\zeta_l)^2]$ are precisely those b such that*

$$\mathrm{Up}_{K_\chi}(b) = \prod_{q|\Delta_{K_\chi/\mathbb{Q}}} \mathrm{Up}_{K_\chi}(q)^{v_q}$$

for $(v_q)_{q|\Delta_{K_\chi/\mathbb{Q}}}$ an element of the left kernel of $\mathrm{Rédei}(K_\chi)$.

(2) *The elements χ' of $\mathrm{Cl}(K_\chi)^\vee[(1-\zeta_l)]$ that are in $(1-\zeta_l)\mathrm{Cl}(K_\chi)^\vee[(1-\zeta_l)^2]$ are precisely those χ' such that*

$$\chi' = \prod_{q|\Delta_{K_\chi/\mathbb{Q}}} \chi_q^{w_q \epsilon_\chi(q)},$$

where $(w_q)_{q|\Delta_{K_\chi/\mathbb{Q}}}$ is in the right kernel of $\mathrm{Rédei}(K_\chi)$.

Proof. This follows upon combining Propositions 5.1 and 5.2. ∎

In Subsections 5.1 and 5.2 we investigate more closely the structure of respectively the left and the right kernel of the Rédei matrix. The resulting characterizations are contained in Corollaries 5.5 and 5.6. In these subsections we additionally provide alternative, and more direct, proofs of these corollaries. The material in Subsection 5.2 relies on the material in Section 4 about central $\mathbb{F}_l$-extensions.

5.1. $(1-\zeta_l)\mathrm{Cl}(K_\chi)[(1-\zeta_l)^2]$

Let $\chi \in \Gamma_{\mu_l}(\mathbb{Q})$. We begin by rewriting $\mathrm{Rédei}(K_\chi)$ as a matrix of symbols coming from cyclic algebras. See Appendix A for the notation and the basic facts used from the theory of such algebras over local and global fields. We use the convention that $A(i,j)$ denotes the element on the i-th row and j-th column of a matrix A.

Proposition 5.4. *For all primes q, q' dividing $\Delta_{K_\chi/\mathbb{Q}}$ we have*

$$\mathrm{Rédei}(K_\chi)(q,q') = j_l^{-1} \circ h_l \circ \eta_{\mathbb{Q}_q}(\{\chi_{q'}^{\epsilon_\chi(q')}, q\}).$$

Proof. This follows immediately from Proposition A.1 combined with the definition of the Rédei matrix $\mathrm{Rédei}(K_\chi)$ and the bilinearity of $(\chi,\theta) \mapsto \{\chi,\theta\}$. ∎

The following important corollary furnishes an interpretation of the left kernel that will be crucial in handling the higher pairings as we shall see in the later sections.

Corollary 5.5. *An element b in $\overline{\mathrm{Cl}}(K_\chi)$ is in $(1-\zeta_l)\mathrm{Cl}(K_\chi)[(1-\zeta_l)^2]$ if and only if b is a norm in K_χ.*

Proof. Fix a prime divisor q of $\Delta_{K_\chi/\mathbb{Q}}$. For now assume that q does not divide b. Then the pairing of χ_q and b is trivial if and only if

$$\prod_{p|b} \eta_{\mathbb{Q}_p}(\{\chi_q, b\}) = 1.$$

Indeed, this follows from Proposition 5.4 and the fact that the result in Proposition A.1 is independent of the choice of uniformizer; the latter observation also follows from a combination of Propositions 4.4 and A.3. From these propositions, we see that the only other place in $\Omega_\mathbb{Q}$ where the cyclic algebra $\{\chi_q, b\}$ could possibly be non-trivial is q. Therefore by Proposition A.2 (Hilbert reciprocity), we learn that

$$\eta_{\mathbb{Q}_q}(\{\chi_q, b\}) = 1.$$

This implies that

$$\eta_{\mathbb{Q}_q}(\{\chi, b\}) = 1$$

for each q that does not divide b by Proposition 4.4. Next assume that q divides b. Denote

$$b' := b/q^{v_{\mathbb{Q}_q}(b)}.$$

From the definition of the pairing we see that χ_q and b have trivial pairing if and only if $\chi_q^{\epsilon_\chi(q)}$ and b have trivial pairing. This is equivalent to

$$\prod_{p|b'} \eta_{\mathbb{Q}_p}(\{\chi_q^{\epsilon_\chi(q)}, b\}) \cdot \Big(\prod_{\substack{p|\Delta_{K_\chi/\mathbb{Q}} \\ p\neq q}} \eta_{\mathbb{Q}_q}(\{\chi_p^{\epsilon_\chi(p)}, b\}) \Big)^{-1} = 1.$$

By Proposition A.2 (Hilbert reciprocity) applied to the first factor, this happens if and only if $\eta_{\mathbb{Q}_q}(\{\chi, b\}) = 1$. Hence the statement that b pairs trivially with all the χ_q is equivalent to

$$\eta_{\mathbb{Q}_q}(\{\chi, b\}) = 1 \quad \text{for each prime divisor } q \text{ of } \Delta_{K_\chi/\mathbb{Q}},$$

which is in turn equivalent to

$$\eta_{\mathbb{Q}_v}(\{\chi, b\}) = 1 \quad \text{for all } v \in \Omega_\mathbb{Q}.$$

Therefore, from Proposition A.2 again, we see that $\mathrm{Up}_{K_\chi}(b)$ is a multiple of $1 - \zeta_l$ in the class group if and only if the cyclic algebra $\{\chi, b\}$ is trivial. Thanks to Corollary A.4, this is equivalent to b being a norm in K_χ, which is precisely the desired statement. ∎

We remark that Corollary 5.5 can be proved directly without the detour through Rédei matrices and class field theory.

5.2. $(1-\zeta_l)\mathrm{Cl}(K_\chi)^\vee[(1-\zeta_l)^2]$

Let $\chi \in \Gamma_{\mu_l}(\mathbb{Q})$. The following characterization of the right kernel of $\mathrm{Rédei}(K_\chi)$ follows easily from Proposition 5.4; it can be proved with an argument analogous to the one given in the proof of Proposition 5.5. Instead, in the rest of this section, we shall opt for a different argument relying on central $\mathbb{F}_l$-extensions. Recall that we identify $\mathbb{F}_l \otimes \mathbb{F}_l$ with $\mathbb{F}_l$ through the map $a \otimes b \mapsto a \cdot b$, where the product is in $\mathbb{F}_l$. This allows us to view the cup of two 1-cocycles in $\mathbb{F}_l$ as a 2-cocycle with values in $\mathbb{F}_l$.

Corollary 5.6. *A character $\chi' \in \mathrm{Cl}(K_\chi)^\vee[1-\zeta_l]$ is in $(1-\zeta_l)\mathrm{Cl}(K_\chi)^\vee[(1-\zeta_l)^2]$ if and only if $((i_l \circ j_l)^{-1} \circ \chi') \cup (j_l^{-1} \circ \chi)$ is trivial in $H^2(G_\mathbb{Q}, \mathbb{F}_l)$.*

Proof. Observe that, thanks to Proposition 4.1, we have

$$((i_l \circ j_l)^{-1} \circ \chi') \cup (j_l^{-1} \circ \chi) = 0 \quad \text{in } H^2(G_\mathbb{Q}, \mathbb{F}_l)$$

if and only if

$$[((i_l \circ j_l)^{-1} \circ \chi') \cup (j_l^{-1} \circ \chi)]_{H^2(\mathrm{Gal}(K_\chi K_{\chi'}/\mathbb{Q}), \mathbb{F}_l)} \in \widetilde{\mathrm{Cent}}_{\mathbb{F}_l}(K_\chi K_{\chi'}/\mathbb{Q}).$$

A straightforward local computation shows that for the primes q dividing $\Delta_{K_\chi/\mathbb{Q}}$ the 2-cocycle

$$h(\chi, \chi') := ((i_l \circ j_l)^{-1} \circ \chi') \cup (j_l^{-1} \circ \chi)$$

is locally trivial if and only if the two characters $(i_l \circ j_l)^{-1} \circ \chi'$ and $j_l^{-1} \circ \chi$ are locally linearly dependent at each such q. In other words, $h(\chi, \chi')$ is locally trivial if and only if it is locally split at all primes dividing the discriminant. Therefore, we conclude by Theorem 4.10 that

$$[h(\chi, \chi')]_{H^2(G_\mathbb{Q}, \mathbb{F}_l)} = 0 \quad \text{implies} \quad \chi' \in (1-\zeta_l)\mathrm{Cl}(K_\chi)[(1-\zeta_l)^2].$$

The converse follows immediately from a combination of Propositions 4.11 and 3.7. ∎

6. Raw cocycles

Let A be a finite $\mathbb{Z}_l[\zeta_l]$-module and let the pair (G, χ) consist respectively of a finite cyclic group of size l and of an isomorphism $\chi : G \to \langle \zeta_l \rangle$. Using χ we turn A into a G-module killed by the norm operator $N_G := \sum_{g \in G} g \in \mathbb{Z}_l[G]$. We write $A \rtimes G$ for the semidirect product with respect to this action. Similarly, through χ, we can view N, introduced in Section 2, as a G-module and we use the symbol $N(\chi)$ to denote the implicit G-module structure, which will play a central role for us.

Whenever we have a quotient map $\widetilde{G} \to G$, we will talk by abuse of language of $N(\chi)$ as a $\widetilde{G}$-module with the induced action. For any profinite group H and any discrete H-module B we denote by $\mathrm{Cocy}(H, B)$ the group of continuous 1-cocycles from H to B. Recall that if $K \leq H$ is a subgroup and $\psi : H \to B$ is an element of $\mathrm{Cocy}(H, B)$, then we call the element $\psi|_K \in \mathrm{Cocy}(K, B)$ the *restriction* of ψ to K. If we have a surjective homomorphism $\pi : \widetilde{H} \twoheadrightarrow H$, then we call the element $\psi \circ \pi \in \mathrm{Cocy}(\widetilde{H}, B)$ the *inflation* of ψ to $\widetilde{H}$.

Proposition 6.1. *Let (G, χ) and A be as above and k a positive integer. Inflation and restriction of 1-cocycles induce a split exact sequence of $\mathbb{Z}_l[\zeta_l]$-modules*

$$0 \to \mathrm{Cocy}(G, N(\chi))[(1-\zeta_l)^k] \to \mathrm{Cocy}(A \rtimes G, N(\chi))[(1-\zeta_l)^k] \to A^\vee[(1-\zeta_l)^k] \to 0.$$

Moreover, $\mathrm{Cocy}(G, N(\chi))[(1-\zeta_l)^k] \simeq_{\mathbb{Z}_l[\zeta_l]} N[(1-\zeta_l)^k]$*, with such an isomorphism arising from the evaluation map $\psi \mapsto \psi(\sigma)$ for σ a fixed non-trivial element of G.*

Proof. Let σ denote a generator of G. Since G is cyclic, we find that the map from $\mathrm{Cocy}(G, N(\chi))[(1-\zeta_l)^k]$ to $N[(1-\zeta_l)^k]$ sending ψ to $\psi(\sigma)$ induces an isomorphism between $\mathrm{Cocy}(G, N(\chi))[(1-\zeta_l)^k]$ and the kernel of the norm Norm_G operating on $N(\chi)[(1-\zeta_l)^k]$, which is the full $N(\chi)[(1-\zeta_l)^k]$. Hence the evaluation map $\psi \mapsto \psi(\sigma)$ induces an isomorphism of $\mathbb{Z}_l[\zeta_l]$-modules between $\mathrm{Cocy}(G, N(\chi))[(1-\zeta_l)^k]$ and $N(\chi)[(1-\zeta_l)^k]$ as claimed.

This implies that the natural inclusion

$$0 \to \mathrm{Cocy}(G, N(\chi))[(1-\zeta_l)^k] \to \mathrm{Cocy}(A \rtimes G, N(\chi))[(1-\zeta_l)^k],$$

induced by inflation of cocycles, is split, since $\mathrm{Cocy}(A \rtimes G, N(\chi))[(1-\zeta_l)^k]$ is finite and certainly killed by $(1-\zeta_l)^k$. Therefore it remains to show that the natural map, induced by restriction of cocycles,

$$\mathrm{Cocy}(A \rtimes G, N(\chi))[(1-\zeta_l)^k] \to A^\vee[(1-\zeta_l)^k],$$

is surjective. Let $\phi \in A^\vee[(1-\zeta_l)^k]$. Consider the map

$$\psi : A \rtimes G \to N(\chi)[(1-\zeta_l)^k]$$

which sends (a, g) to $\phi(a)$. By construction ψ restricts to ϕ. Hence we have to check that ψ is a 1-cocycle. By definition of semidirect product and of the map ψ, we have

$$\psi((a_1, g_1)(a_2, g_2)) = \psi(a_1 + \chi(g_1)a_2, g_1 g_2) = \phi(a_1) + \chi(g_1)\phi(a_2).$$

On the other hand, in order for ψ to be a 1-cocycle for the action of G on A it must satisfy

$$\psi((a_1, g_1)(a_2, g_2)) = \chi(g_1)\psi(a_2, g_2) + \psi(a_1, g_1) = \chi(g_1)\phi(a_2) + \phi(a_1),$$

precisely the same equation as above. This concludes our proof. ∎

The next proposition provides, roughly speaking, a converse to Proposition 6.1. This will be useful later: it will tell us that the Galois group of the field of definition of a 1-cocycle in $N(\chi)$ always splits as a semidirect product.

Proposition 6.2. *Let (G, χ) be as above. Let $\pi : \widetilde{G} \twoheadrightarrow G$ be a surjective homomorphism and moreover let $\psi \in \mathrm{Cocy}(\widetilde{G}, N(\chi))[(1-\zeta_l)^k]$ be a cocycle such that the image of the character $\psi|_{\ker(\pi)}$ is $N(\chi)[(1-\zeta_l)^k]$. Set $H := \{g \in \ker(\pi) : \psi(g) = 0\}$. We have the following facts.*

(1) *The set H is a normal subgroup of $\widetilde{G}$.*
(2) *The assignment*

$$f : \frac{\widetilde{G}}{H} \to N(\chi)[(1-\zeta_l)^k] \rtimes \langle \zeta_l \rangle,$$

defined by the formula $f(g) = (\psi(g), \chi(g))$, is a group isomorphism.

Proof. Let us begin by verifying (1). The restriction of ψ to $\ker(\pi)$ is a group homomorphism. Moreover, for each $g \in \widetilde{G}$ and $h \in \ker(\pi)$ we have

$$\begin{aligned}\psi(ghg^{-1}) &= \chi(g)\psi(hg^{-1}) + \psi(g) = \chi(g)\psi(g^{-1}) + \chi(g)\psi(h) + \psi(g)\\ &= \psi(gg^{-1}) + \chi(g)\psi(h) = \chi(g)\psi(h).\end{aligned}$$

This immediately implies (1). Next take $g_1, g_2 \in \widetilde{G}$. Then

$$\begin{aligned}&(\psi(g_1), \chi(g_1))(\psi(g_2), \chi(g_2))\\ &\qquad = (\psi(g_1), 1)(0, \chi(g_1))(\psi(g_2), 1)(0, \chi(g_1)^{-1})(0, \chi(g_1 g_2))\\ &\qquad = (\psi(g_1) + \chi(g_1)\psi(g_2), \chi(g_1 g_2)) = (\psi(g_1 g_2), \chi(g_1 g_2)).\end{aligned}$$

Hence f is a group homomorphism and it is zero precisely for the g with $\psi(g) = 0$ and $\chi(g) = 1$. This is the definition of H. ∎

Let now $\chi \in \Gamma_{\mu_l}(\mathbb{Q})$. We set $A := \mathrm{Cl}(K_\chi)[(1-\zeta_l)^\infty]$ and $G := \mathrm{Gal}(K_\chi/\mathbb{Q})$. Denote by $H_{\chi,l}$ the largest subextension of the Hilbert class field of K_χ, within $\overline{\mathbb{Q}}$, having degree a power of l. The Artin map gives a canonical identification $\mathrm{Cl}(K_\chi)[(1-\zeta_l)^\infty] = \mathrm{Gal}(H_{\chi,l}/K_\chi)$.

Proposition 6.3. *The surjection $\mathrm{Gal}(H_{\chi,l}/\mathbb{Q}) \twoheadrightarrow \mathrm{Gal}(K_\chi/\mathbb{Q})$ admits a section, yielding an isomorphism*

$$\mathrm{Gal}(H_{\chi,l}/\mathbb{Q}) \simeq_{\mathrm{gr}} \mathrm{Cl}(K_\chi)[(1-\zeta_l)^\infty] \rtimes \mathrm{Gal}(K_\chi/\mathbb{Q})$$

inducing the Artin identification between $\mathrm{Gal}(H_{\chi,l}/K_\chi)$ and $\mathrm{Cl}(K_\chi)[(1-\zeta_l)^\infty]$ when restricted to $\mathrm{Gal}(H_{\chi,l}/K_\chi)$.

Proof. This follows at once from Proposition 3.7. ∎

Proposition 6.4. *Let k be a positive integer. The natural map*

$$\mathrm{Cocy}(\mathrm{Gal}(H_{\chi,l}/\mathbb{Q}), N(\chi))[(1-\zeta_l)^k] \to \mathrm{Cl}(K_\chi)^\vee[(1-\zeta_l)^k]$$

is a split surjection of $\mathbb{Z}_l[\zeta_l]$-modules yielding an isomorphism

$$\mathrm{Cocy}(\mathrm{Gal}(H_{\chi,l}/\mathbb{Q}), N(\chi))[(1-\zeta_l)^k] \simeq \mathrm{Cl}(K_\chi)^\vee[(1-\zeta_l)^k] \oplus N[(1-\zeta_l)^k].$$

Proof. This follows immediately upon combining Propositions 6.3 and 6.1. ∎

Corollary 6.5. *Let k be a positive integer and let $\psi \in \mathrm{Cl}(K_\chi)^\vee[(1-\zeta_l)^k]$. The following are equivalent:*

(1) $\psi \in (1-\zeta_l)\mathrm{Cl}(K_\chi)^\vee[(1-\zeta_l)^{k+1}]$;

(2) *there is a* $\widetilde{\psi} \in (1-\zeta_l)\mathrm{Cocy}(\mathrm{Gal}(H_{\chi,l}/\mathbb{Q}), N(\chi))[(1-\zeta_l)^{k+1}]$ *restricting to* ψ;

(3) *for any* $\widetilde{\psi} \in \mathrm{Cocy}(\mathrm{Gal}(H_{\chi,l}/\mathbb{Q}), N(\chi))[(1-\zeta_l)^k]$ *such that* $\widetilde{\psi}$ *restricts to* ψ *we have* $\widetilde{\psi} \in (1-\zeta_l)\mathrm{Cocy}(\mathrm{Gal}(H_{\chi,l}/\mathbb{Q}), N(\chi))[(1-\zeta_l)^{k+1}]$.

Proof. This is a trivial consequence of Proposition 6.4. ∎

Corollary 6.5 tells us that finding a $(1-\zeta_l)$-lift for an unramified character is the same as finding a $(1-\zeta_l)$-lift for an unramified cocycle representing our character. In turn we now show that finding a $(1-\zeta_l)$-lift for an unramified cocycle is often the same as finding a $(1-\zeta_l)$-lift of the cocycle inflated to the absolute Galois group, providing in total a very convenient criterion for the existence of a $(1-\zeta_l)$-lift of a character in terms of cocycles from $G_\mathbb{Q}$ to $N(\chi)$.

Proposition 6.6. *Let* $\psi \in \mathrm{Cocy}(\mathrm{Gal}(H_{\chi,l}/\mathbb{Q}), N(\chi))[(1-\zeta_l)^k]$ *for some integer* $k \geq 1$. *In case* l *divides* $\Delta_{K_\chi/\mathbb{Q}}$ *we assume that* $\mathrm{Up}_{K_\chi}(l)$ *splits completely in the extension* $L(\psi)K_\chi/K_\chi$, *where we recall that* $L(\psi)$ *denotes the field of definition of* ψ. *Then the following are equivalent:*

(1) $\psi \in (1-\zeta_l)\mathrm{Cocy}(\mathrm{Gal}(H_{\chi,l}/\mathbb{Q}), N(\chi))[(1-\zeta_l)^{k+1}]$;

(2) $\psi \in (1-\zeta_l)\mathrm{Cocy}(G_\mathbb{Q}, N(\chi))[(1-\zeta_l)^{k+1}]$.

Proof. It is a triviality that (1) implies (2). We now show that (2) implies (1). We can assume without loss of generality that ψ restricted to G_{K_χ} surjects onto $N[(1-\zeta_l)^k]$. Let $L := L(\psi)\cdot K_\chi$ and let $\widetilde{\psi}$ be a lift of ψ. Thanks to Proposition 6.2 we find that the map

$$\mathrm{Gal}(L(\widetilde{\psi})/\mathbb{Q}) \to N[(1-\zeta_l)^{k+1}] \rtimes \langle \zeta_l \rangle \tag{6.1}$$

defined by the formula $g \mapsto (\widetilde{\psi}(g), \chi(g))$ is an isomorphism. We see that this map sends $\mathrm{Gal}(L(\widetilde{\psi})/L)$ into $N[1-\zeta_l]$. Therefore we deduce

$$j_l^{-1} \circ i_l^{-1} \circ \widetilde{\psi} \in \Gamma_{\mathbb{F}_l}(L)^{\mathrm{Gal}(L/\mathbb{Q})}.$$

The only primes that ramify in the extension $L/\mathbb{Q}$ are those dividing $\Delta_{K_\chi/\mathbb{Q}}$, since by assumption the extension L/K_χ is contained in H_{K_χ} and hence is unramified. For each q different from l dividing $\Delta_{K_\chi/\mathbb{Q}}$ we find that $\mathrm{Up}_{K_\chi}(q)$ splits completely in $K_\chi L((1-\zeta_l)\psi)$.

As we next explain, from this we conclude that if we restrict $f := (\widetilde{\psi}, \chi)$ to I_q, an inertia subgroup of q, the image of (6.1) is a group of the form $\langle N[1-\zeta_l], \sigma \rangle$, where $\chi(\sigma) \neq 1$. Indeed, since $\chi(I_q) \neq \{1\}$, we see that $f(I_q) \cap N$ is a $\mathbb{Z}_l[\zeta_l]$-submodule; we have $g \in I_q$ with $\chi(g) = \zeta_l$ and conjugation by g on $f(I_q) \cap N$ equals precisely multiplication by ζ_l. If $f(I_q) \cap N$ is trivial we are certainly done. On the other hand, it cannot have size more than l, since $\chi(I_q) \neq \{1\}$ and I_q is of size at most l^2. The only non-trivial $\mathbb{Z}_l[\zeta_l]$-submodule of size l is $N[1-\zeta_l]$. This shows our claim.

Recall that all $\sigma \in N[(1-\zeta_l)^{k+1}] \rtimes \langle \zeta_l \rangle$ with $\chi(\sigma) \neq 1$ have order l. From this, we conclude that the extension is locally split at inertia at q. In particular, in case l divides

$\Delta_{K_\chi/\mathbb{Q}}$, we know by assumption that inertia at l equals the decomposition group, hence the extension is locally split at l. Hence we conclude, by Theorem 4.10, that we can find $\chi' \in \Gamma_{\mathbb{F}_l}(\mathbb{Q})$ such that $L(\widetilde{\psi} + \chi')/L$ is unramified. This ends the proof. ∎

Remark 1. If $k \geq 2$, we claim that $L = L(\psi)$. This fact will be important throughout the paper. Indeed, if n denotes an element of $G_{\mathbb{Q}}$ with $(1-\zeta_l)\psi(n) \neq 0$, then for each $g \in G_\psi$, the group of definition of ψ, we have

$$\psi(g \cdot n) = \chi(g) \cdot n + \psi(g) \quad \text{and} \quad \psi(g) = \psi(\mathrm{id}) = 0,$$

where each equality is justified by the fact that ψ is a cocycle. Therefore we find that $\chi(g) = 1$, thanks to the fact that $(1-\zeta_l)\psi(n) \neq 0$. That means that $G_\psi \subseteq G_{K_\chi}$, which is equivalent to $L(\psi) \supseteq K_\chi$.

Definition 6.7. A *raw cocycle* for χ is a finite sequence $\{\psi_i\}_{0 \leq i \leq j}$ with

$$\psi_i \in \mathrm{Cocy}(\mathrm{Gal}(H_{\chi,l}/\mathbb{Q}), N(\chi))[(1-\zeta_l)^i], \quad (1-\zeta_l)\psi_i = \psi_{i-1} \quad \text{for } 1 \leq i \leq j.$$

The integer j is called the *rank* of the raw cocycle.

7. A reflection principle

7.1. The differential of sums of cocycles

In this section if M denotes a multiset, then we let $\mathrm{Set}(M)$ be the corresponding set and call its elements the elements of M. If $x_0 \in \mathrm{Set}(M)$, we extend the usual operations $-$, $=$, $\neq$ and $|\cdot|$ of sets to multisets in the natural way. Let m be a positive integer and let

$$Y := \prod_{i=1}^{m} Y_i \times \{d\}$$

be a product multiset, where each Y_i is a multiset of l primes that are all 1 modulo l. We further assume that the Y_i are all distinct multisets as i varies in $[m]$, and that the elements in each Y_i are coprime to the integer d, which is also a product of distinct primes q with $q = l$ or q equal to 1 modulo l.

We will identify points in this product space $x \in Y$ with the integer $D_x := (\prod_{i=1}^{m} \pi_i(x))d$ whenever convenient. A subset $C \subseteq Y$ is called a *subcube* in case it is the inverse image of a singleton under the projection of Y on $\prod_{i \in T} Y_i$ for a subset T of $\{1, \ldots, m\}$. The non-negative integer $m - |T|$ is called the *dimension* of C and denoted by $\dim(C)$; it satisfies $l^{\dim(C)} = |C|$.

Define an *amalgama* for Y to be a map

$$\epsilon_Y : \bigcup_{i=1}^{m} \mathrm{Set}(Y_i) \cup \{q \mid d\}_{q \text{ prime}} \to [l-1].$$

If ϵ_Y is an amalgama for Y we obtain for each $D \in Y$ an amalgama $\epsilon_Y(D)$ for D. Conversely, giving an amalgama for Y is the same as assigning an amalgama to each $D \in Y$

in a consistent manner, i.e. taking the same value at a prime whenever that prime divides two different points of Y. In this manner given ϵ_Y, an amalgama for Y, we have for each $D \in Y$ a character

$$\chi_{\epsilon_Y}(D) := \chi_{\epsilon_Y(D)}(D) = \prod_{p|D} \chi_p^{\epsilon_Y(p)}.$$

Next, for ϵ_Y an amalgama for Y, a *raw cocycle* $\Re$ *for* (Y, ϵ_Y) is an assignment that gives to each point $D \in Y$ a raw cocycle for $\chi_{\epsilon_Y}(D)$ (see Definition 6.7) with each point having rank at least $m-1$. We will write $\psi_k(\Re, \chi_{\epsilon_Y}(D))$ for the ψ_k in the raw cocycle of $\chi_{\epsilon_Y}(D)$.

Let $x_0 \in \mathrm{Set}(Y)$. We say that a raw cocycle $\Re$ for (Y, ϵ_Y) is *promising* with respect to x_0 if $\Re(\chi_{\epsilon_Y}(x))$ has rank at least m at each $x \in Y - \{x_0\}$ and

$$\sum_{x \in H} \psi_j(\Re, \chi_{\epsilon_Y}(x)) \in N[1-\zeta_l]$$

for each proper subcube of Y having dimension j not containing x_0. Here we have implicitly inflated our cocycles $\psi_j(\Re, \chi_{\epsilon_Y}(x))$ to $G_{\mathbb{Q}}$, so it makes sense to add them. The rest of this subsection is devoted to computing for each $\sigma, \tau \in G_{\mathbb{Q}}$ the value of

$$d_{x_0}\Big(\sum_{x \in Y : x \neq x_0} \psi_m(\Re, \chi_{\epsilon_Y}(x))\Big)(\sigma, \tau),$$

where $\Re$ is a promising raw cocycle with respect to x_0 and the differential d_{x_0} is taken by considering the sum as a 1-cochain in the $G_{\mathbb{Q}}$-module $N(\chi_{\epsilon_Y}(x_0))$. We remind the reader that Y is a multiset, so that the above sum has to be computed with the corresponding multiplicities. To state the result of this calculation, we need some additional definitions.

Firstly, we assume that each Y_i is written as $\{p_{i0}, \ldots, p_{i(l-1)}\}$ with the convention that the coordinates of x_0 occupy the first $s(i)$ indices for each i. In this manner the multiset Y corresponds bijectively to the set

$$\mathcal{F} := \mathrm{Map}([m], \{0, \ldots, l-1\})$$

with the points of Y giving x_0 in $\mathrm{Set}(Y)$ corresponding to the functions f with $f(i) \leq s(i)-1$ for each i in $[m]$. We let $\mathcal{F}_{x_0}$ be the set of such functions.

For each $j \in [m]$ and $i \in \{0, \ldots, l-1\}$ we define the character

$$\chi_{j,i,\epsilon_Y} := -j_l^{-1} \circ \chi_{p_{j0}}^{\epsilon_Y(p_{j0})} + j_l^{-1} \circ \chi_{p_{ji}}^{\epsilon_Y(p_{ji})}.$$

The following function attached to each $\tilde{f} \in \mathcal{F} - \{0\}$ will play an important role:

$$\chi_{\tilde{f},\epsilon_Y}(\sigma) := \prod_{j \in [m] : \tilde{f}(j) \neq 0} \chi_{j,\tilde{f}(j),\epsilon_Y}(\sigma).$$

Here the product takes place in $\mathbb{F}_l$ and the map is a continuous 1-cochain from $G_{\mathbb{Q}}$ to $\mathbb{F}_l$. For each $\tilde{f} \in \mathcal{F} - \mathcal{F}_{x_0}$ we denote by $H_{\tilde{f}}$ the set of $f \in \mathcal{F}$ such that for all $j \in [m]$, $\tilde{f}(j) \neq 0$ implies $\tilde{f}(j) = f(j)$.

Proposition 7.1. *Let Y, ϵ_Y, x_0 be as above. Let $\mathfrak{R}$ be a raw cocycle for (Y, ϵ_Y) that is promising for x_0. Then*

$$d_{x_0}\Big(\sum_{x\in Y: x\neq x_0} \psi_m(\mathfrak{R}, \chi_{\epsilon_Y}(x))\Big) = \sum_{\tilde{f}\in\mathcal{F}-\{0\}} (-1)^{|\{j:\tilde{f}(j)\neq 0\}|+1}\chi_{\tilde{f}}(\sigma)\Big(\sum_{f\in H_{\tilde{f}}} \psi_{m-|\{j:\tilde{f}(j)\neq 0\}|}(f)(\tau)\Big).$$

Proof. We start with a simple computation

$$\begin{aligned} d_{x_0}(\psi_m(x)) &= \chi_{\epsilon_Y}(x_0)\psi_m(x)(\tau) + \psi_m(x)(\sigma) - \psi_m(x)(\sigma\tau) \\ &= -\chi_{\epsilon_Y}(x)(\sigma)\psi_m(x)(\tau) + \chi_{\epsilon_Y}(x_0)(\sigma)\psi_m(x)(\tau) \\ &= \chi_{\epsilon_Y}(x_0)(\sigma)\left(1 - \frac{\chi_{\epsilon_Y}(x)(\sigma)}{\chi_{\epsilon_Y}(x_0)(\sigma)}\right)\psi_m(x)(\tau). \end{aligned}$$

Put $\chi := \chi_{\epsilon_Y}(x_0)$. Summing over all the ψ_m yields

$$\chi(\sigma)\sum_{f\in\mathcal{F}-\mathcal{F}_{x_0}}\left(1 - \frac{\chi_{\epsilon_Y}(f)(\sigma)}{\chi(\sigma)}\right)\psi_m(f)(\tau), \tag{7.1}$$

where $\chi_{\epsilon_Y}(f)$ is defined to be $\chi_{\epsilon_Y}(x)$ for the unique x that corresponds to f under our bijection. Note that this is not the same as χ_{f,ϵ_Y}. For $f\in\mathcal{F}, \sigma\in G_{\mathbb{Q}}$ define

$$\Sigma(f,\sigma) := \sum_{j\in[m]} \chi_{j,f(j),\epsilon_Y}(\sigma).$$

Observe that for $f\in\mathcal{F}_{x_0}$ the symbol $\psi_m(f)$ is not defined; however, we use the convention that $0\cdot\psi_m(f)$ is always defined and equal to 0, and also $(1-\zeta_l)^i\psi_m(f)$ is defined to be $\psi_{m-i}(f)$ for each integer $0\leq i\leq m$. This will simplify the notation in the coming calculations. With these conventions, the expression (7.1) is equal to

$$\chi(\sigma)\sum_{f\in\mathcal{F}}(1-\zeta_l^{\Sigma(f,\sigma)})\psi_m(f)(\tau). \tag{7.2}$$

Define

$$T^i_{\sigma,f} := \{j\in[m] : \chi_{j,f(j),\epsilon_Y}(\sigma) = i\}.$$

Obviously, $T^i_{\sigma,f}\cap T^j_{\sigma,f} = \emptyset$ for $i\neq j$. We can now rewrite (7.2) as

$$\begin{aligned} &\chi(\sigma)\sum_{f\in\mathcal{F}}\Big(1-\prod_{i=1}^{l-1}\zeta_l^{i|T^i_{\sigma,f}|}\Big)\psi_m(f)(\tau) \\ &\qquad = \chi(\sigma)\sum_{f\in\mathcal{F}}\Big(1-\prod_{i=1}^{l-1}(1+(\zeta_l^i-1))^{|T^i_{\sigma,f}|}\Big)\psi_m(f)(\tau) \\ &\qquad = \chi(\sigma)\sum_{f\in\mathcal{F}} a(f,\sigma)\psi_m(f)(\tau), \end{aligned} \tag{7.3}$$

where

$$a(f,\sigma) := 1 - \sum_{(T_1,\dots,T_{l-1}):\, T_i \subseteq T^i_{\sigma,f}} (-1)^{\sum_{i=1}^{l-1}|T_i|} \prod_{i=1}^{l-1} \Big(\sum_{j=0}^{i-1} \zeta_l^j\Big)^{|T_i|} (1-\zeta_l)^{\sum_{i=1}^{l-1}|T_i|}.$$

We can expand (7.3) as

$$-\chi(\sigma) \sum_{f\in\mathcal{F}} \sum_{\substack{(T_1,\dots,T_{l-1})\neq(\emptyset,\dots,\emptyset)\\ T_i\subseteq T^i_{\sigma,f}}} (-1)^{\sum_{i=1}^{l-1}|T_i|} \prod_{i=1}^{l-1} \Big(\sum_{j=0}^{i-1} \zeta_l^j\Big)^{|T_i|} \psi_{m-\sum_{i=1}^{l-1}|T_i|}(f)(\tau)$$

$$= -\chi(\sigma) \sum_{\substack{(T_1,\dots,T_{l-1})\neq(\emptyset,\dots,\emptyset)\\ \text{all disjoint}}} (-1)^{\sum_{i=1}^{l-1}|T_i|} \prod_{i=1}^{l-1} \Big(\sum_{j=0}^{i-1} \zeta_l^j\Big)^{|T_i|} \cdot \sum_{f\in\mathcal{F}} \mathbf{1}_{\forall i: T_i\subseteq T^i_{\sigma,f}}(\sigma) \psi_{m-\sum_{i=1}^{l-1}|T_i|}(f)(\tau). \tag{7.4}$$

We next make a definition. For $T_\bullet := (T_1,\dots,T_{l-1})$ we say that $\tilde{f}$ is $(T_\bullet,\sigma)$-*good* if

- $T_i \subseteq T^i_{\sigma,\tilde{f}}$ for each $i \in [l-1]$;
- $\tilde{f}(j) = 0$ for each $j \notin \bigcup_{i=1}^{l-1} T_i$.

