TRANSACTIONS OF THE
AMERICAN MATHEMATICAL SOCIETY
Volume 00, Number 0, Pages 000–000
S 0002-9947(XX)0000-0

CATEGORICITY PROPERTIES FOR COMPUTABLE ALGEBRAIC FIELDS

DENIS R. HIRSCHFELDT, KEN KRAMER, RUSSELL MILLER,
AND ALEXANDRA SHLAPENTOKH

Abstract. We examine categoricity issues for computable algebraic fields. We give a structural criterion for relative computable categoricity of these fields, and use it to construct a field that is computably categorical, but not relatively computably categorical. Finally, we show that computable categoricity for this class of fields is Π^0_4-complete.

1. Introduction

Fields were the first class of structures for which the notion of computable categoricity was ever expressed. In their landmark study of effectiveness in field theory, Fröhlich and Shepherdson presented "two explicit fields which are isomorphic but not explicitly isomorphic" [9, Corollary 5.51]. In modern terminology, we would say that these two fields are both computable, and are classically isomorphic but not computably isomorphic. Thus they fail to satisfy the definition of computable categoricity.

Definition 1.1. The *Turing degree* of a countable structure $\mathcal{A}$ is the join of the degrees of the functions and relations of $\mathcal{A}$, or equivalently, the Turing degree of its atomic diagram. A *computable structure* is one with Turing degree $\mathbf{0}$.

A computable structure $\mathcal{A}$ is *computably categorical* if, for every computable structure $\mathcal{B}$ isomorphic to $\mathcal{A}$, there exists a computable isomorphism from $\mathcal{A}$ onto $\mathcal{B}$.

More generally, a computable structure $\mathcal{A}$ is *relatively computably categorical* if, for every structure $\mathcal{B}$ with domain ω that is isomorphic to $\mathcal{A}$, there exists an isomorphism from $\mathcal{A}$ onto $\mathcal{B}$ that is computable in the Turing degree of $\mathcal{B}$.

For these and other definitions from computable model theory, [1] and [15] are excellent sources. The article [33], written by two of the present authors, also serves to introduce these and many related concepts in more detail, and the articles [27]

2010 *Mathematics Subject Classification.* Primary 03D45, Secondary 03C57, 12L99.

The first author was partially supported by Grants # DMS–0801033 and DMS–1101458 from the National Science Foundation. The second author was partially supported by Grant # DMS–0739346 from the National Science Foundation. The third author was partially supported by Grant # DMS–1001306 from the National Science Foundation, by Grant # 13397 from the Templeton Foundation, by the Centre de Recerca Matemática, and by several grants from The City University of New York PSC-CUNY Research Award Program. The fourth author was partially supported by Grants # DMS–0650927 and DMS–1161456 from the National Science Foundation, by Grant # 13419 from the Templeton Foundation, and by an ECU Faculty Senate Summer 2011 Grant.

and [30] present the basic notions about computable fields for readers unfamiliar with them.

For over fifty years since that first result in [9], computable categoricity for fields has remained largely a mystery. For many other classes of structures, mathematicians have found structural definitions equivalent to computable categoricity: see for instance [13], [14], [20], [24], [26], [35], and [36]. As an example, Goncharov and Dzgoev, and independently Remmel, showed that a linear order is computably categorical if and only if it has only finitely many pairs of adjacent elements. (Two distinct elements of a linear order are *adjacent* if there is no element of the order between them.) This criterion is not quite expressible in first-order model theory, since it involves finiteness, but intuitively it is distinctly more "structural" than Definition 1.1. In terms of computational complexity this criterion is Σ^0_3 (and is readily shown to be complete at that level), whereas the statement of Definition 1.1 is Π^1_1, quantifying over all possible (classical) isomorphisms. Indeed, for linear orders, computable categoricity turns out to coincide with relative computable categoricity, and Ash, Knight, Manasse, and Slaman established in [2] that relative computable categoricity is always a Σ^0_3 property. (Unpublished work [4] by Chisholm yields the same result.) On the other hand, although relative computable categoricity clearly implies computable categoricity, it was established independently by Khoussainov and Shore in [19] and by Kudinov in [21] that the two concepts are not equivalent. More recently, Downey, Hirschfeldt, and Khoussainov showed in [5] that relative computable categoricity can be viewed as a kind of uniform version of computable categoricity, although this fact was already implicit in work of Ventsov [40].

For fields, however, only a few significant criteria for computable categoricity (or for its failure) have been discovered. The situation is straightforward when the field is algebraically closed: Ershov showed in [6] that such a field is computably categorical if and only if it has finite transcendence degree over its prime subfield (either $\mathbb{Q}$ or the p-element field $\mathbb{F}_p$, depending on characteristic). Earlier, Fröhlich and Shepherdson [9] had established that all normal algebraic extensions of $\mathbb{Q}$ and of $\mathbb{F}_p$ are computably categorical. These results failed to extend to fields more generally, however: algebraic extensions of $\mathbb{Q}$ that are not computably categorical have been known at least since [6], and Miller and Schoutens recently constructed a computably categorical field of infinite transcendence degree over $\mathbb{Q}$ (see [32]).

The transcendence degree of the field over its prime subfield is soon seen to be of paramount importance in these considerations. For algebraic field extensions F of $\mathbb{Q}$, one can identify each element $x \in F$ to within finitely many possibilities by finding the minimal polynomial of x in $\mathbb{Q}[X]$, and likewise for algebraic extensions of $\mathbb{F}_p$; this fact follows from the existence of *splitting algorithms* for $\mathbb{Q}$ and for each $\mathbb{F}_p$. When one wishes to compute an isomorphism between two such fields, the task of determining an image for x is not completely solved by this knowledge, but its degree of difficulty becomes relatively low; see [29] for the current state of knowledge on this topic. The paper [33], written by two of us, is in many ways a precursor to this paper, and produces a criterion for computable categoricity in case the entire algebraic field F has a splitting algorithm: such an F is computably categorical if and only if its orbit relation is decidable, in which case it is also relatively computably categorical. (The *orbit relation* holds of the pair $\langle a, b\rangle \in F^2$ if and only if some automorphism of F maps a to b.) Below we prove that this criterion does not extend to all computable algebraic fields; indeed both implications fail.

(For further background about splitting algorithms and related concepts, including the *splitting set* and the *root set* of F, we suggest [27], [28], and [33].)

This paper focuses on computable algebraic fields F. We do not assume the existence of a splitting algorithm for F, although our results do apply in the situation where F has a splitting algorithm. That case was mostly explained in [33], however, while here we show that the situation without a splitting algorithm is significantly more difficult. In particular, computable algebraic fields without splitting algorithms can be computably categorical without being relatively computable categorical (see Theorem 4.1). Moreover, the complexity of computable categoricity goes up when the field is not required to have a splitting algorithm: computable categoricity is Π^0_4-complete for algebraic fields (see Theorem 5.4), whereas with a splitting algorithm it is equivalent to relative computable categoricity, hence only Σ^0_3-complete. The increase in complexity is significant, but the switch from Σ to Π is also significant. Indeed, for algebraic fields, Definition 1.1 has complexity Π^0_4, since the property of being isomorphic is only Π^0_2, distinctly simpler than the usual Σ^1_1. (This fact is based on Corollary 2.7.) Therefore, our results show that the standard definition of computable categoricity actually has the minimum possible complexity all by itself, when restricted to algebraic fields: no structural (or other) criterion can improve it. To our knowledge, algebraic fields are the first class of structures for which this has been shown to be the case.

2. Useful Results on Computable Fields

Substantial work on computable algebraic fields and categoricity has appeared recently, giving rise to several useful techniques for constructing computable fields. In this section we review assorted properties of algebraic fields relevant to these techniques, with references to allow the reader to look up their proofs and to see how they were originally used.

The following result, which appears as Lemma 2.10 in [29], will often save us from having to worry about surjectivity as we compute isomorphisms between fields.

Lemma 2.1. *For an algebraic field F, every endomorphism (i.e. every injective homomorphism $g : F \to F$) is an automorphism.*

Corollary 2.2. *If $E \cong F$ are isomorphic algebraic fields, and $f : E \to F$ is a field embedding (by which we mean a field homomorphism with $f(1) \neq 0$), then the image of f is all of F. That is, such an f must be an isomorphism.*

We will use the standard notation for Galois groups: if $F \subseteq K$ is a *Galois extension* (i.e. an algebraic normal separable field extension), then $\mathrm{Gal}(K/F)$, the *Galois group* of K over F, is the group of all field automorphisms of K which restrict to the identity map on F. As we build computable fields, it frequently happens that, having already built a computable field F_s, we wait to see whether a particular function will converge on a particular input. If it does not converge, then F_s itself satisfies a particular requirement $\mathcal{R}_2$ for the construction, whereas if it does converge, we can add more elements to F_s to build the larger field K_2 and satisfy the requirement that way. When considering two distinct requirements, it is useful to be certain that extending F_s to K_2 to satisfy $\mathcal{R}_2$ will not disrupt our plan to build a different extension K_1 if necessary to satisfy a different requirement $\mathcal{R}_1$. Usually, if $\mathrm{Gal}(K_1/F_s) \cong \mathrm{Gal}(E/K_2)$ (where E is the field generated by K_1 and K_2 together), we can avoid the disruption to $\mathcal{R}_1$, and one way to ensure this

isomorphism holds is to make $K_1 \cap K_2 = F_s$. (See [18, p. 243, Exercise 2], for example.)

To achieve this end, we will often use a Galois extension of $\mathbb{Q}$ whose Galois group over $\mathbb{Q}$ is the symmetric group on the roots of a given polynomial, since this choice allows us to adjoin some of these roots immediately and keep others out of the field until needed. (The proof of Theorem 3.5 is a good example of such a construction.) Therefore, it is frequently useful for an extension such as the K_i above to have symmetric Galois group over the current ground field, as this property ensures that it is the splitting field of a polynomial whose roots are essentially all independent of each other. The first theorem for this purpose appeared as Theorem 2.15 in [28], and provides a supply of such extensions. The proof given there was devised by Kevin Keating. Since we also want these extensions not to interfere with each other (and since extensions with large symmetric Galois groups cannot be taken to have relatively prime degrees, which is the most obvious way to avoid such interference), we now extend that theorem to include the linear disjointness of the extensions.

Definition 2.3. Two Galois extensions $E \subseteq K$ and $E \subseteq L$ within a larger field F are *linearly disjoint* if $K \cap L = E$. (This is a particular case of the definition of linear disjointness for algebraic field extensions in general, which requires that K and L together generate an extension whose degree over E is the product $[K : E] \cdot [L : E]$.)

This means that we can add elements of K to E to build F to satisfy one requirement, and close under the field operations, without worrying that these new elements might accidentally force certain elements of L to enter F as well and thereby upset our satisfaction of a different requirement. The simplest case of linear disjointness occurs when the degrees $[K : E]$ and $[L : E]$ are relatively prime: the degree $[K \cap L : E]$ divides both, hence equals 1, so $K \cap L = E$. (Indeed, in this situation K and L need not be Galois extensions of E.)

Proposition 2.4 gives the recursive step for our procedure for building many distinct extensions of $\mathbb{Q}$, each one linearly disjoint from the field generated by all the rest. As explained in [39, §8.10], polynomials over $\mathbb{Q}$ whose Galois group is the symmetric group $\mathcal{S}_n$ can be constructed by using the fact that the only transitive subgroup of $\mathcal{S}_n$ containing a transposition and an $(n-1)$-cycle is $\mathcal{S}_n$ itself. In Proposition 2.4, we force the Galois group to contain such elements by putting together local behavior at suitable primes. In Lemma 3.6 below, we will extend these ideas to a recursive procedure for creating a sequence of polynomials $f_0, f_1, \ldots$ such that $\deg(f_i) = d_i$, $\mathrm{Gal}(\mathbb{Q}(f_i)/\mathbb{Q}) \simeq \mathcal{S}_{d_i}$ and $\mathrm{Gal}(K/\mathbb{Q}) \simeq S_{d_1} \times \cdots \times S_{d_n}$, where K is the compositum of the splitting fields of the f_i's. Thus the splitting field of any f_i is linearly disjoint over $\mathbb{Q}$ from the compositum of the splitting fields of all the others. (Actually, in Lemma 3.6, every d_i will equal 7, but we could have used any computable sequence $\langle d_i \rangle_{i \in \omega}$ instead.)

Proposition 2.4. *Fix any Galois extension $E/\mathbb{Q}$ and any $d > 1$. Then there is a monic irreducible polynomial $f(X)$ in $\mathbb{Z}[X]$ of degree d such that $Gal(K/\mathbb{Q}) \cong Gal(E/\mathbb{Q}) \times \mathcal{S}_d$, where $K = EF$ is the compositum of E and the splitting field F of f over $\mathbb{Q}$. In particular, E and F are linearly disjoint over $\mathbb{Q}$, with $Gal(F/\mathbb{Q}) \cong \mathcal{S}_d$.*

Proof. First we recall some notation and background information. Let $\mathbb{Z}_p$ be the ring of integers in the p-adic field $\mathbb{Q}_p$ and let $\mathbb{F}_p$ denote the field with p elements. By Hensel's Lemma (see e.g. [39, §18.4]), for a monic $h \in \mathbb{Z}[X]$ with mod-(p) reduction $\overline{h} \in \mathbb{F}_p[X]$, if $c \in \mathbb{F}_p$ is a simple root of $\overline{h}$, then there is a root $\alpha \in \mathbb{Z}_p$ of h which,

modulo p, is equal to c itself. If $\overline{h}$ is a product of distinct linear factors over $\mathbb{F}_p$, then h splits completely into linear factors over $\mathbb{Z}_p$, by applying this method to each factor over $\mathbb{F}_p$. This will be used below to satisfy the conditions P and R.

The finite field $\mathbb{F}_{q^d}$ is a Galois extension of $\mathbb{F}_q$ of degree d, with cyclic Galois group generated by the Frobenius automorphism $x \mapsto x^q$. Let $\varphi(X)$ be the minimal polynomial over $\mathbb{F}_q$ for a primitive generator of $\mathbb{F}_{q^d}$. Then $\varphi(X)$ has degree d and splits completely in $\mathbb{F}_{q^d}$, with distinct roots. Now the unique unramified extension L of degree d over $\mathbb{Q}_q$ may be constructed as follows. Choose $\Phi(X) \in \mathbb{Z}[X]$ monic of degree d such that $\Phi \equiv \varphi \pmod{q}$ and let L be the field obtained by adjoining a root of Φ to $\mathbb{Q}_q$. Then L is unramified over $\mathbb{Q}_q$, since the reduction of Φ in $\mathbb{F}_q[X]$ is the separable polynomial φ. (In contrast, 0 is a repeated root in $\mathbb{F}_q$ of the reduction of $X^2 - q$, and the splitting field $\mathbb{Q}_q[\sqrt{q}]$ is *ramified* over $\mathbb{Q}_q$.) Hensel's Lemma shows that Φ splits completely in L and $\mathrm{Gal}(L/\mathbb{Q}_q) \simeq \mathrm{Gal}(\mathbb{F}_{q^d}/\mathbb{F}_q)$ also is cyclic of order d. This will be used in conditions Q and R below.

Now we address Proposition 2.4 itself. The development here follows ideas explained more fully in [22, VII, §2]. The Chebotarev Density Theorem [23, §VIII.4, Thm. 10] guarantees that there are distinct primes $p, q, r \geq d$ completely split in $E/\mathbb{Q}$. (This means that E embeds into each of the fields $\mathbb{Q}_p$, $\mathbb{Q}_q$, and $\mathbb{Q}_r$.) Fixing these primes, we now state the conditions we wish our polynomial f to satisfy, and explain why such an f must exist. Then we will show how the conditions imply the theorem.

- P: $f \equiv (X^2 - \eta)u(X) \bmod p$, for some $\eta \in \mathbb{Z}$ such that $\overline{\eta}$ is not a square in $\mathbb{F}_p$ and some $u(X) \in \mathbb{Z}[X]$ of degree $(d-2)$ such that $\overline{u}(X)$ splits completely into distinct linear factors over $\mathbb{F}_p$.
- Q: f is congruent modulo q to the minimal polynomial of a generator for the unique unramified extension of degree d over $\mathbb{Q}_q$.
- R: $f \equiv X \cdot w(X) \bmod r$, where $w(X) \neq X$ and w is the minimal polynomial of a generator for the unramified extension of degree $d-1$ over $\mathbb{Q}_r$.

Each condition requires that f be congruent to a particular monic polynomial of degree d, modulo one of the distinct primes p, q, and r. So the Chinese Remainder Theorem allows us to choose coefficients for a monic polynomial $f \in \mathbb{Z}[X]$ of degree d satisfying all three of these conditions. Let F be its splitting field over $\mathbb{Q}$, and set $K = EF$, so both $\mathrm{Gal}(F/\mathbb{Q})$ and $\mathrm{Gal}(K/E)$ may be seen as subgroups of $\mathcal{S}_d$. Each of the three conditions will yield a specific element of $\mathrm{Gal}(K/E)$, and the three elements together will imply that $\mathrm{Gal}(K/E)$ is all of $\mathcal{S}_d$. The process is stated in [22, Thm. VII.2.9], and also in Example 7 of the preceding chapter (p. 274). Here we sketch it for the specific case of condition P.

The condition P will yield a transposition in the Galois group $\mathrm{Gal}(F/\mathbb{Q})$. The polynomial $f(X) \in \mathbb{Z}[X]$ of degree d reduces modulo p to $\overline{f}(X)$ of the form given in condition P, and the factorization there (along with the fact that η is not a square modulo p) shows that the splitting field of $\overline{f}(X)$ over $\mathbb{F}_p$ must be a copy of $\mathbb{F}_{p^2}$, with Galois group generated by the automorphism Φ of $\mathbb{F}_{p^2}$ which interchanges the two square roots of $\overline{\eta}$ and fixes each of the other roots of $\overline{f}$. Now $X^2 - \eta$ can be viewed as a polynomial in $\mathbb{Q}_p[X]$, since $\mathbb{Z} \subset \mathbb{Z}_p \subset \mathbb{Q}_p$, and its splitting field L over $\mathbb{Q}_p$ is the unique unramified extension of $\mathbb{Q}_p$ of degree 2, and in fact is the splitting field of f over $\mathbb{Q}_p$, since the roots of f in $\mathbb{F}_p$ yield distinct roots of f in $\mathbb{Z}_p$ by Hensel's Lemma, as argued above. The Galois group $\mathrm{Gal}(L/\mathbb{Q}_p) = \langle \Phi_p \rangle$ is cyclic of order 2, with Φ_p being the lift of Φ from $\mathbb{F}_{p^2}$ to L. This Φ_p must be a transposition, since

it has order 2 and must fix every other root of f. For details, see [22, Proposition VII.2.8]. We argue next that this transposition lifts to a transposition in $\mathrm{Gal}(F/\mathbb{Q})$, by setting $e = 2$ and $g(X) = X^2 - \eta$ in the following lemma.

