



On Computable Field Embeddings and Difference Closed Fields

Matthew Harrison-Trainor, Alexander Melnikov, and Russell Miller

Abstract. We investigate when a computable automorphism of a computable field can be effectively extended to a computable automorphism of its (computable) algebraic closure. We then apply our results and techniques to study effective embeddings of computable difference fields into computable difference closed fields.

1 Introduction

This article is a contribution to effective field theory, where the main objects of study are computable fields. Recall that an algebraic structure is computable if the elements of its domain are associated with natural numbers in such a way that the operations become computable functions upon this domain [Mal61, Rab60]. There are a number of classical results that say that maps between fields can be extended to maps between their algebraic closures. We consider when this can be done effectively, *i.e.*, if all of the fields involved are computable, and we are given a computable map, must there exist a computable extension to the algebraic closures? We obtain both necessary and sufficient conditions on a computable field \mathcal{F} which ensure that these classical theorems hold effectively for the field \mathcal{F} . We also apply our results to computable fields with a distinguished (computable) automorphism; such fields are known as *difference fields*. We investigate the problem of effectively embedding difference fields into computable difference-closed fields (these are existentially closed difference fields, to be discussed). As we will see, the most naive analogy of the well-known results of Rabin [Rab60] and Harrington [Har74] fails for computable difference fields, in all characteristics. Nonetheless, we will find a broad class of fields (including abelian extensions of a prime field) for which a stronger version of the analogous result holds.

1.1 Embeddings into Algebraically Closed Fields

In a pioneering paper, Rabin [Rab60] proved that every computable field \mathcal{F} can be embedded into a computable presentation \mathcal{E} of its algebraic closure by a computable map $\iota: \mathcal{F} \rightarrow \mathcal{E}$. Provided that \mathcal{E} is algebraic over the image $\iota(\mathcal{F})$, we call such an

Received by the editors January 20, 2016; revised November 21, 2016.

Published electronically May 25, 2017.

Author M. H.-T. was partially supported by the Berkeley Fellowship and NSERC grant PGSD3-454386-2014. Author A. M. was partially supported by the Packard Foundation. Author R. M. was supported by NSF grants # DMS-1362206 and DMS-1001306, and by several PSC-CUNY research awards. Some of this work took place at a workshop held by the Institute for Mathematical Sciences of the National University of Singapore.

AMS subject classification: 03D45, 03C57, 12Y05.

Keywords: computable algebra, algebraic field, difference field, extension of automorphism.

embedding ι a *Rabin embedding* of \mathcal{F} into \mathcal{E} , writing $\overline{\mathcal{F}}$ for \mathcal{E} since \mathcal{E} may thus be regarded as an algebraic closure of \mathcal{F} . In what follows it will be important that, in general, the image of \mathcal{F} under the Rabin embedding ι does not have to be a computable subset of $\overline{\mathcal{F}}$. Rabin showed that the problem of deciding the ι -image of \mathcal{F} in $\overline{\mathcal{F}}$ is fully captured by the notion of the *splitting set*. Recall that the splitting set $S_{\mathcal{F}}$ of \mathcal{F} is the set of all polynomials $p \in \mathcal{F}[X]$ which are reducible over \mathcal{F} . If the splitting set of \mathcal{F} is computable, then we say that \mathcal{F} has a *splitting algorithm*. Rabin showed that for each computable field \mathcal{F} , and for each Rabin embedding ι of \mathcal{F} , the image $\iota(\mathcal{F})$ of \mathcal{F} in $\overline{\mathcal{F}}$ is Turing equivalent to the splitting set of \mathcal{F} , which may be undecidable. We note that splitting algorithms had been studied long before Rabin. For instance, in 1882, Kronecker [Kro82] analyzed splitting algorithms for finitely generated extensions of \mathbb{Q} .

1.2 The First Main Result

It is well known that every isomorphic embedding α of a field \mathcal{F} into an algebraically closed \mathcal{K} extends to an embedding β of the algebraic closure of \mathcal{F} into \mathcal{K} . Since we are interested in *effective* embeddings, we ask whether β can always be chosen to be effective. In our notation, with a fixed Rabin embedding ι and an arbitrary computable α , we ask for a computable β such that the following diagram commutes:

$$\begin{array}{ccc} \overline{\mathcal{F}} & \xrightarrow{\beta} & \mathcal{K} \\ \uparrow \iota & \nearrow \alpha & \\ \mathcal{F} & & \end{array}$$

i.e., $\alpha = \beta \circ \iota$. If a computable solution to the above diagram exists for every choice of α and of the computable algebraically closed field \mathcal{K} , then we say that (\mathcal{F}, ι) has the *computable extendability of embeddings property*. Notice, however, that if some Rabin embedding ι of a particular \mathcal{F} has the computable extendability of embeddings property, then so does every other Rabin embedding j of \mathcal{F} (into any computable presentation of $\overline{\mathcal{F}}$): just apply the computable extendability of embeddings property for ι , with j as the α , to get an embedding β_j that extends $j \circ \iota^{-1}$ (and which must be an isomorphism). Then given any other α , the computable extendability of embeddings property for ι yields a β such that $\beta \circ \beta_j^{-1}$ satisfies the computable extendability of embeddings property for j and this α . Therefore, we usually simply say that \mathcal{F} itself has the computable extendability of embeddings property.

The first problem that we address in the paper is the following.

- Find a necessary and sufficient condition for a computable \mathcal{F} to have the computable extendability of embeddings property.

Before we give a necessary and sufficient condition, we discuss a subtlety that would not occur in the classical case. The desired extension β clearly depends on the choice of the Rabin embedding ι . Classically, the dependence on ι is often suppressed, since we can identify \mathcal{F} with its ι -image. However, as noted above, such an identification is generally impossible *effectively*: the membership problem for $\iota(\mathcal{F})$ may be undecidable. To emphasize the dependence on the embedding $\iota: \mathcal{F} \rightarrow \overline{\mathcal{F}}$, we

say that β ι -extends α if it is a solution to the diagram above. Later in the paper we will allow ι to vary, but for now we fix a concrete choice of a Rabin embedding ι .

We may further restrict ourselves and ask for a *uniform procedure*, i.e., a Turing functional, that takes the open diagram of an algebraically closed field \mathcal{K} and an embedding $\alpha: \mathcal{F} \rightarrow \mathcal{K}$ and outputs an embedding of $\overline{\mathcal{F}}$ into \mathcal{K} ι -extending α . For uniform extendability we do not require \mathcal{K} or α to be computable, but we still fix ι . The reader may find it somewhat unexpected that this uniform version is equivalent to the computable extendability of embeddings property.

Theorem 1.1 *Let \mathcal{F} be a computable field together with a computable embedding $\iota: \mathcal{F} \rightarrow \overline{\mathcal{F}}$ of \mathcal{F} into its algebraic closure. Then the following are equivalent:*

- (i) \mathcal{F} has a splitting algorithm;
- (ii) \mathcal{F} has the computable extendability of embeddings property;
- (iii) there exists a Turing functional which, given as its oracle the open diagram of an algebraically closed field \mathcal{K} and an embedding $\alpha: \mathcal{F} \rightarrow \mathcal{K}$, computes an embedding of $\overline{\mathcal{F}}$ into \mathcal{K} ι -extending α .

The property captured by Theorem 1.1 above is also equivalent to an *a priori* weaker uniform extendability condition, namely, the existence of a uniform procedure that takes indices of computable \mathcal{K} and $\alpha: \mathcal{F} \rightarrow \mathcal{K}$ and outputs an index of a computable $\beta: \overline{\mathcal{F}} \rightarrow \mathcal{K}$ extending α . Indeed, this weaker uniform property follows from the uniform extendability condition in Theorem 1.1 and implies the computable extendability property.

In the language of reverse mathematics, Theorem 1.1 would say that in the ω -model *REC* consisting of the computable sets, a field has a unique algebraic closure if and only if that field has a splitting algorithm. Thus, while *RCA*₀ proves that every field with a splitting algorithm has a unique algebraic closure, it is consistent that every other field has more than one algebraic closure. We note that it was already known from work in reverse mathematics (and is easy to see) that in the situation described above there is always a *low* ι -extension of α , and in characteristic zero if \mathcal{F} has a splitting algorithm, then there is a computable extension of α (see [DHS13, Theorem 9] and [FSS83, Theorem 3.3]). In our result we do not restrict ourselves to fields of characteristic 0; the issue that we face in the case of a positive characteristic will be circumvented using purely inseparable extensions (to be defined). We remark that the essential part of our proof of Theorem 1.1 is based on a certain preservation strategy combined with a variation of the Henkin construction; such a combination has not yet been seen in effective algebra.

1.3 The Second Main Result

Another classical result says that every automorphism of a field \mathcal{F} extends to an automorphism of its algebraic closure. In our notation, the diagram

$$\begin{array}{ccc}
 \mathcal{F} & \xrightarrow{\iota} & \overline{\mathcal{F}} \\
 \alpha \downarrow & & \downarrow \beta \\
 \mathcal{F} & \xrightarrow{\iota} & \overline{\mathcal{F}}
 \end{array}$$

always has a solution β such that the diagram commutes, i.e., $\iota \circ \alpha = \beta \circ \iota$. Once again this is dependent on the embedding $\iota: \mathcal{F} \rightarrow \overline{\mathcal{F}}$, and, slightly abusing our terminology, we say that β ι -extends α . We ask when β can be computed effectively. In the setting of automorphisms, it is natural to look at normal algebraic extensions of the prime field (as we will see in Proposition 4.2). In this case, we can apply Theorem 1.1 to fully characterize existence of such ι -extensions in terms of a splitting algorithm; the exact statement will be given in §4 (Corollary 4.1). Although the reader may find Corollary 4.1 interesting in its own right, the above discussed dependence on ι makes it somewhat unsatisfying. Also, as we will discuss in the next subsection, we would like to apply our results to difference fields, and there this dependence on ι is an obstacle. Therefore, in contrast to the situation of the computable extendability of embeddings property above, we would like to allow the embedding ι to vary.

Definition 1.2 We say that a computable field \mathcal{F} has the *computable extendability of automorphisms property* if for every computable automorphism $\alpha: \mathcal{F} \rightarrow \mathcal{F}$ there is a Rabin embedding $\iota: \mathcal{F} \rightarrow \overline{\mathcal{F}}$ and a computable automorphism $\beta: \overline{\mathcal{F}} \rightarrow \overline{\mathcal{F}}$ that ι -extends α .

The second problem we address in the paper is the following.

- Find a necessary and sufficient condition for a computable \mathcal{F} to have the computable extendability of automorphisms property.

As we mentioned above, the computable extendability of automorphisms property is the property which is of interest in constructing embeddings of difference fields into difference closed fields (as we will see in Theorem 1.5). It is not difficult to see that if a normal extension \mathcal{F} of the prime field has a splitting algorithm, then \mathcal{F} has the computable extendability of automorphisms property. Is having a splitting algorithm implied by computable extendability of automorphisms property? Although we do not know if this is true in general (and we conjecture that perhaps not), we give a condition on the Galois group of \mathcal{F} over the prime field—the *non-covering property*—under which the computable extendability of automorphisms property is equivalent to having a splitting algorithm.

Definition 1.3 We say that a group G has the *non-covering property* if for all finite index normal subgroups $M \subsetneq N$ of G and $g \in G$, there is $h \in gN$ such that for all $x \in G$, $x^{-1}hx \notin gM$.

In Lemma 4.5 we will give an equivalent condition in the language of field extensions, using Galois correspondence.

Before we state our second main result, we note that groups with the non-covering property include abelian and simple groups, and the class of profinite groups with the non-covering property is closed under direct products.

Theorem 1.4 Let \mathcal{F} be a computable normal extension of \mathbb{F}_p , for some prime p , such that $\text{Gal}(\mathcal{F}/\mathbb{F}_p)$ has the non-covering property. The following are equivalent:

- \mathcal{F} has a splitting algorithm;
- \mathcal{F} has the computable extendability of automorphisms property;

(iii) \mathcal{F} has the uniform extendability of automorphisms property.

