Homological stability for Hurwitz spaces and the Cohen-Lenstra conjecture over function fields

By Jordan S. Ellenberg, Akshay Venkatesh, and Craig Westerland

Dedicated to Barry Mazur on the occasion of his 75th birthday

Abstract

We prove a homological stabilization theorem for Hurwitz spaces: moduli spaces of branched covers of the complex projective line. This has the following arithmetic consequence: let $\ell > 2$ be prime and A a finite abelian ℓ-group. Then there exists $Q = Q(A)$ such that, for q greater than Q, a positive fraction of quadratic extensions of $\mathbb{F}_q(t)$ have the ℓ-part of their class group isomorphic to A.

Contents

1. Introduction

1.1. *The Cohen-Lenstra heuristics.* Experimental evidence shows clearly that class groups of number fields display interesting biases in their distribution. For instance, class groups of quadratic imaginary fields are much more likely to contain a factor $\mathbb{Z}/9\mathbb{Z}$ than a factor $\mathbb{Z}/3\mathbb{Z} \times \mathbb{Z}/3\mathbb{Z}$. Motivated by this and other examples, Cohen and Lenstra conjectured in [13] that a particular

---

finite abelian group should occur as the class group of a quadratic imaginary field with frequency inversely proportional to its number of automorphisms.

This leads, for instance, to the prediction that the probability that a quadratic imaginary field has class number indivisible by 3 is

$$1 - \prod(1 - 3^{-i}) \sim 0.440 \cdots .$$

The initial motivation for the present paper was to study the Cohen-Lenstra heuristics over function fields, i.e., finite extensions of $\mathbb{F}_q(t)$. The result quoted in the abstract can be stated more quantitatively as follows:

1.2. THEOREM. *Let $\ell > 2$ be prime and A a finite abelian ℓ-group. Write δ^+ (resp. δ^-) for the upper density (resp. lower density) of imaginary[1] quadratic extensions of $\mathbb{F}_q(t)$ for which the ℓ-part of the class group is isomorphic to A. Then $\delta^+(q)$ and $\delta^-(q)$ converge, as $q \to \infty$ with $q \neq 1 \pmod{\ell}$, to $\frac{\prod_{i \geq 1}(1-\ell^{-i})}{|\operatorname{Aut}(A)|}$.*

This is a corollary to Proposition 8.3 and Theorem 8.8. When $q = 1$ (mod ℓ), the method of proof still works; for any fixed positive ℓ-valuation of $q - 1$, the proof yields a distribution that differs from the Cohen-Lenstra distribution. (This is related to Malle's recent observation in [39] that the Cohen-Lenstra heuristics require modification when extra roots of unity are present in the base field.) The description of this distribution in many cases is carried out in the Ph.D. thesis of Garton [30].

In particular, for $q > Q_0(\ell)$, a positive fraction of imaginary quadratic extensions of $\mathbb{F}_q(t)$ have class number divisible by ℓ, and a positive fraction have class number indivisible by ℓ. The infinitude of quadratic extensions of $\mathbb{F}_q(t)$ with class number divisible by ℓ was previously known ([17], [9]) as was the corresponding result for indivisibility by ℓ [29], but in both cases without a positive proportion.

In a different direction, corresponding questions are understood if one studies quadratic field extensions of $\mathbb{F}_q(t)$ with *fixed discriminant degree* and lets $q \to \infty$; see [2], [49].

The essential ingredient in the proof of Theorem 1.2 is, perhaps surprisingly, a theorem in topology — more precisely, a result on stable homology of Hurwitz spaces.

1.3. *Hurwitz spaces.* A *Hurwitz space* is a moduli space for G-covers of the punctured complex plane, where G is a finite group. A thorough definition of these spaces will be given in Section 2; here we content ourselves with a brief description.

_____

[1]By "imaginary" we mean "ramified at ∞."

Hurwitz spaces have vanishing higher homotopy groups; each component has fundamental group isomorphic to a subgroup of the Artin braid group B_n. The group B_n is generated by elements $\sigma_i, 1 \leq i \leq n - 1$, subject to the relations

(1.3.1) $$\begin{cases} \sigma_i \sigma_{i+1} \sigma_i = \sigma_{i+1} \sigma_i \sigma_{i+1}, & 1 \leq i \leq n - 2, \\ \sigma_i \sigma_j = \sigma_j \sigma_i, & |i - j| \geq 2. \end{cases}$$

A Hurwitz space can also be seen as the space of complex points of a *Hurwitz scheme* parametrizing branched covers of $\mathbb{A}^1$. Consequently, the study of Hurwitz spaces lies at the interface of algebraic geometry, topology, and combinatorial group theory.

The majority of the present paper involves only the topology of the Hurwitz space, not its algebro-geometric aspects (e.g., its definition as a scheme over a ring of integers). We therefore start with a purely topological definition of Hurwitz space, in which we replace the complex plane by the unit disc $D = \{(x, y) \in \mathbb{R}^2 : x^2 + y^2 \leq 1\}$.

The Eilenberg-MacLane space $K(B_n, 1)$ has the homotopy type of the configuration space $\mathrm{Conf}_n D$, which parametrizes configurations of n (distinct, unlabeled) points in the interior of the disc. Fixing a point $*$ on the boundary of D, we define the *Hurwitz space* $\mathrm{Hur}_{G,n}$ to be the covering space of $\mathrm{Conf}_n D$ whose fiber above $\{P_1, \dots, P_n\}$ is the set of homomorphisms

$$\pi_1(D - \{P_1, \dots, P_n\}, *) \to G.$$

If $c \subset G$ is a conjugacy class, we denote by $\mathrm{Hur}_{G,n}^c$ the open and closed subspace of $\mathrm{Hur}_{G,n}$ whose fiber over a point of $\mathrm{Conf}_n D$ is the set of homomorphisms sending a loop around each P_j $(1 \leq j \leq n)$ to the conjugacy class c.

The homotopy type of $\mathrm{Hur}_{G,n}$ is then that of the Borel construction $EB_n \times_{B_n} G^n$, where B_n acts on G^n through the *braiding action*

(1.3.2) $$\sigma_j : (g_1, \dots, g_n) \mapsto (g_1, \dots, g_{j-2}, g_{j-1}, g_j g_{j+1} g_j^{-1}, g_j, g_{j+2} \cdots).$$

Similarly, the homotopy type of $\mathrm{Hur}_{G,n}^c$ is that of $EB_n \times_{B_n} c^n$.

1.4. *Stability of homology.* The Hurwitz space is evidently not connected; for example, the braid group action (1.3.2) preserves the subset of c^n consisting of n-tuples with *full monodromy*, i.e., those whose elements generate the whole group G.

Hurwitz proved in [36] that, when $G = S_d$ (the symmetric group on d letters) and c is the conjugacy class of transpositions, the orbits of the B_n-action on

$$\{\mathbf{g} \in c^n : \mathbf{g} \text{ has full monodromy}\}$$

are determined by the *boundary monodromy*; for $\mathbf{g} = (g_1, \dots, g_n)$, this is the product $g_1 \cdots g_n \in S_d$. In geometric terms, the subspace $\mathrm{CHur}_{G,n}^c$ of $\mathrm{Hur}_{G,n}^c$, comprising covers with full monodromy, decomposes into a union of subspaces

indexed by the boundary monodromy of a cover, each of which is connected[2] for all sufficiently large n. This result was used by Severi to establish that the moduli space $\mathcal{M}_g$ of curves of genus g is connected.

By contrast with Hurwitz's connectivity result — which we may think of as a statement about homology in degree zero — very little is known about the higher homology of $\mathrm{Hur}^c_{G,n}$ or $\mathrm{CHur}^c_{G,n}$. The main theorem of this paper is the following stabilization result for the homology of Hurwitz spaces. We write $b_p(X)$ for $\dim H_p(X, \mathbb{Q})$, the pth Betti number of a space X.

THEOREM. *Let G be a finite group and $c \subset G$ a conjugacy class such that*
- *c generates G;*
- *(nonsplitting) for any subgroup $H \leqslant G$, the intersection of c with H is either empty or a conjugacy class of H.*

Then there exist integers $A, B, D > 0$ such that $b_p(\mathrm{Hur}^c_{G,n}) = b_p(\mathrm{Hur}^c_{G,n+D})$ whenever $n \geq Ap + B$.

This theorem is proved as Theorem 6.1 below, with constants A, B, and D that are explicitly computable in terms of the combinatorics of G and c. It is the key input in the proof of Theorem 1.2.

We remark that not even the case $p = 0$ is wholly obvious. Indeed, it is false without the "nonsplitting" condition — for instance, if $G = S_4$ and c is the conjugacy class of transpositions, $b_0(\mathrm{Hur}^c_{G,n})$ has rank at least $n + 1$ coming from the components corresponding to $\mathbf{g} = (12)^i(34)^{n-i}$ for $i = 0, \ldots, n$. Indeed, the nonsplitting enters into our argument *only* to guarantee the validity of Lemma 3.5, whose concern is precisely the stability of the set of connected components, i.e., the $p = 0$ case of homological stability. When (G, c) is *not* nonsplitting, the number of components of $\mathrm{Hur}^c_{G,n}$ always grows without bound.

Unfortunately, the nonsplitting condition is very strong — for instance, it is not satisfied for the case considered by Hurwitz ($G = S_d$ and c the transpositions in S_d) unless $d = 3$. Fortunately, it is satisfied in the cases pertinent to the Cohen-Lenstra heuristics.

In a sense, the fact that the spaces $\mathrm{Hur}^c_{G,n}$ are in general disconnected is one of the central difficulties that is overcome in this paper. It also marks a difference between our result and some (but not all) other results about stable homology, which we recall in Section 1.6 below.

Given this theorem, it is natural to ask whether the stable homology of Hurwitz spaces can be described explicitly. Even the case $p = 0$ (the description of the connected components of $\mathrm{Hur}^c_{G,n}$ for $n \gg 1$) is not obvious; the answer

---

[2]For analogous results for groups other than S_d, we refer the reader to the appendix of [27].

is governed by an argument of Conway, Parker, Fried, and Völklein [27], the ideas of which we make crucial use of in this paper.

In a sequel to the present paper, we will discuss the stable homology of Hurwitz spaces for $p > 0$. As an example of the kind of results we expect, we propose the following conjectural generalization of Hurwitz's theorem to higher homology:

1.5. CONJECTURE. *Suppose G is a symmetric group on more than two letters and c is the conjugacy class of transpositions. Then for any $i \geq 0$, the map*

$$\mathrm{CHur}^c_{G,n} \to \mathrm{Conf}_n D,$$

when restricted to a single component of the domain, induces an isomorphism of rational homology groups in degree i for sufficiently large n.

(Recall, moreover, that so long as $n \geq 2$, the rational homology groups H_j of $\mathrm{Conf}_n(D)$ vanish for $j > 1$ and are one-dimensional for $j = 0, 1$.) This conjecture is motivated by — and implies a form of — Malle's conjecture over function fields, which is to say that both the upper limit δ^+ and the lower limit δ^- in Theorem 1.2 are equal to 1.

1.6. *Some context.* There is already a large body of work in topology concerning homology stabilization for certain "geometrically natural" sequences of manifolds with increasing dimensions. Examples include

(1) the configuration space Conf_n of n points in the plane [12], [4];
(2) the moduli space $\mathcal{M}_g$ of smooth projective curves of genus g [34] (more precisely, the moduli stack; its homology is taken to be orbifold homology);
(3) classifying spaces of arithmetic groups, e.g., the space $\mathrm{BSL}(n, \mathbb{Z})$ [43], [7];
(4) The space of holomorphic mappings $\mathrm{Maps}^d(\Sigma, X)$ of degree d from a Riemann surface Σ to a suitable projective variety X [47], [8].

The Hurwitz spaces $\mathrm{Hur}_{G,n}$ have features in common with all of these examples. On the one hand, their individual components are Eilenberg-MacLane spaces of type $K(\pi, 1)$, as are the first three examples; in such cases, homological stability reduces to a question about group homology, for which there are standard techniques (see Section 1.7).

On the other hand, we may also see $\mathrm{Hur}_{G,n}$ as parametrizing maps from a certain *orbifold* — namely, a sphere on which n points have finite cyclic inertia group — to the classifying space BG, thereby relating it to the fourth example.[3]

---

[3] We learned this point of view from Abramovich, Corti, and Vistoli [1], who use it to define an algebraic compactification of Hurwitz space as a space of *stable* maps from orbifolds to BG.

The results of type 4 in the existing literature require the hypothesis that X is simply connected. The classifying space BG is, of course, not simply connected; this has the effect that the spaces $\mathrm{Hur}_{G,n}$ we consider are typically not connected. This feature turns out to be the source of all the technical difficulty in our paper.

1.7. *The proof of homological stability.* Our method to prove homological stability of Hurwitz spaces is based on the following (by now, standard) setup: Suppose that we are given a sequence $G_1 \subset G_2 \subset \cdots$ of groups and, for each n, a highly connected G_n-simplicial complex X_n, such that the stabilizer of an i-simplex in X_n is precisely G_{n-i-1}. *Then the inclusions $G_n \to G_{n+1}$ tend to induce group homology isomorphisms.* We refer to a paper of Hatcher and Wahl [35] for precise statements of this type.

In our context, the pertinent complex is related to the work of Harer [34] on the homology of the moduli space of curves.

Throughout the method, however, the fact that the spaces $\mathrm{Hur}_{G,n}$ need not be connected proves a difficulty. To handle this, we equip all the higher homology groups with structures of module over the graded ring R formed from the connected components of the Hurwitz spaces. We are then able to reduce all the difficulties to *purely homological* questions about R, which are settled in Section 4.

For the arithmetic applications, it is not sufficient to prove homological stability of Hurwitz spaces; we need the *a priori* stronger statement of homological stability for Hurwitz schemes, moduli schemes over $\mathrm{Spec}\,\mathbb{Z}[\frac{1}{|G|}]$ whose complex points are isomorphic to $\mathrm{Hur}_{G,n}^c$. This requires comparing the cohomology of generic and special fibers of the Hurwitz scheme, which is carried out in Section 7.

1.8. *Analytic number theory over function fields.* Many questions in analytic number theory over $\mathbb{Z}$, when transposed to a function field setting, become questions of the following form:

(1.8.1)   Understand the asymptotics of $|X_n(\mathbb{F}_q)|$, as $n \to \infty$,

where X_n is an algebraic variety over $\mathbb{F}_q$ of dimension growing with n. For example, our analysis of the Cohen-Lenstra heuristics is based on the study of this question for X_n a Hurwitz scheme. We discuss some other examples in Section 1.9.

The philosophy driving this paper can be summed up in the following slogan:

The quantity $|X_n(\mathbb{F}_q)|q^{-\dim X_n}$ should be expected to approach a limit as $n \to \infty$ precisely when the varieties X_n have stable homology.

Of course, one can certainly construct a sequence of varieties $\{X_n\}$ so that $|X_n(\mathbb{F}_q)|q^{-\dim X_n}$ approaches a limit but the homology of X_n is not stable; the slogan is meant to apply just when X_n is a "natural" sequence of moduli spaces.

We now explain how one direction of the above slogan can be demonstrated in practice. The Grothendieck-Lefschetz fixed point formula expresses $|X_n(\mathbb{F}_q)|$ in terms of the action of the Frobenius upon the compactly supported (étale) cohomology of X_n:

(1.8.2)
$$|X_n(\mathbb{F}_q)| = \sum_j (-1)^j \left(\text{trace of Frobenius acting on } H^j_{\text{c,ét}}(X_n \times_{\mathbb{F}_q} \overline{\mathbb{F}_q}, \mathbb{Q}_\ell) \right).$$

For example, if $X_n = \mathbf{P}^n$ with $|X_n(\mathbb{F}_q)| = q^n + q^{n-1} + \cdots + 1$, then the term q^j arises from $H^{2j}_{\text{c,ét}}$ on the right-hand side. More generally, one expects that the dominant terms arise from the compactly supported cohomology in high degree or, what is the same if X_n is smooth, the usual (co)homology in low degree. This leads naturally to asking for some sense in which the low-degree (co)homology of X_n is "controlled."

For instance, suppose that the X_n are smooth of dimension n and geometrically irreducible for large n. Then the only nonvanishing terms on the right-hand sum of (1.8.2) occur in cohomological dimensions $j \leq 2n$, and the contribution of $j = 2n$ is exactly q^n.

To bound the remaining terms, we suppose that there exists a constant C so that

(1.8.3) $\dim H^i_{\text{ét}}(X_n \times_{\mathbb{F}_q} \overline{\mathbb{F}_q}, \mathbb{Q}_\ell) \leq C^i$ for all $n, i \geq 1$.

The Deligne bounds [20] show that the eigenvalues of the action of Frobenius on $H^{2n-i}_{\text{c,et}}(X_n \times_{\mathbb{F}_q} \overline{\mathbb{F}_q}, \mathbb{Q}_\ell)$ are algebraic numbers all of whose complex eigenvalues are bounded above by $q^{n-i/2}$. Now, $\dim(H^{2n-i}_{\text{c,ét}}) \leq C^i$ by (1.8.3) and Poincaré duality, and so the trace of Frobenius on $H^{2n-i}_{\text{c,ét}}$ is bounded above by $q^n \cdot (C/\sqrt{q})^i$. Inserting this bound into the Lefschetz fixed point formula (1.8.2) to handle all terms with $j < 2n$, we arrive at

(1.8.4) $\left| \dfrac{\#X_n(\mathbb{F}_q)}{q^n} - 1 \right| \leq \dfrac{1}{\sqrt{q}/C - 1}$

for all $q > C^2$. In other words, X_n has approximately q^n points over $\mathbb{F}_q$, as one might naively guess. (Indeed, some of the consequences of making this naive guess were discussed by the first and second authors in [23] who, at the time, had no idea that the guess might under some circumstances be correct.)

How might one establish bounds of the form (1.8.3)? Suppose that we can establish the existence of an isomorphism

$$H_i(X_n, \mathbb{Q}_\ell) \to H_i(X_{n+1}, \mathbb{Q}_\ell)$$

for $i \leq n$. (In fact, $i \leq An$ for any positive constant A will be just as good for the type of application discussed in this paper.) Here H_i denotes the singular homology of the complex points of these varieties, equipped with the analytic topology. One immediately obtains the bound $\dim H_i(X_n, \mathbb{Q}_\ell) \leq \dim H_i(X_i, \mathbb{Q}_\ell)$ for $i < n$.[4] In particular, (1.8.3) would then follow from an upper bound of the form

$$(1.8.5) \qquad \dim H^i_{\text{ét}}(X_n \times_{\mathbb{F}_q} \overline{\mathbb{F}_q}, \mathbb{Q}_\ell) \leq C^n$$

for each i, which tends to be much easier; it can be checked given some *a priori* bound on the "complexity" of the variety X_n.

So a theorem about homological stability can, in principle, be used to prove an asymptotic result in analytic number theory over function fields over finite fields. We now present some examples in order to sketch the potential scope of this point of view.

1.9. *Stable homology and analytic number theory over function fields: further examples.* In this section we discuss some other problems that connect analytic number theory of function fields with the homology of a natural sequence of moduli spaces.

(1) The number of squarefree integers in the interval $[X, 2X]$ is asymptotic to $\frac{X}{\zeta(2)}$. Over the rational function field over $\mathbb{F}_q$, the corresponding question is: How many monic squarefree degree-n polynomials are there in $\mathbb{F}_q[t]$?

Set $X_n = \text{Conf}^n \mathbb{A}^1$, the configuration space of n points on $\mathbb{A}^1$, or, equivalently, the space of monic squarefree polynomials of degree n. In this case, one indeed has homological stability [4]: the homology of X_n with $\mathbb{Q}_\ell$-coefficients is nonvanishing only in degrees 0 and 1; a computation with the Lefschetz formula then yields

$$|X_n(\mathbb{F}_q)| = q^n - q^{n-1} = \frac{q^n}{\zeta_{\mathbb{A}^1/\mathbb{F}_q}(2)},$$

which is precisely analogous to the result in the number field case.

(2) A question with no obvious counterpart over a number field is: How does the number of genus g curves over $\mathbb{F}_q$ behave, as $g \to \infty$?

As already mentioned, Harer's theorem gives homological stability for $\mathcal{M}_g$ (as orbifold) as $g \to \infty$. But in this case, there is no bound of the form (1.8.5): the Euler characteristic of $\mathcal{M}_g$ grows superexponentially with g and so, in particular, there is no bound on the Betti numbers in the

---

[4]If the isomorphisms arise from algebraic maps $X_n \to X_{n+1}$ defined over $\mathbb{F}_q$, then one even has isomorphisms of étale cohomology groups compatible with Galois action. We have not pursued this refinement in the present paper.

unstable range analogous to Proposition 2.5. Thus, the homology stabilization does not enforce any regularity on $\frac{|\mathcal{M}_g(\mathbb{F}_q)|}{q^{\dim \mathcal{M}_g}}$, and it is not at all clear this ratio should be expected to approach a limit as $g \to \infty$. (See [18] for a discussion of this case, including the best known upper bounds for $|\mathcal{M}_g(\mathbb{F}_q)|$.)

(3) We expect the problem of counting points of bounded height on varieties over global fields to provide a very general example of the relation between stable homology and analytic number theory.

Over $\mathbb{F}_q(t)$, this problem amounts to counting the number of $\mathbb{F}_q$-points on the space of maps from $\mathbb{P}^1$ to an algebraic variety X; over $\mathbb{C}$, the homological stability for such spaces is example (4) of Section 1.6.

It has been observed by the first two authors [24] that one can "reverse" the reasoning used in this paper, counting points over finite fields via the Hardy-Littlewood method and then applying the Lefschetz fixed point formula to obtain geometric information about $\operatorname{Hol}^d(\mathbb{CP}^1, X)$. In cases where the Hardy-Littlewood method does not apply, there is a notable similarity between the class of varieties X such that $\operatorname{Maps}^d(S^2, X)$ is known to have stable homology, and those where the rational points of bounded height on $X/\mathbb{Q}$ is known to obey the asymptotic prediction of the Batyrev-Manin conjecture.

1.11. *Notation.* If a group G acts on a topological space X, we will use the notation $X /\!\!/ G$ for the Borel construction $X /\!\!/ G := EG \times_G X$, where EG is a contractible G-space with free G action.

If g is an element of a finite group G, we denote the order of g by $|g|$. If $g, h \in G$, we denote by g^h the conjugate $h^{-1}gh$.

We will deal with graded modules $M = \bigoplus M_n$ over various graded rings. We will always understand the grading to be supported on nonnegative integers, i.e., $M_n = 0$ for $n < 0$ for all our graded modules in this paper.

If $M = \oplus M_n$ is a graded module for a graded ring $R = \oplus R_n$, we write $\deg M$ to mean the maximal n such that $M_n \neq 0$. If there is no such n, we say $\deg M = \infty$. The notation $M[k]$ means "M shifted by k" — in other words, $M[k]_n = M_{n-k}$ for all $n \geq k$.

If r is a homogeneous element in R_n, we write $\deg(r) = n$.

