Frequencies of Successive Tuples of Frobenius Classes

Avner Ash, Brandon Bate, and Robert Gross

CONTENTS

In this paper, we consider the sequence of Frobenius conjugacy classes for a Galois extension $K/\mathbb{Q}$, ordered by the increasing sequence of rational primes. For a given K, we look at the frequencies of nonoverlapping consecutive k-tuples in this sequence. We compare these frequencies to what would be expected by the Cebotarev density theorem if there were statistical independence between successive Frobenius classes. We find striking variations of behavior as K varies.

1. INTRODUCTION

For any Galois number field $K/\mathbb{Q}$ of degree d and discriminant Δ, the sequence Frob_2, Frob_3, $\mathrm{Frob}_5, \ldots$ of conjugacy classes in $\mathrm{Gal}(K/\mathbb{Q})$ is defined, except for those rational primes p that are ramified in K. We number the conjugacy classes in $\mathrm{Gal}(K/\mathbb{Q})$ by the integers $0, 1, 2, \ldots, n$, and then consider this sequence $s = s(K)$ of integers, which we will call the *Frobenius sequence of* K. (To keep file sizes of constant length, we arbitrarily assign 0 to the finitely many ramified primes.) For example, if K is a quadratic extension, the Frobenius sequence records the quadratic residue/nonresidue status of successive primes with respect to some modulus.

The Čebotarev density theorem (CDT) says that a given conjugacy class C of cardinality c will appear with limiting frequency $\frac{c}{d}$ in $s(K)$. This paper grew out of the question, given these frequencies, to what extent does $s(K)$ look "random"?

Of course, the Frobenius sequence is completely deterministic. However, one may apply statistical tests to sequences of integers to test whether they deserve the title of "pseudorandom." Dozens of such tests have been proposed, as may be seen by a web search.[1]

We experimented with a number of different approaches, some of which we may report on in another paper. For example, we looked at autocorrelations of a Frobenius sequence. We divided a Frobenius sequence

[1]For example, many tests may be found at http://www.ciphersbyritter.com/RES/RANDTEST.HTM and http://csrc.nist.gov/groups/ST/toolkit/rng/stats_tests.html.

2000 AMS Subject Classification: 11N05, 11K45, 62P99

Keywords: Frobenius classes, pseudorandom sequences

1058-6458/2009 $0.50 per page

into chunks of equal size and used nonparametric tests on various statistics of the chunks. We compared Frobenius sequences for all pairs of different fields having the same Galois group using a comparison χ^2-value for k-tuple frequencies.

In the end, the clearest and most suggestive data we obtained came from the two tests we report on in this paper. The data came from the frequencies of nonoverlapping k-tuples of the Frobenius sequence. First we asked, do successive nonoverlapping k-tuples of the sequence exhibit frequencies that are reasonably like those that would occur in a sequence of the same length of iid random variables, each with the discrete probability distribution given by CDT? See [Knuth 81, pp. 38–45]. We use the χ^2-statistic and the corresponding p-value to interpret the word "reasonably" in the previous sentence.

Second, we apportion nonoverlapping pairs into two classes, of the forms (x, x) and (x, y) with $y \neq x$, and count the frequencies of these two classes. Again we ask, do these two classes exhibit frequencies that are reasonably like those that would occur in a sequence of the same length of iid random variables, each with the discrete probability distribution given by CDT? We call this the "matching test." We use a Z-score to interpret the word "reasonably" in this case.

In the first test, when the p-value exceeds 0.99, we will say that $s(K)$ is exhibiting "nonrandom" behavior. In the second test, we will do so when the absolute value of the Z-score exceeds 3. However, we are not making any claims about actual pseudorandomness, however defined. We use the term "nonrandom" for convenience to describe a situation in which the frequencies deviate largely from those that would have occurred randomly given the CDT frequencies.

We could have counted overlapping k-tuples of the sequence. Work of I. J. Good [Good 53] and P. Billingsley [Billingsley 56] allows for the correct interpretation of these data, and since there are k times as many overlapping k-tuples as nonoverlapping ones, it might be thought that the former would give stronger results. There are three reasons why we did not use overlapping k-tuples.

1. The recommended statistic used in the overlapping k-tuples case is a second-order difference between the so-called ψ^2-values for tuples of size k, $k-1$, and $k-2$. We are more interested in goodness of fit than in hypothesis testing, and this test would obscure matters by mixing these three different-sized tuples. More generally, it seems cleaner to use nonoverlapping k-tuples, which behave independently, to study goodness of fit, whereas overlapping k-tuples are not independent of one another [Knuth 81, p. 60].

2. There is a recommended rule of thumb, recalled below, which tells us how large we can reasonably take k when investigating nonoverlapping k-tuples using the χ^2-test. For example, this rule limits our study of A_5-fields to 1-, 2-, and 3-tuples. Assuming that the same rule of thumb applies to the overlapping k-tuples, it would still not permit the use of 4-tuples in A_5-fields.

