COMPOSITIO MATHEMATICA

Cohen–Lenstra heuristics for étale group schemes and symplectic pairings

Michael Lipnowski and Jacob Tsimerman

Compositio Math. **155** (2019), 758–775.

doi:10.1112/S0010437X19007036

FOUNDATION COMPOSITIO MATHEMATICA

LONDON MATHEMATICAL SOCIETY
EST. 1865

Compositio Math. 155 (2019) 758–775
doi:10.1112/S0010437X19007036

Cohen–Lenstra heuristics for étale group schemes and symplectic pairings

Michael Lipnowski and Jacob Tsimerman

Abstract

We generalize the Cohen–Lenstra heuristics over function fields to étale group schemes G (with the classical case of abelian groups corresponding to constant group schemes). By using the results of Ellenberg–Venkatesh–Westerland, we make progress towards the proof of these heuristics. Moreover, by keeping track of the image of the Weil-pairing as an element of $\wedge^2 G(1)$, we formulate more refined heuristics which nicely explain the deviation from the usual Cohen–Lenstra heuristics for abelian ℓ-groups in cases where $\ell \mid q-1$; the nature of this failure was suggested already in the works of Malle, Garton, Ellenberg–Venkatesh–Westerland, and others. On the purely large random matrix side, we provide a natural model which has the correct moments, and we conjecture that these moments uniquely determine a limiting probability measure.

1. Introduction

In [CL84], Cohen and Lenstra described natural probability measures $m_{\mathrm{CL},u}$ on the set of finite abelian ℓ-groups; the rank u-Cohen–Lenstra measure of every finite abelian ℓ-group A is inversely proportional to $|A|^u \cdot |\mathrm{Aut}(A)|$. The prediction that the distribution of ℓ-parts of class groups of appropriate families of number fields is governed by $m_{\mathrm{CL},u}$ is known as the *Cohen–Lenstra–Martinet conjecture*. Empirically, Cohen and Lenstra observed that $m_{\mathrm{CL},u}$ correctly predicts the distribution of the ℓ-part of class groups of quadratic fields K with $\mathrm{rank}(O_K^\times) = u$, for ℓ an odd prime.[1]

For functions f defined on isomorphism classes of finite abelian ℓ-groups which are absolutely integrable with respect to $m_{\mathrm{CL},u}$, define

$$\mathbf{E}_{n,u}(f) := \lim_{X\to\infty} \frac{\sum_{\deg(K/\mathbb{Q})=n,\mathrm{rank}(O_K^\times)=u,|\mathrm{disc}(K)|<X} f(\mathrm{Cl}(K))}{\sum_{\deg(K/\mathbb{Q})=n,|\mathrm{disc}(K)|<X} 1},$$

assuming the above limit exists. Despite much work on the Cohen–Lenstra heuristics, the only unconditional results known to date are:

- Davenport and Heilbronn's determination of the average size of the 3-torsion subgroup of $\mathrm{Cl}(K)$ for quadratic fields K [DH71]:

$$\mathbf{E}_{2,u}(\#\mathrm{Surj}(\bullet, \mathbb{Z}/3)) = \mathrm{Expectation}_{m_{\mathrm{CL},u}}(\#\mathrm{Surj}(\bullet, \mathbb{Z}/3)),$$

  the cases $u = 0$ and $u = 1$ corresponding to imaginary and real quadratic fields respectively.

Received 25 July 2017, accepted in final form 6 September 2018, published online 20 March 2019.
2010 Mathematics Subject Classification 11G20, 11S31, 20J05, 14H10, 15B52 (primary).
Keywords: curves over finite fields, Cohen–Lenstra heuristics, p-adic random matrices.

[1] If $\ell = 2$, then Gerth conjectured that if one considers the subgroup of the Class group consisting of elements which are squares, the same conjecture holds. This 'degenerate' case turns out to be more accessible and there are unconditional results in this direction by Fouvry and Kulners [FK07], Smith [Smi16], Milovic [Mil18], and Klys [Kly16].

- Bhargava's determination of the average size of the 2-torsion subgroup of $\mathrm{Cl}(K)$ for cubic fields K [Bha05]:

$$\mathbf{E}_{3,u}(\#\mathrm{Surj}(\bullet, \mathbb{Z}/2)) = \mathrm{Expectation}_{m_{\mathrm{CL},u}}(\#\mathrm{Surj}(\bullet, \mathbb{Z}/2)),$$

the cases $u = 1$ and $u = 2$ corresponding to cubic fields admitting one and three real embeddings respectively.

Remarkable recent progress toward the Cohen–Lenstra conjecture has been made for class groups of functions fields of curves over finite fields. In this case, using the methods of étale cohomology and by proving results on homological stability, Ellenberg–Venkatesh–Westerland have obtained unconditional results essentially proving that, for every finite abelian group ℓ-group A, the expectation $\mathbf{E}_{2,0}(\#\mathrm{Surj}(\bullet, A))$ is very close to

$$\mathrm{Expectation}_{m_{\mathrm{CL},0}}(\#\mathrm{Surj}(\bullet, A)).$$

To get a handle in the function field case, one looks at all the geometric ℓ-power torsion points of the Jacobian as a module for the Frobenius operator, of which the class group becomes one small piece. As such, one of our main goals in this paper is to generalize these heuristics in the function field case by remembering the entire action of the Frobenius operator. A convenient language for making this precise is that of étale group schemes.

1.1 Étale group schemes

Let $X/\mathbb{F}_q$ be a (smooth, projective, irreducible) curve over a finite field. The class group of X is naturally $J(X)(\mathbb{F}_q)$ where $J(X)$ is the Jacobian of X. One is then naturally led to ask: what is the distribution of $J(X)(\mathbb{F}_q)[\ell^\infty]$ as X varies in some natural family?[2]

The group $J(X)(\mathbb{F}_q)[\ell^\infty]$ equals the kernel of $1-F$ acting on $J(X)(\overline{\mathbb{F}_q})[\ell^\infty]$, where F denotes the Frobenius operator. The ℓ-part of the class group is thus identified as the '1-eigenspace' of the Frobenius operator. However, there is no need to restrict attention to this particular eigenspace; we could consider all eigenspaces at once. More generally, we consider monic polynomials $P(x) \in \mathbb{Z}_\ell[x]$ for which $P(0)$ is invertible and consider the kernel of $P(F)$. Notably, when we take $P(x) = 1 - x^m$, we recover $J(X)(\mathbb{F}_{q^m})[\ell^\infty]$ together with the action of F.

Since the Jacobian is self-dual, all '$\mathrm{Ker}\, P(F)$ information' is contained in the cokernel of $P(F)$ on the Tate module of the Jacobian, and the latter is the object we actually study.[3] This cokernel is a module with the action of a Frobenius operator. Moreover, because the Frobenius operator acts invertibly, the cokernel may be thought of as an étale group scheme, which is the quotient of $J(X)[\ell^n]$ by $P(F)$, for any large enough n. By using the results of [EVW16] we obtain information about the distribution of these group schemes.

1.2 Symplectic pairings and the Cohen–Lenstra–Martinet heuristics

In the case where the base number field contains roots of unity, Malle [Mal10] presented computational evidence which cast doubt on the Cohen–Lenstra–Martinet heuristics. Malle refined these heuristics, giving a different random model involving the symplectic group, and there has been much evidence that Malle's refinement is correct [Ach06, Gar15]. We present a

[2] One often takes some Hurwitz scheme for the family X varies in, as an analogue of looking at number fields of fixed degree. One could instead the study the statistics of the family of all curves of a fixed genus, but this seems very difficult, and it is not clear what behavior to expect.

[3] Note that one issue which arises now is that the kernel could be infinite, but this should arise very infrequently, so that the distribution we obtain should be supported on finite modules.

refined Cohen–Lenstra heuristic in all cases, which we believe nicely explains these discrepancies. Our refinement involves not just the class group but also a certain naturally defined element of its exterior square. Thus, we get a distribution not just on abelian groups G, but on pairs (G, ω_G) where $\omega_G \in \wedge^2 G$. In the function field setting, this is done as follows.

The Weil-pairing on $J(X)[\ell^n]$ can be thought of as a global section of a certain naturally defined group scheme $\wedge^2 J(X)[\ell^n](1)$ which we define in § 4. This pushes forward, so we naturally get a global section ω of $\wedge^2(J(X)[\ell^n]/(1-F))(1)$.

If $\ell \nmid q-1$ then ω is forced to be 0, and one expects the statistics of $J(X)(\mathbb{F}_q)[\ell^\infty]$, for X varying through many natural families of curves of growing genus, to be governed by the usual Cohen–Lenstra measure. If $\ell | q-1$, however, ω_X is an interesting invariant of the class group of X. It becomes natural to contemplate (Cohen–Lenstra-like) probability distributions on pairs (G, ω), where G is a finite abelian ℓ-group and ω is an element of $\wedge^2 \underline{G}(1)(\mathbb{F}_q)$.[4] In fact, it is natural to combine this decoration with the generalization to arbitrary étale group schemes and this is what we carry out.

1.3 Refined Cohen–Lenstra conjecture

1.3.1 *Modelling the class group.* Ideally, we would like to model the triples

$$(J(X)[\ell^\infty](\overline{\mathbb{F}_q}), \mathrm{Frob}_q, \text{Weil pairing})$$

for X varying through geometric families of curves of growing genus. The groups $J(X)[\ell^\infty](\overline{\mathbb{F}_q})$ are 'too big' to admit meaningful statistics; we instead isolate the pieces $J(X)[\ell^\infty](\mathbb{F}_{q^m}) = \mathrm{Ker}(1 - F^m \mid J(X)[\ell^\infty](\overline{\mathbb{F}_q}))$. These objects are smaller, and we expect that they are governed by a natural limiting probability distribution in reasonable families of curves of growing genus.

