Heuristics for the growth of Mordell-Weil ranks in big extensions of number fields

Barry Mazur, Harvard University
Karl Rubin, UC Irvine

Banff, June 2016

Growth of ranks in cyclic extensions

Fix an elliptic curve E over a number field K.

Question

As L runs through abelian extensions of K, how often is $\mathrm{rank}(E(L)) > \mathrm{rank}(E(K))$?

Growth of ranks in cyclic extensions

Fix an elliptic curve E over a number field K.

Question

As L runs through abelian extensions of K, how often is $\mathrm{rank}(E(L)) > \mathrm{rank}(E(K))$*?*

- By considering the action of $\mathrm{Gal}(L/K)$ on $E(L) \otimes \mathbb{Q}$, it is enough to consider the case where L/K is cyclic.
- General philosophy: it's hard to find L with $E(L) \neq E(K)$.

  For example, suppose L/K is cyclic of degree p, with a large prime p. The representation theory of $\mathbb{Q}[\mathrm{Gal}(L/K)]$ shows that

$$\mathrm{rank}(E(L)) > \mathrm{rank}(E(K)) \Longrightarrow \mathrm{rank}(E(L)) \geq \mathrm{rank}(E(K)) + (p-1).$$

Growth of ranks in cyclic extensions

Let $\mathfrak{f}_{L/K}$ denote the (finite part of the) conductor of L/K, and define

$$M_d(X) := \#\{L/K \text{ cyclic of degree } d : \mathbf{N}\mathfrak{f}_{L/K} < X\},$$
$$N_d(X) := \#\{L/K \text{ cyclic of degree } d : \mathbf{N}\mathfrak{f}_{L/K} < X \text{ and } E(L) \neq E(K)\}.$$

Conjecture (David-Fearnley-Kisilevsky)

If $K = \mathbb{Q}$ then:

1. $\log N_2(X) \sim \log(X)$ *(follows from standard conjectures),*
2. $\log N_3(X) \sim \frac{1}{2}\log(X)$,
3. $\log N_5(X) = o(\log(X))$ *but $N_5(X)$ is unbounded,*
4. $N_p(X)$ *is bounded if p is a prime, $p \geq 7$.*

Growth of ranks in cyclic extensions

Let $\mathfrak{f}_{L/K}$ denote the (finite part of the) conductor of L/K, and define

$$M_d(X) := \#\{L/K \text{ cyclic of degree } d : \mathbf{N}\mathfrak{f}_{L/K} < X\},$$
$$N_d(X) := \#\{L/K \text{ cyclic of degree } d : \mathbf{N}\mathfrak{f}_{L/K} < X \text{ and } E(L) \neq E(K)\}.$$

Conjecture (David-Fearnley-Kisilevsky)

If $K = \mathbb{Q}$ then:

1. $\log N_2(X) \sim \log(X)$ *(follows from standard conjectures),*
2. $\log N_3(X) \sim \frac{1}{2}\log(X)$*,*
3. $\log N_5(X) = o(\log(X))$ *but* $N_5(X)$ *is unbounded,*
4. $N_p(X)$ *is bounded if* p *is a prime,* $p \geq 7$*.*

If $K = \mathbb{Q}$ and p is prime, then $\log M_p(X) \sim \log(X)$.

Motivation for the conjecture is BSD combined with random matrix theory predictions for the vanishing of twisted L-functions.

Growth of ranks: algebraic approach

Theorem (M-R)

Fix an elliptic curve E/K. For every prime p and every $n > 0$, there are infinitely many cyclic extensions L/K of degree p^n such that $E(L) = E(K)$.

(We expect $E(L) = E(K)$ for almost all L/K, when $p > 2$.)

Growth of ranks: algebraic approach

Theorem (M-R)

Fix an elliptic curve E/K. For every prime p and every $n > 0$, there are infinitely many cyclic extensions L/K of degree p^n such that $E(L) = E(K)$.

(We expect $E(L) = E(K)$ for almost all L/K, when $p > 2$.)

Idea of proof: If $[L : K] = p$, then the Weil restriction $\mathrm{Res}_{L/K}E$ decomposes as

$$\mathrm{Res}_{L/K}E \sim A_L \times E$$

with an abelian variety A_L of dimension $p - 1$. Then

$$\mathrm{rank}(E(L)) = \mathrm{rank}(\mathrm{Res}_{L/K}(K)) = \mathrm{rank}(A_L(K)) + \mathrm{rank}(E(K)).$$

Choosing L carefully to have prescribed ramification, we can ensure that $\mathrm{Sel}_p(A_L/K) = 0$, so $\mathrm{rank}(A_L(K)) = 0$, so $\mathrm{rank}(E(L)) = \mathrm{rank}(E(K))$.

Growth of ranks: analytic approach

Question

As L runs through cyclic extensions of K, how often is $\mathrm{rank}(E(L)) > \mathrm{rank}(E(K))$?

Using the Birch & Swinnerton-Dyer conjecture, this is equivalent to the following:

Question

As χ runs through characters of $\mathrm{Gal}(\bar{K}/K)$, how often is $L(E,\chi,1) = 0$?

When $K = \mathbb{Q}$ (which we assume until further notice), this leads to a study of modular symbols.

Modular symbols

Definition

For $r \in \mathbb{Q}$, define the modular symbol $[r]_E$ by

$$[r]_E := \frac{1}{2}\left(\frac{2\pi i}{\Omega_E}\int_{i\infty}^{r} f_E(z)dz + \frac{2\pi i}{\Omega_E}\int_{i\infty}^{-r} f_E(z)dz\right)$$

where f_E is the modular form attached to E, and Ω_E is the real period.

Then $[r]_E \in \mathbb{Q}$, with denominators bounded depending only on $E(\mathbb{Q})_{\mathrm{tors}}$.

Modular symbols

Definition

For $r \in \mathbb{Q}$, define the modular symbol $[r]_E$ by

$$[r]_E := \frac{1}{2}\left(\frac{2\pi i}{\Omega_E}\int_{i\infty}^{r} f_E(z)dz + \frac{2\pi i}{\Omega_E}\int_{i\infty}^{-r} f_E(z)dz\right)$$

where f_E is the modular form attached to E, and Ω_E is the real period.

Then $[r]_E \in \mathbb{Q}$, with denominators bounded depending only on $E(\mathbb{Q})_{\mathrm{tors}}$.

Theorem

For every primitive even Dirichlet character χ of conductor m,

$$\sum_{a \in (\mathbb{Z}/m\mathbb{Z})^\times} \chi(a)[a/m]_E = \frac{\tau(\chi)L(E,\bar{\chi},1)}{\Omega_E}.$$

Here $\tau(\chi)$ is the Gauss sum.

Modular symbols

$$L(E,\chi,1)=0 \iff \sum_{a\in(\mathbb{Z}/m\mathbb{Z})^\times} \chi(a)[a/m]_E = 0. \qquad (*)$$

We want to know how often this happens.

We will try to understand the distribution of the modular symbols $[a/m]_E$, and use that to predict how often the right-hand side of $(*)$ vanishes.

Modular symbols

$$L(E,\chi,1) = 0 \iff \sum_{a \in (\mathbb{Z}/m\mathbb{Z})^\times} \chi(a)[a/m]_E = 0. \qquad (*)$$

We want to know how often this happens.

We will try to understand the distribution of the modular symbols $[a/m]_E$, and use that to predict how often the right-hand side of $(*)$ vanishes.

Our philosophy is that the modular symbols should be randomly distributed . . . except when they're not (i.e., except for certain relations that we understand).

Modular symbols

Let N be the conductor of E. For every $r \in \mathbb{Q}$, modular symbols satisfy the relations:

- $\boxed{[r]_E = [r+1]_E}$ since $f_E(z) = f_E(z+1)$
- $\boxed{[r]_E = [-r]_E}$ by definition
- Atkin-Lehner relation: if w_E is the global root number of E, and $aa'N \equiv 1 \pmod{m}$, then $\boxed{[a'/m]_E = w_E[a/m]_E}$
- Hecke relation: if a prime $\ell \nmid N$ and a_ℓ is the ℓ-th Fourier coefficient of f_E, then $\boxed{a_\ell[r]_E = [\ell r]_E + \sum_{i=0}^{\ell-1}[(r+i)/\ell]_E}$

Modular symbols

We begin by gathering data on modular symbols.

Distribution of modular symbols

Histogram of $\{[a/m]_E : E = 11A1, m = 10007, a \in (\mathbb{Z}/m\mathbb{Z})^\times\}$

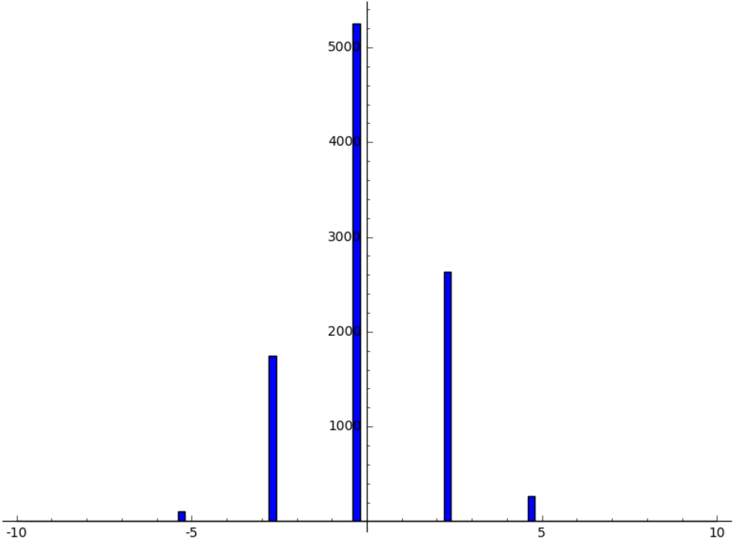

Distribution of modular symbols

Histogram of $\{[a/m]_E : E = 11A1, m = 10007, a \in (\mathbb{Z}/m\mathbb{Z})^\times\}$

