

Mathematical Proceedings of the Cambridge Philosophical Society

VOL. 145

NOVEMBER 2008

PART 3

Math. Proc. Camb. Phil. Soc. (2008), **145**, 513 © 2008 Cambridge Philosophical Society

513

doi:10.1017/S0305004108001588 Printed in the United Kingdom

First published online 11 June 2008

The effect of twisting on the 2-Selmer group

BY SIR PETER SWINNERTON-DYER

DPMMS, University of Cambridge, Wilberforce Road, Cambridge CB3 0WB, UK.

e-mail: hpfs100@dpms.cam.ac.uk

(Received 1 June 2007; revised 9 February 2008)

Abstract

Let Γ be an elliptic curve defined over \mathbf{Q} , all of whose 2-division points are rational, and let Γ_b be its quadratic twist by b . Subject to a mild additional condition on Γ , we find the limit of the probability distribution of the dimension of the 2-Selmer group of Γ_b as the number of prime factors of b increases; and we show that this distribution depends only on whether the 2-Selmer group of Γ has odd or even dimension.



1. Introduction

Let

$$\Gamma : y^2 = (x - c_1)(x - c_2)(x - c_3)$$

be an elliptic curve defined over \mathbf{Q} all of whose 2-division points are rational, where without loss of generality we can assume that all the c_i are integers. Denote by \mathcal{S} the set of places of \mathbf{Q} consisting of 2, ∞ and the odd primes for which Γ has bad reduction. If $b \in \mathbf{Z}$ is square-free and a unit at all places of \mathcal{S} , we denote by Γ_b the elliptic curve

$$\Gamma_b : y^2 = (x - bc_1)(x - bc_2)(x - bc_3)$$

which is the quadratic twist of Γ by b . We denote by \mathcal{B} the union of \mathcal{S} and the primes dividing b . In this paper we investigate the distribution of d_b , the dimension of the 2-Selmer group of Γ_b considered as an \mathbf{F}_2 vector space, as b varies. It will become clear that $d_b \bmod 2$ depends only on the images of b in the $\mathbf{Q}_v^*/\mathbf{Q}_v^{*2}$, where v runs through the places of \mathcal{S} . This is a special case of a result due to Kramer [3].

For the special elliptic curve $y^2 = x^3 - x$ a similar problem was solved by Heath–Brown [2], though he varies b over a different set to ours. He considers all square-free odd b satisfying $0 < b < X$ and lets X tend to infinity. (In his case $\mathcal{S} = \{2, \infty\}$, so that requiring b to be odd is the same as requiring b to be a unit at each prime in \mathcal{S} ; and d_b is even if $b \equiv 1$ or $3 \pmod 8$ and odd if $b \equiv 5$ or $7 \pmod 8$.) We on the other hand consider, for fixed N , the b which are the product of N randomly chosen primes, subject to the condition that the images of b in $\mathbf{Q}_p^*/\mathbf{Q}_p^{*2}$ for each p in \mathcal{S} have pre-assigned values. Since d_b only depends on the quadratic characters of these N primes with respect to one another and on the images of these primes in the $\mathbf{Q}_p^*/\mathbf{Q}_p^{*2}$ for p in \mathcal{S} , we are really only letting b vary over a finite set. We then let N tend to infinity and study the limit of the probability distribution of d_b . It appears that there is no easy way to transfer results from Heath–Brown’s set-up to ours, or vice versa. But it would be extraordinary if the distribution of d_b were not the same in the two cases – and in this sense Heath–Brown’s result for his special elliptic curve is compatible with ours.

In what follows we write

$$\alpha_k = 2^k \beta / \prod_{j=1}^k (2^j - 1) \quad (k = 0, 1, \dots) \tag{1}$$

where $\beta = \prod_{n=0}^{\infty} (1 - 2^{-2n-1})$. It was already proved in [2] that

$$\alpha_0 + \alpha_2 + \alpha_4 + \dots = 1, \quad \alpha_1 + \alpha_3 + \alpha_5 + \dots = 1,$$

so that both the α_v for even v and the α_v for odd v do give a probability distribution. The object of this paper is to prove the following theorem

THEOREM 1. *Suppose that none of the $(c_i - c_j)(c_i - c_k)$ are in \mathbf{Q}^{*2} . Fix the images of b in $\mathbf{Q}_p^*/\mathbf{Q}_p^{*2}$ for all p in \mathcal{S} and let $d \geq 2$ be an integer such that $d \equiv d_b \pmod 2$. Let $\pi_d^{(N)}$ be the probability that $d_b = d$ for that value of N ; then $\pi_d^{(N)} \rightarrow \alpha_{d-2}$ as $N \rightarrow \infty$.*

The condition in the first sentence holds if and only if none of the primitive 4-division points of Γ is rational. If it fails, the methods of this paper will still show that each $\pi_d^{(N)}$ tends to a limit as $N \rightarrow \infty$; but it appears that this limit will depend on Γ and describing it probably requires a substantial subdivision of cases.

Being essentially combinatorial, the methods of this paper still work if we replace \mathbf{Q} by an algebraic number field K . But for the analogue of Theorem 1 to hold requires the further condition that the equation $U^2 + V^2 = -1$ is not soluble in K . Thus for example it suffices to take K real.

2. Two vector-space lemmas

Suppose that ψ is a bilinear form on an \mathbf{F}_2 vector space. We shall say that ψ is *symmetric* (which in this situation is the same as *anti-symmetric*) if $\psi(x, y) = \psi(y, x)$ for all x, y ; and we shall say that ψ is *alternating* if also $\psi(x, x) = 0$. In this section we prove two lemmas which yield our description of 2-descent as a special case; they are stated and proved in a more general form only to simplify the notation. They have already appeared in [5] and [6]; we give a different (and simpler) proof here because we shall later need to appeal to some of the details of it.

LEMMA 1. *Let V be an \mathbf{F}_2 vector space and ψ a non-degenerate alternating bilinear form on V with values in \mathbf{F}_2 , and let W be maximal isotropic in V with respect to ψ . Then V can be decomposed as a direct sum $V = \oplus V_i$ where the V_i are mutually orthogonal, each*

V_i has dimension 2, each $V_i \cap W$ has dimension 1, and the restriction of ψ to any V_i is non-degenerate.

Proof. It follows from the existence of ψ that $\dim(V) = 2n$ and $\dim(W) = n$ for some n . Fix a base w_1, \dots, w_n for W . The map $V \rightarrow \mathbf{F}_2^n$ given by

$$v \mapsto \{\dots, \psi(w_i, v), \dots\}$$

has kernel W , so its image has order 2^n ; in other words, it is onto. Hence we can choose v_1, \dots, v_n so that $\psi(w_j, v_i)$ is 1 if $i = j$ and 0 otherwise. Moreover if $v = \sum \lambda_i v_i$ is in W for some λ_i in \mathbf{F}_2 then

$$0 = \psi(w_j, v) = \sum \lambda_i \psi(w_j, v_i) = \lambda_j.$$

Hence the w_i and v_j are linearly independent and therefore form a base for V . Adding elements of W to each v_i does not alter these properties, and

$$v'_j = v_j + \sum_{i=1}^{j-1} \lambda_i w_i \quad \text{implies} \quad \psi(v'_j, v_i) = \psi(v_j, v_i) + \lambda_i \text{ for } i < j.$$

Hence having chosen v_1, \dots, v_{j-1} we can choose v_j so that also $\psi(v_j, v_i) = 0$ for $i < j$. If V_i is the vector space generated by v_i and w_i then V is the direct sum of the V_i and they are mutually orthogonal. The remaining assertions about the V_i are obvious.