Remark 2. Given $T_\bullet$ there can be *many* functions $\tilde{f}$ that are $(T_\bullet,\sigma)$-good. But for each $\tilde{f}\in\mathcal{F}$ there is at most one ordered choice of $l-1$ disjoint sets $T_\bullet$ such that $\tilde{f}$ is $(T_\bullet,\sigma)$-good; this partition exists if and only if $\chi_{\tilde{f},\epsilon_Y}(\sigma)\neq 0$. It is simply obtained by declaring that $j\in T_i$ if and only if $\chi_{j,\tilde{f}(j),\epsilon_Y}(\sigma) = i$.

We now use the definition of $(T_\bullet,\sigma)$-good to rearrange the sum in equation (7.4) as a sum over subcubes

$$-\chi(\sigma) \sum_{\substack{(T_1,\dots,T_{l-1})\neq(\emptyset,\dots,\emptyset)\\ \text{all disjoint}}} (-1)^{\sum_{i=1}^{l-1}|T_i|} \prod_{i=1}^{l-1} \Big(\sum_{j=0}^{i-1} \zeta_l^j\Big)^{|T_i|} \cdot \sum_{\tilde{f}\in\mathcal{F}-\mathcal{F}_{x_0}} \mathbf{1}_{\tilde{f}\text{ is }(T_\bullet,\sigma)\text{-good}} \cdot \sum_{f\in H_{\tilde{f}}} \psi_{m-\sum_{i=1}^{l-1}|T_i|}(f)(\tau). \tag{7.5}$$

Here $H_{\tilde{f}}$ denotes the set of $f\in\mathcal{F}$ such that $\tilde{f}(j)\neq 0$ implies $\tilde{f}(j) = f(j)$ for all $j\in[m]$. Since the raw cocycle is promising, we can rewrite (7.5) simply as

$$-\sum_{\substack{(T_1,\dots,T_{l-1})\neq(\emptyset,\dots,\emptyset)\\ \text{all disjoint}}} \sum_{\tilde{f}\in\mathcal{F}-\{0\}} (-1)^{|\{j:\,\tilde{f}(j)\neq 0\}|} \chi_{\tilde{f},\epsilon_Y}(\sigma) \mathbf{1}_{\tilde{f}\text{ is }(T_\bullet,\sigma)\text{-good}} \cdot \sum_{f\in H_{\tilde{f}}} \psi_{m-|\{j:\,\tilde{f}(j)\neq 0\}|}(f)(\tau).$$

We now swap the first two sums and apply Remark 2 to obtain

$$\sum_{\tilde{f}\in\mathcal{F}-\{0\}} (-1)^{|\{j:\tilde{f}(j)\neq 0\}|+1}\chi_{\tilde{f},\epsilon_Y}(\sigma)\Big(\sum_{f\in H_{\tilde{f}}}\psi_{m-\{j:\tilde{f}(j)\neq 0\}|}(f)(\tau)\Big),$$

which is the desired expression. ∎

We can further simplify the result of Proposition 7.1 under the following more restrictive assumption. We say that a raw cocycle $\mathfrak{R}$ for (Y,ϵ_Y) is *very promising* if it is promising and moreover

$$\sum_{f\in H_{\tilde{f}_1}}\psi_{m-|\{j:\tilde{f}_1(j)\neq 0\}|}(f)(\tau)=\sum_{f\in H_{\tilde{f}_2}}\psi_{m-|\{j:\tilde{f}_2(j)\neq 0\}|}(f)(\tau)$$

whenever $\tilde{f}_1$ and $\tilde{f}_2$ share the same zero set. In this case for each subset $T\subseteq[m]$ we put

$$\chi_{T,\epsilon_Y}:=\prod_{j\in T}j_l^{-1}\Big(\prod_{p\in Y_j}\chi_p^{\epsilon_Y(p)}\Big)$$

and

$$\psi(\mathfrak{R}([m]-T))(\tau):=\sum_{f\in H_{\tilde{f}}}\psi_{m-|\{j:\tilde{f}(j)\neq 0\}|}(f)(\tau)$$

for any $\tilde{f}\in\mathcal{F}$ such that the set of j with $\tilde{f}(j)=0$ is the set T. The inner product is a product of characters, while the outer product takes place in $\mathbb{F}_l$, yielding a continuous 1-cochain from $G_{\mathbb{Q}}$ to $\mathbb{F}_l$.

Proposition 7.2. *Let Y, ϵ_Y, x_0 be as above. Let $\mathfrak{R}$ be a raw cocycle for (Y,ϵ_Y) that is very promising for x_0. Then*

$$d_{x_0}\Big(\sum_{x\in Y:x\neq x_0}\psi_m(\mathfrak{R},\chi_{\epsilon_Y}(x))(\sigma,\tau)\Big)=\sum_{\emptyset\neq T\subseteq[m]}(-1)^{|T|+1}\chi_{T,\epsilon_Y}(\sigma)\psi(\mathfrak{R}([m]-T))(\tau).$$

The formula in Proposition 7.2 is strongly reminiscent of [22, (2.6)]. One could reach an even closer set-up by taking each of the Y_i such that $|\mathrm{Set}(Y_i)|=2$, one element in $\mathrm{Set}(Y_i)$ has multiplicity 1 and the other has multiplicity $l-1$. In other words, one could work with the usual set-up of 2-power cubes as in [22] at the cost of working with *signed* sums of cocycles.

7.2. Expansion maps

Let m, Y be as in the previous subsection. Let ϵ_Y be an amalgama for Y. The next step is to construct a collection of 1-cochains that have precisely the same recursive formula as in Proposition 7.2. This leads to the following definition. We define a *pre-expansion* for (Y,ϵ_Y) to be a sequence parametrized by the proper subsets of $[m]$,

$$\{\phi_T(Y,\epsilon_Y)\}_{T\subsetneq[m]},$$

consisting of continuous 1-cochains from $G_{\mathbb{Q}}$ to $\mathbb{F}_l$ satisfying

$$(d\phi_T(Y,\epsilon_Y))(\sigma,\tau) = \sum_{\emptyset\neq U\subseteq T} (-1)^{|U|+1}\chi_{U,\epsilon_Y}(\sigma)\phi_{T-U}(Y,\epsilon_Y)(\tau)$$

for each $\sigma,\tau\in G_{\mathbb{Q}}$ and each proper subset T of $[m]$. Here we consider $\mathbb{F}_l$ as a $G_{\mathbb{Q}}$-module with the trivial action. We will assume that $\phi_\emptyset$ is linearly independent from the space of characters spanned by the set $\{\chi_{\{i\},\epsilon_Y}\}_{i\in[m]}$.

A pre-expansion is said to be *promising* if for every $i\in[m]$, every prime $p\in Y_i$ splits completely in the field of definition of $\phi_{[m]-\{i\}}$. Next we define an *expansion* for Y to be

$$\{\phi_T(Y,\epsilon_Y)\}_{T\subsetneq[m]}\cup\{\phi_{[m]}(Y,\epsilon_Y)\},$$

where $\{\phi_T(Y,\epsilon_Y)\}_{T\subsetneq[m]}$ is a pre-expansion for (Y,ϵ_Y) and $\phi_{[m]}(Y,\epsilon_Y)$ is a continuous 1-cochain from $G_{\mathbb{Q}}$ to $\mathbb{F}_l$ satisfying

$$(d\phi_{[m]}(Y,\epsilon_Y))(\sigma,\tau) = \sum_{\emptyset\neq U\subseteq[m]} (-1)^{|U|+1}\chi_{U,\epsilon_Y}(\sigma)\phi_{[m]-U}(Y,\epsilon_Y)(\tau).$$

The maps composing a pre-expansion or an expansion are said to be *good* if their field of definition is unramified above the maximal elementary abelian $\mathbb{F}_l$-extension, which is the field of definition of the map $\chi_{[m],\epsilon_Y}\cdot\phi_\emptyset(Y,\epsilon_Y)$. A pre-expansion or an expansion are said themselves to be good if all their maps are good. The pair (Y,ϵ_Y) is said to be *cooperative* if for any distinct $i,j\in[m]$ the character $\chi_{\{i\},\epsilon_Y}$ is locally trivial at each prime appearing in Y_j and at l.

Proposition 7.3. *Suppose that (Y,ϵ_Y) is cooperative. If $\{\phi_T(Y,\epsilon_Y)\}_{T\subsetneq[m]}$ is a promising good pre-expansion for (Y,ϵ_Y), then it can be completed to a good expansion for (Y,ϵ_Y).*

Proof. Consider the map $\theta: G_{\mathbb{Q}}\times G_{\mathbb{Q}}\to\mathbb{F}_l$ defined by

$$\theta(\sigma,\tau) := \sum_{\emptyset\neq U\subseteq[m]} (-1)^{|U|+1}\chi_{U,\epsilon_Y}(\sigma)\phi_{[m]-U}(Y,\epsilon_Y)(\tau).$$

Simply from the assumption that $\{\phi_T(Y,\epsilon_Y)\}_{T\subsetneq[m]}$ is a pre-expansion, it follows that θ is a 2-cocycle. We will also write θ for the resulting class in $H^2(G_{\mathbb{Q}},\mathbb{F}_l)$. Observe that θ factors through

$$M := \prod_{T\subsetneq[m]} L(\phi_T(Y,\epsilon_Y))$$

and hence defines an element of $H^2(\mathrm{Gal}(M/\mathbb{Q}),\mathbb{F}_l)$. We next show that the $\mathbb{F}_l$-extension of the group $\mathrm{Gal}(M/\mathbb{Q})$ given by the class of θ is actually in $\widetilde{\mathrm{Cent}}_{\mathbb{F}_l}(M/\mathbb{Q})$ and that it can be realized by an *unramified* $\mathbb{F}_l$-extension. We do so by applying Theorem 4.10(1).

Since the expansion is good, we only have to check the primes ramifying in the field of definition of $\chi_{[m],\epsilon_Y}\cdot\phi_\emptyset(Y,\epsilon_Y)$. Locally at these primes the expression defining θ becomes identically zero, because the pre-expansion is good and (Y,ϵ_Y) is cooperative. Hence, by Theorem 4.10, we conclude that indeed $\theta\in\widetilde{\mathrm{Cent}}_{\mathbb{F}_l}(M/\mathbb{Q})$ and that we can realize it as $r_1(\chi)$ for some $\chi\in\Gamma_{\mathbb{F}_l}(M)$ with $L(\chi)/M$ unramified. Thanks to Proposition 4.3

we can write $\theta = d(\phi)$ for some 1-cochain $\phi : G_{\mathbb{Q}} \to \mathbb{F}_l$. One can show with a direct calculation that ϕ restricted to G_M is an element of $\Gamma_{\mathbb{F}_l}(M)^{\mathrm{Gal}(M/\mathbb{Q})}$ with $r_1(\phi) = \theta$. After twisting with an element of $\Gamma_{\mathbb{F}_l}(\mathbb{Q})$ we can assume that ϕ restricts to χ. Finally, a simple calculation (based on the formula for θ) shows that the field of definition $L(\phi)$ is actually $L(\chi)$. Setting $\phi_{[m]}(Y, \epsilon_Y) := \phi$ establishes the proposition. ∎

Our next goal is to determine structural information about the Galois group

$$\mathrm{Gal}(L(\phi_{[m]}(Y, \epsilon_Y))/\mathbb{Q}).$$

To simplify the notation and to enhance generality, in the remainder of this section we assume that we have a collection of characters

$$\{\chi_{\{i\}}\}_{i\in[m]} \cup \{\chi_\emptyset\}$$

living in $\Gamma_{\mathbb{F}_l}(\mathbb{Q})$ and forming altogether an $\mathbb{F}_l$-linearly independent set. For each non-empty subset $T \subseteq [m]$ we denote by χ_T the continuous 1-cochain from $G_{\mathbb{Q}}$ to $\mathbb{F}_l$ defined by

$$\chi_T := \prod_{i\in T} \chi_{\{i\}},$$

where the product is just multiplication in $\mathbb{F}_l$. We assume moreover to have a collection $\{\phi_T\}_{T\subseteq[m]}$ of maps satisfying the above equation of expansion maps,

$$(d\phi_T)(\sigma, \tau) = \sum_{\emptyset\neq U\subseteq T} (-1)^{|U|+1} \chi_U(\sigma)\phi_{T-U}(\tau)$$

for each $T \subseteq [m]$, with furthermore $\phi_\emptyset = \chi_\emptyset$. We proceed to determine the structure of

$$\mathrm{Gal}(L(\phi_{[m]})/\mathbb{Q}).$$

For each non-empty $T \subseteq [m]$ we put $K_T := L(\chi_T)$, i.e. the $\mathbb{F}_l$-elementary extension obtained by adding all the characters $\chi_{\{i\}}$ with $i \in T$. We denote by $\{\sigma_i\}_{i\in[m]}$ the dual basis of $\chi_{\{i\}}$ in $\mathrm{Gal}(K_{[m]}/\mathbb{Q})$. Recall that $\mathrm{Gal}(K_{[m]}/\mathbb{Q})$ acts on $\Gamma_{\mathbb{F}_l}(K_{[m]})$ by conjugation. Observe that clearly $\phi_T|_{G_{K_T}} \in \Gamma_{\mathbb{F}_l}(K_T)$. Take $i \in [m]$ and $T \subseteq [m]$. Let by abuse of notation σ_i denote also any lift of σ_i to $G_{\mathbb{Q}}$; the choice is relevant only to write symbolically meaningful formulas, but since we are going to examine the effect on conjugation on a character, it will be irrelevant for the end result. Observe that for any $\tau \in G_{K_{[m]}}$ we have

$$\begin{aligned}\phi_T(\sigma_i\tau\sigma_i^{-1}) &= \phi_T([\sigma_i, \tau]\tau) = \phi_T([\sigma_i, \tau]) + \phi_T(\tau) - (d\phi_T)([\sigma_i, \tau], \tau)\\ &= \phi_T([\sigma_i, \tau]) + \phi_T(\tau).\end{aligned} \tag{7.6}$$

Since

$$\phi_T([\sigma_i, \tau]) + \phi_T(\tau\sigma_i) - \phi_T(\sigma_i\tau) = (d\phi_T)([\sigma_i, \tau], \tau\sigma_i) = 0,$$

one finds

$$\phi_T([\sigma_i, \tau]) = \phi_T(\sigma_i\tau) - \phi_T(\tau\sigma_i). \tag{7.7}$$

This can in turn be rewritten as

$$\begin{aligned}\phi_T(\sigma_i\tau) - \phi_T(\tau\sigma_i) &= (d\phi_T)(\tau, \sigma_i) - (d\phi_T)(\sigma_i, \tau)\\ &= -(d\phi_T)(\sigma_i, \tau),\end{aligned} \tag{7.8}$$

where in the last identity we make use of the fact that $\tau \in G_{K_{[m]}}$. From (7.6)–(7.8) we find that

$$\phi_T(\sigma_i\tau\sigma_i^{-1}) = \begin{cases}\phi_T(\tau) - \phi_{T-\{i\}}(\tau) & \text{if } i \in T,\\ \phi_T(\tau) & \text{if } i \notin T.\end{cases}$$

Therefore the action is given by the formula

$$\sigma_i \cdot \phi_T = \begin{cases}\phi_T - \phi_{T-\{i\}} & \text{if } i \in T,\\ \phi_T & \text{if } i \notin T.\end{cases}$$

From this formula it follows immediately that $\mathrm{Gal}(L(\phi_{[m]})/K_{[m]})$ is an $\mathbb{F}_l$-vector space whose dual is generated by all the maps ϕ_T. Moreover, we deduce from the above formula that the natural action of $\mathbb{F}_l[\mathrm{Gal}(K_{[m]}/\mathbb{Q})]$ factors through the ideal generated by $\{(\sigma_i - 1)^2\}_{i\in[m]}$. The change of variables $t_i := \sigma_i - 1$ shows that the group ring $\mathbb{F}_l[\mathrm{Gal}(K_{[m]}/\mathbb{Q})]$ is isomorphic to the polynomial ring $\frac{\mathbb{F}_l[t_1,\ldots,t_m]}{(t_1^l,\ldots,t_m^l)}$. Hence $\mathrm{Gal}(L(\phi_{[m]})/K_{[m]})$ is a module over the ring

$$R := \frac{\mathbb{F}_l[t_1, \ldots, t_m]}{(t_1^2, \ldots, t_m^2)}.$$

We next prove that the dual $\mathrm{Gal}(L(\phi_{[m]})/K_{[m]})^\vee$ is a free module of rank 1 over R. It is clear that $\phi_{[m]}$ is a generator. Hence we only need to show that the annihilator ideal of $\phi_{[m]}$ is the zero ideal. Since R is Gorenstein, we see that if the annihilator is not the zero ideal, then it must contain $t_1 \cdot \ldots \cdot t_m$. But, still thanks to the formula, $t_1 \cdot \ldots \cdot t_m$ sends $\phi_{[m]}$ to $\chi_\emptyset$. Since the character $\chi_\emptyset$ is independent of the characters $\chi_{\{i\}}$, we obtain the desired conclusion.

Denote by G the group

$$\frac{\mathbb{F}_l[t_1, \ldots, t_m]}{(t_1^2, \ldots, t_m^2)} \rtimes \mathrm{Gal}(K_{[m]}/\mathbb{Q}),$$

where the implicit action is the natural action of $\mathrm{Gal}(K_{[m]}/\mathbb{Q})$ on $\frac{\mathbb{F}_l[t_1,\ldots,t_m]}{(t_1^2,\ldots,t_m^2)}$. With little extra effort one can show that $\mathrm{Gal}(L(\phi_{[m]})/\mathbb{Q})$ is actually isomorphic to G.

We next examine the map

$$\beta_{m+1}(\phi_{[m]}) : G_{\mathbb{Q}}^{m+1} \to \mathbb{F}_l$$

that sends a vector $(\tau_1, \ldots, \tau_{m+1})$ to

$$\phi_{[m]}([\tau_1, [\tau_2, [\ldots [\tau_m, \tau_{m+1}] \ldots]]]).$$

Using the structure of the group $\mathrm{Gal}(L(\phi_{[m]})/\mathbb{Q})$ one can quite easily establish the formula

$$\beta_{m+1}(\phi_{[m]})(\tau_1,\dots,\tau_{m+1}) = \sum_{\rho\in\mathrm{Sym}\{1,\dots,m\}} \chi_\emptyset(\tau_{m+1}) \prod_{1\le i\le m} \chi_{\rho(i)}(\tau_i) - \sum_{\rho\in\mathrm{Sym}(i\in\{1,\dots,m+1\}-\{m\})} \chi_\emptyset(\tau_m) \prod_{1\le i\le m+1, i\neq m} \chi_{\rho(i)}(\tau_i).$$

For an alternative way to arrive at the same formula one can use the identity

$$\phi_{[m]}([\sigma,\tau]) = \phi_{[m]}(\sigma\tau) - \phi_{[m]}(\tau\sigma) = (d\phi_{[m]})(\tau,\sigma) - (d\phi_{[m]})(\sigma,\tau)$$

and proceed by induction as explained in [22, p. 12]. The formula for β_{m+1} will be of utmost importance in Section 9, since it reduces the task of finding all relations among certain collections of expansion maps, called *governing expansions*, to the task of finding all relations among (suitable functions of) the characters at the base of the expansion.

7.3. Creating unramified cocycles

Let Y and ϵ_Y be as in the previous two subsections and let $x_0 \in \mathrm{Set}(Y)$. In this subsection we build on the previous two subsections to prove that under suitable assumptions, $\psi_1(\mathfrak{R}, \chi_{\epsilon_Y}(x_0)) \in (1-\zeta_l)^{m-1}\mathrm{Cl}(K_{\chi_{\epsilon_Y}(x_0)})^\vee[(1-\zeta_l)^m]$. These assumptions come in two flavors, and we devote a subsection to each.

7.3.1. Minimality. Let $\mathfrak{R}$ be a raw cocycle on (Y,ϵ_Y) that is promising at x_0. Moreover, we assume that Y is non-degenerate; there are no Y_i containing all equal entries. We say that $\mathfrak{R}$ is *minimal* with respect to x_0 if for any subcube H of Y not containing x_0 we have

$$\sum_{x\in H} \psi_{\dim(H)}(\mathfrak{R}, \chi_{\epsilon_Y}(x)) = 0.$$

We have the following fact.

Proposition 7.4. *Let $\mathfrak{R}$ be a promising minimal raw cocycle at (Y,ϵ_Y). Suppose that $\psi_1(\mathfrak{R}, \chi_{\epsilon_Y}(x))$ is constant as x varies in Y. Then there exists $\chi' : G_\mathbb{Q} \to N[1-\zeta_l]$ such that*

$$\psi := -\sum_{x\in Y : x\neq x_0} \psi_m(\mathfrak{R}, \chi_{\epsilon_Y}(x)) + \chi' \in \mathrm{Cocy}(\mathrm{Gal}(H_{\chi,l}/\mathbb{Q}), N(\chi_{\epsilon_Y}(x_0))).$$

One has

$$(1-\zeta_l)^{m-1}\psi = |\{x\in Y : x = x_0\}| \cdot \psi_1(\mathfrak{R}, \chi_{\epsilon_Y}(x_0)),$$

yielding in particular

$$\psi_1(\mathfrak{R}, \chi_{\epsilon_Y}(x_0)) \in (1-\zeta_l)^{m-1}\mathrm{Cl}(K_{\chi_{\epsilon_Y}(x_0)})^\vee[(1-\zeta_l)^m].$$

Proof. We surely have

$$\tilde{\psi} := \sum_{x \in Y : x \neq x_0} \psi_m(\Re, \chi_{\epsilon_Y}(x)) \in \mathrm{Cocy}(G_{\mathbb{Q}}, N(\chi_{\epsilon_Y}(x_0)))$$

thanks to our minimality assumption and Proposition 7.1. Next observe that the proposition is trivially true if $m = 1$. So we can safely assume $m \geq 2$. We claim that

$$(1 - \zeta_l)\tilde{\psi} \in \mathrm{Cocy}(\mathrm{Gal}(H_{\chi,l}/\mathbb{Q}), N(\chi_{\epsilon_Y}(x_0))).$$

Once we show this, the existence of χ' follows from Proposition 6.6. Indeed, if l divides x for some $x \in Y$, then l divides d. In this case we see that $\mathrm{Up}_{K_{\chi_{\epsilon_Y}(x_0)}}(l)$ splits completely in $L((1 - \zeta_l)\tilde{\psi})K_{\chi_{\epsilon_Y}(x_0)}$, which places us in the position to use Proposition 6.6. We are now going to prove the claim. Define

$$L := \prod_{x \in Y : x \neq x_0} L(\psi_m(\Re, \chi_{\epsilon_Y}(x))),$$

where we recall that $L(\psi_m)$ is the field of definition of ψ_m. As $m \geq 2$ we see that L contains $K_{\chi_{\epsilon_Y}(x_0)}$ and ramifies with inertia degree at most l at each prime. Moreover, its ramification locus is contained in the set of primes appearing as coordinates of Y, since all the ψ_m are unramified above their corresponding degree l cyclic extension.

This already implies that the primes dividing $\Delta_{K_{\chi_{\epsilon_Y}(x_0)}/\mathbb{Q}}$ cannot ramify in $L/K_{\chi_{\epsilon_Y}(x_0)}$, since the ramification is already eaten up by $K_{\chi_{\epsilon_Y}(x_0)}/\mathbb{Q}$. We are left with the other primes. For such a prime q, thanks to the minimality assumption, we can always rewrite $(1 - \zeta_l)\tilde{\psi}$ as a sum only over x where the prime q is never used. Therefore the ramification locus of $L((1 - \zeta_l)\tilde{\psi})$ over $\mathbb{Q}$ does not contain any of those q either. Hence we conclude that $(1 - \zeta_l)\tilde{\psi}$ has indeed unramified field of definition above $K_{\chi_{\epsilon_Y}(x_0)}$.

The claim gives χ' satisfying

$$-\tilde{\psi} + \chi' = \psi \in \mathrm{Cocy}(\mathrm{Gal}(H_{\chi,l}/\mathbb{Q}), N(\chi_{\epsilon_Y}(x_0))).$$

Observe that

$$\begin{aligned}(1 - \zeta_l)^{m-1}\psi &= (l^m - |\{x \in Y : x = x_0\}|) \cdot -\psi_1(\Re, \chi_{\epsilon_Y}(x_0)) \\ &= |\{x \in Y : x = x_0\}| \cdot \psi_1(\Re, \chi_{\epsilon_Y}(x_0)),\end{aligned}$$

where in the first equality we make use of the fact that all the ψ_1 are identical, and in the second we make use of the fact that $1 - \zeta_l$ kills each ψ_1 and thus the same is true for l. ∎

Remark 3. If we have the stronger assumption that the sum over *any* proper subcube is trivial for $\Re$, then we can directly conclude that the rank of $\Re$ at x_0 is at least m, since the expression in Proposition 7.4 would be a lift of $\psi_{m-1}(\Re, \chi_{\epsilon_Y}(x_0))$. That said, in our application this will be irrelevant, since raw cocycles are merely a tool to access the pairing $\mathrm{Art}_m(\mathrm{Cl}(K_{\chi_{\epsilon_Y}(x_0)}))$.

7.3.2. Agreement. Let i_a be in $[m]$. We will refer to i_a as the *index of agreement* in $[m]$. Moreover from now on we shall use the notation

$$Y - \{i_a\} := \prod_{i \neq i_a} Y_i.$$

Next we assume to have a good expansion $\{\phi_T(Y-\{i_a\}, \epsilon_Y)\}_{T \subseteq [m]-\{i_a\}}$ for $(Y-\{i_a\}, \epsilon_Y)$, where, by abuse of notation, ϵ_Y denotes also the restriction of ϵ_Y to $Y - \{i_a\}$.

Let $\mathfrak{R}$ be a raw cocycle on (Y, ϵ_Y) that is very promising at x_0. Moreover, we assume that Y is not degenerate, i.e. no Y_i consists of all equal entries. Recall in this case that for $T \subsetneq [m]$ we have introduced the notation $\psi(\mathfrak{R}(T))$ at the end of Section 7.1. We say that $\mathfrak{R}$ *agrees* with a good expansion $\{\phi_T(Y - \{i_a\}, \epsilon_Y)\}_{T \subseteq [m]-\{i_a\}}$ if

- for each $T \subsetneq [m]$ containing i_a one has

$$\psi(\mathfrak{R}(T)) = i_l \circ j_l \circ \phi_{T-\{i_a\}}(Y - \{i_a\}, \epsilon_Y);$$

- for each $T \subsetneq [m]$ not containing i_a one has

$$\psi(\mathfrak{R}(T)) = 0.$$

Proposition 7.5. *Let $\mathfrak{R}$ be a promising raw cocycle at (Y, ϵ_Y). Assume that there exists a character χ such that*

$$\psi_1(\mathfrak{R}, \chi_{\epsilon_Y}(x)) = \chi + \chi_{\pi_{i_a}(x)}^{\epsilon_Y(\pi_{i_a}(x))}$$

for all $x \in Y$. Let $\{\phi_T(Y - \{i_a\}, \epsilon_Y)\}_{T \subseteq [m]-\{i_a\}}$ be a good expansion for (Y, ϵ_Y). Suppose that $\mathfrak{R}$ agrees with this expansion. Then there exists $\chi' : G_{\mathbb{Q}} \to N[1 - \zeta_l]$ such that

$$\begin{aligned}\psi := &- \sum_{x \in Y : x \neq x_0} \psi_m(\mathfrak{R}, \chi_\epsilon(x)) + \chi' + i_l \circ j_l \circ \phi_{[m]-\{i_a\}}(Y - \{i_a\}, \epsilon_Y) \\ &\in \mathrm{Cocy}(\mathrm{Gal}(H_{\chi,l}/\mathbb{Q}), N(\chi_{\epsilon_Y}(x_0))).\end{aligned}$$

One has

$$(1 - \zeta_l)^{m-1}\psi = |\{x \in Y : x = x_0\}| \cdot \psi_1(\mathfrak{R}, \chi_{\epsilon_Y}(x_0)),$$

yielding in particular

$$\psi_1(\mathfrak{R}, \chi_{\epsilon_Y}(x_0)) \in (1 - \zeta_l)^{m-1}\mathrm{Cl}(K_{\chi_{\epsilon_Y}(x_0)})^\vee[(1 - \zeta_l)^m].$$

Proof. The proof is identical to the proof of Proposition 7.4. ∎

7.4. Sum of Artin pairings

In this subsection we upgrade the two results of the previous subsection, showing that, under some additional assumptions, one can also control the sum of the Artin pairings over the cube. We keep the parallel with the previous section, dividing in two the discussion according to the two cases consisting of minimality and agreement. However, in both cases the crucial additional assumption is the following.

Definition 7.6. We say that Y is *consistent* if for all i in $[m]$ and for each prime q dividing d we have that the characters $\chi_{q'}$ with $q' \in Y_i$ are all the same locally at q.

7.4.1. Minimality. Let $\mathfrak{R}$ be a raw cocycle on (Y, ϵ_Y) that is promising at x_0. Suppose also that $m \geq 2$ and that Y is non-degenerate, i.e. there are no Y_i consisting of all equal entries. We now show that under the assumption of Proposition 7.4, and some additional assumptions, we can also obtain a relation among the m-th Artin pairings of the cube.

