EXCEPTIONAL COVERS AND BIJECTIONS ON RATIONAL POINTS

ROBERT M. GURALNICK, THOMAS J. TUCKER, AND MICHAEL E. ZIEVE

ABSTRACT. We show that if $f\colon X \longrightarrow Y$ is a finite, separable morphism of smooth curves defined over a finite field $\mathbb{F}_q$, where q is larger than an explicit constant depending only on the degree of f and the genus of X, then f maps $X(\mathbb{F}_q)$ surjectively onto $Y(\mathbb{F}_q)$ if and only if f maps $X(\mathbb{F}_q)$ injectively into $Y(\mathbb{F}_q)$. Surprisingly, the bounds on q for these two implications have different orders of magnitude. The main tools used in our proof are the Chebotarev density theorem for covers of curves over finite fields, the Castelnuovo genus inequality, and ideas from Galois theory.

1. INTRODUCTION

Let X and Y be normal, geometrically irreducible varieties over $\mathbb{F}_q$, and let $f\colon X \longrightarrow Y$ be a finite, generically étale $\mathbb{F}_q$-morphism. Then f is called an *exceptional cover* if the diagonal is the only geometrically irreducible component of the fiber product $X \times_Y X$ which is defined over $\mathbb{F}_q$.

The prototypical examples of exceptional covers are isogenies of abelian varieties, which are exceptional whenever zero is the only $\mathbb{F}_q$-rational point in the kernel. Other families of examples will be discussed in Section 5.

The primary interest of exceptional covers is that they induce bijections on rational points:

Theorem 1. *If f is exceptional, then f maps $X(\mathbb{F}_q)$ bijectively onto $Y(\mathbb{F}_q)$.*

This result is due to Lenstra (unpublished). Special cases and weaker versions were previously proved by Davenport and Lewis [DL], MacCluer [Mac], Williams [Wi], Cohen [Co], and Fried [Fr, Fr3, FGS]. (See [LMZ] for variants of this result over infinite constant fields.)

2000 *Mathematics Subject Classification.* Primary 11G20, 14G15, Secondary 12F10.

Key words and phrases. Weil bounds, exceptional maps, curves, finite fields, Chebotarev density.

The first author was partially supported by NSF grant DMS-0140578.

Note that, if f is exceptional over $\mathbb{F}_q$, then f is also exceptional over $\mathbb{F}_{q^m}$ for infinitely many m. Thus, f induces a bijection $X(\mathbb{F}_{q^m}) \to Y(\mathbb{F}_{q^m})$ for infinitely many m. This unusual property is the most important feature of exceptional covers.

In the present paper we show that this property characterizes exceptional covers. More precisely, we show (in Prop. 5.6) that f is exceptional if $X(\mathbb{F}_{q^m}) \to Y(\mathbb{F}_{q^m})$ is either injective or surjective for a single sufficiently large m. We can make this completely explicit in case $\dim X = 1$, where it suffices to test a single m larger than an explicit constant depending only on q, the genus of X, and the degree of f:

Theorem 2. *Let X be a curve of genus g_X, and let n be the degree of f.*

(1) *Suppose f maps $X(\mathbb{F}_q)$ injectively into $Y(\mathbb{F}_q)$, and $\sqrt{q} > 2n^2 + 4ng_X$. Then f is exceptional, and therefore bijective on rational points.*
(2) *Suppose f maps $X(\mathbb{F}_q)$ surjectively onto $Y(\mathbb{F}_q)$, and $\sqrt{q} > n!(3g_X + 3n)$. Then f is exceptional, and therefore bijective on rational points.*

Note that the bound in (1) is quite different from the bound in (2): the bound in (1) is a degree-2 polynomial in n, while the bound in (2) depends on $n!$. The reason we get such a better bound under the injectivity assumption is that injectivity is equivalent to the nonexistence of non-diagonal rational points on components of $X \times_Y X$, and these components have genus less than $n^2 + 2g_X n$. There seems to be no curve playing an analogous role for surjectivity, so we are forced to work on the Galois closure of $\mathbb{F}_q(X)/\mathbb{F}_q(Y)$, which may have genus on the order of $n!(g_X + 1)$. We do not know whether this phenomenon is indicative of the true situation, or merely an artifact of our proof. It is possible that there would be counterexamples to (2) if we replaced $n!$ by any polynomial in $\mathbb{Z}[n]$. However, we do not know any examples of non-exceptional maps f which are surjective on $\mathbb{F}_q$-points with $\sqrt{q} > 2n^2 + 4ng_X$.

Our proof of (1) uses the Weil lower bound on the number of rational points on a curve and Castelnuovo's bound on the arithmetic genus of curves in $X \times X$ to show that there are nondiagonal rational points in $X \times_Y X$ when $\sqrt{q} > 2n^2 + 4ng_X$. Our proof of (2) analyzes the decomposition and inertia groups of places of the Galois closure of $\mathbb{F}_q(X)/\mathbb{F}_q(Y)$, using an analog of Chebotarev's density theorem to translate injectivity, surjectivity, and exceptionality into group-theoretic properties which are shown to be equivalent via purely group-theoretic arguments. This proof shows that, if $\sqrt{q} > n!(3g_X+3n)$, then

surjectivity and injectivity of f are equivalent to one another and to exceptionality. By contrast, our proof of (1) does not directly yield surjectivity of f (although surjectivity follows by combining (1) with Theorem 1).

Results along the lines of Theorem 2 were previously proved in the case $g_X = 0$, where of course injectivity and surjectivity are equivalent. Most previous work restricts further to the case where $g_X = 0$ and some point of $Y(\mathbb{F}_q)$ is totally ramified under f. In this case a noneffective version of our result was proved by Davenport and Lewis [DL]. The best known effective version says that, if f is bijective on rational points and $q \geq n^4$, then f is exceptional [FJ, pp. 51–52] (see [LY] for corrections to [FJ]). When $g_X = 0$ our result draws this conclusion under the assumption $q \geq 4n^4$; but it is easy to modify our argument to make use of the ramification assumption and recover the $q \geq n^4$ bound. The effective version of the Davenport-Lewis argument extends at once to the general case $g_X = 0$ (no longer assuming a totally ramified rational point), giving the bound $q \geq 16n^4$ [GM]. Our result improves this to $q \geq 4n^4$.

We prove (1) in the next section, using an argument which is similar in spirit to that of Davenport and Lewis, although with several new ingredients to address difficulties new to the case $g_X > 0$. In Section 4 we prove (2) and Theorem 1, using a Galois-theoretic setup we learned from Lenstra. Our proof of (2) uses an analog of Chebotarev's density theorem, which we prove in Section 3. We conclude in Section 5 with some examples and conjectures.

Let us say a few words about the terminology in this paper. Given a variety W over a field k and an extension k' of k, we let $W(k')$ denote the set of k-morphisms from $\operatorname{Spec} k'$ into W. In particular, $W(k)$ is the set of closed points of W with residue field k. Also $\bar{k}$ denotes an algebraic closure of k. Throughout this paper, all curves are assumed to be projective and geometrically integral. Not all curves are assumed to be smooth; some of the curves we work with in Section 2 may be singular.

Acknowledgments. We would like to thank Hendrik Lenstra for his generous help. In particular, his ideas permeate Sections 3 and 4.

2. Geometry

In this section we use a geometric approach. Our main result concerns maps and curves defined over the field $\mathbb{F}_q$. The first few propositions are valid over an arbitrary ground field k and are stated as such.

Throughout this section, $f\colon X \longrightarrow Y$ is a finite separable morphism between smooth curves, and f is defined over either k or $\mathbb{F}_q$ depending on context. We denote the geometric genus of a curve C as g_C and the arithmetic genus as $p_a(C)$. Finally, by 'component' we always mean geometric component.

