

Random pro- p groups, braid groups, and random tame Galois groups

Nigel Boston ^{*}and Jordan S. Ellenberg [†]

November 19, 2009

1 Introduction

In this article we introduce a heuristic prediction for the distribution of the isomorphism class of $G_S(p)$, the Galois group of the maximal pro- p extension of \mathbb{Q} unramified outside of S , where S is a "random" set of primes. That these groups should exhibit any statistical regularity is not at all obvious; our expectations in this direction are guided by the Cohen-Lenstra conjectures, which (among other things) predict quite precisely how often a fixed finite group appears as the ideal class group of a quadratic imaginary field.

The Cohen-Lenstra conjectures can be obtained in (at least) two ways. On the one hand, the distribution on finite abelian groups suggested by the heuristics has a good claim on being the most natural "uniform distribution" on the category of finite abelian p -groups. On the other hand, as observed by Friedman and Washington ([8], see also [1]) the conjectures can also be recovered via the analogy between number fields and function fields; here one thinks of the class group as the cokernel of $\gamma - 1$ where γ is a p -adic matrix drawn randomly from a suitable subset of the \mathbb{Q}_p -points of an algebraic group. We will show that both heuristic arguments can be generalized to the nonabelian pro- p case, and that both lead to the same prediction, Heuristic 2.4 below.

We conclude by describing some evidence, both theoretical and experimental, that supports (or at least is consistent with) Heuristic 2.4. We pay special attention to the interesting case where $p = 2$ and S consists of two primes congruent to 5 (mod 8). In this case, the heuristic appears to suggest that $G_S(p)$ is infinite 1/16 of the time.

1.1 Notation

When x and y are elements of a group, we use $x \sim y$ to mean "x is conjugate to y." We say a pro- p group Γ is *balanced* if $h_1(\Gamma, \mathbb{F}_p) = h_2(\Gamma, \mathbb{F}_p)$.

^{*}Partially supported by the Claude Shannon Institute, Science Foundation Ireland Grant 06/MI/006 and Stokes Professorship award, Science Foundation Ireland Grant 07/SK/I1252b. We thank Joann Boston for making the figure, Akshay Venkatesh for many useful comments, and Mike Fried for discussion of Heuristic 4.7.

[†]Partially supported by NSF-CAREER Grant DMS-0448750 and a Sloan research fellowship.

2 Statement of the heuristic

If S is a set of primes in \mathbb{Q} , we denote by $G_S(p)$ the Galois group of the maximal p -extension unramified away from S . Our aim in this section is to present a heuristic answer to the question: "When S is a random set of primes, what is the probability that $G_S(p)$ is isomorphic to some specified finite p -group Γ ?"

In order to phrase the question precisely, we need a little notation.

If $S = (\ell_1, \dots, \ell_g)$ is a g -tuple of primes congruent to 1 mod p , we denote by Z_i the closure of $\ell_i^{\mathbb{Z}}$ in \mathbb{Z}_p^* and by W_i the group $\mathbb{Z}_p/(\ell_i - 1)\mathbb{Z}_p$. Note that $G_S(p)^{ab}$ is isomorphic to $W = \bigoplus_{i=1}^g W_i$, a finite abelian p -group of rank g . In this paper, S always denotes an *ordered* finite set of primes.

Definition 2.1. The *type* of S is the sequence of subgroups (Z_1, \dots, Z_g) .

Note that when p is odd, the type of (ℓ_1, \dots, ℓ_g) is determined by the maximal power of p dividing $\ell_i - 1$ for each i ; in particular, the type of S carries the same information as the sequence of groups W_i . When $p = 2$, the type of S determines the W_i , but is not determined by it; for instance, primes ℓ which are 3 (mod 8) and those which are 7 (mod 8) are of different types, but both have $\mathbb{Z}_2/(\ell - 1)\mathbb{Z}_2 \cong \mathbb{Z}/2\mathbb{Z}$. We write $W(Z)$ for the finite p -group attached to a type $Z = (Z_1, \dots, Z_g)$.

Let $Z = (Z_1, \dots, Z_g)$ be a type, and Γ a finite p -group such that $\Gamma^{ab} \cong W(Z)$ and $h_1(\Gamma, \mathbb{F}_p) = h_2(\Gamma, \mathbb{F}_p) = g$. We define $P(Z, \Gamma, X)$ to be the proportion of g -tuples of primes $S = (\ell_1, \dots, \ell_g)$ in $[X, \dots, 2X]$ such that $G_S(p) \cong \Gamma$. Then the behavior of $P(Z, \Gamma, X)$ as X grows can be thought of as the probability that a random g -tuple of primes of type Z has Γ as its maximal unramified Galois group.

In order to state our heuristic estimate for this probability, we need a little more notation.

Write $\pi : \Gamma \rightarrow \Gamma^{ab}$ for the natural projection.

Definition 2.2. Let $A_Z(\Gamma)$ be the number of pairs $((c_1, \dots, c_g), \iota)$ where (c_1, \dots, c_g) is a g -tuple of conjugacy classes in Γ and ι is an involution in Γ (automatically trivial when p is odd) such that:

- $c_i^z = c_i$ for all $z \in Z_i$;
- the elements $\pi(c_1), \dots, \pi(c_g)$ generate Γ^{ab} ;
- the map $W(Z) \rightarrow \Gamma^{ab}$ guaranteed by the two conditions above is an isomorphism;
- if $p = 2$, and ι_i is the unique nontrivial involution in the cyclic subgroup of Γ^{ab} generated by $\pi(c_i)$, we have $\sum_{i=1}^g \iota_i = \pi(\iota)$.

Remark 2.3. When no confusion is likely (i.e. when we are restricting attention to a particular type Z) we will write $A(\Gamma)$ for $A_Z(\Gamma)$.

Heuristic 2.4. $\lim_{X \rightarrow \infty} P(Z, \Gamma, X)$ exists and is equal to $A_Z(\Gamma)/|\text{Aut}(\Gamma)|$.

In the following sections, we explain two justifications for Heuristic 2.4, each one adapted from a justification for the Cohen-Lenstra heuristics.