Theorem 7.7. *Let $\mathfrak{R}$ be a minimal raw cocycle on (Y, ϵ_Y) that is promising at x_0 and suppose that Y is consistent. Suppose that $\psi_1(\mathfrak{R}, \chi_{\epsilon_Y}(x))$ is constant as x varies in Y. Let b be a positive integer whose prime divisors are all divisors of d. Assume that for all $x \in Y$, the element b, viewed as an element of $\overline{\mathrm{Cl}}(K_{\chi_{\epsilon_Y}(x)})$, maps to an element of $(1-\zeta_l)^{m-1}\mathrm{Cl}(K_{\chi_{\epsilon_Y}(x)})[(1-\zeta_l)^m]$. Then*

$$\psi_1(\mathfrak{R}, \chi_{\epsilon_Y}(x_0)) \in (1-\zeta_l)^{m-1}\mathrm{Cl}(K_{\chi_{\epsilon_Y}})[(1-\zeta_l)^m]$$

and furthermore

$$\sum_{x \in Y} \mathrm{Art}_m(\mathrm{Cl}(K_{\chi_{\epsilon_Y}(x)}))(b, \psi_1(\mathfrak{R}, \chi_{\epsilon_Y}(x))) = 0.$$

Proof. We know from Proposition 7.4 that indeed

$$\psi_1(\mathfrak{R}, \chi_{\epsilon_Y}(x_0)) \in (1-\zeta_l)^{m-1}\mathrm{Cl}(K_{\chi_{\epsilon_Y}})[(1-\zeta_l)^m],$$

and moreover we can find an unramified $(1-\zeta_l)^{m-1}$-cocycle lift of the form

$$\psi = -\frac{1}{|\{x \in Y : x = x_0\}|} \cdot \Big(\sum_{x \in Y : x \neq x_0} \psi_m(\mathfrak{R}, \chi_{\epsilon_Y}(x)) + \chi' \Big).$$

It is clear from that proof that the character χ' can ramify at most at the primes occurring as coordinates of a point of Y. Indeed, in that proof, the cocycle $\tilde{\psi}$ has ramification over $\mathbb{Q}$ already contained only at most in such primes, since its field of definition is contained in the compositum of all the ψ_m, for which this last claim is evidently true by their defining property. Recall that in the proof of Proposition 7.4 the character χ' comes from applying Proposition 6.6 and hence Proposition 4.8, where one proceeds by eliminating one by one all the eventual ramifying primes with a character supported precisely in that prime. This substantiates our claim on the shape of χ'.

For each prime divisor q of b, fix an embedding $G_{\mathbb{Q}_q} \subseteq G_{\mathbb{Q}}$ coming from a given fixed embedding $\overline{\mathbb{Q}} \to \overline{\mathbb{Q}_q}$. Thanks to our consistency assumption we see in particular the following crucial fact. For each prime divisor q of b and for each $x \in Y$ we have that

$$\ker(\chi_{\epsilon_Y}(x)) \cap G_{\mathbb{Q}_q}$$

is constantly the same index l subgroup. Let $K_q/\mathbb{Q}_q$ be the corresponding field extension, totally ramified of degree l. Denote by $\tilde{K}_q/K_q$ the unique unramified extension of K_q

of degree l. Recall that this comes with the canonical generator given by Frob_{K_q}, the Frobenius automorphism. Observe that by definition of the Artin pairing we have

$$\sum_{x\in Y} \mathrm{Art}_m(\mathrm{Cl}(K_{\chi_{\epsilon_Y}(x)}))(b, \psi_1(\Re, \chi_{\epsilon_Y}(x))) = \sum_{q|b} \psi(\mathrm{Frob}_{K_q}^{\epsilon_b(q)}) + \sum_{q|b} \sum_{x\in Y-\{x_0\}} \psi_m(\Re, \chi_{\epsilon_Y}(x))(\mathrm{Frob}_{K_q}^{\epsilon_b(q)}),$$

which is simply equal to $\sum_{q|b} \chi'(\mathrm{Frob}_{K_q}^{\epsilon_b(q)})$ by the definition of ψ.

Next, we can decompose χ' as a product of not necessarily distinct characters having conductor a power of a prime that is a coordinate of Y (the power will be precisely 1 if the prime is different from l, and it will be 2 if the prime is equal to l); this has been established above in this proof. If we can show that for such a $\chi_{q'}$ we have

$$\prod_{q|b} \chi_{q'}(\mathrm{Frob}_{K_q}^{\epsilon_b(q)}) = 1,$$

then we are clearly done. To see this pick a point x with $K_{\chi_{\epsilon_Y}(x)}$ ramifying at q'. By definition of the Artin symbol, we have

$$\prod_{q|b} \chi_{q'}(\mathrm{Frob}_{K_q}^{\epsilon_b(q)}) = \chi_{q'}(\mathrm{Up}_{K_{\chi_{\epsilon_Y}(x)}}(b)) = 1,$$

where the last equality follows directly from the assumption on b and Proposition 5.1. ∎

7.4.2. Agreement. Let $\Re$ be a raw cocycle on (Y, ϵ_Y) that is promising at x_0. Suppose also that $m \geq 2$. We also assume that Y is non-degenerate, i.e. there are no Y_i consisting of all equal entries. We now show that under the assumption of Proposition 7.5, and some additional assumptions, we can also obtain a relation among the m-th Artin pairings of the cube. Define

$$M(\phi_{[m]-\{i_a\}}(Y - \{i_a\}, \epsilon_Y)) := \prod_{T\subsetneq[m]-\{i_a\}} L(\phi_T(Y - \{i_a\}, \epsilon_Y)).$$

Theorem 7.8. *Let $\Re$ be a raw cocycle for (Y, ϵ_Y) that is very promising at x_0 and suppose that Y is consistent. Assume that there exists a character χ such that*

$$\psi_1(\Re, \chi_{\epsilon_Y}(x)) = \chi + \chi_{\pi_{i_a}(x)}^{\epsilon_Y(\pi_{i_a}(x))}$$

for all $x \in Y$. Let $\{\phi_T(Y - \{i_a\}, \epsilon_Y)\}_{T\subseteq[m]-\{i_a\}}$ be a good expansion for (Y, ϵ_Y). Suppose that $\Re$ agrees with this expansion. Let b be a positive integer whose prime divisors are all divisors of d. Assume that for all $x \in Y$, the element b, viewed as an element of $\overline{\mathrm{Cl}}(K_{\chi_{\epsilon_Y}(x)})$, maps to an element of $(1-\zeta_l)^{m-1}\mathrm{Cl}(K_{\chi_{\epsilon_Y}(x)})[(1-\zeta_l)^m]$. Then

$$\psi_1(\Re, \chi_{\epsilon_Y}(x_0)) \in (1-\zeta_l)^{m-1}\mathrm{Cl}(K_{\chi_{\epsilon_Y}})[(1-\zeta_l)^m].$$

Moreover, the Frobenius class of each prime q dividing d in $L(\phi_{[m]-\{i_a\}})$ consists of a central element in $\mathrm{Gal}(L(\phi_{[m]-\{i_a\}}(Y-\{i_a\},\epsilon_Y))/M(\phi_{[m]-\{i_a\}}(Y-\{i_a\},\epsilon_Y)))$ *and*

$$\sum_{x\in Y}\mathrm{Art}_m(\mathrm{Cl}(K_{\chi_{\epsilon_Y}(x)}))(b,\psi_1(\Re,\chi_{\epsilon_Y}(x))) = \sum_{q|b}\epsilon_b(q)(i_l\circ j_l\circ\phi_{[m]-\{i_a\}})(Y-\{i_a\},\epsilon_Y)(\mathrm{Frob}_q).$$

Proof. The proof is identical to the proof of Theorem 7.7. ∎

8. Additive systems

Recall that $[d]$ denotes the set $\{1,\dots,d\}$, where d is any integer. In the previous subsections we dealt with just one cube Y. To prove our main theorems, we will use Theorems 7.7 and 7.8 many times in various cubes Y. To facilitate our analysis, we need a flexible notation that can deal with different cubes at the same time. For this reason we now introduce the following notation depending on l, which is similar to the notation of Smith [22, p. 8].

- X will always denote a product set
$$X = X_1\times\cdots\times X_d$$
with X_i disjoint finite sets consisting of primes all equal to 0 or 1 modulo l.
- For $S\subseteq[d]$, we define
$$\overline{X}_S = \Big(\prod_{i\in S}X_i^l\Big)\times\prod_{i\in[d]-S}X_i.$$
Furthermore, write π_i for the projection map from $\overline{X}_S$ to X_i^l if $i\in S$, and the projection map from $\overline{X}_S$ to X_i if $i\in[d]-S$. If $\bar{x}\in\overline{X}_S$, we will sometimes call $\bar{x}$ a *cube*.
- The natural projection maps from X_i^l to X_i are denoted by $\mathrm{pr}_1,\dots,\mathrm{pr}_l$.
- For $S,S_0\subseteq[d]$, we let π_{S,S_0} be the projection map from $\overline{X}_S$ to
$$\Big(\prod_{i\in S\cap S_0}X_i^l\Big)\times\prod_{i\in([d]-S)\cap S_0}X_i$$
given by π_i on each $i\in S_0$. When the set S is clear from context, we will simply write π_{S_0} instead of π_{S,S_0}.
- Let $\bar{x}\in\overline{X}_S$, $T\subseteq S\subseteq[d]$ and put $U:=S-T$. Then we define $\bar{x}(T)$ to be the multiset with underlying set
$$\{\bar{y}\in\overline{X}_T:\pi_{[d]-U}(\bar{y})=\pi_{[d]-U}(\bar{x})\text{ and }\forall i\in U\ \exists j\in[l]:\pi_i(\bar{y})=\mathrm{pr}_j(\pi_i(\bar{x}))\}.$$

We define the *multiplicity* of $\bar{y}$ in $\bar{x}(T)$ to be

$$\prod_{i \in U} |\{j \in [l] : \pi_i(\bar{y}) = \mathrm{pr}_j(\pi_i(\bar{x}))\}|.$$

With this notation one element $\bar{x} \in \overline{X}_S$ corresponds to a cube Y in the previous subsections. This is extremely convenient. For example, with this new notation we can rephrase Theorem 7.8 as

$$\sum_{x \in \bar{z}(\emptyset)} F(x, b) = \sum_{q|b} \epsilon_b(q) \phi_{S,\bar{z}}(\mathrm{Frob}_q),$$

where $\bar{z} \in \overline{X}_S$ and $F(x, b)$ is some Artin pairing. Here we suppress the dependence on the amalgama. We will frequently suppress this dependence in the remainder of the paper, in particular for $\phi_{S,\bar{z}}$.

In the previous sections we have defined expansion maps and raw cocycles. Those objects are rather complicated to work with directly. Instead, we abstract their most important properties in the following combinatorial structure, which we call an l-additive system. The material in this section is an adaptation of Sections 3 and 4 of Smith [22].

Definition 8.1. An *l-additive system* $\mathfrak{A}$ on $X = X_1 \times \cdots \times X_d$ is a tuple

$$(\overline{Y}_S, \overline{Y}_S^{\circ}, F_S, A_S)_{S \subseteq [d]}$$

indexed by the subsets S of $[d]$ with the following properties:

- for all $S \subseteq [d]$, A_S is a finite $\mathbb{F}_l$-vector space and $\overline{Y}_S^{\circ} \subseteq \overline{Y}_S \subseteq \overline{X}_S$;
- for all $S \subseteq [d]$ with $S \neq \emptyset$, we have

$$\overline{Y}_S = \{\bar{x} \in \overline{X}_S : \bar{x}(T) \subseteq \overline{Y}_T^{\circ} \text{ for all } T \subsetneq S\};$$

- for all $S \subseteq [d]$, $F_S : \overline{Y}_S \to A_S$ is a function and

$$\overline{Y}_S^{\circ} = \{\bar{x} \in \overline{Y}_S : F_S(\bar{x}) = 0\};$$

- (Additivity) for $s \in S$ and $\bar{x}_1, \ldots, \bar{x}_{l+1} \in \overline{Y}_S$, if

$$\pi_{[d]-\{s\}}(\bar{x}_1) = \cdots = \pi_{[d]-\{s\}}(\bar{x}_{l+1})$$

and there exist $p_1, \ldots, p_{l-1}, q_1, \ldots, q_l \in X_s$ such that

$$\pi_s(\bar{x}_1) = (p_1, \ldots, p_{l-1}, q_1), \ \ldots, \ \pi_s(\bar{x}_l) = (p_1, \ldots, p_{l-1}, q_l)$$

and $\pi_s(\bar{x}_{l+1}) = (q_1, \ldots, q_l)$, then

$$F_S(\bar{x}_1) + \cdots + F_S(\bar{x}_l) = F_S(\bar{x}_{l+1}).$$

We will sometimes write $\overline{Y}_S(\mathfrak{A})$, $\overline{Y}_S^{\circ}(\mathfrak{A})$, $F_S(\mathfrak{A})$ and $A_S(\mathfrak{A})$ for the data associated to an l-additive system $\mathfrak{A}$.

We remark that the condition $\bar{x}_{l+1} \in \overline{Y}_S$ may be dropped, since it follows from the other conditions. The minimality and agreement conditions from Theorems 7.7 and 7.8 can naturally be encoded in an l-additive system. Although we could already do this now, we postpone this task until Lemma 13.10. Similarly, the existence of the maps $\phi_{S,\bar{x}}$ can be encoded in an l-additive system.

For our analytic techniques to work, it is essential that we can apply Theorems 7.7 and 7.8 to many different cubes $\bar{x}$; the more cubes to which we can apply these theorems, the better. It is for this reason that we need to give a lower bound for the density of $\overline{Y}^\circ_S$ in $\overline{X}_S$ for an l-additive system. Since l-additive systems are purely combinatorial objects, we will state our theorems for general finite sets, not just sets of primes.

Proposition 8.2. *Let $X = X_1 \times \cdots \times X_d$ be a product of finite sets and let $\mathfrak{A}$ be an l-additive system on X. Let δ be the density of $\overline{Y}^\circ_\emptyset$ in X and put*

$$a := \max_{S \subseteq [d]} |A_S|.$$

Then the density of $\overline{Y}^\circ_S$ in $\overline{X}_S$ is lower bounded by $\delta^{l^{|S|}} a^{-(l+1)^{|S|}+1}$ for all subsets S of $[d]$.

Proof. For $S = \emptyset$ this is clear, so from now on we assume that $S \neq \emptyset$. Fix a choice of $s \in S$ for the remainder of the proof. Define for $\bar{x}_0 \in \overline{X}_{S-\{s\}}$

$$V(\bar{x}_0) := \{\bar{y} \in \overline{Y}^\circ_{S-\{s\}} : \pi_{[d]-\{s\}}(\bar{y}) = \pi_{[d]-\{s\}}(\bar{x}_0)\},$$
$$W(\bar{x}_0) := \{\bar{y} \in \overline{Y}^\circ_S : \pi_{[d]-\{s\}}(\bar{y}) = \pi_{[d]-\{s\}}(\bar{x}_0)\}.$$

There are natural injective maps from $V(\bar{x}_0)$ to both X_s and $\overline{X}_{S-\{s\}}$. The former map is given by sending $\bar{y}$ to $\pi_{\{s\}}(\bar{y})$, while the latter is the inclusion $\overline{Y}^\circ_{S-\{s\}} \subseteq \overline{X}_{S-\{s\}}$. Similarly, there are natural injective maps from $W(\bar{x}_0)$ to $V(\bar{x}_0)^l \subseteq X_s^l$ and $\overline{X}_S$. We claim that

$$|W(\bar{x}_0)| \geq \left(|V(\bar{x}_0)|/a^{(l+1)^{|S|-1}}\right)^l. \tag{8.1}$$

If $V(\bar{x}_0)$ is the empty set, (8.1) clearly holds. So suppose that $V(\bar{x}_0)$ is not empty and choose $l-1$ elements $\bar{x}_{1,1}, \ldots, \bar{x}_{1,l-1}$ from $V(\bar{x}_0)$. We define an equivalence relation $\sim_1$ on $V(\bar{x}_0)$ by declaring $\bar{y}_1 \sim_1 \bar{y}_2$ if and only if for all subsets T satisfying $\{s\} \subseteq T \subseteq S$ and $|T| = 1$ and all $\bar{y}_1' \in \bar{y}_1(T - \{s\})$, $\bar{y}_2' \in \bar{y}_2(T - \{s\})$ satisfying $\pi_{[d]-T}(\bar{y}_1') = \pi_{[d]-T}(\bar{y}_2')$ we have

$$F_T(\bar{x}_{1,1}', \ldots, \bar{x}_{1,l-1}', \bar{y}_1') = F_T(\bar{x}_{1,1}', \ldots, \bar{x}_{1,l-1}', \bar{y}_2'), \tag{8.2}$$

where $\bar{x}_{1,1}', \ldots, \bar{x}_{1,l-1}'$ are the unique elements of $\bar{x}_{1,1}(T - \{s\}), \ldots, \bar{x}_{1,l-1}(T - \{s\})$ satisfying

$$\pi_{[d]-T}(\bar{y}_1') = \pi_{[d]-T}(\bar{y}_2') = \pi_{[d]-T}(\bar{x}_{1,1}') = \cdots = \pi_{[d]-T}(\bar{x}_{1,l-1}').$$

Here we remark that the tuple $(\bar{x}_{1,1}', \ldots, \bar{x}_{1,l-1}', \bar{y}_i')$ can naturally be seen as an ele-

ment of $\overline{X}_T$, so (8.2) makes sense. There are at most $a^{l^{|S|-1}}$ equivalence classes. Hence there exists an equivalence class $[\bar{y}]$ with at least $|V(\bar{x}_0)|/a^{l^{|S|-1}}$ elements. Now choose $\bar{x}_{2,1}, \dots, \bar{x}_{2,l-1} \in [\bar{y}]$ and define a new equivalence relation $\sim_2$ on $[\bar{y}]$ by declaring $\bar{y}_1 \sim_2 \bar{y}_2$ if and only if for all subsets T satisfying $\{s\} \subseteq T \subseteq S$ with $|T| = 2$ and all $\bar{y}_1' \in \bar{y}_1(T - \{s\})$, $\bar{y}_2' \in \bar{y}_2(T - \{s\})$ satisfying $\pi_{[d]-T}(\bar{y}_1') = \pi_{[d]-T}(\bar{y}_2')$ we have

$$F_T(\bar{x}_{2,1}', \dots, \bar{x}_{2,l-1}', \bar{y}_1') = F_T(\bar{x}_{2,1}', \dots, \bar{x}_{2,l-1}', \bar{y}_2'), \tag{8.3}$$

where $\bar{x}_{2,1}', \dots, \bar{x}_{2,l-1}'$ are the unique elements of $\bar{x}_{2,1}(T - \{s\}), \dots, \bar{x}_{2,l-1}(T - \{s\})$ satisfying

$$\pi_{[d]-T}(\bar{y}_1') = \pi_{[d]-T}(\bar{y}_2') = \pi_{[d]-T}(\bar{x}_{2,1}') = \cdots = \pi_{[d]-T}(\bar{x}_{2,l-1}').$$

Since the domain of F_T is $\overline{Y}_T$, (8.3) only makes sense if

$$(\bar{x}_{2,1}', \dots, \bar{x}_{2,l-1}', \bar{y}_i') \in \overline{Y}_T.$$

This follows from the construction of $\sim_1$ and additivity.

We inductively proceed until we reach $\sim_{|S|}$. A computation shows that

$$\prod_{\{s\} \subseteq T \subseteq S} |A_T|^{l^{|S|-|T|}} \leq \prod_{i=0}^{|S|-1} a^{\binom{|S|-1}{i} l^i} = a^{(l+1)^{|S|-1}}.$$

Then we find that there is an equivalence class of $\sim_{|S|}$ with at least $\frac{|V(\bar{x}_0)|}{a^{(l+1)^{|S|-1}}}$ elements. Suppose that $\{\bar{y}_1, \dots, \bar{y}_k\}$ is an equivalence class of $\sim_{|S|}$. From additivity we find that $(\bar{y}_{i_1}, \dots, \bar{y}_{i_l}) \in W(\bar{x}_0)$ for all choices of $1 \leq i_1, \dots, i_l \leq k$, where we recall that $W(\bar{x}_0)$ can be identified as a subset of $V(\bar{x}_0)^l$. Hence

$$|W(\bar{x}_0)| \geq \left(|V(\bar{x}_0)|/a^{(l+1)^{|S|-1}}\right)^l,$$

establishing (8.1). Define δ_T to be the density of $\overline{Y}_T^\circ$ in $\overline{X}_T$, so in particular $\delta = \delta_\emptyset$. Also let $\delta_{\bar{x}_0}$ be the density of $V(\bar{x}_0)$ in X_s. Then the density of $V(\bar{x}_0)^l$ in X_s^l is equal to $\delta_{\bar{x}_0}^l$. Since $\overline{Y}_S^\circ$ is the disjoint union of $W(\bar{x}_0)$ over all $\bar{x}_0 \in \pi_{[d]-\{s\}}(\overline{X}_{S-\{s\}})$, it follows from (8.1) that

$$\begin{aligned}
\delta_S &= \sum_{\bar{x}_0} \frac{|W(\bar{x}_0)|}{|\overline{X}_S|} \geq a^{-l\cdot(l+1)^{|S|-1}} \cdot \sum_{\bar{x}_0} \frac{|V(\bar{x}_0)|^l}{|\overline{X}_S|} \\
&= a^{-l\cdot(l+1)^{|S|-1}} \cdot \sum_{\bar{x}_0} \frac{\delta_{\bar{x}_0}^l}{|\pi_{[d]-\{s\}}(\overline{X}_{S-\{s\}})|} \\
&\geq a^{-l\cdot(l+1)^{|S|-1}} \cdot \left(\sum_{\bar{x}_0} \frac{\delta_{\bar{x}_0}}{|\pi_{[d]-\{s\}}(\overline{X}_{S-\{s\}})|}\right)^l.
\end{aligned}$$

We observe that $\delta_{S-\{s\}}$ is the average of $\delta_{\bar{x}_0}$ over all $\bar{x}_0 \in \pi_{[d]-\{s\}}(\overline{X}_{S-\{s\}})$. This shows

$$\delta_S \geq a^{-l\cdot(l+1)^{|S|-1}} \cdot \delta_{S-\{s\}}^l \geq a^{-(l+1)^{|S|}} \cdot \delta_{S-\{s\}}^l. \tag{8.4}$$

Repeated application of (8.4) yields the proposition. ∎

Proposition 8.2 shows that there are many $\bar{x} \in \overline{Y}_S^\circ$. The proof of Proposition 8.2 heavily relies on the special structure of l-additive systems. It will also be important to find $\bar{x}$ with $\bar{x}(\emptyset) \subseteq \overline{Y}_\emptyset^\circ$. Unlike l-additive systems, the set $\overline{Y}_\emptyset^\circ$ has very little structure. Instead we have to rely on Ramsey theory to find such $\bar{x}$.

Proposition 8.3. *Let d be a positive integer and let $X_1, \dots, X_d$ be finite sets all of cardinality at least $n > 0$. Let Y be a subset of $X = X_1 \times \cdots \times X_d$ of density at least $\delta > 0$. Let r be a positive integer satisfying*

$$r \leq n \cdot (2^{-d-1}\delta)^{2r^{d-1}}.$$

Then there are subsets $Z_i \subseteq X_i$ all of cardinality r such that

$$Z_1 \times \cdots \times Z_d \subseteq Y.$$

Proof. This is proven in Proposition 4.1 of Smith [22]. ∎

Before we move on, we explain our strategy for proving our main theorems. In Section 7 we have seen that under suitable conditions on $\bar{x}$,

$$\sum_{x \in \bar{x}(\emptyset)} F(x) = g(\bar{x}), \tag{8.5}$$

where F is a class group pairing and $g(\bar{x})$ is an Artin symbol in a relative governing field. If we could directly get a handle on F, we would be done, but this seems to be completely out of reach with the current methods available. And indeed, (8.5) would be of little help in such a strategy.

Instead, we will take the following approach that uses (8.5) in an essential way. First of all, observe that given g there are many functions F satisfying (8.5). Our goal will be to find one function g for which all F satisfying (8.5) are equidistributed. Obviously, such a conclusion is only possible if we know that (8.5) holds for many $\bar{x}$, and it is here that Propositions 8.2 and 8.3 are essential. Then we use the Chebotarev density theorem to make this function g many times, which allows us to conclude equidistribution of F. Our next definition formalizes these ideas.

Definition 8.4. Let $X_1, \dots, X_d$ be finite non-empty sets, and put $X = X_1 \times \cdots \times X_d$. Let $S \subseteq [d]$ be a set with $|S| \geq 2$. For $Z \subseteq X$ define $\mathbb{F}_l$-vector spaces V and W by

$$V := \{F : Z \to \mathbb{F}_l\}, \quad W := \left\{g : \{\bar{x} \in \overline{X}_S : \bar{x}(\emptyset) \subseteq Z\} \to \mathbb{F}_l\right\}.$$

Let $d : V \to W$ be the linear map given by

$$dF(\bar{x}) = \sum_{x \in \bar{x}(\emptyset)} F(x),$$

where we remind the reader that $\bar{x}(\emptyset)$ is a multiset. Equivalently,

$$dF(\bar{x}) = \sum_{x \in \mathrm{Set}(\bar{x}(\emptyset))} \Big(\prod_{i \in S} |\{j \in [l] : \mathrm{pr}_j(\pi_i(\bar{x})) = \pi_i(x)\}| \Big) F(x).$$

For $\epsilon > 0$ a real number, we say that $F : Z \to \mathbb{F}_l$ is *ϵ-balanced* if for all $a \in \mathbb{F}_l$,

$$(1/l - \epsilon) \cdot |Z| \leq |F^{-1}(a)| \leq (1/l + \epsilon) \cdot |Z|,$$

and F is *ϵ-unbalanced* otherwise. Define $\mathscr{G}_S(Z) := \operatorname{im} d$ and

$$\mathscr{G}_S(\epsilon, Z) := \{g \in \mathscr{G}_S(Z) : g = dF \text{ for some } \epsilon\text{-unbalanced F}\}.$$

Lemma 8.5. *Let X, Z, S and d be as in Definition 8.4 such that $|\pi_{[d]-S}(Z)| = 1$. Further suppose that $\delta > 0$ satisfies*

$$|Z| \geq \delta \cdot |\pi_S(X)|.$$

If $|X_i| \geq n$ for all $i \in S$, then for all $\epsilon > 0$,

$$\frac{|\mathscr{G}_S(\epsilon, Z)|}{|\mathscr{G}_S(Z)|} \leq 2 \cdot l \cdot \exp\big(|\pi_S(X)| \cdot (-\delta \cdot \epsilon^2 + \log l \cdot 2^{|S|+2} \cdot n^{-1/l^{|S|}})\big).$$

Proof. Recall that a cube $\bar{z} \in \overline{X}_S$ is called *degenerate* if there is $i \in S$ such that

$$|\{\mathrm{pr}_1(\pi_{\{i\}}(\bar{z})), \ldots, \mathrm{pr}_l(\pi_{\{i\}}(\bar{z}))\}| = 1.$$

Let Z' be a maximal subset of Z such that all cubes $\bar{z} \in \overline{X}_S$ with $\bar{z}(\emptyset) \subseteq Z'$ are degenerate. Let $F : Z \to \mathbb{F}_l$ be a map with $F(x) \neq 0$ for some $x \in Z - Z'$ and $F(x) = 0$ for all $x \in Z'$. We claim that F is not in the kernel of the linear map $d : V \to W$. Indeed, consider the set $Z' \cup \{x\}$. By construction of Z', we find a non-degenerate $\bar{z} \in \overline{X}_S$ with $\bar{z}(\emptyset) \subseteq Z' \cup \{x\}$ and $x \in \bar{z}(\emptyset)$. Then $dF(\bar{z}) \neq 0$, establishing our claim.

From our claim we deduce that the kernel of d is of size at most $l^{|Z'|}$. On the other hand, Proposition 8.3 with $r = l$ yields

$$|Z'| \leq |\pi_S(X)| \cdot 2^{|S|+2} \cdot n^{-1/l^{|S|}}$$

and hence

$$|\mathscr{G}_S(Z)| \geq l^{|Z|-|Z'|} \geq l^{|Z|} \cdot \exp(-\log l \cdot |\pi_S(X)| \cdot 2^{|S|+2} \cdot n^{-1/l^{|S|}}). \tag{8.6}$$

From Hoeffding's inequality and a straightforward union bound we find that the number of ϵ-unbalanced F is bounded by

$$2 \cdot l^{|Z|+1} \cdot \exp(-2 \cdot \epsilon^2 \cdot |Z|).$$

We conclude that

$$|\mathscr{G}_S(\epsilon, Z)| \leq 2 \cdot l^{|Z|+1} \cdot \exp(-2 \cdot \epsilon^2 \cdot |Z|) \leq 2 \cdot l^{|Z|+1} \cdot \exp(-|\pi_S(X)| \cdot \delta \cdot \epsilon^2). \tag{8.7}$$

The lemma follows upon combining (8.6) and (8.7). ∎

Lemma 8.5 is very much in the spirit of the strategy we outlined earlier. Unfortunately, we do not have equality (8.5) for all $\bar{x} \in \overline{X}_S$ with $\bar{x}(\emptyset) \subseteq \overline{Y}^\circ_\emptyset$. Instead, we will show in Lemma 13.10 that (8.5) holds under the much stronger condition that $\bar{x}(T) \cap \overline{Y}^\circ_T(\mathfrak{A})$ is "large" for all proper subsets T of S, where $\mathfrak{A}$ is a completely explicit l-additive system.

Fortunately, it turns out that $\overline{Y}^\circ_\emptyset$ has some special structure in our application. Namely, in Lemma 13.10 we will prove that for "sufficiently nice" $\bar{x} \in \overline{X}_S$ we have $\bar{x}(\emptyset) \subseteq \overline{Y}^\circ_\emptyset$. Our next definition formalizes what we mean by "sufficiently nice" $\bar{x} \in \overline{X}_S$.

Definition 8.6. For an l-additive system $\mathfrak{A}$ on X define

$$\overline{Z}_S(\mathfrak{A}) := \bigcap_{i \in S} \left\{ \bar{x} \in \overline{X}_S : |\pi_i(\bar{x}(S-\{i\}) \cap \overline{Y}^\circ_{S-\{i\}}(\mathfrak{A}))| \geq \max(1, |\pi_i(\bar{x}(S-\{i\}))|-1) \right\},$$

where π_i of a multiset is defined to be π_i of the underlying set. Let $a \geq 2$ be an integer. Call an l-additive system $\mathfrak{A}$ on X *S-acceptable* if

- $|A_T(\mathfrak{A})| \leq a$ for all subsets T of S;
- if $\bar{x}$ is in $\overline{Z}_S(\mathfrak{A})$, then $\mathrm{Set}(\bar{x}(\emptyset)) \subseteq \overline{Y}^\circ_\emptyset(\mathfrak{A})$.

Before we state the next proposition, we explain why we will vary $\mathfrak{A}$ over all l-additive systems on X in this proposition instead of just the special l-additive system from Lemma 13.10. To prove our equidistribution statements in the final section, we consider a large interval of primes. Using Chebotarev we split this interval into many sets $A_1, \ldots, A_k$.

Then we apply our next proposition to every A_i with $\mathfrak{A}$ equal to the l-additive system from Lemma 13.10 restricted to A_i. Since we have no control over the restriction of this l-additive system to a smaller subset, we simply consider all l-additive systems provided that $\overline{Y}^\circ_\emptyset$ has the special property in Definition 8.6.

Proposition 8.7. *There exists an absolute constant $A > 0$ such that the following holds. Let X and S be as in Definition 8.4. Let $a \geq 2$, $\epsilon > 0$ and define $n := \min_{i \in S} X_i$. Suppose that $|\pi_{[d]-S}(X)| = 1$, $\epsilon < a^{-1}$ and*

$$\log n \geq A \cdot (l \cdot (l+1))^{|S|+3} \cdot \log \epsilon^{-1}.$$

Then there exists $g \in \mathscr{G}_S(X)$ such that for all S-acceptable l-additive systems $\mathfrak{A}$ at S on X and for all $F : \overline{Y}^\circ_\emptyset(\mathfrak{A}) \to \mathbb{F}_l$ satisfying

$$dF(\bar{x}) = g(\bar{x}) \tag{8.8}$$

for all $\bar{x} \in \overline{Z}_S(\mathfrak{A})$, the map F is $\frac{|X|}{|\overline{Y}^\circ_\emptyset(\mathfrak{A})|} \cdot \epsilon$-balanced. In case $|\overline{Y}^\circ_\emptyset(\mathfrak{A})| = 0$, this is to be interpreted as ∞-balanced.

