

Algebra & Number Theory

Volume 9

2015

No. 10

**Polynomial values modulo primes on average and
sharpness of the larger sieve**

Xuancheng Shao



Polynomial values modulo primes on average and sharpness of the larger sieve

Xuancheng Shao

This paper is motivated by the following question in sieve theory. Given a subset $X \subset [N]$ and $\alpha \in (0, \frac{1}{2})$. Suppose that $|X \pmod{p}| \leq (\alpha + o(1))p$ for every prime p . How large can X be? On the one hand, we have the bound $|X| \ll_{\alpha} N^{\alpha}$ from Gallagher's larger sieve. On the other hand, we prove, assuming the truth of an inverse sieve conjecture, that the bound above can be improved (for example, to $|X| \ll_{\alpha} N^{O(\alpha^{2014})}$ for small α). The result follows from studying the average size of $|X \pmod{p}|$ as p varies, when $X = f(\mathbb{Z}) \cap [N]$ is the value set of a polynomial $f(x) \in \mathbb{Z}[x]$.

1. Introduction

For a positive integer N , denote by $[N]$ the set $\{1, 2, \dots, N\}$. The letter p is always used to denote a prime. The primary goal of this paper is to study upper bounds for the sizes of subsets $X \subset [N]$ occupying a small fraction of residue classes modulo many primes p . Gallagher's larger sieve [1971] provides such an upper bound.

Theorem 1.1 (larger sieve). *Let $X \subset [N]$ be a subset and \mathcal{P} be a set of primes. We have*

$$|A| \leq \frac{\sum_{p \in \mathcal{P}} \log p}{\sum_{p \in \mathcal{P}} |X \pmod{p}|^{-1} \log p - \log N}$$

whenever the denominator is positive.

See [Croot and Elsholtz 2004] for some variants of it and references therein for applications. We are particularly interested in the situation when $|X \pmod{p}| \leq \alpha p$ for some fixed $\alpha \in (0, 1)$, and whether the bound provided by the larger sieve is best possible.

Corollary 1.2 (larger sieve, special case). *Let $X \subset [N]$ be a subset and $\alpha \in (0, \frac{1}{2}]$. If $|X \pmod{p}| \leq (\alpha + o(1))p$ for every prime p , then $|X| \ll N^{\alpha+o(1)}$.*

MSC2010: primary 11N35; secondary 11R45, 11R09.

Keywords: Gallagher's larger sieve, inverse sieve conjecture, value sets of polynomials over finite fields.

This is easily deduced from [Theorem 1.1](#) by taking \mathcal{P} to be the set of primes up to $N^{\alpha+o(1)}$. When $\alpha > \frac{1}{2}$, the statement still holds, but is beaten by the bound $|X| \ll_{\alpha} N^{1/2}$ following from the large sieve [[Montgomery 1978](#)]. When $\alpha \leq \frac{1}{2}$, is the bound $|X| \ll N^{\alpha+o(1)}$ sharp? If X is the set of perfect squares up to N , then $|X| \sim N^{1/2}$ and X occupies $(p+1)/2$ residue classes (the quadratic residues) modulo any odd prime p . The question of whether this is the only type of sharp example is usually referred to as the inverse sieve conjecture, informally stated as follows.

Conjecture 1.3 (inverse sieve conjecture, rough form). *Let $X \subset [N]$ be a subset. If $|X \pmod{p}| \leq 0.9p$ for every prime p , then either one of the following two statements holds:*

- (1) *The cardinality of X is extremely small.*
- (2) *The set X possesses algebraic structure.*

See [Conjecture 4.1](#) below for one precise formulation of it. See also [[Croft and Lev 2007](#); [Helfgott and Venkatesh 2009](#); [Walsh 2012](#); [Green and Harper 2014](#)] for more discussions and evidences towards it.

Now assume that $\alpha < \frac{1}{2}$ is fixed. Motivated by the inverse sieve conjecture, we consider the sizes of $X \pmod{p}$ when X is the value set of a polynomial. For a polynomial $f(x) \in \mathbb{Z}[x]$ of degree $d \geq 1$, denote by $f_p \in \mathbb{F}_p[x]$ the reduction of f modulo p . Let $\alpha_p(f) = p^{-1}|f_p(\mathbb{F}_p)|$, the relative size of the value set of $f \pmod{p}$. Define $\alpha(f)$ to be the average of $\alpha_p(f)$ as p varies:

$$\alpha(f) = \lim_{Q \rightarrow \infty} \frac{1}{\pi(Q)} \sum_{p \leq Q} \alpha_p(f).$$

Note the trivial lower bounds $\alpha_p(f) \geq d^{-1}$ and $\alpha(f) \geq d^{-1}$.

Theorem 1.4 (polynomial values modulo primes on average). *Let $f \in \mathbb{Z}[x]$ be a polynomial of degree $d \geq 1$. Then*

$$\lim_{Q \rightarrow \infty} \frac{1}{\pi(Q)} \sum_{p \leq Q} \alpha_p(f)^{-1} \leq \tau(d), \tag{1-1}$$

where $\tau(d)$ is the number of positive divisors of d . Consequently, $\alpha(f) \geq \tau(d)^{-1}$.

Note that for $d \geq 3$ we always have $\tau(d) < d$. Hence, it is reasonable to conjecture that [Corollary 1.2](#) is not sharp whenever α is smaller than (and bounded away from) $\frac{1}{2}$. See the last section in [[Shao 2014](#)] for a preliminary discussion on the simplest case $d = 3$.

Theorem 1.5 (inverse sieve conjecture implies improved larger sieve). *Assume the truth of [Conjecture 4.1](#). Let $X \subset [N]$ be a subset and $\alpha \in (0, 1)$. Let $\epsilon \in (0, 1)$ be a parameter. If $|X \pmod{p}| \leq (\alpha + o(1))p$ for every prime p , then $|X| \ll_{\alpha, \epsilon} N^{1/d}$, where d is the smallest positive integer with $\tau(d) \geq (1 - \epsilon)\alpha^{-1}$.*

For example, when $\alpha = 0.49$, we obtain the upper bound $|X| \ll N^{1/4}$. On the other hand, since $\tau(d) \leq d^{C/\log \log d}$ for some constant $C > 0$, we get $|X| \ll_\alpha N^{\alpha^{c \log \log \alpha^{-1}}}$ for some constant $c > 0$, a huge improvement upon [Corollary 1.2](#) for small α (assuming the truth of the inverse sieve conjecture).

Remark 1.6. Instead of assuming that $|X \pmod p| \leq (\alpha + o(1))p$ for every prime p , knowing this on average over p in an appropriate sense is sufficient for our proof to follow. In this paper, we will focus on the model case in which we assume the pointwise estimate.

In the remainder of this introduction we discuss further the quantities $\alpha_p(f)$ and $\alpha(f)$. Note that (1-1) becomes an equality when $f(x) = x^d$. Indeed, in this case we have $\alpha_p(f) \sim (p - 1, d)^{-1}$, and thus the average of $\alpha_p(f)^{-1}$ is equal to

$$\frac{1}{\phi(d)} \sum_{a \in (\mathbb{Z}/d\mathbb{Z})^\times} (a - 1, d) = \tau(d).$$

Note, however, that in this case the average of $\alpha_p(f)$ is equal to

$$\alpha(f) = \frac{1}{\phi(d)} \sum_{a \in (\mathbb{Z}/d\mathbb{Z})^\times} (a - 1, d)^{-1},$$

which can be evaluated to $\phi(d)/d$ when d is squarefree (and is at least $(\phi(d)/d)^2$ for any d). Since $\phi(d)/d \gg (\log \log d)^{-1}$, the following construction provides polynomials f with smaller $\alpha(f)$.

Theorem 1.7 (polynomials with small value sets modulo primes). *Define a sequence of polynomials $\{f_n\}$ by*

$$f_1(x) = x^2, \quad f_{n+1}(x) = (f_n(x) + 1)^2.$$

Then $\alpha_p(f_n) = a_n$ provided that $p > 2f_{n-1}(0) + 2$ when $n > 1$, where the sequence $\{a_n\}$ is defined by

$$a_1 = \frac{1}{2}, \quad a_{n+1} = a_n - \frac{1}{2}a_n^2.$$

Moreover, we have $a_n \leq 2n^{-1}$ for each n .

See [Remark 2.5](#) below for the reasoning behind this construction of f_n . It is easy to see heuristically why one expects the relation $a_{n+1} = a_n - \frac{1}{2}a_n^2$. Indeed, if we model the value set of $f_n \pmod p$ as a random subset $S \subset \mathbb{F}_p$ with each element $s \in \mathbb{F}_p$ chosen in S with probability a_n independently at random, then for every quadratic residue r , the probability that r can be written as $(s + 1)^2$ for some $s \in S$ is $2a_n - a_n^2$. Hence the expected size of the set $(S + 1)^2$ is $(2a_n - a_n^2)p/2 + O(1)$, as desired.

Since $\deg f_n = 2^n$, we have $\alpha(f_n) \ll (\log(\deg f_n))^{-1}$. We do not know whether this is the best example or whether the bound for $\alpha(f)$ in [Theorem 1.4](#) is sharp. See [Section 6B](#) for more discussion of this.

The investigation of $\alpha_p(f)$ for a fixed prime p has a long history (see [Birch and Swinnerton-Dyer 1959; Cohen 1970]), and explicit formulae for $\alpha_p(f)$ are known in terms of the proportion of fixed-point-free elements in a certain Galois group (see Lemma 5.1 and the remark following). Not surprisingly, the quantity $\alpha(f)$ can also be evaluated in terms of a certain Galois group, and this is recorded in Proposition 6.1. Due to a lack of understanding of the relevant Galois groups, our lower bound for $\alpha(f)$ is instead obtained by studying the number of solutions to $f(x) \equiv f(y) \pmod{p}$ on average as p varies (see Section 2), and it is for this reason that the average of $\alpha_p(f)^{-1}$ naturally shows up.