3 Justification 1: random p -groups with inertia data

There is no uniform distribution on a countably infinite set. Nonetheless, there is a natural "uniform" distribution on the set of isomorphism classes of finite abelian p -groups, following the usual principle

that objects in a category should be assigned a weight inversely proportional to the order of their automorphism group. (See Leinster[13] for a thorough development of this idea.) The sum of $|\text{Aut}(\Gamma)|^{-1}$ as Γ ranges over isomorphism classes of finite abelian p -groups is finite; thus, there is a unique probability distribution on these isomorphism classes such that the mass assigned to Γ is proportional to $|\text{Aut}(\Gamma)|^{-1}$. The Cohen-Lenstra heuristic can be thought of as asserting that the p -part of the class group of a random imaginary quadratic field is a “random” abelian group, in the sense that it follows this “uniform” distribution on finite abelian p -groups. Many of the more general conjectures of Cohen, Lenstra, and Martinet are in the same spirit, predicting that class groups of extensions with more complicated structure (e.g. Galois extensions with Galois group K) are “random” objects in more general categories (e.g. the category of finitely generated $\mathbb{Z}/p\mathbb{Z}[K]$ -modules) in the analogous sense.

We now present an argument in a similar spirit that leads us to Heuristic 2.4.

One might start by trying to construct a probability distribution on finite p -groups where the measure assigned to Γ is proportional to $|\text{Aut}(\Gamma)|^{-1}$. This, however, is easily seen to be incorrect. The relevant difference between the Cohen-Lenstra context and the one treated in the present paper is that, as far as we know, any finite abelian p -group can arise as the p -part of the class group of an imaginary quadratic field. By contrast, not every balanced p -group can be $G_S(p)$ for some set of primes S of a given type. For instance, take $p = 3, g = 2$, and $Z_1 = Z_2 = 1 + 3\mathbb{Z}_3$. Then $G_S(p)$ is a group with abelianization $(\mathbb{Z}/3\mathbb{Z})^2$, and the image of tame inertia at each prime in S lies in a conjugacy class c of Γ satisfying $c^4 = c$. But there exists a group G of order 3^5 which is not generated by any two elements which are conjugate to their fourth powers. Thus, $G_S(p)$ cannot be isomorphic to Γ , and any reasonable heuristic should reflect this by assigning probability 0 to Γ .

In order to avoid problems of this kind, we introduce the notion of a “pro- p -group with local data of type Z .” As in Definition 2.2, we denote by π the natural projection from a group to its abelianization.

Definition 3.1. A *pro- p group with local data of type Z* is a balanced pro- p group Γ endowed with a g -tuple of conjugacy classes (c_1, \dots, c_g) and an involution ι , where ι is an involution in Γ (automatically trivial when p is odd), such that:

- $c_i^z = c_i$ for all $z \in Z_i$;
- the projections $\pi(c_1), \dots, \pi(c_g)$ generate Γ^{ab} ;
- the map $W(Z) \rightarrow \Gamma^{ab}$ guaranteed by the two conditions above is an isomorphism;
- if $p = 2$, and ι_i is the unique nontrivial involution in the cyclic subgroup of Γ^{ab} generated by $\pi(c_i)$, we have $\sum_{i=1}^g \iota_i = \pi(\iota)$.

For each $\ell \neq p$, we fix a generator τ_ℓ of the tame inertia group of $G_{\mathbb{Q}_\ell}$. We also fix a complex conjugation c in $G_{\mathbb{Q}}$. When $S = (\ell_1, \dots, \ell_g)$ is a g -tuple of primes with type Z , the unramified Galois group $G_S(p)$ acquires the structure of pro- p -group with local data of type Z in a natural way; namely, c_i is the conjugacy class of the image of τ_{ℓ_i} in $G_S(p)$ and ι is the image of c .

Finite p -groups with local data of type Z constitute the objects of a category $\mathcal{C}_{p,Z}$ whose morphisms are just homomorphisms of groups preserving c_1, \dots, c_g and ι . Just as the category of finite abelian p -groups is the “natural” home of the p -part of the class group of a quadratic imaginary field, a finite p -group with inertia data is the natural home of $G_S(p)$. Thus, the Cohen-Lenstra philosophy would suggest that the probability that $G_S(p)$ and $(\Gamma, (c_1, \dots, c_g), \iota)$ are isomorphic in $\mathcal{C}_{p,Z}$ should be proportional to $|\text{Aut}_{\mathcal{C}_{p,i}}(\Gamma, (c_1, \dots, c_g), \iota)|^{-1}$. The probability that $G_S(p)$ is isomorphic

to Γ as a group is then obtained by summing over all $\text{Aut}(\Gamma)$ -orbits of inertia data on Γ . This yields precisely the prediction that $P(Z, \Gamma, X)$ is proportional to $A_Z(\Gamma)/|\text{Aut}(\Gamma)|$.

In particular, the fact that the group $G_S(p)$ carries local data of type Z shows that Heuristic 2.4 passes one easy plausibility check:

Proposition 3.2. *If $A_Z(\Gamma) = 0$, there is no g -tuple S of type Z such that $G_S(p)$ is isomorphic to Γ .*

Remark 3.3. This justification also suggests obvious refinements of Heuristic 2.4 taking local data into account; for instance, when $p = 2$, one can ask for what proportion of g -tuples S there is an isomorphism $G_S(p) \cong \Gamma$ taking complex conjugation to a given conjugacy class of involutions in Γ . Furthermore, one could speculate that wildly ramified extensions can be brought into the picture as well by adding to $(c_1, \dots, c_g), \iota$ the data of a homomorphism $G_{\mathbb{Q}_p} \rightarrow \Gamma$. This would bring us closer to the heuristics suggested by Bhargava [2] and Roberts [17] for the number of extensions of global fields with specified Galois groups and restrictions on local behavior. On the other hand, it takes us farther away from the analogy with function fields over finite fields that we explore in the next section.

The class group of a number field is always finite, so the probability distribution is determined by the fact that the sum of the mass of all finite groups is 1. In the present situation, we have no such assurance; for many choices of the type Z , $G_S(p)$ can be either finite or infinite, so we have no principled way of choosing one probability distribution on finite p -groups among all those proportional to $A_Z(\Gamma)/|\text{Aut}(\Gamma)|$. We address this problem in the next section.

4 Justification 2: Random pro- p braids and the analogy with function fields

In this section we give an alternative (though certainly related) justification for Heuristic 2.4, based on the analogy between \mathbb{Z} and the ring $\mathbb{F}_q[t]$ (or, more generally, the coordinate ring of an affine algebraic curve over a finite field.)