In characteristic $p > 0$, all Galois groups are abelian, and so every Galois group has the non-covering property. Thus, in characteristic $p > 0$ the computable extendability of automorphisms property is equivalent to having a splitting algorithm.

1.4 Applications to Difference Closed Fields

Rabin [Rab60] showed that every computable field can be computably embedded into its computable algebraic closure, and Harrington [Har74] later showed that every computable differential field can be computably embedded into a differential closure. We consider the possibility of such a result for fields with a distinguished automorphism; such structures are called *difference fields* [CH99]. An existential closure of such a structure analogous to an algebraically closed field exists and is called a *difference closed field*. (We note that there is no such a thing as *the* difference closure since there might be no “smallest” difference closed field containing a given difference field. The formal definitions will follow later.) In what follows next, we refer to this hypothetical analogous result as the Rabin–Harrington Theorem.

We note that a difference field (\mathcal{F}, σ) may distinguish a rather boring automorphism σ , e.g., the identity, for which the Rabin–Harrington theorem clearly holds. On the other hand, we will see that there exist computable difference fields that do not embed into any computable difference closed field. Thus, the same field may have two different automorphisms, one witnessing the Rabin–Harrington Theorem, and the other witnessing its failure, and finding a satisfactory characterization in this setting seems rather hopeless (yet the reader may try to find one). On the other hand, we are mostly interested in the *properties of the underlying field* which make the Rabin–Harrington Theorem hold, and we are not that much concerned with the properties of some “pathological” automorphism that may witness the failure of the Rabin–Harrington Theorem. Thus, we arrive at the third main question addressed in the paper.

- For which \mathcal{F} does (\mathcal{F}, σ) satisfy the Rabin–Harrington Theorem for all σ ?

Here of course \mathcal{F} is a computable field and σ ranges over all computable automorphisms of \mathcal{F} . We show in Theorem 5.1 that the Rabin–Harrington Theorem holds for difference fields with underlying field \mathcal{F} if and only if \mathcal{F} has the computable extension of automorphisms property. Using our results on extending automorphisms, namely the second main result of the paper (Theorem 1.4), we can find a large class of difference fields which satisfy the Rabin–Harrington Theorem for any interpretation of the distinguished automorphism.

Theorem 1.5 *Let \mathcal{F} be a computable normal extension of \mathbb{F}_p , for some prime p , such that $\text{Gal}(\mathcal{F}/\mathbb{F}_p)$ has the non-covering property. Then the following are equivalent:*

- \mathcal{F} has a splitting algorithm;
- for any computable σ , (\mathcal{F}, σ) can be computably embedded into a computable difference closed field.

Even without the non-covering property, (i) implies (ii). In particular, this theorem gives a complete answer to the third main question of the paper in the case of a normal extension of \mathbb{F}_p for any $p > 0$. On the other hand, Theorem 1.5 will be used to produce various examples of computable difference fields that cannot be embedded into computable difference closed fields. We conclude that the most naive attempt to generalize the results of Rabin and Harrington fails. On the other hand, if we allow the automorphism to vary, we get a complete characterization for a large class of fields.

1.5 The Non-covering Property

Since our main results refer to the non-covering property of Galois groups, we would like to know more about the class of groups having this property. In Subsection 4.4 we study the class of groups that have the non-covering property, with an emphasis on profinite groups. It is not difficult to see that abelian groups and simple groups have the non-covering property (see Lemma 4.5). However, it takes a lot more effort to prove the following theorem.

Theorem 1.6 Let $\{G_i : i \in I\}$ be a collection of profinite groups, each of which has the non-covering property. Then $\prod_{i \in I} G_i$ has the non-covering property.

The proof of this theorem might be of some independent interest to the reader. It filters through Goursat's lemma [Gou89] (to be stated in the proof of Theorem 1.6). We note that our proof uses profiniteness to reduce the case of arbitrarily many direct factors to just two factors, and the proof of the case of just two factors (Lemma 4.9) does not use profiniteness. We leave open whether one can use profiniteness to simplify our proof of Lemma 4.9. We also note that some groups do not have the non-covering property (to be discussed).

1.6 The Structure of the Paper

We will begin in §2 by giving some background on computable fields and difference fields. In §3 we will consider embeddings into algebraically closed fields and the computable extendability of embeddings property, and prove the first main result, Theorem 1.1. In §4 we will consider automorphisms and the computable extendability of automorphisms property. We begin in §4.1 by considering a strengthening of the computable extendability of automorphisms property. In §4.2, we prove the second main result, Theorem 1.4. In §4.3, we study the class of groups with the non-covering property, and in §4.4 we give some applications of Theorem 1.4. In §5 we consider applications to difference fields and the Rabin–Harrington Theorem. Finally, in §6 we state an open problem on the characterization of fields with the computable extendability of automorphisms property.

2 Preliminaries

2.1 Separable and Purely Inseparable Extensions

If \mathcal{F} is a field, a polynomial $f \in \mathcal{F}[X]$ is called *separable* if it has no repeated roots. An element $a \in \mathcal{E}$ of an algebraic field extension \mathcal{E}/\mathcal{F} is called *separable over \mathcal{F}* if its

minimal polynomial over \mathcal{F} is a separable polynomial. An algebraic field extension \mathcal{E}/\mathcal{F} is called *separable* if every element of \mathcal{E} is separable over \mathcal{F} . Recall that if \mathcal{F} is finite or characteristic zero, then it is *perfect*, i.e., every algebraic extension is a separable extension.

An algebraic field extension \mathcal{E}/\mathcal{F} is called *purely inseparable* if $\mathcal{E} \setminus \mathcal{F}$ contains no separable elements. Equivalently, \mathcal{E} is a field of characteristic $p > 0$ and every element of \mathcal{E} is the unique root of a polynomial $X^{p^n} - a = 0$ with $a \in \mathcal{F}$. Given an algebraic field extension \mathcal{E}/\mathcal{F} , the set $\mathcal{F}^s = \{a \in \mathcal{E} : a \text{ is separable over } \mathcal{F}\}$ is the maximal separable extension of \mathcal{F} inside of \mathcal{E} and is called the *separable closure* of \mathcal{F} in \mathcal{E} . The field extension $\mathcal{E}/\mathcal{F}^s$ is purely inseparable. In the special case where $\mathcal{E} = \overline{\mathcal{F}}$ is the algebraic closure of \mathcal{F} , \mathcal{F}^s is called the *separable closure* of \mathcal{F} and is the maximal separable extension of \mathcal{F} .

An algebraic field extension \mathcal{E}/\mathcal{F} is *normal* if every irreducible polynomial in $\mathcal{F}[X]$ that has a root in \mathcal{E} factors completely in $\mathcal{E}[X]$. A normal separable extension \mathcal{E}/\mathcal{F} is called a Galois extension and has associated with it the Galois group $\text{Gal}(\mathcal{E}/\mathcal{F})$ of automorphisms of \mathcal{E} fixing \mathcal{F} . Recall that the Galois group obeys the fundamental theorem of Galois theory: the normal subgroups $H \trianglelefteq \text{Gal}(\mathcal{E}/\mathcal{F})$ correspond to the intermediate normal field extensions.

2.2 Computable Fields

Recall that the *splitting set* $S_{\mathcal{F}}$ of \mathcal{F} is the set of all polynomials $p \in \mathcal{F}[X]$ which are reducible over \mathcal{F} . The splitting set of a field is not necessarily computable (see [Mil08, Lemma 7]), but it is always computably enumerable (c.e.). If the splitting set of \mathcal{F} is computable, then we say that \mathcal{F} has a *splitting algorithm*. Finite fields and algebraically closed fields trivially have splitting algorithms. Kronecker [Kro82] showed that \mathbb{Q} has a splitting algorithm, and also that many other field extensions also have a splitting algorithm.

Theorem 2.1 (Kronecker [Kro82]; see also [vdW70]) *The field \mathbb{Q} has a splitting algorithm. If a computable field \mathcal{F} has a splitting algorithm, and a is transcendental over \mathcal{F} , then $\mathcal{F}(a)$ has a splitting algorithm. If a is separable and algebraic over \mathcal{F} , then $\mathcal{F}(a)$ has a splitting algorithm. Moreover, the splitting algorithm for $\mathcal{F}(a)$ is uniform in the minimal polynomial for a over \mathcal{F} .*

Given a field \mathcal{F} and an element a that is either transcendental over \mathcal{F} , or separable and algebraic over \mathcal{F} , we know that $\mathcal{F}(a)$ has a splitting algorithm. However, the algorithm depends on whether a is transcendental or algebraic. To find a splitting algorithm uniformly, we must know which is the case.

Rabin [Rab60] showed that every computable field \mathcal{F} has a computable algebraic closure $\overline{\mathcal{F}}$, and moreover, there is a computable embedding $\iota: \mathcal{F} \rightarrow \overline{\mathcal{F}}$. We call such an embedding a *Rabin embedding*. Moreover, he characterized the image of \mathcal{F} under this embedding.

Theorem 2.2 (Rabin [Rab60]) *Let \mathcal{F} be a computable field. Then there is a computable algebraically closed field $\overline{\mathcal{F}}$ and a computable field embedding $\iota: \mathcal{F} \rightarrow \overline{\mathcal{F}}$ such*

that $\overline{\mathcal{F}}$ is algebraic over $\iota(\mathcal{F})$. Moreover, for any such $\overline{\mathcal{F}}$ and ι , the image $\iota(\mathcal{F})$ of \mathcal{F} in $\overline{\mathcal{F}}$ is Turing equivalent to the splitting set of \mathcal{F} .

A computable field \mathcal{F} has a *dependence algorithm* if given a and b_1, \dots, b_n , we can compute whether a is algebraically independent over b_1, \dots, b_n . A field has a dependence algorithm if and only if it has a computable transcendence base (see, for example, [HTMM15, Proposition 2.2]). In particular, fields of finite transcendence degree have a dependence algorithm.

Convention By an extension \mathcal{E}/\mathcal{F} of computable fields, we mean that there is a computable embedding of \mathcal{F} into \mathcal{E} .

2.3 Difference Fields

Difference fields were first studied by Ritt in the 1930s. A good reference on the classical algebraic theory of difference fields is the book by Cohn [Coh65]. A difference field is a field \mathcal{F} together with an embedding $\sigma: \mathcal{F} \rightarrow \mathcal{F}$. If σ is onto, (\mathcal{F}, σ) is called *inversive*. As every difference field has a unique inversive closure up to isomorphism, we lose nothing by assuming that all of our difference fields are inversive.

A difference field (\mathcal{F}, σ) is called a *difference closed field* if it is existentially closed in the language of difference fields. Difference closed fields arose in the model theoretic study of difference fields (see [Mac97, CH99]). \mathcal{F} is difference closed if and only if:

- (i) σ is an automorphism of \mathcal{F} ;
- (ii) \mathcal{F} is algebraically closed;
- (iii) for every variety U , every affine variety $V \subseteq U \times \sigma(U)$ which projects generically onto U and $\sigma(U)$, and every algebraic set $W \not\subseteq V$, there is an \mathcal{F} -rational point $a \in U(\mathcal{F})$ such that $(a, \sigma(a)) \in V \setminus W$.

Condition (iii) may be viewed as saying that certain systems of equations and inequations have solutions in \mathcal{F} . Conditions (i)–(iii) axiomatize the theory *ACFA* of *difference closed fields*. *ACFA* is decidable, and moreover the theories *ACFA_p* of difference closed fields of characteristic p are also decidable for any p , including $p = 0$ [CH99, (1.4)]. *ACFA* is the model companion of the theory of difference fields [CH99, (1.4)] and hence every formula is equivalent, modulo *ACFA*, to an existential formula [CH99, (1.6)]. Thus, we have the following fact.