Our first result shows that Castelnuovo's upper bound on the geometric genus of a curve on a split surface is also an upper bound on the arithmetic genus.

Proposition 2.1. *Let C_1 and C_2 be smooth curves and let C' be a curve for which there is a generically injective map $\phi : C' \longrightarrow C_1 \times_k C_2$. For $i = 1, 2$, let g_i be the genus of C_i, let π_i denote projection from $C_1 \times_k C_2$ onto its i-th factor, and let d_i be the degree of the map $\pi_i \circ \phi : C' \longrightarrow C_i$. Then*

$$p_a(\phi(C')) \leq (d_1 - 1)(d_2 - 1) + d_1 g_1 + d_2 g_2. \tag{2.1.1}$$

Proof. We use several results from [Ha, §V.1], which is the source of all references in this proof. For divisors D_1 and D_2, denote the intersection pairing by $D_1.D_2$. By Thm. 1.1, this pairing is symmetric, additive, and depends only on the linear equivalence class of each D_i. Let F_i be a fiber of π_i. Since F_1 is linearly equivalent to any other (disjoint) fiber of π_1, we have $F_1.F_1 = 0$. The adjunction formula (Prop. 1.5) implies $2g_2 - 2 = F_1.K$, where K is the canonical divisor on $C_1 \times_k C_2$. Next, ex. 1.5 says that $K.K = 8(g_1 - 1)(g_2 - 1)$, so $K.K = 2(F_1.K)(F_2.K)$. By ex. 1.9, K is numerically equivalent to $(2g_1 - 2)F_1 + (2g_2 - 2)F_2$.

Let $D = \phi(C')$. By ex. 1.9, $D.D \leq 2d_1d_2$, so ex. 1.3 implies $2p_a(D) - 2 \leq 2d_1d_2 + D.K$. Since $K \equiv (2g_1 - 2)F_1 + (2g_2 - 2)F_2$, we have $D.K = (2g_1 - 2)d_1 + (2g_2 - 2)d_2$. Thus

$$2p_a(D) - 2 \leq 2d_1d_2 + (2g_1 - 2)d_1 + (2g_2 - 2)d_2,$$

and the desired result follows. □

Write $Z = X \times_Y X$, and note that Z embeds naturally into $X \times_k X$ as the locus of points (P, Q) for which $f(P) = f(Q)$.

Proposition 2.2. *Let $(P, Q) \in Z(\bar{k})$ be a point which lies in more than one component of Z. Then f is ramified at both P and Q.*

Proof. If f is unramified at P then f is étale at P, so f is smooth on an open subset U of X containing P. Since the projection $\pi_1\colon Z \longrightarrow X$ is obtained from f by base extension, it follows that $U \times_Y X$ is smooth over X (by [Ha, Prop. III.10.1]) and thus over k. This contradicts the fact that (P, Q) lies in multiple components. □

Corollary 2.3. *Let $f\colon X \longrightarrow Y$ be a finite separable morphism of smooth curves and suppose that f is injective on $X(k)$. Then any component of $X \times_Y X$ other than the diagonal contains at most $(2g_X + 2\deg f - 2)$ k-rational points.*

Proof. Let D be a nondiagonal component of $Z = X \times_Y X$. Since f is injective on $X(k)$, every point in $Z(k)$ has the form (P, P); hence $D(k)$ lies in the support of the intersection of D with the diagonal. It follows that all points in $D(k)$ have the form (P, P) where f ramifies at P. By the Riemann-Hurwitz theorem, the number of such P is at most

$$2g_X - 2 - \deg f(2g_Y - 2) \leq 2g_X + 2\deg f - 2.$$

This completes the proof. □

Our next result generalizes the Weil bound for the number of $\mathbb{F}_q$-rational points on a smooth curve to the case of an arbitrary curve.

Proposition 2.4. *For any curve C over $\mathbb{F}_q$, we have*

$$\mid \#C(\mathbb{F}_q) - q - 1 \mid \leq 2p_a(C)\sqrt{q}.$$

Proof. Let $\tilde{C}$ be the normalization of C. Then $\tilde{C}$ is regular (by [Mat, Thm. 11.2]) and therefore smooth (by [Mat, Thms. 25.2 and 25.3]). The normalization map $\tilde{C} \longrightarrow C$ is an isomorphism away from at most $p_a(C) - g_C$ points of $\tilde{C}$, so

$$\mid \#\tilde{C}(\mathbb{F}_q) - \#C(\mathbb{F}_q) \mid \leq p_a(C) - g_C.$$

Since $\tilde{C}$ is smooth, we also have [We]

$$\mid \#\tilde{C}(\mathbb{F}_q) - q - 1 \mid \leq 2g_C\sqrt{q}.$$

Our result follows. □

Theorem 2.5. *Let $f\colon X \longrightarrow Y$ be a finite separable morphism of degree $n \geq 2$ between smooth curves over $\mathbb{F}_q$. Suppose that f induces an injection from $X(\mathbb{F}_q)$ into $Y(\mathbb{F}_q)$ and that*

$$\sqrt{q} > 2(n-2)^2 + 4(n-1)g_X + 1. \tag{2.5.1}$$

Then f is exceptional.

Proof. We argue by contradiction. Suppose that f is not exceptional. Then there exists a non-diagonal geometric component C of $X \times_Y X$ that is defined over $\mathbb{F}_q$. Since the projection maps from $X \times_Y X$ onto X are generically n-to-1, one sees that the projection maps restricted to C are generically at most $(n-1)$-to-1 (since the diagonal is also a

component of $X \times_Y X$). Proposition 2.1 implies that $p_a(C) \leq (n-2)^2 + 2g_X(n-1)$. Now Proposition 2.4 gives

$$\#C(\mathbb{F}_q) \geq q + 1 - 2p_a(C)\sqrt{q} \geq q + 1 - 2((n-2)^2 + 2g_X(n-1))\sqrt{q}.$$

On the other hand, by Corollary 2.3 we have $\#C(\mathbb{F}_q) \leq 2g_X + 2n - 2$. Finally, it is easily checked that (2.5.1) implies

$$q + 1 - 2((n-2)^2 + 2g_X(n-1))\sqrt{q} - (2g_X + 2n - 2) > 0,$$

so we have our contradiction. □

Remark 2.6. The proof of Theorem 2.5 can be modified to work under the weaker hypothesis that f is injective over non-branch points of Y. Injectivity is used only in Corollary 2.3; since this weaker hypothesis still implies that if (P, Q) is a k-rational point on $X \times_Y X$ then $f(P) = f(Q)$ is a branch point of f, we can replace the bound $(2g_X + 2\deg f - 2)$ from Cor. 2.3 with the bound $(2g_X + 2\deg f - 2)(\deg f - 1)$. This enlarges the bound (2.5.1) slightly (adding 2 to the right hand side is sufficient), but otherwise the reasoning is identical.

3. Chebotarev

In this section we prove analogs of Chebotarev's density theorem for normal varieties over a finite field, which we apply in the next two sections.

Let R be a commutative ring, and let A be a group of automorphisms of R. We denote the fixed ring R^A as B, and we say that R is a Galois extension of B. For a single element $a \in A$ we write R^a instead of $R^{\langle a \rangle}$. Fix a prime $\mathcal{Q}$ in R lying over a prime $\mathcal{P}$ in B, and let $D = D(\mathcal{Q}/\mathcal{P})$ and $I = I(\mathcal{Q}/\mathcal{P})$ denote the decomposition and inertia groups at $\mathcal{Q}$. If $\mathcal{J}$ is a prime ideal in the commutative ring Z, we write $m_{\mathcal{J}}$ for the field of fractions of $Z/\mathcal{J}$.

