

Model Completeness and Relative Decidability

Jennifer Chubb, Russell Miller,* & Reed Solomon

September 1, 2020

Abstract

We study the implications of model completeness of a theory for the effectiveness of presentations of models of that theory. It is immediate that for a computable model \mathcal{A} of a computably enumerable, model complete theory, the entire elementary diagram $E(\mathcal{A})$ must be decidable. We prove that indeed a c.e. theory T is model complete if and only if there is a uniform procedure that succeeds in deciding $E(\mathcal{A})$ from the atomic diagram $\Delta(\mathcal{A})$ for all countable models \mathcal{A} of T . Moreover, if every presentation of a single isomorphism type \mathcal{A} has this property of relative decidability, then there must be a procedure which succeeds uniformly for all presentations of an expansion (\mathcal{A}, \vec{a}) by finitely many new constants. We end with a conjecture about the situation when all models of a theory are relatively decidable.

1 Introduction

The broad goal of computable model theory is to investigate the effective aspects of model theory. Here we will carry out exactly this process with the model-theoretic concept of *model completeness*. This notion is well-known and has been widely studied in model theory, but to our knowledge there has never been any thorough examination of its implications for computability in structures with the domain ω . We now rectify this omission, and find natural and satisfactory equivalents for the basic notion of model completeness of a first-order theory. Our two principal results are each readily stated: that a

*The second author was supported by NSF grant # DMS-1362206, Simons Foundation grant # 581896, and several PSC-CUNY research awards.

computably enumerable theory is model complete if and only if there is a uniform procedure for deciding the elementary diagram of each of its models (on the domain ω) from the corresponding atomic diagram (Theorem 2.5); and that if every presentation of a particular isomorphism type \mathcal{A} individually has this property (which we call *relative decidability*), then for some expansion (\mathcal{A}, \vec{a}) by finitely many constants, there must be a uniform version of the procedure that succeeds for all presentations (Theorem 4.1).

Recall that a structure \mathcal{A} with domain ω (or a decidable subset of ω) is *computable* if its atomic diagram is computable, and is *decidable* if its elementary diagram is computable. This choice of terminology may seem arbitrary, but it is well established, and we will maintain it here. In our notation, $\Delta(\mathcal{A})$ is the atomic diagram of a structure \mathcal{A} with domain ω . In a fixed countable language \mathcal{L} , fix an effective Gödel coding of all atomic sentences in the language $(\mathcal{L}, c_0, c_1, \dots)$ with new constants for the naturals; then $\Delta(\mathcal{A})$ can be viewed as a subset of ω , containing the code numbers of those sentences true in \mathcal{A} , with each c_n representing the domain element n of \mathcal{A} . We assume implicitly that $\Delta(\mathcal{A})$ always respects equality on ω , i.e., that the atomic statement $c_i = c_j$ lies in $\Delta(\mathcal{A})$ only if $i = j$. $E(\mathcal{A})$ is the elementary diagram (a.k.a. the complete diagram) of \mathcal{A} , coded similarly as a subset of ω .

Definition 1.1 A theory T has *quantifier elimination down to \mathcal{C}* , where \mathcal{C} is a class of formulas, if, for every (finitary) formula $\alpha(\vec{x})$, there exists a formula $\gamma(\vec{x})$ in \mathcal{C} , with the same free variables \vec{x} , such that

$$T \models \forall \vec{x} (\alpha(\vec{x}) \leftrightarrow \gamma(\vec{x})).$$

For example, it is well-known that a theory is model complete if and only if it has quantifier elimination down to (finitary) existential formulas, or equivalently, down to universal formulas. (For a proof, see [2, Theorem 3.5.1].) Later on, we will consider quantifier elimination down to existential $L_{\omega_1\omega}$ formulas.

Having introduced the term, we recall the definition.

Definition 1.2 A theory T is *model complete* if, for every model \mathcal{B} of T , every substructure $\mathcal{A} \subseteq \mathcal{B}$ which is itself a model of T is an elementary substructure of \mathcal{B} .