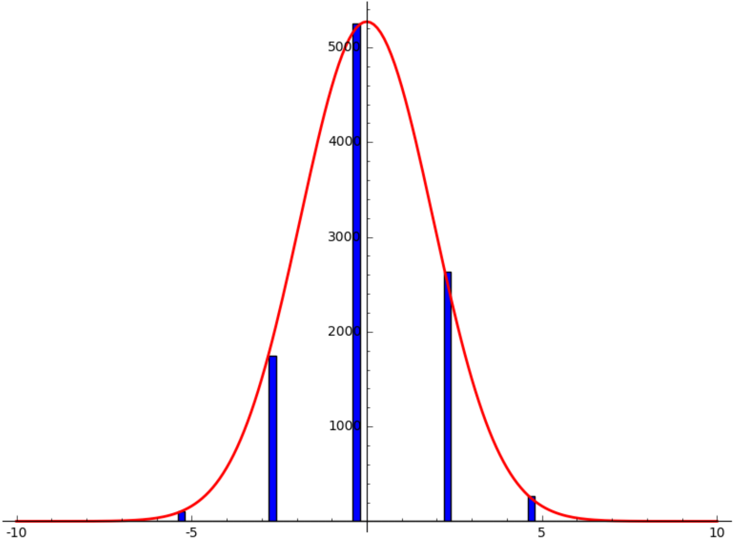

Distribution of modular symbols

Histogram of $\{[a/m]_E : E = 11A1, m = 100003, a \in (\mathbb{Z}/m\mathbb{Z})^\times\}$

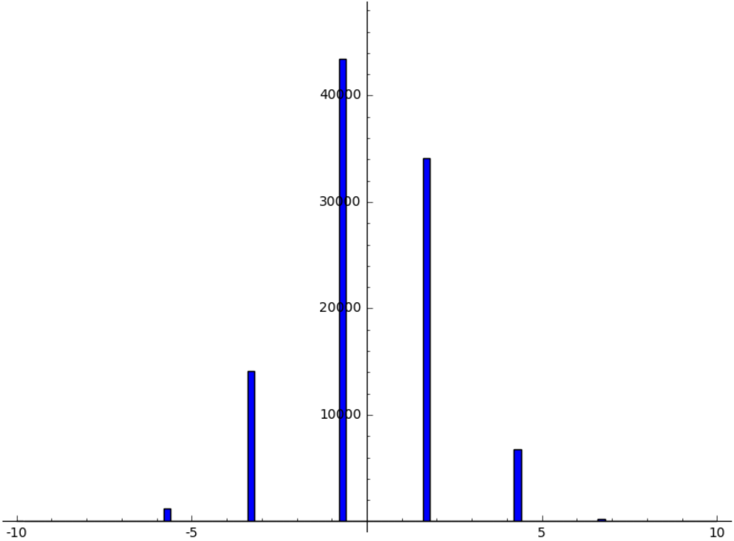

Distribution of modular symbols

Histogram of $\{[a/m]_E : E = 11A1, m = 100003, a \in (\mathbb{Z}/m\mathbb{Z})^\times\}$

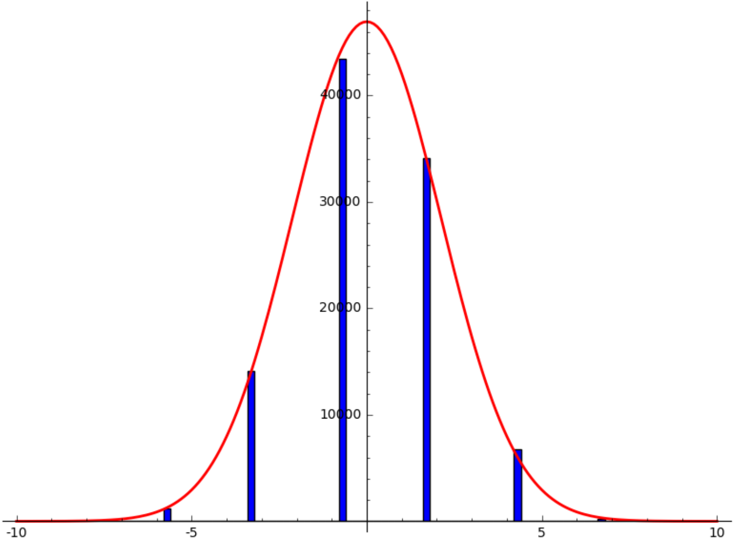

Distribution of modular symbols

Histogram of $\{[a/m]_E : E = 11A1, m = 1000003, a \in (\mathbb{Z}/m\mathbb{Z})^\times\}$

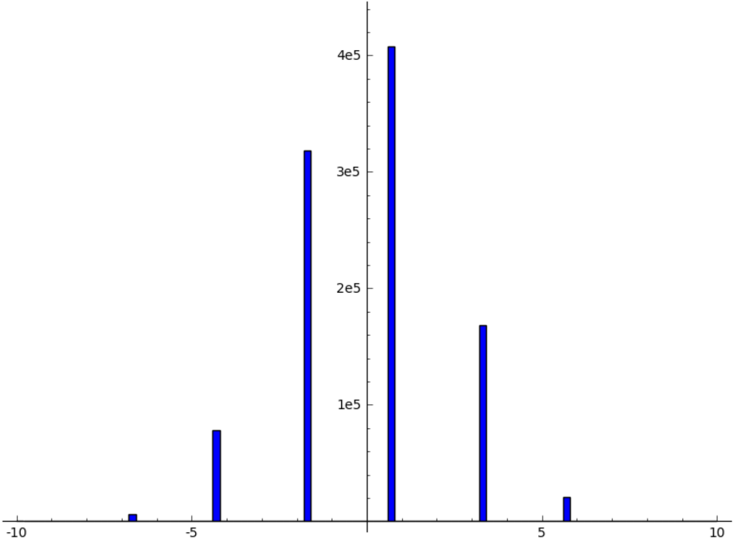

Distribution of modular symbols

Histogram of $\{[a/m]_E : E = 11A1, m = 1000003, a \in (\mathbb{Z}/m\mathbb{Z})^\times\}$

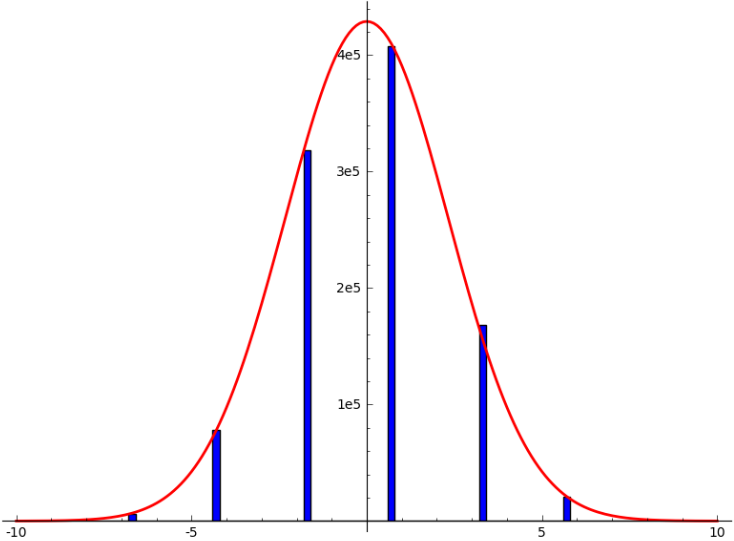

Distribution of modular symbols

Histogram of $\{[a/m]_E : E = 11A1, m = 10000019, a \in (\mathbb{Z}/m\mathbb{Z})^\times\}$

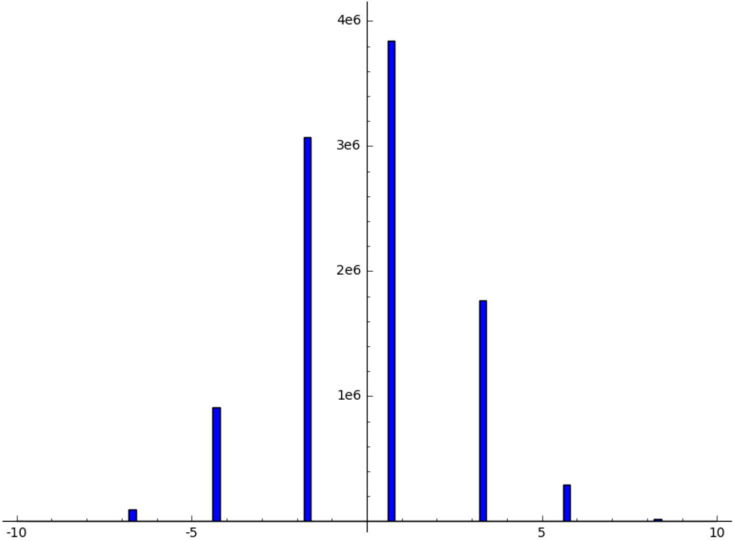

Distribution of modular symbols

Histogram of $\{[a/m]_E : E = 11A1, m = 10000019, a \in (\mathbb{Z}/m\mathbb{Z})^\times\}$

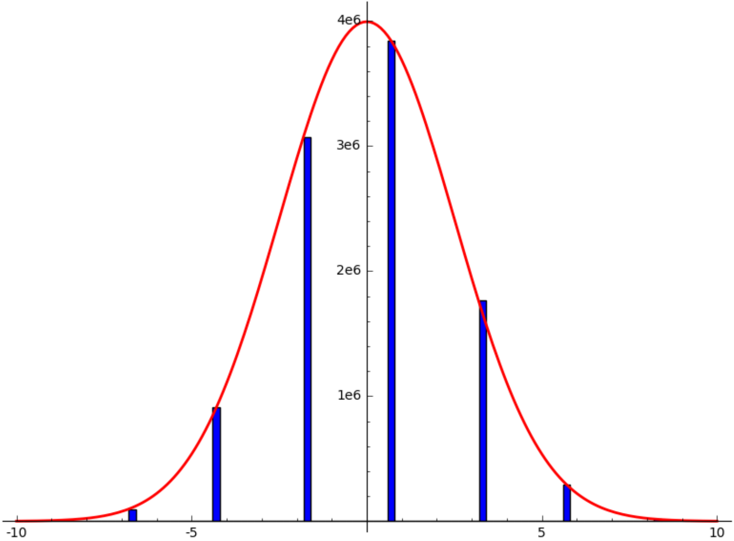

Distribution of modular symbols

This looks like a normal distribution.