2. Definitions

2.1. *Hurwitz spaces.* We begin with a topological definition of Hurwitz spaces. Their interpretation as moduli spaces of branched covers may not be immediate from this definition; we will return to that description afterwards.

Let D be a closed disc with a marked point $*$ on the boundary, and write Conf_n for the configuration space of n unordered, distinct points in the interior of D. Fix a basepoint $c_n = \{P_1, \ldots, P_n\}$ in Conf_n, and recall also that the Artin braid group B_n on n strands is isomorphic both to the mapping class group of the punctured surface $\Sigma := D - \{P_1, \ldots, P_n\}$, and to the fundamental group of Conf_n:

$$\pi_0 \mathrm{Diff}^+(\Sigma, \partial\Sigma) \cong B_n \cong \pi_1(\mathrm{Conf}_n, c_n).$$

Referring to the presentation (1.3.1), the first isomorphism identifies the braid σ_j with the diffeomorphism that exchanges P_j and P_{j+1} by a half Dehn twist along a circle containing only these two punctures. The second isomorphism carries σ_j to a path in Conf_n that switches P_j and P_{j+1}, while leaving immobile the other P_i.

For definiteness, we may take D to be the closed disc centered at 0 of radius $n + 1$, take $P_j = j \in \mathbf{C}$, the diffeomorphism to be the half-Dehn twist around a a circle centered at $j + \frac{1}{2} \in \mathbf{C}$ of radius $3/4$, and the path switching P_j and P_{j+1} to be the path that rotates them both by a half-twist around a circle centered at $j + 1/2$ and of radius $1/2$.

We recall that Conf_n is an Eilenberg-MacLane space[5] $K(B_n, 1)$. Let $\pi = \pi_1(\Sigma, *)$. Fix, for each i, an embedded loop γ_i in Σ, based at $*$, and winding once (counterclockwise) around the puncture P_i, and not winding around any other puncture. It is possible to do this in such a way that the γ_i are not intersecting except at $*$. For definiteness, take γ_i to be a straight path from $*$ to a point P'_i very close to P_i, together with a small loop winding once around P_i based at P'_i.

The γ_i freely generate π; thus we have specified an isomorphism between the free group F_n on n generators and π. Of course, different choices of γ_i

---

[5]Indeed, consider the finite covering space of Conf_n that parametrizes n *ordered* points; this space can be presented as an iterated fibration of punctured discs, and is thus aspherical.

will yield differing isomorphisms. We note, however, that any two choices of generators $\{\gamma_i\}$ and $\{\gamma_i'\}$ are related by a diffeomorphism[6] $h : \Sigma \to \Sigma$ fixing $\partial\Sigma$.

Since the marked point $* \in \partial\Sigma$ is fixed by the action of the diffeomorphism group $\mathrm{Diff}^+(\Sigma, \partial\Sigma)$, there is a natural action of B_n on π, and hence the set $\mathrm{Hom}(\pi, G)$ of homomorphisms from π to any discrete group G. Let $c \subseteq G$ be a conjugacy class or union thereof. Write $\mathrm{Hom}^c(\pi, G)$ for the subset of homomorphisms $f : \pi \to G$ that carry each γ_i into c.

In principle, this subset is dependent upon our choice of generators γ_i. Note, however, that any two loops γ_i, γ_i' around P_i, of the type described above, are conjugate. Thus we may reformulate the condition on f without reference to the choice of $\{\gamma_i\}$ to say that f carries the free homotopy class of a loop around each puncture into c. In this formulation, it is also apparent that this subset is invariant under the action of B_n, since any diffeomorphism of Σ preserves the set of free homotopy classes of loops around the punctures.

Finally, write $\mathrm{Sur}^c(\pi, G)$ for the subset of $\mathrm{Hom}^c(\pi, G)$ consisting of surjective homomorphisms (nonempty only if c generates G), and write $\mathrm{Sur}(\pi, G)$ for the similarly defined subset of $\mathrm{Hom}(\pi, G)$. These are evidently a B_n-invariant subset of $\mathrm{Hom}^c(\pi, G)$.

2.2. *Definition.* Let $\widetilde{\mathrm{Conf}}_n$ be the universal cover of Conf_n, together with a fixed point $\tilde{c}_n$ above c_n; thus B_n acts on $\widetilde{\mathrm{Conf}}_n$ by deck transformations, this action being uniquely specified by requiring the action of $b \in B_n$ on $\tilde{c}_n$ to coincide with the monodromy action of $B_n \simeq \pi_1(\mathrm{Conf}_n, c_n)$. Define the *Hurwitz spaces*

- $\mathrm{Hur}_{G,n} := \widetilde{\mathrm{Conf}}_n \times_{B_n} \mathrm{Hom}(\pi, G)$,
- $\mathrm{Hur}_{G,n}^c := \widetilde{\mathrm{Conf}}_n \times_{B_n} \mathrm{Hom}^c(\pi, G)$,
- $\mathrm{CHur}_{G,n} = \widetilde{\mathrm{Conf}}_n \times_{B_n} \mathrm{Sur}(\pi, G)$,
- $\mathrm{CHur}_{G,n}^c := \widetilde{\mathrm{Conf}}_n \times_{B_n} \mathrm{Sur}^c(\pi, G)$.

These are covering spaces of Conf_n, finite sheeted when G is finite. Moreover,
$$(\text{fiber of } \mathrm{Hur}_{G,n} \to \mathrm{Conf}_n \text{ above } c_n) \cong \mathrm{Hom}(\pi, G) \cong G^n,$$
where the first map sends $\tilde{c}_n \times a \mapsto a \in \mathrm{Hom}(\pi, G)$ and the second map uses the identification, described above, of F_n with π. Similarly for the other space, the fibers above c_n are identified with $\mathrm{Hom}^c(\pi, G), \mathrm{Sur}(\pi, G)$ and $\mathrm{Sur}^c(\pi, G)$, respectively, which in turn are identified with the set of n-tuples of elements of $c \subseteq G$, the elements of c^n that generate G, and the elements of G^n that generate G respectively.

---

[6]The map h may be constructed by gluing together, for each i, diffeomorphisms from the punctured disk bounded by γ_i to the one bounded by γ_i', as well as the exterior of their unions.

In all cases, with respect to these identifications, the monodromy action of $\pi_1(\mathrm{Conf}_n, c_n) \simeq B_n$ on the fiber is identified with the braiding action (1.3.2) of B_n on G^n (see, e.g., [5, eq. (14)]).

Note also that G acts on $\mathrm{Hom}(\pi, G)$ by conjugation in the target; the subsets $\mathrm{Hom}^c(\pi, G), \mathrm{Sur}(\pi, G)$, and $\mathrm{Sur}^c(\pi, G)$ are invariant under this action. When given in terms of sets of n-tuples, the action is by termwise conjugation on the n-tuple. Furthermore, this action commutes with the action of B_n on the domain and so yields an action of G on all of the spaces above.

Finally, we note that $\widetilde{\mathrm{Conf}}_n$ is a contractible space with a free action of B_n, and so we can regard $\mathrm{Hur}_{G,n}$ as the Borel construction $EB_n \times_{B_n} \mathrm{Hom}(\pi, G)$ (and similarly for the other spaces above).

2.3. *Interpretation as moduli spaces.* A *marked n-branched G-cover of the disc* is a quintuple $(Y, p, \bullet, S, \alpha)$, where

- $S \subset D$ is a set of n distinct points in the interior of D,
- $p : Y \to D - S$ is a covering map,
- $\alpha : G \to \mathrm{Aut}(p)$ is a map inducing a simply transitive action of G on each fiber,
- $\bullet$ is a point in the fiber of p above $*$.

Note that we do not restrict ourselves to *connected* covers Y. We say two marked n-branched G-covers Y and Y' are isomorphic if there is a homeomorphism from Y to Y' over $D - S$, compatible with the remaining data.

Then we have bijections between the sets (a), (b), and (c) described below:

(a) points of $\mathrm{Hur}_{G,n}$;
(b) pairs (S, f), where $S \in \mathrm{Conf}_n$, and $f : \pi_1(D - S, *) \to G$ is a homomorphism;
(c) isomorphism classes of marked n-branched G-covers of D.

For the bijection between (a) and (b), regard elements of $\widetilde{\mathrm{Conf}}_n$ above $S \in \mathrm{Conf}_n$ as a homotopy class of paths between c_n and S, in such a fashion that $\tilde{c}_n$ corresponds to the trivial path. Such a path induces an isomorphism of $\pi = \pi_1(D - c_n, *)$ with $\pi_1(D - S, *)$. Thus each point of $\widetilde{\mathrm{Conf}}_n \times \mathrm{Hom}(\pi, G)$ gives a pair $(S, \pi_1(D - S, *) \to G)$, and this descends to the desired bijection.

For the bijection between (b) and (c), start with an n-branched G-cover, and identify the fiber above $*$ with G through the map $g \mapsto g\bullet$; with respect to this identification, the action of $\pi_1(D - S, *)$ on $p^{-1}(*)$ is by right multiplication by G and so defines a homomorphism $\pi_1(D - S, *) \to G$.

Remark. Hurwitz spaces appear in many places in the literature, and the definition admits many variants. We emphasize that our Hurwitz spaces differ from many standard treatments in that we do not restrict our attention to connected G-covers (we reserve the notation CHur for such a Hurwitz space),

and in that we select a marked point in the fiber over $*$. This latter difference means that, by contrast with the notation in some of the literature, the points of our Hurwitz space with some fixed set of branch points $S \subset D$ are in bijection "on the nose" with the homomorphisms from the fundamental group of the punctured disc to G, not with the conjugacy classes of such homomorphisms.

The action of G on $\mathrm{Hur}_{G,n}$ defined in the discussion following Definition 2.2 is given in these terms by moving the marked point, i.e., via the rule

$$g(Y, p, \bullet, S, \alpha) = (Y, p, \alpha(g)\bullet, S, \alpha).$$

Later we shall study the quotient of $\mathrm{Hur}_{G,n}$ by this G-action; this quotient space, which we denote $\mathrm{Hur}_{G,n}/G$, parametrizes n-branched G-covers without the specification of $\bullet$.

The subspace $\mathrm{CHur}_{G,n} \subseteq \mathrm{Hur}_{G,n}$ consists of the space of covers with full monodromy G — in other words, the covers corresponding to $surjective$ homomorphisms $\pi_1(D - S, *) \to G$. The prepended "C" is meant to recall that this space parametrizes $connected$ G-covers of the disc. The space $\mathrm{CHur}_{G,n}$ itself need not be connected in general (cf. Section 1.4).

2.4. $Combinatorial$ $invariants$. $\mathrm{Hur}_{G,n}$ is usually disconnected; i.e., the action of B_n on $\mathrm{Hom}(F_n, G)$ is typically not transitive.

We now describe some invariants of a cover p that are constant on connected components of $\mathrm{Hur}_{G,n}$. By Definition 2.2, it is equivalent to the combinatorial problem of specifying a B_n-invariant function on G^n.

- The $global$ $monodromy$ of p is the image of $\pi_1(D - S, *)$ in G. In combinatorial terms, this is the map $(g_1, \ldots, g_n) \to \langle g_1, \ldots, g_n \rangle$, the subgroup of G generated by the g_i.
- The $boundary$ $monodromy$ of p is the element of G induced by a counterclockwise loop around ∂D. (More precisely, transport around such a loop moves $\bullet$ to a point $g.\bullet$ for a unique $g \in G$.)
  In combinatorial terms, this is the map $(g_1, \ldots, g_n) \mapsto g_1 g_2 \cdots g_n$.
- For each i, the monodromy around a small loop encircling P_i is an element of G, well defined only up to conjugacy. The resulting multiset of n conjugacy classes of G is called the $Nielsen$ $class$ of p.
  Combinatorially, the Nielsen class map associates to $(g_1, \ldots, g_n)$ the multiset obtained by replacing each g_i with its conjugacy class.

Fixing the global monodromy, boundary monodromy, and Nielsen class of a cover specifies a subspace of $\mathrm{Hur}_{G,n}$; although it may be disconnected, there are no "obvious" invariants further separating connected components.

For a conjugacy class $c \subseteq G$, we note that $\mathrm{Hur}_{G,n}^c$ is the subspace of $\mathrm{Hur}_{G,n}$ consisting of covers whose Nielsen class is n copies of c. Our main goal in the present paper is to study the homology groups $H_p(\mathrm{Hur}_{G,n}^c)$, especially

their asymptotic behavior as n grows with G and c held fixed. It is also natural to consider the larger spaces where the monodromy is drawn not from a single conjugacy class c but from a fixed union of conjugacy classes, or for that matter from the whole group. We do not pursue this generalization in the present paper.

2.5. PROPOSITION. $\mathrm{Hur}_{G,n}$ *and* $\mathrm{Hur}_{G,n}^c$ *are both homotopy equivalent to* CW *complexes with at most* $(2|G|)^n$ *cells.*

Proof. Since $\mathrm{Hur}_{G,n}$ and $\mathrm{Hur}_{G,n}^c$ are both coverings of Conf_n with fibers of size $\leq |G|^n$, it suffices to check that Conf_n is homotopy equivalent to a CW complex with $\leq 2^n$ cells. For this see, e.g., [46] or [11]. $\square$

Proposition 2.5 is critical for the arithmetic applications; in the Lefschetz trace formula it is this fact that allows us to neglect the contribution of cohomology classes in the unstable range.

2.6. *Gluing maps.* Arising from the natural inclusions $B_n \times B_m \to B_{n+m}$ and $G^n \times G^m \to G^{n+m}$, we obtain a map on Borel constructions

$$(EB_n \times_{B_n} G^n) \times (EB_m \times_{B_m} G^m) \to EB_{n+m} \times_{B_{n+m}} G^{n+m},$$

which defines (up to homotopy) a gluing map

$$\mathrm{Hur}_{G,n} \times \mathrm{Hur}_{G,m} \longrightarrow \mathrm{Hur}_{G,n+m}.$$

This multiplicative structure is associative up to homotopy since both maps inducing it are; collectively, they make the union of the Hurwitz spaces into an H-space.

Geometrically, these maps associate to a pair of branched covers Y_1, Y_2 of D a new cover, Y_3. Pick two standard disjoint, embedded loops γ_1, γ_2 in D based at $*$; then the restriction of Y_3 to the interior of the region described by each γ_i is isomorphic to Y_i. On the complement of the loops, Y_3 is the trivial G-bundle extension.

Similarly, we have a multiplication

$$\mathrm{Hur}_{G,n}^c \times \mathrm{Hur}_{G,m}^c \longrightarrow \mathrm{Hur}_{G,n+m}^c.$$

We note that these gluing maps are equivariant for the action of G (where G acts on all three factors in the fashion defined after Definition 2.2) when we take the model given by the Borel construction.

3. The ring R of connected components

Let k be a field of characteristic prime to $|G|$. Then the graded ring

$$R = \sum_n H_0(\mathrm{Hur}_{G,n}^c, k)$$

inherits, from the multiplication on Hurwitz spaces (Section 2.6), the structure of a noncommutative k-algebra; moreover, the higher homology of Hurwitz spaces carries the structure of R-module.

3.1. *Definition.* We say the pair (G, c) has the *nonsplitting property* if c generates G and, moreover, for every subgroup H of G, the intersection $c \cap H$ is either empty or a conjugacy class of H.

Our main result in this section is Lemma 3.5, which implies that *if (G, c) has the nonsplitting property, then there exists a central homogeneous element $U \in R$ so that the degree of R/UR is finite.*

Before discussing R, we begin by giving the basic example of nonsplitting pairs:

3.2. Lemma. *Let G be a finite group whose order is $2s$ for s odd. Then there is a unique conjugacy class of involutions $c \subset G$; if c generates G, then (G, c) is nonsplitting.*

Proof. The fact that all involutions are conjugate follows from conjugacy of 2-Sylow subgroups; any subgroup H of G containing an involution has order $2s'$ for s' odd, and the nonsplitting follows from the uniqueness assertion applied to H. $\qquad\square$

A group G as in the lemma is necessarily isomorphic to $G_0 \rtimes (\mathbb{Z}/2\mathbb{Z})$ for some group G_0 of odd order. In fact, these are the only cases of nonsplitting pairs where c is a involution.[7] There are other nonsplitting pairs: for example, $G = A_4$ and c one of the classes of 3-cycles.

For the remainder of this paper all theorems have as a hypothesis that (G, c) has the nonsplitting property.

3.3. *Combinatorial description of R.* The graded ring R has a very concrete description: Let $\tilde{S}$ be the set of tuples of elements from c (of any nonnegative length), and let S be the quotient of $\tilde{S}$ by the action of the braid group. Then S is a semigroup under the operation of concatenation, and R is the semigroup algebra $k[S]$. We let $S_n = c^n/B_n$ be the subset of S consisting of elements of degree n; for $s \in S$ (considered as an element of R) write $\partial s \in G$ for the boundary monodromy of s; if s is represented by $(g_1, \ldots, g_n)$, then $\partial s = g_1 \cdots g_n$.

R is generated over k by degree 1 elements $\{r_g\}_{g \in c}$, subject to the relations

$$(3.3.1) \qquad\qquad r_g r_h = r_{ghg^{-1}} r_g.$$

We occasionally denote r_g by $r(g)$ if the group element in question is too typographically complicated to fit in a subscript. We note that we learned

---

[7]This fact follows from Glauberman's Z^* theorem, as Richard Lyons explained to us; the authors thank `mathoverflow.net` for providing a forum where we could ask about this and be provided with an authoritative reference.

the idea of using the semigroup S to study connected components of Hurwitz spaces from the appendix to [27].

3.4. PROPOSITION. *Let* $g \in c$. *For sufficiently large* n, *every* n-*tuple* $(g_1, \ldots, g_n)$ *in* $\tilde{S}$ *whose elements generate* G *is equivalent under the braid group to an* n-*tuple* $(g, g_2', \ldots, g_n')$, *where* $g_2', \ldots, g_n'$ *generate* G.

This implies stability for the zeroth Betti number: $b_0(\mathrm{CHur}_{G,n}^c)$ is independent of n for sufficiently large n. Regarding $\pi_0(\mathrm{CHur}_{G,n}^c)$ as a subset of S_n, this shows that the map $\pi_0(\mathrm{CHur}_{G,n}^c) \to \pi_0(\mathrm{CHur}_{G,n+1}^c)$ given by adding g at the beginning of an n-tuple is surjective for n sufficiently large. Both sets are finite, so this is eventually a bijection.

Proof. This is well known (see, e.g., [27]), but for completeness we include a proof here. It is clear that, by repeated action of the braid action (1.3.2), we can pull the ith monodromy element to the beginning of the n-tuple; that is, the n-tuple $(g_1, \ldots, g_n)$ is equivalent to $(g_i, g_1', \ldots, g_{n-1}')$ for some $(n-1)$-tuple $(g_1', \ldots, g_{n-1}')$. Write d for the order of an element of c. If $n > d|c|$, some element g' of c occurs at least $d+1$ times in $(g_1, \ldots, g_n)$; we can use the braid action to pull these back to the front, forming an n-tuple

$$\underline{g} = (g', g', \ldots, g', g_1', \ldots, g_{n-d-1}')$$

equivalent under the braid action to $(g_1, \ldots, g_n)$. The elements of this n-tuple generate G, whence $g', g_1', \ldots, g_{n-d-1}'$ generate G.

Now the fact that $(g')^d = 1$ implies that, for any k-tuple $h_1, \ldots, h_k$, the $d+k$-tuple

$$(\underbrace{g', g', \ldots, g'}_{d}, h_1, \ldots, h_k)$$

is equivalent under braiding to

$$(hg'h^{-1}, \ldots, hg'h^{-1}, h_1, \ldots, h_k),$$

where h is the product $h_1 \cdots h_k$; this equivalence is implemented by the braid that winds the first d strands around the last k strands; i.e., we first braid the g's to the right to obtain $(h_1, \ldots, h_k, g', \ldots, g')$ and then braid the hs to the right to obtain the tuple above. Since the braid action can conjugate the d copies of g' by $h_1 \cdots h_k$ and by $h_1 \cdots h_{k-1}$, it can also conjugate those d copies of g' by h_k alone. Repeating this process, we find that

$$(g', g', \ldots, g', h_1, \ldots, h_k)$$

is equivalent to

$$(g'', g'', \ldots, g'', h_1, \ldots, h_k)$$

for any g'' that is conjugate to g' via an element in the group generated by $h_1, \ldots, h_k$.

Now the final $n - d$ elements of $\underline{g}$ generate G, as we have seen. Since g' is conjugate to g via some element of G, the argument above shows that $\underline{g}$ is equivalent under the braid action to

$$(g, g, \ldots, g, g', g'_1, \ldots, g'_{n-d-1}).$$

This proves the proposition. $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\square$

Recall that if $M = \oplus M_n$ is a graded R-module, we write $\deg M$ to mean the maximal n such that $M_n \neq 0$. If there is no such n, we say $\deg M = \infty$. In the following lemma, we prove a finiteness condition on R that will turn out to imply all the homological properties of the category of R-modules that we require for the proof of the main theorem.

3.5. LEMMA. *Suppose that (G, c) has the nonsplitting property. For an integer D, write*

$$U_D = \sum_{g \in c} r_g^{D|g|},$$

so that U_D is in the center of R. Then there exists a D such that the degree of both kernel and cokernel of

$$R \xrightarrow{U_D} R, \quad r \mapsto U_D r$$

are finite.

Proof. The main ingredient is Proposition 3.4. The nonsplitting property for G is used in an essential way, allowing the application of this proposition to the pair $(H, c \cap H)$ for subgroups $H \leq G$.

Within the present proof, we refer to the subset of S_n consisting of braid orbits on n-tuples generating H as $S_n(H)$. We first show that for every subgroup H of G, every element g of $c \cap H$, and every sufficiently large n, the map

$$(3.5.1) \qquad\qquad S_n(H) \to S_{n+|g|}(H), \quad s \mapsto r_g^{|g|} s$$

is bijective. It suffices to show that this map is surjective for large enough n: since all the sets involved are finite, it must then be eventually bijective, which is the assertion to be proved.

If $c \cap H$ is empty, then so too is $S_N(H)$; the claim is vacuously true. If not, take $s \in S_n(H)$; for sufficiently large n, Proposition 3.4 (applied to the group H and conjugacy class $c \cap H$) shows that

$$s \sim r_g s', \quad s' \in S_{n-1}(H).$$

Note it is exactly at this point that the nonsplitting property enters our argument: we used the fact that $c \cap H$ is a single conjugacy class. See also the discussion after the theorem in Section 1.4.

Increasing n as necessary, we can repeat this process $|g|$ times; this shows that $r_g^{|g|}$ induces a surjective map, as desired.