Of course, there is no reason to single out $1 - F^m$; we could instead look at the kernel of $P(F)$ for any polynomial P (or more generally we could look at the kernels of certain ideals in $\mathbb{Z}_\ell[x]$). While we do this in § 3, we keep things somewhat classical and focus in § 1.3.2 on the special case $P(F) = 1 - F^m$.

1.3.2 *A refined distribution.* Consider the family $\mathcal{H}_g$ of hyperelliptic curves $y^2 = x^{2g+1} + a_{2g}x^{2g} + \cdots + a_0$ of odd, genus g hyperelliptic curves. Let $\mathcal{E}_m$ denote the collection of triples (G, F, ω), satisfying:

- G is a finite abelian ℓ-group;
- F is an automorphism of G of order m;
- $\omega \in \wedge^2 G$;
- $F\omega = q\omega$.

The collection $\mathcal{E}_m$ naturally has the structure of a category for which the notion of surjection is sensible. See § 3 for details.

If J denotes the Jacobian of a curve X in the collection $\mathcal{H}_g(\mathbb{F}_q)$, the triple

$$T(X) := (J(\mathbb{F}_{q^m})_\ell, \mathrm{Frob}_q, \text{Weil pairing pushed forward to } J(\mathbb{F}_{q^m})_\ell)$$

belongs to $\mathcal{E}_m$; see § 4 for a more precise description of the constituents of this triple.

Remark 1. If m is divisible by positive integers $m_1, \ldots, m_k$, then the distribution of $T(X)$ determines the joint distribution of the tuple

$$J(X)[\ell^\infty](\mathbb{F}_{q^{m_1}}), \ldots, J(X)[\ell^\infty](\mathbb{F}_{q^{m_k}}).$$

[4] For constant group schemes $\underline{G}$, elements of $\wedge^2 \underline{G}(1)(\mathbb{F}_q)$ are naturally identified with elements of $\wedge^2 G$ satisfying $\omega = q\omega$.

https://doi.org/10.1112/S0010437X19007036 Published online by Cambridge University Press

Understanding the latter joint distribution was our initial motivation for the present paper.

CONJECTURE 1.1. There is a probability measure μ_m on $\mathcal{E}_m$ for which the discrete probability measures

$$\frac{1}{|\mathcal{H}_g(\mathbb{F}_q)|}\sum_{X\in\mathcal{H}_g}\delta_{T(X)}$$

on $\mathcal{E}_m$ weak $*$-converge to μ_m as $g\to\infty$, i.e. the proportion of curves $X\in\mathcal{H}_g(\mathbb{F}_q)$ for which $T(X)$ is isomorphic to T approaches $\mu(T)$ as $g\to\infty$.

Furthermore, μ_m is characterized by the following property: for every $T\in\mathcal{E}_m$, the μ_m-expected number of surjections to T equals 1.

In § 3, we give a putative construction of μ_m via random matrices in the style of Friedman and Washington [FW89]. Furthermore in the case when $\ell\nmid q^m-1$, we prove that our construction indeed produces a probability measure satisfying the uniqueness property from Conjecture 1.1. (see Theorem 3.2, and Conjecture 3.1).

1.4 Results

Our main result is as follows. See § 4 for precise definitions.

THEOREM 1.1 (Corollary 4.7). *Let ℓ be an odd prime. Let G be a finite étale group scheme over $\mathbb{F}_q$ of order ℓ^n, and $\omega_G\in(\wedge^2G)(1)(\mathbb{F}_q)$. For each g, let $\operatorname{Avg}(G,\omega_G,g,q)$ denote the average number of surjections from $\operatorname{Pic}^0(C)[\ell^n]$ to G which push forward the Weil-paring to ω_G, where C varies over hyperelliptic curves of genus g.*

Let $\delta^{\pm}(q,\omega_G)$ be the lower and upper limits of $\operatorname{Avg}(G,\omega_G,g,q)$ as $g\to\infty$. Then as $q\to\infty$ and n stays fixed, $\delta^+(q,\omega_G)$ and $\delta^-(q,\omega_G)$ converge to 1.

Our proof of this theorem closely follows the strategy of [EVW16]. We represent the averages in question in terms of points on a moduli space we construct. These moduli spaces turn out to be twists of the moduli spaces that appear in the work of Ellenberg–Venkatesh–Westerland [EVW16]. We can therefore directly apply their results on cohomology bounds, and the theorem follows from the Lefschetz trace formula once we identify the number of connected components of these moduli spaces.

In §§ 2 and 3 we develop foundational results on Cohen–Lenstra measures in the context of our decorated étale group schemes. We obtain the strongest results in the case where $\omega_G\in\wedge^2G(1)$ is forced to be 0, which happens 'generically'.

THEOREM 1.2 (Theorem 3.2). *Let $P(x)\in\mathbb{Z}_\ell[x]$ be a monic polynomial, such that $P(q)$ is not divisible by ℓ, and assume that ℓ is odd. Let $R=\mathbb{Z}_\ell[x]/P(x)$. There exists a unique probability measure μ, supported on finite R-modules, such that for any finite R-module M, the expected number of surjections from a μ-random module to M is 1. Moreover, μ is supported on precisely the modules of projective dimension 1, and assigns such a module M measure $\mu(M)=c/\#\operatorname{Aut}(M)$ where $c=\prod_{k_j}\prod_{i=1}^{\infty}(1-|k_j|^{-i})$ and the product is over the finite residue fields of R.*

As a consequence of these theorems, we obtain in Proposition 3.4 similar results on limiting measures for our decorated étale group schemes.

As a concrete application of our methods, we prove the following result on the independence of the class group of a hyperelliptic curves and its quadratic twist.

https://doi.org/10.1112/S0010437X19007036 Published online by Cambridge University Press

Theorem 1.3 (Proposition 5.1). *Suppose $\ell \nmid q^2 - 1$ and that $\ell \neq 2$. Let $\epsilon > 0$. Fix a finite set S of finite abelian ℓ-groups. For a curve C over $\mathbb{F}_q$, denote by C^σ the quadratic twist of C.*

There exists $Q(S, \epsilon) \gg 0$ such that if $q, g \geqslant Q(S, \epsilon)$ and $A, B \in S$,

$$\left| \operatorname{Prob}(\operatorname{Jac}(C)(\mathbb{F}_q)_\ell \cong A \text{ and } \operatorname{Jac}(C^\sigma)(\mathbb{F}_q)_\ell \cong B) - \frac{c_R}{\#\operatorname{Aut}_{\mathbb{Z}_\ell}(A)\#\operatorname{Aut}_{\mathbb{Z}_\ell}(B)} \right| < \epsilon,$$

where C varies over hyperelliptic curves of genus g, and c_R is the normalizing constant from Theorem 2.2. i.e. the class groups $\operatorname{Jac}(C)(\mathbb{F}_q)_\ell$ and $\operatorname{Jac}(C^\sigma)(\mathbb{F}_q)_\ell$ behave almost independently for g sufficiently large.

1.5 Plan of the paper

- In § 2 we present a generalization of the usual Cohen–Lenstra measure to rings which are finite over $\mathbb{Z}_\ell$.
- In § 3 we construct a random model for pairs $(G, \omega \in \wedge^2(G)(1))$ where G is a module over a ring $\mathbb{Z}_\ell[F]/P(F)$. We conjecture that our model yields a unique measure with a 'moments equal 1' property, and using our results in § 2 we prove this uniqueness in the case where ω is forced to be 0; in other words, when we don't have to keep track of any symplectic structure, so we can revert to a linearized model.
- In § 4 we use the work of Ellenberg–Venkatesh–Westerland to prove results analogous to theirs in the direction of Cohen–Lenstra for function fields, for our refined distributions.
- In § 5 we present some applications, notably to the independence of the ℓ-part of the class group of a hyperelliptic curve and its quadratic twist.

2. Large random matrices over rings

2.1 Summary

The purpose of this section is to generalize the Cohen–Lenstra measure for finite abelian ℓ-groups to the case of finite R-modules for certain rings R, finite over $\mathbb{Z}_\ell$. This measure has the nice property that for every finite R-module M, the expected number of R-module surjections to M is 1. *The support of this measure is not full*, but on the support the measure of M is proportional to $1/\#\operatorname{Aut} M$.

2.2 The Cohen–Lenstra measure for R-modules

Let R be a finite, local $\mathbb{Z}_p$-algebra, with residue field $\mathbb{F}_R$ such that $\mathbb{Z}_p \subset R$. Let S_R be the set of all finite R-modules and define a measure $\mu_{R,N}$ on S_R as follows. Let $\phi_N : \operatorname{End}_R(R^N) \to S_R$ be defined by $G \to \operatorname{Coker} G$. Then $\mu_{R,N}$ is the pushforward of Haar measure under ϕ_N. Recall that since R is a local ring all projective modules are free. Recall also that we say that a module $M \in S_R$ has projective dimension 1 if it has a projective (free) resolution of length 1: $0 \to F_1 \to F_2 \to M \to 0$. Call T_R the set of modules $M \in S_R$ which occur in the image of ϕ_N for some N. Note that if R is torsion free, T_R coincides with the set of modules of projective dimension 1.

We can give a simple homological criterion for a finite module M to occur in T_R. Define $d_M = \dim_{\mathbb{F}_R} \operatorname{Tor}^1_R(M, \mathbb{F}_R) - \dim_{\mathbb{F}_R} M \otimes_R \mathbb{F}_R$.

Lemma 2.1. *For all finite modules M, $d_M \geqslant 0$, and a module M occurs in T_R if and only if $d_M = 0$.*

Proof. For all finite M, we can find a surjection $R^N \to M$. Thus, we have an exact sequence $0 \to U \to R^N \to M$. Tensoring with $\mathbb{F}_R$ and taking the associated long exact sequence, we see that $\dim_{\mathbb{F}_R} \operatorname{Tor}^1_R(M, \mathbb{F}_R) - \dim_{\mathbb{F}_R} M \otimes_R \mathbb{F}_R = \dim_{\mathbb{F}_R}(U \otimes_R \mathbb{F}_R) - N$. Since $M[1/p] = 0$, U has $R[1/p]$-rank equal to N, and so a minimal generating set for U consists of at least N elements, which means $\dim_{\mathbb{F}_R}(U \otimes_R \mathbb{F}_R) - N \geqslant 0$ by Nakayama's lemma. This shows $d_M \geqslant 0$.