3. We computed the p-values for overlapping k-tuples for many of our test fields, and the results were not such as to give any different impression from the p-values reported in this paper.

When K is abelian, the work of [Rubinstein and Sarnak 94] and [Granville and Martin 06] may often be used to show that $s(K)$ is not pseudorandom in a precise sense. (These authors' results depend on assuming the generalized Riemann hypothesis (GRH) and another hypothesis called the grand simplicity hypothesis, but these seem to be reasonable assumptions.) For example, suppose $K = \mathbb{Q}(\sqrt{-1})$. Let p_i denote the ith prime. Let $X_1 = 0$, $X_i = 1$ if $p_i \equiv 3 \pmod 4$ and $X_i = -1$ if $p_i \equiv 1 \pmod 4$. Let $Y(x) = 1$ if $\sum_{p_i \le x} X_i > 0$ and $Y(x) = 0$ if $\sum_{p_i \le x} X_i \le 0$. Let

$$S(x) = \frac{1}{\log x} \int_2^x \frac{Y(t)\, dt}{t}.$$

Under the assumptions mentioned above, it is proved in [Rubinstein and Sarnak 94] that

$$\lim_{x \to \infty} S(x) \approx 0.9959.$$

Richard Arratia (via the referee of an earlier version of this paper) showed us how this result implies that $s(K)$ almost surely could not be the result of a random series of coin flips, with values in $\{\pm 1\}$. Consider a sequence of random variables $\tilde{X}_i$ modeling fair coin flips with values in $\{\pm 1\}$, and define $\tilde{Y}(x)$ and $\tilde{S}(x)$ analogously to $Y(x)$ and $S(x)$. Then consider the event that $\lim_{x \to \infty} \tilde{S}(x)$ exists and is greater than β. This event belongs to the symmetric σ-algebra, and so by the Hewitt–Savage 0-1 law [Hewitt and Savage 55], it has probability either 0 or 1. By symmetry, if $\beta > 0.5$, the probability must be 0. But if, for example, $\beta = 0.9$, our sequence $s(K)$ instantiates this event. Therefore, almost surely, $s(K)$ could not be the outcome of a sampling of iid variables.

We shall see that the existence of an abelian subfield of K may in fact account for all the "nonrandomness" we observed in the k-tuple frequencies.

This paper is primarily experimental. In contrast to what can be found in [Rubinstein and Sarnak 94], in which a limit of a sum is studied, the sequential properties of Frobenius elements do not seem to be amenable to proof by known analytic techniques. For example, L-functions do not appear to give information concerning frequencies of pairs of consecutive Frobenius classes. (However, see [Pólya 59] for an interesting heuristic concerning consecutive primes.)

We can work only with finite segments of a Frobenius sequence $s = s(K)$. We used two segments, s_{small} and s_{large}. The first fifty million primes made up s_{small} (the fifty millionth prime is 982451653). The first ten million primes larger than 10^{100} made up s_{large}. This was about as large as we could go without making the computation time unreasonably long.

We studied these Frobenius segments for a variety of fields with Galois groups C_2, A_3, S_3, D_5, and A_5. Fields of much larger degree would take too long to compute. The calculations are fairly time-consuming. Each C_2-field took about one hour of computer time to find its s_{small}, and each A_5-field took about one day.

Our findings may be briefly summarized as follows: When K contains an abelian extension (which may be K itself) with small absolute discriminant, we find nonrandom behavior. Otherwise, we tend to see random behavior. In particular, all nine of the A_5-extensions that we considered were random, in the sense that the goodness-of-fit p-values stayed below 0.99 and the Z-scores stayed between -3 and 3. However, there are some nuances that will brought out in Section 3.

The case of S_3-extensions is discussed briefly in Section 4, where we show clearly how the quadratic subfield affects the nonrandom behavior

Our results suggest three questions whose answers will require theoretical understanding of the phenomena we discovered:

1. Does the whole infinite sequence $s(K)$ possess statistical properties similar to those observed for s_{small} and s_{large}? In other words, if we observed random or nonrandom behavior for a field K, does this behavior persist to infinity?

   Note that the best known error term for effective CDT is too weak to predict even the asymptotic statistical behavior of the 1-serial χ^2-statistic (defined in Section 2 below) for $s(K)$. This is true even if we admit the GRH for K. Indeed, let $\pi_C(x)$ denote the number of primes less than or equal to x whose Frobenius class is C. Let $r(x) = \pi_C(x) - (c/d)\pi(x)$, where $\pi(x)$ is the number of primes less than or equal to x. Then it is proved in [Lagarias and Odlyzko 77], under GRH, that $|r(x)| \le \kappa((c/d)x^{1/2}\log(\Delta x^d) + \log\Delta)$ for some explicit constant κ. This is too weak to control the numerator in the formula for χ^2.