A related line of work is on classifying those polynomials $f \in \mathbb{F}_p[x]$ for which $\alpha_p(f)$ is close to the lower bound d^{-1} (for a fixed p). In particular, results in [Gomez-Calderon and Madden 1988] imply that $\alpha_p(f) \geq 2d^{-1} + o(1)$ whenever $p \not\equiv \pm 1 \pmod{d}$.

The rest of this paper is organized as follows. In Section 2 we state a general and quantitative version of Theorem 1.4 for polynomials over arbitrary number fields, and outline the proof strategy, with the details given in Section 3. In Section 4 we state a precise form of the inverse sieve conjecture and deduce Theorem 1.5. In Section 5, Theorem 1.7 is proved by computing relevant Galois groups. Finally, in Section 6, we make some further remarks concerning the larger sieve as well as the quantity $\alpha(f)$.

2. Statement of results and proof strategy

Notation. For a number field K , we denote by \mathbb{O}_K its ring of integers and by Δ_K its (absolute) discriminant. For a prime ideal \mathfrak{p} in \mathbb{O}_K , we use $\kappa_{\mathfrak{p}}$ to denote the residue field $\mathbb{O}_K/\mathfrak{p}$ and $N(\mathfrak{p}) = |\kappa_{\mathfrak{p}}|$ to denote the norm of \mathfrak{p} . For a polynomial f with coefficients in \mathbb{O}_K , we use $f_{\mathfrak{p}}$ to denote its reduction modulo \mathfrak{p} .

As indicated in the introduction, we are mainly interested in studying the sizes of value sets of $f_{\mathfrak{p}}$ for one-variable polynomials f .

Definition 2.1. Let K be a number field, and let $f(x) \in \mathbb{O}_K[x]$ be a polynomial. For any prime ideal \mathfrak{p} in \mathbb{O}_K , define

$$\alpha_{\mathfrak{p}}(f) = N(\mathfrak{p})^{-1} |f_{\mathfrak{p}}(\kappa_{\mathfrak{p}})|.$$

This will be studied via the related quantity that measures the number of solutions of $g_{\mathfrak{p}} = 0$ for (multivariable) polynomials g .

Definition 2.2. Let K be a number field and let $g(\underline{X}) \in \mathbb{O}_K[\underline{X}]$ be a polynomial in n variables $\underline{X} = (X_1, \dots, X_n)$. For any prime ideal \mathfrak{p} in \mathbb{O}_K , define

$$m_{\mathfrak{p}}(g) = N(\mathfrak{p})^{-(n-1)} |\{\underline{X} \in \kappa_{\mathfrak{p}}^n : g_{\mathfrak{p}}(\underline{X}) = 0\}|.$$

To make our result quantitative, we also need a notion that measures the sizes of the coefficients of a polynomial.

Definition 2.3 (heights). Let K be a number field, and let $g(x) \in \mathbb{O}_K[\underline{X}]$ be a polynomial. We define its (absolute logarithmic) height $h(g)$ to be the sum

$$h(g) = \sum_v \max_a \log |a|_v, \tag{2-1}$$

where the sum is over all places v of K and the maximum is taken over all coefficients a of g .

Here $|a|_v$ is the normalized absolute value, so that it does not depend on the choice of the field K . For example, when $f \in \mathbb{Z}[x]$ is primitive, the height $h(f)$ is the logarithm of the (usual archimedean) absolute value of the largest coefficient of f . See [Hindry and Silverman 2000] for basic properties of the height function.

Theorem 2.4. *Let K be a number field and $f \in \mathbb{O}_K[x]$ be a polynomial of degree $d \geq 1$. Let $g \in \mathbb{O}_K[x, y]$ be the polynomial defined by $g(x, y) = f(x) - f(y)$. Let $s(g)$ be the number of irreducible factors of g in $K[x, y]$. Then for any $Q \geq 2$ we have*

$$\sum_{N(\mathfrak{p}) \leq Q} m_{\mathfrak{p}}(g) = s(g) \sum_{N(\mathfrak{p}) \leq Q} 1 + O(Q \exp(-c\sqrt{\log Q}) + h(g))$$

for sufficiently small $c = c(K, d) > 0$.

Proof of Theorem 1.4 assuming Theorem 2.4. First, by an application of the Cauchy-Schwarz inequality, we have $\alpha_{\mathfrak{p}}(f) \geq m_{\mathfrak{p}}(g)^{-1}$. It thus suffices to show that

$$s(g) \leq \tau(d) \tag{2-2}$$

when $K = \mathbb{Q}$. This follows by considering the homogeneous part of degree d of $g(x, y) = f(x) - f(y)$. Indeed, since this homogeneous part is $a(x^d - y^d)$ for some $a \neq 0$, it factors into $\tau(d)$ irreducible factors over \mathbb{Q} (which are cyclotomic polynomials), and thus $g(x, y)$ can be factored into at most $\tau(d)$ irreducible factors over \mathbb{Q} . □

Remark 2.5. The argument above motivates the choice of f_n in Theorem 1.7. Indeed, if a polynomial f is highly decomposable, in the sense that f is the composition of many polynomials (each of which has degree at least 2), then $g(x, y) = f(x) - f(y)$ will necessarily have many irreducible factors, which should lead to small values of $\alpha_{\mathfrak{p}}(f)$. In Proposition 5.3 we deduce another consequence of Theorem 2.4, that indecomposable polynomials have large values of $\alpha_{\mathfrak{p}}$.

We will in fact prove the following more general result, of which Theorem 2.4 is a special case. Two polynomials g_1, g_2 are said to be equivalent if they are scalar multiples of each other. Recall also that $g \in K[\underline{X}]$ is said to be absolutely

(or geometrically) irreducible if it is irreducible and remains irreducible over the algebraic closure of K .

Theorem 2.6 (average number of solutions modulo primes). *Let K be a number field and $g(\underline{X}) \in \mathbb{O}_K[\underline{X}]$ be a polynomial in n variables of total degree $d \geq 1$. Let $s(g)$ be the number of nonequivalent irreducible factors of g in $K[\underline{X}]$. Let L be a Galois extension of K such that g factors into absolutely irreducible factors in $L[\underline{X}]$. Let $C = C(K, n, d) > 0$ be sufficiently large. If $Q \geq \exp(C(\log \Delta_L)^2)$, then*

$$\sum_{N(\mathfrak{p}) \leq Q} m_{\mathfrak{p}}(g) \log N(\mathfrak{p}) = s(g)Q - t(g) \frac{Q^{\beta_0}}{\beta_0} + O(Q \exp(-c\sqrt{\log Q}) + h(g) + \log \Delta_L) \quad (2-3)$$

for sufficiently small $c = c(K, n, d) > 0$, where $t(g) \in [0, s(g)]$, and the second term appears only if the Dedekind zeta function ζ_L has a Siegel zero $\beta_0 \in (\frac{1}{2}, 1)$. Consequently, for $Q \geq \exp(C(\log \Delta_L)^2)$ we have

$$\sum_{N(\mathfrak{p}) \leq Q} m_{\mathfrak{p}}(g) \leq s(g) \sum_{N(\mathfrak{p}) \leq Q} 1 + O(Q \exp(-c\sqrt{\log Q}) + h(g) + \log \Delta_L). \quad (2-4)$$

The bounds for the error terms stem from a quantitative version of the Chebotarev density theorem in [Lagarias and Odlyzko 1977]. Assuming the truth of the generalized Riemann hypothesis (GRH) for ζ_L , we can get a much better error term $O(Q^{1/2}(\log \Delta_L + [L : \mathbb{Q}] \log Q))$, and of course without the Siegel zero term. The unconditional error term, however, is already enough for our application.

Proof of Theorem 2.4 assuming Theorem 2.6. We show that $g(x, y) = f(x) - f(y)$ factors into absolutely irreducible factors over $L = K(\mu_d)$, where μ_d is the group of d -th roots of unity. Indeed, since the homogeneous part of degree d of $f(x) - f(y)$ is $a(x^d - y^d)$ for some nonzero $a \in K$, it factors over L into linear factors. Thus, there is a factorization

$$f(x) - f(y) = \prod_{i=1}^r g_i(x, y)$$

of $f(x) - f(y)$ into absolutely irreducible factors g_1, g_2, \dots, g_r , such that the top-degree part of each g_i is defined over L . We claim that each g_i is defined over L as well. Suppose not. Without loss of generality, assume that some coefficient of g_1 does not lie in L . Let $\tau \in \text{Gal}(\overline{\mathbb{Q}}/L)$ be an automorphism that moves this coefficient. Let $\tau(g_1)$ be the polynomial obtained by applying τ to every coefficient of g_1 . Then $\tau(g_1)$ is also a factor of $f(x) - f(y)$, and thus $\tau(g_1)$ is equivalent to g_i for some $1 \leq i \leq r$. By our choice of τ , $\tau(g_1)$ must be equivalent to g_i for some $i > 1$, and thus g_1 and g_i have equivalent top-degree parts. This contradicts the fact that $x^d - y^d$ has no repeated factors.

Now that the potential Siegel zero β_0 of ζ_L depends only on K and d , the Siegel zero term in (2-3) can be absorbed into the error term, and the conclusion follows easily from partial summation. \square

Remark 2.7. In the argument above we used the fact that polynomials of the form $f(x) - f(y) \in K[x, y]$ factor into absolutely irreducible factors in $L[x, y]$, with $L = K(\mu_d)$. For a general polynomial $g(\underline{X}) \in K[\underline{X}]$ of height $h(g)$, it can be shown that one can take L with $[L : \mathbb{Q}] \leq C$ and $\Delta_L \leq C \exp(Ch(g))$ for some constant $C = C(K, n, d) > 0$. Thus the $\log \Delta_L$ factor in the error term can be removed, and the assumption on Q can be replaced by $Q \geq \exp(Ch(g)^2)$. We will, however, not need this relation between the size of L and the height $h(g)$.