Now, if $M \in T_R$, then we can find an exact sequence $0 \to K \to R^N \to R^N \to M \to 0$ for some R-module K. Thus we can take the U in the above paragraph to be R^N/K, and thus be generated by N elements. Thus, in this case, $d_M = 0$.

Conversely, if $d_M = 0$, then the U in the first paragraph must be generated by N elements and so is a quotient of R^N. Thus $M \in T_R$.

Now, if R is torsion free, then as already mentioned T_R coincides with the set of finite modules of projective dimension 1, which is equivalent to $\operatorname{Tor}^2_R(M, F) = 0$. □

Remark 2. We point out that another natural construction of R-modules, at least in the case $R = \mathbb{Z}_p[F]/(P(F))$, is as follows: one can take a random map $A \in \operatorname{End}(\mathbb{Z}_p^d)$, and consider $\operatorname{Coker} P(A)$ as a module over R, with F acting as A. In fact, this more directly mirrors what occurs in the geometric cases we consider, where $\mathbb{Z}_p^{2g}$ occurs as a Tate module and A as the Frobenius endomorphism. This turns out to be more difficult to study, which is why we focus on the model we have presented. However, one can realize $\operatorname{Coker} P(A)$ in our context as the Cokernel of $F - A$ acting on R^d, since

$$R^d/(F-A) \cong \mathbb{Z}_p[F]^d/(F-A, P(F)) = \mathbb{Z}_p[F]^d/(F-A, P(A)) = \mathbb{Z}_p^d/(P(A)).$$

Theorem 2.2. *The $\mu_{R,N}$ converge (in the weak-* topology) to a probability measure μ_R, supported on T_R, such that for $M \in S_R$ we have $\mu_R(M) = c_R/|\operatorname{Aut}_R(M)|$, where $c_R = \lim_{n\to\infty}(|\mathrm{GL}_n(\mathbb{F}_R)|/|M_n(\mathbb{F}_R)|) = \prod_{i=1}^{\infty}(1 - |\mathbb{F}_R|^{-i})$.*

Proof. Let M be an R-module. If M is not in T_R then by definition M never occurs as the cokernel of an endomorphism G and thus cannot be in the support of $\mu_{R,N}$ for any N. So without loss of generality, suppose $M \in T_R$.

Now let us compute $\mu_{R,N}(M)|\operatorname{Aut}_R(M)|$. Consider the space $\operatorname{End}_R(R^N) \times M^N$, with a choice of Haar measure giving total measure $|M|^N$. We can identify M^N with $\operatorname{Hom}(R^N, M)$. Now consider the subset X consisting of (G, ϕ) so that $\operatorname{Im}(G) = \operatorname{Ker}\phi$. The set of all such G such that $\operatorname{Coker} G \cong M$ has measure $\mu_{R,N}(M)$, and for each such G there are $\operatorname{Aut}_R(M)$ choices of ϕ certifying the isomorphism. Thus, the measure of X is $\mu_{R,N}(M)|\operatorname{Aut}_R(M)|$.

We now compute the measure of X in a different way, by fibering over ϕ instead. Now, since $M \in T_R$ there is an exact sequence

$$R^a \xrightarrow{g} R^a \xrightarrow{f} M \to 0$$

for some a. Let C_f be the kernel of f. Take N to be large relative to the above a. The number of maps from R^N to M is $|M|^N$, and with probability tending to 1 as $N \to \infty$, a random such map ϕ is a surjection. Moreover, with probability tending to 1 some subset of size a of the co-ordinates induces the map $f : R^a \to M$. Whenever this happens, we may make a unipotent change of co-ordinates so that the other $N - a$ co-ordinates all map to 0, and thus the kernel is isomorphic to $C_f \oplus R^{N-a}$. Now the measure of all G whose image is contained in $\operatorname{Ker}\phi$ is $|M|^{-N}$. We need to calculate the measure of the subset of those G that give a surjection. We thus need to compute $\mu_{\text{haar}}(\operatorname{Surj}(R^N, C_f \oplus R^{N-a}))/\mu_{\text{haar}}(\operatorname{Hom}(R^N, C_f \oplus R^{N-a}))$. By Nakayama's lemma,

https://doi.org/10.1112/S0010437X19007036 Published online by Cambridge University Press

it is sufficient to tensor everything with the residue field $\mathbb{F}_R$ of R, so we are reduced to showing that $C_f \otimes_R \mathbb{F}_R \sim \mathbb{F}_R^a$. Since C_f is a quotient of R^a it can be generated by at most a elements. Moreover, as can be seen by tensoring with $\mathbb{Q}_p$, there can be no fewer than a elements in a generating set for C_f. By Nakayama's lemma again, we see that $C_f \otimes_R \mathbb{F}_R \sim \mathbb{F}_R^a$ as desired.

All that remains is to show that the μ_R is really a probability measure (i.e. there is no escape of mass). Note that the above argument shows that $\mu_{N,R}(M) \leqslant 1/|\mathrm{Aut}_R(M)|$. Since μ_R has L^1-norm at most 1, for any $\epsilon > 0$ we can pick a co-finite set $S \subset T_R$ so that $\sum_{M \in S} (1/|\mathrm{Aut}_R(M)|) < \epsilon/2$ and large enough N so that $\sum_{M \notin S} |\mu_{N,R}(M) - \mu_R(M)| < \epsilon/2$, from which it follows that μ_R has L^1-norm at least $1 - \epsilon$. The result follows. $\square$

We can also compute the moments of the measure above. As expected by analogy to classical Cohen–Lenstra heuristics, they are all equal to 1.

Proposition 2.3. *For any finite module M_0,*

$$\sum_{M \in S_R} \#\mathrm{Surj}(M, M_0)\mu_R(M) = 1.$$

It is worth remarking that we do not insist in the above proposition that $M_0 \in T_R$.

Proof. Fix an $N > 0$. Then letting μ_{haar} be the Haar measure on $\mathrm{End}_R(R^N)$ giving total measure 1, we see that

$$\begin{aligned}\sum_{M \in S_R} \#\mathrm{Surj}(M, M_0)\mu_{R,N}(M) &= \int_{\phi \in \mathrm{End}_R(R^N)} \#\mathrm{Surj}(\mathrm{Coker}\,\phi, M_0) d\mu_{\mathrm{haar}} \\ &= \sum_{\psi \in \mathrm{Surj}(R^N, M_0)} \mu_{\mathrm{haar}}(\phi \mid \mathrm{Coker}\,\phi \in \mathrm{Ker}\,\psi) \\ &= \#\mathrm{Surj}(R^N, M_0) \cdot |M_0|^{-N}.\end{aligned}$$

Now, as $N \to \infty$, $\#\mathrm{Sur}(R^N, M_0) \sim |M_0|^N$. Thus,

$$\lim_{N \to \infty} \sum_{M \in S_R} \#\mathrm{Surj}(M, M_0)\mu_{R,N}(M) = 1.$$

By the proof of the theorem above, $\mu_{R,N} \leqslant c_R^{-1}\mu_R$, so the sum converges absolutely, and the result follows. $\square$

We expect that the moments actually determine our measure μ_R. We expect this unique determination property to hold in all cases, though we cannot show it in the case that $\mathbb{F}_R = \mathbb{F}_2$. Our proof is closely related to [EVW16, Lemma 7.2].

Lemma 2.4. *Assume that $\mathbb{F}_R \neq \mathbb{F}_2$. If μ is any measure on S_R such that the expected number of surjections from a μ-random module to M_0 is 1 for any finite M_0, then $\mu = \mu_R$. The same conclusion holds for μ being any function in $L^1(S_R)$.*

Proof. Consider the operator U on the infinite-dimensional Banach space $L^\infty(S_R)$ given by $U_{M,M'} = \#\mathrm{Surj}(M, M')/\#\mathrm{Aut}(M)$. Now, the rows of U have sums c_R^{-1} and so U is indeed an operator on $L^\infty(S_R)$. Moreover, the elements of $U-1$ are positive and have row sums $c_R^{-1}-1$. Now, to estimate c_R, let $q=|\mathbb{F}_R|^{-1}$ and note that by the Euler identity we have

$$c_R=\sum_{n\in\mathbb{Z}}(-1)^n q^{(3n^2-n)/2} > 1-\frac{q}{1-q}=\frac{1-2q}{1-q}.$$

Since $q\leqslant \frac{1}{3}$ by assumption, we conclude that $c_R>\frac{1}{2}$. Thus, the norm of $U-1$ is less than 1 and so U is invertible with inverse $\sum_{j=0}^{\infty}(1-U)^j$.

Now, consider the L^1-function μ. Since the moments to any $M_0\in S_R$ is 1, we must have that $\mu(M_0)\leqslant \#\mathrm{Aut}(M)^{-1}$. Thus the vector V with $V_M=\mu(M)\mathrm{Aut}(M)$ is in $L^\infty(S_R)$. Now the condition on the moments of μ amounts to saying that $UV=1_{S_R}$. Thus, we must have that $V=U^{-1}1_{S_R}$, and so there is a unique such function μ, which must then be μ_R. □

2.3 Remarks on identities

We give an example of an R with torsion where T_R is larger than the set of modules with projective dimension 1. Consider $R=\mathbb{Z}_p[x]/(px,x^2)$. Take $M=R/pR$. Clearly M occurs in T_R. On the other hand, if M had projective dimension 1 then pR would be forced to be projective, and thus free since R is local. However, pR is annihilated by x, and thus cannot be free. In fact, M fits into the exact sequence $0\to\mathbb{F}_p\to R\to R\to M$. Since R is not a regular local ring, $\mathbb{F}_p$ has infinite projective dimension, and thus M does as well.