Fact 2.3 Every computable difference closed field has a computable (full) elementary diagram.

We call a structure with a computable elementary diagram *decidable*; thus every difference closed field is decidable.

3 Extending Embeddings into the Algebraic Closure

We begin by showing that if \mathcal{F} is any computable field with a splitting algorithm, $\iota: \mathcal{F} \rightarrow \overline{\mathcal{F}}$ is a Rabin embedding, and $\alpha: \mathcal{F} \rightarrow \mathcal{K}$ is a computable embedding of \mathcal{F} into an algebraically closed field \mathcal{K} , then there is a computable embedding of $\overline{\mathcal{F}}$ into

\mathcal{K} extending α . In particular, the new results here are in the case of characteristic $p > 0$. The new issue we have to deal with in characteristic $p > 0$ is that Theorem 2.1 fails for non-separable extensions. We begin by finding the separable closure of a field \mathcal{F} within its algebraic closure $\overline{\mathcal{F}}$.

Lemma 3.1 *Let \mathcal{F} be a computable field. Then the separable closure of \mathcal{F} is c.e. If \mathcal{F} has a splitting algorithm, then the separable closure \mathcal{F}^s of \mathcal{F} in $\overline{\mathcal{F}}$ is computable (so that \mathcal{F}^s has a splitting algorithm).*

Proof Embed \mathcal{F} in its algebraic closure $\overline{\mathcal{F}}$. An element $a \in \overline{\mathcal{F}}$ is separable if and only if there is a polynomial $p(X) \in \mathcal{F}[X]$ of degree m with $p(a) = 0$ and with m distinct roots in $\overline{\mathcal{F}}$. Thus the separable closure of \mathcal{F} is c.e. If \mathcal{F} has a splitting algorithm, then given $a \in \overline{\mathcal{F}}$ we can find the minimal polynomial p of a over \mathcal{F} . Then a is separable over \mathcal{F} if and only if p has no repeated roots, which happens if and only if $p'(a) \neq 0$. (Here, $p'(X)$ is the derivative of $p(X)$ with respect to X , treating the coefficients as constants.) So the separable closure of \mathcal{F} is computable. ■

We are now ready to extend an embedding from a field with a splitting algorithm. The main idea is to break the embedding into two steps: first to extend an embedding $\alpha: \mathcal{F} \rightarrow \mathcal{K}$ to an embedding $\beta: \mathcal{F}^s \rightarrow \mathcal{K}$ of the separable closure of \mathcal{F} into \mathcal{K} , and second to note that β extends to a unique embedding of $\overline{\mathcal{F}}$ into \mathcal{K} and that this extension is computable from β .

Theorem 3.2 *Let \mathcal{F} be a computable field and $\iota: \mathcal{F} \rightarrow \overline{\mathcal{F}}$ a Rabin embedding of \mathcal{F} into its algebraic closure. Suppose that \mathcal{F} has a splitting algorithm. Then there is a Turing functional Φ such that whenever $\alpha: \mathcal{F} \rightarrow \mathcal{K}$ is an embedding of \mathcal{F} into an algebraically closed field \mathcal{K} , $\Phi^{\alpha \oplus \mathcal{K}}: \overline{\mathcal{F}} \rightarrow \mathcal{K}$ is an embedding of $\overline{\mathcal{F}}$ into \mathcal{K} ι -extending α .*

Proof Since \mathcal{F} has a splitting algorithm, the image $\iota(\mathcal{F})$ of \mathcal{F} in $\overline{\mathcal{F}}$ is computable. We may identify \mathcal{F} with its image. By Lemma 3.1 the separable closure \mathcal{F}^s of \mathcal{F} is computable as a subset of $\overline{\mathcal{F}}$ and has a splitting algorithm.

Let \mathcal{K} be an algebraically closed field and $\alpha: \mathcal{F} \rightarrow \mathcal{K}$ a field embedding. We will begin by describing a procedure to extend α to an embedding $\beta: \mathcal{F}^s \rightarrow \mathcal{K}$. Let $\{a_1, a_2, \dots\}$ be an enumeration of the elements \mathcal{F}^s . Start with β defined only on \mathcal{F} and ι -extending α . Using the splitting algorithm for \mathcal{F} , find the minimal polynomial $P_1 \in \mathcal{F}[X]$ of a_1 over \mathcal{F} . Find a solution $b_1 \in \mathcal{K}$ to $\alpha(P_1)$. Then define β on $\mathcal{F}(a_1)$ by mapping a_1 to b_1 . Since a_1 is algebraic and separable over \mathcal{F} (and we know its minimal polynomial), we have a splitting algorithm for $\mathcal{F}(a_1)$. The separable closure of $\mathcal{F}(a_1)$ is \mathcal{F}^s . Now find the minimal polynomial $P_2 \in \mathcal{F}[X]$ of a_2 over $\mathcal{F}(a_1)$, and a solution b_2 to $\alpha(P_2)$. Define β on $\mathcal{F}(a_1, a_2)$ by mapping a_2 to b_2 . Note that a_2 is separable over $\mathcal{F}(a_1)$ since

$$\mathcal{F} \subseteq \mathcal{F}(a_1) \subseteq \mathcal{F}(a_1, a_2) \subseteq \mathcal{F}^s$$

and \mathcal{F}^s is a separable algebraic extension of \mathcal{F} . Since a_2 is algebraic and separable over $\mathcal{F}(a_1)$, we have a splitting algorithm for $\mathcal{F}(a_1, a_2)$. Its separable closure is still \mathcal{F}^s . Continuing in this way, we define an embedding $\beta: \mathcal{F}^s \rightarrow \mathcal{K}$ which ι -extends $\alpha: \mathcal{F} \rightarrow \mathcal{K}$.

In characteristic zero, we are done since $\mathcal{F}^s = \overline{\mathcal{F}}$. In characteristic $p > 0$, we can extend β to an embedding $\overline{\mathcal{F}} \rightarrow \mathcal{K}$ in the following manner. Given $b \in \overline{\mathcal{F}}$, find the minimal polynomial $P \in \mathcal{F}^s[X]$ of b over \mathcal{F}^s (recalling that \mathcal{F}^s has a splitting algorithm). Then $P(X)$ is of the form $X^{p^n} - r = 0$ with $r \in \mathcal{F}$. Note that b is the unique solution of $p(X) = 0$, and we can find the unique solution c to $\beta(p)(X) = 0$. Map b to c . This is the unique embedding of $\overline{\mathcal{F}}$ into \mathcal{K} extending β .

The construction was uniform in α and \mathcal{K} , and so we get the desired Turing functional Φ . ■

We are now ready to prove Theorem 1.1, which says that a field \mathcal{F} has a splitting algorithm if and only if it has the computable (or uniform) extendability of embeddings property.

Proof of Theorem 1.1 The implication (i) \Rightarrow (ii) is Theorem 3.2. The implication (ii) \Rightarrow (iii) is immediate. It remains to show the implication (iii) \Rightarrow (i).

Fix $\iota: \mathcal{F} \rightarrow \overline{\mathcal{F}}$, a computable embedding of \mathcal{F} into a computable presentation $\overline{\mathcal{F}}$ of its algebraic closure. Suppose that every computable embedding of \mathcal{F} into a computable algebraically closed field \mathcal{K} ι -extends to a computable embedding of $\overline{\mathcal{F}}$ into \mathcal{K} .

We will attempt to construct a computable field \mathcal{K} and a computable embedding $\alpha: \mathcal{F} \rightarrow \mathcal{K}$ while attempting to diagonalize against all potential computable extensions $\varphi_e: \overline{\mathcal{F}} \rightarrow \mathcal{K}$ (by having $\alpha(a) \neq \varphi_e(\iota(a))$ for some $a \in \mathcal{F}$). We know that the construction must fail, and from this we will conclude that \mathcal{F} has a splitting algorithm.

We construct \mathcal{K} by an effective Henkin-style construction. The Henkin construction will be similar to one that can be used to prove Rabin's theorem that every field embeds into a computable presentation of its algebraic closure. See, for example, [FSS83, Theorem 2.5], where this construction is carried out in reverse mathematics. (Rabin's original proof constructed the algebraic closure using a quotient of a polynomial ring with infinitely many variables.) Let $\mathcal{L}_{\mathcal{F}}$ be the language of fields with constant symbols for the elements of \mathcal{F} , and let T be the consistent theory of algebraically closed fields together with the atomic diagram of \mathcal{F} . By quantifier elimination for the theory of algebraically closed fields, T is a complete theory and hence is decidable. We want to construct a decidable prime model of the theory T , which gives an algebraic closure \mathcal{K} of \mathcal{F} together with an embedding of \mathcal{F} into \mathcal{K} . The embedding $\alpha: \mathcal{F} \rightarrow \mathcal{K}$ will be built as part of the Henkin construction. Constructing a prime model requires a slight modification of the Henkin construction, which is possible in this case—we must also omit the type of an element that is transcendental over \mathcal{F} (see [Mil83] for the general theorem on effectively omitting types).

Let $C = \{c_0, c_1, \dots\}$ be the new constant symbols for the Henkin construction. The domain of \mathcal{K} will be the equivalence classes of some computable equivalence relation on C . Let $\varphi_e: \overline{\mathcal{F}} \rightarrow C$ be a list of partial computable functions which we interpret as the possible computable embeddings $\overline{\mathcal{F}} \rightarrow \mathcal{K}$. Let $\{a_0, a_1, a_2, \dots\}$ be a computable enumeration of \mathcal{F} . We use a_i to denote the constant symbol associated with $a_i \in \mathcal{F}$.

Construction. At each stage s , we define formulas $\delta_0, \dots, \delta_s$ in the language $\mathcal{L}_{\mathcal{F} \cup C}$ which form the partial diagram of \mathcal{K} at stage s . The theory $\Delta = \{\delta_0, \delta_1, \dots\}$ will be a complete theory extending T which is the complete diagram of the model \mathcal{K} (with the domain of \mathcal{K} being the equivalence classes in C by the equivalence relation

$c \sim d \Leftrightarrow \Delta \vdash c = d$). At stage s , let $\psi_s = \delta_0 \wedge \dots \wedge \delta_{s-1}$. We can arrange the construction so that the only constant symbols from \mathcal{F} that appear in δ_s are $\underline{a}_0, \dots, \underline{a}_s$. At stage 0, let δ_0 be $c_0 = c_0$.

At stage $s = 4t + 1$, we try to diagonalize against a φ_e for $e \leq t$. Search for an $e \leq t$ and an $i < s + 5$ such that $\varphi_{e,t}(t(a_i)) = c_j$ and (where $\bar{c} = (c_0, c_1, \dots)$) is the sequence of constants from C that appear in ψ_s):

$$T \not\vdash \forall \bar{x} (\psi_s[\bar{x}/\bar{c}] \Rightarrow \underline{a}_i = x_j).$$

By $\psi_s[\bar{x}/\bar{c}]$, we mean that the variables $\bar{x} = (x_0, x_1, \dots)$ have been substituted for the constants $\bar{c} = (c_0, c_1, \dots)$. This is a bounded search since T is decidable and we only have to search through finitely many \underline{a}_i . If such an e exists, choose the least e such that we have not yet diagonalized against φ_e . Then set δ_s to be the formula $\underline{a}_i \neq c_j$ for that e . If no such e exists, set δ_s to be the formula $c_0 = c_0$.

At stages $s = 4t + 2$, $s = 4t + 3$, and $s = 4t + 4$, we act as in the standard method of constructing a computable prime model (e.g., [Har98, Theorems 5.1, 5.7]), as follows.

At stage $s = 4t + 2$, we add a Henkin witness for δ_t . If δ_t is of the form $(\exists x)\varphi(x)$, then let c_i be a constant which does not appear in ψ_s and let δ_s be $\varphi(c_i)$. Otherwise, set δ_s to be the formula $c_0 = c_0$.

