

Alex Bartel

School of Mathematics & Statistics
University of Glasgow
alex.bartel@glasgow.ac.uk

Research

The Cohen–Lenstra principle

In this article Alex Bartel reviews the revolutionary insights of Henri Cohen and Hendrik Lenstra from the 1980s, which spawned an entire ‘industry’, the work of Manjul Bhargava, which earned him the Fields Medal in 2014, some of the beautiful subsequent work of his school, and finally the recent spectacular breakthroughs of Alexander Smith, which began when he was a PhD student.

Class groups of quadratic number fields

Let us first set the stage, and we could do worse than begin our story with Carl Friedrich Gauss. In his famous treatise *Disquisitiones Arithmeticae*, which appeared in print in 1801, he developed a theory for treating the following question. Let $f(x, y) = ax^2 + bxy + cy^2$, where a, b, c are integers. Such a homogeneous degree 2 polynomial in two variables is called a *binary quadratic form* over the integers. Then which integers n are *represented* by f , meaning are values of $f(x, y)$ as x, y runs through the integers? For example if we have $a = c = 1$ and $b = 0$, then the question is asking which integers are sums of two squares. This special case had already been investigated by Fermat. Other special cases had been considered by Fermat, Legendre, Lagrange, and others, but it was Gauss who, at the age of 21, consolidated that progress into one coherent theory.

Gauss notes that if $f(x, y)$ is a binary quadratic form as above, then for every 2×2 matrix $\begin{pmatrix} a & b \\ c & d \end{pmatrix}$ with integer coefficients and with determinant 1, the binary quadratic form $\begin{pmatrix} a & b \\ c & d \end{pmatrix} f(x, y)$ defined by

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix} f(x, y) = f(ax + cy, bx + dy)$$

represents the exact same integers as f . Indeed, if x and y are integers, then so are $ax + cy$ and $bx + dy$; and since the matrix

has determinant 1, so that its inverse also has integer coefficients, we can apply the above argument with the matrix replaced by its inverse to show the converse: whenever $ax + cy$ and $bx + dy$ are integers, so are x and y . Thus, for example, everything we know about the integers represented by the binary quadratic form $f(x, y) = x^2 + y^2$ translates into knowledge about the quadratic form $g(x, y) = 2x^2 + 2xy + y^2$, because we have $g = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} f$.

The set of 2×2 integer matrices with determinant 1 forms a group under multiplication, denoted by $\mathrm{SL}_2(\mathbb{Z})$. Gauss defines two forms to be equivalent if one can be obtained from the other by applying a suitable matrix in $\mathrm{SL}_2(\mathbb{Z})$ as above. This is clearly an equivalence relation, and as we have just remarked, the set of integers that a quadratic form represents is really a property of the equivalence class of that form. This raises the next question: how can we tell whether two given binary quadratic forms are equivalent? In the above example of the two forms $f = x^2 + y^2$ and $g = 2x^2 + 2xy + y^2$, I convinced you of their equivalence by producing a suitable matrix out of my magician’s hat. What if they had not been equivalent? How much patience would I have had to rummage in that hat of mine before giving up?

Of course (dis)proof by hat-rummaging is not Gauss’s style. Instead Gauss attach-

es to a binary quadratic form a quantity that is invariant under the action of $\mathrm{SL}_2(\mathbb{Z})$: the *discriminant* of a binary quadratic form $ax^2 + bxy + cy^2$ is defined to be the quantity $b^2 - 4ac$. A straightforward calculation shows that equivalent forms have the same discriminant. For example the two forms from the previous paragraph both have discriminant -4 . If we are given two forms that have different discriminants, then we can save ourselves the hat-rummaging and can immediately conclude that they are not equivalent. Unfortunately (or perhaps fortunately), the converse is false, in general: there exist inequivalent forms with the same discriminant. However, we have the following foundational result.

Theorem. *For every integer d , the number of equivalence classes of binary quadratic forms of discriminant d is finite.*

This finite number is called the *class number* of d , traditionally denoted by $h(d)$. Table 1 shows class numbers for the first few negative so-called *fundamental discriminants*. For our purposes it does not matter what exactly a fundamental discriminant is, you may just replace that term by ‘square-free integer’ without erring too badly.