For such rings R, if we instead considered the measure arising from the cokernel of a map $R^{N+d}\to R^N$ we could conceivably get more and more modules M in the support, giving a range of identities. They would be more and more complicated, however. For an R-module M, let $d_M=\dim_{\mathbb{F}_R}\mathrm{Tor}^1_R(M,\mathbb{F}_R)-\dim_{\mathbb{F}_R}M\otimes_R\mathbb{F}_R$. Then we get (by a minor modification of the proof above) the following identities:

$$\sum_{d_M\leqslant d}\frac{\prod_{j=1}^{d-d_M}(1-|\mathbb{F}_R|^{-i})^{-1}}{|M|^d|\#\mathrm{Aut}_R(M)|}=c_R^{-1}.$$

In fact, we can derive a series of *finite* identities from the above. Consider again $R=\mathbb{Z}_p[x]/(px,x^2)$. Then R maps to $\mathbb{Z}_p$, and it is easy to see from the construction that μ_R pushes forward to $\mu_{\mathbb{Z}_p}$. Moreover, M maps to M/xM, and it is easy to show by row and column operations and the fact that x is nilpotent that M is bounded in terms of d_M and M/xM. Thus, we conclude that for each p-group A, we have

$$\sum_{d_M=0,M/xM\sim A}\frac{1}{\#\mathrm{Aut}_R(M)}=\frac{1}{\#\mathrm{Aut}_{\mathbb{Z}_p}(A)}.$$

Of course, one can generalize this to arbitrary local maps $R\to S$ (though perhaps one has to be a bit careful if one wants the sum to remain finite). It is not clear to us, even for the above identity, how to prove it by elementary means.

3. A random model for étale group schemes with a symplectic form

3.1 Étale group schemes and symplectic forms

Let q be a prime power and ℓ a prime not dividing q. Let $P(x)\in\mathbb{Z}_\ell[x]$ be a monic polynomial satisfying $\ell\nmid P(0)$. Consider the collection $\mathcal{E}_P$ of (isomorphism classes of) triples (G,F_G,ω_G) where G is a finite abelian ℓ-group, $\omega_G\in\wedge^2 G$, and F_G is an endomorphism of G for which

https://doi.org/10.1112/S0010437X19007036 Published online by Cambridge University Press

$P(F_G) = 0, F_G(\omega_G) = q\omega_G$. Note that since $P(F_G) = 0$ and $P(0) \in \mathbb{Z}_\ell^\times$ it follows that F_G is an automorphism.

The pair (G, F_G) functorially corresponds to a finite étale group scheme $\mathcal{G}$ over $\mathbb{F}_q$ whose $\overline{\mathbb{F}_q}$ points are isomorphic to G with the Frobenius action corresponding to F_G. We shall construct in § 4 a natural group scheme $\wedge^2\mathcal{G}$ and its Tate twist $\wedge^2\mathcal{G}(1) = \wedge^2\mathcal{G} \otimes \mu_{\ell^\infty}^{-1}$, whose $\overline{\mathbb{F}_q}$ points naturally correspond to $\wedge^2 G$ (once one picks a section of $\mu_{\ell^\infty}(\overline{\mathbb{F}_q})$), and the Frobenius action is given by $q^{-1}F_G$. Thus, the set $\mathcal{E}_P$ naturally corresponds to pairs $(\mathcal{G}, \omega \in (\wedge^2\mathcal{G} \otimes \mu_{\ell^\infty}^{-1})(\mathbb{F}_q))$ where $\mathcal{G}$ is a finite étale ℓ-group scheme over $\mathbb{F}_q$. This is our motivation for studying $\mathcal{E}$, as we are interested in constructing probability distributions on finite étale group schemes, together with a section of $\wedge^2 \otimes \mu_{\ell^\infty}^{-1}$.

3.2 Defining a measure

Let ω be a unimodular symplectic form on $\mathbb{Z}_p^{2g}$. We define $\mathrm{GSp}_{2g}^{(q)}$ to be the coset of Sp_{2g} in GSp_{2g} whose elements scale ω by q. For each positive integer g, there is a map

$$\begin{aligned} \mathrm{GSp}^{(q)}(\mathbb{Z}_\ell^{2g}, \omega) &\to \mathcal{E} \\ F &\mapsto (\mathrm{Coker}(P(F)), F \bmod P(F), \omega \bmod P(F)) \end{aligned}$$

which gives rise to a probability measure μ_g on $\mathcal{E}$ by pushing forward the Haar probability measure on $\mathrm{GSp}_{2g}(\mathbb{Z}_\ell)$.

Theorem 3.1. *Fix $(H, F_H, \omega_H) \in \mathcal{E}$. The μ_g-expected number of equivariant surjections $T : (G, F_G) \twoheadrightarrow (H, F_H)$ for which $T(\omega_G) = \omega_H$ is equal to 0 for $g \leqslant g(H)$, and is equal to 1 for all $g > g(H, \omega_H)$ where $g(H, \omega_H)$ depends only on H.*

Proof. By definition of μ_g, the expected number of such surjections equals the $\mathrm{GSp}_{2g}^{(q)}(\mathbb{Z}_\ell)$-Haar expected number of surjections $T : \mathrm{Coker}(P(F)) \twoheadrightarrow H$, for which F induces F_H, and for which $T(\omega) = \omega_H$. Such surjections are equivalent to the following data:

- a surjection $T : \mathbb{Z}_\ell^{2g} \twoheadrightarrow H$ for which:
  - $T \circ P(F) = 0$;
  - $TF = F_H T$; and
  - $T\omega = \omega_H$.

The condition $T \circ P(F) = 0$ is actually redundant; since $P(F_H) = 0$, the second condition above implies that

$$T \circ P(F) = P(F_H) \circ T = 0.$$

Suppose H is killed by multiplication by ℓ^n. Then the expected number of surjections equals

$$\frac{\#\{(T, F) : T : (\mathbb{Z}/\ell^n)^{2g} \twoheadrightarrow H, F \in \mathrm{GSp}^{(q)}((\mathbb{Z}/\ell^n)^{2g}, \omega), T \circ F = F_H \circ T, T(\omega) = \omega_H\}}{\#\mathrm{GSp}_{2g}^{(q)}(\mathbb{Z}/\ell^n)}. \quad (1)$$

By an analogue of Witt's extension theorem [Mic06, Theorem 2.14], there is an integer $g(H)$, depending only on H, satisfying the following: the image of the mapping $T \mapsto T(\omega)$ from surjections to symplectic forms on H is surjective provided $g > g(H)$. Furthermore, for every $g > g(H)$, every fiber forms a single orbit $\mathcal{O}$ under $\mathrm{Sp}(\mathbb{Z}_\ell^{2g}, \omega)$ (where the symplectic group acts by precomposition). Furthermore, suppose that $TF = F_H T$ and $T(\omega) = \omega_H$. Then for every $g \in \mathrm{Sp}(\mathbb{Z}_\ell^{2g}, \omega)$

$$(Tg)(g^{-1}Fg) = TFg = F_H(Tg)$$

https://doi.org/10.1112/S0010437X19007036 Published online by Cambridge University Press

and

$$Tg(\omega) = T\omega = \omega_H.$$

Therefore, among the pairs (F, T) enumerated in the numerator of (1), the fibers over every T have the same size. Now assuming it exists, fix T_0 satisfying $T_0(\omega) = \omega_H$ (if no such T_0 exists, then the moment is clearly 0). Then

$$\begin{aligned}
&\frac{\#\{(T,F) : T : (\mathbb{Z}/\ell^n)^{2g} \twoheadrightarrow H, F \in \mathrm{GSp}^{(q)}((\mathbb{Z}/\ell^n)^{2g}, \omega), TF = F_H T, T(\omega) = \omega_H\}}{\#\mathrm{GSp}_{2g}^{(q)}(\mathbb{Z}/\ell^n)} \\
&= \frac{\#\{F \in \mathrm{GSp}^{(q)}((\mathbb{Z}/\ell^n)^{2g}, \omega) : T_0 F = F_H T_0\}}{\#\mathrm{GSp}_{2g}^{(q)}(\mathbb{Z}/\ell^n)} \cdot \#\mathcal{O} \\
&= \frac{\#\{F \in \mathrm{GSp}^{(q)}((\mathbb{Z}/\ell^n)^{2g}, \omega) : T_0 F = F_H T_0\}}{\#\mathrm{GSp}_{2g}^{(q)}(\mathbb{Z}/\ell^n)} \cdot \frac{\#\mathrm{Sp}((\mathbb{Z}/\ell^n)^{2g}, \omega)}{\#\mathrm{Stab}_{\mathrm{Sp}((\mathbb{Z}/\ell^n)^{2g}, \omega)}(T_0)} \\
&= \frac{\#\{F \in \mathrm{GSp}^{(q)}((\mathbb{Z}/\ell^n)^{2g}, \omega) : T_0 F = F_H T_0\}}{\#\{g \in \mathrm{Sp}((\mathbb{Z}/\ell^n)^{2g}, \omega) : T_0 g = T_0\}}. \qquad (2)
\end{aligned}$$

The set in the numerator of (2) is either empty or is a torsor for the group in the denominator. Thus we only need to show that a single such F exists.

Now, to show this, consider first any element $F_0 \in \mathrm{Sp}((\mathbb{Z}/\ell^n)^{2g}, \omega)$. Then $F_H^{-1} T_0 F_0$ is a surjection from $((\mathbb{Z}/\ell^n)^{2g}, \omega)$ to (H, ω_H). Thus, by [Mic06, Theorem 2.14], there exists an element $g \in \mathrm{Sp}((\mathbb{Z}/\ell^n)^{2g}, \omega)$ satisfying

$$F_H^{-1} T_0 F_0 g = T_0$$

and therefore $T_0 F_0 g = F_H T_0$. Thus we may take $F = F_0 g$, and this completes the proof. □

3.3 The existence of a limit measure

In light of the results of the previous section, and analogous results for the Cohen–Lenstra measure [EVW16], [Woo17, Theorem 8.2] we make the following conjecture, which roughly says that the moments constitute enough information to recover the full measure in cases of interest.