It is not clear that different choices of $g \in c \cap H$ induce the *same* bijection in (3.5.1). For $g_1, g_2 \in c \cap H$, observe that multiplication by $A := r_{g_1}^{|g_1|}$ and $B := r_{g_2}^{|g_2|}$ commute (as self-maps of S). For sufficiently large n, the map B is a bijection from $S_n(H)$ to $S_{n+|g|}(H)$ and so has an inverse, which we denote $B^{-1} : S_{n+|g|}(H) \to S_n(H)$. So, again for sufficiently large n, $A \circ B^{-1}$ is a permutation of $S_n(H)$. Let D be chosen so that every permutation of every $S_n(H)$ has order dividing D. Then $(A \circ B^{-1})^D = A^D \circ B^{-D}$ induces the identity map on $S_n(H)$ for large enough n. Thus, for such D, the map

$$S_n(H) \xrightarrow{r_g^{D|g|}} S_{n+D|g|}(H)$$

is — for large n — a bijection and is independent of $g \in c \cap H$.

For $m \geq 1$, set $F_m R$ to be the subspace of R generated by elements in $S_n(H)$, as n ranges over nonnegative integers and H ranges over subgroups of order at least m. Note that U_D preserves $F_m R$; we now show, by descending induction on m, that

(3.5.2) $$U_D : F_m R_n \to F_m R_{n+D|g|}$$

is an isomorphism for sufficiently large n. Again, it is sufficient to show that (3.5.2) is surjective for sufficiently large n.

The inductive claim is valid for $m > |G|$ trivially. Now suppose it is true whenever $m > m_0$.

Let H be a subgroup of G of size m_0. It suffices to check that — for large enough N — every $y \in S_N(H)$ belongs to the image of U_D. Again, this is vacuously true if $c \cap H = \emptyset$. Otherwise, by what we have shown, there exists $x \in S_{N-D|g|}(H)$ such that $r_g^{D|g|} x = y$ for all $g \in c \cap H$. Therefore,

$$U_D x = |c \cap H| y + y', \text{ where } y' = \textstyle\sum_{g \in c \backslash (c \cap H)} r_g^{D|g|} x \in F_{m_0+1} R_N.$$

By induction, there exists, for sufficiently large N, $x' \in F_{m_0+1} R_{N-D|g|}$ such that $U_D x' = y'$. Consequently, $U_D(x - x') = |c \cap H| y$. Since $c \cap H$ is a conjugacy class of $H < G$, the size of $c \cap H$ divides $|G|$ and is therefore invertible in k. Thus, we conclude that y is in the image of U_D, as desired.   □

We note that the assumption that $|G|$ is invertible in k was used in a substantial way in this proof. Consequently, we do not expect that rational homological stability for Hurwitz spaces can be improved to an integral result, since Theorem 6.1 depends in a basic way upon this fact. However, the proof of Lemma 3.5 suggests that an integral result may be possible for subspaces of Hurwitz spaces (such as $\mathrm{CHur}_{G,n}^c$).

4. The $\mathcal{K}$-complex associated to an R-module

This section is solely concerned with homological properties of the ring R introduced in Section 3. In particular, we associate to each R-module M a certain Koszul-like complex, the $\mathcal{K}$-complex (Section 4.1). We shall see in Section 5 that the homology of Hurwitz spaces can be inductively expressed in terms of $\mathcal{K}$-complexes formed from homology of *smaller* Hurwitz spaces.

Our main result, Theorem 4.2, is that the higher homology of the $\mathcal{K}$-complex is controlled by its H_0 and H_1. The overall thrust of this section can be roughly summarized by the slogan "R behaves as if it had cohomological dimension 1."

Let us explain why this point of view is useful for proving homological stability for Hurwitz spaces. In most situations where homological stability is understood, one has a sequence of (usually connected) spaces X_n and stabilization maps $f_n : X_n \to X_{n+1}$; the goal is to show that each f_n induces homology isomorphisms in a range of dimensions. Let $X = \sqcup_n X_n$, and consider the homology
$$M_p = H_p(X) = \oplus_n H_p(X_n).$$
Give M_p the structure of a $k[x]$-module by making the indeterminate x act via the stabilization map. M_p admits a grading by the number n, and x acts as a degree 1 operator. Homological stability is rephrased as the statement that x is an isomorphism in sufficiently high degree. Equivalently, we need the quotient and x-torsion
$$\mathrm{Tor}_0^{k[x]}(k, M_p) = M_p/xM_p \quad \text{and} \quad \mathrm{Tor}_1^{k[x]}(k, M_p) = M_p[x]$$
to be concentrated in low degrees.

Approaching homological stability for Hurwitz spaces this way, one immediately runs into a problem: there are many natural stabilization maps, one for each isomorphism class of branched cover of the disk. This more complicated structure is encoded, however, in the ring R of connected components of the Hurwitz space, which replaces $k[x] = H_0(X)$ above. As we have seen in Lemma 3.5, R itself satisfies a form of stability so long as (G, c) is nonsplitting. This fact, combined with control of the homological algebra of R (developed in this section) ultimately gives rise to homological stability for the Hurwitz spaces.

Throughout this section (G, c) is nonsplitting, and we take U to be the central element U_D defined in Lemma 3.5.

4.1. Let M be any graded left R-module. We may define a "Koszul-like" complex ($\mathcal{K}$-complex for short) associated to M, where[8] $\mathcal{K}(M)_q = k[c^q] \otimes_k M[q],$

---

[8] Recall (Section 1.11) that $[i]$ denotes a shift in grading by i.

$q \geq 0$; that is, we have

$$\mathcal{K}(M) := \cdots \to k[c^q] \otimes_k M[q] \to k[c^{q-1}] \otimes_k M[q-1]$$
(4.1.1)
$$\to \cdots \to k[c] \otimes_k M[1] \to M[0],$$

where the differential $\mathcal{K}(M)_{q+1} \to \mathcal{K}(M)_q$ is described by

$$(g_0, \ldots, g_q) \otimes m \mapsto \sum_{i=0}^{q}(-1)^i (g_0, \ldots, \widehat{g_i}, \ldots, g_q) \otimes r(g_i^{g_{i+1}\cdots g_q})m.$$

We equip $k[c^q]$ with the trivial grading, i.e., the one concentrated in degree 0. Then the differentials preserve the grading. Moreover, if $M = R$, each homology group of $\mathcal{K}$ is equipped with the natural structure of a graded *right* R-module. This notion is chosen to model a complex that will arise in our study of the arc complex (Section 5).

4.2. THEOREM. *Suppose (G, c) is nonsplitting; let M be a graded left R-module, and let $h_i = \deg(H_i(\mathcal{K}(M)))$. Then there exists a constant $A_0 = A_0(G, c)$ so that*

(4.2.1)
$$h_q \leq \max(h_0, h_1) + A_0 q \quad (q > 1).$$

Moreover, $M \xrightarrow{U} M$ is an isomorphism in source degree $\geq \max(h_0, h_1) + A_0$. (That is to say, the induced map $M_i \to M_{i+\deg(U)}$ is an isomorphism for $i \geq \max(h_0, h_1) + A_0$.)

Finally, in the case $M = R$, h_0 and h_1 are both finite.

The explicit value of A_0 that comes from the proof is given in (4.5.3); we have not attempted to optimize this value as far as possible, since the precise bound makes no difference to our end goal.

This Theorem is fundamental to the proof of our main result. A basic tool in its proof is comparing the homological algebra of R and the commutative, central subring $k[U]$. The *centrality* of U is vital to our argument and will be used without comment repeatedly.

In what follows, we use the following notation and conventions:

- We denote the two-sided ideal $\oplus_{n>0} R_n$ by $R_{>0}$; we give the field k the structure of R-bimodule by identifying it with $R/R_{>0}$.
- For M a graded left R-module, we denote by $H_i(M)$ the graded left R-module $\mathrm{Tor}_i^R(k, M)$. In particular, $H_0(M) = M/R_{>0}M$.
- We set $\bar{R} = R/UR$; it is an R-bimodule.
- For an R-module M, write $M[U]$ for the U-torsion in M, i.e., the kernel of the "multiplication by U" map $M \xrightarrow{U} M$.
- Finally, recall (Section 1.11) that if M is a graded R-module, we write $\deg(M)$ for the largest degree n such that $M_n \neq 0$, if it exists; otherwise $\deg(M) = \infty$.

Thus, with these conventions, $\deg(\bar{R})$ and $\deg(R[U])$ are both finite, because of Lemma 3.5.

4.3. *Comparing the homological algebra of R and $k[U]$.* Consider the functors

$$\text{Left graded } R\text{-modules} \xrightarrow{\ f\ } \text{Left graded } \bar{R}\text{-modules}$$

$$\xrightarrow{\ g\ } \text{graded } k\text{-vector spaces,}$$

where f sends M to $\bar{R} \otimes_R M$ and g sends $\bar{M}$ to $k \otimes_{\bar{R}} \bar{M}$. Both f and g are right exact and admit left derived functors. Since f carries free left R-modules to free left $\bar{R}$-modules, and each left R-module has a resolution by free left R-modules, we have a spectral sequence

$$(4.3.1) \qquad \text{Tor}_i^{\bar{R}}(k, \text{Tor}_j^R(\bar{R}, M)) \Rightarrow \text{Tor}_{i+j}^R(k, M).$$

4.4. LEMMA. *Let M be a graded left R-module, and let N be a graded right R-module. Then*

$$(4.4.1) \qquad \deg(N \otimes_R M) \leq \deg(N) + \deg(H_0(M)).$$

Proof. When $\deg(N) = \infty$ or $\deg(H_0(M)) = \infty$, the assertion is vacuously true, so we assume both numbers are finite from now on. We will use similar reasoning in the proofs that follow, without explicit mention.

The case of $N = k$ in degree 0 follows at once because $H_0(M) = k \otimes_R M$.

The general case reduces to this. Consider an exact sequence of graded right R-modules $0 \to N_1 \to N_2 \to N_3 \to 0$; if the assertion holds for $N = N_1, N_3$, it holds also for $N = N_2$. The assertions are also unchanged by applying degree shifts to N. Now we proceed, by induction, on the largest degree a in which $N_a \neq 0$. Note that N_a is automatically an R-submodule, isomorphic as R-module to a sum of copies of k, and so the assertion is known for N_a. This gives, in particular, the case $a = 0$ of the induction; and for $a > 0$, we use the exact sequence $N_a \to N \to N/N_a$ and induction. $\square$

4.5. LEMMA. *Let $\bar{M}$ be a left graded $\bar{R}$-module. Then the degree of the module $\text{Tor}_i^{\bar{R}}(k, \bar{M})$ is at most $(\deg \bar{R})i + \deg(\bar{M})$.*

Proof. Let

$$\cdots \to P_2 \to P_1 \to P_0 \to k$$

be a resolution of k by projective right graded $\bar{R}$-modules. We note that P_i can be chosen to be generated in degree at most $i \deg(\bar{R})$. Indeed, this is so for $i = 0$, and we construct P_i by taking the free module on a set of generators for $\ker(P_{i-1} \to P_{i-2})$; by inductive hypothesis, P_{i-1} is supported in degree at most $\deg(\bar{R})i$. Using Lemma 4.4 and the fact that every $\bar{R}$-module is an

R-module, we have

$$\deg(\operatorname{Tor}_i^{\bar{R}}(k, \bar{M})) \leq \deg(P_i \otimes_{\bar{R}} \bar{M})$$
$$\leq \deg(H_0(P_i)) + \deg(\bar{M}) \leq (\deg \bar{R})i + \deg(\bar{M}). \qquad \square$$

In the lemmas that follow, we shall adopt the following notation for M a graded left R-module:

(4.5.1) $$A(M) = \max(\deg M[U], \deg M/UM),$$

(4.5.2) $$\delta(M) = \max\left(\deg \operatorname{Tor}_0^R(\bar{R}, M), \deg \operatorname{Tor}_1^R(\bar{R}, M)\right).$$

In both cases, we allow the value ∞ if the degrees in question are not finite. Note that $A(R) = \max(\deg R[U], \deg R/UR)$ is finite by virtue of Lemma 3.5.

We will also use several constants in the proofs that follow. We summarize them here for convenience:

(4.5.3) $$A_1 = A(R),$$
$$A_2 = A(R) + \deg(U),$$
$$A_0 = 5A_1 + A_2 = 6A(R) + \deg(U).$$

4.6. LEMMA. *Let M be a graded left R-module. Then*

(4.6.1) $$A(M) \leq \delta(M) + A_1,$$

where $A_1 = A(R)$ as in (4.5.3).

Proof. The bound on the degree of $M/UM \simeq \bar{R} \otimes_R M = \operatorname{Tor}_0^R(\bar{R}, M)$ is clear by definition of $\delta(M)$. It remains to bound $\deg M[U]$.

In what follows, UR denotes the two-sided ideal of R generated by U. Let us write $N = (UR) \otimes_R M$. Then N has the structure of a graded left R-module. Now we may regard multiplication by U as a map $M \xrightarrow{U} M$ of degree $\deg(U)$, which can be factored as

$$M \xrightarrow{\alpha} N \xrightarrow{\beta} M,$$

where α is the map $M = R \otimes_R M \to UR \otimes_R M$ given by $s \otimes m \mapsto Us \otimes m$, and $\beta(s \otimes m) = sm$; thus α is of degree $\deg(U)$ and β is of degree 0. Consequently, we get an exact sequence (i.e., exact at the $\ker(\alpha)$ and $M[U]$ terms)

(4.6.2) $$0 \to \ker(\alpha) \to M[U] \xrightarrow{\alpha} \ker(\beta).$$

By tensoring the short exact sequences $UR \hookrightarrow R \twoheadrightarrow \bar{R}$ and $R[U] \hookrightarrow R \twoheadrightarrow UR$ with M, we obtain

$$\ker(\beta) \simeq \operatorname{Tor}_1^R(\bar{R}, M) \text{ and } \ker(\alpha) \leftarrow M \otimes_R R[U].$$

This turns (4.6.2) into the following sequence, exact at the middle term

(4.6.3) $$R[U] \otimes_R M \to M[U] \to \operatorname{Tor}_1^R(\bar{R}, M),$$

where the first map is of degree 0 and the second map is of degree $\deg(U)$. We also observe that, by the definition (4.5.2) of $\delta(M)$, we have $\deg \operatorname{Tor}_1^R(\bar{R}, M) \leq \delta(M)$.

Now

$$\deg(R[U] \otimes_R M) \leq \deg(H_0(M)) + \deg(R[U])$$

by Lemma 4.4. We note that $\deg H_0(M) \leq \deg M/UM \leq \delta(M)$. Thus we arrive at the bound

$$\deg(M[U]) \leq \delta(M) + \deg(R[U])$$
$$\leq \delta(M) + A(R). \qquad \square$$

4.7. LEMMA. *Let M be a graded left R-module; then, with notation as above,*

$$(4.7.1) \qquad \deg \operatorname{Tor}_i^R(\bar{R}, M) \leq \delta(M) + A_1 i,$$

where $A_1 = A(R)$ as before.

Proof. (We thank the referee for the proof that follows, which greatly improves on the prior version.) Firstly, the assertion is clear for $i = 0$ and for $i = 1$ by definition of $\delta(M)$. We proceed now by induction on i.

Note that $R[U]$ is in fact a right $\bar{R}$-module. Construct a resolution of $R[U]$ by free right $\bar{R}$-modules:

$$(4.7.2) \qquad \cdots \to Q_2 \to Q_1 \to Q_0 \twoheadrightarrow R[U].$$

Here, we may suppose that Q_i is generated as a $\bar{R}$-module by elements of degree $\leq \deg(R[U]) + i \deg \bar{R} \leq (i+1)A_1$ by the argument of Lemma 4.5. Combine (4.7.2) with $R[U] \to R \xrightarrow{U} R \to \bar{R}$ to obtain a resolution of $\bar{R}$ by right R-modules

$$\cdots \to Q_2 \to Q_1 \to Q_0 \to R \xrightarrow{U} R \to \bar{R} \to 0.$$

Call this complex $P_\bullet$, so that $P_1 = P_0 = R$ and $P_i = Q_{i-2}$ when $i \geq 2$. We now have a hyperhomology sequence

$$E_{ij}^1 = \operatorname{Tor}_i^R(P_j, M) \implies \operatorname{Tor}_{i+j}^R(\bar{R}, M).$$

That shows that, for $b \geq 2$, $\operatorname{Tor}_b^R(\bar{R}, M)$ admits a filtration, whose associated graded is a subquotient of

$$(4.7.3) \qquad \bigoplus_{u=0}^{b-2} \operatorname{Tor}_u^R(Q_{b-u-2}, M).$$

By construction, Q_{b-u-2} is a free $\bar{R}$-module that was generated in degree at most $(b - u - 1) \cdot A_1$, so the degree of (4.7.3) — and so also of $\operatorname{Tor}_b^R(\bar{R}, M)$ —

is at most the maximum over $0 \le u \le b - 2$ of

$$(b - u - 1)A_1 + \deg \operatorname{Tor}_u^R(\bar{R}, M)$$
$$\le_{(i)} (b - u - 1)A_1 + \delta(M) + A_1 u$$
$$\le bA_1 + \delta(M),$$

where step (i) follows from the induction hypothesis, since $u < b$. $\square$

4.8. LEMMA. *Let $\bar{M}$ be any left $\bar{R}$-module; then $\deg(\bar{M}) \le \deg(k \otimes_{\bar{R}} \bar{M}) + \deg(\bar{R})$.*

Proof. This follows from a form of Nakayama's Lemma: Choose homogeneous elements $x_1, \dots$ of $\bar{M}$ projecting to a k-basis for $k \otimes_R \bar{M}$; the quotient $\mathcal{Q} = \bar{M} / \sum \bar{R}x_i$ is a graded $\bar{R}$-module and satisfies $k \otimes_{\bar{R}} \mathcal{Q} = 0$. We claim $\mathcal{Q}$ is zero; if not, let j be the smallest integer such that jth graded piece of $\mathcal{Q}$ is nonzero. The image of this graded piece in $k \otimes_{\bar{R}} \mathcal{Q}$ cannot be trivial, since this tensor product is the same as the quotient of $\mathcal{Q}$ by $\mathcal{Q}' := \ker(\bar{R} \to k) \cdot \mathcal{Q}$, and $\mathcal{Q}'$ is supported in degrees strictly greater than j. $\square$

4.9. LEMMA. *For any graded left R-module M, we have*

$$\delta(M) \le \max(\deg(H_0(M)), \deg(H_1(M))) + 4A_1.$$

Proof. Now

(4.9.1) $$\deg(\bar{R} \otimes_R M) \le \deg H_0(M) + \deg(\bar{R}),$$

by Lemma 4.4.

We must now bound $\deg \operatorname{Tor}_1^R(\bar{R}, M)$. The spectral sequence (4.3.1) gives an exact sequence

$$\operatorname{Tor}_2^{\bar{R}}(k, \bar{R} \otimes_R M) \to k \otimes_{\bar{R}} \operatorname{Tor}_1^R(\bar{R}, M) \to H_1(M),$$

and applying Lemma 4.8 to the left $\bar{R}$-module $\operatorname{Tor}_1^R(\bar{R}, M)$, we see that

(4.9.2) $$\deg \operatorname{Tor}_1^R(\bar{R}, M) \le \deg \bar{R} + \max(\deg H_1(M), \deg \operatorname{Tor}_2^{\bar{R}}(k, \bar{R} \otimes_R M)).$$

By Lemma 4.5 we have

(4.9.3a) $$\deg \operatorname{Tor}_2^{\bar{R}}(k, \bar{R} \otimes_R M) \le 2 \deg(\bar{R}) + \deg \bar{R} \otimes_R M$$

(4.9.3b) $$\le 3 \deg(\bar{R}) + \deg H_0(M),$$

where we used (4.9.1) again for (4.9.3b). Combining this with (4.9.2) and the fact that $\deg(\bar{R}) \le A_1$ yields the desired bound on $\deg \operatorname{Tor}_1^R(\bar{R}, M)$. $\square$

4.10. PROPOSITION. *Let M be a graded left R-module and N a graded right R-module. Then for all $i \ge 0$,*

$$\deg(\operatorname{Tor}_i^R(N, M)) \le \deg(N) + \max(\deg(H_0(M)), \deg(H_1(M))) + A_1 i + 4A_1.$$

Proof. Just as in the proof of Lemma 4.4, we reduce to the case where $N = k$ in degree 0 and can assume that $\deg H_0(M)$ is finite.

Now

$$\deg(\operatorname{Tor}_i^R(k, M)) \leq \max_{a+b=i} \deg \operatorname{Tor}_a^{\bar{R}}\left(k, (\operatorname{Tor}_b^R(\bar{R}, M))\right)$$

(4.10.1a)
$$\leq_{(i)} \max_{a+b=i} \left(a \deg(\bar{R}) + \deg \operatorname{Tor}_b^R(\bar{R}, M)\right)$$

(4.10.1b)
$$\leq_{(ii)} \max_{a+b=i} \left(a \deg(\bar{R}) + bA_1 + \delta(M)\right)$$

(4.10.1c)
$$\leq_{(iii)} A_1 i + \max(\deg(H_0(M)), \deg(H_1(M))) + 4A(R),$$

where we used Lemma 4.5 for step (i), Lemma 4.7 for step (ii), and Lemma 4.9 together with the fact $\deg(\bar{R}) \leq A_1$ for step (iii). $\qquad\square$

Having set up the basic bounds for the degrees over Tor groups over R and $\bar{R}$, we are now ready to analyze the cohomology of the complex $\mathcal{K}(M)$. As the notation $\mathcal{K}$ indicates, the complex $\mathcal{K}$ is an analog of the Koszul complex, and the following result is an analog of a well-known result in commutative algebra concerning Koszul complexes; see, for example, Theorem 16.4 and following discussion in [40].

4.11. LEMMA. *Each $H_q(\mathcal{K}(R))$ is killed by the right action of $R_{>0}$.*

Proof. For $s = (h_1, \ldots, h_n) \in S$, write $\partial s = h_1 \cdots h_n \in G$, and define

$$S_g(g_0, \ldots, g_q; s) = (g^{(g_0 \cdots g_q \partial s)^{-1}}, g_0, \ldots, g_q; s).$$

Extend S_g linearly to $\mathcal{K}(R)_{q+1} \to \mathcal{K}(R)_{q+2}$. By a routine computation,

$$(S_g d + d S_g)(g_0, \ldots, g_q; s) = (g_0, \ldots, g_q; r(g^{(\partial s)^{-1}})s) = (g_0, \ldots, g_q; sr_g),$$

and so "right multiplication by r_g" is homotopic to zero. $\qquad\square$

4.12. PROPOSITION. $\deg H_q(\mathcal{K}(R)) \leq A_2 + q$, where $A_2 = A(R) + \deg(U)$.

Proof. Multiplication by U induces an endomorphism of the complex $\mathcal{K}(R)$. (Recall that U is central, and thus it does not matter whether this multiplication is taken on the left or the right.) Indeed, in the diagram

$$\begin{array}{ccc}
\mathcal{K}(R)_q & \xrightarrow{d_q} & \mathcal{K}(R)_{q-1} \\
U\uparrow & & \uparrow U \\
\mathcal{K}(R)_q & \xrightarrow{d_q} & \mathcal{K}(R)_{q-1}
\end{array}$$

the vertical arrows induce isomorphisms in source degree $n > A_1 + q$. (Recall that $\mathcal{K}(R)_q$ is just a direct sum of copies of the shift $R[q]$.) This implies that

the map

$$U : \ker d_q \to \ker d_q$$

is an isomorphism in source degree $n > A_1 + q$. In particular, $\ker d_q$ is generated, as right R-module, in degree at most $A_1 + \deg U + q$, and the same is true for its quotient $H_q(\mathcal{K}(R))$. By Lemma 4.11, $H_q(\mathcal{K}(R))$ is killed by $R_{>0}$; it follows that

$$\deg(H_q(\mathcal{K}(R))) \leq A_2 + q,$$

where $A_2 = A_1 + \deg(U)$ is as in (4.5.3). □

4.13. Proposition. *Let M be a left graded R-module. Then*

$$\deg H_q(\mathcal{K}(M)) \leq \max(\deg(H_0(M)), \deg(H_1(M)) + A_1 q + (4A_1 + A_2)$$

for all $q \geq 0$.