At stage $s = 4t + 3$, we satisfy the completeness requirement for the sentence χ_t from some fixed listing $(\chi_t)_{t \in \omega}$ of the sentences in the language $\mathcal{L}_{\mathcal{F} \cup C}$. Let \bar{c} be the constants from C which appear in ψ_s and χ_t . Check whether

$$T \vdash \forall \bar{x} (\psi_s \Rightarrow \chi_t)[\bar{x}/\bar{c}].$$

If this is the case, let δ_s be χ_t . Otherwise, let δ_s be $\neg\chi_t$.

At stage $s = 4t + 4$, we omit the type of an element transcendental over \mathcal{F} . We will have c_t satisfy some polynomial over \mathcal{F} . Let \bar{c} be the constants from C which appear in ψ_s , except for c_t . Search for a polynomial $p(x) \in \mathcal{F}[\bar{X}]$ such that

$$T \not\vdash \forall x \forall \bar{z} (\psi_s[x\bar{z}/c_t\bar{c}] \Rightarrow p(x) \neq 0).$$

Set δ_s to be the formula $p(c_t) = 0$. Some such polynomial p must exist as the type of a transcendental over \mathcal{F} is a non-principal type.

Verification. By the standard Henkin construction arguments, we get a decidable prime model \mathcal{K} whose domain consists of equivalence classes from C . We get a computable embedding α of \mathcal{F} into \mathcal{K} by mapping $a \in F$ to the element of \mathcal{K} labeled by the symbol \underline{a} . Then α extends to an embedding β of $\bar{\mathcal{F}}$ into \mathcal{K} , which we may represent as a computable map $\varphi_e: \bar{\mathcal{F}} \rightarrow C$ (by, say, choosing $\varphi_e(a)$ to be the least element of C in the equivalence class of $\beta(a)$, which we can do computably since the equivalence classes are computable). There is a stage s_0 after which we never diagonalize against an $e' < e$. We never diagonalize against e .

Claim Let $b \in \bar{\mathcal{F}}$ and let t be a stage such that $\varphi_{e,t}(b) \downarrow = c_j$ for some $j \in \omega$. Let $s = 4t + 1$. Then $b \in \iota(\mathcal{F})$ if and only if there is some i such that

$$(*) \quad T \vdash \forall \bar{x} (\psi_s[\bar{x}/\bar{c}] \Rightarrow \underline{a}_i = x_j).$$

Proof Given $(*)$, in \mathcal{K} the constant symbol a_i is interpreted as the equivalence class of c_j . Thus α maps a_i to the equivalence class of c_j . Since β extends α and is one-to-one, $\iota(a_i) = b$.

On the other hand, suppose that $b \in \iota(F)$, say $b = a_i$, and suppose to the contrary that $(*)$ does not hold. We have two cases. First, if $i < s + 5$, then we set δ_s to be the formula $a_i \neq c_j$. Then $\alpha(a_i) \neq c_j = \varphi_e(\iota(a_i))$, which is a contradiction. Second, if $i \geq s + 5$, then let $s' > s$ be the first stage of the form $s' = 4t' + 1$ with $i < s' + 5$. We have $i > s'$ (as if $i \leq s'$, we could have chosen $s' - 4$). Since the only constant symbols from \mathcal{F} that appear in $\psi_{s'}$ are $\underline{a}_0, \dots, \underline{a}_{s'}$, and $i > s'$, \underline{a}_i does not appear in $\psi_{s'}$. Then we have

$$T \not\vdash \forall \bar{x} (\psi_{s'}[\bar{x}/\bar{c}] \Rightarrow \underline{a}_i = x_j).$$

We set $\delta_{s'}$ to be the formula $\underline{a}_i \neq c_j$ which again yields a contradiction. Hence $(*)$ holds. ■

The claim gives us a decision procedure for $\iota(F) \subseteq \overline{\mathcal{F}}$. At any stage s , there are only finitely many constants $c \in C$ mentioned in ψ_s , and hence only finitely many a_i such that we might possibly have $(*)$. So given $b \in \overline{\mathcal{F}}$, compute $s = 4t + 1 \geq s_0$ and j such that $\varphi_{e,t}(b) \downarrow = c_j$, and then check $(*)$ for the finitely many possible a_i to decide whether $b \in \iota(\mathcal{F})$. ■


It was important in Theorem 1.1 that we allow the field \mathcal{K} to vary. This is because if \mathcal{F} is a field of infinite transcendence degree, there may be computable algebraically closed fields of infinite transcendence degree into which \mathcal{F} does not effectively embed. For example, if \mathcal{F} does not have a dependence algorithm but \mathcal{K} does, then there is no computable embedding of \mathcal{F} into \mathcal{K} . If we restrict to the case where \mathcal{F} is an algebraic field, then \mathcal{F} has a computable embedding into every computable algebraically closed field \mathcal{K} . In this particular case we get the following corollary, which we use in §4, where the field \mathcal{K} is fixed.

Corollary 3.3 *Let \mathcal{F} be a computable algebraic field and $\iota: \mathcal{F} \rightarrow \overline{\mathcal{F}}$ a computable embedding of \mathcal{F} into a computable presentation of its algebraic closure. Let \mathcal{K} be a computable algebraically closed field. Then the following are equivalent.*

- (i) \mathcal{F} has a splitting algorithm.
- (ii) There is a Turing functional Φ which takes an embedding $\alpha: \mathcal{F} \rightarrow \mathcal{K}$ to an embedding Φ^α of $\overline{\mathcal{F}}$ into \mathcal{K} extending α .
- (iii) Every computable embedding of \mathcal{F} into \mathcal{K} ι -extends to a computable embedding of $\overline{\mathcal{F}}$ into \mathcal{K} .

Proof By Theorem 1.1, it suffices to show that (iii) in the statement implies that \mathcal{F} has the computable extendability property with respect to $\iota: \mathcal{F} \rightarrow \overline{\mathcal{F}}$. Let $\alpha: \mathcal{F} \rightarrow \mathcal{L}$ be a computable embedding of \mathcal{F} into a computable algebraically closed field \mathcal{L} . We can enumerate, in \mathcal{L} , the algebraic closure of the prime field and this contains the image $\alpha(\mathcal{F})$ of \mathcal{F} . So we may assume that \mathcal{L} is the algebraic closure of its prime field.

We can compute an embedding $j: \mathcal{L} \rightarrow \mathcal{K}$ and let $\alpha^*: \mathcal{F} \rightarrow \mathcal{K}$ be $j \circ \alpha$. By (iii), there is an embedding $\beta^*: \overline{\mathcal{F}} \rightarrow \mathcal{K}$ which ι -extends β .



Since $\overline{\mathcal{F}}$ and $\overline{\mathcal{L}}$ are both algebraic closures of the prime field, the image of β^* is the same as the image of j . So there is an embedding $\beta: \overline{\mathcal{F}} \rightarrow \overline{\mathcal{L}}$ such that $j \circ \beta = \beta^*$. Then β^* ι -extends α . ■

4 Extending Automorphisms of Normal Extensions of the Prime Field

4.1 Strong Extendability of Automorphisms Property

In the setting of automorphisms, it is natural to look at normal algebraic extensions of the prime field (see Proposition 4.2). When \mathcal{F} is such an extension, we get the following corollary of Theorem 1.1, with two strengthenings of the computable extendability of automorphisms property. (We denote the prime field by \mathbb{F}_p even in the case of characteristic $p = 0$.)

Corollary 4.1 *Let \mathcal{F} be a computable normal algebraic extension of the prime field and $\iota: \mathcal{F} \rightarrow \overline{\mathcal{F}}$ an embedding of \mathcal{F} into a computable presentation of its algebraic closure. The following are equivalent.*

- (i) \mathcal{F} has a splitting algorithm.
- (ii) For every computable automorphism $\alpha: \mathcal{F} \rightarrow \mathcal{F}$ of \mathcal{F} , there is a computable automorphism $\beta: \overline{\mathcal{F}} \rightarrow \overline{\mathcal{F}}$ which ι -extends α .
- (iii) There is a uniform procedure which, given any computable automorphism $\alpha: \mathcal{F} \rightarrow \mathcal{F}$ of \mathcal{F} , outputs a computable automorphism $\beta: \overline{\mathcal{F}} \rightarrow \overline{\mathcal{F}}$ which ι -extends α .

Proof of Corollary 4.1 Suppose that \mathcal{F} is a computable normal algebraic field, and $\iota: \mathcal{F} \rightarrow \overline{\mathcal{F}}$ is a Rabin embedding. If \mathcal{F} has a splitting algorithm, then by Corollary 3.3, any automorphism α of \mathcal{F} extends to an automorphism of $\overline{\mathcal{F}}$ (taking $\mathcal{K} = \overline{\mathcal{F}}$ in the statement of the corollary). Indeed, $\iota \circ \alpha$ is a computable embedding of \mathcal{F} into $\overline{\mathcal{F}}$ and hence there is an automorphism β of $\overline{\mathcal{F}}$ which ι -extends $\iota \circ \alpha$; that is, $\beta \circ \iota = \iota \circ \alpha$. So β ι -extends α .

On the other hand, suppose that every automorphism of \mathcal{F} extends to an automorphism of $\overline{\mathcal{F}}$. We will check (iii) of Corollary 3.3 with $\mathcal{K} = \overline{\mathcal{F}}$. Let $\alpha: \mathcal{F} \rightarrow \overline{\mathcal{F}}$ be an

embedding. Since \mathcal{F} is normal, $\alpha(\mathcal{F}) = \iota(\mathcal{F})$. Then $\iota^{-1} \circ \alpha: \mathcal{F} \rightarrow \mathcal{F}$ is an automorphism of \mathcal{F} , and hence extends to an automorphism β of $\overline{\mathcal{F}}$. We have the following diagram:

$$\begin{array}{ccc}
 \overline{\mathcal{F}} & \xrightarrow{\beta} & \overline{\mathcal{F}} \\
 \uparrow \varepsilon & & \uparrow \varepsilon \\
 \iota(\mathcal{F}) & \xrightarrow{\alpha \circ \iota^{-1}} & \iota(\mathcal{F}) \\
 \uparrow \iota & \nearrow \alpha & \uparrow \iota \\
 \mathcal{F} & \xrightarrow{\iota^{-1} \circ \alpha} & \mathcal{F}
 \end{array}$$

Note that $\beta: \overline{\mathcal{F}} \rightarrow \overline{\mathcal{F}}$ ι -extends the embedding α of \mathcal{F} into $\overline{\mathcal{F}}$. ■

Note that we had to use Corollary 3.3 rather than Theorem 1.1, because we needed to fix $K = \overline{\mathcal{F}}$ instead of letting \mathcal{K} be arbitrary.

In Corollary 4.1 we asked for \mathcal{F} to be a normal extension of \mathbb{F}_p . This is required in order to prove the theorem; we will construct an algebraic field which demonstrates that we need the hypothesis of normality in the preceding results. A rigid field automatically satisfies (ii) of Corollary 4.1.

Proposition 4.2 *There is a rigid computable algebraic field \mathcal{F} of characteristic zero with no splitting algorithm.*

Proof Let p_0, p_1, \dots list the primes greater than two. Let $\mathcal{F} = \mathbb{Q}(a_n : n \in \emptyset')$ where a_n is the unique real p_n -th root of 2, and \emptyset' is the Turing jump of the empty set. Since $\mathcal{F} \subseteq \mathbb{R}$, for each $n \in \emptyset'$, a_n is the only p_n -th root of 2 in \mathcal{F} . So every automorphism of \mathcal{F} fixes the a_n , and hence fixes \mathcal{F} . Hence \mathcal{F} is rigid.

We can use an enumeration of \emptyset' to give a computable presentation of \mathcal{F} : \mathcal{F} can be embedded as a c.e. subfield of $\overline{\mathbb{Q}}$ and from this we get a computable presentation of \mathcal{F} .