Remark 2.8. The arguments used in proving Theorem 2.6 can be generalized to study the average behavior of $|V(\mathbb{F}_p)|$ as p varies, for any algebraic variety V defined over \mathbb{Z} . More precisely, let $m = \dim V$. Then the average of $p^{-m}|V(\mathbb{F}_p)|$ as p varies is equal to the number of irreducible components of V .

To finish this section, we sketch the proof of Theorem 2.6. By Lang–Weil, $m_p(g)$ is essentially the number of absolutely irreducible factors of g_p . Factor g into absolutely irreducible factors in $L[\underline{X}]$, and consider the natural action of the Galois group $G = \text{Gal}(L/K)$ on these factors. For almost all primes $\mathfrak{P} \subset \mathbb{O}_L$, these absolutely irreducible factors remain absolutely irreducible modulo \mathfrak{P} , and thus $m_p(g)$ is essentially the number of these factors which are defined over κ_p . This is equal to the number of fixed points of the Frobenius element associated with \mathfrak{P} . By the Chebotarev density theorem, these Frobenius elements are equidistributed in G as \mathfrak{P} varies. Hence the average of $m_p(g)$ is equal to the average number of fixed points of the G -action. By Burnside’s lemma, this is equal to the number of G -orbits, which is exactly the number of irreducible factors $s(g)$ of g . In carrying out this procedure some additional effort is needed to keep track of the explicit dependence on the height of g .

3. Proof of Theorem 2.6

In this section we prove Theorem 2.6. The implied constants appearing in this section are always allowed to depend on K, n, d .

Factor (g) into principal prime ideals in $L[\underline{X}]$:

$$(g) = (g_1)^{e_1} (g_2)^{e_2} \cdots (g_r)^{e_r},$$

where $g_i \in L[\underline{X}]$ is absolutely irreducible, and g_i, g_j are not equivalent when $i \neq j$. Let G be the Galois group $\text{Gal}(L/K)$. For any $1 \leq i \leq r$ and any $\xi \in G$, let $\xi(g_i)$ be the polynomial obtained by applying ξ to all coefficients of g_i . Since $\xi(g_i)$ is also a factor of g , $\xi(g_i)$ is equivalent to g_j for some $1 \leq j \leq r$. Hence ξ acts on

$\{(g_1), \dots, (g_r)\}$ by sending (g_i) to $(\xi(g_i))$. In this way we obtain a G -action on $\{(g_1), \dots, (g_r)\}$.

Lemma 3.1 (Galois descent). *Let E be any field and F be a Galois extension of E . Let $h \in F[\underline{X}]$ be a polynomial. The following two statements are equivalent:*

- (1) *The ideal $(h) \subset F[\underline{X}]$ is fixed by every element of $G = \text{Gal}(F/E)$.*
- (2) *The ideal (h) is defined over E . In other words, there exists a scalar $\alpha \in F^\times$ such that $\alpha h \in E[\underline{X}]$.*

Proof. This is a standard result in the theory of Galois descent. For completeness, we give a proof here. Clearly (2) implies (1). Now assume that (1) holds, so that for each $\xi \in G$, we have $\xi(h) = c_\xi h$ for some $c_\xi \in F^\times$. The scalars $\{c_\xi : \xi \in G\}$ form a 1-cocycle $G \rightarrow F^\times$, and thus by Hilbert’s Theorem 90 we have $c_\xi = \alpha/\xi(\alpha)$ for some $\alpha \in F^\times$. Now that $\xi(h) = \alpha h/\xi(\alpha)$, we conclude that $\xi(\alpha h) = \alpha h$ for each $\xi \in G$. Thus $\alpha h \in E[\underline{X}]$, as desired. □

Lemma 3.2. *Let the notation be as above. The number of orbits of the G -action on $\{(g_1), (g_2), \dots, (g_r)\}$ is equal to $s(g)$.*

Proof. Let $\mathcal{H} = \{h_1, h_2, \dots, h_s\}$ be the set of nonequivalent irreducible factors of g (well defined up to scalars in K), where $s = s(g)$. We construct a bijection between the set of orbits and \mathcal{H} .

Let $\mathcal{O} \subset \{(g_1), (g_2), \dots, (g_r)\}$ be a G -orbit, and let h be the product of those g_i with $(g_i) \in \mathcal{O}$. We claim that (h) is defined over K , and moreover (h) is a prime ideal in $K[\underline{X}]$ (hence $(h) = (h_j)$ for some $1 \leq j \leq s$). In fact, since any $\xi \in G$ permutes the factors in \mathcal{O} , the ideal (h) is fixed by ξ . By Lemma 3.1, the ideal (h) is defined over K . Now let $h' \in K[\underline{X}]$ be a factor of h (with positive degree), and let $\mathcal{O}' \subset \mathcal{O}$ be the set of those $(g_i) \in \mathcal{O}$ dividing h' . For any $(g_i) \in \mathcal{O}'$ and any $\xi \in G$, $\xi(g_i)$ is also a factor of h' and thus $(\xi(g_i)) \in \mathcal{O}'$. This shows that G preserves \mathcal{O}' , and thus $\mathcal{O}' = \mathcal{O}$ and $(h') = (h)$. Hence (h) is a prime ideal.

Conversely, let $h_j \in \mathcal{H}$ be an irreducible factor of g , and let \mathcal{O} be the set of those (g_i) dividing h_j . We claim that \mathcal{O} is a G -orbit, and moreover the product of those ideals in \mathcal{O} is equal to (h_j) . In fact, for any $\xi \in G$ and $(g_i) \in \mathcal{O}$, the polynomial $\xi(g_i)$ is also a factor of h_j . Hence G preserves \mathcal{O} . If $\mathcal{O}' \subset \mathcal{O}$ is a G -orbit, the argument above shows that the product of the ideals in \mathcal{O}' is defined over K . Hence $\mathcal{O}' = \mathcal{O}$ by the irreducibility of h_j . Finally, the argument above also shows that the product of the ideals in \mathcal{O} is defined over K , and is thus equal to (h_j) . □

The following lemma shows that the heights of the factors g_i are controlled by the height of g . Note that the height $h(g_i)$ depends only on the ideal (g_i) since two equivalent polynomials have the same height.

Lemma 3.3 (Gelfond’s inequality). *Let the notation be as above. Then $h(g_i) \leq h(g) + C$ for some constant $C = C(K, n, d) > 0$.*

Proof. See Proposition B.7.3 in [Hindry and Silverman 2000]. □

Let \mathfrak{p} be a prime in \mathbb{O}_K and \mathfrak{P} be a prime in \mathbb{O}_L lying above \mathfrak{p} . For each $1 \leq i \leq r$, let $(g_i) \pmod{\mathfrak{P}}$ be the ideal in $\kappa_{\mathfrak{P}}[\underline{X}]$ obtained by reduction modulo \mathfrak{P} . The following lemma will be used to ensure that $(g_i) \pmod{\mathfrak{P}}$ remains absolutely irreducible for all but finitely many \mathfrak{P} .

Lemma 3.4 (Noether). *Let n, d be positive integers. There exist polynomials ℓ_1, \dots, ℓ_m with integral coefficients depending only on n and d in variables $A_{i_1 \dots i_n}$ ($i_1 + \dots + i_n \leq d$), such that the following statement holds. For any algebraically closed field \bar{F} , a polynomial $f \in \bar{F}[\underline{X}]$ in n variables of total degree at most d with*

$$f(x_1, \dots, x_n) = \sum_{i_1 + \dots + i_n \leq d} a_{i_1 \dots i_n} x_1^{i_1} \cdots x_n^{i_n}$$

is reducible over \bar{F} or has total degree less than d if and only if $\ell_j((a_{i_1 \dots i_n})) = 0$ for each $1 \leq j \leq m$.

Proof. See Theorem 2A in [Schmidt 1976]. □

Lemma 3.5. *Let the notation be as above. There exists a positive integer $E \leq C \exp(Ch(g))$, for some $C = C(K, n, d) > 0$, such that $(g_i) \pmod{\mathfrak{P}}$ is absolutely irreducible for each $1 \leq i \leq r$ whenever $\mathfrak{P} \nmid E$.*

Proof. It suffices to prove the statement for each individual i . Let ℓ_1, \dots, ℓ_m be the polynomials in Lemma 3.4 corresponding to the degree of g_i . After normalizing we may assume that some coefficient of g_i is equal to 1. Thus $h(a) \leq h(g_i)$ for every coefficient a of g_i , where $h(a)$ for $a \in L^\times$ is defined by

$$h(a) = \sum_v \max(\log |a|_v, 0).$$

Since g_i is absolutely irreducible, ℓ_j does not vanish at the coefficient vector of g_i for some $1 \leq j \leq m$; call this nonvanishing value $A \in L \setminus \{0\}$. Since all coefficients of g_i have heights bounded by $h(g_i)$, we have $h(A) = O(h(g_i) + 1) = O(h(g) + 1)$. Therefore, there exists a positive integer $E \leq C \exp(Ch(g))$ such that $A \pmod{\mathfrak{P}}$ is nonzero whenever $\mathfrak{P} \nmid E$. For these \mathfrak{P} , the absolute irreducibility of $g_i \pmod{\mathfrak{P}}$ follows from another application of Lemma 3.4. □

Remark 3.6. The qualitative version of this statement is a special case of a general result in algebraic geometry: if R is a domain with fraction field F and S is a domain finitely generated over R such that the F -algebra $S_F = F \otimes_R S$ is absolutely irreducible over F , then there is a nonempty open subset $U \subset \text{Spec}(R)$ such that the fiber algebra $S_u = k(u) \otimes_R S$ over $k(u)$, the residue field at u , is absolutely irreducible.

Let E be the positive integer from Lemma 3.5. After enlarging E if necessary (but still with $E \leq C \exp(C h(g))$), we may assume that $g_1 \pmod{\mathfrak{P}}, \dots, g_r \pmod{\mathfrak{P}}$ are pairwise inequivalent whenever $\mathfrak{P} \nmid E$.

