NOTES ON ARITHMETIC STATISTICS OVER FUNCTION FIELDS

AARON LANDESMAN

CONTENTS

1. INTRODUCTION

Three of the main conjectures investigated in arithmetic statistics are the Cohen-Lenstra conjectures, Malle's conjecture, and the Poonen-Rains conjectures. Throughout the notes, we will try to explain recent progress toward these conjectures over function fields via homological stability techniques.

The general goal of the notes will be to explain the main ideas of these proofs without getting bogged down in the technical details. For this reason, we will avoid presenting many proofs in maximum generality, but instead try to focus on the simplest cases.

We begin by describing what these three conjectures are about.

1.1. **The Cohen-Lenstra heuristics.**

Let ℓ be an odd prime. The Cohen–Lenstra heuristics, introduced in 1984 by Cohen and Lenstra [CL84], predict that the probability the ℓ-part of the class group $\mathrm{Cl}(\mathbb{Q}(\sqrt{-d}))$ is isomorphic to a given finite ℓ-group G is inversely proportional to $\#\mathrm{Aut}(G)$. More formally,

$$\lim_{X\to\infty} \frac{\#\{d : d \text{ squarefree}, \mathrm{Cl}(\mathbb{Q}(\sqrt{-d}))[\ell] \simeq G, \mathrm{Disc}(\mathbb{Q}(\sqrt{(-d)})) \le X\}}{\#\{d : d \text{ squarefree}, \mathrm{Disc}(\mathbb{Q}(\sqrt{(-d)})) \le X\}} \sim \frac{1}{\#\mathrm{Aut}(G)}.$$

They make essentially the same prediction for the ℓ Sylow subgroup of the class group. Here, the probability is taken over those d so that the discriminant of $\mathbb{Q}(\sqrt{-d})$ is bounded by X, and then one takes a limit as $X \to \infty$. An important consequence of this conjecture is the prediction that the average size of $\mathrm{Cl}(\mathbb{Q}(\sqrt{-d}))[\ell]$ is 2. More precisely, let $IQ_{\le X}$ denote the set of imaginary quadratic fields (of the form $\mathbb{Q}(\sqrt{-d})$ with discriminant having absolute value at most X. Let $RQ_{\le X}$ denote the set of real quadratic fields (of the form $\mathbb{Q}(\sqrt{-d})$ with discriminant at most X.

Conjecture 1.1.1. As K ranges over imaginary quadratic fields of discriminant $\le X$, and ℓ an odd prime number, the average size of ℓ-torsion in the class group of $\mathscr{O}_K$ is 2. That is,

$$\lim_{X\to\infty} \mathbb{E}_{K\in IQ_{\le X}}(\#\mathrm{Cl}(\mathscr{O}_K)[\ell]) = 2.$$

Similarly, as K ranges over real quadratic fields of discriminant less than X, the average size of ℓ torsion in the class group has average size $1 + 1/\ell$:

$$\lim_{X\to\infty} \mathbb{E}_{K\in RQ_{\le X}}(\#\mathrm{Cl}(\mathscr{O}_K)[\ell]) = 1 + 1/\ell.$$

Remark 1.1.2. As we will see repeatedly throughout the course, it is often more natural to write this 2 as $2 = 1 + 1$ where the first 1 corresponds to the trivial element of the class group and the second 1 corresponds to nontrivial elements.

Exercise 1.1.3.
1. Let $sf(X)$ denote the number of integers $0 \leq n \leq X$ which are squarefree (don't have any prime p with $p^2 \mid n$). Compute $\lim_{X \to \infty} \frac{sf(X)}{X}$.
2. Compute the asymptotic growth of the number of imaginary quadratic extensions of $\mathbb{Q}$ of discriminant at most X. You should have an answer of the for $CX^s + o(X)$ for appropriate constants C and s that you should compute.
3. Similarly, find the asymptotic growth of the number of real quadratic extensions.

This essentially computes the "denominator" in Cohen-Lenstra. *Hint:* Quadratic extensions are in bijection with squarefree integers, in the sense that they are all of the form $\mathbb{Q}(\sqrt{d})$. However, the discriminant of such an extension may be either d or $4d$. Use a sieve to count these conditioning on the value of d modulo 4.

There is also a natural function field version of this conjecture, where one replaces $\mathbb{Q}$ by $\mathbb{F}_q(t)$ and one replaces imaginary quadratic extensions with monic hyperelliptic curves, branched over $\infty \in \mathbb{P}^1_{\mathbb{F}_q}$. Despite the vast array of work around this conjecture, apart from $\ell = 3$ (due to Davenport and Heilbronn [DH71] over $\mathbb{Q}$ and the analog over function fields due to Datskovsky and Wright in [DW88]) until the material we discuss in this course, not a single other case had been proven. Over function fields, we have recently been able to compute the average size of ℓ torsion in class groups over $\mathbb{F}_q$, for sufficiently large q, depending on ℓ, and prime to 2ℓ. One of the goals of this course is to explain the ideas behind this proof.

1.2. **Malle's conjecture, counting *G* extensions.**

Another central question in number theory asks: how many degree n extensions of $\mathbb{Q}$ or $\mathbb{F}_q(t)$ are there? Of course there are infinitely many, but we can try to answer how the number of extensions grows as the discriminant grows. It is also natural to count by other invariants, such as the product of primes dividing the discriminant. More generally, for G a finite group how many Galois G extensions are there, for G a finite group? In 2002, Malle [Mal02] proposed a conjecture for the number of such extensions. More generally still, one may ask for the number of G extensions where one prescribes the local inertia at each ramification point of the extension to lie in a certain conjugacy class.

Here is a specific prediction of Malle's conjecture for degree d, S_d extensions:

Conjecture 1.2.1. There is a constant c_d, depending on d, so that the number of degree d, S_d extensions of $\mathbb{Q}$ of discriminant at most X grows as $c_d X$.

Despite a vast array of activity in the past several decades, in some sense, very few cases of Malle's conjecture are known. For example, the only symmetric groups where it is known to hold are S_d for $d \leq 5$, and the cases $d = 4$ and 5 were a significant part of Bhargava's work, leading to his fields medal. Bhargava also separately made an interesting conjecture for the value of c_d. (We will give a proof of a function field version of this conjecture for arbitrary d and q sufficiently large in §19.)

Another primary goal of this course will be to explain how to compute the asymptotic number of G extensions of $\mathbb{F}_q(t)$, for q sufficiently large relative to $\#G$ and q prime to $\#G$.

1.3. **Poonen-Rains heuristics and the minimalist conjecture.** The minimalist conjecture, originally formulated by Goldfeld in 1979 [Gol79] predicts that, among all quadratic twists of a suitable elliptic curve E over $\mathbb{Q}$, half will have rank 0 and half will have rank 1. This conjecture is closely related to the Poonen-Rains heuristics, which predict the distribution of Selmer groups of elliptic curves in such families. Recall that the ℓ-Selmer group is a finite $\mathbb{Z}/\ell\mathbb{Z}$ module which contains as a subgroup $(\mathbb{Z}/\ell\mathbb{Z})^{\operatorname{rk}(E)}$, and so in particular, determining the distribution of these ℓ-Selmer groups will give a lot of information about the rank of elliptic curves. For example, the Poonen-Rains heuristics, introduced in a highly influential 2012 JAMS paper [PR12], predict that the average size of the ℓ Selmer groups of elliptic curves in suitable families is $\ell + 1$, for ℓ a prime number:

Conjecture 1.3.1. For a fixed global field k and prime ℓ, as E varies over all elliptic curves over k ordered by discriminant, the average value of $\operatorname{Sel}_\ell(E)$ is $\ell + 1$.

This and some generalizations are known for $\ell = 2$ in certain families. Its generalization for ℓ replaced by a power of 2 has been recently solved by Alex Smith [Smi22] for many families, and for $\ell \leq 5$ for the family of all elliptic curves was resolved in a series papers, two of which have appeared in Annals papers by Bhargava and Shankar [BS15a, BS15b, BS13a, BS13b].

Another of our primary goals in this course, which we will discuss in the later part of the course, will be to explain how to prove a version of Conjecture 1.3.1 in many quadratic twist families of elliptic curves over $\mathbb{F}_q(t)$ for arbitrary odd ℓ and q sufficiently large, depending on ℓ and prime to 2ℓ.

1.4. **Summarizing tables.** In Table 2, we summarize the current state of the art in function field results relating to the above arithmetic statistics conjectures. In particular, we highlight progress toward what we view as the three most important conjectures in arithmetic statistics: the Cohen-Lenstra conjectures about torsion in class groups, Malle's conjecture about counting

	Arithmetic statistics result	Topological input
(1)	Computation of $\lim_{n\to\infty}\lim_{q\to\infty}\frac{\#X_n(\mathbb{F}_q)}{q^n}$	The 0th homology stabilizes, i.e., $\lim_{n\to\infty} h_0(X_n)$ exists and we can compute its value
(2)	Computation of $\lim_{q\to\infty}\limsup_{n\to\infty}\frac{\#X_n(\mathbb{F}_q)}{q^n}$	The higher homologies stabilize, i.e., $\lim_{n\to\infty} h_i(X_{n,\mathbb{C}})$ exists for all i
(3)	$\lim_{n\to\infty}\frac{\#X_n(\mathbb{F}_q)}{q^n}$ exists, but with unknown value	The higher homologies stabilize Frobenius equivariantly, i.e., $\lim_{n\to\infty}\operatorname{tr}(\mathrm{Frob}_q \mid H^i(X_{n,\overline{\mathbb{F}}_q},\mathbb{Q}_{\ell'}))$ exists for all i
(4)	Computation of $\lim_{n\to\infty}\frac{\#X_n(\mathbb{F}_q)}{q^n}$ for sufficiently large q	A computation of the stable homology determination of $\lim_{n\to\infty} h_i(X_{n,\mathbb{C}})$
(5)	Computation of $\lim_{n\to\infty}\frac{\#X_n(\mathbb{F}_q)}{q^n}$ for all q	A subexponential bound on the trace of Frobenius acting on the unstable homology

TABLE 1. An indication of the required topological input to deduce different arithmetic statistical results, ordered by difficulty

G-covers, and the Poonen-Rains conjectures about Selmer groups of elliptic curves. The rows are labeled $1-5$ corresponding to the 5 limits detailed in the rows in Table 1, which also describes the topological input that would allow one to compute the relevant limits.

Remark 1.4.1. As one can see in Table 2, it still remains to prove all the main conjectures for fixed small q, and complete the bottom line of the table. The various conjectures (Cohen-Lenstra, Malle, and Poonen-Rains) essentially predict that the trace of Frobenius on the stable cohomology is the main contribution to the point counts, and hence the trace of Frobenius on the unstable cohomology must nearly cancel. I speculate that such cancellation must occur in some extremely structured way.

This mysterious cancellation seems to me perhaps the biggest open problem in arithmetic statistics over function fields.

	Cohen-Lenstra	Malle	Poonen-Rains
(1)	Achter [Ach06]	Partial progress by Türkelli [TÏ5] building on Ellenberg-Venkatesh [EV05]	Landesman [Lan21], Feng-Landesman-Rains, [FLR23], and Park-Wang [PW23]
(2)	Ellenberg-Venkatesh-Westerland [EVW16]	Landesman-Levy [LL25a]	Ellenberg-Landesman [EL24]
(3)	Landesman-Levy [LL24]	Landesman-Levy [LL25a]	Ellenberg-Landesman [EL24]
(4)	Landesman-Levy [LL24]	Landesman-Levy [LL25a]	Landesman-Levy [LL25b]
(5)	?	?	?

TABLE 2. Table indicating relevant papers for the 5 categories outlined in Table 1. The question mark categories remain open.

2. THE FUNCTION FIELD PERSPECTIVE

We'll now briefly indicate how the above problems can be translated to the function field setting, and what the extra juice from algebraic geometry and topology we obtain there is.

Before giving examples, let us outline the general strategy, which comes in a few steps

1. Rephrase the arithmetic statistics conjectures in terms of counting the $\mathbb{F}_q$ points of a certain sequence of varieties or stacks.
2. Relate the number of $\mathbb{F}_q$ points of these varieties to to their cohomology, using the Grothendieck-Lefschetz trace formula.
3. Understand their cohomology as best you can.

2.1. Malle rephrasal.

We start with describing how one could approach Malle's conjecture in the function field setting, because it is a bit easier to understand than the approaches to Cohen-Lenstra and Poonen-Rains. The general strategy we will use time and time again is to translate these conjectures into problems where we wish to count the number of $\mathbb{F}_q$ points of certain varieties.

We start with a loose definition of the Hurwitz stack we will be interested in. Later in §4, specifically Definition 4.1.1, we will come back to give a more precise definition.

Definition 2.1.1 (Temporary definition)**.** Let G be a finite group, B be a base scheme on which $\#G$ is invertible, $c \subset G$ a union of conjugacy classes, and n an integer. We define $[\mathrm{Hur}^{G,c}_{n,B}/G]$ to be the moduli stack whose S-points, for S a B scheme parameterize G-Galois covers $X \to \mathbb{P}^1_S$ over S with Galois group G, branched at a degree n subscheme of $\mathbb{A}^1$ with inertia in c. (Arbitrary branching is allowed over ∞.)

Example 2.1.2. For example, in the case $G = S_d$ and c is the set of transpositions, one can identify $[\mathrm{Hur}^{G,c}_{n,B}/G]$ with the stack of degree d covers of $\mathbb{P}^1$ which are simply branched over $\mathbb{A}^1$ but with arbitrary branching allowed over $\infty \in \mathbb{P}^1$.

The use of the above definition is that, using the correspondence between curves and function fields, Malle's conjecture predicts the value of

$$\lim_{n\to\infty} \frac{\#[\mathrm{Hur}^{G,c}_{n,\mathbb{F}_q}/G](\mathbb{F}_q)}{q^n}.$$

(Technically, Malle's conjecture usually assumes the covers are connected, so one only counts the subset of the points of $[\mathrm{Hur}^{G,c}_{n,\mathbb{F}_q}/G]$ corresponding to connected covers, but the above is a reasonable first approximation.) Hence, we have rephrased Malle's conjecture as a problem about counting finite field points of certain stacks. In particular, we will be able to compute this value if we understand the cohomology of Hurwitz stacks sufficiently well, via the relation between finite field point counts and cohomology, discussed in Theorem 3.1.2.

2.2. **Cohen-Lenstra rephrasal.**

The next problem we would like to cast in the light of function field point counting is the Cohen-Lenstra heuristics. For this, we need a slight variant of the definition of Hurwitz spaces. Again, for a first pass, we will start with a somewhat loose definition, but later we'll give a more precise definition.

Definition 2.2.1. Given a scheme B, there is an open subscheme $U \subset \mathbb{A}^n_B$ parameterizing the locus where all coordinates are distinct. There is an action of the symmetric group S_n on U by permuting the coordinates and we define $\mathrm{Conf}_{n,B} := U/S_n$ to be the *configuration space* of n points in $\mathbb{A}^1$ over B. More generally, let $\mathrm{Conf}_{n_1,\ldots,n_v,B} := U/S_{n_1} \times \cdots \times S_{n_v}$ where S_{n_i} acts on the n consecutive coordinates in the range $[n_1 + \cdots + n_{i-1} + 1, n_1 + \cdots + n_i]$. When $B = \operatorname{Spec} R$, for R a ring, we often write this as $\mathrm{Conf}_{n_1,\ldots,n_v,R}$ and when $R = \mathbb{C}$, we abbreviate this to $\mathrm{Conf}_{n_1,\ldots,n_v}$.

Definition 2.2.2 (Temporary imprecise definition of Definition 4.1.13)**.** Let G be a finite group, B be a base scheme on which G is invertible, $c \subset G$ a union of conjugacy classes, and n an integer. We define the *pointed Hurwitz scheme* $\mathrm{Hur}^{G,c}_{n,B}$ to be the G-cover of $[\mathrm{Hur}^{G,c}_{n,B}/G]$ which is a $|c|^n$ finite étale cover of $\mathrm{Conf}_{n,B}$ whose fiber over a geometric point of $\mathrm{Conf}_{n,B}$ can be identified with all $|c|^n$ labelings of the n branch points by elements of c.

We also use use $\mathrm{CHur}_{n,B}^{G,c} \subset \mathrm{Hur}_{n,B}^{G,c}$ to denote the union of components parameterizing covers with geometrically connected source.

Remark 2.2.3. In the case a component of $[\mathrm{Hur}_{n,B}^{G,c}/G]$ parameterizes covers unramified over ∞, the corresponding component of $\mathrm{Hur}_{n,B}^{G,c}$ parameterizes G covers with inertia in c, branched at a degree n subscheme, together with a marked point over ∞. This underlies the general fact that $\mathrm{Hur}_{n,B}^{G,c} \to [\mathrm{Hur}_{n,B}^{G,c}/G]$ is a G Galois cover since the set of possible markings of the point over ∞ is a G-torsor.

Now, what is the relation between this pointed Hurwitz scheme and the Cohen-Lenstra heuristics?

Fix a quadratic field $K/\mathbb{F}_q(t)$, which we will assume is split over ∞, so the fiber of the corresponding curve $C_K \to \mathbb{P}^1_{\mathbb{F}_q}$ has two $\mathbb{F}_q$ points over ∞. (This can be thought of as the analog of a real quadratic extension since there are two maps $\operatorname{Spec}\mathbb{R} \to \operatorname{Spec}\mathscr{O}_K$ for K a real imaginary quadratic extension.)

Definition 2.2.4. A degree d Galois extension L of such a quadratic field K which is split over ∞ is *totally split over* ∞ if $C_L \to C_K \to \mathbb{P}^1_{\mathbb{F}_q}$ is totally split, meaning that it has $2d$ points in the preimage over ∞.

Lemma 2.2.5. *Let ι be the nontrivial element of $\mathbb{Z}/2\mathbb{Z}$, let ℓ be a prime, let $D_{2\ell}$ be the dihedral group of order 2ℓ, and let $c \subset D_{2n}$ denote the conjugacy class of order 2 elements. There is a map $\mathrm{Hur}_{n,B}^{D_{2\ell},c} \to \mathrm{Hur}_{n,B}^{\mathbb{Z}/2\mathbb{Z},\iota}$ induced by the map $D_{2n} \to \mathbb{Z}/2\mathbb{Z}$ which sends a $D_{2\ell}$ Galois cover to the intermediate degree 2 cover. We use $\mathscr{O}_K$ to denote the normalization of $\mathbb{F}_q[t]$ in K. Viewing the quadratic extension $K/\mathbb{F}_q(t)$ split over ∞ (together with a choice of marked point over ∞) as a point of $\mathrm{CHur}_{n,B}^{\mathbb{Z}/2\mathbb{Z},\iota}(\mathbb{F}_q)$, the fiber of $\mathrm{CHur}_{n,B}^{D_{2\ell},c}(\mathbb{F}_q) \to \mathrm{CHur}_{n,B}^{\mathbb{Z}/2\mathbb{Z},\iota}(\mathbb{F}_q)$ over this point has cardinality $\ell(\#\mathrm{Cl}(\mathscr{O}_K)[\ell]-1)$.*

Proof. A result from class field theory [Hay92, 15.6] gives an identification between $\mathrm{Cl}(\mathscr{O}_K)$ and the Galois group of the maximal abelian unramified extension of K split completely over ∞. Now, we can identify elements of $\mathrm{Cl}(\mathscr{O}_K)[\ell]$ with maps $\mathbb{Z}/\ell\mathbb{Z} \to \mathrm{Cl}(\mathscr{O}_K)$, and this set has the same size as the set of maps $\mathrm{Cl}(\mathscr{O}_K) \to \mathbb{Z}/\ell\mathbb{Z}$, which are, by the above correspondence, in bijection with $\mathbb{Z}/\ell\mathbb{Z}$ unramified Galois (not necessarily connected) covers of $\mathscr{O}_K$ split completely over ∞. Such covers can also be viewed as covers of $\mathbb{P}^1$, and since they are unramified over $\mathscr{O}_K$, their inertia over $\mathbb{P}^1$ must lie in c.

Finally, using the description of Remark 2.2.3, we can identify such covers with points of $\mathrm{Hur}_{n,B}^{D_{2\ell},c}$, once we verify these covers have function field M which is Galois over $\mathbb{F}_q(t)$ with Galois group $D_{2\ell}$. We leave this to the reader in the next exercise.

Exercise 2.2.6. Show that $M/\mathbb{F}_q(t)$ is a Galois extension with Galois group $D_{2\ell}$ as follows:

1. Let L/K denote the maximal abelian unramified extension of odd order, completely split over ∞. This is Galois with Galois group equal to the group of odd order elements in $\mathrm{Cl}(\mathscr{O}_K)$ by class field theory. Show that $L/\mathbb{F}_q(t)$ is also Galois. *Hint:* If one chooses an embedding $K \to \overline{\mathbb{F}_q(t)}$, show that the automorphism of K over $\mathbb{F}_q(t)$ preserves L.
2. Show that $\mathrm{Gal}(L/\mathbb{F}_q(t)) \simeq \mathrm{Gal}(K/\mathbb{F}_q(t)) \ltimes \mathrm{Gal}(L/K)$, for some choice of action. *Hint:* Use the Sylow theorem.
3. Show that in fact, the action of $\mathrm{Gal}(K/\mathbb{F}_q(t))$ on $\mathrm{Gal}(L/K)$ is by negation (or inversion) on the abelian group $\mathrm{Gal}(L/K)$. *Hint:* Show that there are no geometrically connected extensions of $\mathbb{F}_q(t)$ of odd order with all ramification degrees dividing 2.
4. Conclude that for any M/K a cyclic $\mathbb{Z}/\ell\mathbb{Z}$ extension with ℓ odd, $M/\mathbb{F}_q(t)$ is Galois and has Galois group $D_{2\ell}$.

Moreover, the map $\mathrm{Cl}(\mathscr{O}_K) \to \mathbb{Z}/\ell\mathbb{Z}$ is surjective (or equivalently nontrivial) if and only if the corresponding cover is connected. Moreover, since the cover is assumed to be split over ∞ it has an $\mathbb{F}_q$ point, and since it is connected it must also be geometrically connected (as connected varieties over $\mathbb{F}_q$ which are not geometrically connected have no $\mathbb{F}_q$ points) That is, the cover corresponds to a point of $\mathrm{CHur}_{n,B}^{D_{2\ell},c}$. We have overcounted such covers by a factor of ℓ since the $\mathbb{F}_q$ points of $\mathrm{CHur}_{n,B}^{D_{2\ell},c}$ count covers with a marked point over ∞ over the specified marked point of C_K over ∞, and there are ℓ such points since the cover has degree ℓ. Altogether, this gives that the fiber of $\mathrm{CHur}_{n,B}^{D_{2\ell},c} \to \mathrm{CHur}_{n,B}^{\mathbb{Z}/2\mathbb{Z},\iota}$ over the point corresponding to $\mathscr{O}_K$ has cardinality $\ell(\#\mathrm{Cl}(\mathscr{O}_K)[\ell]-1)$, where the -1 corresponds to removing the trivial map $\mathrm{Cl}(\mathscr{O}_K) \to \mathbb{Z}/\ell\mathbb{Z}$ and the multiplication by ℓ corresponds to the overcounting described above. □

Now, using Lemma 2.2.5, we see that the average size of the ℓ-torsion in the class group of a real quadratic field can be expressed as the ratio

$$\lim_{n\to\infty, n\equiv 0 \bmod 2} \frac{1}{\ell} \frac{\#\mathrm{CHur}_{n,B}^{D_{2\ell},c}(\mathbb{F}_q)}{\#\mathrm{CHur}_{n,B}^{\mathbb{Z}/2\mathbb{Z},\iota}(\mathbb{F}_q)}, \tag{2.1}$$

which again reduces this version of the Cohen-Lenstra heuristics to point counting on Hurwitz spaces. This ratio is predicted to be $1/\ell$, corresponding to the fact that there is 1 such component both of $\mathrm{CHur}_{n,\mathbb{F}_q}^{D_{2\ell},c}$ and of $\mathrm{CHur}_{n,\mathbb{F}_q}^{\mathbb{Z}/2\mathbb{Z},\iota}$.

Remark 2.2.7. The above is the prediction of Cohen-Lenstra for the real quadratic extensions. It is slightly different from the corresponding prediction of $1 = \ell/\ell$ for the imaginary quadratic case, corresponding to the fact that there are ℓ such components in that case instead of only one.

Exercise 2.2.8. Suppose $\mathbb{F}_q$ is a finite field of characteristic not 2. Compute the number of monic (meaning they have leading coefficient 1) squarefree polynomials of degree n. *Hint:* Any monic polynomial f can be written uniquely as $f = gh^2$ where g and h are monic and g is squarefree.

Exercise 2.2.9. We assume $\mathbb{F}_q$ is a finite field of characteristic not 2. Assume $n > 1$. Show that the number of isomorphism classes degree 2 covers $X \to \mathbb{A}^1_{\mathbb{F}_q}$ with smooth source and discriminant degree q^n (meaning the degree of the relative sheaf of differentials is n) is $2(q^n - q^{n-1})$. Here, two covers $X_1 \to \mathbb{A}^1$ and $X_2 \to \mathbb{A}^1$ are isomorphic if there is an isomorphism $X_1 \to X_2$ so that cover $X_1 \to \mathbb{A}^1$ agrees with the composite $X_1 \to X_2 \to \mathbb{A}^1$. You may assume that every such curve can be written as the vanishing locus of $y^2 - f(x)$ for $f(x)$ a degree n polynomial.

Remark 2.2.10. Usually, when count points, we weight the counts by the automorphisms of the corresponding objects. Therefore, if we consider each curve in Exercise 2.2.9 as having two automorphisms, the count would be $q^n - q^{n-1}$. Therefore, we have computed the "denominator" of Cohen-Lenstra (meaning the denominator of (2.1)) in the function field case.

The next exercise further explores counting objects up to isomorphism.

Exercise 2.2.11. Recall that for G a finite group, BG is the stack (over $\operatorname{Spec} \mathbb{Z}$ so that a map $X \to BG$ is the same as a G-torsor over X.

1. Compute $B(\mathbb{Z}/2\mathbb{Z})(\mathbb{F}_q)$ as a set.
2. When X is a stack, we use $\#X(\mathbb{F}_q) := \sum_{x \in X} \frac{1}{\#\operatorname{Aut}(x)}$. Using this, compute $\#B(\mathbb{Z}/2\mathbb{Z})(\mathbb{F}_q)$.
3. More generally, if G is a finite group, compute $\#BG(\mathbb{F}_q)$. *Hint:* Show there is a bijection between the set of isomorphism classes of G torsors over $\mathbb{F}_q$ and conjugacy classes in G by thinking of a scheme over $\mathbb{F}_q$ in terms of its base change to $\overline{\mathbb{F}}_q$ with an action of Frobenius.

2.3. **Poonen-Rains rephrasal.** Third, we'd like to understand how to express the Poonen-Rains conjectures in terms of counting points on Hurwitz spaces. This is perhaps the trickiest relation to understand, and is one of the main innovations of [EL24]. We will briefly describe the setup here, and come back to it later in the course.

The general idea is that if one starts with an elliptic curve E over $\mathbb{F}_q(t)$, one can form the Néron model $\mathscr{E}$, which is a smooth scheme over $\mathbb{P}^1_{\mathbb{F}_q}$ whose generic fiber is E. For ℓ a prime, it turns out that, in good circumstances, the ℓ-Selmer group can be identified with the étale cohomology group $H^1(\mathbb{P}^1, \mathscr{E}[\ell])$. The latter parameterizes torsors for the group scheme $\mathscr{E}[\ell]$ which is a quasi-finite group scheme of degree ℓ^2 over $\mathbb{P}^1$. In other words, torsors for this group scheme correspond to certain degree ℓ^2 covers of $\mathbb{P}^1$. These covers are not Galois, but it turns out one can pin down their Galois closure to lie the group $\mathrm{AGL}(\mathbb{Z}/\ell\mathbb{Z}) := \mathrm{GL}(\mathbb{Z}/\ell\mathbb{Z}) \rtimes (\mathbb{Z}/\ell\mathbb{Z})^2$. Because the original elliptic curve has certain points of bad reduction, these covers will all have particular types of monodromy at those points of bad reduction. It turns out the average size of Selmer groups in quadratic twist families associated to an elliptic curve E can be identified with the number of $\mathbb{F}_q$ points of certain *Hurwitz modules*, to be defined precisely later in §26, specifically in Definition 26.0.1. The relevant examples of these Hurwitz modules parameterize certain $\mathrm{AGL}(\mathbb{Z}/\ell\mathbb{Z})$ covers of $\mathbb{P}^1$, but have specified local conditions at the points where E has bad reduction.

3. RELATING POINT COUNTING TO HOMOLOGICAL STABILITY

3.1. The Grothendieck-Lefschetz trace formula.

One of the main tools that builds the bridge between the land of topology and land of number theory is the Grothendieck-Lefschetz trace formula. This relates the number of $\mathbb{F}_q$ points to the trace of Frob_q^{-1} on its cohomology.

Definition 3.1.1. Here, Frob_q denotes the *geometric Frobenius*, defined as follows: There is an automorphism $\sigma : \overline{\mathbb{F}}_q \to \overline{\mathbb{F}}_q$ sending $x \mapsto x^q$. This induces a map $1 \times \sigma : X \times_{\mathrm{Spec}\,\mathbb{F}_q} \mathrm{Spec}\,\overline{\mathbb{F}}_q \to X \times_{\mathrm{Spec}\,\mathbb{F}_q} \mathrm{Spec}\,\overline{\mathbb{F}}_q$. Then, geometric Frobenius is the inverse of this automorphism. Note that this is an automorphism over $\mathbb{F}_q$, but it is not an automorphism over $\overline{\mathbb{F}}_q$ because it does not fix $\overline{\mathbb{F}}_q$.

See [Poo17, §7.5.6] for a more detail regarding the definition of geometric Frobenius.

Theorem 3.1.2. *Let X be a finite type smooth Deligne-Mumford stack over $\mathbb{F}_q$, all of whose components have dimension n. We have*

$$\frac{\#X(\mathbb{F}_q)}{q^n} = \sum_{i=0}^{2\dim X} (-1)^i \operatorname{tr}\left(\mathrm{Frob}_q^{-1} | H^i(X_{\overline{\mathbb{F}}_q}, \mathbb{Q}_{\ell'})\right). \tag{3.1}$$

Let's see a few examples to get the hang of how one can use this.

Remark 3.1.3. A useful fact we will use repeatedly is that, for X a smooth finite type scheme of DM stack over $\mathbb{F}_q$, the action of Frob_q on $H^0(X_{\overline{\mathbb{F}}_q}, \mathbb{Q}_\ell)$ is a permutation matrix, obtained by recording how Frobenius permutes the connected components of $X_{\overline{\mathbb{F}}_q}$. (Implicit here is that, $\dim H^0(X_{\overline{\mathbb{F}}_q}, \mathbb{Q}_\ell)$ is the number of connected components of $X_{\overline{\mathbb{F}}_q}$.) In particular, if X is geometrically irreducible, Frob_q acts by 1.

Example 3.1.4. Take $X = \mathbb{P}^1$. Topologically, $\mathbb{P}^1$ is homeomorphic to a 2 sphere, which has cohomology in degrees 0 and 2. It turns out this has the same cohomology of $\overline{\mathbb{F}}_q$. The trace of Frob_q on H^0 is 1 and on H^2 is q. From this it follows that

$$\frac{\#\mathbb{P}^1(\mathbb{F}_q)}{q} = \mathrm{tr}\left(\mathrm{Frob}_q^{-1} \,| H^0(\mathbb{P}^1_{\overline{\mathbb{F}}_q}, \mathbb{Q}_\ell)\right) + \mathrm{tr}\left(\mathrm{Frob}_q^{-1} \,| H^1(\mathbb{P}^1_{\overline{\mathbb{F}}_q}, \mathbb{Q}_\ell)\right) = 1 + q^{-1}.$$

Multiplying by q gives $\#\mathbb{P}^1(\mathbb{F}_q) = 1 + q$. Of course this is one of the hardest ways to prove this. A much easier way is that $\mathbb{P}^1 = \mathbb{A}_1 \cup \infty$, and $\mathbb{A}^1$ has q $\mathbb{F}_q$ points while the point ∞ has a single $\mathbb{F}_q$ point.

Let's see another example.

Example 3.1.5. Take $X = \mathbb{P}^n$. For the next exercise, you may wish to familiarize yourself with the Serre spectral sequence for a fibration (which also has a version in étale cohomology you may assume) and the Gysin exact sequence in étale cohomology (including the twists for the Frobenius action).

Exercise 3.1.6. Compute the cohomology of $\mathbb{P}^n$ in two ways.

1. Using that $\mathbb{P}^n \simeq \mathbb{A}^{n+1}/\mathbb{G}_m$, compute via a spectral sequence that $H^i(\mathbb{P}^n, \mathbb{Q}_\ell)$ is 1-dimensional if $0 \leq i \leq 2n$ is even and 0 otherwise.
2. Using that a hyperplane in $\mathbb{P}^n$ is isomorphic to $\mathbb{P}^{n-1}$ with open complement $\mathbb{A}^n$. Show by induction (using the Gysin exact sequence) that the cohomology of $\mathbb{P}^n$ is as claimed in the previous part. (This is asking you to prove the result of the previous part via a second method.)
3. Using that the Gysin exact sequence coming from the second part respects Frobenius eigenvalues, show that the action of Frob_q on $H^{2i}(\mathbb{P}^n_{\overline{\mathbb{F}}_q}, \mathbb{Q}_\ell)$ is by multiplication by q^i.

Using the above, we can now compute the number of points on $\mathbb{P}^n$ in an extremely difficult fashion, using the Grothendieck-Lefschetz trace formula. We find

$$\frac{\#\mathbb{P}^n(\mathbb{F}_q)}{q^n} = \sum_{i=0}^{n} \mathrm{tr}\left(\mathrm{Frob}_q^{-1} \,| H^{2i}(\mathbb{P}^n_{\overline{\mathbb{F}}_q}, \mathbb{Q}_\ell)\right) = 1 + q^{-1} + q^{-2} + \cdots + q^{-n}.$$

Multiplying by q^n gives $\#\mathbb{P}^n(\mathbb{F}_q) = q^n + q^{n-1} + \cdots + 1$.

Remark 3.1.7. Of course, one can compute $\mathbb{P}^n$ much more easily. For example, one can write $\mathbb{P}^n = \mathbb{A}^n \coprod \mathbb{A}^{n-1} \coprod \cdots \coprod \mathbb{A}^0$ and noting that $\mathbb{A}^n(\mathbb{F}_q) = q^n$.

Exercise 3.1.8. The Grothendieck-Lefschetz trace formula is named after the Lefschetz trace formula, which says that if X is a smooth complex manifold of dimension n and $f : X \to X$ is an automorphism then the number of fixed points of f (counted with multiplicity) is $\sum_{i=1}^{2n}(-1)^i \operatorname{tr}(f|H_i(X, \mathbb{Q}))$.

Now, let $f : S^2 \to S^2$ denote the automorphism given by rotating a sphere 180 degrees about its vertical axis. Clearly this has two fixed points which are the north and south poles. Verify this a different way (and moreover conclude these fixed points have multiplicity 1) using Lefschetz trace formula.

We next define configuration space. A slightly more general definition will be given in Definition 4.1.2.

Definition 3.1.9. Let $U \subset \mathbb{A}^n$ denote the open subset corresponding to n-tuples of points with distinct coordinates and define $\mathrm{Conf}_n := U/S_n$, where the S_n action permutes the coordinates.

Remark 3.1.10. One can also describe Conf_n as the scheme whose S-points, for S a ring, are given by polynomials $f = x^n + a_{n-1}x^{n-1} + \cdots + a_0$ with $a_i \in S$, whose vanishing locus is étale over S. In other words, the $\mathbb{F}_q$ points are monic polynomials over $\mathbb{F}_q$ with no repeated factors.

Let's now use some topology to compute the number of $\mathbb{F}_q$ points of configuration space.

Example 3.1.11. Suppose $n > 1$. It is a fact from topology that the rational cohomology of Conf_n over $\mathbb{C}$ is 1-dimensional in degrees 0 and 1, and 0-dimensional otherwise. (The rational cohomology is the same as that of a circle, but it is integrally the same as $\Omega^2 S^2$; we will come back to this later in Theorem 11.3.10.) Assuming this fact, let us compute $\#\mathrm{Conf}_n(\mathbb{F}_q)$. First, we want to determine the traces of Frobenius. As usual since Conf_n is geometrically irreducible $\operatorname{tr}(\mathrm{Frob}_q^{-1} |H^0(\mathrm{Conf}_{n,\overline{\mathbb{F}}_q}) = 1$. It remains to compute $\operatorname{tr}(\mathrm{Frob}_q^{-1} |H^1(\mathrm{Conf}_{n,\overline{\mathbb{F}}_q}))$. For this, we can use the Gysin exact sequence, by noting that $\mathrm{Conf}_n \subset \mathbb{A}^n$ is an open, with complement a divisor D_n. The

Gysin exact sequence gives a short exact sequence

(3.2)

$$H^{i-2}(D_{n,\overline{\mathbb{F}}_q}, \mathbb{Q}_\ell(-1)) \longrightarrow H^i(\mathbb{A}^n_{\overline{\mathbb{F}}_q}, \mathbb{Q}_\ell) \longrightarrow H^i(\mathrm{Conf}_{n,\overline{\mathbb{F}}_q}, \mathbb{Q}_\ell) \longrightarrow$$

$$H^{(i+1)-2}(D_{n,\overline{\mathbb{F}}_q}, \mathbb{Q}_\ell(-1)) \longrightarrow H^{i+1}(\mathbb{A}^n_{\overline{\mathbb{F}}_q}, \mathbb{Q}_\ell),$$

where the notation $H^i(X, \mathbb{Q}_\ell(i))$ means that Frob_q acts by q^{-i} times its action on $H^i(X, \mathbb{Q}_\ell)$. Since the cohomology of $\mathbb{A}^n$ vanishes in positive degrees, we find $H^1(\mathrm{Conf}_{n,\overline{\mathbb{F}}_q}, \mathbb{Q}_\ell)$ maps isomorphically to $H^{2-2}(D_{n,\overline{\mathbb{F}}_q}, \mathbb{Q}_\ell(-1)) = H^0(D_{n,\overline{\mathbb{F}}_q}, \mathbb{Q}_\ell(-1))$. Therefore, the trace of Frob_q on this is q and hence the trace of Frob_q^{-1} is q^{-1}.

Plugging this in to the Grothendieck-Lefschetz trace formula gives

$$\frac{\#\mathrm{Conf}_n(\mathbb{F}_q)}{q^n} = \mathrm{tr}\left(\mathrm{Frob}_q^{-1} | H^0(\mathrm{Conf}_{n,\overline{\mathbb{F}}_q}, \mathbb{Q}_\ell)\right) - \mathrm{tr}\left(\mathrm{Frob}_q^{-1} | H^1(\mathrm{Conf}_{n,\overline{\mathbb{F}}_q}, \mathbb{Q}_\ell)\right) = 1 - q^{-1}.$$

and so $\#\mathrm{Conf}_n(\mathbb{F}_q) = q^n - q^{n-1}$.

3.2. **The large q limit first.** We'd now like to see how various forms of homological stability can relate to counting points. Another important tool we will need is the following result of Deligne for varieties, which has since also been generalized to Deligne Mumford stacks.

Theorem 3.2.1 (Deligne). *Let X be a smooth finite type Deligne-Mumford stack over $\mathbb{F}_q$. The eigenvalues of Frobenius on $H^i(X_{\overline{\mathbb{F}}_q}, \mathbb{Q}_\ell)$ have absolute value between $q^{i/2}$ and q^i, under any embedding $\mathbb{Q}_\ell \to \mathbb{C}$.*

Remark 3.2.2. If X is moreover proper, then the absolute value turns out to be exactly $q^{i/2}$. However, when X is smooth, these q^i which naturally live in H^{2i} can "leak" into lower cohomology groups, down to H^i. A simple example is Example 3.1.11, where we say that the eigenvalue of Frob_q on H^1 of the smooth-but-not-proper $\mathrm{Conf}_{n,\overline{\mathbb{F}}_q}$ was q, not $q^{1/2}$.

Exercise 3.2.3. Given an elliptic curve E over a finite field $\mathbb{F}_q$, there is a well known bound that $\#E(\mathbb{F}_q) - 1 - q \leq 2\sqrt{q}$.

Verify more generally that if X is a smooth proper curve over $\mathbb{F}_q$ which is geometrically irreducible of genus g, then $|\#X(\mathbb{F}_q) - 1 - q| \leq 2g\sqrt{q}$.

First, a fun consequence is the classical Lang-Weil theorem. There is an easier proof without the Grothendieck-Lefschetz trace formula and Deligne's bounds on the eigenvalues of Frobenius, but let's see an simple proof assuming it.

Notation 3.2.4. Fix a base B which is a localization of a finite over $\operatorname{Spec} \mathbb{Z}$. So, either B is a localization of the spectrum of an order of a ring of integers in a number field or the spectrum of a finite field. When we take limits over q, these will always be over those q so that $\operatorname{Spec} \mathbb{F}_q$ maps to B.

Lemma 3.2.5 (Lang-Weil). *Let $X \to B$ be a smooth map of schemes with geometrically irreducible fibers of pure dimension n. Then* $\lim_{q\to\infty} \frac{\#X(\mathbb{F}_q)}{q^n} = 1$.

Remark 3.2.6. The above statement also holds when the map $X \to B$ is not smooth. One can deduce this more general case by showing that the singular locus has $o(q^n)$ many points, via a similar argument to what we do below.

Proof. By the Grothendieck-Lefschetz trace formula, we have

$$\begin{aligned}\frac{\#X(\mathbb{F}_q)}{q^n} &= \sum_{i=0}^{2n} (-1)^i \operatorname{tr}(\operatorname{Frob}_q^{-1} | H^i(X_{\overline{\mathbb{F}}_q}, \mathbb{Q}_\ell)) \\ &= \operatorname{tr}(\operatorname{Frob}_q^{-1} | H^0(X_{\overline{\mathbb{F}}_q}, \mathbb{Q}_\ell)) + \sum_{i=1}^{2n} (-1)^i \operatorname{tr}(\operatorname{Frob}_q^{-1} | H^i(X_{\overline{\mathbb{F}}_q}, \mathbb{Q}_\ell)).\end{aligned}$$

Therefore, it suffices to show that for each $i > 0$, $\lim_{q\to\infty}(-1)^i \operatorname{tr}(\operatorname{Frob}_q^{-1} | H^i(X_{\overline{\mathbb{F}}_q}, \mathbb{Q}_\ell)) = 0$. Using Deligne's bounds on the eigenvalues of Frobenius, all eigenvalues of Frob_q have absolute values between $q^{i/2}$ and q^i. Therefore, all eigenvalues of $\operatorname{Frob}_q^{-1}$ are between $q^{-i/2}$ and q^{-i}. In particular,

$$\| \operatorname{tr}(\operatorname{Frob}_q^{-1} | H^i(X_{\overline{\mathbb{F}}_q}, \mathbb{Q}_\ell)) \| \leq q^{-i/2} \cdot h^i(X_{\overline{\mathbb{F}}_q}, \mathbb{Q}_\ell),$$

where we use h^i to denote $\dim H^i$. From this it follows that, as $q \to \infty$, the above expression tends to 0, as desired. □

Let's see a simple consequence of Lang Weil.

Corollary 3.2.7. *Suppose X_n is a sequence of smooth finite type varieties over B so that the fibers of X_n have dimension n. Assume there is a constant C for which the fibers have X_n has exactly κ irreducible components for $n > C$, all of which are geometrically irreducible. Then, for any $n > C$,*

$$\lim_{q\to\infty} \frac{\#X_n(\mathbb{F}_q)}{q^n} = \kappa.$$

In particular,

$$\lim_{n\to\infty} \lim_{q\to\infty} \frac{\#X_n(\mathbb{F}_q)}{q^n} = \kappa.$$

Proof. This follows immediately from Lemma 3.2.5 applied separately to each irreducible component of X_n. $\square$

Remark 3.2.8. Note that if X_n is a smooth variety over $\mathbb{F}_{q_0}$ with κ irreducible components that are all geometrically irreducible component, then $h^0(X_n, \mathbb{Q}_\ell) = \kappa$. Therefore, if we know that the 0th cohomology (or homologies) of the X_n stabilize (in a Frobenius equivariant fashion, corresponding to the geometrically irreducible condition) to κ we find that the number of $\mathbb{F}_q$ points of the X_n also stabilize in the $q \to \infty$ limit.

So, the above shows we can interpret the 0th homology of X_n (each of whose irreducible components are geometrically irreducible) as computing the number of $\mathbb{F}_q$ points on X_n as $q \to \infty$.

Exercise 3.2.9. Let $s \in \mathbb{F}_q^\times$ be an element which is not a square. Consider the subscheme $V := V(x^2 - sy^2) \subset \mathbb{P}^2_{x,y,z}$.

(1) Verify V as above is irreducible and $\#V(\mathbb{F}_q) = 1$. Why does this not contradict the generalization of Lemma 3.2.5 to non-smooth maps described in Remark 3.2.6?

(2) There is a version of the Grothendieck-Lefschetz trace formula that applies to singular varieties. It says

$$\#X(\mathbb{F}_q) = \sum_{i=1}^{2\dim X} (-1)^i \operatorname{tr}\left(\mathrm{Frob}_q \,| H^i_c(X, \mathbb{Q}_\ell)\right),$$

where the H^i_c denotes compactly supported cohomology. In the case $X = V(x^2 - sy^2)$ above, explain why the right hand evaluates to 1 (as it must since you showed $V(\mathbb{F}_q) = 1$).

Remark 3.2.10. In particular, the above exercise shows that there can be 1-dimensional varieties over $\mathbb{F}_q$ so that their number of $\mathbb{F}_q$ points is 1 for arbitrarily large q. (E.g. one can take the variety V above and replace q with q^n for n odd so that s will not be a square. So it is possible for the conclusion of Lang-Weil to not be satisfied if the hypotheses are not satisfied.

3.3. **Switching the limits: large height, then large q.**

Next, let us understand what the arithmetic consequence of knowing the higher homologies of X_n stabilize. Because counting problems over number fields don't have any reference to large q limits, the various conjectures we would like to investigate are phrased with a large n limit, but with no large q limit. With that said, it is often fruitful to investigate these conjectures first in the setting of taking a large q limit.

Notation 3.3.1. Let B be a localization of a finite scheme over $\operatorname{Spec}\mathbb{Z}$. Suppose we have a sequence of schemes X_n over a base scheme B with the following properties:

(1) Each fiber of $X_n \to B$ has κ geometrically irreducible components of pure dimension n.
(2) There are numbers h^i_∞ depending on i and constants I, J (but independent of n) so that for $b \to B$ any geometric point, $h^i(X_{n,b}, \mathbb{Q}_\ell) = h^i_\infty$ for $n > Ii + J$
(3) For every i, n there is some constants C, C' so that $h^i(X_{n,b}, \mathbb{Q}_\ell) < C' \cdot C^i$ for every geometric point $b \to B$.

Assuming the above, we can interchange the limits and conclude properties about $X_n(\mathbb{F}_q)$, where one first takes a limit in q, and only second a limit in n.

Lemma 3.3.2. *Suppose $X_n \to B$ is a sequence of varieties as in Notation 3.3.1. Then*

$$\lim_{q\to\infty} \limsup_{n\to\infty} \frac{\#X_n(\mathbb{F}_q)}{q^n} = \lim_{q\to\infty} \liminf_{n\to\infty} \frac{\#X_n(\mathbb{F}_q)}{q^n} = \kappa.$$

Note part of the above statement is that the limits above exist.

Proof. The basic idea is to use the Grothendieck-Lefschetz trace formula and Deligne's bounds. Breaking up the contribution from the 0th homology and the remaining homologies, and arguing similarly to the proof of Corollary 3.2.7, it is enough to show

$$\lim_{q\to\infty} \limsup_{n\to\infty} \sum_{i=1}^{2n} (-1)^i \operatorname{tr}\left(\mathrm{Frob}_q^{-1} | H^i(X_{n,\overline{\mathbb{F}}_q}, \mathbb{Q}_\ell)\right) = 0,$$

and a similar statement for the lim inf. Using the conditions from Notation 3.3.1, we can bound the above in absolute value by

$$\lim_{q\to\infty} \limsup_{n\to\infty} \sum_{i=1}^{\frac{n-J}{I}-1} \| \operatorname{tr}\left(\mathrm{Frob}_q^{-1} | H^i(X_{n,\overline{\mathbb{F}}_q}, \mathbb{Q}_\ell)\|\right) + \sum_{i=\frac{n-J}{I}}^{2n} \| \operatorname{tr}\left(\mathrm{Frob}_q^{-1} | H^i(X_{n,\overline{\mathbb{F}}_q}, \mathbb{Q}_\ell)\right) \|$$

$$\leq \lim_{q\to\infty} \limsup_{n\to\infty} \sum_{i=1}^{\frac{n-J}{I}-1} h^i_\infty \cdot q^{-i/2} + \sum_{i=\frac{n-J}{I}}^{2n} q^{-i/2} C' C^i.$$

Now, the meaning of the double limit is that we can first take q sufficiently large before taking the limit in the second variable. The second term $\sum_{i=\frac{n-J}{I}}^{2n} q^{-i/2} C' C^i$ is a geometric series and evaluates to $C' \left(\frac{C}{q^{1/2}}\right)^{\frac{n-J}{I}} \frac{1}{1-\frac{C}{q^{1/2}}}$.

Hence, we can see this tends to 0 as $n \to \infty$ since $\frac{C}{q^{1/2}} < 1$. For the first term, we can see it tends to $\sum_{i=1}^{\infty} h_{\infty}^{i} \cdot q^{-i/2}$ as $n \to \infty$. We note this converges using the assumption (3) (once $C < q^{1/2}$) since the sum is bounded by a geometric series.

Then, when one further takes the limit as $q \to \infty$, this geometric series tends to 0, as desired. □

Exercise 3.3.3. Give an example of a sequence of geometrically irreducible varieties X_n over $\mathbb{F}_q$ with $\dim X_n = n$ satisfying the hypotheses of Notation 3.3.1 but so that $\lim_{n\to\infty} \frac{\#X_n(\mathbb{F}_q)}{q^n}$ does not exist.

3.4. **A reprise on the large height, then large q limit.**

Remark 3.4.1. Note in the above proof of Lemma 3.3.2, we weren't actually able to show the limit in n existed. The reason we couldn't show this is that we don't know the trace of Frobenius on the cohomologies are consistent. For example, we could imagine a sequence where X_n only has H^0 and H^1, but the trace of Frobenius alternates between $\pm q^{-1}$ on the H^1. Then the number of points would alternate between $1 + q^{-1}$ and $1 - q^{-1}$ and so no limit as $n \to \infty$ would exist, even though these both tend to 1 as $q \to \infty$. Of course, the way to fix this is to ask that the traces of Frobenius on homology also stabilize.

Notation 3.4.2. Suppose we have a sequence X_n as in Notation 3.3.1 and additionally assume that $\operatorname{tr}\left(\mathrm{Frob}_q^{-1} | H^i(X_{n,\overline{\mathbb{F}}_q}, \mathbb{Q}_\ell)\right)$ is independent of n for $n > Ii + J$. Define t_i to be this stable value $\operatorname{tr}\left(\mathrm{Frob}_q^{-1} | H^i(X_{n,\overline{\mathbb{F}}_q}, \mathbb{Q}_\ell)\right)$ when $n > Ii + J$.

Lemma 3.4.3. *For X_n as in Notation 3.4.2, if q is sufficiently large,*

$$\lim_{n\to\infty} \frac{\#X_n(\mathbb{F}_q)}{q^n} = \sum_{i=1}^{\infty} (-1)^i t_i \tag{3.3}$$

Proof. Similarly to the proof of Lemma 3.3.2, for $q^{1/2} > C$, we can use Notation 3.3.1(3) to show that $\sum_{i=0}^{\infty} (-1)^i t_i$ converges, so the above limit makes sense.

Computing similarly to our computation in the proof of Lemma 3.3.2, we find

$$\begin{aligned}
&\lim_{n\to\infty} \frac{\#X_n(\mathbb{F}_q)}{q^n} - \sum_{i=0}^{\infty}(-1)^i t_i \\
&= \lim_{n\to\infty} \sum_{i=1}^{\frac{n-J}{I}-1} (-1)^i \left(\operatorname{tr}\left(\mathrm{Frob}_q^{-1} | H^i(X_{n,\overline{\mathbb{F}}_q}, \mathbb{Q}_\ell)\right) - t_i\right) \\
&\quad + \sum_{i=\frac{n-J}{I}}^{2n} (-1)^i \left(\operatorname{tr}\left(\mathrm{Frob}_q^{-1} | H^i(X_{n,\overline{\mathbb{F}}_q}, \mathbb{Q}_\ell)\right) - t_i\right) \\
&= \lim_{n\to\infty} \sum_{i=\frac{n-J}{I}}^{2n} (-1)^i \left(\operatorname{tr}\left(\mathrm{Frob}_q^{-1} | H^i(X_{n,\overline{\mathbb{F}}_q}, \mathbb{Q}_\ell)\right) - t_i\right).
\end{aligned}$$

Finally, this sum tends to 0 as $n \to \infty$ using Deligne's bounds on the eigenvalues of Frobenius and the assumption that $q^{1/2} > C$ as well as Notation 3.3.1(3) to guarantee that the above is bounded by a decaying geometric series with first term tending to 0 as $n \to \infty$. ☐

Lemma 3.4.4. *For X_n as in Notation 3.4.2,*

$$\lim_{q\to\infty} \lim_{n\to\infty} \frac{\#X_n(\mathbb{F}_q)}{q^n} = \kappa \tag{3.4}$$

The main difference between Lemma 3.4.4 and Lemma 3.3.2 is that the inner limit in Lemma 3.4.4 exists whereas before only the $\limsup_n$ and $\liminf_n$ existed.

Proof. This follows from Lemma 3.3.2 once we show the limit in n exists. Indeed, we showed the limit in n exists in Lemma 3.4.3. ☐

3.5. **The benefit of computing the stable cohomology.** In the prior sections, we explained the benefit of showing that the cohomology stabilizes. However, we couldn't compute the large n limit on its own. We could do so after then computing a large q limit. However, if one can also compute the stable value of the cohomology, one is often able to compute the large n limit without an additional large q limit.

Here is the key lemma:

Lemma 3.5.1. *Suppose we have a sequence of geometrically irreducible varieties X_n as in Notation 3.4.2 and suppose Y_n is another sequence of geometrically irreducible varieties also satisfying the conditions of Notation 3.4.2. Assume also there are* *maps* $X_n \to Y_n$ *inducing isomorphisms on* $H^i(Y_{n,b}, \mathbb{Q}_\ell) \to H^i(X_{n,b}, \mathbb{Q}_\ell)$ *for each geometric point* $b \to B$ *and* $n > Ii + J$*. Then, for* q *sufficiently large,*

$$\lim_{n\to\infty} \frac{\#X_n(\mathbb{F}_q)}{q^n} = \lim_{n\to\infty} \frac{\#Y_n(\mathbb{F}_q)}{q^n}.$$

Equivalently, for q *sufficiently large,* $\lim_{n\to\infty} \frac{\#X_n(\mathbb{F}_q)}{\#Y_n(\mathbb{F}_q)} = 1$.

Proof. The final statement is obtained from the penultimate one by dividing both sides of the penultimate one by its right hand side. We will now prove the penultimate statement. Finally both limits exist by Lemma 3.4.3, which also computes both limits, and shows they are equal since they only depend on the stable homology of the X_n and Y_n, using the following exercise:

Exercise 3.5.2. Suppose $X \to Y$ is a map of finite type Deligne Mumford stacks over $\mathbb{F}_q$ so that the induced map $H^i(Y_{\overline{\mathbb{F}}_q}, \mathbb{Q}_\ell) \to H^i(X_{\overline{\mathbb{F}}_q}, \mathbb{Q}_\ell)$ is an isomorphism. Then

$$\operatorname{tr}\left(\operatorname{Frob}_q^{-1} | H^i(X_{\overline{\mathbb{F}}_q}, \mathbb{Q}_\ell)\right) = \operatorname{tr}\left(\operatorname{Frob}_q^{-1} | H^i(Y_{\overline{\mathbb{F}}_q}, \mathbb{Q}_\ell)\right).$$

□

Example 3.5.3. How will Lemma 3.5.1 help us? It will often come in handy in the case that we understand the stable cohomology of the Y_n. For example, if $Y_n = \operatorname{Conf}_n$, over $B = \operatorname{Spec} \mathbb{F}_{q_0}$ we know from Example 3.1.11 that $\operatorname{Conf}_n(\mathbb{F}_q) = q^n - q^{n-1}$. Therefore, computing the stable cohomology in this way will allow us to compute the large n limit of $\#X_n(\mathbb{F}_q)$. Another way it can help us is that many of the central conjectures in arithmetic statistics regard average sizes of certain objects, see Conjecture 1.1.1 and Conjecture 1.3.1. If we take the X_n to parameterize the objects we want to count and Y_n to parameterize the set of objects we want to average over, then we can often use Lemma 3.5.1 to compute the desired average.

For example, in Conjecture 1.1.1, we would take the X_n to parameterize pairs of a quadratic extension of discriminant q^n and an n-torsion element in the class group of such an extension, while the Y_n would parameterize quadratic extensions of discriminant q^n. Similarly, in Conjecture 1.3.1, the X_n would parameterize pairs of an elliptic curve of height n and an ℓ-Selmer element for such an elliptic curve, while the Y_n would parameterize elliptic curves of height n.

4. HURWITZ SPACES

4.1. **The definition of Hurwitz spaces.**

We begin by carefully defining the Hurwitz spaces we will work with. We start by defining the Hurwitz stack

$[\mathrm{Hur}^{G,c}_{n,B}/G]$. This will be a quotient by an action of G of the pointed Hurwitz space $\mathrm{Hur}^{G,c}_{n,B}$, which we will define next, and this explains why there is a quotient by G in the notation for $[\mathrm{Hur}^{G,c}_{n,B}/G]$.

Definition 4.1.1. Let G be a finite group, $c \subset G$ be a union of conjugacy classes. Let B be a scheme with $\#G$ invertible on B. Assume that c is closed under invertible powering ($g \in c \implies g^t \in c$ for any t prime to $\#G$). Let $[\mathrm{Hur}^{G,c}_{n,B}/G]$ denote the *Hurwitz stack* whose T points $[\mathrm{Hur}^{G,c}_{n,B}/G](T)$ is the groupoid

$$\left(D, i : D \to \mathbb{P}^1_T, X, h : X \to \mathbb{P}^1_T\right)$$

satisfying the following conditions:

(1) D is a finite étale cover of T of degree n.
(2) i is a closed immersion $i : D \subset \mathbb{A}^1_T \subset \mathbb{P}^1_T$.
(3) X is a smooth proper relative curve over T, not necessarily having geometrically connected fibers.
(4) $h : X \to \mathbb{P}^1_T$ is a finite locally free Galois G-cover, (meaning that G acts simply transitively on the geometric generic fiber of h,) which is étale away from $\infty_T \cup i(D) \subset \mathbb{P}^1_T$.
(5) The inertia of $X \to \mathbb{P}^1_T$ over any geometric point of $i(D)$ lies in $c \subset G$.
(6) The morphisms between two points (D_i, i_i, X_i, h_i) for $i \in \{1, 2\}$ are given by (ϕ_D, ψ_X) where $\phi_D : D_1 \simeq D_2$ is an isomorphism so that $i_2 \circ \phi_D = i_1$ and $\psi_X : X_1 \simeq X_2$ is an isomorphism such $h_2 \circ \psi_X = h_1$ and $\psi_X = g^{-1}\psi_X g$ for every $g \in G$.

One can show that the above defined category fibered in groupoids is in fact an algebraic stack. This essentially follows from [ACV03, §1.3.2 and Appendix B], but see [LL24, Remark 2.1.2] for details.

We next give a formal definition of configuration space.

Definition 4.1.2. Given a scheme B, there is an open subscheme $U \subset \mathbb{A}^n_B$ parameterizing the locus where all coordinates are distinct. There is an action of the symmetric group S_n on U by permuting the coordinates and we define $\mathrm{Conf}_{n,B} := U/S_n$ to be the *configuration space* of n points in $\mathbb{A}^1$ over B. More generally, let $\mathrm{Conf}_{n_1,\ldots,n_v,B} := U/S_{n_1} \times \cdots \times S_{n_v}$ where S_{n_i} acts on the n consecutive coordinates in the range $[n_1 + \cdots + n_{i-1} + 1, n_1 + \cdots + n_i]$. When $B = \operatorname{Spec} R$, for R a ring, we often write this as $\mathrm{Conf}_{n_1,\ldots,n_v,R}$ and when $R = \mathbb{C}$, we abbreviate this to $\mathrm{Conf}_{n_1,\ldots,n_v}$.

Exercise 4.1.3. Prove that we have an isomorphism of schemes $\mathrm{Conf}_{n,B} \simeq \mathrm{Hur}^{\mathrm{id},\mathrm{id}}_{n,B}$.

Definition 4.1.4 (Temporary definition)**.** Given a group G, and a union of conjugacy classes $c \subset G$, we now define a space $\mathrm{Hur}^{G,c}_{n,\mathbb{C}}$ which has a G action so that the quotient by this G action is $\mathrm{Hur}^{G,c}_{n,\mathbb{C}}$. We can define this as the finite étale cover of configuration space associated to the action $\pi_1(\mathrm{Conf}_{n,B}) \simeq B_n \to \mathrm{Aut}(c^n)$ where B_n denotes the braid group on n generators and the action the ith generator of B_n on c^n is given by $(x_1, \ldots, x_i, x_{i+1}, \ldots, x_n) \mapsto (x_1, \ldots, x_{i+1}, x_{i+1}^{-1} x_i x_{i+1}, \ldots, x_n)$. Note that this action preserves the product of the elements of c^n.

Exercise 4.1.5.

(1) Prove that the universal cover of configuration space (over the complex numbers) can be described as follows: It parameterizes points in the interior of a rectangle together with arcs joining those points to the bottom edge of the rectangle (which is a contractible subspace we can use in place of a basepoint). The arcs are chosen not to intersect each other. (Another equivalent description is where one considers points in the interior of a disc with a marked point on the boundary and arcs, only intersecting at the boundary point, joining that boundary point to the points in the interior.)
(2) Prove that this universal cover is contractible.
(3) Explain how the braid group acts on this space of points and arcs so that the quotient of the universal cover by this action is identified with configuration space of n points in the interior of the rectangle.