4.1 The action of Frobenius on the fundamental group of an affine curve

Let \mathbb{F}_q be a finite field of characteristic other than p , let C/\mathbb{F}_q be a smooth projective algebraic curve, let x_∞ be a point of $C(\mathbb{F}_q)$, and let U/\mathbb{F}_q be an open subscheme of C not containing x_∞ . Then the étale fundamental group $\pi_1(U)$ fits into an exact sequence

$$1 \rightarrow \pi_1(U_{\overline{\mathbb{F}_q}}) \rightarrow \pi_1(U) \rightarrow G_{\mathbb{F}_q} \rightarrow 1$$

in which the second map admits a section s attached to any choice of tangential base point at x_∞ .

Note that the p -cyclotomic character $\chi_p : \pi_1(U) \rightarrow \mathbb{Z}_p^*$ factors through the projection to $G_{\mathbb{F}_q}$ and sends Frobenius to q . In particular, the image of χ_p is procyclic. The p -cyclotomic character of $G_{\mathbb{Q}}$ is surjective on \mathbb{Z}_p^* ; in particular, when $p = 2$, the image is not procyclic. In order to make our analogy as precise as possible, we assume from now on that p is odd and that q generates \mathbb{Z}_p^* . We will return in section 4.4 to the case $p = 2$.

Remark 4.1. The hypothesis that q generates \mathbb{Z}_p^* is natural when drawing an analogy with \mathbb{Q} ; to generate conjectures about maximal pro- p extensions with restricted ramification of larger number fields K would require a different hypothesis on q reflecting the abelian extensions of \mathbb{Q} contained

in K . In particular, the presence of roots of unity in K ought to change the statistical distribution of the Galois groups, even in the abelian cases considered by Cohen, Lenstra, and Martinet. Such a discrepancy from Cohen-Lenstra predictions has in fact been observed numerically by Malle [14] in the case of number fields and Rozenhart [18] in the case of function fields. Forthcoming work of Garton [9] will explain how one should modify the Cohen-Lenstra-Martiniet heuristics in the presence of roots of unity, based on a function-field argument like the one we present in the present paper.

Let G_∞ be a decomposition group of $\pi_1(U)$ at x_∞ . Write I_∞ for the inertia subgroup of G_∞ . Note that the section $s : G_{\mathbb{F}_q} \rightarrow \pi_1(U)$ factors through G_∞ .

Write $G_U(p)$ for the quotient of $\pi_1(U)$ by the normal subgroup generated by G_∞ . This is the Galois group of the maximal étale cover of U which is totally split at x_∞ . (Note that, in the number field case, the maximal pro- p extension of \mathbb{Q} unramified outside S is indeed totally split at ∞ , because p is odd here.)

The geometric fundamental group $\pi_1(U_{\overline{\mathbb{F}_q}})$ is isomorphic to a free profinite group on N generators, where $N + 1$ is the number of punctures of $U_{\overline{\mathbb{F}_q}}$. Write F for the quotient of $\pi_1(U_{\overline{\mathbb{F}_q}})$ by the normal subgroup generated by I_∞ . The group $s(G_{\mathbb{F}_q})$ acts on $\pi_1(U_{\overline{\mathbb{F}_q}})$ by conjugation, and this action descends to an action on F . Since $G_{\mathbb{F}_q}$ is procyclic, we can describe this action by specifying the action on F of the Frobenius Frob_q in $s(G_{\mathbb{F}_q})$, which is an automorphism $\alpha \in \text{Aut}(F)$.

The group F is freely generated by N elements x_1, \dots, x_N , each one a generator of tame inertia at a puncture, subject to the single relation $x_1 \dots x_N = 1$. The automorphism α has the special property that it sends each x_i to a conjugate of x_j^q for some j . The automorphisms of F with this property lie in the *pro- p braid group* $B_N \subset \text{Aut}(F)$, which we define in the next section.

Now $G_U(p)$ is precisely the quotient of F by relations $\alpha(x_i) = x_i, \forall i$. In particular, $G_U(p)$ is determined by the automorphism α . The main idea driving the heuristics in this paper is that α should be a *random* element drawn from a certain subset of the profinite group B_N under Haar measure. In the following sections we describe the consequences of this heuristic for the behavior of $G_U(p)$.

Remark 4.2. The Cohen-Lenstra heuristics can also be recovered from this point of view, as was first observed by Friedman and Washington [8]; the p -part of the class group of a random quadratic imaginary field is analogous to the \mathbb{F}_q -rational points in $\text{Jac}(X)[p]$, where X/\mathbb{F}_q is a hyperelliptic curve of large genus. This group, in turn, is the cokernel of $\gamma - 1$, where $\gamma \in \text{GSp}_{2g}(\mathbb{Z}_p)$ is a matrix representing the action of Frobenius on $H_1(X, \mathbb{Z}_p)$. In fact, Frobenius lies in the coset $\text{GSp}_{2g}^q(\mathbb{Z}_p)$ of $\text{Sp}_{2g}(\mathbb{Z}_p)$ consisting of matrices which scale the symplectic form by q . Taking γ to be a random element of $\text{GSp}_{2g}(\mathbb{Z}_p)$ yields precisely the Cohen-Lenstra heuristics.

Friedman and Washington took γ to be a random element of $\text{GL}_{2g}(\mathbb{Z}_p)$; the observation that the correct set over which to average is $\text{GSp}_{2g}^q(\mathbb{Z}_p)$ is due to Achter [1].

In fact, as in [1], these heuristic arguments can often be turned into proofs that heuristics of Cohen-Lenstra type are true over function fields “in the large q limit.” In this connection, see also the work of Katz and Sarnak on the relation between random p -adic matrices and the distribution of zeroes of L -functions [11].

