Mod-2 dihedral Galois representations of prime conductor

Kiran S. Kedlaya and Anna Medvedovsky

Kedlaya: Department of Mathematics, University of California, San Diego
kedlaya@ucsd.edu, http://kskedlaya.org/slides/
Medvedovsky: Max-Planck-Institut für Mathematik, Bonn
medved@gmail.com

Algorithmic Number Theory Symposium (ANTS-XIII)
University of Wisconsin, Madison
July 16, 2018

Kedlaya was supported by NSF (grant DMS-1501214) and UC San Diego (Warschawski Professorship). Medvedovsky was supported by an NSF postdoctoral research fellowship (grant DMS-1703834) and has gratefully enjoyed the hospitality of MPIM.

Motivation: Cremona's tables of rational elliptic curves

Over a period of more than two decades, Cremona has tabulated all† elliptic curves over $\mathbb{Q}$ of conductor up to 400000. This table can be accessed in several ways, including the LMFDB (L-Functions and Modular Forms Database; http://www.lmfdb.org).

The rate-limiting step in this computation for a given conductor N is: given the matrix of action of T_p on some basis of $S_2(\Gamma_0(N), \mathbb{Q})$, where p is the smallest prime not dividing N, compute the kernel of $T_p - a_p$ for each integer a_p with $|a_p| \leq 2\sqrt{p}$.

†Initially, Cremona assumed that all elliptic curves over $\mathbb{Q}$ are modular. By 2001, this was known by work of Wiles, Taylor–Wiles, and Breuil–Conrad–Diamond–Taylor.

Motivation: Cremona's tables of rational elliptic curves

Over a period of more than two decades, Cremona has tabulated all† elliptic curves over $\mathbb{Q}$ of conductor up to 400000. This table can be accessed in several ways, including the LMFDB (L-Functions and Modular Forms Database; http://www.lmfdb.org).

The rate-limiting step in this computation for a given conductor N is: given the matrix of action of T_p on some basis of $S_2(\Gamma_0(N), \mathbb{Q})$, where p is the smallest prime not dividing N, compute the kernel of $T_p - a_p$ for each integer a_p with $|a_p| \leq 2\sqrt{p}$.

†Initially, Cremona assumed that all elliptic curves over $\mathbb{Q}$ are modular. By 2001, this was known by work of Wiles, Taylor–Wiles, and Breuil–Conrad–Diamond–Taylor.

Sparse integer linear algebra

One may assume that the matrix of T_p is integral with sparse small entries. For instance, if N is an odd prime (so that $p = 2$), using either the Masser–Oesterlé method of graphs or Birch's method of ternary forms, one gets a matrix with at most three nonzero entries per row, each of absolute value at most 3.

For such matrices, one generically does row reduction by the *multimodular* approach of working modulo a collection of small primes. Cremona prefers to work modulo one word-sized prime.

However, for most N and a_p, the kernel of $T_p - a_p$ is zero. Can one carry out an effective early abort using linear algebra modulo one small prime?

Aside: this strategy is potentially more useful if the matrix of T_p comes separated by Atkin–Lehner eigenspaces and with newforms removed. Birch's method (mostly) does both automatically.

Sparse integer linear algebra

One may assume that the matrix of T_p is integral with sparse small entries. For instance, if N is an odd prime (so that $p = 2$), using either the Masser–Oesterlé method of graphs or Birch's method of ternary forms, one gets a matrix with at most three nonzero entries per row, each of absolute value at most 3.

For such matrices, one generically does row reduction by the *multimodular* approach of working modulo a collection of small primes. Cremona prefers to work modulo one word-sized prime.

However, for most N and a_p, the kernel of $T_p - a_p$ is zero. Can one carry out an effective early abort using linear algebra modulo one small prime?

Aside: this strategy is potentially more useful if the matrix of T_p comes separated by Atkin–Lehner eigenspaces and with newforms removed. Birch's method (mostly) does both automatically.