Let $\mathfrak{p} \nmid E$ be a prime in \mathbb{O}_K and \mathfrak{P} be a prime in \mathbb{O}_L lying above \mathfrak{p} . The decomposition group $G_{\mathfrak{P}} = \text{Gal}(\kappa_{\mathfrak{P}}/\kappa_{\mathfrak{p}})$ acts on the factors $\{g_1 \pmod{\mathfrak{P}}, \dots, g_r \pmod{\mathfrak{P}}\}$ so that $\xi(g_i \pmod{\mathfrak{P}})$ is equivalent to $g_j \pmod{\mathfrak{P}}$ for any $\xi \in G_{\mathfrak{P}}$. Comparing this with the G -action on $\{g_1, \dots, g_r\}$ described at the beginning of this section, we see that they are compatible via the natural inclusion $G_{\mathfrak{P}} \hookrightarrow G$. In particular, both $G_{\mathfrak{P}}$ and G can be viewed as subgroups of the symmetric group on r elements.

For any conjugacy class $[\xi] \subset G$, let $s([\xi])$ be the number of fixed points of any element in $[\xi]$.

Lemma 3.7. *Let the notation be as above. If $\mathfrak{p} \nmid E$ and \mathfrak{p} is unramified in L , then $m_{\mathfrak{p}}(g) = s([\sigma_{\mathfrak{p}}]) + O(N(\mathfrak{p})^{-1/2})$, where $[\sigma_{\mathfrak{p}}]$ is the Frobenius conjugacy class associated to \mathfrak{p} .*

Proof. Let $h \in \{g_1, \dots, g_r\}$. Note that $\sigma_{\mathfrak{P}}$ fixes (h) if and only if $\sigma_{\mathfrak{P}}$ fixes $(h_{\mathfrak{P}})$, and this happens if and only if $(h_{\mathfrak{P}})$ is defined over $\kappa_{\mathfrak{p}}$, by Lemma 3.1. Hence $s([\sigma_{\mathfrak{p}}])$ is exactly the number of nonequivalent absolutely irreducible factors of $g_{\mathfrak{p}}$ in $\kappa_{\mathfrak{p}}[\underline{X}]$.

Now let $h_1, \dots, h_s, h_{s+1}, \dots, h_t$ be the nonequivalent irreducible factors of $g_{\mathfrak{p}}$ in $\kappa_{\mathfrak{p}}[\underline{X}]$, where h_1, \dots, h_s are absolutely irreducible, with $s = s([\sigma_{\mathfrak{p}}])$. Let $V(h_i)$ be the solution set $\{\underline{X} \in \kappa_{\mathfrak{p}}^n : h_i(\underline{X}) = 0\}$. Then

$$m_{\mathfrak{p}}(g) = N(\mathfrak{p})^{-(n-1)} \left(\sum_{i=1}^t |V(h_i)| + O\left(\sum_{i < j} |V(h_i) \cap V(h_j)| \right) \right). \tag{3-1}$$

For $i < j$, since h_i and h_j are nonequivalent, $V(h_i) \cap V(h_j)$ has dimension at most $n - 2$. Thus by the Lang–Weil bound [1954] we have

$$|V(h_i) \cap V(h_j)| \ll N(\mathfrak{p})^{n-2}.$$

(In fact the weaker Schwarz–Zippel estimate is enough here). On the other hand, for $1 \leq i \leq s$, the Lang–Weil bound gives

$$|V(h_i)| = N(\mathfrak{p})^{n-1} (1 + O(N(\mathfrak{p})^{-1/2}))$$

since h_i is absolutely irreducible, and for $s < i \leq t$ we have

$$|V(h_i)| = |V(h_i) \cap V(h_j)| \ll N(\mathfrak{p})^{n-2}$$

for some $s < j \leq t$, with h_j a Galois conjugate of h_i . Combining these estimates together in (3-1) we obtain

$$m_{\mathfrak{p}}(g) = s + O(N(\mathfrak{p})^{-1/2}),$$

as desired. □

We are now ready to evaluate the quantity

$$M_f(Q) = \sum_{N(\mathfrak{p}) \leq Q} m_{\mathfrak{p}}(g) \log N(\mathfrak{p}).$$

By [Lemma 3.7](#), we have

$$M_f(Q) = \sum_{\substack{N(\mathfrak{p}) \leq Q \\ \mathfrak{p} \text{ unramified in } L}} s([\sigma_{\mathfrak{p}}]) \log N(\mathfrak{p}) + O(Q^{1/2} \log Q + \log E + \log \Delta_L).$$

Since $E \leq C \exp(Ch(g))$, we have $\log E = O(h(g) + 1)$. Hence

$$M_f(Q) = \sum_C s(C) \psi_C(Q) + O(Q^{1/2} \log Q + h(g) + \log \Delta_L),$$

where the sum is over all conjugacy classes C in G , and

$$\psi_C(Q) = \sum_{\substack{N(\mathfrak{p}) \leq Q \\ \mathfrak{p} \text{ unramified in } L \\ [\sigma_{\mathfrak{p}}]^m = C}} \log N(\mathfrak{p}).$$

By (a quantitative version of) the Chebotarev density theorem [[Lagarias and Odlyzko 1977](#)], for $Q \geq \exp(C(\log \Delta_L)^2)$ we have

$$\psi_C(Q) = \frac{|C|}{|G|} Q - \frac{|C|}{|G|} \chi_0(C) \frac{Q^{\beta_0}}{\beta_0} + O(Q \exp(-c(\log Q)^{1/2})),$$

where the second term occurs only if the Dedekind zeta function ζ_L has a Siegel zero β_0 , and χ_0 is the real character of a one-dimensional representation of G for which the associated L -function has β_0 as a zero. It follows that

$$M_f(Q) = Q \cdot \frac{1}{|G|} \sum_{\xi \in G} s(\xi) - \frac{Q^{\beta_0}}{\beta_0} \cdot \frac{1}{|G|} \sum_{\xi \in G} s(\xi) \chi_0(\xi) + O(Q \exp(-c(\log Q)^{1/2}) + h(g) + \log \Delta_L).$$

By Burnside’s lemma and [Lemma 3.2](#), we have

$$\frac{1}{|G|} \sum_{\xi \in G} s(\xi) = s(g). \tag{3-2}$$

The equality (2-3) follows by setting

$$t(g) = \frac{1}{|G|} \sum_{\xi \in G} s(\xi) \chi_0(\xi). \tag{3-3}$$

By a change of summation, we can write

$$t(g) = \frac{1}{|G|} \sum_{i=1}^r \sum_{\xi \in G_i} \chi_0(\xi),$$

where $G_i \subset G$ is the subgroup of elements fixing (g_i) . Since χ_0 is a one-dimensional real character, the inner sum is either 0 or $|G_i|$, and in the latter case $\chi_0(\xi) = 1$ for all $\xi \in G_i$. Thus in the sum in (3-3) we may restrict to those ξ with $\chi_0(\xi) = 1$. Comparing this with (3-2), we obtain $t(g) \in [0, s(g)]$ as claimed.

Finally, the inequality (2-4) follows easily from (2-3) by dropping the Siegel zero term and partial summation.

4. Inverse sieve conjecture implies improved larger sieve

In this section we state a precise version of the inverse sieve conjecture and then prove Theorem 1.5. The implied constants here are always allowed to depend on α, ϵ .

Conjecture 4.1 (inverse sieve conjecture). *Let $X \subset [N]$ be a subset and let $\epsilon, \epsilon' > 0$ be real. Assume that for each parameter $Q \geq N^\epsilon$ we have*

$$\sum_{p \leq Q} \frac{|X \pmod p|}{p} \leq (1 - \epsilon')\pi(Q).$$

Then at least one of the following two situations happens:

- (1) (very small size) $|X| \ll_{\epsilon, \epsilon'} N^\epsilon$.
- (2) (algebraic structure) *There exists a polynomial $f(x) \in \mathbb{Q}[x]$ of degree $d \in [2, C]$ and height at most N^C such that $|X \cap f([N])| \geq C^{-1}|X|$, where $C = C(\epsilon, \epsilon')$ is a constant.*

Here, we say that a polynomial $f(x) \in \mathbb{Q}[x]$ has height at most H if $f(x) = A^{-1}f^*(x)$ for some positive integer $A \leq H$ and $f^* \in \mathbb{Z}[x]$ with all coefficients bounded by H in absolute value. This is slightly different from the notion of height used in the statement of Theorem 2.4, in that $h(f)$ is invariant under scalar multiplication but the notion here is not. Note that if a polynomial $f(x) \in \mathbb{Q}[x]$ has height at most H , then $h(f) \ll \log H$.

Remark 4.2. We make a few remarks explaining why some quantitative aspects of this conjecture are reasonable.

- The condition on X essentially says that X misses a positive proportion of residue classes modulo primes p on average, as soon as p exceeds a small positive power of N . With this assumption, we know from the large sieve that $|X| \ll N^{1/2}$ and from the larger sieve that $|X| \ll N^{\alpha+O(\epsilon)}$ if the upper bound $(1 - \epsilon')\pi(Q)$ is

replaced by $\alpha\pi(Q)$. Without the knowledge about $X \pmod p$ for $p \leq N^\epsilon$, one can essentially add to X any N^ϵ extra elements without violating the assumption, but one should still expect to see algebraic structure apart from these extra elements.