Lemma 2.5. *Let $E/\mathbb{Q}$ be Galois, $f(X)$ monic and irreducible in $\mathbb{Z}[X]$, F the splitting field of f over $\mathbb{Q}$, and $K = EF$. Assume further that p is a prime completely split in $E/\mathbb{Q}$, that $e > 1$ is an integer, and that $g(X) \in \mathbb{Z}[X]$ is the minimal polynomial of a generator for the unramified extension of degree e over $\mathbb{Q}_p$. If $\overline{f}(X)$ is the product of $\overline{g}(X)$ and distinct linear factors in $\mathbb{F}_p[X]$, then Gal(K/E) contains an automorphism which cyclically permutes e of the roots of f and fixes each remaining root.*

Proof. We can view $g(X)$ as a polynomial in $\mathbb{Q}_p[X]$, since $\mathbb{Z} \subset \mathbb{Z}_p \subset \mathbb{Q}_p$, and its splitting field L over $\mathbb{Q}_p$ is the unique unramified extension of $\mathbb{Q}_p$ of degree e. The Galois group $\mathrm{Gal}(L/\mathbb{Q}_p)$ is cyclic of order e, generated by some Φ_p which cyclically permutes the roots of g. (Again we refer the reader to [22, Proposition VII.2.8] for details.) Now F is the splitting field over $\mathbb{Q}$ of $f(X)$, and $f(X) \in \mathbb{Q}_p[X]$ via the inclusion $\mathbb{Z} \subset \mathbb{Q}_p$. Also, by assumption E embeds into $\mathbb{Q}_p$. We can extend this embedding to an embedding of K into L by noting that $K = EF$ is generated over E by the roots of $f(X)$, which are all either roots of $g(X)$ or elements of $\mathbb{Q}_p$, since by Hensel's Lemma $f(X)$ is the product of $g(X)$ with linear factors in $\mathbb{Z}_p[X]$. Now the map $\mathrm{Gal}(L/\mathbb{Q}_p) \to \mathrm{Gal}(K/E)$ by restriction (to the image of K within L) is an injective group homomorphism, since L is generated over $\mathbb{Q}_p$ by the roots of f in L. So the restriction of Φ_p to K is the desired element in $\mathrm{Gal}(K/E)$. □

With $e = 2$ and $g(X) = X^2 - \eta$, this lemma, gives us the map $\Phi_p \restriction K$ in $\mathrm{Gal}(K/E)$ and shows it to be a transposition, as required. We next use Lemma 2.5 to satisfy conditions Q and R. For Q, we set $e = d$ and let $g(X)$ be the polynomial shown in condition Q to be congruent modulo q to $f(X)$. The lemma (with q in place of the prime p) shows that $\mathrm{Gal}(K/E)$ contains a cyclic permutation of order d, and therefore must act transitively on the d roots of $f(X)$ in K. (This proves that $f(X)$ is irreducible in $E[X]$.)

Finally we take $g(X)$ to be the polynomial $w(X)$ from condition R, with $e = d - 1$ and with r as the prime in Lemma 2.5. The lemma returns an element of $\mathrm{Gal}(K/E)$ which permutes $(d-1)$ of the roots of f cyclically and fixes the last root. But the existence of such a permutation, along with the transposition supplied by condition P and the transitivity of the group's action, prove that $\mathrm{Gal}(K/E) \cong \mathcal{S}_d$, the symmetric group on the d roots of $f(X)$ in K. (See the Theorem on p. 199 in [39, §8.10].) Therefore $[EF : E] = d! \geq [F : \mathbb{Q}]$, making E and F linearly disjoint over $\mathbb{Q}$ and forcing $\mathrm{Gal}(K/\mathbb{Q}) \cong \mathrm{Gal}(E/\mathbb{Q}) \times \mathcal{S}_d$ as desired. □

Let F be any computable algebraic field. That is, F is an algebraic field extension of its prime subfield Q. The field Q is isomorphic to either the p-element field $\mathbb{F}_p$ if $p = \mathrm{char}(F) > 0$, or else to the field of rational numbers. Every one of these possibilities for Q has a splitting algorithm, and since F is algebraic, this fact forces Q to be computable within F. (An element $x \in F$ lies in Q if and only if its minimal polynomial over Q is linear.)

Officially the domain of F is ω, but since the language of fields contains the symbols 0 and 1 already, we will instead write $x_0, x_1, x_2, \ldots$ for the elements of F. We view F as the union of an infinite chain of finitely generated subfields:

$$Q = F_0 \subseteq F_1 \subseteq F_2 \subseteq \cdots \subseteq F,$$

where $F_s = Q(x_0, \ldots, x_{s-1})$ for every s. The Effective Theorem of the Primitive Element (see [8], or [33, Theorem 3.11]) allows us to compute for each s a single element $z_s \in F_s$ that generates all of F_s; indeed we may assume $z_s = x_t$ for the least t such that $F_s = Q(x_t)$. We may compute the minimal polynomial $q_s(X) \in Q[X]$ of each z_s over Q, and also compute polynomials $p_s \in Q[X_0, \ldots, X_{s-1}]$ such that $p_s(z_0, \ldots, z_{s-1}, X)$ is the minimal polynomial of z_s over F_{s-1}.

Now, following [29] and [33], we define the *automorphism tree for* F to be the following subtree of $\omega^{<\omega}$:

$$I_F = \{\sigma \in \omega^{<\omega} : (\forall s < |\sigma|)\ p_s(x_{\sigma(0)}, \ldots, x_{\sigma(s)}) = 0\}.$$

More generally, for a computable field E isomorphic to F, with domain $\{y_0, y_1, \ldots\}$, we define the *isomorphism tree* for F and E to be

$$I_{FE} = \{\sigma \in \omega^{<\omega} : (\forall s < |\sigma|)\ \tilde{p}_s(y_{\sigma(0)}, \ldots, y_{\sigma(s)}) = 0\},$$

where $\tilde{p}_s \in \tilde{Q}[X_0, \ldots, X_s]$ is the image of p_s when its coefficients are mapped into the prime subfield $\tilde{Q}$ of E by the (unique) isomorphism between these prime subfields. It follows that isomorphisms from F onto E correspond precisely to paths through I_{FE}, and that this correspondence preserves Turing degrees. (For an isomorphism h, the path contains those nodes $\sigma \in I_{FE}$ with $y_{\sigma(i)} = h(z_i)$ for all $i < |\sigma|$; conversely, given a path f through I_{FE}, define $h(z_s) = y_{f(s)}$ for all s.) All of this material is described in detail in [29] and [33].

The isomorphism tree I_{FE} can be defined the same way (but perhaps should be renamed) for any algebraic fields E and F, whether or not they are isomorphic. In this more general case, paths through I_{FE} correspond to field embeddings of F into E. If no such embeddings exist, then I_{FE} must be a finite tree. The following useful corollary generalizes this statement.

Lemma 2.6. *Let F and E be algebraic fields of the same characteristic, and suppose that $x \in F$ and $y \in E$ have the same minimal polynomial over the prime subfields. If there is no field embedding $F \to E$ with $x \mapsto y$, then there exists some finitely generated subfield $F_0 \subseteq F$ containing x such that no field embedding of F_0 into E maps x to y.*

Proof. Build the tree I_{FE} as above, using an enumeration of F that lists x as its first element. Then $\langle y \rangle$ constitutes a node on this tree, through which there is no path. König's Lemma shows that I_{FE} contains only finitely many nodes above $\langle y \rangle$, all of height $< n$, say. Therefore, the field F_0 generated by the first n elements of F (in our enumeration) has no embedding into E with $x \mapsto y$. □

Corollary 2.7. *Two algebraic fields E and F are isomorphic if and only if every finitely generated subfield of each one embeds into the other.*

Proof. Apply Lemma 2.6, with x and y being the zero elements of their fields, to see that each field embeds into the other. Then apply Lemma 2.1 to the composition of the two embeddings. □

Corollary 2.8. *Two algebraic fields E and F are isomorphic if and only if they have the same characteristic and every polynomial over the prime subfield with a root in either E or F also has a root in the other. (Here we extend the unique isomorphism between the prime subfields to an isomorphism of the polynomial rings over those subfields.)*

Proof. For the nontrivial direction, we apply the Theorem of the Primitive Element, which states that each finitely generated subfield (of E, say) is generated by a single element x. Let $p(X)$ be the minimal polynomial of x over the prime subfield Q_E of E. The isomorphism $f : Q_E \to Q_F$ maps $p(X)$ to an irreducible polynomial $p^f(X) \in Q_F[X]$, which by hypothesis must have a root $y \in F$. But then $Q_E(x) \cong Q_E[X]/(p(X)) \cong Q_F[X]/(p^f(X)) \cong Q_F(y)$. So every finitely generated subfield of E embeds into F, and conversely. Corollary 2.7 completes the proof. □

Finally, we must also give some technical definitions for embeddings among fragments of fields, by which we mean subsets of a field that are not necessarily subfields themselves. This notion will arise when we examine all computable copies of a given field: we will need to consider the structure (if any) computed by each partial computable function φ_e, without knowing for certain whether it is a field or not.

Definition 2.9. Let Q be the prime field of a given characteristic, and $\overline{Q}$ a computable presentation of its algebraic closure, with Q regarded as a subfield. Suppose that C is a subset of ω on which two binary operations, which we denote by $+$ and $\cdot$, are partially defined, i.e., for $x, y \in C$, the values of $x + y$ and $x \cdot y$ may be other elements of C, or may be undefined. (These operations are not the usual addition and multiplication on ω; they are simply any binary operations, and we wish to determine whether the set ω forms a field under these operations.) Assume that C contains distinct elements x and y satisfying $x + x = x$ and $y \cdot y = y$; we refer to x as 0 and to y as 1, since they will represent identity elements in this fragment of a field. If these partial operations do not contradict associativity, commutativity, distributivity, cancellation, the identity properties of 0 and 1, or equality of the characteristics of C and Q, then C is a *field fragment* of that characteristic. To contradict associativity, for example, would require the existence of three elements $c_0, c_1, c_2 \in C$ such that the four products $(c_0 \cdot c_1)$, $(c_1 \cdot c_2)$, $(c_0 \cdot c_1) \cdot c_2$, and $c_0 \cdot (c_1 \cdot c_2)$ are all defined, but the last two of these products are distinct. Similarly, a contradiction to cancellation for $\cdot$ would consist of elements $c_0 \neq 0$ and $c_1 \neq c_2$ in C with $(c_0 \cdot c_1)$ and $(c_0 \cdot c_2)$ both defined and equal to each other. All the other properties mentioned are also universal statements (given that 0 and 1 have already been specified), and so a contradiction to any of them is simply an instantiation of the negation. (Notice that we do not check for the existence of inverses.)

For a field fragment C, we define the prime field fragment $Q_0 \subseteq C$ as follows. If Q has positive characteristic, then Q_0 contains 0 and all elements $1, (1 + 1), (1 + 1 + 1), \ldots$ that are defined in C. If Q has characteristic 0, fix some enumeration $\{q_0, q_1, \ldots\}$ of the rationals $\mathbb{Q}$, For each $q_i = \frac{m}{n}$ in turn, we check whether C possesses an element of the form $\frac{1+1+\cdots+1}{1+\cdots+1}$, with m 1's on top and n 1's below, and also possesses an additive inverse of that element. If so, we enumerate both elements into Q_0. If not, then we stop enumerating entirely, so Q_0 consists of only those elements enumerated so far. We make the obvious identification between elements of Q_0 and elements of $\mathbb{Q}$, noting that this identification must be unique, since C is a field fragment.

Suppose that $D = \{x_0, \ldots, x_m\} \subseteq C$. An *embedding of D into $\overline{Q}$* is a function $f : D \to \overline{Q}$ such that for every $i \leq m$, there is some $p_i \in Q_0[X_0, \ldots, X_i]$ with $p_i(x_0, \ldots, x_i) = 0$ in C (specifically, all sums and products in this calculation lie in C), such that $\overline{p_i}(f(x_0), \ldots, f(x_{i-1}), X_n)$ is irreducible in $Q(x_0, \ldots x_{i-1})[X_n]$ and $\overline{p_i}(f(x_0), \ldots, f(x_i)) = 0$ in $\overline{Q}$, where $\overline{p_i} \in \overline{Q}[X_0, \ldots, X_i]$ is the image of p_i via the unique embedding of Q_0 into $\overline{Q}$.

Finally, if C' is another field fragment and $D' \subseteq C'$ is finite, then an *embedding of D' into D* consists of embeddings f of D into $\overline{Q}$ and f' of D' into $\overline{Q}$ such that $\text{range}(f') \subseteq \text{range}(f)$. (We often think of $f^{-1} \circ f'$ as the actual embedding.)

The resulting lemma is easily proven.

Lemma 2.10. *For finite subsets D and D' of finite field fragments C and C', it is decidable whether there exists an embedding of D' into D. The procedure is uniform in the (partial) operations on C and C', but it is necessary to know the exact size of each finite set C, C', D, and D', and to be able to decide the domains of the partial operations.*

3. Relative Computable Categoricity

The definition of a computably categorical structure is often strengthened to consider noncomputable copies, yielding the notion of relative computable categoricity introduced in Definition 1.1. Relative computable categoricity is actually a more natural property than computable categoricity, in the sense that, by results in [2] and [4], it has a fairly simple syntactic characterization: a structure is relatively computably categorical if and only if it has a Σ^0_1 Scott family over finitely many parameters.

Definition 3.1. A Σ^0_1 *Scott family* for a structure $\mathcal{A}$ is a computable sequence $\theta_0(\bar{a}, x_0, \ldots, x_{n_0}), \theta_1(\bar{a}, x_0, \ldots, x_{n_1}), \ldots$ of $\exists$-formulas, where $\bar{a}$ is a finite tuple from $\mathcal{A}$, such that every tuple of elements from $\mathcal{A}$ satisfies at least one θ_i, and for each i, any two tuples satisfying θ_i can be interchanged by an automorphism of $\mathcal{A}$.

In contrast, no such nice characterization of computable categoricity is known. Indeed, while the property of having a Σ^0_1 Scott family is arithmetically a Σ^0_3 property, computable categoricity is known by work in [41] to be at least Π^0_4-hard. (In fact, we will show in Section 4 that it is this hard just within the class of algebraic fields.) In this section and the next two, we add further evidence in favor of the neatness of relative computable categoricity, by considering computable algebraic fields.

Following the notation of [33], we denote the *full orbit relation* on F by

$$A_F = \{\langle a_1, \ldots, a_n, b_1, \ldots, b_n\rangle \in F^{<\omega} : (\exists \sigma \in \text{Aut}(F))(\forall i \leq n)\ \sigma(a_i) = b_i\}.$$

The simple *orbit relation* then is $B_F = A_F \cap F^2$, the binary relation on F of being in the same orbit. By Lemma 2.6, if a pair $\langle x, x'\rangle$ does not lie in B_F, then there is some finitely generated subfield $F_t \subseteq F$ containing x such that no embedding of F_t into F can map x to x'. (Here $Q = F_0 \subseteq F_1 \subseteq \cdots \subseteq F$ can be any chain of subfields with union F. Normally we will set F_s to be the subfield generated by the first $(s-1)$ elements of F.) Of course, if x and x' are not conjugate over Q, then we will realize immediately that $\langle x, x'\rangle \notin B_F$; the interesting case is that in which x and x' are conjugate, yet lie in distinct orbits. Since there are only finitely many conjugates of x in F, the following definition makes sense.

Definition 3.2. Let F be a computable algebraic field, with prime subfield Q and such that F is the union of a chain $Q = F_0 \subseteq F_1 \subseteq \cdots \subseteq F$ of finitely generated subfields. The *orbit function* $h : F \to \omega$ of F defines $h(x)$ to equal the least t such that $x \in F_t$ but, for every $x' \in F$ that is conjugate to x over Q with $\langle x, x'\rangle \notin B_F$, there is no embedding of F_t into F that maps x to x'.

So the function h, on input x, considers all false conjugates x' of x in F: that is, all conjugates of x that do not lie in the orbit of x under automorphisms. (This notion is important for categoricity, since the false conjugates in F correspond to elements of computable copies of F to which we might mistakenly map x.) Therefore, h has the property that, for all $x \in F$,

$$[x \mapsto x' \text{ extends to } F_{h(x)}] \text{ if and only if } \langle x, x' \rangle \in B_F.$$

Moreover, $h(x)$ is the least number with this property. Of course, x might not have any false conjugates in F; in this case $h(x)$ is the least s with $x \in F_s$.

We note that the function h does depend on the choice of chain $Q = F_0 \subseteq F_1 \subseteq \cdots \subseteq F$. Normally we take this to be a computable chain (that is, we can compute a strong index for a finite generating set for each F_s, uniformly in s), and it is readily proven that with the orbit function for one computable chain as an oracle, we can compute an upper bound for the orbit function for any other computable chain. However, for many fields the orbit function will not be computable. Indeed, should it be computable, then F must be computably categorical. More generally, we have the following lemma.

Lemma 3.3. *For a computable algebraic field F, if the orbit function h (for some computable chain $Q = F_0 \subseteq F_1 \subseteq \cdots$ of finitely generated subfields with union F) is computably bounded, then F is relatively computably categorical.*

Proof. Fix some computable function $g : F \to \omega$ such that $g(x) \downarrow \geq h(x)$ for all x. It is straightforward to construct a Σ^0_1 Scott family for F, fixing a primitive generator z_s for each F_s and its minimal polynomial $q_s(Z) \in Q[Z]$. We also fix, for each $s < t$, a polynomial $p_{s,t} \in Q[Z, Y]$ such that $p_{s,t}(z_s, Y)$ is the minimal polynomial of z_t over the subfield $F_s = Q(z_s)$. (This definition can be done effectively.)