We need to argue that for $n \notin \emptyset'$, $a_n \notin \mathcal{F}$. We claim that if $n \notin I$, $a_n \notin \mathbb{Q}(a_i : i \in I)$. Suppose not; then we can find a finite set I and $n \notin I$ such that $a_n \in \mathbb{Q}(a_i : i \in I)$ and for each $j \in I$, $a_j \notin \mathbb{Q}(a_i : i \in I \setminus \{j\})$. Then $\mathbb{Q}(a_i : i \in I)$ is a finite extension of \mathbb{Q} of degree $d = \prod_{i \in I} p_i$. Since p_n does not divide d , $\mathbb{Q}(a_n)$ is not a subfield of $\mathbb{Q}(a_i : i \in I)$. This contradicts the assumption that $a_n \in \mathbb{Q}(a_i : i \in I)$. Thus for $n \notin \emptyset'$, $a_n \notin \mathcal{F}$.

No computable presentation of \mathcal{F} can have a splitting algorithm, as a splitting algorithm would allow us to compute \emptyset' . ■

4.2 Computable Extendability of Automorphisms Property

In Corollary 4.1, we fixed an embedding $\iota: \mathcal{F} \rightarrow \overline{\mathcal{F}}$ and considered only ι -extensions. Now we allow ι to vary. Note that while every computable presentation of $\overline{\mathcal{F}}$ is isomorphic, there may be different computable embeddings $\iota: \mathcal{F} \rightarrow \overline{\mathcal{F}}$ which are not

equivalent up to a computable automorphism of $\overline{\mathcal{F}}$. By Corollary 3.3, if \mathcal{F} is an algebraic field with no splitting algorithm, there are embeddings ι and j of \mathcal{F} into $\overline{\mathcal{F}}$ such that there is no computable automorphism σ of $\overline{\mathcal{F}}$ with $\sigma \circ \iota = j$.

In the introduction, we said that \mathcal{F} had the *computable extendability of automorphisms property* if each computable automorphism of \mathcal{F} had such an extension to $\overline{\mathcal{F}}$. Recall that our interest in the computable extendability of automorphisms property comes from its role in an analogue of Rabin's Theorem in the context of difference closed fields; see Theorem 5.1 which we will prove in the following section.

We can already produce examples of fields without the computable extendability of automorphisms property. We use the fact that every noncomputable c.e. set is the union of two disjoint, computably inseparable c.e. subsets. This is a theorem of Yates, who saw that it followed from a construction of Friedberg; the theorem was subsequently published by Cleave [Cle70].

Proposition 4.3 *For each noncomputable c.e. set C , the field $\mathcal{F} = \mathbb{Q}(\sqrt{p_n} : n \in C)$ (with p_n the n -th prime) does not have the computable extendability of automorphisms property.*

Proof Let A and B be disjoint computably inseparable c.e. sets with $A \cup B = C$. Recalling the classic result (originally due to Besicovitch [Bes40]) that if r and q_1, \dots, q_ℓ are distinct primes, then $\sqrt{r} \notin \mathbb{Q}(\sqrt{q_1}, \dots, \sqrt{q_\ell})$, we define an automorphism α of \mathcal{F} by letting α fix the two square roots of p_n if $n \in A$, but interchange them if $n \in B$. This α is computable, but if ι is any Rabin embedding of \mathcal{F} into some presentation $\overline{\mathcal{F}}$ of its algebraic closure and β is an automorphism of $\overline{\mathcal{F}}$ ι -extending α , then $\{n \in \omega : \beta(\sqrt{p_n}) = \sqrt{p_n}\}$ is a β -computable separation of A from B . ■

However, we would like a more complete description of which fields have, and which do not have, the computable extendability of automorphisms property. We do not obtain a complete description, but we give a characterization in terms of a splitting algorithm for many fields. The idea will be to isolate certain normal extensions \mathcal{K}/\mathbb{F}_p whose subfield structure behaves sufficiently like the field \mathcal{F} in Proposition 4.3 above, allowing us to make a particular diagonalization argument. In diagonalizing against the computable extendability of automorphisms property, we do not have access to a Rabin embedding ι (and it does not seem possible to diagonalize against all possible computable Rabin embeddings). So rather than defining α to diagonalize against β using the image under a fixed ι , we must define α to diagonalize against *all possible* images under all possible ι . In Proposition 4.3, we do this using the fact that if α fixes the square roots of p_n , then so does β for any ι -extension of α under any Rabin embedding ι , and similarly if α interchanges the roots of p_n . In general, we want to have some finite subfield \mathcal{E} of \mathcal{F} , and to define α on \mathcal{E} so that there is no embedding ι under which β ι -extends α . We may have already defined α on some subfield of \mathcal{E} , so we do not have a completely free choice of α . There are some fields where this argument will always work successfully: those with the non-covering property from Definition 1.3. In many other fields, we can find some appropriate subfield which satisfies the required condition, allowing the argument to go through; see Examples 4.11 and 4.13.

Using Galois theory, there is also a field-theoretic characterization of the field extensions whose Galois group has the non-covering property, and it is this characterization that we will use in the proof of Theorem 4.6 (though, in applying the theorem, it will usually be easier to use the group-theoretic characterization). In what follows, it will be helpful to use the language of difference fields to talk about field automorphisms.

Remark 4.4. Let \mathcal{F}/\mathcal{E} be a field extension, α an automorphism of \mathcal{E} , and β an automorphism of \mathcal{F} . Let $\iota: \mathcal{E} \rightarrow \mathcal{F}$ be a field embedding of \mathcal{E} into \mathcal{F} . Then $\beta \iota$ -extends α if and only if ι is an embedding of (\mathcal{E}, α) into (\mathcal{F}, β) as difference fields.

Proof Both are equivalent to having $\beta \circ \iota = \iota \circ \alpha$. ■

Lemma 4.5 Let \mathcal{E}/\mathcal{F} be a separable normal extension. The following are equivalent.

- (i) $\text{Gal}(\mathcal{E}/\mathcal{F})$ has the non-covering property.
- (ii) For all finite normal subextensions $\mathcal{K}_1/\mathcal{F}$ and $\mathcal{K}_2/\mathcal{F}$ with $\mathcal{K}_2 \not\subseteq \mathcal{K}_1$, and every pair of automorphisms σ of \mathcal{K}_1 and τ of \mathcal{K}_2 fixing \mathcal{F} , there is an automorphism α of \mathcal{E} extending σ and incompatible with τ , i.e., (\mathcal{K}_2, τ) does not embed into (\mathcal{E}, α) as a difference field.

The second point is related to the monadic and incompatible extensions of difference fields studied by Cohn [Coh52], Babbitt [Bab62], and Evanovich [Eva73].

Proof We begin by showing (i) \Rightarrow (ii). Let \mathcal{K}_1 and \mathcal{K} be as in (ii). Let σ and τ be automorphisms of \mathcal{K}_1 and \mathcal{K}_2 , respectively, fixing \mathcal{F} . Let $G = \text{Gal}(\mathcal{E}/\mathcal{F})$. Let M be the normal subgroup of automorphisms fixing \mathcal{K}_2 , and N the normal subgroup of automorphisms fixing \mathcal{K}_1 . Since \mathcal{K}_1 and \mathcal{K}_2 are finite extensions, M and N are of finite index. We also have $N \not\subseteq M$. Let $g_1 \in G$ be an automorphism of \mathcal{E} extending σ , and g_2 an automorphism of \mathcal{E} extending τ .

We will argue that there is an $h \in g_1N$ such that for all $z \in G$, $z^{-1}hz \notin g_2M$. Such an h is an automorphism of \mathcal{E} extending σ , and g_2M is the set of automorphisms of \mathcal{E} extending τ . Since for all $x \in G$, $x^{-1}g_1hx \notin g_2M$, (\mathcal{E}, α) is not isomorphic as a difference field to (\mathcal{E}, β) for any extension β of τ .

First, we argue that it suffices to assume $M \not\subseteq N$. Suppose that there is $h' \in g_1NM$ such that for all $z \in G$, $z^{-1}hz \notin g_2M$. Then write $h' = g_1nm$. Suppose for some $z \in G$ that $z^{-1}g_1nz \in g_2M$, say $z^{-1}g_1nz = g_2m'$ with $m' \in M$. Let $m'' \in M$ be such that $z^{-1}mz = m''$. Then $z^{-1}g_1nmz = g_2m'm''^{-1} \in g_2M$. This contradicts the choice of $h' = g_1nm$. So for all $z \in G$, $z^{-1}g_1nz \notin g_2M$. Then $h = g_1n \in g_1N$ is the automorphism of \mathcal{E} that we desire. So we may replace N by NM .

Now we have two cases. First, suppose that there is no $z \in G$ such that $z^{-1}g_1z \in g_2M$. Then $h = g_1$ is as desired.

Second, suppose that for some $c \in M$ and $z \in G$, $z^{-1}g_1z = g_2c$. Then $g_1 = zg_2cz^{-1} = zg_2z^{-1}c'$ for some other $c' \in M$ since M is a normal subgroup. So $zg_2z^{-1} = g_1m$, where $m = (c')^{-1}$. Using the fact that G has the non-covering property, choose $h \in N$ such that for all $x \in G$, $x^{-1}g_1hx \notin g_1M$. We claim that for all $x \in G$, $x^{-1}g_1hx \notin g_2M$. Suppose to the contrary that there is $x \in G$ such that $x^{-1}g_1hx \in g_2M$. Since $x^{-1}g_1hx \in$

g_2M and M is a normal subgroup, $g_1h \in xg_2x^{-1}M$. We have

$$xg_2x^{-1} = (xz^{-1})zg_2z^{-1}(xz^{-1})^{-1} = (xz^{-1})g_1m(xz^{-1})^{-1}.$$

Let $y = (xz^{-1})$. Since $m \in M$ is a normal subgroup, $yg_1my^{-1} = yg_1y^{-1}m'$ for some other $m' \in M$. Thus $g_1h \in yg_1y^{-1}M$ and so $y^{-1}g_1hy \in g_1M$. This contradicts the choice of h . So for all $x \in G$, $x^{-1}g_1hx \notin g_2M$.

The direction (ii) \Rightarrow (i) proceeds simply by the Galois correspondence. Fix finite index normal subgroups $M \not\subseteq N$ of $G = \text{Gal}(\mathcal{E}/\mathcal{F})$ and $g \in G$. Let \mathcal{K}_1 and \mathcal{K}_2 be the fields fixed by N and M , respectively; we have $\mathcal{K}_1 \not\subseteq \mathcal{K}_2$. Let σ be the restriction of g to \mathcal{K}_1 and τ its restriction to \mathcal{K}_2 . There is an automorphism α of \mathcal{E} extending σ and not compatible with τ . So $\alpha \in gN$ and for all extensions β of τ to \mathcal{E} i.e., $\beta \in gM$, (\mathcal{E}, α) is not isomorphic as a difference field to (\mathcal{E}, β) . That is, for all $\gamma \in G$ and $\beta \in gM$, $\gamma \circ \alpha \neq \beta \circ \gamma$. ■

We will restrict our attention to field extensions whose Galois group has the non-covering property, but we will allow the base field to be an extension of \mathbb{F}_p with a splitting algorithm. The main theorem of this section is as follows. (We state it in a slightly more general form than it appears in the introduction.)

Theorem 4.6 *Let \mathcal{E} be a computable normal extension of \mathbb{F}_p and let $\mathcal{F} \subseteq \mathcal{E}$ be a subfield of \mathcal{E} with a splitting algorithm which is also a normal extension of \mathbb{F}_p . Suppose that $\text{Gal}(\mathcal{E}/\mathcal{F})$ has the non-covering property. The following are equivalent.*

- (i) \mathcal{E} has a splitting algorithm.
- (ii) \mathcal{E} has the computable extendability of automorphisms property.
- (iii) \mathcal{E} has the uniform extendability of automorphisms property.