Sparse integer linear algebra

One may assume that the matrix of T_p is integral with sparse small entries. For instance, if N is an odd prime (so that $p = 2$), using either the Masser–Oesterlé method of graphs or Birch's method of ternary forms, one gets a matrix with at most three nonzero entries per row, each of absolute value at most 3.

For such matrices, one generically does row reduction by the *multimodular* approach of working modulo a collection of small primes. Cremona prefers to work modulo one word-sized prime.

However, for most N and a_p, the kernel of $T_p - a_p$ is zero. Can one carry out an effective early abort using linear algebra modulo one small prime?

Aside: this strategy is potentially more useful if the matrix of T_p comes separated by Atkin–Lehner eigenspaces and with newforms removed. Birch's method (mostly) does both automatically.

Sparse integer linear algebra

One may assume that the matrix of T_p is integral with sparse small entries. For instance, if N is an odd prime (so that $p = 2$), using either the Masser–Oesterlé method of graphs or Birch's method of ternary forms, one gets a matrix with at most three nonzero entries per row, each of absolute value at most 3.

For such matrices, one generically does row reduction by the *multimodular* approach of working modulo a collection of small primes. Cremona prefers to work modulo one word-sized prime.

However, for most N and a_p, the kernel of $T_p - a_p$ is zero. Can one carry out an effective early abort using linear algebra modulo one small prime?

Aside: this strategy is potentially more useful if the matrix of T_p comes separated by Atkin–Lehner eigenspaces and with newforms removed. Birch's method (mostly) does both automatically.

A computational experiment...

In order to assess this idea, we tried the following experiment: for every odd prime $N < 500000$, we computed the matrix of action of T_2 on $S_2(\Gamma_0(N), \mathbb{Q})$, reduced mod 2, and tested whether 0 and 1 occur as eigenvalues of the resulting matrix.

Ignoring some sporadic cases with $N \leq 163$, the eigenvalues 0 and 1 occur with the following frequencies:

$N \pmod 8$	1	3	5	7
frequency of eigenvalue 0	16.8%	always	42.2%	17.3%
frequency of eigenvalue 1	always	always	always	47.9%

It would be useful to repeat this experiment with other N, other T_p, reduction modulo other ℓ, possibly even weights $k > 2$...

A computational experiment...

In order to assess this idea, we tried the following experiment: for every odd prime $N < 500000$, we computed the matrix of action of T_2 on $S_2(\Gamma_0(N), \mathbb{Q})$, reduced mod 2, and tested whether 0 and 1 occur as eigenvalues of the resulting matrix.

Ignoring some sporadic cases with $N \leq 163$, the eigenvalues 0 and 1 occur with the following frequencies:

$N \pmod 8$	1	3	5	7
frequency of eigenvalue 0	16.8%	always	42.2%	17.3%
frequency of eigenvalue 1	always	always	always	47.9%

It would be useful to repeat this experiment with other N, other T_p, reduction modulo other ℓ, possibly even weights $k > 2$...

A computational experiment...

In order to assess this idea, we tried the following experiment: for every odd prime $N < 500000$, we computed the matrix of action of T_2 on $S_2(\Gamma_0(N), \mathbb{Q})$, reduced mod 2, and tested whether 0 and 1 occur as eigenvalues of the resulting matrix.

Ignoring some sporadic cases with $N \leq 163$, the eigenvalues 0 and 1 occur with the following frequencies:

$N \pmod 8$	1	3	5	7
frequency of eigenvalue 0	16.8%	always	42.2%	17.3%
frequency of eigenvalue 1	always	always	always	47.9%

It would be useful to repeat this experiment with other N, other T_p, reduction modulo other ℓ, possibly even weights $k > 2$...

… and a different direction

… instead, we try to explain this data in terms of Galois representations. Hereafter, let N be an odd prime and $\mathfrak{m}$ a maximal ideal of $\mathbb{T}_2(N)^\dagger$ containing 2. Note that 0 (resp. 1) occurs as an eigenvalue of the mod-2 reduction of T_2 iff there exists an $\mathfrak{m}$ with $a_2(\mathfrak{m}) = 0$ (resp. $a_2(\mathfrak{m}) = 1$).