These class numbers all appear to be small, but this is an optical illusion, an instance of the ‘law of small numbers’. Gauss computed thousands of class numbers. How he did that is a separate beautiful story, for which we have no time on our purposeful journey to the Cohen–Lenstra heuristics. Based on those computations

d	−3	−4	−7	−8	−11	−15	−19	−20	
$h(d)$	1	1	1	1	1	2	1	2	...

Table 1

Gauss conjectured that as d tends to $-\infty$ through fundamental discriminants, the class numbers $h(d)$ tend to infinity. Moreover he conjectured a complete list of negative fundamental discriminants d for which one has $h(d) = 1$. This conjecture became known as one half of Gauss’s class number 1 problem, and its eventual full resolution over the course of the first half of the twentieth century is yet another beautiful and dramatic story for which we have no time.

Gauss also computed many class numbers of positive fundamental discriminants, and conjectured that, in sharp contrast to the negative ones, infinitely many of these class numbers are equal to 1. This second half of Gauss’s class number 1 problem remains one of the great open problems of algebraic number theory!

From the modern perspective the most important fact about the set of equivalence classes of binary quadratic forms of a given discriminant is that it carries the structure of an abelian group, the so-called *class group* of discriminant d , denoted by Cl_d . Already Gauss defined this group operation, the so-called *composition* of binary quadratic forms. However, today we understand it more conceptually in the context of a vast generalisation of Gauss’s class groups. To a modern number theorist the class group of discriminant d is an invariant attached to the *quadratic field* $\mathbb{Q}(\sqrt{d}) = \{a + b\sqrt{d} : a, b \in \mathbb{Q}\}$. More generally we attach a finite abelian group, the class group, to every *number field*, meaning a field that contains the field $\mathbb{Q}$ of rational numbers and that has finite dimension as a vector space over $\mathbb{Q}$. Gauss’s case is but the special case of dimension 2, but there are, for every $k \in \mathbb{Z}_{\geq 2}$, infinitely many number fields of dimension k over $\mathbb{Q}$, and each one of them has a class group.

Cohen–Lenstra–Martinet heuristics

To recapitulate: as d runs through negative fundamental discriminants, we have an infinite sequence of finite abelian groups, the class groups of the quadratic fields $\mathbb{Q}(\sqrt{d})$. The orders of these groups tend to infinity, but we could ask much finer statistical questions about them. For example, if p is a prime number, what is the average

size of their p -torsion subgroups $\text{Cl}_d[p]$, i.e. of the subgroups consisting of all elements I such that one has $pI = 0$? Do those orders also tend to infinity? Of course such an average might fail to exist even if the orders do not tend to infinity. In a similar vein, each Cl_d , being a finite abelian group, is a direct product over the distinct prime numbers p of subgroups $\text{Cl}_d[p^\infty]$ of order a power of p , the so-called *p -Sylow subgroups* of Cl_d . Thus, if one knows the isomorphism class of $\text{Cl}_d[p^\infty]$ for all prime numbers p , then one knows the isomorphism class of Cl_d . An example of a natural question is: how often are the p -Sylow subgroups of Cl_d , for a given prime number p , cyclic?

Already Gauss determined the order of the 2-torsion subgroup of Cl_d as a simple function of d . This is a beautiful theory in its own right, so-called *genus theory*, which has continued to inspire number theorists to this day. Rather deeper and much more recent is the following theorem [7], which is really where the modern strand of our story begins.

Theorem (Davenport and Heilbronn [7]). *As $X \rightarrow \infty$, the limit of*

$$\frac{\sum_{-X < d < 0} \#\text{Cl}_d[3]}{\sum_{-X < d < 0} 1}$$

exists and is equal to 2, where both sums run over negative fundamental discriminants.