Exercise 4.1.6. Explain why we can think of a point of $\mathrm{Hur}^{G,c}_{n,\mathbb{C}}$ as a point of $\mathrm{Conf}_{n,\mathbb{C}}$ together with a labeling of each of the n branch points by elements of c. Use this to define an action of G on $\mathrm{Hur}^{G,c}_{n,\mathbb{C}}$ by preserving the branch points and letting $g \in G$ send $(x_1, \ldots, x_n) \mapsto (g x_1 g^{-1}, \ldots, g x_n g^{-1})$. Show that the quotient stack by this G action is identified with $[\mathrm{Hur}^{G,c}_{n,\mathbb{C}} / G]$ as defined in Definition 4.1.1.

Exercise 4.1.7. Combining the previous two exercises, and using the correspondence between finite unramified covering spaces, finite sets, and quotients of copies of the universal cover by the fundamental group, give a description of Hurwitz space $\mathrm{Hur}^{G,c}_{n,\mathbb{C}}$ as a quotient of the following space by a B_n action: The space is given by a collection of n points in the interior of the rectangle with points labeled by elements of c and arcs joining each of them to the bottom edge of the rectangle which do not intersect. A generator of the braid group explicitly acts on two nearby points labeled (x, y) by sending the labels to $(y, y^{-1}xy)$ and also moving the ends of the arcs in a half twist around each other.

Remark 4.1.8. Malle's conjecture is about counting G covers, but Cohen-Lenstra is about counting torsion elements in class groups of quadratic fields. Class field theory provides a bijection between such torsion elements and G covers together with a suitable marking over ∞. We now explain what goes into this marking. This will also give a more direct comparison to a topological version of Hurwitz spaces.

One important idea is to work with pointed Hurwitz spaces, which deal with the case of "real quadratic" or "ramified at ∞" quadratic fields. Variants of these were also used in [EL24], but prior work in the context of the Cohen–Lenstra heuristics appears to primarily focus on the case over covers unramified at ∞. These Hurwitz spaces parameterize covers of a stacky $\mathbb{P}^1$, with a root stack at ∞ of order 2. In the case that our base B is $\operatorname{Spec} \mathbb{C}$, this G cover was already defined in Definition 4.1.4.

Remark 4.1.9. We next define these pointed Hurwitz spaces. The heuristic idea will be that we want to mark the sheets over ∞. Topologically, we can define the Hurwitz stack $[\mathrm{Hur}^{G,c}_{n,\mathbb{C}}/G]$ as parameterizing covers $X \to D$, with D a disc, and then the pointed Hurwitz scheme $\mathrm{Hur}^{G,c}_{n,\mathbb{C}}$ will parameterize such covers together with a trivialization over a boundary point of the disc. In order to make precise how to specify this trivialization over the boundary point of a disc in algebro-geometric language, we next discuss root stacks.

Lemma 4.1.10. *Maps $X \to [\mathbb{A}^1/\mathbb{G}_m]$ are in bijection with pairs $(\mathscr{L}, s)$ for $\mathscr{L}$ a line bundle on X and $s : \mathscr{O}_X \to \mathscr{L}$ a section. In particular, any Cartier divisor D induces a map given by $\mathscr{O}_X(D)$ together with the section $\mathscr{O}_X \to \mathscr{O}_X(D)$.*

Proof. A map to $[\mathbb{A}^1/\mathbb{G}_m]$ is equivalent to a $\mathbb{G}_m$ torsor $V \to X$ and a $\mathbb{G}_m$ equivariant map $V \to \mathbb{A}^1$. The data of V is equivalent to the data of a line bundle $\mathscr{L}$ and the data of the map is equivalent to the data of a section of $\mathscr{L}$. For further details, (expressed in a slightly different way,) see [Ols16, Proposition 10.3.7]. □

Definition 4.1.11. Let X be a scheme over B and $D \subset X$ a relative Cartier divisor. For w and integer, the wth *root stack* of X along D is by definition the fiber product $X^{(D,w)}$

$$\begin{tikzcd}
X^{(D,w)} \arrow[r] \arrow[d] & {[\mathbb{A}^1/\mathbb{G}_m]} \arrow[d, "\phi_w"] \\
X \arrow[r, "\phi_D"] & {[\mathbb{A}^1/\mathbb{G}_m]}
\end{tikzcd} \tag{4.1}$$

where the map ϕ_D is induced by the divisor D via Lemma 4.1.10 and the map ϕ_w is induced via the quotient by the $\mathbb{G}_m$ action of the divisor $w \cdot 0 \subset \mathbb{A}^1$.

Remark 4.1.12. One should think of the root stack as adding stackiness of order w at along D and leaving the rest of X unchanged.

Definition 4.1.13. Fix an integer w and define $\mathscr{P}^w$ to be the root stack of order w along ∞ of $\mathbb{P}^1$. It is shown in [Ols16, Theorem 10.3.3(ii)] that the fiber of this root stack over ∞ is the stack quotient $[(\operatorname{Spec}_B \mathscr{O}_B[x]/(x^r)) / \mu_r]$ of the relative spectrum $\operatorname{Spec}_B \mathscr{O}_B[x]/(x^r)$ by μ_r. Let $\widetilde{\infty} : B \to \mathscr{P}^w$ denote the section over σ corresponding to map $B \to [(\operatorname{Spec}_B \mathscr{O}_B[x]/(x^r)) / \mu_r]$ given by the trivial μ_r torsor over B, $\mu_r \to B$, and the μ_r equivariant map $\mu_r \to B \to \operatorname{Spec}_B \mathscr{O}_B[x]/(x^r)$. We use notation as in Definition 4.1.1,

Define the *w-pointed Hurwitz space*, $\left(\mathrm{Hur}^{G,c}_{n,B}\right)^w$, to be the algebraic space whose T points are the set parameterizing data of the form

$$\left(D, h' : X \to \mathscr{P}^w_T, t : T \to X \times_{h', \mathscr{P}^w_T, \widetilde{\infty}_T} T, i : D \to \mathbb{P}^1_T, X, h : X \to \mathbb{P}^1_T\right),$$

where $D, i, X,$ and h satisfy the properties listed in Definition 4.1.1. We also assume the order of inertia of h along ∞ is w and define $\widetilde{\infty}_T$ to be the base change of the section $\widetilde{\infty}$ defined above to T. We additionally impose the condition that h' is a finite locally free G-cover, étale over $\widetilde{\infty}_T$, such that the composition of $h' : X \to \mathscr{P}^w_T$ with the coarse space map $\mathscr{P}^w_T \to \mathbb{P}^1_T$ is h, and $t : T \to X \times_{h', \mathscr{P}^w_T, \widetilde{\infty}_T} T$ is a section of h' over $\widetilde{\infty}_T$.

In general, we define the *pointed Hurwitz space* as $\mathrm{Hur}^{G,c}_{n,B} := \coprod_{w \geq 1} \left(\mathrm{Hur}^{G,c}_{n,B}\right)^w$.

Remark 4.1.14. We note that $\mathrm{Hur}^{G,c}_{n,B}$, defined as an algebraic space, is in fact a scheme. Indeed, $\mathrm{Hur}^{G,c}_{n,B}$ is a finite étale cover of the configuration space of n unordered points in $\mathbb{A}^1_B$, as can be verified in an analogous fashion to [LWZB24, Proposition 11.4].

Exercise 4.1.15. Show that moreover $\mathrm{Hur}^{G,c}_{n,B} \to B$ is smooth and affine of relative dimension n. You may assume the assertions in the above remark.

Notation 4.1.16. We use $\mathrm{CHur}^{G,c}_{n,B}$ to denote the open and closed subscheme of $\mathrm{Hur}^{G,c}_{n,B}$ parameterizing covers $X \to \mathbb{P}^1_T$ so that the geometric fibers of X over T are connected. We use $[\mathrm{CHur}^{G,c}_{n,B} / G]$ to denote the open and closed substack of $[\mathrm{Hur}^{G,c}_{n,B} / G]$ parameterizing covers $X \to \mathbb{P}^1_T$ so that the geometric fibers of X over T are connected.

Notation 4.1.17. We will often work over the complex numbers, and hence it will be convenient to assume $B = \operatorname{Spec} \mathbb{C}$. Therefore, we define $[\mathrm{Hur}^c_n / G] := [\mathrm{Hur}^{G,c}_{n,\operatorname{Spec}\mathbb{C}} / G], \mathrm{Hur}^c_n := \mathrm{Hur}^{G,c}_{n,\operatorname{Spec}\mathbb{C}}, [\mathrm{CHur}^c_n / G] := [\mathrm{CHur}^{G,c}_{n,\operatorname{Spec}\mathbb{C}} / G], \mathrm{CHur}^c_n :=$

$\mathrm{CHur}^{G,c}_{n,\operatorname{Spec}\mathbb{C}}$. If we fix $n = n_1 + \cdots + n_v$, we define $\mathrm{Hur}^c_{n_1,\ldots,n_v}$ to be the union of components of Hur^c_n with n_i points in the ith conjugacy class of c. Similarly, we define $\mathrm{CHur}^c_{n_1,\ldots,n_v} := \mathrm{Hur}^c_{n_1,\ldots,n_v} \times_{\mathrm{Hur}^c_n} \mathrm{CHur}^c_n$.

Exercise 4.1.18. Suppose G is an abelian group and $c \subset G$ is a subset consisting of v elements. From the definition above, there is a map $\mathrm{Hur}^{G,c}_{n_1,\ldots,n_v,\mathbb{C}} \to \mathrm{Conf}_{n_1,\ldots,n_v}$ sending a cover to its branch locus, where one remembers the number of elements of each conjugacy class. Prove this map is an isomorphism.

5. COHEN-LENSTRA IN THE LARGE q LIMIT

Later in these notes, we will see versions of Cohen-Lenstra without taking a large q limit. However, first, we'll discuss the easier case of understanding the large q limit. As mentioned before, and referring to Lemma 3.2.5, understanding this in the large q limit essentially amounts to counting geometrically irreducible components of certain spaces. On the other hand, proving this in the large height limit (without a large q limit) amounts to understanding something additional about their higher cohomology. A paper where much of the material of this section is discussed (albeit in a somewhat different presentation) is [Ach06]. I have also read that it appeared in some unpublished notes by J.K. Yu, though I have never seen them.

Definition 5.0.1. Let $\mathscr{H}^1_g$ denote the pointed Hurwitz scheme which we were denoting $\mathrm{Hur}^{\mathbb{Z}/2\mathbb{Z},\{1\}}_{2g+1,\mathbb{Z}[1/2]}$ in Definition 4.1.13. Informally, this parameterizes moduli stack of double covers $C \to \mathbb{P}^1$ with a ramified point over ∞ with C smooth proper and with geometrically connected fibers, together with a marked point over the residual $\mathbb{Z}/2\mathbb{Z}$ gerbe of a cover over ∞.

Remark 5.0.2. The above corresponds to the imaginary quadratic case of Cohen-Lenstra. We'll avoid discussing the real case here, but later we'll see a (morally stronger) version of the real case in the large height limit in Theorem 7.0.10.

Remark 5.0.3. There is a map $\mathscr{H}^1_g \to \mathrm{Conf}_{2g+1}$ sending a double cover to its branch locus. In fact, this is an isomorphism, essentially because the cover is determined by its branch locus. This also follows from Exercise 4.1.18.

We will start by defining a sheaf on $\mathscr{H}^1_g$ whose $\mathbb{F}_q$ points count torsion elements of class groups.

Definition 5.0.4. Fix an odd prime ℓ. Let $\phi_g : \mathscr{C}^\circ_g \to \mathscr{H}^1_g$ denote the complement of the section over ∞ of the universal curve. So, $\mathscr{C}^\circ_g \to \mathscr{H}^1_g$ is an affine map. Let $\overline{\phi}_g : \mathscr{C}_g \to \mathscr{H}_g^1$ denote the universal curve. Define $X_g^\ell := R^1(\overline{\phi}_g)_* \mu_\ell$, which is a sheaf on $\mathscr{H}_g^1$, which we can identify with a scheme as it is a finite étale cover of $\mathscr{H}_g^1$.

Exercise 5.0.5. Show that the fibers of X_g^ℓ over $\overline{x} \in \mathscr{H}_g^1(\overline{\mathbb{F}}_q)$, corresponding to a hyperelliptic curve $C_{\overline{x}}$, is identified with $H^1(C_{\overline{x}}, \mu_\ell)$. If $\overline{x} \in \mathscr{H}_g^1(\overline{\mathbb{F}}_q)$ is a geometric point over a point $x \in \mathscr{H}_g^1(\mathbb{F}_q)$, let $C_x^\circ = C_x - \infty$, for $\infty \in C_x$ the unique point mapping to $\infty \in \mathbb{P}^1$. Show additionally that $H^1(C_x, \mu_\ell) \simeq H^1(C_x^\circ, \mu_\ell)$. *Hint:* Use proper base change for the first part and the Gysin exact sequence.

For $K/\mathbb{F}_q(t)$ a quadratic extension, let $\mathscr{O}_K$ denote the normalization of $\mathbb{F}_q[t]$ in the function field of K. If this quadratic extension corresponds to the double cover of curves $x \in \mathscr{H}_g^1(\mathbb{F}_q)$ as above, then $\operatorname{Spec} \mathscr{O}_K \simeq C_x^\circ$.

Lemma 5.0.6. *Suppose* $\gcd(q(q-1), \ell) = 1$ *then* $\mathrm{Cl}(K)[\ell] \simeq H^1(C_x, \mu_\ell) \simeq H^1(C_x^\circ, \mu_\ell)$.

Proof. There is an exact sequence

$$0 \longrightarrow \mu_\ell \longrightarrow \mathbb{G}_m \longrightarrow \mathbb{G}_m \longrightarrow 0. \tag{5.1}$$

Taking cohomology yields

$$H^0(C_x, \mathbb{G}_m) \longrightarrow H^0(C_x, \mathbb{G}_m) \longrightarrow H^1(C_x, \mu_\ell) \longrightarrow H^1(C_x, \mathbb{G}_m) \longrightarrow H^1(C_x, \mathbb{G}_m). \tag{5.2}$$

Now, the first map can be identified with the raising to the ℓth power map $\mathbb{F}_q^\times \to \mathbb{F}_q^\times$. Since this is a cyclic group of order $q-1$, which we are assuming is prime to ℓ it is an isomorphism. Hence, the above sequence identifies $H^1(C_x, \mu_\ell) \simeq H^1(C_x, \mathbb{G}_m)[\ell] \simeq \mathrm{Pic}(C_x)[\ell]$. The ℓ torsion in the class group is identified with ℓ torsion in the Picard group because the class group is the Picard group of the once punctured curve $\operatorname{Spec} \mathscr{O}_K = C_x - \infty \simeq C_x^\circ$, for ∞ the unique point over $\infty \in \mathbb{P}^1$. And there is an exact sequence

$$0 \longrightarrow \mathbb{Z} \cdot \infty \longrightarrow \mathrm{Pic}(C_x) \longrightarrow \mathrm{Pic}(\operatorname{Spec} \mathscr{O}_K) \longrightarrow 0. \tag{5.3}$$

Since the groups above are abelian, we find $\mathrm{Pic}(C_x) \simeq \operatorname{Pic} \operatorname{Spec}(\mathscr{O}_K) \times \mathbb{Z}$, so they have the same ℓ torsion. Since these have the same ℓ torsion, this also identifies $H^1(C_x^\circ, \mu_\ell) \simeq H^1(C_x, \mu_\ell)$. □

Lemma 5.0.7. *Let x be a* $\operatorname{Spec} \mathbb{F}_q$ *point of $\mathscr{H}_g^1$ with geometric point $\overline{x}$ over x. Then, for $\pi_g : X_g^\ell \to \mathscr{H}_g^1$ the projection map,* $\pi_g^{-1}(x)(\mathbb{F}_q) = H^1(C_x, \mu_\ell) \simeq H^1(C_x^\circ, \mu_\ell)$, *assuming still* $\gcd(\ell, q(q-1)) = 1$*. Moreover, the fiber* $\pi_g^{-1}(x)(\mathbb{F}_q)$ *can be identified with the invariants of geometric frobenius* Frob_x *acting on* $H^1(C^\circ_{\overline{x}}, \mu_\ell)$ *for* $\overline{x}$ *a geometric point over* x*.*

Proof. Using Exercise 5.0.5, the geometric fiber of X^ℓ_g over a geometric point $\overline{x}$ can be identified with $H^1(C_{\overline{x}}, \mu_\ell)$. There is an action of Frobenius on this vector space, and the $\mathbb{F}_q$ points of the sheaf are the Frobenius invariants. On the other hand, we claim these Frobenius invariants are identified with $H^1(C_x, \mu_\ell) = H^1(C^\circ_x, \mu_\ell)$. To see this, apply the Leray spectral sequence

$$H^p(\mathrm{Gal}(\overline{x}/x), H^q(C_{\overline{x}}, \mu_\ell)) \to H^{p+q}(C_x, \mu_\ell).$$

We get an exact sequence from the low degree terms

(5.4)

$$H^1(\mathrm{Gal}(\overline{x}/x), H^0(C_{\overline{x}}, \mu_\ell)) \longrightarrow H^1(C_x, \mu_\ell) \longrightarrow H^0(\mathrm{Gal}(\overline{x}/x), H^1(C_{\overline{x}}, \mu_\ell)) \longrightarrow$$

$$H^2(\mathrm{Gal}(\overline{x}/x), H^0(C_{\overline{x}}, \mu_\ell)).$$

So, to make the desired identification, we just need to check $H^i(\mathrm{Gal}(\overline{x}/x), H^0(C_{\overline{x}}, \mu_\ell)) = 0$ for $i = 1, 2$. We can identify these groups with $H^i(\mathrm{Gal}(\overline{x}/x), \mu_\ell)$ since $C_{\overline{x}}$ is geometrically connected. The H^2 vanishes because $\mathrm{Gal}(\overline{x}/x)$ is $\widehat{\mathbb{Z}}$, whose homologies with torsion coefficients always vanish in degree 2.

Exercise 5.0.8. Verify $H^1(\mathrm{Gal}(\overline{x}/x), \mu_\ell) = H^1(x, \mu_\ell) = 0$ using the assumption that q and $q-1$ are relatively prime to ℓ.

This completes the proof. □

Definition 5.0.9. We say a quadratic extension K over $\mathbb{F}_q(t)$ has genus g if K is the function field of a genus g curve.

Corollary 5.0.10. *Fix an odd prime* ℓ *and an odd genus* g*. Suppose* q *is a prime power with* $\gcd(q(q-1), \ell) = 1$*. Then,* $\#X^\ell_g(\mathbb{F}_q) = \sum_{K/\mathbb{F}_q(t) \text{quadratic extension of genus } g} \#\mathrm{Cl}(\mathscr{O}_K)[\ell]$*.*

Proof. For each $x \in \mathscr{H}^1_g(\mathbb{F}_q)$, it suffices to identify $\pi_g^{-1}(x)(\mathbb{F}_q)$ with the class group of $\mathrm{Cl}(C^\circ_x)[\ell]$, as then we can sum over all such x. This identification follows by combining Lemma 5.0.7 and Lemma 5.0.6. □

Remark 5.0.11. In the case $\ell = 1$, counting $\#X^1_g(\mathbb{F}_q)$ amounts to counting $\#\mathscr{H}^1_g(\mathbb{F}_q)$. Hence, the Cohen-Lenstra heuristics, which are about counting the limit of

$$\frac{\sum_{K/\mathbb{F}_q(t) \text{quadratic extension of genus g}} \#\mathrm{Cl}(\mathscr{O}_K)[\ell]}{\sum_{K/\mathbb{F}_q(t) \text{quadratic extension of genus g}} 1}$$

as $g \to \infty$ can be rephrased as counting

$$\frac{\# X_g^\ell(\mathbb{F}_q)}{\# \mathscr{H}_g^1(\mathbb{F}_q)}$$

as $g \to \infty$. Instead, we will now discuss evaluating this limit as $q \to \infty$. By Lemma 3.2.5, we equivalently wish to count the number of components of the cover $X_g^\ell \to \mathscr{H}_1^g$.

We next wish to determine the number of components of X_g^ℓ over $\mathbb{F}_q$. The key to doing this is to study the monodromy representation. Fix a geometric point $\overline{x} \to \mathscr{H}_g^1$. Let $V := X_g^\ell|_{\overline{x}}$ denote the fiber. By Exercise 5.0.5, we can identify V with $H^1(C_{\overline{x}}, \mu_\ell)$. We have a monodromy representation

$$\rho_g^\ell : \pi_1(\mathscr{H}_g^1) \to \operatorname{Aut} V.$$

By definition, this comes from the equivalent between representations of the fundamental group of the base, up to conjugacy, and local systems. Intuitively, one can think about this as follows: if one goes around a loop on the base how does this permute the fiber of V when one lifts the path to the cover. We also have a geometric monodromy representation

$$\rho_{g,\overline{\mathbb{F}}_q}^\ell : \pi_1\left((\mathscr{H}_g^1)_{\overline{\mathbb{F}}_q}\right) \to \operatorname{Aut} V.$$

Definition 5.0.12. Fix a finite dimensional vector space V over a field k with a symplectic pairing. This means a bilinear map $\omega : V \times V \to k$ such that $\omega(x, x) = 0$ and B is nondegenerate, meaning that $\omega(x, y) = 0$ for all y implies that $x = 0$. The symplectic group $\operatorname{Sp}(V)$, with implicit dependence on ω, is the group of automorphisms preserving ω. That is, it is the set of T so that $\omega(Tx, Ty) = \omega(x, y)$. The general symplectic group $\operatorname{GSp}(V)$ is the group of automorphisms of V preserving ω up to scaling. That is, there is some constant $\operatorname{mult}(T) \in k^\times$ so that $\operatorname{mult}(T)\omega(x, y) = \omega(Tx, Ty)$. There is a *multiplier map* $\operatorname{GSp}(V) \to k$ sending $T \mapsto \operatorname{mult}(T)$. We also obtain an exact sequence

$$0 \longrightarrow \operatorname{Sp}(V) \longrightarrow \operatorname{GSp}(V) \longrightarrow k^\times. \longrightarrow 0 \tag{5.5}$$

We also use $\operatorname{GSp}^q(V)$ to denote the coset of $\operatorname{Sp}(V)$ in $\operatorname{GSp}(V)$ with multiplier q.

Exercise 5.0.13. Show that any vector space with a symplectic pairing must have even dimension. Show there is a unique symplectic pairing, up to isomorphism, on a vector space of dimension $2g$. *Hint:* Construct a symplectic basis $e_1, \ldots, e_g, f_1, \ldots, f_g$ with $\omega(e_i, e_j) = \omega(f_i, f_j) = 0$ and $\omega(e_i, f_j) = \delta_{ij}$.

Exercise 5.0.14. Show the gth power of the multiplier map of T agrees with the determinant of T.

Remark 5.0.15. In light of the preceding exercise, we use Sp_{2g} to denote the symplectic group $\mathrm{Sp}(V)$ for V a vector space of dimension $2g$ with the unique symplectic pairing. Similarly, we use GSp_{2g} to denote $\mathrm{GSp}_{2g}(V)$ and GSp_{2g}^q to denote the multiplier q subset.

Proposition 5.0.16. *Suppose* $\gcd(q(q-1), \ell) = 1$. *The image of* ρ_g^ℓ *contains the symplectic group* $\mathrm{Sp}_{2g}(\mathbb{Z}/\ell\mathbb{Z})$. *Moreover, the geometric monodromy is* $\mathrm{Sp}_{2g}(\mathbb{Z}/\ell\mathbb{Z})$. *Additionally, for any* $\mathbb{F}_q$ *point* x, *the composition* $\pi_1(x) \to \pi_1(\mathscr{H}_g^1) \to \mathrm{Aut}(V)$ *sends Frobenius to an element of* $\mathrm{GSp}_{2g}(\mathbb{Z}/\ell\mathbb{Z})$ *with multiplier* q.

Sketch. First, for any point x, the Weil pairing is a symplectic pairing on $H^1(C_{\overline{x}}, \mu_\ell)$ which the Frobenius action scales by q. Hence, each Frobenius element acts by an element of $\mathrm{GSp}_{2g}^q(\mathbb{Z}/\ell\mathbb{Z})$. Since Frobenius elements are dense in the fundamental group $\pi_1(\mathscr{H}_g^1)$, this shows the image is contained in $\mathrm{GSp}_{2g}(\mathbb{Z}/\ell\mathbb{Z})$ and also shows the last statement. (We could have also shown the containment by arguing that the Weil pairing exists for the universal curve and argued that it is preserved up to scaling.)

Since we are assuming $q-1$ is prime to ℓ, raising to the qth power is generates $(\mathbb{Z}/\ell\mathbb{Z})^\times$ and so it is enough to show that the image contains $\mathrm{Sp}_{2g}(\mathbb{Z}/\ell\mathbb{Z})$. One can check the multiplier agrees with the cyclotomic character, so, to conclude the proof, we just need to show that the geometric monodromy $\rho_{g,\overline{\mathbb{F}}_q}^\ell(\pi_1((\mathscr{H}_g^1)_{\overline{\mathbb{F}}_q}))$ has image all of $\mathrm{Sp}_{2g}(\mathbb{Z}/\ell\mathbb{Z})$.

There are two approaches to showing this. Let us describe them briefly. The first is to prove it over the complex numbers by explicitly writing down generators of the fundamental group (corresponding to moving two branch points in a half twist around each other) and showing they fill out a basis of transvections which generate the symplectic group. (This is actually not very difficult, and probably easier than it sounds. A good reference for a similar computation is [FM12] and a related statement for hyperelliptic curves was originally shown in [A'C79].) Then one has to compare the monodromy in characteristic 0 to characteristic p, which one can do using the fact that $\mathscr{H}_g^1$ has a normal crossings compactification over $\mathbb{Z}[1/2]$.

A second approach works directly in positive characteristic. A good reference for this is [Hal08, Theorem 5.1]; see also [EL24, Theorem 7.1.5]. It turns out that we can realize the sheaf X_g^ℓ as the middle convolution of the sheaf μ_ℓ. We will avoid getting into the theory of middle convolution here, but let us just explain how this implies the monodromy is the symplectic group.

One can use the theory of middle convolution to show that X_g^ℓ corresponds to an irreducible local system (meaning the image of ρ_g^ℓ acts irreducibly on V) and that that the image is generated by *transvections* which by definition are matrices with a codimension one 1 eigenspace and whose generalized eigenspace is full dimensional. It is then a general theorem in group theory that a subgroup of the symplectic group generated by transvections, acting irreducibly on a vector space over $\mathbb{F}_\ell$, for ℓ an odd prime, is the full symplectic group. □

Corollary 5.0.17. *The Cohen-Lenstra heuristics hold for the $\mathbb{Z}/\ell\mathbb{Z}$ moment in the large q limit. That is, for fixed $g \geq 2$,*

$$\lim_{q\to\infty} \frac{\# X_g^\ell(\mathbb{F}_q)}{\# \mathscr{H}_g^1(\mathbb{F}_q)} = 2,$$

where the limit is taken over prime powers with $\gcd(q(q-1), \ell) = 1$*. Moreover, X_g^ℓ has two components, both of which are geometrically irreducible.*

Proof. Using Lemma 3.2.5, the ratio we are looking for will agree with the number of geometrically irreducible component of X_g^ℓ, since $\mathscr{H}_g^1$ is geometrically irreducible. Next, X_g^ℓ clearly has one component given by the trivial section, corresponding to the trivial element of the class group. Therefore, it's enough to show there's only one other component, which is geometrically irreducible. To check this, we can pass to $\overline{\mathbb{F}}_q$. There, we saw in Proposition 5.0.16 that the image is $\mathrm{Sp}_{2g}(\mathbb{Z}/\ell\mathbb{Z})$. Since this acts on the geometric fiber which is $(\mathbb{Z}/\ell\mathbb{Z})^{2g}$, we just want to see this has two orbits. One orbit is the 0 vector.

Exercise 5.0.18. Show that any two nonzero vectors $v, w \in (\mathbb{Z}/\ell\mathbb{Z})^{2g}$ are related by the action of a matrix in the symplectic group. *Hint:* If you can show there is a unique symplectic vector space of dimension $2g$ you will be able to show this.

This means we can use the monodromy to get from any nonzero point of the fiber over our basepoint to any other, meaning that any two such points are in the same component, and hence X_g^ℓ has only two components. □

This finishes our aimed for description of the Cohen-Lenstra moments in the large q limit. However, recall that the Cohen-Lenstra heuristics actually predict a probability distribution. Let's see how we can also see this fairly easily from the work we've done so far, in the large q limit. It turns out we have the following formulation of Cohen-Lenstra in terms of random matrices. (The original distribution was stated in a different way, but was translated to a statement for random matrices in GL by Friedman and Washington, and then later translated to one about symplectic groups as below by Achter.)

Conjecture 5.0.19 (Cohen-Lenstra)**.** For K a random imaginary quadratic extension of $\mathbb{F}_q$, the probability $\mathrm{Cl}(K)[\ell] \simeq (\mathbb{Z}/\ell\mathbb{Z})^j$ is the probability

$$\lim_{g\to\infty} \mathrm{Prob}_{h\in \mathrm{GSp}^q_{2g}(\mathbb{Z}/\ell\mathbb{Z})}((\ker h - \mathrm{id}) \simeq (\mathbb{Z}/\ell\mathbb{Z})^j).$$

In other words, the probability the ℓ torsion has rank j is equal to the probability that the 1-eigenspace of a random element of the general symplectic group of multiplier ℓ has dimension j.

Remark 5.0.20. There are generalizations of this conjecture to the full class group, but let's just focus on the simpler case of the ℓ torsion for now.

It seems clear there should be some relation to the symplectic group monodromy, but what is it exactly? The relation is given by a theorem, often attributed to Deligne or Katz on equidistribution of Frobenius elements.

Definition 5.0.21. Let X be a geometrically connected finite type scheme over $\mathbb{F}_q$, let G be a profinite group, and let $\lambda : \pi_1(X) \to G$ be a group homomorphism. Let G_0 denote the image of the composition $\pi_1^{\mathrm{geom}}(X) := \pi_1(X_{\overline{\mathbb{F}}_q}) \to \pi_1(X) \to G$ and let $\Gamma := G/G_0$. Then, we define $\mathrm{mult} : G \to \Gamma$ as the natural projection. Because $\pi_1(\operatorname{Spec}\mathbb{F}_q) = \pi_1(X)/\pi_1^{\mathrm{geom}}(X)$, we obtain a resulting map $\pi_1(\operatorname{Spec}\mathbb{F}_q) \to \Gamma$. We let γ_q denote the image in Γ of geometric Frobenius.

Proposition 5.0.22. *Let $\mathcal{X}$ be a smooth affine scheme of finite type over $\mathcal{O}[1/S]$, where $\mathcal{O}$ is a ring of integers in a number field, with geometrically irreducible fibers. For $\mathfrak{q}$ a maximal ideal of $\mathcal{O}[1/S]$ with residue field $\mathbb{F}_q$, write $X := \mathcal{X}|_{\mathcal{O}/\mathfrak{q}}$. Assume that we have a commutative diagram*

$$\begin{tikzcd}
1 \arrow[r] & \pi_1^{\mathrm{geom}}(X) \arrow[r] \arrow[d, "\lambda_0"] & \pi_1(X) \arrow[r, "\deg"] \arrow[d, "\lambda"] & \widehat{\mathbb{Z}} \arrow[r] \arrow[d, "1\mapsto\gamma_q^{-1}"] & 1 \\
1 \arrow[r] & G_0 \arrow[r] & G \arrow[r, "\mathrm{mult}"] & \Gamma \arrow[r] & 1
\end{tikzcd} \tag{5.6}$$

with λ_0 tamely ramified and surjective, G a finite group, and Γ abelian. Suppose $C \subset G$ is a conjugacy-invariant subset. Then

$$\mathrm{Prob}\{x \in X(\mathbb{F}_{q^n}) : \lambda(\mathrm{Frob}_x) \in C\} = \frac{\#C \cap G^{\mathrm{mult}\,\gamma_q^n}}{\#G_0} + O_{\mathcal{X}}\left(\#G\sqrt{\frac{\#C \cap G^{\mathrm{mult}\,\gamma_q^n}}{q^n}}\right).$$

where $G^{\mathrm{mult}\,\gamma_q^n} := \mathrm{mult}^{-1}(\gamma_q^n)$.

We won't discuss the proof much, but it's basically a consequence of a twisted version of the Grothendieck Lefschetz trace formula. It can also be thought of as a version of Chebotarev density, which says Frobenius becomes equidistributed in the monodromy group. The error term can be viewed as the contribution to the point counts from cohomology groups in degree bigger than 0 while the main term has to do with the fact that the coefficient of the trivial character in the indicator function of a conjugacy class of size $\#c$ in G is $\#c/\#G$.

Lemma 5.0.23. *As $q \to \infty$ with fixed $g \geq 2$,*

$$\lim_{q\to\infty} \mathrm{Prob}_{K/\mathbb{F}_q(t) \text{quadratic of genus } g}(\mathrm{Cl}(\mathscr{O}_K)[\ell] \simeq (\mathbb{Z}/\ell\mathbb{Z})^j)$$
$$= \mathrm{Prob}_{h\in \mathrm{GSp}^q_{2g}(\mathbb{Z}/\ell\mathbb{Z})}((\ker h - \mathrm{id}) \simeq (\mathbb{Z}/\ell\mathbb{Z})^j).$$

Proof. Using Lemma 5.0.7 and Lemma 5.0.6, we can identify the class group $\mathrm{Cl}(\mathscr{O}_K)[\ell]$ associated to a point x with the invariants of geometric Frobenius acting on $H^1(C_{\overline{x}}, \mu_\ell)$. The invariants are the same as the 1-eigenspace, so it's enough to know the distribution of the Frobenius elements are uniformly random in GSp^q_{2g}.

In Proposition 5.0.16 we computed the monodromy of the cover to be GSp_{2g} with geometric monodromy Sp_{2g}. It finally follows from Lemma 5.0.23 that Frobenius is equidistributed in GSp^q_{2g}, as we wished to show. □

6. A GLIMPSE INTO THE MAIN HURWITZ SPACE HOMOLOGICAL STABILITY RESULTS

We'll now give a brief glimpse into the main homological stability results that we will be focusing on soon in this course. Following this, we'll see applications to the Cohen-Lenstra heuristics, Malle's conjecture, and the Picard rank conjecture.

Recall that for G a finite group and $c \subset G$ a union of v conjugacy classes, we are using $\mathrm{Hur}^c_{n_1,\ldots,n_v}$ to denote to denote the Hurwitz space parameterizing pointed covers of $\mathbb{P}^1$ with inertia contained in c and n_i points of inertia in c_i in $\mathbb{A}^1$.

The first main result is that the homology of Hurwitz spaces stabilize.

Theorem 6.0.1. *There are constants I and J, depending on c, so that for any $i \geq 0$ and $n_1 > Ii + J$, there is an isomorphism $H_i(\mathrm{CHur}^c_{n_1,\ldots,n_v}, \mathbb{Z}) \to H_i(\mathrm{CHur}^c_{n_1+1,\ldots,n_v}, \mathbb{Z})$ is an isomorphism.*

The second result is that we can compute their stable value when all the n_i are large.

Theorem 6.0.2. *Then are constants I and J, depending on c, so that for any $i \geq 0$ and $n_1, \ldots, n_v > Ii + J$, and any component $Z \subset \mathrm{CHur}^c_{n_1,\ldots,n_v}$, the map $H_i(Z, \mathbb{Z}[1/\#G]) \to H_i(\mathrm{Conf}_{n_1,\ldots,n_v}, \mathbb{Z}[1/\#G])$ is an isomorphism.*

Remark 6.0.3. Much later in the course, in Theorem 18.0.9, we come back to computing the stable value when just one of the n_i is large, but in general it can have different stable homology than that of configuration space.

7. APPLICATION TO COHEN-LENSTRA

Recall that the Cohen-Lenstra conjectures predict the average size of ℓ torsion in class groups of quadratic fields, as mentioned in Conjecture 1.1.1. We saw that Conjecture 1.1.1 predicted

$$\lim_{X\to\infty} \mathbb{E}_{K \in IQ_{\leq X}}(\#\mathrm{Cl}(\mathscr{O}_K)[\ell]) = 2.$$

The function field analog would be

$$\lim_{X\to\infty} \mathbb{E}_{K \in IQ_{\leq X}}(\#\mathrm{Cl}(\mathscr{O}_K)[\ell]) = 2.$$

The Cohen-Lenstra heuristics not only predict the average size of ℓ-torsion in the class group, but more generally predict the full distribution of ℓ-torsion in the class group. In particular, if $H = \mathbb{Z}/\ell\mathbb{Z}$, the above can be phrased as saying that the average number of maps $H \to \mathrm{Cl}(\mathscr{O}_K)$ is 2. Since there is a unique trivial map, this is equivalent to saying that the average number of such injections is 1. Since the number of injections is the same as the number of surjections, we would equivalently conjecture that the average number of surjections from $\mathrm{Cl}(\mathscr{O}_K)$ onto H is 1 when $H = \mathbb{Z}/\ell\mathbb{Z}$. More generally, in the number field case, Cohen and Lenstra conjecture:

Conjecture 7.0.1. For any odd order abelian group H,

$$\lim_{X\to\infty} \mathbb{E}_{K \in IQ_{\leq X}}(\#\mathrm{Surj}(\mathrm{Cl}(\mathscr{O}_K), H)) = 1.$$

Exercise 7.0.2. Assuming Conjecture 7.0.1, for ν an odd integer, calculate

$$\lim_{X\to\infty} \mathbb{E}_{K \in IQ_{\leq X}}(\#\mathrm{Cl}(\mathscr{O}_K)[\nu]))$$

and calculate

$$\lim_{X\to\infty} \mathbb{E}_{K \in IQ_{\leq X}}((\#\mathrm{Cl}(\mathscr{O}_K)[\ell])^2),$$

for ℓ an odd prime.

Remark 7.0.3. Certainly, the distribution of class groups determines the moments. That is, if we know the probability $\mathrm{Cl}(\mathscr{O}_K)[\nu]$ takes on any abelian group, we can compute the number of surjections to any $\mathbb{Z}/\nu\mathbb{Z}$ module H. However, it turns out that knowing the moments in fact determines the distribution as well, in this case. Hence, the full Cohen-Lenstra conjecture regarding the distribution of class groups in imaginary quadratic fields turns out to be equivalent to Conjecture 7.0.1.

Remark 7.0.4. One can similarly generalize this to function fields. This has also be generalized to real quadratic extensions in place of imaginary quadratic extensions. Additionally, it has been generalized to other global fields and to a non-abelian setting.

In the function field case, we can replace $\mathbb{Q}$ with $\mathbb{F}_q(t)$, and make very similar conjectures. In this case, the analog of quadratic fields are hyperelliptic curves, and the ℓ-torsion in their class group corresponds to ℓ-torsion line bundles on the hyperelliptic curve.

We'll now set some notation to state the Cohen-Lenstra conjectures over function fields. To make a precise statement, here is a precise definition of the set of hyperelliptic curves we will consider for Cohen-Lenstra.

Definition 7.0.5. Let $\mathcal{MH}_{n,q}$ denote the set of function fields K of monic hyperelliptic curves $y^2 = f(t)$ for $f \in \mathbb{F}_q[t]$ monic squarefree of degree n. Let $\mathscr{O}_K$ denote $\mathbb{F}_q[t,y]/(y^2 - f(t))$, (the normalization of $\mathbb{F}_q[t]$ in the quadratic extension $K/\mathbb{F}_q(t)$,) and let $\mathrm{Cl}(\mathscr{O}_K)$ denote the set of line bundles on Spec $\mathscr{O}_K$ (or ideal classes in $\mathscr{O}_K$).

Over function fields, by analogy, one would then conjecture

$$\lim_{n\to\infty, n\equiv 1 \bmod 2} \mathbb{E}_{K\in\mathcal{MH}_{n,q}}(\#\mathrm{Surj}(\mathrm{Cl}(\mathscr{O}_K), H)) = 2.$$

A version of this was shown in a certain large q limit by Ellenberg-Venkatesh-Westerland.

Remark 7.0.6. I highly encourage you to take a look at their paper - it was my favorite paper in grad school, and it contained 8 sections, with each section encompassing a different field of math!

Here is their main result.

Theorem 7.0.7 (Ellenberg-Venkatesh-Westerland)**.** *For H a finite abelian group of odd order*

$$\lim_{q\to\infty} \limsup_{\substack{n\to\infty \\ n\equiv 1 \bmod 2}} \mathbb{E}_{K\in\mathcal{MH}_{n,q}}(\#\mathrm{Surj}(\mathrm{Cl}(\mathscr{O}_K), H)) = \lim_{q\to\infty} \liminf_{\substack{n\to\infty \\ n\equiv 1 \bmod 2}} \mathbb{E}_{K\in\mathcal{MH}_{n,q}}(\#\mathrm{Surj}(\mathrm{Cl}(\mathscr{O}_K), H)) = 1,$$

where one takes $q \to \infty$ over odd prime powers so that $q(q-1)$ is relatively prime to the order of H.

Remark 7.0.8. Strictly speaking they showed a slightly weaker version where the group H was an ℓ group, for ℓ an odd prime.

The main input Ellenberg-Venkatesh-Westerland used to prove Theorem 7.0.7 was a homological stability result for Hurwitz spaces. Namely, they proved a version of Theorem 6.0.1 (although they weren't able to prove it for arbitrary groups and conjugacy classes, they could prove it for those relevant to the Cohen-Lenstra heuristics.) Then, they could take this homological stability result as input, and use a version of Lemma 3.3.2 to deduce the above theorem.

Remark 7.0.9. The rough idea Ellenberg-Venkatesh-Westerland used to prove this homological stability result is that they first showed the 0th cohomology stabilizes by counting components, and then showed by induction that if the ith cohomology stabilizes, the $i+1$st cohomology stabilizes as well. This proof combined some standard ideas in homological stability (such as understanding braid group actions on very connected arc complexes) with a number of novel ideas (such as a certain regularity theorem for the homology of Hurwitz spaces). We will see a simplified proof of their result in Theorem 17.0.6.

We were able to prove a stronger version of Cohen Lenstra over fixed finite fields q, so long as q is sufficiently large relative to H.

Theorem 7.0.10 (L-Levy). *For any odd order abelian group H, and q an odd prime power with $\gcd(\#H, q(q-1)) = 1$, there is an integer C depending on H so that for $q > C$,*

$$\lim_{n\to\infty, n\equiv 1 \bmod 2} \mathbb{E}_{K\in\mathcal{MH}_{n,q}}(\#\operatorname{Surj}(\operatorname{Cl}(\mathscr{O}_K), H)) = 1.$$

Moreover,

$$\lim_{n\to\infty, n\equiv 0 \bmod 2} \mathbb{E}_{K\in\mathcal{MH}_{n,q}}(\#\operatorname{Surj}(\operatorname{Cl}(\mathscr{O}_K), H)) = \frac{1}{\#H}.$$

Remark 7.0.11. The assumption that $q-1$ is prime to H is needed to assume there are no extraneous roots of unity dividing the order of H in $\mathbb{F}_q(t)$. This is necessary because the predicted average size changes there are such roots of unity. We prove that the average size is $\wedge^2 H[h]$ if $h := \gcd(\#H, q-1)$ in general.

Remark 7.0.12. More generally still, one might want to formulate a non-abelian Cohen-Lenstra conjecture, which replaces the class group (the maximal *abelian* unramified extension) with the maximal unramified extension. We also prove a suitable version of these conjectures over function fields.

Remark 7.0.13. Theorem 7.0.10 doesn't quite prove the Cohen-Lenstra heuristics over function fields. What is shows is that if you give me a group H, we can compute the H moment over all but finitely many q. However, the

Cohen-Lenstra heuristics predict the value of the H moment for all q. Viewed another way, if you give me a q, they predict all the H moments, but the above theorem can only compute finitely many H moments for a fixed q.

Remark 7.0.14. The result above can be made effective, but the constant is not very good. If one takes $H = \mathbb{Z}/3\mathbb{Z}$, the constant C I we get is around 10^{60}.

As outlined in Lemma 3.5.1, we should be able to deduce this point counting result from our computation of the stable homology in Theorem 6.0.1 compute the stable value of the cohomology.

Example 7.0.15. If one takes $G = \mathbb{Z}/3\mathbb{Z}$, the Hurwitz space parameterizes $\mathbb{Z}/3\mathbb{Z} \rtimes \mathbb{Z}/2\mathbb{Z} = S_3$ covers which are simply branched (the monodromy is a transposition). Then, it turns out there is a unique component Z_n of the Hurwitz space $\mathrm{Hur}^{G,c}_{n,B}$, for n odd, parameterizing connected covers branched at n points (this is related to a classical result of Hurwitz and, independently, of Clebsch, after which Hurwitz spaces were named). Because the cohomology in low degrees of Z_n agrees with that of Conf_n, the number of $\mathbb{F}_q$ points of these two spaces is nearly the same. As one takes $n \to \infty$ with $n \equiv 1 \bmod 2$, $\frac{\# Z_n(\mathbb{F}_q)}{\# \mathrm{Conf}_n(\mathbb{F}_q)} \to 1$. And this ratio can be interpreted as the same ratio appearing in Theorem 7.0.10.

Proof sketch of Theorem 7.0.10. First, we want to relate the moments predicted by Cohen-Lenstra to counting points on Hurwitz spaces. In the case $H = \mathbb{Z}/\ell\mathbb{Z}$, this is done in Lemma 2.2.5.

To describe the case of general H, let $G = \mathbb{Z}/2\mathbb{Z} \rtimes H$ and take $c \subset G$ the conjugacy class of order 2 elements. In the case of more general H an arbitrary odd order abelian group, if $\mathscr{O}_K$ (which choice of trivialization over ∞) corresponds to a point $x : \operatorname{Spec} \mathbb{F}_q \to \mathrm{CHur}^{\mathbb{Z}/2\mathbb{Z},\{1\}}_{n,B}$ then the fiber of the map $\mathrm{CHur}^{G,c}_{n,B}(\mathbb{F}_q) \to \mathrm{CHur}^{\mathbb{Z}/2\mathbb{Z},\{1\}}_{n,B}(\mathbb{F}_q)$ over the point x has order $\#H \cdot \mathrm{Surj}(\mathrm{Cl}(\mathscr{O}_K), H)$ by an argument similar to that given in Lemma 2.2.5; here, the factor of H comes from the fact that for each cover, there will be $\#H$ choices of point to mark over the given marked point of x.

Therefore, the average of $\# \mathrm{Surj}(\mathrm{Cl}(\mathscr{O}_K), H)$ over fields branched at a degree n divisor over $\mathbb{A}^1$ is identified with $\frac{1}{\#H}$ times the ratio $\frac{\# \mathrm{CHur}^{G,c}_{n,\mathbb{Z}[1/\#\#G]}}{\# \mathrm{Conf}_{n,\mathbb{Z}[1/\#G]}}$, as is obtained from the previous observation by summing over all quadratic extensions branched at n points.

In order to compute the above ratio in the limit $n \to \infty$ (restricting either to n even or odd) Now, we wish to apply Lemma 3.5.1, where the X_n are the varieties $\mathrm{CHur}^{G,c}_{n,\mathbb{Z}[1/\#G]}$ and the Y_n are Conf_n, so we have to verify the three conditions from Notation 3.3.1 and the additional condition from Notation 3.4.2. To verify the first condition of Notation 3.3.1, the maps $H^i(Y_{n,b}, \mathbb{Q}_\ell) \to H^i(X_{n,b}, \mathbb{Q}_\ell)$ are isomorphisms for the geometric point $b = \operatorname{Spec} \mathbb{C}$ and $n > Ii + J$ using Theorem 6.0.2. It follows from a general result that Hurwitz spaces have the same cohomology in characteristic 0 and positive characteristic (see §29) and hence we deduce the above statement for all geometric points b. We will see later in Exercise 11.1.6 that the third condition of Notation 3.3.1 is satisfied as well. Moreover, the condition of Notation 3.4.2 will be satisfied for the X_n if it is satisfied for the Y_n by Exercise 3.5.2 and it is satisfied for the Y_n as explained in Example 3.1.11. We have now verified all the conditions except the first from Notation 3.3.1.

Fix $i \in \{0, 1\}$. It then follows that $\lim_{n \to \infty, n \equiv i \bmod 2} \frac{\# \mathrm{CHur}^{G,c}_{n, \mathbb{Z}[1/\#\#G]}}{\# \mathrm{Conf}_{n, \mathbb{Z}[1/\#G]}} = \kappa$, assuming there is some fixed number κ of geometrically irreducible components of $\mathrm{CHur}^{G,c}_{n, \mathbb{Z}[1/\#G]}$ with n large enough (and $n \equiv i \bmod 2$). Since the average size of the class group will be $\frac{\kappa}{\#H}$, as explained earlier in the proof, we will want to show that $\kappa = 1$ when n is even an large enough while $\kappa = \#H$ when n is odd and large enough.

It remains to determine the value of κ and how that implies Theorem 7.0.10. Finally, it has been shown that when $\gcd(q(q-1), \#H) = 1$ and q is odd, there is a unique geometrically irreducible component of $\mathrm{CHur}^{G,c}_{n, \mathbb{Z}[1/\#G]}$ with monodromy g over $\infty \in \mathbb{P}^1$ (see the proof of [LL24, Proposition 5.3.3] for details). In the case n is even, there is a unique possible value of the boundary monodromy (namely the identity) because we require that the cover is unramified over ∞ while in the case n is odd there are $\#H$ many possibilities because it can take on any value of c. This gives the claim that $\kappa = 1$ in the case n is even and $\kappa = \#H$ in the case n is odd.

Example 7.0.16. For example, when $H = \mathbb{Z}/\ell\mathbb{Z}$ and n is odd, we should at least morally believe $\kappa = 1$ using Corollary 5.0.17. (This does not exactly follow from the above because we have not explained the relation between $[\mathrm{CHur}^{G,c}_{n,B} / G]$ and X^ℓ_g, whose $\mathbb{F}_q$ points are both closely related to the class group, but they are not exactly the same scheme.)

□

8. MALLE

Recall Malle's conjecture roughly predicts the growth of the function $\#[\mathrm{CHur}^{G,c}_{n,B} / G](\mathbb{F}_q)$. This is not exactly true, for the following two reasons. One difference is that Malle's conjecture is about connected covers, not geometrically connected covers. The other difference is that the value of

n is not the same as the discriminant, but instead it counts the number of branch points, which is a slightly different but related invariant. However, ignoring these two differences, we can pretend we want to compute the growth of $\#[\mathrm{CHur}^{G,c}_{n,B}/G](\mathbb{F}_q)$. So, using the idea of Lemma 3.5.1, showing the homology of $\#\,\mathrm{CHur}^{n}_{G}\,c$ stabilizes and computing its stable value should allow us to, at least morally, solve Malle's conjecture.

To state Malle's conjecture, we need to define how we count. We will order points by invariants. The two most important examples of counting invariants for us are discriminant and reduced discriminant, but we work in a more general setting of an arbitrary counting invariant.

Definition 8.0.1. Fix a group G let $\mathrm{inv} : G - \mathrm{id} \to \mathbb{Z}_{>0}$ be a function which is constant on conjugacy classes and with $\mathrm{inv}(g) = \mathrm{inv}(g^j)$ for any j relatively prime to $\mathrm{ord}(g)$. Such a function is a *counting invariant*. For $c_i \subset G$ a conjugacy class, we let $\mathrm{inv}(c_i) := \mathrm{inv}(g)$ for any $g \in c_i$. Given a cover of curves $f : X \to \mathbb{P}^1$, corresponding to a point of Hurwitz space, with n_i geometric points whose inertia lies in c_i, (and we do not count any ramification over ∞) we define the invariant of the cover to be $\mathrm{inv}(f) := \sum_i n_i \,\mathrm{inv}(c_i)$.

Example 8.0.2 (Discriminant)**.** Define the invariant $\Delta(g) := \#G - r(g)$, for $r(g)$ the number of orbits of g acting on G. Assume that $\gcd(q, \#G) = 1$. Then, for $f : X \to \mathbb{P}^1$ a cover over $\mathrm{Spec}\,\mathbb{F}_q$, $\Delta(f)$ corresponds to the discriminant of $f^{-1}(\mathbb{A}^1) \to \mathbb{A}^1$ in the sense that, the discriminant of $f^{-1}(\mathbb{A}^1) \to \mathbb{A}^1$ is $q^{\Delta(f)}$. Here, when we say the discriminant of a map g, we mean $q^{\deg \Omega_g}$.

Exercise 8.0.3. Verify the above claim.

Example 8.0.4. Consider the counting invariant $\mathrm{rDisc}(g) := 1$ for every $g \in G - \mathrm{id}$. This is the reduced discriminant.

From the perspective of Hurwitz spaces, the reduced discriminant is a nice invariant to count by, due to the following exercise:

Exercise 8.0.5. Show that a point of $\mathrm{Hur}^{G,c}_{n,\mathbb{Z}[1/\#G]}$ has reduced discriminant n.

To precisely state Malle's conjecture, we need to define certain constants $a(G)$ and $b_M(K, G, c)$, which

Notation 8.0.6. Fix a finite permutation group $G \subset S_d$ and a counting invariant $\mathrm{inv} : G - \mathrm{id} \to \mathbb{Z}_{>0}$. Define $a(G, \mathrm{inv}) := \min_{g \in G - \mathrm{id}} \mathrm{inv}(g)$. Also define $G_{\mathrm{inv}} \subset G - \mathrm{id}$ to be the set of elements attaining the minimum, i.e. with $\mathrm{inv}(g) = a(G, \mathrm{inv})$. More generally, if $X \subset G$, we use X_{inf} to denote $G_{\mathrm{inv}} \cap X$.

Exercise 8.0.7 (Easy exercise)**.** Show that $G_{\mathrm{rDisc}} = G - \mathrm{id}$.

We now give a temporary definition of the constant b_M associated to the discriminant invariant over the field $\mathbb{Q}$, which was used in Malle's conjecture. Later, we will come back to define b_M for more general invariants.

Notation 8.0.8. Fix a finite group G and let $c := G_\Delta \subset G - \mathrm{id}$ denote the subset with minimum discriminant. For example, if we took $G = S_d$, c would correspond to conjugacy class of transpositions. Define an equivalence relation $\sim$ on the conjugacy classes $c_i \subset c$ so that $c_1 \sim c_2$ if they are related by the cyclotomic character. Specifically, suppose g has order t and the s roots of unity lie in K with $s \mid t$, but no further roots of unity of order dividing t lie in K. Then $c_1 \sim c_2$ if, for s and t as above, we have $g \in c_1$ and $g^j \in c_2$ with $j \equiv 1 \bmod s$. Define $b_M(K, G)$ to be the order of the quotient of the set of conjugacy classes in c by the above equivalence relation.

Conjecture 8.0.9 (Malle)**.** Given a finite permutation group G, and a global field K, there are constants $a(G, \Delta), b_M(K, G)$ as defined above so that, the number of extensions of K of discriminant at most X is bounded above and below by constant multiples of $X^{\frac{1}{a(G,\Delta)}}(\log X)^{b_M(K,G)-1}$.

Remark 8.0.10. Some things are known about Malle's conjecture (mostly over $\mathbb{Q}$) but many things remain open. Here are some cases that are known and unknown

1. Abelian groups
2. S_3, S_4, S_5 due to Bhargava in his thesis and leading up to his fields medal, though S_6 is open.
3. D_8 is known but D_{10} is open.
4. Many nilpotent groups and some wreath products of cyclic groups
5. A_4 is open

Recently, Alberts, Lemke Oliver, Wang, and Wood have developed a fairly comprehensive machine to count the number of G extensions of $\mathbb{Q}$ for many groups G. Their results apply to 1665 of the 4953 transitive permutation groups of degree at most 23 (and many more of higher degree) so in this sample, they are batting around .336.

One interesting aspect of Malle's conjecture is that it is incorrect for general groups G.

Theorem 8.0.11 (Klüners)**.** *If $G = \mathbb{Z}/2\mathbb{Z} \ltimes (\mathbb{Z}/3\mathbb{Z})^2$ the value b_M which Malle predicted is wrong for $G \subset S_6$ the embedding corresponding to thinking of G extensions as the composite of a degree 2 followed by a degree 3 cyclic cubic Galois extension.*

Sketch. Let's understand the basic reason that Malle's conjecture is wrong in this example. First, $G \subset S_6$ extensions can be thought of as the composite of a degree 2 extension followed by a degree 3 extension. First, let's find $a(G)$. Thinking of G as acting on a 6 element set, there are elements α, β which are generators of one of the copies of $\mathbb{Z}/3\mathbb{Z} \subset (\mathbb{Z}/3\mathbb{Z})^2 \subset G$, where α acts by permuting the first three elements. The conjugacy classes of α and β have size 2 because the nontrivial element of $\mathbb{Z}/2\mathbb{Z}$ conjugates α to an element acting on the last three points. Then, α, β, together with their conjugates form the set of elements with minimum value of a. There are 4 such elements and two such conjugacy classes. The value of $a(G, \Delta)$ is therefore $6 - 4 = 2$. Then $K = \mathbb{Q}$, since $\mathbb{Q}$ doesn't have cube roots of unity, these two conjugacy classes are identified under the cyclotomic character, because $\alpha^2 = \beta$, as there are no cube roots of unity in $\mathbb{Q}$. Therefore, $b_M(\mathbb{Q}, G) = 1$. So, Malle's conjecture predicts there are asymptotically $X^{\frac{1}{2}}$ many extensions of discriminant at most X.

However, we can now consider the G extensions which are the composite of a quadratic extension and a cyclic cubic extension, so that the quadratic extension is $\mathbb{Q}(\mu_3)$. Then, we are just counting cyclic cubic extension of $\mathbb{Q}(\mu_3)$. One can show using class field theory that there are roughly $X^{1/2} \log X$ extensions of discriminant at most X.

Roughly, speaking the factor of $\log X$ corresponds to the fact that we can specify such extensions by specifying the primes which are ramified in the extension, and for each such prime, specifying it to have inertia type α or inertia type β. We need roots of unity in the base field to be able to make this specification. Morally speaking, what is going wrong with Malle's conjecture is that, over $\mathbb{Q}$, the cyclotomic character is permuting the two conjugacy classes of minimal a value, but by making the quadratic extension be $\mathbb{Q}(\mu_3)$, the two conjugacy classes are no longer permuted by the cyclotomic character, and this yields more extensions.

As another way to say this, there are more G extensions of $\mathbb{Q}$ counted by discriminant, with quadratic subfield $\mathbb{Q}(\mu_3)$ than there are extensions with all other possible quadratic subfields combined. □

Remark 8.0.12. In the function field case, we should think of Klüners' counterexample as follows: We are considering extensions which are the composite of a quadratic and a cyclic cubic extension. We assume $\mathbb{F}_q$ has characteristic prime to $\#G$ and no cube roots of unity in $\mathbb{F}_q$. Note that $\mathbb{F}_q(\mu_3)$ is a quadratic extension of $\mathbb{F}_q$. If we count G covers so that the intermediate quadratic field is not $\mathbb{F}_q(\mu_3)(t)$, we would get the correct asymptotic. However, we can also count cyclic cubic extensions of $\mathbb{F}_q(\mu_3)(t)$. These correspond to connected curves over $\mathbb{P}^1_{\mathbb{F}_q}$ which are geometrically disconnected, because once we adding μ_3 they will have two disconnected components. Hence, Klüners' result is saying that Malle's prediction can only be correct if one restricts to geometrically connected covers.

In light of this Türkelli made a modified conjecture, using motivation from the function field setting.

Conjecture 8.0.13 (Türkelli)**.** There is a constant $b_T(K, G)$ so that Malle's conjecture is true if one replaces $b_M(K, G)$ by $b_T(K, G)$.

Nevertheless, Wang shows Türkelli's conjecture is false when $K = \mathbb{Q}$.

Theorem 8.0.14 (Wang)**.** *If $K = \mathbb{Q}$, there are counterexamples to Türkelli's conjecture, which are of the form $\mathbb{Z}/\ell\mathbb{Z} \ltimes (\mathbb{Z}/m\mathbb{Z})^{\ell}$ for suitable primes ℓ, m.*

We'd now like to prove a version of Malle's conjecture in the function field setting (over sufficiently large finite fields, depending on the group). To do this, we need to define an analog of Malle's b constant. We first do a simple version of this for connected covers.

Notation 8.0.15. Fix a subset $c \subset G$ and a finite field $\mathbb{F}_q$ with $\gcd(q, \#G) = 1$. Assume c is a union of conjugacy classes and that $g \in c$ implies $g^q \in c$. Let $c_1, \ldots, c_v$ denote the set of conjugacy classes of c and define an equivalence relation on the conjugacy classes by generated by $c \sim c^{q^{-1}}$ (or equivalently $c \sim c^q$). Define $b_M(\mathbb{F}_q(t), c)$ to be the number of equivalence classes of conjugacy classes under the above relation.

We saw above that Malle's conjecture is going to be incorrect in general for all covers, since the geometrically disconnected covers can swamp the geometrically connected covers. However, it's plausible it could hold if one restricts to geometrically connected covers. The next result shows this is indeed the case, and also that Türkelli's conjecture holds if one wishes to count all covers.

Theorem 8.0.16 (Landesman-Levy)**.** *For each finite group G, there is a constant C_G so that if $q > C_G$ and $\gcd(q, \#G) = 1$, then for any counting invariant inv, the number of geometrically connected G extensions of $\mathbb{F}_q(t)$ of invariant at most n is bounded above and below by constants times*

$$q^{\frac{n}{a(G,\mathrm{inv})}} n^{b_M(\mathbb{F}_q(t), c_{\mathrm{inv}})-1}.$$

Moreover, Türkelli's conjecture holds in this setting (if one counts all covers instead of just geometrically connected covers and uses b_T instead of b_M).

Remark 8.0.17. If $X = q^n$ then we see $X^{\frac{1}{a}} \log X^{b-1} = q^{n/a} n^{b-1}$ which explains why the expression in Theorem 8.0.16 looks slightly different than that in Conjecture 8.0.9.

Example 8.0.18. In general, the discriminant of any extension of $\mathbb{F}_q(t)$ has a power of q, so the function $f(X)$ given by the number of G extensions of $\mathbb{F}_q(t)$ with discriminant at most X cannot exist (since the value of this function for $X = q^n$ is the same as the value for $X = q^{n+1/2}$).