Many applications of this theorem will have $\mathcal{F} = \mathbb{F}_p$, but the freedom to choose \mathcal{F} will allow us to apply the theorem in situations where $\text{Gal}(\mathcal{E}/\mathbb{F}_p)$ does not have the non-covering property. Producing an example where the theorem cannot be applied seems to be a non-trivial task, and we do not know of any such examples. See §4.4 for some applications of the theorem.

Proof of Theorem 4.6 We already know that the implications (i) \Rightarrow (ii) and (i) \Rightarrow (iii) are true, even given a fixed embedding of \mathcal{E} into $\bar{\mathcal{E}}$. (iii) clearly implies (ii). We must show (ii) \Rightarrow (i).

Suppose that every computable automorphism of \mathcal{E} extends to a computable automorphism of $\bar{\mathcal{E}}$ (via some embedding of \mathcal{E} into $\bar{\mathcal{E}}$). We will attempt to construct a computable automorphism $\alpha \in \text{Gal}(\mathcal{E}/\mathcal{F})$ while diagonalizing against possible computable automorphisms $\varphi_e: \bar{\mathcal{E}} \rightarrow \bar{\mathcal{E}}$ by making sure that the difference field (\mathcal{E}, α) does not embed into the difference field $(\bar{\mathcal{E}}, \varphi_e)$. It suffices to ensure that some difference subfield of (\mathcal{E}, α) does not embed into $(\bar{\mathcal{E}}, \varphi_e)$. We know that the construction must fail, and from this we will conclude that \mathcal{E} has a splitting algorithm.

Note that the field \mathcal{F} has a splitting algorithm and is perfect (since it is an algebraic extension of a perfect field), so any finite algebraic extension of \mathcal{F} has a splitting algorithm which we can determine effectively from a generating set for the extension.

We will require a special enumeration $\{a_1, a_2, \dots\}$ of \mathcal{E} with the following properties:

- (i) For each n , $\mathcal{F}(a_1, \dots, a_n)$ is a normal extension of \mathbb{F}_p .
- (ii) For each n , there are no normal extensions of \mathbb{F}_p which are strictly contained between $\mathcal{F}(a_1, \dots, a_n)$ and $\mathcal{F}(a_1, \dots, a_n, a_{n+1})$.

We can find such an enumeration using the primitive element theorem and Galois theory, as follows. Suppose that we have already defined a_1, \dots, a_n . Given a new element x of \mathcal{E} , first check whether $x \in \mathcal{F}(a_1, \dots, a_n)$ using the splitting algorithm for this field. If x is in $\mathcal{F}(a_1, \dots, a_n)$, we can safely set $a_{n+1} = x$. Otherwise, compute the conjugates $x = x_1, \dots, x_\ell$ of x over \mathbb{F}_p . Search for a single element y such that

$$\mathcal{F}(y) \in \mathcal{F}(a_1, \dots, a_n, x_1, \dots, x_\ell).$$

Such an element exists by the primitive element theorem as $\mathcal{F}(a_1, \dots, a_n, x_1, \dots, x_\ell)$ is a finite separable extension of \mathcal{F} . Since each automorphism of $\mathcal{F}(y)$ is determined by where it maps y , we can now compute the Galois group $\text{Gal}(\mathcal{F}(y)/\mathcal{F})$. We can compute the normal subgroups and hence the normal extensions of \mathcal{F} contained between $\mathcal{F}(a_1, \dots, a_n)$ and $\mathcal{F}(a_1, \dots, a_n, a_{n+1})$. Let

$$\mathcal{F}(a_1, \dots, a_n) \subsetneq \mathcal{K}_1 \subsetneq \dots \subsetneq \mathcal{K}_m = \mathcal{F}(a_1, \dots, a_n, a_{n+1})$$

be a maximal chain of normal extensions of \mathbb{F}_p . We can compute for each \mathcal{K}_i a primitive generator over \mathcal{F} and add these to the enumeration in order (with y chosen as the primitive generator of $\mathcal{K}_m = \mathcal{F}(a_1, \dots, a_n, a_{n+1})$).

Construction. At each stage s , we will have defined an embedding

$$\alpha_s: \mathcal{F}\{a_1, \dots, a_s\} \rightarrow \mathcal{E}$$

fixing \mathcal{F} such that $\alpha_0 \subseteq \alpha_1 \subseteq \dots \subseteq \alpha_s$. Begin with $\alpha_0: \mathbb{F}_p \rightarrow \mathcal{F}$.

At stage $s + 1$, we are given α_s . Use the splitting algorithm for $\mathcal{F}\{a_1, \dots, a_s\}$ to check whether $a_{s+1} \in \mathcal{F}\{a_1, \dots, a_s\}$. If it is, set $\alpha_{s+1} = \alpha_s$. Otherwise, check whether there is $e \leq s$ against which we have not yet diagonalized such that

- (i) $a_i \in \mathcal{F}(a_1, \dots, a_s, a_{s+1}) \setminus \mathcal{F}(a_1, \dots, a_s)$,
- (ii) for each $x \in \bar{\mathcal{E}}$ which satisfies the same minimal polynomial over \mathcal{F} as a_i ,

$$\varphi_{e,s}(x) \downarrow = c$$

for some $c \in \bar{\mathcal{E}}$.

This is a computable search. We have splitting algorithms for $\mathcal{F}(a_1, \dots, a_s, a_{s+1})$ and $\mathcal{F}(a_1, \dots, a_s)$, so we can check (i) for a given a_i . Also, $\varphi_{e,s}(x)$ converges only for x among the first s -many elements of $\bar{\mathcal{E}}$, and we can use our splitting algorithms to compute the finite set of a_i satisfying (i) and also satisfying the same minimal polynomial as some such x .

If there is such an e , choose the least one. Let x_1, \dots, x_n be the conjugates of a_i over \mathcal{F} . By property (ii) of the enumeration,

$$\mathcal{F}(a_1, \dots, a_s, a_{s+1}) = \mathcal{F}(a_1, \dots, a_s, x_1, \dots, x_n).$$

Now we can extend $\varphi_{e,s}$ in a unique way to a computable automorphism of

$$\mathcal{F}(x_1, \dots, x_n).$$

If this automorphism is not the identity on \mathcal{F} , then since \mathcal{F} is normal, $\varphi_{e,s}$ will be incompatible with α no matter how we define α . Suppose that $\varphi_{e,s}$ is the identity on \mathcal{F} . Since $\text{Gal}(\mathcal{E}/\mathcal{F})$ has the non-covering property, by Lemma 4.5 we can extend α_s to an automorphism of α_{s+1} of $\mathcal{F}(a_1, \dots, a_s, a_{s+1})$ that is incompatible with the automorphism $\varphi_{e,s}$ on $\mathcal{F}(x_1, \dots, x_n)$, in the sense that $(\mathcal{F}(x_1, \dots, x_n), \varphi_e)$ does not embed as a difference field into $(\mathcal{F}(a_1, \dots, a_s, a_{s+1}), \alpha_{s+1})$. We can do all of this computably by looking at the actions of the automorphisms on the generators of the fields.

Verification. We get an automorphism $\alpha = \cup_s \alpha_s$ of \mathcal{E} that fixes \mathcal{F} . Now we know that for some e , φ_e is an automorphism of $\bar{\mathcal{E}}$ such that (\mathcal{E}, α) embeds into $(\bar{\mathcal{E}}, \varphi_e)$ as a difference field. We claim that \mathcal{E} has a splitting algorithm. The proof will be to show that we can compute the image of \mathcal{E} in $\bar{\mathcal{E}}$ (since \mathcal{E} is a normal extension of \mathbb{F}_p , this image is unique; we may fix some embedding $\iota: \mathcal{E} \rightarrow \bar{\mathcal{E}}$ and show that the image of \mathcal{E} under ι is computable in $\bar{\mathcal{E}}$).

Let s be a stage after which we never diagonalize against an $e' \leq e$. Fix $x \in \bar{\mathcal{E}}$, and let $x = x_1, x_2, \dots, x_n$ be the conjugates of x over \mathcal{F} . Let $t \geq s$ be a stage by which $\varphi_e(x_i)$ has converged for each i . Since $\mathcal{F}(a_1, \dots, a_t)$ has a splitting algorithm, we can compute its image $\iota(\mathcal{F}(a_1, \dots, a_t))$ in $\bar{\mathcal{E}}$.

Claim $x \in \iota(\mathcal{E})$ if and only if $x \in \iota(\mathcal{F}(a_1, \dots, a_t))$.

Proof If $x \in \iota(\mathcal{F}(a_1, \dots, a_t))$, then $x \in \iota(\mathcal{E})$. On the other hand, suppose that $x \in \iota(\mathcal{E})$, say $x = \iota(a_i)$, and suppose to the contrary that $a_i \notin \mathcal{F}(a_1, \dots, a_t)$. Now, for some $t' > t$, we have $a_i \in \mathcal{F}(a_1, \dots, a_{t'+1}) \setminus \mathcal{F}(a_1, \dots, a_{t'})$. Then at stage $t' + 1$, we define $\alpha_{t'+1} \subset \alpha$ such that $(\mathcal{F}(x_1, \dots, x_n), \varphi_e)$ does not embed into

$$(\mathcal{F}(a_1, \dots, a_{t'+1}), \alpha_{t'+1})$$

as a difference field. Since $\mathcal{F}(x_1, \dots, x_n)$ and $\mathcal{F}(a_1, \dots, a_{t'+1})$ are both normal extensions of \mathbb{F}_p (with the former contained in the latter), (\mathcal{E}, α) cannot embed into $(\bar{\mathcal{E}}, \varphi_e)$. ■

From the claim we get a decision procedure for $\iota(\mathcal{E})$. Given $x \in \bar{\mathcal{E}}$, compute a stage $t \geq s$ at which φ_e converges when applied to all of the conjugates of x over \mathbb{F}_p . Using the splitting algorithm for $\mathcal{F}(a_1, \dots, a_t)$, we check whether $x \in \iota(\mathcal{F}(a_1, \dots, a_t))$ and hence whether $x \in \iota(\mathcal{E})$. ■

4.3 The Non-covering Property

To apply Theorem 4.6, we need a field extension whose Galois group has the non-covering property. We now give some examples of groups with the non-covering property before giving an example of an application of Theorem 4.6.

Lemma 4.7 *The following groups have the non-covering property:*

- (i) *abelian groups,*
- (ii) *simple groups,*
- (iii) *the quaternion group.*

Proof (i) Let G be an abelian group. Let $M \subsetneq N$ be normal subgroups of finite index, and fix $g \in G$. Let h be an element of $g(N \setminus M)$. Then for all $x \in G$, $x^{-1}hx = h \notin gM$.

(ii) Let G be a simple group. Let $M \subsetneq N$ be normal subgroups of finite index, and fix $g \in G$. Then $N = G$ and M is the trivial subgroup. Then $gN = G$; if $g = e$, pick $h \neq e$, and otherwise pick $h = e$. Then $gM = \{g\}$ and h and g are in different conjugacy classes.

(iii) Let $G = \{\pm 1, \pm i, \pm j, \pm k\}$ be the quaternion group. The normal subgroups are $\{1\}$, $\{1, -1\}$, $\{1, -1, i, -i\}$, $\{1, -1, j, -j\}$, $\{1, -1, k, -k\}$, and G . The conjugacy classes are $\{1\}$, $\{-1\}$, $\{i, -i\}$, $\{j, -j\}$, and $\{k, -k\}$. It is easy to see that every coset is a disjoint union of conjugacy classes. Thus, given normal subgroups $M \subsetneq N$ and $g \in G$, there is a conjugacy class in gN which is not in gM ; let h be in this conjugacy class. ■

The example from Proposition 4.3 has an abelian Galois group $\prod_{n \in \omega} C_2$, and hence Proposition 4.3 follows immediately from Theorem 4.6. Also, in characteristic $p > 0$ we have the following theorem.