Distribution of modular symbols

This looks like a normal distribution.

How does the variance depend on m?

Distribution of modular symbols

Plot of variance vs. m, for $E = 11A1$:

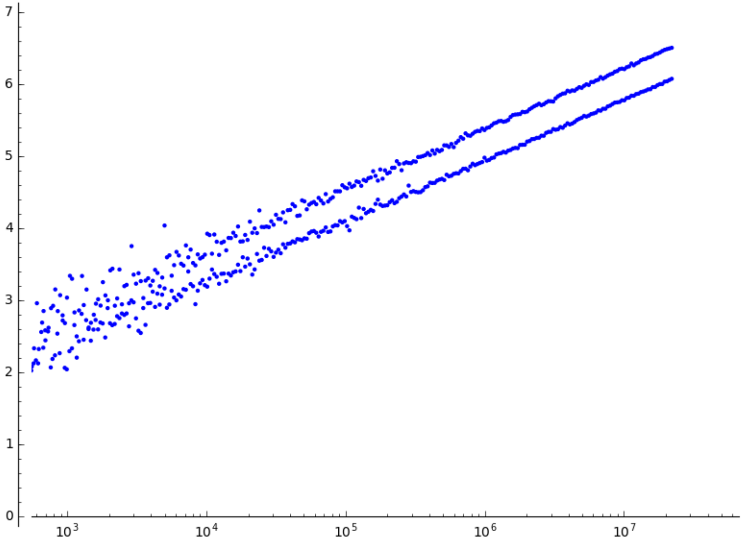

Distribution of modular symbols

Plot of variance vs. m, for $E = 11A1$:

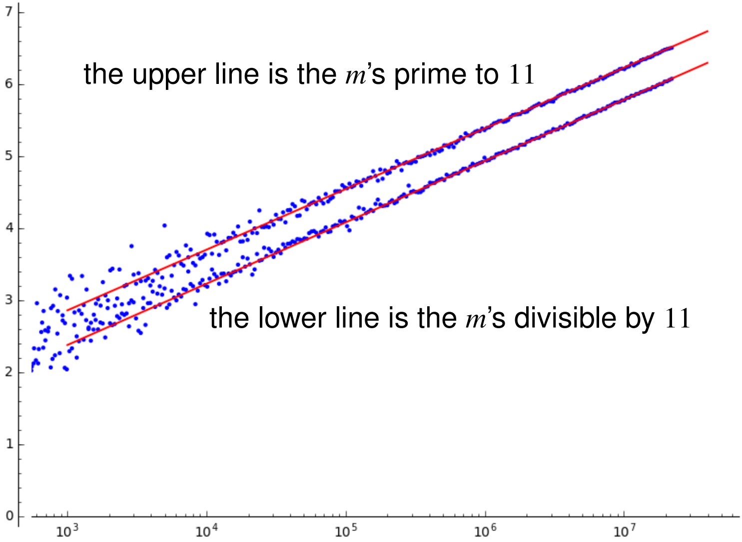

Distribution of modular symbols

Plot of variance vs. m, for $E = 45A1$:

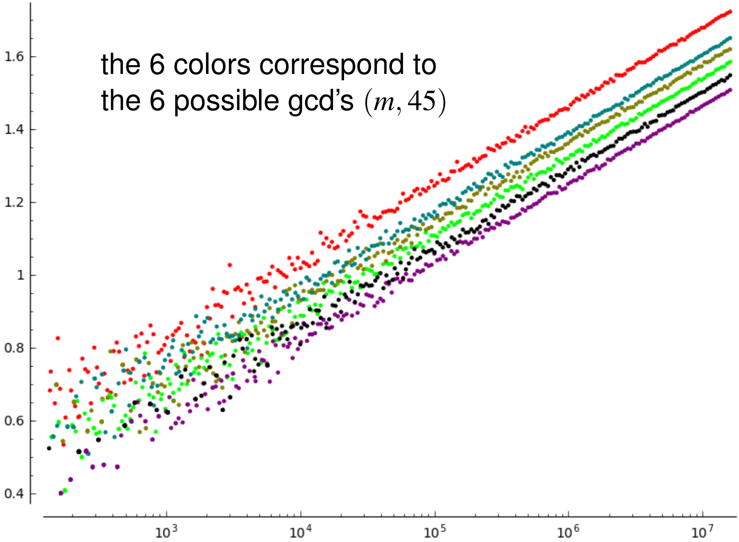

Distribution of modular symbols

It looks like the variance in the distribution of the $[a/m]_E$ converges to

$$\alpha_E \log(m) + \beta_{E,(m,N)}$$

where the slope α_E depends only on E, and $\beta_{E,(m,N)}$ depends on both E and the gcd (m, N).

Distribution of modular symbols

It looks like the variance in the distribution of the $[a/m]_E$ converges to

$$\alpha_E \log(m) + \beta_{E,(m,N)}$$

where the slope α_E depends only on E, and $\beta_{E,(m,N)}$ depends on both E and the gcd (m, N).

What is this constant α_E?

Distribution of modular symbols

Table of α_E, where the variance $\sim \alpha_E \log(m) + \beta_{E,(m,N)}$:

E	α_E	E	α_E
$11A1$	0.366	$1058A1$	1.07
$14A1$	0.108	$1058B1$	3.38
$15A1$	0.180	$1058C1$	0.060
$17A1$	0.189	$1058D1$	95.4
$19A1$	0.290	$1058E1$	0.235
$20A1$	0.043	$1059A1$	0.790
$21A1$	0.123	$1060A1$	0.156
$24A1$	0.062	$1062A1$	0.593
$26A1$	0.206	$1062B1$	0.842
$26B1$	0.038	$1062C1$	0.173
$27A1$	0.092	$1062D1$	1.05
$30A1$	0.038	$1062E1$	0.037
$32A1$	0.040	$1062F1$	11.3
$33A1$	0.220	$1062G1$	0.250

Distribution of modular symbols

Table of α_E, where the variance $\sim \alpha_E \log(m) + \beta_{E,(m,N)}$:

E	α_E	E	α_E
$11A1$	0.366	$1058A1$	1.07
$14A1$	0.108	$1058B1$	3.38
$15A1$	0.180	$1058C1$	0.060
$17A1$	0.189	$1058D1$	95.4
$19A1$	0.290	$1058E1$	0.235
$20A1$	0.043	$1059A1$	0.790
$21A1$	0.123	$1060A1$	0.156
$24A1$	0.062	$1062A1$	0.593
$26A1$	0.206	$1062B1$	0.842
$26B1$	0.038	$1062C1$	0.173
$27A1$	0.092	$1062D1$	1.05
$30A1$	0.038	$1062E1$	0.037
$32A1$	0.040	$1062F1$	11.3
$33A1$	0.220	$1062G1$	0.250

Distribution of modular symbols

Among all isogeny classes with conductor between 1000 and 1100:

largest α_E		largest modular degrees	
E	α_E	E	mod. degree
1015B1	26.0	1012A1	13776
1017G1	15.2	1014E1	6720
1026N1	18.4	1015B1	18720
1045B1	64.8	1020F1	6720
1050N1	35.2	1023A1	6840
1058D1	95.4	1026D1	8568
1062F1	11.3	1050N1	11400
1078B1	33.2	1058D1	19320
1085F1	38.3	1062F1	12672
1089I1	14.4	1085F1	13056

Distribution of modular symbols

Among all isogeny classes with conductor between 1000 and 1100:

largest α_E		largest modular degrees	
E	α_E	E	mod. degree
$1015B1$	26.0	$1012A1$	13776
$1017G1$	15.2	$1014E1$	6720
$1026N1$	18.4	$1015B1$	18720
$1045B1$	64.8	$1020F1$	6720
$1050N1$	35.2	$1023A1$	6840
$1058D1$	95.4	$1026D1$	8568
$1062F1$	11.3	$1050N1$	11400
$1078B1$	33.2	$1058D1$	19320
$1085F1$	38.3	$1062F1$	12672
$1089I1$	14.4	$1085F1$	13056

Distribution of modular symbols

Back to $E = 11A1$, histogram of $[a/m]_E/\sqrt{.366\log(m) + .333}$ for 10^6 random values of a/m with m prime to 11 and $0 < m < 10^{16}$:

The red curve corresponds to a normal distribution with variance 1.

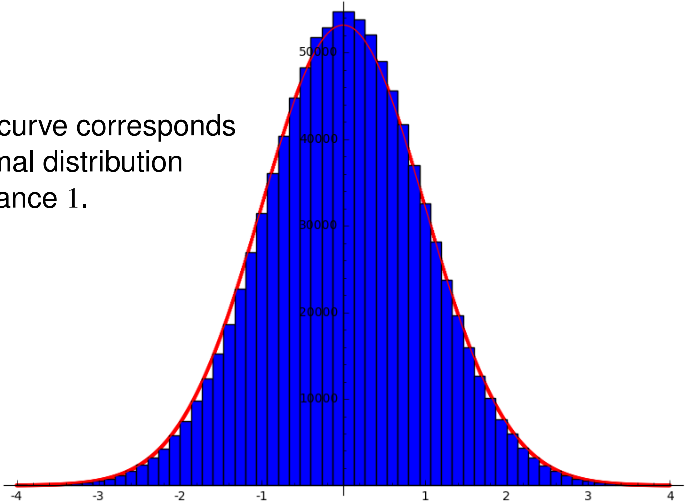

Distribution of modular symbols

OK, it seems pretty convincing that, properly normalized, the $[a/m]_E$ satisfy a normal distribution with variance 1.

What does this tell us about the vanishing of $L(E, \chi, 1)$, for a Dirichlet character χ?