A natural fix to this could be to restrict X to powers of q. However, even this limit will rarely exist. As an example, consider the case $G = \mathbb{Z}/3\mathbb{Z}$ and $q \equiv 2 \bmod 3$. In this case, the discriminant is always a square, which again implies that if we counted by discriminant, and took a limit over all powers of q, the limit would not exist.

One could say the above example is again silly, because we just cooked up the invariant to always force the discriminant to be a square. However, even when the invariant does not obviously preclude certain discriminants from appearing, they can still fail to exist. For example, count $\mathbb{Z}/3\mathbb{Z}$ covers by the reduced discriminant, which is equivalent to counting by the square root of the discriminant in this case.

Then, one can show there are no G extensions of $\mathbb{F}_q(t)$ of reduced discriminant q^n for n odd, since there must always be the same number of geometric points with inertia 1 as with inertia $2 \in \mathbb{Z}/3\mathbb{Z}$. However, there are many extensions of reduced discriminant q^n for n even. In general, there can be many obstructions like this.

Remark 8.0.19 (Comparison with Ellenberg–Tran–Westerland)**.** With the goal of proving Malle's conjecture, there has been substantial progress prior to this paper over function fields. Namely, [ETW17] prove a weak upper bound for the number of G extensions with the correct power of X, but with a value of b that is too large in general. Additionally, they do not obtain a lower bound for the number of G extensions. In contrast, our result Theorem 8.0.16 obtains both upper and lower bounds, as well as the correct power of $\log X$.

Let's now explain the proof of this theorem in the special case of counting geometrically connected covers by the reduced discriminant invariant.

Lemma 8.0.20. *Fix a finite group G and let $c = G - \mathrm{id}$. The number of geometrically irreducible components of $[\mathrm{CHur}^{G,c}_{n,\mathbb{F}_q}/G]$ is bounded above by a polynomial function in n of degree $b_M(\mathbb{F}_q(t), c) - 1$. Moreover, if $c = c_1 \cup \cdots \cup c_v$ is a union of conjugacy classes and $n \equiv 0 \bmod v \cdot \#G^2$, the number of components of $[\mathrm{CHur}^{G,c}_{n,\mathbb{F}_q}/G]$ is bounded below by a polynomial in n of degree $b_M(\mathbb{F}_q(t), c) - 1$.*

Proof. First, let's show the claimed upper bound. For each geometric component, there is a corresponding component of multi-colored configuration space $\mathrm{Conf}_{n_1,\ldots,n_v,\overline{\mathbb{F}}_q}$ on which geometric Frobenius acts. The action sends the conjugacy class c_j to $c_j^{q^{-1}}$ and it permutes the $n_1, \ldots, n_v$ in $b_M(\mathbb{F}_q(t), c)$

many orbits. In order for a component of $\mathrm{Conf}_{n_1,\ldots,n_v,\overline{\mathbb{F}}_q}$ to be invariant under this action, we must have $n_i = n_j$ if c_i and c_j lie in the same orbit. In other words, we can specify only $b_M(\mathbb{F}_q(t), c)$ many values of the indices freely, and the rest are determined. The number of such specifications with $n_1 + \cdots + n_v$ is bounded by a polynomial of degree $b_M(\mathbb{F}_q(t), c) - 1$, where the -1 corresponds to the fact that we want to only count components where $n_1 + \cdots + n_v = n$, (as opposed to those where $n_1 + \cdots + n_v \leq n$, for which, the above shows there is a polynomial of degree $b_M(\mathbb{F}_q(t), c)$ counting these. Now, if one fixes the values of $n_1, \ldots, n_j$, there are a bounded number of components of $[\mathrm{CHur}^{G,c}_{n_1,\ldots,n_v,\mathbb{F}_q}/G]$. One way to see this is from Theorem 6.0.1, which implies this bound using the case $i = 0$ (so $H_0([\mathrm{CHur}^{G,c}_{n_1,\ldots,n_v,\mathbb{F}_q}/G])$ stabilizes so long as one of the n_i are sufficiently large. Since there is a polynomial of degree $b_M(\mathbb{F}_q(t), c) - 1$ bounding the number of components of $\mathrm{Conf}_{n_1,\ldots,n_v,\mathbb{F}_q}$ fixed by Frobenius and there are a bounded number of components of $[\mathrm{CHur}^{G,c}_{n_1,\ldots,n_v,\overline{\mathbb{F}}_q}/G]$ over a fixed component of $\mathrm{Conf}_{n_1,\ldots,n_v,\overline{\mathbb{F}}_q}$ one obtains there is a polynomial of degree $b_M(\mathbb{F}_q(t), c) - 1$ many components of $[\mathrm{CHur}^{G,c}_{n_1,\ldots,n_v,\overline{\mathbb{F}}_q}/G]$. This gives the upper bound.

To produce the same lower bound, it is enough to show that whenever $n_1, \ldots, n_j \equiv 0 \bmod \#G^2$ and are sufficiently large then there is always a component of $[\mathrm{CHur}^{G,c}_{n_1,\ldots,n_v,\overline{\mathbb{F}}_q}/G]$ fixed by Frobenius. It is possible to directly analyze how Frobenius acts on stable components of Hurwitz space, and show this is this case. The analysis is a bit involved, so we won't discuss it here, but a detailed proof is given in [LL25a, Lemma 10.3.1]. □

Proof of a special case of Theorem 8.0.16. We can now use the above to prove Theorem 8.0.16 in the case that the covers are geometrically connected and the invariant is the reduced discriminant.

We want to count the number of geometrically connected extensions of $\mathbb{F}_q(t)$ of reduced discriminant n. Since geometrically connected covers reduced discriminant n correspond to points of $[\mathrm{CHur}^{G,G-\mathrm{id}}_{n,\mathbb{F}_q}/G]$, using Exercise 8.0.5, it is equivalent to count $[\mathrm{CHur}^{G,G-\mathrm{id}}_{n,\mathbb{F}_q}/G](\mathbb{F}_q)$. as n grows. Let $c = G - \mathrm{id}$. Using Lemma 8.0.20, the number of geometrically irreducible components of $[\mathrm{CHur}^{G,G-\mathrm{id}}_{n,\mathbb{F}_q}/G]$ is bounded above and below by a polynomial of degree $b_M(\mathbb{F}_q(t), c) - 1$.

We claim that for any such component $Z \subset [\mathrm{CHur}^{G,G-\mathrm{id}}_{n,\mathbb{F}_q}/G]$ there are constants C_1 and C_2 so that $C_2 q^n \leq \#Z(\mathbb{F}_q) \leq C_1 q^n$. From this claim it will follow that $[\mathrm{CHur}^{G,G-\mathrm{id}}_{n,\mathbb{F}_q}/G](\mathbb{F}_q)$ is bounded above and below by the same constants times $q^n n^{b_M(\mathbb{F}_q(t),c)-1}$, which is our goal. To prove the claim, we use Lemma 3.3.2. The claim then follows from Lemma 3.3.2 taking the X_n to be a sequence of components of $[\mathrm{CHur}^{G,G-\mathrm{id}}_{n,\mathbb{F}_q}/G]$, one for each n, say all of which have n_j sufficiently large for a fixed value of j (upon taking q large enough so that the values of the limsups and liminfs are close to their limiting value in q) once we verify the three hypotheses of this lemma stated in Notation 3.3.1. The first hypothesis is satisfied since we are just assuming X_n has a unique geometrically irreducible component. The third hypothesis will be shown to hold in Exercise 11.1.6.

The second hypothesis is the trickiest to verify, but follows from Theorem 6.0.1.

There are a few issues in applying this directly, since Theorem 6.0.1 gives a bound on the homology of $\mathrm{CHur}^{G,G-\mathrm{id}}_{n,\mathbb{C}}$ and not $[\mathrm{CHur}^{G,G-\mathrm{id}}_{n,\mathbb{F}_q}/G]$. First, it follows from a general result that CHur has a relative normal crossings compactification over $\mathbb{Z}[1/\#G]$ that the homology of $\mathrm{CHur}^{G,G-\mathrm{id}}_{n,\mathbb{C}}$ over the complex numbers is identified with the homology of $\mathrm{CHur}^{G,G-\mathrm{id}}_{n,\mathbb{F}_q}$ over finite fields, as we will see later in Theorem 29.0.1 Second, we can deduce that the homologies of the quotient $[\mathrm{CHur}^{G,G-\mathrm{id}}_{n,\mathbb{F}_q}/G]$ stabilize in the same linear range since the homologies of $\mathrm{CHur}^{G,G-\mathrm{id}}_{n,\mathbb{F}_q}$ are actually G representations and stabilize as a G representation, and so the part of this corresponding to the trivial representation is the homology of $[\mathrm{CHur}^{G,G-\mathrm{id}}_{n,\mathbb{F}_q}/G]$. □

Remark 8.0.21. Let us briefly describe the issues going into generalizing the proof of Theorem 8.0.16 from the case of the reduced discriminant invariant to arbitrary invariants. The main issue is that, when one counts by an arbitrary invariant, one has to carefully count the number of geometrically irreducible components of Hurwitz space on which that invariant has value at most n. If the invariant is not the reduced discriminant, this is no longer as simple as finding the number of points on $\mathrm{Hur}^{G,c}_n$ because points of invariant n no longer have n geometric points where they are ramified. For example, if we count by the usual discriminant and count S_3 covers, viewed as degree 3 covers under the identity inclusion of permutation groups $S_3 \subset S_3$, then degree 3 covers with a degree α branch locus of transpositions and a degree β locus with 3-cycle inertia will have discriminant $\alpha + 2\beta$. This will lie on $\mathrm{Hur}^{G,c}_{\alpha+\beta}$ and so we see the value of $\alpha + \beta$ can be anywhere between half the discriminant (if there are only 3-cycles) and the discriminant (if there are only transpositions).

However, if one is careful about keeping track of this, one can write down a generating series for the number of components and use a Tauberian theorem to determine their order of growth. After this, the remainder of the proof is somewhat similar for counting by an arbitrary invariant.

Remark 8.0.22. Let us know briefly describe how one can also deal with disconnected covers. Every geometrically disconnected cover of $\mathbb{P}^1_{\mathbb{F}_q}$ is a geometrically connected cover of $\mathbb{P}^1_{\mathbb{F}_q^j}$ for some j, so we can fix a value of j corresponding to a quotient $G \to \mathbb{Z}/j\mathbb{Z}$ with kernel N. One can then show that such points can be viewed as points of $[\mathrm{Hur}^{N,N-\mathrm{id}}_{n,\mathbb{F}_q}/G]$, which we can think of parameterizing G covers which geometrically become $\#G/\#N$ copies of an N cover. (Over $\mathbb{F}_q$ they may or may not be connected.) If it was only a matter of counting points on this stack, we would be able to similarly apply our homological stability results. However, this stack also contains disconnected cover over $\mathbb{F}_q$. However, the disconnected covers necessarily lie in the image of a map $[\mathrm{Hur}^{N,N-\mathrm{id}}_{n,\mathbb{F}_q}/N']$ for some subgroup $N \subset N' \subset G$. Therefore, we may perform an inclusion exclusion to count the number of points of $[\mathrm{Hur}^{N,N-\mathrm{id}}_{n,\mathbb{F}_q}/G]$ which are geometrically connected.

9. APPLICATION TO THE PICARD RANK CONJECTURE

One fundamental question in algebraic geometry is to identify the set of line bundles on moduli spaces of interest. For example, the rational Picard group of $\mathscr{M}_g$, the moduli space of smooth proper genus g curves is generated by the hodge bundle and $\mathscr{M}_{g,1}$ has rational Picard group generated by the hodge bundle and the relative canonical divisor of the universal curve.

Similarly, there has been substantial interest in what the rational Picard group of Hurwitz spaces are. This interest is coming from algebraic geometry, where mathematicians typically work with Hurwitz spaces over $\mathbb{P}^1$, meaning that we don't fix the branching behavior over ∞ as we did in §4.1.

Here is an informal definition. We will be working only over the complex numbers for this section.

Definition 9.0.1. Let G be a group and $c \subset G$ a conjugacy class. Let $\mathscr{H}^{G,c}_n := [\mathrm{CHur}^{G,c}_{\mathbb{P}^1,n}/G]$ denote the moduli stack of geometrically connected G covers of $\mathbb{P}^1$ (over the complex numbers) so that the reduced branch divisor has degree n with inertia in c. (This is formally defined similarly to Definition 4.1.1 except that the reduced branch divisor over all of $\mathbb{P}^1$ has degree n instead of fixing the branch divisor over $\mathbb{A}^1$.) In the case $G = S_d$ and $c \subset S_d$ is the set of transpositions, we denote this $\mathscr{H}^d_n$.

Remark 9.0.2. In the paper, we preferred to work with the quotient of this by PGL_2, which is a slightly more standard definition of the Hurwitz stack, but this doesn't make much difference, so we'll opt for the slightly simpler scenario not to quotient by PGL_2.

Example 9.0.3. When $G = \mathbb{Z}/2\mathbb{Z}$, and $c \subset \mathbb{Z}/2\mathbb{Z}$ is the nontrivial element, $\mathscr{H}_n^c$ parameterizes hyperelliptic curves. When $G = S_d$ and $c \subset S_d$ consists of transpositions, $\mathscr{H}_n^{G,c}$ parameterizes simply branched covers of $\mathbb{P}^1$.

Exercise 9.0.4. Show that each component of $\mathscr{H}_n^{G,c} = [\mathrm{CHur}^{G,c}_{\mathbb{P}^1,n}/G]$ has a stratification by two components of $[\mathrm{CHur}^{G,c}/G]$. It has an open substack given by a component of $[\mathrm{CHur}_n^{G,c}/G]$ which parameterizes covers unramified over ∞. The other component is a closed substack and corresponds to a components of $[\mathrm{CHur}^{G,c}_{n-1}/G]$ which is branched over ∞.

Exercise 9.0.5. Show that $\mathscr{H}_n^d$ parameterizes simply branched degree d covers of genus 0 curves.

The formal definition in terms of functor of points can be given similarly to that of Definition 4.1.1 except one requires that the total degree of ramification over $\mathbb{P}^1$ is fixed instead of fixing the degree of ramification over $\mathbb{A}^1$.

It turns out that there are no "obvious" non-torsion line bundles on Hurwitz spaces, which leads to the following conjecture.

Conjecture 9.0.6 (Picard rank conjecture)**.** For all $d > 0$, we have $\mathrm{Pic}(\mathscr{H}_n^d) \otimes \mathbb{Q} \simeq 0$.

Remark 9.0.7. This was stated in [HM06, Conjecture 2.49(1)] (with a quotient by PGL_2). See also the closely related [DE96, Conjecture 3]. We note that the roots of this conjecture extend further back, and a version of it appears in work of Ciliberto from 1986 [Cil86, Conjecture 3.2] and even in work of Enriques from 1919 [Enr19, p. 371].

Remark 9.0.8. Conjecture 9.0.6 has been proven for $d \leq 5$ by Deopurkar-Patel [DP15], but this proof relies on explicit parameterizations of low degree Hurwitz spaces quite similar to those Bhargava used to count number fields of degree at most 5. We also note that Mullane has proven Conjecture 9.0.6 whenever $d \leq g$ (so $n \leq 4d - 2$ for $n = 2g - 2 + 2d$) [Mul23]. However, the case that $d \geq 6$ and $g > d$ remains open.

Theorem 9.0.9 (Landesman-Levy, Asymptotic Picard rank conjecture)**.** *If $g \gg d$, the Picard rank conjecture holds true.*

We can prove a more general version of the asymptotic Picard rank conjecture, which works for arbitrary G, c.

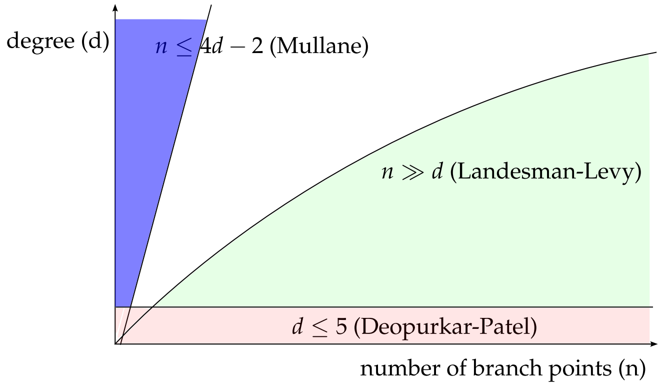

FIGURE 1. This figure depicts known cases of the picard rank conjecture, where $n = 2g + 2d - 2$ is the number of branch points. The $n \leq 4d - 2$ can be reexpressed as $d \leq g$. We computed the constant for our result where it holds roughly for $n > 2^{\binom{d}{2}}$, though it is a bit worse than that.

Theorem 9.0.10. *Let G be a finite group and $c \subset G$ a conjugacy class generating G. For n sufficiently large,* $\mathrm{Pic}(\mathscr{H}_n^{G,c}) \otimes \mathbb{Q} \simeq 0$.

This also works semi-integrally. To explain the precise statement, we need some notation. We also recall that for G a group and c a conjugacy class $H_2(G, c)$ as defined in [Woo21, Definition p. 3] was defined as the quotient of the group homology $H_2(G; \mathbb{Z})$ by the image of all maps $H_2(\mathbb{Z}^2, \mathbb{Z}) \to H_2(G; \mathbb{Z})$ induced by the maps $\mathbb{Z}^2 \to G, (i, j) \mapsto x^i y^j$ for all pairs of commuting $x, y \in c$. Here is our main result toward the Picard rank conjecture.

Theorem 9.0.11. *Let G be a finite group and $c \subset G$ a conjugacy class generating G. For n sufficiently large and each component $Z \subset \mathscr{H}_n^{G,c}$ we have* $\mathrm{Pic}(Z) \otimes \mathbb{Z}[\frac{1}{2\#G}] \simeq \left((\mathbb{Z}/(2n-2)\mathbb{Z}) \otimes \mathbb{Z}[\frac{1}{2\#G}]\right)$.

Let $\mathrm{ord}_{G^{\mathrm{ab}}}(c)$ *denote the order of the image of any element of c in G^{ab}. For n large enough depending on c and divisible by* $\mathrm{ord}_{G^{\mathrm{ab}}}(c)$*, there is a finite number $H_2(G, c)$ so that* $\mathrm{Pic}(\mathscr{H}_n^{G,c}) \otimes \mathbb{Z}[\frac{1}{2\#G}] \simeq \left((\mathbb{Z}/(2n-2)\mathbb{Z}) \otimes \mathbb{Z}[\frac{1}{2\#G}]\right)^{|H_2(G,c)|}$. *If n is sufficiently large and not a multiple of* $\mathrm{ord}_{G^{\mathrm{ab}}}(c)$, $\mathscr{H}_n^{G,c}$ *is empty.*

Remark 9.0.12. The presence of the factor $|H_2(G, c)|$ merely reflects the fact that there are $|H_2(G, c)|$ many components of the Hurwitz space.

There is a branch locus map sending a cover to its branch locus. Since $\mathrm{Pic}(\mathrm{Conf}_{n,\mathbb{P}^1} \otimes \mathbb{Z}[\frac{1}{2\#G}] \simeq \left((\mathbb{Z}/(2n-2)\mathbb{Z}) \otimes \mathbb{Z}[\frac{1}{2\#G}]\right)$ this is expressing the statement that all line bundles on Hurwitz space are pulled back from the corresponding line bundles on configuration space.

Let's explain a proof of Theorem 9.0.9. The key input is Theorem 6.0.2, which roughly says that the first and second rational cohomology groups of Hurwitz space vanish, once the number n of branch points is large enough. And one can relate the Picard group to these two cohomology groups.

Let's now explain this in more detail.

Lemma 9.0.13. *For n sufficiently large $H_i(\mathscr{H}_n^d, \mathbb{Q}) = 0$ for $i \in \{1, 2\}$ and n sufficiently large depending on d.*

In fact, the same proof shows $H_i(\mathscr{H}_n^d, \mathbb{Q}) = 0$ for n large depending on i and d, for i arbitrary.

Proof. This follows from Theorem 6.0.2 as we now explain.

Now, it's a classical fact, due to Hurwitz (and not too difficult to directly prove, that for each even n larger than $2d$, $\mathscr{H}_n^d$ is connected (and it's empty in the case n is odd by Riemann Hurwitz). Using Exercise 9.0.4, there is a stratification of $\mathscr{H}_n^d$ by one component $U \subset [\mathrm{CHur}_n^{G,c}/G]$ and one component of $Z \subset [\mathrm{CHur}_{n-1}^{G,c}/G]$. Moreover, by Theorem 6.0.2, for n sufficiently large the branch locus map induces isomorphisms $H^i(U, \mathbb{Q}) \to H^i(\mathrm{Conf}_{\mathbb{A}^1,n}, \mathbb{Q})$ and $H^i(Z, \mathbb{Q}) \to H^i(\mathrm{Conf}_{\mathbb{A}^1,n-1}, \mathbb{Q})$ for $i \in \{1, 2\}$. Therefore, using the Gysin exact sequence we get a map of exact sequences

(9.1)

$$\begin{tikzcd}
H^i(\mathrm{Conf}_{\mathbb{A}^1,n}, \mathbb{Q}) \arrow[r] \arrow[d] & H^{i-1}(\mathrm{Conf}_{\mathbb{A}^1,n-1}, \mathbb{Q}) \arrow[r] \arrow[d] & H^{i+1}(\mathrm{Conf}_{\mathbb{P}^1,n}, \mathbb{Q}) \arrow[r] \arrow[d] & H^{i+1}(\mathrm{Conf}_{\mathbb{A}^1,n}, \mathbb{Q}) \arrow[r] \arrow[d] & H^i(\mathrm{Conf}_{\mathbb{A}^1,n-1}, \mathbb{Q}) \arrow[d] \\
H^i(U, \mathbb{Q}) \arrow[r] & H^{i-1}(Z, \mathbb{Q}) \arrow[r] & H^{i+1}(\mathscr{H}_n^d, \mathbb{Q}) \arrow[r] & H^{i+1}(U, \mathbb{Q}) \arrow[r] & H^i(Z, \mathbb{Q}).
\end{tikzcd}$$

Since all vertical maps except the middle map induce isomorphisms, the middle map does as well. And so the claim follows because $H^i(\mathrm{Conf}_{\mathbb{P}^1,n}, \mathbb{Q}) = 0$ for $i > 0$. $\square$

Now, we'd like to apply the exponential exact sequence to say that the vanishing of H^1 and H^2 of $\mathscr{H}_n^d$ also implies the vanishing of $\mathrm{Pic}(\mathscr{H}_n^d)$. This is roughly correct, except GAGA won't identify the coherent cohomology of $\mathscr{H}_n^d$ on the Zariski site with the corresponding cohomology on the analytic site. Therefore, we instead argue similarly, but need to pass to a compactification of $\mathscr{H}_n^d$.

Proof of Theorem 9.0.9. Now, fix a smooth proper compactification $\overline{\mathscr{H}}_n^d$ of $\mathscr{H}_n^d$. We'd like to show $\operatorname{Pic}(\mathscr{H}_n^d) \otimes \mathbb{Q} = 0$ for n sufficiently large, depending on d. Let $\partial := \overline{\mathscr{H}}_n^d - \mathscr{H}_n^d$ denote the boundary.

Now, using Lemma 9.0.13, we know $H^1(\mathscr{H}_n^d, \mathbb{Q}) = 0$, and hence the Gysin exact sequence gives us an exact sequence

$$H^{-1}(\partial, \mathbb{Q}) \longrightarrow H^1(\overline{\mathscr{H}}_n^d, \mathbb{Q}) \longrightarrow H^1(\mathscr{H}_n^d, \mathbb{Q}). \tag{9.2}$$

Since the first and last terms are 0, we also obtain $H^1(\overline{\mathscr{H}}_n^d, \mathbb{Q})$. Since the Hodge to de Rham spectral sequence degenerates over the complex numbers, we obtain an injection $H^1(\overline{\mathscr{H}}_n^d, \mathscr{O}) \to H^1(\overline{\mathscr{H}}_n^d, \mathbb{Q})$. Now, the exponential sequence gives an exact sequence

$$H^1(\overline{\mathscr{H}}_n^d, \mathscr{O}) \longrightarrow H^1(\overline{\mathscr{H}}_n^d, \mathscr{O}^\times) = \operatorname{Pic}(\overline{\mathscr{H}}_n^d) \longrightarrow H^2(\overline{\mathscr{H}}_n^d, \mathbb{Z}). \tag{9.3}$$

and importantly, due to GAGA the above analytic sheaf cohomology groups agree with their coherent counterparts. Hence, we obtain an injection $\operatorname{Pic}(\overline{\mathscr{H}}_n^d) \otimes \mathbb{Q} \to H^2(\overline{\mathscr{H}}_n^d, \mathbb{Q})$. Now, the Gysin exact sequence tells us

$$H^0(\partial, \mathbb{Q}) \longrightarrow H^2(\overline{\mathscr{H}}_n^d, \mathbb{Q}) \longrightarrow H^2(\mathscr{H}_n^d, \mathbb{Q}). \tag{9.4}$$

is exact. The last term is 0 by Lemma 9.0.13, and so the first map is a surjection. On the other hand the first map factors through $\operatorname{Pic}(\overline{\mathscr{H}}_n^d) \otimes \mathbb{Q}$ implying $H^0(\partial, \mathbb{Q}) \to \operatorname{Pic}(\overline{\mathscr{H}}_n^d) \otimes \mathbb{Q}$ is also a surjection, as we showed above $\operatorname{Pic}(\overline{\mathscr{H}}_n^d) \otimes \mathbb{Q} \to H^2(\mathscr{H}_n^d, \mathbb{Q})$ is an injection. Finally, the excision exact sequence tells us

$$H^0(\partial, \mathbb{Q}) \longrightarrow \operatorname{Pic}(\overline{\mathscr{H}}_n^d) \otimes \mathbb{Q} \longrightarrow \operatorname{Pic}(\mathscr{H}_n^d) \otimes \mathbb{Q} \longrightarrow 0 \tag{9.5}$$

is exact. The first map being surjective then implies $\operatorname{Pic}(\mathscr{H}_n^d) \otimes \mathbb{Q} \simeq 0$. □

10. A PRIMER ON HIGHER ALGEBRA

The setting for many of the arguments we aim to present related to homological stability is that of higher algebra. The language of higher algebra takes place in the setting of infinity categories.

Remark 10.0.1 (Why are we learning higher algebra?)**.** The need for higher algebra in the ensuing proof is relatively minimal. However, it will be important for us to work in that category. The reason for this is that we want to study the homology of Hurwitz space, $H_i(\mathrm{CHur}^c, \mathbb{Z})$ which is the

ith homology of a chain complex $C_*(\mathrm{CHur}^c, \mathbb{Z})$. In order to compute the stable homology, we will want to perform certain operations on this chain complex, such as localization and quotienting. These operations only make sense (in the way we want them to) in the derived infinity category of $\mathbb{Z}$ modules. For example, there are certain stabilization maps $\alpha_x : \mathrm{CHur}^c_n \to \mathrm{CHur}^c_{n+1}$ corresponding to taking a cover with n branch points and adding a branch point labeled x near the boundary. The stable homology can then be identified with the homology of the localization $C_*(\mathrm{CHur}^c, \mathbb{Z})[\alpha_x^{-1}]$. So, it will be important to make sense of this localization, and that can be done correctly in the derived infinity category of $\mathbb{Z}$ modules. With this in mind, we next describe higher algebra, with the ultimate goal in this section as defining the chains functor, which sends a topological space to an object in the derived infinity category of $\mathbb{Z}$ modules.

Remark 10.0.2 (The vague idea of higher algebra)**.** In case people are familiar with stacks, one may know that they are similar to schemes, but one also wishes to keep track of automorphisms, or "maps between maps." The general idea in higher algebra is that these automorphisms can be viewed as 1-morphisms between a point and itself, but one will also want to include morphisms between morphisms, and morphisms between morphisms between morphisms, and so on.

Another perspective on this is that if one has a topological space Y, one can think of paths as morphisms between points. A morphisms between morphisms would be a homotopy of paths. If one only cares about understanding $\pi_1(Y)$, understanding homotopies is enough. But if one wishes to understand $\pi_i(Y)$ for $i > 1$, it is important to also understand the homotopies between homotopies, and so on.

10.1. **Simplicial sets.**

The formal gadget we will use to capture the data of morphism spaces is known as a simplicial set.

Definition 10.1.1. Let Δ denote the *simplex category* whose objects are nonempty totally ordered finite sets and whose morphisms are weakly order preserving functions between these sets.

Definition 10.1.2. A *simplicial set* is a contravariant functor $X : \Delta \to \mathrm{Set}$ for Set the category of sets. The morphisms between two simplicial sets are given by natural transformations of functors. We denote the category of simplicial sets by sSet.

Remark 10.1.3. Every element of Δ is isomorphic to one of the form $[n] = \{0, 1, \dots, n\}$. Then a map $f : [m] \to [n]$ is a map of sets so that $f(i) \le f(j)$ if $i \le j$. One can express arbitrary such maps in terms of *face maps* $\delta_i : [n-1] \to [n]$, which is the injection not containing i in the image, and the *degeneracy* *maps* $\sigma_i : [n+1] \to [n]$, which is the surjection so that i and $i+1$ both map to i. For X a simplicial set, let $X_n := X([n])$. Applying a functor $X : \Delta \to \mathrm{Set}$ yields face maps $d_i : X_n \to X_{n-1}$ and degeneracy $s_i : X_n \to X_{n+1}$. These face and degeneracy maps satisfy various conditions as in Exercise 10.1.4.

Exercise 10.1.4. Check the following simplicial identities.

1. $d_i d_j = d_{j-1} d_i$ for $i < j$.
2. $d_i s_j = s_{j-1} d_i$ for $i < j$.
3. $d_i s_j = \mathrm{id}$ for $i = j$ or $i = j+1$.
4. $d_i s_j = s_j d_{i-1}$ for $i > j+1$.
5. $s_i s_j = s_{j+1} s_i$ for $i \leq j$.

Exercise 10.1.5. Show that the category of simplicial sets is equivalent to the category of sequences of sets X_n, together with collections of maps $d_i : X_n \to X_{n-1}, 1 \leq i \leq n$, and $s_i : X_n \to X_{n+1}, 1 \leq i \leq n$ satisfying the simplicial identities in Exercise 10.1.4. Hence, simplicial sets can also be defined in terms of face and degeneracy maps.

One of the key examples of simplicial sets are those coming from a topological space.

Exercise 10.1.6. Suppose Y is a topological space. We use Δ^n to denote the n-simplex, viewed here as a topological space (in the future, we will typically view it as a simplicial set). Define a simplicial set $\mathrm{Sing}(Y)$ whose value on Δ^n is given by $\mathrm{Sing}(Y)_n := \mathrm{Hom}(\Delta^n, Y)$, meaning maps of topological spaces $\Delta^n \to Y$. Check this defines a functor and also verify the simplicial identities from Exercise 10.1.4 hold.

10.2. **Spaces.**

Having defined simplicial sets, we will next define spaces. Rather, we will define objects called *Kan complexes*.

Remark 10.2.1. If you talk to a topologist, ask them "Is a space the same thing as a Kan complex?" It seems they will invariably say something like "Yeah, it's fine to think that a space is the same thing as a Kan complex." (This interaction seems to imply that a space is actually something else, but no one has given me a straight answer to what the ideal definition of a space is. The reason for this is that they often think of spaces like sets. If a mathematician asks me what the definition of a set is, I often wouldn't have a very good answer. They often think of spaces in terms of their properties. I.e. they are thingies where you can take their homotopy groups, and various limits and colimits exist, etc.)

In what follows we the following definition of Δ^n.

Definition 10.2.2. Let Δ^n denote the simplicial set whose k simplices are given by sequences $0 \leq a_0 \leq a_1 \leq \cdots \leq a_k \leq n$, with the natural face and degeneracy maps.

Exercise 10.2.3. Show that Δ^n is representable by $[n] := \{0, \ldots, n\}$. That is, if X is a simplicial set, show

$$\operatorname{Hom}(\Delta^n, X) = X([n]).$$

We next define the i horn Λ^n_i loosely speaking, to be the subset of all faces of Δ^n containing the ith vertex.

Definition 10.2.4. The i-horn Λ^n_i is $\cup^n_{j=0, i\neq j} \operatorname{im}(\delta_j : \Delta^{n-1} \to \Delta^n) \subset \Delta^n$. (If you want, it is a good exercise to think through formally what this means. Informally, it consists of all the faces of the n simplex except the n dimensional one and the face consisting of all vertices except i.)

Example 10.2.5. We use notation $[i, j]$ to denote the edge of a simplex connecting vertex i to vertex j. If $n = 2$, the Λ^2_0 consists of the edges $[0, 1]$ and $[0, 2]$ of a triangle, which looks like a horn when drawn. Similarly, Λ^2_1 consists of the edges $[0, 1]$ and $[1, 2]$.

Definition 10.2.6. A *Kan complex* or *space* is a simplicial set X so that any horn $\Lambda^n_i \to X$ for $0 \leq i \leq n$ can be extended to a map

(10.1)
$$\begin{tikzcd}
\Lambda^n_i \arrow[rr] \arrow[dr] & & \Delta^n \arrow[dl] \\
& X &
\end{tikzcd}$$

Why is a Kan complex called a space? The following example may help to explain why.

Exercise 10.2.7. Suppose Y is a topological space. Show the simplicial set $\operatorname{Sing}(Y)$ is a Kan complex.

Remark 10.2.8. It turns out there is also a functor from simplicial sets to topological spaces, which is called the "geometric realization," and this gives a somewhat stronger closer relation between the two. Namely, not all Kan complexes are the nerve of a topological space, but they are all equivalent to the nerve of a "compactly generated weak Hausdorff topological space" for a notion of equivalence we will define later.

Example 10.2.9. However, not every simplicial set is the singular chains of a space. For example, you can take the simplicial set generated by a 0-simplex, and a 1-simplex joining this point to itself. This is equivalent to a circle, but the singular chains on a circle has many more 1-simplices than this does.

For an example of a simplicial set that is not a Kan complex, see Example 10.3.8.

10.3. **Infinity categories.** Having defined spaces (also known as Kan complexes), we next turn to ∞-categories. Sometimes, these are referred to as $(\infty, 1)$ categories, meaning that all k-morphisms for $k \geq 2$ are invertible. We will not have any need for (∞, j) categories for $j > 1$, other than in the next remark, and so we simply refer to these as ∞ categories.

Warning 10.3.1. You may start talking to a homotopy theorist who mentions something about 2-categories. You may be used to thinking about 2-categories as stacks, but beware that homotopy theorists often use the phrase 2-category to mean $(\infty, 2)$ category. Sometimes, the homotopy theorist may then question why you thinking about 2 categories at all, and tell you that you should instead be thinking of them as $(\infty, 1)$ categories.

In order to motivate infinity categories, we consider the following example of the nerve of a category.

Example 10.3.2. Let C be a category. Let C_n denote the set of all composable chains of morphisms $A_0 \to A_1 \to \cdots \to A_n$ for A_i objects of C.

Exercise 10.3.3. Show one can endow $\{C_n\}_{n \geq 0}$ with the structure of a simplicial set.

We call the above simplicial set the *nerve* of C. One can determine the objects of C up to isomorphism from the nerve of C.

Exercise 10.3.4. Let X be a simplicial set which is the nerve of a category. Show that $\mathrm{Hom}(\Delta^n, X) \to \mathrm{Hom}(\Lambda^n_i, X)$ is a bijection for $0 < i < n$.

Exercise 10.3.5. Conversely, show that if X is a simplicial set such that $\mathrm{Hom}(\Delta^n, X) \to \mathrm{Hom}(\Lambda^n_i, X)$ is a bijection for $0 < i < n$, then X is the nerve of a category. *Hint:* What are the objects and morphisms in this category?

Having seen that spaces consist of simplicial sets where all horns can be filled and categories consist of simplicial sets where intermediate horns can be filled uniquely, we define ∞ categories as a common generalization. Namely, those where intermediate horns can be filled.

Definition 10.3.6. An ∞ category is a simplicial set X so that any map of simplicial sets $\Lambda^n_i \to X$ for $0 < i < n$ can be extended to a map

$$\begin{tikzcd} \Lambda^n_i \arrow[rr] \arrow[dr] & & \Delta^n \arrow[dl] \\ & X & \end{tikzcd} \tag{10.2}$$

A map of ∞ categories is a natural transformation of the corresponding simplicial sets.

Example 10.3.7. Any space is an ∞ category because spaces require horn filling for $0 \leq i \leq n$ whereas ∞ categories only require it for $i < i < n$. In particular, for Y a topological space, $\operatorname{Sing}(Y)$ is an ∞ category.

Example 10.3.8. Here is an example of an ∞ category whose nerve is not a space. Consider the category with objects x, y, z, with identity maps as well as maps $f : x \to z$ and $g : y \to z$, but no other maps. Then its nerve is not a space because the horn with vertices going to x, y, z and maps going to f and g cannot be filled to a map from the 2-simplex.

Exercise 10.3.9. Given any category C, show that the nerve of C is an ∞ category.

Fix a category $\mathcal{C}$. Given objects A, B, C of $\mathcal{C}$ and two maps $f : A \to B$ and $g : B \to C$, there is a unique composition $g \circ f : A \to C$. For a general ∞ category, there is no longer a unique such composition, though we are guaranteed that one such composition exists by definition. We next define a notion of homotopy in an ∞ category.

Definition 10.3.10. Let X be an ∞ category. Let A, B, C be three 0 simplices and let $f : A \to B, g : B \to C$ and $h : A \to C$ be three 1-simplices. The notation above means that we can view a 1-simplex as a map of simplicial sets $\Delta^1 \to X$ and the restriction of this to $\Delta^0 = [0] \subset \Delta^1 \to X$ gives A while the restriction to $\Delta^0 = [1] \subset \Delta^1 \to X$ gives B. We say $g \circ f \simeq h$ ($g \circ f$ is *homotopic* to h) if there is a simplex $\Delta^2 \to X$ so that Δ^2 has edge $[0, 1]$ sent to f, edge $[1, 2]$ sent to g and edge $[0, 2]$ sent to h.

Exercise 10.3.11. Suppose X is an ∞ category and f, g, h are 1-simplices as in Definition 10.3.10. If $g \circ f \simeq h$ and $\simeq g \circ f \simeq h'$ show $h \simeq h'$. (Strictly speaking, it might be better to write $h \circ \mathrm{id} \simeq h'$ since the definition of homotopy above is in terms of identifying a composition of two maps with a single map.) Therefore, any two such choices of compositions are homotopic. *Hint:* Use the horn filling property from the definition. Don't forget the degeneracy maps in your simplicial set.

Exercise 10.3.12. Check the preceding exercise concretely in the case that $X = \operatorname{Sing}(Y)$, for Y a topological space by showing that the two paths corresponding to h and h' are homotopic, and identifying homotopies of paths with homotopies in the above sense.

One very basic notion we will need for infinity categories is the notion of an equivalence. This is the analog of the usual notion of isomorphism and roughly says there should be a map the other way whose compositions with the given map are the identity.

Definition 10.3.13. If X is an ∞ category, we refer to an element of $X([0])$, or equivalently a 0-simplex $\Delta^0 \to X$, as an *object* of X. We refer to a 1-simplex $\Delta^1 \to X$, or equivalently an element of $X([1])$ as a *morphism* in X. If a and b are two objects of X, we say $f : a \to b$ is a morphism if f is a morphism in X whose restriction to the point 0 of the simplex Δ^1 is the 0-simplex a and whose restriction to the point 1 of the simplex Δ^1 is the object b.

Definition 10.3.14. If X is an ∞ category, a and b are objects of X and $f : a \to b$ is a morphism, we call f an *equivalence* if there is some $g : b \to a$ and homotopies $\mathrm{id}_a \simeq g \circ f, \mathrm{id}_b \simeq f \circ g$.

Definition 10.3.15. Suppose C and D are simplicial sets. Define the simplicial set $\mathrm{Fun}(C, D)$ so that its n-simplices are given as the set of natural transformations (which we will also call maps) $C \times \Delta^n \to D$.

Proposition 10.3.16. *If C is a simplicial set and D is an ∞ category then* $\mathrm{Fun}(C, D)$ *is an ∞ category.*

Definition 10.3.17. Suppose C and D are two simplicial sets. For $f, g \in \mathrm{Fun}(C, D)$ two 0-simplices, a morphism $\eta : f \to g$ is a map $C \times \Delta^1 \to D$ restricting to $f : C \times 0 \to D$ and to $g : C \times 1 \to D$. Such an η is called a *natural transformation*.

Using the above, we can now define equivalence.

Definition 10.3.18. A map $f : C \to D$ of ∞ categories (recall this means that f is a natural transformation of the corresponding simplicial sets) is an *equivalence* if there is a map $g : D \to C$ so that there are natural transformations $f \circ g \simeq \mathrm{id}_D$ and $g \circ f \simeq \mathrm{id}_C$.

There is an alternate way to define equivalence, via mapping spaces, where we define equivalences as maps which are fully faithful and essentially surjective. We now define some terminology to set this up.

It turns out one can make sense of pullbacks in the category of simplicial sets.

Definition 10.3.19. Let C be an infinity category and $a, b \in C$. Then define $\mathrm{Map}_C(a, b)$, the *mapping space,* is defined as the pullback in the category of simplicial sets

$$\begin{tikzcd}
\mathrm{Map}_C(a,b) \arrow[r] \arrow[d] & \mathrm{Fun}(\Delta^1, C) \arrow[d] \\
\Delta^0 \arrow[r, "{(a,b)}"] & C \times C.
\end{tikzcd} \tag{10.3}$$

where the right vertical map is induced by sending a functor $\Delta^1 \to C$ to the two objects of C given by composition with the two maps $\Delta^0 \to \Delta^1$ associated to 0 and 1.

Exercise 10.3.20. Check that in $\mathrm{Sing}(X)$, $\mathrm{Map}_{\mathrm{Sing}(X)}(a, b)$ is homotopy equivalent to the space of paths from a to b in X, where composition of maps comes from path composition.

Exercise 10.3.21. For X a simplicial set, let $\pi_0(X)$ denote the set of 0 simplices in X modulo the equivalence relation generated by the relations that two 0 simplices are equivalent if they are endpoints of some 1-simplex. Show $\pi_0 \mathrm{Map}_C(a, b)$ is the set of equivalence classes of morphisms from a to b, where, viewing two morphisms as two objects of $\mathrm{Map}_C(a, b)$, the notion of equivalence of these two objects is as defined in Definition 10.3.14. *Hint:* For this, show that all 1-simplices (which one can think of as a homotopy between morphisms) in mapping spaces are invertible.

Definition 10.3.22. Let X be a space and x an object of X, i.e., x is a 0-simplex of X. We call such a pair (X, x) a *pointed space*.

Consider the set of maps $\Delta^n \to X$ sending $\partial\Delta^n \to x$, where $\partial\Delta^n$ is the simplicial set generated by all faces of Δ^n (so it includes all $n-1$ dimensional subsimplices but not the n-dimensional simplex). We define $\pi_n(X, x)$ as the maps $(\Delta^n, \partial\Delta^n) \to (X, x)$ as above, modulo the equivalence relation called *homotopy*, where two such maps are homotopic if there is a map $\Delta^1 \times \Delta^n \to X$ restricting to the first map on $0 \times \Delta^n \to X$, restricting to the second map on $1 \times \Delta^n \to X$, and which sends $\Delta^1 \times \partial\Delta^n$ to x.

Definition 10.3.23. For (X, x) and (Y, y) pointed spaces, a map $f : (X, x) \to (Y, y)$ of pointed space is a homotopy equivalence if f induces isomorphisms $\pi_i(X, x) \to \pi_i(Y, y)$ for all $i \geq 0$. A map $f : X \to Y$ of (unpointed) spaces is a homotopy equivalence if f induces a bijection on components and, for all points $x \in X$ the map $\pi_i(X, x) \to \pi_i(Y, f(x))$ is an isomorphism for all $i \geq 1$.

One may show that the mapping space of two objects in an infinity category is an space, which will be used to make sense of the following definition.

Definition 10.3.24. A map of infinity categories $f : C \to D$ is *fully faithful* if $\mathrm{Map}_C(a, b) \to \mathrm{Map}_D(f(a), f(b))$ are homotopy equivalences for each $a, b \in C$. A map $f : C \to D$ is *essentially surjective* if for every $d \in D$ there is some $c \in C$ so that d is equivalent to $f(c)$.

The following proposition is somewhat technical to prove.

Proposition 10.3.25. *A map $f : C \to D$ is an equivalence (in the sense of Definition 10.3.18) if it is fully faithful and essentially surjective.*

10.4. **The derived infinity category.** The derived infinity category is a refinement of the usual derived category of an abelian category. Objects are bounded below chain complexes of projective modules and the morphisms are also defined similarly, but in the infinity category setting, there are also higher morphisms.

Definition 10.4.1. A *differential graded category* $\mathcal{C}$ consists of the data of

1. A collection of objects of $\mathcal{C}$, denoted $\mathrm{Ob}(\mathcal{C})$
2. For each pair of objects $X, Y \in \mathrm{Ob}(\mathcal{C})$ a chain complex $(\mathrm{Hom}_{\mathcal{C}}(X,Y)_*, \partial)$.
3. For each triple X, Y, Z of objects of $\mathcal{C}$ and each $m, n \in \mathbb{Z}$ a composition law $c_{Z,Y,X} : \mathrm{Hom}_{\mathcal{C}}(Y,Z)_n \times \mathrm{Hom}_{\mathcal{C}}(X,Y)_m \to \mathrm{Hom}_{\mathcal{C}}(X,Z)_{m+n}$ which sends $f \in \mathrm{Hom}_{\mathcal{C}}(X,Y)_m, g \in \mathrm{Hom}_{\mathcal{C}}(X,Z)_n$ to $c_{Z,Y,X}(g,f) := g \circ f$.
4. For each object X an identity morphism $\mathrm{id}_X \in \mathrm{Hom}_{\mathcal{C}}(X,X)_0$.

such that

1. the composition law is associative, i.e. $h \circ (g \circ f) = (h \circ g) \circ f$
2. the composition law is unital i.e., $\mathrm{id} \circ f = f = f \circ \mathrm{id}$
3. the composition law is bilinear and satisfies the Leibniz rule, i.e.,

$$\begin{aligned} g \circ (f + f') &= g \circ f + g \circ f' \\ (g + g') \circ f &= g \circ f + g' \circ f \\ \partial(g \circ f) &= (\partial g) \circ f + (-1)^n g \circ (\partial f). \end{aligned}$$

We will only really care about subcategories of the following differential graded category, the category of chain complexes of R-modules:

Definition 10.4.2. For $\mathcal{A}$ an abelian category, define $\mathrm{Ch}(\mathcal{A})$ as the category whose objects are chain complexes

$$\cdots \to X_i \to X_{i-1} \to \cdots$$

with each X_i in $\mathcal{A}$ and the maps are given as follows: If X and Y are two objects of $\mathrm{Ch}(\mathcal{A})$ which we think of as chain complexes, the set of maps is $\oplus_n \mathrm{Hom}_n(X,Y)$ where we define

$$\mathrm{Hom}_n(X,Y) := \prod_{\ell \in \mathbb{Z}} \mathrm{Hom}(X_\ell, Y_{\ell+n}).$$

This set of homomorphisms can be given the structure of a chain complex as follows: There is a differential $d : \mathrm{Hom}_n(X,Y) \to \mathrm{Hom}_{n-1}(X,Y)$ given by sending $f \in \mathrm{Hom}_n(X,Y)$, to

$$df := d_Y \circ f + (-1)^{n+1} f \circ d_X.$$

In this way, we can enrich the set of homomorphisms with the structure of a chain complex in abelian groups.

For R a ring, we use $\mathrm{Ch}(R)$ to denote $\mathrm{Ch}(\mathrm{Mod}(R))$, for $\mathrm{Mod}(R)$ the abelian category of R modules, which we can view as a differential graded category.

Remark 10.4.3. It may same strange to the reader that $\mathrm{Hom}(X, Y[n]) := \prod_{\ell \in \mathbb{Z}} \mathrm{Hom}(X_\ell, Y_{\ell+n})$ and we do not impose the condition that this is a chain map. The reason for this is that this is the internal hom.

Exercise 10.4.4. Verify that chain complexes $\mathrm{Ch}(R)$ indeed form a differential graded category, by verifying the axioms.

Show that chain maps $f : X \to Y$ (meaning maps f with $f \circ d_X = d_Y \circ f$) are precisely the maps $f \in \mathrm{Hom}_{\mathrm{Ch}(R)}(X, Y)_0$ such that $df = 0$.

Definition 10.4.5. Let $\mathcal{C}$ be a differential graded category. The *differential graded nerve* of $\mathcal{C}$, denoted $N_{dg}(\mathcal{C})$ is the following simplicial set. We define $N_{dg}(\mathcal{C})_n$, which we use as notation for $\mathrm{Hom}_{\mathrm{sSet}}(\Delta^n, N_{dg}(\mathcal{C}))$ to be the set of ordered pairs $(\{X_i\}_{0 \leq i \leq n}, \{f_I\})$ with X_i an object of $\mathcal{C}$ and I ranging over subsets of $[n]$. When $\#I < 2$ we declare $f_I := 0$. Otherwise, if $\#I \geq 2$, we can write $I = \{i_- < i_m < i_{m-1} < \cdots < i_1 < i_+\} \subset [n]$ with $m \geq 0$, $f_I \in \mathrm{Hom}_{\mathcal{C}}(X_{i_-}, X_{i_+})_m$ such that

$$df_I = \sum_{1 \leq j \leq m} (-1)^j (f_{I - \{i_j\}} - f_{i_j < \cdots < i_+} \circ f_{i_- < \cdots < i_j}).$$

Moreover, if $\alpha : [m] \to [n]$ is nondecreasing then $N_{dg}(\mathcal{C})_n \to N_{dg}(\mathcal{C})_m$ is defined by

$$(\{X_i\}_{0 \leq i \leq n}, \{f_I\}) \to \left(\left\{ X_{\alpha(j)} \right\}_{0 \leq j \leq m}, \{g_J\} \right)$$

with

$$g_J = \begin{cases} f_{\alpha(J)} & \text{if } \alpha|_J \text{ is injective} \\ \mathrm{id}_X & \text{if } J = \{j, j'\} \text{ with } \alpha(j) = \alpha(j') = i \\ 0 & \text{otherwise} \end{cases}$$

Exercise 10.4.6. Show that

1. 0-simplices of $N_{dg}(\mathcal{C})$ are in bijection with objects of $\mathcal{C}$
2. 1-simplices of $N_{dg}(\mathcal{C})$ correspond to pairs of objects X, Y with $f \in \mathrm{Hom}_{\mathcal{C}}(X, Y)_0$ so that $df = 0$
3. 2-simplices correspond to $X, Y, Z \in \mathcal{C}$ and $f \in \mathrm{Hom}_{\mathcal{C}}(X, Y)_0, g \in \mathrm{Hom}_{\mathcal{C}}(Y, Z)_0, h \in \mathrm{Hom}_{\mathcal{C}}(X, Z)_0$ with $df = dg = dh = 0$ and $z \in \mathrm{Hom}_{\mathcal{C}}(X, Z)_1$ with $dz = (g \circ f) - h$.

Proposition 10.4.7. *If $\mathcal{C}$ is a differential graded category then $N_{dg}(\mathcal{C})$ is an ∞ category.*

We omit the proof, although this is not too difficult to prove directly from the definition by writing down an explicit formula that shows one can fill inner horns, see [Lur17, Proposition 1.3.1.10].

Definition 10.4.8. For $\mathcal{C}$ a differential graded category, let $N_{dg}(\mathcal{C})$ denote its differential graded nerve. For R a ring, let $\mathrm{Mod}(R)$ denote the abelian category of R modules and let $\mathrm{Proj}(R)_{\geq 0} \subset \mathrm{Ch}(R)_{\geq 0} \subset \mathrm{Ch}(R)$ denote the full subcategory of chain complexes of projective A modules in non-negative degrees. For R a ring, define the *derived* ∞ *category of R-modules* (in non-negative degrees) as $D(R)_{\geq 0} := N_{dg}(\mathrm{Proj}(R)_{\geq 0})$, viewed as an ∞ category via Proposition 10.4.7.

Definition 10.4.9 (The chains functor)**.** Let X be a space. Let $C_*(X)$ denote the complex defined by taking $C_i(X)$ to be the free abelian group generated by $\mathrm{Hom}(\Delta^i, X)$ with differentials induced by the face maps. We view $C_*(X)$ as a complex of free $\mathbb{Z}$ modules, and hence as an object viewed as an object in the derived ∞ category $D(\mathbb{Z})_{\geq 0}$.

More generally, for R a commutative ring, let $C_*(X;R)$ denote the analogous construction where the ith term $C_i(X;R)$ is the free R module generated by $\mathrm{Hom}(\Delta^i, X)$. We view $C_*(X;R)$ as an object in the derived ∞ category $D(R)_{\geq 0}$. When X is a topological space, we use $C_*(X) := C_*(\mathrm{Sing}(X))$.

Fact 10.4.10. For X a space or topological space, there is a way to view $C_*(X;R)$ as something called a "spectrum" which has associated homotopy groups. It turns out that $H_i(X;R) = \pi_i\left(C_*(X;R)\right)$. For the purposes of this course we use this notation so that we match the notation in our papers. It does have a meaning but the meaning does not especially concern us for this course.

Remark 10.4.11. Another type of object that often comes up in higher algebra is *spectra*. We won't really need to work with them, but it might be good to be aware of their existence. They are roughly the initial place in higher algebra where one can do algebra. There's something called the sphere spectrum, which is the initial spectrum. Roughly speaking, you can think that if you take a spectrum, and base change from the sphere spectrum to $\mathbb{Z}$, you get an object in the derived infinity category of $\mathbb{Z}$ modules.

(To flesh out the above a bit using some words we haven't yet encountered, it turns out spectra have a universal property. They are the initial symmetric monoidal stable presentable infinity category, such that the tensor product commutes with colimits. Being stable means that it has an initial and final object 0, and that pushout squares are also pullback squares.)

11. HOMOLOGICAL STABILITY FOR CONFIGURATION SPACES

Before moving onto homological stability for Hurwitz spaces, it is worth meditating on the simpler case of configuration spaces. Despite being fairly classical, these arguments are quite tricky, and there are many mistakes in the literature associated with these arguments. We first show the homology groups of configurations spaces stabilize, and then compute their stable value.

11.1. **Homological stability for configuration spaces.**

First, let's show that the homology groups of configuration spaces stabilize. There are many ways to prove this. The one we will employ uses a particular cell structure on configuration space, called Fox-Neuwirth cells.

Definition 11.1.1. Let Conf_n denote the configuration space of n distinct points in the complex numbers. For $(n_1, \ldots, n_k)$ with $n_1 + \cdots + n_k = n$, the open cell (which is the interior of a closed subset of Conf_n) $\mathrm{FN}_{(n_1,\ldots,n_k)} \subset \mathrm{Conf}_n$ consists of tuples of distinct points $(x_1, \ldots, x_n) \in \mathbb{C}^n$ such that there are real numbers $r_1 < r_2 < \cdots < r_k$ so that the n_j points $x_{n_1+\cdots+n_{j-1}+1}, \ldots, x_{n_1+\cdots+n_j}$ have real coordinate r_j. configuration space.

Example 11.1.2. $\mathrm{FN}_{(1,2,1)}$ is the subset of points $(x, y, z, w) \in \mathbb{C}^4$ where x has the smallest real coordinate, y and z have the same real coordinate, and w has the largest real coordinate.

Remark 11.1.3. We now want to define a cochain complex C^* which computes the cohomology of the one point compactification of configuration space. As you will verify in Exercise 11.1.5, the $\mathrm{FN}_{n_1,\ldots,n_k}$ form a stratification of Conf_n, moreover their closures in the one point compactification of Conf_n give one a chain complex computing the relative cohomology groups $H^i(\mathrm{Conf}_n \cup \infty, \infty; \mathbb{Z}) \simeq H^i_c(\mathrm{Conf}_n; \mathbb{Z})$.

Hence, the Fox-Neuwirth cells give us a cochain complex C^* whose cohomology $H^i(C^*)$ computes $H^i(\mathrm{Conf}_n \cup \infty, \infty; \mathbb{Z})$. Let's be a bit more explicit about what C^* is. In Exercise 11.1.5, you will verify that the dimension of $\mathrm{FN}_{n_1,\ldots,n_k}$ is $n + k$. Then, $C^i = \oplus_{(n_1,\ldots,n_k):n+k=i} \mathbb{Z} e_{n_1,\ldots,n_k}$, where we think of $e_{n_1,\ldots,n_k}$ as being the generator associated to $\mathrm{FN}_{n_1,\ldots,n_k}$, which has dimension $n + k$.

Next, we wish to compute the differential $C^i \to C^{i+1}$. Equivalently, taking $n + k = i + 1$, we want to compute the coefficient in the differentials sending $e_{m_1,\ldots,m_{k-1}}$ to $e_{n_1,\ldots,n_k}$. To understand the coefficient of this differential, we have to understand which cells lie in the closure of others. (This way computes the cohomology, which records the cells $\mathrm{FN}_{m_1,\ldots,m_{k-1}}$ whose closure contains $\mathrm{FN}_{n_1,\ldots,n_k}$. However, it is equivalent to compute the coefficient in the dual chain complex which records when $\mathrm{FN}_{n_1,\ldots,n_k}$ lies in the closure of $\mathrm{FN}_{m_1,\ldots,m_{k-1}}$.) The cell associated with $(n_1, \ldots, n_k)$ in dimension $n+k$ can only specialize into those cells in dimension $n+k-1$ of the form $(n_1, \ldots, n_{j-1}, n_j + n_{j+1}, n_{j+2}, \ldots, n_k)$ where the specialization map comes from letting the real coordinates r_j and r_{j+1} tend toward each other and then shuffling the n_j points having real coordinate r_j with the n_{j+1} points having real coordinate r_{j+1}.

We now want to compute the matrix coefficient in the map $C^i \to C^{i+1}$ going between $e_{n_1,\ldots,n_k}$ and $e_{n_1,\ldots,n_j+n_{j+1},\ldots,n_k}$. Overall, there will be a sum of $\binom{n_j+n_{j+1}}{n_j}$ terms corresponding to the different ways to shuffle the two columns of n_j and n_{j+1} elements together. Namely, if we want to shuffle a collection $(x_1, \ldots, x_\alpha)$ of α elements, with a collection $(y_1, \ldots, y_\beta)$ of β elements. There are $\binom{\alpha+\beta}{\beta}$ ways to order these $\alpha + \beta$ elements while preserving the order of the $x_1, \ldots, x_\alpha$ and the order of the $y_1, \ldots, y_\beta$. We call the group of such permutations (α, β) shuffles, and denote the set of such permutations by $\Sigma_{\alpha,\beta}$. The differential is then given by $S_{\alpha,\beta} := \sum_{\sigma \in \Sigma_{\alpha,\beta}} (-1)^{\mathrm{sgn}(\sigma)}$, where $\mathrm{sgn}(\sigma)$ denotes the sign of the permutation σ.

Summarizing the above, the matrix coefficient of the differential from $e_{n_1,\ldots,n_k}$ to $e_{n_1,\ldots,n_j+n_{j+1},\ldots,n_k}$ $S_{n_j,n_{j+1}}$, and the only nonzero matrix coefficients are of this form.

The above is a bit notationally heavy, so let's see an example that will hopefully clarify what the chain complex computing the homology of configuration space from Fox-Neuwirth cells is.

Example 11.1.4. Let's work out the cochain complex computing $H^i(\mathrm{Conf}_3 \cup \infty, \infty; \mathbb{Z})$ using Fox-Neuwirth cells. The cells we have are $\mathrm{FN}_{1,1,1}$ which has dimension 6, $\mathrm{FN}_{1,2}$ and $\mathrm{FN}_{2,1}$ which have dimension 5 and FN_3 which has dimension 4. Therefore the cochain complex is of the form

$$\mathbb{Z}e_3 = C^4 \to \mathbb{Z}e_{1,2} \oplus \mathbb{Z}e_{2,1} = C^5 \to \mathbb{Z}e_{1,1,1} = C^6.$$

We now want to compute the differentials. The matrix $C^4 \to C^5$ is given by the transpose of $(1, 1)$ and the matrix $C^5 \to C^6$ is $(0, 0)$. This follows directly from the formula given in Remark 11.1.3 but let's also explain more directly how to compute it.

First, let's think about the matrix coefficient from $e_{2,1}$ to $e_{1,1,1}$. If we think of $e_{1,1,1}$ as corresponding to points x, y, z with distinct real coordinates and $e_{2,1}$ as triples x, y, z where x and y have the same real coordinate, then the differential keeps track of ways to degenerate x and y to have the same real coordinate. There are two such ways: either x goes above y or x goes beneath y. Keeping track of orientations, one can see these appear with opposite signs, and so the matrix coefficient is 0.

Next, let's think about the matrix coefficient from e_3 to $e_{2,1}$. Then, $e_{2,1}$ corresponds to triples x, y, z where x and y have the same real coordinate but z has a smaller real coordinate, while e_3 corresponds to triples where all 3 have the same real coordinate. There are three ways to degenerate corresponding to the cases that z goes above x and y, z goes between x and y, or z goes below x and y. If we choose orientations suitably, the total sum of these contributions is $1 - 1 + 1 = 1$. The other matrix coefficients can be computed similarly so this gives the claimed cochain complex.

Finally, using Poincaré duality, we can identify the cochain complex C^{6-*} whose ith cohomology computes $H^{6-i}(\mathrm{Conf}_3 \cup \infty, \infty; \mathbb{Z})$ with a chain complex C_* explicitly of the form $C_2 \to C_1 \to C_0$ whose ith homology computes $H_i(\mathrm{Conf}_3; \mathbb{Z})$.

Exercise 11.1.5.
(1) Check that the dimension of the cell $\mathrm{FN}_{(n_1,\ldots,n_k)}$ is $n + k$.
(2) Verify the Fox-Neuwirth cells form a stratification of configuration space and their closures consist of cells for the one point compactification of configuration space. In particular, show that the chain complex associated to Fox-Neuwirth cells computes the relative cohomology group $H^i(\mathrm{Conf}^n \cup \infty, \infty)$.