- The conclusion $|X \cap f([N])| \geq C^{-1}|X|$ is equivalent to the seemingly weaker one $|X \cap f(\mathbb{Q})| \geq C^{-1}|X|$, after a suitable modification of the polynomial f which does not increase its height too much. To see that the interval $[N]$ can be replaced by \mathbb{Z} , note that the set $J = \{n \in \mathbb{Z} : 1 \leq f(n) \leq N\}$ is the union of at most d intervals and has size at most dN . Since $X \cap f(\mathbb{Z}) = X \cap f(J)$, there is an interval $I \subset J$ with $|X \cap f(I)| \geq d^{-1}|X \cap f(\mathbb{Z})|$, and we may assume that $I \subset [N]$ after a translation. To see that $f(\mathbb{Z})$ can be replaced by $f(\mathbb{Q})$, note that if $f(x) \in \mathbb{Z}$ for some $x \in \mathbb{Q}$ then the denominator of x must divide some positive integer B depending on the coefficient of f . Then $f(\mathbb{Q}) \cap \mathbb{Z} \subset f^*(\mathbb{Z}) \cap \mathbb{Z}$, where f^* is defined by $f^*(x) = f(x/B)$.
- The conclusion that $f(\mathbb{Z})$ captures a positive proportion of X cannot be replaced by the stronger one that $f(\mathbb{Z})$ captures almost all of X . Indeed, it is possible for X to be the union of $f(\mathbb{Z})$ for several distinct polynomials f .

If $|X \pmod p| \leq \alpha p$ for small α , repeated applications of [Conjecture 4.1](#) allow us to strengthen it by requiring the degree d to be fairly large.

Proposition 4.3 (inverse sieve conjecture in the larger sieve regime). *Assume the truth of [Conjecture 4.1](#). Let $X \subset [N]$ be a subset. Let $\alpha \in (0, 1)$ and $\epsilon \in (0, \alpha)$ be real. Assume that $|X \pmod p| \leq (\alpha + o(1))p$ for each prime p . Then at least one of the following two situations happens:*

- (1) (very small size) $|X| \ll_\epsilon N^\epsilon$.
- (2) (algebraic structure) *There exists a polynomial $f(x) \in \mathbb{Q}[x]$ of degree $d \in [2, C]$ and height at most N^C such that $|X \cap f(\mathbb{Z})| \geq C^{-1}|X|$, where $C = C(\epsilon)$ is a constant. Moreover, we may ensure that $\tau(d) \geq (1 - \epsilon)\alpha^{-1}$.*

Proof. Suppose that $|X| \gg N^\epsilon$. We will apply [Conjecture 4.1](#) iteratively to construct a sequence of polynomials f_1, f_2, \dots, f_k and a sequence of sets $X_0 = X, X_1, X_2, \dots, X_k$, with $k = O(1)$, such that the following conditions hold:

- (1) $\deg f_i = d_i \in [2, C]$, and $\tau(d_1 d_2 \dots d_k) \geq (1 - \epsilon)\alpha^{-1}$;
- (2) the height of f_i is at most $N^{O(1)}$ for each $1 \leq i \leq k$;
- (3) $X_i \subset [N]$ and $|X_i| \gg |X_{i-1}|$ for each $1 \leq i \leq k$;
- (4) $f_i(X_i) \subset X_{i-1}$ for each $1 \leq i \leq k$.

Suppose first that these objects are constructed. Let $f = f_1 \circ f_2 \circ \dots \circ f_k$. By property (1), the degree d of f is $O(1)$ and satisfies $\tau(d) \geq (1 - \epsilon)\alpha^{-1}$. By

property (2), the height of f is $N^{O(1)}$. By property (3), we have $|X_k| \gg |X|$. By property (4), we have $f(X_k) \subset X \cap f([N])$. Hence

$$|X \cap f([N])| \geq |f(X_k)| \gg |X_k| \gg |X|,$$

as desired.

It thus remains to construct f_1, \dots, f_k and X_1, \dots, X_k . Suppose that f_j, X_j with $j < i$ are already chosen for some $i \geq 1$ satisfying the required properties (2)–(4) and $\deg f_j = d_j \in [2, C]$. We will construct f_i and X_i from these. Let $F = f_1 \circ \dots \circ f_{i-1}$ if $i > 1$ and let F be the identity map if $i = 1$. Let D be the degree of F . We may assume that $\tau(D) < (1 - \epsilon)\alpha^{-1}$, since we may stop the iteration otherwise. By property (4), we have $F(X_{i-1}) \subset X$.

Let $F = A^{-1}F^*$ with $A \leq N^C$ a positive integer and $F^* \in \mathbb{Z}[x]$ a polynomial whose coefficients are all bounded by N^C . Let $G \in \mathbb{Z}[x, y]$ be the polynomial defined by $G(x, y) = F^*(x) - F^*(y)$. Let $p \nmid A$ be a prime. For each $r \in \mathbb{Z}/p\mathbb{Z}$, let $v_p(r)$ be the number of $x \in \mathbb{Z}/p\mathbb{Z}$ with $F(x) \equiv r \pmod p$. Then

$$\begin{aligned} |X_{i-1} \pmod p| &\leq \sum_{r \in F(X_{i-1}) \pmod p} v_p(r) \\ &\leq |X \pmod p|^{1/2} \left(\sum_r v_p(r)^2 \right)^{1/2} \\ &\leq (\alpha + o(1))^{1/2} m_p(G)^{1/2} p, \end{aligned}$$

by Cauchy–Schwarz, the assumption that $|X \pmod p| \leq (\alpha + o(1))p$, and the definition of $m_p(G)$ in [Definition 2.2](#). For any $Q \geq N^\epsilon$, we then have

$$\begin{aligned} \sum_{p \leq Q} \frac{|X_{i-1} \pmod p|}{p} &\leq (\alpha + o(1))^{1/2} \sum_{p \leq Q} m_p(G)^{1/2} + O(\log A) \\ &\leq (\alpha + o(1))^{1/2} \pi(Q)^{1/2} \left(\sum_{p \leq Q} m_p(G) \right)^{1/2} + O(\log N). \end{aligned}$$

Now apply [Theorem 2.4](#) (and recall (2-2)) to obtain

$$\begin{aligned} \sum_{p \leq Q} \frac{|X_{i-1} \pmod p|}{p} &\leq [(\alpha + o(1))\tau(D)]^{1/2} \pi(Q) + O(Q \exp(-c(\log Q)^{1/2}) + Q^{1/2} \log N). \end{aligned}$$

Since $\tau(D) < (1 - \epsilon)\alpha^{-1}$, the first term above is at most $(1 - \epsilon/2)\pi(Q)$, and thus X_{i-1} satisfies the hypotheses in [Conjecture 4.1](#), with ϵ replaced by $\epsilon/3$ and N sufficiently large. Since $|X_{i-1}| \gg N^\epsilon$, we must be in the algebraic case. Let $f_i \in \mathbb{Q}[x]$ be a polynomial of degree $d_i \in [2, C]$ and height at most N^C such that $|X_{i-1} \cap f_i([N])| \gg |X_{i-1}|$, and let $X_i \subset [N]$ be chosen so that $f_i(X_i) \subset X_{i-1}$

and $|X_i| \gg |X_{i-1}|$. This completes the inductive construction. Finally, since the quantity $\tau(d_1 d_2 \cdots d_i)$ strictly increases with i , the process terminates after $O(1)$ iterations. \square

Proof of Theorem 1.5. Apply Proposition 4.3 to conclude that either $|X|$ is very small and we are done, or else there exists a polynomial $f(x) \in \mathbb{Q}[x]$ of degree $d \in [2, C]$ and height at most N^C such that $|X \cap f([N])| \geq C^{-1}|X|$. Moreover, we have $\tau(d) \geq (1 - \epsilon)\alpha^{-1}$. Hence

$$|X| \ll |X \cap f([N])| \leq |[N] \cap f([N])| \ll N^{1/d},$$

where the last inequality follows from a result of Walsh [2015], which removes the ϵ term from the exponent appearing in [Bombieri and Pila 1989; Heath-Brown 2002]. \square

5. Polynomials with small value sets modulo primes

In this section we prove Theorem 1.7. First we state a result connecting the quantity $\alpha_p(f)$ to a Galois group. For a polynomial $f(x) \in \mathbb{F}_p[x]$ of degree d , denote by R_f the set of roots in $\overline{\mathbb{F}_p}(t)$ of the polynomial $f(x) - t$. Define

$$G_f = \text{Gal}(\mathbb{F}_p(R_f)/\mathbb{F}_p(t)), \quad G_f^* = \text{Gal}(\overline{\mathbb{F}_p}(R_f)/\overline{\mathbb{F}_p}(t)).$$

In other words, G_f and G_f^* are the Galois groups of the splitting field of $f(x) - t$ over $\mathbb{F}_p(t)$ and $\overline{\mathbb{F}_p}(t)$, respectively. It is easy to see that G_f^* is a normal subgroup of G_f with G_f/G_f^* cyclic. In fact, G_f/G_f^* is isomorphic to $\text{Gal}(\mathbb{F}_p(R_f) \cap \overline{\mathbb{F}_p}/\mathbb{F}_p)$. For any subset $\Xi \subset G_f$, we use $\alpha(\Xi)$ to denote the proportion of elements in Ξ with at least one fixed point under the natural action on R_f .

Lemma 5.1 [Cohen 1970]. *Let $f(x) \in \mathbb{F}_p[x]$ be a polynomial of degree $d \geq 1$. Let σG_f^* be the coset which is the Frobenius generator of the cyclic quotient G_f/G_f^* , considered under its isomorphism with $\text{Gal}(\mathbb{F}_p(R_f) \cap \overline{\mathbb{F}_p}/\mathbb{F}_p)$. Then*

$$\alpha_p(f) = \alpha(\sigma G_f^*) + O_d(p^{-1/2}).$$

In particular, if $G_f = G_f^$ then*

$$\alpha_p(f) = \alpha(G_f) + O_d(p^{-1/2}).$$

Remark 5.2. In [Cohen 1970] this is deduced from a function field version of the Chebotarev density theorem. The Galois groups G_f and G_f^* above can be interpreted in terms of finite étale Galois coverings of $\mathbb{P}^1(\mathbb{F}_p)$. In this way Lemma 5.1 becomes a 0-dimensional special case of Deligne’s equidistribution theorem. See [Kowalski 2010] for an excellent survey on this topic. This function field version of the Chebotarev density theorem and related equidistribution results play an important

role in proving function field analogues of certain classical analytic number theory conjectures [Bank et al. 2015; Andrade et al. 2015; Entin 2014].