Theorem 4.8 *Let \mathcal{E} be a computable normal extension of \mathbb{F}_p in characteristic $p > 0$. The following are equivalent.*

- (i) \mathcal{E} has a splitting algorithm.
- (ii) \mathcal{E} has the computable extendability of automorphisms property.
- (iii) \mathcal{E} has the uniform extendability of automorphisms property.

Proof The Galois group of every normal extension \mathcal{K}/\mathbb{F}_p in characteristic $p > 0$ is abelian and hence has the non-covering property. Theorem 4.6 finishes the proof. ■

We can also take arbitrary products of Galois groups with the non-covering property and produce another group with the non-covering property. We must assume that the groups are profinite, but as every Galois group is profinite, this is not a restriction. See [FJ08] for an introduction to profinite groups.

Now we can prove Theorem 1.6.

Proof of Theorem 1.6 We reduce the proof to the case of a product of two groups. If $M \subsetneq N$ are normal subgroups of $\prod_{i \in I} G_i$ of finite index, then M contains a finite intersection of the groups

$$\widehat{G}_i = \{(x_j)_{j \in I} : x_i = e\}.$$

The intersection of all of the \widehat{G}_i is the trivial group, so $\bigcap \widehat{G}_i \subseteq M$. Moreover, it is easy to check that these groups are open in the profinite topology of the profinite group $\prod_{i \in I} G_i$ (which is just the product topology) and hence they are closed as well. As the profinite topology is compact, M contains $\widehat{G}_{i_1} \cap \dots \cap \widehat{G}_{i_n}$ for some i_1, \dots, i_n .

Let $M', N' \subseteq G_{i_1} \times \dots \times G_{i_n}$ be the projection of M and N to these indices; M' and N' are normal subgroups. Then $(\prod_{i \in I} G_i)/M \cong (G_{i_1} \times \dots \times G_{i_n})/M'$. We will prove in the following lemma that $G_{i_1} \times \dots \times G_{i_n}$ has the non-covering property, and we can use this to check (for M and N) that $\prod_{i \in I} G_i$ has the non-covering property.

Lemma 4.9 *Let G and H be groups which both have the non-covering property. Then $G \times H$ has the non-covering property.*

Proof Let $M \not\subseteq N$ be normal subgroups of $G \times H$. Let π_1 and π_2 be the projections onto G and onto H , respectively.

Case 1: We have $\pi_1(M) \not\subseteq \pi_1(N)$. Let $a = (a_1, a_2) \in G \times H$ and $b = (b_1, b_2) \in G \times H$ be arbitrary. Choose $g = a_1 g' \in a_1 \pi_1(N)$ such that for all $x \in G$, $x^{-1} g x \notin b_1 \pi_1(M)$. Let $h' \in H$ be such that $(g', h') \in N$, and let $h = a_2 h'$. Then $f = (g, h) \in aN$ is such that for all $z = (x, y) \in G \times H$, $z^{-1} f z \notin bM$.

Case 2: We have $\pi_2(M) \not\subseteq \pi_2(N)$. Similar to Case 1.

Case 3: $\pi_1(M) = \pi_1(N)$ and $\pi_2(M) = \pi_2(N)$. Define $M_1 \subseteq G$ and $M_2 \subseteq H$ by

$$M_1 = \{x \in G : (x, e) \in M\} \text{ and } M_2 = \{y \in H : (e, y) \in M\}.$$

Then $M_1 \times M_2 \subseteq M$. Define N_1 and N_2 similarly. We have $M_1 \subseteq N_1$ and $M_2 \subseteq N_2$.

Claim 1 M_1 and N_1 are normal subgroups of G and M_2 and N_2 are normal subgroups of H .

Proof We show that M_1 is a normal subgroup of G . Let $m \in M_1$ and $x \in G$. Let $x' = (x, e)$ and $m' = (m, e)$. Then, since M is a normal subgroup of $G \times H$, $(x^{-1} m x, e) = x'^{-1} m' x' \in M$. Hence $x^{-1} m x \in M_1$. ■

Claim 2 $M_1 \not\subseteq N_1$ and $M_2 \not\subseteq N_2$.

Proof We use Goursat's Lemma.

Lemma ([Gou89]) Let G_1 and G_2 be groups. Let H be a subgroup of $G_1 \times G_2$ such that the projections $\pi_1: H \rightarrow G_1$ and $\pi_2: H \rightarrow G_2$ are surjective. Let N_1 and N_2 be the kernels of π_2 and π_1 respectively; N_1 can be identified as a normal subgroup of G_1 , and N_2 as a normal subgroup of G_2 . Then the image of H in $G_1/N_1 \times G_2/N_2$ is isomorphic to the graph of an isomorphism between G_1/N_1 and G_2/N_2 .

By Goursat's Lemma, the image of M in $\pi_1(M)/M_1 \times \pi_2(M)/M_2$ is the graph of an isomorphism $\pi_1(M)/M_1 \cong \pi_2(M)/M_2$. The same is true with M replaced by N . Since $\pi_1(M) = \pi_1(N)$, $\pi_2(M) = \pi_2(N)$, and $M \not\subseteq N$, we must have $M_1 \not\subseteq N_1$ and $M_2 \not\subseteq N_2$. ■

Claim 3 $[G, \pi_1(M)] \subseteq M_1$ and $[H, \pi_2(M)] \subseteq M_2$. Thus $[G \times H, \pi_1(M) \times \pi_2(M)] \subseteq M_1 \times M_2$.

Proof Let $g \in G$ and $m \in \pi_1(M)$. Let $g' = (g, e)$ and $m' = (m, e)$. Since M is a normal subgroup of $G \times H$, $[g', m'] = ([g, m], e) \in M$. Thus $[g, m] \in M_1$. ■

Fix $g \in G \times H$ for which we will show that there is $h \in gN$ such that for all $x \in G \times H$, $x^{-1} h x \notin gM$. This will finish the proof of Lemma 4.9. Since $N_2 \not\subseteq M_2$, we can choose $b \in \pi_2(g)M_2$ such that for all $y \in H$, $y^{-1} b y \notin \pi_2(g)M_2$. Choose $a = \pi_1(g)$. Then $(a, b) \in gN$. Suppose that $(x, y) \in G \times H$ is such that $(x^{-1} a x, y^{-1} b y) \in gM$. Let $m \in M$ be such that $x^{-1} a x = \pi_1(gm)$.

Claim 4 $\pi_1(m) \in M_1$.

Proof Suppose to the contrary that $\pi_1(m) \notin M_1$. Let $m_1 = \pi_1(m)$ and $g_1 = a = \pi_1(g)$. We have $x^{-1}g_1x = g_1m_1$. Let K be the subgroup of G generated by M_1 and m_1 . Since M_1 is a normal subgroup of G , each element of K can be written in the form km_1^ℓ for some $k \in M_1$ and $\ell \in \mathbb{N}$. Then K is a normal subgroup of G since $[G, m_1] \in M_1$. If $m_1 \notin M_1$, then M_1 is a proper subgroup of K . So there is $h \in K$ such that for all $z \in G$, $z^{-1}g_1hz \notin g_1M_1$. Let r be such that $m_1^r = e$ and let $h = km_1^\ell$ with $k \in M_1$ and $\ell < r$. Then since $[x, m_1] \in M_1$, $x^{-(r-\ell)}g_1hx^{r-\ell} \in x^{-(r-\ell)}g_1x^{r-\ell}m_1^\ell M_1 = g_1m_1^r M_1 = g_1M_1$. This is a contradiction which proves the claim. ■

Since $\pi_1(m) \in M_1$, we have $(e, \pi_2(m)) = m - (\pi_1(m), e) \in M$, and so $\pi_2(m) \in M_2$. But $y^{-1}by = \pi_2(gm) \notin \pi_2(g)M_2$, a contradiction. This completes the proof of Lemma 4.9. ■

This completes the proof of Theorem 1.6. ■

4.4 Examples

We can apply Theorem 1.6 to construct groups having the non-covering property from the groups in Lemma 4.7. In all cases, we know that if the field \mathcal{E} has a splitting algorithm, then it has the computable extendability of automorphisms property. We begin by noting that there exist groups without the non-covering property:

Proposition 4.10 *The groups S_3 , D_8 , and A_4 do not have the non-covering property.*

Proof For S_3 , let $M = \{e\}$ and let N be the normal subgroup of rotations. Let g be a reflection. Then gN is the set of all reflections, and all reflections are conjugate.

Write $D(8) = \{e, a, a^2, a^3, x, ax, a^2x, a^3x\}$. Let $M = \{e\}$, $N = \{e, a^2\}$, and $g = a$. Then $aM = \{a\}$ and $aN = \{a, a^3\}$. We have $x^{-1}ax = a^3$.

For A_4 , let $M = \{e\}$ and let N be the normal subgroup of A_4 isomorphic to $C_2 \times C_2$. Let g be the permutation $(1, 2, 3)$. Then gN consists of $(1, 2, 3)$, $(1, 4, 2)$, $(2, 4, 3)$, and $(1, 3, 4)$, all of which are conjugate. ■

Even if $\text{Gal}(\mathcal{E}/\mathbb{F}_p)$ does not have the non-covering property, we can still sometimes apply Theorem 4.6 either by finding the right field \mathcal{F} as in the statement of the theorem, or using Lemma 4.12 below with a subfield \mathcal{F} and applying Theorem 4.6 to the field extension \mathcal{F}/\mathbb{F}_p . The following two examples illustrate these methods. We begin with a field extension \mathcal{E}/\mathbb{Q} whose Galois group does not have the non-covering property, but we can use the freedom in choosing the field \mathcal{F} in the statement of Theorem 4.6 to apply the theorem.

Example 4.11 Let $\mathcal{E} = \mathbb{Q}(\omega, \sqrt[3]{p_n} : n \in \mathcal{O}')$, where ω is a primitive cube root of unity. Note that $\text{Gal}(\mathcal{E}/\mathbb{Q})$ does not have a forking lattice of subgroups for the same reason as S_3 , because its Galois group is $\text{Gal}(\mathcal{E}/\mathbb{Q}) = \prod_{i \in \omega} C_3 \rtimes C_2$ with C_2 acting on C_3 by inverting elements. Here we need to know that the intersection of the fields $\mathbb{Q}(\omega, \sqrt[3]{p_n} : n \in U)$ and $\mathbb{Q}(\omega, \sqrt[3]{p_n} : n \in V)$ for U and V disjoint is the field $\mathbb{Q}(\omega)$. See [Mor53].

Let $\mathcal{F} = \mathbb{Q}(\omega)$. Then \mathcal{F} has a splitting algorithm, and $\text{Gal}(\mathcal{E}/\mathcal{F}) = \prod_{i \in \omega} C_3$, which is abelian. Since \mathcal{E} does not have a splitting algorithm, by Theorem 4.6 it does not have the computable extension of automorphisms property.

The following lemma will allow us to consider a subextension of \mathcal{E} ; this will be useful when the Galois group of the extension does not have the non-covering property, but it has a quotient which does.

Lemma 4.12 *Let $\mathcal{E} \supseteq \mathcal{F} \supseteq \mathbb{F}_p$ be computable algebraic extensions such that \mathcal{E} is a normal extension of \mathbb{F}_p . Suppose that given $x \in \mathcal{E}$, we can compute the minimal polynomial of x over \mathcal{F} . Then if \mathcal{E} has the computable extendability of automorphisms property, \mathcal{F} does as well.*

Proof This follows from the fact that we can computably extend an automorphism of \mathcal{F} to an automorphism of \mathcal{E} in the style of Theorem 3.2 and uses the fact that \mathcal{F} is a perfect field. ■

We now have an example where we apply this lemma together with Theorem 4.6.

Example 4.13 This example is quite complicated. The idea is to produce a field extension whose Galois group is $\prod_{n \in \omega} S_3$, but which does not have a splitting algorithm.