Proof. We note that $\mathcal{K}(M) = \mathcal{K}(R) \otimes_R M$; then the "universal coefficients" spectral sequence

$$\mathrm{Tor}_i^R(H_{q-i}(\mathcal{K}(R)), M) \Rightarrow H_q(\mathcal{K}(M))$$

shows that $H_q(\mathcal{K}(M))$ is filtered by subquotients of $\mathrm{Tor}_i^R(H_{q-i}(\mathcal{K}(R)), M)$; the degree of this Tor-group is bounded by Propositions 4.10 and 4.12:

$$\deg \mathrm{Tor}_i^R(H_{q-i}(\mathcal{K}(R)), M)$$
$$\leq \max(\deg H_0(M), \deg H_1(M)) + (A_2 + q - i) + A_1 i + 4A_1.$$

This gives the result. □

4.14. *Proof of Theorem* 4.2. The last sentence has already been proved (Proposition 4.12).

We are going to show that

(4.14.1) $$\deg H_i(M) \leq \deg H_i(\mathcal{K}(M)) \quad (i = 0, 1);$$

then (recall that $h_i = \deg H_i(\mathcal{K}(M))$) the bound

$$h_q \leq \max(h_0, h_1) + (5A_1 + A_2)q$$

follows from Proposition 4.13. Since we are taking $A_0 = 5A_1 + A_2$ (see (4.5.3)), this will prove (4.2.1) in the theorem. As for the the assertion about $M \xrightarrow{U} M$, Lemma 4.6 asserts that U is an isomorphism in source degree at least $\delta(M) + A_1 + 1$; applying Lemma 4.9 and (4.14.1) that number is

$$\leq \max(h_0, h_1) + 4A_1 + A_1 + 1.$$

Since $A_0 \geq 5A_1 + 1$, we have proved that $M \xrightarrow{U} M$ is an isomorphism in source degree at least $\max(h_0, h_1) + A_0$ as desired.

Therefore, it remains to verify (4.14.1). The case $i = 0$ of (4.14.1) follows from the fact that $H_0(M) = H_0(\mathcal{K}(M))$.

For the H_1 inequality, we will factor the map $k[c] \otimes M[1] \to M$ from the final terms of the Koszul complex (4.1.1) as follows:

$$k[c] \otimes_k M[1] \xrightarrow{\alpha} R_{>0} \otimes_R M \xrightarrow{\beta} M,$$

where the first map α sends $g \otimes m$ to $r_g \otimes m$ for $g \in c$, and the second map β sends $i \otimes m$ to im. Now α is surjective since elements r_g generate $R_{>0}$ as a R-module. Also, α is degree-preserving because the degree of $g \otimes m$ in $k[c] \otimes_R M[1]$ equals $\deg(m) + 1$ whereas the degree of $r_g \otimes m$ in $R_{>0} \otimes_R M$ also equals $\deg(m) + 1$.

Let $\mathfrak{k}$ be the kernel of $\beta \circ \alpha : k[c] \otimes M[1] \to M$. Then there is a sequence

$$k[c^2] \otimes_k M[2] \xrightarrow{d} \mathfrak{k} \to H_1(\mathcal{K}(M)) \to 0,$$

which is exact at the middle and on the right. Here d is the differential from the $\mathcal{K}$-complex (4.1.1).

It follows that $\mathfrak{k}$ is generated as k-vector space by the image of $k[c^2] \otimes_k M[2] \xrightarrow{d} k[c] \otimes_k M[1]$, together with terms in degree at most $\deg H_1(\mathcal{K}(M))$.

But the composite map

$$k[c^2] \otimes_k M[2] \xrightarrow{d} k[c] \otimes_k M[1] \xrightarrow{\alpha} R_{>0} \otimes_R M$$

is *zero*: it sends $(g_1, g_2) \otimes m$ first to $g_1 \otimes r(g_2)m - g_2 \otimes r(g_1^{g_2})m$, and then to $(r(g_1)r(g_2) - r(g_2)r(g_1^{g_2})) \otimes m$, which is zero in $R_{>0} \otimes_R M$.

Therefore, $\alpha(\mathfrak{k})$ is spanned, as k-vector space by elements of degree $\leq \deg H_1(\mathcal{K}(M))$; since α was surjective, that means that $\ker(R_{>0} \otimes_R M \to M)$ is supported in degree $\leq \deg H_1(\mathcal{K}(M))$.

But by tensoring the exact sequence $R_{>0} \to R \to k$ with M, we find an isomorphism $H_1(M) \simeq \ker(R_{>0} \otimes_R M \to M)$. So we have shown $\deg H_1(M) \leq \deg H_1(\mathcal{K}(M))$. This is the case $i = 1$ of (4.14.1). $\qquad\square$

5. The arc complex

In this section, we shall prove, as previously promised, that the homology of Hurwitz spaces can be computed in terms of $\mathcal{K}$-complexes formed from homology of *smaller* Hurwitz spaces.

Define the graded left R-module

(5.0.2)
$$M_p = \bigoplus_n H_p(\mathrm{Hur}^c_{G,n}, k);$$

the R-module structure arises from the gluing on Hurwitz spaces (Section 2.6), and the grading is in the n-variable.

5.1. PROPOSITION. *There exists a homological spectral sequence E^1_{qp} converging to $H_{q+p}(\mathrm{Hur}^c_{G,n}, k)$ in dimensions $q + p < n - 2$. Moreover, each row (E^1_{*p}, d_1) is isomorphic to the nth graded piece of $\mathcal{K}(M_p)$; that is to say,*

$$E^1_{qp} = nth \; graded \; piece \; of \; \mathcal{K}(M_p)_{q+1}, \quad p, q \geq 0, \quad p + q < n - 2.$$

5.2.   This spectral sequence arises (in a similar way to [34], [35]) by considering the action of the braid group B_n on (the geometric realization of) a highly connected simplicial complex $\mathcal{A}$ (a variant of the *arc complex*), and filtering the complex by the dimension of simplices. We give the geometric construction of $\mathcal{A}$ in Section 5.5.

In the present subsection, we give a combinatorial model of $\mathcal{A}$. This will be helpful for proofs. Namely, we construct a semisimplicial set $\mathbf{A}$ (i.e., we describe a set $\mathbf{A}_q$ of q-simplices for each q and give consistent face maps $\partial_i : \mathbf{A}_q \to \mathbf{A}_{q-1}$), and it will follow from Proposition 5.6 that in fact the geometric realization of $\mathbf{A}$ is homeomorphic to that of $\mathcal{A}$.

Fix n. Let L_q be the subgroup of B_n (presented as in (1.3.1)) generated by $\sigma_{q+2}, \ldots, \sigma_{n-1}$; it is abstractly isomorphic, then, to B_{n-q-1}. Note that if $q \geq n-2$, we understand L_q to be the trivial group. Let $\mathbf{A}_q = B_n/L_q$ (as a B_n-set). Define the faces of the q-simplex bL_q by the formula

$$\partial_i(bL_q) = bs_{q,i}L_{q-1}, \ \ 0 \leq i \leq q,$$

where $s_{q,i} = \sigma_{i+1}\sigma_{i+2}\cdots\sigma_q$, and we interpret $s_{q,q} = 1$. Note that since L_q actually commutes with σ_j for $j \leq q$, this is independent of the choice of b representing the coset bL_q.

5.3. PROPOSITION. $\mathbf{A}$ *is a semisimplicial set*; *that is, the semisimplicial identity* $\partial_i\partial_j = \partial_{j-1}\partial_i$ *holds for* $i < j$.

Proof. This is a computation in the braid group. For $j = q$, it is the identity $bs_{q-1,i}L_{q-2} = bs_{q,i}L_{q-2}$, which follows because $s_{q,i} = s_{q-1,i}\sigma_q$ and $\sigma_q \in L_{q-2}$. For $j < q$,

$$
\begin{aligned}
\partial_i\partial_j(bL_q) &= b\sigma_{j+1}\cdots\sigma_q\sigma_{i+1}\cdots\sigma_q L_{q-2} \\
&= b\sigma_{i+1}\cdots\sigma_{j-1}\sigma_{j+1}\cdots\sigma_q\sigma_j\cdots\sigma_q L_{q-2} \\
&= b\sigma_{i+1}\cdots\sigma_{j-1}\sigma_{j+1}\sigma_j\sigma_{j+2}\cdots\sigma_q\sigma_{j+1}\cdots\sigma_q L_{q-2} \\
&=_{(a)} b\sigma_{i+1}\cdots\sigma_{j+1}\sigma_j\sigma_{j+1}^{-1}\sigma_{j+2}\cdots\sigma_q\sigma_{j+1}\cdots\sigma_q L_{q-2},
\end{aligned}
$$

where at step (a) the braid relation was used in the form

$$\sigma_{j+1}\sigma_j = (\sigma_j\sigma_{j+1}\sigma_j)\sigma_{j+1}^{-1}.$$

Define, for $m = j, \ldots, q-1$, the element

$$x_m := \sigma_{i+1}\cdots\sigma_{m+1}\sigma_j\cdots\sigma_m\sigma_{m+1}^{-1}\underbrace{\sigma_{m+2}\cdots\sigma_q}\sigma_{m+1}\cdots\sigma_q,$$

where the underbraced product is understood to be empty in the case when $m = q - 1$. The defining relations in the braid group yield the recursion $x_m = x_{m+1}$, so

$$\partial_i\partial_j(bL_q) = bx_jL_{q-2} = bx_{q-1}L_{q-2} = b\sigma_{i+1}\cdots\sigma_q\sigma_j\cdots\sigma_{q-1}L_{q-2} = \partial_{j-1}\partial_i(bL_q)$$

as desired. $\qquad\square$

It may be helpful to note that the set of vertices of the q-simplex bL_q is then given by

$$bL_0, b\sigma_1 L_0, b\sigma_2 \sigma_1 L_0, \ldots, b\sigma_q \cdots \sigma_1 L_0.$$

Here, the jth vertex is obtained from the iterated face map

$$b\sigma_j \cdots \sigma_1 L_0 = \partial_0 \circ \cdots \circ \partial_{j-1} \partial_{j+1} \circ \cdots \circ \partial_q (bL_q).$$

One can check that the lone vertex missing in the face $\partial_j(bL_q) = bs_{q,j}L_{q-1}$ is indeed $b\sigma_j \cdots \sigma_1 L_0$.

PROPOSITION. *The geometric realization of $\mathbf{A}$ is $(n-2)$-connected.*

This is a combinatorial reformulation of a result of Hatcher-Wahl. We give the proof in Section 5.5.

Now consider the complex $\mathbf{A} \times c^n$, considered with the product B_n-action. Recall that for any topological space Z endowed with a B_n-action, we write $Z /\!\!/ B_n$ for the Borel construction $EB_n \times_{B_n} Z$; the homology of this space is the B_n-equivariant homology of Z.

It follows from the proposition that the natural map

$$H_p((\mathbf{A} \times c^n) /\!\!/ B_n) \to H_p(c^n /\!\!/ B_n) = H_p(\mathrm{Hur}^c_{G,n})$$

is an isomorphism in degrees $p < n - 2$. Note further that

$$H_*((\mathbf{A}_q \times c^n) /\!\!/ B_n) \cong H_*(c^n /\!\!/ L_q) \cong H_*(c^{q+1} \times \mathrm{Hur}^c_{G,n-q-1}),$$

so $H_p((\mathbf{A}_q \times c^n) /\!\!/ B_n)$ is identified with $\mathcal{K}(M_p)_{q+1}$.

Now, $(\mathbf{A} \times c^n) /\!\!/ B_n$ is filtered by the simplicial structure on $\mathbf{A}$. The resulting spectral sequence is of the form

$$H_p((\mathbf{A}_q \times c^n) /\!\!/ B_n) = E^1_{qp} \implies H_{p+q}((\mathbf{A} \times c^n) /\!\!/ B_n),$$

and the target is isomorphic to $H_{p+q}(\mathrm{Hur}^c_{G,n})$ when $p + q < n - 2$.

5.4. LEMMA. *The differential d_1 in the spectral sequence $E^1_{qp} = H_p(c^n /\!\!/ L_q)$ is the alternating sum*

$$d_1 = \sum_{i=0}^q (-1)^i \partial_i = \sum_{i=0}^q (-1)^i [s_{q,i}^{-1}]$$

of the maps $[s_{q,i}^{-1}] : H_*(c^n /\!\!/ L_q) \to H_*(c^n /\!\!/ L_{q-1})$ induced by $s_{q,i}^{-1} : c^n \to c^n$:

$$s_{q,i}^{-1} : (g_0, \ldots, g_q, g_{q+1}, \ldots, g_{n-1})$$
$$\mapsto (g_0, \ldots, \widehat{g_i}, g_{i+1}, \ldots, g_q, g_i^{g_{i+1}\cdots g_q}, g_{q+1}, \ldots, g_{n-1}).$$

Compare with (4.1.1) to get Proposition 5.1.

Proof. The face maps in the semi-simplicial space $EB_n \times_{B_n} (\mathbf{A} \times c^n)$ are induced by those in $\mathbf{A}$, as above. Identifying $\mathbf{A}_q = B_n/L_q$, this becomes $\partial_i : EB_n \times_{L_q} c^n \to EB_n \times_{L_{q-1}} c^n$ given by

$$\partial_i(e, \mathbf{g}) = (es_{q,i}, \ s_{q,i}^{-1}\mathbf{g}),$$

where $0 \leq i \leq q$ and $\mathbf{g} \in c^n$.

Since L_q is a subgroup of B_n, we have a natural identification $H_*(c^n /\!/ L_q) = H_*(EB_n \times_{L_q} c^n)$. Since $s_{q,i}$ and L_q commute, the maps $EB_n \times_{L_q} c^n \to EB_n \times_{L_{q-1}} c^n$, $(e, \mathbf{g}) \mapsto (es_{q,i}, \ s_{q,i}^{-1}\mathbf{g})$, and $(e, \mathbf{g}) \mapsto (e, \ s_{q,i}^{-1}\mathbf{g})$ are freely homotopic, giving the result. $\square$

5.5. *The arc complex.* We now prove that the geometric realisation of the complex $\mathbf{A}$ defined in the previous section is indeed $(n-2)$-connected, by identifying it with a geometric construction (the "arc complex") of Hatcher and Wahl.

Let Σ be, as in Section 2.1, an n-punctured disc. By an *arc* on Σ we mean a smooth path τ in D from $*$ to one of the punctures P_j with the following properties:

- τ is a smooth embedding $\tau : [0,1] \to D$ satisfying $\tau(0) = * \in \partial\Sigma$ and $\tau(1) = P_j$;
- 0 is the only element of the domain carried to the boundary of Σ, and 1 the only element carried into the set $\{P_i\}$ of punctures;
- the derivative $\tau'(0)$ is not tangent to the boundary.

The *arc complex* $\mathcal{A}$ is the simplicial complex whose vertex set consists of all isotopy classes of such arcs, and whose faces are collections of such isotopy classes that have representatives that intersect only at $*$, where they have distinct tangent vectors. We denote by $\mathcal{A}_q$ the q-simplices of $\mathcal{A}$, i.e., the set of isotopy classes of $(q+1)$-tuples of arcs intersecting only at $*$.

We may define a partial ordering on the set of vertices: two arcs are comparable if they span an edge; then the ordering of the pair is given by the counterclockwise ordering of their tangent vectors at $*$. This extends naturally to a total ordering on the set of vertices spanning a face; in this way, $\mathcal{A}$ becomes an ordered simplicial complex. This ordering equips the collection of sets $\mathcal{A}_*$ with the structure of a semisimplicial set: The face maps $d_i : \mathcal{A}_q \to \mathcal{A}_{q-1}$ ($i = 0, \ldots, q$) are defined by

$$d_i(\gamma_0, \ldots, \gamma_q) = (\gamma_0, \ldots, \widehat{\gamma_i}, \ldots, \gamma_q),$$

where $\gamma_0, \ldots, \gamma_q$ are arranged in increasing order. In other words, the semisimplicial structure is given by deletion of arcs.

5.6. PROPOSITION. $\mathcal{A}$ is $(n-2)$-connected. There exists an action of B_n on $\mathcal{A}$ that is transitive on q-simplices for each q. Moreover, $\mathcal{A}$ and $\mathbf{A}$ are B_n-equivariantly isomorphic.

Proof. The connectivity assertion is a special case of Proposition 7.2 (and Definition 3.4) of Hatcher and Wahl [35]. (In fact, $\mathcal{A}$ is contractible — we refer the reader to the recent [15] for a careful proof of this fact.) Let Diff_n be the group of orientation-preserving diffeomorphisms of D that fix the boundary pointwise and fix the set of punctures (setwise). Then Diff_n acts on the set of isotopy classes of arcs on Σ (i.e., the vertices of $\mathcal{A}$). In fact, Diff_n acts on the whole of the complex $\mathcal{A}$, since diffeomorphisms preserve the nonintersection condition. This evidently descends to an action of $B_n = \pi_0(\text{Diff}_n)$ on $\mathcal{A}$.

Now the transitivity of B_n, as well as the final assertion — which amounts to a computation of the stabilizer of a simplex, see (5.6.1) below — are very similar to [48, Prop. 2.2], which proves precisely the same result in the context that the *endpoints* of all the arcs also coincide. The proof of [48, Prop. 2.2] also applies in this context. For completeness, we will recall the main steps of this proof below.

Transitivity: We will be particularly brief about transitivity, since it is the easier part. Given two q-simplices $\overline{\tau}$ and $\overline{\tau}'$, we choose systems of $q+1$ nonintersecting arcs $(\tau_0, \ldots, \tau_q)$ and $(\tau_0', \ldots, \tau_q')$ representing them. While the proof of [48, Prop. 2.2] transcribes almost verbatim to this context, we will sketch an alternate direct argument.

Replacing the τ_i, τ_i' by isotopic curves, and using the fact that the ordering of τ_i, τ_i' coincide near $*$, we may suppose that all τ_i and τ_i' are linear in a small neighbourhood of 0 (i.e., straight near $*$), and that $\tau_i(t) = \tau_i'(t)$ for $0 \leq t \leq \varepsilon$ and some $\varepsilon > 0$. Applying a suitable diffeomorphism, and using the fact that the braid group B_n surjects to S_n, we may assume that $\tau_i(1) = \tau_i'(1)$ for each i. Fix a smooth increasing function h on $[0,1]$ such that $h(t) = t$ for $t \leq 2\varepsilon/3$ and $h(1) = \varepsilon$. Then $\tau_i \circ h = \tau_i' \circ h$.

It is now sufficient to show that there exists an "arc-retracting" diffeomorphism, i.e., a diffeomorphism F that carries τ_i to the arc $\tau_i \circ h$ while fixing all punctures P_i that are not of the form $\tau_i(1)$ $(0 \leq i \leq q)$. Once this is done, and a similar diffeomorphism F' constructed for the τ_i', then $F' \circ F^{-1}$ gives the desired diffeomorphism carrying τ_i to τ_i' and fixing all remaining punctures. Let R be the rectangle $[\varepsilon/2, 1.01] \times [-1, 1]$. Choosing a normal vector field to each τ_i, we may find a collection of embeddings $g_i : R \to D$, carrying $(t, 0)$ to $\tau_i(t)$ for $\frac{\varepsilon}{2} \leq t \leq 1$, and such that the $g_i(R)$ are pairwise disjoint as well as disjoint from all other punctures — in other words, $g_i(R)$ is an explicit tubular neighbourhood of $\tau_i|_{[\varepsilon/2, 1]}$. Our assertion is reduced to the claim that there exists a diffeomorphism $G : R \to R$, trivial in a neighbourhood of the boundary,

and carrying $(x, 0)$ to $(h(x), 0)$ for $\varepsilon/2 \le x \le 1$. It is routine to write down such a diffeomorphism explicitly.

Computation of stabilizer: Order the punctures $P_1, \ldots, P_n$ in such a way that the straight line segments $[*, P_i]$ from $*$ to P_i are in counterclockwise order around $*$. The *standard q-simplex* v_q is the one consisting of linear arcs from $*$ to $P_1, \ldots, P_{q+1}$. We claim that the stabilizer of v_q in B_n is L_q.

Granting that, we define a B_n-equivariant semisimplicial map $\mathbf{A} \to \mathcal{A}$ on q-simplices by the bijection

(5.6.1) $$B_n/L_q \to \mathcal{A}_q \quad \text{given by} \quad b \mapsto b \cdot v_q.$$

To see that this is semisimplicial, use the fact that $s_{q,i} v_{q-1} = d_i(v_q)$. (Both consist of the q straight line segments from $*$ to $P_1, \ldots, \widehat{P_{i+1}}, \ldots, P_{q+1}$.) The proposition then follows.

Therefore, it remains only to check that the stabilizer of v_q in B_n is L_q. Now, L_q is generated by the Dehn twists σ_i $(i \ge q + 2)$ that involve only the punctures P_i and P_{i+1}, so it is apparent that L_q does in fact stabilize v_q. Additionally, there is a map from the mapping class group

$$\mathrm{Mod}(\Sigma \setminus v_q) \to \mathrm{Mod}(\Sigma) = B_n$$

that extends diffeomorphisms by the identity on a neighborhood of v_q. Write Σ' for the complement of $n - q - 1$ points in D. Then $\mathrm{Mod}(\Sigma')$ is isomorphic to B_{n-q-1}. Picking a diffeomorphism $\Sigma' \to \Sigma \setminus v_q$ defines an isomorphism $B_{n-q-1} \cong \mathrm{Mod}(\Sigma \setminus v_q)$. In fact, we may choose this diffeomorphism so that the composition

$$B_{n-q-1} = \mathrm{Mod}(\Sigma') \cong \mathrm{Mod}(\Sigma \setminus v_q) \to \mathrm{Mod}(\Sigma) = B_n$$

carries $\sigma_j \in B_{n-q-1}$ to $\sigma_{j+q+1} \in B_n$. The image of this map is precisely L_q.