4.2 The pro- p braid group and random pro- p groups

Let p be an odd prime, and let F be the pro- p group generated by elements x_1, \dots, x_N subject to the single relation $x_1 \dots x_N = 1$. We define the *pure pro- p braid group* P_N to be the subgroup of $\text{Aut}(F)$ consisting of automorphisms α such that $\alpha(x_i) \sim x_i$ for all i . For each $(q, \sigma) \in \mathbb{Z}_p^* \times S_N$ we denote by $B_N(q, \sigma)$ the set of automorphisms α of F such that

$$\alpha(x_i) \sim x_{\sigma(i)}^q$$

for all i . Then the pro- p braid group B_N is defined by

$$B_N = \bigcup_{q, \sigma} B_N(q, \sigma)$$

The pro- p braid group fits into an exact sequence

$$1 \rightarrow P_N \rightarrow B_N \rightarrow \mathbb{Z}_p^* \times S_N \rightarrow 1$$

where the preimage of $(q, \sigma) \in S_N \times \mathbb{Z}_p^*$ is $B_N(q, \sigma)$.

Remark 4.3. The definition of the pro- p braid group is due to Ihara [10]. More precisely, his pro- p braid group on N strands is the image of our B_N in $\text{Out}(F)$.

Remark 4.4. The natural map $B_N \rightarrow \mathbb{Z}_p^*$ is surjective, so $B_N(q, \sigma)$ is nonempty; but we are not aware of any purely group-theoretic proof of this fact. Rather, one observes (again following Ihara) that when F is identified with the pro- p geometric fundamental group of an N -punctured genus 0 algebraic curve over \mathbb{Q} , the resulting action of $G_{\mathbb{Q}}$ on F induces a map $G_{\mathbb{Q}} \rightarrow B_N$ whose composition with the projection to \mathbb{Z}_p^* is the cyclotomic character.

We are now ready to define the main object of study of the present paper: namely, a “random pro- p group” $\mathbf{G}(q, \sigma)$.

Definition 4.5. We denote by $\mathbf{G}(q, \sigma)$ the quotient of F by the relations $\alpha(x_i) = x_i, i = 1 \dots N$, where α is a random element of $B_N(q, \sigma)$ in Haar measure.

4.3 Statistical properties of $\mathbf{G}(q, \sigma)$

Write k_1, \dots, k_g for the lengths of the cycles of σ , and choose representatives i_1, \dots, i_g of the cycles. Write g_j for the image of x_{i_j} in $\mathbf{G}(q, \sigma)$. Then one sees that $g_j \sim g_j^{q^{k_j}}$. We also note that

$$\mathbf{G}(q, \sigma)^{ab} = \sum_{j=1}^g \mathbb{Z}_p / (q^{k_j} - 1) \mathbb{Z}_p.$$

We denote this finite abelian group by $W(q, \sigma)$.

Proposition 4.6. *Let g be the number of cycles of σ . Then $\mathbf{G}(q, \sigma)$ is a g -generated, g -related pro- p group.*

Proof. $\mathbf{G}(q, \sigma)$ is clearly generated by the g elements g_j , and is specified by the g relations $\alpha(x_{i_j})x_{i_j}^{-1}$. So $h_2(\mathbf{G}(q, \sigma)) \leq g$; since $\mathbf{G}(q, \sigma)^{ab}$ is finite, we also have $h_2(\mathbf{G}(q, \sigma)) \geq g$. \square

Let Γ be some finite g -generated p -group, and let c_1, \dots, c_g be conjugacy classes in Γ such that

- c_1, \dots, c_g generate Γ ;
- $c_i^{q^{k_i}} = c_i$ for all i .

We denote by (Γ, \mathbf{c}) the data of a pro- p group Γ together with a g -tuple of conjugacy classes as above. Given a permutation σ together with a set i_1, \dots, i_g of cycle representatives for σ , we write

$$\text{Epi}(\mathbf{G}(q, \sigma), (\Gamma, \mathbf{c}))$$

for the set of surjective homomorphisms $\mathbf{G}(q, \sigma) \rightarrow \Gamma$ which send g_j to the conjugacy class c_j for each j . (The notation is slightly misleading insofar as $\text{Epi}(\mathbf{G}(q, \sigma), (\Gamma, \mathbf{c}))$ depends not only on σ but on the choice of cycle representatives, but this ambiguity will not concern us.)

Similarly, we write $\text{Epi}(F, (\Gamma, \mathbf{c}))$ for the set of surjective homomorphisms $F \rightarrow \Gamma$ which send $x_{\sigma^m(i_j)}$ to the conjugacy class $c_j^{q^m}$ for each j . In particular, the pullback of an element of $\text{Epi}(\mathbf{G}(q, \sigma), (\Gamma, \mathbf{c}))$ to F always lies in $\text{Epi}(F, (\Gamma, \mathbf{c}))$.

The set $\text{Epi}(F, (\Gamma, \mathbf{c}))$ carries an action of $P_N \subset \text{Aut}(F)$ by composition on the left. Our first aim is to make an educated guess about the transitivity of this action.

Heuristic 4.7. *Suppose Γ is balanced. Then the action of P_N on $\text{Epi}(F, (\Gamma, \mathbf{c}))$ is transitive when N is sufficiently large relative to g .*

Remark 4.8. A transitivity theorem of this kind for the action of the full braid group has been proven by Fried and Völklein [7, Appendix]. More precisely, the orbits of the full braid group are identified with a quotient of the Schur multiplier $H_2(\Gamma, \mathbb{Z})$; in this case, the Schur multiplier is trivial because Γ is balanced.

We are now ready to present the second justification for Heuristic 2.4.

Proposition 4.9. *Suppose Γ is a balanced finite p -group with abelianization isomorphic to $W(q, \sigma)$, and assume Heuristic 4.7. Then the probability that $\mathbf{G}(q, \sigma)$ is isomorphic to G is $A_Z(\Gamma)/|\text{Aut}(\Gamma)|$.*

Proof. The set $\text{Epi}(\mathbf{G}(q, \sigma), (\Gamma, \mathbf{c}))$ is the set of α -fixed points on $\text{Epi}(F, (\Gamma, \mathbf{c}))$, where α is a random element of the coset $B_N(q, \sigma)$ of P_N in B_N . By Burnside's lemma, the average number of fixed points of γ is the number of fixed points of $(q, \sigma) \in \mathbb{Z}_p^* \times S_N$ in its action on the P_N -orbits on $\text{Epi}(F, (\Gamma, \mathbf{c}))$. By Heuristic 4.7, there is only one such orbit. We conclude that

$$\mathbb{E}|\text{Epi}(\mathbf{G}(q, \sigma), (\Gamma, \mathbf{c}))| = 1.$$

Lemma 4.10. *Every surjection from $\mathbf{G}(q, \sigma)$ to Γ is an isomorphism.*

Proof. The long exact sequence in group cohomology yields an exact sequence

$$0 \rightarrow H^1(\Gamma, \mathbb{Q}_p/\mathbb{Z}_p)/pH^1(\Gamma, \mathbb{Q}_p/\mathbb{Z}_p) \rightarrow H^2(\Gamma, \mathbb{Z}/p\mathbb{Z}) \rightarrow H^2(\Gamma, \mathbb{Q}_p/\mathbb{Z}_p)[p] \rightarrow 0.$$

The first two terms have dimension g , so the third term vanishes; hence, so does $H^2(\Gamma, \mathbb{Q}_p/\mathbb{Z}_p)$.