The following result is standard and easy (e.g., see [Bo, Thm. 2, p. 331]).

Lemma 3.1. *Suppose that $m_{\mathcal{P}}$ is perfect. Then*

1. *A is transitive on the set of primes of R lying over $\mathcal{P}$;*
2. *$m_{\mathcal{Q}}/m_{\mathcal{P}}$ is a finite Galois extension of degree $[D : I]$; and*
3. *$D/I \cong \mathrm{Gal}(m_{\mathcal{Q}}/m_{\mathcal{P}})$.*

The next result is also known (e.g., see [Wa]), but we include a proof for the sake of completeness. Let H be a subgroup of A, let $U = R^H$, and let $\mathcal{S}$ be the set of left cosets of H in A. We may view A as a group of permutations of the set $\mathcal{S}$.

Lemma 3.2. *The number of primes $\mathcal{J} \subset U$ lying over $\mathcal{P}$ such that $m_{\mathcal{J}} = m_{\mathcal{P}}$ is equal to the number of common orbits of D and I on $\mathcal{S}$. In particular, if $I = 1$ then the number of primes $\mathcal{J} \subset U$ lying over $\mathcal{P}$ such that $m_{\mathcal{J}} = m_{\mathcal{P}}$ is equal to the number of fixed points of D on $\mathcal{S}$.*

Proof. For $a \in A$, let $\mathcal{Q}' = a^{-1}\mathcal{Q}$ and $\mathcal{J} = \mathcal{Q}' \cap U$. Since $H \cap a^{-1}Da$ and $H \cap a^{-1}Ia$ are the decomposition and inertia groups for $\mathcal{Q}'$ over $\mathcal{J}$, we have

$$[m_{\mathcal{J}} : m_{\mathcal{P}}] = \frac{[m_{\mathcal{Q}'} : m_{\mathcal{P}}]}{[m_{\mathcal{Q}'} : m_{\mathcal{J}}]} == \frac{[a^{-1}Da : a^{-1}Ia]}{[H \cap a^{-1}Da : H \cap a^{-1}Ia]}.$$

Using the fact that $|H \cap a^{-1}Ma| = |M||H|/|MaH|$ for any subgroup M of G, we thus obtain

$$[m_{\mathcal{J}} : m_{\mathcal{P}}] = [D : I] \cdot \frac{|DaH|}{|D||H|} \cdot \frac{|I||H|}{|IaH|} = \frac{|DaH|}{|IaH|},$$

which is equal to 1 if and only if DaH is a common orbit of D and I.

For $b \in A$, we have $(b^{-1}\mathcal{Q}) \cap U = \mathcal{J}$ if and only if $DaH = DbH$. Thus, we achieve the desired result by summing over all orbits DbH of D on $\mathcal{S}$. □

Let W be a normal variety over the finite field k, and let V be a normal variety over a finite extension ℓ of k. Let $\rho : V \longrightarrow W$ be a finite, generically étale map of k-schemes. Write K and L for the fields of rational functions on W and V, so that ρ induces an inclusion $K \hookrightarrow L$. Assume that L/K is Galois, and put $A = \mathrm{Gal}(L/K)$ and $G = \mathrm{Gal}(L/K.\ell)$ (here $K.\ell$ denotes the compositum of K and ℓ in L). Then $A/G \cong \mathrm{Gal}(\ell/k)$ is cyclic. Pick $a \in A$ with $\langle aG \rangle = A/G$. Let t be an extension of k such that $[t : k] = \#\langle a \rangle$. Note that t contains ℓ. Pick an automorphism $\tilde{a}$ of the compositum $L.t$ such that $\tilde{a}|_L = a$ and $t^{\tilde{a}} = k$; such an automorphism exists because $\ell^a = k$. Then $(L.t)^{\tilde{a}} \supseteq R_i^a \supseteq B_i$, and k is algebraically closed in $(L.t)^{\tilde{a}}$.

Galoisness of L/K implies that $V^A = W$, in the sense that W admits an affine cover $M_i = \operatorname{Spec} B_i$ such that $\rho^{-1}(M_i) = \operatorname{Spec} R_i$ and $R_i^A = B_i$. Then each R_i is normal, so each B_i is as well ([Bo, V.1.9]). Furthermore, R_i is the integral closure of B_i in L since R_i is normal and integral over B_i. The ring $T_i = R_i.t$ is mapped to itself by $\tilde{a}$. Let V_t be the variety obtained from V by base-extension from ℓ to t, and let $V_t^{\tilde{a}}$ be the quotient variety of V_t obtained by piecing together the fixed rings $T_i^{\tilde{a}}$.

The degree of the field $L.t$ over the field of fractions of $T_i^{\tilde{a}}$ is equal to $\#\langle \tilde{a} \rangle = \#\langle a \rangle = [t : k]$, so $T_i^{\tilde{a}}.t$ has field of fractions $L.t$. Now, $T_i^{\tilde{a}}.t$ and T_i are both normal, because R and $T_i^{\tilde{a}}$ are normal ([Gr, 6.7.4]). Both

T_i and $T_i^{\tilde{a}}.t$ are integral over $T_i^{\tilde{a}}$ as well, so we must have $T_i = T_i^{\tilde{a}}.t$. Since T_i is a finite Galois extension of both $T_i^{\tilde{a}}$ and $B_i.t$, it is also a finite Galois extension of $T_i^{\tilde{a}} \cap B_i.t = B_i$.

We define the degree of a maximal ideal $\mathcal{I}$ in any of the rings B, R, T_i, $T_i^{\tilde{a}}$ to be $[m_{\mathcal{I}} : k]$. Let $J \in V_t^{\tilde{a}}(k)$ and let $\mathcal{J}$ be the corresponding degree one maximal ideal in some $T_i^{\tilde{a}}$. Then $T_i/T_i\mathcal{J} \cong k \otimes_\ell t \cong t$, so $T_i\mathcal{J}$ is the unique prime in T_i lying over $\mathcal{J}$. Thus, the map $\phi_i : \mathcal{J} \mapsto T_i\mathcal{J} \cap R_i$ gives a well-defined map from degree one maximal ideals in $T_i^{\tilde{a}}$ to maximal ideals in R_i.

Lemma 3.3. *Let $\mathcal{Q}$ be a maximal ideal in R_i that lies over a degree-one maximal ideal $\mathcal{P}$ of B_i. If $\langle aI(\mathcal{Q}/\mathcal{P})\rangle = D(\mathcal{Q}/\mathcal{P})/I(\mathcal{Q}/\mathcal{P})$, then there are exactly $[m_{\mathcal{Q}} : \ell]$ degree one maximal ideals $\mathcal{J} \in T_i^{\tilde{a}}$ such that $\phi_i(\mathcal{J}) = \mathcal{Q}$. Otherwise, there are no degree one maximal ideals $\mathcal{J} \in T_i^{\tilde{a}}$ such that $\phi_i(\mathcal{J}) = \mathcal{Q}$.*