It therefore suffices to show that if $b \in B_n$ stabilizes v_q, then b is isotopic to a diffeomorphism that is in the image of this map $\mathrm{Mod}(\Sigma \setminus v_q) \to \mathrm{Mod}(\Sigma)$. This can be done as in [48, pp. 552–553]: Take a diffeomorphism ϕ representing b. *A priori* b fixes only the isotopy class of each arc in v_q. We may replace ϕ by an isotopic ϕ' that fixes the arcs in v_q pointwise, inductively using the isotopy extension theorem [42]. We briefly summarize the key point of the proofs of the proof in our setting, closely following [48] but indicating the minor differences.

A direct application of the isotopy extension theorem implies that ϕ is isotopic to a diffeomorphism that fixes the (linear) arc a_1 from $*$ to P_1. Replace ϕ by this isotopic diffeomorphism. Now the arc a_2 is isotopic to $\phi(a_2)$, by assumption, but the image of this isotopy need not be disjoint from $\phi(a_1) = a_1$, so one cannot simply proceed as in the first step. Let $H : [0, 1] \times [0, 1] \to D$ be an isotopy, so that $H(0, -)$ and $H(1, -)$ correspond to the arcs a_2 and $\phi(a_2)$ respectively, and $H(-, 0) = *, H(-, 1) = P_2$. Adjust H to be transverse

to a_1. The preimage $H^{-1}(a_1)$ of the arc a_1 under H is a union of circles and intervals. Note that these components intersect the boundary of $[0,1] \times [0,1]$ only along $[0,1] \times 0$. This is a difference from the setting of [48], and means that we can ignore the case of a "nonempty union of intervals" mentioned at the bottom of page 552, loc. cit. In the present setting, every component of $H^{-1}(a_1)$ other than $[0,1] \times 0$ is contained in the interior of $[0,1] \times [0,1]$ and is thus a circle. So each component of $H^{-1}(a_1)$ bounds a disk in $[0,1] \times [0,1]$; restricting H to this disk defines an element of the relative homotopy group $\pi_2(\Sigma, a_1)$, and the triviality of that group implies that we can assume that H can be replaced by H' for which $(H')^{-1}(a_1) = (H')^{-1}(*)$. At this point, we can again use the isotopy extension theorem to ensure that ϕ fixes a_2, and we proceed inductively. $\qquad\square$

6. Homological stability for Hurwitz spaces

We now prove the main theorem of the paper, that the homology of Hurwitz spaces stabilize under the nonsplitting condition; it is by now an easy consequence of the main results of the prior three sections.

6.1. THEOREM. *Suppose (G,c) satisfies the nonsplitting condition, and let k be a field in which $|G|$ is invertible. Then there exist constants A, B (depending on G) such that the map*

$$(6.1.1) \qquad U : H_p(\mathrm{Hur}^c_{G,n}, k) \to H_p(\mathrm{Hur}^c_{G,n+\deg U}, k)$$

(see Lemma 3.5 and the start of Section 4 for definition of U) is an isomorphism whenever $n > Ap + B$.

The same assertion holds for the restricted maps

$$(6.1.2) \qquad U : H_p(\mathrm{CHur}^c_{G,n}, k) \to H_p(\mathrm{CHur}^c_{G,n+\deg U}, k).$$

Our proof will actually give a range $n > A'p + B'$ with different constants for the restricted maps (6.1.2) of the second assertion, but one can then simply replace A by $\max(A, A')$ and similarly for B.

Proof. Let M_p be the graded R-module corresponding to the pth homology of Hurwitz spaces, taken with k-coefficients, as defined in (5.0.2). We prove, by increasing induction on p, that for all $q \geq 0$, we have

$$\deg H_q(\mathcal{K}(M_p)) \leq A_2 + A_0(3p + q)$$

from which the result follows, for suitable B, by the second assertion in Theorem 4.2.

For $p = 0$, we have $M_0 = R$; the inductive assumption follows in this case from Proposition 4.12, since

$$A_2 + A_0(3p + q) = A_2 + A_0 q \geq A_2 + q$$

in this case.

Now suppose the statement holds for $p < P$. Consider the leftmost part of $\mathcal{K}(M_P)$, i.e.,

$$(6.1.3) \qquad M_P \xleftarrow{e} k[c] \otimes M_P[1] \xleftarrow{f} k[c^2] \otimes M_P[2].$$

The map e is an edge morphism in the spectral sequence of Proposition 5.1, whereas f is identified with the differential $d_1 : E^1_{1P} \to E^1_{0P}$. More generally, in the spectral sequence for $\mathrm{Hur}^c_{G,n}$, we have for each $q > 0$ that $E^2_{q,p}$ is the nth graded piece of $H_{q+1}(\mathcal{K}(M_p))$.

The inductive hypothesis implies that (6.1.3) is exact at the middle and left terms in degrees greater than

$$A_2 + 3A_0 P.$$

To see this, we note that the inductive hypothesis ensures that for $j > 1$, $E^2_{j,P+1-j}$, or in other words the nth graded piece of $H_{j+1}(\mathcal{K}(M_{P+1-j}))$, vanishes in degrees above

$$(6.1.4) \qquad A_2 + A_0(3P + 4 - 2j).$$

Thus, once $n > A_2 + 3A_0 P$, there are no differentials d_j for any $j > 1$ going into or out of E^2_{0P}. Thus, $E^2_{0P} = E^\infty_{0P}$ in the spectral sequence for $H_P(\mathrm{Hur}^c_{G,n}, k)$ in the range $n > A_2 + 3A_0 P$.

Further, for $j > 0$, $E^2_{j,P-j} = 0$ vanishes in degrees above

$$(6.1.5) \qquad A_2 + A_0(3P + 1 - 2j).$$

So, once $n > A_2 + A_0(3P - 1)$, we have that all the graded pieces of $H_P(\mathrm{Hur}^c_{G,n}, k)$ other than $E^\infty_{0P} = E^2_{0P}$ vanish. Hence, $e : \mathrm{coker}(f) = E^2_{0P} \to M_P$ is an isomorphism in degrees above $A_2 + 3A_0 P$.

In other terms,

$$\deg H_0(\mathcal{K}(M_P)), \ \deg H_1(\mathcal{K}(M_P)) \leq A_2 + 3A_0 p.$$

Now apply Theorem 4.2; it implies that, for $q \geq 2$, $H_q(\mathcal{K}(M_P))$ vanishes in degrees strictly above $A_2 + 3A_0 p + A_0 q$, which is precisely the inductive hypothesis to be proved.

Now we address the final assertion of the theorem — namely, that the same result holds for the space of connected covers.

For each subgroup Q of G that is generated by $c \cap Q$, write $\mathrm{Hur}^{Q,c}_{G,n}$ for the union of connected components of $\mathrm{Hur}^c_{G,n}$ whose global monodromy is exactly equal to Q — i.e., arising from elements $(x_1, \ldots, x_n) \in c^n$ with $\langle x_1, \ldots, x_n \rangle = Q$. Note that we can identify $\mathrm{Hur}^{Q,c}_{G,n}$ with $\mathrm{CHur}^{c \cap Q}_{Q,n}$.

What we have proved so far in this theorem applies equally well with G replaced by Q, for (Q, c) is still nonsplitting. By increasing induction on the order of $|Q|$ — and possibly increasing the constants A, B — we may suppose that

$$U_Q : H_p(\mathrm{Hur}^{Q,c}_{G,n}, k) \to H_p(\mathrm{Hur}^{Q,c}_{G,n+\deg(U)}, k)$$

is an isomorphism for $n \geq Ap + B$ for every proper subgroup Q. Here U_Q is defined similarly to U, but "relative to Q," that is to say, $U_Q = \sum_{g \in c \cap Q} r_g^{D|g|}$. Note that on $H_p(\mathrm{Hur}_{G,n}^{Q,c})$, we have

(6.1.6) $$U = \frac{|c|}{|c \cap Q|} U_Q \text{ modulo } \bigoplus_{Q' \supsetneq Q} H_p(\mathrm{Hur}_{G,n+\deg(U)}^{Q',c}, k).$$

What we have already proved shows that

$$U : H_p(\mathrm{Hur}_{G,n}^{G,c}, k) \to H_p(\mathrm{Hur}_{G,n+\deg(U)}^{G,c}, k)$$

is injective for $n \geq Ap + B$, so it remains to verify surjectivity in that same range. Take $x \in H_p(\mathrm{Hur}_{G,n+\deg(U)}^{G,c}, k)$; we may write $x = Uy$ for some $y \in H_p(\mathrm{Hur}_{G,n}^{c}, k)$. Write $y = \sum_Q y_Q$, where $y_Q \in H_p(\mathrm{Hur}_{G,n}^{Q,c}, k)$.

By an increasing induction on the size of Q, we see that $y_Q = 0$ if $Q \neq G$. This is obviously true when $|Q| = 1$. Next, if $m < |G|$ and we know the assertion is true for all $|Q| < m < |G|$, then take Q of size m. By (6.1.6) we get

$$0 = Q\text{-component of } Uy = \frac{|c|}{|c \cap Q|} U_Q y_Q,$$

and by inductive assumption U_Q is an isomorphism. Therefore $y_Q = 0$.

This induction on the size of Q has shown that $y \in H_p(\mathrm{Hur}_{G,n}^{G,c}, k)$, and we are done. □

In the arithmetic applications to follow, we will be concerned with the quotients $\mathrm{Hur}_{G,n}^{c} / G$ under the G-action introduced previously (see page 740, after Definition 2.2). These spaces are easily seen to stabilize in homology as well.

6.2. COROLLARY. *Suppose that G is center-free and (G, c) satisfies the nonsplitting condition. Let k, U, A, B be as in Theorem 6.1. Then the map*

$$U : H_p(\mathrm{Hur}_{G,n}^{c} / G, k) \to H_p(\mathrm{Hur}_{G,n+\deg U}^{c} / G, k)$$

is an isomorphism whenever $n > Ap + B$. This restricts to an isomorphism $U : H_p(\mathrm{CHur}_{G,n}^{c} / G, k) \to H_p(\mathrm{CHur}_{G,n+\deg U}^{c} / G, k)$ in the same range.

Proof. Recall that the action of G on $\mathrm{Hur}_{G,n}^{c}$ is induced by the B_n-equivariant action of G on c^n given by termwise conjugation. This is evidently not free, so the same is true for the action of G on $\mathrm{Hur}_{G,n}^{c}$. However, it is apparent from this description that the stabilizers G_x of points $x \in \mathrm{Hur}_{G,n}^{c}$ are locally constant. Namely, when x lies on a component whose global monodromy is H, the stabilizer G_x is the centralizer of H in G. In particular, since G is center-free by hypothesis, the action of G on $\mathrm{CHur}_{G,n}$ is free. Thus the quotient $q : EG \times_G \mathrm{Hur}_{G,n}^{c} \to \mathrm{Hur}_{G,n}^{c} / G$ is a fibre bundle with fibre over the class of x equivalent to BG_x. Since the order of G (and hence G_x) is invertible in k, q is an isomorphism in $H_*(-, k)$, so $H_*(\mathrm{Hur}_{G,n}^{c} / G, k)$ is isomorphic to the G-coinvariants in $H_*(\mathrm{Hur}_{G,n}^{c}, k)$.

The operator U, considered as a class in $H_0(\mathrm{Hur}^c_{G,\deg U})$, is fixed by the action of G (at least if U is chosen as in Lemma 3.5). It then follows from the G-equivariance of the gluing maps that the isomorphism in Theorem 6.1 is one of $k[G]$-modules. Taking G-coinvariants on both sides yields the desired isomorphism.   □

7. **Homological stability for Hurwitz schemes**

So far, we have considered Hurwitz spaces as topological spaces parametrizing continuous branched covers of the disc. In order to apply our theorems to arithmetic questions, we need to identify those topological spaces with the complex points of moduli schemes ("Hurwitz schemes") defined over arithmetic bases.

In order to apply the topological results of the first part of the paper to counting problems over finite fields, we will need to show that the étale cohomology of Hurwitz schemes is "the same" in characteristic 0 and characteristic p. This would follow immediately if the Hurwitz schemes were smooth and proper; since they are not proper, more work is required, involving the use of a compactification whose boundary has normal crossings.

Throughout this section, G denotes a finite group. In our main application, G will be of the form $A \rtimes \mathbb{Z}/2\mathbb{Z}$, where A is a nontrivial finite ℓ-group, and $\mathbb{Z}/2\mathbb{Z}$ acts on A by inversion.

For the reader's convenience, in Table 1 we summarize the various spaces (more precisely, schemes over $\mathrm{Spec}(\mathbb{Z})$ or schemes over $\mathrm{Spec}(R), R := \mathbb{Z}[\frac{1}{|G|}]$) to be introduced. For example, Conf'_n denotes the configuration space of n distinct points on $\mathbb{P}^1$, and $\mathsf{H}_{G,n}$ will denote the Hurwitz space of tame G-covers of $\mathbb{P}^1$ (see Section 7.1) that are branched at n points.

space	configurations of ...	Hurwitz space of G-covers with corresponding branching
Conf'_n (Section 7.1)	n points on $\mathbb{P}^1$	$\mathsf{H}_{G,n}$ (Section 7.1)
Conf_n (Section 7.3)	n points on $\mathbb{A}^1$	$\mathsf{Hn}_{G,n}$ (Section 7.3)
PConf_n (Section 7.5)	n *labelled* points on $\mathbb{A}^1$	$\mathsf{PHn}_{G,n}$ (p.769)

Table 1.

Note that a G-cover "branched at n points on $\mathbb{A}^1$" may or may not be branched at ∞, i.e., it may have either n or $n+1$ branch points on $\mathbb{P}^1$, and the associated Hurwitz spaces will be correspondingly disconnected.

We write (e.g.) $X/\mathrm{Spec}\,\mathbb{Z}$ or simply $X/\mathbb{Z}$ to denote a scheme X equipped with its canonical morphism to $\mathrm{Spec}(\mathbb{Z})$. For such a scheme, we will denote by $X/\mathbb{F}_q$ the scheme $X \times_{\mathrm{Spec}(\mathbb{Z})} \mathrm{Spec}(\mathbb{F}_q)$, together with the natural morphism to $\mathrm{Spec}(\mathbb{F}_q)$.

7.1. Our basic reference for Hurwitz schemes is the paper of Romagny and Wewers [44].

If $P_n/\operatorname{Spec}\mathbb{Z}$ is the projective n-space with coordinates $a_0 : a_1 : \ldots : a_n$, which we think of as parametrizing binary forms $a_0 X^n + a_1 X^{n-1}W + \cdots + a_n W^n$ up to scaling, we define $\mathsf{Conf'}_n/\operatorname{Spec}\mathbb{Z}$ to be[9] the open subscheme of P_n where the discriminant $\Delta(\sum_{i=0}^n a_i X^{n-i}W^i)$ does not vanish. Then, for any field k, $\mathsf{Conf'}_n(k)$ is the set of squarefree k-rational degree-n divisors on $\mathbb{P}^1$.

Let k be a field. Then by a *tame G-cover of $\mathbb{P}^1$ over k* we shall mean a triple (X, f, ϕ), where

- X is a smooth proper geometrically connected curve X/k;
- $f : X \to \mathbb{P}^1$ is *tame*: that is, there exists a reduced divisor D on $\mathbb{P}^1$ such that f is étale over $\mathbb{P}^1 - D$, and such that the ramification of f over each geometric point of D is nontrivial and prime to the characteristic of k;
- f is *Galois*: that is, $\operatorname{Aut}(f)$ acts transitively on the geometric fibers of f;
- ϕ is an isomorphism from G to $\operatorname{Aut}(f)$.

We say a tame G-cover has n branch points if n is the degree of the unique divisor D satisfying the condition in the second item above, and we call such a D the *branch locus* of f.

Romagny and Wewers [44, Th. 4.11] construct a scheme $\mathsf{H}_{G,n}$ with the following properties:

- $\mathsf{H}_{G,n}$ is a scheme over $R = \mathbb{Z}[\frac{1}{|G|}]$, endowed with a finite étale morphism

$$\pi : \mathsf{H}_{G,n} \to \mathsf{Conf'}_n/\operatorname{Spec} R.$$

- For every algebraically closed field k with characteristic prime to $|G|$, there is an $\operatorname{Aut}(k)$-equivariant bijection between $\mathsf{H}_{G,n}(k)$ and the set of isomorphism classes of tame G-covers of $\mathbb{P}^1$ over k with n branch points. If $x \in \mathsf{H}_{G,n}(k)$ corresponds to a G-cover f, the image of x in $\mathsf{Conf'}_n(k)$ is the point parametrizing the branch locus of f in $\mathbb{P}^1$. If G is center-free, the two statements above hold for an arbitrary field k, not only algebraically closed fields.

In fact, Romagny and Wewers construct a Hurwitz scheme over $\operatorname{Spec}\mathbb{Z}$, but for our present purpose we only need the open subscheme lying over $\operatorname{Spec} R$.

The above assertions are all included in the statement of [44, Th. 4.11, Cor. 4.12], with the exception of the finiteness of π, which is [44, Rem. 4.15(ii)]. Because the proof is not given in full there, we explain it briefly. Each geometric fiber of the map from $\mathsf{H}_{G,n}$ to $\mathsf{Conf'}_n/R$, say above the section $s : \operatorname{Spec} k \to \mathsf{Conf'}_n/R$ with k algebraically closed, is in bijection with the

---

[9] The prime in $\mathsf{Conf'}_n$ is because it represents a configuration space for the *projective* line, whereas our previous Conf_n is homotopy equivalent to a configuration space for the affine line.

set of G-covers f over k with a fixed branch locus $D \subset \mathbb{P}^1$ parametrized by s. This set of G-covers, in turn, is in bijection with the set of conjugacy classes of surjective homomorphisms from $\pi_1^{et}(\mathbb{P}^1 - D, \bar{x})$ to G, where $\bar{x}$ is a geometric basepoint on $\mathbb{P}^1 - D$. Because $|G|$ is prime to the characteristic of k, this number of surjections is actually independent of the choice of the fiber s; see [31, Exp. XIII, Cor. 2.12]. So all geometric fibers of π have the same cardinality, whence π is finite by [22, Lemme 1.19, II].

7.2. *Remark.* A discussion of the functor represented by $\mathsf{H}_{G,n}$, and also a construction of a representing stack in the case when G has nontrivial center, is given in Wewers' (unpublished) thesis [50] (see also [44, Cor 2.2]). We do not need this for our purposes.

7.3.   The scheme $\mathsf{H}_{G,n}$ is not exactly the right one for our purposes. First of all, we want to study G-covers of $\mathbb{A}^1$, not of $\mathbb{P}^1$. To this end, let Conf_n be the closed subscheme of $\mathsf{Conf'}_{n+1}$ cut out by $a_0 = 0$; that is, we force the associated divisor to contain ∞ (i.e., the point $X = 1$, $W = 0$). Conf_n also embeds as an open subscheme of $\mathsf{Conf'}_n$ via

(7.3.1) $$[a_0 : \cdots : a_{n+1}] \mapsto [a_1 : \cdots : a_{n+1}]$$

identifying it with the open subscheme where the first coordinate is nonzero.

We now define $\mathsf{Hn}_{G,n}$ to be the disjoint union of $\mathsf{H}_{G,n+1} \times_{\mathsf{Conf'}_{n+1}} \mathsf{Conf}_n$ and $\mathsf{H}_{G,n} \times_{\mathsf{Conf'}_n} \mathsf{Conf}_n$, where the maps from Conf_n are the closed inclusion (respectively, the open inclusion) described above. This somewhat convoluted definition is necessary because in our topological definition of Hurwitz spaces, the branched cover may be either ramified at ∞ or not, these two cases corresponding to the two fiber products above. In the language of the first section of the paper, the second component parametrizes those branched covers whose boundary monodromy is trivial.

Suppose that c is a rational union of conjugacy classes in G. (Recall that a union of conjugacy classes is called *rational* when $g \in c \implies g^N \in c$ for each N relatively prime to the order of G.) We say a tame G-cover $f : X \to \mathbb{P}^1$ has monodromy of type c if the image of a generator of tame inertia at each branch point of f other than ∞ lies in c. (Because c is rational, the choice of generator does not matter.) Then there is a closed and open subscheme $\mathsf{Hn}^c_{G,n} \subset \mathsf{Hn}_{G,n}$ parametrizing tame G-covers with monodromy of type c; in the proof of [44, Th. 4.11], this appears as a disjoint union of some subset of the collection of schemes denoted there as $\mathcal{H}'_\mu$.

7.4. LEMMA. *Suppose G is center-free. The complex manifold $\mathsf{Hn}_{G,n}(\mathbb{C})$ is homeomorphic to the topological space $\mathrm{CHur}_{G,n}/G$ of Definition 2.2; similarly, $\mathsf{Hn}^c_{G,n}(\mathbb{C})$ is homeomorphic to $\mathrm{CHur}^c_{G,n}/G$.*

Proof. See [44, Th. 4.11(iii)]. We explicate the map in our case: A $\mathbb{C}$-point of $\mathsf{Hn}_{G,n}(\mathbb{C})$ is by definition a point of Conf_n — that is to say, a subset $S \subset \mathbb{C}$ of size n — together with a tame G-cover of $\mathbb{P}^1_{\mathbb{C}}$ branched either at S or at $S \cup \{\infty\}$.

By comparison of étale and topological π_1 [31, Exp. XII, Th. 5.1], and also using the same reasoning used to construct a bijection of (b) and (c) of Section 2.3, to give such a tame G-cover is the same as giving a conjugacy class of surjections $\pi_1(\mathbb{A}^1(\mathbb{C}) - S, x_0) \twoheadrightarrow G$. Here, x_0 is an arbitrarily chosen basepoint.

Fix, once and for all, a homeomorphism of the interior of D with $\mathbb{A}^1(\mathbb{C})$. This allows us to identify any subset $S \subset \mathbb{C}$ as before with a subset $S' \subset D$ of size n; i.e., it induces an identification $\mathsf{Conf}_n(\mathbb{C}) \simeq \mathrm{Conf}_n$. Moreover, we have an identification, canonical up to conjugacy,

$$\pi_1(\mathbb{A}^1(\mathbb{C}) - S, x_0) \simeq \pi_1(D - S', *).$$

In particular, there is a canonical identification of conjugacy classes of surjections $f : \pi_1(\mathbb{A}^1(\mathbb{C}) - S, x_0) \twoheadrightarrow G$ with conjugacy classes of surjections $g : \pi_1(D - S', *) \twoheadrightarrow G$. For such a surjection g, the pair (S', g) defines, as in Section 2.3, a point of $\mathrm{CHur}_{G,n} /G$: the quotient by G arises from the fact that we have only a conjugacy class of surjections.