Conjecture 3.1. The measures μ_g converge to a measure μ on $\mathcal{E}$, such that the expected number of surjections from a μ-random element to any element in $\mathcal{E}$ is 1. Moreover, this property characterizes μ.

We devote the rest of this section to proving Conjecture 3.1 in a couple special cases. Most notably, we can use the results of § 2 to prove the conjecture in the case where the symplectic structure 'doesn't come up'. In that case we can use the much easier additive model in § 2 as opposed to the model with symplectic matrices. We 'get rid of' the symplectic structure as follows: if $P(q)$ is not divisible by ℓ, then since ω_G is killed by both $P(q)$ and a power of ℓ it is forced to be 0, so $\mathcal{E}$ is equivalent to the category of finite $\mathbb{Z}_\ell[x]/P(x)$ modules.

Theorem 3.2. *In the notation above, assume that $P(q)$ is not divisible by ℓ, and assume that ℓ is odd. Then Conjecture 3.1 holds. Moreover, μ is supported on precisely the $R = \mathbb{Z}_\ell[x]/P(x)$ modules of projective dimension 1, and assigns such a module M measure $\mu(M) = c/\#\mathrm{Aut}(M)$ where $c = \prod_{k_j} \prod_{i=1}^{\infty} (1 - |k_j|^{-i})$ and the product is over the finite residue fields of R.*

Proof. Assume first that R is a local ring. Then note that we have already constructed one measure satisfying the above hypothesis on moments in Theorem 2.2, and Lemma 2.4 guarantees that these moments specify a unique measure. Hence it is sufficient to show that the μ_g converge to a measure with 1 expected surjection to each finite R module. Note that by Theorem 3.1 the μ_g satisfy $\mu_g(M) \leqslant 1/\#\mathrm{Aut}\, M$ for any finite R-module M. Letting S_R be the set of finite R-modules as in § 2, there is an operator U on $L^\infty(S_R)$ given by $U_{M,M'} = \#\mathrm{Surj}(M, M')/\#\mathrm{Aut}(M)$. Now the vector $V_g \in L^\infty(S_R)$ given by $V_M = \mu_g(M)\#\mathrm{Aut}(M)$ has L^∞ norm bounded by 1. Further, by Theorem 3.1 the product UV_g is a vector W_g consisting of 0 and 1 entries, whose entries each eventually become 1 as g increases. Thus, we can write $V_g = U^{-1}(W_g)$. Since as in Lemma 2.4 the operator U^{-1} is bounded, it can be represented as an infinite matrix with rows in L^1 with uniformly bounded L^1 norm. Since the W_g have L^∞ norm bounded by 1 and each entry eventually stabilizes, we can conclude that the μ_g converge to μ in the weak-* topology, as desired.

Now, even if R is not local, it is ℓ-adically complete and R/ℓ is artinian, so R is a product of local rings $R = \prod_j R_j$. It follows that we can take $\mu = \prod_j \mu_{R_j}$ and this measure will have all the correct moments, and the exact same proof as in the previous paragraph shows that the μ_g converge to μ. Thus it only remains to show that μ is determined by its moments. Note here that the exact same proof as in the local case won't work, since the constant c (from the theorem statement) could be less than $\frac{1}{2}$. Inducting on the number of local rings that R is a product of, we may write $R = R_1 \times R_2$ where R_1 and R_2 both have the property that the corresponding measures μ_{R_i} are determined by their moments. Now, suppose that m is any other measure with the correct moments. For an R_1-module M_1 and an R_2-module M_2^0 we let

$$a(M_1, M_2^0) := \sum_{M_2} m(M_1 \times M_2)\#\mathrm{Surj}(M_2, M_2^0).$$

Then it follows that for each M_2^0, M_1^0,

$$\sum_{M_1} a(M_1, M_2^0)\#\mathrm{Surj}(M_1, M_1^0) = 1.$$

Thus, by our induction assumption for R_2 it follows that $a(M_1, M_2) = \mu_{R_1}(M_1)$. Now, by our induction assumption for R_1 we learn that

$$m(M_1 \times M_2) = \mu_{R_1}(M_1) \times \mu_{R_2}(M_2) = \mu(M_1 \times M_2)$$

as desired. □

In the case where the symplectic structure is present, we do not even have a good conjecture as to what the limiting measure in Conjecture 3.1 should be. It is natural to guess that it is proportional to the inverse of the size of the automorphism group, where now one only takes automorphisms if they preserve ω_G, but this does not agree with computations of Garton [Gar15]! We think it would be very interesting to at least develop a plausible heuristic.

3.4 Moments approximately 1 implies approximately Cohen–Lenstra measure

Fix a finite subset $S' \subset \mathcal{E}$. Let Conf_g denote the moduli space of g distinct, unordered unlabelled points in $\mathbb{A}^1$. Let $\mathcal{C} \to \mathrm{Conf}_g(\mathbb{A}_1)$ denote the associated family of hyperelliptic curves. For $x \in \mathrm{Conf}_g(\mathbb{A}^1)(\mathbb{F}_q)$, let F_x denote Frobenius acting on the ℓ-adic Tate-module of $\mathrm{Jac}(C_x)$. For every g, let ν_g be the discrete probability measure

$$\nu_g = \frac{1}{\#\mathrm{Conf}_g(\mathbb{F}_q)} \sum_{x \in \mathrm{Conf}_g(\mathbb{F}_q)} \delta_{\mathrm{Coker}\, P(F_x)}.$$

Building on the work of [EVW16], we will show in § 4 that for any $\delta > 0$, there is some $Q(S', \delta), G(S', \delta) \gg 0$ such that for all $q \geqslant Q(S', \delta)$, provided $g \geqslant G(S', \delta)$, then

$$\text{Expectation}_{\nu_g}(\#\text{Surj}(\bullet, A)) \in [1 - \delta, 1 + \delta] \quad \text{for all } A \in S'. \tag{3}$$

Geometry gives us access to moments, and we would like to recover as much information about the measures ν_g as we can from a large set of approximate moments as in (3).

DEFINITION 1. Let $R = \mathbb{Z}_\ell[x]/P(x)$. Let A be a finite R-module. An *enlargement* A' of A is an R-module admitting a surjection onto A whose kernel is a simple R-module. An *s-enlargement* B of A is a finite R-module admitting a surjection onto A whose kernel has R-length equal to s.

Say that R has the *few enlargements property* if for every finite R-module A, the number of isomorphism classes of s-enlargements of A is subexponential in s.

LEMMA 3.3. *If R is a product of maximal orders, then R satisfies the few enlargements property.*

Proof. This follows exactly as in the argument from [EVW16, Lemma 8.4]. □

PROPOSITION 3.4. *Suppose that ℓ does not divide $P(q)$, and $R = \mathbb{Z}_\ell[x]/P(x)$ has the few enlargements property. Let ν be a probability measure on $\mathcal{E}$. Fix a finite subset $S \subset \mathcal{E}$. Fix $\epsilon > 0$. There exist $\delta > 0$ and a finite subset $S' \subset \mathcal{E}$ satisfying*

$$\begin{gathered}\text{Expectation}_{\nu}(\#\text{Surj}(\bullet, A')) \in [1 - \delta, 1 + \delta] \quad \textit{for all } A' \in S' \\ \implies |\nu(A) - \mu_R(A)| < \epsilon \quad \textit{for all } A \in S.\end{gathered}$$

Proof. The hypothesis $\ell \nmid P(q)$ ensures that the symplectic form equals 0. The argument from [EVW16, Proposition 8.3] carries over verbatim to the present context. □

4. Moments of étale group schemes via the Lefschetz trace formula

In this section we define moduli spaces over $\mathbb{F}_q$, whose $\mathbb{F}_q$-points correspond to surjections from torsion sub-group schemes of Jacobians of hyperelliptic curves to étale group schemes $\mathcal{G}$ together with a section of $\wedge^2\mathcal{G}(1)$, and prove Theorem 1.1. In particular, we identify the rationally defined geometric components of the moduli spaces considered in [EVW16] with the set $\wedge^2\mathcal{G}(1)(\mathbb{F}_q)$.

4.1 Multilinear algebra for étale group schemes

Let S be a scheme. Let G/S be a finite étale group scheme.

PROPOSITION 4.1. *Let G/S be a finite, commutative étale group scheme. There exists a finite étale group scheme $\wedge^2 G/S$ and a morphism $\iota : G \times_S G \to \wedge^2 G$ satisfying the following universal property.*

(a) *The morphism ι is biadditive, i.e. for all S-schemes T and all $x, y, z \in G(T)$,*

$$\iota(x + y, z) = \iota(x, z) + \iota(y, z) \quad \text{and} \quad \iota(z, x + y) = \iota(z, x) + \iota(z, y).$$

(b) *The morphism ι is alternating, i.e. for all S-schemes T and all $v \in G(T)$,*

$$\iota(v, v) = 0 \in (\wedge^2 G)(T).$$

(c) *The morphism ι is universal with respect to the properties (a), (b) if $f : G \times_S G \to H$ is a biadditive, alternating morphism to commutative group scheme H/S, there is a unique S-group scheme morphism π for which $f = \pi \circ \iota$.*

https://doi.org/10.1112/S0010437X19007036 Published online by Cambridge University Press

Proof. Let $\underline{A}_S$ denote the constant group scheme on the finite abelian group A. Let H/S be another group scheme. A morphism $\underline{A}_S \to H$ of group schemes is determined by a collection of sections $s_a \in H(S)$ indexed by $a \in A$ satisfying $s_{a+b} = s_a + s_b$ for all $a, b \in A$. Let $\underline{a} \in \underline{A}(S)$ denote the constant section determined by $a \in A$. The morphism

$$\begin{aligned} \underline{A}_S \times \underline{A}_S &\to \underline{\wedge^2 A}_S \\ (\underline{a}, \underline{b}) &\mapsto \underline{a \wedge b} \end{aligned}$$

is biadditive, alternating, and satisfies the desired universal property by the universal property of $\wedge^2$ for abelian groups.