Write K for the kernel of the map $\mathbf{G}(q, \sigma) \rightarrow \Gamma$. The inflation-restriction sequence gives

$$H^1(\Gamma, \mathbb{Q}_p/\mathbb{Z}_p) \rightarrow H^1(\mathbf{G}(q, \sigma), \mathbb{Q}_p/\mathbb{Z}_p) \rightarrow H^1(K, \mathbb{Q}_p/\mathbb{Z}_p)^\Gamma \rightarrow H^2(\Gamma, \mathbb{Q}_p/\mathbb{Z}_p) = 0.$$

Our hypothesis on Γ^{ab} implies that the first map is an isomorphism; we conclude that $H^1(K, \mathbb{Q}_p/\mathbb{Z}_p)^\Gamma = 0$, which implies that K is trivial. \square

Summing over all \mathbf{c} of type Z , we find that the expected number of isomorphisms from $G(q, \sigma)$ to Γ is exactly $A_Z(G)$. The number of such isomorphisms is either 0 or $|\text{Aut}(\Gamma)|$; the desired result follows. \square

We now explain how Heuristic 2.4 follows from Proposition 4.9. The N punctures of the affine curve U/\mathbb{F}_q (excepting x_∞) can be thought of as g closed points; we refer to this set of points as S and define the type S as in the first section. The hypothesis that S is of some specified type Z can be thought of as a condition on σ ; namely, that $W(q, \sigma) \cong W(Z)$. So the probability that

$G_U(p) \cong G$ can be heuristically estimated as the probability that $\mathbf{G}(q, \sigma) \cong G$, conditional on the fact that $W(q, \sigma) \cong W(Z)$. By Proposition 4.9, This probability is $A_Z(\Gamma)/|\text{Aut}(\Gamma)|$ for each such σ .

More generally, the argument of this section suggests the following: Let Γ be a p -group such that the action of the pure braid group P_N on $\text{Epi}(F, (\Gamma, \mathbf{c}))$ is transitive; then the expected number of epimorphisms from $G_S(p)$ to Γ is $A_Z(\Gamma)$. (We emphasize that we do not presently know any examples where the action of P_N is intransitive.)

In particular, suppose Γ is a group of p -class k , with the property that the existence of a surjection from a balanced group G to Γ implies that Γ is the quotient of G by the $(k+1)$ -th term in its lower p -central series. (We will encounter such examples in the data presented in the following sections.) Then the probability that the p -class k quotient of $G(q, \sigma)$ is isomorphic to Γ should be $A_Z(\Gamma)/|\text{Aut}(G)|$.

4.4 The case $p = 2$

The geometric analogy in the case $p = 2$ is somewhat more difficult to justify, since there is no choice of q making the cyclotomic character $\chi : G_{\mathbb{F}_q} \rightarrow \mathbb{Z}_2^*$ surjective. Our belief in Heuristic 2.4 in this case stands on three legs:

- The argument via Justification 1 that $\lim P(Z, \Gamma, X)$ should be proportional to $A_Z(\Gamma)/|\text{Aut}(\Gamma)|$ for all p ;
- The argument via Justification 2 that, in case p is odd, the constant of proportionality should be 1;
- The case $p = 2$ is the easiest to test experimentally; as we shall see in the following section, Heuristic 2.4 appears to agree very well with the experimental data and with provable asymptotics when such are available.

5 Evidence

Let p be an odd prime and S a set of g primes that are all $1 \pmod{p}$. Then the Galois group of the maximal p -extension of \mathbb{Q} unramified outside S is a pro- p group with generator rank g and relator rank g . The same result holds for $p = 2$ if we allow ramification at ∞ .

Given a pro- p group Γ with generator rank and relator rank both equal to g , we consider the set of g -tuples of primes, all $1 \pmod{p}$, such that the corresponding Galois group is isomorphic to Γ . In all cases considered so far, this set appears to have density predicted by Heuristic 2.4. In some cases this result on the density can be proven, whereas in others we give experimental evidence.

5.1 Groups with abelianization $\mathbb{Z}/p\mathbb{Z} \times \mathbb{Z}/p\mathbb{Z}$ (p odd)

In this case, S consists of two primes, say q, r , that are $1 \pmod{p}$ but not $1 \pmod{p^2}$ (*).

(i) The smallest 2-generator 2-relator p -group Γ is unique of order p^3 (denoted `SmallGroup(p3, 4)` in the MAGMA database). As shown on p.129 of [12], the Galois group of the maximal p -extension of \mathbb{Q} unramified outside q, r is isomorphic to Γ if and only if q is not a p th power \pmod{r} or r is not a p th power \pmod{q} . Among ordered pairs of primes satisfying (*), the natural density of such pairs is then $1 - 1/p^2$.

As for the heuristic, there are $p^2 - 1$ conjugacy classes that are outside the Frattini subgroup of Γ and that are closed under taking $(p + 1)$ th powers. Working in $\mathbb{Z}/p\mathbb{Z} \times \mathbb{Z}/p\mathbb{Z}$, we see that $p(p + 1)(p - 1)^2$ ordered pairs of such conjugacy classes generate Γ^{ab} and hence Γ . On the other hand, the automorphism group of Γ has order $(p - 1)p^3$ and so $A(\Gamma)/|\text{Aut } \Gamma| = 1 - 1/p^2$.