Proof. Let $\mathscr{G}_S(\epsilon, a, X)$ be the set of those $g \in \mathscr{G}_S(X)$ that fail to satisfy the conclusion of Proposition 8.7. We claim that

$$\frac{|\mathscr{G}_S(\epsilon, a, X)|}{|\mathscr{G}_S(X)|} \leq \exp(-|X| \cdot \epsilon^{6+(l+1)^{|S|+3}}), \tag{8.9}$$

which immediately yields the proposition. Define

$$\delta := |\overline{Y}^{\circ}_{\emptyset}(\mathfrak{A})|/|X|.$$

Let $g \in \mathscr{G}_S(\epsilon, a, X)$. Then, from the definition of $\mathscr{G}_S(\epsilon, a, X)$, there exists an S-acceptable l-additive system $\mathfrak{A}$ on X and a $\delta^{-1}\epsilon$-unbalanced $F : \overline{Y}^{\circ}_{\emptyset}(\mathfrak{A}) \to \mathbb{F}_l$ satisfying $dF(\bar{x}) = g(\bar{x})$ for all $\bar{x} \in \overline{Z}_S(\mathfrak{A})$. For $f : [l-1] \to \overline{Y}^{\circ}_{\emptyset}(\mathfrak{A})$ and $x \in \overline{Y}^{\circ}_{\emptyset}(\mathfrak{A})$ we let $c(f, x)$ be the unique element of $\overline{X}_S$ satisfying

$$\pi_i(f(j)) = \mathrm{pr}_j(\pi_i(c(f, x))) \quad \text{and} \quad \pi_i(x) = \mathrm{pr}_l(\pi_i(c(f, x)))$$

for $i \in S$ and $j \in [l-1]$. Next define

$$Z_S(\mathfrak{A}, f) := \{x \in X : \text{ writing } \bar{x} := c(f, x), \text{ we have } \bar{y} \in \overline{Y}^{\circ}_{T}(\mathfrak{A}) \text{ for all } T \subsetneq S \\ \text{and all } \bar{y} \in \bar{x}(T) \text{ satisfying } f([l-1]) \cap \bar{y}(\emptyset) \neq \emptyset\}.$$

Note that $x \in Z_S(\mathfrak{A}, f)$ implies $x \in \overline{Y}^{\circ}_{\emptyset}(\mathfrak{A})$. There is a natural injective map from $\overline{Y}^{\circ}_{S}(\mathfrak{A})$ to

$$\coprod_{f : [l-1] \to \overline{Y}^{\circ}_{\emptyset}(\mathfrak{A})} Z_S(\mathfrak{A}, f),$$

where $\coprod$ denotes disjoint union. This map is given by sending $\bar{y}$ to the pair (f, x), where $f : [l-1] \to \overline{Y}^{\circ}_{\emptyset}(\mathfrak{A})$ and $x \in Z_S(\mathfrak{A}, f)$ are uniquely determined by

$$\pi_i(f(j)) = \mathrm{pr}_j(\pi_i(\bar{y})) \quad \text{and} \quad \pi_i(x) = \mathrm{pr}_l(\pi_i(\bar{y}))$$

for $i \in S$ and $j \in [l-1]$. We conclude that

$$|\overline{Y}^{\circ}_{S}(\mathfrak{A})| \leq \sum_{f : [l-1] \to \overline{Y}^{\circ}_{\emptyset}(\mathfrak{A})} |Z_S(\mathfrak{A}, f)| \tag{8.10}$$

$$\leq |X|^{l-1} \cdot \max_{f : [l-1] \to \overline{Y}^{\circ}_{\emptyset}(\mathfrak{A})} |Z_S(\mathfrak{A}, f)|. \tag{8.11}$$

Since F is $\delta^{-1}\epsilon$-unbalanced, it follows that $\epsilon/2 \leq \delta$. Hence the density of $\overline{Y}^{\circ}_{\emptyset}(\mathfrak{A})$ in X is at least $\epsilon/2$. From Proposition 8.2 we see that the density of $\overline{Y}^{\circ}_{S}(\mathfrak{A})$ in $\overline{X}_S$ is lower bounded by

$$\delta^{l^{|S|}} a^{-(l+1)^{|S|}+1} \geq (\epsilon/2)^{l^{|S|}} \epsilon^{(l+1)^{|S|}+1} \geq \epsilon^{(l+1)^{|S|}+3}. \tag{8.12}$$

Upon combining (8.10) and (8.12) we find that there exists $f_1 : [l-1] \to \overline{Y}^{\circ}_{\emptyset}(\mathfrak{A})$ such that $Z_S(\mathfrak{A}, f_1)$ has density at least $\epsilon^{(l+1)^{|S|}+3}$ in X. If the complement $\overline{Y}^{\circ}_{\emptyset}(\mathfrak{A}) - Z_S(\mathfrak{A}, f_1)$ has density at least $\epsilon/2$ in X, we can repeat this argument with the S-acceptable l-additive system $\mathfrak{A}'$ on X given by

$$\overline{Y}^{\circ}_{\emptyset}(\mathfrak{A}') := \overline{Y}^{\circ}_{\emptyset}(\mathfrak{A}) - Z_S(\mathfrak{A}, f_1)$$

and the same maps F_T and groups A_T as for $\mathfrak{A}$. Hence we find a sequence of functions $f_1, \dots, f_r : [l-1] \to \overline{Y}^\circ_\emptyset(\mathfrak{A})$ such that

$$Z_S(\mathfrak{A}, f_j) - Z_S(\mathfrak{A}, f_{j-1}) - \cdots - Z_S(\mathfrak{A}, f_1)$$

has density at least $\epsilon^{(l+1)^{|S|+3}}$ in X for $j \geq 1$ and so that

$$\overline{Y}^\circ_\emptyset(\mathfrak{A}) - Z_S(\mathfrak{A}, f_r) - \cdots - Z_S(\mathfrak{A}, f_1)$$

has density at most $\epsilon/2$ in X. Define

$$Z'_S(\mathfrak{A}, j) := Z_S(\mathfrak{A}, f_j) - Z_S(\mathfrak{A}, f_{j-1}) - \cdots - Z_S(\mathfrak{A}, f_1).$$

Then there exists a j such that F is $\epsilon/2$-unbalanced when restricted to $Z'_S(\mathfrak{A}, j)$. If $\bar{x} \in \overline{X}_S$ satisfies $\mathrm{Set}(\bar{x}(\emptyset)) \subseteq Z'_S(\mathfrak{A}, j)$, (8.8) combined with the additivity of dF and g imply $dF(\bar{x}) = g(\bar{x})$. From this we deduce that $g \in \mathscr{G}_S(\epsilon, a, X)$ implies

$$g|_{\{\bar{x} \in \overline{X}_S : \mathrm{Set}(\bar{x}(\emptyset)) \subseteq Z'_S(\mathfrak{A}, j)\}} \in \mathscr{G}_S(\epsilon/2, Z'_S(\mathfrak{A}, j))$$

for some S-acceptable l-additive system $\mathfrak{A}$ on X and some j. We see from Lemma 8.5 that the number of $g \in \mathscr{G}_S(X)$ with $g|_{\{\bar{x} \in \overline{X}_S : \mathrm{Set}(\bar{x}(\emptyset)) \subseteq Z'_S(\mathfrak{A}, j)\}} \in \mathscr{G}_S(\epsilon/2, Z'_S(\mathfrak{A}, j))$ is bounded by

$$2 \cdot l \cdot |\mathscr{G}_S(X)| \cdot \exp\big(|X| \cdot (-\epsilon^{4+(l+1)^{|S|+3}} + \log l \cdot 2^{|S|+2} \cdot n^{-1/l^{|S|}})\big). \tag{8.13}$$

For A sufficiently large we can simplify (8.13) as

$$|\mathscr{G}_S(X)| \cdot \exp(-|X| \cdot \epsilon^{5+(l+1)^{|S|+3}}). \tag{8.14}$$

Let us now give an upper bound for the number of subsets E of X such that there exists an S-acceptable l-additive system $\mathfrak{A}$ on X and $f : [l-1] \to \overline{Y}^\circ_\emptyset(\mathfrak{A})$ satisfying $E = Z_S(\mathfrak{A}, f)$. A straightforward computation shows that

$$Z_S(\mathfrak{A}, f) = \{x \in X : \pi_{S-\{i\}}(x) \in \pi_{S-\{i\}}(Z_S(\mathfrak{A}, f)) \text{ for all } i \in S\}.$$

Therefore, $Z_S(\mathfrak{A}, f)$ is determined by the sets $\pi_{S-\{i\}}(Z_S(\mathfrak{A}, f))$ as i varies through S. From this, we obtain the following upper bound for the number of possible sets E:

$$2^{|X| \cdot \sum_{i \in S} \frac{1}{|X_i|}} \leq 2^{|X| \cdot |S| \cdot n^{-1}}.$$

Hence there are at most $2^{r \cdot |X| \cdot |S| \cdot n^{-1}}$ sequences $Z'_S(\mathfrak{A}, f_1), \dots, Z'_S(\mathfrak{A}, f_r)$ and at most r choices of j. Multiplying this with the bound from (8.14) we conclude that

$$|\mathscr{G}_S(\epsilon, a, X)| \leq r \cdot 2^{r \cdot |X| \cdot |S| \cdot n^{-1}} \cdot |\mathscr{G}_S(X)| \cdot \exp\big(-|X| \cdot \epsilon^{5+(l+1)^{|S|+3}}\big). \tag{8.15}$$

Using $r \leq \epsilon^{-(l+1)^{|S|+3}}$ and (8.15) we infer

$$\begin{aligned} |\mathscr{G}_S(\epsilon, a, X)| &\leq |\mathscr{G}_S(X)| \cdot \exp\big(|X| \cdot (-\epsilon^{5+(l+1)^{|S|+3}} + r \cdot |S| \cdot n^{-1})\big) \\ &\leq |\mathscr{G}_S(X)| \cdot \exp\big(-|X| \cdot \epsilon^{6+(l+1)^{|S|+3}}\big) \end{aligned}$$

for A sufficiently large. This establishes (8.9), completing our proof. ∎

9. Governing expansions

We will heavily use the notation introduced at the beginning of Section 8; recall that this notation implicitly depends on l. Let d be a positive integer and let $X_1, \dots, X_d$ be disjoint sets of primes q that are either 1 modulo l or equal to l. Fix a subset $\overline{Y}_\emptyset$ and a function $f : X_1 \sqcup \cdots \sqcup X_r \to [l-1]$, where we remind the reader that $\sqcup$ denotes disjoint union.

If $\bar{x} \in \overline{X}_S$, we obtain an amalgama ϵ for the cube $\bar{x}$ by restricting f to $\bar{x}$. Fix furthermore an integer $i_a \in [d]$. The coming definition will depend implicitly on the choice of f, and we shall suppress this dependence in the notation. A collection of sets

$$\{\overline{Y}_S\}_{i_a \in S \subseteq d}$$

with $\overline{Y}_S \subseteq \overline{X}_S$, together with a collection of continuous 1-cochains

$$\{\phi_{S,\bar{x}} : G_{\mathbb{Q}} \to \mathbb{F}_l\}_{\bar{x} \in \overline{Y}_S}$$

for each S containing i_a, is said to be a *governing expansion* $\mathscr{G}$ if it satisfies the following requirements:

(1) for each $\bar{x} \in \overline{Y}_{i_a}$ we have

$$\phi_{\{i_a\},\bar{x}} = j_l^{-1} \circ \sum_{i=1}^{l} \chi_{\mathrm{pr}_i(\pi_{i_a}(\bar{x}))}^{f(\mathrm{pr}_i(\pi_{i_a}(\bar{x})))};$$

(2) if $\bar{x}_1, \bar{x}_2 \in \overline{Y}_S$ are cubes with the same multisets X_i for each $i \in S$, then $\phi_{S,\bar{x}_1} = \phi_{S,\bar{x}_2}$;

(3) if $i_a \in S$, then for all $\bar{x} \in \overline{Y}_S$ and all subsets T satisfying $i_a \in T \subseteq S$ we have $\bar{x}_T \in \overline{Y}_T$ for any choice of $\bar{x}_T \in \bar{x}(T)$; moreover, the collection $\{\phi_{T,\bar{x}_T}\}_{i_a \in T}$ is a good expansion for all choices of $\bar{x}_T \in \bar{x}(T)$ (here the collection of subsets of $[m]$ not containing i_a is naturally identified with the collection of subsets of $[m]$ containing i_a);

(4) for each rational prime q that ramifies in $\prod_{\bar{x}} L(\phi_{S,\bar{x}})/\mathbb{Q}$ we choose a generator σ_q of an inertia subgroup; we assume that $\phi_{S,\bar{x}}(\sigma_q) = 0$ for each $\bar{x}$;

(5) if $\bar{x} \in \overline{X}_S$, then $\bar{x} \in \overline{Y}_S$ if and only if the following two conditions are satisfied:
- $\bar{x} \in \overline{Y}_T$ for all subsets T satisfying $i_a \in T \subseteq S$ and all $\bar{x}_T \in \bar{x}(T)$;
- for each $i \in S$, $\mathrm{pr}_1(\pi_i(\bar{x})), \dots, \mathrm{pr}_l(\pi_i(\bar{x}))$ split completely in the field $L(\phi_{S-\{i\},\bar{x}_{S-\{i\}}})$.

These requirements are very similar to those imposed in [22, pp. 10–12], but not completely the same. Using the calculations done in Section 7.2 and following the same proof strategy explained in [22, Proposition 2.3] one obtains the following important fact.

Proposition 9.1. *If $\mathscr{G}$ is a governing expansion, then the assignment $\bar{x} \mapsto \phi_{S,\bar{x}}(\mathscr{G})$ is additive for each S (see Definition 8.1).*

Fix a set S satisfying $i_a \in S \subseteq [d]$. Denote by $\mathcal{A}(\overline{Y}_S, \mathbb{F}_l)$ the $\mathbb{F}_l$-vector space of all additive maps from $\overline{Y}_S$ to $\mathbb{F}_l$. Following the same proof strategy as in [22, Proposition 2.4] one obtains the following proposition.

Proposition 9.2. *Let $\mathcal{G}$ be a governing expansion. The assignment*

$$\sigma \mapsto (\bar{x} \mapsto \phi_{S,\bar{x}}(\mathcal{G})(\sigma))_{\bar{x} \in \overline{Y}_S(\mathcal{G})}$$

gives an isomorphism between the group

$$\mathrm{Gal}\Big(\prod_{\bar{x} \in \overline{Y}_S(\mathcal{G})} L(\phi_{S,\bar{x}}(\mathcal{G})) \Big/ \prod_{i_a \in T \subsetneq S} \prod_{\bar{x} \in \overline{Y}_T(\mathcal{G})} L(\phi_{T,\bar{x}}(\mathcal{G})) \Big)$$

and the space $\mathcal{A}(\overline{Y}_S(\mathcal{G}), \mathbb{F}_l)$ of additive maps.

It will be important in our main application to recognize that for the product space $X = X_1 \times \cdots \times X_d$ the space of additive maps $\mathcal{A}(\overline{X}_S, \mathbb{F}_l)$ is equal to $\mathcal{G}_S(X)$ as defined in Definition 8.4.

Proposition 9.3. *The image of the map $d : \mathrm{Map}(X, \mathbb{F}_l) \to \mathrm{Map}(\overline{X}_S, \mathbb{F}_l)$ is equal to $\mathcal{A}(\overline{X}_S, \mathbb{F}_l)$. Furthermore, the dimension of $\mathcal{A}(\overline{X}_S, \mathbb{F}_l)$ is*

$$\prod_{i \in S} (|X_i| - 1) \cdot \prod_{j \in [d]-S} |X_j|.$$

Proof. It is a triviality that

$$\mathcal{G}_S(X) \subseteq \mathcal{A}(\overline{X}_S, \mathbb{F}_l). \tag{9.1}$$

We will now establish that $\mathcal{A}(\overline{X}_S, \mathbb{F}_l) \subseteq \mathcal{G}_S(X)$. To do so we pick once and for all a point $x_0 \in X$. We define

$$\mathrm{Max}(x_0) := \{x \in X : \exists i \in S \text{ with } \pi_i(x) = \pi_i(x_0)\}.$$

We observe that $\mathrm{Max}(x_0)$ does not contain any product sets of the form

$$\prod_{i \in S} Y_i \times \prod_{j \in [d]-S} \{y_j\},$$

where each Y_i has precisely two elements. We claim that $\mathrm{Max}(x_0)$ is a maximal subset of X with the above property. Indeed, take any $y \in X - \mathrm{Max}(x_0)$. Then $\pi_i(y) \neq \pi_i(x_0)$ for each i in S. Now define

$$Y(y) := \prod_{i \in S} \{\pi_i(y), \pi_i(x_0)\} \times \prod_{j \in [d]-S} \{\pi_j(y)\},$$

which is clearly contained in $\mathrm{Max}(x_0) \cup \{y\}$. This shows the claim.

Next observe that the size of the complement of $\mathrm{Max}(x_0)$ is trivially

$$\prod_{i\in S}(|X_i|-1)\cdot\prod_{j\in[d]-S}|X_j|.$$

Hence

$$|\mathrm{Max}(x_0)| = |X| - \prod_{i\in S}(|X_i|-1)\cdot\prod_{j\in[d]-S}|X_j|.$$

Following the proof of Proposition 8.5, we find that the kernel of d has dimension at most $|\mathrm{Max}(x_0)|$. This gives

$$\dim_{\mathbb{F}_l}\mathscr{G}_S(X) \geq \prod_{i\in S}(|X_i|-1)\cdot\prod_{j\in[d]-S}|X_j|. \tag{9.2}$$

Finally, consider the set

$$\mathrm{Min}(x_0) := \{\bar{x}\in\overline{X}_S : \pi_S(\bar{x}(\emptyset)) \text{ contains } \pi_S(x_0) \text{ with multiplicity } (l-1)^{|S|}\}.$$

It is not difficult to show that an additive function is completely determined by its restriction to $\mathrm{Min}(x_0)$. Hence we conclude that

$$\dim_{\mathbb{F}_l}\mathscr{A}(\overline{X}_S,\mathbb{F}_l) \leq \prod_{i\in S}(|X_i|-1)\cdot\prod_{j\in[d]-S}|X_j|. \tag{9.3}$$

The proposition follows upon combining (9.1)–(9.3). ∎

Proposition 9.4. *Let $X := X_1\times\cdots\times X_d$ and let $S\subseteq[d]$ with $|S|\geq 2$. Assume that there is some constant A such that $|X_i| = A$ for all $i\in S$. Further assume that there is a governing expansion $\mathfrak{G}$ on X such that $\overline{X}_S = \overline{Y}_S(\mathfrak{G})$. Define*

$$F(X) := \prod_{i\in S}\prod_{p\in X_i}\overline{\mathbb{Q}}^{\ker(\chi_p)}\mathbb{Q}(\zeta_l,\sqrt[l]{p})\prod_{i_a\in T\subsetneq S}\prod_{\bar{x}\in\overline{X}_T}L(\phi_{T,\bar{x}}(\mathfrak{G})),$$

where χ_p is any character from $G_{\mathbb{Q}}$ to $\langle\zeta_l\rangle$ of conductor dividing p^∞. Then the degree of $F(X)$ depends only on A and $|S|$, and we denote it $d(A,|S|)$. If P is a set of primes all equal to 0 or 1 modulo l and disjoint from $\bigcup_{i\in S}X_i$, then moreover

$$\Big(\prod_{\bar{x}\in\overline{Y}_S(\mathfrak{G})}L(\phi_{S,\bar{x}}(\mathfrak{G}))\Big)\cap\Big(\prod_{p\in P}\overline{\mathbb{Q}}^{\ker(\chi_p)}\mathbb{Q}(\zeta_l,\sqrt[l]{p})\Big) = \mathbb{Q}.$$

Let $X' := X_1'\times\cdots\times X_d'$ with the same conditions as for X and suppose that $|X_i\cap X_i'| = 1$ for all $i\in S$. Then the degree of $F(X)F(X')$ is

$$\frac{d(A,|S|)^2}{l^{2|S|}}.$$

10. Prime divisors

In the previous sections it has been very beneficial to work with product spaces of the shape $X = X_1 \times \cdots \times X_r$, where the X_i are disjoint non-empty sets of primes equal to 0 or 1 modulo l. Let $S_r(N, l)$ be the set of squarefree integers of size at most N with exactly r prime divisors all equal to 0 or 1 modulo l. Then there is a natural injective map from X to $S_r(\infty, l)$. To prove our analytic results, we will not work with all product spaces X, but only those that are sufficiently nice. By carefully studying $S_r(N, l)$ we are able to show that most product spaces X have the nice properties we need. The material in this section is directly based on Section 5 of Smith [22].

Definition 10.1. Let $N \geq e^{e^{10 \cdot l}}$ be a large real number and let r be an integer satisfying

$$1 \leq r \leq 2 \log\log N. \tag{10.1}$$

For $n \in S_r(N, l)$ we write $(p_1, \ldots, p_r)$ for the prime divisors of n with $p_1 < \cdots < p_r$.

- If $D_1 > 100$ is a real number, we say that n is *comfortably spaced above* D_1 if for all $i < r$ satisfying $p_i > D_1$,

$$l^{200} D_1 < l^{200} p_i < p_{i+1}.$$

- Let $C_0 > 1$ be a real number. We call n C_0-*regular* if for all $i \leq \frac{1}{3} r$,

$$\left| i - \frac{r \log\log p_i}{\log\log N} \right| < C_0^{1/5} \cdot \max(i, C_0)^{4/5}.$$

We say $X = X_1 \times \cdots \times X_r \subseteq S_r(N, l)$ is C_0-*regular* if some $n \in X$ is C_0-regular.

For a general squarefree integer n, there is a well-known heuristic model for the values of $\log\log p_i$. This heuristic predicts that the values of $\log\log p_i$ for $i = 1, \ldots, r$ behave like a Poisson point process of intensity 1. It is not hard to see that this heuristic breaks down for small and large values of i, but it is nevertheless a solid heuristic: see for example the work of Granville [9].

The heuristic model needs to be slightly modified in our setting. Recall that, loosely speaking, a typical squarefree integer n has roughly $\log\log n$ prime divisors on average with standard deviation $\sqrt{\log\log n}$. We require only very weak conditions on r in Definition 10.1 far outside the typical range. This makes the correction factor $r/\log\log N$ in the definition of C_0-regular necessary.

Assuming the heuristic model, it is an exercise in probability theory to show that most integers are comfortably spaced above D_1 and C_0-regular. This is done in Proposition 5.2 in Smith [22]. Remarkably enough, this proposition is then used to establish the analogous result for the integers.

We will now show that almost all $n \in S_r(N, l)$ are comfortably spaced above D_1 and C_0-regular following the strategy of Smith [22]. Since we are following Smith's strategy, our first goal is to generalize Proposition 5.2 of Smith [22].

Proposition 10.2. *Let $L > 2$ be a real number and let $r \geq 1$ be an integer. Suppose that $X_1, \dots, X_r \sim U(0, L)$ are independent, uniformly distributed random variables. Define $U_{(i)}$ to be the i-th order statistic of $X_1, \dots, X_r$. For a real number $C_0 > 0$, we say that $X_1, \dots, X_r$ are C_0-regular if for all $1 \leq i \leq r$,*

$$\left| i - \frac{rU_{(i)}}{L} \right| < C_0^{1/5} \cdot \max(i, C_0)^{4/5}.$$

Then there is an absolute constant $c > 0$ such that

$$\mathbb{P}(X_1, \dots, X_r \text{ is not } C_0\text{-regular}) = O(\exp(-c \cdot C_0)).$$

Proof. Define $L' := rL/L = r$ and $X_i' := rX_i/L$. Now apply Proposition 5.2 of Smith [22]. ∎

Having established Proposition 10.2, we are ready to study $S_r(N, l)$. In our proofs, we will frequently encounter the following integral

$$I_r(u) := \int_{\substack{t_1, \dots, t_r \geq 1 \\ t_1 + \dots + t_r \leq u}} \frac{dt_1}{t_1} \cdot \dots \cdot \frac{dt_r}{t_r},$$

which was first studied by Ramanujan. It is this integral that provides the connection between $S_r(N, l)$ and the heuristic model. Note that $I_r(u)$ is trivially bounded by $(\log u)^r$. Our next lemma gives a better bound for $I_r(u)$ in some ranges of u and r.

Lemma 10.3. *Let γ be the Euler–Mascheroni constant. Let $u \geq 3$ be a real number and let $r \geq 1$ be an integer. Set $\alpha := r/\log u$. Then*

$$\left| I_r(u) - \frac{e^{-\gamma\alpha}}{\Gamma(1+\alpha)} (\log u)^r \right| = O\left((\alpha + 1)(\log u)^r \frac{(\log\log u)^3}{\log u} \right).$$

Proof. This is Lemma 5.1 of Smith [22]. ∎

Define $S_r'(N, l)$ to be the subset of $S_r(N, l)$ consisting of those integers that are not divisible by l. For technical reasons, it turns out to be more convenient to work with $S_r'(N, l)$.

Theorem 10.4. *Let N and r be as in Definition 10.1.*

- *Suppose that $D_1 > 100$. Then*

$$\frac{|\{n \in S_r'(N, l) : n \text{ is not comfortably spaced above } D_1\}|}{|S_r'(N, l)|} = O\left(\frac{1}{\log D_1} + \frac{1}{\log\log N} \right).$$

- *There is an absolute constant $c > 0$ such that for all $C_0 > 0$,*

$$\frac{|\{n \in S_r'(N, l) : n \text{ is not } C_0\text{-regular}\}|}{|S_r'(N, l)|} = O\big(\exp(-c \cdot C_0) + \exp(-c r^{1/3})\big).$$

Proof. This is a mostly straightforward generalization of Theorem 5.4 in Smith [22]. Define, for any real $x > 1$,

$$F_l(x) := \sum_{\substack{p \le x \\ p \equiv 1 \bmod l}} \frac{1}{p}.$$

We use the Siegel–Walfisz theorem and partial summation to obtain absolute constants $A, c > 0$ such that for all $x \ge e^l$,

$$\left| F_l(x) - \frac{1}{\varphi(l)} \log\log x - B_1(l) \right| \le A \cdot e^{-c\sqrt{\log x}},$$

where $B_1(l)$ is a computable constant in terms of l. In particular, there is a constant $A(l) > 0$ depending only on l and an absolute constant $c > 0$ such that for all $x \ge 1.5$

$$\left| F_l(x) - \frac{1}{\varphi(l)} \log\log x - B_1(l) \right| \le A(l) \cdot e^{-c\sqrt{\log x}}.$$

Denote by $\mathcal{P}_l$ the set of primes equal to 1 modulo l and let T be a subset of $\mathcal{P}_l^r$. Then we define $\mathrm{Grid}(T) \subseteq \mathbb{R}^r$ by

$$\mathrm{Grid}(T) := \bigcup_{(p_1,\dots,p_r)\in T} \prod_{1\le i\le r} \left[\varphi(l) \cdot \left(F_l(p_i) - \frac{1}{p_i} - B_1(l) \right), \varphi(l) \cdot (F_l(p_i) - B_1(l)) \right].$$

Here $\prod$ is to be interpreted as the Cartesian product of intervals. To facilitate our analysis of $S_r'(N, l)$ we define

$$S_r'(N, D, l) := \{n \in S_r'(N, l) : p \mid n \Rightarrow p > D\},$$

where $D > 1.5$ is a real number. Let $\mathbb{R}^r_{\ge B}$ be the subset of $\mathbb{R}^r$ with all coordinates at least B. Define, for $u > 0$ a real number and $r \ge 1$ an integer,

$$V_r(u, D, l) := \{(x_1, \dots, x_r) \in \mathbb{R}^r_{\ge \varphi(l)(F_l(D) - B_1(l))} : e^{x_1} + \cdots + e^{x_r} \le u\}.$$

Put

$$T_r(N, D, l) := \{(p_1, \dots, p_r) \in \mathcal{P}_l^r : p_1 \cdot \ldots \cdot p_r < N,\ p_i > D\}.$$

There exists a constant $\kappa(l)$ depending only on $A(l)$ and c such that

$$\exp\bigl(x + A(l)\exp(-c \cdot e^{x/2})\bigr) - \exp(x) \le \kappa(l).$$

This implies that for a good choice of $A(l)$ and c,

$$V_r(\log N - r\kappa(l), D, l) \subseteq \mathrm{Grid}(T_r(N, D, l)) \subseteq V_r(\log N + r\kappa(l), D, l). \tag{10.2}$$

A change of variables shows

$$\mathrm{Vol}(V_r(u, D, l)) = I_r(e^{-\varphi(l)(F_l(D) - B_1(l))}u). \tag{10.3}$$

Setting $B(D,l) := e^{\varphi(l)(F_l(D)-B_1(l))}$, we deduce from (10.2) and (10.3) that

$$I_r\left(\frac{\log N - r\kappa(l)}{B(D,l)}\right) \leq \varphi(l)^r \sum_{\substack{p_1\cdot\ldots\cdot p_r<N\\ p_1,\ldots,p_r\in\mathcal{P}_l\\ p_1,\ldots,p_r>D}} \frac{1}{p_1\cdot\ldots\cdot p_r} \leq I_r\left(\frac{\log N + r\kappa(l)}{B(D,l)}\right).$$

Now assume that $3\cdot B(D,l)\leq \log N$ and $r^2\leq \log N$. Using the classical differential equation $I_r'(u) = r/u\cdot I_{r-1}(u-1)$ due to Ramanujan and the trivial bound for $I_{r-1}(u-1)$, we obtain

$$\varphi(l)^r \sum_{\substack{p_1\cdot\ldots\cdot p_r<N\\ p_1,\ldots,p_r\in\mathcal{P}_l\\ p_1,\ldots,p_r>D}} \frac{1}{p_1\cdot\ldots\cdot p_r} = I_r\left(\frac{\log N}{B(D,l)}\right) + O\left(\frac{r^2}{\log N}\cdot(\log\log N - \log B(D,l))^{r-1}\right).$$

We introduce the following sums:

$$F_r'(N,D,l) := \varphi(l)^r \sum_{\substack{p_1\cdot\ldots\cdot p_r<N\\ p_1,\ldots,p_r\in\mathcal{P}_l\\ p_1,\ldots,p_r>D}} \frac{1}{p_1\cdot\ldots\cdot p_r},$$

$$G_r'(N,D,l) := \varphi(l)^r \sum_{\substack{p_1\cdot\ldots\cdot p_r<N\\ p_1,\ldots,p_r\in\mathcal{P}_l\\ p_1,\ldots,p_r>D}} \log(p_1\cdot\ldots\cdot p_r),$$

$$H_r'(N,D,l) := \varphi(l)^r \sum_{\substack{p_1\cdot\ldots\cdot p_r<N\\ p_1,\ldots,p_r\in\mathcal{P}_l\\ p_1,\ldots,p_r>D}} 1.$$

Put

$$u := \frac{\log N}{B(D,l)}.$$

Until now we have only assumed that $u\geq 3$ and $r^2\leq \log N$. We additionally suppose that $\log\log N \geq 1.1\log\log D$. Under these three assumptions we claim that

$$G_r'(N,D,l) = rN\cdot I_{r-1}(u) + O\left(\frac{r^4N}{\log N}\cdot(\log u)^{r+3}\right), \tag{10.4}$$

$$H_r'(N,D,l) = \frac{rN}{\log N}\cdot I_{r-1}(u) + O\left(\frac{r^4N}{(\log N)^2}\cdot(\log u)^{r+3}\right). \tag{10.5}$$

Recall that we have already shown that

$$F_r'(N,D,l) = I_r(u) + O\left(\frac{r^2}{\log N}\cdot(\log u)^{r-1}\right). \tag{10.6}$$

Let us start with (10.4). It will be convenient to abbreviate $p_1 \cdot \ldots \cdot p_{r-1}$ as P. Then

$$\begin{aligned} G'_r(N,D,l) &= \varphi(l)^r \sum_{\substack{p_1\cdot\ldots\cdot p_r<N\\ p_1,\ldots,p_r\in\mathcal{P}_l\\ p_1,\ldots,p_r>D}} \log(p_1\cdot\ldots\cdot p_r) = r\varphi(l)^r \sum_{\substack{P<N/D\\ p_1,\ldots,p_{r-1}\in\mathcal{P}_l\\ p_1,\ldots,p_{r-1}>D}} \sum_{\substack{p\in\mathcal{P}_l\\ p>D}}^{N/P} \log p \\ &= r\varphi(l)^r \sum_{\substack{P<N/D\\ p_1,\ldots,p_{r-1}\in\mathcal{P}_l\\ p_1,\ldots,p_{r-1}>D}} \left(\frac{N}{\varphi(l)P}\cdot(1+O(e^{-c\sqrt{\log N/P}})) - \sum_{\substack{p\in\mathcal{P}_l\\ p<D}} \log p\right) \\ &= rN\cdot F'_{r-1}(N/D,D,l) - r\varphi(l)\cdot H'_{r-1}(N/D,D,l)\cdot \sum_{\substack{p\in\mathcal{P}_l\\ p<D}} \log p \\ &\quad + rN\varphi(l)^{r-1} \sum_{\substack{P<N/D\\ p_1,\ldots,p_{r-1}\in\mathcal{P}_l\\ p_1,\ldots,p_{r-1}>D}} \frac{1}{P}\cdot O(e^{-c\sqrt{\log N/P}}). \end{aligned} \tag{10.7}$$

To simplify (10.7) we first attack

$$rN\varphi(l)^{r-1} \sum_{\substack{P<N/D\\ p_1,\ldots,p_{r-1}\in\mathcal{P}_l\\ p_1,\ldots,p_{r-1}>D}} \frac{1}{P}\cdot O(e^{-c\sqrt{\log N/P}}).$$

Define $N_0 := Ne^{-(c^{-1}\log\log N)^2}$. Then splitting this sum into two ranges depending on $P < N_0$ or $P > N_0$ yields

$$\begin{aligned} rN\varphi(l)^{r-1} &\sum_{\substack{P<N_0\\ p_1,\ldots,p_{r-1}\in\mathcal{P}_l\\ p_1,\ldots,p_{r-1}>D}} \frac{1}{P}\cdot O(e^{-c\sqrt{\log N/P}}) \\ &\quad + rN\varphi(l)^{r-1} \sum_{\substack{N_0<P<N/D\\ p_1,\ldots,p_{r-1}\in\mathcal{P}_l\\ p_1,\ldots,p_{r-1}>D}} \frac{1}{P}\cdot O(e^{-c\sqrt{\log N/P}}). \end{aligned}$$

The former sum is bounded by

$$O(rNe^{-c\sqrt{\log N/N_0}} F'_{r-1}(N_0,D,l)), \tag{10.8}$$

while the latter is bounded by

$$O\big(rNe^{-c\sqrt{\log D}}(F'_{r-1}(N,D,l) - F'_{r-1}(N_0,D,l))\big). \tag{10.9}$$

A careful computation using (10.6) shows that both (10.8) and (10.9) are within the error term of (10.4) for our choice of N_0. To further simplify (10.7) we look at

$F'_{r-1}(N/D, D, l)$. Because of our assumptions on N and D coupled with (10.6) we can put $F'_{r-1}(N, D, l) - F'_{r-1}(N/D, D, l)$ in the error term of (10.4). So far we have shown

$$\begin{aligned} G'_r(N, D, l) &= rN \cdot F'_{r-1}(N, D, l) \\ &\quad - r\varphi(l) \cdot H'_{r-1}(N/D, D, l) \cdot \sum_{\substack{p \in \mathcal{P}_l \\ p < D}} \log p \\ &\quad + O\left(\frac{r^4 N}{\log N} \cdot (\log u)^{r+3}\right). \end{aligned} \tag{10.10}$$

To finish the proof of (10.4), we have to deal with the term $-r\varphi(l) \cdot H'_{r-1}(N/D, D, l)$. If we carefully go through the proof of (10.10), we see that

$$G'_r(N, D, l) = O(rN \cdot F'_{r-1}(N, D, l)) \tag{10.11}$$

without any restrictions on D and N. Then partial summation combined with (10.11) shows

$$\begin{aligned} H'_r(N, D, l) &= \frac{G'_r(N, D, l)}{\log N} + \int_D^N \frac{G'_r(x, D, l)}{x(\log x)^2}\, dx \\ &= \frac{G'_r(N, D, l)}{\log N} + O\left(\frac{rN}{(\log N)^2} \cdot (\log u)^{r-1}\right). \end{aligned}$$

Plugging this in (10.10) and using (10.11) again establishes (10.4) and hence (10.5). We have the trivial upper bound

$$|S'_r(N, D, l)| \leq \frac{H'_r(N, D, l)}{\varphi(l)^r \cdot r!}.$$

To give a lower bound for $|S'_r(N, D, l)|$ in terms of $H'_r(N, D, l)$, we observe that

$$\begin{aligned} \frac{H'_r(N, D, l) - \varphi(l)^r \cdot r! \cdot |S'_r(N, D, l)|}{H'_r(N, D, l)} &= O\left(\sum_{\substack{p \in \mathcal{P}_l \\ p > D}} \frac{1}{p^2}\right) \\ &= O\left(\frac{1}{D}\right). \end{aligned}$$

This implies that

$$|S'_r(N, D, l)| \geq \frac{H'_r(N, D, l)}{\varphi(l)^r \cdot r!} \cdot O\left(\frac{1}{D}\right).$$

Hence upon taking $D > D_0(l)$ for some constant $D_0(l)$ depending only on l, we get

$$|S'_r(N, D, l)| \geq \frac{H'_r(N, \max(D, D_0(l)), l)}{2 \cdot \varphi(l)^r \cdot r!}.$$

Altogether we have proven the following claim.