Proof of Theorem 1.7. Recall that the sequence of polynomials $\{f_n\}$ is defined by

$$f_1(x) = x^2, \quad f_{n+1}(x) = (f_n(x) + 1)^2.$$

Write $G_n = G_{f_n}$, $G_n^* = G_{f_n}^*$, and $R_n = R_{f_n}$. Since any root $\alpha \in R_n$ satisfies either $f_{n-1}(\alpha) = -1 + \sqrt{t}$ or $f_{n-1}(\alpha) = -1 - \sqrt{t}$, we may decompose R_n as the union $R_n = R_n^+ \cup R_n^-$, with

$$R_n^\pm = \{\alpha \in R_n : f_{n-1}(\alpha) = -1 \pm \sqrt{t}\}.$$

Note that both $\text{Gal}(\mathbb{F}_p(R_n^+)/\mathbb{F}_p(\sqrt{t}))$ and $\text{Gal}(\mathbb{F}_p(R_n^-)/\mathbb{F}_p(\sqrt{t}))$ are isomorphic to G_{n-1} , and similarly both $\text{Gal}(\overline{\mathbb{F}}_p(R_n^+)/\overline{\mathbb{F}}_p(\sqrt{t}))$ and $\text{Gal}(\overline{\mathbb{F}}_p(R_n^-)/\overline{\mathbb{F}}_p(\sqrt{t}))$ are isomorphic to G_{n-1}^* .

Let H_n and H_n^* be the normal subgroup of G_n and G_n^* that fixes \sqrt{t} , so that $[G_n : H_n] = [G_n^* : H_n^*] = 2$. Since H_n preserves both R_n^+ and R_n^- , we get an embedding $\iota_n : H_n \hookrightarrow G_{n-1} \times G_{n-1}$ by setting the first and second components of $\iota_n(\xi)$ to be the images of ξ under the two quotient maps $H_n \rightarrow \text{Gal}(\mathbb{F}_p(R_n^+)/\mathbb{F}_p(\sqrt{t}))$ and $H_n \rightarrow \text{Gal}(\mathbb{F}_p(R_n^-)/\mathbb{F}_p(\sqrt{t}))$, respectively. Similarly, we also get an embedding $\iota_n^* : H_n^* \hookrightarrow G_{n-1}^* \times G_{n-1}^*$.

We show, by induction on n , that when $p > 2f_{n-1}(0) + 2$, the embeddings ι_n and ι_n^* are in fact isomorphisms, and moreover $G_n = G_n^*$ for each n . The base case is clear. Now assume that $G_{n-1} = G_{n-1}^*$. To see that ι_n^* is surjective, by Lemma 15 in [Fried 1970] it suffices to verify that for each $\lambda \in \overline{\mathbb{F}}_p$, at most one of the two values $-1 + \sqrt{\lambda}$ and $-1 - \sqrt{\lambda}$ is a branch point of f_{n-1} . By definition, the set of branch points of f_{n-1} is

$$\{f_{n-1}(x) : x \in \overline{\mathbb{F}}_p, f'_{n-1}(x) = 0\}.$$

This is easily computed to be the set

$$\{f_1(0), f_2(0), \dots, f_{n-1}(0)\} = \{0, 1, 4, 25, \dots\}.$$

When $p > 2f_{n-1}(0) + 2$, it is indeed the case that at most one of $-1 + \sqrt{\lambda}$ and $-1 - \sqrt{\lambda}$ can lie in this set for any λ . This shows that

$$H_n^* \cong G_{n-1}^* \times G_{n-1}^* \cong G_{n-1} \times G_{n-1}.$$

Moreover, since $H_n^* \subset H_n \subset G_{n-1} \times G_{n-1}$, we conclude that $H_n^* = H_n$, and thus $G_n^* = G_n$ as well. This completes the induction step.

With the structure of G_n in hand, it is now a simple matter to write down the recursive relation

$$\alpha_p(f_n) = \frac{1}{2}[1 - (1 - \alpha_p(f_{n-1}))^2] = \alpha_p(f_{n-1}) - \frac{1}{2}\alpha_p(f_{n-1})^2,$$

provided that $p > 2f_{n-1}(0) + 2$. In fact, if $\xi \in G_n$ has a fixed point, then ξ must fix \sqrt{t} and thus lie in H_n , and moreover at least one of the two components of $\iota_n(\xi)$ has a fixed point. Finally, the bound $a_n \leq 2n^{-1}$ follows from a standard induction argument. \square

We mentioned in [Remark 2.5](#) the reasoning behind making f_n highly decomposable. To end this section, we show that decomposability is quite essential in order for α_p to be small. The proof uses [Theorem 2.4](#) together with results in [[Fried 1970](#)] (similar arguments are also used in [[Guralnick and Wan 1997](#)]). We say that a polynomial f is indecomposable if it cannot be written as a composition of two polynomials of degree at least 2.

Proposition 5.3. *Let $f(x) \in \mathbb{Z}[x]$ be an indecomposable polynomial of degree $d \geq 1$. Then the average value of $\alpha_p(f)^{-1}$ as p varies is at most 2. Consequently, $\alpha(f) \geq \frac{1}{2}$.*

Proof. Let G be the Galois group of the splitting field of $f(x) - t$ over $\mathbb{Q}(t)$, viewed as a subgroup of the symmetric group S_d on d letters via its action on the d roots of $f(x) - t$. Since f is indecomposable, G is primitive [[Fried 1970](#), Lemma 2]. Moreover, G contains a d -cycle [[ibid.](#), Lemma 3]. Hence either d is prime or G is doubly transitive [[ibid.](#), Lemma 7]. In either case, the conclusion follows from [Theorem 2.4](#), since $\tau(d) = 2$ when d is prime and $(f(x) - f(y))/(x - y) \in \mathbb{Q}[x, y]$ is irreducible when G is doubly transitive [[ibid.](#), Lemma 14]. \square

6. Further remarks

6A. More on the sharpness of Gallagher’s larger sieve. Gallagher’s larger sieve in its general form as stated in [Theorem 1.1](#) has the optimal bound. Indeed, if we take $A \subset [N]$ to be any subset with cardinality Q and take \mathcal{P} be the set of all primes between Q and N , the general form of the larger sieve gives the sharp bound $|A| \ll Q$, because the numerator is about N and the denominator is about N/Q . This shows that any potential improvement to [Corollary 1.2](#) must incorporate the ill-distribution modulo many *small* primes.

Under the assumption of [Corollary 1.2](#), one may go over the argument in the proof of the larger sieve to find out what happens if $|A|$ is close to N^α . Indeed, in the typical proof of Gallagher’s larger sieve, one uses the upper and lower bounds

$$\frac{|X|^2}{\alpha} \log Q \leq \sum_{x, x' \in X} \sum_{\substack{p|x-x' \\ p \leq Q}} \log p \leq |X|^2 \log N + |X|Q, \tag{6-1}$$

where Q is about N^α .

If the upper bound is (almost) sharp, then almost all of the nonzero differences $x - x'$ should be Q -smooth, meaning that they do not have prime divisors larger

than Q . For a random integer n , it is reasonable to expect that

$$\sum_{\substack{p|n \\ p \leq Q}} \log p \approx \sum_{p \leq Q} \frac{\log p}{p} \sim \log Q. \tag{6-2}$$

If this indeed holds for almost all differences $x - x'$, then one can take Q to be any small power of N and deduce from (6-1) that $|X| \ll N^\epsilon$.

Now consider the situation when X is the set of d -th powers up to N . Because of the factorization

$$a^d - b^d = \prod_{\ell|d} \Phi_\ell(a, b),$$

where Φ_ℓ is the cyclotomic polynomial of degree $\phi(\ell)$, we cannot expect (6-2) to be true for $n = a^d - b^d$. However, it is still reasonable to expect that each factor $\Phi_\ell(a, b)$ satisfies (6-2). If so, then we obtain an upper bound in (6-1) with $\log N$ there replaced by $\tau(d) \log Q$, which in turn implies that $\tau(d) \geq \alpha^{-1}$. This is consistent with the conclusion of [Theorem 1.5](#).

On the other hand, making this heuristic rigorous could be extremely hard. For example, it is an open problem to obtain a bound better than $|X| \ll N^{1/2}$ for $X \subset [N]$ with all nonzero differences $x - x'$ ($x, x' \in X$) N^κ -smooth, where $\kappa > 0$ is sufficiently small (see [\[Elsholtz and Harper 2015\]](#)).

There are versions of Gallagher’s larger sieve over arbitrary number fields [\[Ellenberg et al. 2009; Zywinia 2010\]](#). One can ask similar questions about their sharpness in this general setting, and use [Theorem 2.4](#) to formulate an improved larger sieve conjecture. We will not do so here since the case over \mathbb{Z} is already quite interesting.

6B. Computing $\alpha(f)$ via Galois groups. The main result of this paper computes the average of $m_p(f)$ as p varies, as a consequence of the Chebotarev density theorem. It is natural to ask if one can compute $\alpha(f)$, the average of $\alpha_p(f)$ as p varies, directly, especially since we do have such a formula for each individual $\alpha_p(f)$ as in [Lemma 5.1](#).

Proposition 6.1. *Let K be a number field and $f(x) \in \mathbb{C}_K[x]$ be a monic polynomial of degree d . Let $G = \text{Gal}(K(R)/K(t))$, where R is the set of roots of $f(x) - t$. Let $\alpha(G)$ be the proportion of elements in G with at least one fixed point under the natural action on R . Then*

$$\lim_{Q \rightarrow \infty} \frac{1}{\pi(Q)} \sum_{N(p) \leq Q} \alpha_p(f) = \alpha(G).$$

In other words, $\alpha(f) = \alpha(G)$.

Remark 6.2. Unfortunately, we are unable to use this interpretation to obtain good lower bounds on $\alpha(f)$, but see [\[Guralnick and Wan 1997\]](#) for an example where

large values of $\alpha_p(f)$ are studied via Galois groups. On the other hand, we feel that any possible improvement to the bound $\alpha(f) \geq \tau(d)^{-1}$ is likely to come from studying the Galois group G .