For more general finite étale group schemes G/S, the desired $\wedge^2 G$ may be constructed by descent. Let $\{U_\bullet \to S\}$ be an étale cover trivializing the finite étale group scheme G. The above already constructs $\iota_{U_1} : G_{U_1} \times_{U_1} G_{U_1} \to \wedge^2 G_{U_1}$ and $\iota_{U_2} : G_{U_2} \times_{U_2} G_{U_2} \to \wedge^2 G_{U_2}$. Then $(\wedge^2 G_{U_1})_{U_1 \times_S U_2}$ and $(\wedge^2 G_{U_2})_{U_1 \times_S U_2}$ both satisfy the universal property defining $\wedge^2 G_{U_1 \times_S U_2}$. Thus, there is a unique isomorphism $\iota_{U_1, U_2} : (\wedge^2 G_{U_1})_{U_1 \times_S U_2} \xrightarrow{\sim} (\wedge^2 G_{U_2})_{U_1 \times_S U_2}$ commuting with the structure morphisms $(\iota_{U_1})_{U_2}$ and $(\iota_{U_2})_{U_1}$. A second application of the universal property shows that these isomorphisms satisfy the cocycle condition on triple overlaps. By étale descent, $\{\iota_{U_\bullet} : G_{U_\bullet} \times_{U_\bullet} G_{U_\bullet} \to \wedge^2 G_{U_\bullet}\}$ descends to a biadditive, alternating morphism $\iota : G \times_S G \to \wedge^2 G$.

Let $f : G \times_S G \to H$ be a biadditive, alternating map. By the universal property, every $f_{U_\bullet} : G_{U_\bullet} \times_{U_\bullet} G_{U_\bullet} \to H_{U_\bullet}$ factors uniquely through $\wedge^2 G_{U_\bullet} \xrightarrow{\pi_\bullet} H_{U_\bullet}$. By the universal property of $\wedge^2$, the morphisms $\pi_\bullet$ must agree on double overlaps: $\iota_{U_1, U_2} \circ (\pi_1)_{U_1 \times_S U_2} = (\pi_2)_{U_1 \times_S U_2}$. By étale descent for morphisms, $\pi_\bullet$ descends uniquely to a morphism $\pi : \wedge^2 G \to H$ satisfying $\pi \circ \iota = f$. $\square$

A completely analogous argument allows one to make any tensorial construction for finite étale group schemes. The key point is that universal properties from linear algebra induce descent data that allow one to étale-localize the construction to the case of constant group schemes, for which the construction is simple. We single out the following special case for later use.

Proposition 4.2. *Let G_1/S and G_2/S be finite commutative étale group schemes. There exists a finite commutative étale group scheme* $\mathrm{Hom}(G_1, G_2)/S$ *equipped with a morphism $e : G_1 \times_S$* $\mathrm{Hom}(G_1, G_2) \to G_2$ *satisfying the following universal property.*

(a) *The morphism e is biadditive, i.e. for all S-schemes T and all $x, y \in G_1(T)$ and $\alpha, \beta \in$* $\mathrm{Hom}(G_1, G_2)$,

$$e(x + y, \alpha) = e(x, \phi) + e(y, \alpha) \quad \textit{and} \quad e(x, \alpha + \beta) = e(x, \alpha) + e(x, \beta).$$

(b) *The morphism e is universal with respect to the property (a): if H/S is a commutative group scheme and $f : G_1 \times_S H \to G_2$ is biadditive, there is a unique S-group scheme morphism $\pi :$* $\mathrm{Hom}(G_1, G_2) \to H$ *for which $e = f \circ (1, \pi)$.*

Furthermore, $\mathrm{Hom}(G_1, G_2)$ *represents the functor on S-schemes*

$$T \mapsto \mathrm{Hom}_{T\text{-group schemes}}((G_1)_T, (G_2)_T).$$

4.2 Generalities on moduli spaces

Let $\mathcal{A} \to V/S$ be a family of principally polarized, g-dimensional abelian varieties over S. Suppose that ℓ is invertible on S. Let G/S be a finite étale group scheme annihilated by ℓ^n. We claim the moduli problem

$$V_G(T) = \{A \in V(T), T\text{-group morphism } \phi : \mathcal{A}_A[\ell^n] \twoheadrightarrow G_T\} \quad \text{for all } T/S$$

https://doi.org/10.1112/S0010437X19007036 Published online by Cambridge University Press

is representable. To see this, consider the finite étale group S-scheme $\mathrm{Hom}(\mathcal{A}[\ell^n], G_V)$, and consider the subscheme Y of $\mathcal{A}[\ell^n] \times \mathrm{Hom}(\mathcal{A}[\ell^n], G_V)$ mapping to the origin in G_V. The group scheme Y is finite étale over $\mathrm{Hom}(\mathcal{A}[\ell^n], G_V)$, so V_G is just the subscheme over which Y is of degree $\ell^{2ng}/|G|$, which is a union of connected components of $\mathrm{Hom}(\mathcal{A}[\ell^n], G_V)$.

The morphism $V_G \to V$ is thus finite étale.

4.2.1 *The Weil pairing morphism to* $\wedge^2 G \otimes \mu_{\ell^n}^{-1}$. Let T be an S-scheme. Let A/T be a principally polarized abelian T-scheme. Let ℓ be invertible on S. We may naturally regard the Weil pairing $w_{\ell^n}(A) : A[\ell^n] \times_T A[\ell^n] \to \mu_{\ell^n}/T$ as an element of $\mathrm{Hom}(\wedge^2 A[\ell^n], \mu_{\ell^n})(T)$.

Finite étale group schemes locally isomorphic to $\underline{\mathbb{Z}/\ell^n}_S$ form an abelian group under tensor product with identity $\underline{\mathbb{Z}/\ell^n}_S$ and inverse $H^{-1} := \mathrm{Hom}(H, \underline{\mathbb{Z}/\ell^n}_S)$. We let $H^m := H^{\otimes m}$ and $H^{-n} := (H^{-1})^{\otimes n}$.

Consider the multilinear map $A[\ell^n]^4 \to \mu_{\ell^n} \otimes \mu_{\ell^n}$ given on sections by

$$(a, b, c, d) \to w_{\ell^n}(a, c) \otimes w_{\ell^n}(b, d) \cdot [w_{\ell^n}(b, c) \otimes w_{\ell^n}(a, d)]^{-1}.$$

By the universal property for $\wedge^2$, this induces a pairing $\wedge^2 A[\ell^n] \times \wedge^2 A[\ell^n] \to \mu_{\ell^n}^2$. One can check on the level of points that this pairing is perfect. Thus we can naturally identify $\wedge^2 A[\ell^n]$ with the Cartier dual of $\wedge^2 A[\ell^n] \otimes \mu_{\ell^n}^{-1}$. It follows that we may naturally regard the Weil pairing w_{ℓ^n} as an element of $(\wedge^2 A[\ell^n] \otimes \mu_{\ell^n}^{-1})(T)$. From now on we write $H(m)$ for $H \otimes \mu_{\ell^n}^{-m}$.

Lemma 4.3. *Let $\mathcal{A} \to V/S$ be a family of g-dimensional principally polarized abelian varieties over S. The morphism*

$$\begin{aligned} V_G &\xrightarrow{\pi} (\wedge^2 G)(1) \\ (A, \phi) &\mapsto \phi(w_{\ell^n}(A)) \end{aligned}$$

is functorial and hence algebraic. If $g \geqslant c(G)$, where the constant $c(G)$ depends only on G, the morphism π is surjective on geometric points.

Proof. Let $y \in (\wedge^2 G)(1)$ be a geometric point. Let $A \in V(S)$ be an arbitrary abelian scheme. Over the algebraically closed residue field $k(y)$, the group schemes $A[\ell^n]_{k(y)}, \mu_{\ell^n}$ and $G_{k(y)}$ become constant, isomorphic to $\underline{(\mathbb{Z}/\ell^n)^{2g}}, \underline{\mathbb{Z}/\ell^n}$ and $B = G(k(y))$ respectively.

The Weil pairing $\omega \in (\wedge^2 A[\ell^n](1))(k(y))$ is non-degenerate. Surjections $A[\ell^n]_{k(y)} \to G_{k(y)}$ are equivalent to surjections of finite abelian groups $(\mathbb{Z}/\ell^n)^{2g} = A[\ell^n](k(y)) \twoheadrightarrow G(k(y)) = B$.

Let $\omega_B \in \wedge^2 B$ correspond[5] to the geometric point $y \in \wedge^2 G(1)$. By [Mic06, Proposition 2.14], there is some constant $c(G)$ such that if $g \geqslant c(G)$, there exists some surjection $\phi : (\mathbb{Z}/\ell^n)^{2g} \twoheadrightarrow A$ for which $\phi(\omega) = \omega_A$. The result follows. □

4.3 Geometric monodromy and connected components

Proposition 4.4. *Let k be a field. Let $\mathcal{A} \to V/k$ be a family of g-dimensional principally polarized abelian varieties with universal Weil pairing ω. Let G/k be a finite étale commutative group scheme. Suppose that for every geometric point $\overline{z} \in V$, the action of the geometric monodromy group $\pi_1(V, \overline{z}) = \mathrm{Gal}(k(\overline{\eta})/k(\eta))$ on $\mathcal{A}[\ell^n](k(\overline{z})) \cong (\mathbb{Z}/\ell^n)^{2g}$ has image equal to the full symplectic group $\mathrm{Sp}(\mathcal{A}[\ell^n](k(\overline{z})), \omega_{k(\overline{z})})$. There is a constant $c(G)$ such that if $g \geqslant c(G)$, the following hold:*

- *$\pi : V_G \to (\wedge^2 G)(1)$ is surjective on geometric points;*
- *for every geometric point $y \in (\wedge^2 G)(1)$, the fiber $\pi^{-1}(y)$ is connected.*

[5] This is a well-defined correspondence upon fixing a generator of μ_{ℓ^n}.