It is equivalent to require that, for every $\mathcal{A} \subseteq \mathcal{B}$ as described, every existential statement $\exists \vec{x} \varphi(\vec{a}, \vec{x})$ which holds in \mathcal{B} (of a tuple \vec{a} from \mathcal{A}) also holds in \mathcal{A} .

Our two principal theorems will be proven in Sections 2 and 4, respectively. In between, in Section 3, we offer examples of theories that are or are not model complete; these might be perused before Section 2 by the reader desiring a refresher, but they also lead in well to Section 4. At the end, in Section 5, we offer a conjecture regarding the situation of a theory for which every model is relatively decidable, but not uniformly so. After the initial submission of this article, Harrison-Trainor resolved our conjecture; a reference is provided.

Arguments in this article always use countable structures, whether we say so or not. A *presentation* of a countable structure is simply a structure isomorphic to it, whose domain is the set ω of nonnegative integers. This enables our Gödel coding above to make sense in all cases. Finally, all languages considered here will be computable languages, with at most countably many symbols. We believe that, for a noncomputable countable language \mathcal{L} , all the results would go through without difficulty if one simply relativized the statements and arguments to the Turing degree of \mathcal{L} .

2 Basics of Model Completeness

We begin with the following result, which was remarked in [4, Prop. 6.7].

Proposition 2.1 *Let T be a model complete theory, and assume T is c.e. Then every computable model of T is decidable.*

Proof. Fix any formula $\varphi(x_0, \dots, x_{n-1})$ and any $\vec{a} \in \omega^n$ for which we wish to determine whether $\varphi(\vec{a})$ holds in \mathcal{A} . As noted above, T has the property that, for this (and every) formula $\varphi(y_0, \dots, y_{n-1})$, there is a universal formula in the same variables which is provably equivalent (under T) to φ . Since T is c.e., we may therefore search until we find quantifier-free formulas $\alpha(x_0, \dots, x_{n-1}, y_0, \dots, y_m)$ and $\beta(\vec{x}, \vec{y})$ and a deduction showing

$$T \vdash \forall \vec{x} [(\varphi(\vec{x}) \leftrightarrow \forall \vec{y} \alpha(\vec{x}, \vec{y})) \ \& \ ((\neg \varphi(\vec{x})) \leftrightarrow \forall \vec{y} \beta(\vec{x}, \vec{y}))].$$

Model-completeness of T ensures that this search will terminate.

Now, in the computable structure \mathcal{A} , we search for some tuple $\vec{b} \in \omega^m$ such that either $(\neg \alpha(\vec{a}, \vec{b}))$ or $(\neg \beta(\vec{a}, \vec{b}))$ holds in \mathcal{A} . By our choice of α and β , this search must also terminate. When it does, we know which of the formulas $\forall \vec{y} \alpha(\vec{a}, \vec{y})$ and $\forall \vec{y} \beta(\vec{a}, \vec{y})$ fails to hold in \mathcal{A} , and from this we determine which of $\varphi(\vec{a})$ and $(\neg \varphi(\vec{a}))$ holds in \mathcal{A} . ■

The proposition generalizes easily to the following full theorem. (It would be natural here for \mathbf{c} to be the Turing degree of an axiomatization of the theory T .)

Theorem 2.2 *Let T be a model complete theory, and \mathbf{c} a Turing degree. Assume T is \mathbf{c} -c.e. Then, for every Turing degree $\mathbf{d} \geq \mathbf{c}$, every \mathbf{d} -computable model of T is \mathbf{d} -decidable. ■*

We ask whether the converse of this theorem holds: if a c.e. theory T has the property that every \mathbf{d} -computable model of T is \mathbf{d} -decidable, must T be model complete? (Likewise for \mathbf{c} -c.e. theories where this holds for all $\mathbf{d} \geq \mathbf{c}$.) Alternatively, there might be some version involving $L_{\omega_1\omega}$ formulas, perhaps computable ones, which would be equivalent to model completeness.

The first answer is that the converse does not hold.