LEMMA 2. *Let the V_i be n vector spaces over \mathbf{F}_2 , each equipped with a non-degenerate alternating bilinear form ψ_i with values in \mathbf{F}_2 , and for each i let W_i be maximal isotropic in V_i . Denote by ψ the sum of the ψ_i , which is a non-degenerate alternating bilinear form on $V = \bigoplus V_i$, and let U be maximal isotropic in V with respect to ψ . Then there exist maximal isotropic subspaces $K_i \subset V_i$ such that $V = U \oplus K$ and*

$$W = (U \cap W) \oplus (K \cap W) \tag{2}$$

where $W = \bigoplus W_i$ and $K = \bigoplus K_i$.

Suppose also that on each V_i there is a function ϕ_i with values in \mathbf{F}_2 which satisfies

$$\phi_i(\xi + \eta) = \phi_i(\xi) + \phi_i(\eta) + \psi_i(\xi, \eta) \tag{3}$$

for any ξ, η in V_i , and let ϕ on V be the sum of the ϕ_i . Assume that ϕ is trivial on U and ϕ_i is trivial on W_i . Then we can further ensure that ϕ_i is trivial on K_i and therefore ϕ is trivial on K .

Proof. If any V_i has dimension greater than 2, by Lemma 1 we can decompose it as a direct sum of mutually orthogonal subspaces of dimension 2, on each of which the restriction of the bilinear form ψ_i is non-degenerate and each of which meets W_i in a subspace of dimension 1. This only reduces our freedom to choose the K_i , and the triviality of ϕ_i on the old K_i will follow from its triviality on the new and smaller K_i by (3). Thus we can assume that every V_i has dimension 2 and every W_i has dimension 1.

If \mathcal{N} is any subset of $\{1, \dots, n\}$ write $W_{\mathcal{N}} = \bigoplus_{i \in \mathcal{N}} W_i$ and similarly for $K_{\mathcal{N}}$. Choose \mathcal{M} maximal among the subsets \mathcal{N} for which $W_{\mathcal{N}} \cap U$ is trivial; such subsets \mathcal{N} do exist because the empty set is one of them. Let \mathcal{R} be the complement of \mathcal{M} . If r is in \mathcal{M} let α_r be the nontrivial element of W_r , so that $\phi_r(\alpha_r) = 0$ by hypothesis, and choose $K_r = W_r$. Thus $K_{\mathcal{M}} = W_{\mathcal{M}}$. If r is not in \mathcal{M} , by the maximality of \mathcal{M} we can find a nontrivial element γ_r of $W_{\mathcal{M} \cup \{r\}} \cap U$. Let β_r be the projection of γ_r on V_r ; then β_r must be the nontrivial element

of W_r , which incidentally implies that γ_r is unique. Let α'_r and $\alpha''_r = \alpha'_r + \beta_r$ be the other nontrivial elements of V_r . By (3) we have

$$\phi_r(\alpha''_r) - \phi_r(\alpha'_r) = \phi_r(\beta_r) + \psi(\alpha'_r, \beta_r) = 1.$$

Choose α_r to be that one of α'_r and α''_r which satisfies $\phi_r(\alpha_r) = 0$, and let K_r be the vector space generated by α_r . (If we drop the second paragraph in the statement of the Lemma, we can choose α_r to be either α'_r or α''_r .) The α_r for r in \mathcal{M} and the γ_r for r not in \mathcal{M} are linearly independent elements of W , so they span W ; it follows that the γ_r span $U \cap W$ and the α_r for r in \mathcal{M} span $K \cap W$. This proves (2). We have arranged that each $\phi_r(\alpha_r) = 0$, and $\phi_r(0) = 0$ because 0 is in U ; so $\phi = 0$ on each K_r .

We shall need the following remarks in §4. A necessary and sufficient condition for x to be in U is that it is orthogonal to every element of U . Choose a base u_1, \dots, u_n for U ; then $W_{\mathcal{N}} \cap U$ is trivial if and only if the equations

$$\sum_{i \in \mathcal{N}} \lambda_i \psi(u_j, w_i) = 0 \quad (j = 1, \dots, n)$$

with λ_i in \mathbf{F}_2 have no nontrivial solution. To test this, and therefore to find the candidates for \mathcal{M} , it is only necessary to know the $\psi(u_j, w_i)$. Again, suppose that r is not in \mathcal{M} and let w_r be the nontrivial element of W_r . By the maximality of \mathcal{M} , there must exist λ_{ir} in \mathbf{F}_2 for all i in \mathcal{M} such that $w_r + \sum_{i \in \mathcal{M}} \lambda_{ir} w_i$ is in U ; and these λ_{ir} are unique. They can be obtained by solving the equations

$$\psi(u_j, w_r) + \sum_{i \in \mathcal{M}} \lambda_{ir} \psi(u_j, w_i) = 0 \quad (j = 1, \dots, n);$$

hence to determine the λ_{ir} it is again sufficient to know the $\psi(u_j, w_r)$ and the $\psi(u_j, w_i)$ for i in \mathcal{M} .

3. An algorithm for 2-descent

In this section we recapitulate the most recent version of 2-descent on curves of the form Γ_b ; this was first described in [1], where full proofs can be found. To make the account intelligible appears to require a historical survey of how the process has developed. The basic version of 2-descent, which goes back to Fermat, is as follows. To any rational point (x, y) on Γ_b there correspond rational m_1, m_2, m_3 with $m_1 m_2 m_3 = m^2 \neq 0$ such that the three equations

$$m_i y_i^2 = x - bc_i \quad \text{for } i = 1, 2, 3 \tag{4}$$

are simultaneously soluble. We can multiply the m_i by non-zero squares, so that for example we can require them to be square-free integers; indeed one should really think of them as elements of $\mathbf{Q}^*/\mathbf{Q}^{*2}$, with a suitable interpretation of the equations which involve them. Denote by $\mathcal{C}(\mathbf{m})$ the curve given by the three equations (4), where $\mathbf{m} = (m_1, m_2, m_3)$. Looking for solutions of Γ_b is the same as looking for quadruples x, y_1, y_2, y_3 which satisfy (4) for some \mathbf{m} . For this purpose we need only consider the finitely many \mathbf{m} for which the m_i are units at all primes outside \mathcal{B} ; for if any m_i is divisible to an odd power by some prime p not in \mathcal{B} then Γ_b is already insoluble in \mathbf{Q}_p .