Proof. Suppose that $\langle aI(\mathcal{Q}/\mathcal{P})\rangle = D(\mathcal{Q}/\mathcal{P})/I(\mathcal{Q}/\mathcal{P})$. Since $a\mathcal{Q} = \mathcal{Q}$, we must have $\tilde{a}T_i\mathcal{Q} = T_i\mathcal{Q}$. Thus $\tilde{a}$ acts on $T_i/T_i\mathcal{Q} \cong m_{\mathcal{Q}} \otimes_\ell t$ by acting as a acts on $m_{\mathcal{Q}}$ and as $\tilde{a}$ acts on ℓ. The primes in T_i lying over $\mathcal{Q}$ correspond to the primes in $m_{\mathcal{Q}} \otimes_\ell t$. Now, since a generates $\mathrm{Gal}(m_{\mathcal{Q}}/k)$ and $\tilde{a}$ generates $\mathrm{Gal}(t/k)$, there is a map $\psi : m_{\mathcal{Q}} \hookrightarrow t$ such that $\psi(ax) = \tilde{a}\psi(x)$. Then the $[m_{\mathcal{Q}} : \ell]$ primes in $m_{\mathcal{Q}} \otimes_\ell$ correspond to the kernels of the maps $p_j : m_{\mathcal{Q}} \otimes_\ell t \longrightarrow t$ given by $p_j(u \otimes v) = (\psi(a^{[\ell:k]j}u)v)$ for $0 \le j \le [m_{\mathcal{Q}} : \ell] - 1$. Since the kernel of p_j is the set of all $\sum_n (u_n \otimes v_n)$ such that $\sum_n \psi(a^{[\ell:k]j}u_n)v_n = 0$, the kernel of p_j is preserved by the action of $\tilde{a}$, so $\tilde{a}\mathcal{Q}' = \mathcal{Q}'$ for all $\mathcal{Q}'$ lying over $\mathcal{Q}$. Writing $\mathcal{J} = \mathcal{Q}' \cap T_i^{\tilde{a}}$, we then have $\langle \tilde{a} \rangle = D(\mathcal{Q}'/\mathcal{J})$ since T_i is unramified over $T_i^{\tilde{a}}$. For each of these $[m_{\mathcal{Q}} : \ell]$ maximal ideals $\mathcal{J}$, we have $\phi_i(\mathcal{J}) = \mathcal{Q}$.

Now, suppose that $\langle aI(\mathcal{Q}/\mathcal{P})\rangle \neq D(\mathcal{Q}/\mathcal{P})/I(\mathcal{Q}/\mathcal{P})$. Let $\mathcal{Q}'$ be a maximal ideal in T_i such that $\mathcal{Q}' \cap R_i = \mathcal{Q}$ and let $\mathcal{J} = \mathcal{Q}' \cap T_i^{\tilde{a}}$. If $a \notin D(\mathcal{Q}/\mathcal{P})$, then $\tilde{a}\mathcal{Q}' \neq \tilde{a}\mathcal{Q}'$, so there is more than one prime in T_i lying over $\mathcal{J}$, which means that $\mathcal{J}$ cannot have degree one. If $a \in D(\mathcal{Q}/\mathcal{P})$ but does not generate $D(\mathcal{Q}/\mathcal{P})/I(\mathcal{Q}/\mathcal{P})$, then $\mathcal{Q} \cap R_i^a$ has degree greater than one, by (3) of Lemma 3.1, so $\mathcal{J}$ does also, since $\mathcal{J}$ lies over $\mathcal{Q} \cap R_i^a$. □

If Q and P are closed points of V and W with $\rho(Q) = P$, we denote the decomposition and inertia groups of Q over P as $D(Q/P)$ and $I(Q/P)$, respectively. Clearly, these are the same as $D(\mathcal{Q}/\mathcal{P})$ and $I(\mathcal{Q}/\mathcal{P})$ where $\mathcal{Q}$ is a prime in some R_i corresponding to Q and $\mathcal{P}$ is a prime in B_i such that $\mathcal{Q} \cap B_i = \mathcal{P}$. Similarly, we define m_Q to be $m_{\mathcal{Q}}$.

Proposition 3.4. *With notation as above,*

$$\sum_{P\in W(k)} \sum_{\substack{\rho(Q)=P \\ \langle aI(Q/P)\rangle = D(Q/P)}} [m_Q : \ell] = \#V_t^{\tilde{a}}(k) \tag{3.4.1}$$

where $I(Q/P)$ is the inertia group of Q over P.

Proof. The maps ϕ_i patch together to form a map ϕ from $J(k)$ to closed points in V; indeed if we let ϕ be the map that takes a point $J \in W_t^{\tilde{a}}(k)$ to the closed point Q of W lying under the unique closed point of W_t that lies over J, then ϕ agrees with ϕ_i on each affine piece $\operatorname{Spec} W_t^{\tilde{a}}$. The proposition thus follows from Lemma 3.3. □

Corollary 3.5. *Suppose that V and W are nonsingular and projective. Let $r = \dim V$ and let $b_0, \ldots, b_{2r}$ be the Betti numbers (see* [Ha, p. 451 and 456]*) of V. Then*

$$\left| \Big(\sum_{P\in W(k)} \sum_{\substack{\rho(Q)=P \\ \langle aI(Q/P)\rangle = D(Q/P)}} [m_Q : \ell] \Big) - (\#k)^r \right| \le \left| \sum_{i=0}^{2r-1} (-1)^i b_i (\#k)^{i/2} \right|$$

Proof. Since $T_i^{\tilde{a}}.t = T_i$ on each affine piece $\operatorname{Spec} T_i$ of V_t, we see that $V_t^{\tilde{a}}$ with the base extended from k to t is isomorphic to V_t (i.e., $\left(V_t^{\tilde{a}}\right)_t \cong V_t$). Thus, $V_t^{\tilde{a}}$ is also nonsingular ([Gr, 6.7.4]), and V, V_t, and $V_t^{\tilde{a}}$ all have the same Betti numbers. Thus, applying the Weil bound ([De], see also [Ha, Appendix 3] for an overview) to $V_t^{\tilde{a}}(k)$ in (3.4.1) gives the desired result. □

Proposition 3.4 also gives rise to a generalization of the effective Chebotarev density theorem for curves that Murty and Scherk proved in [MS] (see also [FJ, Chapter 5]). Let $\mathcal{V}$ denote the set of all unramified points in $V(\bar{\ell})$ that lie over points in $W(k)$; let $\mathcal{V}_a$ denote the set of all points in $\mathcal{V}$ that correspond to closed points Q of V such that $I(Q/\rho(Q))$ is trivial and $\langle a\rangle = D(Q/\rho(Q))$. Note that counting points in $V(\bar{\ell})$ is different from counting closed points; each closed point Q on V corresponds to $[m_Q : \ell]$ distinct points in $V(\bar{\ell})$.

Corollary 3.6. *Suppose that V and W are nonsingular and projective. Let $r = \dim V$, let $b_0, \ldots, b_{2r}$ be the Betti numbers of V, let $c_0, \ldots, c_{2r}$ be the Betti numbers of W, and let U be the ramification locus of ρ in*

W (thought of as a subscheme of W). Then

$$\begin{aligned}\left|\#\mathcal{V}_a - \frac{\#\mathcal{V}}{\#G}\right| \leq & (\#G)(\#U(k)) + \left|\sum_{i=0}^{2r-1}(-1)^i b_i(\#k)^{i/2}\right| \\ & + \left|\sum_{i=0}^{2r-1}(-1)^i c_i(\#k)^{i/2}\right|.\end{aligned} \tag{3.6.1}$$

Proof. If $I(Q/P)$ is trivial than ρ does not ramify at Q, so Q does not lie over a point in the ramification locus of ρ. As noted above, each such Q corresponds to $[m_Q : \ell]$ points in $\mathcal{V}_a$. Letting $\mathcal{U}_a$ denote the set of $J \in V_t^{\tilde{a}}(k)$ lying over points in $U(k)$ and applying Proposition 3.4, we obtain $\#\mathcal{V}_a = \#V_t^{\tilde{a}}(k) - \#\mathcal{U}_a$, Since the degree of $V_t^{\tilde{a}}$ over W is $\#G$, we have $\#\mathcal{U}_a \leq (\#G)(\#U(k))$. Thus, the Weil bound for $V_t^{\tilde{a}}(k)$ yields

$$\begin{aligned}&(\#k)^r - \left|\sum_{i=0}^{2r-1}(-1)^i b_i(\#k)^{i/2}\right| - (\#G)(\#U(k)) \\ &\leq \mathcal{V}_a \\ &\leq (\#k)^r + \left|\sum_{i=0}^{2r-1}(-1)^i b_i(\#k)^{i/2}\right|.\end{aligned} \tag{3.6.2}$$