We classify $\mathfrak{m}$ based on the projective image of the corresponding modular mod-2 Galois representation $\rho : G_\mathbb{Q} \to \mathrm{GL}_2(\bar{\mathbb{F}}_2)$:

- *reducible*,
- *dihedral*,
- *exceptional* (A_4, S_4, A_5), or
- *big-image* ($\mathrm{PSL}_2(\mathbb{F}_q)$) or larger, excluding previous cases for small q).

For rational newforms, only reducible and dihedral cases occur because $\mathrm{SL}_2(\mathbb{F}_2) \cong D_3$. We analyze these cases thoroughly; this explains the "always" entries in the table, but only partly explains the other frequencies.

† This Hecke algebra omits T_2; write $a_2(\mathfrak{m})$ for the eigenvalue at 2.

… and a different direction

… instead, we try to explain this data in terms of Galois representations. Hereafter, let N be an odd prime and $\mathfrak{m}$ a maximal ideal of $\mathbb{T}_2(N)^\dagger$ containing 2. Note that 0 (resp. 1) occurs as an eigenvalue of the mod-2 reduction of T_2 iff there exists an $\mathfrak{m}$ with $a_2(\mathfrak{m}) = 0$ (resp. $a_2(\mathfrak{m}) = 1$).

We classify $\mathfrak{m}$ based on the projective image of the corresponding modular mod-2 Galois representation $\rho : G_{\mathbb{Q}} \to \mathrm{GL}_2(\bar{\mathbb{F}}_2)$:

- *reducible*,
- *dihedral*,
- *exceptional* (A_4, S_4, A_5), or
- *big-image* ($\mathrm{PSL}_2(\mathbb{F}_q)$ or larger, excluding previous cases for small q).

For rational newforms, only reducible and dihedral cases occur because $\mathrm{SL}_2(\mathbb{F}_2) \cong D_3$. We analyze these cases thoroughly; this explains the "always" entries in the table, but only partly explains the other frequencies.

†This Hecke algebra omits T_2; write $a_2(\mathfrak{m})$ for the eigenvalue at 2.

… and a different direction

… instead, we try to explain this data in terms of Galois representations. Hereafter, let N be an odd prime and $\mathfrak{m}$ a maximal ideal of $\mathbb{T}_2(N)^\dagger$ containing 2. Note that 0 (resp. 1) occurs as an eigenvalue of the mod-2 reduction of T_2 iff there exists an $\mathfrak{m}$ with $a_2(\mathfrak{m}) = 0$ (resp. $a_2(\mathfrak{m}) = 1$).

We classify $\mathfrak{m}$ based on the projective image of the corresponding modular mod-2 Galois representation $\rho : G_{\mathbb{Q}} \to \mathrm{GL}_2(\bar{\mathbb{F}}_2)$:

- *reducible*,
- *dihedral*,
- *exceptional* (A_4, S_4, A_5), or
- *big-image* ($\mathrm{PSL}_2(\mathbb{F}_q)$ or larger, excluding previous cases for small q).

For rational newforms, only reducible and dihedral cases occur because $\mathrm{SL}_2(\mathbb{F}_2) \cong D_3$. We analyze these cases thoroughly; this explains the "always" entries in the table, but only partly explains the other frequencies.

†This Hecke algebra omits T_2; write $a_2(\mathfrak{m})$ for the eigenvalue at 2.

First steps

A reducible $\mathfrak{m}$ occurs iff 2 is an Eisenstein prime for N. By Mazur, this occurs iff 2 divides the numerator of $\frac{N-1}{12}$, yielding:

Lemma

If $N \equiv 1 \pmod 8$, then there is exactly one reducible $\mathfrak{m}$, for which $a_2(\mathfrak{m}) = 1$. Otherwise, there are no reducible $\mathfrak{m}$.

For $\mathfrak{m}$ dihedral, ρ has kernel G_L where $L/\mathbb{Q}$ is a D_3-extension. For $K/\mathbb{Q}$ the quadratic subfield of L, we say that $\mathfrak{m}$ is *K-dihedral*.