Provided one treats the m_i as elements of $\mathbf{Q}^*/\mathbf{Q}^{*2}$, the triples \mathbf{m} form an abelian group under componentwise multiplication:

$$\mathbf{m}' \times \mathbf{m}'' \longmapsto \mathbf{m}'\mathbf{m}'' = (m'_1 m''_1, m'_2 m''_2, m'_3 m''_3).$$

The \mathbf{m} for which $\mathcal{C}(\mathbf{m})$ is everywhere locally soluble form a finite subgroup, called the 2-Selmer group. This is computable, and it contains the group of those \mathbf{m} for which $\mathcal{C}(\mathbf{m})$ is actually soluble in \mathbf{Q} . This smaller group is $\Gamma_b(\mathbf{Q})/2\Gamma_b(\mathbf{Q})$, where $\Gamma_b(\mathbf{Q})$, the group of rational points on Γ_b , is the Mordell-Weil group of Γ_b . The quotient of the 2-Selmer group by this smaller group is ${}_2\text{III}$, the group of those elements of the Tate-Safarevic group which are killed by 2. The process of going from the curve Γ_b to the set of curves $\mathcal{C}(\mathbf{m})$, or the finite subset which is the 2-Selmer group, is a 2-descent, and the curves $\mathcal{C}(\mathbf{m})$ themselves are called 2-coverings.

We now put this process into more modern language. In what follows, italic capitals will always denote vector spaces over \mathbf{F}_2 and p will be either a finite prime or ∞ . Write

$$Y_p = \mathbf{Q}_p^*/\mathbf{Q}_p^{*2}, \quad Y_B = \bigoplus_{p \in \mathcal{B}} Y_p.$$

Let V_p denote the vector space of all triples (μ_1, μ_2, μ_3) with each μ_i in Y_p and $\mu_1\mu_2\mu_3 = 1$; and write $V_B = \bigoplus_{p \in \mathcal{B}} V_p$. This is the best way to introduce these spaces, because it preserves symmetry; but the reader should note that the prevailing custom in the literature is to define V_p as $Y_p \times Y_p$, which is isomorphic to the V_p defined above but not in a canonical way. Next, write $X_B = \mathfrak{o}_B^*/\mathfrak{o}_B^{*2}$ where \mathfrak{o}_B^* is the group of nonzero rationals which are units outside \mathcal{B} ; and let U_B be the image in V_B of the group of triples (m_1, m_2, m_3) such that the m_i are in X_B and $m_1m_2m_3 = 1$. It is known that the map $X_B \rightarrow Y_B$ is an embedding and $\dim U_B = 1/2 \dim V_B$; both these depend on the requirement that \mathcal{B} contains 2 and ∞ . Finally, if (x, y) is a point of Γ_b defined over \mathbf{Q}_p other than a 2-division point then the image of the map $(x, y) \mapsto (x - bc_1, x - bc_2, x - bc_3)$ is in V_p . This map is the Kummer map $\Gamma_b(\mathbf{Q}_p) \rightarrow V_p$, which is a homomorphism. We denote its image by W_p ; clearly W_p is the set of those triples \mathbf{m} for which (4) is soluble in \mathbf{Q}_p . We can supply the images of the 2-division points by continuity; for example the image of $(bc_1, 0)$ is

$$((c_1 - c_2)(c_1 - c_3), b(c_1 - c_2), b(c_1 - c_3)), \tag{5}$$

and the image of the point at infinity is the trivial triple $(1, 1, 1)$, which is also the product of the three triples like (5). The 2-Selmer group of Γ_b can now be identified with $U_B \cap W_B$ where $W_B = \bigoplus_{p \in \mathcal{B}} W_p$; for as was noted above, (4) is soluble at every prime outside \mathcal{B} if and only if the elements of \mathbf{m} are in X_B .

The next major step was taken by Tate. He introduced the bilinear form e_p on $V_p \times V_p$, defined for $\mathbf{m}' = (m'_1, m'_2, m'_3)$ and $\mathbf{m}'' = (m''_1, m''_2, m''_3)$ by

$$e_p(\mathbf{m}', \mathbf{m}'') = (m'_1, m''_1)_p + (m'_2, m''_2)_p + (m'_3, m''_3)_p.$$

Here $(u, v)_p$ is the additive Hilbert symbol with values in \mathbf{F}_2 , defined by

$$(u, v)_p = \begin{cases} 0 & \text{if } ux^2 + vy^2 = 1 \text{ is soluble in } \mathbf{Q}_p, \\ 1 & \text{otherwise.} \end{cases}$$

The Hilbert symbol is symmetric and additive in each argument:

$$(u, v)_p = (v, u)_p \quad \text{and} \quad (u_1u_2, v)_p = (u_1, v)_p + (u_2, v)_p.$$

Effectively it is a replacement for the quadratic residue symbol, with the advantage that it treats the places 2 and ∞ in just the same way as any other prime. Its key property is the Hilbert product formula

$$\sum_p (u, v)_p = 0,$$

where the sum is taken over all p including ∞ ; the left hand side is meaningful because $(u, v)_p = 0$ whenever p is an odd prime at which u and v are units. If p is an odd prime and u is prime to p , we shall write

$$\chi_p(u) = (u, p)_p$$

for the quadratic residue symbol with values in \mathbf{F}_2 .

The bilinear form e_p is non-degenerate and alternating on $V_p \times V_p$; thus $e_B = \sum_{p \in B} e_p$ is a non-degenerate alternating bilinear form on $V_B \times V_B$. It is known from class field theory that U_B is a maximal isotropic subspace of V_B . Tate showed that W_p is a maximal isotropic subspace of V_p , and therefore W_B is a maximal isotropic subspace of V_B . Thus

$$\dim W_B = \dim U_B = \frac{1}{2} \dim V_B; \tag{6}$$

and hence the 2-Selmer group of Γ_b can be identified with both the left and the right kernel of the restriction of e_B to $U_B \times W_B$.

For both aesthetic and practical reasons, one would like to show that this restriction is skew-symmetric – and preferably even that it is alternating. But to make such a statement meaningful we need an isomorphism between U_B and W_B ; and though they have the same structure as vector spaces it is not obvious that there is a natural isomorphism between them. The way round this obstacle was first shown in [1]. It requires the construction inside each V_p of a maximal isotropic subspace K_p such that $V_B = U_B \oplus K_B$ where $K_B = \bigoplus_{p \in B} K_p$. Assuming that such spaces K_p can be constructed, we let $t_B : V_B \rightarrow U_B$ be the projection along K_B and define

$$U'_B = U_B \cap (W_B + K_B), \quad W'_B = W_B / (W_B \cap K_B) = \bigoplus_{p \in B} W'_p$$

where $W'_p = W_p / (W_p \cap K_p)$. The map t_B induces an isomorphism

$$\tau_B : W'_B \longrightarrow U'_B,$$

and the bilinear function e_B induces a bilinear function

$$e'_B : U'_B \times W'_B \longrightarrow \mathbf{F}_2.$$

The bilinear functions $\theta^b_B : U'_B \times U'_B \rightarrow \mathbf{F}_2$ and $\theta^\sharp_B : W'_B \times W'_B \rightarrow \mathbf{F}_2$ defined respectively by

$$\theta^b_B : \mathbf{u}'_1 \times \mathbf{u}'_2 \mapsto e'_B(\mathbf{u}'_1, \tau_B^{-1}(\mathbf{u}'_2)) \quad \text{and} \quad \theta^\sharp_B : \mathbf{w}'_1 \times \mathbf{w}'_2 \mapsto e'_B(\tau_B \mathbf{w}'_1, \mathbf{w}'_2) \tag{7}$$

are symmetric. Here the images of $\mathbf{w}'_1 \times \mathbf{w}'_2$ under the second map and of $\tau_B \mathbf{w}'_1 \times \tau_B \mathbf{w}'_2$ under the first map are the same. The 2-Selmer group of Γ_b is isomorphic to both the left and the right kernel of e'_B , and hence also to the kernels of the two maps (7).