Proof. Write $E = K(R)$. Let $G^* = \text{Gal}(\overline{K}(R)/\overline{K}(t))$. Let $L = E \cap \overline{K}$ be the algebraic closure of K in E , so that $E = L(R)$ and $G^* = \text{Gal}(E/L(t))$. By the primitive element theorem, there exists $\theta \in E$ such that $E = L(t, \theta)$. Suppose that θ satisfies the relation

$$h_m(t)\theta^m + \dots + h_1(t)\theta + h_0(t) = 0,$$

where $m = [E : L(t)]$ and $h_m(t), \dots, h_1(t), h_0(t)$ are relatively prime polynomials over L . Let $h \in L[t, y]$ be the two-variable polynomial defined by

$$h(t, y) = h_m(t)y^m + \dots + h_1(t)y + h_0(t).$$

Clearly h is a minimal polynomial of θ , and thus h is irreducible. By the definition of L , the polynomial h is also absolutely irreducible.

Let $\mathfrak{p} \subset \mathbb{O}_K$ be a prime in K and $\mathfrak{P} \subset \mathbb{O}_L$ be a prime in L lying above \mathfrak{p} . By [Lemma 3.5](#), $h_{\mathfrak{P}} \in \kappa_{\mathfrak{P}}[t, y]$ remains absolutely irreducible for all but finitely many \mathfrak{P} . Let $\theta_{\mathfrak{P}} \in \overline{\kappa_{\mathfrak{P}}}(t)$ be an element satisfying $h_{\mathfrak{P}}(t, \theta_{\mathfrak{P}}) = 0$, so that $E_{\mathfrak{P}} = \kappa_{\mathfrak{P}}(t, \theta_{\mathfrak{P}})$ is a degree- m field extension of $\kappa_{\mathfrak{P}}(t)$ with $E_{\mathfrak{P}} \cap \overline{\kappa_{\mathfrak{P}}} = \kappa_{\mathfrak{P}}$. Since $E/L(t)$ is Galois, all roots of $h(t, y)$ in $\overline{L}(t)$ lie in E . This implies that all roots of $h_{\mathfrak{P}}(t, y)$ in $\overline{\kappa_{\mathfrak{P}}}(t)$ lie in $E_{\mathfrak{P}}$ for all but finitely many \mathfrak{P} , and thus $E_{\mathfrak{P}}/\kappa_{\mathfrak{P}}(t)$ is also Galois. Note that there is a natural isomorphism $G^* = \text{Gal}(E/L(t)) \cong \text{Gal}(E_{\mathfrak{P}}/\kappa_{\mathfrak{P}}(t))$, since an element in either Galois group is determined by its image of θ or $\theta_{\mathfrak{P}}$.

Now we look at the polynomial $f(x) - t$. Since it factors into linear factors over E , its reduction $f_{\mathfrak{p}}(x) - t$ factors into linear factors over $E_{\mathfrak{P}}$ for all but finitely many \mathfrak{P} . By an abuse of notation, we will continue to write R for the set of roots of $f_{\mathfrak{p}}(x) - t$ in $\overline{\kappa_{\mathfrak{P}}}(t)$. Therefore the splitting fields $\kappa_{\mathfrak{P}}(R)$ and $\kappa_{\mathfrak{p}}(R)$ are contained in $E_{\mathfrak{P}}$. On the other hand, since $\theta \in K(R)$ and $L \subset K(R)$, we have $\theta_{\mathfrak{P}} \in \kappa_{\mathfrak{p}}(R)$ and $\kappa_{\mathfrak{P}} \subset \kappa_{\mathfrak{p}}(R)$ for all but finitely many \mathfrak{P} . This shows that $\kappa_{\mathfrak{p}}(R) = E_{\mathfrak{P}}$.

Let $\sigma_{\mathfrak{P}}G^*$ be the coset which is the inverse image of the Frobenius automorphism $\sigma_{\mathfrak{P}}$ under the quotient map

$$\text{Gal}(E_{\mathfrak{P}}/\kappa_{\mathfrak{p}}(t)) \rightarrow \text{Gal}(\kappa_{\mathfrak{P}}(t)/\kappa_{\mathfrak{p}}(t)) = \text{Gal}(\kappa_{\mathfrak{P}}/\kappa_{\mathfrak{p}}),$$

which has kernel $\text{Gal}(E_{\mathfrak{P}}/\kappa_{\mathfrak{P}}(t)) = G^*$. By [Lemma 5.1](#), we have

$$\alpha_{\mathfrak{p}}(f) = \alpha(\sigma_{\mathfrak{P}}G^*) + O_d(N(\mathfrak{p})^{-1/2}).$$

Note that the quantity $\alpha(\sigma_{\mathfrak{P}}G^*)$ does not depend on the choice of \mathfrak{P} . Via the inclusion $\text{Gal}(\kappa_{\mathfrak{P}}/\kappa_{\mathfrak{p}}) \hookrightarrow \text{Gal}(L/K)$, we may view $\sigma_{\mathfrak{P}}$ as an element in $\text{Gal}(L/K)$

and $\sigma_{\mathfrak{p}}G^*$ as a coset in G . By the Chebotarev density theorem, the cosets $\sigma_{\mathfrak{p}}G^*$ become equidistributed in G as \mathfrak{p} varies. Therefore $\alpha(f) = \alpha(G)$ as desired. \square

For a generic polynomial of degree d , the Galois group G in Proposition 6.1 is the full symmetric group S_d . Indeed, using the large sieve inequality, Gallagher [1973] obtained a precise bound on the number of exceptional polynomials (with coefficients bounded by a parameter) whose Galois group is not S_d . This bound has since been improved by Dietmann [2013]. By Proposition 6.1, this implies that

$$\alpha(f) = \alpha(S_d) = 1 - \frac{1}{2} + \frac{1}{6} - \frac{1}{24} + \dots + \frac{(-1)^{d-1}}{d!}$$

for a typical f of degree d . Moreover, this quantity tends to $1 - e^{-1}$ as $d \rightarrow \infty$.

For $d \leq 4$, we have the following sharp lower bounds.

Proposition 6.3 (polynomials of small degree). *For a positive integer d , let α_d be the smallest possible value of $\alpha(f)$, where $f \in \mathbb{Q}[x]$ is a polynomial of degree d . Then $\alpha_2 = \frac{1}{2}$, $\alpha_3 = \frac{2}{3}$, and $\alpha_4 = \frac{3}{8}$.*

Proof. For $d = 2$ this is obvious. Suppose that $d \in \{3, 4\}$. Let G be the Galois group as in Proposition 6.1. We claim that $G \neq \mathbb{Z}/d\mathbb{Z}$, the cyclic group of order d . In fact, for $t \in \mathbb{Z}$ sufficiently large, the polynomial $f(x) = t$ has at least one real root and at least one nonreal root. Let $\alpha \in \mathbb{R}$ be a real root of $f(x) = t$. Then the splitting field of $f(x) - t$ contains properly the subfield $\mathbb{Q}(\alpha)$, and thus has degree larger than d over \mathbb{Q} . This shows that the Galois group of $f(x) - t$ is not $\mathbb{Z}/d\mathbb{Z}$ for all t sufficiently large. The fact that $G \neq \mathbb{Z}/d\mathbb{Z}$ then follows from Hilbert’s irreducibility theorem. Now that $G \subset S_d$ is transitive and $G \neq \mathbb{Z}/d\mathbb{Z}$, the only possibilities are $G = S_3$ when $d = 3$ and $G \in \{S_4, A_4, D_4\}$ when $d = 4$. The conclusion follows by computing $\alpha(G)$ for these choices of G . \square

Not surprisingly, the nature of α_d depends not only on the size of d , but also the arithmetic of d (see Proposition 5.3). In general, given a transitive subgroup $G \subset S_d$, we do not know how to tell whether G can be realized as a Galois group as in Proposition 6.1. This is reminiscent of the inverse Galois problem over $K(t)$, but here we require the polynomial to take the shape $f(x) - t$ for some $f(x) \in K[x]$. We refer the interested reader to the book [Serre 2008] and references therein for background and known results on the classical inverse Galois problem.

Acknowledgements

Thanks to Brian Conrad for help with proofs and many useful comments, to Kannan Soundararajan for helpful discussions, to Akshay Venkatesh for asking a question that led to this paper, and to the anonymous referee for helpful suggestions.