Lemma

If $\mathfrak{m}$ is dihedral, then it is either $\mathbb{Q}(\sqrt{N})$-dihedral or $\mathbb{Q}(\sqrt{-N})$-dihedral.

We say $\mathfrak{m}$ is *ordinary* if $a_2(\mathfrak{m}) \neq 0$ and *supersingular* if $a_2(\mathfrak{m}) = 0$; for an elliptic curve, this corresponds to having ordinary or supersingular reduction at 2.

First steps

A reducible $\mathfrak{m}$ occurs iff 2 is an Eisenstein prime for N. By Mazur, this occurs iff 2 divides the numerator of $\frac{N-1}{12}$, yielding:

Lemma

If $N \equiv 1 \pmod 8$, then there is exactly one reducible $\mathfrak{m}$, for which $a_2(\mathfrak{m}) = 1$. Otherwise, there are no reducible $\mathfrak{m}$.

For $\mathfrak{m}$ dihedral, ρ has kernel G_L where $L/\mathbb{Q}$ is a D_3-extension. For $K/\mathbb{Q}$ the quadratic subfield of L, we say that $\mathfrak{m}$ is *K-dihedral*.

Lemma

If $\mathfrak{m}$ is dihedral, then it is either $\mathbb{Q}(\sqrt{N})$-dihedral or $\mathbb{Q}(\sqrt{-N})$-dihedral.

We say $\mathfrak{m}$ is *ordinary* if $a_2(\mathfrak{m}) \neq 0$ and *supersingular* if $a_2(\mathfrak{m}) = 0$; for an elliptic curve, this corresponds to having ordinary or supersingular reduction at 2.

First steps

A reducible $\mathfrak{m}$ occurs iff 2 is an Eisenstein prime for N. By Mazur, this occurs iff 2 divides the numerator of $\frac{N-1}{12}$, yielding:

Lemma

If $N \equiv 1 \pmod 8$, then there is exactly one reducible $\mathfrak{m}$, for which $a_2(\mathfrak{m}) = 1$. Otherwise, there are no reducible $\mathfrak{m}$.

For $\mathfrak{m}$ dihedral, ρ has kernel G_L where $L/\mathbb{Q}$ is a D_3-extension. For $K/\mathbb{Q}$ the quadratic subfield of L, we say that $\mathfrak{m}$ is *K-dihedral*.

Lemma

If $\mathfrak{m}$ is dihedral, then it is either $\mathbb{Q}(\sqrt{N})$-dihedral or $\mathbb{Q}(\sqrt{-N})$-dihedral.

We say $\mathfrak{m}$ is *ordinary* if $a_2(\mathfrak{m}) \neq 0$ and *supersingular* if $a_2(\mathfrak{m}) = 0$; for an elliptic curve, this corresponds to having ordinary or supersingular reduction at 2.

First steps

A reducible $\mathfrak{m}$ occurs iff 2 is an Eisenstein prime for N. By Mazur, this occurs iff 2 divides the numerator of $\frac{N-1}{12}$, yielding:

Lemma

If $N \equiv 1 \pmod 8$, then there is exactly one reducible $\mathfrak{m}$, for which $a_2(\mathfrak{m}) = 1$. Otherwise, there are no reducible $\mathfrak{m}$.

For $\mathfrak{m}$ dihedral, ρ has kernel G_L where $L/\mathbb{Q}$ is a D_3-extension. For $K/\mathbb{Q}$ the quadratic subfield of L, we say that $\mathfrak{m}$ is *K-dihedral*.

Lemma

If $\mathfrak{m}$ is dihedral, then it is either $\mathbb{Q}(\sqrt{N})$-dihedral or $\mathbb{Q}(\sqrt{-N})$-dihedral.

We say $\mathfrak{m}$ is *ordinary* if $a_2(\mathfrak{m}) \neq 0$ and *supersingular* if $a_2(\mathfrak{m}) = 0$; for an elliptic curve, this corresponds to having ordinary or supersingular reduction at 2.