It is advantageous to choose the K_p so as to make U' and W' small, and to make θ^b and θ^\sharp alternating. Since $U'_B \supset U_B \cap W_B$, the best we can hope for is $U'_B = U_B \cap W_B$; we obtain this by satisfying the stronger requirement

$$W_B = (U_B \cap W_B) \oplus (K_B \cap W_B). \tag{8}$$

For suppose that (8) holds; then $W_B + K_B = (U_B \cap W_B) + K_B$ and it follows immediately that

$$U'_B = U_B \cap (W_B + K_B) = U_B \cap W_B. \tag{9}$$

The motivation for (8) is that we want to make $W_B \cap K_B$ as large as possible – that is, to choose K_B so that as much of it as possible is contained in W_B . But because K_B must be complementary to U_B , only the part of W_B which is complementary to $W_B \cap U_B$ is available for this purpose.

Since the 2-Selmer group $U_B \cap W_B$ is identified with the left and right kernels of each of the functions (7), if (9) holds then these functions are trivial and therefore alternating. The formal statement of all this is as follows.

LEMMA 3. We can choose maximal isotropic subspaces $K_p \subset V_p$ for each p in \mathcal{B} so that $V_B = U_B \oplus K_B$. We can further ensure that

$$W_B = (U_B \cap W_B) \oplus (K_B \cap W_B),$$

which implies $U'_B = U_B \cap W_B$. If so, the functions θ_B^\flat and θ_B^\sharp defined in (7) are trivial.

This is just Lemma 2 in a different notation, together with the fact that (8) implies (9). The K_p constructed in the proof of Lemma 3 (which are not unique) are explicitly described in the proof of Lemma 2. But the other properties of the K_p chosen in this way are not at all obvious. Hence it is advantageous to consider other recipes for choosing the K_p , for which (8) does not hold but we can still prove that the functions (7) are alternating.

For this purpose we write \mathcal{B} as the disjoint union of \mathcal{B}' and \mathcal{B}'' , where we shall always suppose that 2 and ∞ are both in \mathcal{B}' . For any odd prime p we denote by T_p the subset of V_p consisting of those triples (μ_1, μ_2, μ_3) with $\mu_1\mu_2\mu_3 = 1$ for which each μ_i is in $\mathfrak{o}_p^*/\mathfrak{o}_p^{*2}$ – that is, each μ_i is the image of a p -adic unit. The main point of the following theorem is that for p in \mathcal{B}'' it enables us to replace the definition of K_p used in the proof of Lemma 3 by the simpler choice $K_p = T_p$. How one chooses \mathcal{B}'' depends on the particular application which one has in mind.

THEOREM 2. Let \mathcal{B} be the disjoint union of $\mathcal{B}' \supset \{2, \infty\}$ and \mathcal{B}'' . We can construct maximal isotropic subspaces $K_p \subset V_p$ such that $V_B = U_B \oplus K_B$,

$$W_{\mathcal{B}'} = (U_{\mathcal{B}'} \cap W_{\mathcal{B}'}) \oplus (K_{\mathcal{B}'} \cap W_{\mathcal{B}'}) \tag{10}$$

and $K_p = T_p$ for all p in \mathcal{B}'' ; and (10) implies that $U'_{\mathcal{B}'} = U_{\mathcal{B}'} \cap W_{\mathcal{B}'}$. Moreover

$$U'_B = J_*U'_{\mathcal{B}'} \oplus \tau_B W'_{\mathcal{B}''} = J_*U'_{\mathcal{B}'} \oplus \left(\bigoplus_{p \in \mathcal{B}''} \tau_B W'_p \right), \tag{11}$$

and the restriction of θ_B^\flat to $J_*U'_{\mathcal{B}'} \times J_*U'_{\mathcal{B}'}$ is trivial.

If \mathcal{B}' also contains all the odd primes p such that the $v_p(c_i - c_j)$ are not all congruent mod 2, then we can choose the K_p for p in \mathcal{B}' so that also θ_B^\flat is alternating on U'_B .

The appearance of $J_*U'_{\mathcal{B}'}$ in and just after (11) calls for some explanation. Let \mathbf{u} be any element of $U_{\mathcal{B}'}$; then \mathbf{u} is in U_B . Moreover, for p in \mathcal{B}'' the image of \mathbf{u} in V_p is in $T_p = K_p$ and therefore in $K_p + W_p$; hence \mathbf{u} is in U'_B . In this way we define a map $U'_{\mathcal{B}'} \rightarrow U'_B$ which is clearly an injection and which we denote by J_* . Moreover $j_*\tau_B = \tau_B$ on $W'_{\mathcal{B}'} \subset W'_B$. To prove Theorem 2 we construct the K_p for p in \mathcal{B}' according to the recipe given in the proof of Lemma 2, which involves the functions ϕ_i . For \mathbf{m} in V_p we take $\phi_p(\mathbf{m})$ to be any one of the expressions

$$(m_i(c_i - c_j)(c_i - c_k), m_j(c_j - c_i)(c_j - c_k))_p,$$

whose values are easily shown to be equal. That the ϕ_p have the requisite properties is proved in [5].

4. *Some preliminary lemmas*

From now on we shall assume that Γ is fixed, as are the classes of b in $\mathbf{Q}_p^*/\mathbf{Q}_p^{*2}$ for each prime p in \mathcal{S} . We usually also suppose that we have fixed N , the number of primes p_1, \dots, p_N which divide b ; but we temporarily regard the p_i themselves as unknowns. Denote by G the multiplicative commutative group generated by $\mathfrak{o}_{\mathcal{S}}^*$ and the p_i . The components of a triplet \mathbf{u} will always be elements of G/G^2 ; if u_j is such a component we shall say that p_i divides u_j if and only if some (and therefore each) representative of u_j in G is divisible by an odd power of p_i . For p in \mathcal{S} the class of b in $\mathbf{Q}_p^*/\mathbf{Q}_p^{*2}$ determines W_p ; so we can fix the decomposition of V_p in accordance with Lemma 1. We shall denote it by $V_p = \bigoplus V_{p_i}$; but we shall not need to be more specific about the V_{p_i} . If however p divides b it will be useful to make the decomposition explicit. In this case W_p consists of $(1, 1, 1)$ and the three triples like (5), for these all lie in W_p and are distinct and we know from (6) that W_p has order 4. Following the construction in the proof of Lemma 1, where now $\psi = e_p$, we can take

$$\begin{aligned} \mathbf{w}_{p2} &= (b(c_2 - c_1), (c_2 - c_1)(c_2 - c_3), b(c_2 - c_3)), & \mathbf{v}_{p2} &= (v, v, 1), \\ \mathbf{w}_{p3} &= (b(c_3 - c_1), b(c_3 - c_2), (c_3 - c_1)(c_3 - c_2)), & \mathbf{v}_{p3} &= (v, 1, v), \end{aligned} \tag{12}$$

where v is a quadratic non-residue of p . Thus $V_p = V_{p2} \oplus V_{p3}$ where V_{p_i} is the vector space generated by \mathbf{v}_{p_i} and \mathbf{w}_{p_i} and $W_{p_i} = V_{p_i} \cap W_p$ is generated by \mathbf{w}_{p_i} . Note that if p_i is not in \mathcal{M} in the notation of the proof of Lemma 2 then α_{p_i} is \mathbf{v}_{p_i} .