For each z_s, we have a formula $\gamma_s(Z)$ saying:

$$q_s(Z) = 0 \ \& \ \exists y \ p_{s,g(s)}(Z, y) = 0.$$

By Definition 3.2, and since $g(z_s) \geq h(z_s)$, every z satisfying $\gamma_s(z)$ lies in the orbit of z_s. For a general tuple of elements $x_0, \ldots, x_n \in F$, we find the least s with all $x_i \in F_s = Q(z_s)$, find polynomials $r_i(Z) \in Q[Z]$ with $r_i(z_s) = x_i$, and define the formula $\delta_{\vec{x}}(X_0, \ldots, X_n)$:

$$\exists z [\gamma_s(z) \ \& \ (\forall i \leq n) X_i = r_i(z)].$$

The Scott family is then just the set $\mathfrak{S} = \{\delta_{\vec{x}}(\vec{X}) : \vec{x} \in F^{<\omega}\}$, In fact, it is a particularly nice Scott family, since it involves no parameters from F. (Of course, polynomial equations with parameters from Q can be expressed entirely in terms of the constant symbols 0 and 1 and the operations of addition and multiplication.)

Since relative computable categoricity is equivalent to having a Σ^0_1 Scott family, this definition completes the proof of Lemma 3.3. However, to illuminate the use of the orbit function further, we will also give a direct proof of relative computable categoricity of F, by constructing a computable isomorphism from F onto an arbitrary computable field $\widetilde{F}$ isomorphic to F. Our construction relativizes easily to the degree of a noncomputable field $\widetilde{F}$. Of course, the prime subfield $Q = F_0$ is computable within the algebraic field F and has a unique embedding f_0 into $\widetilde{F}$; this f_0 is computable and has image $\tilde{Q}$, the prime subfield of $\widetilde{F}$. Moreover, f_0 is known

to extend to an isomorphism from F onto $\widetilde{F}$, since the two fields are assumed to be isomorphic and every isomorphism between them must restrict to f_0.

Now assume inductively that we have constructed an embedding $f_s : F_s \to \widetilde{F}$ that extends to some (not necessarily computable) embedding ρ of F into $\widetilde{F}$. Let $\tilde{q}_{s+1}$ and $\tilde{p}_{s+1,g(s+1)}$ be the images of the polynomials q_{s+1} and $p_{s+1,g(s+1)}$ under the map $f_0 : Q \to \tilde{Q}$ on their coefficients. To extend f_s to F_{s+1}, compute $g(s+1)$ and search for any two elements $\tilde{y}, \tilde{z} \in \widetilde{F}$ such that $\tilde{q}_{s+1}(\tilde{z}) = 0$ and $\tilde{p}_{s+1,g(s+1)}(\tilde{z}, \tilde{y}) = 0$, and such that the map $z_{s+1} \mapsto \tilde{z}$ would send z_s to $f_s(z_s)$. (Of course, $z_s \in F_{s+1} = Q(z_{s+1})$, so the choice of an image for z_{s+1} uniquely determines the image of z_s. Also, notice that these conditions are satisfied when $\tilde{y} = \rho(z_{h(s+1)})$ and $\tilde{z} = \rho(z_{s+1})$, so we must eventually find some $\tilde{y}$ and $\tilde{z}$ as desired, although they will not necessarily be $\rho(z_{h(s+1)})$ and $\rho(z_{s+1})$.) Define $f_{s+1}(z_s)$ to be this $\tilde{z}$, thus defining f_{s+1} on all of F_{s+1}. We claim that this f_{s+1} extends f_s and is a field embedding of F_{s+1} into $\widetilde{F}$, and also that f_{s+1} itself extends to some embedding of all of F into $\widetilde{F}$.

First, since $\tilde{q}_{s+1}(\tilde{z}) = 0$ and $q_{s+1}(z_{s+1}) = 0$, and since q_{s+1} and $\tilde{q}_{s+1}$ are irreducible in $Q[Z]$ and $\tilde{Q}[Z]$, respectively, we know that

$$F_{s+1} = Q(z_{s+1}) \cong Q[Z]/(q_{s+1}(Z)) \cong \tilde{Q}[Z]/(\tilde{q}_{s+1}(Z)) \cong \tilde{Q}(\tilde{z}),$$

via the map $z_{s+1} \mapsto \tilde{z}$. Therefore f_{s+1} really is a field embedding of F_{s+1} into $\widetilde{F}$. Moreover, we checked that $z_{s+1} \mapsto \tilde{z}$ sends z_s to $f_s(z_s)$, and so $f_s \subseteq f_{s+1}$. It remains to see that f_{s+1} extends to a field embedding $\alpha : F \to \widetilde{F}$. To prove this fact, notice that with $\tilde{p}_{s+1,g(s+1)}(\tilde{z}, \tilde{y}) = 0$, we must have $p_{s+1,g(s+1)}(z, y) = 0$ as well, where $z = \rho^{-1}(\tilde{z})$ and $y = \rho^{-1}(\tilde{y})$. Since $g(s+1) \geq h(s+1)$, this fact means that $p_{s+1,h(s+1)}(z, Y)$ must have a root in F as well. But by Definition 3.2, we then have $\langle z_{s+1}, z\rangle \in B_F$, so there is some automorphism β of F with $\beta(z_{s+1}) = z$. But then

$$\rho(\beta(z_{s+1})) = \rho(z) = \tilde{z},$$

and so $(\rho \circ \beta)$ is an embedding of F into $\widetilde{F}$ that extends f_{s+1}, as required.

Therefore, the union $f = \bigcup_s f_s$ is a well-defined computable function with domain $\bigcup_s F_s = F$. Since every f_s is a field embedding, so is f. But with $\widetilde{F} \cong F$, a field embedding of F into $\widetilde{F}$ must in fact be an isomorphism, by Corollary 2.2. This conclusion proves Lemma 3.3. □

We can also prove the converse of Lemma 3.3, yielding the full equivalence.

Theorem 3.4. *For a computable algebraic field F, the following are equivalent.*

1. *F is relatively computably categorical.*
2. *F has a Σ^0_1 Scott family.*
3. *F has a Σ^0_1 Scott family using no parameters from F.*
4. *The orbit function h for F (with respect to some computable chain $Q = F_0 \subseteq F_1 \subseteq \cdots$) is computably bounded.*
5. *The orbit function h for F (with respect to every computable chain) is computably bounded.*

Proof. (5) $\implies$ (4) is immediate, and the equivalence of (1) and (2) was established (for all computable structures, not just fields) by Ash, Knight, Manasse, and Slaman in [2], and independently by Chisholm in [4]. Lemma 3.3 shows (4) $\implies$ (1), and its proof also explained how (4) yields (3), which in turn clearly implies

(2). So we now prove (2) $\implies$ (5). Fix any computable chain with union F, and fix a Σ^0_1 Scott family $\mathfrak{S}$ for F. By the Effective Theorem of the Primitive Element, we may assume that $\mathfrak{S}$ uses (at most) a single parameter a from F. Since a has finitely many conjugates over Q in the algebraic field F, we may assume that we know all elements $a = b_0, b_1, \ldots, b_m$ satisfying $\langle a, b\rangle \in B_F$. Also, for each $b \in F$ conjugate to a but with $\langle a, b\rangle \notin B_F$, Corollary 2.6 yields an s such that no embedding of F_s into F maps a to b. So we may fix some s_0 so large that all conjugates of a in F lie in F_{s_0}, but that $b_0, \ldots, b_m$ are the only possible images of a under embeddings of F_{s_0} into F.

We now compute the value $g(z_s)$ for a primitive generator z_s of F_s, for arbitrary $s > s_0$. For each $i \leq m$, we list out the formulas of $\mathfrak{S}$ and search through F for witnesses for these existential formulas, until we find a formula

$$\gamma_i(Z) = \exists \vec{x} \delta_i(Z, \vec{x}, a) \in \mathfrak{S},$$

an n_i, and a tuple $\vec{v}_i$ from F such that $\delta_i(z_s, \vec{v}_i, b_i)$ holds in F. It is important to notice that in our search, we have replaced the parameter a by b_i. (Of course $b_0 = a$, so one search involving a does still take place.) Now for each i, there is (at least one) automorphism ψ_i of F with $\psi_i(b_i) = a$. By the definition of Scott family, there is a formula $\gamma(Z) = \exists \vec{x} \delta(Z, \vec{x}, a)$ in $\mathfrak{S}$ for which $\gamma(\psi_i(z_s))$ holds; that is, $\delta(\psi_i(z_s), \vec{v}, a)$ holds for some $\vec{v}$. But then $\delta(z_s, \psi_i^{-1}(\vec{v}), b_i)$ also holds, since ψ_i is an automorphism. Therefore eventually our search halts and produces a formula $\gamma_i(Z)$, a tuple $\vec{v}_i$, and an n_i. Define $g(z_s)$ to be the least number $\geq s$ such that for all $i \leq m$, the entire finite tuple $\vec{v}_i$ is contained in the subfield $F_{g(z_s)}$. This definition completes our computation of g, on every $s > s_0$; now we must show that $g(z_s) \geq h(z_s)$.

So consider any conjugate z of z_s, such that $\langle z_s, z\rangle \notin B_F$. Now for each i, $\gamma_i(\psi_i(z_s))$ holds in F (with ψ_i as chosen above). But all $\psi_i(z_s)$ lie in the orbit of z_s, and all $\psi_i(z)$ lie in the orbit of z. Since these two orbits are distinct, the definition of Scott family shows that $\gamma_i(\psi_j(z))$ must be false in F for each $i, j \leq m$. So there is no tuple $\vec{x}$ from F for which any $\delta_i(\psi_j(z), \vec{x}, a)$ holds, and hence (by applying the automorphism ψ_j^{-1}) no tuple $\vec{x}$ for which any $\delta_i(z, \vec{x}, b_j)$ holds either.

Now if $p_{s,g(s)}(z, Y)$ had a root y in F (where $p_{s,g(s)}$ is as in the proof of Lemma 3.3), then there would be an isomorphism from $F_{g(z_s)}$ onto $F(y)$: start with a field embedding ψ of F_s into F that sends z_s to its F_0-conjugate z, and then extend by setting $\psi(z_{g(z_s)}) = y$ (which is still a field embedding of $F_{g(z_s)}$ into F, since $z_{g(z_s)}$ has minimal polynomial $p_{s,g(z_s)}(z_s, Y)$ over F_s and y has minimal polynomial $p_{s,g(z_s)}(\psi(z_s), Y)$ over $\psi(F_s)$). However, then the quantifier-free formula

$$\delta_0(z, \psi(\vec{v}_0), \psi(a))$$

must hold in F, since $\delta_0(z_s, \vec{v}_0, a)$ held in $F_{g(z_s)}$. By our choice of s_0, and since $s \geq s_0$, we have $\psi(a) = b_j$ for some j, contradicting our conclusion above that $\delta_i(z, \vec{x}, b_i)$ fails for all tuples $\vec{x}$ in F.

So, for all conjugates z of z_s with $\langle z_s, z\rangle \notin B_F$, the polynomial $p_{s,g(z_s)}(z, Y)$ has no root in F. By the minimality of $h(z_s)$ in Definition 3.2, this fact forces $g(z_s) \geq h(z_s)$. The argument above only defined g on inputs z_s with $s > s_0$, but it requires just finitely much more information to set $g(z_s) = h(z_s)$ for all $s \leq s_0$. Then one simply defines $g(x)$ on arbitrary $x \in F$ by $g(x) = g(z_s)$, where s is minimal such that $x \in F_s$. This definition works because, if $F_{g(x)}$ has an embedding α into F with $x' = \alpha(x)$, then by our construction, $\langle z_s, \alpha(z_s)\rangle \in B_F$. Hence the map $\alpha \restriction F_s$

extends to an automorphism of F, and since $x \in F_s$, that automorphism maps x to $\alpha(x) = x'$. So the computable function g does indeed bound h. □

As mentioned in the introduction, it was shown in [33] that if the algebraic field F has a splitting algorithm, then F is (relatively) computably categorical if and only if B_F is computable. We now show that the situation is different for a general algebraic field F. Using Scott families, we easily show that relative computable categoricity implies that B_F is Σ_1. The challenging part is to show that the converse is false, and here we actually strengthen that result by making B_F (and also the full orbit relation A_F) computable.

Theorem 3.5. *Let $\mathfrak{M}$ be any computable, relatively computably categorical structure. Then the full orbit relation of $\mathfrak{M}$ (i.e., the set of pairs of tuples $\langle \vec{a}, \vec{b} \rangle$ such that some automorphism of $\mathfrak{M}$ maps $\vec{a}$ to $\vec{b}$) is computably enumerable. However, there exists a computable algebraic field F that is not computably categorical, yet has A_F computable.*

Proof. Suppose that $\mathfrak{M}$ is a computable, relatively computably categorical structure, and denote its full orbit relation by $A_{\mathfrak{M}}$. Let $\mathfrak{S}$ be a Σ^0_1 Scott family for $\mathfrak{M}$. For any pair $\langle \vec{a}, \vec{b} \rangle$ of n-tuples from $\mathfrak{M}$, search through all formulas in $\mathfrak{S}$ with exactly n free variables, and all tuples of possible witness elements for each formula. Enumerate $\langle \vec{a}, \vec{b} \rangle$ into $A_{\mathfrak{M}}$ if ever we find a single formula in $\mathfrak{S}$ satisfied by both these tuples.

Of course, every $\vec{a}$ satisfies some formula in $\mathfrak{S}$, and if $\langle \vec{a}, \vec{b} \rangle \in A_{\mathfrak{M}}$, then $\vec{b}$ satisfies that same formula, so the pair is enumerated. Conversely, by the definition of Scott family, any two tuples satisfying the same formula must be images of each other under automorphisms of $\mathfrak{M}$.

We describe a simple version of the basic module used to construct the field F and its computable copy $\widetilde{F}$ which together prove the second statement. To ensure that a single φ_e is not an isomorphism from F onto $\widetilde{F}$, we use the cube roots of 2. Each field starts with one $\sqrt[3]{2}$, called θ_0 and $\tilde{\theta}_0$, respectively. In order to be an isomorphism, $\varphi_e(\theta_0)$ must converge to $\tilde{\theta}_0$. If this convergence ever happens, then we adjoin two more cube roots θ_1 and θ_2 of 2 to F, and likewise in $\widetilde{F}$. In F we also *tag* the original θ_0, by finding a polynomial $q \in \mathbb{Q}[X, Y]$ such that we can adjoin a root of $q(\theta_0, Y)$ to F without adjoining any roots of either $q(\theta_1, Y)$ or $q(\theta_2, Y)$; thus these two conjugates of θ_0 are not tagged. In $\widetilde{F}$ we adjoin a root of $q(\tilde{\theta}_1, Y)$. Thus the two fields remain isomorphic, but only via isomorphisms mapping θ_0 to $\tilde{\theta}_1$.

Moreover, no matter what the outcome of this basic module, every computable field $E \cong F$ has computable orbit relation B_E. Certainly every pair $\langle x, x \rangle \in B_E$, and if our program for computing B_E is ever given a pair $\langle \theta, \theta' \rangle$ of distinct cube roots of 2, then it searches for a third such cube root θ'' and also for a tag, i.e. a root of $q(\theta, Y)$ or of $q(\theta', Y)$ or of $q(\theta'', Y)$. These must exist, since the existence of the second cube root means that the basic module must have performed the action described above. Once our program finds the tag, it knows which two of these roots lie in the same orbit, and therefore can answer correctly whether $\langle \theta, \theta' \rangle \in B_E$.

In fact, such tagging requires an elaborate algebraic proof of the existence of the appropriate polynomials. (For instance, letting $q(X, Y)$ be $(Y^2 - X)$ would not work, because any square root of θ_0 would generate square roots of θ_1 and θ_2 as well.) Moreover, if any basic module adjoins to F all three cube roots of some

element (such as 2), then every x with a cube root F must have three cube roots in F. Thus this basic module does not extend readily to infinitely many requirements, so we resort to a similar strategy using roots of polynomials with symmetric Galois groups $\mathcal{S}_7$ of order 7.

Viewing such a Galois group as the symmetric permutation group on the seven roots of some polynomial, we see that its symmetry allows us to adjoin the sum of any subset S (with $1 < |S| < 6$) of the seven roots we like while keeping the individual roots out of the field. For instance, if x and y are two of these roots, then the subfield generated by $(x + y)$ is the fixed field of the subgroup of $\mathcal{S}_7$ containing those permutations fixing $\{x, y\}$ setwise, and as this subgroup fixes no individual root, the subfield contains none of the roots. Moreover, $\mathcal{S}_7$ is a sufficiently large group to allow us both to adjoin specific elements and to tag them. By analogy to the basic module above, think of $\theta_0 = x + y$ as the sum of two roots x and y, and $\theta_1 = u + v$ as the sum of two others u and v. We can tag either one later if necessary by adjoining x or u. If there were only four roots in total, then adjoining $x + y$ would have forced $u + v$ to enter the field as well. With only five, adjoining both $(x + y)$ and $(u + v)$ would force the fifth root to enter the field, since every permutation fixing both $\{x, y\}$ and $\{u, v\}$ would have to fix the fifth root. With only six, it would force the sum of the other two roots to enter the field, where it would be conjugate to θ_0 and θ_1, in the manner of θ_2 above; this would not ruin the argument, but it would complicate it, so we use seven roots instead.

Lemma 3.6. *There exists a computable infinite sequence $p_0(X), p_1(X), \ldots$ of polynomials of degree 7 in $\mathbb{Q}[X]$ such that for every $e \in \omega$, if K_e is the subfield of $\overline{\mathbb{Q}}$ generated by the splitting fields of all $p_i(X)$ with $i \neq e$, then the splitting field P_e of $p_e(X)$ over K_e is the symmetric group on the seven roots of $p_e(X)$.*

Proof. Apply Proposition 2.4 repeatedly, to get $\mathbb{Q} = E_0 \subset E_1 \subset E_2 \subset \cdots$ and polynomials $p_0, p_1, \ldots$, all of degree 7 in $\mathbb{Z}[X]$, such that each E_{e+1} is the compositum of E_e with the splitting field P_e of p_e over $\mathbb{Q}$ and such that $\mathrm{Gal}(P_e/\mathbb{Q}) \cong \mathcal{S}_7$ for all e. Notice that if P_e were not linearly disjoint from K_e (as defined in Definition 2.3), then for some j, the field P_e would fail to be linearly disjoint from the subfield $K_e \cap E_{j+1}$, which (for the least such j) would contradict the linear disjointness of P_j from E_j. □

We fix a sequence of polynomials $p_e(X)$ as described in Lemma 3.6. In the following construction, using a fixed computable copy $\overline{\mathbb{Q}}$ of the algebraic closure of $\mathbb{Q}$, we let x_e, y_e, u_e, and v_e be the four $<$-least roots of p_e in $\overline{\mathbb{Q}}$. Our fields F and $\widetilde{F}$ will both be computably enumerable subfields of $\overline{\mathbb{Q}}$, hence computably isomorphic (by taking pullbacks) to computable algebraic fields of characteristic 0. The requirements, which never injure one another, are

$$\mathcal{R}_e : \ \varphi_e \text{ is not an isomorphism from } F \text{ onto } \widetilde{F},$$

for every $e \in \omega$.