Let q_0, q_1, \dots be a list of infinitely many distinct primes in the arithmetic progression $4n + 27$, and let a_n be such that $4a_n + 27 = q_n$. Let \mathcal{E} be the splitting field over \mathbb{Q} of the polynomials $\{x^3 + a_n x + a_n : n \in \mathcal{O}'\}$. Let ω_n be a primitive element for the splitting field of $x^3 + a_n x + a_n$, so that $\mathcal{E} = \mathbb{Q}(\omega_n : n \in \mathcal{O}')$. Each of these polynomials has discriminant $D_n = -4a_n^3 - 27a_n^2 = -a_n^2 q_n < 0$, and hence $\mathbb{Q}(\omega_n)$ has Galois group S_3 . We claim that the Galois group of \mathcal{E} is $\prod_{n \in \omega} S_3$. It suffices to show that given m and n_1, \dots, n_ℓ all distinct that $\mathbb{Q}(\omega_m)$ and $\mathbb{Q}(\omega_{n_1}, \dots, \omega_{n_\ell})$ are disjoint. Suppose not; then there is a non-trivial subfield \mathcal{K} of $\mathbb{Q}(\omega_m)$ which is contained in $\mathbb{Q}(\omega_{n_1}, \dots, \omega_{n_\ell})$. We may assume that $\mathcal{K} = \mathbb{Q}(\sqrt{D_m}) = \mathbb{Q}(\sqrt{-q_m})$. Then $\sqrt{D_m} \in \mathbb{Q}(\sqrt{D_{n_1}}, \dots, \sqrt{D_{n_\ell}})$, which is a contradiction, since $q_m, q_{n_1}, \dots, q_{n_\ell}$ are distinct primes. Now \mathcal{E} does not have a splitting algorithm, but $\prod_{n \in \omega} S_3$ does not have a forking lattice of subgroups.

Now let $\mathcal{F} = \mathbb{Q}(\sqrt{-q_n} : n \in \mathcal{O}')$. Then \mathcal{F} does not have a splitting algorithm. By Theorem 4.6, \mathcal{F} does not have the computable extension of isomorphisms property, and hence by Lemma 4.12, \mathcal{E} does not have the computable extension of automorphisms property.

We do not know of any examples in which one cannot use either a direct application of Theorem 4.6 or one of the methods in these two examples.

5 Applications to Difference Closed Fields

We will conclude this paper by applying our results to difference closed fields. The main idea will be to note that (\mathcal{F}, σ) embeds into a computable difference closed field if and only if there is an embedding ι of \mathcal{F} into $\overline{\mathcal{F}}$ and an automorphism τ of $\overline{\mathcal{F}}$ such that $\tau \iota$ -extends σ . In the one direction, this will follow from an effective Henkin

construction, while, on the other hand, it will follow from the fact that the algebraic closure of the prime field can be enumerated in any difference closed field.

Theorem 5.1 *Let \mathcal{F} be a computable extension of \mathbb{F}_p , and σ a computable automorphism of \mathcal{F} . Then the following are equivalent.*

- (i) (\mathcal{F}, σ) embeds computably into a computable difference closed field.
- (ii) There is a computable embedding $\iota: \mathcal{F} \rightarrow \overline{\mathcal{F}}$ of \mathcal{F} into a computable presentation of its algebraic closure and a computable automorphism τ of $\overline{\mathcal{F}}$ which ι -extends σ .

Proof We begin by proving (i) \Rightarrow (ii). Suppose that there is a computable difference closed field (\mathcal{K}, ρ) into which (\mathcal{F}, σ) embeds. We can enumerate in \mathcal{K} the algebraic closure $\overline{\mathcal{F}}$ of \mathcal{F} (which is also the algebraic closure of the prime field) and the restriction τ of ρ to $\overline{\mathcal{F}}$ (recall that every computable presentation of the algebraic closure of \mathcal{F} is computably isomorphic to every other computable presentation). Then since (\mathcal{F}, σ) embeds into (\mathcal{K}, ρ) and is algebraic over \mathbb{F}_p , its image is in $(\overline{\mathcal{F}}, \tau)$. Then τ is an extension of σ to $\overline{\mathcal{F}}$ via this embedding.

We now prove (i) \Rightarrow (ii). The completions of ACFA are given by the possible actions of the automorphism σ on the algebraic closure of the prime field $\overline{\mathbb{F}_p}$ (see [CH99, (1.4)]). Let ι be a computable embedding of \mathcal{F} into $\overline{\mathcal{F}}$ and τ an ι -extension of σ to $\overline{\mathcal{F}}$. Let $\mathcal{L}_{\overline{\mathcal{F}}}$ be the language of difference fields together with names for the constants of $\overline{\mathcal{F}}$. Let T be the consistent theory axiomatized by ACFA together with the existential diagram of $(\overline{\mathcal{F}}, \sigma)$. Then T contains a completion of ACFA, and since every formula is equivalent to an existential formula modulo ACFA, T is complete. Moreover, T is recursively axiomatizable and hence computable. So T has a decidable model (\mathcal{K}, ρ) . Using the embedding $\iota: \mathcal{F} \rightarrow \overline{\mathcal{F}}$, we get an embedding of the difference field (\mathcal{F}, σ) into (\mathcal{K}, ρ) . ■

We can use this, together with the examples from the previous section, to see that Rabin's Theorem on the existence of computable algebraic closures (and its analogue in differentially closed fields due to Harrington [Har74]) does not hold in the context of difference closed fields:

Corollary 5.2 *There exist computable difference fields which cannot be effectively embedded into any computable difference closed field. Moreover, there is a counterexample in every characteristic.*

Proof In characteristic zero, apply the previous corollary to the field from Proposition 4.3, and in characteristic $p > 0$, by Corollary 4.8, we can use any normal extension of \mathbb{F}_p with no splitting algorithm. ■

Corollary 5.3 *The analogue of Rabin's Theorem holds for difference fields with underlying field \mathcal{F} if and only if \mathcal{F} has the computable extension of automorphisms property.*

A set is *low* if its Turing jump is as low as possible, i.e., Turing equivalent to \emptyset' . We note that every computable difference field does embed into a *low* difference closed field.

Fact 5.4 (essentially [FSS83]) *Every computable difference field embeds (by a map of low degree) into a low difference closed field.*

Proof Let (\mathcal{F}, σ) be a computable difference field. Let $\iota: \mathcal{F} \rightarrow \overline{\mathcal{F}}$ be a computable embedding of \mathcal{F} into its algebraic closure. Then there is a low automorphism τ of $\overline{\mathcal{F}}$ extending σ (see [FSS83]). The theory *ACFA* together with the action of τ on $\overline{\mathcal{F}}$ is a complete low theory, and an effective Henkin construction produces a low model as in Theorem 5.1. ■

In Theorem 4.6, we showed that for a field whose Galois group has the non-covering property, having a splitting algorithm is equivalent to the computable extendability of automorphisms property. We do not know in general whether these are equivalent. We leave the following question open.

Question 5.5 *For a normal extension \mathcal{F} of \mathbb{Q} , is the computable extendability of automorphisms property equivalent to having a splitting algorithm?*

References

- [Bab62] Albert E. Babbitt, Jr., *Finitely generated pathological extensions of difference fields*. Trans. Amer. Math. Soc. **102**(1962), 63–81. <http://dx.doi.org/10.1090/S0002-9947-1962-0133326-0>
- [Bes40] Abram S. Besicovitch, *On the linear independence of fractional powers of integers*. J. London Math. Soc. **15**(1940), 3–6. <http://dx.doi.org/10.1112/jlms/s1-15.1.3>
- [CH99] Zoé Chatzidakis and Ehud Hrushovski, *Model theory of difference fields*. Trans. Amer. Math. Soc. **351**(1999), no. 8, 2997–3071. <http://dx.doi.org/10.1090/S0002-9947-99-02498-8>
- [Cle70] John P. Cleave, *Some properties of recursively inseparable sets*. Z. Math. Logik Grundlagen Math. **16**(1970), 187–200. <http://dx.doi.org/10.1002/malq.19700160208>
- [Coh52] Richard M. Cohn, *Extensions of difference fields*. Amer. J. Math. **74**(1952) 507–530. <http://dx.doi.org/10.2307/2372012>
- [Coh65] ———, *Difference algebra*. Interscience Publishers John Wiley & Sons, New York, 1965.
- [DHS13] François G. Dorais, Jeffrey Hirst, and Paul Shafer, *Reverse mathematics and algebraic field extensions*. Computability **2**(2013), no. 2, 75–92.
- [Eva73] Peter Evanovich, *Algebraic extensions of difference fields*. Trans. Amer. Math. Soc. **179**(1973), 1–22. <http://dx.doi.org/10.1090/S0002-9947-1973-0314809-4>
- [FJ08] Michael D. Fried and Moshe Jarden, *Field arithmetic*. Third edition. Ergebnisse der Mathematik und ihrer Grenzgebiete, 11. Springer-Verlag, Berlin, 2008.
- [FSS83] Harvey M. Friedman, Stephen G. Simpson, and Rick L. Smith, *Countable algebra and set existence axioms*. Ann. Pure Appl. Logic **25**(1983), no. 2, 141–181. [http://dx.doi.org/10.1016/0168-0072\(83\)90012-X](http://dx.doi.org/10.1016/0168-0072(83)90012-X)
- [Gou89] Edouard Goursat, *Sur les substitutions orthogonales et les divisions régulières de l'espace*. Ann. Sci. École Norm. Sup. (3) **6**(1889), 9–102.
- [Har74] Leo Harrington, *Recursively presentable prime models*. J. Symbolic Logic **39**(1974), 305–309. <http://dx.doi.org/10.2307/2272643>
- [Har98] Valentina S. Harizanov, *Pure computable model theory*. In: Handbook of recursive mathematics, Vol. 1. Stud. Logic Found. Math., 138. North-Holland, Amsterdam, 1998, pp. 3–114.
- [HTMM15] Matthew Harrison-Trainor, Alexander Melnikov, and Antonio Montalbán, *Independence in computable algebra*. J. Algebra **443**(2015), 441–468. <http://dx.doi.org/10.1016/j.jalgebra.2015.06.004>
- [Kro82] Leopold Kronecker, *Grundzüge einer arithmetischen Theorie der algebraischen Größen*. J. Reine Angew. Math. **92**(1882), 1–122. <http://dx.doi.org/10.1515/crll.1882.92.1>
- [Mac97] Angus Macintyre, *Generic automorphisms of fields*. Ann. Pure Appl. Logic **88**(1997), no. 2–3, 165–180. [http://dx.doi.org/10.1016/S0168-0072\(97\)00020-1](http://dx.doi.org/10.1016/S0168-0072(97)00020-1)
- [Mal61] Anatoly I. Mal'cev. *Constructive algebras. I*. Uspehi Mat. Nauk **16**(1961), no. 3 (99), 3–60.

- [Mil83] Terrence Millar, *Omitting types, type spectrums, and decidability*. J. Symbolic Logic 48(1983), no. 1, 171–181. <http://dx.doi.org/10.2307/2273331>
- [Mil08] Russell Miller, *Computable fields and Galois theory*. Notices Amer. Math. Soc. 55(2008), no. 7, 798–807.
- [Mor53] Louis J. Mordell, *On the linear independence of algebraic numbers*. Pacific J. Math. 3(1953), 625–630. <http://dx.doi.org/10.2140/pjm.1953.3.625>
- [Rab60] Michael O. Rabin, *Computable algebra, general theory and theory of computable fields*. Trans. Amer. Math. Soc. 95(1960), 341–360.
- [vdW70] Bartel L. van der Waerden, *Algebra*. Vol 1. Frederick Ungar, New York, 1970.

Group in Logic and the Methodology of Science, University of California, Berkeley, USA
e-mail: matthew.h-t@berkeley.edu

The Institute of Natural and Mathematical Sciences, Massey University, New Zealand
e-mail: alexander.g.melnikov@gmail.com

Dept. of Mathematics, Queens College, & Ph.D. Programs in Mathematics & Computer Science, Graduate Center, City University of New York, USA
e-mail: Russell.Miller@qc.cuny.edu