Exercise 11.1.6 (Important exercise)**.**
(1) Show that there are fewer than 2^n Fox-Neuwirth cells in the stratification of Conf_n.
(2) Deduce that there is a cell structure for Hur^c_n with fewer than $2^n |c|^n$ cells.
(3) Conclude that $\dim H_i(\mathrm{Hur}^c_n, \mathbb{Q}) \leq (2|c|)^n$.
(4) Using homological stability for Hurwitz spaces Theorem 6.0.1, and the previous part, show there is a constant C (depending on the conjugacy class c) so that $H_i(\mathrm{CHur}^c_{n_1,\ldots,n_k}, \mathbb{Q}) \leq C^i$. (Note that the previous part bounds the left hand side by C^n but your task is to bound it by C^i.)

Remark 11.1.7. By the above exercise, the chain complex associated to Fox-Neuwirth cells computes $H^i(\mathrm{Conf}^n \cup \infty, \infty)$.

Since we have isomorphisms $H_{2n-i}(\mathrm{Conf}_n) \simeq H^i_c(\mathrm{Conf}_n) \simeq H^i(\mathrm{Conf}_n \cup \infty, \infty)$ coming from Poincaré duality and the definition of compactly supported cohomology, we can compute the cohomology of configuration space in terms of Fox-Neuwirth cells. Here, the boundary maps are obtained by letting two columns collide with each other and keeping track of all the ways to shuffle the two adjacent columns, with suitable signs.

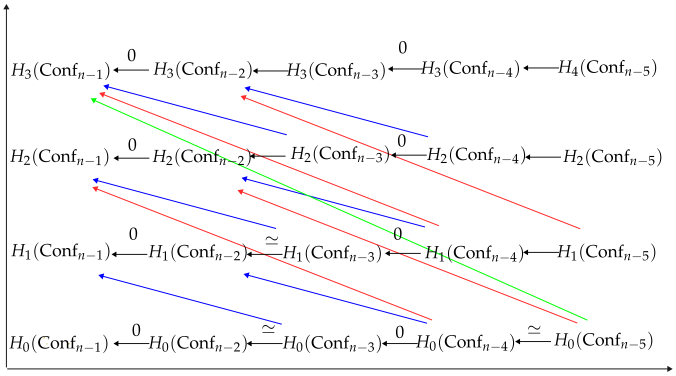

FIGURE 2. This depicts the spectral sequence for configuration space coming from the Fox Neuwirth stratification of configuration space.

Theorem 11.1.8. *The cohomology of configuration space stabilizes. That is, for any fixed i and for n sufficiently large depending on i,* $H_i(\mathrm{Conf}_n;\mathbb{Z}) \simeq H_i(\mathrm{Conf}_{n+1};\mathbb{Z})$.

Proof. Let C^{2n-*} denote the chain complex whose basis elements are Fox-Neuwirth cells for Conf_n, as described above in Remark 11.1.3. This cochain complex can be identified with a chain complex C_*, which we denote C_*^n to emphasize its dependence on the n in Conf_n. Then, C_*^n computes the homology of Conf_n as explained in Remark 11.1.7. Let $F_*^i \subset C_*^n$ denote the subcomplex consisting of those cells C_*^n with $n_1 \geq i+1$. Note that this is a subcomplex because the differentials are obtained by letting two real coordinates collide, and so if we require that there are at least n_1 elements in the first column, this will be closed under the differential. From this filtration, we obtain a spectral sequence with E_1 page $H_j(F_*^i/F_*^{i+1}) \implies H_{i+j}(\mathrm{Conf}_n;\mathbb{Z})$. We will next see $H_i(F_*^i/F_*^{i+1}) \simeq H_i(\mathrm{Conf}_{n-(i+1)};\mathbb{Z})$. So, this spectral sequence is depicted in Figure 2.

Now, the ith graded part F_*^i/F_*^{i+1} can be thought of as those cells for configuration space where the first column has exactly $i+1$ elements (instead of at least $i+1$) and the differentials involving the first column all vanish. This is exactly the chain complex C_*^{n-i} computing the homology of Conf_{n-i}. (We can think of this as corresponding to the rightmost $n-i$ points in the original configuration space Conf_n.)

Exercise 11.1.9. Show the maps on the E_1 page $H_j(F_*^{i+1}/F_*^{i+2}) \to H_j(F_*^i/F_*^{i+1})$ are given by 0 when i is even and the stabilization map (which is the map induced on homology by adding a point near ∞) when i is odd. *Hint:* One way to see this is to think through the differentials in the Fox-Neuwirth chain complex. If one merges the left column with i elements with a column having a single element, there are $i+1$ ways to shuffle the two elements, and the signs appearing in the shuffles alternate between 1 and -1. Summing the signs gives either 0 or 1 depending on the cases above. (You can also see this from the explicit formula for the differentials in terms of the sum $S_{i,1}$ from Remark 11.1.3; see also Example 11.1.4.)

Now, inductively assuming that the homologies $H_j(\mathrm{Conf}_n;\mathbb{Z})$ for $j < k$ have stabilized, we will prove H_k stabilizes. The spectral sequence above identifies $H_k(\mathrm{Conf}_n;\mathbb{Z})$ with the E_∞ page of the spectral sequence, whose terms on the E_1 page consist of $H_k(\mathrm{Conf}_{n-1};\mathbb{Z}), H_{k-1}(\mathrm{Conf}_{n-2};\mathbb{Z}), \ldots, H_0(\mathrm{Conf}_{n-k};\mathbb{Z})$. By inductive assumption, the E_1 differentials induce isomorphisms $H_{k-i}(\mathrm{Conf}_{n-i-1};\mathbb{Z}) \leftarrow H_{k-i}(\mathrm{Conf}_{n-i-2};\mathbb{Z})$ when i is odd and an isomorphism $H_{k-i}(\mathrm{Conf}_{n-i};\mathbb{Z}) \leftarrow H_{k-i}(\mathrm{Conf}_{n-i-1};\mathbb{Z})$ when i is even. This implies that on the E_2 page, we get that the kth diagonal is $H_k(\mathrm{Conf}_{n-1};\mathbb{Z}), 0, \ldots, 0$. Note that $H_k(\mathrm{Conf}_{n-1};\mathbb{Z})$ survives to the E_2 page because the map from $H_k(\mathrm{Conf}_{n-2};\mathbb{Z})$ is 0 by Exercise 11.1.9. Finally, it remains to show $H_k(\mathrm{Conf}_{n-1};\mathbb{Z})$ survives to the E_∞ page of the spectral sequence. To see this, the maps on the E_r page come from the spot where $H_{k-r}(\mathrm{Conf}_{n-r-2};\mathbb{Z})$ was on the E_1 page. However, by inductive assumption, we can assume these already stabilized, and so die on the E_2 page, and hence we obtain that the boundary map $H_k(\mathrm{Conf}_{n-1};\mathbb{Z}) \to H_k(\mathrm{Conf}_n;\mathbb{Z})$ is an isomorphism.

Strictly speaking, to make the above precise, we need to keep track of an explicit range so that we can properly run the above induction. We leave this to the reader in Exercise 11.1.10. □

Exercise 11.1.10. Work carefully the the proof of Theorem 11.1.8 to obtain an explicit linear range for which the homology of configuration space stabilizes. In particular, show that $H_i(\mathrm{Conf}_{n-1};\mathbb{Z}) \simeq H_i(\mathrm{Conf}_n;\mathbb{Z})$ is an isomorphism when $i \leq \frac{n-1}{2}$. *Hint:* The two tightest cases to worry about are what happens on the E_2 page with the map going from the locations of $H_{k-1}(\mathrm{Conf}_{n-3};\mathbb{Z})$ and $H_{k-2}(\mathrm{Conf}_{n-4};\mathbb{Z})$ on the E_1 page to the location of $H_k(\mathrm{Conf}_{n-1};\mathbb{Z})$ on the E_1 page.

11.2. **The bar construction.**

Definition 11.2.1. A topological monoid is a topological space M together with an an associative map $\cdot : M \times M \to M$ with an identity element $\mathrm{id} \in M$ satisfying $\mathrm{id} \cdot m = m = m \cdot \mathrm{id}$ for all $m \in M$.

It turns out configuration space is not quite a monoid, but it is "up to homotopy" see Exercise 11.3.9. However, one can make a Moore variant of configuration space which records a time parameter, and this will be a topological monoid, see Example 11.2.9.

Definition 11.2.2. Let M be a topological monoid, X be a right M module and Y be a left M module. (That X is a right M module means there is an action map $X \times M \to X$ so that $x \cdot (m \cdot n) = (x \cdot m) \cdot n$.) With X, Y, M as in Definition 11.2.2. Then, the bar construction, which we will notate as $X \otimes_M Y$ or $\mathrm{Bar}(X, M, Y)$, is defined as follows: (We will use these interchangeably but the first is more commonly thought of as the homotopy type of the space, and the second is more commonly thought of as the particular construction we give below.)

It is the colimit (or geometric realization) of

$$\cdots \to X \times M^3 \times Y \to X \times M^2 \times Y \to X \times M \times Y \to X \times Y \tag{11.1}$$

where really there are $n+1$ maps from $X \times M^n \times Y \to X \times M^{n-1} \times Y$ given by either the leftmost element of M acting on X or Y or multiplying two adjacent elements of M. There are also n maps in the reverse direction $X \times M^{n-1} \times Y \to X \times M^n \times Y$ given by including and taking the identity element of the monoid in position n. When we say it is the geometric realization, we mean one can take the geometric realization of this complex, which means that we take the colimit of the spaces along these maps, or equivalently we take the disjoint union and glue along these maps. Let us now spell out concretely what this means.

For M a topological monoid, X is a right M module and Y is a left M module, we can form the bar construction $\mathrm{Bar}(X, M, Y)$, which is a model for the tensor product $X \otimes_M Y$ (and we will often denote it as this tensor product). Consider pairs of tuples of points $((m_1, \ldots, m_n), (t_1, \ldots, t_n), x, y)$ with $m_i \in M$ and $0 \leq t_1 \leq t_2 \leq \cdots \leq t_n \leq 1$, $x \in X$, $y \in Y$. Then, $X \otimes_M Y$ is the quotient of the above space modulo the following equivalence relations:

1. If $t_1 = 0$ then the above tuple is identified with $((m_2, \ldots, m_n), (t_2, \ldots, t_n), x \cdot m_1, y)$ (so m_1 acts on x on the left of the interval)
2. If $t_n = 1$ the above is identified with $((m_1, \ldots, m_{n-1}), (t_1, \ldots, t_{n-1}), x, m_n \cdot y)$ (so m_n acts on y on the right edge of the interval)
3. If $t_i = t_{i+1}$, then the above is equivalent to
$$((m_1, \ldots, m_i \cdot m_{i+1}, \ldots, m_n), (t_1, \ldots, t_i, \widehat{t_{i+1}}, \ldots, t_n), x, y)$$
(so points at the same time can multiply in the monoid)
4. If $m_i = \mathrm{id}$ is the identity element in the monoid, the above is equivalent to $((m_1, \ldots, m_{i-1}, m_{i+1}, \ldots, m_n), (t_1, \ldots, t_{i-1}, t_{i+1}, \ldots, t_n), x, y)$, so identity elements can disappear.

Remark 11.2.3. Loosely speaking, we can think of the bar construction as being represented by points along an interval, which can fall of the edges or multiply when they hit each other.

Remark 11.2.4. Why do we use the notation $\otimes$? There is a sense in which one can think of this as a tensor product. If you are used to seeing tensor products of modules over a ring, you may recall $X \otimes_M Y$ is identified with the quotient of $X \otimes Y$ by the relation $(xm, y) \sim (x, my)$. In other words, it is the coequalizer of the two maps $X \otimes M \otimes Y \to X \otimes Y$ given by either multiplying M on the left or the right. Said another way it is $\mathrm{Tor}^0_M(X, Y)$. However, when we are dealing with topological spaces in the setting of higher algebra, we also want to remember the higher tor groups or higher homotopies, and so we need to take the whole complex as in the definition of the bar construction to model $X \otimes_M Y$.

Example 11.2.5. Consider the case $X = Y = \mathrm{pt}$ with the action of the monoid on X and Y being trivial (everything sends the point to itself). Then, $X \otimes_M Y$ is commonly called BM and is the classifying space of that monoid. This can be given a model by the bar construction where we can think of collections of points in M on the interval, which can fall off the edges of the interval (which corresponds to them acting trivially on the X and Y.

Remark 11.2.6. More generally, one can make sense of the bar construction when X is just a $\mathbb{E}_1$ algebra (i.e. an associative algebra in spaces) and X and Y are right and left modules for M. This will be useful if one wants to use the usual configuration space (on which multiplication is only associative up to homotopy, but not strictly associative).

Example 11.2.7. If we take $M = G$ a finite group, then G forms a monoid (and even a group) via multiplication. Take $X = Y = \mathrm{pt}$ so that the bar construction is the (perhaps familiar) BG. We can form a model for the bar construction given by

$$\cdots \to G^3 \to G^2 \to G \to \mathrm{pt}\,.$$

If you have seen group homology, you may be familiar with this construction, as the cohomology of the above complex computes the group homology of G, or, equivalently the homology of BG.

Remark 11.2.8. More generally, if G acts on some topological space Y, then we can form the bar construction $\mathrm{pt} \otimes_G Y$ with

$$\cdots \to G^3 \times Y \to G^2 \times Y \to G \times Y \to Y$$

It turns out that one can identify this with the homotopy quotient $Y//G$. For example, if G acts freely on Y, this is just the usual quotient space Y/G.

Let's now see a couple examples of bar constructions that are relevant to configuration space.

Example 11.2.9. Take M to be the topological monoid of pairs (x, t) where $t \in \mathbb{R}_{\geq 0}$, and there is some $n \in \mathbb{Z}_{\geq 0}$ so that x is a configuration $x = (x_1, \ldots, x_n)$ of n points in the interior of the rectangle $(0, t) \times (0, 1)$. This is a so-called "Moore" variant of configuration space of points in the plane. See Figure 3 for a visual depiction of the multiplication.

Consider the bar construction BM. We can think of this as configurations of points in a rectangle. (This is identification is certainly not trivial, but hopefully it is believable, where, if one starts with an element of the bar construction, one can think of it as a collection of rectangles at various points of the interval, and one can think of this construction as gluing these rectangles together.) Here, points are not allowed to collide with each other, but the can collide with the left or right boundaries of the rectangle, in which case they disappear.

Example 11.2.10. Let's now take M to again be the configuration space monoid from the previous Example 3.1.11. However, let's take X to be the left module $\mathbb{Z}$ and Y to be the right module $\mathbb{Z}$ where the action of $(x = (x_1, \ldots, x_n), t) \in M$ acts by addition by n on both modules above. Then, there is a model for the bar construction associated to $X \otimes_M Y$ above as configurations of points in a rectangle, together with integer labels of the left and right edges of the rectangle. But now, when a point hits the left or right edge, it disappears and increments that edge label by 1. Points are also allowed to emerge from the left or right edge, in which case the label decreases by 1. In this case, the sum of the left label, the number of points in the rectangle, and the right label are constant in any given connected component.

11.3. **The stable homology of configuration spaces.** Having shown configuration spaces satisfy homological stability, we next compute the stable homology via a procedure called "scanning." This will be the model example for computing the stable homology, and it will be helpful to understand this argument carefully, as we will do many similar but more complicated scanning arguments in the future.

The other ingredient, in addition to scanning, is the group completion theorem. First, we need to define loop spaces.

Definition 11.3.1. Let X, Y, Z be spaces with maps $X \to Y$ and $Y \to Z$, the homotopy pullback $X \times_Y Z$ is defined to be the space whose n-simplices $\Delta^n \to X \times_Y Z$ are given by the data of maps $\Delta^n \to X, \Delta^n \to Z$ together with a homotopy between the compositions $\Delta^n \to X \to Y$ and $\Delta^n \to Y \to Z$. This gives the maps from n simplices, and one can use this to construct natural face and degeneracy maps to give $X \times_Y Z$ the structure of a Kan complex. One can verify that $X \times_Y Z$ is a space (assuming X, Y, and Z are spaces).

Exercise 11.3.2. Show that one can give the following equivalent definition of the pullback $X \times_Y Z$ in spaces. Namely, for X, Y, Z spaces, $X \times_Y Z$ is the pullback in the category of Kan complexes

$$\begin{tikzcd}
X \times_Y Z \arrow[r] \arrow[d] & X \times Z \arrow[d] \\
\mathrm{Map}(\Delta^1, Y) \arrow[r] & Y \times Y,
\end{tikzcd} \tag{11.2}$$

where $\mathrm{Map}(\Delta^1, Y)$ is the simplicial set whose n simplices are given by maps $\Delta^n \times \Delta^1 \to Y$.

For now, we only need the special case of pullbacks that X and Z are points.

Definition 11.3.3. In general, for (X, x) a pointed space, we use ΩX to denote the loop space of X, which is the fiber product $* \times_X *$. The points of this can be thought of as a homotopy between the basepoint and itself, i.e., a loop at the basepoint.

For X a pointed space, we use $\Omega_0 X$ to denote the component of ΩX containing the basepoint (the constant loop at x).

Exercise 11.3.4.
1. Show that the 0 simplices of ΩX are in bijection with maps $\Delta^1 \to X$ sending both end points to the marked point x of X (which you can think of as a map $\Delta^0 \to X$.
2. Conclude $\pi_0(\Omega X) \simeq \pi_1(X, x)$.
3. Can you figure out the relation between the homotopy groups of X and ΩX in general? (No need to write up a proof of this, but just make a guess.) *Hint:* Already the case of $\pi_1(\Omega X, x)$ is a lot trickier. One wants to identify maps from $\Delta^1 \times \Delta^1$ constant on the boundary with maps from Δ^2 constant on the boundary. One can do this by covering the boundary by two triangles, though it is quite involved.

The following theorem (and its generalizations) is our main tool for computing the stable value of homology of configuration spaces.

Theorem 11.3.5 (Group completion theorem)**.** *Suppose M is a topological monoid with $\pi_0(M) = \mathbb{N}$ and use M_j to denote the component of M indexed by $j \in \mathbb{N}$.*

Then $\operatorname{colim}_{j \in \mathbb{N}} H_*(M_j; \mathbb{Z}) \simeq H_*(\Omega_0 BM; \mathbb{Z})$, *where the transition maps are given by multiplication by powers of some fixed choice of $m \in M_1$.*

Remark 11.3.6. Let us spell out more concretely what the above is saying. In the above setting, we can identify $H_*(\Omega_0 BM; \mathbb{Z})$ as the stable homology of M. Spelled out, the left hand side of the group completion theorem becomes

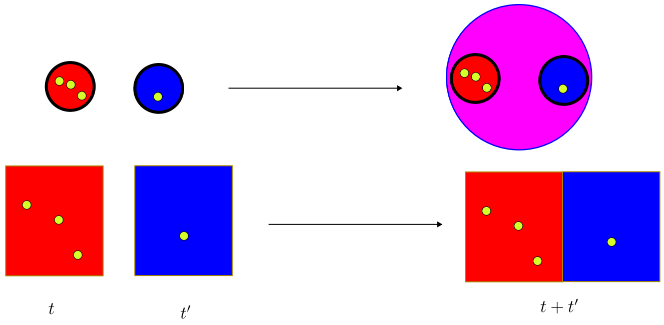

FIGURE 3. This is a picture of the monoid structure on configuration space. The first line depicts the monoid structure on configuration space, which is not quite a topological monoid because it is not associative, but it is a monoid up to homotopy. It depicts the map $\mathrm{Conf}_3 \times \mathrm{Conf}_1 \to \mathrm{Conf}_4$. The second picture shows the Moore variant where we keep track of a time parameter, which is a topological monoid.

the colimit $\cdots H_*(M;\mathbb{Z}) \to H_*(M;\mathbb{Z}) \to \cdots$ where the multiplication operation is by some representative of the element $m \in M_1$. In other words, the colimit consists of elements in some $H_i(M_j;\mathbb{Z})$, modulo the relation that $x \in H_i(M_j;\mathbb{Z})$ is identified with y if they agree after multiplication by m^j for some j. If the homology stabilizes, by M_{α_i}, this means the maps $H_i(M_\beta;\mathbb{Z}) \to H_i(M_{\beta+1};\mathbb{Z})$ are all isomorphisms for $\beta > \alpha_i$, so this colimit agrees with $H_i(M_{\alpha_i};\mathbb{Z})$. So, the group completion theorem is saying one can identify the stable homology of M with the homology of $\Omega_0 BM$.

Remark 11.3.7. In general, even if the homology doesn't stabilize, one can still speak of the stable homology, which just means the above colimit.

Example 11.3.8. For example, we can take M to be the monoid $\mathrm{Conf} := \cup_{n\in\mathbb{N}} \mathrm{Conf}_n$, or we can also take M to be the monoid $\cup_{g\in\mathbb{N}} B\,\mathrm{Mod}_g^1$. (This latter example is closely related to the moduli space of curves and is of much interest throughout mathematics, though we won't explain how to make this into a monoid in these notes.) Strictly speaking, Conf_n does not quite have the structure of a topological monoid The monoidal operation on Conf_n is roughly given by taking $X \in \mathrm{Conf}_n$ and $Y \in \mathrm{Conf}_m$ two configurations in a disc, and embedding the two discs into a bigger disc to obtain an element of Conf_{n+m}. See Figure 3. Technically speaking, this is not actually a topological monoid (though it is up to homotopy so one can make sense of it as a monoid in an appropriate infinity categorical setting). Instead, if one wished to work with an actual topological monoid, one can use the Moore variant from Example 11.2.9.

Exercise 11.3.9. Check the above described monoidal operation on Conf does not actually make Conf into a topological monoid, since the operation will not be associative. Explain why this issue is remedied by working with the Moore variant of configuration space from Example 11.2.9 (This will mean that our upcoming arguments should mostly be interpreted up to homotopy, and to get actual rigorous arguments, one will often need to "strictify" the heuristic arguments we give in the scanning argument.)

Concretely, in the case $M = \mathrm{Conf}_n := \cup_{n \geq 0} \mathrm{Conf}_n$ is the configuration space of points in $\mathbb{R}^2$ (or in the disc) the above is saying that the stable homology of Conf_n, i.e., $H_i(\mathrm{Conf}_n; \mathbb{Z})$ for n sufficiently large (which was shown to stabilize in Theorem 11.1.8) agrees with $H_i(\Omega B \,\mathrm{Conf}; \mathbb{Z})$.

The next theorem computes the stable homology of configuration space. Namely, it shows it is that of $\Omega^2_0 S^2$ which can be completely understood using homotopy theory (which we won't go into in this course).

Theorem 11.3.10. *There is an equivalence* $\Omega B \,\mathrm{Conf} \simeq \Omega^2 S^2$.

Remark 11.3.11. One may still wonder how to compute the homology of $\Omega^2 S^2$. It turns out that $\Omega^2_0 S^2$ has the rational homology of a circle, which is to say $H_i(\Omega^2_0 S^2; \mathbb{Q}) = H_i(S^1; \mathbb{Q})$ which is $\mathbb{Q}$ is $i \in \{0, 1\}$ and 0 otherwise. However, $\Omega^2_0 S^2$ does have a lot of torsion homology.

Proof. Recall that a space is *grouplike* if $\pi_0 X$ is a group.

Fact 11.3.12. In general, for any grouplike space X, we can identify $\Omega B X \simeq X$. (This is a generalization of the fact that for G a group, $\Omega B G = G$; the idea of the proof of this is essentially that $\pi_i(BG) = G$ if $i = 1$ and 0 otherwise, and Ω shifts homotopy groups by 1 so $\pi_i \Omega B G = G$ if $i = 0$ and 0 otherwise, implying $G \to \Omega B G$ is an equivalence.)

Now, since $B\,\mathrm{Conf}$ is grouplike as its component group is a point, the above implies $\Omega B \,\mathrm{Conf} \simeq \Omega(\Omega B) B \,\mathrm{Conf}$. Hence, to complete the proof, it suffices to identify $B^2 \,\mathrm{Conf} := B(B \,\mathrm{Conf}) \simeq S^2$. This is done in Lemma 11.3.18. $\square$

We now build up to checking $B^2 \,\mathrm{Conf} \simeq S^2$ in four steps. The first two involve identifying $B\,\mathrm{Conf}$ and the next two involve identifying $B^2\,\mathrm{Conf}$ similarly.

Lemma 11.3.13. $B\,\mathrm{Conf}$ *can be identified with the space of points in a rectangle where points can fall off or emerge from the left or right sides (but not the top or bottom).*

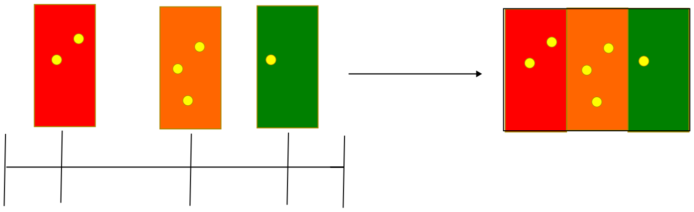

FIGURE 4. This is a picture of the first part of the scanning argument, which takes a point of $B\operatorname{Conf}$, thought of as a collection of configurations of points, and puts them together to get a single point of configuration space of points in a rectangle, but where the points are allowed to fall off the ends.

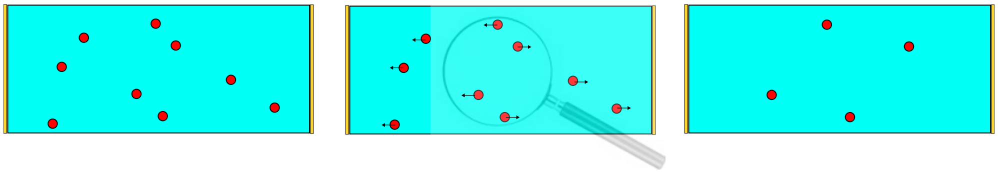

FIGURE 5. This is a picture of the scanning argument of Lemma 11.3.15 where one pushes the points out horizontally to arrange that there is only one in each row.

Proof sketch. The idea is to view a point of the bar construction $((m_1, \ldots, m_n), (t_1, \ldots, t_n))$ as a collection of n configurations in a rectangle, with the ith such configuration has r_i points, and is centered in a small interval around t_i. This can alternatively be described as a single configuration of $r_1 + \cdots + r_n$ points in a rectangle. See Figure 4. (This procedure is part of what is sometimes called scanning.) The equivalence relations from the definition of the bar construction precisely translate to the condition that points are free to fall off the left or right ends of the rectangle.

Exercise 11.3.14 (strenuous exercise, optional)**.** The above is clearly not completely rigorous but gets across the idea. To make it completely rigorous requires some work, and in particular one has to show that one obtains a homeomorphism respecting the topologies on both sides.

□

Lemma 11.3.15. *$B\operatorname{Conf}$ can be identified with the space of points in a rectangle where point can fall off or emerge from the left or right sides (but not the top or bottom), and, moreover, there are no two point with the same vertical coordinate.*

sketch. First, using Lemma 11.3.13, we can put ourself in the setting described there, and we only need to show we can arrange that there are no two points

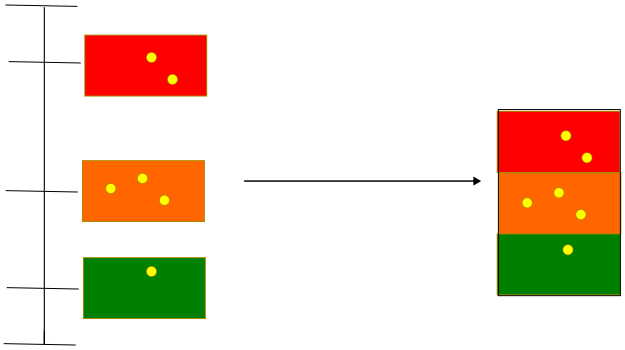

FIGURE 6. This is a picture depicting depicting the proof idea of Lemma 11.3.17.

with the same vertical coordinate. This can be achieved by zooming in to an interval close to the middle of the rectangle. Either there will be no points in a small neighborhood of the middle of the rectangle, in which case the resulting configuration will be empty. If there are points remaining, we can zoom in sufficiently far so that there is at most one point remaining in each vertical coordinate. (If there are two, there is some definite distance between them and we can zoom in enough so at least one falls off.) Thus, we have achieved the condition that there is at most one with any given vertical coordinate. (This procedure of zooming in is sometimes referred to as scanning.) See Figure 5.

Exercise 11.3.16 (very strenuous exercise, optional)**.** Fill in the details of the above proof. *Hint/Reference:* One possible reference for a very similar statement and argument is [BDPW23, Theorem 5.2.3]. □

Having created a model for $B\,\mathrm{Conf}$ we next follow a very similar procedure to make a model for $B^2\,\mathrm{Conf}$.

Lemma 11.3.17. $B^2\,\mathrm{Conf}$ *can be identified with the space of points in a rectangle where points can fall off any of the four sides, and there are no two points with the same vertical coordinate.*

Proof. This is very similar to the proof of Lemma 11.3.13. See Figure 6 for a visualization. If C is the space described in Lemma 11.3.15, then we can identify $B(B\,\mathrm{Conf}) \simeq BC$, and we can think of a point of this as a collection of n configurations as in C, stacked vertically, which can either fall off the top or bottom. This is precisely the above description in the lemma statement. □

Lemma 11.3.18. $B^2\,\mathrm{Conf} \simeq S^2$.

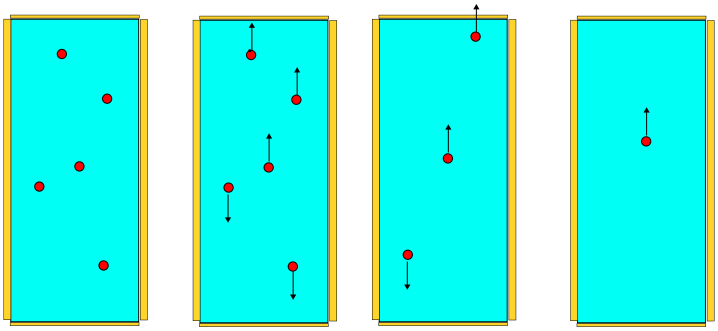

FIGURE 7. This is a picture depicting depicting the proof idea of Lemma 11.3.18.

Proof. Using Lemma 11.3.17 can be identified with the space of points in a rectangle where points can fall off any of the four sides and there are never 2 points with the same vertical coordinate. Then, by performing a similar "zooming in" homotopy to that described in Lemma 11.3.15, one can zoom in on the horizontal coordinate $1/2$ and arrange that there is at most a single horizontal coordinate occupied by a point in the configuration. See Figure 6. This implies there is at most 1 point in the configuration. The resulting space of ≤ 1 point in a disc is the 1-point compactification of the disc (the additional point corresponding to the empty configuration) which is precisely S^2. $\square$

Exercise 11.3.19. Let $\mathrm{Conf}_{n,\mathbb{R}^k}$ denote the configuration space of n distinct points in $\mathbb{R}^k$ and let $\mathrm{Conf}_{\mathbb{R}^k} := \cup_{n\geq 0}\mathrm{Conf}_{n,\mathbb{R}^k}$.

(1) Using arguments similar to the above, show more generally that $B^k\,\mathrm{Conf}_{\mathbb{R}^k} \simeq S^k$. Note that above we used Conf for $\mathrm{Conf}_{\mathbb{R}^2}$.
(2) Use the first part to show and the group completion theorem to show the stable homology of $\mathrm{Conf}_{\mathbb{R}^k}$ is identified with the homology of $\Omega_0^k S^k$.

(As usual, $\Omega_0^k S^k$ denotes the component of $\Omega^k S_k$ containing the basepoint which can be thought of as the constant map from S^k to S^k at the basepoint.)

The next exercise can be interpreted as saying that, as $k \to \infty$, the above map $\mathrm{Conf}_{n,\mathbb{R}^k} \to BS_n$ becomes closer and closer to being a homotopy equivalence.

Exercise 11.3.20. (1) With notation as in the previous exercise, show that the map $\mathrm{Conf}_{n,\mathbb{R}^k} \to BS_n$ (coming from the fact that unordered configuration space is a quotient of ordered confirmation space by the action of S_n permuting the marked points) induces an isomorphism on homology groups H_i for $i \leq k-2$.

(2) Additionally, show the map is also an isomorphism on homotopy groups π_i for $i \leq k-2$. To do this, identify the homotopy fiber of the map with $\mathrm{Pconf}_{n,\mathbb{R}^k}$. Show this is simply connected and its first $k-2$ homology groups vanish. Deduce from a version of the Hurewicz theorem that its first homotopy groups vanish. Conclude $\mathrm{Conf}_{n,\mathbb{R}^k} \to BS_n$ also induces isomorphisms on homotopy groups π_i for $i \leq k-2$.

Exercise 11.3.21. Assume $H_i(S_n;\mathbb{Z})$ stabilizes to for i fixed and $n \to \infty$ to some $H_i(S_\infty;\mathbb{Z})$. (For us, S_∞ does not have any meaning, and this is just notation for this stable value of these homology groups.) Also assume that $H_i(\Omega^k_0 S^k;\mathbb{Z})$ stabilizes as $k \to \infty$ to some $H_i(\Omega^\infty_0 \mathbb{S};\mathbb{Z})$.

Assuming these two stabilizations, and using the previous exercises, show that we have an isomorphism of $\mathbb{Z}$ modules $H_i(\Omega^\infty_0 \mathbb{S};\mathbb{Z}) \simeq H_i(S_\infty;\mathbb{Z})$.

Remark 11.3.22. In case you are familiar with the sphere spectrum, the above exercise says that the stable homology of symmetric groups is the same as the stable homology of the sphere spectrum. This is known as the Barratt–Priddy theorem or the Barratt–Priddy–Quillen theorem.

11.4. **A more advanced group completion theorem.** Although we didn't need it for configuration spaces, in the future, it will be useful to have the following more advanced version of the group completion theorem.

Definition 11.4.1. Of A_* is a graded ring, a multiplicatively closed subset $S \subset A_*$ satisfies the *left Ore condition* if the following two conditions hold:

(1) for every $s \in S$ and $a \in A_*$, there is $s' \in S, a' \in A_*$ with $as' = sa'$ in A_*.
(2) If $a \in A_*, t \in S$ satisfy $at = 0$ then there is $s \in S$ with $sa = 0$.

We also recall the notion of localization. This is just the standard notion of localization from an introductory commutative algebra course.

Definition 11.4.2. For A_* a graded ring, and N_* a graded A_* module, and $S \subset A_*$ a subset satisfying the left Ore condition, the localization $N_*[S^{-1}]$ (also sometimes notated $S^{-1}N_*$) is the the quotient of $S \times N_*$ by the equivalence relation generated by $(s,n) \sim (s',n')$ if there are $a, a' \in A_*$ so that $as = a's' \in S$ and $an = a'n'$.

For M a monoid, we use $\pi_0 M$ to denote its monoid of components. For example, if $M = \mathrm{Conf}$ then $\pi_0 M = \mathbb{N}$, with points of Conf_n lying in the component $n \in \mathbb{N}$. Since $\pi_0 M$ acts on $H_*(M;\mathbb{Z})$, we can localize $H_*(M;\mathbb{Z})$ at the multiplicative subset $\pi_0 M$, which we notate as $H_*(M;\mathbb{Z})[(\pi_0 M)^{-1}]$. (There is no higher algebra going on here, by localization, we just mean localization in the above sense.)

Theorem 11.4.3. *Let M be a graded monoid so that $\pi_0 M \subset H_*(M;\mathbb{Z})$ satisfies the left Ore condition. Then,*

$$H_*(M;\mathbb{Z})[(\pi_0 M)^{-1}] \simeq H_*(\Omega BM;\mathbb{Z}).$$

Exercise 11.4.4. Show Theorem 11.3.5 follows from Theorem 11.4.3. *Hint:* It may be helpful to know that $\pi_0(\Omega BM)$ is the group completion of $\pi_0 M$, meaning that it is the initial group to which $\pi_0 M$ maps.

12. BACKGROUND ON RACKS

Prior to this, we have been focusing on homological stability for Hurwitz spaces associated to c, where $c \subset G$ is a union of conjugacy classes in a group. It turns out that to define Hurwitz spaces, one only needs a notion of how c acts on itself by conjugation. There is no need to have the additional group structure of multiplication and inversion. This is because the action of the braid group on c^n is given by elementary generators sending $(g,h) \mapsto (h, h^{-1}gh)$. Racks are a type of algebraic object which encodes this information.

Definition 12.0.1. A *rack* is a set c with a binary operation $\triangleright : c \times c \to c$ such that the following two properties hold.

(1) For every $x \in c$, the map $\phi_x : c \to c, y \mapsto x \triangleright y$ is a bijection.
(2) For any $x, y, z \in c$, $x \triangleright (y \triangleright z) = (x \triangleright y) \triangleright (x \triangleright z)$.

Definition 12.0.2. Given a rack c, there is an action of c on itself, and a *component* of c is an orbit of c under this action.

Example 12.0.3. An important class of racks is those coming from unions of conjugacy classes in a group. Say $c \subset G$ is a union of conjugacy classes. In this case we take $x \triangleright y := x^{-1}yx$, for $x, y \in c$. On the right hand side, $x^{-1}yx$ is viewed as multiplication in the group. This is a rack because conjugation by $x \in c$ induces a bijection on c and the second condition holds because $x^{-1}(y^{-1}zy)x = (x^{-1}yx)^{-1}(x^{-1}zx)(x^{-1}yx)$.

If c generates G, then the components of c are precisely its conjugacy classes. The reason for this is that if x and y lie in the same conjugacy class, there is some element $g \in G$ so that $x = g^{-1}yg$, and then by writing g as a product of elements in c, we see x and y lie in the same component of c.

Exercise 12.0.4 (Important exercise)**.** Show that the following gives an equivalent definition of a rack. A rack is equivalent to the data of a set c with an action map $\triangleright : c \times c \to c, (a,b) \mapsto a \triangleright b$ such that for all $n \geq 1$ and all $1 \leq i \leq n-1$, the operation

$$\sigma_i : c^n \to c^n$$
$$(x_1, \ldots, x_{i-1}, x_i, x_{i+1}, x_{i+2}, \ldots, x_n) \mapsto (x_1, \ldots, x_{i-1}, x_{i+1}, x_{i+1} \triangleright x_i, x_{i+2}, \ldots, x_n)$$

defines an action of the braid group B_n, generated by $\sigma_1, \dots, \sigma_{n-1}$, on c^n.

Exercise 12.0.5. View the map $\triangleright : c \times c \to c$ as a map $\phi : c \to \mathrm{Hom}(c, c), x \mapsto \phi_x$. Show that Definition 12.0.1(2) is equivalent to the condition that ϕ is equivariant for the actions of c, where $z \in c$ acts on the source via $x \mapsto z \triangleright x$ and the action of $z \in c$ on the target is given by sending $\alpha(\bullet) \mapsto z \triangleright \alpha((z\triangleright)^{-1}\bullet)$, for $\alpha \in \mathrm{Hom}(c, c)$. Show that Definition 12.0.1(1) is the condition that these endomorphisms ϕ_x are automorphisms.

Exercise 12.0.6. Give an example of a rack with two elements that is not a union of conjugacy classes in a group.

Exercise 12.0.7. A *quandle* is a rack with the property that $x \triangleright x = x$. The example from Exercise 12.0.6 will not be a quandle. Give an example of a quandle with three elements that is not a union of conjugacy classes in a group. *Hint:* It turns out there is a unique such quandle, and it is called the Joyce quandle. To show it is not a subset of a group, note that if x, y are elements of a group $x^{-1}yx = y$ if and only if $y^{-1}xy = x$ if and only if x and y commute.

Notation 12.0.8. For c a rack and $g, g_1, \dots, g_n \in c$, we will use the notation

$$(g_1 \cdots g_n) \triangleright g := (g_n \triangleright (\cdots (g_2 \triangleright (g_1 \triangleright g)) \cdots)).$$

Note the reversal of order in the iteration of the g_i, which is compatible with the convention from Example 12.0.3.

Definition 12.0.9. The *reduced structure group* G_c^0 of a rack c is the subgroup of $\mathrm{Aut}(c)$, the automorphism group of the rack c, generated by the automorphisms $y \mapsto x \triangleright y$ for all $x \in c$. It follows from the definition of a rack that these are rack automorphisms.

Exercise 12.0.10. If G is a group and $c \subset G$ is a union of conjugacy classes in G. What is G_c^0? *Hint:* Even if $c = G$, G_c^0 may fail to equal G.

As mentioned, the natural setting to define Hurwitz spaces over the complex numbers is in the setting of racks which defines them as a finite étale cover of configuration space.

Definition 12.0.11. Let c be a rack. Upon identifying $\pi_1(\mathrm{Conf}_n) \simeq B_n$, we can identify finite étale covers of Conf_n with maps from B_n to finite sets. Define the *pointed Hurwitz scheme over* $\mathbb{C}$ to be the finite étale cover $\mathrm{Hur}_n^c \to \mathrm{Conf}_n$ corresponding to the map $B_n \to \mathrm{Aut}(c^n)$ associated to the rack, using the formulation in Exercise 12.0.4.

If c is a rack with components $c_1, \dots, c_v$ and $n_1 + \cdots + n_v = n$, let $c^{n_1, \dots, n_v} \subset c^n$ denote the subset such that there are n_i elements in component c_i. Then we define $\mathrm{Hur}^c_{n_1,\ldots,n_v} \to \mathrm{Conf}_{n_1,\ldots,n_v}$ to be the finite étale cover corresponding to the map $B_n \to \mathrm{Aut}(c^{n_1,\ldots,n_v})$. There is a subset $(c^n)^\circ \subset c^n$ parameterizing n tuples of elements in c whose actions together generate G^0_c. We let CHur^c_n denote the finite étale cover of Conf_n corresponding to the map $B_n \to \mathrm{Aut}((c^n)^\circ)$. Finally, we let $\mathrm{CHur}^c_{n_1,\ldots,n_v}$ denote the finite étale cover of $\mathrm{Conf}_{n_1,\ldots,n_v}$ corresponding to the map $B_n \to \mathrm{Aut}((c^n)^\circ \cap c^{n_1,\ldots,n_v})$. This is a union of components of $\mathrm{Hur}^c_{n_1,\ldots,n_v}$.

Example 12.0.12. Suppose $c \subset G$ is a union of v conjugacy classes $c_1 \cup \cdots \cup c_v$ in G. Then the complex points of Hur^c_n parameterize G covers of $\mathbb{C}$ with n branch points whose inertia lies in $c \subset G$ with trivialization of the G cover in some subset $[k, \infty]$ for some $k \in \mathbb{R}$ (where the choice of k is not part of the data of a point in the space). Similarly, CHur^c_n parameterizes covers as above whose source is connected. The complex points of $\mathrm{Hur}^c_{n_1,\ldots,n_v}$ parameterize covers with n_i branch points whose inertia lies in c_i, and $\mathrm{CHur}^c_{n_1,\ldots,n_v}$ parameterizes such covers whose source is connected.

Exercise 12.0.13. Let c be the Joyce quandle. Explicitly this has three elements x, y, z where $x \triangleright y = z, x \triangleright z = y$ and all other actions are trivial, i.e. $s \triangleright t = t$ if (s, t) is not (x, y) or (x, z). Write $c = c_1 \cup c_2$ with $c_1 = \{x\}, c_2 = \{y, z\}$ as a union of its two components. Show that points of $\mathrm{Hur}^c_{n_1,n_2}$ can be interpreted as double covers $f : X \to \mathbb{A}^1$ together with n_2 marked points in X at which f is unramified and n_1 marked points of X at which f is ramified. *Possible cryptic hint:* Interpret how a generator of the braid group acts on the Hurwitz space as how x, y, z act on each other when one branch point winds around another branch point. Another way to think about this is that this has a map to the two colored configuration space, and the fiber keeps track of when the element of the second component is y or z.

Remark 12.0.14. One reason to sometimes prefer working with groups is that there is a natural moduli interpretation of Hurwitz spaces associated to a group, which lets one define them over $\mathbb{Z}[\#G^{-1}]$. As described above, it is less clear whether there is such a natural moduli interpretation for Hurwitz spaces associated to racks. This could potentially be answered by the following question. This is in interesting open problem that I think should be worked out, and would have natural applications to arithmetic statistics.

The above exercise describes Hurwitz spaces associated to the Joyce quandle c as G^0_c covers together with a marking of some of the points in the fiber of the cover.

Question 12.0.15. If c is any rack, can Hur^c_n be interpreted as G^0_c covers together certain marked points in the source of the cover.

I conjecture the answer to this question is yes. It may actually be doable (I haven't thought carefully about it) for you to prove this conjecture is correct (or find a counterexample).

13. THE MAIN RESULTS ON THE STABLE HOMOLOGY

Having shown the homology of Hurwitz spaces stabilizes, we next aim to compute what it stabilizes to. The general statement is as follows.

Recall we use G_c^0 to denote the structure group of a rack c. So if $c \subset G$ is a union of conjugacy classes which generates G and G has trivial center, then $G_c^0 = G$.

In general, the homology of Hur_n^c will not stabilize in n since the number of its components may grow as n grows. Instead, we have the following result. Recall $\mathrm{CHur}^c_{n_1,\ldots,n_v}$ denotes the union of connected components of Hur_n^c with $n_1 + \cdots + n_v = n$ parameterizing connected covers with n_i points in c_i, for c_i the ith component of c. (Recall the components of a rack are the orbits of the rack acting on itself.)

Notation 13.0.1. Let c be a rack. We can identify the connected components of the pointed Hurwitz space Hur_n^c with orbits of the B_n action on c^n. Under this identification, we use $\alpha_x \in \pi_0 \mathrm{Hur}_1^c$ as notation for the component with a single branch point labeled x and so $\alpha_{x_1} \cdots \alpha_{x_n} \in \pi_0 \mathrm{Hur}_n^c$ corresponds to the B_n orbit of $(x_1, \ldots, x_n)$.

Theorem 13.0.2. *Let c be a rack with components $c = c_1 \cup \cdots \cup c_v$. There are constants I and J so that for $n_1 > Ii + J$, and any $x \in c_1$, the map $\mathrm{CHur}^c_{n_1,\ldots,n_v} \to \mathrm{CHur}^c_{n_1+1,n_2,\ldots,n_v}$ induced by multiplication by α_x (sending a cover to the same cover together with one more branch point labeled x near ∞) induces an isomorphism $H_i(\mathrm{CHur}^c_{n_1,\ldots,n_v}; \mathbb{Z}) \simeq H_i(\mathrm{CHur}^c_{n_1+1,n_2,\ldots,n_v}; \mathbb{Z})$ when $n_1 > Ii + J$.*

We repeat this in Theorem 16.0.1 and prove it there.

We can also compute the stable value of this homology group. For this, the reader may wish to recall the notion of the reduced structure group G_c^0 from Definition 12.0.9.

Theorem 13.0.3. *Let c be a rack with components $c = c_1 \cup \cdots \cup c_v$. There are constants I and J so that For $Z \subset \mathrm{CHur}^c_{n_1,\ldots,n_v}$ a component, the map $\mathrm{CHur}^c_{n_1,\ldots,n_v} \to \mathrm{Conf}_{n_1,\ldots,n_v}$ (which comes from the definition of CHur^c, but in the case c is a subset of a group corresponds to sending a cover to its branch locus) induces an isomorphism $H_i(Z; \mathbb{Z}[1/\#G_c^0]) \simeq H_i(\mathrm{Conf}_{n_1,\ldots,n_v}; \mathbb{Z}[1/\#G_c^0])$ when $n_1, \ldots, n_v > Ii + J$.*

The above result on the stable homology requires that all n_i are sufficiently large, whereas the earlier theorem tells us the homology stabilizes once any one of the n_i are sufficiently large. This leaves open the question of what the stable value of the homology is when only n_1 is large. We give an answer in Theorem 18.0.9 which then implies Theorem 13.0.3.

14. STABILITY OF A QUOTIENT

The starting point for our proof of homological stability is to show a certain quotient of Hurwitz space has homology which stabilizes.

To state the result precisely, we need some notation.

Notation 14.0.1. Let c be a rack. Let $\mathrm{Hur}^c := \cup_n \mathrm{Hur}^c_n$ and let $A_c := C_*(\mathrm{Hur}^c)$, where C_* denotes the chains functor defined in Definition 10.4.9. For $\alpha \in \pi_0 A_c$ we write A_c/α to denote the object in the derived infinity category $A_c \xrightarrow{\times\alpha} A_c$ (the mapping cone of α). Fix a rack c and choose a total order $x_1, \ldots, x_{|c|}$ of the elements of c. Define the A_c bimodule

$$A_c/(\alpha_c^{\mathrm{ord}_c(c)}) := A_c/\alpha_{x_1}^{\mathrm{ord}_c(x_1)} \otimes_{A_c} A_c/\alpha_{x_2}^{\mathrm{ord}_c(x_2)} \otimes_{A_c} \cdots \otimes_{A_c} A_c/\alpha_{x_{|c|}}^{\mathrm{ord}_c(x_{|c|})}.$$

(Above we used a notation for tensor products. This is a natural monoidal operation in the derived infinity category. If you'd like one way to define the tensor product is as the bar construction from (11.1) we discussed earlier.)

Remark 14.0.2. The reason we are quotienting by the $\mathrm{ord}_c(x_i)$ power of x_i is so that $x_i^{\mathrm{ord}_c(x_i)} y = y x_i^{\mathrm{ord}_c(x_i)}$ which will allow us to view the quotient as a bimodule instead of just a 1-sided module.

Definition 14.0.3. If X is a graded space or graded chain complex, we say X is *bounded in a linear range* if there are constants $I \geq 0$ and J so that $\pi_i(X_n) = 0$ for for $n > Ii + J$. (Recall that when X is a chain complex, $\pi_i(X_n) = H_i(X_n; \mathbb{Z})$ by convention, see Fact 10.4.10.)

Theorem 14.0.4. *For c a rack, and $A_c := C_*(\mathrm{Hur}^c)$, as defined above, then $A_c/(\alpha_c^{\mathrm{ord}_c(c)})$ is bounded in a linear range.*

We will prove a similar, but slightly simpler result in Theorem 17.0.6. The same idea used to prove that also works to prove this and we leave it as an exercise in Exercise 17.0.10. We won't further explain the details of this statement, but the interested reader can see [LL25a, Theorem 3.2.6].

Exercise 14.0.5. Using that there is a map $\sigma : C_*(\mathrm{Conf}_n) \to C_*(\mathrm{Conf}_{n+1})$ which induces an isomorphism on ith homology once $\frac{n-1}{2} \geq i$, show that $C_*(\mathrm{Conf})/\sigma$ is bounded in a linear range. *Hint:* In this hint, we use the notation ΣX to denote the shift of X by 1 (so the ith homology of ΣX is the $i-1$st homology of X). In general, $\Sigma(A/\sigma)$ has a two step filtration whose associated graded pieces are $\ker\sigma$ and $\Sigma(\operatorname{coker}\sigma)$ This is coming from the fact that there is an exact sequence in the derived category

$$0 \longrightarrow \ker\sigma \longrightarrow \Sigma(A/\sigma) \longrightarrow \Sigma\operatorname{coker}\sigma \longrightarrow 0 \tag{14.1}$$

which comes from the exact sequence of chain complexes

$$0 \longrightarrow \ker\sigma \longrightarrow A \xrightarrow{\sigma} A \longrightarrow \operatorname{coker}\sigma \longrightarrow 0 \tag{14.2}$$

15. SHOWING THE HOMOLOGY STABILIZES FOR TRANSPOSITIONS IN S_3

Let's first explain how to prove this in the special case that $G = S_3$ and $c = \{x, y, z\} \subset S_3$ is the set of transpositions. A lot of things are simpler in this case, though many of the essential aspects of the proof are already present.

We use α_x to denote the element of $\pi_0 \operatorname{Hur}^c_1$ corresponding to the element x, and similarly for y and z. The starting point is Theorem 14.0.4 which says that the homology of $C_*(\operatorname{Hur}^c)/(\alpha_x^2, \alpha_y^2, \alpha_z^2)$ is bounded in a linear range. This means concretely that its ith homology vanishes for n large relative to i (i.e. when $n > Ii + J$). The idea will then be to remove each of the three elements we are quotienting by, and compute the stable homology at each step. The fact that $C_*(\operatorname{Hur}^c)/(\alpha_x^2, \alpha_y^2, \alpha_z^2)$ is bounded in a linear range implies that α_z^2 induces isomorphisms $C_*(\operatorname{Hur}^c)/(\alpha_x^2, \alpha_y^2) \to C_*(\operatorname{Hur}^c)/(\alpha_x^2, \alpha_y^2)$ in a linear range.

The idea will be to compute the stable value of this cohomology. There is an obvious source of the cohomology of $C_*(\operatorname{Hur}^c)/(\alpha_x^2, \alpha_y^2)$. Namely, there is a subrack of c generated by the singleton z. We'll refer to this subrack as $z \subset c$. There is a map $\operatorname{Hur}^z \to \operatorname{Hur}^c$ sending a cover with branch labels all in z to the same cover, since $z \in c$. It will be convenient to have a retraction, i.e. a map in the other direction. This doesn't exist yet, since we don't know where to send labels x or y. We can rectify this by adding a disjoint basepoint.

Notation 15.0.1. For X a space, we use X_+ to denote the space $X \coprod *$, the pointed space which is the disjoint union of X and a basepoint $*$.

We then obtain a map $\operatorname{Hur}^c_+ \to \operatorname{Hur}^z_+$ which sends a cover labeled by only zs to itself and it sends every cover either with an x or y to the basepoint. Using $\widetilde{C}_*$ to denote reduced chains, we have $\widetilde{C}_*(\operatorname{Hur}^c_+) \simeq C_*(\operatorname{Hur}^c)$, so it will be enough to understand reduced chains of these Hurwitz spaces with disjoint basepoints.

In order to set up our next calculation, we need the notion of localization.

Definition 15.0.2. Let R be a ring object in the derived infinity category of $\mathbb{Z}$ modules. Fix $r \in \pi_0 R$ and let $\mathrm{Ch}(R[r^{-1}]) \subset \mathrm{Ch}(R)$ denote the full subcategory consisting of R modules M so that r acts by an equivalence on M as above. The inclusion $\mathrm{Ch}(R[r^{-1}]) \subset \mathrm{Ch}(R)$ has a left adjoint called *localization* $\mathrm{Ch}(R) \subset \mathrm{Ch}(R[r^{-1}])$ which sends an R module M to the module we denote by $M[r^{-1}]$. (We will not prove this adjoint exists, but just take it as a fact. However, you may wish to see the next exercise.)

Example 15.0.3. Consider an associative ring R an element $r \in \pi_0 R$ and a module M as above. Suppose that r is $\mathbb{E}_2$ central for R, meaning that there is an R-bimodule maps $R \to R$ sending 1 to r. (Loosely, you can think of this as saying that r suitably commutes with left and right multiplication by elements of R.) In the context of this course, we will only use this in the case that R is a Hurwitz space, and $r \in \pi_0 R$ is a central element of the discrete monoid $\pi_0 R$, in which case one can construct a homotopy between left and right multiplication by r on R by "moving r around the other element," (see [LL24, Lemma 4.3.1]) implying r is $\mathbb{E}_2$ central. Then, the localization $M[r^{-1}]$ can be identified with the colimit $M_\infty := \mathrm{colim}(M \to M \to M \to \cdots)$ where the maps are all multiplication by r. The fact that r is $\mathbb{E}_2$ central ensures that the above maps $M \to M$ by multiplication by r are all module maps, because we can use $\mathbb{E}_2$ centrality to identify $s \cdot r$ with $r \cdot s$.

Let us now give a sketch for why $M_\infty = M[r^{-1}]$, we want to show that multiplication by r induces an equivalence on M_∞ and that if N is any module on which r acts invertibly $\mathrm{Hom}_R(M, N) \to \mathrm{Hom}_R(M_\infty, N)$ is an equivalence. The first statement can be deduced using the Yoneda to define an inverse map.

If we are given some map $X \to M_\infty$, we can view it as a map $f : X \to M$ in, say the ith coordinate for i sufficiently large, which yields a map $rf : X \to M$ in the $i+1$ coordinate. We then send this map to the map $f : X \to M$ in the $i+1$ coordinate, which is r times our original map. To conclude, we want to show $\mathrm{Hom}_R(M, N) \to \mathrm{Hom}_R(M_\infty, N)$ is an equivalence. The inverse to this equivalence is simply given by sending $M_\infty \to N$ to the composite $M \to M_\infty \to N$. The tricky part is actually defining the above map. However, we can do so as follows: given $M \to N$, we get a map $M_\infty \to N_\infty$, and then we can identify $N_\infty \simeq N$ using that N_∞ is a colimit of equivalences $r : N \to N$.

Remark 15.0.4. In what follows, we will often localize some hurwitz space at α_x for $x \in c$. This is the same as localizing at $\alpha_x^{\mathrm{ord}_c(x)}$, which is $\mathbb{E}_2$ central, and so the localization may be computed explicitly using the above example.

The key calculation we now need to verify is the following:

Proposition 15.0.5. *The map* $\mathrm{Hur}^c_+ \to \mathrm{Hur}^z_+$ *induces an equivalence* $\widetilde{C}_*(\mathrm{Hur}^c_+)/(\alpha_x^2, \alpha_y^2)[\alpha_z^{-1}] \to \widetilde{C}_*(\mathrm{Hur}^z_+)/(\alpha_x^2, \alpha_y^2)[\alpha_z^{-1}]$.

In other words, we can interpret $C_*(\mathrm{Hur}^c)/(\alpha_x^2, \alpha_y^2)[\alpha_z^{-1}]$ as the part of the homology of $C_*(\mathrm{Hur}^c)/(\alpha_x^2, \alpha_y^2)$ which is stable with respect to multiplication by α_z. Then, the result says that the α_z stable homology in $C_*(\mathrm{Hur}^c)/(\alpha_x^2, \alpha_y^2)$ completely comes from the α_z stable homology in $C_*(\mathrm{Hur}^z)/(\alpha_x^2, \alpha_y^2)$.

In what follows we use the general notion of bar constructions associated to modules over an $\mathbb{E}_1$ algebra from Remark 11.2.6.

Remark 15.0.6. A ring R in the derived category ∞ category is connective if $\pi_i R = 0$ for $i < 0$. (Recall we are using $\pi_i R$ as synonymous notation for $H_i(R)$.)

Lemma 15.0.7. *Suppose* $R \to S$ *is a map of connective associative* ($\mathbb{E}_1$) *rings which induces a bijection* $\pi_0 R \to \pi_0 S$ *and an equivalence* $\pi_0 R \otimes_R \pi_0 R \to \pi_0 S \otimes_S \pi_0 S$. *Then* $R \to S$ *is an equivalence.*

Proof. The proof is a standard computation in descent. Namely, consider the map $R \to \pi_0 R$.

Lemma 15.0.8. *For* R *an* $\mathbb{E}_1$ *ring, there is an equivalence between* R *and the limit of the complex*

$$\pi_0 R \to \pi_0 R \otimes_R \pi_0 R \to \pi_0 R \otimes_R \pi_0 R \otimes_R \pi_0 R \to \cdots,$$

whose nth term is $\pi_0 R^{\otimes_R n+1}$, *and there are really* $n+1$ *maps from* $\pi_0 R^{\otimes_R n} \to \pi_0 R^{\otimes_R n+1}$ *corresponding to including into the complement of the ith factor.*

Remark 15.0.9. Intuitively, you can think of this as a statement in descent, where $\pi_0 R$ is thought of as a cover of R. Then, the n-fold tensor product above corresponds to an n-fold fiber product. In usual descent for schemes, one can check an equivalence between two schemes if one checks there is an equivalence on the 1 and 2-fold covers. For stacks, one can check there is an equivalence if one checks the 1, 2, and 3 fold covers. In this setting, it is enough to check on n-fold covers for all n.

Proof. Let I denote the fiber $I \to R \to \pi_0 R$. Note that $\pi_0 I = 0$ using the sequence

$$\pi_1(\pi_0 R) \longrightarrow \pi_0(I) \longrightarrow \pi_0(R) \longrightarrow \pi_0(\pi_0 R). \tag{15.1}$$

Since $\pi_0(\pi_0 R) = \pi_0 R$, the last map is an equivalence. Since $\pi_0 R$ is concentrated in homological degree 0, $\pi_1(\pi_0 R) = 0$ and so we conclude $\pi_0 I = 0$. This implies that I is concentrated in homological degree at least 1 and so

$I^{\otimes_R k}$ is concentrated in homological degree k. Therefore, it suffices to show that $I^{\otimes_R k}$ is equivalent to the limit

$$R \to \pi_0 R \to \pi_0 R \otimes_R \pi_0 R \to \cdots \to \pi_0 R^{\otimes_R k-1} \to \pi_0 R^{\otimes_R k}$$

where, again, there are i maps from the $i-1$th tensor power to the i tensor power coming from including away from the jth coordinate for $1 \leq j \leq i$.