With no injury involved, we may define our fields quite simply. The c.e. subfield F contains precisely the following elements of $\overline{\mathbb{Q}}$:

- $(x_e + y_e)$, for all $e \in \omega$;
- $(u_e + v_e)$, for all $e \in \omega$ for which $\varphi_e(x_e + y_e)\!\downarrow = x_e + y_e$;
- x_e (hence also y_e), for all $e \in \omega$ for which $\varphi_e(x_e + y_e)\!\downarrow = x_e + y_e$;
- and all elements of $\overline{\mathbb{Q}}$ generated by these.

$\widetilde{F}$ contains precisely the following elements of $\overline{\mathbb{Q}}$:

- $(x_e + y_e)$, for all $e \in \omega$;
- u_e and v_e (hence also $(u_e + v_e)$), for all $e \in \omega$ with $\varphi_e(x_e + y_e)\!\downarrow = x_e + y_e$;
- and all elements of $\overline{\mathbb{Q}}$ generated by these.

We readily define an isomorphism ρ from F onto $\widetilde{F}$. If $\varphi_e(x_e+y_e)\!\downarrow = x_e+y_e$, then $\rho(x_e) = u_e$, $\rho(y_e) = v_e$, and $\rho(u_e + v_e) = x_e + y_e$. Otherwise $\rho(x_e + y_e) = x_e + y_e$. (In fact, there are 2^ω-many isomorphisms, since the first case occurs for infinitely many e, and when it does, $\rho(x_e) = v_e$ and $\rho(y_e) = u_e$ is also possible; moreover, by Lemma 3.6, the choices of $\rho(x_e)$ may be made independently for all the different e for which the first case occurs.)

Let $F_e \subseteq F$ be generated by $\{(x_e + y_e), (u_e + v_e), x_e\} \cap F$, and $\widetilde{F}_e \subseteq \widetilde{F}$ by $\{(x_e + y_e), u_e, v_e\} \cap \widetilde{F}$. Let $L_e = K_e \cap F$ be the subfield of F generated by all those generators of F with indices $\neq e$ (here K_e is as in Lemma 3.6), and similarly $\widetilde{L}_e = K_e \cap \widetilde{F}$. Now by Lemma 3.6, $K_e \cap F_e = K_e \cap \widetilde{F}_e = \mathbb{Q}$ for all e, so the splitting field of $p_e(X)$ over L_e has Galois group $\mathcal{S}_7$, the symmetric group on the seven roots of $p_e(X)$. (It has this Galois group over the larger ground field K_e from the lemma, and over $\mathbb{Q}$ itself, hence also over all intermediate fields, including L_e.) This fact also shows that every automorphism σ of F fixes each F_e setwise, since it must map F_e into the intersection of F with the splitting field P_e of $p_e(X)$ over $\mathbb{Q}$, and this intersection is just F_e itself. (One says that F_e is *normal within* F.)

Now $\mathrm{Gal}(P_e/\mathbb{Q})$ is the symmetric group $\mathcal{S}_7$ on the seven roots of $p_e(X)$. We know that $\mathrm{Gal}(P_e/\mathbb{Q}(x_e + y_e))$ contains exactly those permutations in $\mathcal{S}_7$ that either interchange x_e with y_e or fix both. The $\mathbb{Q}$-conjugates of $(x_e + y_e)$ in P_e are all sums of two distinct roots of p_e, and no other such sum can be fixed by all these permutations, so $(x_e + y_e)$ has no conjugates in $\mathbb{Q}(x_e + y_e)$. If $(u_e + v_e)$ ever enters F, then it is a conjugate of $(x_e + y_e)$, but then $x_e \in F$ as well, so $\mathrm{Gal}(P_e/F_e)$ then contains those permutations that fix the elements x_e and y_e and the set $\{u_e, v_e\}$. No other conjugate of $(x_e + y_e)$ is fixed by all those permutations, and since there is an element of $\mathrm{Gal}(P_e/F_e)$ interchanging u_e with v_e, neither u_e nor v_e (nor any of the remaining three roots of $p_e(X)$ in $\overline{\mathbb{Q}}$) lies in F_e.

So, if $\varphi_e(x_e + y_e)\!\downarrow = x_e + y_e$, then for every automorphism σ of F, we must have $\sigma(x_e) \in \{x_e, y_e\}$, and hence $\sigma(x_e + y_e) = x_e + y_e$, leaving $\sigma(u_e + v_e) = u_e + v_e$. Thus $\sigma \restriction F_e$ has only two possibilities: the identity, and the map $\sigma_e \in \mathrm{Aut}(F_e)$ interchanging x_e with y_e. Note that this σ_e fixes $(x_e + y_e)$, and also fixes $(u_e + v_e)$.

On the other hand, if $\varphi_e(x_e + y_e)$ diverges or converges to any value $\neq x_e + y_e$, then $(x_e + y_e)$ has no conjugates in F except itself, and generates F_e. So in this case $\sigma \restriction F_e$ must be the identity.

It follows that B_F is computable. First, of course, every pair $\langle z, z\rangle$ lies in B_F. By the discussion above, $\langle x_e + y_e, z\rangle \in B_F$ if and only if $z = x_e + y_e$, and likewise $\langle u_e + v_e, z\rangle \in B_F$ if and only if $z = u_e + v_e$. Next, for any root z in F of any p_e, we know that $\varphi_e(x_e + y_e)\!\downarrow = x_e + y_e$, since F could not contain such a root otherwise; hence $\langle z, z'\rangle \in B_F$ if and only if $z' \in \{x_e, y_e\}$. Thus our decision procedure for B_F can compute the orbit of every generator of F_e. So, for an arbitrary pair $\langle z, z'\rangle \in F^2$, it can express z in terms of the generators of finitely many F_e, compute the orbits of each of these generators, and use this information to determine whether $\langle z, z'\rangle \in B_F$. (We do use here the remark above that for any $e_0, \ldots, e_n$ with all $x_{e_i} \in F$, the choice of images of the x_{e_i} may be made independently for all $i \leq n$.)

Because F is a field, computability of B_F implies computability of A_F, as follows. Given any tuples $\vec{a}, \vec{b} \in F^n$, the Effective Theorem of the Primitive Element allows us to identify a single element $a \in F$ such that $\mathbb{Q}(a) = \mathbb{Q}(\vec{a})$. Having found some $q \in \mathbb{Q}[X_1, \ldots, X_n]$ with $a = q(\vec{a})$, we set $b = q(\vec{b})$. In turn, each $a_i = q_i(a)$ for some $q_i \in \mathbb{Q}[X]$. Now $\langle \vec{a}, \vec{b} \rangle \in A_F$ if and only if $\langle a, b \rangle \in B_F$ and, for each $i \leq n$, $b_i = q_i(b)$. For full details, we refer the reader to [33], where the Effective Theorem of the Primitive Element appears as Theorem 3.11.

However, for any e, we know that $\varphi_e(x_e + y_e) \downarrow = x_e + y_e$ if and only if no isomorphism $\rho : F \to \widetilde{F}$ fixes $(x_e + y_e)$. In particular, if $\varphi_e(x_e + y_e)$ diverges or converges to any other value, then $(x_e + y_e)$ has no $\mathbb{Q}$-conjugate in either F or $\widetilde{F}$; whereas if $\varphi_e(x_e + y_e) \downarrow = x_e + y_e$, then x_e and y_e lie in F, but the only two roots of p_e in $\widetilde{F}$ are u_e and v_e, whose sum is $\neq \varphi_e(x_e + y_e)$. In both these cases, therefore, φ_e is not an isomorphism from F onto $\widetilde{F}$. Since this fact holds for all e, the field F is not computably categorical. □

4. Computable Categoricity

Relative computable categoricity immediately implies computable categoricity, and for many theories T, the converse also holds of all computable models of T. This is the case for the theories of linear orders, Boolean algebras, and trees (in the language of partial orders). Indeed, in [10] Goncharov showed it to hold for all computable structures $\mathfrak{M}$ for which the set of Σ_2 sentences in the elementary diagram of $\mathfrak{M}$ is decidable. On the other hand, in [21], Kudinov proved that decidability of the Σ_1 fragment of the elementary diagram does not suffice, by producing examples of such 1-*decidable* structures that are computably categorical but not relatively computably categorical. Such structures are widely considered to be unusual, and among classes of structures for which characterizations of computable categoricity are known (linear orders, Boolean algebras, trees, etc.), computably categoricity always implies relative computable categoricity. Here we address this question for algebraic fields, proving that computable categoricity of an algebraic field can fail to relativize.

Theorem 4.1. *There exists a computable, computably categorical algebraic field F whose orbit relation B_F is not Σ^0_2. It follows from Theorem 3.5 that F is not relatively computably categorical.*

Proof. Every computable algebraic field has Π^0_2 orbit relation B_F, by Lemma 2.6. Therefore, it suffices to construct such a field F that is computably categorical, but for which B_F is not Δ^0_2.

Our construction of F takes place on a tree T, and satisfies two types of requirements. Every node β at level $2e$ of T is dedicated to satisfying requirement $\mathcal{C}_e$ for computable categoricity:

$\mathcal{C}_e$: If φ_e decides the atomic diagram of a field C_e isomorphic to F, then there is a computable isomorphism $C_e \to F$.

Such a β is also called a $\mathcal{C}_e$-node. It has two outcomes $\cong$ and $\not\cong$, and hence two immediate successors, which we order:

$$\beta^\frown\langle \cong \rangle \prec \beta^\frown\langle \not\cong \rangle.$$

For the β on the true path of T, the $\cong$ outcome will hold infinitely often if and only if C_e really is isomorphic to F, in which case we must be sure to satisfy $\mathcal{C}_e$. To guess at whether $C_e \cong F$, we use Corollary 2.7. Each time we see both that a new larger subfield of C_e embeds into F, and that a new larger subfield of F embeds into C_e, we make the outcome $\cong$ eligible. The *counter* c_β is used for this purpose: at stage s, it will be maximal such that the first $c_{\beta,s}$ elements of $C_{e,s}$ can be embedded together into F_s, and the first $c_{\beta,s}$ elements of F_s can be embedded together into $C_{e,s}$.

Every node α at level $2e+1$ of the tree is dedicated to satisfying the requirement $\mathcal{R}_e$ destroying limit-computability of B_F:

$$\mathcal{R}_e : \ \exists w \lim_t \varphi_e(w,t) \neq B_F(w).$$

An $\mathcal{R}_e$-node α has only one outcome. It acts so as to satisfy its requirement by observing the behavior of its function $\varphi_e(\langle x_\alpha, y_\alpha\rangle, t)$ for a particular pair $\langle x_\alpha, y_\alpha\rangle$ of elements of F, and ensures that this pair lies in B_F if and only if the limiting value of φ_e on the pair says otherwise.

The construction somewhat resembles that of Theorem 4.1 of [29], which built a computable algebraic field that is not $\mathbf{0}'$-categorical. We will use here the same principal tool given for that construction, and the same notation. In particular, when $h(X)$ is a polynomial with coefficients in $\mathbb{Q}[\sqrt{p}]$, we will write $h^-(X)$ to denote the image of this polynomial under the nontrivial automorphism of $\mathbb{Q}[\sqrt{p}]$, with $\sqrt{p}$ mapped to $-\sqrt{p}$.

Proposition 4.2 (Proposition 2.15 in Miller [29])**.** *For any fixed prime p, let F be the field $\mathbb{Q}[\sqrt{p}]$. Then for every odd prime number d, there exists a polynomial $h(X) \in F[X]$ of degree d with the following properties.*

- *h and h^- are both irreducible in the polynomial ring $F[X]$.*
- *The splitting field of h over F has Galois group isomorphic to $\mathcal{S}_d$, the symmetric group on the d roots of h, and the same holds for h^-. (Since $\mathcal{S}_d$ acts transitively on the roots, this property implies the preceding one.)*
- *The splitting field of $h(X)$ over the splitting field of $h^-(X)$ also has Galois group isomorphic to $\mathcal{S}_d$ (and vice versa). In particular, each of $h(X)$ and $h^-(X)$ is irreducible over the splitting field of the other.*

Moreover, uniformly in p, d, and any computable presentation of F, it is computable whether an arbitrary $h(X) \in F[X]$ satisfies these properties.

For an $\mathcal{R}_e$-node α, x_α and y_α will be the two square roots of a specific prime number p_α. At each stage, α will use two polynomials $h_{b_\alpha-1,\alpha}$ and $h_{b_\alpha,\alpha}$, provided by the Proposition for this p_α, of distinct prime degrees, with b_α keeping count of these polynomials. Both $h_{b_\alpha-1,\alpha}$ and $h_{b_\alpha,\alpha}$ will have roots in F, but neither $h^-_{b_\alpha-1,\alpha}$ and $h^-_{b_\alpha,\alpha}$ will have a root. Thus, x_α is tagged in two different ways to distinguish it from y_α. Whenever $\varphi_e(\langle x_\alpha, y_\alpha\rangle, t)$ equals 0 for a new larger t, α takes a step toward making $\langle x_\alpha, y_\alpha\rangle$ lie in B_F: it adjoins a root of $h^-_{b_\alpha-1,\alpha}$ to F, so that y_α now has this tag, just like x_α. At the same time, though, α chooses a new $h_{b_\alpha+1,\alpha}$, giving x_α a new tag which y_α lacks, and increments b_α so as to keep track of the current tags. Therefore, the only way $\langle x_\alpha, y_\alpha\rangle$ can end up in B_F is if this step is repeated infinitely often, in which case $\lim_t \varphi_e(\langle x_\alpha, y_\alpha\rangle, t) \neq 1$. If the limit is 1, then this step is repeated only finitely often, and x_α is tagged in some way in which y_α never is, so that $\langle x_\alpha, y_\alpha\rangle \notin B_F$. The tagging of x_α by two separate polynomials ensures

that higher-priority $\mathcal{C}$-nodes can always distinguish x_α from y_α, so that they can build their computable isomorphisms, and that once they have built them, their guesses in their fields C_e will remain correct about which node corresponds to x_α and which to y_α: at least one tag will always be present, keeping their computation correct, even while the other tag is removed and redefined.

We start with z_0 as the multiplicative identity, so that F_0 is a copy of $\mathbb{Q}$. All nodes are initialized at stage 0, which means that all counters $c_{\beta,0}$ for $\mathcal{C}$-nodes β and all primes p_α, counters b_α, and potential field elements x_α and y_α for $\mathcal{R}$-nodes α are undefined.

To arrange our stages, we fix the bijection $\omega \times \omega \to \omega$ according to the listing of $\omega \times \omega$ as follows:

$$\langle 0,0\rangle, \langle 0,1\rangle, \langle 1,1\rangle, \langle 0,2\rangle, \langle 1,2\rangle, \langle 2,2\rangle, \langle 0,3\rangle, \ldots.$$

At all stages $s+1$ with $s = \langle 0, n\rangle$, the root node is eligible. At all stages $s+1$ with $s = \langle m, n\rangle$ and $m \neq 0$, the node eligible at the preceding stage will have designated one of its immediate successors (of length m) to be eligible at this stage. Each eligible $\mathcal{C}$-node chooses a successor to be eligible, while for an eligible $\mathcal{R}$-node, its unique successor is eligible.

At each stage $s+1$ where a $\mathcal{C}_e$-node β is eligible, let s' be the most recent stage (if any) at which β was eligible. If there has been no such stage s' yet, or if β has been initialized since that stage, then we set $c_{\beta,s+1} = 1$ and make $\beta\hat{}\langle \not\cong\rangle$ eligible at stage $s+1$. Otherwise, we consider the function $\varphi_{e,s}$ as the (partial) characteristic function for the atomic diagram of a structure C_e in the language of fields. If $\varphi_{e,s}$ enumerates any facts inconsistent with the field axioms, or if the initial segment of β on which $\varphi_{e,s}$ converges is the same as that on which $\varphi_{e,s'}$ converged, then $\beta\hat{}\langle \not\cong\rangle$ is eligible at substage $t+1$. Otherwise, we consider the fragment $C_{e,s}$ of a field as described by $\varphi_{e,s}$. Lemma 2.10 ensures that we can check whether both of the following hold (as defined in Definition 2.9).

- If D is the finite subset $\{0, 1, \ldots, c_{\beta,s'}\}$ of the field fragment $\{0, 1, \ldots, s\}$ of F_s, then D embeds into $C_{e,s}$.
- The following finite subset D' of the field fragment $C_{e,s}$ has an embedding into the field fragment $\{0, \ldots, s\}$ of F_s. The set D' contains the elements $0, 1, \ldots, c_{\beta,s'}$, and also contains all elements $x \in C_{e,s}$ for which there exists an α with $\beta\hat{}\langle\cong\rangle \subseteq \alpha$ and $i \in \omega$ and a currently defined polynomial of the form $(X^2 - p_{\alpha,s+1})$ or one of the forms $h_{i,\alpha,s}$ or $h^-_{i,\alpha,s}$, such that when the coefficients of this polynomial are mapped into $C_{e,s}$, x is a root in $C_{e,s}$ of the resulting polynomial there. (Notice that $(X^2 - p_{\alpha,s+1})$ has coefficients in $\mathbb{Q}$, hence has a unique image in $C_{e,s}$, or no image at all if the necessary elements of $\mathbb{Q}$ have not appeared in C_e by stage s. The other two polynomials, $h_{i,\alpha,s}$ and $h^-_{i,\alpha,s}$, have coefficients in $\mathbb{Q}(p_{\alpha,s+1})$, but interchanging the two square roots of $p_{\alpha,s+1}$ only interchanges these two polynomials with each other. So the description above names at most three polynomials in $C_{e,s}[X]$ for each α, without ambiguity.)

If so, then we set $c_{\beta,s+1} = c_{\beta,s'} + 1$, initialize all nodes to the right of $\beta\hat{}\langle\cong\rangle$ (including $\beta\hat{}\langle\not\cong\rangle$ and all its successors), and make $\beta\hat{}\langle\cong\rangle$ eligible at stage $s+2$. If either of these fragments fails to embed, then we leave $c_{\beta,s+1} = c_{\beta,s'}$ and make $\beta\hat{}\langle\not\cong\rangle$ eligible.