Distribution of θ-coefficients

Suppose $L/\mathbb{Q}$ has conductor m, so there is a canonical surjection

$$\rho_L : (\mathbb{Z}/m\mathbb{Z})^\times \cong \mathrm{Gal}(\mathbb{Q}(\boldsymbol{\mu}_m)/\mathbb{Q}) \twoheadrightarrow \mathrm{Gal}(L/\mathbb{Q}).$$

Define

$$c_g := \sum_{a \in \rho_L^{-1}(g)} [a/m]_E \quad \text{for } g \in \mathrm{Gal}(L/\mathbb{Q}),$$
$$\theta_L := \sum_{g \in \mathrm{Gal}(L/\mathbb{Q})} c_g g \in \mathbb{Q}[\mathrm{Gal}(L/\mathbb{Q})].$$

Then for all $\chi : \mathrm{Gal}(L/\mathbb{Q}) \hookrightarrow \mathbb{C}^\times$,

$$\chi(\theta_L) = \frac{\tau(\chi) L(E, \bar{\chi}, 1)}{\Omega_E}.$$

We want to know how often this vanishes.

Distribution of θ-coefficients

If $[L:\mathbb{Q}] = d$, then each θ-coefficient c_g is a sum of $\varphi(m)/d$ modular symbols. We (think we) know how the modular symbols are distributed, but are they independent? If so, then the

$$\frac{c_g}{\sqrt{(\alpha_E \log(m) + \beta_{E,(m,N)})(\varphi(m)/d)}}$$

should satisfy a normal distribution with variance 1.

There are two ways to get data to test this:

- choose d large (so there are many data points) and try one or more m,
- choose any d, and try many different m.

Distribution of θ-coefficients, large d

$E = 11A1$, $m = 25035013$, L is the field of degree 5003 in $\mathbb{Q}(\boldsymbol{\mu}_m)$:

The red curve is the expected normal distribution.

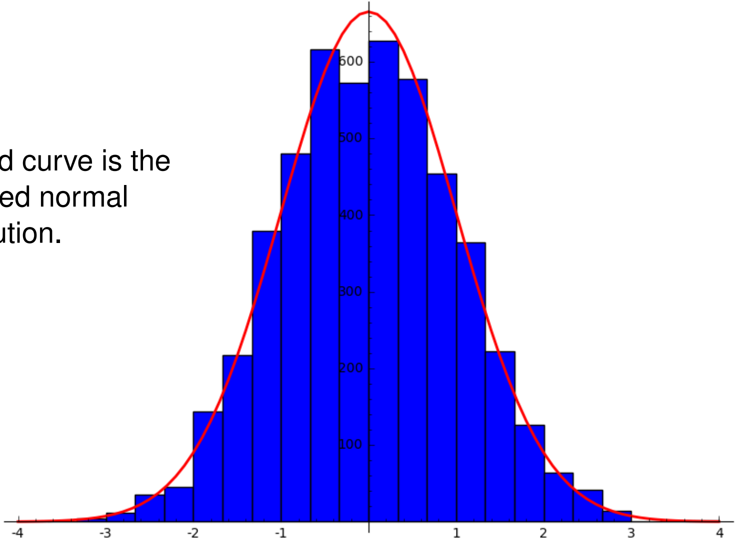

Distribution of θ-coefficients, large d

$E = 11A1$, $m = 49063009$, L is the field of degree 7001 in $\mathbb{Q}(\boldsymbol{\mu}_m)$:

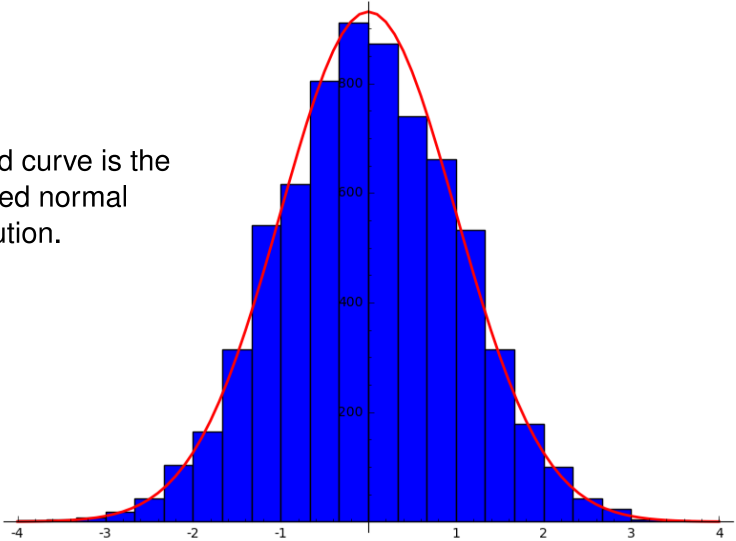

The red curve is the expected normal distribution.

Distribution of θ-coefficients, $d = 3$

$E = 11A1$, $m \equiv 1 \pmod 3$, $L \subset \mathbb{Q}(\boldsymbol{\mu}_m)$, $[L : \mathbb{Q}] = 3$,

$10000 < m < 20000$:

The red curve is the expected normal distribution.

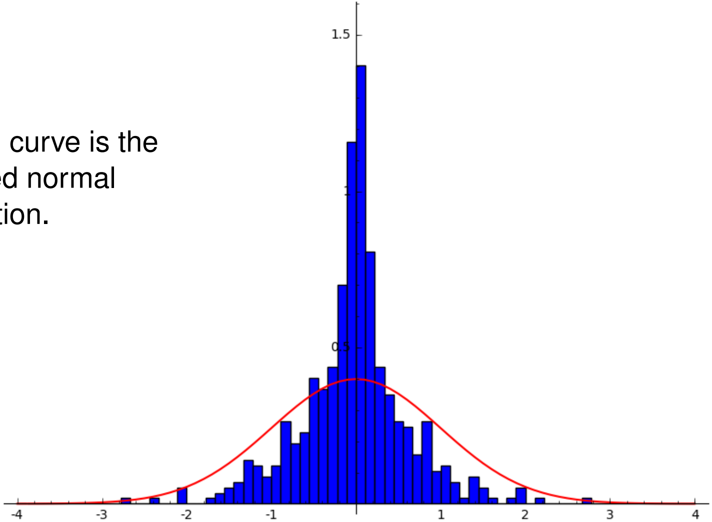

Distribution of θ-coefficients, $d = 3$

$E = 11A1$, $m \equiv 1 \pmod 3$, $L \subset \mathbb{Q}(\boldsymbol{\mu}_m)$, $[L : \mathbb{Q}] = 3$,

$20000 < m < 40000$:

The red curve is the expected normal distribution.

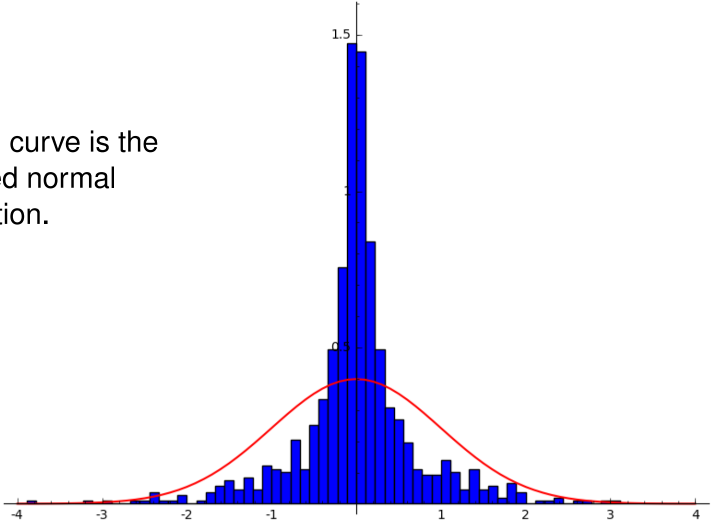

Distribution of θ-coefficients, $d = 3$

$E = 11A1$, $m \equiv 1 \pmod 3$, $L \subset \mathbb{Q}(\boldsymbol{\mu}_m)$, $[L : \mathbb{Q}] = 3$,

$40000 < m < 80000$:

The red curve is the expected normal distribution.

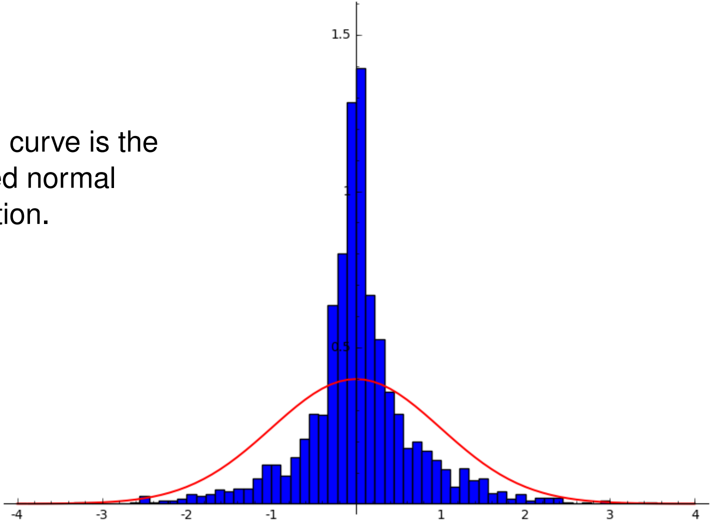

Distribution of θ-coefficients, $d = 3$

$E = 11A1$, $m \equiv 1 \pmod 3$, $L \subset \mathbb{Q}(\boldsymbol{\mu}_m)$, $[L : \mathbb{Q}] = 3$,

$80000 < m < 160000$:

The red curve is the expected normal distribution.