In this way we decompose W as a direct sum of 1-dimensional subspaces; we temporarily write the nontrivial elements of these subspaces as $\mathbf{w}_1, \dots, \mathbf{w}_n$. Choose a base $\mathbf{u}_1, \dots, \mathbf{u}_n$ for $U_{\mathcal{B}}$. Once we fix the values of all the Hilbert symbols $(\alpha, \beta)_p$ where p is in \mathcal{B} and each of α, β runs through -1 and all the primes in \mathcal{B} , we shall know all the $e_p(\mathbf{u}_i, \mathbf{w}_j)$. By the remarks in the last paragraph of §2, these determine the possible \mathcal{M} ; and once \mathcal{M} is chosen, it determines the K_p and therefore the map $t_{\mathcal{B}}$ and finally the 2-Selmer group. The values of the Hilbert symbols $(\alpha, \beta)_p$ described above only depend on:

- (i) the classes of -1 and the p_i in the $\mathbf{Q}_p^*/\mathbf{Q}_p^{*2}$ for primes p in \mathcal{S} , where the product of the classes of the p_i must be the class of b ;
- (ii) the $\chi_{p_i}(p_j)$ and the $\chi_{p_i}(-1)$ for $i \neq j$, subject to the law of quadratic reciprocity or equivalently to the Hilbert product formula.

We call these values the *structure constants* associated with Γ_b ; we can choose $(1/2)N(N - 1) + (N - 1)(\#\mathcal{S})$ of them independently. To each of the allowable choices of the structure constants we assign the same probability. Thus if Γ, N and the images of b in the $\mathbf{Q}_p^*/\mathbf{Q}_p^{*2}$ for p in \mathcal{S} are given, it makes sense to talk about the probability distribution of d_b ; this gives a precise meaning to the statement that the p_i are randomly chosen primes. For this purpose we regard N as fixed. However, I have not been able to determine the probability distribution of the d_b for fixed N . In this paper I only address the easier problem of finding the limit of this distribution as $N \rightarrow \infty$.

Suppose that \mathcal{B} is the disjoint union of \mathcal{B}' and \mathcal{B}'' . Henceforth we shall assume that \mathcal{B}' contains \mathcal{S} , so that in particular 2 and ∞ are in \mathcal{B}' . Write M for the number of primes in \mathcal{B}' which divide b ; in this section we shall assume that M is fixed. The left kernel of $e_{\mathcal{B}}$ restricted to $U_{\mathcal{B}'} \times W_{\mathcal{B}'}$ is $U_{\mathcal{B}'} \cap W_{\mathcal{B}'}$, which consists of those elements of $U_{\mathcal{B}'}$ for which the corresponding 2-covering is locally soluble at each place in \mathcal{B}' . We shall denote it by $Z_{\mathcal{B}'}$; it is independent of the choice of the K_p , and its dimension, which we shall denote by $d(\mathcal{B}')$, only depends on the choice of the Hilbert symbols $(\alpha, \beta)_p$ as above. Indeed $Z_{\mathcal{B}'}$ only depends on the choice of those Hilbert symbols which do not depend on any p_i in \mathcal{B}'' . Provided \mathcal{B}'' is

not empty, $(1/2)M(M - 1) + M(\#S)$ of them can be chosen independently. Thus for fixed M the probability distribution of $d(\mathcal{B}')$ is well defined. In the notation of Theorem 2 $Z_{\mathcal{B}'}$ is also the kernel of the restriction of θ^b to $j_*U'_{\mathcal{B}'}$. In particular $Z_{\mathcal{B}}$ can be identified with the 2-Selmer group of Γ_b .

The reason for the next four lemmas is as follows. In Section 7 we study the effect on $Z_{\mathcal{B}'}$ of moving a prime q' from \mathcal{B}'' to \mathcal{B}' . Provided that q' is not the last prime in \mathcal{B}'' , $Z_{\mathcal{B}'}$ does not depend on q' . Thus we can regard q' as a random prime, and the probability distribution of $d(\mathcal{B}' \cup \{q'\})$ will only depend on $Z_{\mathcal{B}'}$. In the general case, which in the notation of Section 7 is when the elements of \mathcal{T} are independent, the probability distribution of $d(\mathcal{B}' \cup \{q'\})$ will only depend on that of $d(\mathcal{B}')$ and can be described explicitly. To complete the proof of the main theorem, we need to show that when M is large the probability of not being in the general case is small. When we use the “ O ” notation, the implied constant will depend only on $\#S$. It will turn out that if $Z_{\mathcal{B}'}$ is not in the general case it must contain at least one triple of the kind described in one of these four lemmas.

Let $\mathbf{u} = (u_1, u_2, u_3)$ be in $U_{\mathcal{B}}$ and let p be a prime dividing b . If the u_i are all units at p , the condition that the 2-covering associated with \mathbf{u} should be locally soluble at p is

$$\chi_p(u_1) = \chi_p(u_2) = 0. \tag{13}$$

If the u_i are not all units at p , suppose for example that u_1 is a unit at p and u_2, u_3 are not. Now the condition that the 2-covering associated with \mathbf{u} should be locally soluble at p is

$$\chi_p(u_1) = \chi_p((c_1 - c_2)(c_1 - c_3)), \quad \chi_p(bu_2/p^2) = \chi_p(c_1 - c_2). \tag{14}$$

Note that the second condition here involves the image of b in $\mathbf{Q}_p^*/\mathbf{Q}_p^{*2}$.

LEMMA 4. *The probability that $Z_{\mathcal{B}'}$ contains an element of U_S other than $(1, 1, 1)$ is $O(2^{-M})$.*

Proof. Let $\mathbf{u} = (u_1, u_2, u_3)$ be an element of U_S other than $(1, 1, 1)$. Without loss of generality we can suppose that $u_1 \neq 1$. For \mathbf{u} to be in $Z_{\mathcal{B}'}$ it is necessary that $\chi_p(u_1) = 0$ for each p in $\mathcal{B}' \setminus S$. But the only possible dependence relation among the $\chi_p(u_1)$ is that coming when $\mathcal{B}' = \mathcal{B}$ from the fact that the product of the p_i is b . Hence the probability that $\chi_p(u_1) = 0$ for all such p is at most 2^{1-M} , and so the probability that \mathbf{u} is in $Z_{\mathcal{B}'}$ is at most 2^{1-M} . There are $2^{2\#S} - 1$ elements of U_S other than $(1, 1, 1)$; so the probability that at least one of them is in $Z_{\mathcal{B}'}$ is less than $2^{2\#S+1-M}$.