(ii) The next smallest 2-generator 2-relator 3-group with abelianization $\mathbb{Z}/3\mathbb{Z} \times \mathbb{Z}/3\mathbb{Z}$ and nonzero probability of arising is a unique group Γ of order 3^5 . There is another 2-generator 2-relator group of order 3^5 , but it is not generated by elements conjugate to their 4th power. For such groups our conjecture is trivially true, since it follows that $A(\Gamma) = 0$ and also that Γ lacks elements that would play the role of tame inertia generators.

Suppose that q, r are primes that are $1 \pmod{3}$ but not $1 \pmod{9}$. By the method of Boston and Leedham-Green [5], the Galois group of the maximal p -extension of \mathbb{Q} unramified outside q, r is isomorphic to Γ if and only if two of the Sylow 3-subgroups of the ray class groups of modulus qr of its 4 cubic subextensions are isomorphic to $\mathbb{Z}/3\mathbb{Z} \times \mathbb{Z}/3\mathbb{Z} \times \mathbb{Z}/3\mathbb{Z}$ and two isomorphic to $\mathbb{Z}/3\mathbb{Z} \times \mathbb{Z}/9\mathbb{Z}$. This implies in particular that q is a cube mod r and r is a cube mod q , a condition which is satisfied for $1/9$ of all pairs (q, r) . Heuristic 2.4 asserts that $2/9$ of these will have $G_{\{q,r\}}(3)$ isomorphic to Γ , since $A(\Gamma)/|\text{Aut } \Gamma| = 2/81$.

We tested 58 pairs with q a cube mod r and r a cube mod q ; a MAGMA computation shows that 13 of these pairs satisfy the condition on ray class groups of subextensions equivalent to $G_{\{q,r\}}(3) \cong \Gamma$. The empirical density $13/58 \times 1/9 = 0.0249$ compares well with the predicted $2/81 = 0.0247$.

5.2 Groups with abelianization $\mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}$

In this case, S consists of two primes q and r both congruent to $3 \pmod{4}$. Order q and r such that q is a square mod r . Then, as shown in Theorem 2.1 of Boston-Perry [6], the Galois group $G_{\{q,r\}}(2)$ is semidihedral of order 2^k , where 2^k is the largest power of 2 dividing $q^2 - 1$. In particular, $G_{\{q,r\}}(2)$ is determined by the type of (q, r) . This agrees with the heuristic, which gives $A_Z(\Gamma)/|\text{Aut } \Gamma| = 1$ for each semidihedral group Γ .

5.3 Groups with abelianization $\mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/4\mathbb{Z}$

In this case the two primes q, r are $3 \pmod{4}$ and $5 \pmod{8}$. Let us suppose further that q is $3 \pmod{8}$, which fixes the type.

(i) The smallest such 2-relator group is the modular group of order 16. The Galois group of the maximal 2-extension unramified outside $\{q, r\}$ is isomorphic to this group if and only if q is not a square \pmod{r} . This occurs with probability $1/2$, as Heuristic 2.4 predicts.

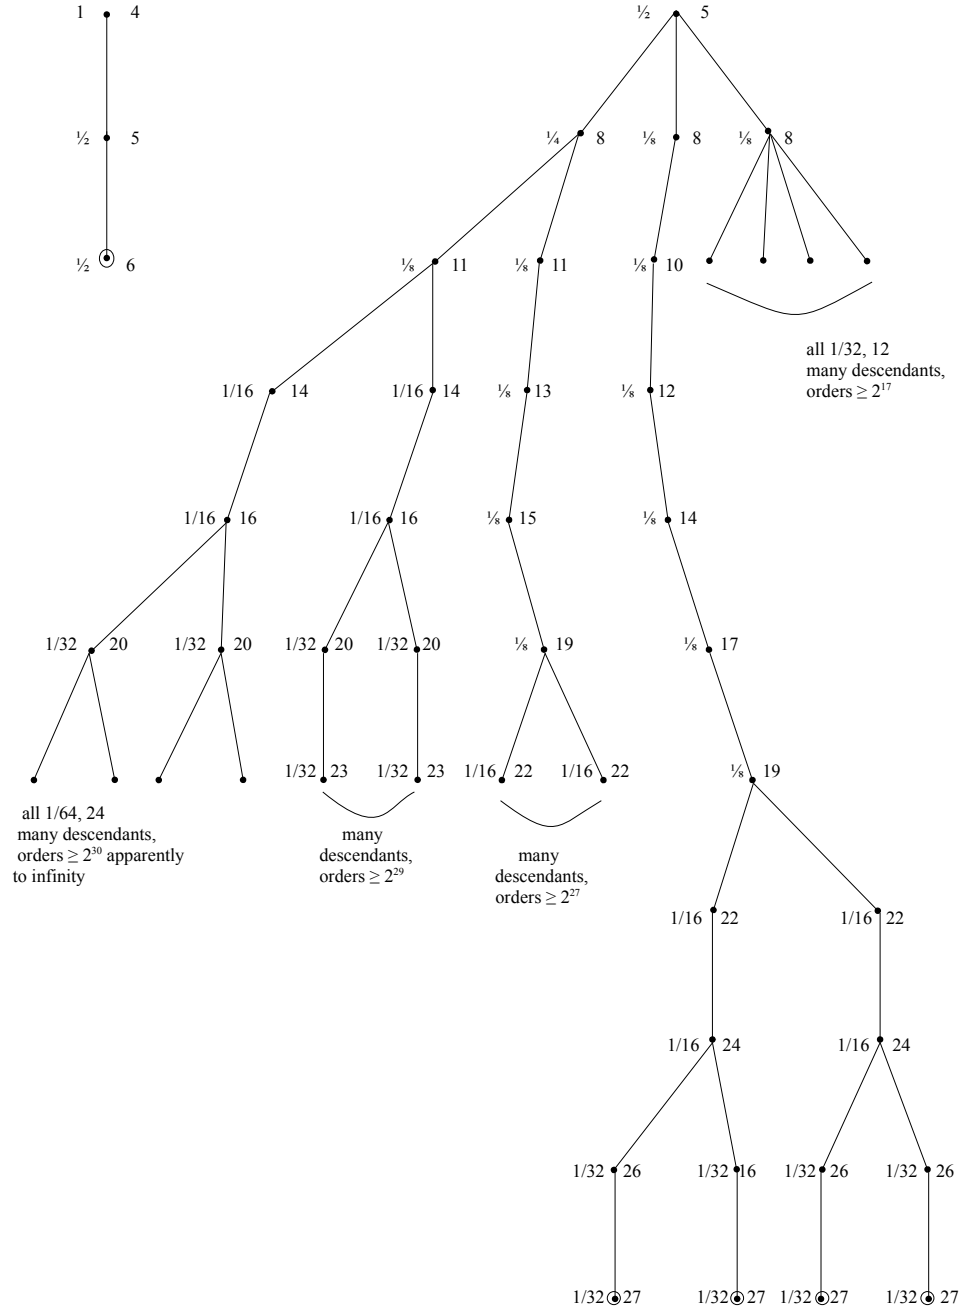
(ii) The next smallest such 2-relator group with nonzero probability of arising has order 128 (SmallGroup(128, 87) in the MAGMA database). As shown in Boston-Perry [6], this arises if and only if q is a 4th power \pmod{r} . This occurs with probability $1/4$ and agrees with Heuristic 2.4.