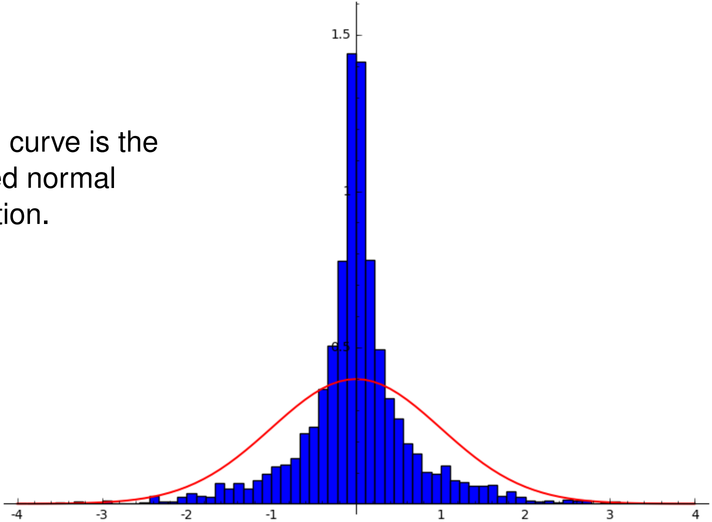

Distribution of θ-coefficients, $d = 3$

$E = 11A1$, $m \equiv 1 \pmod 3$, $L \subset \mathbb{Q}(\boldsymbol{\mu}_m)$, $[L : \mathbb{Q}] = 3$,

$160000 < m < 320000$:

The red curve is the expected normal distribution.

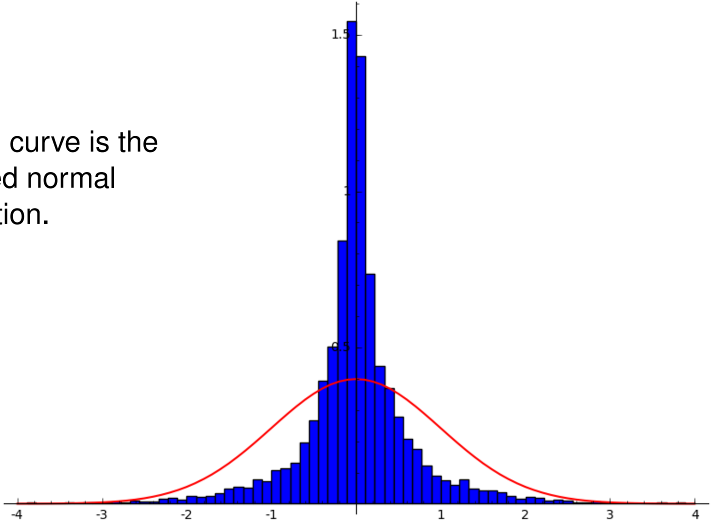

Distribution of θ-coefficients, $d = 3$

$E = 11A1$, $m \equiv 1 \pmod 3$, $L \subset \mathbb{Q}(\boldsymbol{\mu}_m)$, $[L:\mathbb{Q}] = 3$,

$320000 < m < 640000$:

The red curve is the expected normal distribution.

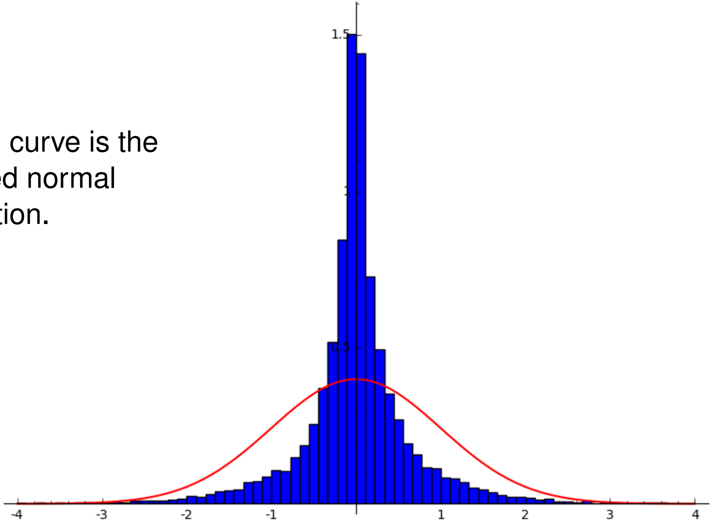

Distribution of θ-coefficients, $d = 3$

$E = 11A1$, $m \equiv 1 \pmod 3$, $L \subset \mathbb{Q}(\boldsymbol{\mu}_m)$, $[L : \mathbb{Q}] = 3$,

$10000 < m < 640000$:

The red curve is the expected normal distribution.

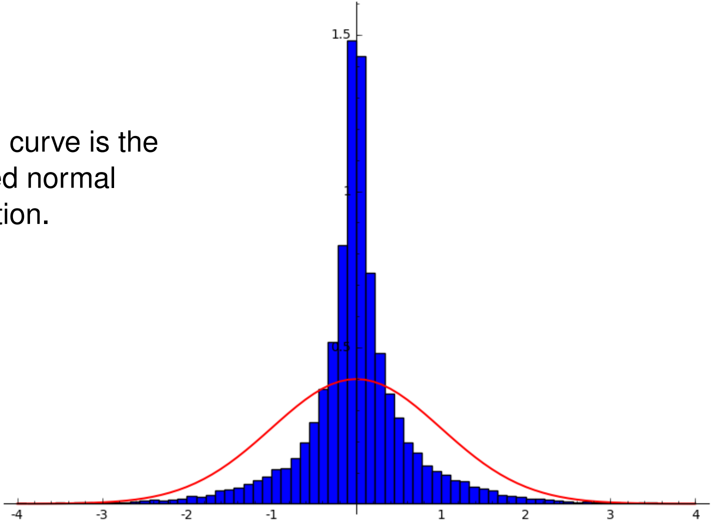

Distribution of θ-coefficients, $d = 3$

Doesn't look normal when $d = 3$.

Fix

- m prime, $m \equiv 1 \pmod 3$,
- cubic $L \subset \mathbb{Q}(\boldsymbol{\mu}_m)$,
- $H \subset (\mathbb{Z}/m\mathbb{Z})^\times$ the subgroup of cubes,

Then the coefficients of θ_L are the three sums

$$c_b := \sum_{a \in bH} [a/m]_E, \quad b \in (\mathbb{Z}/m\mathbb{Z})^\times/H.$$

The pictures seem to say that the $[a/m]_E$ are not independently distributed among these three sums.

Distribution of θ-coefficients, $d = 3$

$$c_b := \sum_{a \in bH} [a/m]_E, \quad b \in (\mathbb{Z}/m\mathbb{Z})^\times/H.$$

Recall that if $aa'N \equiv 1 \pmod{m}$, then $[a/m]_E = w_E[a'/m]_E$. Therefore

$$c_{bN} = w_E c_{b^{-1}N}.$$

So there are really only 2 coefficients, and one of them is zero if $w_E = -1$.

Distribution of θ-coefficients, small d

$E = 11A1$, $m \equiv 1 \pmod{d}$, $L \subset \mathbb{Q}(\boldsymbol{\mu}_m)$, $[L : \mathbb{Q}] = d$,

$d = 3$

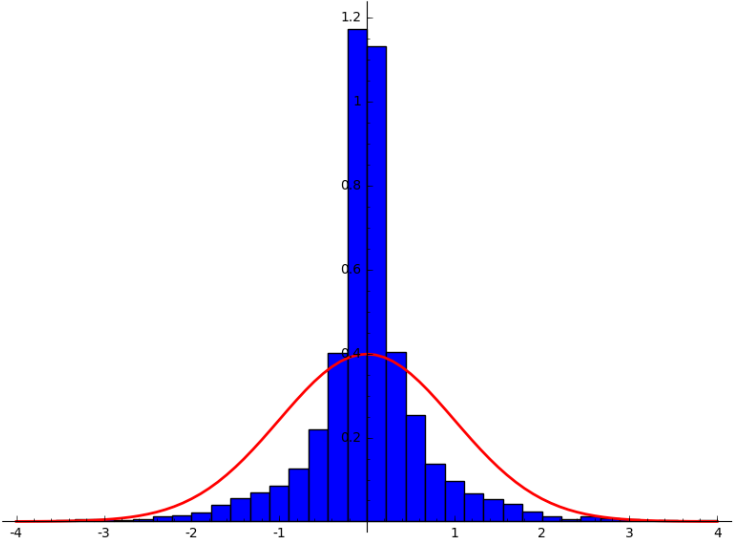

Distribution of θ-coefficients, small d

$E = 11A1$, $m \equiv 1 \pmod{d}$, $L \subset \mathbb{Q}(\boldsymbol{\mu}_m)$, $[L:\mathbb{Q}] = d$,

$d = 5$

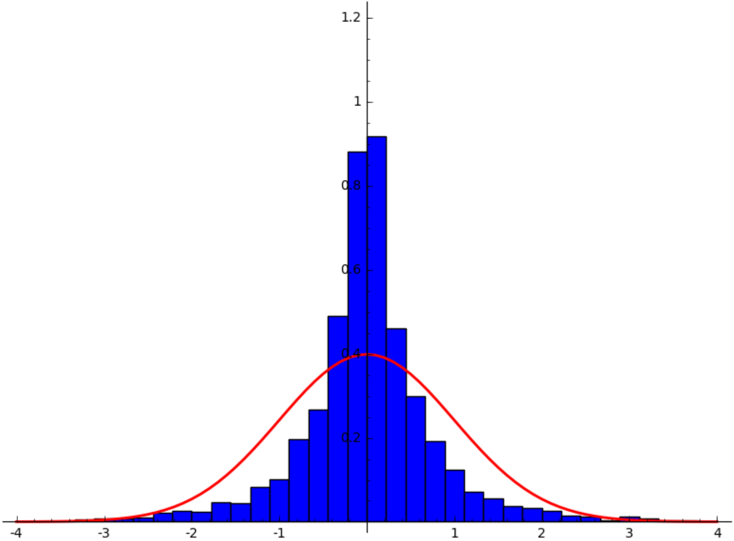

Distribution of θ-coefficients, small d

$E = 11A1$, $m \equiv 1 \pmod{d}$, $L \subset \mathbb{Q}(\boldsymbol{\mu}_m)$, $[L:\mathbb{Q}] = d$,