The case $k = 1$ is the definition of I, so let's next look a $k = 1$. Then, we can reformulate the colimit we are trying to compute as the colimit

$$\begin{tikzcd} R \arrow[r] \arrow[d] & \pi_0 R \arrow[d] \\ \pi_0 R \arrow[r] & \pi_0 R \otimes_R \pi_0 R \end{tikzcd} \tag{15.2}$$

or equivalently the colimit of the diagram

$$\begin{tikzcd} R \otimes_R R \arrow[r] \arrow[d] & R \otimes_R \pi_0 R \arrow[d] \\ \pi_0 R \otimes_R R \arrow[r] & \pi_0 R \otimes_R \pi_0 R. \end{tikzcd} \tag{15.3}$$

The fiber of the left hand vertical map is $I \otimes_R R$ and the fiber of the right hand vertical map is $I \otimes_R \pi_0 R$, so the colimit of the above diagram is the same as the fiber of $I \otimes_R R \to I \otimes_R \pi_0 R$, which is precisely $I \otimes_R I = I^{\otimes_R 2}$. The proof in the general case is similar, and can be demonstrated by induction. For example, let us just run through the $k = 3$ case. Here, the colimit we wish to show is $I^{\otimes_R 3}$ is the colimit of the cube with front and back faces

$$\begin{tikzcd} R \otimes_R R \otimes_R R \arrow[r] \arrow[d] & \pi_0 R \otimes_R R \otimes_R R \arrow[d] & R \otimes_R R \otimes_R \pi_0 R \arrow[r] \arrow[d] & \pi_0 R \otimes_R R \otimes_R \pi_0 R \arrow[d] \\ R \otimes_R \pi_0 R \otimes_R R \arrow[r] & \pi_0 R \otimes_R \pi_0 R \otimes_R R & R \otimes_R \pi_0 R \otimes_R \pi_0 R \arrow[r] & \pi_0 R \otimes_R \pi_0 R \otimes_R \pi_0 R \end{tikzcd} \tag{15.4}$$

We showed above that the colimit of the front face is $I \otimes_R I \otimes_R R$ while the colimit of the back face is $I \otimes_R I \otimes_R \pi_0 R$. So, the total colimit is the fiber of the map $I \otimes_R I \otimes_R R \to I \otimes_R I \otimes_R \pi_0 R$ which is $I \otimes_R I \otimes_R I$. The general case for k larger is similar and can be proven by induction. □

Using Lemma 15.0.8, in order to identify R and S, it suffices to identify $\pi_0 R \simeq \pi_0 S$ and $\pi_0 R^{\otimes_R n+1} \simeq \pi_0 S^{\otimes_S n+1}$. We are assuming the case $n = 1$ holds, and from this we can deduce the case $n > 1$ by iteratively applying the $n = 1$ case. For example, to see the case $n = 2$, using the $n = 1$ case twice we see

$$\begin{aligned}(\pi_0 R \otimes_R \pi_0 R) \otimes_R \pi_0 R &\simeq (\pi_0 R \otimes_S \pi_0 R) \otimes_R \pi_0 R \\ &\simeq \pi_0 R \otimes_S (\pi_0 R \otimes_R \pi_0 R) \\ &\simeq \pi_0 R \otimes_S (\pi_0 R \otimes_S \pi_0 R) \\ &\simeq \pi_0 S \otimes_S \pi_0 S \otimes_S \pi_0 S.\end{aligned}$$

□

A similar argument (whose proof we omit) also allows us to prove a slightly upgraded version of the above lemma which will help us deal with quotients:

Lemma 15.0.10. *Suppose $R \to S$ is a map of connective associative rings which is surjective on π_0, suppose $I_S \subset \pi_0 S$ is a two-sided ideal and $I_R \subset \pi_0 R$ is the preimage of I_S. Let M be a bounded below right R module so that I_R acts nilpotently on $\pi_i M$ for each i. If*

$$\pi_0 R/I_R \otimes_R \pi_0 R/I_R \to \pi_0 S/I_S \otimes_S \pi_0 S/I_S$$

is an equivalence then $M \simeq M \otimes_R S$ is also an equivalence.

Proof of Proposition 15.0.5. We apply Lemma 15.0.10 with $R = C_*(\mathrm{Hur}^c)[\alpha_z^{-1}], S = C_*(\mathrm{Hur}^z)[\alpha_z^{-1}]$, $I_S = 0$ and $I_R = (\alpha_x, \alpha_y)$, and $M = C_*(\mathrm{Hur}^c)/(\alpha_x^2, \alpha_y^2)[\alpha_z^{-1}]$. First, note that α_x and α_y act nilpotently on $\pi_i M$ for all i because M involves a quotient by α_x^2 and α_y^2. Since we can identify

$$C_*(\mathrm{Hur}^z)/(\alpha_x^2, \alpha_y^2)[\alpha_z^{-1}] \simeq C_*(\mathrm{Hur}^c)/(\alpha_x^2, \alpha_y^2)[\alpha_z^{-1}] \otimes_{C_*(\mathrm{Hur}^c)} C_*(\mathrm{Hur}^z),$$

it remains only to verify the hypothesis of the lemma. That is, we only need to show

$$\pi_0 R/I_R \otimes_R \pi_0 R/I_R \to \pi_0 S/I_S \otimes_S \pi_0 S/I_S$$

is an equivalence. Explicitly, $I_S = 0$ and $I_R = (\alpha_x, \alpha_y)$, and hence $\pi_0 S/I_S \simeq C_*(\mathbb{Z}) \simeq \widetilde{C}_*(\mathbb{Z})$ and we also have $\pi_0 R/I_R \simeq \widetilde{C}_*(\mathbb{Z}_+)$, where the $\mathbb{Z}$ keeps track of the number of elements labeled z and the $+$ is the component of the basepoint. We use $\widetilde{C}$ for the reduced chains functor (the cofiber of $C_*(\mathrm{pt}) \to C_*(X)$). In particular, $\widetilde{C}_*(X_+) = C_*(X)$. Spelling this out, we want to show there the map

$$\widetilde{C}_*(\mathbb{Z}_+) \otimes_{\widetilde{C}_*(\mathrm{Hur}^c_+)} \widetilde{C}_*(\mathbb{Z}_+) \to \widetilde{C}_*(\mathbb{Z}_+) \otimes_{\widetilde{C}_*(\mathrm{Hur}^z_+)} \widetilde{C}_*(\mathbb{Z}_+)$$

is a homology equivalence. (Technically, a direct application of Lemma 15.0.10 would have us tensor over $R = C_*(\mathrm{Hur}^z)[\alpha_z^{-1}]$, but since α_z^{-1} acts invertibly on $\mathbb{Z}_+$, we can identify $\widetilde{C}_*(\mathbb{Z}_+) \otimes_{\widetilde{C}_*(\mathrm{Hur}^c_+)[\alpha_z^{-1}]} \widetilde{C}_*(\mathbb{Z}_+) \simeq \widetilde{C}_*(\mathbb{Z}_+) \otimes_{\widetilde{C}_*(\mathrm{Hur}^c_+)} \widetilde{C}_*(\mathbb{Z}_+)$.) The above map is identified with a map

$$\widetilde{C}_*(\mathbb{Z}_+ \otimes_{\mathrm{Hur}^c_+} \mathbb{Z}_+) \to \widetilde{C}_*(\mathbb{Z}_+ \otimes_{\mathrm{Hur}^z_+} \mathbb{Z}_+),$$

where the tensor products above are taken in pointed spaces. Moreover, it suffices to do this on the level of spaces, before taking chains. That is, it suffices to show

$$\mathbb{Z}_+ \otimes_{\mathrm{Hur}^c_+} \mathbb{Z}_+ \to \mathbb{Z}_+ \otimes_{\mathrm{Hur}^z_+} \mathbb{Z}_+$$

is a homology equivalence , which is the content of Proposition 15.0.11. □

Proposition 15.0.11. *We have that*

$$\mathbb{Z}_+ \otimes_{\mathrm{Hur}^c_+} \mathbb{Z}_+ \to \mathbb{Z}_+ \otimes_{\mathrm{Hur}^z_+} \mathbb{Z}_+$$

is a homology equivalence.

Proof. We can identify $\mathbb{Z}_+ \otimes_{\mathrm{Hur}^c_+} \mathbb{Z}_+$ explicitly with configurations in a rectangle, where the points are labeled by elements of c, and so that they can fall off the left or the right. If a z falls of either side, it increments the number in the corresponding $\mathbb{Z}$ by 1, and if x or y falls of the configuration is sent to the base point. Moreover, when a point labeled by α passes above β it gets conjugated by β. Moreover, via a scanning argument as in Lemma 11.3.15, we can assume there is at most a single point with any given vertical coordinate.

The map $\mathrm{Hur}^c_+ \to \mathrm{Hur}^z_+$ has a section $\mathrm{Hur}^z_+ \to \mathrm{Hur}^c_+$ given by sending a cover with branching only labeled by elements of z to itself, by viewing $\{z\}$ as a subset of c. Under the above identification, the complement consists of consists of the subspace of $\mathbb{Z}_+ \otimes_{\mathrm{Hur}^c_+} \mathbb{Z}_+$ corresponding to covers with some branch point labeled by either x or y. To conclude, it is enough to produce a nullhomotopy of this subspace. This nullhomotopy is constant at the basepoint, and, on all other configurations, is given by locating the lowest point labeled x or y and pulling a point labeled z out from the left directly above the point labeled x or y and moving it across the rectangle to the right at constant speed. Assuming this map is continuous, it will indeed be a nullhomotopy, because when the z passes above the single x or y below it, it will get conjugated to y or x, and so when it then hits the right hand side, it will send the configuration to the basepoint.

In order to check the claimed map is continuous, we need to verify that small perturbations of the configuration will result in the same map. First, any point moves (and possibly falls off the rectangle), apart from the lowest x or y, our homotopy would have been the same, pulling the z above the lowest x or y, which we are assuming does not move. Therefore, the only thing to worry about is whether the lowest x or y moves and falls off. However, in this case, our configuration would be sent to the basepoint, which is also compatible with our nullhomotopy.

Remark 15.0.12. There is actually an important technical point above that needs to be addressed. If one simply moves the point z "directly above" the lowest point labeled x or y, we have not exactly specified the height at which we move it. We need to be more precise here. One way to try to fix this is to choose a specific ε and if the lowest x or y is at height h, move the z at height $h + \frac{\varepsilon}{2}$. However, it could be that there is some other point at height $h + \varepsilon/4$, and then this wouldn't be the homotopy we wanted, since we want there to be no point between the z we are moving and the height h of the lowest x or y.

To get around this issue, one can choose a specific $\varepsilon > 0$ and work with the subset of the Hurwitz space where any two vertical points have spacing at least ε. The usual Hurwitz space is the union over $\varepsilon > 0$ of these ε Hurwitz spaces. If one then moves z at height $h + \varepsilon/2$ as above, it lies in an $\varepsilon/2$ spaced Hurwitz space, and will be directly above the lowest x or y because we are assuming vertical heights of points in the initial configuration have spacing ε. It is a general fact that it is enough to produce homotopies on these ε spaced Hurwitz spaces; see [LL24, Lemma A.4.6] for the technical details.

$\square$

Exercise 15.0.13. Suppose n is an odd number and c is the conjugacy class of order 2 elements in D_{2n}, the dihedral group of order $2n$. Choose any $z \in c$ and define $\mathbb{Z}_+$ as a module for Hur^c_+ by letting z act by incrementing the integer (it sends $n \mapsto n+1$ and fixes the basepoint) and any $x \in c - z$ act by sending everything to the base point (it sends any $n \in \mathbb{Z}_+$ to the basepoint). Argue similarly to the above that we have an equivalence $\mathbb{Z}_+ \otimes_{\mathrm{Hur}^c_+} \mathbb{Z}_+ \to \mathbb{Z}_+ \otimes_{\mathrm{Hur}^z_+} \mathbb{Z}_+$.

An important part of this exercise is thinking through the continuity of the homotopy.

The next exercise elucidates an important simplification in the S_3 case that does not work in general.

Exercise 15.0.14. Suppose now that we take c to be the conjugacy class of transpositions in S_4. Why does the above homotopy no longer work to identify $\mathbb{Z}_+ \otimes_{\mathrm{Hur}^c_+} \mathbb{Z}_+$ and $\mathbb{Z}_+ \otimes_{Hur^z_+} \mathbb{Z}_+$?

Exercise 15.0.15. In the proof of Proposition 15.0.11, we defined a nullhomotopy where we find the lowest element not labeled z and move a z above it. One could imagine trying to instead define a homotopy where we move the lowest element not labeled z to the left (instead of introducing a new point that we move above it). When this point hits the left, the configuration will become the base point. Explain why this alternate proposed homotopy is not continuous. *Hint:* Consider a sequence of configuration where the lowest point not labeled z tends to the right boundary.

So now that we've understood the stable homology of $C_*(\mathrm{Hur}^c_+)/(\alpha_x^2, \alpha_y^2)$, we'd like to remove the other two elements we're quotienting by. However, we run into an issue. Namely, we'd like to remove the quotient by α_y^2, just as we removed α_x^2, but we no longer have that the stable homology of $C_*(\mathrm{Hur}^c)/(\alpha_x^2, \alpha_y^2)$ is 0. Rather, it is that of $C_*(\mathrm{Hur}^z)/(\alpha_x^2, \alpha_y^2)$. If instead it were 0, we would be all set to try and remove the next element of the quotient. But thinking about things some more, we realize that $C_*(\mathrm{Hur}^z)/(\alpha_x^2, \alpha_y^2)$ wasn't something we were very interested in anyway. If we think of elements of Hur^c as degree 3 covers, Hur^z corresponds to the components parameterizing disconnected degree 3 covers, which are the disjoint union of of a degree 1 cover and a degree 2 cover where all the branching of the degree 2 cover is labeled by z. However, it is natural to focus on the connected covers. That is, instead of considering Hur^c, we can consider the union of components $\mathrm{CHur}^c \subset \mathrm{Hur}^c$ parameterizing connected covers. Then, the same argument as in Proposition 15.0.5 shows:

Proposition 15.0.16. $\widetilde{C}_*(\mathrm{CHur}^c_+)/(\alpha_x^2, \alpha_y^2)[\alpha_z^{-1}] = 0$.

Remark 15.0.17. Instead of repeating the argument, we can deduce Proposition 15.0.16 directly from Proposition 15.0.5 by filtering the two sides by their components parameterizing connected degree 3 covers. Since there are no connected degree 3 covers associated to Hur^z, this will show the stable homology of $\widetilde{C}_*(\mathrm{CHur}^c_+)/(\alpha_x^2, \alpha_y^2)[\alpha_z^{-1}]$ vanishes.

Why is the above helpful for computing the homology of $\widetilde{C}_*(\mathrm{CHur}^c_+)/(\alpha_x^2, \alpha_y^2)$? We have only computed it after inverting α_z. The next lemma says that in order to show the homology of $\widetilde{C}_*(\mathrm{CHur}^c_+)/(\alpha_x^2, \alpha_y^2)$ is bounded in a linear range, it is enough to show the quotient by α_z^2 is bounded in a linear range and it vanishes after inverting α_z^2 (which is the same as inverting α_z).

For the next lemma statement, it will be useful to recall $\alpha_x^{\mathrm{ord}_c(x)}$ is $\mathbb{E}_2$ central in $C_*(\mathrm{Hur}^c)$ for any $x \in c$ as in Example 15.0.3.

Lemma 15.0.18. *Suppose that X is a graded associative ring in the derived infinity category and $v : X \to X$ is a $\mathbb{E}_2$ central map of degree $|v|$ in the sense that it sends the jth graded part to the $(j + |v|)$th graded part for some $|v| \in \mathbb{N}$. Suppose X/v is bounded in a linear range and $X[v^{-1}] = 0$. Then X is bounded in a linear range.*

Proof. We have a sequence of homotopy groups that is exact in the middle

$$\pi_{i+1}((X/v)_{j+|v|}) \to \pi_i X_j \xrightarrow{v} \pi_i X_{j+|v|}.$$

Because X/v is bounded in a linear range, say $\pi_i((X/v)_j) = 0$ for $j > f(i)$ for some linear function f. Then, the multiplication by v is injective if $j + |v| > f(i+1)$. Using Example 15.0.3, we can identify $\operatorname{colim}_n \pi_i X_{j+n|v|} = \pi_i X_j[v^{-1}]$. So, in such degrees as above, the map $\pi_i X_j \to \operatorname{colim}_n \pi_i X_{j+n|v|} = \pi_i X_j[v^{-1}] = 0$ is injective. Hence, $\pi_i X_j = 0$ when $j + |v| > f(i+1)$, so X is again bounded in a linear range. □

Hence, from the above, we actually obtain that $\widetilde{C}_*(\mathrm{CHur}^c_+)/(\alpha_x^2, \alpha_y^2)$ is bounded in a linear range.

Now we are all set to remove another element from the quotient. Namely, one can check localizations and chains commute with quotients (see [LL24, Lemma 3.4.4] for a proof of the fact that localizations commute with quotients) and so $\widetilde{C}_*(\mathrm{CHur}^c_+)/(\alpha_x^2, \alpha_y^2)[\alpha_z^{-1}] \simeq \left(\widetilde{C}_*(\mathrm{CHur}^c_+)/(\alpha_x^2)[\alpha_z^{-1}]\right)/(\alpha_y^2))$. α_y^2 induces an isomorphism $\widetilde{C}_*(\mathrm{CHur}^c_+)/(\alpha_x^2)[\alpha_z^{-1}] \to \widetilde{C}_*(\mathrm{CHur}^c_+)/(\alpha_x^2)[\alpha_z^{-1}]$. If we knew the stable value of this cohomology were 0, we would obtain that this is bounded in a linear range. One can argue this stable value is 0 by using that y and z both act invertibly, but they also generate x, which acts nilpotently. This idea is formalized in the next lemma.

Lemma 15.0.19. *We have* $\widetilde{C}_*(\mathrm{CHur}^c_+)/(\alpha_x^2)$ *is bounded in a linear range.*

Proof. Using Lemma 15.0.18, it suffices to show $\widetilde{C}_*(\mathrm{CHur}^c_+)/(\alpha_x^2)[\alpha_y^{-1}] \simeq 0$. Moreover, since the above map is an equivalence if and only if it is an equivalence both after quotienting by α_z^2 and inverting α_z^2 (see [LL24, Lemma 3.3.4]) it is enough to check $\widetilde{C}_*(\mathrm{CHur}^c_+)/(\alpha_x^2, \alpha_z^2)[\alpha_y^{-1}]$ and $\widetilde{C}_*(\mathrm{CHur}^c_+)/(\alpha_x^2)[\alpha_y^{-1}, \alpha_z^{-1}]$ are 0.

The first vanishing follows from Proposition 15.0.16 (by switching y and z).

The second vanishing follows from the simple observation that $x = y^{-1}zy$, which means $\alpha_y^{-1}\alpha_z\alpha_y = \alpha_x$. We are quotienting by α_x^2, α_x will act nilpotently on $\widetilde{C}_*(\mathrm{CHur}^c_+)/(\alpha_x^2)[\alpha_y^{-1}, \alpha_z^{-1}]$. (See [LL24, Lemma 3.5.2] for a proof of this believable fact.) On the other hand, since we are inverting both α_y and α_z, $\alpha_y^{-1}\alpha_z\alpha_y$ will act invertibly on this. Therefore, we have some element acting both nilpotently and invertibly on $\widetilde{C}_*(\mathrm{CHur}^c_+)/(\alpha_x^2)[\alpha_y^{-1}, \alpha_z^{-1}]$, which implies it must be 0. □

We can rephrase the above result as follows:

Corollary 15.0.20. *Let $c \subset S_3$ denote the set of transpositions. There are constants I, J so that the map α_x induces an isomorphism $H_i(\mathrm{CHur}^c_n) \to H_i(\mathrm{CHur}^c_{n+1})$ when $n > Ii + J$.*

Proof. All in all, Lemma 15.0.19 implies the multiplication by α_x^2 induces isomorphisms $\widetilde{C}_*(\mathrm{CHur}^c_+) \to \widetilde{C}_*(\mathrm{CHur}^c_+)$ in some linear range, or equivalently α_x induces isomorphisms on homology in a linear range, which is the statement we are trying to show. $\square$

Exercise 15.0.21 (Involved exercise)**.** For this exercise, you may assume Exercise 15.0.13. Let ℓ be an odd prime. Let $c \subset D_{2\ell}$ denote the subset of order 2 elements (which can be thought of as acting by reflections on an ℓ-gon). Show that CHur^c satisfies homological stability in the sense above that there are constants I and J (depending on ℓ) so that for any $x \in c$, the map $H_i(\mathrm{CHur}^c_n) \to H_i(\mathrm{CHur}^c_{n+1})$ induced by α_x is an isomorphism for $n > Ii + J$. (This generalizes the above argument where we considered the case $\ell = 3$ where $S_3 = D_6$.)

As we saw earlier in the course, the Hurwitz space associated to dihedral groups in the above exercise is closely related to the Cohen-Lenstra heuristics. We will next see that when $\ell = 2$, the homological stability properties of the Hurwitz space behave quite differently. This reflects the fact that the ℓ torsion in class groups behave quite differently when ℓ is odd or even.

Exercise 15.0.22. Similarly to the above exercise, suppose $\ell = 2$, so $c \subset D_4 \simeq \mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}$ is the subset of order 2 elements. Compute the dimension of $H_0(\mathrm{CHur}^c_n)$ and use this to prove that there do not exist constants I and J so that $H_i(\mathrm{CHur}^c_n) \to H_i(\mathrm{CHur}^c_{n+1})$ is an isomorphism for $n > Ii + J$.

Assuming Theorem 13.0.2 give a corrected homological stability statement for $H_i(\mathrm{CHur}^c_n)$.

Remark 15.0.23. Incidentally, this gives a proof of homological stability for CHur^c, but it shows slightly more. Namely, it shows that all of the operators α_x, α_y, and α_z (we only showed this for the first but it holds for the others by symmetry) induce isomorphisms on homology of CHur^c in a linear range. This is in contrast to Theorem 17.0.6, (coming in a future section,) which shows that the operator $U = \alpha_x^2 + \alpha_y^2 + \alpha_z^2$ induces isomorphism on $C_*(\mathrm{Hur}^c)$ in a linear range. If we work with all of Hur^c, α_x will not induce isomorphisms on homology because it will not send anything to the components Hur^z (which contain only zs and no xs).

The above shows that $\alpha_x, \alpha_y, \alpha_z$ all induce isomorphisms on $C_*(\mathrm{CHur}^c)$ past some linear range.

16. PROVING THE HOMOLOGY STABILIZES IN GENERAL

The computation of the stable homology in general is quite similar to the degree 3 case. Let's now outline how it goes, focusing on the new aspects. For now, we will only verify that the homology stabilizes. For the computation of what it stabilizes to, we will prove a more general version later in Theorem 21.0.1 where we compute the stable value when just one of the components has many points.

Theorem 16.0.1. *Let c be a finite rack with components $c_1, \ldots, c_v$. There are constants I, J so that for any $x \in c_1$, the map α_x induces an isomorphism $H_i(\mathrm{CHur}^c_{n_1,\ldots,n_v}) \to H_i(\mathrm{CHur}^c_{n_1+1,\ldots,n_v})$ when $n_1 > Ii + J$.*

We give the proof in the remainder of this section. The result is implied by Lemma 16.0.6, as is explained immediately prior to its statement in §16.0.5.

Recall that our starting point is Theorem 14.0.4 which says that $A_c/(\alpha_c^{\mathrm{ord}(c)})$ is bounded in a linear range with respect to the grading induced by any component of c. Choose a subset $x_1, \ldots, x_v$ with $x_i \in c_i$. We use $CA_c := C_*(\mathrm{CHur}^c)$.

16.0.2. *Rephrasing boundedness of a quotient as stability.* Our goal is to show $CA_c/(\alpha_{x_1}^{\mathrm{ord}(x_1)}, \ldots, \alpha_{x_v}^{\mathrm{ord}(x_v)})$ is also bounded in a linear range with respect to the grading induced by any component of c. Since the x_i all live in different components, the next lemma Lemma 16.0.3 implies that $CA_c/(\alpha_{x_1}^{\mathrm{ord}(x_1)})$ will be bounded in a linear range (by taking the bigrading in that lemma where the first grading in the number of branch points in component c_1 and the second bigrading is the number of branch points in all other components), which implies that $\alpha_{x_1}^{\mathrm{ord}(x_1)}$ induces isomorphisms $H_i(\mathrm{CHur}^c_{n_1,\ldots,n_v}) \to H_i(\mathrm{CHur}^c_{n_1+\mathrm{ord}(x_1),\ldots,n_v})$ in a linear range. Finally, if $\alpha_{x_1}^{\mathrm{ord}(x_1)}$ induces isomorphisms, one can also deduce α_{x_1} induces isomorphisms in a linear range, since a composite of injective maps which is an isomorphism are individually isomorphisms (and we apply this to $\alpha_{x_1}^{\mathrm{ord}(x_1)}$ being a composite of $\mathrm{ord}(x_1)$ iterates of α_{x_1}).

Lemma 16.0.3. *If X is a bigraded space and $v : X[0,1] \to X$ is a map that shifts the second grading by 1, if X/v is bounded in a linear range (with respect to the first grading) then X is also bounded in a linear range (with respect to the first grading).*

Proof. The cofiber sequence

$$(X/v^{n-1})[0,1] \longrightarrow X/v^n \longrightarrow X/v \tag{16.1}$$

shows by induction that X/v^n is bounded in a linear range (with respect to the first grading). Since X agrees with X/v^n in the first n degrees for the second grading, we see that X is bounded in a linear range (with the same constants as for X/v) because all X/v^n are all bounded in the same linear range. ☐

Now, choose a set $V \subset c$ containing one element of each component of c. We claim

$$CA_c/(\alpha_V^{\mathrm{ord}(V)}) \tag{16.2}$$

is bounded in a linear range. We will prove Theorem 14.0.4 by downward induction on the size of V.

Lemma 16.0.4. *$CA_c/(\alpha_c^{\mathrm{ord}(c)})$ is bounded in a linear range.*

Proof. When $V = c$, we know

$$A_c/(\alpha_c^{\mathrm{ord}(c)})$$

is bounded in a linear range. We want to get from here to showing $CA_c/(\alpha_c^{\mathrm{ord}(c)})$ is also bounded in a linear range. Moreover, A_c has a filtration whose associated graded parts are of the form $CA_{c'}$ for subracks $c' \subset c$. By induction on the size of c, we know $CA_{c'}/(\alpha_{c'}^{\mathrm{ord}(c')})$ is bounded in a linear range (and there are only finitely many subracks of c, so we can take uniform constants defining the linear range) and since every $x \in c - c'$ acts trivially on $CA_{c'}/(\alpha_{c'}^{\mathrm{ord}(c')})$, we also deduce each $CA_{c'}/(\alpha_c^{\mathrm{ord}(c)})$ is bounded in a linear range. It follows that $CA_c/(\alpha_c^{\mathrm{ord}(c)})$ is bounded in a linear range. ☐

16.0.5. *Deducing the stability theorem Theorem 16.0.1.* Now, suppose we know (16.2) holds for all sets of size larger than V, which contain at least one element in each component of V. We know that for each $y \in c - V$, $CA_c/(\alpha_V^{\mathrm{ord}(V)}, \alpha_y^{\mathrm{ord}(y)})$ is bounded in a linear range. If we can also show $CA_c/(\alpha_V^{\mathrm{ord}(V)})[\alpha_y^{-1}] = 0$ then we will know $CA_c/(\alpha_V^{\mathrm{ord}(V)})$ is bounded in a linear range by Lemma 15.0.18. Hence, to prove Theorem 16.0.1, it suffices to prove the following lemma.

Lemma 16.0.6. *Let $y \in c - V$. Supposing by induction on $|V|$ that* (16.2) *holds for all V' with $|V'| > |V|$ we have $CA_c/(\alpha_V^{\mathrm{ord}(V)})[\alpha_y^{-1}] = 0$.*

Proof. Using that localization commutes with quotients by central elements, by induction on $|V|$, taking $z \in V - y$ and $V' := V \cup z$, we have

$$CA_c/(\alpha_V^{\mathrm{ord}(V)})[\alpha_y^{-1}]/(\alpha_z^{\mathrm{ord}(z)}) = CA_c/(\alpha_{V'}^{\mathrm{ord}(V')})[\alpha_y^{-1}] = 0.$$

which implies, by Lemma 15.0.18 that it suffices to show $CA_c/(\alpha_V^{\mathrm{ord}(V)})[\alpha_y^{-1}, \alpha_z^{-1}] = 0$. Repeating this argument iteratively implies that it is enough to show $CA_c/(\alpha_V^{\mathrm{ord}(V)})[\alpha_y^{-1}, y \in c - V] = 0$. Hence, the result follows from Lemma 16.0.7 and Lemma 16.0.9. □

Lemma 16.0.7. *If $V \subset c$ is not a subrack, we have $CA_c/(\alpha_V^{\mathrm{ord}(V)})[\alpha_y^{-1}, y \in c - V] = 0$.*

Proof. If $c - V$ is not a subrack, we have some $x, y \in c - V$ with $\alpha_x^{-1}\alpha_y\alpha_x = \alpha_z$ with $z \in V$. But then α_x, α_y act invertibly while α_z acts nilpotently on $CA_c/(\alpha_V^{\mathrm{ord}(V)})[\alpha_y^{-1}, y \in c - V]$ so it must be 0. □

The next lemma is really the key step to prove, and it also came up for us as the most difficult step of the S_3 case when we needed to show the vanishing where we invert a single element and quotient by the other two in Proposition 15.0.5. The proof of the next lemma is nearly the same as Proposition 15.0.5, though there will be some small technical additions, where we will need to consider the normalizer of a subrack. (In the case of $z \subset c$ for c the transpositions in S_3, z was its own normalizer.)

Definition 16.0.8. For $c' \subset c$ a subrack, the normalizer of c' in c, $N_c(c')$ is the set of element $x \in c$ so that $x \triangleright y \in c'$ for every $y \in c'$.

Lemma 16.0.9. *If $c - V \subset c$ is a subrack containing an element of each component of c, we have $CA_c/(\alpha_V^{\mathrm{ord}(V)})[\alpha_y^{-1}, y \in c - V] = 0$.*

Proof. The key to this proof will be to consider the map

$$f_{c,c-V} : A_c/(\alpha_V^{\mathrm{ord}(V)})[\alpha_y^{-1}, y \in c - V] \to A_{N_c(c-V)}/(\alpha_V^{\mathrm{ord}(V)})[\alpha_y^{-1}, y \in c - V]$$

We will show this map is an equivalence. Once we verify this is an equivalence, let us explain why this will conclude the proof. Indeed, the map is compatible with the grading that takes the "connected" part of these modules, i.e. the part of them that generates the rack c. On the source of $f_{c,c-V}$ this just gives us $CA_c/(\alpha_V^{\mathrm{ord}(V)})[\alpha_y^{-1}, y \in c - V]$. However, on the other hand, let us observe that $N_c(c - V) \subsetneq c$. This follows from the next easy lemma.

Lemma 16.0.10. *Suppose $c' \subset c$ is a subrack which is not a union of components of c. Then $N_c(c') \neq c$.*

Proof. By assumption, there is some component $c'' \subset c$ not contained in c' but which meets c'. Hence there is some $x \in c'' \cap c'$ and some y with $y \triangleright x \notin c'$. Therefore, $y \notin N_c(c')$. □

There is thus some $x \in c$ which sends $y \in c - V$ to $z \in V$, which implies that $x \notin N_c(c - V)$. Hence, the connected part of the source of $f_{c,c-V}$ is sent to 0 on the target, since it vanishes in the top piece of the filtration, as no collection of elements in $N_c(c - V)$ can generate all of c. This gives the desired vanishing.

To conclude, let us explain why $f_{c,c-V}$ is an equivalence. Again, this is quite similar to Proposition 15.0.5, so we will be somewhat terse on the details. The point is that we can apply the descent lemma, Lemma 15.0.10 in the case the map of rings is $A_c[\alpha_x^{-1}, x \in c - V] \to A_{N_c(c-V)}[\alpha_x^{-1}, x \in c - V]$, I_S is the ideal generated by $\alpha_x, x \in V \cap N_c(c - V)$, I_R is the ideal generated by $\alpha_x, x \in V$ and M is $A_c/(\alpha_x^{\mathrm{ord}(x)}, x \in V)[\alpha_x^{-1}, x \in c - V]$. In this case, $\pi_0 R/I_R$ and $\pi_0 S/I_S$ can both be identified with $\pi_0 A_{c-V}[\alpha_y^{-1}, y \in c - V]$. This reduces us to verifying

$$\begin{aligned}&\pi_0 A_{c-V}[\alpha_y^{-1}, y \in c - V] \otimes_{A_c[\alpha_y^{-1} y \in c - V]} \pi_0 A_{c-V}[\alpha_y^{-1}, y \in c - V]\\ &\to \pi_0 A_{c-V}[\alpha_y^{-1}, y \in c - V] \otimes_{A_{N_c(c-V)}[\alpha_y^{-1} y \in c - V]} \pi_0 A_{c-V}[\alpha_y^{-1}, y \in c - V]\end{aligned}$$

is an equivalence. Since $\pi_0 A_{c-V}$ is by definition the 0th reduced homology of $\pi_0 \operatorname{Hur}_+^{c-V}$, upon taking $c - V = c'$, this is precisely the content of Proposition 16.0.12 □

Exercise 16.0.11. Recall that in Exercise 15.0.14, you explained why, in the case c was the set of transpositions in S_4, $V = c - (12)$, the map $A_c/(\alpha_V^{\mathrm{ord}(V)})[\alpha_y^{-1}, y \in c - V] \to A_{(12)}/(\alpha_V^{\mathrm{ord}(V)})[\alpha_y^{-1}, y \in c - V]$ was not an equivalence.

Explain why if one replaces (12) by $N_c(12)$, the rest of the proof of Proposition 15.0.11 goes through to instead show $A_c/(\alpha_V^{\mathrm{ord}(V)})[\alpha_y^{-1}, y \in c - V] \to A_{N_c((12))}/(\alpha_V^{\mathrm{ord}(V)})[\alpha_y^{-1}, y \in c - V]$ is an equivalence. (The point is that in the S_3 case, (12) happened to be its own normalizer.) In particular, explain what benefit we get from using the normalizer of (12) in place of just (12). *Hint:* The solution is spelled out in the Proposition 16.0.12, but we recommend thinking about this before reading Proposition 16.0.12.

The proof of the following is quite similar to Proposition 15.0.11. We encourage the reader to go back and first absorb the proof of that if they have not yet done so.

Proposition 16.0.12. *Let c be a rack, $c' \subset c$ be a subrack with normalizer $N_c(c')$ and $X_+ := \pi_0 \operatorname{Hur}^{c'}[\alpha_{c'}^{-1}]$. Then*

$$X_+ \otimes_{\operatorname{Hur}_+^c} X_+ \to X_+ \otimes_{\operatorname{Hur}_+^{N_c(c')}} X_+$$

is a homology equivalence.

Proof. As with Proposition 15.0.11 we will construct a nullhomotopy. Using scanning, we can identify both sides with spaces consisting of an element of X_+ on the left and right side, together with elements of c (or $N_c(c')$) in the middle, with at most one in each row. When the points hit the boundary, they act on it, and any $x \in c - c'$ acts on X_+ by sending that element to the basepoint. Hence, we can consider the subspace of $X_+ \otimes_{\mathrm{Hur}^c_+} X_+$ with some labeled point v of the configuration not lying in $N_c(c')$, and produce a nullhomotopy of that subspace.

Actually, since we only need to show the map is a homology equivalence, we will define a filtration, and construct a nullhomotopy on each graded part of that filtration. (In general, to show a map of spaces is a homology equivalence, it suffices to do it on each graded part of a filtration, though this does not suffice if one wishes to show it is a homotopy equivalence.) The ith step of the filtration F^i is given by those configurations so that one of the lowest i labeled elements lies in $c - N_c(c')$. The associated graded is then configurations where the $i-1$ lowest elements lie in $N_c(c')$ and the ith lowest lies in $c - N_c(c')$. By constructing a nullhomotopy on the associated graded of this filtration, we mean that we construct a nullhomotopy on F^i, the ith step of these filtrations, but where we identify the F^{i-1} with the basepoint.

Now, choose a total ordering on c'. Also, as in Remark 15.0.12, we may assume that any two points in our configuration have height differing by at least ε. There is necessarily some element $x \in c'$ we can slide out from the left, move across the configuration directly above v (at height $\varepsilon/2$ higher than v, so that when it hits the right hand side it will not be labeled by something in c. (The reason such an x exists is that when x passes above an element, it gets conjugated by that element, while if x passes below an element, nothing happens to x. Therefore, we can choose x so that when it gets to passing by v, it is sent from something in c' to something outside of c', and then when it passes above further elements of c' it will remain outside of c'. Finally, this point will hit the right boundary as something not in c' and so send the configuration to the basepoint.) We now choose x to be the smallest such element in the total ordering. On the basepoint, this nullhomotopy is constant at the basepoint.

It remains to check that the above nullhomotopy is continuous. This way, we need to check that it is compatible with operation that let points of the configuration act on the boundaries. First, if any point beneath the ith point acts on the boundary, the resulting point will lie in F^{i-1} because the $i-1$th point from the bottom will lie in $c - N_c(c')$. Therefore, we can identify such a point with the basepoint in the associated graded of our filtration, and the configuration is sent to the basepoint.

If the ith point acts on the left or right boundary, since the ith label does not lie in c', the configuration is again sent to the basepoint. Finally, if any point above the ith point is sent to the boundary, we claim we will have the same nullhomotopy where we continue to move x across at constant speed.

Exercise 16.0.13. Verify the above claim. What this entails showing the following. Suppose the point at location $i + \alpha$ from the hits the left boundary for $\alpha > 0$ and let x be the minimal element, under the total ordering on c', we can move across to send the configuration to the basepoint. Let x' be the minimal element that we could move across before the point at $i + \alpha$ hit the boundary. Then, you must show $x = x'$.

The above shows the map is continuous, so we are done. □

Exercise 16.0.14. For this exercise, you may assume that you are working in an ϵ-spaced bar construction, in the sense that any two points in the bar construction have height differing by at least ε.

Suppose c is the set of transpositions in S_4 and take c' to be the subrack consisting of only (12). Explain why the filtration in Proposition 16.0.12 is not necessary to prove it in this case by showing the homotopy of sliding an element of c across the rectangle above the lowest element not in $N_c(12)$ is already continuous (without needing to pass to a filtration). If you did Exercise 16.0.11, you should have already verified this.

Exercise 16.0.15. Given an example of a rack c such that the use of the filtration in the proof of Proposition 16.0.12 is necessary. , explain through an example why the above equivalence $A_c/(\alpha_V^{\mathrm{ord}(V)})[\alpha_y^{-1}, y \in c - V] \to A_{N_c((12))}/(\alpha_V^{\mathrm{ord}(V)})[\alpha_y^{-1}, y \in c - V]$ needs the filtration introduced in Proposition 16.0.12. I.e. show that if one just tries to slide elements across and doesn't work on the associated graded of the filtration, the homotopy sliding elements across will not be continuous.

17. THE ORIGINAL HOMOLOGICAL STABILITY ARGUMENT (REMIX)

The original homological stability argument presented in [EVW16] was novel for its time, but also incredibly difficult to grok (in my opinion). Since then, if one is willing to stomach a modicum of higher algebra a substantial simplification has been presented in [RW20]. Ishan and I have found a further simplification, which I present now.

Remark 17.0.1. Loosely speaking, Randal-Williams argument introduces several unnecessary steps where he goes back and forth between bounding

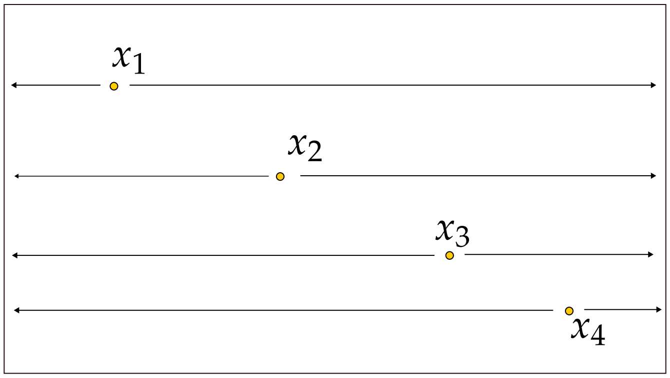

FIGURE 8. This is a picture of one of the points in a typical cell of dimension 4. Points can move left to right staying on their fixed row, and overall this traces out a copy of S^4 when we collapse the boundary of this cell corresponding to points hitting the left and right boundaries of the rectangle.

modules over $R := C_*(\mathrm{Hur}^c)$ and bounding modules over $\pi_0 R$. However, we are able to just directly work with modules over R, without going back and forth. See [LL25a, Remark 3.0.1] for more details on the improvement.

At this point, the reader may wish to recall the definition of being bounded in a linear range from Definition 14.0.3. In this section, for c a fixed rack, we will write $A := C_*(\mathrm{Hur}^c)$.

Lemma 17.0.2. *For c any rack, and $A := C_*(\mathrm{Hur}^c)$, let A act on k by letting each α_x for $x \in c$ act by 0. Then, $\pi_i((k \otimes_A k)_n) = 0$ for $n > i$. In particular $k \otimes_A k$ is bounded in a linear range.*

Proof. The proof follows from a scanning argument similar to that presented in Lemma 11.3.15. An important difference here is that we are taking a different augmentation map, by which we mean that in this case, when points fall off, they send the configuration to the basepoint, whereas before they just fell off and nothing happened. (There reason for this difference is related to the fact that we are showing the homology stabilizes instead of computing the stable value.) First, $k \otimes_A k$ can be identified with $\widetilde{C}_*(*_+ \otimes_{\mathrm{Hur}^c_+} *_+)$. So, we are interested in bounding the homology of $*_+ \otimes_{\mathrm{Hur}^c_+} *_+$. Using a similar scanning argument, we can identify points of $*_+ \otimes_{\mathrm{Hur}^c_+} *_+$ with collections of points in a rectangle labeled by elements of c with at most one in each row, so that if x_1 labeled by s passes to the right over x_2 labeled by t then x_1 gets the label $t \triangleright s$, while if a point falls off the boundary, the configuration goes to the basepoint. See Figure 8 for a picture of a point in one of the cells.

Overall, there will be $|c|^n$ cells in grading n, corresponding to specifying n points in distinct vertical coordinates. These each constitute an n-sphere, which are glued to the basepoint along its boundary since each can be identified with $[0,1]^n$ corresponding to the vertical positions of the n points in the rectangle (recalling that the vertical positions are fixed). Since this space is built out of only n cells in grading n, $\pi_i(X_n) = 0$ for $i > n$. $\square$

Definition 17.0.3. A *rack c* is *non-splitting* if every subrack $c' \subset c$ has a single component.

Example 17.0.4. A conjugacy class in a group is non-splitting if and only if its intersection with every subgroup is either empty or a single conjugacy class.

Exercise 17.0.5. Show that a rack is non-splitting if and only if $H_0(\mathrm{Hur}^c_n)$ stabilizes as n grows. *Hint:* You may use that if c is a rack with a single component, then $H_0(\mathrm{CHur}^c_n)$ stabilizes as n grows, which follows from Theorem 13.0.2 (though it is possible to prove in a much simpler direct way).

For the next result, it may be useful to recall Fact 10.4.10, so the homotopy groups of $C_*(X)$ are used as alternate notation for the homology of X.

Theorem 17.0.6. *Suppose c is non-splitting and consists of order d elements. Let* $U = \sum_{x \in c} \alpha^d_x$ *be the degree d stabilization operator and* $A := C_*(\mathrm{Hur}^c)$. *There are constants I and J so that* $U : H_i(\mathrm{Hur}^c_n) \to H_i(\mathrm{Hur}^c_{n+d})$ *induces an isomorphism once* $n > Ii + J$.

Remark 17.0.7. The idea to prove this result goes through the notion of $\mathbb{E}_1$ cells. Essentially, the $\mathbb{E}_1$ cells of an algebra R are the generators of the homotopy groups of $k \otimes_R k$ and the $\mathbb{E}_1$ cells of an R module M are the generators of the homotopy groups of $k \otimes_R M$. The point of the theorem is roughly that knowing bounds for the $\mathbb{E}_1$ cells of A and for the cells of A/U over A, together with the nilpotency condition that A/U acts nilpotently on U will in total give bounds for the grading A/U, or equivalently the grading where U induces an isomorphism on A.

Proof. Consider the quotient A/U, i.e., the cone $A \xrightarrow{U} A$. Our goal can be rephrased as showing that A/U is bounded in a linear range, as this then implies U induces isomorphisms on homotopy groups of A in a linear range. The assumption that c is nonsplitting implies $U : H_0(\mathrm{Hur}^c_n) \to H_0(\mathrm{Hur}^c_{n+d})$ is an isomorphism for n sufficiently large. (See [EL24, Proposition A.3.1] for a proof, but this is not very important to understand the idea, you can just assume that we are in a situation where there is some operator U which induces an isomorphism on $H_0(\mathrm{Hur}^c_n) \to H_0(\mathrm{Hur}^c_{n+d})$ for n sufficiently large. This is closely related to Exercise 17.0.5, as showing the 0th homology stabilizes is certainly a prerequisite for the existence of a map inducing an isomorphism on the 0th homology group.) In particular, $\pi_0(A/U)$ has bounded grading since it vanishes as soon as the map U is surjective.

The first idea is now to use nilpotency of the action of U on A/U to deduce that if the generators of $\pi_i(A/U)$ are bounded, then $\pi_i(A/U)$ is also bounded.

Lemma 17.0.8. *There is a fixed number R so that the augmentation ideal $I = \ker(\pi_0 A \to k)$ kills $\pi_i(A/U)$ for each $i \geq 0$. In fact, we may take $R = N + 2d$ for $d = \deg U$ and N the smallest number so that for any $n > N$, $U : (\pi_0 A)_n \to (\pi_0 A)_{n+d}$ is surjective. In particular, if $\pi_i(A/U)$ is generated in grading at most C, $\pi_i(A/U)$ vanishes in grading above $C + R$.*

Proof. The final statement follows from the first one because every element of a $\pi_0 A$ module is either a generator or some I multiple of a generator.

We now prove the first statement. For any element $x \in \pi_0(A)$, x^2 acts by 0 on A/x (see [LL24, Lemma 3.5.2]) so in particular, U^2 acts by 0 on A/U. This implies that a power R of the augmentation ideal $I = \ker(\pi_0 A \to k)$ kills A/U, and hence also kills $\pi_i(A/U)$. To see why, note that the non-splitting condition implies that the map $U : (\pi_0 A)_n \to (\pi_0 A)_{n+d}$ is a surjective for $n > N$, for an appropriate constant N. Hence, for I the augmentation ideal, any element in I^{N+2d} is in the image of U^2, and so acts by 0 on A/U. □

We now prove by induction on i that

(1) $\pi_i(A/U)$ is generated as a $\pi_0 A$ module in grading at most $(R+2)i + R$

(2) $\pi_j(k \otimes_A \pi_i(A/U))$ vanishes in gradings larger than $j + (R+2)i + 2R$.

If we show condition (1) above, it follows from Lemma 17.0.8 that $\pi_i(A/U)$ vanishes in grading above $(R+2)i + 2R$, which will conclude our proof.

We already explained above that (1) holds when $i = 0$ because $\pi_0(A/U)$ can be identified with the cokernel of multiplication by U on $\pi_0(A)$ (no derived things are going on here). In particular, it is a quotient of $\pi_0(A)$ so it is generated in grading 0 by the image of the unit in $\pi_0(A)$.

Next, let us explain why (1) for a fixed value of i implies (2) for the same value of i and arbitrary $j \geq 0$. We will start by showing (1) for fixed i implies (2) for that same value of i. That is, we show $\pi_j(k \otimes_A \pi_i(A/U))$ vanishes in gradings at most $j + (R+2)i + 2R$, assuming $\pi_i(A/U)$ vanishes in grading at most $(R+2)i + R$. In order to bound the gradings appearing in $\pi_j(k \otimes_A \pi_i(A/U))$ we can filter $\pi_i(A/U)$ by its grading so that A (which already acts through $\pi_0 A$ in fact acts trivially on it. Hence, it suffices to bound the gradings appearing in $(k \otimes_A k) \otimes_k \pi_i(A/U)$. Recall we are assuming $\pi_i(A/U)$ is generated in gradings at most $(R+2)i + R$. Using Lemma 17.0.8,

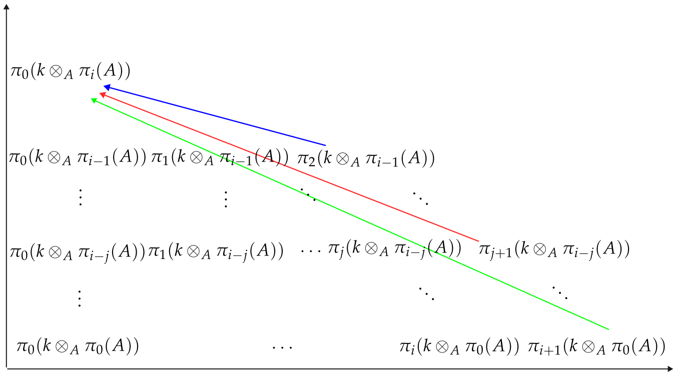

FIGURE 9. This is a picture of the spectral sequence $\pi_j(k \otimes_A \pi_{i-j}(A/U)) \implies k \otimes_A (A/U)$. Once the grading is big enough, the diagonal and super diagonal terms in rows beneath the top row i will vanish, implying $\pi_0(k \otimes_A \pi_i(A))$ vanishes as well.

it follows $\pi_i(A/U)$ is concentrated in gradings at most $(R+2)i + 2R$. Next, since $\pi_j(k \otimes_A k)$ vanishes in grading greater than j by Lemma 17.0.2, we find that $(k \otimes_A k) \otimes_k \pi_i(A/U)$. vanishes in grading more than $j + (R+2)i + 2R$, and hence $\pi_j(k \otimes_A \pi_i(A/U))$ does as well.

We next show why the vanishing of $\pi_j(k \otimes_A \pi_{i'}(A/U))$ in grading at most $j + (R+2)i' + 2R$ for $i' < i$ also implies the vanishing of $\pi_0(k \otimes_A \pi_i(A/U))$ in gradings $(R+2)i + R$. (Note that this is similar to (2), but slightly stronger, since the bound here is R less than the bound in (2).) This is a spectral sequence argument as in [RW20, Proposition 8.1], which was in turn inspired by the argument in [EVW16, Theorem 6.1].

The point is that there is a spectral sequence $\pi_j(k \otimes_A \pi_{i-j}(A/U)) \implies \pi_i(k \otimes_A (A/U))$. coming from filtering A/U by its homotopy groups. This converges to $k \otimes_A A/U \simeq k/U \simeq k_0 \oplus \Sigma^1 k_d$, which is just two copies of k since U acts by 0 on k, one in grading 0 and homological degree 0, and the other in grading d and homological degree 1, since U acts by 0 on k. See Figure 9.

Therefore, this spectral sequence converges to 0 in grading larger than d. We see that for $j > 0$, $\pi_j(k \otimes_A \pi_{i-j}(A/U))$ all vanish for grading more than $(R+2)i + R$ because $(R+2)i + R > j + (R+2)(i-j) + 2R$ for $j > 0$, and these terms vanish above the latter grading by our inductive assumption. Note here $R > d$ so the spectral sequenced is converging to 0 in this grading range, as explained above.

This implies that the term starting at $\pi_0(k \otimes_A \pi_i(A/U))$ is the only term on the diagonal on the diagonal among the $\pi_j(k \otimes_A \pi_{i-j}(A/U))$ on the starting

E_2 page of the spectral sequence. Therefore, all the differentials mapping into $\pi_j(k \otimes_A \pi_{i-j}(A/U))$ must jointly surject onto it in grading more than $(R+2)i+R$. So, it is enough to show all the differentials mapping into it have vanishing source. Indeed, the differentials come from terms of the form $\pi_{j+1}(k \otimes_A \pi_{i-j}(A/U))$ for $j > 0$. These terms vanish in grading more than $(R+2)i+R$ because $(R+2)i+R \geq (j+1)+(R+2)(i-j)+2R$ for $j > 0$, and again the latter terms vanish by induction.

Finally, we wish to deduce (1) for the value i from (2) for i and $j = 0$. This follows from Lemma 17.0.9, which we prove next. $\square$

The next lemma is essentially explaining why, if you can bound the $\mathbb{E}_1$ cells of a module, you can bound the generators of the module; see Remark 17.0.7. Recall a module is *connective* by definition if it has no negative homotopy groups.

Lemma 17.0.9. *Suppose M is a connective graded $\pi_0 A$ module and the augmentation ideal $I = \ker(\pi_0 A \to k)$ acts nilpotently on M with $\pi_0(k \otimes_A M)$ vanishing in grading above C. Then M is generated as a $\pi_0 A$ module in grading at most C.*

Proof. We can identify $\pi_0(k \otimes_A M)$ with $\pi_0(k \otimes_{\pi_0 A} M)$ and so this vanishes above grading C. The map $M \to \pi_0(k \otimes_{\pi_0 A} M)$ is surjective and the target vanishes in grading at most C. We want to show the source also vanishes in grading at most C. Choose generators of the target, and lift them arbitrarily to elements of the source. Let $M' \subset M$ denote the submodule spanned by these generators. It suffices to show $M' = M$, or equivalently $M/M' = 0$. (Note that M and M' are just usual discrete modules and so this quotient just means the usual quotient, there is no derived stuff going on here.) By construction, the map $\pi_0(k \otimes_{\pi_0 A} M') \to \pi_0(k \otimes_{\pi_0 A} M)$ is surjective, so its cokernel vanishes. On the other hand, the cokernel of this map is identified with $\pi_0(k \otimes_{\pi_0 A} (M/M'))$. So $\pi_0(k \otimes_{\pi_0 A} (M/M')) = 0$. Now, using the long exact sequence on homotopy groups coming from tensoring $I \to \pi_0 A \to k$ with M/M' we get an exact sequence

(17.1)

$$\pi_0(I \otimes_{\pi_0 A} M/M') \longrightarrow \pi_0(\pi_0 A \otimes_{\pi_0 A} M/M') \longrightarrow 0 = \pi_0(k \otimes_{\pi_0 A} M/M').$$

We can then identify $\pi_0 A \otimes_{\pi_0 A} M/M' \simeq M/M'$, which is a discrete module and so we also have $\pi_0(M/M') = M/M'$. Hence, the image of the first map is identified with $I(M/M') \subset M/M'$. From the surjectivity of this map, we conclude $I(M/M') = M/M'$. Since I acts nilpotently on M by assumption, we find $M/M' = 0$. $\square$

Exercise 17.0.10 (Involved exercise)**.** Let $A := C_*(\mathrm{Hur}^c; k)$. Let $I := \ker(\pi_0 A \to k)$. Recall notation from Notation 14.0.1.

(1) In the proof above, we used that $k \otimes_A A/U \simeq k/U = k_0 \oplus \Sigma^1 k_d$ in concentrated in finitely many gradings (namely 0 and d). Show similarly that $k \otimes_A (A/\alpha_c^{\mathrm{ord}(c)})$ is concentrated in finitely many gradings and homological gradings.

(2) You may take for granted that there is a constant κ so that I^κ acts by 0 on $A/(\alpha_c^{\mathrm{ord}(c)})$. (This is not so difficult, and is something you could reasonably show yourself. See [LL25a, Lemma 3.2.7] for a proof.) Using this in place of Lemma 17.0.8, and the first part of this problem, show $A/(\alpha_c^{\mathrm{ord}(c)})$ is bounded in a linear range. *Hint:* The argument will be similar to the above, but you will have to adjust the constants in the proof.

18. THE STABLE VALUE OF HOMOLOGY IN ALL DIRECTIONS

Above, in Theorem 13.0.3, we stated what the stable homology of Hurwitz spaces looks like assuming our cover has sufficiently many points of each monodromy type. In order to prove this, we prove a more general result describing the stable homology when some component has few monodromy types.

The typical example is as follows:

Example 18.0.1. Consider $c := S_3 - \mathrm{id}$, so Hur_n^c parameterizes degree 3 covers, (or, equivalently, S_3 covers) branched at n points. Then, c has two conjugacy classes, 3 cycles $(123), (132)$ and transpositions $(12), (13), (23)$. The 3 cycles correspond to totally ramified points, so the fibers of the degree 3 cover have a unique point. The transpositions correspond to simply ramified points, where each such fiber has two points, one branched and one unramified. Above, we were able to compute the stable homology of components with many transpositions and many 3 cycles. However, we may be interested in the stable homology of components with one 3 cycle and many transpositions. This is what we will now address.

To express the stable homology in general directions, we will need the notion of a quotient rack.

Definition 18.0.2. Let c be a rack with operation $\triangleright : c \times c \to c$. A subset $c' \subset c$ so that $x \triangleright y \in c'$ if $x, y \in c'$ is called a *subrack*. A subrack $c' \subset c$ is *normal* if $x \triangleright y \in c'$ for $x \in c, y \in c'$. In other words, the normalizer of c' is all of c.

Given a normal subrack $c' \subset c$, one can form the quotient rack c/c', whose elements consist of $c/\sim$ for $\sim$ the equivalence relation generated by $x \sim y \triangleright x$ for $x \in c, y \in c'$. In other words, we identify two elements $x, z \in c$ of the rack in the quotient rack c/c' if there is a sequence of elements

$y_1, \dots, y_j \in c'$ so that $z = (y_1 \cdots y_j) \triangleright x$, i.e. x and z are related by the action of a sequence of elements in c'.

Exercise 18.0.3. For c a rack and $c' \subset c$ a normal subrack, show that c/c' is again a rack. That is, one needs to verify that if $x, y \in c$ and $x', y' \in c$ are such that x' is in the same c' orbit of x and y' is in the same c' orbit of y then $x \triangleright y$ is in the same c' orbit as $x' \triangleright y'$. *Hint:* You will need to use normality of $c' \subset c$.

Definition 18.0.4. The *trivial rack with k elements* is a rack c with k elements so that for any $x, y \in c$, $x \triangleright y = y$.

Example 18.0.5. Consider the rack $c = S_3 - \mathrm{id}$. Let $c' \subset c$ denote the conjugacy class of transpositions. Then c/c' is the two element trivial rack. To see this, simply note that the action of transpositions on themselves is transitive, and the action of transpositions on 3-cycles is also transitive. Hence, each of the two conjugacy classes get identified to a single point, and they act trivially on each other, so c/c' is the two element trivial rack.

Exercise 18.0.6. Consider the rack $c = S_3 - \mathrm{id}$ and let $c' \subset c$ denote the conjugacy class of 3-cycles. Show that c/c' is the Joyce quandle of Exercise 12.0.7.

Exercise 18.0.7. Let c be a rack and $c' \subset c$ be a normal subrack which generates c in the sense that the reduced structure group of c is generated by the action of c' on c. Show that c/c' is a trivial rack of order equal to the number of components of c.

Exercise 18.0.8. For c a rack and $c' \subset c$ a normal subrack, show that c and c/c' have the same number of components.

The point of introducing quotient racks is the following:

Theorem 18.0.9. *Let c be a rack with components $c = c_1 \cup \cdots \cup c_v$. Let $c' \subset c$ denote the first j components of c. There are constants I and J so that for $Z \subset \mathrm{CHur}^c_{n_1,\dots,n_v}$ a component which maps to a corresponding component $Z' \subset \mathrm{CHur}^{c/c'}_{n_1,\dots,n_v}$. Then, the map $Z \to Z'$ induces an isomorphism $H_i(Z, \mathbb{Z}[1/\#G^0_c]) \simeq H_i(Z', \mathbb{Z}[1/\#G^0_c])$ when $n_1, \dots, n_j > Ii + J$.*

We will restate and prove this in Theorem 21.0.1.

Exercise 18.0.10. For c a rack, show we can identify $\mathrm{Hur}^{c/c}_{n_1,\dots,n_v} \simeq \mathrm{Conf}_{n_1,\dots,n_v}$.

Remark 18.0.11. Note that this is a generalization of Theorem 13.0.3 because if $c' = c$, the quotient rack c/c is a trivial rack of order equal to the number of components of c by Exercise 18.0.7. Then, we can identify $\mathrm{Hur}^{c/c}_{n_1,\dots,n_v} \simeq \mathrm{Conf}_{n_1,\dots,n_v}$, implying Theorem 13.0.3 is a special case of Theorem 18.0.9.

19. APPLICATION: BHARGAVA'S CONJECTURE

Bhargava's conjecture, [Bha07, Conjecture 1.2], predicts the asymptotic growth of the number of degree d number fields with Galois group S_d, as a function of the discriminant. For culture, before continuing, we recall the statement of Bhargava's conjecture.

Conjecture 19.0.1 ([Bha07, Conjecture 1.2])**.** Let $N_d(X)$ denote the number of number fields of degree d having discriminant with absolute value at most X. Let $q(n,k)$ denote the number of partitions of n into at most k parts. Let $r_2(S_d)$ denote the number of elements of order either 1 or 2 in S_d. Then,

(19.1)
$$\lim_{X\to\infty}\frac{N_d(X)}{X}=\frac{r_2(S_d)}{2d!}\prod_{p \text{ prime}}\left(\frac{\sum_{k=0}^{d}q(k,d-k)-q(k-1,d-k+1)}{p^k}\right).$$

This has been proven for $d \leq 5$, the $d=3$ case was shown by Davenport and Heilbronn while $d=4,5$ was shown by Bhargava. It is wide open for $d \geq 6$.

Now, we'd like to describe a proof of a function field analog of this conjecture. Computing the stable homology in all directions has a number of natural applications in arithmetic statistics. As a very simple example, with the simpler result Theorem 13.0.3, we were able to obtain good asymptotics for the total number of G covers in Theorem 8.0.16. However, we couldn't count S_3 covers with one 3 cycle and many transpositions. Theorem 18.0.9 allows us to now do this. A well known conjecture of Bhargava predicts the number of S_d extension of $\mathbb{Q}$ when counted by discriminant. We now prove a function field version of this conjecture, for q large relative to d.

Remark 19.0.2. Before counting the number of S_d extensions of $\mathbb{F}_q(t)$, it is natural to ask why one needs the the more sophisticated version Theorem 18.0.9 computing the stable homology in all directions, instead of the more pedestrian Theorem 13.0.3. Indeed, it might seem that one only needs to count the number of components with many branch points of each type, as there will be relatively few other components. This would indeed be true if we were counting by reduced discriminant. However, Bhargava's conjecture is specifically about counting by the discriminant of the extension.

Let's think about the example where $G = S_3$. Then, if we count by reduced discriminant, we will be counting covers with n_1 transpositions and n_2 three cycles, and $n_1 + n_2 = n$ is growing. In this case, there will be roughly $n+1$ components all of dimension n, and so each contribute roughly the same point count. This would allow us to count by reduced discriminant.

However, now, suppose we count by the usual discriminant. Since each 3 cycle contributes twice as much to inertia as a transposition, we are counting covers with $n_1 + 2n_2 = n$. Now, when we count these with fixed value of n, these will have varying dimensions; the dimension is $n_1 + n_2$. The biggest dimension is when $n_1 = n, n_2 = 0$ and the smallest dimension is when n_1 is 0 or 1 and n_2 is roughly $n/2$, in which case the dimension will roughly be $n/2$. You should think that the biggest contribution comes from $(n_1 = n, n_2 = 0)$, which has roughly q^n points, the next biggest comes from $(n_1 = n - 2, n_1 = 1)$ which has roughly q^{n-1} points. After that, terms continue to drop off and contribute less. From this perspective, we can see that roughly q^{n-1} out of a total of roughly q^n points come from components where n_1 is large and $n_2 = 1$. This contributes a proportion roughly $q^{n-1}/q^n = 1/q$ to this term will show up in our point count. In particular, we will need to understand the stable homology with 1 three cycle and many transpositions.

Here is the statement of the function field version of Bhargava's conjecture.

Notation 19.0.3. For $d \geq 2$, write $S_d - \mathrm{id} = c_1 \cup \cdots \cup c_v$ as a disjoint union of its non-identity conjugacy classes, and let c_1 denote the conjugacy class of transpositions. We fix q a prime power, relatively prime to $d! = |S_d|$.