(The point of this stage is that, if β is eligible at infinitely many stages, then $\beta^\frown\langle\cong\rangle$ will be eligible at infinitely many stages if and only if φ_e is total and is the characteristic function of the atomic diagram of a field C_e such that every finitely generated subfield of C_e embeds into F and vice versa. By Lemma 2.6, this latter condition is equivalent to $C_e \cong F$.)

At each stage $s+1$ at which an $\mathcal{R}_e$-node α is eligible, we let s' be the most recent stage at which α was eligible and executed either Step 1 or Step 3 (below). In Step 1, α chooses its prime p_α, adjoins its square roots x_α and y_α, and sets up the first two tags on x_α. In Step 2, it waits for all $\mathcal{C}$-nodes β with $\beta^\frown\langle\cong\rangle \subseteq \alpha$ to complete their computations on the current tags, since α cannot add any new tags until this has been done. Finally, in Step 3, α gets to check whether $\lim_t \varphi_e(\langle x_\alpha, y_\alpha\rangle, t)$ has taken any further steps towards equalling 0; if so, then it executes the appropriate action with its tags, while if not, it does nothing. Here are the full descriptions of the steps.

(1) If this is the first stage at which α has been eligible (so the stage s' does not exist), or if α has been initialized since stage s', then we let $p_{\alpha,s+1}$ be the least odd prime number not yet chosen as $p_{\gamma,t}$ at any previous stage t for any node γ. We adjoin to F_s a new element $x_{\alpha,s+1}$ satisfying $x_{\alpha,s+1}^2 = p_{\alpha,s+1}$. Of course, this action also adjoins a second square root of $p_{\alpha,s+1}$, and we name this second element $y_{\alpha,s+1}$. Then we choose the least two odd prime numbers not yet used in the construction, and, for each of these two prime numbers d, search for a polynomial as described in Proposition 4.2 of this degree d for the prime $p_{\alpha,s+1}$. Let $h_{0,\alpha,s+1}(X)$ and $h_{1,\alpha,s+1}(X)$ be the two polynomials we find, and set $b_{\alpha,s+1} = 1$. We adjoin to $F_s(x_{\alpha,s+1})$ one root of $h_{0,\alpha,s+1}$ and one root of $h_{1,\alpha,s+1}$, to form F_{s+1}. This action completes this stage.

(2) Otherwise, $p_{\alpha,s'}$ and $b_{\alpha,s'}$ are already defined, as are $x_{\alpha,s'}$ and $y_{\alpha,s'}$ and polynomials $h_{0,\alpha,s'}(X), \ldots, h_{b_{\alpha,s'},\alpha,s'}(X)$. We check whether, for every $i \leq e$, either the $\mathcal{C}_i$-node $\beta \subset \alpha$ has $\beta^\frown\langle\ncong\rangle \subseteq \alpha$ or else the field fragment $C_{i,s}$ contains a root of the minimal polynomial of the least primitive generator of $F_{s'}$ over $\mathbb{Q}$. If this is not the case, then we do nothing at this stage. If it is the case, then we execute Step 3 below.

(3) Let $u \geq 0$ be maximal with the property that all $\varphi_{e,s}(\langle x_{\alpha,s'}, y_{\alpha,s'}\rangle, t)$ with $t < u$ converge, and let u' be the corresponding maximum with s' in place of s. If none of $\varphi_{e,s}(\langle x_{\alpha,s'}, y_{\alpha,s'}\rangle, u'), \ldots, \varphi_{e,s}(\langle x_{\alpha,s'}, y_{\alpha,s'}\rangle, u-1)$ equals 0, or if $u' = u$, then we do nothing. Otherwise we set $b_{\alpha,s+1} = b_{\alpha,s'} + 1$, choose the least odd prime d not yet seen in the construction, and find a polynomial $h_{b_{\alpha,s+1},\alpha,s+1}(X)$ satisfying Proposition 4.2 for this d and for $p_{\alpha,s'}$. We adjoin to F_s one root of this $h_{b_{\alpha,s+1},\alpha,s+1}$, and also one root of $h^-_{b_{\alpha,s+1}-2,\alpha,s'}$. (In Lemma 4.3 below, we show that each of these polynomials is irreducible over F_s, and indeed over the root of the other, so that this suffices to define F_{s+1}.) Of course, $b_{\alpha,s+1} - 2 = b_{\alpha,s} - 1$, so this has added a tag for $y_{\alpha,s+1}$ for that old h-polynomial, as well as a tag for x_α for the new h-polynomial $h_{b_{\alpha,s+1},\alpha,s+1}$.

Whichever step was executed, we then end this stage, with the unique successor of α eligible at the next stage.

Since this process was effective and every F_{s+1} was an algebraic extension (proper or not) of the preceding F_s, we have constructed a computable algebraic field $F =$

$\bigcup_s F_s$ of characteristic 0. We claim that every requirement $\mathcal{C}_e$ and $\mathcal{R}_e$ is satisfied by our construction, so that F is the field needed to establish the theorem.

The $\mathcal{C}$-nodes β always make one of their two successors eligible, and $\mathcal{R}$-nodes α always make their unique successor eligible. Therefore, the set containing the leftmost node at each level that is eligible infinitely often forms a path through the tree, called the *true path* P. Each requirement corresponds to a unique node on P, which will be the node causing that requirement to be satisfied. If $\beta \in P$ is a $\mathcal{C}$-node, then $\beta^\frown\langle \not\cong \rangle \in P$ if and only if $\lim_s c_{\beta,s}$ exists and is finite; if $\beta^\frown\langle \cong \rangle \in P$, then $c_{\beta,s} \to \infty$.

We start with an analysis of the strategy of the $\mathcal{R}_e$-node $\alpha \in P$, for any fixed e, starting with the last stage at which this α is initialized. Let s_0 be the first subsequent stage at which it is eligible. Then whenever α is eligible after s_0, it has two particular h_α polynomials that have roots, but such that the corresponding h^-_α polynomials do not have roots. We say that $x_\alpha = x_{\alpha,s_0}$ is *tagged* by these polynomials, while y_{α,s_0} is not (yet) tagged by them. Suppose that α has the correct guess about which $\mathcal{C}$-nodes preceding it correspond to fields isomorphic to F, and that $\varphi_e(\langle x_\alpha, y_\alpha\rangle, t)\downarrow$ for all t (since otherwise $\mathcal{R}_e$ is trivially satisfied). If there are infinitely many t for which $\varphi_e(\langle x_\alpha, y_\alpha\rangle, t) = 0$, then the limit of φ_e on $\langle x_\alpha, y_\alpha\rangle$ can only equal 0 or fail to exist, yet both x_α and y_α do end up each tagged by all the polynomials $h_{n,\alpha}$, and thus will lie in the same orbit. Conversely, if there are only finitely many such t, then $\lim_t \varphi_e(\langle x_\alpha, y_\alpha\rangle, t) \neq 0$ (and the limit may not exist at all), but in this case only finitely many polynomials $h_{n,\alpha,s}$ were ever defined, and the last two still tag x_α without tagging y_α. So in this case $\langle x_\alpha, y_\alpha\rangle \notin B_F$. Thus in both cases, $\mathcal{R}_e$ will be satisfied. Of course, to complete this argument, we must show that the tags really do work the way we claimed here, and in particular that no extraneous tags were introduced by the actions of other nodes.

So consider the elements adjoined to F_s at a specific stage $s+1$, with an $\mathcal{R}$-node α eligible at this stage. If α is in Step 1, then it first adjoins a square root $x_{\alpha,s+1}$ of its prime $p_{\alpha,s+1}$. This extension has degree 2, since we chose a $p_{\alpha,s+1}$ that does not already have a square root in F. Then the stage adjoins roots of $h_{0,\alpha,s+1}$ and $h_{1,\alpha,s+1}$, which were chosen to be irreducible over $\mathbb{Q}(x_\alpha)$ with degrees that are both new prime numbers, and indeed are irreducible over F_s as well, by Lemma 4.3 below. Thus the root of $h_{i,\alpha,s+1}$ generates an extension of degree $\deg(h_{i,\alpha,s+1})$, for each i, and since these degrees are prime to each other, they generate linearly disjoint field extensions of $F_s(x_{\alpha,s+1})$ (that is, field extensions whose intersection equals just $F_s(x_{\alpha,s+1})$). So the degree $[F_{s+1} : F_s]$ is the product of these two primes and 2.

Now suppose α is in Step 3 at stage $s+1$, and adjoins to F_s one root of $h_{b_{\alpha,s+1},\alpha,s+1}$ and one root r^- of $h^-_{\alpha,b_{\alpha,s+1}-2,s+1}$. The former is irreducible over F_s and has a new large prime as its degree, and its root thus generates an extension of that degree. We also claim that the root r^- of $h^-_{\alpha,b_{\alpha,s+1}-2,s+1}$ generates a further extension of degree $\deg(h^-_{\alpha,b_{\alpha,s+1}-2,s+1})$. The following lemma justifies these claims.

Lemma 4.3. *The following holds for every stage s. First, for the α and i (if any) such that $h_{i,\alpha,s+1}(X)$ is first defined at stage $s+1$, the polynomial $h_{i,\alpha,s+1}(X)$ is irreducible in $F_s(x_{\alpha,s+1})[X]$. Second, for any α and i such that $h_{i,\alpha,s}(X)$ has a root in F_s but $h^-_{i,\alpha,s}(X)$ does not, $h^-_{i,\alpha,s}(X)$ is irreducible in $F_s[X]$.*

Indeed, at a stage at which α enumerates two roots (apart from $x_{\alpha,s+1}$ itself) into F_{s+1}, using Step 1 or Step 3, the minimal polynomial of each of these roots (either $h_{i,\alpha,s+1}$ or $h^-_{i,\alpha,s+1}$) remains irreducible over the extension of F_s by the other root.

Proof. We prove the first two statements simultaneously, by induction on s. First suppose that $h_{i,\alpha,s+1}(X)$ is defined at stage $s+1$, by α in Step 1 or Step 3, with a root r adjoined to F_s. By Proposition 4.2, $[\mathbb{Q}(x_{\alpha,s+1}, r) : \mathbb{Q}(x_{\alpha,s+1})] = d$, the degree of $h_{i,\alpha,s+1}(X)$, and therefore d divides $[F_s(x_{\alpha,s+1}, r) : \mathbb{Q}(x_{\alpha,s+1})]$. However, the prime degree d was never used for any h-polynomials except $h_{i,\alpha,s+1}$. Now we use our inductive hypothesis on previous stages, noting that since the elements adjoined by $\mathcal{R}$-nodes γ at previous stages t were roots of irreducible polynomials $h_{j,\gamma,t}$ or $h^-_{j,\gamma,t}$ (or square roots of primes), those adjoinments created extensions of prime degrees distinct from d. Therefore, d must divide $[F_s(x_{\alpha,s+1}, r) : F_s]$, and since $x_{\alpha,s+1}$ has degree either 1 or 2 over F_s, we have that d divides $[F_s(x_{\alpha,s+1}, r) : F_s(x_{\alpha,s+1})]$. On the other hand, r is a root of $h_{i,\alpha,s+1}(X)$, which has degree d, and so $d = [F_s(x_{\alpha,s+1}, r) : F_s(x_{\alpha,s+1})]$, forcing $h_{i,\alpha,s+1}(X)$ to be the minimal polynomial of r over $F_s(x_{\alpha,s+1})$. Therefore $h_{i,\alpha,s+1}(X)$ is irreducible in $F_s(x_{\alpha,s+1})[X]$.

Next, suppose that $h_{i,\alpha,s+1}$ has a root in F_{s+1} but $h^-_{i,\alpha,s+1}$ does not. Let $d = \deg(h^-_{i,\alpha,s+1})$, and fix a root $r \in F_{s+1}$ of $h_{i,\alpha,s+1}$. Set $E = F_{s+1}(r^-)$, where r^- is a root of $h^-_{i,\alpha,s+1}$. (If $h^-_{i,\alpha,s+1}$ is reducible, then r^- may be a root of any of its irreducible factors in $F_{s+1}[X]$, and the argument below will apply.) E thus contains roots of both $h_{i,\alpha,s+1}$ and $h^-_{i,\alpha,s+1}$, and by Proposition 4.2, d^2 must divide $[E : \mathbb{Q}(x_{\alpha,s+1})]$, since d divides both $[\mathbb{Q}(r, x_{\alpha,s+1}) : \mathbb{Q}(x_{\alpha,s+1})]$ and $[\mathbb{Q}(r^-, r, x_{\alpha,s+1}) : \mathbb{Q}(r, x_{\alpha,s+1})]$. However, the prime degree d was never used for any h-polynomials except $h_{i,\alpha,s+1}$ and $h^-_{i,\alpha,s+1}$. Using the inductive hypothesis once again, we see that among all $[F_{t+1} : F_t]$ with $t < s$, the only one divisible by d is the one for the stage $t+1$ with $r \in F_{t+1} - F_t$; moreover, for this t, we have that $[F_{t+1} : F_t]$ is divisible by d but not by d^2. It follows that d must divide $[E : F_s]$, and therefore the minimal polynomial of r^- over F_s has degree divisible by d. But r^- is a root of $h^-_{i,\alpha,s+1}$, which itself has degree d, and so this is the minimal polynomial of r^- over F_s. Thus $h^-_{i,\alpha,s+1}$ is irreducible over F_s. This conclusion completes the induction.

Finally, considering the two roots enumerated into F_{s+1} by α, we note that their minimal polynomials over F_s have distinct prime degrees. Therefore, the field extensions generated by each are linearly disjoint: their intersection is F_s. It follows that neither extension can cause the other minimal polynomial to factor (see e.g. [28, Lemma 2.12]; or just extend the argument from the induction above). □

Corollary 4.4. *Fix any i, α, and t. Then F has a root of $h^-_{i,\alpha,t}(X)$ if and only if the node α itself adjoins such a root by entering Step 3 at some stage $s+1$ with $b_{\alpha,s+1} = i+2$ and with $h_{i,\alpha,t} = h_{i,\alpha,s+1}$.*

Proof. By Lemma 4.3, at any stage $s+1$, nodes $\gamma \neq \alpha$ enumerate only roots of polynomials that are irreducible over F_s and have prime degrees distinct from the degree of $h_{i,\alpha,t}$. The same holds for the node α itself at stages $s+1$ such that α is initialized between that stage and stage t, or such that α adjoins roots of polynomials $h_{j,\alpha,s+1}$ or $h^-_{j,\alpha,s+1}$ with $j \neq i$. Finally, when α adjoins a root of $h_{i,\alpha,t}$, Proposition 4.2 shows that no root of $h^-_{i,\alpha,t}$ can result. □

To see that requirement $\mathcal{R}_e$ is satisfied, let α be the $\mathcal{R}_e$-node on the true path P, i.e. the leftmost node at level $2e+1$ that is eligible at infinitely many stages.

For simplicity, write x and y for $\lim_s x_{\alpha,s}$ and $\lim_s y_{\alpha,s}$, write h_i for $\lim_s h_{i,\alpha,s}$, etc. Now for any $\mathcal{C}_e$-node $\beta \subseteq \alpha$, either $\beta^\frown\langle \not\cong \rangle \subseteq \alpha$, or else $\beta^\frown\langle \cong \rangle$ was eligible in between every pair of stages at which α was eligible. In the latter case, $\beta^\frown\langle \cong \rangle$ was eligible infinitely often, and so the field C_e must be isomorphic to F. Therefore, α cannot simply execute Step 2 at cofinitely many stages; it executes Step 1 at the first stage at which it is eligible, and enters Step 3 at infinitely many stages after that.

Suppose first that $\lim_t \varphi_e(\langle x, y\rangle, t)$ exists and equals 1. Then at all but finitely many of the stages at which α goes through Step 3, it does nothing, and so we have a finite limit $b = \lim_s b_{\alpha,s}$. Set $d = \deg(h_b)$. Now $h_b(X)$ has a root r in F, while $h_b^-(X)$ does not, by Lemma 4.3. Therefore $\langle x, y\rangle \notin B_F$, since any automorphism of F mapping x to y would have to map r to some root of h_b^-. So in this case $\mathcal{R}_e$ is satisfied.

On the other hand, if $\lim_t \varphi_e(\langle x, y\rangle, t) = 0$, then α executed Step 3 infinitely many times, and so $b_{\alpha,s} \to \infty$ as $s \to \infty$. Thus, for every b, both h_b and h_b^- have a root in F. We claim that every subfield F_s containing x has an embedding into F mapping x to y. Once this claim is established, Lemma 2.6 will show that there exists a field embedding of F into itself mapping x to y, which will prove that $\langle x, y\rangle \in B_F$, since by Lemma 2.1 this embedding must be an automorphism. The key to proving this claim is the following standard fact from field theory.

Lemma 4.5. *Let $K \subseteq L$ be a field extension generated by a single $x \in L$ that is algebraic over K, and let $f : K \to E$ be a field embedding. Fix the minimal polynomial $h \in K[X]$ of x over K, and let h^f be its image in $E[X]$ under the map f on its coefficients. Then f extends to an embedding of L into E if and only if E contains a root of $h^f(X)$.*

To prove the claim that every subfield F_s containing x has an embedding into F mapping x to y, we show how to extend such embeddings from F_s to F_{s+1}. Assume by induction that f is an embedding of F_s into F with $f(x) = y$ (and hence $f(y) = x$), with $f(x_{\gamma,t}) = x_{\gamma,t}$ for all $\gamma \neq \alpha$ (and also with $\gamma = \alpha$ for stages t before the last initialization of α), with $f(r) = r$ for every generator r adjoined by any $\mathcal{R}$-node $\gamma \neq \alpha$ or by α before its last initialization, and such that, for every root $r \in F_s$ of any $h_{i,\alpha,s}(X)$, we have that $f(r)$ is a root of $h^-_{i,\alpha,s}(X)$, and for every root $r^- \in F_s$ of any $h^-_{i,\alpha,s}(X)$, we have that $f(r^-)$ is a root of $h_{i,\alpha,s}(X)$. (Of course, this assumption applies only to elements of F_s, which might not include x or others of the above.) We may assume that some $\mathcal{R}$-node γ is eligible at stage $s+1$, since otherwise $F_{s+1} = F_s$.