To date, no such theorem is known where 3 is replaced by any other odd prime number. As for the question how often the Sylow subgroups of class groups of negative discriminants are cyclic, we do not know the answer, but numerically it seems like they are so very often. For example the 3-Sylow subgroups appear to be cyclic about 0.98 of the time — an observation that must have appeared absolutely baffling to mathematicians prior to the Cohen–Lenstra heuristic, for there seems to be no number theoretic reason for this preponderance. After all, it stands in sharp contrast to the paucity of cyclic groups among all isomorphism classes of abelian

groups of order a power of 3: for every $r \in \mathbb{Z}_{\geq 1}$, there is exactly one isomorphism class of cyclic groups of order 3^r , while the total number of isomorphism classes of abelian groups of that order is the partition number of r , which grows faster than any polynomial in r .

The major insight of Cohen and Lenstra was that the above statistical phenomena, far from pointing to some hidden number theoretic structure, actually mirror what one would observe if the p -Sylow subgroups of the class groups, for odd primes p , were entirely random! But what does a random abelian group look like?

The idea of Cohen and Lenstra was this: first suppose that you fix the order O of your group in advance, and you construct a random group of order O by filling in an $O \times O$ multiplication table at random — throw away your attempt and just try again if your multiplication table happens to not describe a group. Then, once you have obtained one realisation of a particular group, a simultaneous permutation of the rows and columns, which simply amounts to a relabelling of the elements, results in an isomorphic group. Thus, you would expect every group to be realised by $O!$ different multiplication tables. However, this is not quite right, since if a group has symmetries, i.e. automorphisms, then certain permutations of the rows and columns of its multiplication table will result in an *identical* multiplication table, rather than a different table describing an isomorphic group. More precisely, if one group, G_1 say, has t times more automorphisms than another, G_2 say, then there are t times fewer distinct multiplication tables realising a group isomorphic to G_1 than one isomorphic to G_2 . At this point Cohen and Lenstra made a leap of imagination: one can try to apply the same probabilistic reasoning to groups of *varying* orders. Let $\text{Aut } G$ denote the group of automorphisms of a group G . If we fix a prime number p , then it turns out that the sum $\sum_A 1/\#\text{Aut } A$ over a full set of representatives A of isomorphism classes of finite abelian p -groups (meaning abelian groups of order a power of p) converges, to c_p , say. Therefore one can define a probability distribution on that set of isomorphism class representatives. Such a random abelian p -group is isomorphic to a given group A with probability $c_p^{-1}/\#\text{Aut } A$. This is an instance of what I refer to in the title as the Cohen–Lenstra principle.

Heuristic (The Cohen–Lenstra principle). *If a concrete instance of an algebraic object is constructed at random, then it will be isomorphic to a given object X with probability inversely proportional to $\#\text{Aut } X$.*

A visually intuitive example of this principle, which I urge the reader to work through, is that of random graphs: start with three vertices, and for each pair of vertices flip a coin to determine whether or not to draw an edge between them. Now compute the probability of obtaining the complete graph, and compare it with the probability of obtaining a graph with exactly one edge. Explain the difference in the terms just outlined.

Cohen and Lenstra conjectured [5] that for odd primes p , the p -Sylow subgroups $\text{Cl}_d[p^\infty]$ of the class groups for $d < 0$ ‘look random’ in the above sense.

Conjecture (Cohen and Lenstra [5]). *Let p be an odd prime number, and let A be a finite abelian p -group. Then the limit*

$$\lim_{X \rightarrow \infty} \frac{\sum_{-X < d < 0} \mathbf{1}(\text{Cl}_d[p^\infty] \cong A)}{\sum_{-X < d < 0} 1}$$

exists and is equal to $c_p^{-1}/\#\text{Aut } A$, where both sums run over the negative fundamental discriminants and $\mathbf{1}$ denotes the characteristic function, which takes value 1 if the condition is satisfied and 0 otherwise.