Proposition 2.3 *The theory $T = \text{Th}(\omega, S)$ is not model complete, yet for every \mathbf{d} , every \mathbf{d} -computable model is \mathbf{d} -decidable.*

Proof. Note that the theory T is computable because $\text{Th}(\omega, S, 0)$ is a computable theory which is a conservative extension of T . Therefore, for any $\{S\}$ -formula φ , $\varphi \in T$ if and only if $\varphi \in \text{Th}(\omega, S, 0)$.

If $\mathcal{A} = (\omega, S)$ and $\mathcal{B} = (\omega - \{0\}, S)$, then \mathcal{A} and \mathcal{B} are both models of T , with $\mathcal{A} \subseteq \mathcal{B}$, yet \mathcal{B} is not an elementary substructure of \mathcal{A} : the element 1 is a successor in \mathcal{A} , but not in \mathcal{B} . Therefore, T is not model complete.

However, whenever a model \mathcal{C} of T is \mathbf{d} -computable, there is a decision procedure for its elementary diagram. Indeed, all models of the model complete theory $\text{Th}(\omega, 0, S)$ have this property, and one simply applies the same procedure to \mathcal{C} using the unique element of \mathcal{C} with no predecessor as the constant. ■

This proposition points to the true kernel of the converse. The process of deciding the elementary diagram of a model of $\text{Th}(\omega, S)$ from its atomic diagram was nonuniform: it required knowledge of the unique non-successor element, which cannot be determined uniformly from the atomic diagram. On the other hand, the procedure described in Theorem 2.2 was uniform. Therefore, we modify our conjecture about the converse to require uniformity.

Definition 2.4 A structure \mathcal{A} is *relatively decidable* if $E(\mathcal{A}) \leq_T \Delta(\mathcal{A})$. A class \mathfrak{S} of structures is *uniformly relatively decidable* if there exists a single

Turing functional Γ such that, for every $\mathcal{A} \in \mathfrak{S}$, the function $\Gamma^{\Delta(\mathcal{A})}$ is the characteristic function of the elementary diagram $E(\mathcal{A})$.

The structure \mathcal{A} is *uniformly relatively decidable* if the class of all structures (with domain ω) isomorphic to \mathcal{A} is uniformly relatively decidable. A theory T is *uniformly relatively decidable* if the class of all models of T with domain ω is uniformly relatively decidable.

To prove equivalence between this concept and model completeness, we use a broader form of quantifier elimination.

Theorem 2.5 *For a computably enumerable theory T (in particular, for any theory with a decidable axiom set), the following are equivalent:*

1. T is model complete.
2. T is uniformly relatively decidable.
3. T has effective quantifier elimination down to Σ_1^c formulas, that is, down to computable infinitary Σ_1 formulas.
4. T has effective quantifier elimination down to finitary Σ_1 formulas.

Proof. (1) and (4) are well-known to be equivalent; see for instance [2, Theorem 3.5.1]. (The effectiveness of the quantifier elimination is not stated there, but a c.e. theory with quantifier elimination always has effective quantifier elimination.) Moreover, the proof of Proposition 2.1 actually shows that (4) implies (2), since the procedure given there for deciding the elementary diagram of \mathcal{A} from $\Delta(\mathcal{A})$ is uniform for all models \mathcal{A} of T .

It is also quickly seen that (3) implies (1). Recall Definition 1.2, and let \mathcal{A} and \mathcal{B} be models of T with $\mathcal{A} \subseteq \mathcal{B}$. Now for each formula $\varphi(\vec{x})$, we have a Σ_1^c formula $\gamma(\vec{x})$ equivalent to $(\neg\varphi(\vec{x}))$ in all models of T . For each tuple \vec{a} from \mathcal{A} such that $\mathcal{B} \models \varphi(\vec{a})$, it is immediate that $\mathcal{B} \not\models \gamma(\vec{a})$, hence that $\mathcal{A} \not\models \gamma(\vec{a})$ (because Σ_1^c formulas true in \mathcal{A} must hold in superstructures of \mathcal{A}), hence that $\mathcal{A} \models \varphi(\vec{a})$. Thus \mathcal{A} is an elementary substructure of \mathcal{B} .