2. What is the exact role of the discriminant of K or of its abelian Galois subfields in predicting randomness? We do not know whether the product of the prime factors of the discriminant or the root discriminant is more or less important than the discriminant itself. On the other hand, the signature of K plays an obscure role, if any.

3. Is the existence of an abelian subfield the only source of nonrandom behavior for K?

We observed in many cases that a lack of randomness seems to occur because $s(K)$ is trying too hard to satisfy the CDT frequencies. That is, in these cases, if (a, b) is a consecutive pair of terms in s, then $a \neq b$ is more likely than would be predicted by the CDT frequencies. We measured this phenomenon using the "matching test" described above. The results are described in Section 3.

We actually computed statistics for about one hundred fields. To save space, we include in this paper results for just 58 fields, which provide an adequate sample of the various kinds of behavior we found.

These computations were mostly performed on Macintosh computers using the programming language C++, NTL (Number Theory Library), and GSL (GNU Scientific Library). Some of the number-theoretic computations were performed using PARI.

In practice, there are at least two ways in which our implementation differs slightly from that described above. First, to save computer time, in some cases, Frob_p was assigned on the basis of the factorization of f modulo p (and hence not necessarily to 0 if $p \mid \Delta_f$), where K is the splitting field of the monic integral polynomial f and Δ_g denotes the discriminant of a polynomial g. The number of ramified primes in the fields we studied is very small compared with the total number of entries in the portion of the Frobenius sequence in question, so the arbitrary coding of ramified primes is harmless. Second, our segment of ten million large primes actually consisted of numbers that were only very probably prime.

2. PROTOCOLS

Fix a Galois number field K with $n+1$ conjugacy classes and Frobenius sequence $s = s(K)$. Let s_* be a finite segment of this sequence of length N.

Consider the space of all sequences of length N with terms drawn from $\{0, 1, \ldots, n\}$. On this space we place the probability measure determined by the probabilities for each conjugacy class specified by CDT for K, under the assumption of independence of each term. In other words, if c_i is the number of elements in the ith conjugacy class and d is the degree of K, we set $p_i = c_i/d$. Then $\mathrm{Prob}(u_1, \ldots, u_N) = \prod_j p_{u_j}$.

There are some random variables on this space whose distributions are known to a high degree of accuracy if N is large. In this paper, we compute statistics that follow either the normal or χ^2-distributions to a good approximation.

We used (1) the χ^2-statistic to measure the goodness of fit of the frequencies of nonoverlapping k-tuples of s_* to the frequencies that would be predicted by CDT and the assumption of independence of each term in s_* from the other terms. From the χ^2-value we computed the p-value. We used (2) the Z-score of a "matching" test to measure the goodness of fit of nonoverlapping pairs s_* under the same assumption. For brevity we use the term "nonrandom" when the goodness of fit is poor, as measured by $p \geq 0.99$ or $|Z| \geq 3$.

In test (1) we calculate the χ^2-statistic by the formula

$$\chi^2 := \sum_{0 \leq i \leq n} \frac{(N_i - E_i)^2}{E_i},$$

where N_i is the number of times i occurs in s_*, and E_i is the expected number of times i occurs. Thus, $N = \sum N_i$ and $E_i = Np_i$.

The number of degrees of freedom is one less than the number of conjugacy classes, i.e., n. We then calculate a p-value using the χ^2-distribution [Bulmer 79, pp. 154–158]. Our convention is to have a p-value of 0 indicate a perfect adherence to CDT probability densities and to have p-values approaching 1 indicate a strong deviation from the CDT densities. Thus, for truly random sequences, a p-value exceeding 0.99 would occur 1% of the time.

The test just described is generally called the *frequency test*, although we will also refer to it as the 1-*serial test*. Let $k \geq 1$. We can generalize this test by redefining N_i and E_i to represent the number of times and the expected number of times the ith k-tuple $(a_1, \ldots, a_k)$ appears in the sample of nonoverlapping consecutive k-tuples of s_*: $(s_1, \ldots, s_k), (s_{k+1}, \ldots, s_{2k}), \ldots$. In this case, $E_i = \frac{N}{k} \prod_{j=1}^{k} p_{a_j}$. We refer to this generalization as the k-*serial test*. We use nonoverlapping k-tuples for reasons explained in the introduction.