First steps

A reducible $\mathfrak{m}$ occurs iff 2 is an Eisenstein prime for N. By Mazur, this occurs iff 2 divides the numerator of $\frac{N-1}{12}$, yielding:

Lemma

If $N \equiv 1 \pmod 8$, then there is exactly one reducible $\mathfrak{m}$, for which $a_2(\mathfrak{m}) = 1$. Otherwise, there are no reducible $\mathfrak{m}$.

For $\mathfrak{m}$ dihedral, ρ has kernel G_L where $L/\mathbb{Q}$ is a D_3-extension. For $K/\mathbb{Q}$ the quadratic subfield of L, we say that $\mathfrak{m}$ is *K-dihedral*.

Lemma

If $\mathfrak{m}$ is dihedral, then it is either $\mathbb{Q}(\sqrt{N})$-dihedral or $\mathbb{Q}(\sqrt{-N})$-dihedral.

We say $\mathfrak{m}$ is *ordinary* if $a_2(\mathfrak{m}) \neq 0$ and *supersingular* if $a_2(\mathfrak{m}) = 0$; for an elliptic curve, this corresponds to having ordinary or supersingular reduction at 2.

Dihedral maximal ideals

Let K be one of $\mathbb{Q}(\sqrt{N})$ or $\mathbb{Q}(\sqrt{-N})$. Assume also that $N > 3$.

Theorem

(a) *The number of ordinary K-dihedral $\mathfrak{m}$ is $\frac{h(K)^{\mathrm{odd}}-1}{2}$, where $h(K)$ is the order of the class group and odd denotes the odd part.*

(b) *Of these, the number with $a_2(\mathfrak{m}) = 1$ is $\frac{h(K)^{\mathrm{odd,2\text{-}split}}-1}{2}$, where 2-split means divide by the subgroup generated by an ideal above 2.*

Theorem

(a) *If $N \equiv 3 \pmod 8$ and $K = \mathbb{Q}(\sqrt{-N})$, then there are exactly $h(K)$ supersingular K-dihedral $\mathfrak{m}$.*

(b) *If $N \equiv 5 \pmod 8$, $K = \mathbb{Q}(\sqrt{N})$, and the fundamental unit of K is $\equiv 1 \pmod{2\mathcal{O}_K}$, then there are $h(K)$ supersingular K-dihedral $\mathfrak{m}$.*

(c) *In all other cases, no supersingular K-dihedral maximal ideals exist.*

Dihedral maximal ideals

Let K be one of $\mathbb{Q}(\sqrt{N})$ or $\mathbb{Q}(\sqrt{-N})$. Assume also that $N > 3$.

Theorem

(a) *The number of ordinary K-dihedral $\mathfrak{m}$ is $\frac{h(K)^{\mathrm{odd}}-1}{2}$, where $h(K)$ is the order of the class group and odd denotes the odd part.*

(b) *Of these, the number with $a_2(\mathfrak{m}) = 1$ is $\frac{h(K)^{\mathrm{odd,2\text{-}split}}-1}{2}$, where 2-split means divide by the subgroup generated by an ideal above 2.*

Theorem

(a) *If $N \equiv 3 \pmod 8$ and $K = \mathbb{Q}(\sqrt{-N})$, then there are exactly $h(K)$ supersingular K-dihedral $\mathfrak{m}$.*

(b) *If $N \equiv 5 \pmod 8$, $K = \mathbb{Q}(\sqrt{N})$, and the fundamental unit of K is $\equiv 1 \pmod{2\mathcal{O}_K}$, then there are $h(K)$ supersingular K-dihedral $\mathfrak{m}$.*

(c) *In all other cases, no supersingular K-dihedral maximal ideals exist.*

Aside: corollaries for elliptic curves (part 1)

As a bonus, we recover some results on elliptic curves over $\mathbb{Q}$. Note that elliptic curves of conductor N with a rational 2-torsion point are completely classified (they are *Neumann–Setzer(–Hadano) curves*).