Consider first the case where either $\gamma \neq \alpha$ or α is initialized after stage $s+1$. Now γ may adjoin $x_{\gamma,s+1}$ and a root of each of $h_{0,\gamma,s+1}$ and $h_{1,\gamma,s+1}$ to F_s. If so, then $x^2_{\gamma,s+1} = p_{\gamma,s+1}$, and each of these roots satisfies an irreducible polynomial over $\mathbb{Q}(x_{\gamma,s+1})$. So we set $f(x_{\gamma,s+1}) = x_{\gamma,s+1}$, which extends f to $F_s(x_{\gamma,s+1})$, by Lemma 4.5, and also set f to be the identity on the roots of these h-polynomials (which works for the same reason). Alternatively, γ may have been in Step 3 and have adjoined a root of $h_{b_{\gamma,s+1},\gamma,s+1}(X)$ and a root of $h^-_{b_{\gamma,s+1}-2,\gamma,s+1}(X)$. By Lemma 4.3, both are irreducible over $F_s(x_{\gamma,s+1})$. But by assumption, f restricts to the identity on the coefficients of both (which all lie in $\mathbb{Q}(x_{\gamma,s+1})$), and so again we can extend f to these roots just by taking the identity map on them.

Now consider the case where $\gamma = \alpha$ and α is never initialized after stage $s+1$. First, if α adjoins $x_{\alpha,s+1}$ to F_s, then its negative $y_{\alpha,s+1}$ also appears in F_{s+1} and is

conjugate to $x_{\alpha,s+1}$ over F_s, so Lemma 4.5 allows us to define $f(x_{\alpha,s+1}) = y_{\alpha,s+1}$. In either Step 1 or Step 3, whenever α defines a new polynomial $h_{i,\alpha,s+1}(X)$ and adjoins a root r of it, we know (by our assumption that $b_{\alpha,s} \to \infty$) that eventually α will also adjoin a root r^- of $h^-_{i,\alpha,s+1}(X)$ to F, and so we define $f(r)$ to be that r^-. By Lemma 4.5, this definition does extend f to a field embedding on $F_s(r)$. Likewise, if α is in Step 3 and adjoins some root r^- of $h^-_{b_{\alpha,s+1}-2,\alpha,s+1}$ to F_s, we know by Lemma 4.3 that $h^-_{b_{\alpha,s+1}-2,\alpha,s+1}$ is the minimal polynomial of this r^- over F_s, and so it is safe to set $f(r^-)$ to equal r, since r is a root of the image $h_{b_{\alpha,s+1}-2,\alpha,s+1}$ of $h^-_{b_{\alpha,s+1}-2,\alpha,s+1}$ under the map f on its coefficients. Thus in all cases we have extended f from F_s to a field embedding of F_{s+1} into F, with $f(x) = y$ whenever $x \in F_s$. This fact proves the claim, and completes our argument that requirement $\mathcal{R}_e$ is satisfied.

Turning to the $\mathcal{C}$-requirements, for any e, we let β be the node of length $2e$ on P, i.e. the leftmost node of that length that is eligible at infinitely many stages. Of course, β works for the requirement $\mathcal{C}_e$. Suppose first that $\beta^\frown\langle\cong\rangle$ is never eligible after some stage s_0. Then F contains only finitely many elements enumerated by nodes $\alpha \supseteq \beta^\frown\langle\cong\rangle$, and moreover $\lim_s c_{\beta,s} = c_{\beta,s_0}$ is finite. But the elements $\{0, 1, \ldots, c_{\beta,s_0}\}$ and those enumerated by these α together generate a subfield of F that must not embed into C_e, or else the subfield of C_e generated by $\{0, \ldots, c_{\beta,s_0}\}$ does not embed into F, since otherwise $\beta^\frown\langle\cong\rangle$ would have become eligible again. So in this case $F \ncong C_e$, satisfying $\mathcal{C}_e$.

Therefore we may assume that $C_e \cong F$ and that $\beta^\frown\langle\cong\rangle \in P$. We construct a computable isomorphism from F onto C_e as follows. First, let s_0 be the last stage at which any node to the left of β is eligible. We may start by assuming that we know the restriction $g_0 = f \restriction F_{s_0}$ of the given isomorphism $f : F \to C_e$, since this knowledge requires only finitely much information, namely the images of the finitely many generators of F_{s_0}.

Now let $s_1 < s_2 < \cdots$ be all stages $> s_0$ at which $\beta^\frown\langle\cong\rangle$ is eligible. (We can compute this sequence, of course.) We extend each g_n to the finite field extension $F_{s_{n+1}}$ of F_{s_n} in turn, as follows. If an element x was adjoined to F by a node α to the right of $\beta^\frown\langle\cong\rangle$, then α is initialized at stage s_{n+1}, so we simply check how many elements that α enumerated into F before stage s_{n+1}. In particular, for any s with $s_n < s < s_{n+1}$, let $t < s_{n+1}$ be the greatest stage before the next initialization of α. If α enumerated a root of some polynomial $h_{i,\alpha,s}$, we check whether it also enumerated a root of $h^-_{i,\alpha,s}$ by stage t or not. This will be the case for finitely many i, but eventually we will reach an i for which F_t contains a root of $h_{i,\alpha,t}$ but no root of $h^-_{i,\alpha,t}$. (Indeed, the same holds for $i+1$ as well, since α always keeps two tags on x_α which y_α does not yet have.) Fixing this i, we find both square roots of $p_{\alpha,t}$ in C_e, and find a root of $h_{i,\alpha,t}$ over one of those square roots; we then map $x_{\alpha,s}$ to the conjugate with this root (and the root of $h_{i,\alpha,s}$ to the root itself, and likewise for $i+1$), and map $y_{\alpha,s}$ to the other conjugate. This mapping also then determines, for each j such that the polynomials $h_{j,\alpha,t}$ and $h^-_{j,\alpha,t}$ both have roots in $F_{s_{n+1}}$, where these roots should be mapped.

By our choice of s_0, the only other nodes that can enumerate any element into F between stages s_n and s_{n+1} are nodes α with $\beta^\frown\langle\cong\rangle \subseteq \alpha$. So next we suppose that such an α enters Step 1 at stage $s+1$, with $s_n \leq s < s_{n+1}$, and adjoins $x_{\alpha,s+1}$ and roots r and r' of the two polynomials $h_{0,\alpha,s+1}$ and $h_{1,\alpha,s+1}$. Recall that

these polynomials both have coefficients in the field $\mathbb{Q}(\sqrt{p_{\alpha,s+1}})$. We wait for both square roots of $p_{\alpha,s+1}$ to appear in C_e, which must happen eventually, since by assumption $F \cong C_e$. Once they have appeared, each one gives rise to an image in $C_e[X]$ of the polynomial $h_{0,\alpha,s+1} \in F_{s+1}[X]$, since either square root of $p_{\alpha,s+1}$ can be used as the square root in $h_{0,\alpha,s+1}$. As soon as either of these two polynomials in $C_e[X]$ acquires a root in C_e, we define $g_{n+1}(r)$ to equal that root, and define $g_{n+1}(x_{\alpha,s+1})$ to equal the square root of $p_{\alpha,s+1}$ that gave rise to the polynomial that has this root. We also consider the polynomial $h_{1,\alpha,s+1}(X)$ in $C_e[X]$ defined using $g_{n+1}(x_{\alpha,s+1})$, and wait for this polynomial to acquire a root in C_e, which then becomes the value of $g_{n+1}(r')$. All of these events must eventually happen, since $F \cong C_e$. It remains to show that this definition of g_{n+1} actually does extend to an embedding of F into C_e.

Notice first that at every subsequent stage t, the polynomial $h_{b_{\alpha,t}-1,\alpha,t}$ has a root r in F_t and $h_{b_{\alpha,t},\alpha,t}$ has a root r' there, but neither $h^-_{b_{\alpha,t}-1,\alpha,t}$ nor $h^-_{b_{\alpha,t},\alpha,t}$ has any root in F_t. This $b_{\alpha,t}$ stays fixed from one stage to the next (starting with $b_{\alpha,s+1} = 1$), except for stages at which α enters Step 3. At such stages, C_e must contain images $g_{n+1}(r)$ and $g_{n+1}(r')$, since we defined g_{n+1} on r and r' as soon as we found those roots in C_e. Also, no roots of the g_{n+1}-images of $h^-_{b_{\alpha,t}-1,\alpha,s+1}$ and $h^-_{b_{\alpha,t},\alpha,s+1}$ have appeared yet, because by the construction for the node β, such roots would prevent $\beta\hat{}\langle\cong\rangle$ from becoming eligible (and so F would never have enumerated roots of $h^-_{b_{\alpha,t}-1,\alpha,s+1}$ and $h^-_{b_{\alpha,t},\alpha,s+1}$, and thus F would not have been isomorphic to C). In Step 3, α adjoins to F_s a root of $h^-_{b_{\alpha,t}-1,\alpha,t}$, but $h^-_{b_{\alpha,t},\alpha,t}$ still has no root, and a new polynomial $h_{b_{\alpha,t}+1,\alpha,t}$ is defined, with a root r'' in F_{t+1} but such that $h^-_{b_{\alpha,t}+1,\alpha,t}$ has no root there. So the situation remains the same, except that one of the two holding polynomials has been replaced by a new one. Before α can be eligible again, C_e must acquire an image for r'', but cannot acquire any root for the g_{n+1}-image of $h^-_{b_{\alpha,t},\alpha,t}$.

Having understood the above, we consider three cases.

(1) If α is initialized at some stage $t+1 > s+1$, then the g_{n+1}-image of $h_{b_{\alpha,t},\alpha,t}$ has a root in C_e, but the g_{n+1}-image of $h^-_{b_{\alpha,t},\alpha,t}$ will never acquire one. Therefore, our choice of $g_{n+1}(x_{\alpha,s+1})$ was correct.
(2) If α is never initialized after stage $s+1$ but is only eligible at finitely many stages, let t be the last stage at which it is eligible. The exact same analysis applies here as for the case when α is re-initialized.
(3) Otherwise α is never again initialized, but is eligible infinitely often. In this case, α must enter Step 3 infinitely many times (since $F \cong C_e$ precludes it from staying in Step 2 forever), and so $\langle x_\alpha, y_\alpha\rangle \in B_F$, as discussed above. Therefore, either of the square roots of p_α in C_e can be the image of x_α under an isomorphism. So in this case either choice for $g_{n+1}(x_{\alpha,s+1})$ would have been correct.

Thus our definition of g_{n+1} was correct for every node α going through Step 1. It remains to define g_{n+1} on all elements adjoined to F at any stage $s+1$ between stages s_n and s_{n+1} by $\mathcal{R}$-nodes α in Step 3 of the construction. But this definition is simple, because for such an α, the value $x_{\alpha,s+1}$ must already have been defined, and we have already defined $g_{n+1}(x_{\alpha,s+1})$. Therefore, when $h_{b_{\alpha,s+1},\alpha,s+1}$ is given a root r in F, we know the image of $h_{b_{\alpha,s+1},\alpha,s+1}$ in $C_e[X]$ under the map g_{n+1} on its coefficients, and we wait for this image to acquire a root in C_e, which then

becomes $g_{n+1}(r)$. Likewise, we know the image of $h_{b_{\alpha,s+1}-2,\alpha,s+1}(X)$ under g_{n+1}, and so we may wait for it to acquire a root in C_e, then define g_{n+1} to map the root of $h_{b_{\alpha,s+1}-2,\alpha,s+1}$ in F to this root in C_e. Since $C_e \cong F$, and since $g_{n+1}(x_{\alpha,s+1})$ is correctly defined, such roots must appear.

Thus we have extended g_{n+1} to all elements adjoined by any $\mathcal{R}$-node between stages s_n and s_{n+1}, so we have defined our computable embedding g_{n+1} on all of $F_{s_{n+1}}$. It is clear that this process can then continue to $F_{s_{n+2}}$ and beyond, so that $g = \bigcup_n g_n$ is a computable embedding of F into C_e. But since we know that $C_e \cong F$, Lemma 2.1 shows that g is a computable isomorphism, and so $\mathcal{C}_e\Sigma^0_1$ is satisfied.

The satisfaction of the requirement $\mathcal{R}_e$ shows that $\lim_s \varphi_e(\cdot, s)$ is not the characteristic function of B_F, and so all these requirements together prove that B_F is not Δ^0_2, hence not Σ^0_2. On the other hand, the satisfaction of the $\mathcal{C}$-requirements shows that F is computably categorical, since every computable field isomorphic to F has an atomic diagram decidable by some φ_e, meaning that it is the field C_e, which was made computably isomorphic to F by the requirement $\mathcal{C}_e$. These conclusions complete the proof of Theorem 4.1. □

5. Complexity of Computable Categoricity

Ostensibly, computable categoricity is a Σ^1_1 property, since its definition involves the existence of (classical) isomorphisms, hence involves quantifying over functions from ω to ω. However, for those classes of structures for which an exact complexity is known, it has always turned out to be far less complex than Σ^1_1. For instance, a computable linear order $\mathcal{L}$ is computably categorical if and only if $\mathcal{L}$ contains only finitely many pairs of adjacent points, and this condition can be expressed as a Σ^0_3 formula in the (computable) order relation on $\mathcal{L}$. Indeed, for arbitrary computable structures $\mathfrak{M}$, the statement "$\mathfrak{M}$ has a Σ^0_1 Scott family" is Σ^0_3, and so relative computable categoricity is always a Σ^0_3 property.

For algebraic fields, the very fact of being isomorphic is nowhere near Σ^1_1. Corollary 2.7 shows that for algebraic fields E and F, being isomorphic is Π^0_2, since for any finitely generated subfield F_0 we can effectively find a primitive generator of F_0, and then find the minimal polynomial of that generator over the prime subfield of F_0, so that the embeddability of F_0 into E reduces to the existence in E of a root of that minimal polynomial (translated from the prime subfield of F to that of E, of course). Thus, algebraic fields E and F over the same prime subfield Q are isomorphic if and only if

$$(\forall p(X) \in Q[X]) \left[(\exists x \in E\ p(x) = 0) \iff (\exists y \in F\ p(y) = 0) \right].$$

If we write C_e for the field (if any) whose atomic diagram has characteristic function φ_e, as in the proof of Theorem 4.1, then we can discuss various complexities exactly.

Proposition 5.1. *All of the following sets are Π^0_2-complete.*

- ***Fld*** $= \{e : C_e$ *is a field*$\}$.
- ***AlgFld*** $= \{e : C_e$ *is an algebraic field*$\}$.
- $\{\langle e, i\rangle : C_e$ *and* C_i *are isomorphic algebraic fields*$\}$.
- $\{i : C_i$ *is isomorphic to the field* $C_e\}$*, where* C_e *is any fixed algebraic field.*

Proof. Π^0_2 definitions of all these sets except **Fld** are readily produced, given Corollary 2.7 and our discussion above. Saying that φ_e is the characteristic function

of the atomic diagram of a field requires saying that φ_e is total with range $\{0,1\}$ (a Π^0_2 property) and that the field axioms are satisfied by this diagram. As usually stated, most of the field axioms are Π^0_2, but the existence of an identity element for each operation appears to be Σ^0_2, and the existence of inverses (stated below for multiplication) appears to be Σ^0_3:

$$\exists c \forall x \exists y (x + x = x \text{ or } x \cdot y = y \cdot x = c).$$

This sentence can be reduced to a Π^0_2 statement simply by having constant symbols for 0 and 1 in the signature, but it is worth noting that even without such constant symbols, the field axioms are still Π^0_2.

Lemma 5.2. *A structure $\mathfrak{M}$ in the signature with $+$ and $\cdot$ is a field if and only if these two operations are both associative and commutative, $\cdot$ distributes over $+$, and the following hold:*

$$\exists x \exists y (x \neq y) \ \&\ \forall x \forall y \exists u (x + u = y) \ \&\ \forall x \forall y \exists u (x + x \neq x \implies x \cdot u = y).$$

Thus the field axioms can be expressed as a single first-order $\forall\exists$ sentence.

Proof. The forwards implication is immediate, so assume that the given axioms hold. Fix any single x, and apply the middle axiom to get a u with $x + u = x$. But now for any y, we have some v with $x + v = y$ and hence, given associativity and commutativity,

$$y + u = (x + v) + u = v + (x + u) = v + x = y,$$

so that this u is actually an additive identity element 0. The given axiom for addition then yields additive inverses. But once we have these, we see that $x+x = x$ implies $x = 0$, so there must exist a y with $y+y \neq y$ (lest $\mathfrak{M}$ have only one element). Then we repeat for multiplication the same argument as for addition, using this y to get the identity element. □

The Π^0_2-completeness of the sets in Proposition 5.1 is mostly an elementary exercise. One easily shows that $\mathbf{Inf} = \{e : |W_e| = \infty\}$ (where W_e is the eth c.e. set) is 1-reducible to **Fld**, for instance, just by fixing a computable field F and, on a given input e, building the characteristic function of the decision procedure for F one element at a time, each time we get further evidence that $e \in \mathbf{Inf}$ (i.e., each time a new element enters W_e). It is worth noting that each of the other three sets is Π^0_2-complete (under 1-reductions) within the class **Fld**. (The relevant definition can be found in [3, Defn. 1.2].) For instance, there is a computable injective function f such that $\forall e (f(e) \in \mathbf{Fld})$, but the field $C_{f(e)}$ is algebraic if and only if $e \in \mathbf{Inf}$. (Start building the field $\mathbb{Q}(X_0, X_1, \ldots)$ of infinite transcendence degree, one element at a time, and when e gets its n-th chip, turn X_n into a rational number itself, so large that it has not yet been ruled out by the finitely many elements currently in $C_{f(e)}$.) □

Since classical isomorphism is so easily expressed for algebraic fields, the complexity of computable categoricity for the class becomes much simpler than Σ^1_1.

Proposition 5.3. *For algebraic fields, the property of being computably categorical is Π^0_4.*

Proof. We simply write out the definition of computable categoricity and apply Proposition 5.1. The computable algebraic field $F = C_e$ is computably categorical if and only if:

$$(\forall i)[(i \in \mathbf{Fld} \ \&\ C_i \cong C_e) \implies \exists j (\varphi_j \text{ is an isomorphism} : C_i \to C_e)].$$

The statements $i \in \mathbf{Fld}$ and $C_i \cong C_e$ are both Π^0_2. For φ_j to be an isomorphism, it must be total (which is Π^0_2) and must preserve the field structure:

$$\forall x \forall y[\varphi_j(x+y) = \varphi_j(x) + \varphi_j(y) \ \& \ \varphi_j(x \cdot y) = \varphi_j(x) \cdot \varphi_j(y)],$$

which is Π^0_1 once we know that φ_j is total. (For φ_j to have image ω is also Π^0_2, but in fact is not needed here, by Corollary 2.2.) □

Our main theorem for this section is the complementary property: that for computable algebraic fields, computable categoricity is Π^0_4-hard, and therefore Π^0_4-complete. This theorem is substantially different from previously known results about the complexity of computable categoricity for specific classes of structures, and thus serves to distinguish algebraic fields from all those other classes. In particular, all previously known cases were Σ^0_n-complete for some n, usually for $n = 3$, so Π^0_4-completeness suggests that something distinctly different is happening here.