This discussion has constructed a continuous function

$$\mathsf{Hn}_{G,n}(\mathbb{C}) \to \mathrm{CHur}_{G,n} /G,$$

covering the homeomorphism $\mathsf{Conf}_n(\mathbb{C}) \simeq \mathrm{Conf}_n$ made above. Moreover, this function induces a bijection between fibers of $\mathsf{Hn}_{G,n}(\mathbb{C}) \to \mathsf{Conf}_n(\mathbb{C})$ and fibers of $\mathrm{CHur}_{G,n} /G \to \mathrm{Conf}_n$. Therefore it is a homeomorphism. The map "with monodromy c" is obtained by restriction. $\square$

7.5. Our basic tool for comparing cohomology in characteristic 0 and characteristic p is a suitable compactification of configuration space. We denote by $\mathsf{PConf}_n/\mathrm{Spec}\,\mathbb{Z}$ the complement, in $\mathbb{A}^n$, of the divisors $\{z_i = z_j\}$ (where $1 \le i < j \le n$). This is a hyperplane complement whose points over any ring A are the ordered n-tuples $(z_1, \ldots, z_n)$ of sections in $\mathbb{A}^1(A)$ that are disjoint — i.e., that satisfy $z_i - z_j \in A^\times$ whenever $i \ne j$. There is a map from PConf_n to Conf_n sending $(z_1, \ldots, z_n)$ to the polynomial $\prod_i(X - z_i W)$.

7.6. LEMMA. *For all $n \ge 2$, PConf_n has a compactification X_n that is a smooth and proper scheme over $\mathrm{Spec}\,\mathbb{Z}$, and such that the complement $D_n = X_n \backslash \mathsf{PConf}_n$ is a relative normal crossings divisor.*

Proof. We remark first that the inclusion of PConf_n into $\mathbb{P}^n$ is a smooth proper compactification, but the complement is not normal crossings.

In the proof, we will make use of the moduli stack $\overline{\mathcal{M}}_{0,n+1}$ of stable $(n+1)$-pointed genus-0 curves. It is a proper smooth scheme over $\operatorname{Spec}\mathbb{Z}$ and it contains the open subscheme $\mathcal{M}_{0,n+1}$ corresponding to nonsingular curves, and the complement is a relative normal crossings divisor. These facts are in Knudsen [38, Th 2.7], except for the fact that $\overline{\mathcal{M}}_{0,n+1}$ is a scheme. This follows from the identification (explained in the proof of loc. cit., following from [38, Prop 2.1]) of $\overline{\mathcal{M}}_{0,n+1}$ with the universal curve over $\overline{\mathcal{M}}_{0,n}$, together with the fact that $\overline{\mathcal{M}}_{0,3} = \operatorname{Spec}(\mathbb{Z})$ is a scheme.

There is a natural map from $\mathsf{PConf}_n \subset \mathbb{A}^n$ to $\mathcal{M}_{0,n+1}$ sending a set of points $p_1, \ldots, p_n$ to $(\mathbb{P}^1; \infty, p_1, \ldots, p_n)$, i.e., to the $n+1$-pointed curve defined by $\mathbb{P}^1$ together with the sections defined by $\infty, p_1, \ldots, p_n$.

Let Aff be the group scheme of affine linear transformations of $\mathbb{A}^1$; that is, Aff is the group scheme of upper triangular matrices $\left(\begin{smallmatrix} A & B \\ 0 & 1 \end{smallmatrix}\right)$, corresponding to the transformation $z \mapsto Az + B$. We can map $\mathsf{PConf}_n \subset \mathbb{A}^n$ to Aff via $(p_1, \ldots, p_n) \in \mathsf{PConf}_n \mapsto (A, B) = (p_2 - p_1, p_1)$.

The resulting (product) map

$$F : \mathsf{PConf}_n \longrightarrow \mathcal{M}_{0,n+1} \times \mathrm{Aff}$$

is an isomorphism: For any base scheme S, the S-points of PConf_n are collections $p_1, \ldots, p_n \in \mathbb{A}^1(S) = \Gamma(S, \mathcal{O}_S)$, with the property that $p_i - p_j$ are units everywhere on S, and the map F sends this to

$$(\mathbb{P}^1_S; \infty, p_1, p_2, p_3, \ldots, p_n) \in \mathcal{M}_{0,n+1}(S) \times (p_2 - p_1, p_1) \in \mathrm{Aff}(S).$$

That this is a bijection follows from the fact that $\mathcal{M}_{0,3} \simeq \operatorname{Spec}(\mathbb{Z})$ which, in particular, implies that if one fixes $u, v \in \Gamma(S, \mathcal{O}_S)$ with $u - v \in \Gamma(S, \mathcal{O}_S^\times)$, any point $\mathcal{M}_{0,n+1}(S)$ is *uniquely representable* by $(\mathbb{P}^1_S, \infty, u, v, q_3, \ldots, q_n)$ for suitable sections q_i.

Since Aff is isomorphic, as scheme, to $\mathbb{G}_m \times \mathbb{G}_a$, it has a compactification by $\mathbb{P}^1 \times \mathbb{P}^1$ whose complement is normal crossings.

So $X_n = \overline{\mathcal{M}}_{0,n+1} \times \mathbb{P}^1 \times \mathbb{P}^1$ satisfies the requirements of the lemma. $\square$

7.7. PROPOSITION. *Let A be a Henselian discrete valuation ring, whose quotient field has characteristic zero. Let $\bar{\eta}$ resp. $\bar{s}$ be a geometric generic (resp. special) point of $\operatorname{Spec}A$. Let X be a scheme proper and smooth over $\operatorname{Spec}A$ and $D \subset X$ a reduced normal crossings divisor relative[10] to $\operatorname{Spec}A$.*

Let $U := X - D$, and let $\pi : U' \to U$ be a finite étale cover. Let $\mathcal{G}$ be a finite group that acts compatibly on U' and U, both actions covering the trivial $\mathcal{G}$-action on $\operatorname{Spec}A$. Then $H^i_{\text{ét}}(U'_{\bar{\eta}}, \mathbb{Z}/\ell\mathbb{Z})$ and $H^i_{\text{ét}}(U'_{\bar{s}}, \mathbb{Z}/\ell\mathbb{Z})$ are isomorphic as $\mathcal{G}$-modules for all i and all primes ℓ invertible in A.

---

[10]Recall that this means that D is — étale locally on X — isomorphic to a union of coordinate hyperplanes in an affine space over A.

Proof. We are grateful to a referee for providing a sketch of this argument, which substantially simplifies our original proof.

Let $\mathcal{F} = \pi_*(\mathbb{Z}/\ell)$. It is a locally constant sheaf of $\mathbb{Z}/\ell$-modules on U, and it is tamely ramified [31, Exp. XIII, 2.3(c)] along D — automatic because the generic point of D has characteristic zero.

Let K be the sheaf $j_!\mathcal{F}$ on X. Then

$$H^i(X_{\bar{s}}, K) = H^i_c(U_{\bar{s}}, \mathcal{F}) = H^i_c(U'_{\bar{s}}, \mathbb{Z}/\ell\mathbb{Z}),$$

and exactly the same assertion holds with $\bar{s}$ replaced by $\bar{\eta}$. The specialization map (as a reference, see [26, Ch. III, §3]) $H^i_c(U'_{\bar{s}}, \mathbb{Z}/\ell\mathbb{Z}) \to H^i_c(U'_{\bar{\eta}}, \mathbb{Z}/\ell\mathbb{Z})$ is $\mathcal{G}$-equivariant, by the functoriality of the constructions of loc. cit., and is identified with the corresponding specialization map for X and K. We prove it is an isomorphism:

The specialization map between cohomologies of special and generic fiber fits into a triangle involving vanishing cycles (see [21, Exp. XIII, eq. 2.1.8.9])

$$H^i(X_{\bar{s}}, K) \to H^i(X_{\bar{\eta}}, K) \to H^i(X_{\bar{s}}, \mathbf{R}\Phi_{\bar{\eta}}K) \xrightarrow{[1]},$$

where H^i denotes hypercohomology in case of a complex of sheaves. The final term *vanishes* by [21, Exp. XIII, 2.1.11], and we conclude that the specialization map furnishes an isomorphism $H^i_c(U'_{\bar{s}}, \mathbb{Z}/\ell\mathbb{Z}) \to H^i_c(U'_{\bar{\eta}}, \mathbb{Z}/\ell\mathbb{Z})$.

Because both $U'_{\bar{s}}$ and $U'_{\bar{\eta}}$ are smooth varieties, Poincaré duality converts (7.7.1) to the desired $\mathcal{G}$-equivariant isomorphism of usual (not compactly supported) cohomology

$$(7.7.1) \qquad\qquad H^j(U'_{\bar{s}}, \mathbb{Z}/\ell\mathbb{Z}) \xleftarrow{\sim} H^j(U'_{\bar{\eta}}, \mathbb{Z}/\ell\mathbb{Z}). \qquad\qquad \square$$

With Proposition 7.7 in hand we can now prove the desired upper bound for the Betti numbers of Hurwitz spaces over finite fields.

7.8. PROPOSITION. *Suppose that (G, c) satisfies the nonsplitting condition, where G is center-free and c is a rational conjugacy class generating G. Then there exists $C(G, c)$ depending only on (G, c) so that*

$$(7.8.1) \qquad\qquad \dim H^i_{\text{ét}}(\mathsf{Hn}^c_{G,n}/\bar{\mathbb{F}}_q, \mathbb{Q}_\ell) \leq C(G, c)^{i+1}$$

for all i, n, so long as $\ell > \max(|G|, q, n)$.

We note that the restriction on ℓ is irrelevant for the application to Cohen-Lenstra heuristics in Section 8.

Proof. Write $\mathsf{PHn}^c_{G,n}$ for the Cartesian product $\mathsf{Hn}^c_{G,n} \times_{\mathsf{Conf}_n} \mathsf{PConf}_n$. Applying Proposition 7.7, together with the comparison of étale and analytic cohomology [6, Th. 4.4, Exposé XI], to $A = W(\mathbb{F}_q)$, $U = \mathsf{PConf}_n$, $X = X_n$

as in Lemma 7.6, $U' = \mathsf{PHn}^c_{G,n}$, and $\mathcal{G} = S_n$ acting by permuting points on PConf_n, we find that

$$H^i(\mathsf{PHn}^c_{G,n}(\mathbb{C}), \mathbb{Z}/\ell\mathbb{Z}) \cong H^i_{\text{ét}}(\mathsf{PHn}^c_{G,n}/\overline{\mathbb{F}}_q, \mathbb{Z}/\ell\mathbb{Z}) \quad \text{(iso of S_n-modules)}.$$

If $\ell > n$, then the S_n-invariants on the mod ℓ cohomology of $\mathsf{PHn}^c_{G,n}$ recovers the mod ℓ cohomology of $\mathsf{Hn}^c_{G,n}$. Thus, supposing $\ell > n$, we obtain an isomorphism

$$(7.8.2) \qquad H^i(\mathsf{Hn}^c_{G,n}(\mathbb{C}), \mathbb{Z}/\ell\mathbb{Z}) \cong H^i_{\text{ét}}(\mathsf{Hn}^c_{G,n}/\overline{\mathbb{F}}_q, \mathbb{Z}/\ell\mathbb{Z}).$$

Now Lemma 7.4 and Corollary 6.2, together with an application of duality to pass between homology and cohomology, give a stability property for the left-hand side; thus we get the corresponding property for the right-hand side too:

$$(7.8.3) \qquad H^p_{\text{ét}}(\mathsf{Hn}^c_{G,n}/\overline{\mathbb{F}}_q, \mathbb{Z}/\ell\mathbb{Z}) \cong H^p_{\text{ét}}(\mathsf{Hn}^c_{G,n+D}/\overline{\mathbb{F}}_q, \mathbb{Z}/\ell\mathbb{Z})$$

whenever $n > Ap + B$ and $\ell > \max(q, |G|, n + D)$; recall that the constant D was introduced in Lemma 3.5. Now (7.8.1) with $\mathbb{Z}/\ell\mathbb{Z}$ coefficients in place of $\mathbb{Q}_\ell$ follows from Proposition 2.5.

The statement with $\mathbb{Q}_\ell$-coefficients follows: The dimension of cohomology with $\mathbb{Q}_\ell$-coefficients is bounded above by the rank (i.e., number of generators) of cohomology with $\mathbb{Z}_\ell$-coefficients, and the latter is bounded above by the dimension of $\mathbb{Z}/\ell\mathbb{Z}$-cohomology by the universal coefficient sequence. $\qquad\square$

7.9. *Remark.* The requirement that c generates G is technically unnecessary: if c fails to generate G, then $\mathsf{Hn}^c_{G,n}$ is empty.

Note that we have *not* shown that the isomorphism implicit in (7.8.3) is equivariant for the action of Frobenius on source and target. Doing so would allow us to show that $\delta^+(q) = \delta^-(q)$ in Theorem 1.2, so that we could talk about limits rather than limits inferior and superior. The reason for the deficit is that our stabilization map U is constructed in an essentially nonalgebraic way. Although this problem can perhaps be remedied (cf. [32, §4] for the corresponding issue in the case of moduli spaces of curves), we have not pursued this course in the present paper. As we record in Conjecture 1.5, we believe that apart from "obvious classes" both sides of the isomorphism in Proposition 7.8 are 0, making the Frobenius equivariance vacuous.

8. The Cohen-Lenstra heuristics

The *Cohen-Lenstra heuristics* are a family of conjectures proposed by the two named authors [13] concerning the distribution of class groups of quadratic number fields among all finite abelian groups. In fact, the phrase nowadays incorporates an even broader family of conjectures, worked out by Cohen, Lenstra, and Martinet [14], about class groups of number fields of all degree,

with conditions on Galois group, and so forth. They make sense over any global field.

In this section, we shall prove Theorem 1.2, which sheds some light on the Cohen-Lenstra heuristics over rational function fields over finite fields.

8.1. Let $\mathcal{L}$ be the set of isomorphism classes of finite abelian ℓ-groups. The *Cohen-Lenstra distribution* is a probability distribution on $\mathcal{L}$: the μ-mass of the (isomorphism class of) A equals

$$\prod_{i \geq 1}(1 - \ell^{-i}) \cdot |\operatorname{Aut}(A)|^{-1}.$$

The measure μ can be alternately described (see [28]) as the distribution of the cokernel of a random map $\mathbb{Z}_\ell^N \to \mathbb{Z}_\ell^N$ (*random* according to the additive Haar measure on the space of such maps), as $N \to \infty$. From this latter description, we see that the expected number of surjections from a μ-random group into a fixed abelian ℓ-group A_0 equals 1.

In fact, this last remark *characterizes* μ. Writing $\operatorname{Sur}(G_1, G_2)$ for the set of surjections from the group G_1 to the group G_2, we have

8.2. LEMMA. *If ν is any probability measure on $\mathcal{L}$ for which the expected number of surjections from a ν-random group to A_0 always equals 1 — i.e.,*

$$\sum_{B \in \mathcal{L}} \nu(B) \cdot |\operatorname{Sur}(B, A_0)| = 1$$

for all $A_0 \in \mathcal{L}$ — then $\nu = \mu$.

Proof. Indeed, the assumption gives, for every abelian ℓ-group A,

(8.2.1) $\qquad |\operatorname{Aut}(A)| \cdot \nu(A) + \sum_{B \in \mathcal{L}, B \neq A} |\operatorname{Sur}(B, A)| \cdot \nu(B) = 1.$

Here $|\operatorname{Sur}(B, A)|$ denotes the number of surjections from (a representative for) B to A.

First of all, (8.2.1) forces

$$|\operatorname{Aut}(A)| \cdot \nu(A) \leq 1.$$

Inserting this upper bound back into (8.2.1), we obtain the *lower* bound

$$|\operatorname{Aut}(A)| \cdot \nu(A) \geq 1 - \beta,$$

where $\beta := \sum_{B \neq A} \frac{|\operatorname{Sur}(B,A)|}{|\operatorname{Aut}(B)|} = (\prod_{i \geq 1}(1 - \ell^{-i})^{-1} - 1)$, the latter equality from the fact that (8.2.1) holds for $\nu = \mu$.

Proceeding in this fashion, we find that $\nu(A) \cdot |\operatorname{Aut}(A)|$ is bounded above and below by alternating partial sums of the series $1 - \beta + \beta^2 - \cdots$ and consequently $\nu(A) \cdot |\operatorname{Aut}(A)| = \frac{1}{1+\beta}$ for every A, as desired. $\qquad\square$

This result admits a more quantitative form. If ν is a probability measure on $\mathcal{L}$, and $A \in \mathcal{L}$, write $\langle \operatorname{Sur}(-, A) \rangle_\nu$ for the expected number of surjections from a ν-random group to A.

8.3. PROPOSITION. *Suppose given $\epsilon_0 > 0$ and a finite subset $L \subset \mathcal{L}$. Then there exists $\delta > 0$ and a finite subset $L' \subset \mathcal{L}$ such that, if ν is any probability measure on $\mathcal{L}$ for which $\langle \mathrm{Sur}(-, A) \rangle_\nu \in [1 - \delta, 1 + \delta]$, for any $A \in L'$, then also $|\nu(A) - \mu(A)| \leq \epsilon_0$ for any $A \in L$.*

The proof will require the following:

8.4. LEMMA. *Given $\epsilon > 0$ and $A \in \mathcal{L}$, there exist a constant $c(A)$ and a finite subset $M \subset \mathcal{L}$ so that, whenever $|X| > c(A)$,*

$$|\mathrm{Sur}(X, A)| \leq \epsilon \frac{\sum_{A' \in M} |\mathrm{Sur}(X, A')|}{|M|}.$$

Proof. Call an *enlargement* of A any group A' that admits a surjection $A' \twoheadrightarrow A$ with kernel of size ℓ. We claim that for any abelian ℓ-group X surjecting onto A, with $|X| > |A|$, there exists an enlargement A' such that

(8.4.1) $|\mathrm{Sur}(X, A')| \geq (\ell - 1)|\mathrm{Sur}(X, A)|.$

Certainly there exists an enlargement A' such that $\mathrm{Sur}(X, A')$ is nonempty (take a suitable quotient of X). Fix such an A', and fix a surjection $\pi : A' \twoheadrightarrow A$ and $f \in \mathrm{Sur}(X, A)$. We examine lifts $\tilde{f} : X \to A'$ of f (with respect to π). Then (8.4.1) follows from the fact that the number of such *surjective* lifts $\tilde{f}$ is at least $\ell - 1$:

(i) If A' is not isomorphic to $A \times \mathbb{Z}/\ell\mathbb{Z}$, any such lift $\tilde{f}$ is surjective, and the set of lifts is a principal homogeneous space under $\mathrm{Hom}(X, \mathbb{Z}/\ell\mathbb{Z})$.

(ii) If A' is isomorphic to $A \times \mathbb{Z}/\ell\mathbb{Z}$: Note that $\ker(f)$ cannot be contained in ℓX; if it were, then f induces an isomorphism $X/\ell \to A/\ell$, but then the ℓ-rank of X and A coincide, and then $\mathrm{Sur}(X, A')$ would be empty, contradicting our choice of A'.

Thus there exists a homomorphism $\varphi : X \to \mathbb{Z}/\ell\mathbb{Z}$ that is nontrivial on $\ker(f)$, and (f, φ) gives a surjection $X \to A'$ that lifts f. Since there are at least $(\ell - 1)$ choices for φ, there are at least $\ell - 1$ surjective lifts of f.

We now iterate (8.4.1): Call an *s-enlargement* of A any group A' that admits a map $A' \twoheadrightarrow A$ with kernel of size ℓ^s. Then, for any abelian ℓ-group X surjecting onto A and of size larger than $\ell^s|A|$, we see that there exists an *s*-enlargement A' so that

$$|\mathrm{Sur}(X, A')| \geq (\ell - 1)^s|\mathrm{Sur}(X, A)|.$$

On the other hand, the number of (isomorphism classes of) *s*-enlargements is bounded by the number of partitions of $s + m$, where $\ell^m = |A|$. Since $p(s + m)(\ell - 1)^{-s} \to 0$ as $s \to \infty$, the statement of the lemma follows, taking M to be the set of *s*-enlargements of A and $c(A) = \ell^s|A|$. $\square$

We will now prove Proposition 8.3. Recall that we say a sequence of measures ν_k on $\mathcal{L}$ weakly converges to a limit ν_∞ if one has convergence of integrals $\int f\nu_k \to \int f\nu_\infty$ for each compactly supported continuous function; for measures on the discrete space $\mathcal{L}$, this is equivalent to asking that $\nu_k(A) \to \nu_\infty(A)$ for every $A \in \mathcal{L}$. Any sequence ν_k has (by a diagonal argument) a weakly convergent subsequence. However, the limit need not be a probability measure; it may assign $\mathcal{L}$ a mass that is strictly less than 1.

Proof. Let L'_k be the subset of $\mathcal{L}$ comprising groups with $|A| \leq k$. If the assertion were false, there would be some measure ν_k that "does not work" for $L' = L'_k, \delta = 1/k$ — that is to say,

(1) $|\langle \mathrm{Sur}(-, A)_{\nu_k} \rangle - 1| \leq 1/k$ for all $A \in L'_k$;
(2) $|\nu_k(A) - \mu(A)| > \epsilon_0$ for some $A \in L$.

Passing to a weakly convergent subsequence, we obtain measures ν_k having the following property:

$$(8.4.2) \qquad \lim_{k \to \infty} \langle \mathrm{Sur}(-, A) \rangle_{\nu_k} = 1,$$

for every fixed $A \in \mathcal{L}$, but ν_k weakly converge to a measure $\nu_\infty \neq \mu$. We will deduce a contradiction.

Fix $\epsilon > 0$ arbitrary. This is not related to ϵ_0 in the statement above. Indeed, we will apply Lemma 8.4 with this value of ϵ and then let ϵ approach zero at the end of the argument.

We claim $\langle \mathrm{Sur}(-, A) \rangle_{\nu_\infty} = 1$. Indeed, this expectation is ≤ 1 by Fatou's lemma; on the other hand, with $c = c(A)$ as in the statement of Lemma 8.4,

$$(8.4.3) \quad \langle \mathrm{Sur}(-, A) \rangle_{\nu_\infty} = \sum_{|B| \leq c} \nu_\infty(B)|\mathrm{Sur}(B, A)| + \sum_{|B| > c} \nu_\infty(B)|\mathrm{Sur}(B, A)|$$

$$\geq \lim_k \sum_{|B| \leq c} \nu_k(B)|\mathrm{Sur}(B, A)|$$

$$= 1 - \lim_k \sum_{|B| > c} \nu_k(B)|\mathrm{Sur}(B, A)|.$$

By Lemma 8.4,

$$(8.4.4) \qquad \sum_{|B| > c} \nu_k(B)|\mathrm{Sur}(B, A)| \leq \epsilon |M|^{-1} \sum_{|B| > c, A' \in M} \nu_k(B)|\mathrm{Sur}(B, A')|.$$

Now, by assumption, for any $A' \in M$ and any $k > |A'|$,

$$\sum_{|B| > c} \nu_k(B)|\mathrm{Sur}(B, A')| \leq \langle \mathrm{Sur}(-, A') \rangle_{\nu_k} \leq 1 + 1/k,$$

and using (8.4.4) and passing to the limit, we get

$$\limsup_k \sum_{|B| > c} \nu_k(B)|\mathrm{Sur}(B, A)| \leq \epsilon.$$

Thus by (8.4.3) and the prior discussion, we get $\langle \mathrm{Sur}(-, A) \rangle_{\nu_\infty} \in [1 - \epsilon, 1]$. Since ϵ was arbitrary,

$$\langle \mathrm{Sur}(-, A) \rangle_{\nu_\infty} = 1.$$

Applying this conclusion with A trivial, we see that ν_∞ is a probability measure; now Lemma 8.2 shows $\nu_\infty = \mu$, a contradiction.   □

8.5.   The Cohen-Lenstra heuristics in the simplest case — as formulated in [13] — assert that, for $\ell \neq 2$, the ℓ-part of class groups of imaginary quadratic extensions of $\mathbb{Q}$ — when ordered by discriminant — approach μ in distribution. Precisely: amongst the set S_X of imaginary quadratic fields of discriminant less than X, the fraction for which the ℓ-part of the class group is isomorphic to A approaches $\mu(A)$, as $X \to \infty$.