Theorem (Setzer, 1975; Kida, 2003; Hadano, 1974)

- *If $N \equiv 1, 7 \pmod 8$ and $3 \nmid h(\mathbb{Q}(\sqrt{\pm N}))$, then every elliptic curve of conductor N has a rational 2-torsion point.*
- *If $3 \nmid h(\mathbb{Q}(\sqrt{\pm N})), h(\mathbb{Q}(\sqrt{(-1)^{(N-1)/2}N}), 2)$, then every elliptic curve of conductor N has a rational 2-torsion point. Here $h(\bullet, 2)$ denotes a ray class number of modulus 2.*
- *If $3 \nmid h(\mathbb{Q}(\sqrt{\pm N})), h(\mathbb{Q}(\sqrt{\pm 2N}))$, then every elliptic curve of conductor $2N$ has a rational 2-torsion point. Moreover, no such curves occur if $N \equiv 3, 5 \pmod 8$.*

Aside: corollaries for elliptic curves (part 1)

As a bonus, we recover some results on elliptic curves over $\mathbb{Q}$. Note that elliptic curves of conductor N with a rational 2-torsion point are completely classified (they are *Neumann–Setzer(–Hadano) curves*).

Theorem (Setzer, 1975; Kida, 2003; Hadano, 1974)

- *If $N \equiv 1, 7 \pmod 8$ and $3 \nmid h(\mathbb{Q}(\sqrt{\pm N}))$, then every elliptic curve of conductor N has a rational 2-torsion point.*
- *If $3 \nmid h(\mathbb{Q}(\sqrt{\pm N})), h(\mathbb{Q}(\sqrt{(-1)^{(N-1)/2} N}), 2)$, then every elliptic curve of conductor N has a rational 2-torsion point. Here $h(\bullet, 2)$ denotes a ray class number of modulus 2.*
- *If $3 \nmid h(\mathbb{Q}(\sqrt{\pm N})), h(\mathbb{Q}(\sqrt{\pm 2N}))$, then every elliptic curve of conductor $2N$ has a rational 2-torsion point. Moreover, no such curves occur if $N \equiv 3, 5 \pmod 8$.*

Aside: corollaries for elliptic curves (part 1)

As a bonus, we recover some results on elliptic curves over $\mathbb{Q}$. Note that elliptic curves of conductor N with a rational 2-torsion point are completely classified (they are *Neumann–Setzer(–Hadano) curves*).

Theorem (Setzer, 1975; Kida, 2003; Hadano, 1974)

- *If $N \equiv 1, 7 \pmod 8$ and $3 \nmid h(\mathbb{Q}(\sqrt{\pm N}))$, then every elliptic curve of conductor N has a rational 2-torsion point.*
- *If $3 \nmid h(\mathbb{Q}(\sqrt{\pm N})), h(\mathbb{Q}(\sqrt{(-1)^{(N-1)/2}N}), 2)$, then every elliptic curve of conductor N has a rational 2-torsion point. Here $h(\bullet, 2)$ denotes a ray class number of modulus 2.*
- *If $3 \nmid h(\mathbb{Q}(\sqrt{\pm N})), h(\mathbb{Q}(\sqrt{\pm 2N}))$, then every elliptic curve of conductor $2N$ has a rational 2-torsion point. Moreover, no such curves occur if $N \equiv 3, 5 \pmod 8$.*

Aside: corollaries for elliptic curves (part 1)

As a bonus, we recover some results on elliptic curves over $\mathbb{Q}$. Note that elliptic curves of conductor N with a rational 2-torsion point are completely classified (they are *Neumann–Setzer(–Hadano) curves*).

Theorem (Setzer, 1975; Kida, 2003; Hadano, 1974)

- *If* $N \equiv 1, 7 \pmod 8$ *and* $3 \nmid h(\mathbb{Q}(\sqrt{\pm N}))$, *then every elliptic curve of conductor* N *has a rational 2-torsion point.*
- *If* $3 \nmid h(\mathbb{Q}(\sqrt{\pm N})), h(\mathbb{Q}(\sqrt{(-1)^{(N-1)/2}N}), 2)$, *then every elliptic curve of conductor* N *has a rational 2-torsion point. Here* $h(\bullet, 2)$ *denotes a ray class number of modulus 2.*
- *If* $3 \nmid h(\mathbb{Q}(\sqrt{\pm N})), h(\mathbb{Q}(\sqrt{\pm 2N}))$, *then every elliptic curve of conductor* $2N$ *has a rational 2-torsion point. Moreover, no such curves occur if* $N \equiv 3, 5 \pmod 8$.