Proof. Let $\omega_0 \in \wedge^2(\mathbb{Z}/\ell^n)^{2g}$ be non-degenerate. Let A be a finite abelian ℓ-group. Consider the map

$$\Phi : \mathrm{Surjections}((\mathbb{Z}/\ell^n)^{2g}, A) \to \wedge^2 A$$
$$T \mapsto T(\omega_0).$$

By [Mic06, Proposition 2.14], there is a constant $c(A)$ such that if $g \geqslant c(A)$, Φ is surjective and forms a single orbit under the symplectic group $\mathrm{Sp}((\mathbb{Z}/\ell^n)^{2g}, \omega_0)$.

Set $c(G) = c(B)$ where $B = G(\overline{k})$ for any algebraic closure $\overline{k}/k$. By Lemma 4.3, the map π is surjective on geometric points. Let $\overline{y} \in \pi^{-1}(y)$ be a geometric point. Let $\overline{z} = \pi_G(\overline{y}) \in V$, where $\pi_G : V_G \to V$ is the forgetful map.

Note that the fiber $\pi_G^{-1}(\overline{z})$ equals

$$\mathrm{Surjections}(\mathcal{A}[\ell^n](k(\overline{z})), B) \cong \mathrm{Surjections}((\mathbb{Z}/\ell^n)^{2g}, B);$$

the symbol $\cong$ means that there is an isomorphism $\mathcal{A}[\ell^n](k(\overline{z})) \to (\mathbb{Z}/\ell^n)^{2g}$ which is equivariant for the action of $\mathrm{Sp}(\mathcal{A}[\ell^n](k(\overline{z})), \omega_{k(\overline{z})})$ on the left and of $\mathrm{Sp}((\mathbb{Z}/\ell^n)^{2g}, \omega_0)$ on the right.

The points of $\pi_G^{-1}(\overline{z})$ lying over y, corresponding to $\omega_B \in \wedge^2 B$, can be identified with

$$\{T \in \mathrm{Surjections}((\mathbb{Z}/\ell^n)^{2g}, B) : T(\omega_0) = \omega_B\}.$$

By the above remarks, our assumption that the image of $\pi_1(V, \overline{z})$ in $\mathrm{Sp}(\mathcal{A}[\ell^n](k(\overline{z})), \omega_{k(\overline{z})})$ is surjective implies that $\pi_1(V, \overline{z})$ acts transitively on $\pi^{-1}(\overline{z})$. It follows that V_G is geometrically connected. ◻

Corollary 4.5. *Let k be a finite field of characteristic p. Let $\ell' \neq \ell, p$ be a prime. Let G/k be a finite étale group scheme. Same notation and hypotheses as in Proposition 4.4. Let $c(G)$ be the constant from Proposition 4.4 and assume that $g \geqslant c(G)$. Let $\overline{k}/k$ be an algebraic closure. The map $\pi : V_G \to (\wedge^2 G)(1)$ induces a $\mathrm{Gal}(\overline{k}/k)$-equivariant isomorphism*

$$H^0_{\text{ét}}((\wedge^2 G(1))_{\overline{k}}, \mathbb{Q}_{\ell'}) \xrightarrow{\pi^*} H^0_{\text{ét}}((V_G)_{\overline{k}}, \mathbb{Q}_{\ell'}).$$

In particular,

$$\mathrm{tr}(\mathrm{Frob}_{\overline{k}/k} | H^0_{\text{ét}}((V_G)_{\overline{k}}, \mathbb{Q}_{\ell'})) = \#((\wedge^2 G)(1))(k).$$

Proof. Equivariance of π^* under $\mathrm{Gal}(\overline{k}/k)$ follows because π is defined over k. The map π^* induces an isomorphism because $\pi^{-1}(y)$ is connected for every geometric point $y \in \wedge^2 G$, by Proposition 4.4. ◻

4.4 Comparison between moduli spaces of abelian varieties with level structure and Ellenberg–Venkatesh–Westerland moduli spaces of covers

Let S be any base. Let Conf_n/S be the moduli space of n distinct unlabelled points in $\mathbb{A}^1/S$. Let $\mathcal{C} \to \mathrm{Conf}_n$ be the associated family of hyperelliptic curves over S. Let $\mathcal{A} \to \mathrm{Conf}_n/S$ denote the relative Jacobian of $\mathcal{C}/\mathrm{Conf}_n$. There is an associated Torelli map $\mathcal{J} : \mathcal{C}/\mathrm{Conf}_n \to \mathcal{A}/\mathrm{Conf}_n$.

Let B be a finite abelian group of odd order. Let $\underline{B}_{\mathrm{Conf}_n}$ be the associated constant group scheme. Let

$$\pi_{\underline{B}} : \mathrm{Conf}_{n,\underline{B}} \to \mathrm{Conf}_n$$

be the finite étale cover described in § 4.2. Let $\mathcal{D} := \mathcal{A}_{\mathrm{Conf}_{n,\underline{B}}}$ and let $\phi : \mathcal{D}[\ell^n]/\mathrm{Conf}_{n,\underline{B}} \twoheadrightarrow \underline{B}/\mathrm{Conf}_{n,\underline{B}}$ be the associated universal quotient. There is an associated finite étale cover over $\mathrm{Conf}_{n,\underline{B}}$

$$\mathcal{D}/\mathrm{Ker}\,\phi \to \mathcal{D}/\mathcal{D}[\ell^n] \cong \mathcal{D} = \mathcal{A}_{\mathrm{Conf}_{n,\underline{B}}},$$

https://doi.org/10.1112/S0010437X19007036 Published online by Cambridge University Press

where the isomorphism $\cong$ from the second map is the inverse of the projection isomorphism $\mathcal{D} \xrightarrow{\sim} \mathcal{D}/\mathcal{D}[\ell^n]$. Pulling back this finite étale cover by the Torelli map $\mathcal{J}_{\mathrm{Conf}_{n,\underline{B}}} : \mathcal{C}_{\mathrm{Conf}_{n,\underline{B}}} \to \mathcal{A}_{\mathrm{Conf}_{n,\underline{B}}}$ defines a finite étale cover $\mathcal{C}' \to \mathcal{C}_{\mathrm{Conf}_{n,\underline{B}}}$ of $\mathrm{Conf}_{n,\underline{B}}$ with abelian Galois group B. This finite étale cover defines a morphism

$$\Phi : \mathrm{Conf}_{n,\underline{B}} \to \mathrm{Hn}^{\mathrm{c}}_{B\rtimes\langle\pm1\rangle,n}/\mathrm{Conf}_n,$$

where $\mathrm{Hn}^{\mathrm{c}}_{B\rtimes\langle\pm1\rangle,n}$ denotes the moduli space of $B \rtimes \langle\pm1\rangle$-covers of $\mathbb{P}^1$ ramified at ∞ and having monodromy in the conjugacy class of involutions at n-finite punctures. This is the moduli space considered by [EVW16]. We refer the reader to [EVW16] and [RW06] for details on the algebraic construction of $\mathrm{Hn}^{\mathrm{c}}_{B\rtimes\langle\pm1\rangle,n}$. In particular, we emphasize that both $\mathrm{Conf}_{n,\underline{B}}/\mathrm{Conf}_n$ and $\mathrm{Hn}^{\mathrm{c}}_{B\rtimes\langle\pm1\rangle,n}/\mathrm{Conf}_n$ are finite étale.

PROPOSITION 4.6. *Let $S = \overline{\mathbb{F}}_q$ be an algebraic closure of the finite field $\mathbb{F}_q$. The morphism $\Phi : \mathrm{Conf}_{n,\underline{B}}/\mathrm{Conf}_n \to \mathrm{Hn}^{\mathrm{c}}_{B\rtimes\langle\pm1\rangle,n}/\mathrm{Conf}_n$ described above is an isomorphism.*

Proof. Because $\mathrm{Conf}_{n,\underline{B}}/\mathrm{Conf}_n$ and $\mathrm{Hn}^{\mathrm{c}}_{B\rtimes\langle\pm1\rangle,n}/\mathrm{Conf}_n$ are finite étale, the morphism Φ is necessarily finite étale. By [EVW16, Proposition 8.7], Φ induces a bijection $\mathrm{Conf}_{n,\underline{B}}(\overline{\mathbb{F}}_q) \xrightarrow{\Phi} \mathrm{Hn}^{\mathrm{c}}_{B\rtimes\langle\pm1\rangle,n}(\overline{\mathbb{F}}_q)$. It follows that ϕ must have degree 1 and is thus an isomorphism. $\square$

COROLLARY 4.7. *Assume now that ℓ is odd. Let G be a finite étale group scheme over $\mathbb{F}_q$ of order ℓ^n, and $\omega_G \in (\wedge^2 G)(1)(\mathbb{F}_q)$. For each g, define*

$$\mathrm{Avg}(G, \omega_G, g, q) := \frac{\#\{\phi \in \mathrm{Surj}(\mathrm{Pic}^0(C)[\ell^n], G), \phi_*(\omega_{C,\ell^n}) = \omega_G\}}{\#\mathrm{Conf}_g(\mathbb{F}_q)}$$

where ω_{C,ℓ^n} is the weil-pairing.

Let $\delta^{\pm}(q, \omega_G)$ be the lower and upper limits of $\mathrm{Avg}(G, \omega_G, g, q)$ as $g \to \infty$. Then as $q \to \infty$ and n stays fixed, $\delta^+(q, \omega_G)$ and $\delta^-(q, \omega_G)$ converge to 1.