When using the k-serial test, we follow the rule of thumb that the test should not be employed if $E_i < 5$ for any i. Empirical evidence shows that following this rule works well in actual statistical practice [Knuth 81, p. 42]. For example, suppose that $N = 5 \cdot 10^7$. In the case that $[K : \mathbb{Q}] = 2$, there will be 2^k possible k-tuples, each occurring with predicted probability 2^{-k}. By the rule of thumb, we can allow k to be as large as 19: We have $2^{-19}5 \cdot 10^7/19 \approx 5.02$, while $2^{-20}5 \cdot 10^7/20 \approx 2.38$. If $\mathrm{Gal}(K/\mathbb{Q}) = S_3$, we can allow k to be only as large as 7, while if $\mathrm{Gal}(K/\mathbb{Q}) = A_5$, the maximum allowable value of k is 3. In general, for each Galois group, we performed the k-serial test for $k = 1, \ldots, k_{\max}$, where $k_{\max}$ is the largest k allowed by our rule of thumb. However, we will include in Table 2 results only for $1 \leq k \leq 3$, since we never observed any interesting phenomena when using other values of k that were not already apparent for these values.

The *matching test* (2) is an ad hoc test that we developed because of an empirical observation: In many cases, a conjugacy class occurred twice in succession less often than expected by CDT. This test relies on the normal distribution and seeks nonrandom behavior that might occur because there is more alternation between conjugacy classes in s_* than would be expected for a random process.

The test starts with a finite segment $s^\dagger$ of a Frobenius sequence. We then replace $s^\dagger$ with the sequence s^* of half its length as follows. We divide $s^\dagger$ into consecutive nonoverlapping pairs. If the two entries of a pair are the same, we place an M in the corresponding position of s^*; otherwise, we place an N. In other words, M = match and N = nonmatch.

If $s^\dagger$ were randomly and independently generated according to the CDT, then the outcome s^* would be a sequence of outcomes of a Bernoulli trial, with the probability of M being $p = \sum_{0 \leq i \leq n} p_i^2$. If X counts the number of matches in s^*, then we can measure how likely it is that the sequence of M's and N's was drawn from such a Bernoulli trial by computing the Z-score:

$$Z := \frac{X - p\ell}{\sqrt{p(1-p)\ell}},$$

where ℓ is the length of s^*.

For a given Frobenius sequence $s(K)$, let s_{small} denote its first $5 \cdot 10^7$ terms and let s_{large} denote the portion of

s corresponding to the 10^7 consecutive primes starting with the smallest prime larger than 10^{100}. We performed tests (1) and (2) for both $s_* = s_{\text{small}}$ and $s_* = s_{\text{large}}$, and listed the p-values and Z-scores in Table 2.

The higher the absolute value of a given Z-score, the less random behavior is to be imputed to $s^\dagger$. For example, a Z-score greater than 3 would be expected less than about 0.3% of the time if the sequence were truly random.

3. DATA ANALYSIS

Broadly speaking, if K or an abelian subfield of K has small absolute discriminant, its Frobenius sequence $s(K)$ is more likely to have high p-values and Z-scores in our tests. The effect of the discriminant is particularly noticeable in the quadratic case. In each pair of quadratic fields generated by $x^2 \pm a$, the one with the lower discriminant never showed less nonrandomness than the other. In fact it is surprising that multiplying the discriminant by 4 sometimes makes a big difference. The examples of fields 7, 8, 9, and 10 show that the signature of the field seems not to be a factor here.

The A_3-extensions versus the C_2-extensions generally show a much stronger tendency to nonrandomness as a function of their root discriminants. For example, compare fields 7 and 19. Since both extensions are abelian and the root discriminant seems like a reasonable basis of comparison, we have no explanation for this behavior. As a function of the absolute discriminant, it seems even more strange. Fields 27 and 28 show that a huge discriminant need not result in much lower p-values than a merely big discriminant.

The elementary observations about S_3-extensions in Section 4 show that there are S_3-extensions with arbitrarily large discriminant exhibiting nonrandom behavior because of a quadratic subfield with small discriminant. However, the fields 39 and 40 make a curious pair of S_3-fields. The quadratic subfield of the second has a much larger discriminant than the quadratic subfield of the first, and yet the second field has considerably larger p-scores than the first, showing a much poorer fit to the serial frequencies than would be in accordance with the CDT frequencies.

In the last example, the second S_3-field, 40, itself also had a much larger discriminant than field 39. In the case of the D_5-extensions we have examples for which the field's own discriminant is less important than that of its quadratic subfield. Field 41 has a discriminant that is nearly 40000 times larger than the discriminant of field 43. However, it shows nonrandomness on the 2-serial test for both s_{small} and s_{large} and on the 3-serial test for s_{large}, while field 43 has p-scores below 0.99 for all three tests. The probable explanation for this observation is that the quadratic subfield of field 41 has a slightly smaller discriminant (in absolute value) than the quadratic subfield of field 43 (520 and -824, respectively).

All of the A_5-fields have relatively low p-values and Z-scores, thus showing random behavior, as expected, since there is no abelian subfield. Moreover, there seems to be no strong correlation between the size of the discriminant and the p-values or Z-scores.

We can make an interesting comparison of fields with different Galois groups but with roughly equal root discriminants. Compare the behavior on the 2-serial and 3-serial tests of fields 6, 8, 34, 43, 50, and 52. They all have root discriminants between 28 and 47. All of them have roughly similar p-scores, except that those of the A_5-fields are decidedly lower.