Aside: corollaries for elliptic curves (part 1)

As a bonus, we recover some results on elliptic curves over $\mathbb{Q}$. Note that elliptic curves of conductor N with a rational 2-torsion point are completely classified (they are *Neumann–Setzer(–Hadano) curves*).

Theorem (Setzer, 1975; Kida, 2003; Hadano, 1974)

- *If $N \equiv 1, 7 \pmod 8$ and $3 \nmid h(\mathbb{Q}(\sqrt{\pm N}))$, then every elliptic curve of conductor N has a rational 2-torsion point.*
- *If $3 \nmid h(\mathbb{Q}(\sqrt{\pm N})), h(\mathbb{Q}(\sqrt{(-1)^{(N-1)/2} N}), 2)$, then every elliptic curve of conductor N has a rational 2-torsion point. Here $h(\bullet, 2)$ denotes a ray class number of modulus 2.*
- *If $3 \nmid h(\mathbb{Q}(\sqrt{\pm N})), h(\mathbb{Q}(\sqrt{\pm 2N}))$, then every elliptic curve of conductor $2N$ has a rational 2-torsion point. Moreover, no such curves occur if $N \equiv 3, 5 \pmod 8$.*

Aside: corollaries for elliptic curves (part 2)

The preceding results are derived using a totally different approach: following Ogg, one considers the formula

$$\Delta = -16(4A^3 + 27B^2)$$

for the discriminant of a short Weierstrass equation $y^2 = x^3 + Ax + B$ as an S-unit equation. This is then combined with a study of the splitting field of $x^3 + Ax + B$.

The latter is the 2-division field, so one can reinterpret much of the analysis in terms of mod-2 Galois representations; the results would be similar to what we have done. However, our point of view adapts readily to mod-ℓ representations for $\ell > 2$; the case $\ell = 3$ is likely to be particularly amenable. (The class numbers will be replaced by something else; see below.)

Aside: corollaries for elliptic curves (part 2)

The preceding results are derived using a totally different approach: following Ogg, one considers the formula

$$\Delta = -16(4A^3 + 27B^2)$$

for the discriminant of a short Weierstrass equation $y^2 = x^3 + Ax + B$ as an S-unit equation. This is then combined with a study of the splitting field of $x^3 + Ax + B$.

The latter is the 2-division field, so one can reinterpret much of the analysis in terms of mod-2 Galois representations; the results would be similar to what we have done. However, our point of view adapts readily to mod-ℓ representations for $\ell > 2$; the case $\ell = 3$ is likely to be particularly amenable. (The class numbers will be replaced by something else; see below.)

About those frequencies

Recall the table I showed earlier:

$N \pmod 8$	1	3	5	7
frequency of $a_2(\mathfrak{m}) = 0$	16.8%	always	42.2%	17.3%
frequency of $a_2(\mathfrak{m}) = 1$	always	always	always	47.9%

Of the "always" entries, the case $N \equiv 1 \pmod 8$ comes from reducible $\mathfrak{m}$; the others come from genus theory plus $h(\mathbb{Q}(\sqrt{-N})) > 1$ for $N > 163$.

As for the remaining probabilities, compare these to the lower bounds coming from Cohen–Lenstra heuristics:†

$N \pmod 8$	1	3	5	7
frequency of $a_2(\mathfrak{m}) = 0$	none	always	33.3%	none
frequency of $a_2(\mathfrak{m}) = 1$	always	always	always	43.1%

†These are for the 3-torsion in quadratic class groups, and so are partially accessible by work of Davenport–Heilbronn, Bhargava–Shankar–Tsimerman, Taniguchi–Thorne.