If $K/\mathbb{F}_q(t)$ is an extension and $\mathscr{O}_K$ is the normalization of $\mathbb{F}_q[t]$ in K, we say $K/\mathbb{F}_q(t)$ has discriminant equal to the discriminant of $\mathscr{O}_K$ over $\mathbb{F}_q[t]$, which we define to be $q^{\deg \Omega_{\mathscr{O}_K/\mathbb{F}_q[t]}}$. We use $\Delta(\mathbb{F}_q(t), S_d - \mathrm{id}, q^n)$ to denote the number of degree d, S_d extensions K of $\mathbb{F}_q(t)$ such that $\mathscr{O}_K$ that of discriminant q^n and, for $c_i \subset S_d$ a conjugacy class, we use $\Delta(c_i) := d - r(g)$, where $r(g)$ is the number of orbits of any $g \in c_i$ on $\{1, \ldots, d\}$.

Definition 19.0.4. Let $\sigma(n_1, \ldots, n_v)$ denote the number of conjugacy classes of S_d so that the projection of that conjugacy class to the abelianization $S_d^{\mathrm{ab}} \simeq \mathbb{Z}/2\mathbb{Z}$ agrees with the projection of $n_1 c_1 + \cdots + n_v c_v$ to S_d^{ab}.

The following theorem is our function field version of Bhargava's conjecture, which counts the number of S_d extensions of $\mathbb{F}_q(t)$ by discriminant.

Theorem 19.0.5. *Using notation from Notation 19.0.3 and Definition 19.0.4, if q is sufficiently large depending on d, we have*

$$\Delta(\mathbb{F}_q(t), S_d - \mathrm{id}, q^n) = \left\| \sum_{\substack{n_1, \ldots, n_v \\ \sum_{i=1}^{v} n_i \Delta(c_i) = n}} \sigma(n_1, \ldots, n_v) \# \mathrm{Conf}_{n_1, \ldots, n_v, \mathbb{F}_q}(\mathbb{F}_q) \right\| + o(q^n). \tag{19.2}$$

The reader familiar with Bhargava's conjecture in the number field case may be used to seeing it expressed in terms of an Euler product. The next exercise explains why one can view the point counts on configuration space as an Euler product.

Exercise 19.0.6. (1) Show the number of $\mathbb{F}_q$ points of Conf_n is the coefficient of t^n in $\prod_{x \in \mathbb{A}^1}(1 + t^{\deg(x)})$, where t is a formal variable.

(2) Find a similar formula to express $\#\mathrm{Conf}_{n_1,\ldots,n_k}(\mathbb{F}_q)$ as a coefficient in an Euler product over closed points of $\mathbb{A}^1$.

(3) Find a similar Euler product whose coefficients compute the number of pairs of monic polynomials, both of degree n over $\mathbb{F}_q$, which share no common factors with each other (though individually they may have repeated factors).

(4), Bonus Can you evaluate the coefficients in the expression from the first part? It turns out to be a fairly simple rational function in q. If you haven't seen this before, you can look up the zeta function of a variety. Similarly, can you evaluate the expression from the third part?

Now, let's explain how the proof of Bhargava's conjecture follows from Theorem 18.0.9.

Lemma 19.0.7. *Let* $c := S_d - \mathrm{id}$ *and let* $c_1 \cup \cdots \cup c_v = c$ *with* c_1 *the conjugacy class of transpositions. There are* $\sigma(n_1, \ldots, n_v)$ *components of* $[\mathrm{CHur}^c_{n_1,\ldots,n_v} / S_d]$ *for* n_1 *sufficiently large.*

Suppose q *has characteristic prime to* $d!$*. Then, the same holds over* $\mathbb{F}_q$*. I.e., there are* $\sigma(n_1, \ldots, n_v)$ *components of* $[\mathrm{CHur}^c_{n_1,\ldots,n_v,\mathbb{F}_q} / S_d]$ *with which is moreover geometrically irreducible.*

Proof. The proof is a bit orthogonal from the main subject of these notes, so we will only sketch the idea of the proof. The first step is to determine the components of $\mathrm{CHur}^c_{n_1,\ldots,n_v}$ with specified boundary monodromy. The boundary monodromy means that we take the product of the elements appearing as monodromies over $\mathbb{A}^1$. The abelianization of S_d is $\mathbb{Z}/2\mathbb{Z}$ and if $n = n_1 + \cdots + n_v$ the product of n_i elements of c_i must have specified image in $\mathbb{Z}/2\mathbb{Z}$, (so the boundary monodromy will either be in A_d or not) but this is the only constraint on the boundary monodromy for n_1 sufficiently large. One can then show that if one fixes the boundary monodromy in the specified coset of A_d, there will be a unique component of $\mathrm{CHur}^c_{n_1,\ldots,n_v}$ with that boundary monodromy. To see this, one can deduce an upper bound on the number of components as a certain quotient of $H_2(S_d, \mathbb{Z})$. The group $H_2(S_d, \mathbb{Z})$ has size at most 2 and the quotient ends up always being trivial, which ends up implying there is a unique component as claimed.

Then, quotienting by S_d, we identify the different conjugacy classes of boundary monodromies. Therefore, the number of components is precisely $\sigma(n_1, \ldots, n_v)$, which is the number of conjugacy classes in S_d, subject to the constraint that they have the correct image in the abelianization, as described above.

Using a general comparison theorem between Hurwitz spaces over $\mathbb{C}$ and $\overline{\mathbb{F}}_q$ when q is prime to $\#G$, as discussed in Theorem 29.0.1, there is a bijection between components of $[\mathrm{CHur}^c_{n_1,\ldots,n_v}/S_d]$ and components of $[\mathrm{CHur}^c_{n_1,\ldots,n_v,\overline{\mathbb{F}}_q}/S_d]$ with specified conjugacy class of boundary monodromy. From this, we obtain $[\mathrm{CHur}^c_{n_1,\ldots,n_v,\overline{\mathbb{F}}_q}/S_d]$ also has $\sigma(n_1,\ldots,n_v)$ many components.

To conclude, we want to show the same holds over $\mathbb{F}_q$, instead of $\overline{\mathbb{F}}_q$. Indeed, the Frobenius action fixes the conjugacy class of boundary monodromy. Since there is a unique component of $[\mathrm{CHur}^c_{n_1,\ldots,n_v,\overline{\mathbb{F}}_q}/S_d]$ with specified conjugacy class of boundary monodromy, the Frobenius action must fix each such component over $\overline{\mathbb{F}}_q$. Hence, each such component of $[\mathrm{CHur}^c_{n_1,\ldots,n_v,\mathbb{F}_q}/S_d]$ over $\mathbb{F}_q$ is geometrically irreducible. □

Proof idea of Theorem 19.0.5. Let $c = S_d - \mathrm{id}$. We want to count the number of $\mathbb{F}_q$ points of $[\mathrm{Hur}^c/S_d]$ corresponding to connected covers having discriminant at most q^n for q sufficiently large and show it is given by

$$\left\| \sum_{\substack{n_1,\ldots,n_v \\ \sum_{i=1}^v n_i\Delta(c_i)=n}} \sigma(n_1,\ldots,n_v)\#\mathrm{Conf}_{n_1,\ldots,n_v,\mathbb{F}_q}(\mathbb{F}_q) \right\| + o(q^n).$$

To do this, we first show that $o(q^n)$ extensions have geometric Galois group strictly contained in S_d using Theorem 8.0.16 to bound the total number of such extensions. We can similarly use the same result to bound the number of geometrically connected S_d extensions with $n_1 < n/2$. Therefore, it suffices to compute the number of extensions with $n_1 > n/2$. We then know that the stable cohomology of $\mathrm{CHur}^c_{n_1,\ldots,n_v}$ is identified with the stable cohomology of $\mathrm{CHur}^{c/c_1}_{n_1,\ldots,n_v}$ for n_1 sufficiently large by Theorem 18.0.9. We can also bound the dimension of the unstable cohomology of these hurwitz spaces using Exercise 11.1.6. Therefore, by Lemma 3.5.1, up to a small error term, we can identify the point counts of $\mathrm{CHur}^{c/c_1}_{n_1,\ldots,n_v}$ with $\sigma(n_1,\ldots,n_k)$ times that of $\mathrm{Conf}_{n_1,\ldots,n_v}$ using that there are $\sigma(n_1,\ldots,n_k)$ geometrically irreducible components by Lemma 19.0.7. This gives the desired result. □

20. COMPUTING THE STABLE HOMOLOGY IN THE DEGREE 3 CASE

Our next goal is to prove Theorem 18.0.9, computing the stable homology of Hurwitz space. We will explain how to verify it in Theorem 20.2.1 in this section when $c \subset S_3$ consists of transpositions. We will take up the general case in the next section.

How can we prove the above result on computing the stable homology?

One possible answer is " the group completion theorem." Namely, $\pi_0 \operatorname{CHur}^c$ is generated by $\alpha_x, \alpha_y, \alpha_z$, and so the group completion theorem identifies $C_*(\operatorname{CHur}^c)[\alpha_x^{-1}, \alpha_y^{-1}, \alpha_z^{-1}] \simeq C_*(\Omega B \operatorname{CHur}^c)$.

Exercise 20.0.1. Using Theorem 13.0.2 in the case c is a rack with a single component. Identify the homology of $C_*(\operatorname{CHur}^c)[\alpha_c^{-1}]$ with a sum of $\mathbb{Z}$ many copies of the stable homology of CHur^c. That is, identify the ith homology of $C_*(\operatorname{CHur}^c)[\alpha_c^{-1}]$ with $\oplus_{\mathbb{Z}} H_i(\operatorname{CHur}^c_\infty)$, where $H_i(\operatorname{CHur}^c_\infty)$ denotes the vector space $H_i(\operatorname{CHur}^c_n)$ for $n > Ii + J$ with I and J the constants coming from the homological stability theorem. *Possible hint:* The reason there are $\mathbb{Z}$ many copies comes from the fact that the source has an additional grading in n, the number of branch points, while the target doesn't. However if one fixes the grading, this is saying that the above localization computes the stable homology.

By the above exercise, we can identify the stable homology with the (one of the graded pieces of) the homology of $C_*(\operatorname{CHur}^c)[\alpha_x^{-1}, \alpha_y^{-1}, \alpha_z^{-1}]$. We will only sketch this approach as we prefer a different one which will be more suitable for generalizations and relies on fewer outside references.

Remark 20.0.2. If you have been following developments in arithmetic statistics, you may be aware of the retracted paper [EVW12], which claimed to compute the stable homology of Hur^c in the above setting. However, they had a serious mistake. The mistake was precisely that they computed $C_*(\operatorname{CHur}^c)[\alpha_x^{-1}, \alpha_y^{-1}, \alpha_z^{-1}]$, but didn't show that this was identified with the stable homology of $C_*(\operatorname{CHur}^c)$. The main issue they missed is that they didn't address the content of Proposition 15.0.16 at all. Without that, there could have been a lot of additional stable homology on which x and y acted by 0 and on which z acted invertibly. In this way, one can view some version of Proposition 15.0.16 as the key obstruction which accounts for the error in the retracted paper [EVW12] of Ellenberg-Venkatesh-Westerland.

20.1. First method for the stable homology: the group completion theorem.

We first briefly sketch the original method for computing the stable homology of Hurwitz spaces, we will be somewhat brief because following this, we will present a second method more similar to the above presentation of showing that the homology stabilizes. However, we thought it would be nice to briefly indicate what goes into applying the group completion theorem. The main reference we learned this material from is [RW20, §5] especially [RW20, Corollary 5.4] and this argument is also spelled out in [LL24, Proposition 4.5.1]

Recall we are trying to compute $C_*(\mathrm{CHur}^c)[\alpha_x^{-1}, \alpha_y^{-1}, \alpha_z^{-1}] \simeq C_*(\Omega B\,\mathrm{CHur}^c)$. The group completion theorem identifies this with $C_*(\Omega B\,\mathrm{CHur}^c)$. First, using that multiplication $\alpha_x^2\alpha_y^2\alpha_z^2 : \mathrm{Hur}^c \to \mathrm{Hur}^c$ factors through CHur^c one can show $\Omega B\,\mathrm{CHur}^c \simeq \Omega B\,\mathrm{Hur}^c$, so we wish to compute $C_*(\Omega B\,\mathrm{Hur}^c, \mathbb{Z}[1/6])$. Next, a result of Etingof-Grana [EG03, Theorem 4.2] can be used to deduce $C_*(B\,\mathrm{Hur}^c, \mathbb{Z}[1/6]) \simeq C_*(B\,\mathrm{Conf}, \mathbb{Z}[1/6])$. From this, one wishes to show that each path component of $\Omega B\,\mathrm{Hur}^c$ has the same $\mathbb{Z}[1/6]$ homology as $B\,\mathrm{Conf}$. If we knew $B\,\mathrm{Hur}^c$ were homotopic to $B\,\mathrm{Conf}$, then we could deduce each component of their loops spaces has the same homology. Unfortunately, this is not the case. However, it turns out to be the case that the fundamental group of both these spaces acts trivially on their homotopy groups [FRS07, Proposition 5.2] and this turns out to imply that each component of their path spaces has the same $\mathbb{Z}[1/6]$ homology.

20.2. **The more self-contained method.** Above, we sketched one argument for computing the stable homology, which relied heavily on the results [EG03, Theorem 4.2] and [FRS07, Proposition 5.2]. This may be a bit unsatisfying because it doesn't really explain very well what is going on. We now present an alternative proof, which will be needed anyway for our application to the Poonen Rains conjectures later in the course. This alternative proof also has the virtues that it doesn't rely on prior results and is closely connected to our method for showing the homology stabilizes. That is, a variant of the method for showing the homology stabilizes will also let us compute its stable value. To start, let us state precisely the statement regarding the stable value of the homology.

Theorem 20.2.1. *For each component $Z \subset \mathrm{Hur}^c_n$ the map $\mathrm{Hur}^c_n \to \mathrm{Conf}_n$ induces an isomorphism $H_i(Z, \mathbb{Z}[1/6]) \to H_i(\mathrm{Conf}_n, \mathbb{Z}[1/6])$ for n sufficiently large relative to i.*

Remark 20.2.2. Combining this with Corollary 15.0.20 shows that Theorem 18.0.9 holds for $c \subset S_3$ transpositions. That is, the above map is an isomorphism not only when n is large enough relative to i, but in fact when $n > Ii + J$ for certain constants I and J.

The proof of Theorem 20.2.1 is extremely similar to that of Proposition 15.0.5. Recall the basic idea was to create a nullhomotopy by pulling a suitable element across the rectangle. Here, we will create a similar nullhomotopy, but we will have to work with sums of elements.

Sketch. We use α_c^{-1} as notation for inverting all of $\alpha_x, \alpha_y, \alpha_z$. By Remark 15.0.23, see also Exercise 20.0.1, all of $\alpha_x, \alpha_y, \alpha_z$ act invertibly on the stable homology. Therefore, it suffices to show $\mathrm{Hur}^c[\alpha_c^{-1}] \to \mathrm{Conf} \times_{\pi_0 \mathrm{Conf}} \pi_0\,\mathrm{Hur}^c[\alpha_c^{-1}]$ is a homology equivalence with $\mathbb{Z}[1/6]$ coefficients. Using Lemma 15.0.7, since these two spaces have the same π_0 by construction, it is enough to show

$$\pi_0 \operatorname{Hur}^c \otimes_{\operatorname{Hur}^c} \pi_0 \operatorname{Hur}^c[\alpha_c^{-1}] \simeq \pi_0 \operatorname{Hur}^c \otimes_{(\operatorname{Conf} \times_{\pi_0 \operatorname{Conf}} \pi_0 \operatorname{Hur}^c)} \pi_0 \operatorname{Hur}^c[\alpha_c^{-1}].$$

is a homology equivalence with $\mathbb{Z}[1/6]$ coefficients.

Using Lemma 20.2.6 below, this is equivalent to the homology equivalence

$$\begin{aligned}&(\pi_0 \operatorname{Hur}^c \otimes_{\operatorname{Hur}^c} \pi_0 \operatorname{Conf})[\alpha_c^{-1}] \times_{\pi_0 \operatorname{Conf}[\alpha_c^{-1}]} \pi_0 \operatorname{Hur}^c[\alpha_c^{-1}] \\ &\simeq (\pi_0 \operatorname{Conf} \otimes_{\operatorname{Conf}} \pi_0 \operatorname{Conf})[\alpha_c^{-1}] \times_{\pi_0 \operatorname{Conf}[\alpha_c^{-1}]} \pi_0 \operatorname{Hur}^c[\alpha_c^{-1}]\end{aligned}$$

with $\mathbb{Z}[1/6]$ coefficients. Equivalently, it is enough to show it is an equivalence on each component, and so we only need to prove the homology equivalence

$$(\pi_0 \operatorname{Hur}^c \otimes_{\operatorname{Hur}^c} \pi_0 \operatorname{Conf})[\alpha_c^{-1}] \simeq (\pi_0 \operatorname{Conf} \otimes_{\operatorname{Conf}} \pi_0 \operatorname{Conf})[\alpha_c^{-1}] \tag{20.1}$$

with $\mathbb{Z}[1/6]$ coefficients. For this equivalence, we again will produce a suitable nullhomotopy. We can change the basis x, y, z to the basis $\frac{x+y+z}{3}, x-y, x-z$. We call the first an averaged element and the latter two anti-averaged. (What comes next isn't completely rigorous, see Remark 20.2.3, but it can be made so.) Now, we can think of the fiber product as parameterizing configurations in the rectangle, labeled by one of the three types of elements above. Note this uses that we have inverted 6 (really we only needed to invert 3) to make sense of $\frac{x+y+z}{3}$. Now, we can think of $\pi_0 \operatorname{Conf}$ as keeping track of the number of $\frac{x+y+z}{3}$.

We now prove the above equivalence (20.1) in two steps, Lemma 20.2.4 where we first average the right module and then Lemma 20.2.5 where we complete the proof. □

Remark 20.2.3. The reason the proof above was only a sketch is that the nullhomotopy we will describe next is not really rigorous since we haven't explained how to make sense of pulling linear combinations of points across, and the equivalence is not going to exist at the level of spaces because we have to invert 6. It should be possible to make rigorous sense of this, or, alternatively, one can directly use the idea of this proof to construct a chain nullhomotopy.

Lemma 20.2.4. *There is a homology equivalence*

$$(\pi_0 \operatorname{Hur}^c \otimes_{\operatorname{Hur}^c} \pi_0 \operatorname{Conf})[\alpha_c^{-1}] \simeq (\pi_0 \operatorname{Conf} \otimes_{\operatorname{Hur}^c} \pi_0 \operatorname{Conf})[\alpha_c^{-1}]$$

with $\mathbb{Z}[1/6]$ *coefficients.*

Proof. We can think of elements of the bar construction above as consisting of points where there is an element of $\pi_0 \operatorname{Hur}^c$ on the left elements of Hur^c

in the rectangle and an element of $\pi_0 \operatorname{Conf}$ on the right. However, we will want to work in a slightly more general setting where we really allow linear combinations of these elements to occur in each place. So, for example, on the left module, there is a map $\pi_0 \operatorname{Hur}^c[\alpha_c^{-1}] \to \pi_0 \operatorname{Conf}[\alpha_{c/c}^{-1}]$ where the target is identified with the integers and the fibers of the map have three elements. For example the fiber over 0 is generated by $\mathrm{id}, \alpha_x\alpha_y^{-1}$ and $\alpha_x\alpha_z^{-1}$ and the fiber over 1 (corresponding to one point in configuration space) is $\alpha_x, \alpha_y, \alpha_z$. Therefore, there is a section to the map $\pi_0 \operatorname{Hur}^c[\alpha_c^{-1}] \to \pi_0 \operatorname{Conf}[\alpha_{c/c}^{-1}]$ which is a map $\pi_0 \operatorname{Conf}[\alpha_{c/c}^{-1}] \to \pi_0 \operatorname{Hur}^c[\alpha_c^{-1}]$ so that the composite map sends m to $m + m\alpha_x\alpha_y^{-1} + m\alpha_x\alpha_z^{-1}$. Because this is averaged over the fibers of the projection, the elements x, y, z will all act the same way on it. For example, if $m = \alpha_x$ then the composite of the projection and section will send it to $\alpha_x + \alpha_y + \alpha_z$ and you can see that if one multiplies this by any of α_x, α_y, or α_z one gets the same linear combination of elements of $\pi_0 \operatorname{Hur}^c[\alpha_c^{-1}]$. For example, multiplying by α_x one gets $\alpha_x^2 + \alpha_x\alpha_y + \alpha_x\alpha_z$ and multiplying by α_y one gets $\alpha_y^2 + \alpha_y\alpha_x + \alpha_y\alpha_z$ and these are equivalent because one can multiply the first by $\alpha_y^2\alpha_x^{-2}$ to get the second and the element $\alpha_y^2\alpha_x^{-2}$ is equivalent to the identity.

In any case, we now describe a "homotopy" which will send any element of the left hand side to its average value. We can separately perform three homotopies: The three homotopies first moves α_x along the top from right to left and then after it hits the right hand side it moves α_t along the top from left to right, for $t \in \{x, y, z\}$. Overall, all three only change the left value of m; the first has no effect, the second sends $m \mapsto m\alpha_x\alpha_y^{-1}$ and the third sends $m \mapsto m\alpha_x\alpha_z^{-1}$. If we could sum over these three values, we would have accomplished sending m to its averaged value

$$m \mapsto \frac{m + m\alpha_x\alpha_y^{-1} + m\alpha_x\alpha_z^{-1}}{3}$$

which would be the desired homotopy.

There are several issues with the above. The main one is that it is not continuous because if one moves an α_y to the left boundary, it can turn the α_x into an α_z. The way to fix this is to symmetrize the top by allowing 9 homotopies instead of just 3. These nine go by first sending α_s to the left and then α_t to the right, and now the collection of these 9 homotopies will be invariant under things hitting the boundary below. These 9 movements are visualized in Figure 10.

So we then just average over these 9 homotopies so that the total map will indeed replace m by its averaged value.

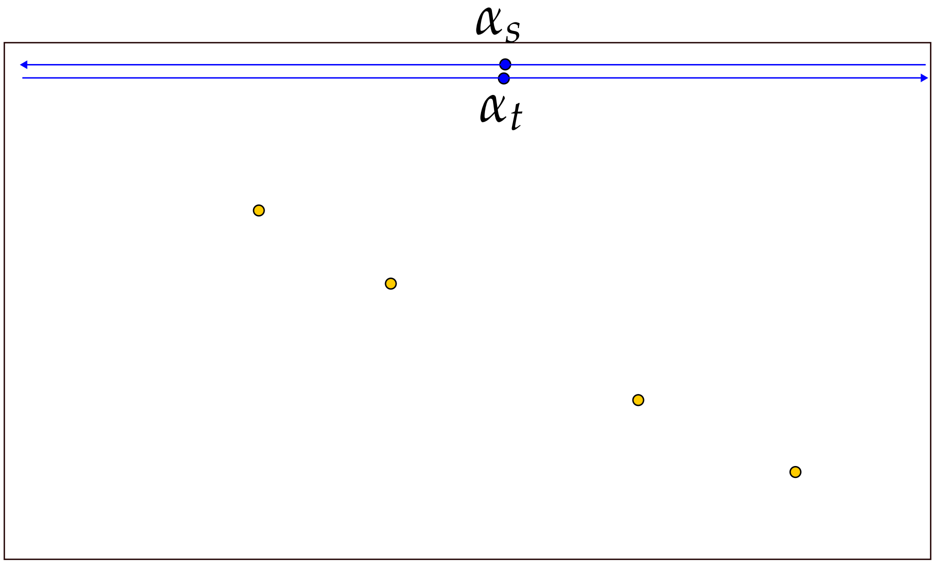

FIGURE 10. Here, we depict the first step of the nullhomotopy, making the label on the left averaged.

The above is not really possible at the level of spaces where one has to divide by 3, but it can be made sense of either at the level of chain complexes. □

Lemma 20.2.5. *In the setting above, we have*

$$(\pi_0 \operatorname{Hur}^c \otimes_{\operatorname{Hur}^c} \pi_0 \operatorname{Conf})[\alpha_c^{-1}] \simeq (\pi_0 \operatorname{Conf} \otimes_{\operatorname{Conf}} \pi_0 \operatorname{Conf})[\alpha_c^{-1}]$$

Proof. We have already shown we can replace the source with $(\pi_0 \operatorname{Conf} \otimes_{\operatorname{Hur}^c} \pi_0 \operatorname{Conf})[\alpha_c^{-1}]$ in Lemma 20.2.4. So we now want to replace the elements in the middle of the rectangle with their averaged values, which will be $\frac{x+y+z}{3}$. To do this, we can change basis so the labels are either $x-y, x-z$ or $\frac{x+y+z}{3}$. We will start at the top, and find the highest element that is not averaged, i.e. it is either $x-y$ or $x-z$. Say this is element i from the top. We then simultaneously perform several homotopies: The structure group of the rack in this case is S_3. If we could replace $x-y$ by $\frac{1}{\# S_3}\sum_{g \in S_3} g \triangleright (x-y)$ this would be 0 because x and y would both get replaced by $\frac{x+y+z}{3}$ and so the difference would be 0. So, we just need to perform homotopies that enact every element of S_3. If we want conjugation by $x = (12)$, we just move x above the ith element $x-y$ and x below the ith element. This is visualized in Figure 11.

Enacting multiplication by $x \cdot y$ is a bit trickier, but it is not so bad. We can first move y above the ith element and then below the ith element, and, following that, move x above the ith element and then below the ith

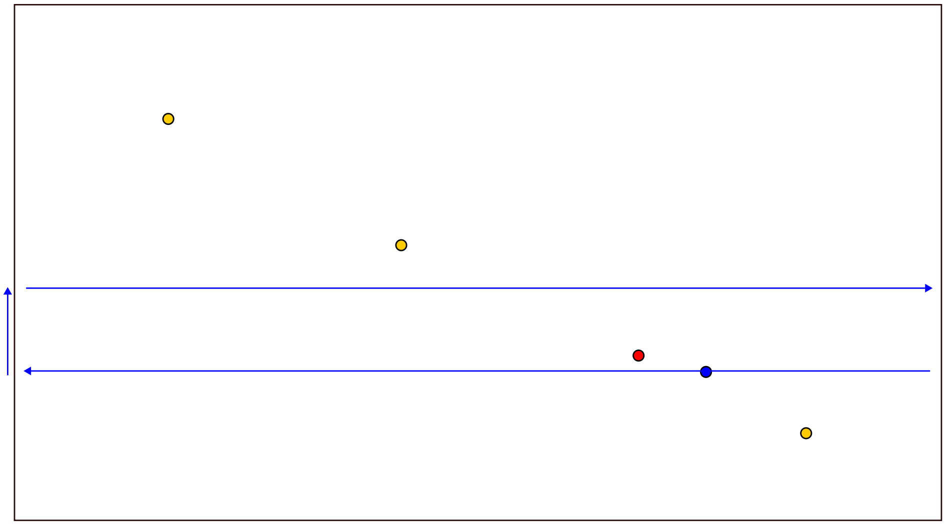

FIGURE 11. The red dot depicts the highest antiaveraged element. We move the blue dot around it to enact multiplication by the value of the value of the blue dot on the red dot. Writing elements of the reduced structure group as products allows us to enact multiplication by any element of the reduced structure group by a sequence of these moves, and then we average over elements of the reduced structure group.

element. So, for each element of S_3, we choose a word in x, y, z whose action corresponds to that element, and then we move that around the ith element. We then average over the 6 elements in S_3, which may seem to send the above element to 0.

However, this map will not quite be continuous. The reason is as explained in the proof of Lemma 20.2.4. Above, we have chosen a word for each element, which may not be invariant under elements below hitting the boundary. However, we can fix this by adding in more words, which are the orbits of the words we chose under the action of elements hitting the boundary below. In this case, there will be at most 6 representatives for each word, where for each word $w_1 \cdots w_t$ we add in $(g^{-1}w_1g)\cdots(g^{-1}w_tg)$ for each $g \in S_3$. So there will be at most 36 terms in the homotopy. Now, summing these up and averaging will be invariant. So, if there is any antiaveraged element, i.e. $x - y$ or $x - z$ this homotopy will kill it, and we will obtain only elements where all elements in the rectangle are averaged in the end. □

Lemma 20.2.6. *We can identify*

$$\pi_0 \operatorname{Hur}^c \otimes_{\operatorname{Hur}^c} \pi_0 \operatorname{Hur}^c[\alpha_c^{-1}] \simeq (\pi_0 \operatorname{Hur}^c \otimes_{\operatorname{Hur}^c} \pi_0 \operatorname{Conf})[\alpha_c^{-1}] \times_{\pi_0 \operatorname{Conf}[\alpha_c^{-1}]} \pi_0 \operatorname{Hur}^c[\alpha_c^{-1}]$$

and

$$\begin{aligned} &\pi_0 \operatorname{Hur}^c \otimes_{(\operatorname{Conf} \times_{\pi_0 \operatorname{Conf}} \pi_0 \operatorname{Hur}^c)} \pi_0 \operatorname{Hur}^c[\alpha_c^{-1}] \\ &\simeq (\pi_0 \operatorname{Conf} \otimes_{\operatorname{Conf}} \pi_0 \operatorname{Conf})[\alpha_c^{-1}] \times_{\pi_0 \operatorname{Conf}[\alpha_c^{-1}]} \pi_0 \operatorname{Hur}^c[\alpha_c^{-1}] \end{aligned}$$

Proof. This is one of the most technical parts of the proof so we won't belabor it too much, but just try to give the idea. Let's just explain the second identification. We have

$$\begin{aligned} &\pi_0 \operatorname{Hur}^c \otimes_{(\operatorname{Conf} \times_{\pi_0 \operatorname{Conf}} \pi_0 \operatorname{Hur}^c)} \pi_0 \operatorname{Hur}^c[\alpha_c^{-1}] \\ &\simeq \left(\pi_0 \operatorname{Conf} \times_{\pi_0 \operatorname{Conf}} \pi_0 \operatorname{Hur}^c\right) \otimes_{(\operatorname{Conf} \times_{\pi_0 \operatorname{Conf}} \pi_0 \operatorname{Hur}^c)} \left(\pi_0 \operatorname{Conf} \times_{\pi_0 \operatorname{Conf}} \pi_0 \operatorname{Hur}^c\right) [\alpha_c^{-1}] \\ &\simeq (\pi_0 \operatorname{Conf} \otimes_{\operatorname{Conf}} \pi_0 \operatorname{Conf})\, [\alpha_c^{-1}] \times_{\left(\pi_0 \operatorname{Conf} \otimes_{\pi_0 \operatorname{Conf}} \pi_0 \operatorname{Conf}\right)[\alpha_c^{-1}]} \left(\pi_0 \operatorname{Hur}^c \otimes_{\pi_0 \operatorname{Hur}^c} \pi_0 \operatorname{Hur}^c\right) [\alpha_c^{-1}] \\ &\simeq (\pi_0 \operatorname{Hur}^c \otimes_{\operatorname{Conf}} \pi_0 \operatorname{Conf})\, [\alpha_c^{-1}] \times_{\pi_0 \operatorname{Conf}[\alpha_c^{-1}]} \pi_0 \operatorname{Hur}^c[\alpha_c^{-1}]. \end{aligned}$$

The second step involved commuting tensor products and pullbacks. This is not formal, and [BF06, Theorem B.4], whose hypotheses are nontrivial to verify in this case (one needs to verify a certain map of simplicial sets is a Kan fibration, namely the map from the right factor to the bottom factor in the third line above. One can do directly, but its a bit combinatorially involved.) More details are spelled out in [LL25b]. □

21. PROOF OF THE STABLE HOMOLOGY IN ALL DIRECTIONS

We just saw above how Theorem 18.0.9 was useful to verify Bhargava's conjecture. We'd next like to explain the proof of Theorem 18.0.9. Let us now restate the theorem.

Theorem 21.0.1. *Let c be a rack with components $c = c_1 \cup \cdots \cup c_v$. Let $c' \subset c$ denote the first j components of c. There are constants I and J so that for $n_1 > Ii + J$ and $Z \subset \operatorname{CHur}^c_{n_1,\ldots,n_v}$ a component which maps to a corresponding component $\overline{Z} \subset \operatorname{CHur}^{c/c'}_{n_1,\ldots,n_v}$. Then, the map $Z \to Z'$ induces an isomorphism $H_i(Z, \mathbb{Z}[1/\#G_c^0]) \simeq H_i(Z', \mathbb{Z}[1/\#G_c^0])$ when $n_1, \ldots, n_j > Ii + J$.*

The proof of the above result is very similar to Theorem 20.2.1.

We already know the cohomology of $\operatorname{CHur}^c_{n_1,\ldots,n_v}$ stabilizes for n_1 sufficiently large and so we only want to identify its stable value with that of configuration space. We can rephrase our goal as follows. Consider the map $\operatorname{CHur}^c \to \operatorname{CHur}^{c/c'} \times_{\pi_0 \operatorname{CHur}^{c/c'}} \pi_0 \operatorname{CHur}^c$. Effectively, this takes each component of CHur^c and replaces it with the corresponding component of $\mathrm{CHur}^{c/c'}$. We want to show this is an equivalence on components with n_1 sufficiently large. Equivalently, we can rephrase our goal as simply showing that $\mathrm{CHur}^c[\alpha_{c'}^{-1}] \to \mathrm{CHur}^{c/c'}[\alpha_{c'}^{-1}] \times_{\pi_0 \mathrm{CHur}^{c/c'}[\alpha_{c'}^{-1}]} \pi_0 \mathrm{CHur}^c[\alpha_{c'}^{-1}]$ is a homology equivalence inverting $\# G_c^0$ in the coefficients. Actually, we will show the slightly stronger statement that

$$\mathrm{Hur}^c[\alpha_{c'}^{-1}] \to \mathrm{Hur}^{c/c'}[\alpha_{c'}^{-1}] \times_{\pi_0 \mathrm{Hur}^{c/c'}[\alpha_{c'}^{-1}]} \pi_0 \mathrm{Hur}^c[\alpha_{c'}^{-1}]$$

is a homology equivalence inverting $\# G_c^0$ in the coefficients. This is slightly stronger because we replace CHur with Hur, so we remove the condition that the covers are connected. To verify this is an equivalence, we can check it after tensoring over Hur^c with $\pi_0 \mathrm{Hur}^c$, using Lemma 15.0.7. That is, taking $R = \mathrm{Hur}^c[\alpha_{c'}^{-1}]$ and $S := \mathrm{Hur}^{c/c'}[\alpha_{c'}^{-1}] \times_{\pi_0 \mathrm{Hur}^{c/c'}[\alpha_{c'}^{-1}]} \pi_0 \mathrm{Hur}^c[\alpha_{c'}^{-1}]$, since R and S have the same π_0 by construction, it is enough to check

$$\begin{aligned} &\pi_0 \mathrm{Hur}^c[\alpha_{c'}^{-1}] \otimes_{\mathrm{Hur}^c[\alpha_{c'}^{-1}]} \pi_0 \mathrm{Hur}^c[\alpha_{c'}^{-1}] \\ &\simeq \pi_0 \mathrm{Hur}^c[\alpha_{c'}^{-1}] \otimes_{\mathrm{Hur}^{c/c'}[\alpha_{c'}^{-1}] \times_{\pi_0 \mathrm{Hur}^{c/c'}[\alpha_{c'}^{-1}]} \pi_0 \mathrm{Hur}^c[\alpha_{c'}^{-1}]} \pi_0 \mathrm{Hur}^c[\alpha_{c'}^{-1}] \end{aligned}$$

is a homology equivalence inverting $\# G_c^0$ coefficients.

We will now simplify both these sides of the above equation:

Lemma 21.0.2. *We can identify*

$$\begin{aligned} &\pi_0 \mathrm{Hur}^c[\alpha_{c'}^{-1}] \otimes_{\mathrm{Hur}^c[\alpha_{c'}^{-1}]} \pi_0 \mathrm{Hur}^c[\alpha_{c'}^{-1}] \\ &\simeq \left(\pi_0 \mathrm{Hur}^c \otimes_{\mathrm{Hur}^c} \pi_0 \mathrm{Hur}^{c/c'}\right) [\alpha_{c'}^{-1}] \times_{\pi_0 \mathrm{Conf}[\alpha_{c'}^{-1}]} \pi_0 \mathrm{Hur}^c[\alpha_{c'}^{-1}] \end{aligned}$$

and

$$\begin{aligned} &\pi_0 \mathrm{Hur}^c[\alpha_{c'}^{-1}] \otimes_{\mathrm{Hur}^{c/c'}[\alpha_{c'}^{-1}] \times_{\pi_0 \mathrm{Hur}^{c/c'}[\alpha_{c'}^{-1}]} \pi_0 \mathrm{Hur}^c[\alpha_{c'}^{-1}]} \pi_0 \mathrm{Hur}^c[\alpha_{c'}^{-1}] \\ &\simeq \left(\pi_0 \mathrm{Hur}^{c/c'} \otimes_{\mathrm{Hur}^{c/c'}} \pi_0 \mathrm{Hur}^{c/c'}\right) [\alpha_{c'-1}] \times_{\pi_0 \mathrm{Conf}[\alpha_{c'}^{-1}]} \pi_0 \mathrm{Hur}^c[\alpha_{c'}^{-1}] \end{aligned}$$

The proof of the above lemma is quite technical, but is essentially the same as that of Lemma 20.2.6. We will omit further details of its proof.

We wanted to identify the left hand sides of the equations in Lemma 21.0.2 on homology (after inverting $|G_c^0|$), and hence it suffices to identify the right hand sides. Note that both the right hand sides have the same base and right parts of the fiber product appearing, so it is enough to identify the left parts of the fiber products. That is, it is enough to show the following.

Proposition 21.0.3. *We have*

(21.1)
$$\left(\pi_0 \operatorname{Hur}^c \otimes_{\operatorname{Hur}^c} \pi_0 \operatorname{Hur}^{c/c'}\right)[\alpha_{c'}^{-1}, |G^0_{c'}|^{-1}] \simeq \left(\pi_0 \operatorname{Hur}^{c/c'} \otimes_{\operatorname{Hur}^{c/c'}} \pi_0 \operatorname{Hur}^{c/c'}\right)[\alpha_{c'^{-1}}, |G^0_{c'}|^{-1}]$$

The idea for proving this is quite similar to the idea used for proving Theorem 20.2.1.

Remark 21.0.4. Again, the proof we outline here won't be completely rigorous, because we will imagine we are free to invert G^0_c but still work with spaces. Therefore, the details we give in the proof won't completely make sense, but just try to think through them as best you can. A meta-exercise could be to make rigorous sense of the following argument on the level of spaces. Alternatively, you can check out the implementation we give of this on the level of chains in our paper [LL25b, Proposition 7.2.8]. But personally, I think this space level argument is more enlightening, despite not being fully rigorous.

Definition 21.0.5. Let $U : k\{c\} \to k\{c\}$ be the operator given by $U(x) := \sum_{g \in G^0_{c'}} gx$ for $x \in c$, and extend U by linearity to all of $k\{c\}$. An element v of $k\{c\}$ is *averaged* if $U(v) = v$ and *antiaveraged* if $U(v) = 0$. Let $G^{c'}_{\pi_0} := \ker(\pi_0 \operatorname{Hur}^{c'}[\alpha_{c'}^{-1}] \to \pi_0 \operatorname{Hur}^{c'/c'}[\alpha^{-1}_{c'/c'}])$. Also, let $G^{c'}_c$ be the subgroup of the structure group G^0_c of c generated by the actions of elements in $c' \subset c$. Similarly, there is an operator U on $\pi_0 \operatorname{Hur}^c[\alpha_{c'}^{-1}]$ which sends $m \mapsto \frac{1}{|G^{c'}_{\pi_0}|} \sum_{x \in G^{c'}_{\pi_0}} m \cdot x$. For $v \in k\{\pi_0 \operatorname{Hur}^c[\alpha_{c'}^{-1}]\}$ we say v is *averaged* if $U^{\pi_0}_{c'}(v) = v$ and *antiaveraged* if $U^{\pi_0}_{c'}(v) = 0$.

Example 21.0.6. In the case $c = \{x, y, z\}$ is the transpositions in S_3, the only averaged elements are linear multiples of $\frac{x+y+z}{3}$. The antiaveraged elements include things like $x - y, x - z, x + y - 2z$.

Proof. Imagine that we had the following space level model of the two spaces appearing in the statement of Proposition 21.0.3. The source consists of rectangles labeled by configurations of points which are linear combinations of Hur^c in the middle, allowing $\#G^0_c$ in the denominators, which are either averaged or antiaveraged, and also with labels in $\pi_0 \operatorname{Hur}^c[\alpha_{c'}^{-1}]$ on the left and right sides. We also impose that the element on the right side is averaged. Points are allowed to collide into the boundaries, in which case they act on the boundaries.

The target consists of a similar description, but we only allow averaged elements in the middle, bottom, and left. (The above descriptions are exactly true before we invert $\#G^0_c$.)

We now wish to construct a homotopy that kills any configuration where there is some antiaveraged element, either on the left boundary or on the interior of the rectangle. We first describe a homotopy that makes the left boundary averaged. The idea for what we are doing here was depicted in Figure 10. For each element $g \in G^{c'}_{\pi_0}$, write it as a product $g = y_\ell z_\ell^{-1} y_{\ell-1} z_{\ell-1}^{-1} \cdots y_1 z_1^{-1}$. Then, we can enact multiplication by g on the configuration by moving y_i across to the right, then z_i across to the left, for i traversing from ℓ down to 1. Averaging this over all $g \in G^{c'}_{\pi_0}$ will not be continuous, since if one moves some element of c' to the left or right boundary, it can change the way we write g as a product. Namely, if we move x to the right, it can send the expression for $x \triangleright g$ to the expression $g = (x \triangleright y_\ell)(x \triangleright z_\ell)^{-1} \cdots (x \triangleright y_1)(x \triangleright z_1)^{-1}$. Therefore, if we average over all g and all expressions for g of the form $(x \triangleright y_\ell)(x \triangleright z_\ell)^{-1} \cdots (x \triangleright y_1)(x \triangleright z_1)^{-1}$, the homotopy described above will then be continuous. There is one more small issue, which is that it is only continuous for elements $x \in c'$, and not for $x \in c - c'$. However, we can deal with this issue by introducing a filtration by the number of elements in $c - c'$ in the interior of the rectangle. We can verify the homology equivalence on the associated graded of this filtration.

Next, we describe the homotopy that makes the elements in the middle of the rectangle averaged. The idea for what we are doing here was depicted in Figure 11. This is extremely similar to that making the left boundary averaged. Namely, for each $g \in G^{c'}_c$, we write $g = w_1 \cdots w_t$. We consider a point of Hurwitz space where we allow labelings in $k\{c\}$ and choose it so that the left and right boundaries are averaged, the top $i-1$ elements are averaged, and the ith element $v_i \in k\{c\}$ is antiaveraged. When we then move w_1 below the ith element and then back directly above the ith element, it will send $v_i \mapsto w_1 \triangleright v_i$. Repeating this for $w_2, \ldots, w_t$ will cumulatively send $v_i \mapsto g \triangleright v_i$, while maintaining the rest of the labels. Averaging over all $g \in G^{c'}_c$ will then send $v_i \mapsto U_{c'}(v_i) = 0$ since v_i is antiaveraged.

This will not be continuous for two reasons. First, the expressions for g will not be invariant under the action of elements of c'. We can fix this by throwing in expressions of the form $(x \triangleright w_1) \cdots (x \triangleright w_t)$. for $x \in G^{c'}_c$ and averaging over those instead. Second, this will not be invariant under the action of elements of $k\{c - c'\}$. Again, we can fix this by filtering by the number of elements of $k\{c - c'\}$. □

22. INTRODUCTION TO THE POONEN-RAINS HEURISTICS

Having had a significant discussion about the Cohen-Lenstra heuristics and the closely related Malle's conjecture, we next turn to the Poonen-Rains heuristics. Loosely speaking, so far, we have understood Hurwitz spaces, which we can think of as certain topological monoids where one can take two covers of a disc and form another cover of the same disc by putting those two discs in a bigger disc. The Poonen-Rains conjectures are related not to these topological monoids, but to modules over these monoids. Therefore, it is natural to try to prove a homological stability theorem for Hurwitz space modules. Before explaining how this goes, let's review the Poonen-Rains heuristics.

Definition 22.0.1. An elliptic curve over a field K is a smooth proper curve E of genus 1 with a marked point $\operatorname{Spec} K \to E$.

Theorem 22.0.2 (Mordell-Weil). *For K a global field, and E an elliptic curve over K, $E(K)$ is a finitely generated abelian group.*

Definition 22.0.3. With E an elliptic curve over a global field, by the Mordell-Weil theorem, we can write $E(K) = \mathbb{Z}^{\operatorname{rk} E} \oplus T$ for T a finite torsion group. The number $\operatorname{rk} E$ is called the *rank* of E.

Before the Poonen-Rains heuristics, which involve conjectures for the distribution of the rank of E and also the Selmer groups of E, Goldfeld made a conjecture just on the distribution of the rank of E. These involve quadratic twist families, which we next define.

Definition 22.0.4. Recall that for L/M a finite locally free extension of schemes and X a scheme over L the Weil restriction $\operatorname{Res}_{L/M} X$ is functor over M scheme defined by $(\operatorname{Res}_{L/M} X)(T) = X(T \times_M L)$ for any M scheme T.

Remark 22.0.5. In good situations (such as when X is quasi-projective over L) the Weil restriction above is representable by a scheme. I.e., there exists a scheme, which we also call $\operatorname{Res}_{L/M} X$, whose T points are also given by $X(T \times_M L)$. See [BLR90, §7.6, Theorem 4].

Definition 22.0.6. Let E be an elliptic curve over a global field K. For each quadratic extension K'/K the *quadratic twist* of E along K'/K is $(\operatorname{Res}_{K'/K} E)/E$, where $E \to \operatorname{Res}_{K'/K} E_{K'}$ is the natural map coming from the adjunction between Weil restriction from K' to K and the base change from K' to K.

In the function field case, we define the height of the quadratic twist of E along K'/K to be the degree of $\Omega_{C_{K'}/C_K}$ for C_K the smooth proper curve over a finite field $\mathbb{F}_q$ with function field K. (It might be more natural to define this to be the log height and define $\operatorname{disc}(K'/K)$ to be the height, which also makes sense over number fields, but we will mostly be concerned with function fields in these notes. It also might be more natural to omit height contributions from certain chosen infinite places, but we won't worry about this minor optimization in these notes.) The *family of quadratic twists of E*

of height h is the set of quadratic twists of E which have height h and have $\Omega_{C_{K'}/C_K}$ prime to the discriminant of E over K.

Definition 22.0.7. Let C be a smooth proper curve over a base B. Let $Z \subset C$ be a divisor étale over B and $U := C - Z$. Define $\mathrm{QTwist}^n_{U/B}$ to be the stack parameterizing double covers of C branched over a degree n divisor disjoint from Z.

We call this the *stack of quadratic twists* because if E is an elliptic curve over U, we can form the universal quadratic twist of E over $\mathrm{QTwist}^n_{U/B}$ whose fiber over a geometric point of $\mathrm{QTwist}^n_{U/B}$ corresponding to a double cover $f : X \to U$ is the quadratic twist of E along that quadratic extension, given by $\mathrm{Res}_{X/U}(E \times_U X)/E$.

Remark 22.0.8. When we study the quadratic twists as above in relation to the Poonen-Rains heuristics, we will typically consider an elliptic curve E which is smooth and proper over U, so we can also think of $Z \subset C$ as above as corresponding to the discriminant of the elliptic curve E, or, more generally a subscheme of C that contains the discriminant.

Exercise 22.0.9. Show $\mathrm{rk}_K((\mathrm{Res}_{K'/K} E)/E) = \mathrm{rk}_{K'} E_{K'} - \mathrm{rk}_K E$.

Exercise 22.0.10. Suppose we are given an elliptic curve E over $\mathbb{F}_q(t)$. If the characteristic is not 2, we can write E in the form $E : y^2 = x^3 + C(t)x^2 + A(t)x + B(t)$. Moreover, when the characteristic is not 3, we can take $C(t) = 0$.

(1) Show that, in this form, the quadratic twist family is identified with the set
$$\{E_f := f(t)y^2 = x^3 + C(t)x^2z + A(t)xz^2 + B(t)z^3\}$$
for $f(t) \in \mathbb{F}_q[t]$ squarefree and prime to the discriminant of E.
(2) Show that the height of f is the smallest even number bigger than or equal to the degree of f. (Though it would also be reasonable to redefine the height to just be the degree of f in this case.)
(3) There is a small redundancy above, in that multiple squarefree polynomials correspond to the same quadratic twist. Characterize this redundancy.

Conjecture 22.0.11 (Goldfeld, minimalist conjecture)**.** Fix an elliptic curve E over a global field K. When one orders the quadratic twists of E by height, half have rank 0, half have rank 1, and 0% have rank > 1.

Strictly speaking, the above conjecture of Goldfeld turns out to be a bit too general. (It is not the version Goldfeld originally stated.)

Remark 22.0.12. It is possible that every quadratic twist of a given curve has Selmer rank of a fixed parity. (The Selmer rank is technically different from the rank, but it is conjecturally equivalent.) Hence, quadratic twists of such a curve do not satisfy the minimalist conjecture. A specific example is given by the elliptic curve $y^2 = \lambda(\lambda-1)x(x-1)(x-\lambda)$, over $\mathbb{F}_q(\lambda)$, where q is a prime which is 1 mod 4. This is a variant of the Legendre family. See [EL24, Remark 1.4.1] for further details.

The Poonen-Rains heuristics are about refining the statistical model for elliptic curves, and capturing more data relevant to them, involving their Selmer group. Recall that for ν a positive integer, the ν Selmer group is a finite abelian $\mathbb{Z}/\nu\mathbb{Z}$ module associated to an elliptic curve over a global field.

Definition 22.0.13. Let K be a function global field, and let v index the closed points of C_K, the smooth proper curve with function field K. For E an elliptic curve over K, the multiplication by ν exact sequence induces the sequences on étale cohomology
(22.1)

$$\begin{tikzcd}
0 \arrow[r] & E(K)/\nu E(K) \arrow[r] \arrow[d] & H^1(\operatorname{Spec} K, E[\nu]) \arrow[r, "\beta"] \arrow[d] \arrow[dr, "\alpha"] & H^1(\operatorname{Spec} K, E)[\nu] \arrow[r] \arrow[d, "\gamma"] & 0 \\
0 \arrow[r] & \prod_{v \in C_K} E(K_v)/\nu E_v(K_v) \arrow[r] & \prod_v H^1(\operatorname{Spec} K_v, E_v[\nu]) \arrow[r] & \prod_v H^1(\operatorname{Spec} K_v, E_v)[\nu] \arrow[r] & 0.
\end{tikzcd}$$

The ν-**Selmer** group of E is

$$\operatorname{Sel}_\nu(E) := \ker \alpha$$

The **Tate-Shafarevich** group is

$$Ш(E) := \ker\left(H^1(\operatorname{Spec} K, E) \to \prod_v H^1(\operatorname{Spec} K_v, E_v)\right).$$

Remark 22.0.14. Once can make a similar definition for number fields where one lets v range over places of K (which is the analog of the closed points of C_K in the function field case).

Exercise 22.0.15. Using the definitions above, show there is an exact sequence

$$0 \longrightarrow E(K)/\nu E(K) \longrightarrow \operatorname{Sel}_\nu(E) \longrightarrow Ш(E)[\nu] \longrightarrow 0 \tag{22.2}$$

Where $Ш(E)[\nu]$ denotes the ν torsion elements in $Ш(E)$.

Remark 22.0.16. By the above exercise, the Selmer group interpolates the rank of E and the Tate-Shafarevich group of E, which is an important invariant of E whose finiteness would imply the Birch and Swinerton-Dyer conjecture over function fields.

Now, there is the simplest case of the Poonen-Rains conjecture on Selmer groups.

Conjecture 22.0.17 (Version of Bhargava–Shankar [BS13a, Conjecture 4] and Poonen–Rains [PR12, Conjecture 1.4(b)])**.** Let ℓ be a prime number. Let $B = \operatorname{Spec} \mathbb{F}_q$ and let U be a smooth curve over B. When elliptic curves in a quadratic twist family are ordered by height,

$$\lim_{n\to\infty} \mathbb{E}_{f \in \mathrm{QTwist}^n_{U/B}(\mathbb{F}_q)}(\#\operatorname{Sel}_\ell(E_f)) = \ell + 1.$$

Remark 22.0.18. Throughout these notes, we will mostly focus on quadratic twist families. However, the Poonen-Rains conjectures were originally made for the universal family parameterizing all elliptic curves. The reason we focus more on quadratic twist families here is that we can say more about the relevant stable homology groups. However, at least the number of stable components (so the relevant 0th homology groups) relevant to the Poonen-Rains heuristics are understood in [Lan21] and [FLR23].

Remark 22.0.19. There is a growing literature on this in the universal family case. Over $\mathbb{Q}$ for $\ell \leq 5$, Bhargava and Shankar have verified this. Bhargava-Shankar-Swaminathan also computed the second moment of 2-Selmer groups, meaning that they computed the average value of $\#\operatorname{Sel}_2(E)^2$. When $\ell = 3$ there is work of de Jong over function fields and when $\ell = 2$ Hồ, Lê Húng and Ngô, as well as separate work by Achenjang.

For quadratic twist families, there is work when $\ell = 2$ by many people and powers of 2 by Smith, but little for odd ℓ.

Here is the generalization to the composite case:

Conjecture 22.0.20 (Version of Bhargava–Shankar [BS13a, Conjecture 4] and Poonen–Rains [PR12, Conjecture 1.4(b)])**.** Let ν be a positive integer. Let $B = \operatorname{Spec} \mathbb{F}_q$ and let U be a smooth curve over B. When elliptic curves in a quadratic twist family are ordered by height,

$$\lim_{n\to\infty} \mathbb{E}_{f \in \mathrm{QTwist}^n_{U/B}(\mathbb{F}_q)}(\#\operatorname{Sel}_\nu(E_f)) = \sum_{s \mid \nu} s.$$

The following observation of Bhargava provides some relation between the Poonen-Rains conjectures and the minimalist conjecture.

Corollary 22.0.21. *Let $P_q^{\leq d}$ denote the proportion of elliptic curve of rank ≥ 2 over $\mathbb{F}_q(t)$ of height up to d. If Conjecture 22.0.20 were true,*

$$\lim_{d \to \infty} P_q^{\leq d} = 0$$

That is, 100% *of elliptic curves over $\mathbb{F}_q(t)$ have rank at most* 1*.*

Proof. Suppose some proportion ϵ have rank ≥ 2. Since $\nu^{\operatorname{rk} E} \leq \# \operatorname{Sel}_\nu(E)$, we have,

$$\varepsilon \nu^2 \leq \operatorname{Avg}(\nu^{\operatorname{rk} E}) \leq \operatorname{Avg}\left(\# \operatorname{Sel}_\nu(E)\right) = \sum_{s \mid \nu} s.$$

For primes ν, this says $\varepsilon \nu^2 \leq \nu + 1$, and now let ν grow arbitrarily large. □

Here is an exercise which illustrates what we can deduce about the rank if we only know the average size of the ℓ Selmer group for a single prime ℓ, instead of for all primes ℓ.

Exercise 22.0.22. Suppose we have a family of elliptic curves over which the average size of the ℓ-Selmer group tends to $\ell + 1$ as the height tends to ∞. Show the average rank of elliptic curves in this family is at most $\frac{\ell+1}{\ell}$.

Remark 22.0.23 (Three heuristics for the average size of Selmer groups)**.** When computing the average size of Selmer groups, it is natural to ask if there is some deeper reason for why the average size of ν-Selmer groups should be $\sum_{s \mid \nu} s$.

(1) In [BS13a, Conjecture 4], the conjecture is based on the fact that the Tamagawa number PGL_s is s, and their average size is revealed to be the sum of the Tamagawa numbers of PGL_s for $s \mid \nu$. The is used in the proof of [BS15a, BS15b, BS13b], [HLHN14], and [Ach23].

(2) A monodromy perspective (which we will discuss more later) reveals the average size of ν-Selmer groups is the number of orbits of a certain orthogonal group $O(Q)$ associated to a quadratic form Q on a rank $12d - 4$ free $\mathbb{Z}/\nu\mathbb{Z}$ module V. Such orbits are in bijection with geometric components (i.e., irreducible components over an algebraic closure) of a moduli space for Selmer elements, which we call the ν-Selmer space. There are $\sum_{s \mid \nu} s$ such orbits. It is easiest to see this in the case $\nu = \ell$ is prime, and we want to show there are $\ell + 1$ orbits. There is the 0 vector and then the ℓ level sets, corresponding to those $v \in V$ with $Q(V) = t$ for $t \in \mathbb{Z}/\ell\mathbb{Z}$.

(3) Yet a third heuristic appears in [dJ02] for 3-Selmer groups, in [Vak01] for 2-Selmer groups, and in [dJF11, Theorem 5.4] for ν-Selmer groups. These works suggest that the average size of the ν-Selmer group

should equal the number of balanced (also called rigid) rank s projective bundles over $\mathbb{P}^1$ for $s \mid \nu$. Indeed, the balanced rank m projective bundles are all of the form $\operatorname{Proj}_{\mathbb{P}^1} \operatorname{Sym}^\bullet (\mathscr{O}^{\oplus a} \oplus \mathscr{O}(-1)^{\oplus m-a})$ for $1 \leq a \leq m$, and so there are m total such bundles. Altogether, there are $\sum_{s|\nu} s$ such bundles as s ranges over the divisors of n.

(4) Fix a global field K, and elliptic curve E over K and a prime ℓ. A fourth heuristic has to do with modelling the Selmer group as an intersection of two maximal isotropic subspaces of a $2r$ dimensional vector space V over $\mathbb{Z}/\ell\mathbb{Z}$ in the limit that $r \to \infty$. This was the original heuristic introduced by Poonen and Rains, and is based off the fact that the Selmer group is actually an intersection of two such subspaces, albeit in the infinite dimensional space which is the restricted product $\prod_v'(H^1(K_v, E_v[\ell]), H^1(\mathscr{O}_v, \mathscr{E}_v[\ell]))$ for v ranging over places of the global field K and $\mathscr{E}_v$ the Néron model of E over $\mathscr{O}_v$.

In fact, the Poonen-Rains conjectures can be generalized to not only predict the average sizes of Selmer groups, but even give a distribution. The conjectures of Poonen-Rains were generalized in the article [BKL$^+$15], whose authors have initials BKLPR, so we call the ensuing distribution the BKLPR distribution, as described in the next definition.

Definition 22.0.24. Let O_{2r} denote an orthogonal group of rank $2r$ over the finite field $\mathbb{F}_\ell$ and define the BKLPR distribution, $\mathrm{Sel}_\ell^{BKLPR}$ by

$$\operatorname{Prob}(\dim \mathrm{Sel}_\ell^{BKLPR} = \alpha) = \lim_{r \to \infty} \operatorname{Prob}_{M \in \mathrm{O}_{2r}}(\dim \ker(M - \mathrm{id}) = \alpha).$$

Moreover, one can predict the joint distribution of rank and Selmer groups, by stipulating that the rank is 0 when the random matrix M lies in $\mathrm{SO}(\mathbb{Z}/\ell\mathbb{Z})$ and the rank is 1 when M does not lie in $\mathrm{SO}(\mathbb{Z}/\ell\mathbb{Z})$.

23. POONEN-RAINS IN THE LARGE q LIMIT

In order to prove versions of the Poonen-Rains heuristics, we first want to set up a parameter space whose $\mathbb{F}_q$ points correspond to Selmer elements. Then we can try to count those $\mathbb{F}_q$ points. We will be quite brief in this section, skipping many of the details, because it is rather technical. We suggest the reader consult §5 for a similar sequence of computations in a simpler setting. Throughout, we will assume $\operatorname{Char} \mathbb{F}_q$ is not 2 or 3.

The first key observation is the following:

Lemma 23.0.1. *Let $B = \operatorname{Spec} \mathbb{F}_q$, $C \to B$ be a smooth, proper, geometrically connected curve and $j : U \subset C$ be the complement of an étale divisor. Suppose E is a nonisotrivial elliptic curve over U so that E has reduced (squarefree) discriminant over C and suppose ℓ is prime to $6q$. Then $\mathrm{Sel}_\ell(E) \simeq H^1(C, j_*E[\ell])$.*

The proof of this is somewhat technical and a bit orthogonal to the point of the course. The motivated reader can find a proof in [Ces16, Proposition 5.4(c)], since the Tamagawa number factors appearing there are all 1 as we are assuming E has reduced discriminant. A reader not familiar with Selmer groups should also just feel free to take $H^1(C, j_*E[\ell])$ as the definition of the ℓ Selmer group E.

With the above in mind, our goal will be to construct a sheaf on $\mathrm{QTwist}^n_{U/B}$ whose fiber over a point x is identified with $H^1(C_x, j_*E_x[\ell])$, for E_x the elliptic curve over U_x associated to the point x, as we next define.

Definition 23.0.2. We start with an elliptic curve E over a curve $U \subset C \to B$ as in Definition 22.0.7. Let $\mathscr{F} := E[\ell]$, which is a $\mathbb{Z}/\ell\mathbb{Z}$ étale sheaf on U. For each point $x \in \mathrm{QTwist}^n_{U/B}$, we can associate the curve $C_x := C \times_B x$, U_x which is the complement in C_x of Z and the ramification locus of the double cover associated to x, and an étale double cover $X_x \to U_x$. Moreover, we can associate an elliptic curve $E_x \to U_x$ which is the quadratic twist of the elliptic curve E given by $\mathrm{Res}_{X_x/U_x}(E|_{U_x} \times_{U_x} X_x)/E|_{U_x}$.

Said a bit more formally, Over $\mathrm{Conf}_{n,U/B}$ there is a universal curve $\mathscr{U}^n_B$ whose fiber over a point $D \subset U$ is $U - D$. More precisely, if $\mathscr{C}^n_B := C \times_B \mathrm{Conf}_{n,U/B}$, and $\mathscr{D}^n_B \subset \mathscr{C}^n_B$ is the universal divisor over $\mathrm{Conf}_{n,U/B}$, then $\mathscr{U}^n_B := \mathscr{C}^n_B - \mathscr{D}^n_B - (\mathscr{C}^n_B \times_C Z)$. We can construct a sheaf $\mathscr{F}^n_B$ on $\mathscr{U}^n_B \times_{\mathrm{Conf}_{n,U/B}} \mathrm{QTwist}^n_{U/B}$ whose fiber over a point $x \in \mathrm{QTwist}^n_{U/B}$ is the ℓ torsion $E_x[\ell]$, for E_x the quadratic twist of E associated to x.