No number field of the more than one hundred that we tested ever returned a p-value above 0.95 for the 1-serial test. For s_{small}, the largest p-value observed in the 1-serial test was 0.93, for field 16; for s_{large}, it was 0.94 (field 46). No other field had a p-value above 0.85 on the 1-serial test on s_{small}. If the p-values for all one hundred fields were distributed randomly, we would expect around ten of them to be above 0.9. So the goodness of fit of the singleton Frobenius classes to the CDT frequencies is *better* than would be expected from true randomness, as measured across all the fields we tested.

Table 2 shows that two of the D_5-extensions, fields 41 and 42, exhibited nonrandomness in the k-serial tests. Because of the delicacy of distinguishing the two conjugacy classes into which the 5-cycles fall, it is interesting to consider for these fields a restricted subsequence consisting only of those Frobenius elements that fall into those two conjugacy classes. The resulting subsequence is approximately 40% as long as the initial sequence. The serial tests on this subsequence of s_{small} (not reported in detail here) showed that field 41 did not exhibit nonrandomness in this sense for this subsequence, while field 42 produced $p > 0.999$ for the 3-serial test.

Overall, we observe a tendency for nonrandom behavior to dissipate as we move from the small to the large primes. However, in many cases, especially when the relevant discriminants are small, p values above 0.99 are observed for both s_{small} and s_{large}. Among the quadratic fields, notable cases of dissipation may be seen in fields 6 and 11. Dissipation is more often seen in the A_3-fields, especially field 25. Surprising cases of "antidissipation" occur in the D_5-fields 46 and 47.

Similarly, in Table 2, we observe that almost all of the time, the result of the matching test decreased in absolute value between s_{small} and s_{large}. If the sign changed, it almost always changed from negative to positive. The counterexamples are few and difficult for us to explain. Field 24 has matching tests of -0.53 and -2.95 on s_{small} and s_{large} respectively, and field 46 behaves similarly, except with positive Z-scores. Field 42 is particularly notable, in that the matching test scores are 6.60 and -8.86. The Z-scores that are largest in absolute value are all negative, but fields 41 and 42 exhibit relatively large positive Z-scores for s_{small}.

4. A CASE STUDY

To see the effect of an abelian subfield of small discriminant, consider the case of polynomials of the form $x^3 \pm a$. For such fields, we assign Frob_p to conjugacy class 0 if it is in the class of 3-cycles in S_3; we assign Frob_p to class 1 if it is in the class of 2-cycles; and we assign Frob_p to class 2 if it is in the class of the identity. If $p \mid a$, we assign Frob_p to class 0.

Table 1 contains the counts of nonoverlapping pairs of consecutive classes for the sequence s_{small} for the polynomial $x^3 - 2$.

The reader can compute that the χ^2-statistic corresponding to this table is approximately $3 \cdot 10^5$, which produces a p-value extremely close to 1.

However, we can use this table to draw more inferences. A Frobenius element Frob_p will be labeled with a 1 if and only if $p \equiv 2 \pmod 3$. This observation is true for any polynomial of the form $x^3 \pm a$. Therefore, any polynomial of the form $x^3 \pm a$ will have a χ^2-statistic larger than $(5563851 - 6250000)^2/6250000 \approx 75328$. With eight degrees of freedom, any χ^2-statistic larger than 21 will give a p-value above 0.99, so our observations guarantee a p-value very close to 1 independent of a, regardless of the eight other frequency counts. A similar observation holds for the sequence s_{large}.

The explanation for this phenomenon is that the quadratic subfield of $x^3 \pm a$ is always $\mathbb{Q}(\sqrt{-3})$. Because that field yields a p-value so close to 1, all of the cubic fields that contain it as a subfield must yield a p-value close to 1.

i	j	$N_{i,j}$	i	j	$N_{i,j}$	i	j	$N_{i,j}$
0	0	2471736	0	1	4622447	0	2	1235947
1	0	4628097	1	1	5563851	1	2	2311420
2	0	1236999	2	1	2311345	2	2	618158

TABLE 1. Frequency counts for $x^3 - 2$.

5. TABLES OF DATA

In all, we gathered data on more than one hundred fields: 51 quadratic fields, 12 A_3-extensions, 22 S_3-extensions, and 9 each of D_5- and A_5-extensions. To save space, we restrict consideration to a selection of fields showing the range of behavior we observed. In Table 2, the first column contains the ordinal of the polynomial and the polynomial itself. (In Table 3, the polynomial is referred to using the ordinal.)