(iii) The next smallest such 2-relator groups with nonzero probability are the two groups Γ_1, Γ_2 of order 2^{19} found by Boston and Leedham-Green [5]. They showed there that the Galois group of the maximal 2-extension unramified outside $\{q, r\}$ is isomorphic to Γ_1 or Γ_2 if and only if q is a square but not a 4th power mod r and the Sylow 2-subgroup of the ray class group of modulus qr of $\mathbb{Q}(\sqrt{-qr})$ is isomorphic to $\mathbb{Z}/2 \times \mathbb{Z}/16$. The heuristic suggests that each group should arise $1/16$ of the time, so that this case should arise $1/8$ of the time.

The experimental evidence agrees - we test it by letting $q = 5$ and p run through all primes < 100000 that are $19 \pmod{40}$ (so that, as assumed, p is $3 \pmod{8}$ and a square but not a 4th power \pmod{q}). In general, the Sylow 2-subgroup of the ray class group is $\mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/2^n\mathbb{Z}$ with

$n \geq 4$. The cases $n = 4, 5, 6, 7$ appear respectively 301, 151, 74, 42 times, which clearly suggests that half of the $1/4$ of cases remaining from (i) and (ii) above have $n = 4$.

5.4 Groups with abelianization $\mathbb{Z}/4\mathbb{Z} \times \mathbb{Z}/4\mathbb{Z}$



Here the two primes q, r are $5 \pmod{8}$, which fixes the type.

The figure gives a subtree of the O'Brien tree [16] of 2-generator 2-groups. The vertices of his tree are isomorphism classes of 2-generator 2-groups. It has a root, namely $\mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}$, and the groups at distance $n - 1$ from the root are those of 2-class n . A group of 2-class n is connected to a group H of 2-class $n + 1$ if and only if it is isomorphic to $H/P_n(H)$, where $P_n(H)$ is the last but one term of its lower 2-central series.

Our subtree consists of those groups G that have abelianization $\mathbb{Z}/4\mathbb{Z} \times \mathbb{Z}/4\mathbb{Z}$, have nonzero "mass" $A(G)$ (its value is given to the left of each vertex), and are 2-class quotients of some (possibly infinite) 2-relator pro-2 group. We call such groups *viable*. This last matter can be detected by checking that the 2-multiplicator rank of G minus its nuclear rank is at most 2. If the nuclear rank of G is zero, then it is 2-relator. In that case it has no descendants and is a terminal vertex. Otherwise, if the nuclear rank of G is r , then (except for the root) its immediate descendants all have order $2^r|G|$. This is important for us since, except for the groups of 2-class 2, one descendant of a group cannot be a quotient of another and we do indeed (assuming transitivity of the pure braid group action) have a distribution on a rooted tree. Note that, at least under this assumption, one might expect $A(\Gamma)$ to be the number of surjections from $G_{\{q,r\}}(2)$ to Γ . The exponent of the order of a group is placed to the right of the vertex.

The subtree falls into 4 parts. By Part I, we mean the isolated twig to the left. Let $\Gamma_i (i = 1, 2, 3)$ denote the three groups of 2-class 3 and order 256. Parts II, III, and IV will refer to the subtrees with roots Γ_2, Γ_3 , and Γ_1 respectively. Two of the abelianizations of the index 2 subgroups of each group are isomorphic to $\mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/4\mathbb{Z} \times \mathbb{Z}/4\mathbb{Z}$, whereas the third is $\mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/4\mathbb{Z} \times \mathbb{Z}/4\mathbb{Z}$ for Γ_1 and $\mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/8\mathbb{Z}$ for the other two groups. Ray class groups of quadratic fields tell us that Γ_1 arises if one of q, r is a 4th power mod the other but not vice versa, which does occur with density $1/4$, and that Γ_2 or Γ_3 arises otherwise, occurring with density $1/4$. These match the predicted values.

5.4.1 Part I

The smallest 2-relator group in the subtree has order 64 (SmallGroup(64,28) in the MAGMA database) and arises if and only if q is not a square mod r . This occurs with probability $1/2$, as predicted by Heuristic 2.4.

5.4.2 Part II

The nuclear rank of Γ_2 is 2, indicating that its descendant tree should not be too complicated. Using MAGMA to produce its viable descendants we find that its subtree is finite, terminating in four groups of order 2^{27} and 2-class 12. These have not appeared in the literature before. They each have mass $1/32$ according to the heuristic. Among the abelianizations of their index 2 subgroups is still $\mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/8\mathbb{Z}$. On the other hand, the viable descendants of Γ_3 all have larger abelianizations of index 2 subgroups. It follows that if neither q nor r is a 4th power mod the other, then the Galois group of the maximal 2-extension of \mathbb{Q} unramified outside $\{q, r\}$ is one of these 4 groups of order 2^{27} . Moreover, the total mass they carry, namely $1/8$, agrees with our conjecture, which furthermore suggests that each of the 4 groups should occur equally often as q, r vary.

5.4.3 Part III

The descendant tree of Γ_3 contains groups with nuclear rank 5 and higher, which makes it prohibitive to calculate to any depth. By the process of elimination above, it and its descendants correspond to

q, r that are both 4th powers mod the other. Since in this case the corresponding ray class groups include one isomorphic to $\mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/2^n\mathbb{Z}$ for any $n \geq 4$, there will be infinitely many possible Galois groups, but as conjectured in [3], the evidence suggests that these groups are all finite.