About those frequencies

Recall the table I showed earlier:

$N \pmod 8$	1	3	5	7
frequency of $a_2(\mathfrak{m}) = 0$	16.8%	always	42.2%	17.3%
frequency of $a_2(\mathfrak{m}) = 1$	always	always	always	47.9%

Of the "always" entries, the case $N \equiv 1 \pmod 8$ comes from reducible $\mathfrak{m}$; the others come from genus theory plus $h(\mathbb{Q}(\sqrt{-N})) > 1$ for $N > 163$.

As for the remaining probabilities, compare these to the lower bounds coming from Cohen–Lenstra heuristics:[†]

$N \pmod 8$	1	3	5	7
frequency of $a_2(\mathfrak{m}) = 0$	none	always	33.3%	none
frequency of $a_2(\mathfrak{m}) = 1$	always	always	always	43.1%

[†] These are for the 3-torsion in quadratic class groups, and so are partially accessible by work of Davenport–Heilbronn, Bhargava–Shankar–Tsimerman, Taniguchi–Thorne.

About those frequencies

Recall the table I showed earlier:

$N \pmod 8$	1	3	5	7
frequency of $a_2(\mathfrak{m}) = 0$	16.8%	always	42.2%	17.3%
frequency of $a_2(\mathfrak{m}) = 1$	always	always	always	47.9%

Of the "always" entries, the case $N \equiv 1 \pmod 8$ comes from reducible $\mathfrak{m}$; the others come from genus theory plus $h(\mathbb{Q}(\sqrt{-N})) > 1$ for $N > 163$.

As for the remaining probabilities, compare these to the lower bounds coming from Cohen–Lenstra heuristics:†

$N \pmod 8$	1	3	5	7
frequency of $a_2(\mathfrak{m}) = 0$	none	always	33.3%	none
frequency of $a_2(\mathfrak{m}) = 1$	always	always	always	43.1%

†These are for the 3-torsion in quadratic class groups, and so are partially accessible by work of Davenport–Heilbronn, Bhargava–Shankar–Tsimerman, Taniguchi–Thorne.

How to finish the analysis

To close the gaps between the two tables, one would need to also analyze exceptional and big-image representations; the presence of these should be related to the existence of G-extensions of $\mathbb{Q}$ unramified outside $2N$ for $G \subseteq \mathrm{GL}_2(\mathbb{F}_q)$ (where q is a power of 2).

In some cases (notably, when such an extension is unramified over a quadratic field), this is governed by a suitable analogue of the Cohen–Lenstra heuristics (see the work of Wood). We hope that such heuristics will suffice to explain the frequencies in the original table.

This also applies in case $\ell > 2$ (in which case q is a power of ℓ). Again, the case $\ell = 3$ seems particularly attractive.

How to finish the analysis

To close the gaps between the two tables, one would need to also analyze exceptional and big-image representations; the presence of these should be related to the existence of G-extensions of $\mathbb{Q}$ unramified outside $2N$ for $G \subseteq \mathrm{GL}_2(\mathbb{F}_q)$ (where q is a power of 2).

In some cases (notably, when such an extension is unramified over a quadratic field), this is governed by a suitable analogue of the Cohen–Lenstra heuristics (see the work of Wood). We hope that such heuristics will suffice to explain the frequencies in the original table.

This also applies in case $\ell > 2$ (in which case q is a power of ℓ). Again, the case $\ell = 3$ seems particularly attractive.

How to finish the analysis

To close the gaps between the two tables, one would need to also analyze exceptional and big-image representations; the presence of these should be related to the existence of G-extensions of $\mathbb{Q}$ unramified outside $2N$ for $G \subseteq \mathrm{GL}_2(\mathbb{F}_q)$ (where q is a power of 2).

In some cases (notably, when such an extension is unramified over a quadratic field), this is governed by a suitable analogue of the Cohen–Lenstra heuristics (see the work of Wood). We hope that such heuristics will suffice to explain the frequencies in the original table.

This also applies in case $\ell > 2$ (in which case q is a power of ℓ). Again, the case $\ell = 3$ seems particularly attractive.