Theorem 5.4. *For computable algebraic fields, the property of being computably categorical is Π^0_4-complete.*

Proof. With Proposition 5.3 already proven, it remains to show hardness. Let S be any Π^0_4-complete set, such as the complement of $\emptyset^{(4)}$. Since the set **Inf** is Π^0_2-complete, we may express S by fixing some 1-1 total computable function f for which:

$$S = \{n \in \omega : \forall a \exists b (f(n,a,b) \in \mathbf{Inf})\} = \{n : \forall a \exists b \ |W_{f(n,a,b)}| = \infty\}.$$

It will simplify our construction to assume that every set $W_{f(n,a,b)}$ contains the element 0, and that at each single stage, at most one set $W_{f(n,a,b)}$ receives a new element.

We will describe a 1-1 total computable function that maps each $n \in \omega$ to the index for some computable algebraic field F, which will be computably categorical if and only if $n \in S$. The output of this function is the program that uses the following construction (which is uniform in n) to build a computable field. At the end of the construction, we will demonstrate that the computable algebraic field F that it built is computably categorical if $n \in S$, but not otherwise.

The construction of F is performed on a tree T, in a style reminiscent of that in the proof of Theorem 4.1, adapted to incorporate the question of whether $\forall a \exists b \ f(n,a,b) \in \mathbf{Inf}$. As there, we let C_e denote the structure (in the language of fields) whose atomic diagram is decided by the partial function φ_e. The tree T for the construction will consist of two types of nodes. We now describe the basic modules used by each type to satisfy its requirement.

Every node β at level $2e$ of T is a *categoricity node*, or *$\mathcal{C}$-node*, dedicated to satisfying requirement $\mathcal{C}_e$ for computable categoricity for F:

$$\mathcal{C}_e : \ C_e \cong F \implies \exists \text{ a computable isomorphism } g_e : C_e \to F.$$

Such a $\mathcal{C}_i$-node β has two outcomes, $\cong$ and $\ncong$, ordered with $\cong \prec \ncong$. The outcome $\cong$ denotes that the hypothesis of $\mathcal{C}_e$ turned out to be true: $C_e \cong F$. In this case, the β on the true path at level $2e$ will produce the computable isomorphism g_β required, since no node above it or to its right will ever add anything to F that could cause problems for its isomorphism. This process is much the same as that performed by the categoricity nodes in the tree for Theorem 4.1. The outcome $\ncong$ denotes the negation of the outcome $\cong$, in which case $\mathcal{C}_e$ holds automatically.

Every node α at level $2a+1$ of the tree is a *non-categoricity node*, or $\mathcal{R}$*-node*, trying to construct a computable field $E_\alpha \cong F$ to satisfy the opposite requirement:

$$\mathcal{R}_a : \ [\forall b \ f(n,a,b) \notin \mathbf{Inf}] \implies [\forall b \ \varphi_b : E_\alpha \to F \text{ is not an isomorphism}].$$

The construction will build the computable fields E_α for every $\mathcal{R}$-node α, all isomorphic to F. An $\mathcal{R}_i$-node α has outcomes ordered in order type ω:

$$0 \ \prec \ 1 \ \prec \ 2 \ \prec \cdots$$

If α lies on the true path, then for the least $b \in \omega$ (if any) such that $f(n,a,b) \in \mathbf{Inf}$, the node $\alpha^\frown\langle b\rangle$ will be the leftmost successor eligible infinitely often. If there is no such b, then the hypothesis of $\mathcal{R}_a$ is satisfied, and in fact the true path will end at α; in this case, the field E_α built by α will prove that F is not computably categorical.

The $\mathcal{R}_i$-node α runs the following basic module simultaneously for all $b \in \omega$, although whenever $W_{f(n,a,b)}$ receives a new element, α restarts its strategy for every $b' \geq b$. For each b, α starts by adjoining one witness element to E_α (with a corresponding witness adjoined to F) and waits for φ_b to map the witness in E_α to the witness in F, which is its unique possible image there. If φ_b does so, then α adds a new element to F to "tag" the witness there. It waits until all categoricity nodes β with $\beta^\frown\langle\cong\rangle \subseteq \alpha$ have mapped the witness and its tag to an appropriate image, then adjoins a second witness to F, conjugate to the original witness there, and likewise adjoins a second witness to E_α. However, in E_α, α tags the second witness instead of the first. Therefore, assuming no further tags nor conjugates of the two witnesses ever appear in F, φ_b cannot be an isomorphism, since it mapped the untagged witness in E_α to the tagged witness in F.

All through this process (and forever after), α keeps watching to see if $W_{f(n,a,b)}$ receives any more elements. If it ever does, then α terminates its procedure for b and for all $b' > b$, makes $\alpha^\frown\langle b\rangle$ eligible and begins its entire process over again with a new witness (which is the root of a completely new minimal polynomial). Therefore, α precludes φ_b from being an isomorphism only if $f(n,a,b) \notin \mathbf{Inf}$. If every $f(n,a,b) \notin \mathbf{Inf}$, then all of α's basic modules succeed, leaving E_α isomorphic to F but not computably isomorphic to it.

At stage 0, we begin with $F_0 = \mathbb{Q}$ and also all fields $E_{\alpha,0} = \mathbb{Q}$. All nodes are initialized, so that all values mentioned below for each node are undefined at stage 0.

The stages are ordered as in the construction in Theorem 4.1, so that the root is eligible at every stage $\langle 0,k\rangle + 1$, and at each stage $\langle l,k\rangle + 1$, some node at level l is eligible and (if $l < k$) chooses a node at level $(l+1)$ to be eligible at the following stage $\langle l+1,k\rangle + 1$.

At stage $s+1$, suppose that the $\mathcal{C}_e$-node β is eligible. Let s' be the greatest stage $\leq s$ at which either β was initialized or the node $\beta^\frown\langle\cong\rangle$ was eligible. If the length of agreement between F_s and $C_{e,s}$ (as defined in the proof of Theorem 4.1) is no greater than the domain of $g_{\beta,s'}$, then we do nothing at this stage, and make $\beta^\frown\langle\ncong\rangle$ eligible at the next stage. If the length of agreement has increased, then $\beta^\frown\langle\cong\rangle$ will be eligible at the next stage. At this stage, we define the map $g_{\beta,s+1}$ to extend the map $g_{\beta,s'}$ to the next element of the field C_e. (By assumption, this must be a partial field embedding.) This completes the stage.

At a stage $s+1$ at which an $\mathcal{R}_a$-node α is eligible, we again let s' be the greatest stage $\leq s$ at which α either was initialized or was eligible. Fix the least b_0 for which

$W_{f(n,a,b_0),s'} \neq W_{f(n,a,b_0),s+1}$. (If there is no such b_0, then find the least $t > s+1$ for which $(\exists b_0) W_{f(n,a,b_0),s'} \neq W_{f(n,a,b_0),t}$, and choose that b_0. Since all the sets $W_{f(n,a,b)}$ are nonempty and only one can receive an element at any given stage, we eventually find such a b_0.) The node $\alpha^\frown\langle b_0\rangle$ will be eligible at the next stage.

If α was initialized at stage s', then we simply set both $E_{\alpha,s+1}$ and F_{s+1} to equal F_s, and end this substage. If α was not initialized at stage s', then we execute the following instructions.

For each $b \geq b_0$, we initialize the α-strategy for b, by making $p_{\alpha,b,s+1}$ and all related roots, witnesses, and tags undefined. First, however, for each $b \geq b_0$ for which α is currently waiting to perform Step 3 (so $x_{\alpha,b,s'} \in F_s$, but $E_{\alpha,s}$ does not yet contain any element $u_{\alpha,b,s'}$), we adjoin $\tilde{x}_{\alpha,b,s'}$ to $E_{\alpha,s}$ (and then make $\tilde{x}_{\alpha,b,s+1}$ undefined, along with all other roots and tags). This ensures that $E_{\alpha,s+1}$ becomes isomorphic to F_{s+1} once again (except possibly for certain tags for α-strategies for values $b' < b_0$; such tags might still lie in F_s but have no images in $E_{\alpha,s+1}$).

For each $b < b_0$, we proceed according to the following steps.

(1) If no polynomial $p_{\alpha,b,s'}(X)$ is currently defined, then we use Proposition 2.4 to choose a polynomial $p_{\alpha,b,s+1}(X) \in \mathbb{Q}[X]$ of degree 7, whose Galois group (over the splitting field of the product of all p-polynomials used so far in the construction, i.e. all $p_{\alpha',b',t}(X)$ with $t \leq s$) is the symmetric group $\mathcal{S}_7$ on its seven roots. (Here we regard $\mathbb{Q}$ as a subfield of F_s, so that this polynomial lies in $F_s[X]$.) We define $x_{\alpha,b,s+1}$ and $y_{\alpha,b,s+1}$ to be two roots of $p_{\alpha,b,s+1}(X)$, but at this step we only adjoin their sum $(x_{\alpha,b,s+1} + y_{\alpha,b,s+1})$ to F_s, forming F_{s+1} and leaving the roots themselves for possible later use. Likewise, we adjoin the sum $(\tilde{x}_{\alpha,b,s+1} + \tilde{y}_{\alpha,b,s+1})$ of two roots of $\tilde{p}_{\alpha,b,s+1}(X)$ to every field $E_{\alpha',s}$, for every $\mathcal{R}$-node α' including $E_{\alpha,s}$ itself. So each such field $E_{\alpha',s+1}$ remains isomorphic to F_{s+1} (unless $E_{\alpha',s} \not\cong F_s$). We define $q_{\alpha,b,s+1}(X) \in \mathbb{Q}[X]$ to be the minimal polynomial of $(x_{\alpha,b,s+1} + y_{\alpha,b,s+1})$ over $\mathbb{Q}$. Roots of $q_{\alpha,b,s+1}(X)$ will be called *witnesses* being used for α and b, in their respective fields F and E_α. (In Step 3, a second witness may be adjoined to each of F and E_α.)

(2) If $x_{\alpha,b,s'}$ and $y_{\alpha,b,s'}$ are already defined but $(x_{\alpha,b,s'} + y_{\alpha,b,s'})$ has not yet been tagged (as below), then we check whether $\varphi_{b,s}(x_{\alpha,b,s'} + y_{\alpha,b,s'}) \downarrow = (\tilde{x}_{\alpha,b,s+1} + \tilde{y}_{\alpha,b,s+1})$. If not, then we do nothing at this stage. If so, then we adjoin $x_{\alpha,b,s'}$ to F_s, calling it a *tag* for the witness $(x_{\alpha,b,s'} + y_{\alpha,b,s'})$. To preserve isomorphisms, we also adjoin $\tilde{x}_{\alpha,b,s'}$ to $E_{\alpha',s}$ for every $\mathcal{R}$-node α' except α, keeping $E_{\alpha',s+1} \cong F_{s+1}$ (unless $E_{\alpha',s} \not\cong F_s$). Thus we leave $E_{\alpha,s+1} = E_{\alpha,s} \not\cong F_{s+1}$, with no tag adjoined to $E_{\alpha,s}$.

(3) If $x_{\alpha,b,s'} \in F_s$ already, and $E_{\alpha,s'}$ contains no corresponding tag, then we check whether, for every $\mathcal{C}_e$-node β with $\beta^\frown\langle\cong\rangle \subseteq \alpha$, the domain of $g_{\beta,s+1}$ contains $x_{\alpha,b,s'}$ and the field fragment $C_{e,s+1}$ contains exactly one witness for α and b. If not, then we do nothing. If so, then we define $u_{\alpha,b,s+1}$ and $v_{\alpha,b,s+1}$ to be new roots of $p_{\alpha,b,s+1}(X)$, adjoin their sum $(u_{\alpha,b,s+1} + v_{\alpha,b,s+1})$ to F_s as a new witness, and likewise adjoin a new witness $(\tilde{u}_{\alpha,b,s+1} + \tilde{v}_{\alpha,b,s+1})$, the sum of two new roots of $\tilde{p}_{\alpha,b,s'}(X)$, to every $E_{\alpha',s}$ with $\alpha' \neq \alpha$. To $E_{\alpha,s}$ we adjoin the two new roots $\tilde{u}_{\alpha,b,s+1}$ and $\tilde{v}_{\alpha,b,s+1}$ of $\tilde{p}_{\alpha,b,s'}(X)$; this also adjoins their sum, of course, as a new witness, and leaves $F_{s+1} \cong E_{\alpha,s+1}$, but only via isomorphisms mapping $(x_{\alpha,b,s'} + y_{\alpha,b,s'})$ to $(\tilde{u}_{\alpha,b,s+1} + \tilde{v}_{\alpha,b,s+1})$, since these are the witnesses in

their respective fields that now have tags. This situation will be preserved forever (unless either α is initialized or some $W_{f(n,a,b')}$ with $b' \leq b$ later receives a new element), and so φ_b cannot be an isomorphism from F onto E_α.

(4) If none of the foregoing conditions applies, then α has satisfied $\mathcal{R}_a$, and we do nothing at this stage.

Having completed these steps for every $b < b_0$, we have finished this stage.

When stage $s+1$ is completed, we initialize every node to the right of the node eligible at that stage (exactly as we did for α-strategies for each $b \geq b_0$ in the construction for $\mathcal{R}$-nodes). For a $\mathcal{C}$-node β, initialization simply means that $g_{\beta,s+1}$ becomes the empty function. For an $\mathcal{R}$-node α, and for every $b \in \omega$, we make all polynomials, roots, and tags associated with α undefined at stage $s+1$, and we also make $E_{\alpha,s+1}$ undefined. This completes stage $s+1$.

It is clear that this construction builds a computable algebraic field F, uniformly in n, and that this field is the extension of $\mathbb{Q}$ generated by various witnesses and tags adjoined by assorted $\mathcal{R}$-nodes. We claim that F is computably categorical if and only if $n \in S$, which is to say, if and only if for every a there is some b with $f(n,a,b) \in \mathbf{Inf}$. As usual, the proof is based on the *true path* P through T, i.e. the set of all nodes in T that are eligible at infinitely many stages, but initialized only finitely many times.

Suppose first that $n \in S$. Now every $\mathcal{C}$-node β makes one of its two successors eligible whenever β itself is eligible. Moreover, an $\mathcal{R}_a$-node α on P will make its successor $\alpha^\frown\langle b\rangle$ eligible infinitely often, where b is minimal such that $f(n,a,b) \in \mathbf{Inf}$, while for every $b' < b$, $\alpha^\frown\langle b'\rangle$ will be eligible only finitely often. With $n \in S$, this means that P will be an infinite path through T, picking out the least b corresponding to each a at the $\mathcal{R}_a$-node α, and picking out $\beta^\frown\langle\cong\rangle$ or $\beta^\frown\langle\ncong\rangle$ above a $\mathcal{C}_e$-node β according as $C_e \cong F$ or not.

Now the list of fields C_e includes every computable presentation of every computably presentable field. So, if F is isomorphic to an arbitrary computable field E (via an isomorphism f, say), then that E is precisely equal to some C_e. We claim that the $\mathcal{C}_e$-node β on P allows us to compute an isomorphism g from F onto C_e. First, let s_0 be a stage after which β is never initialized (so that no node to the left of β is ever again eligible). Now for every $\mathcal{R}$-node $\alpha \subset \beta$, fix $b_\alpha \in \omega$ such that $\alpha^\frown\langle b_\alpha\rangle \subseteq \beta$. Each of these $\alpha^\frown\langle b_\alpha\rangle$ is initialized only finitely often, and the construction makes it clear that each one, after its final initialization, adjoins only finitely many elements to F: at most two witnesses and one tag. Therefore, there exists a stage $s_1 \geq s_0$ after which no $\alpha \subset \beta$ ever again adjoins any elements to F. Since the field F_{s_1} is finitely generated, $f\restriction F_{s_1}$ is computable from the images of its generators, which constitute finitely much information. Hence we may set $g\restriction F_{s_1} = f\restriction F_{s_1}$.

It remains to define g on elements adjoined by other $\mathcal{R}$-nodes α. If α lies to the right of β, then whenever α adjoins any element to F at some stage s in its strategy for some b, we simply wait until the next stage at which α is initialized. Once this stage is complete, α never again adjoins any elements from the splitting field of $p_{\alpha,b,s}(X)$, and so once that stage is reached, we may find images for these elements in C_e (since $C_e \cong F$) and define g to map them there. (Of course, this uses Proposition 2.4 and the choice of the p-polynomials to show that every such splitting field is linearly disjoint from the compositum of all the others, and that

therefore these values for g do not interfere with the construction of g on any other splitting field.)

Finally, suppose $\beta \subset \alpha$. Of course we do not know whether such an α lies on P or to its left or right. However, when that α adjoins its first witness $(x_{\alpha,b,s} + y_{\alpha,b,s})$ to F at some stage s for the α-strategy for some b, we simply look for the first root of $q_{\alpha,b,s}(X)$ to appear in C_e, and let g map the first witness to that root. (Since $C_e \cong F$, such a root must eventually appear in C_e, and by linear disjointness, this extension of g is still a field embedding.) If the α-strategy for b never moves beyond Step 1, then F contains no more elements of the splitting field of $p_{\alpha,b,s}(X)$, and so this is sufficient. If it continues to Step 2 and adds the tag $x_{\alpha,b,s'}$ to F at some stage $s' > s$, then we wait for such a tag to appear in C_e and define it to be $g(x_{\alpha,b,s'})$. Notice that even if α eventually adjoins a second witness to F at a later stage s'', the first witness to appear in C_e *must* be the one with the tag. This follows from Step 3 of the construction for $\mathcal{R}$-nodes, in which α waits until C_e contains exactly one witness node and also contains a tag for that node. If C_e acquired a second witness before it acquired the tag for the first one, then the construction would never have adjoined the second witness to F, and C_e would not be isomorphic to F, contrary to hypothesis. So C_e must have produced the tag for $g(x_{\alpha,b,s} + y_{\alpha,b,s})$ before adjoining any second witness, and therefore it was safe for us to define g as we did on the first witness in F. When (and if) F acquires a tag for its first witness (in Step 2), C_e must subsequently acquire a tag for its own first witness (in order to be isomorphic to F), and then the second witness $(u_{\alpha,b,s''} + v_{\alpha,b,s''})$ to appear in F (if α should execute Step 3 in its strategy) will be matched by an (untagged) witness in C_e, to which g maps the second witness in F. Thus we can compute the value of this g on every generator of F, and so g is a computable field embedding of F into C_e. But with $C_e \cong F$ by assumption, Corollary 2.2 shows that this g is then an isomorphism. Hence F is computably categorical.