Similarly, we obtain

$$\begin{aligned}&(\#G)\left((\#k)^r - \left|\sum_{i=0}^{2r-1}(-1)^i c_i(\#k)^{i/2}\right| - (\#U(k))\right) \\ &\leq \mathcal{V} \\ &\leq (\#G)\left((\#k)^r + \left|\sum_{i=0}^{2r-1}(-1)^i c_i(\#k)^{i/2}\right|\right),\end{aligned} \tag{3.6.3}$$

by using the Weil bound for $W(k)$. Dividing (3.6.3) by $\#G$ and subtracting it from (3.6.2) yields (3.6.1). □

Remark 3.7. When V and W are smooth curves, Corollary 3.6 is a slight improvement of [MS, Theorem 1]. Note that in this case, the ramification locus corresponds to a finite set of points in $W(\bar{k})$. To make Corollary 3.6 completely explicit in the higher-dimensional case, one must use bounds on $U(k)$ (such as those that come from applying the Weil bounds to desingularizations of the components of U, for example).

Recall that g_C denotes the genus of a curve C.

Corollary 3.8. *Suppose that W and V are smooth projective curves. Let $\mathcal{U}$ be a finite subset of $W(k)$, and pick $a \in A$ with $\langle aG \rangle = A/G$. If*

$$\sqrt{\#k} \geq 2g_V + \sqrt{(\#G)(\#\mathcal{U})}, \tag{3.8.1}$$

then there is a closed point Q on V lying over a point $P \in W(k) \setminus \mathcal{U}$ such that $a \in D(Q/P)$ and $\langle aI(Q/P) \rangle = D(Q/P)/I(Q/P)$.

Proof. The Weil bound says that

$$V_t^{\tilde{a}}(k) \geq \#k + 1 - 2g_{V_t^{\tilde{a}}}\sqrt{\#k}.$$

As in Corollary 3.5, we have $g_{V_t^{\tilde{a}}} = g_V$. Now, (3.8.1) implies that

$$V_t^{\tilde{a}}(k) \geq 1 + \sqrt{\#k}\sqrt{(\#G)(\#\mathcal{U})} \geq (\#G)(\#\mathcal{U}).$$

Since the number of closed points of $V_t^{\tilde{a}}$ that lie over points in $\mathcal{U}$ is at most $([L.t : K])(\#\mathcal{U}) = (\#G)(\#\mathcal{U})$, it follows that there is a $J \in V_t^{\tilde{a}}(k)$ that lies over a point $P \in W(k) \setminus \mathcal{U}$. □

When V is singular, it is difficult to get something as uniform as Corollary 3.5, since we are not able to apply the Weil bound. It follows from the older estimate of Lang-Weil ([LW, Theorem 1]), however, that for any variety Z of dimension r over k, there is a constant δ (depending on Z) such that for all extensions k' of k, one has

$$|Z(k') - (\#k')^r| \leq \delta(\#k')^{r-\frac{1}{2}}.$$

This can be proved by induction on the dimension of Z. If $r = 0$, then Z is a point and we're done. Otherwise, let Z' be an affine subset of Z and let $\overline{Z'}$ be a projective closure of Z'. Applying the Lang-Weil estimate to $\overline{Z'}$ and the inductive hypothesis to $\overline{Z'} \setminus Z'$ and $Z \setminus Z'$ finishes the proof.

This allows us to treat the case of a single map $\rho : V \longrightarrow W$ with k and ℓ varying. Let k' be an extension of k and let $A_{K.(k' \cap \ell)}$ be the subgroup of A fixing $K.(k' \cap \ell)$. Each element in $a \in A_{K.(k' \cap \ell)}$ extends to an element $a'_k \in \mathrm{Gal}(L.k'/K.k')$ that acts as a does on L and acts trivially on k'. Since

$$\#\,\mathrm{Gal}(L.\ell/K.k') = \frac{\#A}{[k' \cap \ell : k]} = A_{K.(k' \cap \ell)},$$

every element of $\mathrm{Gal}(L.k'/K.k')$ is equal to $a_{k'}$ for some $a \in A_{K.(k' \cap \ell)}$. For convenience, we denote $\mathrm{Gal}(L.k'/K.k')$ as $A_{k'}$ and $\mathrm{Gal}(L.k'/(K.k'.\ell))$ as $G_{k'}$. We let $\rho_{k'}$ denote ρ with its base extended to k'; we have $\rho_{k'} : V_{k'.\ell} \longrightarrow W_{k'}$.

Corollary 3.9. *Let $r = \dim V$. There is a constant δ such that for any finite extension k', we have*

$$\left| \Big(\sum_{P \in W(k')} \sum_{\substack{\rho_{k'}(Q) \in P \\ \langle \sigma I(Q/P) \rangle = D(Q/P)}} [m_Q : (\ell.k')] \Big) - (\#k')^r \right| \leq \delta (\#k')^{r - \frac{1}{2}}$$

for any $\sigma \in A_{k'}$ such that $\langle \sigma G_{'k'} \rangle = A_{k'}/G_{k'}$.

Proof. Let $\sigma \in \mathrm{Gal}(L.k'/K.k')$. We may write $t = a_{k'}$ for some $a \in A_{K.(k' \cap k)}$ as we saw above. We define t and $\tilde{a}$ be before and define $\tilde{\sigma}$ to be the automorphism of $L.k'.t$ that acts as σ does on $L.k'$ and acts as $\tilde{a}$ does on t. Then $V^{\tilde{\sigma}}_{t.k'}$ is isomorphic to $\left(V^{\tilde{a}}_t\right)_{k'}$, so, by (3.4.1), we have

$$\begin{aligned} &\left| \Big(\sum_{P \in W(k')} \sum_{\substack{\rho_{k'}(Q) \in P \\ \langle \sigma I(Q/P) \rangle = D(Q/P)}} [m_Q : (\ell.k')] \Big) - (\#k')^r \right| \\ &= \left| V^{\tilde{\sigma}}_{t.k'}(k') - (\#k')^r \right| \\ &\leq \delta_a (\#k')^{r - \frac{1}{2}}, \end{aligned}$$

for some constant δ_a depending only on a. Letting δ be the maximum of all of the δ_a gives the desired result. □

4. Exceptionality

Let X and Y be normal, geometrically irreducible varieties over $\mathbb{F}_q$, and let $f : X \longrightarrow Y$ be a finite, generically étale $\mathbb{F}_q$-morphism. In this section we give Lenstra's Galois-theoretic proof that exceptionality of f implies bijectivity of the induced map $X(\mathbb{F}_q) \to Y(\mathbb{F}_q)$. We then use the same Galois-theoretic setup to show that, if X and Y are curves and q is sufficiently large compared to n and g_X, then injectivity, surjectivity, and exceptionality are equivalent.

We begin with some notation. We will view the function field $\mathbb{F}_q(Y)$ as a subfield of $\mathbb{F}_q(X)$, via the inclusion $\mathbb{F}_q(Y) \hookrightarrow \mathbb{F}_q(X)$ induced by f. Since f is generically étale, the extension $\mathbb{F}_q(X)/\mathbb{F}_q(Y)$ is separable. Let Ω be the Galois closure of this extension. Let $\mathbb{F}_{q^k}$ be the algebraic closure of $\mathbb{F}_q$ in Ω. Put $A = \mathrm{Gal}(\Omega/\mathbb{F}_q(Y))$ and $G = \mathrm{Gal}(\Omega/\mathbb{F}_{q^k}(Y))$; then G is a normal subgroup of A and $A/G \cong \mathrm{Gal}(\mathbb{F}_{q^k}/\mathbb{F}_q)$ is cyclic of order k. Let $H = \mathrm{Gal}(\Omega/\mathbb{F}_q(X))$. We may view A as a group of permutations of the set $\mathcal{S}$ of left cosets of H in A. Note that G acts transitively on $\mathcal{S}$, and that $n := \#\mathcal{S}$ is the degree of f.