Proof. First, note that $\mathrm{Avg}(G, \omega_G, g, q) \cdot |\mathrm{Conf}_g(\mathbb{F}_q)|$ is simply equal to the number of points on the subscheme Y of $\mathrm{Conf}_{g,G}$ which maps to ω_G under the natural map to $(\wedge^2 G)(1)$. By a result of Yu [Yu97, Hal08], the monodromy condition in Lemma 4.4 is satisfied, so Y is geometrically connected. Moreover, by the discussion above the ℓ'-adic cohomology of $\mathrm{Conf}_{g,G}$ is the same as that of $\mathrm{Hn}^{c}_{B,g}$ where $B = G(\overline{\mathbb{F}_q}) \rtimes \mathbb{Z}/2\mathbb{Z}$ and c is the conjugacy class of all involutions. Thus, by [EVW16, Lemma 7.8] we see that for all $i > 0$, there is an integer $C(\ell^n)$ satisfying $\dim H^i(\mathrm{Conf}_{g,G}, \mathbb{Q}_{\ell'}) \leqslant C(\ell^n)^{i+1}$. The same bound therefore holds on the cohomology of Y. Thus, by the Lefschetz trace formula we get that for $q > 2C(\ell^n)^2$, $\#Y(\mathbb{F}_q) = q^n(1 + O(C(G, \ell^n)/\sqrt{q}))$. The result follows since $|\mathrm{Conf}_g(\mathbb{F}_q)| = q^n - q^{n-1}$. $\square$

Remark 3. Note that if we stopped keeping track of ω_G and only cared about étale group scheme, the number of surjections to a group scheme G approaches $\#(\wedge^2 G)(1)(\mathbb{F}_q)$. If G is a constant group scheme $\underline{B}$, this amounts to counting elements of $\wedge^2 B$ which are killed by $q - 1$. This is consistent with the random model considered by Garton, and explains the failure of ordinary Cohen–Lenstra heuristics to hold if $q \not\equiv 1 \bmod \ell$.

https://doi.org/10.1112/S0010437X19007036 Published online by Cambridge University Press

5. Applications

5.1 Joint distribution

Fix positive integers $n_1, \dots, n_k$. What is the joint distribution of the finite abelian ℓ-groups $A(\mathbb{F}_{p^{n_1}})_\ell, \dots, A(\mathbb{F}_{p^{n_k}})_\ell$ as A varies through a family V_g of g-dimensional principally polarized abelian varieties?

5.1.1 *Étale group schemes refine joint moments.* Let $n = \mathrm{lcm}(n_1, \dots, n_k)$. Fix $G_1, G_2, \dots, G_k$ finite abelian ℓ-groups. Let $M_n(A)$ denote the $\mathbb{Z}_\ell[x]/(x^n-1)$-module $A(\mathbb{F}_{p^n})_\ell$ with its natural Frobenius action. Let $S_{G_1,\dots,G_\ell}$ denote the set of isomorphism classes of $\mathbb{Z}_\ell[x]/(x^n-1)$-modules M for which

$$M[x^{n_1}-1] \cong G_1, \dots, M[x^{n_k}-1] \cong G_k.$$

Then

$$\mathrm{Prob}_{A\in V_g(\mathbb{F}_p)}(A(\mathbb{F}_{p^{n_1}})_\ell \cong G_1, \dots, A(\mathbb{F}_{p^{n_k}})_\ell \cong G_k) = \sum_{M\in S_{G_1,\dots,G_k}} \mathrm{Prob}_{A\in V_g(\mathbb{F}_p)}(M_n(A) \cong M).$$

So the distribution of $A(\mathbb{F}_{p^n})_\ell$ as a $\mathbb{Z}_\ell[x]/(x^n-1)$-module is a strictly more refined statistic than the joint distribution of $A(\mathbb{F}_{p^{n_1}}), \dots, A(\mathbb{F}_{p^{n_k}})$.

5.2 Results for the universal family of hyperelliptic curves

In this subsection, we spell out the consequences of our main theorems for the universal family of hyperelliptic curves in one special case. For $x \in \mathrm{Conf}_g(\mathbb{F}_q)$, let C_x denote the associated hyperelliptic curve and let C_x^σ denote its quadratic twist.

For the ring $R = \mathbb{Z}_\ell[x]/(x^2-1)$ and finite R-module M, let $M^\pm$ denote the ±1-eigenspaces of multiplication by x.

Proposition 5.1. *Suppose $\ell \nmid q^2-1$ and that $\ell \neq 2$. Let $\epsilon > 0$. Fix a finite set S of finite abelian ℓ-groups. Let $M_{A,B}$ denote the unique R-module whose $+1$-eigenspace equals A and whose -1-eigenspace equals B. There exists $Q(S,\epsilon) \gg 0$ such that if $q, g \geqslant Q(S,\epsilon)$ and $A, B \in S$,*

$$\left|\mathrm{Prob}_{x\in\mathrm{Conf}_g(\mathbb{F}_q)}(\mathrm{Jac}(C_x)(\mathbb{F}_q)_\ell \cong A \text{ and } \mathrm{Jac}(C_x^\sigma)(\mathbb{F}_q)_\ell \cong B) - \frac{c_R}{\#\mathrm{Aut}_R(M_{A,B})}\right| < \epsilon,$$

where c_R is the normalizing constant from Theorem 2.2.

That is, the class groups $\mathrm{Jac}(C_x)(\mathbb{F}_q)_\ell$ and $\mathrm{Jac}(C_x^\sigma)(\mathbb{F}_q)_\ell$ behave almost independently for g sufficiently large.

Proof. Note that

$$\mathrm{Jac}(C_x)(\mathbb{F}_q)_\ell \cong A \quad \text{and} \quad \mathrm{Jac}(C_x^\sigma)(\mathbb{F}_q)_\ell \cong B \iff \mathrm{Jac}(C_x)(\mathbb{F}_{q^2})_\ell \cong M_{A,B}.$$

By Propositions 3.4 and 4.7, for q sufficiently large we have

$$\left|\mathrm{Prob}(\mathrm{Jac}(C_x)(\mathbb{F}_{q^2})_\ell \cong M_{A,B}) - \frac{c_R}{\#\mathrm{Aut}_R(M_{A,B})}\right| < \epsilon.$$

Because R splits as a product,

$$\frac{c_R}{\#\mathrm{Aut}_R(M_{A,B})} = \frac{c_{\mathbb{Z}_\ell}}{\#\mathrm{Aut}_{\mathbb{Z}_\ell}(A)} \cdot \frac{c_{\mathbb{Z}_\ell}}{\#\mathrm{Aut}_{\mathbb{Z}_\ell}(B)}.$$

The result follows. □

Acknowledgements
It is a pleasure to thank Krishnaswami Alladi for explaining how to prove the Cohen–Lenstra identities using iterated Durfee's squares. We would also like to thank Melanie Wood for helpful comments which improved the exposition of this paper.

References

Ach06 J. Achter, *The distribution of class groups of function fields*, J. Pure Appl. Algebra **204** (2006), 316–333.

Bha05 M. Bhargava, *The density of discriminants of quartic rings and fields*, Ann. of Math. (2) **162** (2005), 1031–1063.

CL84 H. Cohen and H. Lenstra, *Heuristics on class groups of number fields*, in *Number theory, Noordwijkerhout 1983, Noordwijkerhout, 1983*, Lecture Notes in Mathematics, vol. 1068 (Springer, Berlin, 1984), 33–62.

DH71 H. Davenport and H. Heilbronn, *On the density of discriminants of cubic fields. II*, Proc. R. Soc. Lond. Ser. A **322** (1971), 405–420.

EVW16 J. S. Ellenberg, A. Venkatesh and C. Westerland, *Homological stability for Hurwitz spaces and the Cohen–Lenstra conjecture over function fields*, Ann. of Math. (2) **183** (2016), 729–786.

FK07 E. Fouvry and J. Klüners, *On the 4-rank of class groups of quadratic number fields*, Invent. Math. **167** (2007), 455–513.

FW89 E. Friedman and L. Washington, *On the distribution of divisor class groups of curves over a finite field*, in *Théorie des nombres, Quebec, PQ, 1987* (de Gruyter, Berlin, 1989), 227–239.

Gar15 D. Garton, *Random matrices, Cohen–Lenstra heuristics, and roots of unity*, Algebra Number Theory **9** (2015), 149–171.

Hal08 C. Hall, *Big symplectic or orthogonal monodromy modulo* ℓ, Duke Math. J. **141** (2008), 179–203.

Kly16 J. Klys, *The distribution of* p*-torsion in degree* p *cyclic fields*, Preprint (2016), arXiv:1610.00226v1.

Mal10 G. Malle, *On the distribution of class groups of number fields*, Experiment. Math. **19** (2010), 465–474.

Mic06 A. A. G. Michael, *Finite abelian actions on surfaces*, Topology Appl. **153** (2006), 2591–2612.

Mil18 D. Z. Milovic, *On the 8-rank of narrow class groups of* $\mathbb{Q}(\sqrt{-4pp}), \mathbb{Q}(\sqrt{-8pq})$*, and* $\mathbb{Q}(\sqrt{8pq})$, Int. J. Number Theory **14** (2018), 2165–2193.

RW06 M. Romagny and S. Wewers, *Hurwitz spaces*, in *Groupes de Galois arithmétiques et différentiels*, Séminaire et Congrès, vol. 13 (Société Mathématique de France, Paris, 2006), 313–341.

Smi16 A. Smith, *Governing fields and statistics for 4-Selmer groups and 8-class groups*, Preprint (2016), arXiv:1607.07860v1.

Woo17 M. M. Wood, *The distribution of sandpile groups of random graphs*, J. Amer. Math. Soc. **30** (2017), 915–958.

Yu97 J.-K. Yu, *Toward a proof of the Cohen–Lenstra conjecture in the function field case*, Preprint (1997).

Michael Lipnowski michael.lipnowski@mcgill.ca
Department of Mathematics, McGill University, Burnside Hall, Room 1005,
805 Sherbrooke Street West, Montreal, QC H3A 0B9, Canada

Jacob Tsimerman jacobt@math.toronto.edu
Department of Mathematics, University of Toronto, Bahen Centre, Room 6290,
40 St. George Street, Toronto, ON M5S 2E4, Canada