In fact, Cohen and Lenstra made a similar prediction not just for characteristic functions, but for any ‘reasonable’ $\mathbb{C}$ -valued function on the class of isomorphism classes of finite abelian p -groups: the average of such a function over the set $\text{Cl}_d[p^\infty]$ with d running over the negative fundamental discriminants should be equal to its expected value with respect to the probability distribution we have just discussed. The meaning of the word ‘reasonable’ was not further specified in the original paper. This indeterminacy is the reason why the term ‘Cohen–Lenstra heuristic’ is usually used, rather than ‘conjecture’. See [3], however, for several proposals of what ‘reasonable function’ might mean.

In the same paper Cohen and Lenstra formulated a conjecture for class groups of positive fundamental discriminants. Somewhat distressingly (to the authors,

not just to us), that conjecture departs from what we have been calling the Cohen–Lenstra principle, and assigns to a finite abelian p -group A the probability weight $1/(\#\text{Aut } A \cdot \#A)$.

In [6] the Cohen–Lenstra heuristic was generalised by Cohen–Martinet from quadratic fields to very general families of number fields. We will not state the Cohen–Martinet heuristic, but will instead make some remarks. Firstly, if F is a number field, then the group of its automorphisms acts on the class group Cl_F , and Cohen and Martinet realised that one should model the class group not just as a mere group, but a group together with those specific symmetries, i.e. as a module over a suitable ring. Accordingly, one should expect the probability weights to involve not all group automorphisms but only those respecting the special symmetries. However, the actual weights in the Cohen–Martinet heuristic are not merely inversely proportional to sizes of symmetry-respecting automorphism groups of the class groups, but are even more complicated than in the positive d case of the original Cohen–Lenstra heuristic. A second additional issue that arises in this generalisation is the question how to order families of general number fields, since they are no longer just parametrised by square-free integers d . Cohen and Martinet enumerated their fields by absolute value of the discriminant, the by-far most often used invariant for enumerating number fields for over a century. In fact, their heuristic also generalises the Cohen–Lenstra heuristic in a different direction: in place of the base field $\mathbb{Q}$ they take an arbitrary number field, and predict the behaviour of so-called *relative class groups* in families of extensions of that number field. Finally, the observant reader will have noticed that we said very little about the 2-Sylow subgroups of the class groups of quadratic fields. There are multiple reasons for this, the most straight forward being that already Gauss knew that the 2-torsion subgroups of class groups of quadratic fields do not ‘look random’. We shall return to this point later. Analogously, Cohen and Martinet also exclude the p -Sylow subgroup for certain ‘bad’ primes p from their heuristic. Exactly which primes p are bad and which ones are good in any given family of number fields was already a subject of some speculation in [6]. Today the question of what is going

on at the bad primes is among those at the frontier of current research on the Cohen–Lenstra–Martinet heuristics, and we shall return to it.

Reformulation of Friedman–Washington

Notice that in our justification of the Cohen–Lenstra heuristic via the model of random multiplication tables, the fact that class groups are always abelian played no role. Here is an alternative procedure for producing a random *abelian* group: pick a large integer r , start with the free abelian group on r generators, and quotient out r random relations. In other words, let the random group be the cokernel of a random element of the group $M_r(\mathbb{Z})$ of $r \times r$ integer matrices. Actually, this does not work so well, because the group $M_r(\mathbb{Z})$ is not compact, and carries no suitable probability distribution, so instead one replaces $\mathbb{Z}$ with the ring $\mathbb{Z}_p$ of p -adic integers for a chosen prime p . We do not expect any familiarity with p -adic integers: the reader is welcome to think of $M_r(\mathbb{Z})$, but imagine that we are only picking out the p -part of the cokernel. The cokernel of an element of $M_r(\mathbb{Z}_p)$ is always a pro- p group — just think ‘made up of finite p -groups’. Moreover, the group $M_r(\mathbb{Z}_p)$ has a Haar measure, and the cokernel of a Haar-random matrix is finite with probability 1. This procedure defines a discrete probability distribution $\mathcal{P}_r$ on the class of isomorphism classes of finite abelian p -groups. Of course for any fixed r , most finite abelian p -groups cannot be cokernels or such a matrix, only those that can be generated by r elements can occur. However Friedman and Washington [10], guided by an analogy between number fields and function fields over finite fields, have discovered that as $r \rightarrow \infty$, the sequence $\mathcal{P}_r$ converges in distribution, and the limiting distribution is the Cohen–Lenstra distribution! Even more surprisingly, Matchett Wood showed in [18] that instead of Haar measure, one could take almost *any* distribution in which the entries of the matrices are independent of each other, as long as one avoids an obviously bad situation of these entries falling into fixed congruence classes modulo p too often.