We first give the standard (transparent) group-theoretic translation of exceptionality.

Lemma 4.1. *f is exceptional if and only if the only A-orbit on $\mathcal{S} \times \mathcal{S}$ which is also a G-orbit is the diagonal.*

The next two lemmas provide a group-theoretic counting argument which we use to relate exceptionality to injectivity and surjectivity. These are variants of [Co, Lemma 6] and [FGS, Lemma 13.1].

Lemma 4.2. *Let A be a finite group acting on a finite set $\mathcal{T}$, let G be a normal subgroup of A with A/G cyclic, and let aG be a generator of A/G. Then the number of A-orbits on $\mathcal{T}$ which are also G-orbits equals*

$$\frac{1}{\#G} \sum_{\alpha \in aG} \#\mathcal{T}^\alpha,$$

where $\mathcal{T}^\alpha$ denotes the set of fixed points of α on $\mathcal{T}$.

Proof. By examining the different A-orbits separately, we may assume that A is transitive on $\mathcal{T}$. If G is transitive then ag has a fixed point for some $g \in G$ (since, if a maps $s \mapsto t$, we can choose g mapping $t \mapsto s$); conversely, if some ag has a fixed point then the G-orbit containing this fixed point must also be an A-orbit, hence (since A is transitive) must equal $\mathcal{T}$. Thus G fails to be transitive if and only if both sides of the equation are zero. So assume that A and G are both transitive (so they have precisely one common orbit).

Put

$$\mathcal{V} = \{(\alpha, t) \in aG \times \mathcal{T} \colon \alpha(t) = t\}.$$

Let A_t be the stabilizer of t in A (and similarly for G). On the one hand, if $\alpha \in aG$ fixes t, then $A_t \cap aG = \alpha G_t$, so there are $\#G_t$ elements in $A_t \cap aG$; hence $\#\mathcal{V} = (\#G_t)(\#\mathcal{T}) = \#G$. On the other hand, $\#\mathcal{V} = \sum_{\alpha \in aG} \#\mathcal{T}^\alpha$, and the result follows. □

Lemma 4.3. *Let A be a finite group acting on a finite set $\mathcal{S}$, and let G be a transitive normal subgroup of A with A/G cyclic. Then the following properties are equivalent:*

1. *The only A-orbit on $\mathcal{S} \times \mathcal{S}$ which is also a G-orbit is the diagonal.*
2. *Every $a \in A$ with $\langle aG \rangle = A/G$ has a unique fixed point in $\mathcal{S}$.*
3. *Every $a \in A$ with $\langle aG \rangle = A/G$ has at most one fixed point in $\mathcal{S}$.*
4. *Every $a \in A$ with $\langle aG \rangle = A/G$ has at least one fixed point in $\mathcal{S}$.*

Proof. Since G is transitive on $\mathcal{S}$, by applying the previous lemma to $\mathcal{T} = \mathcal{S}$ we see that the average of $\#\mathcal{S}^\alpha$ (over all α in a generating coset of A/G) is 1. Thus (2), (3), and (4) are all equivalent to one another.

Applying the previous lemma to $\mathcal{T} = \mathcal{S} \times \mathcal{S}$ shows that (1) is equivalent to the average of $(\#\mathcal{S}^\alpha)^2$ being 1. Since $(\#\mathcal{S}^\alpha)^2 \geq \#\mathcal{S}^\alpha$, with equality if and only if $\#\mathcal{S}^\alpha$ is 0 or 1, it follows that (1) is equivalent to having every $\#\mathcal{S}^\alpha \leq 1$. Hence (1) and (3) are equivalent, which completes the proof. □

Combining these three lemmas with Lemma 3.2 yields a proof that exceptionality implies bijectivity.

Proposition 4.4 (Lenstra)**.** *If f is exceptional then f is bijective on rational points.*

Proof. Let $P \in Y(\mathbb{F}_q)$. By the definition of finiteness ([Ha, p. 84]), there is an affine open subset $M = \operatorname{Spec} B$ of Y with $P \in M(\mathbb{F}_q)$ such that $f^{-1}(M)$ is affine and can be written as $\operatorname{Spec} U$ for a ring U that is finite (and therefore integral) over B. Since X is normal, U must be the integral closure of B in $\mathbb{F}_q(X)$. Let R be the integral closure of B in Ω. As at the beginning of the section, we let $A = \operatorname{Gal}(\Omega/\mathbb{F}_q(Y))$, let $H = \operatorname{Gal}(\Omega/\mathbb{F}_q(X))$, and let $\mathcal{S}$ be the set of left cosets of H in A. Then $R^H = U$ and $R^A = B$. Let $\mathcal{P}$ be the prime in B corresponding to P and let D and I be the decomposition and inertia groups at some prime $\mathcal{Q}$ of R lying over $\mathcal{P}$. Then D/I is cyclic and (since $\mathcal{P}$ has degree one) $DG = A$. Since f is exceptional, Lemmas 4.1 and 4.3 show that every $a \in A$ with $\langle aG \rangle = A$ has a unique fixed point in $\mathcal{S}$. Since $A = DG$ and $I \subseteq D \cap G$, every $d \in D$ with $\langle dI \rangle = D$ also satisfies $\langle dG \rangle = A$ and hence has a unique fixed point in $\mathcal{S}$. By Lemma 4.2, D and I have a unique common orbit on $\mathcal{S}$. Thus, Lemma 3.2 implies there is exactly one maximal ideal $\mathcal{J}$ in U lying over $\mathcal{P}$ such that $m_{\mathcal{J}} = m_{\mathcal{P}}$, which means there is a unique point $J \in X(\mathbb{F}_q)$ such that $f(J) = P$. Hence, f is bijective on rational points. □

We now restrict to the case $\dim X = 1$, and prove the converse to Proposition 4.4 for sufficiently large q. To give an explicit bound on q we need the following estimate on the genus of the Galois closure Ω of $\mathbb{F}_q(X)/\mathbb{F}_q(Y)$. Here $n = \deg f$.

Lemma 4.5. *The genus of Ω satisfies*

$$g_\Omega \leq 1 + \#G \cdot \frac{g_X - 1 - (n-2)(g_Y - 1)}{2} \leq 1 + n! \cdot \frac{g_X + n - 3}{2}.$$

Proof. Combining Riemann-Hurwitz with Hilbert's formula for the degree of the ramification divisor gives the formula

$$2g_X - 2 = n(2g_Y - 2) + \sum_{y \in Y} \sum_{i \geq 0} \frac{n - \#(\mathcal{S} \backslash G_i(y))}{[G_0(y) : G_i(y)]},$$

where $G_i(y)$ denotes the i-th higher ramification group (in the lower numbering) of $\Omega/\mathbb{F}_q(Y)$ at a point of Ω lying over y, and $\mathcal{S}\backslash U$ denotes the set of orbits of U on $\mathcal{S}$. But any group U of permutations on $\mathcal{S}$ trivially satisfies $(\#U)(n-(\#\mathcal{S}\backslash U)) \geq 2(\#U-1)$, so we get

$$\begin{aligned}\frac{g_\Omega - 1}{\#G} &= g_Y - 1 + \frac{1}{2}\sum_{y\in Y}\sum_{i\geq 0}\frac{\#G_i(y)-1}{\#G_0(y)} \\ &\leq g_Y - 1 + \frac{1}{4}\sum_{y\in Y}\sum_{i\geq 0}\frac{n-\#(\mathcal{S}\backslash G_i(y))}{[G_0(y):G_i(y)]} \\ &= g_Y - 1 + \frac{g_X-1}{2} - n\frac{g_Y-1}{2}.\end{aligned}$$

The result follows. □

Remark 4.6. The above lemma bounds g_Ω in terms of $\#G$, and then applies the trivial bound $\#G \leq n!$. However, it is often the case that $\#G$ is much smaller than $n!$. For instance, if G is primitive (meaning that, geometrically, the cover $X \to Y$ does not have a proper subcover) and G is not S_n or A_n, then $\#G$ is vastly less than $n!$. Indeed, with explict exceptions, the order of the group will be polynomial in n. Without exceptions, one knows that a primitive subgroup of S_n which doesn't contain A_n must have order less than 4^n [PS].