LEMMA 5. *Suppose that $M < N$ and that $(c_1 - c_2)(c_1 - c_3)$ is not in \mathbf{Q}^{*2} . Then the probability that there is an element of the form $\mathbf{u} = (1, u_0, u_0)$ in $Z_{\mathcal{B}'}$ other than $(1, 1, 1)$ is $O((\frac{3}{4})^M)$.*

Proof. For any fixed u_0 in $X_{\mathcal{B}}$ other than 1, let \mathcal{B}^\sharp be the set of primes in $\mathcal{B}' \setminus S$ which divide u_0 , and let \mathcal{B}^\flat be the complement of \mathcal{B}^\sharp in $\mathcal{B}' \setminus S$. The conditions (13) now take the form $\chi_p(u_0) = 0$ for p in \mathcal{B}^\flat . Write $n = \#\mathcal{B}^\sharp$. The $M - n$ conditions obtained from (13) and the $2n$ conditions (14) are independent, because for p in \mathcal{B}^\flat the condition $\chi_p(u_0) = 0$ is the only one which involves p , for p in \mathcal{B}^\sharp the second condition (14) is the only one which involves $\chi_p(q)$ for any q in \mathcal{B}'' , and the various first conditions (14) are independent. So the probability of a particular \mathbf{u} being in $Z_{\mathcal{B}'}$ is at most 2^{-M-n} . For a given value of n there are

$2^{\#S} M! / n!(M - n)!$ possible \mathbf{u} ; so the probability that some such \mathbf{u} is in $Z_{B'}$ is at most

$$\sum_{n=0}^M \frac{M!}{n!(M - n)!} 2^{\#S - M - n} = 2^{\#S - M} \left(\frac{3}{2}\right)^M.$$

LEMMA 6. *Suppose that $M < N$; then the probability that there is an element $\mathbf{u} = (u_1, u_2, u_3)$ in $Z_{B'}$ with some u_i in X_S but not equal to 1 is $O((3/4)^M)$.*

Proof. Choose \mathbf{u} in $U_{B'}$ where to fix ideas we shall suppose that u_1 is in X_S and not equal to 1. Let B^{\sharp} be the set of primes in $B' \setminus S$ which divide u_2 and let B^{\flat} be the complement of B^{\sharp} in $B' \setminus S$. Write $n = \#B^{\sharp}$. If $n = 0$ the probability that \mathbf{u} is in $Z_{B'}$ is $O(2^{-M})$ by Lemma 4. If $n > 0$ we have $2(M - n)$ conditions coming from (13) with p in B^{\flat} and n conditions coming from the second condition (14) with p in B^{\sharp} . All these are independent, because for p in B^{\sharp} the corresponding condition (14) is the only one which involves $\chi_p(q)$ for any q in B' and for p in B^{\flat} the two conditions derived from (13) are the only ones which involve p , and they are clearly independent. For a given value of n and a given u_1 there are $2^{\#S} M! / n!(M - n)!$ possible \mathbf{u} ; so the probability that at least one such \mathbf{u} is in $Z_{B'}$ is at most

$$O(2^{-M}) + \sum_{n=1}^M \frac{M!}{n!(M - n)!} 2^{\#S - 2M + n} = O(2^{-M}) + 2^{\#S - 2M} (3^M - 1).$$

Since there are $2^{\#S} - 1$ possible u_1 , this completes the proof.

LEMMA 7. *Suppose that $M < N$; then the probability that there are distinct elements $\mathbf{u}' = (u'_1, u'_2, u'_3)$ and $\mathbf{u}'' = (u''_1, u''_2, u''_3)$ in $Z_{B'}$ with $u'_1 = u''_1$ and no component equal to 1 is $O((15/16)^M)$.*

Proof. Choose $\mathbf{u}', \mathbf{u}''$ in $U_{B'}$ with $u'_1 = u''_1$ and with no component equal to 1. By Lemma 6 we can assume that none of these components is in X_S . In general there are eight possible types of prime p which are in $B' \setminus S$: if p divides u_0 where $u_0 = u'_1 = u''_1$ then p divides one of u'_2 and u'_3 and also one of u''_2 and u''_3 , while if p does not divide u_0 then it divides both or neither of u'_2 and u'_3 and also both or neither of u''_2 and u''_3 . For each such prime p there are four conditions for local solubility at p derived from (13) and (14); but in general these will not all be independent. To express them without going into too much detail, we adopt the convention that A will denote a well-determined product of some of the $c_i - c_j$, which need not be the same from one appearance to the next.

If $p|u_0$ let i, j be such that u'_i and u''_j are units at p ; then we can write the conditions in the form

$$\chi_p(u'_i) = \chi_p(A), \quad \chi_p(u''_j) = \chi_p(A), \quad \chi_p(bu_0/p^2) = \chi_p(A). \tag{15}$$

In the third equation there are two distinct formulae for A , but it is possible for the quotient of their values to be a square; otherwise we obtain a further condition. This certainly happens when $i = 2, j = 1$; now the two formulae for A are $c_1 - c_2$ and $c_2 - c_1$, so that we obtain the additional condition $\chi_p(-1) = 0$. If p divides all of u'_2, u'_3, u''_1 and u''_3 then we can write the conditions in the form

$$\chi_p(u_0) = \chi_p(A), \quad \chi_p(bu'_3/p^2) = \chi_p(A), \quad \chi_p(u'_3u''_3/p^2) = \chi_p(A). \tag{16}$$

If for example p divides u'_2 and u'_3 but not u''_1 or u''_3 then we can write the conditions in the form

$$\chi_p(u_0) = 0, \quad \chi_p(u''_1) = 0, \quad \chi_p(bu'_2/p^2) = \chi_p(A). \tag{17}$$

If p divides none of u'_2, u'_3, u''_1, u''_3 then we can write the conditions in the form

$$\chi_p(u_0) = 0, \quad \chi_p(u'_3) = 0, \quad \chi_p(u''_3) = 0. \tag{18}$$

We begin with two special cases. The first is when $u'_3/u''_1 = u''_3/u'_2$ is in X_S ; now only four of the types of prime p listed above can occur. Let

- \mathcal{B}_1 be the set of primes in $\mathcal{B}' \setminus S$ which divide u_0, u'_2 and u''_3 ,
- \mathcal{B}_2 be the set of primes in $\mathcal{B}' \setminus S$ which divide u_0, u'_3 and u''_1 ,
- \mathcal{B}_3 be the set of primes in $\mathcal{B}' \setminus S$ which divide u'_2, u'_3, u''_1 and u''_3 ,
- \mathcal{B}_4 be the set of primes in $\mathcal{B}' \setminus S$ which divide none of u'_2, u'_3, u''_1 and u''_3 .