Recall that K is the splitting field of the given polynomial. The table is divided horizontally by Galois group $G = \text{Gal}(K/\mathbb{Q})$, which will be indicated for each subtable. The next three columns are labeled $k = 1, 2, 3$. (Even though none of these fields produced a large p value for $k = 1$, we include the data, since it is of interest to see just how well the CDT frequencies themselves are observed.) These contain the results of the k-serial tests on s_{small}. The next three columns give the same data for the sequence s_{large}. To save space, we give the p-value multiplied by 100, truncated. A p-value of 0.99 or higher is indicated in italics with *99*. The final two columns, both labeled "Matching," give the Z-score for the matching test run on s_{small} and s_{large}. A negative score means that there were fewer matches than expected.

The S_3- and D_5-extensions both have quadratic subfields that must play a role in the predictability of the distribution of the Frobenius elements. Table 3 contains the discriminants Δ of both the splitting field K and of the unique quadratic subfield L, along with the root discriminant of K, $\left|\Delta_{K/\mathbb{Q}}\right|^{1/[K:\mathbb{Q}]}$.

6. APPENDIX

It is not obvious how to distinguish the two conjugacy classes of 5-cycles in A_5 and D_5 (and similarly for higher degrees). We summarize the approach found in [Roberts 04]: Suppose that K is the splitting field for an irreducible quintic polynomial $f(x)$ over $\mathbb{Q}$. Suppose that p is a rational prime, and $f(x)$ is irreducible in $\mathbf{F}_p[x]$. In this case, Frob_p must be a 5-cycle.

We know that Δ_f must be a perfect integer square. Let $\delta_\pm$ be the two integer square roots of Δ_f.

In the field $\mathbf{F}_p[x]/(f(x))$, we let $\alpha_1 = x + (f(x))$, and define $\alpha_k = \alpha_{k-1}^p$ for $k = 2, \ldots, 5$. We compute

$$\overline{\delta} = \prod_{1 \le i < j \le 5} (\alpha_j - \alpha_i).$$

	k (s_{small})			k (s_{large})			Matching	
Polynomial	1	2	3	1	2	3	s_{small}	s_{large}
Quadratic extensions								
1. x^2-2	29	*99*	*99*	33	*99*	*99*	−391.54	−25.23
2. x^2+2	10	*99*	*99*	53	*99*	*99*	−391.81	−26.63
3. x^2-5	19	*99*	*99*	10	*99*	*99*	−478.46	−27.62
4. x^2+5	17	*99*	*99*	56	*99*	*99*	−268.59	−20.70
5. x^2-541	29	*99*	*99*	43	*99*	*99*	−8.85	−4.97
6. x^2+541	50	*99*	*99*	5	0	7	3.34	0.27
7. x^2-1987	41	24	0	89	63	86	−0.24	0.54
8. x^2+1987	1	95	*99*	16	19	95	−2.84	0.10
9. x^2-3581	76	40	*99*	46	24	30	−0.70	−0.37
10. x^2+3581	74	96	31	69	24	24	2.60	−0.07
11. x^2-7933	71	*99*	*99*	84	66	71	−4.01	0.84
12. x^2+7933	61	50	9	28	7	25	0.80	−0.38
13. $x^2-3\cdot5\cdot7\cdot11\cdot13$	22	76	*99*	60	36	*99*	−2.03	0.67
14. $x^2+3\cdot5\cdot7\cdot11\cdot13$	5	*99*	*99*	24	98	81	5.68	−3.17
15. $x^2-3\cdot5\cdot7\cdot11\cdot13\cdot17$	24	12	39	56	71	71	0.47	1.31
16. $x^2+3\cdot5\cdot7\cdot11\cdot13\cdot17$	93	89	47	73	91	97	0.89	2.30
A_3-extensions								
17. x^3-3x+1	0	*99*	*99*	27	*99*	*99*	−502.25	−37.11
18. x^3-7x+7	0	*99*	*99*	15	*99*	*99*	−590.98	−37.52
19. $x^3-877x+877$	19	*99*	*99*	72	*99*	*99*	−7.66	−4.01
20. $x^3-379281x+144488$	20	*99*	*99*	79	24	78	1.78	−0.30
21. $x^3-14979x+84881$	0	*99*	*99*	15	72	87	−1.97	0.06
22. $x^3-73713x+24571$	79	*99*	93	85	64	88	−0.83	−1.63
23. $x^3-292789x+1171156$	53	98	99	44	20	94	0.78	−0.42
24. $x^3-106347x+815327$	36	96	*99*	62	99	4	−0.53	−2.95
25. $x^3-53247x+3070577$	58	*99*	*99*	11	8	11	−0.25	0.29
26. $x^3-87871x+1142323$	4	40	*99*	91	56	57	1.37	−0.72
27. $x^3-1073774599x+1073774599$	7	12	32	42	11	55	0.31	0.48
28. $x^3-31432849x+31432849$	14	7	30	31	42	0	0.42	−0.32
S_3-extensions								
29. x^3-2	11	*99*	*99*	63	*99*	*99*	−438.35	−23.57
30. x^3-5	5	*99*	*99*	20	*99*	*99*	−437.95	−24.60
31. x^3+x-1	0	*99*	*99*	38	*99*	*99*	−122.52	−14.74
32. x^3+x+4	27	*99*	*99*	45	*99*	*99*	14.73	−4.73
33. x^3+7x+4	4	*99*	*99*	4	97	*99*	4.96	−3.06
34. x^3+5x+7	10	94	*99*	66	67	83	−1.87	−0.83
35. $x^3-8x+15$	14	88	32	82	57	41	−2.33	0.86
36. $x^3-16x+29$	14	52	*99*	56	15	47	−0.30	1.16
37. $x^3-14x+27$	32	64	98	29	11	61	0.26	0.21
38. $x^3-20x+29$	16	27	51	17	93	27	1.20	0.04
39. $x^3-47x+73$	14	42	20	22	1	20	0.95	−0.55
40. $x^3+383x+398$	44	71	74	7	22	45	−0.81	1.73
D_5-extensions								
41. $x^5-20x^3-20x^2+15x+8$	8	*99*	*99*	58	*99*	*99*	19.28	−0.51
42. $x^5-7x^3-5x^2+17x+17$	9	*99*	*99*	92	*99*	*99*	6.60	−8.86
43. $x^5-4x^3+4x^2-x+8$	11	93	96	70	98	87	3.13	−1.06
44. $x^5-20x^3-15x^2+10x+4$	16	11	53	16	31	31	1.04	−1.03
45. $x^5-19x^3-19x^2+14x+7$	36	6	18	10	8	24	1.06	1.27
46. $x^5-19x^3-14x^2+10x+4$	8	28	2	94	86	77	0.19	2.69
47. $x^5-18x^3-18x^2+13x+6$	15	5	5	26	91	89	0.09	−0.28
48. $x^5-18x^3-13x^2+10x+4$	2	92	98	8	70	38	−1.04	1.74
49. $x^5-17x^3-17x^2+12x+5$	41	93	90	51	5	2	2.24	−0.89
A_5-extensions								
50. $x^5+10x^3-10x^2+35x-18$	6	24	65	56	5	0	1.01	−0.61
51. $x^5-5x^3-11x^2-17x-13$	37	26	7	65	44	95	−1.20	0.14
52. $x^5-4x^3+x^2-2x+9$	9	2	24	18	36	72	0.08	0.35
53. $x^5-18x^3-2x^2+20x+8$	39	82	83	90	73	91	0.29	−1.74
54. $x^5-18x^3+3x^2+14x+1$	24	3	1	42	33	10	−0.31	−0.21
55. $x^5-17x^3-9x^2+4x+1$	34	81	89	87	93	9	0.81	−0.33
56. $x^5-17x^3+20x^2+20x+27$	71	25	53	90	69	96	−0.88	−0.08
57. $x^5-16x^3-2x^2+20x+8$	1	3	17	6	31	61	0.08	1.52
58. $x^5-16x^3+11x^2+20x+1$	50	28	30	10	27	14	−1.91	−0.34