$d = 7$

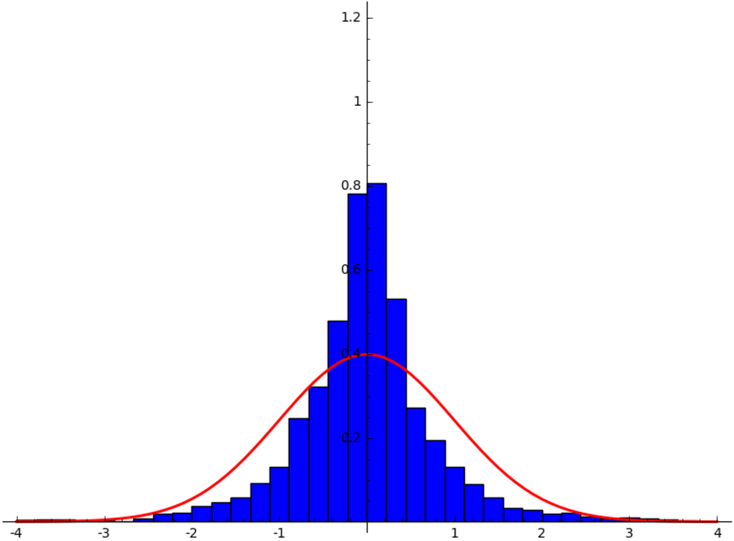

Distribution of θ-coefficients, small d

$E = 11A1$, $m \equiv 1 \pmod{d}$, $L \subset \mathbb{Q}(\boldsymbol{\mu}_m)$, $[L : \mathbb{Q}] = d$,

$d = 11$

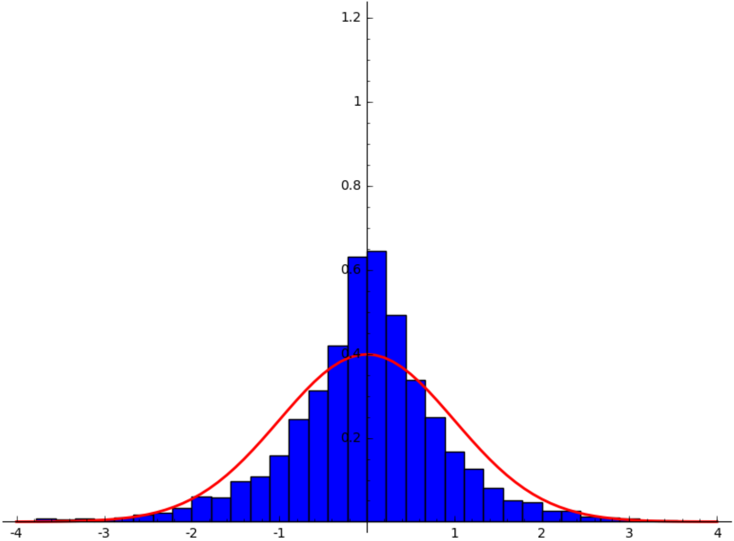

Distribution of θ-coefficients, small d

$E = 11A1$, $m \equiv 1 \pmod{d}$, $L \subset \mathbb{Q}(\boldsymbol{\mu}_m)$, $[L : \mathbb{Q}] = d$,

$d = 13$

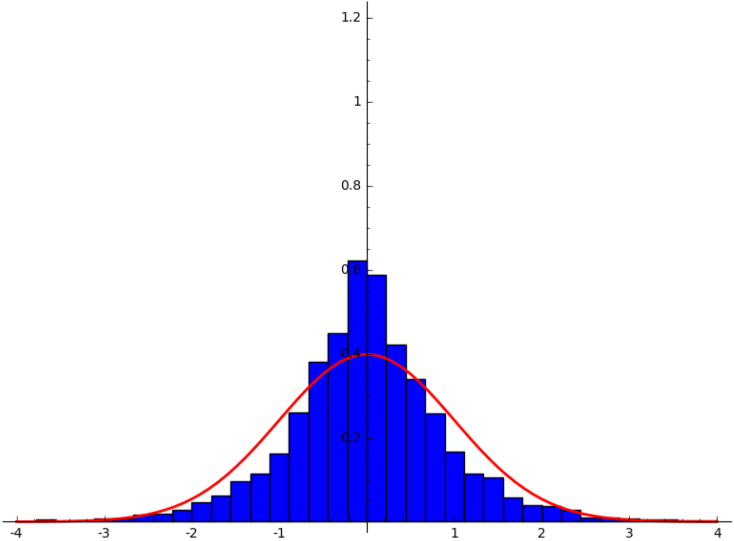

Distribution of θ-coefficients, small d

$E = 11A1$, $m \equiv 1 \pmod{d}$, $L \subset \mathbb{Q}(\boldsymbol{\mu}_m)$, $[L : \mathbb{Q}] = d$,

$d = 17$

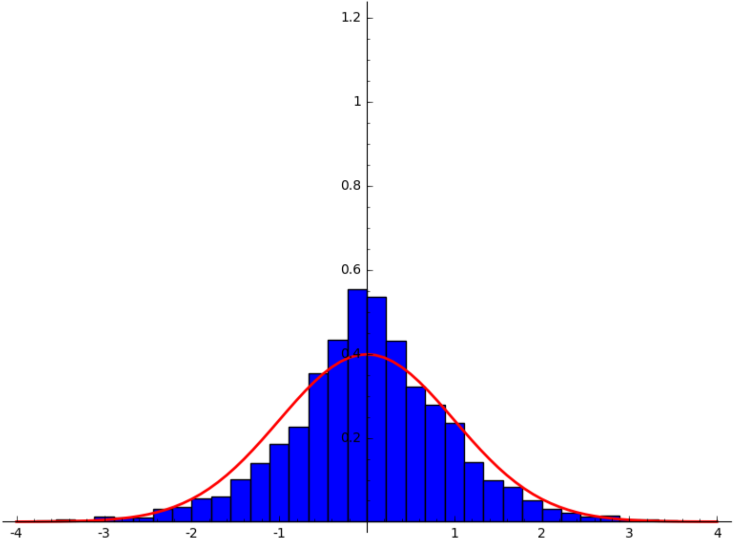

Distribution of θ-coefficients, small d

$E = 11A1$, $m \equiv 1 \pmod{d}$, $L \subset \mathbb{Q}(\boldsymbol{\mu}_m)$, $[L:\mathbb{Q}] = d$,

$d = 23$

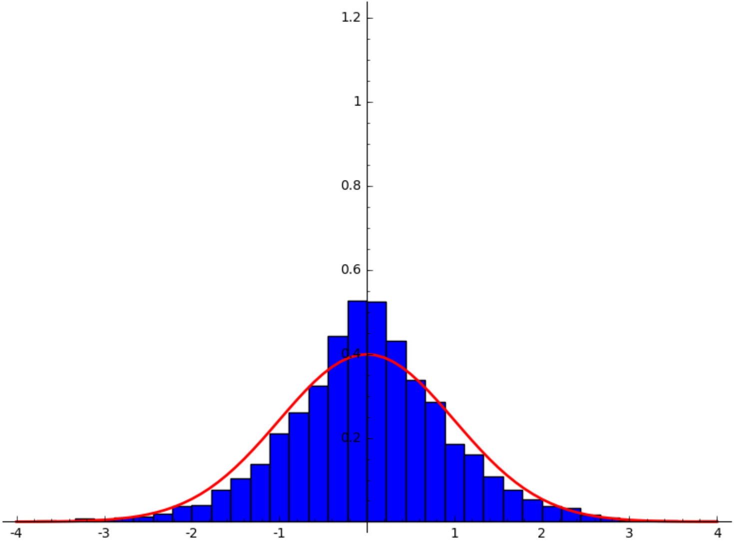

Distribution of θ-coefficients, small d

$E = 11A1$, $m \equiv 1 \pmod{d}$, $L \subset \mathbb{Q}(\boldsymbol{\mu}_m)$, $[L : \mathbb{Q}] = d$,

$d = 31$

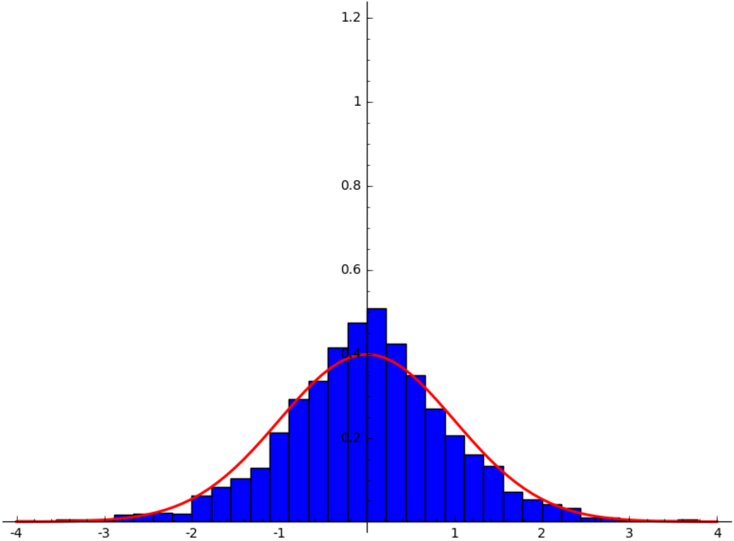

Distribution of θ-coefficients, small d

$E = 11A1$, $m \equiv 1 \pmod{d}$, $L \subset \mathbb{Q}(\boldsymbol{\mu}_m)$, $[L:\mathbb{Q}] = d$,

$d = 41$

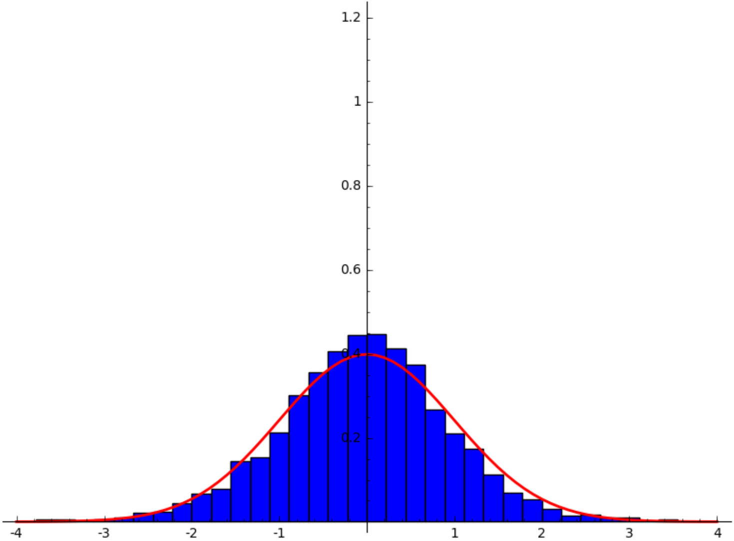

Distribution of θ-coefficients, small d

$E = 11A1$, $m \equiv 1 \pmod{d}$, $L \subset \mathbb{Q}(\boldsymbol{\mu}_m)$, $[L:\mathbb{Q}] = d$,