Note also that if $g_X > 1$ then we have the lower bound

$$g_\Omega \geq 1 + \#G \cdot (g_X - 1)/n,$$

so our upper bound has the right order of magnitude in this situation.

Combining the above lemmas with Corollary 3.8 yields the main result of this section:

Theorem 4.7. *Let $f : X \longrightarrow Y$ be a finite, separable morphism of smooth projective curves and let n be the degree of f. Suppose that $\sqrt{q} \geq n!(3g_X+3n)$ and that f is either injective or surjective on rational points. Then f is exceptional, and is bijective on rational points.*

Proof. Let a be an element of A such that $\langle aG\rangle = A/G$. Let $\mathcal{U}$ be the set of points in $Y(\mathbb{F}_q)$ over which f ramifies. By Riemann-Hurwitz, $\#\mathcal{U} \leq (2g_X + 2n - 2)$. Combining this bound with Lemma 4.5 and the inequality $[\Omega : \mathbb{F}_q(Y)] \leq n!$, we see that

$$2g_\Omega + \sqrt{[\Omega:\mathbb{F}_q(Y)](\#\mathcal{U})} < n!(3g_X + 3n).$$

Now Corollary 3.8 implies there is a closed point Q of Ω which lies over a point $P \in Y(\mathbb{F}_q)$ such that Q/P is unramified and its decomposition group is generated by a. By Lemma 3.2, the number of points in $X(\mathbb{F}_q)$ lying over P equals the number of points of $\mathcal{S}$ fixed by a.

Thus, surjectivity of f implies property (4) of Lemma 4.3, and injectivity of f implies property (3). By Lemma 4.3 and Lemma 4.1, if f is either surjective or injective then f is exceptional, and hence (by Proposition 4.4) f is bijective. □

Remark 4.8. The above result (and its proof) remains valid under the weaker hypothesis that the map $X(\mathbb{F}_q) \to Y(\mathbb{F}_q)$ is either injective or surjective over non-branch points of Y. More generally, let $\mathcal{U}$ be any subset of $Y(\mathbb{F}_q)$ which includes all points of $Y(\mathbb{F}_q)$ over which f is ramified. The above proof shows that either injectivity or surjectivity of $X(\mathbb{F}_q) \setminus f^{-1}(\mathcal{U}) \to Y(\mathbb{F}_q) \setminus \mathcal{U}$ implies exceptionality of f (and hence bijectivity of $X(\mathbb{F}_q) \to Y(\mathbb{F}_q)$), so long as $\sqrt{q} \geq 2g_\Omega + \sqrt{\#G \cdot \#\mathcal{U}}$.

5. Examples and Further Directions

We first give examples of covers of curves $X \to Y$ over $\mathbb{F}_q$ which are surjective but not injective on rational points, as well as examples which are injective but not surjective. In these examples, the degree n of the cover is small relative to q, but the genus of X is as large as q.

Example 5.1. Let q be an odd prime power and let $n > 1$ divide $(q-1)/2$. Pick $a, \gamma \in \mathbb{F}_q^*$ with a an n-th power. Let X be the normalization of the affine curve

$$y^n = \gamma \prod_{t \in \mathbb{F}_q^* \setminus a} (x - t),$$

so X has genus $(n-1)(q-3)/2$. Let $f : X \longrightarrow \mathbb{P}^1_{\mathbb{F}_q}$ be the morphism induced by projection onto the x-coordinate. Then f is totally ramified over all the rational points of $\mathbb{P}^1_{\mathbb{F}_q}$ (including infinity) except for 0 and a. Moreover, $\prod_{t \in \mathbb{F}_q^* \setminus a}(0-t) = a^{-1}$ and $\prod_{t \in \mathbb{F}_q^* \setminus a}(a-t) = -a^{-1}$ are both n-th powers in $\mathbb{F}_q^*$. Thus, if γ is an n-th power in $\mathbb{F}_q^*$, then both $x = 0$ and $x = a$ split completely under f, so f is surjective on rational points but not injective. If γ is not an n-th power, then neither $x = 0$ nor $x = a$ is the image of an $\mathbb{F}_q$-point under f, so f is injective on rational points but not surjective. □

Our next two examples are functions of the same shape, all of which are bijective but only some of which are exceptional. The exceptional ones come from the largest known family of exceptional functions.

Example 5.2. We consider the case of degree-5 maps between genus-0 curves. It follows from Dickson's results [Di] that, if such a map is totally ramified over some $\mathbb{F}_q$-rational point and $q > 13$, then injectivity on rational points implies exceptionality. Here are examples showing the necessity of the ramification hypothesis: the rational function $(x^5 -$

$ax)/(x^4 - b)$ is non-exceptional and bijective on $\mathbb{F}_q$-points if (q, a, b) is either $(17, 10, 3)$ or $(29, 13, 4)$. □

Example 5.3. The above example has the same shape as some exceptional maps. Namely, if k is any field containing a primitive fourth root of unity i and a nonsquare b, then $f(x) = (x^5 - b(4i - 3)x)/(x^4 - b)$ is an exceptional map $\mathbb{P}^1 \to \mathbb{P}^1$ over k. These examples come from the construction in [Fr2] as follows. Let $\varphi\colon E \to E'$ be the 5-isogeny between the elliptic curves $E\colon w^2 = x^3 + xb(1 + 2i)/4$ and $E'\colon v^2 = u^3 + ub(1+2i)^5/4$ such that the nontrivial elements in the kernel of φ are the pairs (x, w) with $x^2 = -b/4$ and $w^2 = xbi/2$. Map E and E' to $\mathbb{P}^1$ by taking the quotient by the automorphism (of curves) $P \mapsto (0, 0) - P$. Then φ induces a map $\mathbb{P}^1 \to \mathbb{P}^1$ which is easily seen to be our f. More generally, the largest known supply of exceptional maps $\mathbb{P}^1 \to \mathbb{P}^1$ are maps induced from isogenies of elliptic curves (cf. [Fr2, GMS]). □

We next give a large class of exceptional covers of curves.

Example 5.4. Let C be a curve on an abelian variety A over $\mathbb{F}_q$, let $\varphi\colon A \to A$ be the multiplication-by-d map, and suppose d is coprime to both q and $\#A(\mathbb{F}_q)$. Suppose furthermore that $\varphi^{-1}(C)$ is geometrically irreducible. Then the map $\varphi^{-1}(C) \to C$ is exceptional. This follows from the fact that the induced map from $\varphi^{-1}(C)$ to C is bijective over any extension of ℓ of $\mathbb{F}_q$ that does not contain the field of definition of any of the points in $A(\overline{\mathbb{F}}_q)$ having order a nontrivial divisor of d.

We conclude by discussing possible higher-dimensional analogs of Theorem 2.