Lemma 10.5. *Let $N \geq e^{e^{10 \cdot l}}$ and $D > 1.5$ be real numbers satisfying $\log\log N \geq 1.1 \log\log D$ and $3 \cdot B(D,l) \leq \log N$. Assume that*

$$1 \leq r \leq \epsilon \log\log N \tag{10.12}$$

for some $\epsilon > 0$. Then there are positive constants $A_2(\epsilon, l)$ and $A_3(\epsilon, l)$ depending only on ϵ and l such that

$$\frac{A_2(\epsilon,l)}{\varphi(l)^r} \cdot \frac{N}{\log N} \cdot \frac{(\log u)^{r-1}}{(r-1)!} < |S_r'(N,D,l)| < \frac{A_3(\epsilon,l)}{\varphi(l)^r} \cdot \frac{N}{\log N} \cdot \frac{(\log u)^{r-1}}{(r-1)!}.$$

We stress that condition (10.12) is essential for the correctness of Lemma 10.5. Indeed, in general it is not clear that the main term in (10.5) dominates the error term. However, if r satisfies (10.12), we can use Lemma 10.3 to show that the main term does dominate the error term. Take N and D as in the previous lemma and suppose that r satisfies (10.12). If $k \leq r$ is an integer we define $S_{r,k}'(N,D,l)$ to be the subset of $S_r'(N,D,l)$ with exactly k prime divisors smaller than

$$N_1 := \exp(\sqrt{\log N \cdot B(D,l)}).$$

Lemma 10.6. *Let $N \geq e^{e^{10 \cdot l}}$ and $D > 1.5$ be real numbers satisfying $\log\log N \geq 2 \log\log D$ and $3 \cdot B(D,l) \leq \log N$. If r satisfies (10.1), there exists an absolute constant $c > 0$ such that*

$$\Big| \bigcup_{|0.5r-k| \leq r^{2/3}} S_{r,k}'(N,D,l) \Big| / |S_r'(N,D,l)| = O(\exp(-c r^{1/3})).$$

Now suppose that k satisfies $|0.5r - k| \leq r^{2/3}$. Let

$$\mathrm{Tuples}_k(N_1,D,l) := \{(p_1,\ldots,p_k) \in \mathcal{P}_l^k : D < p_1 < \cdots < p_k < N_1\}$$

and let T_1 and T_2 be subsets of $\mathrm{Tuples}_k(N_1,D,l)$. Define $S_{r,k}'(N,D,l,T)$ to be the subset of $S_{r,k}'(N,D,l)$ for which the k smallest prime factors $p_1,\ldots,p_k$ lie in T. Then

$$\frac{|S_{r,k}'(N,D,l,T_1)|}{|S_{r,k}'(N,D,l,T_2)|} = O\left(\frac{\mathrm{Vol}(\mathrm{Grid}(T_1))}{\mathrm{Vol}(\mathrm{Grid}(T_2))}\right).$$

Proof. We start with the easy formula

$$|S_{r,k}'(N,D,l)| = \sum_{\substack{p_1,\ldots,p_k \in \mathcal{P}_l \\ D<p_1<\cdots<p_k<N_1}} \left| S_{r-k}'\left(\frac{N}{p_1 \cdot \ldots \cdot p_k}, N_1, l\right)\right|. \tag{10.13}$$

If N is sufficiently large, we have $N_1^k < \sqrt{N}$. For convenience we will write $P := p_1 \cdot \ldots \cdot p_k$. Then for all choices of P,

$$\left(\frac{\log\log N - \log B(D,l)}{\log\log(N/P) - \log B(D,l)}\right)^{r-1} = O(1).$$

We also have, for all choices of P,

$$\frac{\log N}{\log(N/P)} = O(1).$$

Suppose that $r - k > 0$. An appeal to Lemma 10.5 yields constants $A_4(l)$ and $A_5(l)$ depending only on l with the property

$$\frac{A_4(l)}{P} \cdot |S'_{r-k}(N, N_1, l)| < |S'_{r-k}(N/P, N_1, l)| < \frac{A_5(l)}{P} \cdot |S'_{r-k}(N, N_1, l)|. \quad (10.14)$$

Plugging (10.14) into (10.13) shows

$$|S'_{r,k}(N, D, l)| = O\left(\frac{N}{\log N} \cdot \frac{2^{-r+1}(\log\log N - \log B(D, l))^{r-1}}{\varphi(l)^r (r-k-1)!k!}\right) \quad (10.15)$$

for $r - k > 0$. In case $r = k$ we have the trivial bound $|S'_{r,k}(N, D, l)| \leq N_1^r$, so we can remove this case from the union. Now we are in a position to apply Hoeffding's inequality to (10.15), which proves the first part of the lemma. The second part quickly follows from (10.14). ∎

Lemma 10.6 directly relates $S'_{r,k}(N, D, l)$ and the heuristic model introduced at the beginning of the section. Since we have already dealt with the heuristic model in Proposition 10.2, the second part of Theorem 10.4 is now straightforward. Indeed, we first restrict to those k for which $|0.5r - k| \leq r^{2/3}$. Take $T_2 := \mathrm{Tuples}_k(N_1, 1.5, l)$ and take T_1 to be the subset of T_2 that is not C_0-regular. There exists an absolute constant $c > 0$ such that

$$\mathrm{Vol}(\mathrm{Grid}(T_2)) \geq \frac{c}{k!} \cdot (\log\log N_1)^k.$$

Furthermore, there is some constant $\kappa(l) > 0$, depending only on l, such that no element of $\mathrm{Grid}(T_1)$ is $C_0 - \kappa(l)$-regular. Proposition 10.2 shows that

$$\mathrm{Vol}(\mathrm{Grid}(T_1)) = O\left(\frac{\exp(-c' \cdot C_0)}{k!} \cdot (\log\log N_1)^k\right)$$

for some absolute constant $c' > 0$. This proves the second part of Theorem 10.4.

It remains to prove the first part of Theorem 10.4. If $r \leq 2$, we can prove Theorem 10.4 directly, so suppose that $r \geq 3$. The number of $n \in S'_r(N, l)$ that are not comfortably spaced above D_1 is bounded by

$$\sum_{p > D_1}^{N} \sum_{q > p}^{l^{200}p} |S'_{r-2}(N/(pq), l)|. \quad (10.16)$$

We split (10.16) into two ranges depending on $p < N^{1/4}$ and $p \geq N^{1/4}$. First suppose that $p < N^{1/4}$. Then we use Lemma 10.5 to bound (10.16) by

$$O\left(|S'_{r-2}(N, l)| \cdot \sum_{p > D_1}^{N} \sum_{q > p}^{l^{200}p} \frac{1}{pq}\right) = O\left(\frac{|S'_r(N, l)|}{\log D_1}\right).$$

Now suppose that $p \geq N^{1/4}$. From the crude bound $|S'_{r-2}(N/(pq), l)| \leq N/(pq)$ we deduce

$$\sum_{p>N^{1/4}}^{N} \sum_{q>p}^{l^{200}p} |S'_{r-2}(N/(pq), l)| = O\left(\frac{N}{\log N}\right) = O\left(\frac{|S'_r(N, l)|}{\log\log N}\right),$$

since $r > 1$. This implies the theorem. ∎

Finally, we use Theorem 10.4 to prove that most integers in $S_r(N, l)$ are comfortably spaced above D_1 and C_0-regular.

Theorem 10.7. *Let N and r be as in Definition 10.1.*

- *Suppose that $D_1 > 100$. Then*

$$\frac{|\{n \in S_r(N, l) : n \text{ is not comfortably spaced above } D_1\}|}{|S_r(N, l)|} = O\left(\frac{1}{\log D_1} + \frac{1}{\log\log N}\right).$$

- *There is an absolute constant $c > 0$ such that for all $C_0 > 0$,*

$$\frac{|\{n \in S_r(N, l) : n \text{ is not } C_0\text{-regular}\}|}{|S_r(N, l)|} = O\big(\exp(-c \cdot C_0) + \exp(-c r^{1/3})\big).$$

Proof. This is an easy consequence of Theorem 10.4. ∎

11. Rédei matrices

To prove our main results, we will arrange cyclic degree l extensions in product spaces called boxes. These boxes provide the combinatorial structure that allows us to apply the results from the previous sections. Let us start by giving a precise definition of a box.

Definition 11.1. Let $\max(100, l) < D_1$ be a real number and let $1 \leq k \leq r$ be integers. Choose a sequence of primes

$$p_1 < \cdots < p_k < D_1$$

all equal to 0 or 1 modulo l. Also choose real numbers

$$D_1 < t_{k+1} < \cdots < t_r.$$

For $i \geq k + 1$ define

$$t'_i := \left(1 + \frac{1}{e^{i-k} \cdot \log D_1}\right) \cdot t_i.$$

We assume that $t_i \geq l^{100} t'_{i-1}$ for all $k + 1 \leq i \leq r$, where by definition $t'_k := p_k$. Put

$$X := X_1 \times \cdots \times X_r,$$

where $X_i := \{p_i\}$ for $i \leq k$ and X_i is the set of primes equal to 1 modulo l in the interval (t_i, t'_i) for $i \geq k + 1$. Finally, choose a function $f : X_1 \amalg \cdots \amalg X_r \to [l - 1]$. We will say that the set X as constructed above is a *box*.

In Smith [22], the transition from squarefree integers to boxes is done by appealing to his Proposition 6.9. Instead of degree 2 extensions of $\mathbb{Q}$, we need to keep track of cyclic degree l extensions and this is substantially more difficult. One key difference is that there are precisely $l-1$ characters $\chi : G_{\mathbb{Q}} \to \langle \zeta_l \rangle$ with the same field of definition, while our algebraic results use only one of the $l-1$ characters.

It will therefore be important to keep track of the characters we have chosen, and this is the reason for introducing f. For $x \in X$ we define the character

$$\chi_{x,f} := \prod_{1 \leq i \leq r} \chi_{\pi_i(x)}^{f(\pi_i(x))},$$

where χ_p is the character of conductor dividing p^∞ that we fixed in Subsection 2.2. Given a box X we also define $\mathrm{Field}(X)$ to be the set of cyclic degree l number fields K over $\mathbb{Q}$ satisfying

- for all $1 \leq i \leq r$ there is exactly one prime $p \in X_i$ such that K is ramified at p;
- K is unramified at all primes p that are not in any X_i.

Given X and f there is a natural map $i_f : X \to \mathrm{Field}(X)$ that sends x to the field fixed by the kernel of $\chi_{x,f}$. There is also a natural map $j : \mathrm{Field}(X) \to S_r(\infty, l)$ given by sending K to the radical of D_K. The composition $j \circ i_f : X \to S_r(\infty, l)$ is the natural inclusion from X to $S_r(\infty, l)$ that we have seen in Section 10.

Finally, define $\mathrm{Field}(N, r, l)$ to be the set of cyclic degree l number fields K over $\mathbb{Q}$ with the following properties:

- the radical of D_K is at most N;
- D_K has exactly r prime divisors.

Our first proposition is a variant of Smith's Proposition 6.9 that can deal with the more complicated structure of our new boxes.

Proposition 11.2. *Let l be a prime and let $N \geq D_1 > \max(100, l)$ with $\log\log N \geq 2\log\log D_1$. Suppose that r satisfies (10.1) and let W be a subset of $S_r(N, l)$ that is comfortably spaced above D_1. Let $\epsilon > 0$ be such that*

$$|W| > (1-\epsilon) \cdot |S_r(N, l)|.$$

Let V be a subset of $\mathrm{Field}(N, r, l)$. Assume that there exists $\delta > 0$ such that for all boxes X with $j(\mathrm{Field}(X)) \cap W \neq \emptyset$ and $\mathrm{Field}(X) \subseteq \mathrm{Field}(N, r, l)$ we have

$$(\delta - \epsilon) \cdot |\mathrm{Field}(X)| < |V \cap \mathrm{Field}(X)| < (\delta + \epsilon) \cdot |\mathrm{Field}(X)|.$$

Then

$$|V| = \delta \cdot |\mathrm{Field}(N, r, l)| + O\left(\left(\epsilon + \frac{1}{\log D_1}\right) \cdot |\mathrm{Field}(N, r, l)|\right).$$

Proof. Let $0 \leq k \leq r$ be an integer. Let T_k be the set of tuples $(p_1, \dots, p_k)$ such that

$$p_1 < \cdots < p_k < D_1$$

and all the p_i are primes equal to 0 or 1 modulo l. Define T_k' to be the set of tuples $(t_{k+1}, \ldots, t_r)$ satisfying

$$D_1 < t_{k+1} < \cdots < t_r,$$

where the t_i are real numbers with $t_{i+1} \geq 2t_i$. Given $\mathbf{t} \in T_k$ and $\mathbf{t}' \in T_k'$ there is a natural way to construct a box $X(\mathbf{t}, \mathbf{t}')$. Take

$$T_k'' := \{(\mathbf{t}, \mathbf{t}') \in T_k \times T_k' : j(\mathrm{Field}(X(\mathbf{t}, \mathbf{t}'))) \cap W \neq \emptyset \\ \text{and } \mathrm{Field}(X(\mathbf{t}, \mathbf{t}')) \subseteq \mathrm{Field}(N, r, l)\}.$$

Now consider

$$\sum_{\mathbf{t} \in T_k} \int_{\substack{\mathbf{t}' \in T_k' \\ (\mathbf{t}, \mathbf{t}') \in T_k''}} \frac{|V \cap \mathrm{Field}(X(\mathbf{t}, \mathbf{t}'))|}{t_{k+1} \cdot \ldots \cdot t_r} \, dt_{k+1} \cdot \ldots \cdot dt_r. \tag{11.1}$$

Let $K \in V$ with $j(K) := (q_1, \ldots, q_r) \in W$. Note that $K \in \mathrm{Field}(X(\mathbf{t}, \mathbf{t}'))$ if and only if

$$(q_1, \ldots, q_k) = \mathbf{t}$$

and for all $k + 1 \leq i \leq r$,

$$t_i < q_i < \left(1 + \frac{1}{e^{i-k} \cdot \log D_1}\right) \cdot t_i.$$

If K also satisfies

$$q_1 \cdot \ldots \cdot q_r < N \cdot \prod_{i=k+1}^{r} \left(1 + \frac{1}{e^{i-k} \cdot \log D_1}\right)^{-1}, \tag{11.2}$$

we see that the contribution of K to (11.1) is equal to

$$\prod_{i=k+1}^{r} \log\left(1 + \frac{1}{e^{i-k} \cdot \log D_1}\right). \tag{11.3}$$

If $j(K)$ is outside W or $j(K)$ does not satisfy (11.2), we see that the contribution of K to (11.1) is bounded by (11.3). Put

$$H_r(N, l) := \sum_{\substack{p_1 \cdot \ldots \cdot p_r < N \\ p_i \equiv 0,1 \bmod l}} 1.$$

Due to (10.5) and Lemma 10.5 we have, for all $c \in (0, 0.5)$,

$$\frac{1}{r!} \cdot \big(H_r(N, l) - H_r((1 - c) \cdot N, l)\big) = O\left(c + \frac{(\log \log N)^4}{\log N}\right) \cdot |S_r(N, l)|.$$

Hence the number of elements in $S_r(N,l)$ failing (11.2) is bounded by $O(|S_r(N,l)|/\log D_1)$. In particular, we can bound the number of $K \in \mathrm{Field}(N,r,l)$ with $j(K)$ failing (11.2) by $O(|\mathrm{Field}(N,r,l)|/\log D_1)$. Finally, consider

$$\sum_{k\geq 0}^{r}\prod_{i=k+1}^{r}\left(1+\frac{1}{e^{i-k}\cdot\log D_1}\right)^{-1} \cdot \sum_{\mathbf{t}\in T_k}\int_{(\mathbf{t},\mathbf{t}')\in T_k''}^{\mathbf{t}'\in T_k'} \frac{|V\cap \mathrm{Field}(X(\mathbf{t},\mathbf{t}'))|}{t_{k+1}\cdot\ldots\cdot t_r}\,dt_{k+1}\cdot\ldots\cdot dt_r. \tag{11.4}$$

Since the contribution of any $K \in V$ to (11.1) is always bounded by (11.3), we have an upper bound for (11.4) given by $|V|$. On the other hand, if $K \in V \cap j^{-1}(W)$ and satisfies (11.2), the contribution of K to (11.1) is equal to (11.3). This yields a lower bound for (11.4), namely

$$|V\cap j^{-1}(W)| - O\left(\frac{|\mathrm{Field}(N,r,l)|}{\log D_1}\right).$$

Using our assumption

$$(\delta-\epsilon)\cdot|\mathrm{Field}(X)| < |V\cap\mathrm{Field}(X)| < (\delta+\epsilon)\cdot|\mathrm{Field}(X)|$$

for all boxes X with $j(\mathrm{Field}(X))\cap W \neq \emptyset$ and $\mathrm{Field}(X)\subseteq \mathrm{Field}(N,r,l)$, we can again obtain upper and lower bounds for (11.4). Indeed, we have an upper bound for (11.4) given by $(\delta+\epsilon)\cdot|\mathrm{Field}(N,r,l)|$ and a lower bound for (11.4) given by

$$\delta\cdot|\mathrm{Field}(N,r,l)| - O\left(\left(\epsilon+\frac{1}{\log D_1}\right)\cdot|\mathrm{Field}(N,r,l)|\right).$$

Combining the various lower and upper bounds finishes the proof of the proposition. ∎

The usefulness of Proposition 11.2 lies in the fact that it allows us to deduce equidistribution of $\mathrm{Field}(N,r,l)$ from equidistribution of product spaces $\mathrm{Field}(X)$. However, our algebraic results work for a product space of the shape $i_f(X)$ and not for the full set $\mathrm{Field}(X)$. To work around this issue, the identity

$$|\{f : K\in i_f(X)\}| = |\{f : K'\in i_f(X)\}|$$

for all $K, K' \in \mathrm{Field}(X)$ will be pivotal in the next sections.

In the coming sections it will also be important to have very fine control over r. Until this point we have only assumed that r satisfies (10.1), but we will now introduce the much stronger requirement

$$|r-\log\log N| \leq (\log\log N)^{2/3}. \tag{11.5}$$

Our next theorem shows that (11.5) is usually satisfied.

Theorem 11.3. *Recall that* $\mathrm{Field}(N,l)$ *is the set of cyclic degree* l *number fields* K *over* $\mathbb{Q}$ *with* $\mathrm{rad}(D_K) \leq N$. *Then*

$$\left||\mathrm{Field}(N,l)| - \bigcup_{r \text{ satisfies (11.5)}} |\mathrm{Field}(N,r,l)|\right| = O\left(\frac{|\mathrm{Field}(N,l)|}{(\log\log N)^c}\right) \tag{11.6}$$

for some absolute constant $c > 0$.

Proof. Note that the map $j : \mathrm{Field}(N,r,l) \to S_r(N,l)$ is $(l-1)^{r-1}$-to-1. This observation combined with Lemma 10.5 shows that

$$|\{K \in \mathrm{Field}(N,l) : \omega(D_K) \leq \log\log N - (\log\log N)^{2/3}\}| = O\left(\frac{N}{(\log\log N)^c}\right)$$

for some absolute constant $c > 0$. From [20, Lemma 2.2] we infer that $N = O(|\mathrm{Field}(N,l)|)$. We conclude that our error term fits in the error term of the theorem. To deal with the case

$$\omega(D_K) \geq \log\log N + (\log\log N)^{2/3},$$

we take a different approach. Indeed, Lemma 10.5 does not directly apply, since condition (10.12) may not be satisfied. In the classical paper [10] it is proven that there are absolute constants $C > 0$ and $K > 0$ such that

$$\pi_k(x) < \frac{Kx}{\log x}\frac{(\log\log x + C)^k}{k!},$$

where $\pi_k(x)$ is the number of squarefree integers $\leq x$ and with exactly k prime divisors. A straightforward generalization of their argument proves

$$|S_r(N,l)| < \frac{KN}{(l-1)^r \log N}\frac{(\log\log N + C)^r}{r!}$$

for some absolute constants $C > 0$ and $K > 0$. Then a small computation finishes the proof of (11.6). ∎

For $\alpha \in \mathbb{Z}[\zeta_l]$ and $\mathfrak{n}$ an ideal of $\mathbb{Z}[\zeta_l]$ with $(\mathfrak{n}, 1-\zeta_l) = (1)$ we write

$$\left(\frac{\alpha}{\mathfrak{n}}\right)_{\mathbb{Z}[\zeta_l],l}$$

for the l-th power residue symbol in $\mathbb{Z}[\zeta_l]$. We assume that the reader is familiar with the basic properties of the power residue symbol. Suppose that $p \neq l$. Then, given χ_p, there is a unique prime ideal $\mathfrak{p}$ of $\mathbb{Z}[\zeta_l]$ satisfying

$$\chi_p(\mathrm{Frob}(q)) = \left(\frac{q}{\mathfrak{p}}\right)_{\mathbb{Z}[\zeta_l],l} = \frac{\mathrm{Frob}(\mathfrak{p})(\sqrt[l]{q})}{\sqrt[l]{q}}$$

for all primes q. We will now define a generalized Rédei matrix.

Definition 11.4. Take P to be a set of prime numbers 1 modulo l. Choose a function $f : P \sqcup X_1 \sqcup \cdots \sqcup X_r \to [l-1]$. Let X be any box with $X_1, \dots, X_r$ disjoint from the set P. Put

$$M := \{(i,j) : 1 \leq i, j \leq r,\ i \neq j\}.$$

We also define

$$M_{P,1} := [r] \times P, \quad M_{P,2} := P \times [r].$$

For an assignment $a : M \sqcup M_{P,1} \sqcup M_{P,2} \to \langle \zeta_l \rangle$, we define $X(a)$ to be the set of tuples $(x_1, \dots, x_r) \in X$ satisfying

$$\begin{aligned}
\chi_{x_j}^{f(x_j)}(\mathrm{Frob}(x_i)) &= a(i,j) && \text{for all } (i,j) \in M,\\
\chi_{p}^{f(p)}(\mathrm{Frob}(x_i)) &= a(i,p) && \text{for all } (i,p) \in M_{P,1},\\
\chi_{x_j}^{f(x_j)}(\mathrm{Frob}(p)) &= a(p,j) && \text{for all } (p,j) \in M_{P,2}.
\end{aligned}$$

Note that $X(a)$ depends on the choice of the f. To make this more explicit we will sometimes write $X(a, f)$. We can think of the assignment a as an analogue of the classical Rédei matrix. In our final section we need to treat $X(a, f)$ as fixed. For this reason we would like to prove that $X(a, f)$ is of the expected size. Unfortunately, this turns out to be rather hard. The reason is that we made one choice of χ_p, but we could just as well have chosen χ_p^s for some integer s with $(s, l) = 1$. This creates substantial difficulties, for example, when dealing with sums of the type

$$\sum_{X < p < Y} \chi_p(\mathrm{Frob}(q))$$

for fixed q. Indeed, by changing many of the χ_p to χ_p^s it is easy to make such a sum unbalanced. Instead we will prove something weaker, but still sufficient for our application. There is one final obstacle to deal with: for $i \leq k$ the set X_i consists of only one element. Hence we must restrict our attention to a special set of a.

Definition 11.5. Let X be a box and let $a : M \sqcup M_{P,1} \sqcup M_{P,2} \to \langle \zeta_l \rangle$. For $1 \leq i \leq k$, let x_i be the unique element of X_i. We say that *a agrees with X at stage k* if

$$\begin{aligned}
\chi_{x_j}^{f(x_j)}(\mathrm{Frob}(x_i)) &= a(i,j) && \text{for all } (i,j) \in M \text{ with } i, j \leq k;\\
\chi_{p}^{f(p)}(\mathrm{Frob}(x_i)) &= a(i,p) && \text{for all } (i,p) \in M_{P,1} \text{ with } i \leq k;\\
\chi_{x_j}^{f(x_j)}(\mathrm{Frob}(p)) &= a(p,j) && \text{for all } (p,j) \in M_{P,2} \text{ with } j \leq k.
\end{aligned}$$

We stress that this notion implicitly depends on the choice of f. Whenever we need to make the choice of f explicit, we will say instead that *X agrees with a and f*.

Clearly, if a does not agree with X, we have $X(a) = \emptyset$. It will be convenient to define

$$g(l, P, k) := l^{|M|+|M_{P,1}|+|M_{P,2}|-k(k-1)-2k|P|},$$

so that $g(l, P, k)$ is the number of a's that agree with a given X. Let K be a Galois extension of $\mathbb{Q}$. Then we define

$$X(a, f, K) := \{x \in X(a, f) : \pi_i(x) \text{ splits completely in } K \text{ for all } k+1 \le i \le r\}.$$

Also define, for a set P of primes,

$$L(P) := \prod_{p \in P} \overline{\mathbb{Q}}^{\ker(\chi_p)} \mathbb{Q}(\zeta_l, \sqrt[l]{p}).$$

We can now show that $X(a, f, K)$ is of the expected size for most choices of f.

Theorem 11.6. *Assume GRH and let l be an odd prime. There are constant $A(l), A'(l) > 0$ such that the following holds. Let X be a box with $D_1 > A(l)$. Let P be a set of primes equal to 1 modulo l disjoint from all the X_i. Write x_i for the unique element of X_i for $1 \le i \le k$ and fix a function $g : P \cup \{x_1, \dots, x_k\} \to [l-1]$. Suppose that the assignment $a : M \sqcup M_{P,1} \sqcup M_{P,2} \to \langle \zeta_l \rangle$ agrees with X at stage k. Let K be a Galois extension of $\mathbb{Q}$ of degree n_K that is disjoint from $L(P \cup x)$ for all $x \in X$. Define B to be the maximum of the primes in P and X_{k+1}. Assume that*

$$l^{2|P|+2k} \cdot (n_K \cdot \log B + \log |\Delta_{K/\mathbb{Q}}|) \le t_{k+1}^{1/8}. \tag{11.7}$$

Then the proportion of f's in $\mathrm{Map}(X_{k+1} \sqcup \cdots \sqcup X_r, [l-1])$ with

$$\left| |X(a, f, K)| - \frac{|X|}{n_K^{r-k} g(l, P, k)} \right| > A'(l) \cdot \frac{|X|}{n_K^{r-k} g(l, P, k)} \cdot x_k^{-1/4} \tag{11.8}$$

is at most $O\big(e^{-2|X_{k+1}|^{1/8}}\big)$.