In view of what we have just proved, this is equivalent to the validity of the following assertion for all abelian ℓ-groups A: the *average number of surjections* from the class group of a varying quadratic field to A equals 1. Explicitly,

$$(8.5.1) \qquad \lim_{X \to \infty} \frac{\sum_{K \in S_X} |\mathrm{Sur}(\mathrm{Cl}_K, A)|}{|S_X|} = 1.$$

In the formulation (8.5.1), there are results for specific A: (8.5.1) is true for $A = \mathbb{Z}/3\mathbb{Z}$ by work of Davenport and Heilbronn; the corresponding assertion is even known over an arbitrary global field by work of Datskovsky and Wright [16]; and a natural variant for $A = \mathbb{Z}/4\mathbb{Z}$ is a theorem of Fouvry and Klüners [25].

If we replace $\mathbb{Q}$ by $\mathbb{F}_q(t)$, there is a fair amount of work ([2], [49]) on the different problem (with no obvious analog over a number field) in which we fix the discriminant degree and take a limit as $q \to \infty$.

8.6.   Let $\mathbb{F}_q$ be a finite field, and let $K = \mathbb{F}_q(t)$. Let ℓ be an odd prime not dividing q, let A be a nontrivial finite abelian ℓ-group, and define

$$G := A \rtimes (\mathbb{Z}/2\mathbb{Z}),$$

where the nontrivial element of $\mathbb{Z}/2\mathbb{Z}$ acts on A by inversion. Let c be the conjugacy class in G consisting of all involutions. Then c generates G. Moreover, the pair (G, c) is nonsplitting by Lemma 3.2. For brevity, we write simply $\mathbf{X}_n$ for the Hurwitz scheme $\mathsf{Hn}^c_{G,n} \times_{\mathrm{Spec}(R)} \mathrm{Spec}\,\mathbb{F}_q$ that parametrizes (Sections 7.1 and 7.3) tame G-covers of the affine line, branched at n points, all of whose ramification is of type c.

8.7. PROPOSITION. *Let n be an odd integer. There is a bijection between $\mathbf{X}_n(\mathbb{F}_q)$ and the set of isomorphism classes of pairs (L, α), where L is a quadratic extension of K of discriminant degree $n + 1$ ramified at ∞, and α is a*

surjective homomorphism

$$\alpha : \mathrm{Cl}_L \to A.$$

Here, Cl_L *is the class group of* $\mathcal{O}_L$, *the integral closure of* $\mathbb{F}_q[t]$ *inside* L. *Two pairs* $(L, \alpha), (L', \alpha')$ *are isomorphic if there exists a* K-*isomorphism* $f : L \to L'$ *with* $f^*\alpha' = \alpha$.

Note that (L, α) and (L, β) are isomorphic if and only if α, β are interchanged by the automorphism of L/K, i.e., (see discussion below) if and only if $\beta = \pm\alpha$.

Note that, under the assumptions, $\mathbb{F}_q$ is automatically algebraically closed inside L. If we denote by C_L the smooth proper curve over $\mathbb{F}_q$ associated to L, then Cl_L is identified with the group of degree zero $\mathbb{F}_q$-rational divisors on C_L up to equivalence: an ideal of $\mathcal{O}_L$ defines a divisor D_I on C_L, and then $I \mapsto D_I - \deg(D_I).\infty$ descends to the desired isomorphism, where ∞ denotes the unique closed point of C_L above ∞ on $\mathbb{P}^1$. Also $\mathrm{Cl}_L = \mathrm{Cl}(\mathcal{O}_L)$ is isomorphic to the $\mathbb{F}_q$-points of the Jacobian of C_L (see, e.g., [45, Th. C, (ii)]).

Proof. For L a quadratic extension of K as in the proposition statement, we let σ be the nontrivial automorphism of L over K, and we let J_L be the Galois group of the maximal abelian everywhere unramified extension E/L with pro-ℓ Galois group.

Note that E/K is Galois, with $\mathrm{Gal}(E/L) \simeq J_L$ as a normal subgroup. The subgroup $\langle x + \sigma(x) : x \in J_L \rangle$ is a normal subgroup of $\mathrm{Gal}(E/K)$. Let F_L be the fixed field of this subgroup. The extension F_L/K is also Galois, and its Galois group fits in an extension

$$1 \to J'_L \to \mathrm{Gal}(F_L/K) \to \mathrm{Gal}(L/K) \to 1,$$

where J'_L is the quotient of J_L by all elements $x + \sigma(x)$ with $x \in J_L$; i.e. J'_L is the largest quotient group of J_L on which σ acts by -1.

Since $\mathrm{Gal}(L/K) \cong \mathbb{Z}/2\mathbb{Z}$ and $\ell \neq 2$, this sequence splits as a semidirect product

$$(8.7.1) \qquad\qquad \mathrm{Gal}(F_L/K) = J'_L \rtimes \langle \tau \rangle,$$

where τ is any involution in $\mathrm{Gal}(F_L/K)$.

Class field theory yields a short exact sequence

$$(8.7.2) \qquad\qquad (\mathrm{Cl}_L)_\ell \hookrightarrow J_L \twoheadrightarrow \widehat{\mathbb{Z}}_\ell,$$

where we have written $(\mathrm{Cl}_L)_\ell$ for the Sylow ℓ-subgroup of Cl_L. Now σ acts compatibly on all terms of (8.7.2); the action on $\widehat{\mathbb{Z}}_\ell$ is trivial, and its action on Cl_L is by negation: for any fractional ideal I of $\mathcal{O}_L$, the product $I \cdot \sigma(I)$ is the extension of an ideal from K, and thus principal.

Thus, the sequence (8.7.2) induces, via the snake lemma, a canonical isomorphism between $J_L/(1+\sigma)J_L$ and $(\mathrm{Cl}_L)_\ell/(1+\sigma)(\mathrm{Cl}_L)_\ell$ or, in other words,

$$(8.7.3) \qquad\qquad\qquad J'_L \xrightarrow{\sim} (\mathrm{Cl}_L)_\ell.$$

Therefore, beginning with (L,α) as in the statement of the theorem, we obtain by composing α with (8.7.3) a surjection $f_\alpha : J'_L \twoheadrightarrow A$; by (8.7.1) this extends to $g_\alpha : \mathrm{Gal}(F_L/K) \to G$, by sending τ to any involution in G. Since all involutions in G are conjugate under A, the extension g_α is unique up to A-conjugacy.

Let F_α be the fixed field of $\ker(g_\alpha)$; it is a Galois extension of K, equipped with an isomorphism $\mathrm{Gal}(F_\alpha/K) \simeq G$ that is defined up to A-conjugacy. Moreover, $\mathbb{F}_q$ is algebraically closed inside F_α. To say the same geometrically, given (L,α) we have associated a geometrically connected curve Y_α (namely, the curve associated to F_α) together with a map $Y_\alpha \to \mathbb{P}^1$ and an isomorphism $g_\alpha : \mathrm{Aut}(Y/\mathbb{P}^1) \to G$.

The ramified places of F_α/K are the same as those of L/K, because F_α/L is everywhere unramified. Finally, any inertia group I_v above such a ramified place v satisfies $g_\alpha(I_v) = \langle g \rangle$ for some $g \in c$; indeed, $g_\alpha(I_v)$ is generated by a single element of G that maps to the nontrivial element of $\mathbb{Z}/2\mathbb{Z}$.

Although g_α is only well defined up to A-conjugacy, the isomorphism class of the G-cover of $\mathbb{P}^1$ defined by (Y_α, g_α) *does not change* if we conjugate g_α by A. In other words, we have defined a tame G-cover of $\mathbb{P}^1/\mathbb{F}_q$ branched at n points of $\mathbb{A}^1$, and its ramification is all of type c. By the description of the Hurwitz scheme in Section 7.1, this is equivalent to a point of $\mathbf{X}_n(\mathbb{F}_q)$.

Therefore, our discussion yields a map

$$(8.7.4) \qquad\qquad \{(L,\alpha)\} \text{ up to isomorphism} \longrightarrow \mathbf{X}_n(\mathbb{F}_q).$$

This map is bijective: Let $X \to \mathbb{P}^1$ be a tame G-cover of $\mathbb{P}^1$, branched at n points of $\mathbb{A}^1$. The quotient map $X \to X/A$ is étale above $\mathbb{A}^1$ — this follows from a local computation, using the fact that the ramification is of type c. Therefore, the degree 2 quotient map $X/A \to \mathbb{P}^1$ is ramified at an odd number of points of $\mathbb{A}^1$, and it must therefore also be ramified at ∞. The monodromy of $X \to \mathbb{P}^1$ above ∞ is then a cyclic subgroup that projects surjectively on $\mathbb{Z}/2\mathbb{Z}$; such a subgroup must be of order 2, and so $X \to X/A$ is étale everywhere. Thus we get an inverse to (8.7.4) by sending this tame G-cover to the pair (L,α), where L is the function field of X/A and $\alpha : \mathrm{Cl}_L \to A$ is the map arising from applying class field theory to the étale cover $X \to X/A$. $\qquad\square$

Let $\mathfrak{S}_n$ be the set of quadratic extensions of K of the form $L = K(\sqrt{f(t)})$, where $f(t)$ is a squarefree polynomial of odd degree n. To exhaust $\mathfrak{S}_n$, it is sufficient to let f range through a set of representatives for squarefree polynomials up to the multiplication action of $(\mathbf{F}_q^*)^2$. The number of *monic* squarefree

polynomials of degree n with coefficients in $\mathbf{F}_q$ is equal [10] to $q^n - q^{n-1}$, from where we deduce

$$(8.7.5) \qquad |\mathfrak{S}_n| = 2(q^n - q^{n-1}).$$

In what follows, we will average over fields in $\mathfrak{S}_n$, and let $n \to \infty$. This is the analog of "dyadic averages" in analytic number theory, and we find it to be aesthetically preferable in the case of a function field. However, it is easy to deduce similar results for averages over sets such as $\coprod_{m \leq n} \mathfrak{S}_m$; for example, from Theorem 8.8 below, one immediately deduces the corresponding statement with $\mathfrak{S}_n$ replaced by $\coprod_{m \leq n} \mathfrak{S}_m$.

Write $m_A(L) = |\mathrm{Sur}(\mathrm{Cl}(\mathcal{O}_L), A)|$. Then — in view of Proposition 8.3 — the following theorem implies Theorem 1.2.

8.8. THEOREM. *Let ℓ be an odd prime not dividing q or $q - 1$, and A an ℓ-group. There is a constant $B(A)$ such that*

$$(8.8.1) \qquad \left| \frac{\sum_{L \in \mathfrak{S}_n} m_A(L)}{|\mathfrak{S}_n|} - 1 \right| \leq B(A)/\sqrt{q}$$

for all n, q with $\sqrt{q} > B(A)$ and n an odd integer greater than $B(A)$.

Here is the explicit argument that this implies Theorem 1.2: Let A_0 be any fixed abelian ℓ-group, and let $\epsilon > 0$. For a given n, let ν_n be the probability measure[11] on $\mathcal{L}$ with $\nu_n(A)$ equal to the fraction of $L \in \mathfrak{S}_n$ with $\mathrm{Cl}(\mathcal{O}_L) \simeq A$.

Apply Proposition 8.3 to the measure ν_n and with $L = \{A_0\}$; the corollary gives "as output" a finite list L' of abelian ℓ-groups and $\delta > 0$ with the property that

$$\text{if } \left| \frac{\sum_{L \in \mathfrak{S}_n} m_A(L)}{|\mathfrak{S}_n|} - 1 \right| < \delta \text{ for all } A \in L', \text{ then } |\nu_n(A_0) - \mu(A_0)| < \epsilon.$$

Now, notation as in Theorem 8.8, let Q be chosen so that $Q > B(A)$ and $B(A)/\sqrt{Q} < \delta$ for every $A \in L'$; we see that if $q > Q$ and $n > Q$, then we have $|\nu_n(A_0) - \mu(A_0)| \leq \epsilon$. Thus, for any $q > Q$, the upper and lower densities discussed in Theorem 1.2 are both bounded between $\mu(A_0) - \epsilon$ and $\mu(A_0) + \epsilon$. Since ϵ was arbitrary, the result follows.

Proof. In the proof that follows, we use H^i to denote ith étale cohomology and H^i_c to denote the corresponding compactly supported cohomology group.

Note that if A is the trivial group, the left-hand side of (8.8.1) is zero. It is enough to treat the case that A is nontrivial.

---

[11] We do not know that the limit $\lim_n \nu_n$ is a probability measure, but we are not using that.

By Proposition 8.7, and the fact that (L, α) and (L, β) are isomorphic if and only if $\alpha = \pm\beta$, we know that

$$\sum_{L \in \mathfrak{S}_n} m_A(L) = 2|\mathbf{X}_n(\mathbb{F}_q)|,$$

and as noted in (8.7.5), we have $|\mathfrak{S}_n| = 2(q^n - q^{n-1})$. It will suffice, then, to show that

(8.8.2) $$\left| \frac{|\mathbf{X}_n(\mathbb{F}_q)|}{q^n} - 1 \right| \leq \frac{B(A)}{\sqrt{q}}$$

when n and q are sufficiently large relative to A.

Denote by $\bar{\mathbf{X}}_n$ the base change of $\mathbf{X}_n$ to $\bar{\mathbb{F}}_q$, i.e., $\bar{\mathbf{X}}_n = \mathsf{Hn}^c_{G,n} \times_{\mathrm{Spec}(R)} \mathrm{Spec}(\bar{\mathbb{F}_q})$. Fix a sufficiently large prime ℓ. By (7.8.1) and Poincaré duality, for the smooth n-dimensional variety $\bar{\mathbf{X}}_n$, we get the existence of a constant $C(A)$ such that

$$\dim H^{2n-i}_c(\bar{\mathbf{X}}_n, \mathbb{Q}_\ell) = \dim H^i(\bar{\mathbf{X}}_n; \mathbb{Q}_\ell) \leq C(A)^i$$

for all $i > 0$ — just take $C(A) = C(G, c)^2$ in the notation of (7.8.1).

Deligne has proven [19] that every eigenvalue of the geometric Frobenius Frob_q (i.e., if we fix a projective embedding of $\mathbf{X}_n$ over $\mathbb{F}_q$, this is the operation that raises coordinates to the qth power) on *compactly supported* H^j_c of a smooth variety is bounded above, in absolute value, by $q^{j/2}$. Consequently,

$$\left| q^{-n} \sum_{j < 2n} (-1)^j \mathrm{Tr}\left(\mathrm{Frob}_q \,|\, H^j_c(\bar{\mathbf{X}}_n; \mathbb{Q}_\ell) \right) \right| \leq q^{-n} \sum_{j=0}^{2n-1} q^{j/2} \dim H^j_c(\bar{\mathbf{X}}_n; \mathbb{Q}_\ell)$$

$$\leq q^{-n} \sum_{j=0}^{2n-1} C(A)^{2n-j} q^{j/2}$$

$$\leq \sum_{k=1}^{\infty} \left(\frac{C(A)}{\sqrt{q}} \right)^k.$$

The last quantity above is at most $2\frac{C(A)}{\sqrt{q}}$ as long as $\frac{C(A)}{\sqrt{q}} \leq 1/2$; in other words, taking $B(A)$ to be $2C(A)$, we have

(8.8.3) $$\left| q^{-n} \sum_{j < 2n} (-1)^j \mathrm{Tr}\left(\mathrm{Frob}_q \,|\, H^j_c(\bar{\mathbf{X}}_n; \mathbb{Q}_\ell) \right) \right| \leq \frac{B(A)}{\sqrt{q}}$$

whenever $\sqrt{q} > B(A)$.

We now claim that, for sufficiently large n (this notion depending only on A)

$$\mathrm{Tr}\left(\mathrm{Frob}_q \,|\, H^{2n}_c(\bar{\mathbf{X}}_n; \mathbb{Q}_\ell) \right) = q^n.$$

By Poincaré duality, this is equivalent to the statement that there is exactly one $\mathbb{F}_q$-rational connected component of $\bar{\mathbf{X}}_n$. This, together with the Lefschetz trace formula (1.8.2) and the bound (8.8.3), will imply (8.8.2).

Let η be the generic point of $\mathsf{Conf}_n \times_{\mathrm{Spec}\, R} \mathrm{Spec}\, \mathbb{F}_q$, and write K for the function field of η. Then the finite étale cover $\mathbf{X}_n \to \mathsf{Conf}_n \times_{\mathrm{Spec}\, R} \mathrm{Spec}\, \mathbb{F}_q$ is determined by its geometric generic fiber Σ together with the action of $\mathrm{Gal}(\bar{K}/K)$ on that fiber. The latter group sits in a sequence

$$\mathrm{Gal}(\bar{K}/\overline{\mathbb{F}_q}K) \to \mathrm{Gal}(\bar{K}/K) \to \mathrm{Gal}(\overline{\mathbb{F}_q}/\mathbb{F}_q).$$

Then the desired conclusion (that there is exactly one $\mathbb{F}_q$-rational connected component of $\bar{\mathbf{X}}_n$) is precisely the statement that only one $\mathrm{Gal}(\bar{K}/\overline{\mathbb{F}_q}K)$-orbit on Σ is preserved by the action of $\mathrm{Gal}(\overline{\mathbb{F}_q}/\mathbb{F}_q)$.

We prove this by expressing Σ in a different way, allowing us to make contact with the existing literature on monodromy in families of hyperelliptic curves.

Recall (Section 7.3) the definition of Conf_n as the subscheme of Conf'_{n+1} with $a_0 = 0$. Let C be the smooth hyperelliptic curve over K birational to the plane curve

$$Y^2 = a_1 X^n + a_2 X^{n-1}W + \cdots + a_{n+1}W^n.$$

Choose k sufficiently large so that $\ell^k A = 0$. Let V be the ℓ^k-torsion points of the Jacobian $\mathrm{Jac}(C)$ over $\bar{K}$. Then $V \simeq (\mathbb{Z}/\ell^k\mathbb{Z})^{2g}$ (with $g = \lfloor \frac{n-1}{2} \rfloor$, which, under our standing hypothesis that n is odd, equals $\frac{n-1}{2}$) and we have a *monodromy homomorphism*

$$\mu : \mathrm{Gal}(\bar{K}/K) \to \mathrm{Aut}(V).$$

Now consider the set $\mathrm{Sur}(V, A)$ of surjective homomorphisms from V to A. This set carries a natural action of $\mathrm{Gal}(\bar{K}/K)$ derived from μ.

Since Conf_n is a moduli scheme for degree-n squarefree divisors on $\mathbb{A}^1$, there is a universal such divisor on $\mathbb{A}^1/\mathsf{Conf}_n$, which restricts to a canonical degree-n squarefree (i.e., reduced) divisor D on $\mathbb{A}^1/\bar{K}$. The set $\mathbf{X}_n(\bar{K})$ of tame G-covers of $\mathbb{A}^1/\bar{K}$ branched at D (which is to say Σ) is naturally identified by the argument of Proposition 8.7 with $\mathrm{Sur}(V, A)$, equivariantly for the action of $\mathrm{Gal}(\bar{K}/K)$ on both sides. (In the proof of Proposition 8.7, we replace the statement of class field theory, which enters at (8.7.2), with the fact (see, e.g., [37, (2.4)]; in the case at hand, this is just Kummer theory) that the abelian étale extensions of $C/\bar{K}$ with Galois group A are classified by surjections $\mathrm{Jac}(C)[\ell^k](\bar{K}) \twoheadrightarrow A$.)

It thus suffices to show that only one $\mathrm{Gal}(\bar{K}/\overline{\mathbb{F}_q}K)$-orbit on $\mathrm{Sur}(V, A)$ is preserved by the action of $\mathrm{Gal}(\overline{\mathbb{F}_q}/\mathbb{F}_q)$ (again, for n large enough).

The action of $\mathrm{Gal}(\bar{K}/\overline{\mathbb{F}_q}K)$ on V preserves the Weil pairing $V \times V \to \mathbb{Z}/\ell^k\mathbb{Z}(1)$, which we write as $\langle v_1, v_2 \rangle$ for $v_1, v_2 \in V$. For $m \in (\mathbb{Z}/\ell^k)^\times$, write

$\mathrm{GSp}_m(V)$ for all automorphisms $\alpha \in \mathrm{Aut}(V)$ that satisfy $\langle \alpha(v_1), \alpha(v_2) \rangle = m \langle v_1, v_2 \rangle$; write $\mathrm{Sp}(V) = \mathrm{GSp}_1(V)$ for automorphisms preserving $\langle -, - \rangle$ and $\mathrm{GSp}(V)$ for $\bigcup_m \mathrm{GSp}_m(V)$. Thus,

$$\mu(\mathrm{Gal}(\bar{K}/K)) \subset \mathrm{GSp}(V) \text{ and } \mu(\mathrm{Gal}(\bar{K}/\overline{\mathbb{F}_q}K)) \subset \mathrm{Sp}(V).$$

Moreover, if F is an element of $\mathrm{Gal}(\bar{K}/K)$ lying over Frobenius in $\mathrm{Gal}(\overline{\mathbb{F}_q}/\mathbb{F}_q)$, then $\mu(F)$ lies in $\mathrm{GSp}_q(V)$.

Jiu-Kang Yu has proved [51] that $\mu(\mathrm{Gal}(\bar{K}/\overline{\mathbb{F}_q}K)) = \mathrm{Sp}(V)$ for large enough g in this case (again using that $\ell \neq 2$; for $\ell = 2$, the monodromy group is in fact smaller.)[12] For other proofs of Yu's (unpublished) result, see Achter–Pries [3, Th. 3.4] and Hall [33, Th. 4.1]. This "big monodromy" theorem simplifies the situation considerably: The geometric components of $\bar{\mathbf{X}}_n/\overline{\mathbb{F}}_q$ are therefore in bijection with $\mathrm{Sp}(V)$-orbits on $\mathrm{Sur}(V, A)$, and an orbit O is defined over $\mathbb{F}_q$ if and only if the stabilizer in $\mathrm{GSp}(V)$ of some $x \in O$ (equivalently: every $x \in O$) has nontrivial intersection with $\mathrm{GSp}_q(V)$.