Conjecture 5.5. Let $f\colon X \longrightarrow Y$ be a finite, generically étale map of degree $n \geq 2$ between two smooth projective varieties of dimension r defined over $\mathbb{F}_q$. Then there exists a constant C, depending only on n, the dimensions of X and Y, and the ℓ-adic Betti numbers $b_1, \ldots, b_{2r-1}$ of X, such that if $q > C$ and f induces an injection or a surjection from $X(\mathbb{F}_q)$ to $Y(\mathbb{F}_q)$, then f is exceptional and gives a bijection from $X(\mathbb{F}_q)$ to $Y(\mathbb{F}_q)$.

We can prove Conjecture 5.5 for maps $f : \mathbb{P}^m \longrightarrow \mathbb{P}^m$ (we mean the implication of exceptionality since injectivity and surjectivity are equivalent) although we are not able to give a simple formula for C. Here is a sketch of the proof. If D is a geometric component of $\mathbb{P}^m \times_f \mathbb{P}^m$ defined over $\mathbb{F}_q$, then D is birational to a subvariety of $\mathbb{P}^{2m}$ of dimension m and degree at most $(2n)^m$. Thus, by Lang-Weil [LW], there is a constant C_1, depending only on n and m, such that $\#D(\mathbb{F}_q) \geq q^m - C_1 q^{m-1/2}$. Arguing as in Proposition 2.2, we also see that, since f is injective, $\#D(\mathbb{F}_q) \leq R(\mathbb{F}_q)$ where R is the ramification locus of f. Since R is a divisor of degree at most $2mn$ on $\mathbb{P}^m$, there is a constant

C_2, depending only on m and n such that $\#R(\mathbb{F}_q) \leq C_2 q^{m-1}$ (again by Lang-Weil), which contradicts our earlier lower bound on $\#D(\mathbb{F}_q)$ when q is sufficiently large.

Unfortunately, proving Conjecture 5.5 in general seems to be much more complicated since we cannot use Lang-Weil and are forced instead to attempt to control Betti numbers of various varieties that arise. One possibility, suggested by Lucien Szpiro, is to directly prove the equivalence of injectivity and surjectivity in higher dimensions by examining the induced maps from curves in X to curves in Y.

The best we can do for maps between general varieties is the following non-explicit version of Conjecture 5.5, where we allow the constant C to depend on the map f.

Proposition 5.6. *Let $f : X \longrightarrow Y$ be a finite separable map between normal varieties over $\mathbb{F}_q$. If f induces an surjective or injective map from $X(\mathbb{F}_{q^m})$ to $Y(\mathbb{F}_{q^m})$ for infinitely many m, then f is exceptional.*

Proof. Let D_f denote the ramification locus of f in X. Let $W = X \setminus D_f$, and let V be the normalization of W in Ω (the Galois closure of $\mathbb{F}_q(X)$ over $\mathbb{F}_q(Y)$)). Let $A_m = \mathrm{Gal}(\mathbb{F}_{q^m}(\Omega)/\mathbb{F}_{q^m}(W))$ and let $G_m = \mathrm{Gal}(\ell(\Omega)/\ell(W))$, where ℓ is the closure of $\mathbb{F}_{q^m}$ in $\Omega.\mathbb{F}_{q^m}$. By Corollary 3.9, there exists M such that for any $m \geq M$ and any $\sigma \in A_m$ such that $\langle \sigma G_m \rangle = A_m/G_m$, there is a point $P \in W(\mathbb{F}_{q^m})$ and a closed point Q on $V_{\mathbb{F}_{q^m}}$ such that $\langle \sigma I(Q/P) \rangle = D(Q/P)/I(Q/P)$. For such m, injectivity and surjectivity each imply exceptionality by Lemma 4.3. □

We note that Fried ([Fr]) has previously proved Proposition 5.6 above in the case of maps from affine space to itself.

References

[Bo] N. Bourbaki, Elements of Mathematics, Commutative Algebra, Hermann, Paris, France, 1972.

[Co] S. D. Cohen, *The distribution of polynomials over finite fields*, Acta Arith. **17** (1970), 255–271.

[DL] H. Davenport and D. J. Lewis, *Notes on congruences. I*, Quart. J. Math. Oxford **14** (1963), 51–60.

[De] P. Deligne, *La conjecture de Weil. I*, Inst. Hautes Études Sci. Publ. Math. No. 43 (1974), 273–307.

[Di] L. E. Dickson, *The analytic representation of substitutions on a power of a prime number of letters, with a discussion of the linear group*, Ann. of Math. **11** (1896), 65–120.

[Fr] M. Fried, *On a theorem of MacCluer*, Acta Arith. **25** (1974), 121–126.

[Fr2] ———, *Galois groups and complex multiplication*, Trans. Amer. Math. Soc. **235** (1978), 141–163.

[Fr3] ———, *Global construction of general exceptional covers*, in: Finite Fields: Theory, Applications, and Algorithms, American Mathematical Society, Providence (1994), 69–100.

[FGS] M. Fried, R. Guralnick and J. Saxl, *Schur covers and Carlitz's conjecture*, Israel J. Math. **82** (1993), 157–225.

[FJ] M. Fried and M. Jarden, Field Arithmetic, Springer-Verlag, Berlin, 1986.

[GM] J. von zur Gathen and K. Ma, *The computational complexity of recognizing permutation functions*, Comput. Complexity **5** (1995), 76–97.

[Gr] A. Grothendieck, *Éléments de géométrie algébrique. IV. Étude locale des schémas et des morphismes de schémas. II*, Inst. Hautes Études Sci. Publ. Math. No. 24 (1967), 1–231.

[GMS] R. Guralnick, P. Mueller and J. Saxl, *The rational function analogue of a question of Schur and exceptionality of permutation representations*, Mem. Amer. Math. Soc. **162**, no. 773 (2003), 1–79.

[Ha] R. Hartshorne, Algebraic Geometry, Springer-Verlag, New York, 1977.

[LW] S. Lang and A. Weil, *Number of points on varieties in finite fields*, Amer. J. Math. **76** (1954), 819–827.

[LY] D. Leep and C. Yeomans, *The number of points on a singular curve over a finite field*, Arch. Math. **63** (1994), 420–426.

[LMZ] H. W. Lenstra, Jr., D. Moulton, M. Zieve, *Exceptional covers*, in preparation.

[Mac] C. R. MacCluer, *On a conjecture of Davenport and Lewis concerning exceptional polynomials*, Acta Arith. **12** (1967), 289–299.

[Mat] H. Matsumura, Commutative Ring Theory, Cambridge University Press, Cambridge, 1986.

[MS] V. K. Murty and J. Scherk, *Effective versions of the Chebotarev density theorem for function fields*, C.R. Acad. Sci. Paris Sér. I Math. **319** (1994), 523–528.

[PS] C. Praeger and J. Saxl, *On the orders of primitive permutation groups*, Bull. London Math. Soc. **12** (1980), 303–307.

[Wa] B. L. van der Waerden, *Die Zerlegungs- und Trägheitsgruppe als Permutationsgruppen*, Math. Ann. **111** (1935), 731–733.

[We] A. Weil, Variétés Abéliennes et Courbes Algébriques, Hermann, Paris, 1948.

[Wi] K. S. Williams, *On exceptional polynomials*, Canad. Math. Bull. **11** (1968), 279–282.

Department of Mathematics, University of Southern California, Los Angeles, CA 90089–1113, USA.

E-mail address: guralnic@math.usc.edu

Department of Mathematics, Hylan Building, University of Rochester, Rochester, NY 14627, USA.

E-mail address: ttucker@math.rochester.edu

Center for Communications Research, 805 Bunn Drive, Princeton, NJ 08540-1966, USA.

E-mail address: zieve@math.rutgers.edu

URL: http://www.math.rutgers.edu/∼zieve