Exercise 23.0.3. Write out a precise definition of $\mathscr{F}^n_B$ by working in families. *Hint:* You will need to use the universal étale double cover of $\mathscr{U}^n_B \times_{\mathrm{Conf}_{n,U/B}} \mathrm{QTwist}^n_{U/B}$.

Let $j : \mathscr{U}^n_B \times_{\mathrm{Conf}_{n,U/B}} \mathrm{QTwist}^n_{U/B} \xrightarrow{j} \mathscr{C}^n_B \times_{\mathrm{Conf}_{n,U/B}} \mathrm{QTwist}^n_{U/B} \xrightarrow{\lambda} \mathrm{QTwist}^n_{U/B}$. Then define $\mathrm{Sel}^{\ell}_{\mathscr{F}^n_B}$ to be the algebraic space associated to the sheaf $R^1\lambda_*(j_*(\mathscr{F}^n_B))$. We sometimes also notate this as $\mathrm{Sel}_{E[\ell]^n_B}$.

Lemma 23.0.4. *Suppose E is an elliptic curve over U over B with squarefree discriminant so that 6ℓ is in invertible on B. We obtain a Selmer sheaf* $\mathrm{Sel}_{E[\ell]^n_B}$ *on* $\mathrm{QTwist}^n_{U/B}$*. Then, the Selmer sheaf is a finite étale cover of* $\mathrm{QTwist}^n_{U/B}$*. In particular, it commutes with base change on* $\mathrm{QTwist}^n_{U/B}$.

Sketch. The proof of this lemma too is rather technical and so we only provide a brief sketch. The motivated reader can find a proof in [EL24, Proposition 5.2.1]. The basic idea is that we can realize $\mathrm{Sel}_{\mathscr{F}^n_B}$ as the image of a map of sheaves $R^1\lambda_*(j_!(\mathscr{F}^n_B)) \to R^1(\lambda \circ j)_*(\mathscr{F}^n_B)$. One can then show $R^1\lambda_*(j_!(\mathscr{F}^n_B))$ commutes with base change by proper base change and moreover that it is locally constant using that its swan conductor is trivial, using our assumption that 6 is invertible on B, and any elliptic curve over such a base has trivial swan conductor (since it can only have wild inertia in characteristics 2 and 3. Finally, one can verify $R^1(\lambda \circ j)_*(\mathscr{F}_B^n)$ is also locally constant constructible using Poincaré duality to relate it to $R^1\lambda_*(j_!(\mathscr{F}_B^n))$. It follows that the Selmer space is also locally constant constructible as it is the image of a map of locally constant constructible sheaves. □

Lemma 23.0.5. *Suppose E is an elliptic curve over U over B with squarefree discriminant so that 6ℓ is in invertible on B. Let $\pi : \mathrm{Sel}_{E[\ell]_B^n} \to \mathrm{QTwist}_{U/B}^n$ denote the projection map. Then for $x \in \mathrm{QTwist}_{U/B}^n(\mathbb{F}_q)$ we have $\pi^{-1}(x)(\mathbb{F}_q) = H^1(C_x, j_*E_x[\ell]) \simeq \mathrm{Sel}_\ell(E_x)$, for E_x the quadratic twist of E over the quadratic twist C_x of C corresponding to the point x.*

Proof.

Exercise 23.0.6. Verify the first isomorphism $\pi^{-1}(x)(\mathbb{F}_q) = H^1(C_x, j_*E_x[\ell])$. *Hint:* The explanation will be very similar to Lemma 5.0.7. As input you will need $H^0(C_{\overline{x}}, j_*E_{\overline{x}}[\ell]) = 0$ which you may take for granted (and follows from the assumptions on E).

The second equality $H^1(C_x, j_*E_x[\ell]) \simeq \mathrm{Sel}_\ell(E_x)$ follows from Lemma 23.0.1. □

Definition 23.0.7. Suppose V is a vector space with a quadratic form $q : V \to k$. We use $\mathrm{O}(V)$ to denote the group of automorphisms of the vector space preserving the quadratic form, meaning those T so that $q(Tx) = q(x)$. We use $\mathrm{SO}(V)$ to denote the subgroup of $\mathrm{O}(V)$ of matrices with determinant 1. (Technically, this is not the right definition in characteristic 2, but we will work away from characteristic 2.)

Proposition 23.0.8. *Suppose E is an elliptic curve over U over B with squarefree discriminant so that 6ℓ is in invertible on B. Then, for V a geometric fiber of $\mathrm{Sel}_{\mathscr{F}_{\mathbb{F}_q}^n} \to \mathrm{QTwist}_{U/B}^n$, we obtain a representation $\rho_{E[\ell]_B^n} : \pi_1(\mathrm{QTwist}_{U/B}^n) \to \mathrm{Aut}\, V$ which is contained in an orthogonal group $\mathrm{O}(V)$ with underlying vector space V and the image has index at most 2 in $\mathrm{O}(V)$ and is not contained in $\mathrm{SO}(V)$.*

Proof. Similar to Proposition 5.0.16, the computation of the monodromy is somewhat involved and uses the theory of middle convolution. For details, one can consult [Hal08, Theorem 6.4] or [EL24, Theorem 7.1.5]. One important point here is that the assumption on the discriminant being squarefree implies $E[\ell]$ is irreducible as a rank two $\mathbb{Z}/\ell\mathbb{Z}$ module over U, which implies that Since $\mathrm{Sel}_{\mathscr{F}_{\mathbb{F}_q}^n}$ is obtained from a twist of this sheaf via middle convolution, we also obtain $\mathrm{Sel}_{\mathscr{F}_{\mathbb{F}_q}^n}$ is irreducible.

□

Corollary 23.0.9. *Suppose E is an elliptic curve over U over B with squarefree discriminant so that 6ℓ is in invertible on B. For n sufficiently large,* $\mathrm{Sel}_{\mathscr{F}^n_{\mathbb{F}_q}}$ *has $\ell+1$ components, all of which are geometrically irreducible for n sufficiently large.*

Hence, for such n,

$$\lim_{q\to\infty} \frac{\#\,\mathrm{Sel}_{\mathscr{F}^n_B}(\mathbb{F}_q)}{\#\,\mathrm{QTwist}^n_{U/B}(\mathbb{F}_q)} = \ell+1,$$

for q ranging over those $\mathbb{F}_q$ with maps to B.

Proof. To deduce that there are $\ell+1$ components, one only needs to count the number of orbits of the orthogonal group (or the relevant index 2 subgroup of the orthogonal group on its underlying vector space over $\mathbb{Z}/\ell\mathbb{Z}$). There is the 0 vector, and then the ℓ level sets given by $\{v \in V - 0 : q(v) = x\}$ for each $x \in \mathbb{Z}/\ell\mathbb{Z}$. This accounts for the $\ell+1$ orbits in the case that the monodromy is the full orthogonal group, and the case that the monodromy group has index 2 is similar.

The point counting statement then follows from Lemma 3.2.5. □

Exercise 23.0.10. Generalize the above to show that, in the large q limit, the average size of the ν selmer group is $\sum_{m|\nu} m$, for ν an arbitrary positive integer.

For the purposes of this problem, you may assume that the monodromy representation has image $\mathrm{O}(V)$ (even though this is not in general true), for V a free $\mathbb{Z}/\nu\mathbb{Z}$ module which is the geometric fiber of $\mathrm{Sel}_{E[\nu]^n_B} \to \mathrm{QTwist}^n_{U/B}$. You may also assume the elementary fact the orthogonal group $\mathrm{O}(V)$ acts transitively on primitive vectors $w \in V$ with $q(w) = t$ for $t \in \mathbb{Z}/\nu\mathbb{Z}$; here, w is primitive if $w \neq sw'$ with $s \mid \nu$ and $s > 1$.

Remark 23.0.11. Similar to Lemma 5.0.23, one can also deduce a version of the full BKLPR distribution in the large q limit, as in Definition 22.0.24 from Proposition 23.0.8 using Proposition 5.0.22. Recall Definition 22.0.24 said that a random curve has selmer distribution which looks like the kernel of a random orthogonal group element. If the monodromy group were the orthogonal group, this would be relatively straightforward. The fact that the monodromy group can have index two in the orthogonal group makes this a little more complicated, but it ends up being true if one first takes the large q limit and then the large height limit. We worked this out in the case of the universal family in [FLR23, Theorem 1.1] and in the case of quadratic twist families in [EL24, Lemma 9.1.5]. We note that a lot of the work there was to deal with the case that ℓ was replaced with a composite number, and most of the content in the ℓ prime case was covered in [FLR23, §4].

24. APPLICATION OF HOMOLOGICAL STABILITY TO POONEN-RAINS

Before getting to the homological stability result, we will start with assuming it, and providing an application to Poonen-Rains.

Notation 24.0.1. Fix a smooth proper geometrically connected curve C over a finite field $\mathbb{F}_q$ of odd characteristic. Let $U \subset C$ be a nonempty open subscheme with nonempty complement $Z := C - U$.

Fix a prime ℓ and an elliptic curve $E \to U$. For $x \in \mathrm{QTwist}^n_{U/\mathbb{F}_q}(\mathbb{F}_{q^j})$, let E_x denote the corresponding to a quadratic twist of E.

Theorem 24.0.2. *Choose q with $\operatorname{char} \mathbb{F}_q > 3$ and ℓ a prime number relatively prime to $6q$. With notation as in Notation 24.0.1, suppose E is a nonconstant elliptic curve with squarefree discriminant. There is a constant C_ℓ depending on ℓ (but not on E) so that if $j > C_\ell$,*

$$\lim_{\substack{n\to\infty \\ n \text{ even}}} \frac{\sum_{x\in \mathrm{QTwist}^n_{U/\mathbb{F}_q}(\mathbb{F}_{q^j})} \#\mathrm{Sel}_\ell(E_x)}{\sum_{x\in \mathrm{QTwist}^n_{U/\mathbb{F}_q}(\mathbb{F}_{q^j})} 1} = \ell + 1. \tag{24.1}$$

Proof assuming later results. To prove this theorem, we will need the following ingredients:

1. Elements of Selmer groups on quadratic twists of E are in bijection with $\mathbb{F}_q$ points of the Selmer space $\mathrm{Sel}_{\mathscr{E}[\ell]^n_B}$ (explained in Lemma 23.0.5)
2. The number of stable components of the Selmer space over $\mathbb{F}_q$ is $\ell + 1$ and they are all geometrically connected (explained in Proposition 23.0.8).
3. The Selmer space $\mathrm{Sel}_{\mathscr{E}[\ell]^n_B}$ lives over $\mathrm{QTwist}^n_{U/B}$ which in turn lives over $\mathrm{Conf}_{n,U/B}$. The resulting map $\mathrm{Sel}_{\mathscr{E}[\ell]^n_B} \to \mathrm{Conf}_{n,U/B}$ induces an isomorphism on stable homology Proposition 27.3.4. This is really by far the most difficult step needing our topological tools.
4. The trace of Frobenius on the stable cohomology of $\mathrm{Conf}_{n,U_{\overline{\mathbb{F}}_q}}$ also stabilizes. This is verified in Lemma 25.5.1.
5. The unstable cohomology of $\mathrm{Sel}_{\mathscr{E}[\ell]^n_B}$ is exponentially bounded, see Corollary 26.1.3. This shows that $h^i(\mathrm{Sel}_{\mathscr{E}[\ell]^n_B}, \mathbb{Q}_\ell)$ is bounded by C^n for some constant C. However, using that the homology also stabilizes, one can deduce it is also bounded by C^i, for a different constant C. (This is analogous to the argument intended to be given in Exercise 11.1.6.)

After establishing the above five ingredients, the first two of which we already saw, Theorem 24.0.2 follows from Lemma 3.5.1. Note that to apply this, we are really applying it twice, the first time compares the point counts of $\mathrm{Sel}_{\mathscr{E}[\ell]^n_B}$ with those of $\mathrm{Conf}_{n,U/B}$ and the second identifies compares the point counts of $\mathrm{QTwist}^n_{U/B} \simeq \mathrm{Sel}_{\mathscr{E}[1]^n_B}$, with $\mathrm{Conf}_{n,U/B}$. In the end, the $\ell+1$ comes from the fact that $\mathrm{Sel}_{\mathscr{E}[\ell]^n_B}$ has $\ell+1$ stable geometric components, while $\mathrm{QTwist}^n_{U/B}$ has only 1. $\square$

We now concentrate on explaining the the final three ingredients used in the proof of Theorem 24.0.2. First, in § 25, we explain how to use log geometry to show the traces of Frobenius also stabilize.

Then, in §26, we introduce Hurwitz modules and explain how to bound the unstable cohomology in §26.1. This will involve relating the selmer sheaf to a Hurwitz space.

Finally, in § 27 we will then explain why the stable cohomology can be related to that of configuration space, assuming a topological result, whose proof we'll then describe in §28.

25. STABILITY OF FROBENIUS TRACES ON CONFIGURATION SPACE AND LOG GEOMETRY

Assuming the cohomology of configuration space stabilizes with respect to the operator of "adding a point at ∞" we wish to show the traces of Frobenius on the cohomology also stabilize. Ultimately, we wish to show this for arbitrary curves $U/\mathbb{F}_q$, but for now, let's just focus on the case $U = \mathbb{A}^1$, and call $\mathrm{Conf}_{n,k}$ the configuration space of n points in $\mathbb{A}^1$ over $\operatorname{Spec} k$. We'll come back to the case that U is an arbitrary curve later.

Said another way, there are maps $\alpha : H^i(\mathrm{Conf}_{n,\overline{\mathbb{F}}_q}) \to H^i(\mathrm{Conf}_{n+1,\overline{\mathbb{F}}_q})$ which induce isomorphisms for n large enough, and we want to show that these induce isomorphisms of Frobenius modules.

Remark 25.0.1. Of course it is not difficult to directly compute the trace of Frobenius on the cohomology of configuration space of $\mathbb{A}^1$. The point of what we are doing here is that we are explaining a method in the simplest case which generalizes to many other cases where one cannot directly compute the trace of Frobenius so easily.

Suppose these isomorphisms came from a map of schemes $\mathrm{Conf}_{n,\overline{\mathbb{F}}_q} \to \mathrm{Conf}_{n+1,\overline{\mathbb{F}}_q}$. Then Frob_q would act compatibly on these schemes, and hence the map α would induce a *Frobenius equivariant* isomorphism.

Unfortunately, it is difficult to produce such a map of schemes. And I suspect no such map exists. The new advance I will discuss today is that, however, there is a map of log schemes inducing this map.

25.1. **The stabilization map for configuration space.** Let's now describe the stabilization map α. One obtains α by comparison between cohomology of $\overline{\mathbb{F}}_q$ and $\mathbb{C}$, so we will just describe the map over $\mathbb{C}$. Here, one takes a disc with n marked points and one boundary component, and glues on a cylinder with 1 marked point, attaching one boundary component of the cylinder to the boundary of the disc. This defines a map $\mathrm{Conf}_n \to \mathrm{Conf}_{n+1}$. We want to identify this topological map with a map of log schemes. Now we have to explain what log schemes are.

25.2. **Background on log schemes.**

Definition 25.2.1. For the purposes of today's talk, a log scheme is a tuple $(X, (L_i, \sigma_i)_{i=1}^r)$ where X is a scheme, L_i are line bundles on X and $\sigma_i : \mathscr{O}_X \to L_i$ are sections of L_i. A strict map of log schemes $f : (X, (L_i, \sigma_i)_{i=1}^n) \to (Y, (M_i, \tau_i)_{i=1}^r)$ is a map $f : X \to Y$ with isomorphisms $f^* M_i \simeq L_i$ which identifies $f^* \tau_i$ with σ_i.

There is a good theory of étale cohomology for log schemes.

Goal 25.2.2. We would like to construct a map of log schemes so that the induced map on cohomology can be identified, over the complex numbers, with α.

For this, we need a few more facts about the cohomology of log schemes.

Example 25.2.3. Suppose X is a scheme and $D \subset X$ is a simple normal crossing divisor. If $L = \mathscr{O}_X(D)$ and σ is the associated section vanishing on D, the log scheme (X, L, σ) has the same cohomology as the scheme $X - D$.

Example 25.2.4. Suppose X is a scheme, then the cohomology of $(X, \mathscr{O}_X, 0)$ can be identified with the cohomology of $X \times \mathbb{G}_m$ (as long as the coefficients are invertible on the base).

Over the complex numbers, the examples above are both special cases of a more general description of the cohomology of a log scheme.

Definition 25.2.5. Namely, given a log scheme (X, L, σ) over $\mathbb{C}$, one can form the associated *Kato-Nakayama* space, which is a real analytic space whose cohomology agrees with that of the log scheme. In the case $n = 0$ the Kato-Nakayama space is simply the usual analytification. In general, it is the real oriented blow up, which is a subset of an S^{2r-1} bundle over X. Affine locally, we can locally trivialize all the line bundles, in which case it is a subspace of $X \times S^{2r-1}$ determined by a functions $f : X \to \mathbb{R}^{2r}$ associated to the $\sigma_1, \ldots, \sigma_{2k} i$ given by

$$\mathbb{L}_f^{\mathbb{R}} X := \left\{ (x, (w_1, \ldots, w_{2r})) \in X \times S^{2r-1} : f(x_i) w_j = f(x_j) w_i \right\}.$$

25.3. **Describing the log structures for the gluing map.** We'd now like to describe the relevant log schemes which give the gluing map α. Let Conf_n denote the configuration space of n points in $\mathbb{A}^1$. Another way to think of this, which will be more convenient for today, is that it is the configuration space of n red points in $\mathbb{A}^1$ and a distinct blue point at ∞. We now take a partial compactification of Conf_n, which we call $\overline{\mathrm{Conf}}_n$. This partial compactification contains stable maps $P \to \mathbb{P}^1$ with a degree n divisor on P, and a degree 1 divisor mapping to ∞. What this means concretely is that P has genus 0 with one component mapping isomorphically to $\mathbb{P}^1$, and the other genus 0 components are contracted under the above map. Whenever 2 points collide with each other (including the possibility that a blue point collides with the red point at ∞, the stable map degenerates to a curve of the form $X \cup \mathbb{P}^1$ where X has genus 0 and the two points that collided now lie on X.

There is then a natural inclusion $\mathrm{Conf}^{n-1} \to \overline{\mathrm{Conf}}_n$ given by sending a configuration of $n-1$ points $Z \subset \mathbb{P}^1$ not meeting infinity to the stable map $Z' \subset P = \mathbb{P}^1 \cup_{\infty,0} X \to \mathbb{P}^1$ where the notation means that P is a disjoint union of two copies of $\mathbb{P}^1$, with X being contracted and $\mathbb{P}^1$ mapping isomorphically. The point ∞ in the first $\mathbb{P}^1$ is glued to 0 in X. We now mark an additional red point at 1 on X and the blue point lies at ∞ on X. Altogether, this gives a stable map with n red points created from one with $n-1$ red points. This is pictured in Figure 12.

Note, however, the above map does not induce a map $\mathrm{Conf}_{n-1} \to \mathrm{Conf}_n$ because the image $\mathrm{Conf}^{n-1} \to \overline{\mathrm{Conf}}_n$ precisely lies in the complement $\overline{\mathrm{Conf}}_n - \mathrm{Conf}_n$.

We now upgrade this map of schemes to a map of log schemes that will realize our gluing operator α. To do so, we have to specify line bundles on $\overline{\mathrm{Conf}}_n$.

Definition 25.3.1. We define $\overline{\mathrm{Conf}}_n^{\log}$ to be the log scheme with underlying space $\overline{\mathrm{Conf}}_n$ together with the line bundle $\mathscr{O}_{\overline{\mathrm{Conf}}_n}(\mathrm{Conf}_{n-1})$ with the section σ associated to the divisor of singular curves (where some component under the map $P \to \mathbb{P}^1$ is contracted).

We define $\mathrm{Conf}_{n-1}^{\log}$ to be the log scheme $(\mathrm{Conf}^{n-1}, \mathscr{O}_X, 0)$.

There is a morphism of log schemes $\mathrm{Conf}^{n-1,\log} \to \overline{\mathrm{Conf}}_n\mathrm{log}$ given by the inclusion of schemes described above, using the fact that the line bundle $\mathscr{O}(\mathrm{Conf}^{n-1})$ pulls back to the trivial line bundle under the map $\mathrm{Conf}^{n-1} \to \overline{\mathrm{Conf}}_n$ with the 0 section.

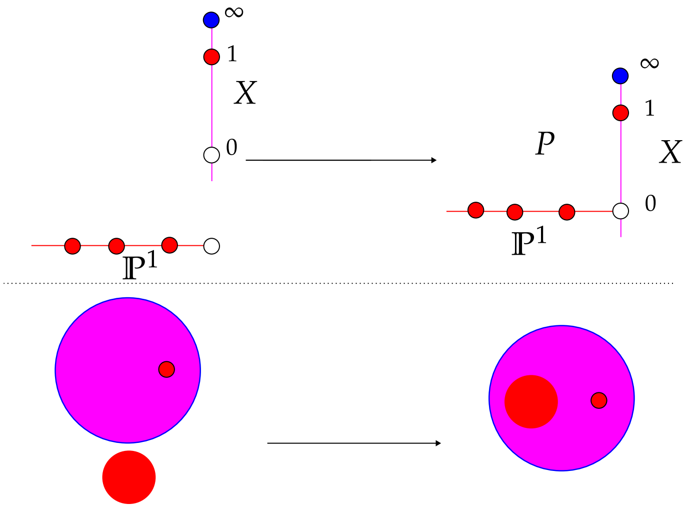

FIGURE 12. The top half of the diagram pictures the algebraic gluing map on the configuration space, corresponding to the stabilization map. This is meant to be an algebraic incarnation of the topological structure pictured coming from the little discs operad in the bottom half of the diagram.

Exercise 25.3.2. Check carefully $\mathrm{Conf}_{n-1}^{\log} \to \overline{\mathrm{Conf}_n}^{\log}$ is indeed a map of log schemes. *Hint:* What do you need to verify to show this? What are the line bundles on Conf_n?

Remark 25.3.3. We will see that the composition of maps (which only make sense on the left of spaces and not on the level of log schemes; indeed S^1 is a circle and not a log scheme)

$$\mathrm{Conf}^{n-1} \to \mathrm{Conf}^{n-1} \times S^1 \simeq \mathrm{Conf}_{n-1}^{\log} \to \overline{\mathrm{Conf}_n}^{\log} \simeq \mathrm{Conf}_n$$

is identified with the gluing map α.

There is a map $\gamma : \mathrm{Conf}_{n-1}^{\log} \to \overline{\mathrm{Conf}_n}^{\log}$ and a map There is also a map of log schemes $\beta : \mathrm{Conf}^{n-1,\log} \to \mathrm{Conf}^{n-1}$ (the latter has trivial log structure given by no line bundles; there is always a unique map from a scheme with any log structure to the same scheme with the trivial log structure). Combining these maps, on cohomology, we get

$$\gamma : H^i(\overline{\mathrm{Conf}_n}^{\log}) \to H^i(\mathrm{Conf}_{n-1}^{\log})$$
$$\beta : H^i(\mathrm{Conf}^{n-1}) \to H^i(\mathrm{Conf}_{n-1}^{\log})$$

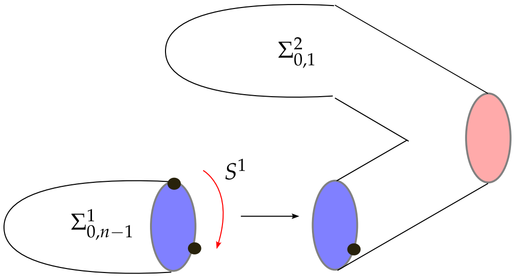

FIGURE 13. A figure depicting the gluing construction described in the proof of Proposition 25.3.4.

Proposition 25.3.4. *One can define a canonical section* ε *to* β*, and the composite of* γ *with* ε *yields a map* $H^i(\mathrm{Conf}^{n,\log}) \to H^i(\mathrm{Conf}^{n-1})$*. Over the complex numbers, this can be identified, via taking Kato-Nakayama spaces, with the map induced by the topological gluing map on cohomology which takes* $\Sigma^1_{0,n-1}$ *(genus* 0*,* $n-1$ *punctures, and* 1 *boundary component) and glues it along a boundary component to* $\Sigma^1_{0,1}$*.*

Proof. We can identify the Kato-Nakayama space of $\overline{\mathrm{Conf}_n}^{\log}$ with a moduli space of genus 0 marked curves, so that whenever the curve becomes nodal, we also specify a circle above the node and a direction on the circle. Its Kato-Nakayama space deformation retracts onto that of Conf_n by Example 25.2.3. We can identify the Kato-Nakayama space of $\mathrm{Conf}^{\log}_{n-1}$ simply as $S^1 \times \mathrm{Conf}^{n-1}$, and the corresponding map of Kato-Nakayama spaces can be thought of moduli-theoretically by taking $\Sigma^1_{0,n-1}$ (corresponding to a point of Conf^{n-1}) and gluing it to a fixed $\Sigma^2_{0,1}$ by gluing their boundaries (both of which can be identified with a copy of S^1) via a twist by a specified direction in S^1. Upon choosing a fixed direction in S^1, this restricts to the gluing map $\mathrm{Conf}_{n-1} \to \mathrm{Conf}_n$ from topology. We picture this in Figure 13.

Topologically, we can choose a section $\mathrm{Conf}^{n-1} \to S^1 \times \mathrm{Conf}^{n-1}$ (although this does not come from a map of log schemes) corresponding to choosing a fixed direction in S^1. This section induces the map ε in the statement on cohomology. Composing with γ gives the resulting map on cohomology $H^i(\mathrm{Conf}^{n,\log}) \to H^i(\mathrm{Conf}^{n-1})$ then identifies it with the above gluing map. □

25.4. **Summary of the main result.** Overall, we have a commutative diagram
(25.1)

$$\begin{tikzcd}
H^i(\mathrm{Conf}_n) \arrow[r] \arrow[d] & H^i(\overline{\mathrm{Conf}}_n^{\log}) \arrow[r, "\gamma"] \arrow[d] & H^i(\mathrm{Conf}_{n-1}^{\log}) \arrow[r, "\epsilon"] \arrow[d] & H^i(\mathrm{Conf}_{n-1}) \arrow[d] \\
H^i(\mathrm{Conf}_{n,\overline{\mathbb{F}}_q}) \arrow[r] & H^i(\overline{\mathrm{Conf}}_{n,\overline{\mathbb{F}}_q}^{\log}) \arrow[r] & H^i(\mathrm{Conf}_{n-1,\overline{\mathbb{F}}_q}^{\log}) \arrow[r] & H^i(\mathrm{Conf}_{n-1,\overline{\mathbb{F}}_q})
\end{tikzcd}$$

where the top row occurs over the complex numbers and the bottom row occurs over $\overline{\mathbb{F}}_q$. Above, we used the important Example 25.2.3 for the top left isomorphism.

One can check that the resulting vertical specializations maps are isomorphisms using that the relevant configuration spaces can be identified with the complement of a normal crossings divisor inside a smooth proper scheme (using the Fulton-MacPherson compactification). Using Proposition 25.3.4, we can then identify the top map with the topological gluing map.

This implies that we get a correspondence of log schemes over $\overline{\mathbb{F}}_p$ which induces a Frobenius equivariant map on cohomology $H^i(\overline{\mathrm{Conf}}_{n,\overline{\mathbb{F}}_p}^{\log}) \to H^i(\mathrm{Conf}_{\overline{\mathbb{F}}_p}^{n-1,\log})$ as in the above diagram. Hence, we have proven the following:

Proposition 25.4.1. *The stabilization map* $\alpha : H^i(\mathrm{Conf}_n) \to H^i(\mathrm{Conf}_{n-1})$ *corresponds to a map* $\alpha_{\overline{\mathbb{F}}_q} : H^i(\mathrm{Conf}_{n,\overline{\mathbb{F}}_q}) \to H^i(\mathrm{Conf}_{n-1,\overline{\mathbb{F}}_q})$ *as in Equation 25.1, which is equivariant for frobenius. In particular, the Frobenius traces on these two groups agree once* n *is sufficiently large so that* α *is an isomorphism.*

The following exercise is somewhat tricky, but if you can solve it, you will have quite a good grasp on the above ideas.

Exercise 25.4.2. In a similar way, consider the Harer stabilization map, which glues a genus g surface with 1 boundary component to a genus 1 surface with 2 boundary components, to form a genus $g+1$ surface with 1 boundary component. Realizes this as a equivalent to a certain map of Kato-Nakayama spaces associated to log schemes. Deduce that the Harer stability map is equivariant for the resulting action of Frobenius.

25.5. **Configuration spaces on higher genus curves.**

Lemma 25.5.1. *Let* C *be a smooth proper curve over a dvr* B *with generic characteristic* 0 *and closed point* $\operatorname{Spec} \mathbb{F}_q$. *Suppose that* $Z \subset C$ *is an étale divisor containing a section* $\sigma : B \to C$ *and let* $U := C - Z$. *Suppose the stabilization map over the complex number* $\alpha : H^i(\mathrm{Conf}_{n,U_{\mathbb{C}}}) \to H^i(\mathrm{Conf}_{n-1,U_{\mathbb{C}}})$ *induces an isomorphism* *for i sufficiently large (which we will see later is always true). Then the trace of* Frob_q^{-1}*, for* Frob_q *geometric Frobenius, on* $H^i(\mathrm{Conf}_{n,U_{\overline{\mathbb{F}}_q}})$ *and* $H^i(\mathrm{Conf}_{n-1,U_{\overline{\mathbb{F}}_q}})$ *agree.*

Remark 25.5.2. This is shown in [Pet17, Theorem 1.2](2). We give a different proof here using log geometry. We present this proof because the same idea can be applied in many other problems in arithmetic statistics to show the trace of Frobenius stabilizes on cohomology.

Proof. Similarly to our proof of Proposition 25.3.4, we create a correspondence of log schemes realizing this map on cohomology. We create a compactification of $\mathrm{Conf}_{n,U/B}$ which we call $\overline{\mathrm{Conf}}_{n,U/B}$ whose T points consists of a degree 1 map $P \to C \times_B T$ with a degree n divisor $Z \subset P$, étale over T and a degree 1 divisor $\widetilde{\sigma}$ mapping isomorphically to σ_T in C. Concretely, this means there is one component of P isomorphic to C, mapping isomorphically to C and the remaining components have genus 0. We give $\overline{\mathrm{Conf}}_{n,U/B}$ a log structure $\overline{\mathrm{Conf}}^{\log}_{n,U/B}$ associated to its divisor of parameterizing stable maps with singular source (having more than 1 component). There is a map $\mathrm{Conf}_{n-1,U/B} \to \overline{\mathrm{Conf}}_{n,U/B}$ given by sending a configuration $Z \subset U \subset C$ to the stable map $P := C \cup \mathbb{P}^1 \to C$ where C is glued along σ to 0 in $\mathbb{P}^1$, and there is an additional point $\widetilde{\sigma}$ at $\infty \in \mathbb{P}^1$, and we also specify the degree n étale divisor given as the union of Z with $1 \in \mathbb{P}^1$. This two has a log structure $\mathrm{Conf}^{\log}_{n-1,U/B}$ given by the trivial line bundle with the 0 section.

Exercise 25.5.3. Similarly to Proposition 25.3.4 complete the proof of this lemma. The idea is to show that there are Frobenius equivariant maps on cohomology $H^i(\mathrm{Conf}_{n,U_{\overline{\mathbb{F}}_q}}) \simeq H^i(\overline{\mathrm{Conf}}^{\log}_{n,U_{\overline{\mathbb{F}}_q}}) \to H^i(\mathrm{Conf}^{\log}_{n-1,U_{\overline{\mathbb{F}}_q}}) \to H^i(\mathrm{Conf}_{n-1,U_{\overline{\mathbb{F}}_q}})$ so that the composite can be identified with the stabilization map over $\mathbb{C}$.

□

Exercise 25.5.4. Generalize Lemma 25.5.1 to show the result still holds even when $Z \to B$ doesn't have a section σ. *Hint:* Reduce to the case $Z \to B$ does have a section by passing to an étale cover of degree $(\deg Z)!$ where one orders all the points of Z. Use the lemma to deduce the statement for this cover and then conclude via a transfer argument.

26. HURWITZ MODULES

In order to understand the cohomology of Selmer space, we will introduce Hurwitz modules, which are essentially modules for Hurwitz spaces. One way to think about this is that Hur^c has a monoid structure because one can take s points in a disc and t points in a disc, and put those discs in a bigger disc to get a disc with $s+t$ points. However, if you have s points in a punctured disc and t points in a punctured disc, you can't put them together because you'd then get a twice punctured disc. Instead, what you can do is take s points in a punctured disc and join them with t points in a disc to get $s+t$ points in a punctured disc. In this sense, configuration spaces of points in a punctured disc are a module for configuration spaces of points in a disc, and one can do something similar with Hurwitz spaces. Here is the formal definition.

Definition 26.0.1. Let $\Sigma^1_{g,f}$ denote a topological surface of genus g with f punctures and one boundary component. Let $\mathrm{Conf}_n^{\Sigma^1_{g,f}}$ denote the configuration space of n points in the interior of $\Sigma^1_{g,f}$ Let $B_n^{\Sigma^1_{g,f}} := \pi_1(\mathrm{Conf}_n^{\Sigma^1_{g,f}})$ denote the surface braid group associated to n points on $\Sigma^1_{g,f}$. Fix a rack c. A *Hurwitz module over c* is a triple $S = (\Sigma^1_{g,f}, \{T_n\}_{n\in\mathbb{Z}_{\geq 0}}, \{\psi_n\}_{n\in\mathbb{Z}_{\geq 0}})$ where $g, f \in \mathbb{Z}$, T_0 is a set, $T_n := c^n \times T_0$, and $\psi_n : B_n^{\Sigma^1_{g,f}} \times T_n \to T_n$ is a left action of the surface braid group on T_n such that for $0 \leq i \leq n$, such that the diagram

$$\begin{tikzcd}
(B_i^{\Sigma^1_{0,0}} \times B_{n-i}^{\Sigma^1_{g,f}}) \times (c^i \times (c^{n-i} \times T_0)) \arrow[r] \arrow[d] & c^i \times (c^{n-i} \times T_0) \arrow[d] \\
B_n^{\Sigma^1_{g,f}} \times c^n \times T_0 \arrow[r] & c^n \times T_0
\end{tikzcd} \tag{26.1}$$

commutes; the maps in the above diagram are defined as follows. The horizontal map is induced by the action of $B_i^{\Sigma^1_{0,0}} \simeq B_n$ on c^i from the definition of c and the action maps defining the bijective Hurwitz module. The left vertical map comes from the inclusion $B_i^{\Sigma^1_{0,0}} \times B_{n-i}^{\Sigma^1_{g,f}} \subset B_n^{\Sigma^1_{g,f}}$ associated to including $\Sigma^1_{0,0}$ into $\Sigma^1_{g,f}$.

Given a Hurwitz module S as above, we call T_n the n-set of S. In particular, when $n=0$, T_0 is the 0-set of S.

Exercise 26.0.2. Suppose $S = (\Sigma^1_{g,f}, \{T_n\}_{n\in\mathbb{Z}_{\geq 0}}, \{\psi_n\}_{n\in\mathbb{Z}_{\geq 0}})$ and $S' = (\Sigma^1_{g,f}, \{T_n\}_{n\in\mathbb{Z}_{\geq 0}}, \{\psi'_n\}_{n\in\mathbb{Z}_{\geq 0}}$ are two Hurwitz modules over the same rack c, with the same values of g, f and the same T_n. Assume moreover that ψ_1 agrees with ψ'_1. Show that ψ_n can be identified with ψ'_n for all $n > 1$.

The above notion of Hurwitz modules seems too general for the proofs of many of our main results, and we will mostly work in the slightly more restricted setting of bijective Hurwitz modules.

Definition 26.0.3. Fix a rack c. A *bijective Hurwitz module* over c is a Hurwitz module $S = (\Sigma^1_{g,f}, \{T_n\}_{n\in\mathbb{Z}_{\geq 0}}, \{\psi_n\}_{n\in\mathbb{Z}_{\geq 0}})$ such that the maps $B_1^{\Sigma^1_{g,f}} \times c \times T_0 \xrightarrow{\psi_1} c \times T_0 \to c$, and $B_1^{\Sigma^1_{g,f}} \times c \times T_0 \xrightarrow{\psi_1} c \times T_0 \to T_0$, induce maps $B_1^{\Sigma^1_{g,f}} \times T_0 \to \operatorname{Aut}(c)$ and $B_1^{\Sigma^1_{g,f}} \times c \to \operatorname{Aut}(T_0)$. For $\gamma \in B_1^{\Sigma^1_{g,f}}$ and $t \in T_0$, we denote the first map by $\sigma_t^\gamma : c \to c$ and for $\gamma \in B_1^{\Sigma^1_{g,f}}$ and $x \in c$ we denote the second map by $\tau_x^\gamma : T_0 \to T_0$.

Example 26.0.4. One important class of examples of bijective Hurwitz modules are obtained by taking G to be a finite group, $c \subset G$ a union of conjugacy classes, and taking its 0 set T_0 to be the set of maps $\operatorname{Hom}(\pi_1(\Sigma^1_{g,f}), G)$. In general, T_n is the set of maps $\operatorname{Hom}(\pi_1(\Sigma^1_{g,f+n}), G)$ sending loops around the n punctures to elements of c and this inherits an action of $B_n^{\Sigma^1_{g,f}}$ via its action on $\pi_1(\Sigma^1_{g,f+n})$.

Let's get a little more concrete with the above example. One can think of T_n this as choosing an element of G for each of the f punctures, an element of c for each of the n points, and an element of G for each of the $2g$ standard homology loops for the surface, so this set has cardinality $\#G^{2g+f}\#c^n$. By moving points in the configuration around the surface, these labels get modified, and the Hurwitz module keeps track of this modification. For example, when a point labeled by h passes around a puncture labeled by k, this corresponds to a full twist, and so (h,k) changes after a half twist to $(k, k^{-1}hk)$ and then after a full twist to $(k^{-1}hk, k^{-1}h^{-1}khk)$.

Exercise 26.0.5.
1. For c a rack, show that the Hurwitz space Hur^c itself is isomorphic to a Hurwitz module. That is, show there is some S so that $\operatorname{Hur}^c \simeq \operatorname{Hur}^{c,S}$. (What are the values of g, f, T_0 associated to S in this case?)
2. Fix a rack c. In the case $g = f = 0$ (so $\Sigma^1_{g,f} = \Sigma^1_{0,0}$), give a complete description of all such Hurwitz modules over c.

Just as it was important to split up racks into components for analyzing the stable homology of Hurwitz spaces, it will also be convenient to split up Hurwitz modules into their corresponding components, which we define next.

Definition 26.0.6. For c a rack and S a bijective Hurwitz module over c, an *S-component* of c is a subset $z \subset c$ which is a minimal nonempty subset of c closed under the action of c on itself and closed under the action of $B_1^{\Sigma^1_{g,f}} \times T_0$ on c.

Example 26.0.7. In the case $g = f = 0$ and T_0 is a singleton, so that $\mathrm{Hur}^{c,S} = \mathrm{Hur}^c$, the S-components of c are the same as the components of c (in our usual sense) because the action of $B_1^{\Sigma^1_{g,f}} \times T_0$ on c is trivial as $B_1^{\Sigma^1_{g,f}} \times T_0 = \pi_1(\Sigma^1_{0,0})$ is trivial.

Exercise 26.0.8. Give an example of a rack c and a bijective Hurwitz module S over c where c has two components but only one S-component.

We next introduce notation for the schemes over the complex numbers which are naturally associated to Hurwitz modules.

Definition 26.0.9. Let c be a rack and $S = (\Sigma^1_{g,f}, \{T_n\}_{n \in \mathbb{Z}_{\geq 0}}, \{\psi_n\}_{n \in \mathbb{Z}_{\geq 0}})$ be a bijective Hurwitz module over c. Let $\mathrm{Conf}_n^{\Sigma^1_{g,f}}$ denote the configuration space parameterizing n distinct points on the interior of $\Sigma^1_{g,f}$. Upon identifying $B_n^{\Sigma^1_{g,f}} \simeq \pi_1(\mathrm{Conf}_n^{\Sigma^1_{g,f}})$, we can view the bijective Hurwitz module as yielding an action $B_n^{\Sigma^1_{g,f}} \to \mathrm{Aut}(c^n \times T_0)$. Define $\mathrm{Hur}_n^{c,S}$ as the topological space which is the finite unramified cover of $\mathrm{Conf}_n^{\Sigma^1_{g,f}}$ corresponding to the above action. In particular, this covering space has degree $|c|^n \cdot |T_0|$. Suppose c has S-components $c_1, \ldots, c_v$. Suppose $n_1 + \cdots + n_v = n$ and let let $S^{n_1, \ldots, n_v} \subset c^n \times T_0$ denote the subset such that there are n_i points with labels in c_i. Then let $\mathrm{Hur}^{c,S}_{n_1,\ldots,n_v}$ denote the finite unramified cover of $\mathrm{Conf}^{\Sigma^1_{g,f}}_{n_1,\ldots,n_v}$ corresponding to the map $B_n^{\Sigma^1_{g,f}} \to \mathrm{Aut}(S^{n_1,\ldots,n_v})$.

Example 26.0.10. In the case $g = f = 0$, we can take $T_0 = *$ and we obtain $\mathrm{Hur}^{c,S}$ recovers the usual Hurwitz space Hur^c.

26.1. **Bounding the unstable cohomology.**

The bound on the unstable cohomology stems from the following:

Lemma 26.1.1. *The space* $\mathrm{Conf}_n^{\Sigma^1_{g,f}}$ *has a* 1*-point compactification with fewer than* 2^{2g+f+n} *cells.*

Proof. Recall that in the case $g = 0, f = 0$, we had Fox Neuwirth cells which kept track of how many points had the same vertical coordinate. (Technically, before, we worked with horizontal coordinates, but it is equivalent to rotate

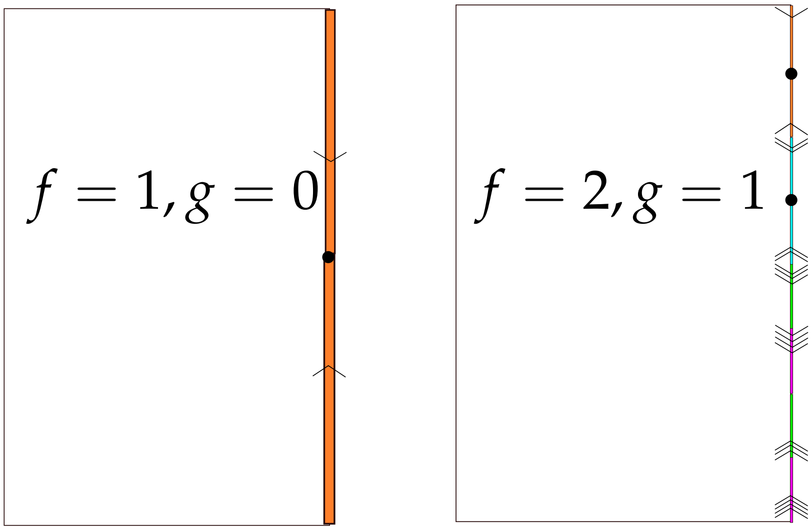

FIGURE 14. The left hand side pictures the quotient rectangle when $g = 0, f = 1$, the right hand side pictures the quotient when $g = 1, f = 2$.

the picture and consider vertical coordinates instead.) There were at most 2^n cells because we can specify the subset of points which have a different vertical coordinate from the previous one, and there are at most 2^n subsets of n points. Next, let's consider the case that $f = 1, g = 0$. Then, imagine a square $[0, 1] \times [0, 1]$. Puncture the square at the point $(1, 1/2)$ and identify the intervals $1 \times [0, 1/2)$ with $1 \times (1/2, 1)$ with opposite orientations so that the quotient surface is isomorphic to $\Sigma^1_{0,1}$. Now, we can describe cells by the number of points on the right boundary, together with the usual Fox-Neuwirth cells for the remaining points. In this case, we again specify which points have new vertical coordinates, but there are 2 cases, either the last vertical coordinate lies on the right boundary or not. This multiplies the previous upper bound 2^n by a factor of 2, so we get 2^{n+1}. See Figure 14. If we have f and g arbitrary, we can create $\Sigma^1_{g,f}$ as a quotient of a square by puncturing the right boundary f times and dividing it into $2f + 4g$ segments and making suitable identifications.

For the f punctures, we identify segments directly above and below the puncture, while for the $4g$ segments, we identify segments with value 1 and 3 mod 4 with opposite orientations and segments with value 2 and 0 mod 4 with opposite orientations.

Exercise 26.1.2. With the above quotient description of $\Sigma^1_{g,f}$, show there are at most 2^{n+f+2g} cells if one keeps track of which points lie on different horizontal coordinates, and also which points lie in the $f + 2g$ identified intervals on the right hand side of the quotient of the punctured square.

□

In fact, it turns out the number of cells is subexponential, but we won't need that.

Corollary 26.1.3. *Let c be a rack, S be a bijective Hurwitz module, with 0 set T_0. $\mathrm{Hur}_n^{c,S}$ has a one point compactification with at most $|c|^n 2^{2g+f+n}|T_0|$ cells. In particular, the number of cells grows exponentially in n, where the base of the exponent is bounded above by $2|c|$.*

Proof. There is a map $\mathrm{Hur}_n^{c,S} \to \mathrm{Conf}_n^{\Sigma^1_{g,f}}$ of degree $|c|^n|T_0|$ since the fiber of this map is identified with $T_n = c^n \times T_0$, and the cover is obtained from this by the monodromy action defining $\mathrm{Hur}_n^{c,S}$ as a cover of $\mathrm{Conf}_n^{\Sigma^1_{g,f}}$ via the identification $B_n^{\Sigma^1_{g,f}} \simeq \pi_1(\mathrm{Conf}_n^{\Sigma^1_{g,f}})$. Hence, the result follows from Lemma 26.1.1 □

27. STABLE COHOMOLOGY OF HURWITZ MODULES

27.1. **The theorem showing the homology of bijective Hurwitz modules stabilizes.** First, we state the result saying that the homology of bijective Hurwitz modules stabilizes. In order to state this, we need a notion of $\mathrm{CHur}^{c,S}$ associated to a Hurwitz module which we give next. The definition is a little technical, but it captures the idea that the component in the connected part should not come from a smaller rack or a smaller module.

Definition 27.1.1. Let c be a rack and S be a Hurwitz module over c. Let $c' \subset c$ be a subrack. We say a bijective Hurwitz module S' over c' is a *subset* of S over c if there is an inclusion $T_0' \subset T_0$ which induces commuting diagrams

$$\begin{tikzcd} B_n^{\Sigma^1_{g,f}} \times T_n' \arrow[r] \arrow[d] & T_n' \arrow[d] \\ B_n^{\Sigma^1_{g,f}} \times T_n \arrow[r] & T_n. \end{tikzcd} \tag{27.1}$$

We write $(c', S') \subset (c, S)$ to indicate that S' is a subset of S.

Exercise 27.1.2. Verify that any two such subsets (c', S_1) and (c', S_2) are contained in a bigger subset. So the above maximal subset is well defined.

Construction 27.1.3. *Given a rack c and a bijective Hurwitz module S over c, we put a doubly filtered structure on* $\mathrm{Hur}^{c,S}$*. We define* $F_{*,*}\,\mathrm{Hur}^{c,S}$ *to be a bifiltration on* $\mathrm{Hur}^{c,S}$ *defined as follows.*

Suppose $c'' \subset c$ *and* S'' *is a bijective Hurwitz module over* c'' *which is a subset of the bijective Hurwitz module S in the sense of Definition 27.1.1.*

We then define the (i,j)*th part of the bifiltration* $F_{i,j}\,\mathrm{Hur}^{c,S}$ *to be the union of all components contained in some* $\mathrm{Hur}^{c'',S''}$ *for* $(c'',S'') \subset (c,S)$ *with* $|c''| \geq i$ *and* $|T_0''| \geq j$ *for* T_0'' *the* 0*-set of* S''*. Define* $\mathrm{CHur}^{c,S} := F_{|c|,|T_0|}\,\mathrm{Hur}^{c,S}$*.*

We use $A_{c,S} := C_*(\mathrm{Hur}^{c,S};\mathbb{Z})$*. We use* $F_{*,*}A_{c,S}$ *to denote the associated bifiltration obtained from* $F_{*,*}\,\mathrm{Hur}^{c,S}$ *by taking chains. We will also view* $F_{*,*}A_{c,S}$ *as giving a bifiltration on* $A_{c,S}$ *as an* A_c *module. If* T_0 *is the* 0*-set of S, define* $CA_{c,S} := F_{|c|,|T_0|}A_{c,S}$*.*

Exercise 27.1.4 (Easy exercise)**.** With the definition above, show that $\mathrm{CHur}^{c,S}$ is the union of components of $\mathrm{Hur}^{c,S}$ which are not contained in $\mathrm{Hur}^{c',S'}$ for any subset $(c',S') \subset (c,S)$

Exercise 27.1.5. Suppose $S = (\Sigma^1_{g,f}, \{T_n\}_{n\geq 0}, \{\psi_n\}_{n\geq 0})$ with $g = f = 0$ and T_0 a point. In this case, show that $\mathrm{Hur}^{c,S} = \mathrm{Hur}^c$. Given an example to show that in general $\mathrm{CHur}^{c,S} \neq \mathrm{CHur}^c$. (This inequality probably means we should have chosen a different definition of CHur^c so that these match, but unfortunately, this is not the definition we chose in the papers.)

Here is the result on stable homology:

Theorem 27.1.6. *Let c be a finite rack and let S be a bijective Hurwitz module over c. Let* $c_1,\ldots,c_v$ *denote the S-components of c. Using notation from Definition 26.0.3, there are constants I and J, depending on* $|c_1|$ *and the maximal order of an element of* c_1 *acting on c, so that for any* $i \geq 0$ *and* $n_1 > Ii + J$*, any element* $x \in c_1$ *induces an isomorphism* $H_i(\mathrm{CHur}^{c,S}_{n_1,\ldots,n_v};\mathbb{Z}) \to H_i(\mathrm{CHur}^{c,S}_{n_1+1,\ldots,n_v};\mathbb{Z})$*.*

We explain the proof in §28.2.2.

27.2. **The theorem on the stable homology of bijective Hurwitz modules.**

We next want to state our result computing the stable value of the homology of bijective Hurwitz modules. Just as we saw that the stable homology of Hurwitz spaces identify their homology with the homology of a quotient rack in Theorem 21.0.1, the stable homology of Hurwitz modules will also be identified with the homology associated to a quotient rack. In order to state our main result on the stable homology of Hurwitz modules, we need the notion of a quotient of a Hurwitz module:

Definition 27.2.1. If c is a rack and $c' \subset c$ is a subrack, and let $S = (\Sigma^1_{g,f}, \{T_n\}_{n \in \mathbb{Z}_{\geq 0}}, \{\psi_n\}_{n \in \mathbb{Z}_{\geq 0}})$ be a bijective Hurwitz module over c. Suppose $c' \subset c$ is normal and closed under the action of $B_1^{\Sigma^1_{g,f}} \times T_0$ on c. Define the bijective Hurwitz module $S/c' = (\Sigma^1_{g,f}, \{\overline{T}_n\}_{n \in \mathbb{Z}_{\geq 0}}, \{\overline{\psi}_n\}_{n \in \mathbb{Z}_{\geq 0}})$ over c/c' as follows. Take $\overline{T}_0$ to denote the quotient of T_0 by the equivalence relation generated by $s \sim s'$ if there is some $\gamma \in B_n^{\Sigma^1_{g,f}}$ and $x_1, \ldots, x_n \in c'$ with $\psi_n(\gamma, x_1, \ldots, x_n, s) = (y_1, \ldots, y_n, s')$ such that y_i and x_i have the same image in c'/c'. Then, take $\overline{T}_n := (c/c')^n \times \overline{T}_0$. Finally, for $x \in T_n$, we use $\overline{x}$ to denote its image in $\overline{T}_n$ and for $\gamma \in B_n^{\Sigma^1_{g,f}}$ define $\overline{\psi}_n((\gamma, \overline{x})) := \overline{\psi_n((\gamma, x))}$.

Exercise 27.2.2 (Easy exercise)**.** Show that, in the above definition any S-component $c' \subset c$ is both normal and closed under the action of $B_1^{\Sigma^1_{g,f}} \times T_0$ on c. Hence, we may make sense of the quotient of c by an S-component.

Exercise 27.2.3 (Tricky exercise)**.** From the above definition, it is not clear that the quotient of a Hurwitz module is again a Hurwitz module, because it is not clear the operation $\overline{\psi}_n$ is well defined. Verify this is indeed well defined.

In the paper, we only found a computational approach to directly verifying the above by writing out how generators act. However, it seems likely there should be a conceptual approach to this that we didn't find. See if you can find one!

With the definition of quotients of Hurwitz modules given, the main result on the stable homology of Hurwitz modules now says that their homology stabilizes and the stable homology of each component in the c' direction is identified with that of the corresponding component of $\mathrm{CHur}^{c/c', S/c'}$.

Theorem 27.2.4. *Let c be a finite rack and S a bijective Hurwitz module over c as in Definition 26.0.3. Let $c_1, \ldots, c_v$ denote the S-components of c. There are constants I and J, depending only on $|c_1|$ and the minimal order of an element of c_1 acting on c, so that for any $i \geq 0$ and $n_1 > Ii + J$, and any component $Z \subset \mathrm{CHur}^{c,S}_{n_1, \ldots, n_v}$ mapping to a component $Z' \subset \mathrm{CHur}^{c/c_1, S/c_1}_{n_1, \ldots, n_v}$, the map $H_i(Z, \mathbb{Q}) \to H_i(Z', \mathbb{Q})$ is an isomorphism.*

We explain the proof idea in §28.4.

Remark 27.2.5. It is possible to list a finite set of primes to invert, which we do in [LL25b, Theorem 1.4.9] but in order to simplify the statement, we just work rationally here.

27.3. **Computing the stable homology of Selmer spaces.** We'd like to relate Selmer spaces to Hurwitz spaces. The basic idea is that given an element of Selmer space, which corresponds to an element of $H^1(C, j_*E[\ell])$, we can view this as a torsor for the sheaf $j_*E[\ell]$. This corresponds to a generically finite cover of C, and its normalization is a finite cover. What is the covering group? If we take the Galois closure of $E[\ell]$, we get a $\mathrm{SL}_2(\mathbb{Z}/\ell\mathbb{Z})$ cover. However, this is a torsor for that, and so we get a cover for a group $\mathrm{ASL}_2(\mathbb{Z}/\ell\mathbb{Z}) := (\mathbb{Z}/\ell\mathbb{Z})^2 \rtimes \mathrm{SL}_2(\mathbb{Z}/\ell\mathbb{Z})$. Here is a formal definition:

Definition 27.3.1. The *affine special linear group* is $\mathrm{ASL}_2(\mathbb{Z}/\ell\mathbb{Z}) := (\mathbb{Z}/\ell\mathbb{Z})^2 \rtimes \mathrm{SL}_2(\mathbb{Z}/\ell\mathbb{Z})$, where the action of $\mathrm{ASL}_2(\mathbb{Z}/\ell\mathbb{Z})$ on $(\mathbb{Z}/\ell\mathbb{Z})^2$ is via the standard action of matrices on their underlying free rank $\mathbb{Z}/\ell\mathbb{Z}$ module of rank 2.

Remark 27.3.2. By definition, $\mathrm{ASL}_2(\mathbb{Z}/\ell\mathbb{Z})$ sits in an exact sequence
(27.2)

$$\begin{tikzcd} 0 \arrow[r] & (\mathbb{Z}/\ell\mathbb{Z})^2 \arrow[r, "\iota"] & \mathrm{ASL}_2(\mathbb{Z}/\ell\mathbb{Z}) \arrow[r, "\Pi"] & \mathrm{SL}_2(\mathbb{Z}/\ell\mathbb{Z}) \arrow[r] & 0 \end{tikzcd}$$

with inclusion map ι and quotient map Π. With this presentation, $\mathrm{ASL}_2(\mathbb{Z}/\ell\mathbb{Z})$ can be explicitly described as those matrices of the form
(27.3)

$$\mathrm{ASL}_2(\mathbb{Z}/\ell\mathbb{Z}) \simeq \left\{ \begin{pmatrix} m_{11} & m_{12} & v_1 \\ m_{21} & m_{22} & v_2 \\ 0 & 0 & 1 \end{pmatrix} = \begin{pmatrix} M & v \\ 0 & 1 \end{pmatrix} \in \mathrm{GL}_3(\mathbb{Z}/\ell\mathbb{Z}) : M \in \mathrm{SL}_2(\mathbb{Z}/\ell\mathbb{Z}), v \in (\mathbb{Z}/\ell\mathbb{Z})^2 \right\}.$$

Here is the definition of the Hurwitz module which will be isomorphic to Selmer space over the complex numbers

Definition 27.3.3. Let $B = \operatorname{Spec} \mathbb{C}$ and let $\ell > 3$ be a prime. Suppose E is an elliptic curve as in Notation 24.0.1 and assume $E[\ell]$ corresponds to a representation $\rho_\ell : \pi_1(\Sigma^1_{g,f}) \to \mathrm{SL}_2(\mathbb{Z}/\ell\mathbb{Z})$ coming from the fact that the Galois closure of $E[\ell]$ is an $\mathrm{SL}_2(\mathbb{Z}/\ell\mathbb{Z})$ cover. (This assumption will be satisfied whenever E is nonconstant elliptic curve with squarefree discriminant as follows from [Zyw14, Proposition 2.7] using that the number $\mathscr{L}_E$ there is 1 by assumption that E has squarefree discriminant and we are assuming ℓ is not 2 or 3.) Let $c \subset \mathrm{ASL}_2(\mathbb{Z}/\ell\mathbb{Z})$ denote the conjugacy class of elements g with $\Pi(g) = -\mathrm{id}$. Let α_i, β_i denote a standard basis of homology for Σ^1_g and let $\gamma_1, \dots, \gamma_{f+1}$ be the loops around the f punctures and boundary component of $\Sigma^1_{g,f}$. Let T_0 be the subset of $\phi \in \mathrm{Hom}(\pi_1(\Sigma^1_{g,f}), \mathrm{ASL}_2(\mathbb{Z}/\ell\mathbb{Z}))$ such that

(1) $\Pi(\phi(\alpha_i)) = \pm\rho_\ell(\alpha_i)$
(2) $\Pi(\phi(\beta_i)) = \pm\rho_\ell(\beta_i)$

(3) $\phi(\gamma_i) = (\rho_\ell(\gamma_i), v_i)$ with $v_i \in \operatorname{im}(\rho(\gamma_i) - \mathrm{id})$.

Proposition 27.3.4. *Suppose E is an elliptic curve over U over B with squarefree discriminant so that 6ℓ is in invertible on B. Using notation as in Definition 23.0.2 and notation for c, S as in Definition 27.3.3, for $\mathscr{F} = E[\ell]$, and $B = \mathbb{C}$ there is an isomorphism $\mathrm{Sel}_{\mathscr{F}^n_B} \to [\mathrm{Hur}^{c,S} \ / \ \mathrm{ASL}_2(\mathbb{Z}/\ell\mathbb{Z})]$.*

Proof. As mentioned above, we can identify elements of $\mathrm{Sel}_{\mathscr{F}^n_B}$ with certain $\mathrm{ASL}_2(\mathbb{Z}/\ell\mathbb{Z})$ covers as follows. Any element of $H^1(C, j_*E[\ell])$, can be viewed this as a torsor for the sheaf $j_*E[\ell]$. Since the Galois closure of $E[\ell]$ will be $\mathrm{SL}_2(\mathbb{Z}/\ell\mathbb{Z})$ cover. However, since taking H^1 corresponds to torsors, we get a cover for a group $\mathrm{ASL}_2(\mathbb{Z}/\ell\mathbb{Z}) := (\mathbb{Z}/\ell\mathbb{Z})^2 \rtimes \mathrm{SL}_2(\mathbb{Z}/\ell\mathbb{Z})$. It remains to identify the local monodromies of such covers with those appearing in T_n. First, the data of a quadratic twist ramified along a degree n divisor, which corresponds to a homomorphism $\pi_1(\Sigma^1_{g,f+n}) \to \mathbb{Z}/2\mathbb{Z}$, of which there are 2^{2g}, since we have to send the generators around the f punctures to the trivial element, the generators around the n points to the nontrivial element, and the generators of homology can go anywhere. This corresponds to the conditions that

$$\Pi(\phi(\alpha_i)) = \pm\rho_\ell(\alpha_i), \tag{27.4}$$

$$\Pi(\phi(\beta_i)) = \pm\rho_\ell(\beta_i), \tag{27.5}$$

$$\Pi(\phi(\gamma_i)) = \rho_\ell(\gamma_i), \tag{27.6}$$

$$\Pi(\phi(\delta_i)) = -\mathrm{id}, \tag{27.7}$$

for δ_i the loops around the n punctures. This will give a sheaf on $U = C - Z$ which is a torsor for $E[\ell]$, and we need to understand the condition for such a sheaf to extend across the punctures and become a torsor for $j_*E[\ell]$. Let $M_i = \rho_\ell(\gamma_i)$. At the points s_i corresponding to the f punctures, and let p_i correspond to the n punctures. In general the cokernel $E[\ell] \to j_*E[\ell]$ will be a torsion sheaf supported on the points $s_i \cup p_i$. At a point p with monodromy M, the cokernel will have rank $\operatorname{im} M - \mathrm{id}$. The point here is essentially that if the monodromy is the identity, the torsor will extend, and if the monodromy has rank 1, a 1 dimensional subspace of the torsor will extend and we will have to specify the monodromy in the other direction. When there is no 1 eigenspace, we have to specify the total monodromy for the torsor at the point, (such as over all of the n points where the monodromy is $-\mathrm{id}$) which lives in $(\mathbb{Z}/\ell\mathbb{Z})^2$.

One can check this bijection holds as the n points vary, and hence defines a map of quotients of étale covers of configuration space. The reason we have to quotient by $\mathrm{ASL}_2(\mathbb{Z}/\ell\mathbb{Z})$ is because the Selmer space really parameterizes isomorphism classes of torsors, which means we consider them up to global conjugation by their covering group. □

In order to understand the stable homology of Selmer space, the following will give a description of the stable homology of the corresponding coefficient system.