A sloppy but easily memorisable way of summarising these findings is that random abelian groups obtained from random multiplication tables also look like groups with many commuting generators subject to equally many random relations.



Photo: Archives of the Mathematisches Forschungsinstitut Oberwolfach

Hendrik Lenstra and Henri Cohen in 2005 at the Oberwolfach workshop on ‘Explicit Methods in Number Theory’

One can even recover the more complicated Cohen–Lenstra distribution for positive discriminants in this random matrix model: those groups look like groups with many commuting generators subject to ‘many + 1’ random relations.

The analogy between number fields and function fields is extremely fruitful in the area of arithmetic statistics, and we shall return to it.

Subsequent developments

The Cohen–Lenstra–Martinet heuristics have spawned so much exciting activity that it will be impossible to mention, let alone describe all of it, and I apologise to anyone whose work could have been mentioned here but has been omitted.

Already in 1987, several years before the Cohen–Martinet extension of the Cohen–Lenstra heuristic, Gerth III attacked the problem of bad primes [11]. In the case of quadratic fields the prime 2 is bad, because thanks to Gauss’s genus theory we understand a part of $\text{Cl}_d[2^\infty]$, namely $\text{Cl}_d[2]$ (the reader may object that this is an idiosyncratic understanding of the word ‘bad’). Gerth III consequently conjectured, based on an actual theorem that we will not state here, that the groups $\text{Cl}_d[2^\infty]/\text{Cl}_d[2]$ do behave like Cohen–Lenstra–random groups.

There is a generalisation of genus theory to arbitrary number fields, which explains a certain ‘piece’ of the class group as being not random looking. The prime numbers that this phenomenon makes bad are those dividing the $\mathbb{Q}$ -dimension of the so-called Galois closure of the field. Gerth conjectured that if that dimension

is a prime number, then ‘everything apart from the genus piece’ behaves according to the Cohen–Lenstra heuristic. To incorporate genus theory into the Cohen–Lenstra heuristic for general number fields is still an open problem.

In [8] Ellenberg, Venkatesh, and Westerland pioneered an ingenious topological method, hinging on a homological stability result for so-called *Hurwitz spaces*, for proving theorems on class groups of function fields over finite fields. This general method has since then been successfully applied for proving more results in the area, some of which we will mention below.

Malle, pursuing a suggestion of Lenstra, noticed [14] through extensive numerical experimentation that if a prime number p divides the order of the group of roots of unity in the base field, then the behaviour of the p -Sylow subgroups of the class groups of the extensions of that base field seems to deviate from the Cohen–Lenstra–Martinet predictions. He conjectured that this is not a mere artifact of unreliable numerical data, but rather that the model needs to be adjusted in those cases. Malle did not offer a heuristic, but proposed alternative formulae in some situations, which exhibited much better agreement with the data. It then became an urgent problem to formulate a general conjecture, beyond the special cases treated by Malle, and ideally also one that would explain the deviation in a structural manner, rather than just quantify it. This was achieved by Lipnowski, Sawin, and Tsimerman [12] in the case of quadratic extensions: they identified a very subtle additional du-

ality on the relevant pieces of the class groups, accounted for by roots of unity in the base field, formulated a model for a random piece of data consisting of a group with this additional duality structure, and proved a result towards their conjecture in the function field setting, employing the Ellenberg–Venkatesh–Westerland machine. More recently, Matchett Wood and Sawin formulated a comprehensive conjecture in great generality [15]. Like Malle, they did not offer a heuristic model, but instead the conjecture is suggested by a theorem in the function field setting: using work of Liu, Matchett Wood and Zureick-Brown [13] on a non-abelian generalisation of the Cohen–Lenstra heuristic, in which once again the Hurwitz spaces method played a crucial role, Matchett Wood and Sawin computed the so-called moments of a probability distribution that occurs in the function field setting, proved that these moments determine a unique distribution, and conjectured that that distribution also governs class groups of number fields in the presence of roots of unity in the base field. Thus, in the span of only a few years our understanding of this particular type of ‘bad prime’ has made enormous leaps, but is still not complete: we are still lacking a model, generalising the Lipnowski–Sawin–Tsimerman model from the case of quadratic extensions, that would explain the Matchett Wood–Sawin probability weights.