References

- [Andrade et al. 2015] J. C. Andrade, L. Bary-Soroker, and Z. Rudnick, “Shifted convolution and the Titchmarsh divisor problem over $\mathbb{F}_q[t]$ ”, *Philos. Trans. A* **373**:2040 (2015), 20140308, 18. [MR 3338116](#)
- [Bank et al. 2015] E. Bank, L. Bary-Soroker, and L. Rosenzweig, “Prime polynomials in short intervals and in arithmetic progressions”, *Duke Math. J.* **164**:2 (2015), 277–295. [MR 3306556](#) [Zbl 06416949](#)
- [Birch and Swinnerton-Dyer 1959] B. J. Birch and H. P. F. Swinnerton-Dyer, “Note on a problem of Chowla”, *Acta Arith.* **5** (1959), 417–423. [MR 22 #4675](#) [Zbl 0091.04301](#)
- [Bombieri and Pila 1989] E. Bombieri and J. Pila, “The number of integral points on arcs and ovals”, *Duke Math. J.* **59**:2 (1989), 337–357. [MR 90j:11099](#) [Zbl 0718.11048](#)
- [Cohen 1970] S. D. Cohen, “The distribution of polynomials over finite fields”, *Acta Arith.* **17** (1970), 255–271. [MR 43 #3234](#) [Zbl 0209.36001](#)
- [Croot and Elsholtz 2004] E. S. Croot, III and C. Elsholtz, “On variants of the larger sieve”, *Acta Math. Hungar.* **103**:3 (2004), 243–254. [MR 2005c:11113](#) [Zbl 1060.11056](#)
- [Croot and Lev 2007] E. S. Croot, III and V. F. Lev, “Open problems in additive combinatorics”, pp. 207–233 in *Additive combinatorics*, CRM Proc. Lecture Notes **43**, Amer. Math. Soc., Providence, RI, 2007. [MR 2009d:11038](#) [Zbl 1183.11005](#)
- [Dietmann 2013] R. Dietmann, “Probabilistic Galois theory”, *Bull. Lond. Math. Soc.* **45**:3 (2013), 453–462. [MR 3065016](#) [Zbl 06176870](#)
- [Ellenberg et al. 2009] J. S. Ellenberg, C. Elsholtz, C. Hall, and E. Kowalski, “Non-simple abelian varieties in a family: geometric and analytic approaches”, *J. Lond. Math. Soc.* (2) **80**:1 (2009), 135–154. [MR 2010f:11093](#) [Zbl 1263.11064](#)
- [Elsholtz and Harper 2015] C. Elsholtz and A. J. Harper, “Additive decompositions of sets with restricted prime factors”, *Trans. Amer. Math. Soc.* **367**:10 (2015), 7403–7427. [MR 3378834](#) [Zbl 06479361](#)
- [Entin 2014] A. Entin, “On the Bateman–Horn conjecture for polynomials over large finite fields”, preprint, 2014. [arXiv 1409.0846](#)
- [Fried 1970] M. Fried, “On a conjecture of Schur”, *Michigan Math. J.* **17** (1970), 41–55. [MR 41 #1688](#) [Zbl 0169.37702](#)
- [Gallagher 1971] P. X. Gallagher, “A larger sieve”, *Acta Arith.* **18** (1971), 77–81. [MR 45 #214](#) [Zbl 0231.10028](#)
- [Gallagher 1973] P. X. Gallagher, “The large sieve and probabilistic Galois theory”, pp. 91–101 in *Analytic number theory* (St. Louis, MO, 1972), Proc. Sympos. Pure Math. **24**, Amer. Math. Soc., Providence, R.I., 1973. [MR 48 #11020](#) [Zbl 0279.10036](#)
- [Gomez-Calderon and Madden 1988] J. Gomez-Calderon and D. J. Madden, “Polynomials with small value set over finite fields”, *J. Number Theory* **28**:2 (1988), 167–188. [MR 89d:11111](#) [Zbl 0634.12014](#)
- [Green and Harper 2014] B. Green and A. J. Harper, “Inverse questions for the large sieve”, *Geom. Funct. Anal.* **24**:4 (2014), 1167–1203. [MR 3248483](#) [Zbl 1316.11085](#)
- [Guralnick and Wan 1997] R. Guralnick and D. Wan, “Bounds for fixed point free elements in a transitive group and applications to curves over finite fields”, *Israel J. Math.* **101** (1997), 255–287. [MR 98j:12002](#) [Zbl 0910.11053](#)
- [Heath-Brown 2002] D. R. Heath-Brown, “The density of rational points on curves and surfaces”, *Ann. of Math.* (2) **155**:2 (2002), 553–595. [MR 2003d:11091](#) [Zbl 1039.11044](#)

- [Helfgott and Venkatesh 2009] H. A. Helfgott and A. Venkatesh, “How small must ill-distributed sets be?”, pp. 224–234 in *Analytic number theory*, Cambridge Univ. Press, 2009. [MR 2010d:11087](#) [Zbl 1217.11073](#)
- [Hindry and Silverman 2000] M. Hindry and J. H. Silverman, *Diophantine geometry*, Graduate Texts in Mathematics **201**, Springer, New York, 2000. [MR 2001e:11058](#) [Zbl 0948.11023](#)
- [Kowalski 2010] E. Kowalski, “Some aspects and applications of the Riemann hypothesis over finite fields”, *Milan J. Math.* **78**:1 (2010), 179–220. [MR 2011g:11229](#) [Zbl 1271.11113](#)
- [Lagarias and Odlyzko 1977] J. C. Lagarias and A. M. Odlyzko, “Effective versions of the Chebotarev density theorem”, pp. 409–464 in *Algebraic number fields: L-functions and Galois properties* (Durham, UK, 1975), edited by A. Fröhlich, Academic Press, London, 1977. [MR 56 #5506](#) [Zbl 0362.12011](#)
- [Lang and Weil 1954] S. Lang and A. Weil, “Number of points of varieties in finite fields”, *Amer. J. Math.* **76** (1954), 819–827. [MR 16,398d](#) [Zbl 0058.27202](#)
- [Montgomery 1978] H. L. Montgomery, “The analytic principle of the large sieve”, *Bull. Amer. Math. Soc.* **84**:4 (1978), 547–567. [MR 57 #5931](#) [Zbl 0408.10033](#)
- [Schmidt 1976] W. M. Schmidt, *Equations over finite fields: an elementary approach*, Lecture Notes in Mathematics **536**, Springer, Berlin, 1976. [MR 55 #2744](#) [Zbl 0329.12001](#)
- [Serre 2008] J.-P. Serre, *Topics in Galois theory*, 2nd ed., Research Notes in Mathematics **1**, A K Peters, Wellesley, MA, 2008. [MR 2008i:12010](#) [Zbl 1128.12001](#)
- [Shao 2014] X. Shao, “On an inverse ternary Goldbach problem”, preprint, 2014. To appear in *Amer. J. Math.* [arXiv 1404.6022](#)
- [Walsh 2012] M. N. Walsh, “The inverse sieve problem in high dimensions”, *Duke Math. J.* **161**:10 (2012), 2001–2022. [MR 2954623](#) [Zbl 06063227](#)
- [Walsh 2015] M. N. Walsh, “Bounded rational points on curves”, *Int. Math. Res. Not.* **2015**:14 (2015), 5644–5658. [MR 3384452](#)
- [Zywina 2010] D. Zywina, “Hilbert’s irreducibility theorem and the larger sieve”, preprint, 2010. [arXiv 1011.6465](#)

Communicated by Andrew Granville

Received 2014-12-17

Revised 2015-07-19

Accepted 2015-08-17

xuancheng.shao@maths.ox.ac.uk *Mathematical Institute, University of Oxford,
Radcliffe Observatory Quarter, Woodstock Road, Oxford,
OX2 6GG, United Kingdom*

Algebra & Number Theory

msp.org/ant

EDITORS

MANAGING EDITOR

Bjorn Poonen
Massachusetts Institute of Technology
Cambridge, USA

EDITORIAL BOARD CHAIR

David Eisenbud
University of California
Berkeley, USA

BOARD OF EDITORS

Georgia Benkart	University of Wisconsin, Madison, USA	Susan Montgomery	University of Southern California, USA
Dave Benson	University of Aberdeen, Scotland	Shigefumi Mori	RIMS, Kyoto University, Japan
Richard E. Borcherds	University of California, Berkeley, USA	Raman Parimala	Emory University, USA
John H. Coates	University of Cambridge, UK	Jonathan Pila	University of Oxford, UK
J-L. Colliot-Thélène	CNRS, Université Paris-Sud, France	Anand Pillay	University of Notre Dame, USA
Brian D. Conrad	Stanford University, USA	Victor Reiner	University of Minnesota, USA
Hélène Esnault	Freie Universität Berlin, Germany	Peter Sarnak	Princeton University, USA
Hubert Flenner	Ruhr-Universität, Germany	Joseph H. Silverman	Brown University, USA
Sergey Fomin	University of Michigan, USA	Michael Singer	North Carolina State University, USA
Edward Frenkel	University of California, Berkeley, USA	Vasudevan Srinivas	Tata Inst. of Fund. Research, India
Andrew Granville	Université de Montréal, Canada	J. Toby Stafford	University of Michigan, USA
Joseph Gubeladze	San Francisco State University, USA	Ravi Vakil	Stanford University, USA
Roger Heath-Brown	Oxford University, UK	Michel van den Bergh	Hasselt University, Belgium
Craig Huneke	University of Virginia, USA	Marie-France Vignéras	Université Paris VII, France
Kiran S. Kedlaya	Univ. of California, San Diego, USA	Kei-Ichi Watanabe	Nihon University, Japan
János Kollár	Princeton University, USA	Efim Zelmanov	University of California, San Diego, USA
Yuri Manin	Northwestern University, USA	Shou-Wu Zhang	Princeton University, USA
Philippe Michel	École Polytechnique Fédérale de Lausanne		

PRODUCTION

production@msp.org

Silvio Levy, Scientific Editor


See inside back cover or msp.org/ant for submission instructions.

The subscription price for 2015 is US \$255/year for the electronic version, and \$440/year (+\$55, if shipping outside the US) for print and electronic. Subscriptions, requests for back issues and changes of subscribers address should be sent to MSP.

Algebra & Number Theory (ISSN 1944-7833 electronic, 1937-0652 printed) at Mathematical Sciences Publishers, 798 Evans Hall #3840, c/o University of California, Berkeley, CA 94720-3840 is published continuously online. Periodical rate postage paid at Berkeley, CA 94704, and additional mailing offices.

ANT peer review and production are managed by EditFLOW[®] from MSP.

PUBLISHED BY

 **mathematical sciences publishers**
nonprofit scientific publishing

<http://msp.org/>

© 2015 Mathematical Sciences Publishers

Algebra & Number Theory

Volume 9 No. 10 2015

Equivariant torsion and base change	2197
MICHAEL LIPNOWSKI	
Induction parabolique et (φ, Γ)-modules	2241
CHRISTOPHE BREUIL	
On the normalized arithmetic Hilbert function	2293
MOUNIR HAJLI	
The abelian monoid of fusion-stable finite sets is free	2303
SUNE PRECHT REEH	
Polynomial values modulo primes on average and sharpness of the larger sieve	2325
XUANCHENG SHAO	
Bounds for Serre's open image theorem for elliptic curves over number fields	2347
DAVIDE LOMBARDO	
On 0-cycles with modulus	2397
AMALENDU KRISHNA	