Lemma 27.3.5. *Suppose c is a rack with a single component and* $S = (\Sigma^1_{g,f}, \{T_n\}_{n\in\mathbb{Z}_{\geq 0}}, \{\psi_n\}_{n\in\mathbb{Z}_{\geq 0}})$ *is a bijective Hurwitz module over c. Then, every component of* $\mathrm{Hur}_n^{c/c,S/c}$ *maps isomorphically to* $\mathrm{Conf}_n^{\Sigma^1_{g,f}}$.

Proof. Let $S/c := (\Sigma^1_{g,f}, \{\overline{T}_n\}_{n\in\mathbb{Z}_{\geq 0}}, \{\overline{\psi}_n\}_{n\in\mathbb{Z}_{\geq 0}})$. Start with an element $(x_1, \dots, x_n, s) \in T_n$ mapping to an element $(z_1, \dots, z_n, t) \in \overline{T}_n$. The statement of the lemma is equivalent to the statement that every element of $B_n^{\Sigma^1_{g,f}}$ acts trivially on $(z_1, \dots, z_n, t)$. Suppose we have some path $\gamma \in B_n^{\Sigma^n_{g,f}}$ so that $\psi_n(\gamma, x_1, \dots, x_n, s) = (x'_1, \dots, x'_n, s')$. Then, we wish to show x_i is equivalent to x'_i in c/c and s is equivalent to s' in $\overline{T}_0$. Since c has a single component x_i and x'_i lie in the same component, so are equivalent in c/c. Finally, s is equivalent to s' in $\overline{T}_0$ as is immediate from the definition of $\overline{T}_0$, using that x_i and x'_i lie in the same component of c. □

Using the above, let's deduce the description of the stable homology of Selmer space.

Corollary 27.3.6. *Suppose E is a nonisotrivial elliptic curve over U over B with squarefree discriminant so that* 6ℓ *is in invertible on* $B = \operatorname{Spec} \mathbb{F}_q$. *For each component* $Z \subset \mathrm{Sel}_{\mathscr{F}^n_B}$, *the projection map* $Z \to \mathrm{Conf}_{n,U/B}$ *induces an isomorphism on stable cohomology with* $\mathbb{Q}_\ell$ *coefficients. That is, there are constants I and J so that for* $n > Ii + J$, $H^i(\mathrm{Conf}_{n,U/B}, \mathbb{Q}_\ell) \to H^i(Z, \mathbb{Q}_\ell)$ *is an isomorphism.*

Proof. We can compare to characteristic 0 using the existence of normal crossings compactifications of Hurwitz spaces; we explain the idea for Hurwitz spaces over $\mathbb{P}^1$ in see §29 and Theorem 29.0.1, but a similar proof works over arbitrary curves; see [EL24, Corollary B.1.4] for a reference for the normal crossings compactification, which implies that the cohomology in positive characteristic agrees with that in characteristic 0. So, it's enough to understand the statement over $B = \operatorname{Spec} \mathbb{C}$. Moreover, using Proposition 27.3.4, we can identify $\mathrm{Sel}_{\mathscr{F}^n_B}$ as a quotient of a certain bijective Hurwitz module $\mathrm{Hur}^{c,S}$. We have a factorization $\mathrm{Hur}_n^{c,S} \to [\mathrm{Hur}_n^{c,S} \,/\, \mathrm{ASL}_2(\mathbb{Z}/\ell\mathbb{Z})] \to \mathrm{Sel}_{\mathscr{F}^n_B} \to [\mathrm{Conf}_{n,U/B} \,/\, \mathrm{ASL}_2(\mathbb{Z}/\ell\mathbb{Z})] \to \mathrm{Conf}_{n,U/B}$. The first two maps are finite étale and the last map induces an isomorphism on cohomology because the action of $\mathrm{ASL}_2(\mathbb{Z}/\ell\mathbb{Z})$ is trivial on configuration space. Hence, it's enough to identify the stable rational homology of each component of Hurwitz space with that of configuration space. Now, Theorem 27.2.4, shows each component of $\mathrm{CHur}^{c,S}$ has the same stable homology as each component of $\mathrm{CHur}^{c/c,S/c}$. One may use the squarefree discriminant and nonisotrivial hypotheses on E to show that in this case, $\mathrm{CHur}^{c,S}$ is all of $\mathrm{Hur}^{c,S}$, (essentially because the elements around the f punctures will already generate the covering group,) and $\mathrm{CHur}^{c/c,S/c}$ is all of $\mathrm{Hur}^{c/c,S/c}$. This identifies the stable homology of each component of $\mathrm{Hur}^{c,S}$ with the corresponding component of $\mathrm{Hur}^{c/c,S/c}$, and so we conclude because the stable homology of $\mathrm{Hur}_n^{c/c,S/c}$ is identified with that of $\mathrm{Conf}_n^{\Sigma^1_{g,f}}$ by Lemma 27.3.5. □

28. PROVING THE TOPOLOGICAL THEOREMS FOR HURWITZ SPACE MODULES

We'd now like to explain proofs of Theorem 27.1.6 and Theorem 27.2.4. These proofs of the module result will be similar to the case of Hurwitz spaces, but every step will be a little trickier. However, it will give us a good chance to practice the techniques we were learning. Morally, it should be at least as difficult, because if we take the trivial Hurwitz module S over a rack c where $T_0 = *$ and $g = f = 0$, we recover $\mathrm{Hur}^{c,S} = \mathrm{Hur}^c$.

Remark 28.0.1. As the argument for this is already quite complicated, we will take every opportunity to simplify it that we can. In particular, we will often specialize our proofs to the case $g = 0, f = 1$. Often there is more to keep track of in the other cases, but this will significantly simplify notation, especially because the fundamental group of the punctured disc is $\mathbb{Z}$, generated by a single element.

28.1. **Stability of a quotient.**

We start with discussing the proof that the homology of $\mathrm{Hur}^{c,S}$ stabilizes. As in the case of Hurwitz spaces, before showing this stabilizes, we show that the homology of a quotient stabilizes. The general strategy will be to first consider a quotient of the homology by elements of c and show that quotient stabilizes, and then remove elements from the quotient one at a time. Our first result toward this is Proposition 28.1.2 below, which is the result saying that the homology of the quotient stabilizes.

Notation 28.1.1. Fix a rack c and let $\alpha_x \in \pi_0 \mathrm{Hur}^c_1$ denote the corresponding element. Fix an ordering of the elements of c. For $X \subset c$ a subset, we write $\alpha_X := \{\alpha_x : x \in X\}$ and $\alpha^i_X := \{\alpha^i_X : x \in X\}$. We use $C_*(\mathrm{Hur}^{c,S}_+)/(\alpha_X^{\mathrm{ord}(X)})$ to denote the tensor product $C_*(\mathrm{Hur}^c_+)/(\alpha_{x_{i_1}}^{\mathrm{ord}(x_{i_1})}) \otimes_{C_*(\mathrm{Hur}^c_+)} \cdots \otimes_{C_*(\mathrm{Hur}^c_+)}$

$C_*(\mathrm{Hur}^c_+)/(\alpha_{x_{i_{|X|}}}^{\mathrm{ord}(x_{i_{|X|}})}) \otimes_{C_*(\mathrm{Hur}^c_+)} C_*(\mathrm{Hur}^{c,S})$ for $i_1, \dots, i_{|X|}$ the indices of the elements of X taken in order of the ordering on c.

At this point, it may be useful to recall the notion of being bounded in a linear range, as defined in Definition 14.0.3.

Proposition 28.1.2. *The chains of the quotient* $M := \widetilde{C}_*(\mathrm{Hur}^{c,S}_+)/(\alpha_c^{\mathrm{ord}(c)})$ *is bounded in a linear range.*

Proof. The proof is quite similar to the proof of Theorem 14.0.4 and Theorem 17.0.6, but let us say a little bit about what goes into it. One can axiomatize the ingredients needed to prove the quotient stabilizes, and it amounts to the following four ingredients in this case.

Proposition 28.1.3. *Let* $A := \widetilde{C}_*(\mathrm{Hur}^c_+)$ *viewed as a graded* $\mathbb{Z}$ *module and assume*

1. *$k \otimes_A k$ is bounded in a linear range*
2. *$k \otimes_A M$ is bounded in a linear range*
3. *If $I \subset \pi_0 A$ denotes the kernel of the augmentation map $\pi_0 A \to k$, there is some uniform $t \in \mathbb{N}$ so that I^{t+1} acts by 0 on $\pi_i M$ for each i.*

Then, M is bounded in a linear range.

We will not explain the proof of this proposition, which again is similar to the proof of Theorem 14.0.4 and Theorem 17.0.6, but instead we will explain how to verify the four hypotheses.

Most of the above points are fairly easy to verify: We saw in Lemma 17.0.2 that $k \otimes_A k$ is bounded in a linear range. We similarly have that $k \otimes_A M$ is bounded in a linear range, as we explain in Lemma 28.1.4. For the final condition (3), we want to know that a power of I acts by 0 on the homology of the quotient $C_*(\mathrm{Hur}^{c,S}_+)/(\alpha_c^{\mathrm{ord}(c)})$. The point is that $\alpha_x^{2\,\mathrm{ord}(x)}$ acts by 0 on $C_*(\mathrm{Hur}^{c,S}_+)/\alpha_x^{\mathrm{ord}(x)}$, and so if x_i is the ith element among the members of c in order, then $\alpha_{x_i}^{2^i\,\mathrm{ord}(x_i)}$ acts by 0 on $C_*(\mathrm{Hur}^{c,S}_+)/\alpha_c^{\mathrm{ord}(c)}$. So if we take a uniformly high enough power I^{t+1} of I (say we can take t larger than $\sum_{i=1}^{|c|} 2^i\,\mathrm{ord}(x_i)$), which is generated by the α_x, some $\alpha_{x_i}^{2^i\,\mathrm{ord}(x_i)}$ must show up, and so I^{t+1} acts by 0. □

The last result we need as input for the proof of Proposition 28.1.2 is the following, which verifies condition (2) in Proposition 28.1.3.

Lemma 28.1.4. *We have* $*_+ \otimes_{\mathrm{Hur}^c_+} \left(\mathrm{Hur}^{c,S}_+ /(\alpha_c^{\mathrm{ord}(c)})\right)$ *is bounded in a linear range.*

FIGURE 15. The left hand side pictures the scanned rectangle when $g = 0, f = 1$, the right hand side pictures the scanned rectangle when $g = 1, f = 2$. On the left hand side, the cells of dimension n correspond to placing n points at distinct heights on the orange rectangle and sliding them back and forth. On the right hand side, there are now 4 rectangles we can place points in, and the n-cells correspond to placing n points in total among the four rectangles, all at distinct heights

Lemma 28.1.4 follows from Lemma 28.1.5 below because once we show $*_+ \otimes_{\mathrm{Hur}^c_+} \left(\mathrm{Hur}^{c,S}_+\right)$ is bounded in a linear range, by quotienting by the actions of finitely many $\alpha_x^{\mathrm{ord}_c(x)}$, we also obtain the iterated quotient $*_+ \otimes_{\mathrm{Hur}^c_+} \left(\mathrm{Hur}^{c,S}_+ / (\alpha_c^{\mathrm{ord}(c)})\right)$ is bounded in a linear range.

Lemma 28.1.5. *The n-cells of $*_+ \otimes_{\mathrm{Hur}^c_+} \left(\mathrm{Hur}^{c,S}_+\right)$ are concentrated in grading n. In particular, $*_+ \otimes_{\mathrm{Hur}^c_+} \left(\mathrm{Hur}^{c,S}_+\right)$ is bounded in a linear range.*

Proof. The key for this is to use the quotient model of configuration space/Hurwitz space described in the proof of Lemma 26.1.1. Let's first think about the case $g = 0, f = 1$.

Then, we can think of $\Sigma^1_{g,f}$ as a quotient of a rectangle with 1 puncture on the right, where we identify $1 \times [0, 1/2)$ with $1 \times (1/2, 1]$ with opposite orientations. See Figure 15. The loop around the puncture corresponds to sending an element at height $3/4$ to the right, in which case $(0, 3/4)$ gets identified with $(0, 1/4)$ and comes back to the left.

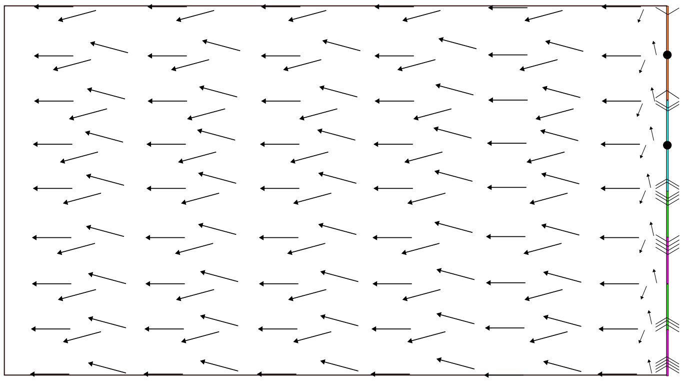

FIGURE 16. Here, we depict the part of the scanning argument where we create $2g + f$ separate rectangles. This is the case $f = 2, g = 1$.

Remark 28.1.6. In general, one can use a scanning argument to identify $*_+ \otimes_{\mathrm{Hur}^c_+} \mathrm{Hur}^{c,S}_+$ with collections of points in the rectangle so that no two have the same vertical coordinate h and also if one has vertical coordinate h no other has vertical coordinate $1 - h$, where points are labeled by elements of c. Moreover, there are no points with vertical coordinate $1/2$. We still have that when elements pass over each other, their labels change, but now the label changes and it acts on T_0 whenever it passes through the right boundary via the quotient defining $\Sigma^1_{0,1}$.

The way to achieve this scanning model follows two steps. First we use a general scanning argument to identify points of the tensor product with collections of labeled points as above. Then, we can apply a homotopy to push all elements outward from the puncture, so that in particular, there are no points with vertical coordinate $1/2$. Once no points have vertical coordinate $1/2$, the space is topologically a rectangle, but with height h identified with height $1 - h$, and so we can apply the usual scanning argument to assume there is one in each row. This corresponds to pushing points either to the left above $1/2$ at height h or to the left below $1/2$ at height $1 - h$, which is why no point will be in height $1 - h$ and h at the same time after scanning. A similar procedure works for general g and f. For example, the case $f = 2, g = 1$ is depicted in Figure 16

Overall, Remark 28.1.6 gives cells for $*_+ \otimes_{\mathrm{Hur}^c_+} \mathrm{Hur}^{c,S}_+$ indexed by tuples in $c^n \times T_0$ where the points all have heights between $1/2 < h < 1$. The cells in degree n all correspond to n-spheres, since they are identified with $[0, 1]^n$, where the ith coordinate indicates the location of the ith point on the horizontal line it can move across (which involves both the y coordinates h

and $1-h$, which are identified). We identify the boundary of this cube with the basepoint, corresponding to when the point hits the left boundary, either above $1/2$ for the 0 boundary or below $1/2$ for the 1 boundary. Then, we see this space is built out of n-cells in degree n and hence π_i of its nth graded part vanishes when $i > n$, meaning it is bounded in a linear range.

When g, f are general, a similar argument works, but the scanning argument will produce $2g + f$ rectangles.

Exercise 28.1.7. How many n-cells will $*_+ \otimes_{\mathrm{Hur}^c_+} \mathrm{Hur}^{c,S}_+$ have in the case of genus g and f punctures. You should express your answer in terms of $|c|$ and $|T_0|$.

□

28.2. **Showing the homology of bijective Hurwitz modules stabilizes.** We now have to carry out the "unquotienting" procedure. That is, we showed $\widetilde{C}_*(\mathrm{Hur}^{c,S}_+ /(\alpha_c^{\mathrm{ord}(c)}))$ stabilizes to 0 in a linear range. We next want to remove elements from the quotient to show the homology of $\mathrm{Hur}^{c,S}_+$ itself stabilizes in a linear range, thereby proving Theorem 27.1.6. The procedure of the proof is similar to that of the monoid case of Hurwitz spaces, which we explained in Theorem 16.0.1. Hence, we will be somewhat brief.

The starting point is that we showed in Proposition 28.1.2 that the quotient $M := \widetilde{C}_*(\mathrm{Hur}^{c,S}_+ /(\alpha_c^{\mathrm{ord}(c)}))$ stabilizes to 0 in a linear range. Recall the filtration on $\mathrm{Hur}^{c,S}$ from Construction 27.1.3. Using a similar argument to Lemma 16.0.4, we can filter $\widetilde{C}_*(\mathrm{Hur}^{c,S}_+ /(\alpha_c^{\mathrm{ord}(c)}))$ by its connected parts to deduce the following

Lemma 28.2.1. $\widetilde{C}_*(\mathrm{CHur}^{c,S}_+ /(\alpha_c^{\mathrm{ord}(c)}))$ *is bounded in a linear range.*

28.2.2. *Proof of Theorem 27.1.6, assuming Lemma 28.2.3.* Next, consider a subset $V \subset c$ containing at least on element of each S-component of c. (Recall S components are minimal unions of components of c which are permuted under the action of T_0.) We consider the quotient $\widetilde{C}_*(\mathrm{CHur}^{c,S}_+ /(\alpha_V^{\mathrm{ord}(V)}))$ and we aim to show it vanishes by downward induction on the size of V. The case that $V = |c|$ was covered in Lemma 28.2.1. We will assume by induction that $\widetilde{C}_*(\mathrm{CHur}^{c,S}_+ /(\alpha_V^{\mathrm{ord}(V)}, \alpha_y^{\mathrm{ord}(y)}))$ is bounded in a linear range for each $y \in c - V$. To show $\widetilde{C}_*(\mathrm{CHur}^{c,S}_+ /(\alpha_V^{\mathrm{ord}(V)}))$ it is therefore enough to show $\widetilde{C}_*(\mathrm{CHur}^{c,S}_+ /(\alpha_V^{\mathrm{ord}(V)}))[\alpha_y^{-1}] = 0$ using Lemma 15.0.18. This will be shown in Lemma 28.2.3 below. Overall, this allows us to prove $\widetilde{C}_*(\mathrm{CHur}^{c,S}_+ /(\alpha_V^{\mathrm{ord}(V)})) = 0$. Now, knowing this in the case that V contains exactly one element in each S component allows us to conclude the proof using Lemma 16.0.3: The point is that if we fix a single component, this implies multiplication by a power of that element induces an isomorphism on homology, and hence that element itself will induce an isomorphism on homology in a slightly larger degree. (See §16.0.2 for a similar deduction of the theorem Theorem 16.0.1.) □

To prove the homology of Hurwitz space modules stabilize, we therefore only need to prove the next lemma:

Lemma 28.2.3. *Suppose V contains an element in each S-component of c, and $y \in c - V$. Assume inductively that $\widetilde{C}_*(\mathrm{CHur}_+^{c,S} / (\alpha_{V'}^{\mathrm{ord}(V')})) = 0$ for all V' with $|V'| > |V|$ then $\widetilde{C}_*(\mathrm{CHur}_+^{c,S} / (\alpha_V^{\mathrm{ord}(V)}))[\alpha_y^{-1}] = 0$*

Proof. The proof is similar to Lemma 16.0.6 so we will be brief. As argued there, it suffices to prove $\widetilde{C}_*(\mathrm{CHur}_+^{c,S} / (\alpha_V^{\mathrm{ord}(V)}))[\alpha_y^{-1} : y \in c - V] = 0$. We prove this in Lemma 28.2.4 and Lemma 28.2.6. □

The next lemma's proof is very similar to Lemma 16.0.7, so we leave it as an exercise.

Lemma 28.2.4. *If $V \subset c$ contains an element in each S-component of c, and V is not a subrack, then $\widetilde{C}_*(\mathrm{CHur}_+^{c,S} / (\alpha_V^{\mathrm{ord}(V)}))[\alpha_y^{-1} : y \in c - V] = 0$.*

Exercise 28.2.5. Give a proof of the above lemma.

Hence, we are reduced to the crucial statement, which is as follows:

Lemma 28.2.6. *If $V \subset c$ contains an element in each S-component of c, and V is a subrack, then $\widetilde{C}_*(\mathrm{CHur}_+^{c,S} / (\alpha_V^{\mathrm{ord}(V)}))[\alpha_y^{-1} : y \in c - V] = 0$.*

In order to prove the above lemma, we need the following notion of a certain module S' associated to the normalizer of a subrack. For this, the reader may wish to recall the notion of a subset from Definition 27.1.1.

Definition 28.2.7. Let c be a rack and S be a bijective Hurwitz module over c. For $c' \subset c$ a subrack, define $S_{c'}$ to be the bijective Hurwitz module over c' which is maximal among all subsets, $(c', S_{c'}) \subset (c, S)$.

Suppose c is a rack, $c' \subset c$ is a subrack. Let $S = (\Sigma^1_{g,f}, \{T_n\}_{n \in \mathbb{Z}_{\geq 0}}, \{\psi_n\}_{n \in \mathbb{Z}_{\geq 0}})$ be a bijective Hurwitz module over c and $S_{c'} = (\Sigma^1_{g,f}, \{T'_n\}_{n \in \mathbb{Z}_{\geq 0}}, \psi'_n)$ be the system over c' defined above. Then define

$$S' := (\Sigma^1_{g,f}, \{N_c(c')^n \times T'_0\}_{n \in \mathbb{Z}_{\geq 0}}, \{\psi_n|_{N_c(c')^n \times T'_0}\}_{n \in \mathbb{Z}_{\geq 0}}). \tag{28.1}$$

It turns out this is a bijective Hurwitz module, though it is somewhat tricky to verify this.

Proof of Lemma 28.2.6. Let S' be the bijective Hurwitz module obtained from S associated to $N_c(V)$ as defined in (28.1). We consider the map

$$f_{S,S'} : \widetilde{C}_*(\mathrm{Hur}_+^{c,S} / (\alpha_V^{\mathrm{ord}(V)}))[\alpha_y^{-1} : y \in c - V] \to \widetilde{C}_*(\mathrm{Hur}_+^{N_c(V),S'} / (\alpha_V^{\mathrm{ord}(V)}))[\alpha_y^{-1} : y \in c - V]$$

The map will take the connected part of the source (which we want to show vanishes) to the connected part of the target. The connected part of the target vanishes by the following lemma:

Lemma 28.2.8. *Suppose $c' \subset c$ is a subrack which is not a union of S-components of c and let $(N_c(c'), S') \subset (c, S)$ be the associated subset as in* (28.1). *Then we cannot have equality $N_c(c') = c$ and $S' = S$ as bijective Hurwitz modules.*

Proof. By Lemma 16.0.10, we must have that $c' \subset c$ is a union of components of c. Suppose T_0' is the 0 set of S' and T_0 is the 0 set of S. By definition of the S-components of c, there must be some $t \in T_0$, $x \in c'$, and $\gamma \in B_1^{\Sigma^1_{g,f}}$ so that $\sigma_t^\gamma(x) \notin c'$. (Recall σ_t^γ is part of the structure of the bijective Hurwitz module defined in Definition 26.0.3, which records the action of $B_1^{\Sigma^1_{g,f}} \times T_0$ on c.) Therefore, $t \notin T_0'$ and so $T_0' \neq T_0$ and hence $(c', S') \subsetneq (c, S)$. □

Hence, to conclude our proof, we only need to check $f_{S,S'}$ is an equivalence. Note that since V is a subrack of c, $c' := c - V$ is also a subrack. Take $R = \widetilde{C}_*(\mathrm{Hur}_+^c), S = \widetilde{C}_*(\mathrm{Hur}_+^{N_c(c')})$, (where $\widetilde{C}$ is the reduced chains functors) I_S the ideal generated by the classes of α_x for $x \in V$, and I_R be the ideal generated by the classes of α_x for $x \in V$, and $M = \widetilde{C}_*(\mathrm{Hur}_+^{c,S} / (\alpha_V^{\mathrm{ord}(V)}))[\alpha_y^{-1} : y \in c - V])$. A general result in descent (which goes into the proof of Lemma 15.0.10) says that in order to identify M with $S \otimes_R M$ it suffices to identify

$$\pi_0(S)^{\otimes_S n} \otimes_S (S \otimes_R M) \simeq \pi_0(R)^{\otimes_R n} \otimes_R M \tag{28.2}$$

for each n, so long as I_R acts nilpotently on $\pi_i(M)$ for each i and I_S acts nilpotently on $\pi_i(S \otimes_R M)$ for each i. We note that the nilpotency of these actions is satisfied since each α_x for $x \in V$ acts nilpotently on M, (since α_x appears in the set we are quotienting M by,) and hence linear combinations of products of the α_x will also act nilpotently.

Therefore, plugging in the values of M, R, S to (28.2) we only need to show

$$C_*(\pi_0 \mathrm{Hur}_+^{c'}[\alpha_{c'}^{-1}]^{\otimes_{\mathrm{Hur}_+^c} n} \otimes_{\mathrm{Hur}_+^c} \mathrm{Hur}_+^{c,S} / (\alpha_V^{\mathrm{ord}(V)}))[\alpha_y^{-1} : y \in c - V])$$

$$\to C_*(\pi_0 \mathrm{Hur}_+^{c'}[\alpha_{c'}^{-1}]^{\otimes_{\mathrm{Hur}_+^{N_c(c')}} n} \otimes_{\mathrm{Hur}_+^{N_c(c')}} \mathrm{Hur}_+^{N_c(V),S'} / (\alpha_V^{\mathrm{ord}(V)}))[\alpha_y^{-1} : y \in c - V])$$

The $n = 1$ case is a direct consequence of Proposition 28.2.9 below. For the case of general n, we can reduce to the $n = 1$ case by using Proposition 16.0.12 to the first n factors. □

One of the key steps in the proof of the monoid version Theorem 16.0.1 was Proposition 16.0.12. The analog for modules is the following, where $N_c(c')$ denotes the set of elements $x \in c$ so that $x \triangleright y \in c'$ for $y \in c'$.

Proposition 28.2.9. *Let c be a rack, $c' \subset c$ be a subrack, S a bijective Hurwitz module over c and S' the corresponding bijective Hurwitz module over $N_c(c')$, as defined in* (28.1). *Then*

(28.3)

$$\left(\pi_0 \operatorname{Hur}^{c'}\right)[\alpha_{c'}^{-1}] \otimes_{\operatorname{Hur}^c_+} \operatorname{Hur}^{c,S}_+ \simeq \left(\pi_0 \operatorname{Hur}^{c'}\right)[\alpha_{c'}^{-1}] \otimes_{\operatorname{Hur}^{N_c(c')}_+} \operatorname{Hur}^{N_c(c'),S'}_+$$

Proof. The proof is quite similar to that of Proposition 16.0.12. Namely, we can use the scanning model for $\left(\pi_0 \operatorname{Hur}^{c'}\right)[\alpha_{c'}^{-1}] \otimes_{\operatorname{Hur}^c_+} \operatorname{Hur}^{c,S}_+$ given by a rectangle with a label in $\left(\pi_0 \operatorname{Hur}^{c'}\right)[\alpha_{c'}^{-1}]$ on the left and with points labeled by c on the interior, along with an the label of an element $t \in T_0$. This was described earlier in Remark 28.1.6

Definition 28.2.10. An *allowable move* in the above scanning model consists of taking an element $x \in c'$, moving it across the rectangle until it reaches the right hand side in which case it is identified with a lower point on the right hand side and then it is moved backwards into the left boundary. An *allowable path* refers to the path traversed by the point, so it is essentially the same as an allowable move except it does not record the value of $x \in c'$. The *left output* of an allowable move is the corresponding element of c that hits the left boundary after this procedure.

Note allowable moves always go from higher to lower, so their starting point is the higher point.

We picture some possible allowable paths in Figure 17

There is a similar scanning model for $\left(\pi_0 \operatorname{Hur}^{c'}\right)[\alpha_{c'}^{-1}] \otimes_{\operatorname{Hur}^{N_c(c')+}} \operatorname{Hur}^{N_c(c'),S'}_+$ where now the interior only contains elements of $N_c(c')$ and a label from T'_0, the 0 set of S'.

One can show that if we have some point of the source of (28.3) which is not in the target, then there is some sequence of allowable moves so that the left output of the last one is not in c'. With a little more work, it is even possible to show a single allowable move which accomplishes this. The idea of the nullhomotopy is simply to pick an ordering on c' and then to make the lowest allowable move that accomplishes this, and if there are two such allowable moves then make the move where we move the element $x \in c'$ which is earlier in the ordering. If nothing hits the boundary, this will be sent to the basepoint since the left output will not be in c' and hence act on the left boundary by sending it to the basepoint.

FIGURE 17. Here we visualize some allowable paths in a configuration with $f = 2, g = 1$ and two yellow points present in the configuration.

The only issue above is that the homotopy described will not be continuous. However, we will now define a bifiltration so that it is continuous on the associated graded of the bifiltration, and this will be enough to check the claimed statement on the level of homology.

The first grading of the filtration counts the number of elements j_1 below the lowest height h at which we can start an allowable move whose left output leaves c'. The second filtration counts the number of elements j_2 in the rectangle so that if we treated them as lying in c when we move them around to the right boundary, they land in c'.

Now, we check the homotopy is continuous on the associated graded of the above filtration when points in the configuration hit the boundary. If the point hits the left boundary above h, it won't effect the allowable move we can make so the homotopy will be continuous in this case. Second, suppose the point hits the left boundary below h. We can assume the action of the element on the boundary is via an element of c', as otherwise the configuration will go to the basepoint. If it is via a point going to the lower of its two equivalent heights (meaning that two heights are equivalent in the scanning model of the rectangle made by identifying two points on the right hand side), it leaving will decrease the second filtration while fixing the first. On the other hand, if it hits the boundary at the higher of its two equivalent heights, then it decreases the value of the first filtration j_1 while preserving the second. In either case, it decreases the filtration, so we can ignore these moves as the send the configuration to the basepoint. This shows the homotopy is continuous on the associated graded of a filtration, so we obtain the claimed equivalence on homology. □

28.3. **Computing the stable homology of Hurwitz space modules.** Having proven Theorem 27.1.6 that the homology of bijective Hurwitz space modules stabilizes, it only remains to compute the stable value and prove Theorem 27.2.4.

The idea is similar to the proof of Theorem 18.0.9 where we want to average all of the elements appearing in the rectangle.

28.4. **Proof sketch of Theorem 27.2.4.** This is one of the most technical parts of the argument, and so we will not give all the details, which you can find in the paper, but try to emphasize the key ideas.

First, we set up a comparison map as follows. Let c' be an S-component of c. $\mathrm{Hur}^{c,S}[\alpha_{c'}^{-1}] \to \mathrm{Hur}^{c/c',S/c'}[\alpha_{c'/c'}^{-1}] \times_{\pi_0 \mathrm{Hur}^{c/c',S/c'}[\alpha_{c'/c'}^{-1}]} \pi_0 \mathrm{Hur}^{c,S}[\alpha_{c'}^{-1}]$ which we want to show induces a rational equivalence on homology. We view these both as Hur^c modules and use descent along the map $\pi_0 \mathrm{Hur}^c[\alpha_{c'}^{-1}] \to \mathrm{Hur}^c[\alpha_{c'}^{-1}]$. That is via a descent argument, it suffices to check these two modules are equivalent after taking iterated fiber products of $\left(\pi_0 \mathrm{Hur}^c[\alpha_{c'}^{-1}]\right)$ over $\mathrm{Hur}^c[\alpha_{c'}^{-1}]$; that is, it suffices to check

$$\begin{aligned}&\left(\pi_0 \mathrm{Hur}^c[\alpha_{c'}^{-1}]\right)^{\otimes_{\mathrm{Hur}^c[\alpha_{c'}^{-1}]} n+1} \otimes_{\mathrm{Hur}^c[\alpha_{c'}^{-1}]} \mathrm{Hur}^{c,S}[\alpha_{c'}^{-1}] \\ &\to \left(\pi_0 \mathrm{Hur}^c[\alpha_{c'}^{-1}]\right)^{\otimes_{\mathrm{Hur}^c[\alpha_{c'}^{-1}]} n+1} \otimes_{\mathrm{Hur}^c[\alpha_{c'}^{-1}]} \left(\mathrm{Hur}^{c/c',S/c'}[\alpha_{c'/c'}^{-1}] \times_{\pi_0 \mathrm{Hur}^{c/c',S/c'}[\alpha_{c'/c'}^{-1}]} \pi_0 \mathrm{Hur}^{c,S}[\alpha_{c'}^{-1}]\right)\end{aligned}$$

for all nonnegative integers n. It suffices to make the desired identification when $n = 0$ as the other cases will then follow by tensoring with $\left(\pi_0 \mathrm{Hur}^c[\alpha_{c'}^{-1}]\right)^{\otimes_{\mathrm{Hur}^c[\alpha_{c'}^{-1}]} n}$.

We next want to reduce to the proposition below. To make this reduction, we need two ingredients, which we don't explain further. First, we need an argument similar to that proving Lemma 21.0.2 and Lemma 20.2.6 (where we will have to commute tensor products and pullbacks). Second, we need an additional argument where one needs to show that if multiplication by some element of $\pi_0 \mathrm{Hur}^c[\alpha_{c'}^{-1}]$ fixes a component of $\pi_0 \mathrm{Hur}^{c,S}[\alpha_{c'}^{-1}]$ then it acts trivially on the rational homology of $\pi_0 \mathrm{Hur}^c[\alpha_{c'}^{-1}] \otimes_{\mathrm{Hur}^c} \mathrm{Hur}^{c,S}$. Granting these two ingredients (you can see the proof of [LL25b, Theorem 8.3.3] for the first and [LL25b, Proposition 8.3.1] for the second) we reduce to showing the following equivalence:

Proposition 28.4.1. *Let c be a rack, S be a bijective Hurwitz module over c and $c' \subset c$ be an S component. The map*

$$\begin{aligned}&\left(\pi_0 \operatorname{Hur}^c[\alpha_{c'}^{-1}]\right) \otimes_{\operatorname{Hur}^c[\alpha_{c'}^{-1}]} \operatorname{Hur}^{c,S}[\alpha_{c'}^{-1}] \\ &\to \left(\pi_0 \operatorname{Hur}^c[\alpha_{c'}^{-1}]\right) \otimes_{\operatorname{Hur}^c[\alpha_{c'}^{-1}]} \left(\operatorname{Hur}^{c/c',S/c'}[\alpha_{c'/c'}^{-1}] \times_{\pi_0 \operatorname{Hur}^{c/c',S/c'}[\alpha_{c'/c'}^{-1}]} \pi_0 \operatorname{Hur}^{c,S}[\alpha_{c'}^{-1}]\right)\end{aligned}$$

induces a rational homology equivalence.

Proof sketch. As usual, we can create a section to the map and show that the complement of the section is contractible.

To simplify matters, we will focus on the case $f = 1$ and $g = 0$, so the surface has only a single puncture. Our scanned model consists of a module $M := \pi_0 \operatorname{Hur}^c[\alpha_{c'}^{-1}] \otimes_{\mathbb{Z}} \mathbb{Q}$ labeling the left hand side, a collection of points on the interior in $\mathbb{Q}\{c\}$ and a label $s \in \mathbb{Q}\{T_0\}$.

We will slightly further simplify matters by assuming $c = c'$ has a single component with $c = c'$. While this hides a substantial amount of difficulty, we remark it is the only case we use for our application to Poonen-Rains (although there we typically need $f > 1$ since the elliptic curve will have more than 1 place of bad reduction). The reason for working with this case is that it makes the description of $S/c = S/c'$ much simpler. Namely, we just take the equivalence relation generated by $s \sim s'$ if they are related by the action of some element of c moving around the puncture.

The complement to the section consists of linear combinations of configurations where top $i - 1$ points are averaged and the ith point is antiaveraged, or all n points are averaged but s is antiaveraged, meaning that s is a linear combination of elements of T_0 in the same orbit of the action of $c = c'$.

We'll first discuss the case that $i = n$. Here, we essentially want to do a homotopy where we move a point around the boundary of the rectangle. Consider the finite group $G_{T_0}^{c'}$, the automorphisms of T_0 induced by the action of $c' = c$. Any element of $k\{T_0\}$ fixed by this group necessarily is averaged, i.e. is a linear combination of orbits for the c' action on T_0. Therefore, we just need to enact the action of this group. We can choose a representative for each element as a product of elements of c', and then move each of these elements around the boundary of the rectangle one by one, and then average them. The motion of a single element is depicted in blue in Figure 18.

As usual, these representatives will not be closed under the action of c', hitting the boundary, but we can add in their translates by the c' action to assume that they are. Overall, this allows us to nullhomotope this part of the complex.

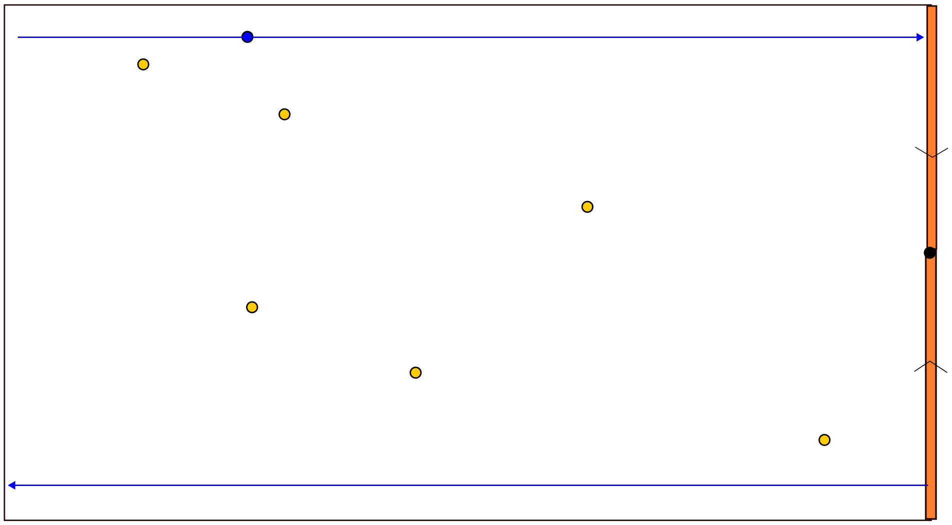

FIGURE 18. This depicts the action of moving a single element around the rectangle. A sum of iterations of this results in averaging the label for T_0.

We conclude by considering the case $i < n$. Here, the ith element $v_i \in k\{c'\}$ is antiaveraged. So, if we apply $U_{c'}(v_i) = \frac{1}{|G_c^{c'}|}\sum_{g \in G_c^{c'}} g \triangleright v_i$ we get 0. Thus, we only want to enact this averaging procedure. As usual, we can try to do this by writing elements of g as products of elements of c', and moving them on the allowable paths directly below and above v_i. One such movement is depicted in Figure 19.

As usual, this will not be continuous when things below and above v_i hit the boundary, but we can try to fix this by adding in translates of the words by elements of c'. This works when elements below and above v_i hit the boundary, but there is a small issue when v_i hits the boundary. When v_i hits the boundary by moving left, the result is just 0 since v_i is antiaveraged. However, when v_i hits the boundary by moving right, it can be nonzero. In this case, however, there is only one more observation we have to make. After v_i hits the boundary by moving right, it is a linear combination of elements which are related by the action of moving elements of c' around the boundary of the rectangle. Namely, if $v_i = u - w$, for example, then it will be the difference of moving u around the rectangle's boundary and moving w around the rectangle's boundary. Hence, we can kill this element via a further nullhomotopy similar to the case $i = n$ by averaging over moving elements around the boundary of the rectangle. This finally concludes the construction of the nullhomotopy and our sketch of the proof. $\square$

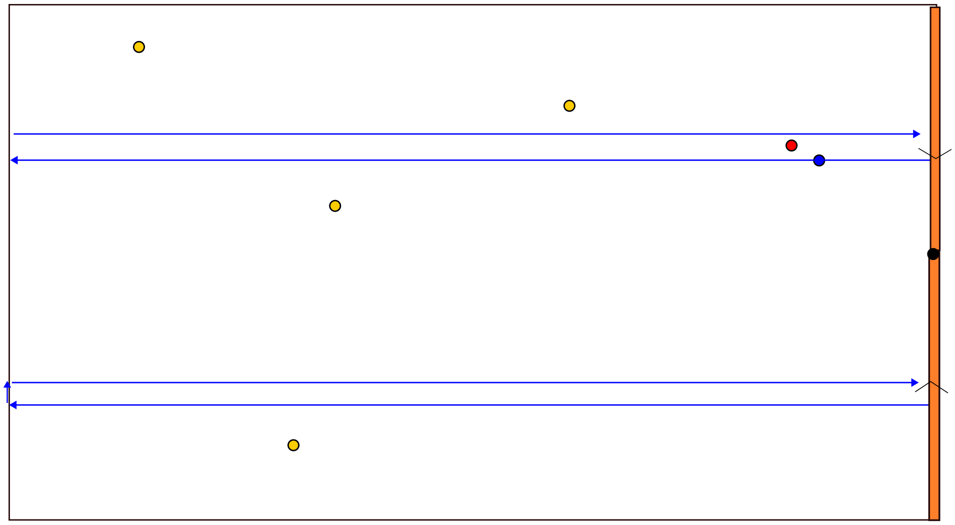

FIGURE 19. This depicts the homotopy used to kill configurations with antiaveraged elements. In the picture, the yellow dots are averaged and the red dots are antiaveraged. We move the blue dot around the highest antiaveraged element, then iterate to get every element of the relevant structure group, and average over the group.

29. COMPARING COHOMOLOGY OF HURWITZ SPACES IN DIFFERENT CHARACTERISTICS

One key result we used throughout these notes was the comparison of cohomology of Hurwitz spaces in characteristic 0 and characteristic p. Here is the statement:

Theorem 29.0.1. *Let G be a group, and $c := G - \mathrm{id}$. Then $H^i(\mathrm{Hur}^{G,c}_{n,\overline{\mathbb{F}}_q}, \mathbb{Q}_\ell) \simeq H^i(\mathrm{Hur}^{G,c}_{n,\mathbb{C}}, \mathbb{Q}_\ell)$.*

Definition 29.0.2. If X is a smooth scheme over a base B, a relative normal crossing compactification of X over B is a smooth proper scheme Y containing X as a fiberwise dense open subscheme so that $D = Y - X$ is a normal crossing divisor, meaning that around every point of D, étale locally D is isomorphic to a union of hyperplanes meeting transversely, i.e. over B it is étale locally of the form $V(x_1, \ldots, x_n)$ in $\mathbb{A}^{n+m}$ for some $m \geq 0$.

This follows from a general result that if we have a smooth scheme (or DM stack) $X \to B$ with a relative normal crossings compactification then the cohomology of any two geometric fibers of $X \to B$ are isomorphic. You can think about this general result first in the smooth and proper case over $B = \operatorname{Spec}\mathbb{C}$. Then, it simply amounts to Ehresmann's theorem, that is smooth proper map is locally a product in the analytic topology, so any two fibers will have isomorphic cohomology. to show that each stratum One can prove a version of Ehresmann's theorem over general bases using the formalism of vanishing cycles (and some results which also go into the smooth base change theorem). And one can also generalize this to the relative normal crossings case. Therefore, the key ingredient we need to prove the comparison on cohomology is that Hurwitz spaces have a relative normal crossings compactification.

Theorem 29.0.3. *Let $B := \mathbb{Z}[1/\#G]$, G be a group, and $c := G - \mathrm{id}$. Then $\mathrm{Hur}^{G,c}_{n,B}$ has a relative normal crossings compactification over B.*

Remark 29.0.4. A much more general version of Hurwitz spaces has relative normal crossings compactifications, where one considers covers of higher genus, punctured curves, instead of just $\mathbb{A}^1$, which we think of as a once punctured $\mathbb{P}^1$. This generalization applies, for example, to constructing a compactification of $\mathrm{Conf}_{n,U/B}$ which was needed in for our application to the Poonen-Rains heuristics. In the case the curve is proper, and G is trivial, the compactification we construct agrees with the Fulton-MacPherson compactification of configuration space.

The proof is somewhat technical and involves a bit of log geometry, so we will just sketch the main idea of the proof. We start by constructing the relative normal crossings compactification. We want to study G covers of $\mathbb{A}^1$ branched at n points. We think of these as G covers of $\mathbb{P}^1$ branched at $n+1$ points, but with ∞ as a special point. We view $[\mathbb{P}^1/G]$ as a quotient stack with the trivial G action on $\mathbb{P}^1$. We let $K_{n+1}([\mathbb{P}^1/G], 1)$ denote the stack of balanced twisted stable maps whose S points are described as follows: Suppose we are given a map $S \to B$. Then, $K_{n+1}([\mathbb{P}^1/G], 1)(S)$ is the groupoid of representable maps $h : \mathcal{X} \to [\mathbb{P}^1/G]$ from an $n+1$ pointed balanced twisted curve (where twisted means the curve can be stacky and also have stacky structures at the nodes and balanced is a condition about how the nodes are glued together, formal definitions can be found in [ACV03]) such that

1. The coarse space of $\mathcal{X}$ is a curve $f : X \to S$
2. The fibers of f have genus 0
3. The map f factors through a map $X \to \mathbb{P}^1$ which is an isomorphism on one component and contracts the other components.

There is an evaluation map $\pi : K_{n+1}([\mathbb{P}^1/G], 1) \to \mathbb{P}^1$ which sends the stable map to the value of the $n+1$st section, which we think of as mapping to ∞. Now, define

$$K_{n+1}([\mathbb{P}^1/G], \infty, 1) := \pi^{-1}(\infty) = K_{n+1}([\mathbb{P}^1/G], \infty, 1) \times_{\pi, \mathbb{P}^1, \infty} B.$$

which just takes the closed substack where the $n+1$th point maps to ∞. The open subset of this stack where the first n points are distinct from each other and from ∞ is an S_n cover of $\mathrm{Hur}^{G,c}_{n,B}$. Therefore, in order to prove Theorem 29.0.3, it suffices to prove the following:

Theorem 29.0.5. *The stack $K_{n+1}([\mathbb{P}^1/G], \infty, 1)$ is smooth and proper, and it is a relative normal crossings compactification of the open locus where then n points are distinct from each other and ∞.*

Proof. The proof of this is not the part that gets rather technical with log geometry, but the idea is not very difficult. Properness is automatic from general results about twisted stable maps and the real issue is to show that the stack is smooth and that it is a relative normal crossings compactification of the specified open. Recall being a relative normal crossings compactification means that the complement of the open looks étale locally like a union of hyperplanes.

In order to prove something is smooth, the typical strategy is to use deformation theory and to show that the obstructions to deformations vanish. So that would be a good strategy to show smoothness of the compactification. But how can one show the boundary divisor is relative normal crossings? It is conceivable that one could show this by giving a moduli interpretation to the boundary in terms of how many points collide. It seems like it should be possible but a bit tricky to pin this down exactly. However, there is another approach. It turns out there is a notion of a map being log smooth. Normal crossings compactifications are log smooth. And it turns out that in good situations (such as the current situation when the characteristic monoid is given by a power of the natural numbers) being log smooth implies the map is given by a relative normal crossings compactification.

Therefore, it is enough to check $K_{n+1}([\mathbb{P}^1/G], \infty, 1)$ is log smooth over B, where the log structure is given by the divisor parameterizing singular curves. To do this, we analyze the deformation theory. In general, given a map of smooth curves $f : X \to Y$, the deformation theory is governed by the cotangent complex, which is $\mathbb{L} := [f^*\Omega_Y \xrightarrow{\alpha} \Omega_X]$ in degree -1 and 0. To check the moduli space of these covers is smooth, one wants to check $\mathrm{Ext}^2(\mathbb{L}, \mathscr{O}_X) = 0$. In this case, α is injective, so it is equivalent to its cokernel $\operatorname{coker} \alpha$, so we just want to check $\mathrm{Ext}^2(\operatorname{coker} \alpha, \mathscr{O}_X) = 0$. This holds by Serre duality since it is identified with $H^{-1}(\mathscr{F} \otimes \omega_X) = 0$.

Now, one wants to generalize the above computation to the log setting. Here, it turns out the log deformation theory is governed by the log cotangent sheaf. Consider a twisted stable map $\mathcal{X} \to [\mathbb{P}^1/G]$. which ends up being the follow the sheaf on the stacky curve $h : \mathcal{X} \to [\mathbb{P}^1/G]$ from the definition of $K_{n+1}([\mathbb{P}^1/G], \infty, 1)$. Then the log cotangent complex is $\mathbb{L}^{\log} := h^* \left(\omega_{[\mathbb{P}^1/G]}([\infty/G])\right) \to \omega_{\mathcal{X}}(\mathcal{D} + \mathcal{E})$ for $\mathcal{D}$ the divisor corresponding to the first n sections and $\mathcal{E}$ the divisor corresponding to the last section. So our goal is to show $\operatorname{Ext}^2(\mathbb{L}^{\log}, \mathscr{O}_X) = 0$. In turn, if $f : X \to \mathbb{P}^1$ is the coarse space of $\mathcal{X}$ and D, E are the divisors corresponding to the first n and last section, and we let $\mathbb{L}_f^{\log} := f^* \left(\omega_{[\mathbb{P}^1/G]}([\infty/G])\right) \xrightarrow{\phi} \omega_X(D + E)$, then one can identify $\operatorname{Ext}^2(\mathbb{L}^{\log}, \mathscr{O}_{\mathcal{X}}) \simeq \operatorname{Ext}^2(\mathbb{L}_f^{\log}, \mathscr{O}_X)$. So we only need show the latter vanishes. Now, the map ϕ has a kernel supported on the contracted components and a finite cokernel Q. We have $\operatorname{Ext}^2(Q, \mathscr{O}_X) = 0$ for the same reason as above while one can separately show that $\operatorname{Ext}^2(K[1], \mathscr{O}_X) = \operatorname{Ext}^1(K, \mathscr{O}_X) = 0$ by showing that K is sufficiently positive, as we explain a bit more. Let $Y \subset X$ be the smallest subset of components containing all the marked points such that Y is connected. Let $W \subset X$ denote the union of components not in Y. Then one can calculate $K = i_* \mathscr{O}_W(-(Y \cap W))$ for $i : W \to X$ the inclusion. Using Serre duality, we can calculate

$$\begin{aligned}
\operatorname{Ext}^2(i_*\mathscr{O}_W(-(Y \cap W))[1], \mathscr{O}_X) &\simeq \operatorname{Ext}^1(i_*\mathscr{O}_W(-(Y \cap W)), \mathscr{O}_X) \\
&\simeq H^0(X, i_*\mathscr{O}_W(-(Y \cap W)) \otimes \omega_X)^\vee \\
&\simeq H^0(W, \mathscr{O}_W(-(Y \cap W)) \otimes \omega_X|_W)^\vee \\
&\simeq H^0(W, \mathscr{O}_W(-(Y \cap W)) \otimes \omega_W(Y \cap W))^\vee \\
&\simeq H^1(W, \mathscr{O}_W) \\
&= 0.
\end{aligned}$$

□

30. ACKNOWLEDGEMENTS

I thank Ishan Levy for embarking on these projects with me and for figuring out with me most of the material present in these notes. Many of the novel ideas here are due to him. I would like to thank Xinyu Fang for pointing out quite a number of typos, errors, and confusing points in these notes. I also thank Niven Achenjang for a number of helpful comments. I thank Daniel Hu, Jack Miller, and Melanie Wood, for pointing out some typos. I would like to thank Ishan Levy, Akhil Matthew, and Andy Senger for help understanding the basics of higher algebra discussed in these notes.

REFERENCES

[A'C79] N. A'Campo. Tresses, monodromie et le groupe symplectique. *Comment. Math. Helv.*, 54(2):318–327, 1979.

[Ach06] Jeffrey D. Achter. The distribution of class groups of function fields. *J. Pure Appl. Algebra*, 204(2):316–333, 2006.

[Ach23] Niven Achenjang. The average size of 2-selmer groups of elliptic curves in characteristic 2. *arXiv preprint arXiv:2310.08493v2*, 2023.

[ACV03] Dan Abramovich, Alessio Corti, and Angelo Vistoli. Twisted bundles and admissible covers. volume 31, pages 3547–3618. 2003. Special issue in honor of Steven L. Kleiman.

[BDPW23] Jonas Bergström, Adrian Diaconu, Dan Petersen, and Craig Westerland. Hyperelliptic curves, the scanning map, and moments of families of quadratic L-functions. *arXiv preprint arXiv:2302.07664*, 2023.

[BF06] Aldridge K Bousfield and Eric M Friedlander. Homotopy theory of γ-spaces, spectra, and bisimplicial sets. In *Geometric Applications of Homotopy Theory II: Proceedings, Evanston, March 21–26, 1977*, pages 80–130. Springer, 2006.

[Bha07] Manjul Bhargava. Mass formulae for extensions of local fields, and conjectures on the density of number field discriminants. *Int. Math. Res. Not. IMRN*, (17):Art. ID rnm052, 20, 2007.

[BKL$^+$15] Manjul Bhargava, Daniel M. Kane, Hendrik W. Lenstra, Jr., Bjorn Poonen, and Eric Rains. Modeling the distribution of ranks, Selmer groups, and Shafarevich-Tate groups of elliptic curves. *Camb. J. Math.*, 3(3):275–321, 2015.

[BLR90] Siegfried Bosch, Werner Lütkebohmert, and Michel Raynaud. *Néron models*, volume 21 of *Ergebnisse der Mathematik und ihrer Grenzgebiete (3) [Results in Mathematics and Related Areas (3)]*. Springer-Verlag, Berlin, 1990.

[BS13a] Manjul Bhargava and Arul Shankar. The average number of elements in the 4-selmer groups of elliptic curves is 7. *arXiv preprint arXiv:1312.7333v1*, 2013.

[BS13b] Manjul Bhargava and Arul Shankar. The average size of the 5-Selmer group of elliptic curves is 6, and the average rank is less than 1. *arXiv preprint arXiv:1312.7859v1*, 2013.

[BS15a] Manjul Bhargava and Arul Shankar. Binary quartic forms having bounded invariants, and the boundedness of the average rank of elliptic curves. *Ann. of Math. (2)*, 181(1):191–242, 2015.

[BS15b] Manjul Bhargava and Arul Shankar. Ternary cubic forms having bounded invariants, and the existence of a positive proportion of elliptic curves having rank 0. *Ann. of Math. (2)*, 181(2):587–621, 2015.

[Ces16] Kestutis Cesnavicius. Selmer groups as flat cohomology groups. *J. Ramanujan Math. Soc.*, 31(1):31–61, 2016.

[Cil86] Ciro Ciliberto. Rationally determined line bundles on families of curves. In *The curves seminar at Queen's, Vol. IV (Kingston, Ont., 1985–1986)*, volume 76 of *Queen's Papers in Pure and Appl. Math.*, pages Exp. No. D, 48. Queen's Univ., Kingston, ON, 1986.

[CL84] H. Cohen and H. Lenstra. Heuristics on class groups of number fields. *Number Theory Noordwijkerhout 1983*, pages 33–62, 1984.

[DE96] Steven Diaz and Dan Edidin. Towards the homology of Hurwitz spaces. *J. Differential Geom.*, 43(1):66–98, 1996.

[DH71] H. Davenport and H. Heilbronn. On the Density of Discriminants of Cubic Fields. II. *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences*, 322(1551):405–420, 1971.

[dJ02] A. J. de Jong. Counting elliptic surfaces over finite fields. *Mosc. Math. J.*, 2(2):281–311, 2002. Dedicated to Yuri I. Manin on the occasion of his 65th birthday.

[dJF11] A. J. de Jong and Robert Friedman. On the geometry of principal homogeneous spaces. *Amer. J. Math.*, 133(3):753–796, 2011.

[DP15] Anand Deopurkar and Anand Patel. The picard rank conjecture for the hurwitz spaces of degree up to five. *Algebra & Number Theory*, 9(2):459–492, 2015.

[DW88] Boris Datskovsky and David J. Wright. Density of discriminants of cubic extensions. *J. Reine Angew. Math.*, 386:116–138, 1988.

[EG03] Pavel Etingof and Matıas Grana. On rack cohomology. *Journal of pure and applied algebra*, 177(1):49–59, 2003.

[EL24] Jordan S. Ellenberg and Aaron Landesman. Homological stability for generalized Hurwitz spaces and Selmer groups in quadratic twist families over function fields. *arXiv preprint arXiv:2310.16286v2*, 2024.

[Enr19] Federigo Enriques. Questioni numerative e loro significato nella geometria sopra le curve algebriche. *Rend. Mat. Acc. Lincei*, 28:370 – 374, 1919.

[ETW17] Jordan S. Ellenberg, TriThang Tran, and Craig Westerland. Fox-neuwirth-fuks cells, quantum shuffle algebras, and malle's conjecture for function fields. *arXiv preprint arXiv:1701.04541v2*, 2017.

[EV05] Jordan S. Ellenberg and Akshay Venkatesh. Counting extensions of function fields with bounded discriminant and specified Galois group. In *Geometric methods in algebra and number theory*, volume 235 of *Progr. Math.*, pages 151–168. Birkhäuser Boston, Boston, MA, 2005.

[EVW12] Jordan S. Ellenberg, Akshay Venkatesh, and Craig Westerland. Homological stability for Hurwitz spaces and the Cohen-Lenstra conjecture over function fields, II. *arXiv preprint arXiv:1212.0923v1*, 2012.

[EVW16] Jordan S. Ellenberg, Akshay Venkatesh, and Craig Westerland. Homological stability for Hurwitz spaces and the Cohen-Lenstra conjecture over function fields. *Ann. of Math. (2)*, 183(3):729–786, 2016.

[FLR23] Tony Feng, Aaron Landesman, and Eric M. Rains. The geometric distribution of Selmer groups of elliptic curves over function fields. *Math. Ann.*, 387(1-2):615–687, 2023.

[FM12] Benson Farb and Dan Margalit. *A primer on mapping class groups*, volume 49 of *Princeton Mathematical Series*. Princeton University Press, Princeton, NJ, 2012.

[FRS07] Roger Fenn, Colin Rourke, and Brian Sanderson. The rack space. *Trans. Amer. Math. Soc.*, 359(2):701–740, 2007.

[Gol79] Dorian Goldfeld. Conjectures on elliptic curves over quadratic fields. In *Number theory, Carbondale 1979 (Proc. Southern Illinois Conf., Southern Illinois Univ., Carbondale, Ill., 1979)*, volume 751 of *Lecture Notes in Math*, pages pp 108–118. Springer, Berlin, 1979.

[Hal08] Chris Hall. Big symplectic or orthogonal monodromy modulo l. *Duke Math. J.*, 141(1):179–203, 2008.

[Hay92] David R. Hayes. A brief introduction to Drinfeld modules. In *The arithmetic of function fields (Columbus, OH, 1991)*, volume 2 of *Ohio State Univ. Math. Res. Inst. Publ.*, pages 1–32. de Gruyter, Berlin, 1992.

[HLHN14] Q. P. Hồ, V. B. Lê Hùng, and B. C. Ngô. Average size of 2-Selmer groups of elliptic curves over function fields. *Math. Res. Lett.*, 21(6):1305–1339, 2014.

[HM06] Joe Harris and Ian Morrison. *Moduli of curves*, volume 187. Springer Science & Business Media, 2006.

[Lan21] Aaron Landesman. The geometric average size of Selmer groups over function fields. *Algebra Number Theory*, 15(3):673–709, 2021.

[LL24] Aaron Landesman and Ishan Levy. The Cohen-Lenstra moments over function fields via the stable homology of non-splitting Hurwitz spaces. *arXiv preprint arXiv:2410.22210v3*, 2024.

[LL25a] Aaron Landesman and Ishan Levy. Homological stability for Hurwitz spaces and applications. *arXiv preprint arXiv:2503.03861v2*, 2025.

[LL25b] Aaron Landesman and Ishan Levy. The stable homology of Hurwitz modules and applications. *arXiv preprint arXiv:2510.02068v1*, 2025.

[Lur17] Jacob Lurie. Higher Algebra. `https://www.math.ias.edu/~lurie/papers/HA.pdf`, 2017. Accessed: 2024-10-24.

[LWZB24] Yuan Liu, Melanie Matchett Wood, and David Zureick-Brown. A predicted distribution for Galois groups of maximal unramified extensions. *Invent. Math.*, 237(1):49–116, 2024.

[Mal02] Gunter Malle. On the distribution of Galois groups. *J. Number Theory*, 92(2):315–329, 2002.

[Mul23] Scott Mullane. The Hurwitz space Picard rank conjecture for $d > g - 1$. *Adv. Math.*, 430:Paper No. 109211, 27, 2023.

[Ols16] Martin Olsson. *Algebraic spaces and stacks*, volume 62 of *American Mathematical Society Colloquium Publications*. American Mathematical Society, Providence, RI, 2016.

[Pet17] Dan Petersen. A spectral sequence for stratified spaces and configuration spaces of points. *Geom. Topol.*, 21(4):2527–2555, 2017.

[Poo17] Bjorn Poonen. *Rational points on varieties*, volume 186 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, RI, 2017.

[PR12] Bjorn Poonen and Eric Rains. Random maximal isotropic subspaces and Selmer groups. *J. Amer. Math. Soc.*, 25(1):245–269, 2012.

[PW23] Sun Woo Park and Niudun Wang. On the Average of p-Selmer Ranks in Quadratic Twist Families of Elliptic Curves Over Global Function Fields. *International Mathematics Research Notices*, page rnad095, 05 2023.

[RW20] Oscar Randal-Williams. Homology of Hurwitz spaces and the Cohen-Lenstra heuristic for function fields [after Ellenberg, Venkatesh, and Westerland]. *Astérisque*, (422):Exp. No. 1164, 469–497, 2020.

[Smi22] Alexander Smith. The distribution of ℓ^∞-selmer groups in degree ℓ twist families I. *arXiv preprint arXiv:2207.05674v2*, 2022.

[T15] Seyfi Türkelli. Connected components of Hurwitz schemes and Malle's conjecture. *J. Number Theory*, 155:163–201, 2015.

[Vak01] Ravi Vakil. Twelve points on the projective line, branched covers, and rational elliptic fibrations. *Math. Ann.*, 320(1):33–54, 2001.

[Woo21] Melanie Matchett Wood. An algebraic lifting invariant of Ellenberg, Venkatesh, and Westerland. *Res. Math. Sci.*, 8(2):Paper No. 21, 13, 2021.

[Zyw14] David Zywina. The inverse Galois problem for orthogonal groups. *arXiv preprint arXiv:1409.1151v1*, 2014.