Write $n_i = \#\mathcal{B}_i$. We have $M = n_1 + n_2 + n_3 + n_4$ because $\mathcal{B}' \setminus S$ is the disjoint union of the \mathcal{B}_i . For any p in \mathcal{B}' at least two of the three conditions associated with p in the appropriate one of (15), (16) or (18) are independent; and the only way in which all three can fail to be independent is if the derived condition for $\chi_p(u'_3/u''_1)$ is trivial. But going back to the exact form of (14), we see that the first condition (15) implies

$$\begin{aligned} \chi_p(u'_3/u''_1) &= \chi_p((c_1 - c_2)(c_2 - c_3)) \text{ for } p \text{ in } \mathcal{B}_1, \\ \chi_p(u'_3/u''_1) &= \chi_p(u''_3/u'_2) = \chi_p((c_1 - c_2)(c_3 - c_1)) \text{ for } p \text{ in } \mathcal{B}_2; \end{aligned}$$

and the conditions (18) imply $\chi_p(u'_3/u''_1) = 0$ for p in \mathcal{B}_4 . Since

$$(c_1 - c_2)(c_2 - c_3) + (c_1 - c_2)(c_3 - c_1) = -(c_1 - c_2)^2$$

the two terms on the left cannot both be in \mathbf{Q}^{*2} . Hence for at least one of $\mathcal{B}_1, \mathcal{B}_2$ and \mathcal{B}_4 the three conditions associated with each p in that \mathcal{B}_i are indeed independent.

Thus we have retained at least $2M + \min(n_1, n_2, n_4)$ conditions. I claim that all but at most six of these are independent. To prove this, we choose a prime p_1^* in $\mathcal{B}' \setminus S$ which divides one but not the other of u_0 and u'_3 ; this is possible since u_0/u'_3 is not in X_S . If for example p_1^* divides u_0 choose a further prime p_2^* in $\mathcal{B}' \setminus S$ which divides u'_3 . If $p_2^* | u_0$ replace the condition $\chi_p(u_0) = 0$ for p in \mathcal{B}_4 by $\chi_p(u_0 u'_3) = 0$ and note that p_1^* divides $u_0 u'_3$ but p_2^* does not. Once we drop the conditions for which p is p_1^* or p_2^* the first and third conditions (15) for p in $\mathcal{B}_1 \cup \mathcal{B}_2$, the first and second conditions (16) for p in \mathcal{B}_3 , and the first and second conditions (18) for p in \mathcal{B}_4 , together with the set of nontrivial conditions on $\chi_p(u'_3/u''_1)$ just obtained, are independent. For if q is in \mathcal{B}' each of the third conditions (15) and the second conditions (16) involves a $\chi_p(q)$ which appears in no other condition; so they cannot be involved in any dependency conditions. Each of the remaining conditions, other than those in the final set, involves a $\chi_p(p_i^*)$ which appears in no other condition, and none of the $\chi_{p_i^*}(p)$ appear at all. Hence they too are not involved in any dependency conditions. The remaining conditions are clearly independent. Hence we have at least $\mu = 2M - 6 + \min(n_1, n_2, n_4)$ independent conditions. For given n_1, \dots, n_4 there are $2^{3\#S} M! / \prod (n_i!)$ possible pairs $\mathbf{u}', \mathbf{u}''$ of the kind we are currently considering; so the probability that some such pair is in $Z_{\mathcal{B}'}$ is at most

$$\sum \frac{M!}{\prod (n_i!)} 2^{3\#S - \mu} < \sum \frac{M!}{\prod (n_i!)} 2^{3\#S - 2M + 6} (2^{-n_1} + 2^{-n_2} + 2^{-n_4})$$

where each sum is taken over all acceptable n_1, \dots, n_4 . Each of the three sums on the right is equal to $3 \cdot 2^{3\#S + 6} (7/8)^M$.

The second special case is when $u'_2/u'_1 = u'_3/u'_3$ is in X_S . Again only four of the eight types of prime p listed above can occur. Let

- \mathcal{B}_1 be the set of primes in $\mathcal{B}' \setminus \mathcal{S}$ which divide u_0, u'_3 and u''_3 ,
- \mathcal{B}_2 be the set of primes in $\mathcal{B}' \setminus \mathcal{S}$ which divide u_0, u'_2 and u''_1 ,
- \mathcal{B}_3 be the set of primes in $\mathcal{B}' \setminus \mathcal{S}$ which divide u'_2, u'_3, u''_1 and u''_3 ,
- \mathcal{B}_4 be the set of primes in $\mathcal{B}' \setminus \mathcal{S}$ which divide none of u'_2, u'_3, u''_1 and u''_3 .

Write $n_i = \#\mathcal{B}_i$. This time the additional nontrivial conditions are the conditions $\chi_p(-1) = 0$ for p in \mathcal{B}_1 , which were derived just after (15). The remainder of the argument is essentially as before, except that we now have at least $2M - 6 + n_1$ independent conditions, and the probability that some pair $\mathbf{u}', \mathbf{u}''$ of this kind is in $Z_{\mathcal{B}'}$ is again $O((7/8)^M)$.

Now suppose that we are not in either of these special cases, so that each of the eight types of prime p listed above can potentially occur. For each such prime p we have three conditions, listed in the relevant one of (15) to (18), and as in the second special case we also have the condition $\chi_p(-1) = 0$ when p divides u_0, u'_2 and u''_3 . Using arguments similar to those described in detail in the proof for the first special case, we find that if we delete the conditions associated with any of a certain bounded number of primes then the remaining conditions are independent. We now use the identity

$$\sum \frac{M!}{\prod (n_i!)} 2^{-n_i} = \left(\frac{15}{2}\right)^M$$

where the sum is taken over all non-negative n_1, \dots, n_8 with $\sum n_i = M$; thus the probability that there is some such pair $\mathbf{u}', \mathbf{u}''$ in $Z_{\mathcal{B}'}$ is $O((15/16)^M)$.

5. Proof of the main theorem

Now suppose temporarily that the primes p in $\mathcal{B}' \setminus \mathcal{S}$ and the images of b in the corresponding $\mathbf{Q}_p^*/\mathbf{Q}_p^{*2}$ are known. Suppose further that \mathcal{B}'' is not empty. We next study how Z is changed when \mathcal{B}' is replaced by $\mathfrak{B} = \mathcal{B}' \cup \{q'\}$, where q' is in \mathcal{B}'' . An immediate observation is that $Z_{\mathfrak{B}} \cap U_{\mathcal{B}'} \subset Z_{\mathcal{B}'}$. Choose the K_p as in Theorem 2; then $U'_{\mathfrak{B}} = U'_{\mathcal{B}'} \oplus \tau_{\mathcal{B}}W_{q'}$ and $U'_{\mathcal{B}'} = Z_{\mathcal{B}'}$. With the obvious base for $U'_{\mathfrak{B}}$, the restriction of θ^b to $j_*U'_{\mathfrak{B}} \times j_*U'_{\mathfrak{B}}$ is given by a matrix of the form

$$\begin{pmatrix} 0 & A \\ {}^t A & C \end{pmatrix} \tag{19}$$

where the shape of A is $d(\mathcal{B}') \times 2$ and C is alternating. Being alternating, the matrix (19) must have even rank; and it is easy to see that its rank is 4 if A has rank 2, 0 if $A = 0$ and $C = 0$, and 2 otherwise. Thus $d(\mathfrak{B}) - d(\mathcal{B}')$ is -2 in the first case, 2 in the second case and 0 in the third case. Note that if $Z_{\mathfrak{B}}$ necessarily has two generators which involve q' we must be in the second case.