So far we have focussed on conjectures in the Cohen–Lenstra–Martinet setting of class groups of number fields, and have alluded to theorems in the function field setting. But what do we actually *know* about the statistical properties of class groups of number fields?

We have mentioned Gauss’s genus theory results on the 2-torsion of class groups of quadratic fields and the Davenport–Heilbronn theorem on the 3-torsion of quadratic fields. Both have been generalised in various directions, the outcome usually being the determination of the average of some ‘reasonable’ function on the class of isomorphism classes of abelian groups, when evaluated on a natural sequence of class groups of number fields. All these generalisations only see the n -torsion of class groups for some fixed n . Most notably, the geometry-of-numbers technique used by Davenport and Heilbronn has seen a huge revival and extension thanks to the Fields Medal

winning work of Bhargava, who determined, for example, the average size of 2-torsion of class groups of cubic fields [4], but also applied those techniques to counting number fields themselves and, most famously, to bounding average ranks of elliptic curves. The school that Bhargava built up took those techniques further still, and has considered many beautiful generalisations of the original Cohen–Lenstra–Martinet question, which would merit their own survey. Meanwhile, Fouvry and Klüners confirmed in [9] Gerth’s modification of the Cohen–Lenstra heuristic for the 4-torsion of class groups of quadratic fields.

Then in 2017 Alexander Smith, at the time a PhD student at Harvard, in a pre-print that has since then been substantially revised and extended [16, 17], made a gigantic leap forward: he determined the full distribution of the 2-Sylow subgroups of class groups of quadratic fields, thus confirming Gerth’s modification for the entire 2-Sylow subgroups of those class groups! In the process he introduced a host of new ideas that have since been extended to other settings, and will, no doubt, remain of central importance in the area for the foreseeable future.

In the meantime, in [3] Lenstra and I showed that there are infinitely many families of number fields and infinitely many reasonable functions such that the function has an average over the corresponding family of class groups and that that average can be computed to any desired precision

in finite time. Apart from Smith’s result, this was the first instance of infinitely many reasonable functions for which we knew the average on class groups to exist. However, rather than confirming the Cohen–Lenstra–Martinet conjectures this way, we exhibited a specific function whose average is *not* equal to the predicted expected value, thus disproving the conjectures! It turns out that the discriminant, despite being by the far the most common invariant for enumerating number fields, is deficient in a particular way. We proposed a correction to the heuristics, which amounts to enumerating number fields in certain alternative ways. In the same work, however, we found a second, entirely unrelated problem with the heuristics. We have mentioned above that for general number fields F the class group should be understood not merely as a group, but as a group together with an action of $\text{Aut } F$. Lenstra and I showed that this structure is quite subtle, and that certain such modules, despite having positive probability in the Cohen–Lenstra–Martinet model, can actually never occur as class groups. Finding a satisfactory correction to this defect was harder than to the first one just mentioned, and we proposed one together with Johnston in [1].

To bring this story full circle, recall that as we observed above, the Cohen–Lenstra principle does not seem to apply to class groups in any situation other than that of quadratic fields with negative discriminant. It turns out that this is because the

class group is the wrong object to apply it to! Lenstra and I showed in [3] that if one instead considers the so-called *Arakelov class group* Pic_F^0 of a number field F , a more complicated invariant that ‘knows’ everything about the class group, then the postulate that Pic_F^0 is isomorphic to a suitable object X with probability inversely proportional to $\#\text{Aut } X$ in fact implies the Cohen–Lenstra–Martinet heuristics, in particular recovering the otherwise mysterious looking probability weights for the class groups. The big caveat is that Arakelov class groups of number fields typically have infinitely many automorphisms, so quite a bit of work [2] goes into turning the previous sentence from non-sense into rigorous mathematics.