Proof. Here and later on we implicitly extend our function $f \in \mathrm{Map}(X_{k+1} \sqcup \cdots \sqcup X_r, [l-1])$ to $f \in \mathrm{Map}(P \sqcup X_1 \sqcup \cdots \sqcup X_r, [l-1])$ by using the function g. Define, for $k+1 \le i \le r$,

$$s_i := \frac{|X_i|}{n_K \cdot l^{2|P|+2i-2}}, \quad s_{r+1} = 0.$$

We claim that there exist constants $A_1(l), A_2(l) > 0$ depending only on l such that the proportion of f's with

$$\left| |X(a, f, K)| - \frac{|X|}{n_K^{r-k} g(l, P, k)} \right| > A_1(l) \cdot \frac{|X|}{n_K^{r-k} g(l, P, k)} \cdot \Big(\sum_{k+1 \le i \le r} t_i^{-1/4} \Big) \tag{11.9}$$

is bounded by

$$A_2(l) \cdot e^{-2 s_{k+1}^{1/4}}.$$

Once we establish the claim, we immediately deduce (11.8) and the theorem. We proceed by downward induction on k with base case $k = r$. In the base case (11.9) is trivial, so henceforth we shall assume $k < r$. Define the number field

$$L := L(P \cup \{x_1, \dots, x_k\}).$$

There is an isomorphism between $\mathrm{Gal}(L/\mathbb{Q}(\zeta_l))$ and $(\langle\zeta_l\rangle \times \langle\zeta_l\rangle)^{|P|+k}$ given by

$$\sigma \mapsto (\chi_p(\sigma), \sigma(\sqrt[l]{p})/\sqrt[l]{p})$$

on each coordinate, where p runs through P and $x_1, \dots, x_k$. Then the Galois group of L over $\mathbb{Q}$ is naturally isomorphic to

$$\mathrm{Gal}(L/\mathbb{Q}) \simeq (\langle\zeta_l\rangle \times \langle\zeta_l\rangle)^{|P|+k} \rtimes \mathbb{F}_l^*.$$

If $s \in \mathbb{F}_l^*$, the $\mathbb{F}_l^*$-action is given by

$$(x, y) \mapsto (x, y^s)$$

on every copy of $\langle\zeta_l\rangle \times \langle\zeta_l\rangle$. We denote this automorphism of $(\langle\zeta_l\rangle \times \langle\zeta_l\rangle)^{|P|+k}$ by T_s. Using the classical formula

$$\Delta_{L/\mathbb{Q}} = N_{\mathbb{Q}(\zeta_l)/\mathbb{Q}}(\Delta_{L/\mathbb{Q}(\zeta_l)}) \cdot \Delta_{\mathbb{Q}(\zeta_l)/\mathbb{Q}}^{[L:\mathbb{Q}(\zeta_l)]},$$

we obtain the bound

$$\log|\Delta_{L/\mathbb{Q}}| \ll (|P| + k) \cdot l^{2|P|+2k} \cdot \log B \ll l^{2|P|+2k} \cdot \log B.$$

This implies

$$\begin{aligned}\log|\Delta_{KL/\mathbb{Q}}| &\leq n_K \log|\Delta_{L/\mathbb{Q}}| + (l-1) \cdot l^{2|P|+2k} \log|\Delta_{K/\mathbb{Q}}| \\ &\ll l^{2|P|+2k} \cdot (n_K \cdot \log B + \log|\Delta_{K/\mathbb{Q}}|) \leq t_{k+1}^{1/8},\end{aligned}$$

where the last inequality is just (11.7). We know that

$$\mathrm{Gal}(KL/\mathbb{Q}) \simeq \mathrm{Gal}(K/\mathbb{Q}) \times \mathrm{Gal}(L/\mathbb{Q}),$$

and we let p_1 and p_2 be the natural projection maps. Let C be a conjugacy class of $\mathrm{Gal}(KL/\mathbb{Q})$ with $p_1(C) = \mathrm{id}$ and $p_2(C) \subseteq \mathrm{Gal}(L/\mathbb{Q}(\zeta_l))$. Then $C = \{(\mathrm{id}, T_s\sigma) : s \in \mathbb{F}_l^*\}$ for some $\sigma \in \mathrm{Gal}(L/\mathbb{Q}(\zeta_l))$. The Chebotarev density theorem yields, conditional on GRH [16],

$$\sum_{\substack{t_{k+1} < p < t'_{k+1} \\ p \equiv 1 \bmod l \\ \mathrm{Frob}(p) = C}} 1 = \frac{|C|}{n_K \cdot l^{2|P|+2k}}(|X_{k+1}| + O(t_{k+1}^{5/8})).$$

Recall that we have chosen prime ideals $\mathfrak{p}$ in $\mathbb{Z}[\zeta_l]$ above every $p \equiv 1 \bmod l$ in the range $t_{k+1} < p < t'_{k+1}$. Hence we obtain

$$\sum_{\substack{t_{k+1} < p < t'_{k+1} \\ p \equiv 1 \bmod l \\ \mathrm{Frob}(\mathfrak{p}) \in C}} 1 = \frac{|C|}{n_K \cdot l^{2|P|+2k}}(|X_{k+1}| + O(t_{k+1}^{5/8})). \tag{11.10}$$

Let $a : M \amalg M_{P,1} \amalg M_{P,2} \to \langle \zeta_l \rangle$ be an assignment that agrees with X at stage k. Our next step is to attach a conjugacy class C to a, for which we will use (11.10). There exists exactly one element $(a_{1i}, a_{2i})_{1 \le i \le |P|+k}$ satisfying

$$a(k+1,i) = a_{1i} \quad \text{for all } 1 \le i \le k, \qquad a(i,k+1) = a_{2i} \quad \text{for all } 1 \le i \le k,$$
$$a(k+1,p) = a_{1p} \quad \text{for all } p \in P, \qquad a(p,k+1) = a_{2p} \quad \text{for all } p \in P.$$

Define C to be the conjugacy class of $\mathrm{Gal}(KL/\mathbb{Q})$ with $p_1(C) = \mathrm{id}$ and

$$(a_{1i}, a_{2i})_{1 \le i \le |P|+k} \in p_2(C).$$

Take $\sigma := (a_{1i}, a_{2i})_{1 \le i \le |P|+k}$. Then we say that f is *balanced* if

$$\left| \sum_{\substack{t_{k+1} < p < t'_{k+1},\ p \equiv 1 \bmod l \\ p_1(\mathrm{Frob}(p)) = \mathrm{id},\ p_2(\mathrm{Frob}(\mathfrak{p})) = \sigma}} 1 - \sum_{\substack{t_{k+1} < p < t'_{k+1},\ p \equiv 1 \bmod l \\ \mathrm{Frob}(\mathfrak{p}) \in C}} \frac{1}{|C|} \right| \le \left(\frac{|C||X_{k+1}|}{n_K \cdot l^{2|P|+2k}} \right)^{5/8}, \tag{11.11}$$

and f is *unbalanced* otherwise. We deduce from Hoeffding's inequality that the proportion of unbalanced f's is bounded by

$$(A_2(l)/2) \cdot e^{-2s_{k+1}^{1/4}}$$

for a good choice of $A_2(l)$. Take $p \in X_{k+1}$ with $p_1(\mathrm{Frob}(p)) = \mathrm{id}$ and $p_2(\mathrm{Frob}(\mathfrak{p})) = \sigma$. Define the box

$$X_p := X_1 \times \cdots \times X_k \times \{p\} \times X_{k+2} \times \cdots \times X_r.$$

Then we have the decomposition

$$|X(a, f, K)| = \sum_{\substack{t_{k+1} < p < t'_{k+1} \\ p \equiv 1 \bmod l \\ p_1(\mathrm{Frob}(p)) = \mathrm{id} \\ p_2(\mathrm{Frob}(\mathfrak{p})) = \sigma}} |X_p(a, f, K)|. \tag{11.12}$$

To apply the induction hypothesis we must check that the estimate (11.7) is still valid. This is a straightforward computation and an appeal to the induction hypothesis gives

$$\left| |X_p(a, f, K)| - \frac{|X_p|}{n_K^{r-k-1} g(l, P, k+1)} \right| \le A_1(l) \cdot \frac{|X_p|}{n_K^{r-k-1} g(l, P, k+1)} \cdot \Big(\sum_{k+2 \le i \le r} t_i^{-1/4} \Big) \tag{11.13}$$

except for a proportion of f's bounded in magnitude by

$$A_2(l) \cdot e^{-2s_{k+2}^{1/4}}.$$

We say that f is *exceptional* if f is unbalanced or fails (11.13) for some choice of X_p. Then the total proportion of exceptional f's is bounded by

$$\frac{A_2(l)}{2} \cdot e^{-2s_{k+1}^{1/4}} + A_2(l) \cdot |X_{k+1}| \cdot e^{-2s_{k+2}^{1/4}} \le A_2(l) \cdot e^{-2s_{k+1}^{1/4}},$$

if D_1 is sufficiently large. We employ the triangle inequality to bound the LHS of (11.9) as follows:

$$\left||X(a,f,K)| - \sum_{\substack{t_{k+1}<p<t'_{k+1},\ p\equiv 1 \bmod l\\ p_1(\mathrm{Frob}(p))=\mathrm{id},\ p_2(\mathrm{Frob}(\mathfrak{p}))=\sigma}} \frac{|X_p|}{n_K^{r-k-1} g(l,P,k+1)}\right| + \left| \sum_{\substack{t_{k+1}<p<t'_{k+1},\ p\equiv 1 \bmod l\\ p_1(\mathrm{Frob}(p))=\mathrm{id},\ p_2(\mathrm{Frob}(\mathfrak{p}))=\sigma}} \frac{|X_p|}{n_K^{r-k-1} g(l,P,k+1)} - \frac{|X|}{n_K^{r-k} g(l,P,k)}\right|. \tag{11.14}$$

If f is not exceptional, we deduce from (11.12) and (11.13) that the first term of (11.14) is bounded by

$$\frac{A_1(l)|X|/|X_{k+1}|}{n_K^{r-k-1} g(l,P,k+1)} \cdot \Big(\sum_{k+2\le i\le r} t_i^{-1/4} \Big) \cdot \sum_{\substack{t_{k+1}<p<t'_{k+1}\\ p\equiv 1 \bmod l\\ p_1(\mathrm{Frob}(p))=\mathrm{id}\\ p_2(\mathrm{Frob}(\mathfrak{p}))=\sigma}} 1. \tag{11.15}$$

For sufficiently large D_1, we use (11.10) and (11.11) to bound the second term of (11.14) and to bound (11.15). Then a straightforward computation completes the proof of the theorem. ∎

12. Klys revisited

Suppose that $1 \le k \le r, s$ are integers. Define

$$P(r,s,l,j) := \frac{|\{A \in \mathrm{Mat}(r,s,\mathbb{F}_l) : \dim(\ker(A)) = j\}|}{|\mathrm{Mat}(r,s,\mathbb{F}_l)|}.$$

Fix $M \in \mathrm{Mat}(k,k,\mathbb{F}_l)$. Let $\mathrm{Mat}(r,s,\mathbb{F}_l,M)$ be the subset of $\mathrm{Mat}(r,s,\mathbb{F}_l)$ consisting of those matrices A satisfying $A(i,j) = M(i,j)$ for all $1 \le i,j \le k$. Then we set

$$Q(r,s,l,M,j) := \frac{|\{A \in \mathrm{Mat}(r,s,\mathbb{F}_l,M) : \dim(\ker(A)) = j\}|}{|\mathrm{Mat}(r,s,\mathbb{F}_l,M)|}.$$

We are interested in the difference $P(r,r-1,l,j) - Q(r,r-1,l,M,j)$ as r goes to infinity, independent of the choice of M.

Lemma 12.1. *Suppose that $r \ge 2k$. Then*

$$|P(r,r-1,l,j) - Q(r,r-1,l,M,j)| \le 2k \cdot l^{2k-r}.$$

Proof. For a matrix A, let $a_1, \dots, a_k$ denote its first k columns. We define

$$P(r,s,l,j,k) := \frac{|\{A \in \mathrm{Mat}(r,s,\mathbb{F}_l) : \dim(\ker(A)) = j \text{ and } \dim(a_1,\dots,a_k) = k\}|}{|\mathrm{Mat}(r,s,\mathbb{F}_l)|},$$

$$Q(r,s,l,M,j,k) := \frac{|\{A \in \mathrm{Mat}(r,s,\mathbb{F}_l,M) : \dim(\ker(A)) = j \text{ and } \dim(a_1,\dots,a_k) = k\}|}{|\mathrm{Mat}(r,s,\mathbb{F}_l,M)|}.$$

Then

$$P(r,r-1,l,j) - P(r,r-1,l,j,k) \le 1 - P(r,k,l,0)$$

and

$$Q(r,r-1,l,M,j) - Q(r,r-1,l,M,j,k) \le 1 - P(r-k,k,l,0)$$

due to our assumption $r \ge 2k$. We observe that $P(r,r-1,l,j,k) = Q(r,r-1,l,M,j,k)$. Combining this with the previous two inequalities gives

$$|P(r,r-1,l,j) - Q(r,r-1,l,M,j)| \le 2 - 2P(r-k,k,l,0). \tag{12.1}$$

Using the classical formula for $P(r-k,k,l,0)$, we obtain

$$P(r-k,k,l,0) = \frac{\prod_{j=0}^{k-1}(l^{r-k} - l^j)}{l^{k(r-k)}} = \prod_{j=0}^{k-1}(1 - l^{j+k-r})$$
$$\ge (1 - l^{2k-r})^k \ge 1 - k \cdot l^{2k-r},$$

where the last inequality follows from Bernoulli's inequality. Inserting this in (12.1) ends the proof of our theorem. ∎

With this lemma we have done all the preparatory work needed for understanding the $(1-\zeta_l)^2$-rank when K varies in $\mathrm{Field}(N,l)$. Recall that we have defined a matrix $\mathrm{Rédei}(K)$ in Section 5. It will be useful to observe that an assignment $a : M \to \langle \zeta_l \rangle$ uniquely determines a Rédei matrix and vice versa. Proposition 5.1 implies that the $(1-\zeta_l)^2$-rank of K is equal to $r - 1 - \mathrm{rank}_{\mathbb{F}_l} \mathrm{Rédei}(K)$ (see also [24, Theorem 1]). Our next theorem is similar to Theorem 4 in Klys [12], but has the benefit of providing an error term.

Theorem 12.2. *Assume GRH and let l be an odd prime. Let* $\mathrm{Field}(N,l,j)$ *be the subset of* $\mathrm{Field}(N,l)$ *consisting of those fields K with $(1-\zeta_l)^2$-rank equal to j. Then*

$$\left| \lim_{s\to\infty} P(s,s-1,l,j) \cdot |\mathrm{Field}(N,l)| - |\mathrm{Field}(N,l,j)| \right| = O\left(\frac{|\mathrm{Field}(N,l)|}{(\log\log N)^c} \right)$$

for some absolute constant $c > 0$.

Proof. By Theorem 11.3 we know that almost all $r := \omega(D_K)$ satisfy (11.5). Hence it suffices to prove that there exists an absolute constant $c > 0$ with

$$\left|\lim_{s\to\infty} P(s, s-1, l, j)\cdot|\mathrm{Field}(N, r, l)| - |\mathrm{Field}(N, r, l, j)|\right| = O\left(\frac{|\mathrm{Field}(N, r, l)|}{(\log\log N)^c}\right),$$

where $\mathrm{Field}(N, r, l, j)$ is defined in the obvious way. An easy computation shows that we may replace $\lim_{s\to\infty} P(s, s-1, l, j)$ with $P(r, r-1, l, j)$. Put

$$D_1 := \log N, \quad C_0 := \tfrac{1}{10}\log\log\log N.$$

Let W be the subset of $S_r(N, l)$ that is comfortably spaced above D_1 and C_0-regular. By Proposition 11.2 and Theorem 10.7 it is enough to show

$$\left|P(r, r-1, l, j)\cdot|\mathrm{Field}(X)| - |\mathrm{Field}(X)\cap\mathrm{Field}(N, r, l, j)|\right| = O\left(\frac{|\mathrm{Field}(X)|}{(\log\log N)^c}\right)$$

for all boxes X with $j(\mathrm{Field}(X))\cap W \neq \emptyset$ and $\mathrm{Field}(X)\subseteq\mathrm{Field}(N, r, l)$. Let X be such a box and write $x_1, \dots, x_k$ for the unique elements of $X_1, \dots, X_k$. Fix a function $g : \{x_1, \dots, x_k\} \to [l-1]$. Then

$$|\mathrm{Field}(X)| = \frac{1}{W(X)}\sum_f |i_f(X)|, \tag{12.2}$$

where $W(X)$ is a weight depending only on X and the sum is taken over all f in the set $\mathrm{Map}(X_{k+1}\sqcup\cdots\sqcup X_r, [l-1])$. We implicitly extend f to $\mathrm{Map}(X_1\sqcup\cdots\sqcup X_r, [l-1])$ using our function g. We have another identity

$$|\mathrm{Field}(X)\cap\mathrm{Field}(N, r, l, j)| = \frac{1}{W(X)}\sum_f |i_f(X)\cap\mathrm{Field}(N, r, l, j)|. \tag{12.3}$$

Due to (12.2) and (12.3) it suffices to establish

$$\frac{1}{W(X)}\sum_f \left|P(r, r-1, l, j)\cdot|i_f(X)| - |i_f(X)\cap\mathrm{Field}(N, r, l, j)|\right| = O\left(\frac{|\mathrm{Field}(X)|}{(\log\log N)^c}\right). \tag{12.4}$$

Define $g(l, \emptyset, k, j)$ to be the number of functions a such that

- a agrees with X at stage k;
- the Rédei matrix A associated to a has kernel of rank j.

Then we claim

$$\left|\frac{g(l, \emptyset, k, j)}{g(l, \emptyset, k)} - P(r, r-1, l, j)\right| = O\left(\frac{1}{\sqrt{\log N}}\right). \tag{12.5}$$

We have

$$g(l, \emptyset, k) = l^{r^2-r-k^2+k}. \tag{12.6}$$

Let $M \in \mathrm{Mat}(k, k, \mathbb{F}_l)$ agree with X, i.e.

$$\zeta_l^{M(i,j)} = \chi_{x_j}(\mathrm{Frob}(x_i)) \quad \text{for all } 1 \le i, j \le k \text{ with } i \ne j,$$

where x_i and x_j are the unique elements of X_i and X_j respectively. Then

$$g(l, \emptyset, k, j) = \sum_{\substack{M \in \mathrm{Mat}(k,k,\mathbb{F}_l) \\ M \text{ agrees with } X}} Q(r, r-1, l, M, j) l^{r^2-r-k^2}. \tag{12.7}$$

We combine (12.6) and (12.7) to deduce

$$\frac{g(l, \emptyset, k, j)}{g(l, \emptyset, k)} = \frac{1}{l^k} \sum_{\substack{M \in \mathrm{Mat}(k,k,\mathbb{F}_l) \\ M \text{ agrees with } X}} Q(r, r-1, l, M, j). \tag{12.8}$$

Since our box X is C_0-regular, we are in a position to apply Lemma 12.1 to the sum in (12.8). Using once more that X is C_0-regular, we see that k is roughly $\log r$. Hence we can fit the difference $|P(r, r-1, l, j) - Q(r, r-1, l, M, j)|$ in the error of (12.5), thus establishing (12.5). Because of (12.4) and (12.5) it remains to prove

$$\frac{1}{W(X)} \sum_f \left| \frac{g(l, \emptyset, k, j)}{g(l, \emptyset, k)} \cdot |i_f(X)| - |i_f(X) \cap \mathrm{Field}(N, r, l, j)| \right| = O\left(\frac{|\mathrm{Field}(X)|}{(\log\log N)^c} \right).$$

We observe that $i_f(X)$ is equal to the disjoint union of $X(a, f)$ over a. An application of Theorem 11.6 finishes the proof. ∎

13. Proof of Theorem 1.2

The goal of this section is to prove Theorem 1.2, which will follow from a combination of Theorem 7.7, Theorem 7.8 and Proposition 8.7. Unfortunately, these results are only valid under very strong conditions. Hence most of the work in this section are reduction steps. Before we start the proof of Theorem 1.2, we need a definition.

Definition 13.1. Let N be a large real and let X be a box. Put

$$D_1 := e^{(\log\log N)^2}, \quad C_0 := \tfrac{1}{10} \log\log\log N.$$

Define W to be the subset of $S_r(N, l)$ that is comfortably spaced above D_1 and C_0-regular. We say that X is a *nice box* for N if

- r satisfies (11.6);
- $j(\mathrm{Field}(X)) \cap W \ne \emptyset$;
- $\mathrm{Field}(X) \subseteq \mathrm{Field}(N, l)$.

Proposition 13.2. *Assume GRH and let l be an odd prime. There are $c, A, N_0 > 0$ such that for all $N > N_0$, all nice boxes X for N, all integers $m \geq 2$ and all sequences $n_2 \geq \cdots \geq n_{m+1} \geq 0$ of integers, we have*

$$\left|\left|\mathrm{Field}(X) \cap \bigcap_{k=2}^{m+1} D_{l,k}(n_k)\right| - P(n_{m+1}|n_m) \cdot \left|\mathrm{Field}(X) \cap \bigcap_{k=2}^{m} D_{l,k}(n_k)\right|\right| \leq \frac{A|\mathrm{Field}(X)|}{(\log\log N)^{\frac{c}{m^2(l^2+l)^m}}}.$$

Proof that Proposition 13.2 implies Theorem 1.2. Theorem 11.3 implies that we need only consider $\mathrm{Field}(N, r, l)$ with r satisfying (11.6). Now apply Proposition 11.2 with W as in Definition 13.1 and use the lower bound for W established in Theorem 10.7. ∎

For the remainder of this paper, a will always denote an assignment from M to $\langle \zeta_l \rangle$. We let A be the *Rédei matrix* associated to a, i.e. A is the unique matrix with entries $a(i, j)$ and the property

$$\prod_{j=1}^{r} a(i, j) = 1 \quad \text{for all } 1 \leq i \leq r,$$

which uniquely specifies $a(i, i)$. In this section it is essential to keep track of the characters we have chosen. If S is a subset of $[r]$, we define

$$\mathrm{Ch}(S) := \mathrm{Map}\Big(\coprod_{i \in S} X_i, [l-1]\Big).$$

Furthermore, we set

$$W(X, S) := |\{f \in \mathrm{Ch}(S) : K \in i_f(X)\}|,$$

where K is any field in $i_f(X)$. Note that this does not depend on the choice of K.

Definition 13.3. Let V be the $\mathbb{F}_l$-vector space $\mathbb{F}_l^r$, which we think of as column vectors. Given the assignment $a : M \to \langle \zeta_l \rangle$ and the associated Rédei matrix A, we define

$$D_{a,2} := \{v \in V : v^T A = 0\}, \quad D_{a,2}^{\vee} := \{v \in V : Av = 0\},$$

where we think of A as having entries in $\mathbb{F}_l$ through the isomorphism j_l^{-1}. Put

$$n_{\max} := \left\lfloor \sqrt{\frac{c'}{m^2(l^2+l)^m} \log\log\log N} \right\rfloor, \quad n_2 := -1 + \dim_{\mathbb{F}_l} D_{a,2}$$

with c' a small constant depending only on l. Define $R := (1, \ldots, 1) \in \mathbb{F}_l^r$ and

$$\alpha := |\{j \in [r] : r/4 \leq j \leq r/3\}|.$$

We say that the assignment $a : M \to \langle \zeta_l \rangle$ is *generic* if

- $n_2 \le n_{\max}$;
- for all $i \in \mathbb{F}_l$, all $T_1 \in D_{a,2}$ and all $T_2 \in D_{a,2}^\vee$ such that $T_1 \neq 0$ or $T_2 \notin \langle R \rangle$,

$$\left| |\{j \in [r] : r/4 \le j \le r/3 \text{ and } \pi_j(T_1 + T_2) = i\}| - \frac{\alpha}{l} \right| \le \frac{r}{(\log\log N)^{1/10}}, \quad (13.1)$$

where π_j is the projection on the j-th coordinate.

During our proof we will fix all previous Artin pairings, and then prove that the m-th Artin pairing is equidistributed. We formalize this in the following definition.

Definition 13.4. Fix an assignment $a : M \to \langle \zeta_l \rangle$. Let $m \ge 2$ and choose filtrations of $\mathbb{F}_l$-vector spaces

$$D_{a,2} \supseteq \cdots \supseteq D_{a,m}, \qquad D_{a,2}^\vee \supseteq \cdots \supseteq D_{a,m}^\vee$$

with $R \in D_{a,m}^\vee$. For $2 \le k \le m$ we define an integer n_k by

$$n_k := -1 + \dim_{\mathbb{F}_l} D_{a,k}.$$

For $2 \le k < m$, choose a bilinear pairing

$$\mathrm{Art}_k : D_{a,k} \times D_{a,k}^\vee \to \mathbb{F}_l$$

with left kernel $D_{a,k+1}$ and right kernel $D_{a,k+1}^\vee$. We call the set $\{\mathrm{Art}_k\}_{2 \le k < m}$ a *sequence of valid Artin pairings*. Given a sequence of valid Artin pairings, we define

$$X(a, f, i) := \left\{ x \in X(a, f) : \text{the Artin pairing of } x \text{ agrees with } \{\mathrm{Art}_k\}_{2 \le k \le i} \right\}.$$

Proposition 13.5. *Assume GRH and let l be an odd prime. There are $c, A, N_0 > 0$ such that for all $N > N_0$, all nice boxes X for N, all generic assignments $a : M \to \langle \zeta_l \rangle$ that agree with X, all integers $m \ge 2$, all sequences of valid Artin pairings $\{\mathrm{Art}_k\}_{2 \le k < m}$ and a valid Artin pairing Art_m, and for $S := [r] - [k]$, we have*

$$\frac{1}{W(X,S)} \sum_{f \in \mathrm{Ch}(S)} \left| |X(a, f, m)| - l^{-n_m(n_m+1)} \cdot |X(a, f, m-1)| \right| \le \frac{A}{W(X,S)} \sum_{f \in \mathrm{Ch}(S)} \frac{|X(a,f)|}{(\log\log N)^{\frac{c}{m(l^2+l)^m}}}.$$

Proof that Proposition 13.5 implies Proposition 13.2. Note that $X(a, f)$ is in fact a slight abuse of notation, since this is only defined for $f \in \mathrm{Ch}([r])$. However, one of the assumptions in the statement of Proposition 13.5 is that a agrees with X, and this involves a choice of a function g in $\mathrm{Map}(\{x_1, \ldots, x_k\}, [l-1])$, where $x_1, \ldots, x_k$ are the unique elements of $X_1, \ldots, X_k$. Hence, whenever we write $f \in \mathrm{Ch}(S)$, we mean the function f extended to $\mathrm{Ch}([r])$ using g. We have the identities

$$\left| \mathrm{Field}(X) \cap \bigcap_{k=2}^{m+1} D_{l,k}(n_k) \right| = \frac{1}{W(X,S)} \sum_{f \in \mathrm{Ch}(S)} \left| i_f(X) \cap \bigcap_{k=2}^{m+1} D_{l,k}(n_k) \right|$$

and

$$\left|\mathrm{Field}(X) \cap \bigcap_{k=2}^{m} D_{l,k}(n_k)\right| = \frac{1}{W(X,S)} \sum_{f \in \mathrm{Ch}(S)} \left|i_f(X) \cap \bigcap_{k=2}^{m} D_{l,k}(n_k)\right|.$$

Hence the LHS of Proposition 13.2 is upper bounded by

$$\frac{1}{W(X,S)} \sum_{f \in \mathrm{Ch}(S)} \left|\left|i_f(X) \cap \bigcap_{k=2}^{m+1} D_{l,k}(n_k)\right| - P(n_{m+1}|n_m) \cdot \left|i_f(X) \cap \bigcap_{k=2}^{m} D_{l,k}(n_k)\right|\right|. \tag{13.2}$$

Since $i_f(X)$ is the disjoint union of $X(a, f, m)$, the quantity in (13.2) is at most

$$\frac{1}{W(X,S)} \sum_{a} \sum_{f \in \mathrm{Ch}(S)} \sum_{\{\mathrm{Art}_k\}_{2 \le k \le m}} \left||X(a,f,m)| - l^{-n_m(n_m+1)} \cdot |X(a,f,m-1)|\right|, \tag{13.3}$$

where the first sum is over all a that agree with X and the last sum is over all sequences of valid Artin pairings with

$$\dim_{\mathbb{F}_l} D_{a,k} = n_k + 1 \quad \text{for all } 2 \le k \le m+1.$$

We split the sum in (13.3) into two parts depending on the genericity of a. If a is generic, we use Proposition 13.5 to bound the sum. There are at most

$$l^{mn_2(n_2+1)} \le l^{mn_{\max}(n_{\max}+1)}$$

sequences of valid Artin pairings, so the sum is within the error of Proposition 13.2. If a is not generic, we employ the trivial bound to (13.3) inducing an error of size at most

$$\frac{2}{W(X,S)} \sum_{a \text{ not generic}} \sum_{f \in \mathrm{Ch}(S)} |X(a,f)|. \tag{13.4}$$

The sum over f is easily bounded by Theorem 11.6. So it remains to count the number of a's that agree with X and are not generic, which is a purely combinatorial problem. We first deal with the a for which $n_2 > n_{\max}$. These a are easily bounded using the ideas from the proof of Theorem 12.2.

Now consider the assignments $a : M \to \langle \zeta_l \rangle$ not satisfying (13.1). We have to estimate

$$\frac{|\{a \text{ assignment} : a \text{ agrees with } X \text{ and fails (13.1)}\}|}{|\{a \text{ assignment} : a \text{ agrees with } X\}|},$$

which is clearly upper bounded by

$$2^{k^2} \frac{|\{a \text{ assignment} : a \text{ fails (13.1)}\}|}{|\{a \text{ assignment}\}|}. \tag{13.5}$$

We first count the number of pairs (T_1, T_2) that fail (13.1) with $T_1 \neq 0$ and T_2 linearly independent of R. From Hoeffding's inequality we deduce that the proportion of such pairs is at most

$$O\left(\exp\left(\frac{-r}{12\cdot(\log\log N)^{1/5}}\right)\right).$$

Now observe that the number of a's for which $T_1 \in D_{a,2}$ and $T_2 \in D_{a,2}^\vee$ does not depend on the pair (T_1, T_2) provided that $T_1 \neq 0$ and that T_2 is linearly independent of R. Hence we get the desired upper bound for (13.5). We still need to deal with the case $T_1 = 0$ and $T_2 \in \langle R\rangle$. In both cases we apply Hoeffding's inequality once more, and proceed along the same lines. This proves the proposition. ∎

For generic a, our next goal is to find sets S for which we can apply Theorems 7.7 and 7.8. Following Smith [22], we call such sets S *variable indices*.

Definition 13.6. Let $a : M \to \langle\zeta_l\rangle$ be an assignment and let $m \geq 2$ be an integer. For the rest of this paper, fix a basis $w_{2,1}, \ldots, w_{2,n_2+1}$ for $D_{a,2}$ and fix a basis $w_{1,1}, \ldots, w_{1,n_2}, R$ for $D_{a,2}^\vee$ in such a way that for all $2 \leq k \leq m$, $w_{2,1}, \ldots, w_{2,n_k+1}$ is a basis for $D_{a,k}$ and $w_{1,1}, \ldots, w_{1,n_k}, R$ is a basis for $D_{a,k}^\vee$. Let $1 \leq j_1 \leq n_m + 1$ and $1 \leq j_2 \leq n_m$ be integers. We say that $S(j_1, j_2) \subseteq [r]$ is a *set of variable indices for* (j_1, j_2) if there are integers $i_1(j_1, j_2)$ and $i_2(j_1, j_2)$ with the following properties:

- $|S(j_1, j_2)| = m + 1$;
- $i_1(j_1, j_2), i_2(j_1, j_2) \in S(j_1, j_2)$;
- $S(j_1, j_2)$ lies in the zero set of $w_{1,j}$ for all $j \leq n_2$ other than j_1;
- $S(j_1, j_2)$ lies in the zero set of $w_{2,j}$ for all $j \leq n_2 + 1$ other than j_2;
- $\{i \in [r] : \pi_i(w_{1,j_2}) \neq 0\} \cap S(j_1, j_2) = \{i_1(j_1, j_2)\}$;
- $\{i \in [r] : \pi_i(w_{2,j_1}) \neq 0\} \cap S(j_1, j_2) = \{i_2(j_1, j_2)\}$;
- $S(j_1, j_2) - \{i_2(j_1, j_2)\} \subseteq [r/3, r/4]$.

With this definition in place, we are ready to find variable indices for generic a. We do so with the following lemma.

Lemma 13.7. *Let $a : M \to \langle\zeta_l\rangle$ be a generic assignment. If $w_1, \ldots, w_d \in D_{a,2}$ are linearly independent and also $w_{d+1}, \ldots, w_e, R \in D_{a,2}^\vee$ are linearly independent, then for all $\mathbf{v} \in \mathbb{F}_l^e$,*

$$\left|\left|\{i \in [r] : r/4 \leq i \leq r/3 \text{ and } \pi_i(w_j) = \pi_j(\mathbf{v}) \text{ for all } 1 \leq j \leq e\}\right| - \frac{\alpha}{l^e}\right| \leq \frac{100^e \cdot r}{(\log\log N)^{1/10}}.$$

Proof. The case $e = 1$ follows easily from the genericity condition on a. We start with the special case $d = 1, e = 2$. Define, for $\mathbf{x} \in \mathbb{F}_l^2$,

$$g(\mathbf{x}) = |\{i \in [r] : r/4 \leq i \leq r/3 \text{ and } \pi_i(w_j) = \pi_j(\mathbf{x}) \text{ for all } 1 \leq j \leq 2\}|.$$

We use (13.1) with the pair (T_1, T_2) in the set

$$\{(w_1, \beta w_2) : \beta \in \mathbb{F}_l\} \cup \{(0, w_2)\}.$$

Then, if $\mathbf{x}_1, \dots, \mathbf{x}_l$ lie on an affine line in $\mathbb{F}_l^2$, we have

$$\left| \left(\sum_{i=1}^{l} g(\mathbf{x}_i) \right) - \frac{\alpha}{l} \right| \leq \frac{r}{(\log\log N)^{1/10}}. \tag{13.6}$$

Now take an element $\mathbf{v} \in \mathbb{F}_l^2$. Let $L(\mathbf{v})$ be the collection of affine lines in $\mathbb{F}_l^2$ through $\mathbf{v}$. We use (13.6) for all elements in $L(\mathbf{v})$ to deduce that

$$l \left| g(\mathbf{v}) - \frac{\alpha}{l^2} \right| = \left| \sum_{L(\mathbf{v})} \sum_{\mathbf{x} \in L(\mathbf{v})} \left(g(\mathbf{x}) - \frac{\alpha}{l} \right) \right| \leq \sum_{L(\mathbf{v})} \left| \sum_{\mathbf{x} \in L(\mathbf{v})} \left(g(\mathbf{x}) - \frac{\alpha}{l} \right) \right| \leq \frac{(l+1) \cdot r}{(\log\log N)^{1/10}}.$$

Since the special cases $d = 0$, $e = 2$ and $d = e = 2$ are trivial, this settles the case $e = 2$. An easy induction establishes the lemma for all $e > 2$. ∎

Fix two integers $1 \leq j_1 \leq n_m + 1$ and $1 \leq j_2 \leq n_m$. We will now demonstrate how to find variable indices $S(j_1, j_2)$ for (j_1, j_2) using Lemma 13.7. We apply Lemma 13.7 with $w_{2,1}, \dots, w_{2,n_2+1} \in D_{a,2}$ and $w_{1,1}, \dots, w_{1,n_2}, R \in D_{a,2}^{\vee}$, so $d = n_2 + 1$ and $e = 2n_2 + 1$. We let $\mathbf{v} \in \mathbb{F}_l^e$ be the unique vector satisfying $\pi_j(\mathbf{v}) = 1$ for $j = j_1$ and $j = d + j_2$, and furthermore $\pi_j(\mathbf{v}) = 0$ for all other j. With these choices we can choose $S(j_1, j_2)$ to be any subset of

$$\{i \in [r] : r/4 \leq i \leq r/3 \text{ and } \pi_i(w_j) = \pi_j(\mathbf{v}) \text{ for all } 1 \leq j \leq e\}$$

with $|S(j_1, j_2)| = m + 1$. Then we must have

$$m + 1 \leq \frac{\alpha}{l^{2n_2+1}} - \frac{100^{2n_2+1} \cdot r}{(\log\log N)^{1/10}}. \tag{13.7}$$

We always have

$$m < \log\log\log\log N,$$

since otherwise Theorem 1.2 is trivial. If N is sufficiently large, this implies (13.7). Having found our variable indices, we are ready for our next reduction step.