TABLE 2. k-serial tests.

Field	$\Delta_{K/\mathbb{Q}}$	$\lvert\Delta_{K/\mathbb{Q}}\rvert^{1/[K:\mathbb{Q}]}$	$\Delta_{L/\mathbb{Q}}$
Quadratic extensions			
1.	2^3	2.8	
2.	-2^3	2.8	
3.	5	2.2	
4.	$-2^2 5$	4.5	
5.	541	23.3	
6.	$-2^2 541$	46.5	
7.	$2^2 1987$	89.2	
8.	-1987	44.6	
9.	3581	59.8	
10.	$-2^2 3581$	119.7	
11.	7933	89.1	
12.	$-2^2 7933$	178.1	
13.	$2^2 3\cdot 5\cdot 7\cdot 11\cdot 13$	245.1	
14.	$-3\cdot 5\cdot 7\cdot 11\cdot 13$	122.5	
15.	$2^2 3\cdot 5\cdot 7\cdot 11\cdot 13\cdot 17$	1010.5	
16.	$-3\cdot 5\cdot 7\cdot 11\cdot 13\cdot 17$	505.2	
A_3-extensions			
17.	3^4	4.3	
18.	7^2	3.7	
19.	877^2	91.6	
20.	18061^2	688.4	
21.	4993^2	292.1	
22.	$3^4 24571^2$	3656.9	
23.	$7^2 151^2 277^2$	4409.3	
24.	$3^4 35449^2$	4669.1	
25.	17749^2	680.4	
26.	$7^2 12553^2$	1976.5	
27.	$7951^2 135049^2$	1048597.3	
28.	$7^2 67^2 67021^2$	99599.2	
S_3-extensions			
29.	$-2^4 3^7$	5.7	-3
30.	$-3^7 5^4$	10.5	-3
31.	-31^3	5.6	-31
32.	$-2^6 109^3$	20.9	$-2^2 109$
33.	$-11^3 41^3$	21.2	$-11\cdot 41$
34.	-1823^3	42.7	-1823
35.	-4027^3	63.5	-4027
36.	-6323^3	79.5	-6323
37.	-8707^3	93.3	-8707
38.	9293^3	96.4	9293
39.	271409^3	521.0	271409
40.	$-2^9 28625557^3$	15132.9	$-2^3 28625557$
D_5-extensions			
41.	$2^{15} 5^{13} 13^5$	82.6	$2^3\cdot 5\cdot 13$
42.	$-7^5 17^5$	10.9	$-7\cdot 17$
43.	$-2^{15} 103^5$	28.7	$-2^3 103$
44.	$5^{13} 257^5$	129.9	$5\cdot 257$
45.	$7^5 1483^5$	101.9	$7\cdot 1483$
46.	$2^{15} 23^5 149^5$	165.6	$2^3\cdot 23\cdot 149$
47.	$2^{10} 2027^5$	90.0	$2^2 2027$
48.	23173^5	152.2	23173
49.	$47^5 131^5$	78.5	$47\cdot 131$
A_5-extensions			
50.	$2^{90} 5^{96}$	37.1	
51.	$11^{30} 587^{40}$	232.5	
52.	$17^{30} 29^{40}$	38.9	
53.	$2^{40} 881^{40}$	145.9	
54.	$7^{30} 7523^{40}$	1015.8	
55.	$2^{40} 5^{40} 677^{40}$	357.9	
56.	$43^{40} 487^{30}$	270.9	
57.	$2^{40} 7^{40} 331^{40}$	278.0	
58.	$61^{40} 733^{40}$	1259.8	