The scope of this article has only allowed us to give a very cursory overview of the huge amount of activity around the Cohen–Lenstra heuristic and its generalisations. We have not said much about the rich connections to elliptic curves, the parallels with low-dimensional manifolds and their fundamental groups, nor about the fascinating non-abelian generalisations, which are, I would say, enough to elevate the Friedman–Washington result on cokernels of random matrices to a general principle, a sibling of the Cohen–Lenstra principle. But I hope that what we did say has been enough to convince the reader of the current verity of the discipline that Cohen and Lenstra birthed and of the exciting future that awaits it. 

References

- Alex Bartel, Henri Johnston and Hendrik W. Lenstra, Jr., Arakelov class groups of random number fields, *Math. Ann.* (2024), to appear.
- Alex Bartel and Hendrik W. Lenstra, Jr., Comensurability of automorphism groups, *Compos. Math.* 153(2) (2017), 323–346.
- Alex Bartel and Hendrik W. Lenstra, Jr., On class groups of random number fields, *Proc. Lond. Math. Soc.* (3) 121(4) (2020), 927–953.
- Manjul Bhargava, The density of discriminants of quartic rings and fields, *Ann. of Math.* 162 (2005), 1031–1063.
- H. Cohen and H. W. Lenstra, Jr., Heuristics on class groups of number fields, in *Number Theory, Noordwijkerhout 1983*, Lecture Notes in Math., Vol. 1068, Springer, 1984, pp. 33–62.
- Henri Cohen and Jacques Martinet, Étude heuristique des groupes de classes des corps de nombres, *J. Reine Angew. Math.* 404 (1990), 39–76.
- H. Davenport and H. Heilbronn, On the density of discriminants of cubic fields. II, *Proc. Roy. Soc. London Ser. A* 322(1551) (1971), 405–420.
- Jordan S. Ellenberg, Akshay Venkatesh and Craig Westerland, Homological stability for Hurwitz spaces and the Cohen–Lenstra conjecture over function fields, *Ann. of Math.* (2) 183(3) (2016), 729–786.
- Étienne Fouvry and Jürgen Klüners, On the 4-rank of class groups of quadratic number fields, *Invent. Math.* 167(3) (2007), 455–513.
- Eduardo Friedman and Lawrence C. Washington, On the distribution of divisor class groups of curves over a finite field, in *Théorie des nombres (Quebec, PQ, 1987)*, De Gruyter, 1989, pp. 227–239.
- Frank Gerth, III, Densities for ranks of certain parts of p -class groups, *Proc. Amer. Math. Soc.* 99(1) (1987), 1–8.
- Michael Lipnowski, Will Sawin and Jacob Tsimerman, Cohen–Lenstra heuristics and bilinear pairings in the presence of roots of unity (2020), arXiv:2007.12533 [math.NT].
- Yuan Liu, Melanie Matchett Wood and David Zureick-Brown, A predicted distribution for galois groups of maximal unramified extensions, *Invent. Math.* (2024), to appear.
- Gunter Malle, Cohen–Lenstra heuristic and roots of unity, *J. Number Theory* 128(10) (2008), 2823–2835.
- Will Sawin and Melanie Matchett Wood, Conjectures for distributions of class groups of extensions of number fields containing roots of unity (2023), arXiv:2301.00791 [math.NT].
- Alexander Smith, The distribution of ℓ^∞ -selmer groups in degree ℓ twist families I (2022), arXiv:2207.05674 [math.NT].
- Alexander Smith, The distribution of ℓ^∞ -selmer groups in degree ℓ twist families II (2023), arXiv:2207.05143 [math.NT].
- Melanie Matchett Wood, Random integral matrices and the Cohen–Lenstra heuristics, *Amer. J. Math.* 141(2) (2019), 383–398.