Proposition 13.8. *Assume GRH and let l be an odd prime. There are $c, A, N_0 > 0$ such that for all $N > N_0$, all nice boxes X for N, all generic assignments $a : M \to \langle \zeta_l \rangle$ that agree with X, all integers $m \geq 2$, all sequences $\{\mathrm{Art}_k\}_{2 \leq k < m}$ of valid Artin pairings, a valid Artin pairing Art_m and a non-zero multiplicative character F from $\mathrm{Mat}(n_m + 1, n_m, \mathbb{F}_l)$ to $\langle \zeta_l \rangle$, we have*

$$\frac{1}{W(X,S)} \sum_{f \in \mathrm{Ch}(S)} \left| \sum_{x \in X(a,f,m-1)} F(\mathrm{Art}(x, f, m)) \right| \leq \frac{A}{W(X,S)} \sum_{f \in \mathrm{Ch}(S)} \frac{|X(a,f)|}{(\log\log N)^{\frac{c}{m(l^2+l)^m}}},$$

where $S := [r] - [k]$ and $\mathrm{Art}(x, f, m)$ is the m-th Artin pairing of the field $i_f(x)$.

Proof that Proposition 13.8 implies Proposition 13.5. This is straightforward. ∎

Before we can make our final reduction step, we must restrict ourselves to rather special product spaces. Our next definition will make this precise.

Definition 13.9. Let $a : M \to \langle \zeta_l \rangle$ be a generic assignment. Define $S_{\mathrm{var}} := S(j_1, j_2) - \{i_2(j_1, j_2)\}$. For each $i \in S_{\mathrm{var}}$, let Z_i be subsets of X_i of cardinality

$$M_{\mathrm{box}} := \lfloor (\log\log N)^{\frac{1}{10m}} \rfloor,$$

which is at least l for sufficiently large N. Put

$$Z := \prod_{i \in S_{\mathrm{var}}} Z_i.$$

We say that Z is *well-governed* if there is a governing expansion $\mathfrak{G}$ on Z such that $i_1(j_1, j_2) = i_a(\mathfrak{G})$ and furthermore all $\bar{x} \in \overline{Z}_{S_{\mathrm{var}}}$ satisfy $\bar{x} \in \overline{Y}_{S_{\mathrm{var}}}(\mathfrak{G})$. Set

$$M(Z) := \prod_{\bar{x} \in \overline{Z}_{S_{\mathrm{var}}}} L(\phi_{S_{\mathrm{var}}, \bar{x}}), \quad M_\circ(Z) := \prod_{S \subsetneq S_{\mathrm{var}}} \prod_{\bar{x} \in \overline{Z}_S} L(\phi_{S, \bar{x}}).$$

Define

$$M^\circ(Z) := L(Z \cup \{x_1, \ldots, x_k\}) M_\circ(Z).$$

For $i > k$ not in S_{var}, we define $X_i(M^\circ(Z))$ to be the subset consisting of the $p \in X_i$ satisfying

- p splits completely in the extension $M_\circ(Z)/\mathbb{Q}$;
- for all $j \in S_{\mathrm{var}}$ and all $x \in Z_j$,

$$\chi_x^{f(x)}(\mathrm{Frob}(p)) = a(i, j), \quad \chi_p^{f(p)}(\mathrm{Frob}(x_j)) = a(j, i);$$

- for all $j \in [k]$,

$$\chi_{x_j}^{f(x_j)}(\mathrm{Frob}(p)) = a(i, j), \quad \chi_p^{f(p)}(\mathrm{Frob}(x_j)) = a(j, i).$$

Take

$$\widetilde{Z} := \prod_{1 \le i \le k} X_i \times Z \times \prod_{\substack{i > k \\ i \notin S_{\mathrm{var}}}} X_i(M^\circ(Z)).$$

We call $\widetilde{Z}$ a *satisfactory product space* if

- Z is well-governed;
- $x_1, \ldots, x_k$ split completely in the extension $M_\circ(Z)/\mathbb{Q}$;
- for all distinct $i, j \in S_{\mathrm{var}}$, all $x_i \in \pi_i(Z)$ and all $x_j \in \pi_j(Z)$,

$$\chi_{x_j}^{f(x_j)}(\mathrm{Frob}(x_i)) = a(i, j).$$

We remind the reader that our equidistribution comes from a combination of Theorem 7.7, Theorem 7.8 and Proposition 8.7. We will apply these results to satisfactory product spaces $\widetilde{Z}$. In order to do so, we need to construct an l-additive system $\mathfrak{A}$ on $\widetilde{Z}$ that satisfies the conditions of Proposition 8.7. This is done in our next lemma.

Lemma 13.10. *Let l be an odd prime and let $\widetilde{Z}$ be a satisfactory product space. Let P be an element of $\pi_{[r]-S(j_1,j_2)}(\widetilde{Z})$ and define*

$$\widetilde{Z}(P) := \{P\} \times Z \times X_{i_2(j_1,j_2)}(M^\circ(Z)).$$

Let F be a non-zero multiplicative character from $\mathrm{Mat}(n_m + 1, n_m, \mathbb{F}_l)$ to $\langle \zeta_l \rangle$ that depends on the entry (j_1, j_2). There exists an l-additive system $\mathfrak{A}$ on $\widetilde{Z}(P)$ such that

- $\overline{Y}^\circ_\emptyset(\mathfrak{A}) = X(a, f, m-1) \cap \widetilde{Z}(P)$;
- *$\mathfrak{A}$ is $S(j_1, j_2)$-acceptable (see Definition 8.6) with $|A_T(\mathfrak{A})|$ bounded by $l^{n_2(n_2+m+1)}$;*
- *for all $\bar{x} \in \overline{Z}_{S(j_1,j_2)}(\mathfrak{A})$ (see once more Definition 8.6),*

$$d\,\widetilde{F}(\bar{x}) = j_l^{-1}(F(j_1, j_2)) \cdot \pi_{i_2(j_1,j_2)}(w_{2,j_2}) \cdot \phi_{S_{\mathrm{var}},\bar{z}}(\mathrm{Frob}(p_1) \cdot \ldots \cdot \mathrm{Frob}(p_l)),$$

where $p_i := \mathrm{pr}_i(\pi_{i_2(j_1,j_2)}(\bar{x}))$ for $i \in [l]$, $\widetilde{F}(x) := F(\mathrm{Art}(x, f, m))$ and $\bar{z}$ is any element of $\bar{x}(S_{\mathrm{var}})$.

Proof. We will start by constructing an l-additive system $\mathfrak{A}$ and then verify the required properties. To do this, we need to introduce the concept of acceptable ramification, which is based on Smith [22, p. 33]. We write W for $\widetilde{Z}(P)$. If $w \in D^\vee_{a,2}$, we define a raw cocycle $\mathfrak{R}(w)$ for (W, f) to be a choice of raw cocycle for each $x \in W$ such that the ψ_1 of this raw cocycle is equal to the character naturally associated to w. Here and later we shall often suppress the dependence on f. Now choose a raw cocycle $\mathfrak{R}(w_{1,j})$ for (W, f), where j runs through $1, \ldots, n_2$. Define M to be the compositum of all the $L(\psi_k(\mathfrak{R}, x))$, where $\mathfrak{R}$ is one of the raw cocycles $\mathfrak{R}(w_{1,j})$, $x \in W$ and $k \leq \min(m, \mathrm{rk}(\mathfrak{R})(x))$.

Let M' be the compositum of M with $\phi_{S_{\mathrm{var}},\bar{x}}$ for all $\bar{x} \in \overline{W}_{S_{\mathrm{var}}}$. If p ramifies in M', then the ramification degree of p is l, and we choose an element $\sigma_p \in \mathrm{Gal}(M'/\mathbb{Q})$ such that $\langle \sigma_p \rangle$ is the inertia group of some prime dividing p in M'. We assume that the $\phi_{S,\bar{x}}$ are normalized in such a way that $\phi_{S,\bar{x}}(\sigma_p) = 0$ for all p ramifying in M. Here and for the rest of the proof, $\phi_{S,\bar{x}}$ is defined to be the zero map if S does not contain $i_1(j_1, j_2)$. Let S be a subset of $S(j_1, j_2)$, $\bar{x} \in \overline{W}_S$ and suppose that

$$\mathrm{rk}(\mathfrak{R}(w_{1,j})) \geq |S| + 1 \quad \text{for all } x \in \bar{x}(\emptyset).$$

We say that $\mathfrak{R}(w_{1,j})$ is *acceptably ramified* at $(\bar{x}, i)$ for $i \in S_{\mathrm{var}} - S$ if

$$\sum_{x \in \bar{x}(\emptyset)} \psi_{|S|+1}(\mathfrak{R}(w_{1,j}), x)(\sigma_{\pi_i(\bar{x})}) = 0.$$

We inductively define a subset $\overline{Y}^\circ_S(\mathfrak{A})$ in $\overline{Y}_S(\mathfrak{A})$. If $S = \emptyset$, we have already done this by the first property of $\mathfrak{A}$ in the lemma. Now suppose that S is a subset of S_{var} of cardinality

at most $|S_{\mathrm{var}}| - 1$. If S does not satisfy these two conditions, we let $F_S(\mathfrak{A})$ be the zero map and $\overline{Y}^\circ_S(\mathfrak{A}) = \overline{Y}_S(\mathfrak{A})$.

Let $\bar{x} \in \overline{Y}_S(\mathfrak{A})$. If $j \neq i_1(j_1, j_2)$, we know that

$$\psi(\Re(w_{1,j}), \bar{x}) := \sum_{x \in \bar{x}(\emptyset)} \psi_{|S|}(\Re(w_{1,j}), x).$$

is a character. If instead $j = i_1(j_1, j_2)$, then $\psi(\Re(w_{1,i_1(j_1,j_2)}), \bar{x}) - \phi_{S,\bar{x}}$ is a character. We call this character $\chi(\bar{x}, j)$ in both cases. We define $\overline{Y}^\circ_S(\mathfrak{A})$ to be those $\bar{x} \in \overline{Y}_S(\mathfrak{A})$ such that $\chi(\bar{x}, j) = 0$ for all j and furthermore $\Re(w_{1,j})$ is acceptably ramified for all j and $i \in S_{\mathrm{var}} - S$. We will now describe how to encode this as a map $F_S(\mathfrak{A})$ and we do so for each j separately.

The acceptable ramification can be encoded with an additive map to $\mathbb{F}_l^{|S_{\mathrm{var}}-S|}$. To deal with $\chi(\bar{x}, j)$, we remark that the acceptable ramification conditions at the stages $S - \{i\}$ ensure that $\chi(\bar{x}, j)$ is an unramified character above $K_{\chi_{x,f}}$ for all $x \in \bar{x}(\emptyset)$. In particular, $\chi(\bar{x}, j)$ is supported outside the primes in S.

Let p be a prime ramifying in $K_{\chi_{x,f}}$ and not in $\pi_S(x)$. We claim that the prime above p splits in the extension $K_{\chi(\bar{x},j)}K_{\chi_{x,f}}/K_{\chi_{x,f}}$. By construction of $\overline{Y}^\circ_\emptyset$ we have

$$\mathrm{rk}(\Re(w_{1,j}))(x) > |S| \quad \text{for each } x \in \bar{x}(\emptyset).$$

This implies that p has residue field degree 1 in $\psi_S(\Re(w_{1,j}), x)$. Since $\overline{Y}^\circ_\emptyset$ is also contained in $X(a, f)$ by construction, we conclude that p has residue field degree 1 in the compositum of the $\psi_S(\Re(w_{1,j}), x)$ over all $x \in \bar{x}(\emptyset)$ and $\phi_{S,\bar{x}}$. Since $K_{\chi(\bar{x},j)}$ is contained in this compositum, we have proven our claim.

Hence, in order to test if $\chi(\bar{x}, j)$ is zero, we merely have to check that the primes in $\pi_S(x)$ split completely in $K_{\chi(\bar{x},j)}K_{\chi_{x,f}}/K_{\chi_{x,f}}$ and furthermore that $\chi(\bar{x}, j)$ is zero in the vector space $D^\vee_{a,2}$. We make this precise in the following way. It follows from Lemma 13.7 that there exists a subset $A \subseteq [r] - S(j_1, j_2)$ and a bijection $g : [n_2 + 1] \to A$ such that

$$\pi_{g(k_1)}(w_{2,k_2}) = \delta_{k_1,k_2} \text{ for all } k_1, k_2 \in [n_2 + 1]$$

with δ_{k_1,k_2} the Kronecker delta function. Fix an $x \in \bar{x}(\emptyset)$ and define the map $F_S(\mathfrak{A})$ that sends $\chi(\bar{x}, j)$ to

$$\chi(\bar{x}, j)(\mathrm{Frob}(\mathrm{Up}_{K_{\chi_{x,f}}}(z))), \quad \chi(\bar{x}, j)(\sigma_{\pi_{g(z')}(x)}).$$

Here z runs through the primes in $\pi_S(x)$ and z' runs through $[n_2 + 1]$. The key property is that this does not depend on x, which follows from the fact that $\overline{Y}^\circ_\emptyset$ is contained in $X(a, f)$ and $\chi(\bar{x}, j)$ is supported outside the primes in S. Once this is established, it is not hard to show that $F_S(\mathfrak{A})$ is additive. This describes our l-additive system $\mathfrak{A}$. We will now demonstrate that $\mathfrak{A}$ is $S(j_1, j_2)$-acceptable.

Let $\bar{x} \in \overline{Z}_{S(j_1,j_2)}(\mathfrak{A})$. If $\pi_i(\bar{x}(S(j_1, j_2) - \{i\}))$ consists of only one element for some $i \in S$, the condition $\bar{x}(\emptyset) \subseteq \overline{Y}^\circ_\emptyset(\mathfrak{A})$ is trivially satisfied. So now suppose that

$$|\pi_i(\bar{x}(S(j_1, j_2) - \{i\}))| > 1$$

for all $i \in S$. Let $\bar{z}_{i,m_1}, \dots, \bar{z}_{i,m_i}$ be a distinct list of elements of $\bar{x}(S(j_1, j_2) - \{i\})$. By assumption we have

$$m_i \geq |\pi_i(\bar{x}(S(j_1, j_2) - \{i\}))| - 1$$

for all $i \in S$. In case we have strict inequality for some $i \in S$, it is immediate that $\mathfrak{A}$ is $S(j_1, j_2)$-acceptable. So suppose that we have equality for all $i \in S$, and let x_0 be the unique element outside all of the $\bar{z}_{i,m}(\emptyset)$. We have to show that $x_0 \in X(a, f, m-1)$.

We start by checking that x_0 is in $X(a, f)$, so take two distinct indices $i, i' \in S$. Take the $l-1$ other points, with multiplicity, in the set $\bar{x}(\emptyset)$ having the property

$$\pi_{[r]-i'}(x_0) = \pi_{[r]-i'}(x).$$

All these $l-1$ points are in $X(a, f)$. Now the splitting conditions coming from the existence of $\phi_{S,\bar{x}}$ show that x_0 must be in $X(a, f)$ as well. It remains to prove that $\mathrm{Art}(x, f, i)$ is equal to the Artin pairing Art_i for all $2 \leq i \leq m-1$.

So take such an i and let S a subset of $S(j_1, j_2)$ with $i+1$ elements not containing $i_1(j_1, j_2)$. Take any cube $\bar{y} \in \bar{x}(S)$ with $x_0 \in \bar{y}(\emptyset)$. We apply Theorem 7.7 with $\bar{y}$, S and for all raw cocycles $\mathfrak{R}(w_{1,j})$ and all b corresponding to some $w_{2,j'}$. Since $i_1(j_1, j_2) \notin S$, we have minimality in all cases. From Theorem 7.7 we deduce that $\mathrm{Art}(x_0, f, m)$ is equal to Art_i.

Finally, we must check that $\mathfrak{A}$ satisfies the third property listed in the lemma. But this follows from an application of Theorems 7.7 and 7.8. ∎

Proposition 13.11. *Assume GRH and let l be an odd prime. There are $c, A, N_0 > 0$ such that for all $N > N_0$, all nice boxes X for N, all generic assignments $a : M \to \langle \zeta_l \rangle$ that agree with X, all integers $m \geq 2$, all sequences $\{\mathrm{Art}_k\}_{2 \leq k < m}$ of valid Artin pairings, a valid Artin pairing Art_m, a non-zero multiplicative character F from $\mathrm{Mat}(n_m + 1, n_m, \mathbb{F}_l)$ to $\langle \zeta_l \rangle$ that depends on the entry (j_1, j_2) and all satisfactory product spaces $\widetilde{Z}$, and for $S := [r] - [k] - S_{\mathrm{var}}$ we have*

$$\frac{1}{W(X,S)} \sum_{f \in \mathrm{Ch}(S)} \Big| \sum_{x \in \widetilde{Z} \cap X(a,f,m-1)} F(\mathrm{Art}(x, f, m)) \Big| \leq \frac{A}{W(X,S)} \sum_{f \in \mathrm{Ch}(S)} \frac{|\widetilde{Z} \cap X(a, f)|}{(\log\log N)^{\frac{c}{m(l^2+l)^m}}}.$$

Proof that Proposition 13.11 implies Proposition 13.8. Define, for $i > k$,

$$X_i(a) := \{x \in X_i : \chi_x^{f(x)}(\mathrm{Frob}(x_j)) = a(j,i) \text{ and } \chi_{x_j}^{f(x_j)}(\mathrm{Frob}(x)) = a(i,j) \text{ for all } 1 \leq j \leq k\}$$

and

$$X_{\mathrm{var}} := \prod_{i \in S_{\mathrm{var}}} X_i(a).$$

We let V_{var} be the subset of X_{var} that is consistent with a, i.e. $P \in V_{\mathrm{var}}$ if and only if for all distinct $i, j \in S_{\mathrm{var}}$,

$$\chi_{\pi_j(P)}^{f(\pi_j(P))}(\mathrm{Frob}(\pi_i(P))) = a(i,j).$$

Set

$$R := \lfloor \exp\big(\exp\big(\tfrac{1}{10}\log\log N\big)\big)\rfloor.$$

We let $Z_{\mathrm{var}}^1, \ldots, Z_{\mathrm{var}}^t$ be a longest sequence of subsets of X_{var} with the following properties:

- for all $1 \le s \le t$,
$$Z_{\mathrm{var}}^s = \prod_{i \in S_{\mathrm{var}}} Z_i^s$$
for some subset Z_i^s of $X_i(a)$ of cardinality M_{box};
- each Z_{var}^s is a subset of V_{var}, and any element of V_{var} is in at most R different Z_{var}^s;
- $|Z_{\mathrm{var}}^s \cap Z_{\mathrm{var}}^{s'}| \le 1$ for all distinct s and s';
- Z_{var}^s is well-governed and $x_1, \ldots, x_k$ split completely in the extension $M_\circ(Z_{\mathrm{var}}^s)/\mathbb{Q}$.

We make the important remark that the sequence $Z_{\mathrm{var}}^1, \ldots, Z_{\mathrm{var}}^t$ depends only on the choice of characters for X_i with $i \in S_{\mathrm{var}}$. Hence if $f, f' \in \mathrm{Ch}([r]-[k])$ restrict to the same function in $\mathrm{Ch}(S_{\mathrm{var}})$, we may and will take the same sequence $Z_{\mathrm{var}}^1, \ldots, Z_{\mathrm{var}}^t$ for f and f'.

Define $V_{\mathrm{var}}^{\mathrm{bad}}$ to be the subset of points in V_{var} that are in fewer than R of the Z_{var}^s and let δ be the density of $V_{\mathrm{var}}^{\mathrm{bad}}$ in X_{var}. Our goal is to give an upper bound for δ using the tools from Section 8. Using a straightforward greedy algorithm, we can find a subset W of $V_{\mathrm{var}}^{\mathrm{bad}}$ of density at least $\delta/(RM_{\mathrm{box}}^m)$ such that $|W \cap Z_{\mathrm{var}}^s| \le 1$ for all s.

Our next step is to construct a "nice" l-additive system $\mathfrak{A}$ on X_{var} with $\overline{Y}_\emptyset^\circ = W$. Using this l-additive system $\mathfrak{A}$ we are going to find a well-governed subset of W if δ is sufficiently large. Since this is impossible by construction of $Z_{\mathrm{var}}^1, \ldots, Z_{\mathrm{var}}^t$, we obtain the desired upper bound for δ. We will now define the l-additive system $\mathfrak{A}$.

First suppose S is a singleton $\{j\}$. Then we define $F_j(\bar{x})$ to be

$$j_l^{-1} \circ \chi_p\big(\mathrm{Frob}\ \mathrm{pr}_1(\pi_j(\bar{x})) \cdot \ldots \cdot \mathrm{Frob}\ \mathrm{pr}_l(\pi_j(\bar{x}))\big),$$

where p runs through the primes in $\pi_{S_{\mathrm{var}}-\{j\}}(\bar{x})$ and l. Hence we can take $A_S(\mathfrak{A}) := \mathbb{F}_l^m$. Now suppose that S is such that $|S| > 1$ and $i_1(j_1, j_2) \in S$. In this case we define $F_S(\bar{x})$ as

$$\phi_{S,\bar{x}}(\mathrm{Frob}\ p),$$

where p runs through $x_1, \ldots, x_k$ and the primes in $\pi_{S_{\mathrm{var}}-S}(\bar{x})$, so that $A_S(\mathfrak{A}) := \mathbb{F}_l^{m-|S|+k}$. We remark that both χ_p and $\phi_{S,\bar{x}}$ implicitly depend on f. By Proposition 8.2 the density of $\overline{Y}_{S_{\mathrm{var}}}^\circ$ in X_{var}^l is at least

$$\delta' := \left(\frac{\delta}{RM_{\mathrm{box}}^m l^{m+k}}\right)^{(l+1)^m}.$$

We will now use some of the techniques and notation from the proof of Proposition 8.7. For $g : [l-1] \to \overline{Y}^{\circ}_{\emptyset}(\mathfrak{A})$ and $x \in \overline{Y}^{\circ}_{\emptyset}(\mathfrak{A})$ we have defined an element $c(g,x)$ of $\overline{X}_{S_{\mathrm{var}}}$. Also define

$$Z(\mathfrak{A}, S_{\mathrm{var}}, g) := \{x \in \overline{Y}^{\circ}_{\emptyset}(\mathfrak{A}) : \text{writing } \bar{x} := c(g,x), \text{ we have } \bar{x} \in \overline{Y}^{\circ}_{S_{\mathrm{var}}}(\mathfrak{A})\}.$$

There is a natural injective map from $\overline{Y}^{\circ}_{S_{\mathrm{var}}}(\mathfrak{A})$ to

$$\coprod_{g:[l-1]\to\overline{Y}^{\circ}_{\emptyset}(\mathfrak{A})} Z(\mathfrak{A}, S_{\mathrm{var}}, g),$$

where $\coprod$ denotes disjoint union. This map is given by sending $\bar{y}$ to the pair (g,x), where $g : [l-1] \to \overline{Y}^{\circ}_{\emptyset}(\mathfrak{A})$ and $x \in Z(\mathfrak{A}, S_{\mathrm{var}}, g)$ are uniquely determined by

$$\pi_i(g(j)) = \mathrm{pr}_j(\pi_i(\bar{y})) \quad \text{and} \quad \pi_i(x) = \mathrm{pr}_l(\pi_i(\bar{y}))$$

for $i \in S_{\mathrm{var}}$ and $j \in [l-1]$. We conclude that

$$|\overline{Y}^{\circ}_{S_{\mathrm{var}}}(\mathfrak{A})| \le \sum_{g:[l-1]\to\overline{Y}^{\circ}_{\emptyset}(\mathfrak{A})} |Z(\mathfrak{A}, S_{\mathrm{var}}, g)| \le |X_{\mathrm{var}}|^{l-1} \cdot \max_{g:[l-1]\to\overline{Y}^{\circ}_{\emptyset}(\mathfrak{A})} |Z(\mathfrak{A}, S_{\mathrm{var}}, g)|.$$

This implies that $Z(\mathfrak{A}, S_{\mathrm{var}}, g)$ has density at least δ' in X_{var} for a good choice of g. Now the key observation is that there are no subsets Z_i of $X_i(a)$ all of cardinality M_{box} with

$$\prod_{i\in S_{\mathrm{var}}} Z_i \subseteq Z(\mathfrak{A}, S_{\mathrm{var}}, g),$$

since then we would be able to extend $Z^1_{\mathrm{var}}, \dots, Z^t_{\mathrm{var}}$ to a longer sequence. Hence we can apply the contrapositive of Proposition 8.3. This yields

$$\delta'^{2} M_{\mathrm{box}}^{m-1} < 2^{m+2} \Big(\min_{i\in S_{\mathrm{var}}} |X_i(a)| \Big)^{-1}.$$

For almost all f we know that $|X_i(a)|$ is of the expected size due to Theorem 11.6. For these f we get the desired upper bound for δ and for the remaining f we employ the trivial bound. It follows from Theorem 11.6 that the contribution of those $x \in X(a,f)$ with $\pi_{S_{\mathrm{var}}}(x) \in V^{\mathrm{bad}}_{\mathrm{var}}$ fits in the error term of Proposition 13.8. Define, for $x \in X(a,f)$,

$$\Lambda(x) := |\{1 \le s \le t : x \in \widetilde{Z^s_{\mathrm{var}}}\}|.$$

We will compute the first and second moments of $\Lambda(x)$. Let $y \in V_{\mathrm{var}}$ and let $d(M_{\mathrm{box}}, m)$ be the degree of $M^{\circ}(Z^s_{\mathrm{var}})$ over $L(y, x_1, \dots, x_k)$. Then $d(M_{\mathrm{box}}, m)$ does not depend on s or y by Proposition 9.4. For $y \in Z^s_{\mathrm{var}}$, Theorem 11.6 implies that the proportion of f's violating

$$\left| |X(a,f) \cap \widetilde{Z^s_{\mathrm{var}}} \cap \pi^{-1}_{S_{\mathrm{var}}}(y)| - \frac{|X(a,f) \cap \pi^{-1}_{S_{\mathrm{var}}}(y)|}{d(M_{\mathrm{box}}, m)^{r-k-m}} \right| \le \frac{A'(l)}{\log N} \frac{|X(a,f) \cap \pi^{-1}_{S_{\mathrm{var}}}(y)|}{d(M_{\mathrm{box}}, m)^{r-k-m}}$$

for some $1 \le s \le t$ is within the error of the proposition, where $A'(l)$ is a sufficiently large constant. Therefore

$$\begin{aligned}\sum_{x\in X(a,f)} \Lambda(x) &= \sum_{y\in V_{\mathrm{var}}} \sum_{\substack{x\in X(a,f)\\ \pi_{S_{\mathrm{var}}}(x)=y}} \sum_{1\le s\le t} \mathbf{1}_{x\in \widetilde{Z}^s_{\mathrm{var}}}\\ &= \sum_{y\in V_{\mathrm{var}}} \sum_{1\le s\le t} |X(a,f)\cap \widetilde{Z}^s_{\mathrm{var}} \cap \pi^{-1}_{S_{\mathrm{var}}}(y)|\\ &= \sum_{y\in V_{\mathrm{var}}} \sum_{\substack{1\le s\le t\\ y\in Z^s_{\mathrm{var}}}} \frac{|X(a,f)\cap \pi^{-1}_{S_{\mathrm{var}}}(y)|}{d(M_{\mathrm{box}},m)^{r-k-m}}\\ &\quad + O\bigg(\frac{1}{\log N}\,\frac{|X(a,f)\cap \pi^{-1}_{S_{\mathrm{var}}}(y)|}{d(M_{\mathrm{box}},m)^{r-k-m}}\bigg)\\ &= \frac{R|X(a,f)|}{d(M_{\mathrm{box}},m)^{r-k-m}} + O\bigg(\frac{R}{\log N}\,\frac{|X(a,f)|}{d(M_{\mathrm{box}},m)^{r-k-m}}\bigg).\end{aligned}$$

An application of Proposition 9.4 shows that for distinct s and s',

$$[M^\circ(Z^s_{\mathrm{var}})M^\circ(Z^{s'}_{\mathrm{var}}) : L(y,x_1,\dots,x_k)] = d(M_{\mathrm{box}},m)^2,$$

where we use

$$|Z^s_{\mathrm{var}} \cap Z^{s'}_{\mathrm{var}}| \le 1.$$

If $y \in Z^s_{\mathrm{var}} \cap Z^{s'}_{\mathrm{var}}$ for distinct s and s', we infer from Theorem 11.6 that

$$\begin{aligned}&\bigg||X(a,f)\cap \widetilde{Z}^s_{\mathrm{var}} \cap \widetilde{Z}^{s'}_{\mathrm{var}} \cap \pi^{-1}_{S_{\mathrm{var}}}(y)| - \frac{|X(a,f)\cap \pi^{-1}_{S_{\mathrm{var}}}(y)|}{d(M_{\mathrm{box}},m)^{2(r-k-m)}}\bigg|\\ &\qquad\le \frac{A'(l)}{\log N}\,\frac{|X(a,f)\cap \pi^{-1}_{S_{\mathrm{var}}}(y)|}{d(M_{\mathrm{box}},m)^{2(r-k-m)}}\end{aligned}$$

except for a vanishingly small proportion of f's that fit in the error term. Then we can compute the second moment of $\Lambda(x)$ in exactly the same way as we computed the first moment, and this yields

$$\begin{aligned}\sum_{x\in X(a,f)} \Lambda(x)^2 &= \frac{R|X(a,f)|}{d(M_{\mathrm{box}},m)^{r-k-m}} + \frac{(R^2-R)|X(a,f)|}{d(M_{\mathrm{box}},m)^{2(r-k-m)}}\\ &\quad + O\bigg(\frac{R^2}{\log N}\,\frac{|X(a,f)|}{d(M_{\mathrm{box}},m)^{2(r-k-m)}}\bigg)\\ &= \frac{R^2|X(a,f)|}{d(M_{\mathrm{box}},m)^{2(r-k-m)}} + O\bigg(\frac{R^2}{\log N}\,\frac{|X(a,f)|}{d(M_{\mathrm{box}},m)^{2(r-k-m)}}\bigg).\end{aligned}$$

Now use Chebyshev's inequality to finish the proof of the proposition. ∎

Proof of Proposition 13.11. Let P be an element of $\pi_{[r]-S(j_1,j_2)}(\widetilde{Z} \cap X(a, f))$ and make a choice of characters for all primes in P. Then it suffices to prove

$$\frac{1}{W(X, \{i_2(j_1, j_2)\})} \sum_{f \in \mathrm{Ch}(\{i_2(j_1,j_2)\})} \left| \sum_{x \in \widetilde{Z} \cap X(a,f,P,m-1)} F(\mathrm{Art}(x, f, m)) \right|$$
$$\leq \frac{A}{W(X, \{i_2(j_1, j_2)\})} \sum_{f \in \mathrm{Ch}(\{i_2(j_1,j_2)\})} \frac{|\widetilde{Z} \cap X(a, f, P)|}{(\log\log N)^{\frac{c}{m(l^2+l)^m}}}$$

for all P, where $X(a, f, P, m-1)$ and $X(a, f, P)$ are the subsets of $X(a, f, m-1)$ and $X(a, f)$ equal to P on $[r] - S(j_1, j_2)$. Define $Z := \pi_{S_{\mathrm{var}}}(\widetilde{Z})$. Then Proposition 9.3 implies

$$\mathrm{Gal}(M(Z)/M_\circ(Z)) \simeq \mathscr{G}_{S_{\mathrm{var}}}(Z),$$

where the isomorphism sends σ to the map $\bar{x} \mapsto \phi_{S_{\mathrm{var}},\bar{x}}(\sigma)$. This is well-defined by our assumption that Z is well-governed. Furthermore, the Galois group $\mathrm{Gal}(M(Z)/M_\circ(Z))$ is the center of $\mathrm{Gal}(M(Z)/\mathbb{Q})$.