Next, suppose that $n \notin S$, and fix the least a such that no b satisfies $f(n, a, b) \in \mathbf{Inf}$. Now as argued above, each node on the true path P at any level $\leq 2a$ will have a successor on P. When we reach the $\mathcal{R}_a$-nodes at level $2a+1$, however, the $\alpha \in P$ at that level will have no successor eligible infinitely often, since $(\forall b) f(n, a, b) \notin \mathbf{Inf}$. We claim that instead, the field E_α built by this α after its last initialization is isomorphic to F, yet not computably isomorphic to F. Since E_α is clearly a computable field (given finitely much information, namely the last stage at which α was initialized), this will show that F is not computably categorical.

To see that $F \cong E_\alpha$, we begin at the first stage s_0 at which α is eligible after its last initialization. At this stage E_{α,s_0} is defined to be F_{s_0} itself. At all subsequent stages, the construction (for every node α', not just α) never adjoins an element to F without adjoining a corresponding element to E_α. The only exceptions to this rule are performed by α itself, at Step 2 of its strategies for various values of b: in Step 2 at those stages s, α adjoins $x_{\alpha,b,s}$ to F (which already contained the witness $(x_{\alpha,b,s} + y_{\alpha,b,s})$) without adjoining any element to E_α (which already contained a witness element $(\tilde{x}_{\alpha,b,s} + \tilde{y}_{\alpha,b,s})$ of its own). But at all subsequent stages, α will attempt to execute Step 3 for this b. It will not be allowed to do so as long as any $\mathcal{C}_e$-node β with $\beta^\frown\langle\cong\rangle \subseteq \alpha$ prevents it, which occurs if that C_e fails to contain exactly one witness for the α-strategy for b, along with a tag for that witness. However, if this C_e prevented it forever in this manner, then C_e would not be isomorphic to F, contradicting the fact that such a $\beta^\frown\langle\cong\rangle$ must lie on P. Therefore, eventually each

of the finitely many $\mathcal{C}$-nodes below α gives permission for α to execute Step 3 in its strategy for b. In doing so, α adjoins to E_α a new tagged witness, and adjoins to F a new untagged witness. Moreover, by linear disjointness, no more elements of the splitting field of $p_{\alpha,b,s}(X)$ ever again enter either F or E_α, Thus the witnesses and tags in E_α and F can be paired up perfectly, and so indeed E_α is isomorphic to F.

Finally, suppose that some φ_b were an isomorphism from F onto E_α. Then, at some stage t_b after which $W_{f(n,a,b)}$ receives no more elements, the construction will have adjoined a first witness element $(x_{\alpha,b,t_b} + y_{\alpha,b,t_b})$ to F for b. The isomorphism φ_b must map it to the corresponding witness $(\tilde{x}_{\alpha,b,t_b} + \tilde{y}_{\alpha,b,t_b})$ adjoined to E_α at the same stage, since these elements have no other conjugates in their fields at that stage, and none are ever added unless φ_b maps the witness in F to that in E_α. But once it does, α executes Step 2, adjoining a tag for the witness in F, and then (as we saw just above) eventually executes Step 3 and adjoins a new tagged witness in E_α and a new untagged witness in F. Therefore, φ_b maps the tagged witness $(x_{\alpha,b,t_b} + y_{\alpha,b,t_b})$ in F to the untagged witness $(\tilde{x}_{\alpha,b,t_b} + \tilde{y}_{\alpha,b,t_b})$ in E_α, and so φ_b is not an isomorphism after all. Since this holds for every b, F is not computably categorical. This completes the proof of Theorem 5.4. □

At first glance, the foregoing proof appears to be a standard $\emptyset''$ construction, using the true path P through a computable tree. However, a $\emptyset''$ oracle is not in fact enough to compute P. It can compute the successor on P of any $\mathcal{C}$-node $\beta \in P$, and it can compute the successor of an $\mathcal{R}$-node $\alpha \in P$ provided that α has one. However, P may actually end at α (in which case E_α is the computable field showing that F is not computably categorical), and this situation holds if and only if $\forall b f(n,a,b) \notin \mathbf{Inf}$, which is a Π^0_3 condition. So in fact, to compute P and recognize when it terminates (if ever), a $\emptyset'''$ oracle is required.

6. Conclusions and Questions

The ultimate goal of this project was to provide a structural characterization for computable categoricity for algebraic fields. The main question, therefore, is the extent to which we have achieved this goal. Admittedly, the goal itself is somewhat vague: what constitutes a structural characterization? A first-order property in model theory would be the ideal result, but this goal seems beyond reach.

For illumination on this question, consider the characterization of computably categorical linear orders $\mathcal{L}$ as those with only finitely many adjacencies. This property is not expressible in first-order languages, as one quickly proves using the Compactness Theorem. It is also readily seen to be a Σ^0_3-complete property, and so, in terms of complexity, we know exactly the level of difficulty of deciding computable categoricity for computable linear orders.

Notice also that, because computable categoricity implies relative computable categoricity for linear orders, another equivalent characterization would be the existence of a Σ^0_1 Scott family for $\mathcal{L}$. This property is also Σ^0_3-complete, for linear orders as for computable structures in general, and could also be taken as a characterization of computable categoricity. However, it is vastly less satisfying than the characterization by the number of adjacencies: the latter feels much more "structural." To quantify this, we note that the characterization using adjacencies can be expressed as a computable $L_{\omega_1\omega}$ formula (that is, with countable conjunctions and disjunctions allowed) in the language of linear orders. In addition, the proof of the equivalence of the latter to computable categoricity makes it clear exactly how

the property of finitely many adjacencies corresponds to computable categoricity, much more clear than can be said of the characterization by Scott families. So we consider the characterization by adjacencies to be the better characterization.

Since the initial consideration of computable categoricity for fields by Fröhlich and Shepherdson in [9], the problem of characterizing computable categoricity for fields has not given much ground. Without offering specific justification, we suspect that the results in this article are as good as one is likely to get in the case of algebraic fields. As far as complexity, that statement can be quantified: Π^0_4-completeness of computable categoricity for algebraic fields, demonstrated in Theorem 5.4, pinpoints the complexity of the notion. Likewise, of course, the characterization for relative computable categoricity turned out to be Σ^0_3 (as it must, by the work in [2] and [4]) and complete at that level (as commonly happens for relative computable categoricity). As usual, the characterization by Scott families is unsatisfying, and we consider Theorem 3.4 to be a significant step forward, since it equates this characterization to the more structural notion in items (4) and (5) of that theorem. It is not clear that any more satisfactory characterization of relative computable categoricity is likely to be discovered.

For computable categoricity, we likewise consider Theorem 5.4 to be substantial progress. Nevertheless, the result still feels less satisfactory. The property given in Proposition 5.3 is really just the definition of computable categoricity, in the specific context of algebraic fields. Theorem 5.4 then shows that one cannot do better, in terms of complexity, and we consider it important to recognize that in this context, Definition 1.1 can achieve the minimum possible complexity, simply by replacing the notion of classical isomorphism by an equivalent statement (namely the condition from Corollary 2.8). We believe that this is the first known instance of this phenomenon. However, it still does not seem impossible that a "more structural" characterization might be found.

We attach additional importance to Theorem 5.4 because of the new level of complexity it exhibits. Previous characterizations of computable categoricity for standard classes of computable structures have generally shown it to be Σ^0_3-complete (and equivalent to relative computable categoricity): this situation holds for linear orders, Boolean algebras, trees (as partial orders), and ordered abelian groups, for example. Relative computable categoricity is widely viewed as a "nicer" property, largely because of its straightforward syntactic characterization in [2] and [4], and it was already known that computable categoricity has strictly higher complexity than relative computable categoricity in many well-known classes of structures, such as graphs, partial orders, groups, and rings. In [41], White showed that for computable graphs, computable categoricity is Π^0_4-hard, and [16] allows the complexity result to be carried over to the other well-known classes mentioned (although it only proves computable categoricity to be Π^0_4-hard in those classes, not necessarily Π^0_4-complete). The fact that computable categoricity turned out to be Π^0_4-complete for algebraic fields took us rather by surprise, as this is the first everyday class of mathematical structures in which it turned out to be a Π^0_n-complete property (as opposed to Σ^0_n-complete) for any n at all. Indeed, to our knowledge, algebraic fields are the first standard class of structures for which the complexity of computable categoricity has been determined and has turned out not to be Σ^0_3-complete.

(For careful readers, we point out a small error in the final paragraph of [41], where it is asserted that computable categoricity is Π^0_3-complete for the class of

algebraically closed fields. In fact, for such fields, Ershov [6] showed it to be equivalent to the property of having finite transcendence degree, which is Σ^0_3-complete and is also equivalent to relative computable categoricity for such fields. Likewise, as of the writing of [41], all other known index sets for computable categoricity were Σ^0_3, not Π^0_3 as stated there.)

It should be noted that the class of algebraic fields is not first-order definable: every axiom set that holds in all algebraic fields will hold in certain non-algebraic fields as well. This fact might help explain the unusual level of complexity. In characteristic 0, our theorems carry over to fields of finite transcendence degree over $\mathbb{Q}$, since essentially all constructions can be carried out after replacing $\mathbb{Q}$ by a purely transcendental subfield $\mathbb{Q}(X_1, \ldots, X_k)$ over which F is algebraic. (Alternatively, for transcendence degree k, just enrich the signature by k constants, with axioms saying that they are algebraically independent over $\mathbb{Q}$.) For fields of infinite transcendence degree, the question of computable categoricity is not trivial: most such fields are not computably categorical, but the work of Miller and Schoutens in [32] proved the existence of a computably categorical field of infinite transcendence degree. One would guess that computable categoricity has even higher complexity for the class of all fields; it certainly cannot become any lower than Π^0_4, since algebraic fields form a subclass.

Finally, we mention computable dimension for algebraic fields. Goncharov defined the *computable dimension* of a computable structure to be the number of computable presentations of that structure, up to computable isomorphism. He showed that every cardinal from 1 through ω can be the computable dimension of a computable structure. (See [11] and [12] for these and related results.) However, by far the most common computable dimensions are 1 (which is equivalent to computable categoricity) and ω, and for many classes of structures, these are the only possible computable dimensions: linear orders, Boolean algebras, and trees, for example. At one time, we believed that we had proven this to be true of algebraic fields as well. However, a problem subsequently was found in the proof of the theorem (from another article) which we used to show the impossibility of finite computable dimension > 1. A recent result in [31] shows that in computable fields of infinite transcendence degree over $\mathbb{Q}$, all computable dimensions are possible. However, the question of finite computable dimension for computable algebraic fields remains open. We conjecture it to be impossible in purely quadratic algebraic extensions of $\mathbb{Q}$ (i.e., those algebraic extensions F such that for every finitely generated subfield E of F, the degree of E over $\mathbb{Q}$ is a power of 2), and also in the related class containing all computable finite-branching subtrees of $\omega^{<\omega}$. It is known that the latter class includes structures which are computably categorical but not relatively so (by a proof very similar to that in Section 4), and so this conjecture would imply that the condition of computable categoricity without relative computable categoricity need not entail the existence of structures of finite computable dimension > 1: the one pathology can occur without the other.

References

1. C.J. Ash & J.F. Knight; *Computable Structures and the Hyperarithmetical Hierarchy* (Amsterdam: Elsevier, 2000).
2. C.J. Ash, J.F. Knight, M.S. Manasse, & T.A. Slaman; Generic copies of countable structures, *Annals of Pure and Applied Logic* **42** (1989), 195–205.

3. W. Calvert, V. Harizanov, J.F. Knight, & S. Miller; Index sets for computable structures, *Algebra and Logic* **45** (2006), 306–325.
4. J. Chisholm; On intrinsically 1-computable trees, unpublished MS.
5. R.G. Downey, D.R. Hirschfeldt, & B. Khoussainov; Uniformity in computable structure theory, *Algebra and Logic* **42** (2003), 318–332.
6. Yu.L. Ershov; Theorie der Numerierungen, *Zeits. Math. Logik Grund. Math.* **23** (1977), 289–371.
7. Yu.L. Ershov & S.S. Goncharov, Constructive fields, Section 2.5 in *Constructive Models* (New York: Kluwer Academic/Plenum Press, 2000).
8. M.D. Fried & M. Jarden, *Field Arithmetic* (Berlin: Springer-Verlag, 1986).
9. A. Fröhlich & J.C. Shepherdson; Effective procedures in field theory, *Phil. Trans. Royal Soc. London, Series A* **248** (1956) 950, 407–432.
10. S.S. Goncharov; Autostability and computable families of constructivizations, *Algebra and Logic* **14** (1975), 647–680 (Russian), 392–409 (English translation).
11. S.S. Goncharov; Nonequivalent constructivizations, *Proc. Math. Inst. Sib. Branch Acad. Sci.* (Novosibirsk: Nauka, 1982).
12. S.S. Goncharov; Autostable models and algorithmic dimensions, *Handbook of Recursive Mathematics*, vol. 1 (Amsterdam: Elsevier, 1998), 261–287.
13. S.S. Goncharov & V.D. Dzgoev; Autostability of models, *Algebra and Logic* **19** (1980), 45–58 (Russian), 28–37 (English translation).
14. S.S. Goncharov, S. Lempp & R. Solomon; The computable dimension of ordered abelian groups, *Advances in Mathematics* **175** (2003) 1, 102–143.
15. V.S. Harizanov; Pure computable model theory, *Handbook of Recursive Mathematics*, vol. 1 (Amsterdam: Elsevier, 1998), 3–114.
16. D.R. Hirschfeldt, B. Khoussainov, R.A. Shore, & A.M. Slinko; Degree spectra and computable dimensions in algebraic structures, *Annals of Pure and Applied Logic* **115** (2002), 71–113.
17. D.R. Hirschfeldt, B. Khoussainov, & R.I. Soare; On computable isomorphisms of graphs, to appear.
18. N. Jacobson; *Basic Algebra I* (New York: W.H. Freeman & Co., 1985).
19. B. Khoussainov & R.A. Shore; Computable isomorphisms, degree spectra of relations, and Scott families, *Annals of Pure and Applied Logic* **93** (1998), 153-193.
20. N. Kogabaev, O. Kudinov, & R.G. Miller; The computable dimension of I-trees of infinite height, *Algebra and Logic* **43** (2004) 6, 393–407.
21. O.V. Kudinov; An autostable 1-decidable model without a computable Scott family of $\exists$ formulas, *Algebra and Logic* **35** (1996), 255–260 (English translation).
22. S. Lang; *Algebra,* revised third edition (Springer-Verlag, 2002).
23. S. Lang; *Algebraic Number Theory,* second edition (Springer-Verlag, 1994).
24. S. Lempp, C. McCoy, R.G. Miller, & R. Solomon; Computable categoricity of trees of finite height, *Journal of Symbolic Logic* **70** (2005), 151–215.
25. G. Metakides & A. Nerode; Effective content of field theory, *Annals of Mathematical Logic* **17** (1979), 289–320.
26. R. Miller; The computable dimension of trees of infinite height, *Journal of Symbolic Logic* **70** (2005), 111–141.
27. R. Miller, Computable fields and Galois theory, *Notices of the American Mathematical Society* **55** (August 2008) 7, 798–807.
28. R. Miller; Is it harder to factor a polynomial or to find a root?, *Transactions of the American Mathematical Society*, **362** (2010) 10, 5261-5281.
29. R. Miller; $\boldsymbol{d}$-Computable categoricity for algebraic fields, *The Journal of Symbolic Logic* **74** (2009) 4, 1325–1351.
30. R. Miller, Computability and differential fields: a tutorial, to appear in *Differential Algebra and Related Topics: Proceedings of the Second International Workshop*, eds. L. Guo & W. Sit. Also available at `qcpages.qc.cuny.edu/~rmiller/research.html`.
31. R. Miller, J. Park, B. Poonen, H. Schoutens, & A. Shlapentokh; A computable functor from graphs to fields, in preparation.
32. R. Miller & H. Schoutens; Computably categorical fields via Fermat's Last Theorem, *Computability* **2** (2013) 51–65.
33. R. Miller & A. Shlapentokh; Computable categoricity for algebraic fields with splitting algorithms, to appear in the *Transactions of the American Mathematical Society.*

34. M. Rabin; Computable algebra, general theory, and theory of computable fields, *Transactions of the American Mathematical Society* **95** (1960), 341–360.
35. J.B. Remmel; Recursively categorical linear orderings, *Proceedings of the American Mathematical Society* **83** (1981), 387–391.
36. J.B. Remmel; Recursive isomorphism types of recursive Boolean algebras, *Journal of Symbolic Logic* **46** (1981), 572–594.
37. R.I. Soare; *Recursively Enumerable Sets and Degrees* (New York: Springer-Verlag, 1987).
38. V. Stoltenberg-Hansen & J.V. Tucker; Computable rings and fields, in *Handbook of Computability Theory*, ed. E.R. Griffor (Amsterdam: Elsevier, 1999), 363–447.
39. B.L. van der Waerden; *Algebra*, volume I, trans. F. Blum & J.R. Schulenberger (New York: Springer-Verlag, 1970 hardcover, 2003 softcover).
40. Y.G. Ventsov; Effective choice for relations and reducibilities in classes of constructive and positive models, *Algebra and Logic* **31** (1992), 63–73.
41. W.M. White; On the complexity of categoricity in computable structures, *Mathematical Logic Quarterly* **49** (2003) 6, 603–614.

Department of Mathematics, University of Chicago, 5734 S. University Ave., Chicago, IL 60637 U.S.A.
E-mail address: drh@math.uchicago.edu

Department of Mathematics, Queens College – C.U.N.Y., 65-30 Kissena Blvd., Flushing, New York 11367 U.S.A.; Ph.D. Program in Mathematics, C.U.N.Y. Graduate Center, 365 Fifth Avenue, New York, NY 10016 U.S.A.
E-mail address: kkramer@qc.cuny.edu

Department of Mathematics, Queens College – C.U.N.Y., 65-30 Kissena Blvd., Flushing, New York 11367 U.S.A.; Ph.D. Programs in Mathematics and Computer Science, C.U.N.Y. Graduate Center, 365 Fifth Avenue, New York, NY 10016 U.S.A.
E-mail address: Russell.Miller@qc.cuny.edu

East Carolina University, Department of Mathematics, Greenville, NC 27858 U.S.A.
E-mail address: shlapentokha@ecu.edu