$d = 53$

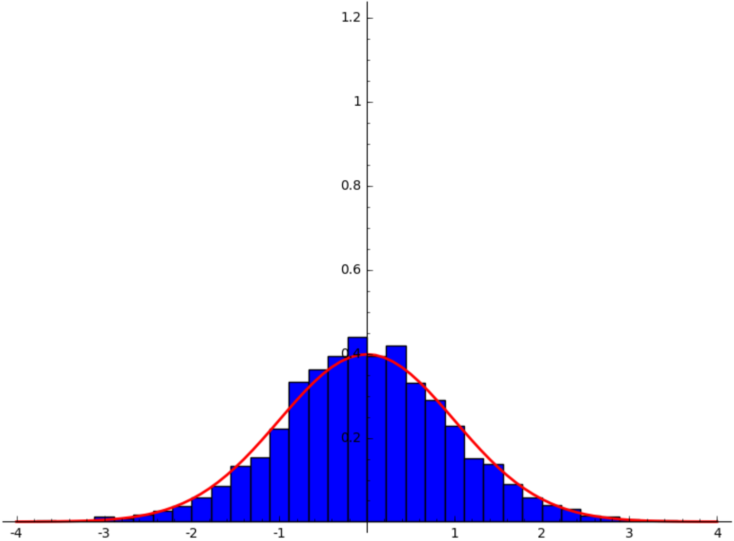

Distribution of θ-coefficients, small d

$E = 11A1$, $m \equiv 1 \pmod{d}$, $L \subset \mathbb{Q}(\boldsymbol{\mu}_m)$, $[L : \mathbb{Q}] = d$,

$d = 97$

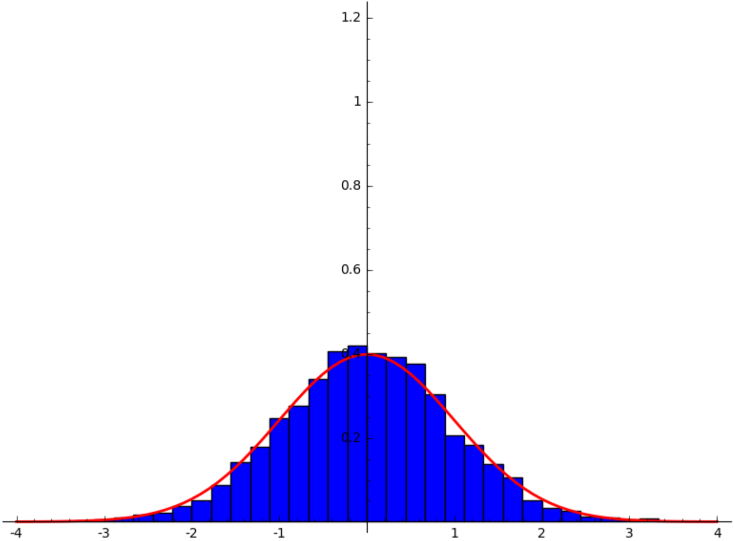

Distribution of θ-coefficients, small d

$E = 11A1$, $m \equiv 1 \pmod{d}$, $L \subset \mathbb{Q}(\boldsymbol{\mu}_m)$, $[L:\mathbb{Q}] = d$,

$d = 293$

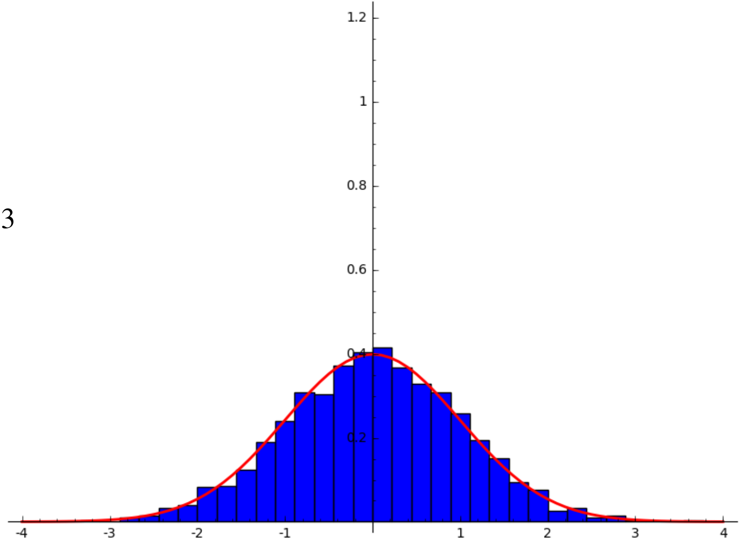

Oversimplified picture

Suppose $L/\mathbb{Q}$ is a cyclic extension of degree d and conductor m.

Very roughly, θ_L lies in a cube of side $\sqrt{\alpha_E \log(m)\varphi(m)/d}$ in the d-dimensional lattice $\mathbb{Z}[\mathrm{Gal}(L/\mathbb{Q})]$.

Oversimplified picture

Suppose $L/\mathbb{Q}$ is a cyclic extension of degree d and conductor m.

Very roughly, θ_L lies in a cube of side $\sqrt{\alpha_E \log(m)\varphi(m)/d}$ in the d-dimensional lattice $\mathbb{Z}[\mathrm{Gal}(L/\mathbb{Q})]$.

Suppose $\chi : \mathrm{Gal}(L/\mathbb{Q}) \twoheadrightarrow \boldsymbol{\mu}_d$ is a faithful character. Then

$$L(E,\chi,1) = 0 \iff \theta_L \in \ker(\chi : \mathbb{Z}[\mathrm{Gal}(L/\mathbb{Q})] \to \mathbb{C}).$$

That kernel is a sublattice of codimension $\varphi(d)$, so we might expect the "probability" that $L(E,\chi,1) = 0$ should be about

$$\left(\frac{C_E}{\sqrt{\log(m)\varphi(m)/d}}\right)^{\varphi(d)}$$

for some constant C_E.

Oversimplified picture

Suppose $L/\mathbb{Q}$ is a cyclic extension of degree d and conductor m.

Very roughly, θ_L lies in a cube of side $\sqrt{\alpha_E \log(m)\varphi(m)/d}$ in the d-dimensional lattice $\mathbb{Z}[\mathrm{Gal}(L/\mathbb{Q})]$.

Suppose $\chi : \mathrm{Gal}(L/\mathbb{Q}) \twoheadrightarrow \boldsymbol{\mu}_d$ is a faithful character. Then

$$L(E,\chi,1) = 0 \iff \theta_L \in \ker(\chi : \mathbb{Z}[\mathrm{Gal}(L/\mathbb{Q})] \to \mathbb{C}).$$

That kernel is a sublattice of codimension $\varphi(d)$, so we might expect the "probability" that $L(E,\chi,1) = 0$ should be about

$$\left(\frac{C_E}{\sqrt{\log(m)\varphi(m)/d}}\right)^{\varphi(d)}$$

for some constant C_E.

This goes to zero *very* fast as d and m grow.

Oversimplified picture

This isn't quite right:

- The previous argument ignores the Atkin-Lehner relation, which "pairs up" the coefficients and forces θ_L into a sublattice of $\mathbb{Z}[\mathrm{Gal}(L/\mathbb{Q})]$ with rank approximately $d/2$. Taking this into account changes the expectation to

$$\left(\frac{C_E\, d}{\log(m)\varphi(m)}\right)^{\varphi(d)/4}.$$

Oversimplified picture

This isn't quite right:

- The previous argument ignores the Atkin-Lehner relation, which "pairs up" the coefficients and forces θ_L into a sublattice of $\mathbb{Z}[\mathrm{Gal}(L/\mathbb{Q})]$ with rank approximately $d/2$. Taking this into account changes the expectation to

$$\left(\frac{C_E\, d}{\log(m)\varphi(m)}\right)^{\varphi(d)/4}.$$

- The distribution of the (normalized) θ_L is not uniform in a box, and we don't fully understand what the correct distribution is. Fortunately, for applications, it doesn't seem to matter very much what the distribution is, only that there is one.

Distribution of θ-coefficients

Fix d and let $G := \boldsymbol{\mu}_d$. Suppose $L/\mathbb{Q}$ is a cyclic extension of degree d. Suppose (only to avoid dealing with lots of cases)

- d is odd,
- the conductor m of $L/\mathbb{Q}$ is prime to the conductor N of E,
- the root number $w_E = +1$.

For every faithful character $\chi : \mathrm{Gal}(L/\mathbb{Q}) \xrightarrow{\sim} G$, define

$$\theta^*_{L,\chi} := \frac{\chi(\rho_L(N)^{(d+1)/2}\theta_L)}{\sqrt{(\alpha_E \log(m) + \beta_{E,1})(\varphi(m)/d)}} = \sum_{g \in G} c_{L,\chi,g} \cdot g \in \mathbb{R}[G]$$

which has been normalized for size and so that the Atkin-Lehner relation gives $c_{L,\chi,g} = c_{L,\chi,g^{-1}}$. Then

$$\theta^*_{L,\chi} \in \mathbb{R}[G]^+$$

where the superscript "$+$" denotes the space fixed under $g \mapsto g^{-1}$.