We formally apply Proposition 8.7 with the space $X := Z \times [M_{\mathrm{box}}]$. There is a natural bijection between $\mathscr{G}_S(X)$ and the set of maps g from $[M_{\mathrm{box}}]^l$ to $\mathrm{Gal}(M(Z)/M_\circ(Z))$ with

$$g(i_1, \dots, i_{l-1}, k_1) + g(i_1, \dots, i_{l-1}, k_2) + \cdots + g(i_1, \dots, i_{l-1}, k_l) = g(k_1, \dots, k_l).$$

For a prime $p \in X_{i_2(j_1,j_2)}$, let $\mathfrak{p}$ be the prime ideal in $\mathbb{Z}[\zeta_l]$ above p corresponding to the character $\chi_p^{f(p)}$. Given primes $p_1, \dots, p_{M_{\mathrm{box}}}$ in $X_{i_2(j_1,j_2)}$ one can construct such a function g by defining

$$g(i_1, \dots, i_l) = j_l^{-1}(F(j_1, j_2)) \cdot \pi_{i_2(j_1,j_2)}(w_{2,j_2}) \cdot (\mathrm{Frob}(\mathfrak{p}_{i_1}) + \cdots + \mathrm{Frob}(\mathfrak{p}_{i_l})).$$

Proposition 8.7 gives us a specific function g_0 such that we have equidistribution for all acceptable l-additive systems $\mathfrak{A}$, all choices of $\overline{Y}^\circ_\emptyset$ and all choices of $\widetilde{F}$ with

$$d\,\widetilde{F}(\bar{x}) = g_0(\bar{x}).$$

We are now going to use Theorem 11.6 to partition $X_{i_2(j_1,j_2)}$ into sets of size M_{box} with the property that they all give the function g_0, i.e.

$$g_0(i_1, \dots, i_l) = j_l^{-1}(F(j_1, j_2)) \cdot \pi_{i_2(j_1,j_2)}(w_{2,j_2}) \cdot (\mathrm{Frob}(\mathfrak{p}_{i_1}) + \cdots + \mathrm{Frob}(\mathfrak{p}_{i_l})). \tag{13.8}$$

Now define

$$X_{i_2(j_1,j_2)}(\sigma) = \{p \in X_{i_2(j_1,j_2)} : \mathrm{Frob}(\mathfrak{p}) = \sigma\},$$

where σ is an element of $\mathrm{Gal}(M(Z)L(P \cup Z)/\mathbb{Q})$ that restricts to the element of $\mathrm{Gal}(M^\circ(Z)/\mathbb{Q})$ corresponding to $X_{i_2(j_1,j_2)}(M^\circ(Z))$. Then Theorem 11.6 implies

$$\left| |X_{i_2(j_1,j_2)}(\sigma)| - \frac{|X_{i_2(j_1,j_2)}(M^\circ(Z))|}{l^{(M_{\mathrm{box}}-1)^m} \cdot l^{2r-2k-2m-2}} \right| \leq \frac{A'(l)}{\log N} \frac{|X_{i_2(j_1,j_2)}(M^\circ(Z))|}{l^{(M_{\mathrm{box}}-1)^m} \cdot l^{2r-2k-2m-2}} \tag{13.9}$$

for all but very few f's, where $A'(l)$ is a sufficiently large constant. Now take any p_1. Then given $\mathrm{Frob}(\mathfrak{p}_1)$, there is a unique choice of $\mathrm{Frob}(\mathfrak{p}_2), \dots, \mathrm{Frob}(\mathfrak{p}_{M_{\mathrm{box}}})$ such that (13.8) is satisfied; in fact $\mathrm{Frob}(\mathfrak{p}_2), \dots, \mathrm{Frob}(\mathfrak{p}_{M_{\mathrm{box}}})$ is simply a linear function of $\mathrm{Frob}(\mathfrak{p}_1)$ and the fixed function g_0. From this observation and (13.9), we conclude that $X_{i_2(j_1,j_2)}(M^\circ(Z)L(P \cup Z))$ can be partitioned into sets A of size M_{box} such that (13.8) is valid except for a small set that fits in the error.

If A is such a set of size M_{box}, we have an explicit bijection between A and $[M_{\mathrm{box}}]$ coming from our choice of $p_1, \dots, p_{M_{\mathrm{box}}}$. We first restrict the l-additive system $\mathfrak{A}$ constructed in Lemma 13.10 to $Z \times A$ and then use this bijection to get an l-additive system on $Z \times [M_{\mathrm{box}}]$. This gives the desired equidistribution for F on $Z \times A$, and hence also on $\widetilde{Z}$. ∎

14. Equidistribution in $(\mathcal{G}_{\mathbb{Z}_l[\zeta_l]}, \mu^1_{\mathrm{C.L.}})$

In Proposition 3.4 we have shown that $O^*_{K_\chi} \otimes_{\mathbb{Z}} \mathbb{Z}_l$ is a free module over $\mathbb{Z}_l[\zeta_l]$ of rank 1 for $\chi \in \Gamma_{\mu_l}(\mathbb{Q})$, where the action of $\mathbb{Z}_l[\zeta_l]$ is as usual induced by the Galois action through χ; observe that after tensoring with $\mathbb{Z}_l$ the norm operator acts trivially. Thus, by analogy with real quadratic fields, it is natural to expect that, as χ varies in $\Gamma_{\mu_l}(\mathbb{Q})$, the $\mathbb{Z}_l[\zeta_l]$-module $(1-\zeta_l)\mathrm{Cl}(K_\chi)[(1-\zeta_l)^\infty]$ should equidistribute in the Cohen–Lenstra probability space $(\mathcal{G}_{\mathbb{Z}_l[\zeta_l]}, \mu^1_{\mathrm{C.L.}})$ which is defined as follows.

The set $\mathcal{G}_{\mathbb{Z}_l[\zeta_l]}$ consists of the set of isomorphism classes of $\mathbb{Z}_l[\zeta_l]$-modules with finite cardinality. To each $A \in \mathcal{G}_{\mathbb{Z}_l[\zeta_l]}$ we give weight

$$\mu^1_{\mathrm{C.L.}}(A) := \frac{\eta^1_\infty(l)}{|A| \cdot |\mathrm{Aut}_{\mathbb{Z}_l[\zeta_l]}(A)|},$$

where $\eta^1_\infty(l) := \prod_{i=2}^\infty (1 - 1/l^i)$. Later in this section we will prove that this formula defines a probability measure. The goal of this section is to show that this statistical model is equivalent to the statistical model for the sequence of ranks established in Theorem 1.2. In particular with the material of this section one sees that Theorem 1.2 implies Theorem 1.1.

Let $\mathcal{D}$ be the set of non-increasing functions $f : \mathbb{Z}_{\geq 1} \to \mathbb{Z}_{\geq 0}$ that are eventually 0. Recall that the map

$$\mathrm{rk} : \mathcal{G}_{\mathbb{Z}_l[\zeta_l]} \to \mathcal{D}$$

defined by $A \mapsto (i \mapsto \mathrm{rk}_{(1-\zeta_l)^i} A)$ is a bijection of sets. The next proposition gives the pushforward of $\mu^1_{\mathrm{C.L.}}$ under the bijection rk. For $0 \leq j \leq n$, recall that $P(j|n)$ is the probability that an $n \times (n+1)$ matrix with entries in $\mathbb{F}_l$ has rank $n - j$. Moreover, we set

$$\eta_n(l) := \prod_{i=1}^n \left(1 - \frac{1}{l^i}\right).$$

Then we have the following proposition.

Proposition 14.1. *Let j be a positive integer. Let $i_1 \geq \ldots \geq i_j \geq 0$ be a sequence of integers. Then*

$$\mu^1_{\mathrm{C.L.}}\left(\mathrm{rk}^{-1}(\{f \in \mathcal{D} : f(1) = i_1, \ldots, f(j) = i_j\})\right) = \frac{\eta_\infty(l)}{l^{i_1(i_1+1)}\eta_{i_1}(l)\eta_{i_1+1}(l)} \cdot \prod_{1 \leq k < j} P(i_{k+1}|i_k).$$

The rest of this section is devoted to the proof of Proposition 14.1. First recall that for each $A \in \mathcal{G}_{\mathbb{Z}_l[\zeta_l]}$, the measure $\mu^1_{\mathrm{C.L.}}(A)$ can be obtained as the limit

$$\lim_{N\to\infty} \mu_{\mathrm{Haar}}\left(\left\{(v_1, \ldots, v_{N+1}) \in (\mathbb{Z}_l[\zeta_l]^N)^{N+1} : \frac{\mathbb{Z}_l[\zeta_l]^N}{\langle v_1, \ldots, v_{N+1}\rangle} \simeq A\right\}\right) = \mu^1_{\mathrm{C.L.}}(A).$$

The following argument is essentially due to Friedman and Washington [5], who dealt with the analogous case of $N \times N$ matrices with coefficients in $\mathbb{Z}_l$ corresponding to the case of imaginary quadratic number fields.

For each positive integer N, denote by $\mathcal{L}_{A,N}$ the set of $\mathbb{Z}_l[\zeta_l]$-submodules L of $\mathbb{Z}_l[\zeta_l]^N$ satisfying

$$\frac{\mathbb{Z}_l[\zeta_l]^N}{L} \simeq A.$$

We have

$$\mu_{\mathrm{Haar}}\left(\left\{(v_1, \ldots, v_{N+1}) \in (\mathbb{Z}_l[\zeta_l]^N)^{N+1} : \frac{\mathbb{Z}_l[\zeta_l]^N}{\langle v_1, \ldots, v_{N+1}\rangle} \simeq A\right\}\right) = \sum_{L \in \mathcal{L}_{A,N}} \mu_{\mathrm{Haar}}(\{(v_1, \ldots, v_{N+1}) \in (\mathbb{Z}_l[\zeta_l]^N)^{N+1} : \langle v_1, \ldots, v_{N+1}\rangle = L\}).$$

We further have

$$\mu_{\mathrm{Haar}}(\{(v_1, \ldots, v_{N+1}) \in (\mathbb{Z}_l[\zeta_l]^N)^{N+1} : \langle v_1, \ldots, v_{N+1}\rangle = L\}) = \prod_{i=2}^{N+1}\left(1 - \frac{1}{l^i}\right) \cdot \frac{1}{|A|^{N+1}}$$

and the following simple formula for $|\mathcal{L}_{A,N}|$:

$$|\mathcal{L}_{A,N}| = |\mathrm{Epi}_{\mathbb{Z}_l[\zeta_l]}(\mathbb{Z}_l[\zeta_l]^N, A)|/|\mathrm{Aut}_{\mathbb{Z}_l[\zeta_l]}(A)|.$$

But observe that

$$|\mathrm{Epi}_{\mathbb{Z}_l[\zeta_l]}(\mathbb{Z}_l[\zeta_l]^N, A)|/|A|^N \to 1$$

as N goes to infinity. This gives

$$\lim_{N\to\infty} \mu_{\mathrm{Haar}}\left(\left\{(v_1, \ldots, v_{N+1}) \in (\mathbb{Z}_l[\zeta_l]^N)^{N+1} : \frac{\mathbb{Z}_l[\zeta_l]^N}{\langle v_1, \ldots, v_{N+1}\rangle} \simeq A\right\}\right) = \mu^1_{\mathrm{C.L.}}(A).$$

It is not difficult to show the slightly refined conclusion that the convergence also holds if we take a subset of $\mathcal{G}_{\mathbb{Z}_l[\zeta_l]}$, which thus shows that $\mu^1_{\text{C.L.}}$ is a probability measure. Using this we can show Proposition 14.1 by first computing the pushforward of $\mu^1_{\text{C.L.}}$ by rk at stage N, i.e. the N-th approximation of the pushforward. Sending N to infinity will yield the desired conclusion.

In the notation of Proposition 14.1, let us begin with $j = 1$ and fix an integer $i_1 \geq 0$. Observe that the probability that $f(1) = i_1$ at stage N is given by the probability that the reduction of $v_1, \ldots, v_{N+1}$ modulo $(1-\zeta_l)$ generates a subspace of dimension $N - i_1$. Splitting the probability by the contribution coming from each subspace of dimension $N - i_1$ one gets

$$|\text{subspaces of dimension } N - i_1 \text{ in } \mathbb{F}_l^N| \cdot \frac{\mathbb{P}((w_1, \ldots, w_{N+1}) \in \mathbb{F}_l^{N-i_1} \text{ generate})}{l^{i_1(N+1)}}.$$

This we can rewrite as

$$\frac{\mathbb{P}((w_1, \ldots, w_{N+1}) \in \mathbb{F}_l^{N-i_1} \text{ generate})}{\mathrm{Aut}_{\mathbb{F}_l}(\mathbb{F}_l^{i_1})} \cdot \frac{1}{l^{i_1}} \cdot \frac{|\mathrm{Epi}(\mathbb{F}_l^N, \mathbb{F}_l^{i_1})|}{l^{i_1 N}}.$$

Again the factor $\frac{|\mathrm{Epi}(\mathbb{F}_l^N, \mathbb{F}_l^{i_1})|}{l^{i_1 N}}$ approaches 1 as N goes to infinity. Moreover

$$\mathbb{P}((w_1, \ldots, w_{N+1}) \in \mathbb{F}_l^{N-i_1} \text{ generate}) = \prod_{i=i_1+2}^{N+1} \left(1 - \frac{1}{l^i}\right).$$

Plugging in and sending N to infinity yields the case $j = 1$ for Proposition 14.1. We now continue by induction to compute the N-th approximation for any $N > i_1$. Observe that whether $(v_1, \ldots, v_{N+1})$ is giving an A with $f(1) = i_1, \ldots, f(j) = i_j$ can be decided completely by the image of $(v_1, \ldots, v_{N+1})$ modulo $(1-\zeta_l)^j$. Hence we proceed to show that if we *fix* the image of $(v_1, \ldots, v_{N+1})$ modulo $(1-\zeta_l)^j$, then the N-th probability that $f(j+1) = i_{j+1}$, conditional on the image modulo $(1-\zeta_l)^j$ being fixed, is always $P(i_{j+1}|i_j)$. From this the desired conclusion follows immediately.

Since the image modulo $(1-\zeta_l)^j$ has been fixed, we fix a subset $\mathcal{B}$ of $[N+1]$ such that $\{v_i\}_{i\in\mathcal{B}}$ forms a minimal set of generators for the image modulo $(1-\zeta_l)^j$. By construction the set $\mathcal{B}$ has size $N - i_j$. By multiplying each element of $\mathcal{B}$ with suitable powers of $1-\zeta_l$ we obtain a subset of

$$V := \frac{(1-\zeta_l)^j \mathbb{Z}_l[\zeta_l]^N}{(1-\zeta_l)^{j+1} \mathbb{Z}_l[\zeta_l]^N}$$

generating a space of dimension $N - i_j$, which we call V'. Note that V' is fixed as $(v_1, \ldots, v_{N+1})$ varies among vectors with fixed image modulo $(1-\zeta_l)^j$. Since $\mathcal{B}$ is a minimal set of generators modulo $(1-\zeta_l)^j$, we see that there is a natural map F that sends the $i_j + 1$ vectors $(v_i)_{i\notin\mathcal{B}}$ into the i_j-dimensional $\mathbb{F}_l$-vector space

$$\frac{\frac{(1-\zeta_l)^j \mathbb{Z}_l[\zeta_l]^N}{(1-\zeta_l)^{j+1} \mathbb{Z}_l[\zeta_l]^N}}{V'}.$$

It is seen at once that if the image of F spans a space of dimension k, then the resulting A will satisfy $f(j+1) = i_j - k$. Moreover, it is easy to see that each vector is obtained equally often through F. Thus we obtain the desired conclusion.

Appendix A. Cyclic algebras

In this small appendix we collect several basic facts that are used in this paper coming from local and global class field theory along with some more general facts about cyclic algebras over general fields.

Let K be any field. Denote by K^{sep} a fixed separable closure of K and by G_K the group of K-algebra automorphisms of K^{sep}. Let $\chi : G_K \to \mathbb{C}^*$ be a continuous character and define

$$K(\chi) := (K^{\mathrm{sep}})^{\ker(\chi)}.$$

Let n be the degree of $K(\chi)$ over K and let θ be in K^*. Following the notation from [26, Ch. 9] we denote by $\{\chi, \theta\}$ the twisted polynomial ring $K(\chi)\langle\beta\rangle$ with the relations

$$\beta^n = \theta \quad \text{and} \quad \beta\lambda\beta^{-1} = \chi^{-1}\left(\exp\left(\frac{2\pi \mathrm{i}}{\mathrm{ord}(\chi)}\right)\right)(\lambda),$$

which is a cyclic algebra. Denote by Φ_χ the unique map from G_K to $\mathbb{R}$ such that $\mathrm{Im}(\Phi_\chi) \subseteq [0, 1)$ and

$$\exp(2\pi \mathrm{i} \cdot \Phi_\chi) = \chi.$$

The map Φ_χ is locally constant with values in the set

$$\left\{0, \frac{1}{\mathrm{ord}(\chi)}, \ldots, \frac{\mathrm{ord}(\chi) - 1}{\mathrm{ord}(\chi)}\right\}.$$

Denote by $\widetilde{\Phi_\chi}$ the map $\mathrm{ord}(\chi) \cdot \Phi_\chi$. The map $\widetilde{\Phi_\chi}$ is locally constant with values in $\{0, \ldots, \mathrm{ord}(\chi) - 1\}$. Observe that since χ is a character, for each $\sigma, \tau \in G_K$ the element

$$\Phi_\chi(\sigma) + \Phi_\chi(\tau) - \Phi_\chi(\sigma\tau)$$

is an integer. This allows us to define a 2-cocycle $h_{\{\chi,\tau\}}$ of G_K with values in K^* by the formula

$$(\sigma, \tau) \mapsto \theta^{\Phi_\chi(\sigma)+\Phi_\chi(\tau)-\Phi_\chi(\sigma\tau)}.$$

If we could separate the three values on the exponent of $h_{\{\chi,\tau\}}$ we would obtain trivially a coboundary, for this reason we already know that the above formula defines a 2-cocycle. But the three terms in general cannot be separated, since it is only the total sum that is an integer. This observation will be useful in Proposition A.3. The reason why we introduced this particular 2-cocycle is that the class of $h_{\{\chi,\theta\}}$ in Br_K is precisely the class of $\{\chi, \theta\}$. This fact is established in [26]. Recall the following fundamental fact from [26, p. 223].

Proposition A.1. *Let K be a local field. There is a unique isomorphism*

$$\eta_K : \mathrm{Br}_K \to \mu_\infty(\mathbb{C})$$

such that for any continuous unramified character $\chi : G_K \to \mathbb{C}^*$ and any uniformizer π of K we have

$$\eta_K(\{\chi, \pi\}) = \chi(\mathrm{Frob}_K\, \pi).$$

The map η_K actually equals $\exp(2\pi \mathrm{i} \cdot \mathrm{inv}_K)$ (for the definition of inv_K see [21]). We shall use η_K instead of inv since our main reference is [26]. It is defined also for $K = \mathbb{R}$ or $K = \mathbb{C}$ being trivial in the latter case and being the unique isomorphism between $\mathrm{Br}_\mathbb{R}$ and $\langle -1 \rangle$ in the former. Recall the following reformulation of Hilbert's reciprocity law, whose proof can be found in [26, p. 255].

Proposition A.2. (Hilbert reciprocity law). *Let K be a number field and let $\alpha \in \mathrm{Br}_K$. Then $\eta_{K_v}(\alpha)$ is trivial for all but finitely many values of $v \in \Omega_{K_v}$. It is trivial at all places if and only if α itself is trivial. Moreover,*

$$\prod_{v \in \Omega_K} \eta_{K_v}(\alpha) = 1.$$

Let l be an odd prime. We now recall a relation between cyclic algebras and cup products in case there are l-th roots of unity. We shall confine ourselves to classes killed by l, since this is the relevant case in our application. For more general results the reader can consult [8]. For a field K provided with a distinguished generator ζ_l for $\mu_l(K)$ we shall use precisely the same symbolic formulas introduced in Section 2.2. Moreover, for such a K and for an element $\theta \in K^*$ we denote by $\chi_\theta : G_K \to \mathbb{F}_l$ the unique continuous character such that for each $\beta \in K^{\mathrm{sep}}$ with $\beta^l = \theta$ we have

$$\sigma(\beta) = (j_l \circ \chi_\theta(\sigma))\beta.$$

In what follows, when we consider the cup product, the trivial Galois modules $\mathbb{F}_l \otimes \mathbb{F}_l$ and $\mathbb{F}_l$ are identified with the isomorphism $a \otimes b \mapsto a \cdot b$. In particular the cup product of two characters χ_1, χ_2 in $\mathbb{F}_l$ is literally just the product map $(\sigma, \tau) \mapsto \chi_1(\sigma)\chi_2(\tau)$.

Proposition A.3. *Suppose K is equipped with an element ζ_l of multiplicative order l. Let χ be a continuous homomorphism from G_K to $\mathbb{F}_l$. Let θ be in K^*. Then in Br_K we have*

$$h_{\{j_l \circ \chi, \theta\}} = j_l \circ (\chi \cup \chi_\theta).$$

Proof. We divide the 2-cocycle $h_{\{j_l \circ \chi, \theta\}}$ by the 1-coboundary

$$(\sigma, \tau) \mapsto \frac{\sigma(\beta^{\widetilde{\Phi}_\chi(\tau)})}{\beta^{\widetilde{\Phi}_\chi(\sigma\tau) - \widetilde{\Phi}_\chi(\sigma)}},$$

where β is any element of K^{sep} with $\beta^l = \theta$ and $\widetilde{\Phi}_\chi$ is shorthand for $\widetilde{\Phi_{j_l \circ \chi}}$. Now using the formula

$$\sigma(\beta^{\widetilde{\Phi}_\chi(\tau)}) = \zeta_l^{\chi_\theta(\sigma)\chi(\tau)} \beta^{\widetilde{\Phi}_\chi(\tau)},$$

we find that the cocycle we are considering is $j_l \circ (-\chi_\theta \cup \chi)$. Recalling that, in cohomology, the cup is antisymmetric we conclude the proof immediately. ∎

We in particular deduce the following corollary, which holds in greater generality (see [26]), than for cyclic degree l characters and without any restriction on the characteristic. However this generality is the one we need and on the other hand we propose an unusual argument based only on the material of Section 4.

Corollary A.4. *Suppose* $\mathrm{char}(K) \neq l$. *Let* χ *be a continuous character from* G_K *to* $\mathbb{F}_l$. *Let* θ *be in* K^*. *Then* $\{\chi, \theta\}$ *is trivial in* Br_K *if and only if* θ *is a norm from* $K(\chi)$.

Proof. We show how to reduce to the case that $\mu_l(K) \neq \{1\}$. Once that is done, we reach the desired conclusion by an application of Proposition 4.12. Since $\mathrm{char}(K) \neq l$ we can adjoin to K an element ζ_l from K^{sep} with multiplicative order l. Thanks to the co-restriction map, we see that $\{\chi, \theta\}$ is trivial in $\mathrm{Br}_{K(\zeta_l)}$ if and only if $\{\chi, \theta\}$ is trivial in Br_K. Here we use the fact that $[K(\zeta_l) : K]$ divides $l - 1$ and hence is coprime to l.

It remains to prove that θ is a norm from $K(\zeta_l)(\chi)$ if and only if θ is a norm from $K(\chi)$. Suppose that θ is a norm from $K(\zeta_l)(\chi)$, say

$$\theta = N_{K(\zeta_l)(\chi)/K(\zeta_l)}(\gamma)$$

for some $\gamma \in K(\zeta_l)(\chi)$. Then we see that

$$\theta^{[K(\zeta_l):K]} = N_{K(\zeta_l)(\chi)/K}(\gamma)$$

and hence

$$\theta^{[K(\zeta_l):K]} = N_{K(\chi)/K}(N_{K(\zeta_l)(\chi)/K(\chi)}(\gamma)).$$

Using once more that $[K(\zeta_l) : K] \mid l - 1$, we conclude that θ is a norm from $K(\chi)$. The other direction is more general. Now suppose $\theta = N_{K(\chi)/K}(\gamma)$ for some $\gamma \in K(\chi)$. Observe that $\mathrm{Gal}(K(\zeta_l)(\chi)/K(\zeta_l))$ maps injectively into a normal subgroup of $\mathrm{Gal}(K(\chi)/K)$. Fix a set $\mathcal{S}$ of representatives for the quotient of this normal subgroup. Write $\gamma' := \prod_{g \in \mathcal{S}} g(\gamma)$. Then

$$N_{K(\zeta_l)(\chi)/K(\zeta_l)}(\gamma') = \prod_{h \in \mathrm{Gal}(K(\zeta_l)(\chi)/K(\zeta_l))} h(\gamma') = \prod_{g \in \mathrm{Gal}(K(\chi)/K)} g(\gamma) = \theta.$$

This shows the other direction. ∎

We end this section by recalling how the field of definition of the character $\chi_q \in \Gamma_{\mu_l}(\mathbb{Q})$, introduced in Section 2.2, looks locally at q. We begin by recalling the following basic fact.

Proposition A.5. *Let* K *be a local field and* d *a positive integer coprime to the size of the residue field of* K. *Let* $f(x)$ *be a degree* d *Eisenstein polynomial over* K. *Then*

$$K[x]/f(x) \simeq_{K\text{-alg}} K\left[\sqrt[d]{-f(0)}\right].$$

Proof. Without loss of generality we may assume that f is monic. In $K[x]/f(x)$ we have

$$x^d = -f(0) + \sum_{i=1}^{d-1} a_i x^i,$$

where the $-a_i$ are the various coefficients of $f(x)$. Dividing out by $-f(0)$ we obtain

$$\frac{x^d}{-f(0)} = 1 + \sum_{i=1}^{d-1} \frac{a_i}{-f(0)} x^i.$$

Since $f(x)$ is Eisenstein, $\frac{a_i}{-f(0)}$ is still integral. Hence $\frac{a_i}{-f(0)} x^i$ is in the maximal ideal of $O_{K[x]/f(x)}$ for each i between 1 and $d-1$. This implies

$$1 + \sum_{i=1}^{d-1} \frac{a_i}{-f(0)} x^i \in U_1(K[x]/f(x)).$$

The topological group $U_1(K[x]/f(x))$ is a $\mathbb{Z}_l$-module, where l is the residue characteristic of K. In particular it is d-divisible, since d is coprime to l. Therefore we conclude that

$$1 + \sum_{i=1}^{d-1} \frac{a_i}{-f(0)} x^i$$

is a d-th power and hence also $-f(0)$ is a d-th power, since it is the ratio of two d-th powers. Finally, the polynomial $T^d + f(0)$ is again Eisenstein; so if we pick $\beta \in K[x]/f(x)$ with $\beta^d = -f(0)$ we see that $K(\beta) = K[x]/f(x)$ and

$$K(\beta) \simeq_{K\text{-alg}} K\big[\sqrt[d]{-f(0)}\big],$$

which is the desired isomorphism. ∎

Therefore we conclude the following fact, which also follows from class field theory as we shall see in the second proof.

Corollary A.6. *Let q be a prime that is* 1 *modulo l and let χ_q be as in Section* 2.2. *Then*

$$(K_{\chi_q})_{\mathrm{Up}_{K_{\chi_q}}(q)} \simeq_{\mathbb{Q}_q\text{-alg}} \mathbb{Q}_q(\sqrt[l]{q}).$$

First proof. If we denote by $\Phi_q(T)$ the q-th cyclotomic polynomial, we observe that $\Phi_q(T+1)$ is Eisenstein of degree $q-1$. Moreover, if evaluated at 0, $\Phi_q(T+1)$ is equal to q. Therefore we conclude by Proposition A.5 that $\mathbb{Q}_q(\zeta_q)$ completed at $(1-\zeta_q)$ is the extension $\mathbb{Q}_q(\sqrt[q-1]{-q})$. In particular its unique degree l subextension is given by $\mathbb{Q}_q(\sqrt[l]{-q})$. Since l is odd, the minus sign is irrelevant and the conclusion follows. ∎

Second proof. Still by looking at the polynomial $\Phi_q(T)$, we see that $-q$ is a norm locally from $\mathbb{Q}_q(\zeta_q)$. Hence $-q$ is also a norm from the degree l subextension, which is of odd degree, so q is a norm from the degree l subextension. On the other hand, since q is 1 modulo l, we find that the extension $\mathbb{Q}_q(\sqrt[l]{q})$ is cyclic of degree l. Moreover, taking the norm of $-\sqrt[l]{q}$ we see that q is a norm also in this extension. It is not difficult to deduce from this that the two fields in the isomorphism have the same norm group. The conclusion follows from local class field theory. ∎

Acknowledgments. We are very grateful to Alexander Smith for answering many questions regarding his paper. We would also like to thank Stephanie Chan, Jan-Hendrik Evertse, Hendrik Lenstra, Djordjo Milovic, Peter Stevenhagen and Mark Watkins for various useful discussions. The first author greatly appreciates the hospitality of the Max Planck Institute for Mathematics during two visits.

References

[1] Cohn, H., Lagarias, J. C.: On the existence of fields governing the 2-invariants of the class-group of $\mathbf{Q}(\sqrt{dp})$ as p varies. Math. Comp. **41**, 711–730 (1983) Zbl 0523.12002 MR 717716

[2] Davenport, H., Heilbronn, H.: On the density of discriminants of cubic fields. II. Proc. Roy. Soc. London Ser. A **322**, 405–420 (1971) Zbl 0212.08101 MR 491593

[3] Fouvry, E., Klüners, J.: Cohen–Lenstra heuristics of quadratic number fields. In: Algorithmic Number Theory, Lecture Notes in Comput. Sci. 4076, Springer, Berlin, 40–55 (2006) Zbl 1143.11352 MR 2282914

[4] Fouvry, E., Klüners, J.: On the 4-rank of class groups of quadratic number fields. Invent. Math. **167**, 455–513 (2007) Zbl 1126.11062 MR 2276261

[5] Friedman, E., Washington, L. C.: On the distribution of divisor class groups of curves over a finite field. In: Théorie des nombres (Québec, PQ, 1987), de Gruyter, Berlin, 227–239 (1989) Zbl 0693.12013 MR 1024565

[6] Gerth, F., III: The 4-class ranks of quadratic fields. Invent. Math. **77**, 489–515 (1984) Zbl 0533.12004 MR 759260

[7] Gerth, F., III: Densities for certain l-ranks in cyclic fields of degree l^n. Compos. Math. **60**, 295–322 (1986) MR 869105

[8] Gille, P., Szamuely, T.: Central Simple Algebras and Galois Cohomology. Cambridge Stud. Adv. Math. 101, Cambridge Univ. Press, Cambridge (2006) Zbl 1137.12001 MR 2266528

[9] Granville, A.: Prime divisors are Poisson distributed. Int. J. Number Theory **3**, 1–18 (2007) Zbl 1118.11005 MR 2310491

[10] Hardy, G. H., Ramanujan, S.: The normal number of prime factors of a number n. Quart. J. Math. **48**, 76–92 (1917)

[11] Heath-Brown, D. R.: The size of Selmer groups for the congruent number problem. II. Invent. Math. **118**, 331–370 (1994) Zbl 0815.11032 MR 1292115

[12] Klys, J.: The distribution of p-torsion in degree p cyclic fields. Algebra Number Theory **14**, 815–854 (2020) Zbl 1450.11118 MR 4114057

[13] Koymans, P., Milovic, D.: On the 16-rank of class groups of $\mathbb{Q}(\sqrt{-2p})$ for primes $p \equiv 1$ mod 4. Int. Math. Res. Notices **2019**, 7406–7427 Zbl 07154613 MR 4039017

[14] Koymans, P., Milovic, D. Z.: Spins of prime ideals and the negative Pell equation $x^2 - 2py^2 = -1$. Compos. Math. **155**, 100–125 (2019) Zbl 1443.11232 MR 3880026

[15] Koymans, P., Milovic, D. Z.: Joint distribution of spins. arXiv:1809.09597 (2018)

[16] Lagarias, J. C., Odlyzko, A. M.: Effective versions of the Chebotarev density theorem. In: Algebraic Number Fields: L-functions and Galois Properties (Durham, 1975), Academic Press, London, 409–464 (1977) Zbl 0362.12011 MR 0447191

[17] Li, C.: 2-Selmer groups, 2-class groups and rational points on elliptic curves. Trans. Amer. Math. Soc. **371**, 4631–4653 (2019) Zbl 1448.11108 MR 3934463

[18] Michailov, I. M.: Four non-abelian groups of order p^4 as Galois groups. J. Algebra **307**, 287–299 (2007) Zbl 1171.12004 MR 2278055

[19] Milovic, D.: On the 16-rank of class groups of $\mathbb{Q}(\sqrt{-8p})$ for $p \equiv -1 \bmod 4$. Geom. Funct. Anal. **27**, 973–1016 (2017) Zbl 1381.11103 MR 3678506

[20] Pollack, P.: The smallest inert prime in a cyclic number field of prime degree. Math. Res. Lett. **20**, 163–179 (2013) Zbl 1300.11009 MR 3126729

[21] Serre, J.-P.: Local Fields. Grad. Texts in Math. 67, Springer, New York (1979) Zbl 0423.12016 MR 554237

[22] Smith, A.: 2^∞-Selmer groups, 2^∞-class groups, and Goldfeld's conjecture. arXiv:1702.02325 (2017)

[23] Stevenhagen, P.: Class groups and governing fields. PhD thesis, Univ. of California, Berkeley (1988) Zbl 0701.11056 MR 2636917

[24] Stevenhagen, P.: Rédei-matrices and applications. In: Number Theory (Paris, 1992–1993), London Math. Soc. Lecture Note Ser. 215, Cambridge Univ. Press, Cambridge, 245–259 (1995) Zbl 0830.11039 MR 1345183

[25] Weibel, C. A.: An Introduction to Homological Algebra. Cambridge Stud. Adv. Math. 38, Cambridge Univ. Press, Cambridge (1994) Zbl 0797.18001 MR 1269324

[26] Weil, A.: Basic Number Theory. Classics Math., Springer, Berlin (1995) Zbl 0823.11001 MR 1344916