That evidence came from performing the experiment of looking at millions of pro-2 groups with presentation of the form $\langle x, y | x^a = x^5, y^b = y^5 \rangle$ (as a and b vary through words in x and y) and by saving those whose 2-class quotients appeared to grow unboundedly and whose finite index subgroups had finite abelianization (within reasonable computational limits). All the groups that passed these filters had abelianizations of all index 2 subgroups isomorphic to $\mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/4\mathbb{Z} \times \mathbb{Z}/4\mathbb{Z}$; in other words, it appears that whenever $G_{\{q,r\}}(2)$ is infinite, it is a descendant of Γ_1 . We have established above that all descendants of Γ_2 are finite; we conjecture that, likewise, all descendants of Γ_3 are finite.

5.4.4 Part IV

The first author also suggested in [3] that all viable 2-relator descendants of Γ_1 should be infinite, and that the sequence of orders of their 2-class quotients should be a particular one (A001461 in Sloane's database of sequences). Our new investigations indicate that this must be modified to say that this is true for exactly one quarter of these cases.

What happens is that Γ_1 has two viable descendants, each order 2^{11} and carrying mass $1/8$. One of these has two viable descendants, each of order 2^{14} and carrying mass $1/16$. Denote by Γ the first of these groups of order 2^{14} . Extensive experiments indicate that every pro-2 group with presentation of the form $\langle x, y | x^a = x^5, y^b = y^5 \rangle$ and an element of order two outside its commutator subgroup (playing the role of complex conjugation) and whose 2-class 4 quotient is Γ is infinite and has 2-class tower whose orders match the sequence A001461. Moreover, it appears that every descendant of Γ_1 which is not a descendant of Γ is finite.

The descendant tree of the second groups of order 2^{11} and 2^{14} can be pursued for quite a distance. The experiment above indicates that these subtrees will eventually terminate.

6 Conservation of mass

For a fixed prime p and positive integer g , consider O'Brien's rooted tree whose nodes are isomorphism classes of finite g -generator p -groups. We prune this by saving only those groups of a particular type (which fixes their abelianization), that could arise as a p -class quotient of a g -relator group (their p -multiplier and nuclear ranks differ by at most g), and that have nonzero mass.

It can happen that one such group of a given p -class is a quotient of another - for instance, if F is the free pro-2 group on 2 generators, then $F/P_2(F)$ of order 32 has $\mathbb{Z}/4\mathbb{Z} \times \mathbb{Z}/4\mathbb{Z}$ as a quotient but both are of p -class 2. This tends however to be rare, since if the difference between the p -multiplier rank and nuclear rank of Γ equals g , the same will be true of its immediate descendants and they will all have the same order, namely $|\Gamma|p^r$, where r is the nuclear rank of G [15]. This means that (again, given some hypothesis on transitivity of pure braid group actions) we have produced a probability distribution on the rooted tree with root Γ , by assigning the mass $A_{\mathbb{Z}}(\Gamma)/|\text{Aut } \Gamma|$ to each vertex Γ .

Two interesting questions arise. Is it possible to have a point mass, i.e. an infinite end along which the mass is bounded away from 0? Second, does the accumulated mass of FAb groups or ones without infinite p -adic analytic quotients amount to 100%? These both have positive answers with respect to the measure given in [4].

References

- [1] J. Achter. *Results of Cohen-Lenstra type for quadratic function fields*. In *Computational Arithmetic Geometry*, eds. K. Lauter and K. Ribet, Contemporary Mathematics (463), American Mathematical Society, 2008.
- [2] M. Bhargava. *Mass Formulae for Extensions of Local Fields, and Conjectures on the Density of Number Field Discriminants*. International Mathematics Research Notices (2007) Vol. 2007 : article ID rnm052, 20 pages, doi:10.1093/imrn/rnm052.
- [3] N. Boston. *Reducing the Fontaine-Mazur conjecture to group theory*. Proceedings of Thompson’s 70th Birthday Conference, in “Progress in Galois Theory”, (2005), 39–50.
- [4] N. Boston. *Random pro- p groups and random Galois groups*. Annales des Sciences mathématiques du Québec. 32, no 2.
- [5] N. Boston and C. R. Leedham-Green. *Explicit computation of Galois p -groups unramified at p* . J. Algebra 256 (2002), 402–413.
- [6] N. Boston and D. Perry. *Maximal 2-extensions with restricted ramification*. J. Algebra 232 (2000), 664–672.
- [7] M. Fried and H. Völklein. *The inverse Galois problem and rational points on moduli spaces*. Math. Annalen, vol. 290, no.1 (1991), 771–800.
- [8] E. Friedman and L. Washington. *On the distribution of divisor class groups of curves over a finite field*. Théorie des nombres (Quebec, PQ, 1987), de Gruyter, Berlin, 1989, pp. 227–239.
- [9] D. Garton. *Random matrices and the Cohen-Lenstra statistics for global fields with roots of unity*. In progress.
- [10] Y. Ihara. *Profinite braid groups, Galois representations, and complex multiplications*. Ann. of Math. 123 (1986), 43–106.
- [11] N. Katz and P. Sarnak. *Random matrices, Frobenius eigenvalues, and monodromy*. Colloquium Publications 45, American Mathematical Society, 1999.
- [12] H. Koch. *Galois theory of p -extensions*, with a foreword by I. R. Shafarevich, translated from the 1970 German original by Franz Lemmermeyer with a postscript by the author and Lemmermeyer. Springer Monographs in Mathematics. Springer-Verlag, Berlin, (2002).
- [13] T. Leinster. *The Euler characteristic of a category*. Doc. Math. 13 (2008), 21–49.
- [14] G. Malle. *Cohen-Lenstra heuristic and roots of unity*. J. Number Theory 128 (2008), 2823–2835.
- [15] H. Nover. *Computation of Galois groups of 2-class towers*. Ph.D. Thesis, University of Wisconsin, 2009.
- [16] E.A.O’Brien. *The p -group generation algorithm*. J. Symbolic Comput. 9 (1990), 677–698.
- [17] D. Roberts. *Wild partitions and number theory*. Article 07.6.6. Journal of Integer Sequences, Vol. 10 (2007).
- [18] P. Rozenhart. *Fast Tabulation of Cubic Function Fields*. Ph.D. Thesis, University of Calgary, 2009.