Distribution of θ-coefficients

Define

$$S(X) := \{(L,\chi) : \text{conductor}(L/\mathbb{Q}) < X,\ \chi : \mathrm{Gal}(L/\mathbb{Q}) \xrightarrow{\sim} G\}.$$

Questions

As $X \to \infty$,

1. *for each $g \in G$, does the distribution of the $c_{L,\chi,g} \in \mathbb{R}$ for $(L,\chi) \in S(X)$ converge to a bounded function f_g?*
2. *is the bound on f_g independent of d?*
3. *does the distribution of the $\theta^*_{L,\chi} \in \mathbb{R}[G]^+$ for $(L,\chi) \in S(X)$ converge to $\prod_{\{g,g^{-1}\}} f_g$?*

Distribution of θ's

Questions and Remarks

As $X \to \infty$,

1. *for each $g \in G$, does the distribution of the $c_{L,\chi,g} \in \mathbb{R}$ for $(L,\chi) \in S(X)$ converge to a bounded function f_g?*

   The data make this look plausible.

2. *is the bound on f_g independent of d?*

   Since the f_g seem to get closer and closer to a fixed normal distribution as d grows, this seems plausible too.

3. *does the distribution of the $\theta^*_{L,\chi} \in \mathbb{R}[G]^+$ for $(L,\chi) \in S(X)$ converge to $\prod_{\{g,g^{-1}\}} f_g$?*

   The third question is equivalent to asking that the coefficients be independent.

Distribution of θ's

Questions and Remarks

As $X \to \infty$,

1. *for each $g \in G$, does the distribution of the $c_{L,\chi,g} \in \mathbb{R}$ for $(L, \chi) \in S(X)$ converge to a bounded function f_g?*

   The data make this look plausible.

2. *is the bound on f_g independent of d?*

   Since the f_g seem to get closer and closer to a fixed normal distribution as d grows, this seems plausible too.

3. *does the distribution of the $\theta^*_{L,\chi} \in \mathbb{R}[G]^+$ for $(L, \chi) \in S(X)$ converge to $\prod_{\{g,g^{-1}\}} f_g$?*

   The third question is equivalent to asking that the coefficients be independent.

Distribution of θ's

Questions and Remarks

As $X \to \infty$,

1. *for each $g \in G$, does the distribution of the $c_{L,\chi,g} \in \mathbb{R}$ for $(L,\chi) \in S(X)$ converge to a bounded function f_g?*

   The data make this look plausible.

2. *is the bound on f_g independent of d?*

   Since the f_g seem to get closer and closer to a fixed normal distribution as d grows, this seems plausible too.

3. *does the distribution of the $\theta^*_{L,\chi} \in \mathbb{R}[G]^+$ for $(L,\chi) \in S(X)$ converge to $\prod_{\{g,g^{-1}\}} f_g$?*

   The third question is equivalent to asking that the coefficients be independent.

Distribution of θ's

Questions and Remarks

As $X \to \infty$,

1. *for each $g \in G$, does the distribution of the $c_{L,\chi,g} \in \mathbb{R}$ for $(L,\chi) \in S(X)$ converge to a bounded function f_g?*

   The data make this look plausible.

2. *is the bound on f_g independent of d?*

   Since the f_g seem to get closer and closer to a fixed normal distribution as d grows, this seems plausible too.

3. *does the distribution of the $\theta^*_{L,\chi} \in \mathbb{R}[G]^+$ for $(L,\chi) \in S(X)$ converge to $\prod_{\{g,g^{-1}\}} f_g$?*

   The third question is equivalent to asking that the coefficients be independent.

The heuristic

If the answer to these questions is "Yes", we get a heuristic estimate:

Heuristic

There is a constant C_E, depending only on E, such that

$$\mathrm{Prob}[L(E,\chi,1)=0] \le \left(\frac{C_E d}{\log(m)\varphi(m)}\right)^{\varphi(d)/4}$$

where d is the order of χ and m its conductor.

This should hold for all χ of order greater than 2.

Consequences of the heuristic

Heuristic

$\mathrm{Prob}[L(E,\chi,1)=0] \leq \left(\frac{C_E d}{\log(m)\varphi(m)}\right)^{\varphi(d)/4}$.

Example ($d=3$)

$$\sum_{\chi \text{ order } 3,\text{ conductor } < X} \mathrm{Prob}[L(E,\chi,1)=0] \ll \sum_{m=2}^{X} \frac{1}{(\log(m)\varphi(m))^{1/2}} \ll \sqrt{X}.$$

Example ($d=5$)

$$\sum_{\chi \text{ order } 5,\text{ conductor } < X} \mathrm{Prob}[L(E,\chi,1)=0] \ll \sum_{m=2}^{X} \frac{1}{\log(m)\varphi(m)} \ll \log X.$$

These are consistent with the prediction of David-Fearnley-Kisilevsky.

Consequences of the heuristic

Heuristic

$\text{Prob}[L(E,\chi,1)=0] \leq \left(\frac{C_E d}{\log(m)\varphi(m)}\right)^{\varphi(d)/4}$.

Example ($d=7$)

$$\sum_{\chi \text{ of order } 7} \text{Prob}[L(E,\chi,1)=0] \ll \sum_{m=2}^{\infty} \frac{1}{(\log(m)\varphi(m))^{3/2}} < \infty.$$

This is consistent with the prediction of David-Fearnley-Kisilevsky.

Consequences of the heuristic

Heuristic

$\mathrm{Prob}[L(E,\chi,1)=0] \le \left(\frac{C_E d}{\log(m)\varphi(m)}\right)^{\varphi(d)/4}$.

Proposition

Suppose $t : \mathbb{Z}_{>0} \to \mathbb{R}$ *is a function, and* $t(d) \gg \log(d)$. *Then*

$$\sum_{d\,:\,t(d)>1} \ \sum_{\chi \text{ order } d} \left(\frac{C_E d}{\log(m)\varphi(m)}\right)^{t(d)} \quad \textit{converges.}$$

Consequences of the heuristic

Heuristic

$\mathrm{Prob}[L(E,\chi,1)=0] \le \left(\frac{C_E d}{\log(m)\varphi(m)}\right)^{\varphi(d)/4}.$

Proposition

Suppose $t:\mathbb{Z}_{>0}\to\mathbb{R}$ *is a function, and* $t(d)\gg\log(d)$*. Then*

$$\sum_{d\,:\,t(d)>1}\ \sum_{\chi \text{ order } d}\left(\frac{C_E d}{\log(m)\varphi(m)}\right)^{t(d)} \quad \textit{converges.}$$

Applying this with $t(d)=\varphi(d)/4$ shows

Heuristic

$$\sum_{d\,:\,\varphi(d)>4}\ \sum_{\chi \text{ order } d}\mathrm{Prob}[L(E,\chi,1)=0] \quad \textit{converges.}$$

Consequences of the heuristic

This leads to:

Conjecture

Suppose $L/\mathbb{Q}$ is an abelian extension with only finitely many subfields of degree 2, 3, or 5 over $\mathbb{Q}$.

Then for every elliptic curve $E/\mathbb{Q}$, we expect that $E(L)$ is finitely generated.

Consequences of the heuristic

This leads to:

Conjecture

Suppose $L/\mathbb{Q}$ is an abelian extension with only finitely many subfields of degree 2, 3, or 5 over $\mathbb{Q}$.

Then for every elliptic curve $E/\mathbb{Q}$, we expect that $E(L)$ is finitely generated.

For example, these conditions hold when L is:

- the $\hat{\mathbb{Z}}$-extension of $\mathbb{Q}$,
- the maximal abelian ℓ-extension of $\mathbb{Q}$, for $\ell \geq 7$.

Extensions and generalizations

Extension to other base fields: Suppose now that K is a number field and E is an elliptic curve over K. In this case there can be characters χ of $\mathrm{Gal}(\bar{K}/K)$ such that $L(E,\chi,1)$ vanishes because its root number is -1.

If for all *other* χ we have

$$\mathrm{Prob}[L(E,\chi,1)=0] \ll \left(\frac{C_E d_\chi}{\log(m_\chi)\varphi(m_\chi)}\right)^{\varphi(d_\chi)/4}$$

where d_χ is the order of χ and m_χ is the norm of its conductor, then we get a similar conclusion:

Extensions and generalizations

Conjecture

Suppose L/K is an abelian extension with only finitely many subfields of degree 2, 3, or 5.

Then for every elliptic curve E/K, if we exclude those characters with root number -1, then we expect $L(E,\chi,1) = 0$ for only finitely many other characters χ of $\mathrm{Gal}(L/K)$.

Extensions and generalizations

Conjecture

Suppose L/K is an abelian extension with only finitely many subfields of degree 2, 3, or 5.

Then for every elliptic curve E/K, if we exclude those characters with root number -1, then we expect $L(E, \chi, 1) = 0$ for only finitely many other characters χ of $\mathrm{Gal}(L/K)$.

Conjecture

Suppose $L/\mathbb{Q}$ is an abelian extension with only finitely many subfields of degree 2, 3, or 5.

Then for every elliptic curve E/L, we expect that $E(L)$ is finitely generated.

Extensions and generalizations

Studying p-Selmer: Instead of asking how often $L(E, \chi, 1) = 0$, we can ask how often $L(E, \chi, 1)/\Omega_E$ is divisible by (some prime above) p. By the Birch & Swinnerton-Dyer conjecture, this will tell us about the p-Selmer group $\mathrm{Sel}_p(E/L)$.

It seems reasonable to expect that if the θ-coefficents $c_{L,\chi,g}$ are not all the same (mod p), then they are equidistributed (mod p).

Extensions and generalizations

Studying p-Selmer: Instead of asking how often $L(E,\chi,1)=0$, we can ask how often $L(E,\chi,1)/\Omega_E$ is divisible by (some prime above) p. By the Birch & Swinnerton-Dyer conjecture, this will tell us about the p-Selmer group $\mathrm{Sel}_p(E/L)$.

It seems reasonable to expect that if the θ-coefficents $c_{L,\chi,g}$ are not all the same (mod p), then they are equidistributed (mod p).

For example, this leads to the following:

Conjecture

Let S be a finite set of rational primes, not containing p. Let L be the compositum of the cyclotomic $\mathbb{Z}_\ell$-extensions of $\mathbb{Q}$ for $\ell \in S$. If E is an elliptic curve over $\mathbb{Q}$ whose mod p representation is irreducible, then $\dim_{\mathbb{F}_p} \mathrm{Sel}_p(E/L)$ is finite.

The heuristic does *not* predict finite p-Selmer rank when S is infinite.