Suppose first that $M < N - 1$, so that \mathcal{B}'' contains at least two primes; then we can choose the q in the proofs of the previous three lemmas to be different to the q' of the last paragraph. If a particular row of the matrix (19), other than one of the last two, corresponds to $\mathbf{u} = (u_1, u_2, u_3)$, then it follows from (12) that the corresponding row of A is $(\chi_{q'}(u_2), \chi_{q'}(u_3))$. Assume that none of the $(c_i - c_j)(c_i - c_k)$ is in \mathbf{Q}^{*2} . Write $\theta = 15/16$ and denote by \mathcal{T} the set of $2d(\mathcal{B}') + 1$ elements consisting of the entries in A and one of the non-diagonal entries in C . It follows from the construction in the proof of Lemma 2 that for example $\tau_{\mathcal{B}}w_{q'2}$ in the notation of (12) has the form $(q'u_1, u_2, q'u_3)$ where the u_i are in $X_{\mathcal{B}'}$. Hence each

non-diagonal entry of C has the form $\chi_{q'}(bu/q')$ for some u in $X_{B'}$. Since all the entries in A have the form $\chi_{q'}(u')$ for some u' in $X_{B'}$, no dependency relation among the elements of \mathcal{T} can involve the non-diagonal element of C . If the elements of \mathcal{T} are independent when considered as functions of q' , then the probability distribution of $d(\mathfrak{B})$ is given by $\pi(d(B'), d(\mathfrak{B}))$ where

$$\pi(i, j) = \begin{cases} 2^{-2i-1} & \text{if } j = i + 2, \\ 3 \cdot 2^{-i} - 5 \cdot 2^{-2i-1} & \text{if } j = i, \\ 1 - 3 \cdot 2^{-i} + 2^{1-2i} & \text{if } j = i - 2, \\ 0 & \text{otherwise.} \end{cases} \tag{20}$$

If we revert to the situation where the members of $B' \setminus S$ are random primes, and denote by $P(d, M)$ the probability that $d(B') = d$ for some pre-assigned integer d , then

$$\sum_{d=0}^{\infty} \left| P(d, M + 1) - \sum_r \pi(r, d) P(r, M) \right| \tag{21}$$

is bounded by twice the probability that the elements of \mathcal{T} are not independent. But if the elements of \mathcal{T} are not independent, then $Z_{B'}$ must come under one of the cases considered in Lemmas 4 to 7. The probability of this happening is therefore $O(\theta^M)$.

If instead $M = N - 1$ then $Z_{\mathfrak{B}}$ contains the three elements like (5) which correspond to the three 2-division points. Hence the final sentence of the last paragraph but one applies, and we have

$$d_b - d(B') = d(\mathfrak{B}) - d(B') = 2.$$

If we exclude this last step, what we have here is an approximation, increasingly good as M increases, to one of two Markov processes. The states in each Markov process correspond to the values of $d(B')$, so the values are the even non-negative integers for one chain and the odd positive integers for the other. The transition probabilities are given by (20). Note that if α_i is given by (1) then

$$\alpha_j = \pi(j - 2, j)\alpha_{j-2} + \pi(j, j)\alpha_j + \pi(j + 2, j)\alpha_{j+2},$$

so that the α_j provide an invariant distribution in the sense of Markov chain theory. Provided $M < N - 1$, the process of replacing B' by \mathfrak{B} is a stochastic process whose limit can by abuse of language be described as one of the two Markov processes above. Because the $P(d, M)$ do not depend on N provided $M < N - 1$, and because we are only interested in behaviour as $N \rightarrow \infty$, we can now forget about N and the condition $M < N - 1$ and simply study the behaviour of $P(d, M)$ as $M \rightarrow \infty$. Under this simplification, Theorem 1 is equivalent to $P(d, M) \rightarrow \alpha_d$.

To make this argument precise, denote by $Q(r, M, n)$ the probability that the process is in state r after n steps if for each d the probability that it starts in state d is $P(d, M)$. Because (21) is $O(\theta^M)$ we have for each $r \geq 0$

$$\begin{aligned} & \sum_{d=0}^{\infty} |P(d, M + r + 1) - Q(d, M, r + 1)| \\ & \leq \sum_{d=0}^{\infty} \sum_{e=0}^{\infty} |\pi(e, d)\{P(e, M + r) - Q(e, M, r)\}| + O(\theta^{M+r}). \end{aligned}$$

By reversing the order of summation we see that the double sum is equal to

$$\sum_{e=0}^{\infty} |P(e, M+r) - Q(e, M, r)|.$$

Since by definition $Q(d, M, 0) = P(d, M)$, it follows that as $r \rightarrow \infty$

$$\limsup \sum_{d=0}^{\infty} |P(d, M+r) - Q(d, M, r)| \leq \sum_{r=0}^{\infty} O(\theta^{M+r}) = O(\theta^M). \quad (22)$$

In the standard terminology of Markov chain theory, each of the two Markov chains given by (20) is irreducible and aperiodic and has an invariant distribution given by the relevant α_d ; so the fundamental theorem of Markov chain theory (see for example [4, theorem 1.8.3.]) shows that

$$Q(d, M, r) \rightarrow \alpha_d \text{ as } r \rightarrow \infty.$$

Thus it follows from (22) that as $r \rightarrow \infty$

$$\limsup |P(d, M+r) - \alpha_d| = O(\theta^M).$$

This is enough to show that $P(d, M) \rightarrow \alpha_d$ as $M \rightarrow \infty$.

REFERENCES

- [1] J.-L. COLLIOT-THÉLÈNE, A. N. SKOROBOGATOV and SIR P. SWINNERTON-DYER. Hasse principle for pencils of curves of genus one whose Jacobians have rational 2-division points. *Invent. Math.* **134** (1998), 579–650.
- [2] R. HEATH-BROWN. The size of Selmer groups for the congruent number problem, II. *Invent. Math.* **118** (1994), 331–370.
- [3] K. KRAMER. Arithmetic of elliptic curves upon quadratic extension. *Trans. Amer. Math. Soc.* **264** (1981), 121–135.
- [4] J. R. NORRIS. *Markov Chains* (Cambridge, 1997).
- [5] A. N. SKOROBOGATOV and SIR P. SWINNERTON-DYER. 2-descent on elliptic curves and rational points on certain Kummer surfaces. *Adv. Math.* **198** (2005), 448–483.
- [6] SIR P. SWINNERTON-DYER. 2-descent through the ages. in *Ranks of Elliptic Curves and Random Matrix Theory* (ed. J. B. Conrey *et al.*), *London Math. Soc. Lecture Note Ser.* **341** (Cambridge University Press, 2007), 345–356.