We claim that, for sufficiently large n, there is a *unique* $\mathrm{Sp}(V)$-orbit on $\mathrm{Sur}(V, A)$ defined over $\mathbb{F}_q$. This can be reduced to a corresponding "linear algebra" statement with V replaced by $\mathbb{Z}_\ell^{2g}$ as follows: Write T for the full Tate module of $\mathrm{Jac}(C)$, so that $T/\ell^k \simeq V$. Because we chose k so that $\ell^k A = 0$, the pullback under $T \to V$ identifies $\mathrm{Sur}(V, A) \overset{\sim}{\to} \mathrm{Sur}(T, A)$. By smoothness, the map $\mathrm{GSp}_q(T) \to \mathrm{GSp}_q(V)$ is surjective. Therefore, our desired conclusion follows from the subsequent lemma.   □

8.9. LEMMA. *Let V be a finite free $\mathbb{Z}_\ell$-module of rank $2g$, equipped with a perfect symplectic pairing $\omega : V \times V \to \mathbb{Z}_\ell$. Let A be a finite abelian ℓ-group and $q \in \mathbb{Z}_\ell^\times$ be such that $q - 1$ is invertible in $\mathbb{Z}_\ell$. Define O as the set of all surjections $V \to A$ whose stabilizer, inside $\mathrm{GSp}(V)$, intersects $\mathrm{GSp}_q(V)$ nontrivially:*
(8.9.1)
$$O = \{f : V \to A \text{ surjective, and there exists } h \in \mathrm{GSp}_q(V) \text{ with } f \circ h = f\}.$$

Then, for g sufficiently large, O is nonempty, and $\mathrm{Sp}(V)$ acts transitively on O.

Proof. We shall use the following four facts, all of which remain valid for any finite rank free $\mathbb{Z}_\ell$-module M with a nondegenerate symplectic form ("nondegenerate" means that the symplectic form induces an isomorphism $M \to \mathrm{Hom}_{\mathbb{Z}_\ell}(M, \mathbb{Z}_\ell)$):

(i) any two maximal isotropic $\mathbb{Z}_\ell$-submodules of V are conjugate to one another under $\mathrm{Sp}(V)$;

---

[12] In fact, he proved this as part of a program to study the Cohen-Lenstra conjecture over function fields, just as we do; his theorem on monodromy allows him to prove a result in the $q \to \infty$ limit as alluded to above.

(ii) if a direct sum decomposition $V = V_1 \oplus V_2$ is orthogonal for ω, the restriction $\omega|V_j$ is nondegenerate for $j = 1, 2$;

(iii) V admits a decomposition $V_+ \oplus V_-$, where both V_+, V_- are maximal isotropic;

(iv) given a decomposition $V = A \oplus B$, where both A, B are isotropic for ω, then A, B are maximal isotropic.

For (i), one can argue by extending a free $\mathbb{Z}_\ell$-basis $\{x_1, \ldots, x_g\}$ for a maximal isotropic $\mathbb{Z}_\ell$-submodule to a standard symplectic basis for V: by nondegeneracy, choose a basis y_j with $\langle x_i, y_j \rangle = \delta_{ij}$, and then successively modify y_j by a combination $\sum_{i \leq j} a_i x_i$ so that $\langle y_i, y_j \rangle = 0$. For (ii), we note that a "degenerate vector" in V_1, i.e., a vector $v_1 \in V_1$ that satisfies $\langle v_1, w \rangle \in \ell\mathbb{Z}_\ell$ for all $w \in V_1$, would also be degenerate when considered as a vector in $V_1 \oplus V_2$. (iii) is immediate from Corollary 3.5 of [41, Ch. 1]. (iv) follows from the corresponding fact for symplectic forms over fields.

To ensure that O is nonempty when g is sufficiently large, write V as the direct sum of two maximal isotropic subspaces $V_+ \oplus V_-$. The automorphism of V that acts as q on V_+ and 1 on V_- lies in $\mathrm{GSp}_q(V)$, and it fixes any surjection from V to A factoring through projection to V_-. Such a surjection exists as long as $g \geq \dim_{\mathbb{F}_\ell} A/\ell A$.

It remains to verify that $\mathrm{Sp}(V)$ acts transitively on O. Take $f \in O$; there exists $h \in \mathrm{GSp}_q(V)$ so that the image of $h - 1$ is contained in the kernel of f. Let V_1 (resp. V_q) be the sum of generalized eigenspaces of h on V for all eigenvalues that reduce to 1 (resp. q) in $\bar{\mathbb{F}}_\ell$. By this we mean, more precisely, the following: set $\overline{V} = V \otimes \overline{\mathbb{Q}}_\ell$, and then set

$$V_1 = V \cap \bigoplus_{|\lambda - 1| < 1} \overline{V}_\lambda,$$

where $\overline{V}_\lambda$ is the generalized λ-eigenspace; V_q is defined similarly. Equivalently, V_1 (resp. V_q) consists of all $v \in V$ for which $(h-1)^n v \to 0$ (resp. $(h-q)^n v \to 0$) as $n \to \infty$.

Let W be the sum of all other generalized eigenspaces of h on V, i.e., for all eigenvalues λ that satisfy $|(\lambda - 1)(\lambda - q)| = 1$; in other words,

$$(8.9.2) \qquad\qquad W = \cap_{n=1}^\infty (h-1)^n (h-q)^n V.$$

Then, since q and 1 are distinct in $\mathbb{F}_\ell$, we have

$$(8.9.3) \qquad\qquad V = V_1 \oplus V_q \oplus W.$$

Indeed, given $v \in V$, we can certainly write $v = v_1 + v_q + w$, where v_1, v_q, w lie (respectively) in the $\mathbb{Q}_\ell$-spans of V_1, V_q, W. By applying a large power of $(h-1)(h-q)$ we deduce that $(h-1)^n(h-q)^n w \in W$ for sufficiently large n. But $(h-1)(h-q)$ is easily seen to be invertible on W, so in fact $w \in W$. It

follows then that $v_1 + v_q \in V$, and proceeding similarly we see $v_1 \in V_1, v_q \in V_q$, yielding (8.9.3).

Moreover, $V_1 \oplus V_q$ is orthogonal to W, and both V_1 and V_q are isotropic: to see this, let x be an element of V_1 and y an element of $V_1 \oplus W$. Then $(h-1)^n x$ approaches 0 as $n \to \infty$, while for all n, there exists $z_n \in V$ such that $y = (h-q)^n z_n$. Now

$$\begin{aligned} \omega(x,y) &= \omega(x, h(h-q)^{n-1} z_n) - \omega(x, q(h-q)^{n-1} z_n) \\ &= \omega(x, h(h-q)^{n-1} z_n) - \omega(hx, h(h-q)^{n-1} z_n) \\ &= \omega((1-h)x, (h-q)^{n-1} x_n), \text{ with } x_n = h z_n. \end{aligned}$$

Iterating, we see that $\omega(x,y)$ lies in $\omega((h-1)^n x, V)$; this being the case for all n, we have $\omega(x,y) = 0$. The proof that V_q is orthogonal to $V_q \oplus W$ is exactly the same.

By (ii), W is nondegenerate. By (iii), we can express $W = W_+ \oplus W_-$ as the sum of two isotropic submodules; since $V = (W_+ \oplus V_1) \oplus (W_- \oplus V_q)$ and both summands are isotropic, they are by (iv) both maximal isotropic. In particular, $W_- \oplus V_q$ is a maximal isotropic submodule of V, which furthermore lies in the image of $h-1$, and thus belongs to $\ker(f)$.

Now fix a decomposition $V = V_+ \oplus V_-$ into isotropic submodules, both free of rank g over $\mathbb{Z}_\ell$. Modifying f by an element of $\mathrm{Sp}(V)$, we may assume by (i) that f factors through the projection $V \longrightarrow V_+$. Since every automorphism of $\mathrm{GL}(V_+)$ is induced by an element of $\mathrm{Sp}(V)$, we are reduced to checking that any two surjections $V_+ \to A$ are conjugate under $\mathrm{GL}(V_+)$.

We must show that any two surjections $f_1, f_2 : \mathbb{Z}_\ell^g \to A$ are conjugate under $\mathrm{GL}_g(\mathbb{Z}_\ell)$. Fix $x_1, \ldots, x_k \in A$ such that the classes of x_i form a basis for $A/\ell A$ as a $\mathbb{Z}/\ell$-vector space. Then $x_1, \ldots, x_k$ generate A. Lift $x_1, \ldots, x_k$ to $y_1, \ldots, y_k \in \mathbb{Z}_\ell^g$. The y_i are linearly independent modulo ℓ, and so we can extend $y_1, \ldots, y_k$ to a $\mathbb{Z}_\ell$-basis $y_1, \ldots, y_g$ for $\mathbb{Z}_\ell^g$ where $f(y_i) = 0$ for $i > k$: simply extend arbitrarily to a basis, and then modify the y_i's for $i > g$ by linear combinations of $y_1, \ldots, y_k$ to ensure that $f(y_i) = 0$ for $i > k$. Similarly, lift $x_1, \ldots, x_k$ via f_2 to $y_1', \ldots, y_g'$; the automorphism of $\mathbb{Z}_\ell^g$ carrying y_i to y_i' then carries f_1 to f_2. $\qquad\square$

References

[1] D. ABRAMOVICH, A. CORTI, and A. VISTOLI, Twisted bundles and admissible covers, *Comm. Algebra* **31** (2003), 3547–3618, special issue in honor of Steven L. Kleiman. MR 2007376. Zbl 1077.14034. http://dx.doi.org/10.1081/AGB-120022434.

[2] J. D. ACHTER, Results of Cohen-Lenstra type for quadratic function fields, in *Computational Arithmetic Geometry*, *Contemp. Math.* **463**, Amer. Math. Soc.,

Providence, RI, 2008, pp. 1–7. MR 2459984. Zbl 1166.11018. http://dx.doi.org/10.1090/conm/463/09041.

[3] J. D. ACHTER and R. PRIES, The integral monodromy of hyperelliptic and trielliptic curves, *Math. Ann.* **338** (2007), 187–206. MR 2295509. Zbl 1129.11027. http://dx.doi.org/10.1007/s00208-006-0072-0.

[4] V. I. ARNOL′D, The cohomology ring of the group of dyed braids, *Mat. Zametki* **5** (1969), 227–231. MR 0242196. Zbl 0277.55002. http://dx.doi.org/10.1007/BF01098313.

[5] E. ARTIN, Theory of braids, *Ann. of Math.* **48** (1947), 101–126. MR 0019087. Zbl 0030.17703. http://dx.doi.org/10.2307/1969218.

[6] M. ARTIN, A. GROTHENDIECK, and J. L. VERDIER (eds.), *Théorie des topos et cohomologie étale des schémas. Tome* 1: *Théorie des topos*, *Lecture Notes in Mathematics*, *Vol.* 269, Springer-Verlag, Berlin-New York, 1972, Séminaire de Géométrie Algébrique du Bois-Marie 1963–1964 (SGA 4), avec la collaboration de N. Bourbaki, P. Deligne et B. Saint-Donat. MR 0354652. Zbl 0234.00007. http://dx.doi.org/10.1007/BFb0061319.

[7] A. BOREL, Stable real cohomology of arithmetic groups, *Ann. Sci. École Norm. Sup.* **7** (1974), 235–272 (1975). MR 0387496. Zbl 0316.57026. Available at http://www.numdam.org/item?id=ASENS_1974_4_7_2_235_0.

[8] C. P. BOYER, J. C. HURTUBISE, and R. J. MILGRAM, Stability theorems for spaces of rational curves, *Internat. J. Math.* **12** (2001), 223–262. MR 1823576. Zbl 1110.58303. http://dx.doi.org/10.1142/S0129167X01000721.

[9] D. BYEON, Indivisibility of class numbers of imaginary quadratic function fields, *Acta Arith.* **132** (2008), 373–376. MR 2413359. Zbl 1170.11042. http://dx.doi.org/10.4064/aa132-4-6.

[10] L. CARLITZ, The arithmetic of polynomials in a Galois field, *Proc. Natl. Acad. Sci. USA* **17** (1931), 120–122. Zbl 0001.12401. Available at http://www.pnas.org/content/17/2/120.full.pdf.

[11] R. CHARNEY and M. W. DAVIS, Finite $K(\pi, 1)$s for Artin groups, in *Prospects in Topology* (Princeton, NJ, 1994), *Ann. of Math. Stud.* **138**, Princeton Univ. Press, Princeton, NJ, 1995, pp. 110–124. MR 1368655. Zbl 0930.55006.

[12] F. R. COHEN, T. J. LADA, and J. P. MAY, *The Homology of Iterated Loop Spaces*, *Lecture Notes in Math.* **533**, Springer-Verlag, New York, 1976. MR 0436146. Zbl 0334.55009. http://dx.doi.org/10.1007/BFb0080464.

[13] H. COHEN and H. W. LENSTRA, JR., Heuristics on class groups of number fields, in *Number theory*, *Noordwijkerhout* 1983 (Noordwijkerhout, 1983), *Lecture Notes in Math.* **1068**, Springer-Verlag, New York, 1984, pp. 33–62. MR 0756082. Zbl 0558.12002. http://dx.doi.org/10.1007/BFb0099440.

[14] H. COHEN and J. MARTINET, Étude heuristique des groupes de classes des corps de nombres, *J. Reine Angew. Math.* **404** (1990), 39–76. MR 1037430. Zbl 0699.12016. http://dx.doi.org/10.1515/crll.1990.404.39.

[15] C. DAMIOLINI, The braid group and the arc complex, 2013. Available at http://www.math.leidenuniv.nl/scripties/DamioliniMaster.pdf?8.

[16] B. Datskovsky and D. J. Wright, Density of discriminants of cubic extensions, *J. Reine Angew. Math.* **386** (1988), 116–138. MR 0936994. Zbl 0632.12007. http://dx.doi.org/10.1515/crll.1988.386.116.

[17] M. Daub, J. Lang, M. Merling, A. M. Pacelli, N. Pitiwan, and M. Rosen, Function fields with class number indivisible by a prime ℓ, 2011, pp. 339–359. MR 2847264. Zbl 1263.11098. http://dx.doi.org/10.4064/aa150-4-2.

[18] A. J. DeJong and N. Katz, Counting the number of curves over a finite field, 2001, preprint.

[19] P. Deligne, La conjecture de Weil. I, *Inst. Hautes Études Sci. Publ. Math.* **43** (1974), 273–307. MR 0340258. Zbl 0287.14001. http://dx.doi.org/10.1007/BF02684373.

[20] P. Deligne, La conjecture de Weil. II, *Inst. Hautes Études Sci. Publ. Math.* **52** (1980), 137–252. MR 0601520. Zbl 0456.14014. http://dx.doi.org/10.1007/BF02684780.

[21] P. Deligne and N. Katz (eds.), *Groupes de Monodromie en Géométrie Algébrique. II*, *Lecture Notes in Math.* **340**, Springer-Verlag, New York, 1973, Séminaire de Géométrie Algébrique du Bois-Marie 1967–1969 (SGA 7 II). MR 0354657. Zbl 0258.00005. http://dx.doi.org/10.1007/BFb0060505.

[22] P. Deligne and M. Rapoport, Les schémas de modules de courbes elliptiques, in *Modular Functions of One Variable*, *II* (Proc. Internat. Summer School, Univ. Antwerp, Antwerp, 1972), *Lecture Notes in Math.* **349**, Springer-Verlag, New York, 1973, pp. 143–316. MR 0337993. Zbl 0281.14010. http://dx.doi.org/10.1007/978-3-540-37855-6.

[23] J. S. Ellenberg and A. Venkatesh, Counting extensions of function fields with bounded discriminant and specified Galois group, in *Geometric Methods in Algebra and Number Theory*, *Progr. Math.* **235**, Birkhäuser, Boston, 2005, pp. 151–168. MR 2159381. Zbl 1085.11057. http://dx.doi.org/10.1007/0-8176-4417-2_7.

[24] J. S. Ellenberg and A. Venkatesh, Letter to A. J. de Jong, 2008.

[25] É. Fouvry and J. Klüners, On the 4-rank of class groups of quadratic number fields, *Invent. Math.* **167** (2007), 455–513. MR 2276261. Zbl 1126.11062. http://dx.doi.org/10.1007/s00222-006-0021-2.

[26] E. Freitag and R. Kiehl, *Étale Cohomology and the Weil Conjecture*, *Ergeb. Math. Grenzgeb.* **13**, Springer-Verlag, New York, 1988, translated from the German by Betty S. Waterhouse and William C. Waterhouse, with an historical introduction by J. A. Dieudonné. MR 0926276. Zbl 0643.14012. http://dx.doi.org/10.1007/978-3-662-02541-3.

[27] M. D. Fried and H. Völklein, The inverse Galois problem and rational points on moduli spaces, *Math. Ann.* **290** (1991), 771–800. MR 1119950. Zbl 0763.12004. http://dx.doi.org/10.1007/BF01459271.

[28] E. Friedman and L. C. Washington, On the distribution of divisor class groups of curves over a finite field, in *Théorie des Nombres* (Quebec, PQ, 1987), de Gruyter, Berlin, 1989, pp. 227–239. MR 1024565. Zbl 0693.12013. http://dx.doi.org/10.1515/9783110852790.227.

[29] C. Friesen, Class number divisibility in real quadratic function fields, *Canad. Math. Bull.* **35** (1992), 361–370. MR 1184013. Zbl 0727.11045. http://dx.doi.org/10.4153/CMB-1992-048-5.

[30] D. Garton, Random matrices, the Cohen-Lenstra heuristics, and roots of unity, *Algebra Number Theory* **9** (2015), 149–171. MR 3317763. Zbl 06424744. http://dx.doi.org/10.2140/ant.2015.9.149.

[31] A. Grothendieck, *Revêtements étales et Groupe Fondamental. Fasc. II*: *Exposés* 6, 8 à 11, *Séminaire de Géométrie Algébrique* **1960/61**, Institut des Hautes Études Scientifiques, Paris, 1963. MR 0217088.

[32] R. Hain and E. Looijenga, Mapping class groups and moduli spaces of curves, in *Algebraic Geometry—Santa Cruz* 1995, *Proc. Sympos. Pure Math.* **62**, Amer. Math. Soc., Providence, RI, 1997, pp. 97–142. MR 1492535. Zbl 0914.14013.

[33] C. Hall, Big symplectic or orthogonal monodromy modulo l, *Duke Math. J.* **141** (2008), 179–203. MR 2372151. Zbl 1205.11062. http://dx.doi.org/10.1215/S0012-7094-08-14115-8.

[34] J. L. Harer, Stability of the homology of the mapping class groups of orientable surfaces, *Ann. of Math.* **121** (1985), 215–249. MR 0786348. Zbl 0579.57005. http://dx.doi.org/10.2307/1971172.

[35] A. Hatcher and N. Wahl, Stabilization for mapping class groups of 3-manifolds, *Duke Math. J.* **155** (2010), 205–269. MR 2736166. Zbl 1223.57004. http://dx.doi.org/10.1215/00127094-2010-055.

[36] A. Hurwitz, Über Riemann'sche Flächen mit gegebenen Verzweigungspunkten, *Math. Ann.* **39** (1891), 1–60. JFM 23.0429.01.

[37] N. M. Katz and S. Lang, Finiteness theorems in geometric classfield theory, *Enseign. Math.* **27** (1981), 285–319 (1982), with an appendix by Kenneth A. Ribet. MR 0659153. Zbl 0495.14011.

[38] F. F. Knudsen, The projectivity of the moduli space of stable curves. II. The stacks $M_{g,n}$, *Math. Scand.* **52** (1983), 161–199. MR 0702953. Zbl 0544.14020. Available at http://www.mscand.dk/article/view/12001/10017.

[39] G. Malle, Cohen-Lenstra heuristic and roots of unity, *J. Number Theory* **128** (2008), 2823–2835. MR 2441080. Zbl 1225.11143. http://dx.doi.org/10.1016/j.jnt.2008.01.002.

[40] H. Matsumura, *Commutative Ring Theory, Cambridge Stud. Adv. Math.* **8**, Cambridge Univ. Press, Cambridge, 1986, translated from the Japanese by M. Reid. MR 0879273. Zbl 0603.13001.

[41] J. Milnor and D. Husemoller, *Symmetric Bilinear Forms, Ergeb. Math. Grenzgeb.* **73**, Springer-Verlag, New York, 1973. MR 0506372. Zbl 0292.10016.

[42] R. S. Palais, Local triviality of the restriction map for embeddings, *Comment. Math. Helv.* **34** (1960), 305–312. MR 0123338. Zbl 0207.22501. Available at http://retro.seals.ch/digbib/view?pid=com-001:1960:34::24.

[43] D. Quillen, Finite generation of the groups K_i of rings of algebraic integers, in *Algebraic K-theory, I: Higher K-theories* (Proc. Conf., Battelle Memorial Inst., Seattle, Wash., 1972), *Lecture Notes in Math.* **341**, Springer-Verlag, New York, 1973, pp. 179–198. MR 0349812. Zbl 0355.18018.

[44] M. Romagny and S. Wewers, Hurwitz spaces, in *Groupes de Galois Arithmétiques et Différentiels*, *Sémin. Congr.* **13**, Soc. Math. France, Paris, 2006, pp. 313–341. MR 2316356. Zbl 1156.14314.

[45] M. Rosen, S-units and S-class group in algebraic function fields, *J. Algebra* **26** (1973), 98–108. MR 0327777. Zbl 0265.12003. http://dx.doi.org/10.1016/0021-8693(73)90036-7.

[46] M. Salvetti, Topology of the complement of real hyperplanes in $\mathbf{C}^N$, *Invent. Math.* **88** (1987), 603–618. MR 0884802. Zbl 0594.57009. http://dx.doi.org/10.1007/BF01391833.

[47] G. Segal, The topology of spaces of rational functions, *Acta Math.* **143** (1979), 39–72. MR 0533892. Zbl 0427.55006. http://dx.doi.org/10.1007/BF02392088.

[48] N. Wahl, Homological stability for mapping class groups of surfaces, in *Handbook of Moduli. Vol. III* (Somerville, MA), *Adv. Lect. Math.* (*ALM*) **26**, Int. Press, 2013, pp. 547–583. MR 3135444. Zbl 1322.57016.

[49] L. C. Washington, Some remarks on Cohen-Lenstra heuristics, *Math. Comp.* **47** (1986), 741–747. MR 0856717. Zbl 0627.12002. http://dx.doi.org/10.2307/2008187.

[50] S. Wewers, Construction of Hurwitz spaces, 1998, Thesis, U. Duisburg-Essen. Zbl 0925.14002.

[51] J.-K. Yu, Toward a proof of the Cohen-Lenstra conjecture in the function field case, *preprint* (1997).

University of Wisconsin, Madison, WI
E-mail: ellenber@math.wisc.edu

Stanford University, Stanford, CA
E-mail: akshay@math.stanford.edu

University of Minnesota, Minneapolis, MN
E-mail: cwesterl@umn.edu