TABLE 3. Field discriminants

In fact, $\overline{\delta}$ must be an element of $\mathbf{F}_p$, and $\overline{\delta} \equiv \delta_{\pm} \pmod{p}$. We assign Frob_p to one conjugacy class or the other according to whether $\overline{\delta} \equiv \delta_+$ or $\overline{\delta} \equiv \delta_-$.

ACKNOWLEDGMENTS

We thank David Rohrlich, Richard Arratia and especially Warren Sinnott for their helpful contributions. We also thank the referees of an earlier version of this paper.

The first and second authors wish to thank the National Science Foundation for support of this research through NSF grant DMS-0455240.

REFERENCES

[Billingsley 56] Patrick Billingsley. "Asymptotic Distributions of Two Goodness of Fit Criteria." *Ann. Math. Statist.* 27 (1956), 1123–1129.

[Bulmer 79] M. G. Bulmer. *Principles of Statistics*, second edition. New York: Dover, 1979.

[Good 53] I. J. Good. "The Serial Test for Sampling Numbers and Other Tests for Randomness." *Proc. Cambridge Philos. Soc.* 49 (1953), 276–284.

[Granville and Martin 06] Andrew Granville and Greg Martin. "Prime Number Races." *Amer. Math. Monthly* 113:1 (2006), 1–33.

[Hewitt and Savage 55] E. Hewitt and L. J. Savage. "Symmetric Measures on Cartesian Products." *Trans. Amer. Math. Soc.* 80:1 (1955), 470–501.

[Knuth 81] Donald E. Knuth. *The Art of Computer Programming*, vol. 2, second edition. Upper Saddle River, NJ: Addison-Wesley, 1981.

[Lagarias and Odlyzko 77] J. C. Lagarias and A. M. Odlyzko. "Effective Versions of the Chebotarev Density Theorem." In *Algebraic Number Fields: L-Functions and Galois Properties (Proc. Sympos., Univ. Durham, Durham, 1975)*, pp. 409–464. London: Academic Press, 1977.

[Pólya 59] G. Pólya. "Heuristic Reasoning in the Theory of Numbers." *Amer. Math. Monthly* 66 (1959), 375–384.

[Press et al. 96] William H. Press, Saul A. Teukolsky, William T. Vetterling, and Brian P. Flannery. *Numerical Recipes in Fortran 77 and Fortran 90*, second edition. Cambridge: Cambridge University Press, 1996.

[Roberts 04] David P. Roberts. "Frobenius Classes in Alternating Groups." *Rocky Mountain J. Math.* 34:4 (2004), 1483–1496.

[Rubinstein and Sarnak 94] Michael Rubinstein and Peter Sarnak. "Chebyshev's Bias." *Experiment. Math.* 3:3 (1994), 173–197.

Avner Ash, Department of Mathematics, Boston College, Chestnut Hill, MA 02467-3806 (ashav@bc.edu)

Brandon Bate, Department of Mathematics, Rutgers University, Piscataway, NJ 08854-8019 (brandonbate@gmail.com)

Robert Gross, Department of Mathematics, Boston College, Chestnut Hill, MA 02467-3806 (gross@bc.edu)

Received August 27, 2007; accepted in revised form June 4, 2008.