It remains to show that (2) implies (3). Assume that T is uniformly relatively decidable, via some functional Γ . Consider the arbitrary (finitary) formula $\alpha(x_0, \dots, x_n)$. For simplicity, we will assume for the time being that α includes the conditions $x_i \neq x_j$ for all $i < j \leq n$. At the end of the proof, we show how to remove this assumption.

We claim that in every model of T , the formula $\alpha(\vec{x})$ will be equivalent to the following Σ_1^c formula $\beta_\alpha(\vec{x})$:

$$\bigvee_{\sigma \in H_\alpha} \exists y_{n+1}, \dots, y_{m_\sigma} \gamma_\sigma(\vec{x}, \vec{y}),$$

where the set H_α is c.e., uniformly in α , as we now explain. First, given any $\sigma \in 2^{<\omega}$, let $\gamma_\sigma(c_0, \dots, c_{m_\sigma})$ be the finitary quantifier-free formula

$$\left(\bigwedge_{\sigma(\ulcorner \psi \urcorner)=1} \psi \right) \wedge \left(\bigwedge_{\sigma(\ulcorner \psi \urcorner)=0} (\neg \psi) \right) \wedge \left(\bigwedge_{i < j < |\sigma|} c_i \neq c_j \right),$$

using our Gödel numbering of the atomic formulas in the language of T with new constants c_0, c_1, \dots . Here σ is viewed as a possible initial segment of an atomic diagram, although of course for many strings σ , γ_σ may already contradict T . Next, let

$$H_\alpha = \{ \sigma \in 2^{<\omega} : \Gamma^\sigma(\ulcorner \alpha(c_0, \dots, c_n) \urcorner) \downarrow = 1 \}.$$

So σ lies in H_α if γ_σ is enough information for the procedure Γ to conclude that $\alpha(\vec{c})$ holds in those models \mathcal{A} of T (if any) whose atomic diagrams specify that γ_σ holds in \mathcal{A} .

Finally, when defining $\beta_\alpha(\vec{x})$, we quantified over certain y_i ; we now explain how. Each γ_σ is a quantifier-free formula involving some constants c_0, \dots, c_{m_σ} , quite possibly with $m_\sigma > n$. In $\beta_\alpha(\vec{x})$, each remaining c_i (with $i > n$) is replaced by y_i , with an existential quantification $\exists y_i$ added in front of γ_σ . Thus the formula $\beta_\alpha(\vec{x})$ involves \vec{x} but no other free variables, nor any constants.

It is clear that H_α is a c.e. set, and therefore that $\beta_\alpha(\vec{x})$ is a Σ_1^c formula. Moreover, the process is uniform in α . It remains to show that, for every model \mathcal{A} of T and every \vec{a} from \mathcal{A} , $\alpha(\vec{a})$ holds in \mathcal{A} if and only if $\beta_\alpha(\vec{a})$ holds there. Suppose first that $\mathcal{A} \models \alpha(\vec{a})$. Fix a permutation f of ω satisfying:

- $f(a_i) = i$ for all $i \leq n$; and
- $f(a) = a$ for all but finitely many $a \in \omega$.

(Recall that we are assuming for the time being that $\alpha(\vec{x})$ includes the conditions $x_i \neq x_j$ for $i \neq j$. This is necessary for f to exist.) Define the structure

\mathcal{A}_f so that f is an isomorphism from \mathcal{A} onto \mathcal{A}_f . Then $\mathcal{A}_f \models \alpha(0, \dots, n)$, and therefore $\Gamma^{\Delta(\mathcal{A}_f)}(\ulcorner \alpha(0, \dots, n) \urcorner) \downarrow = 1$. Let σ be the initial segment of $\Delta(\mathcal{A}_f)$ as long as the use of this computation, so that $\sigma \in H_\alpha$. But this means that $\mathcal{A}_f \models \gamma_\sigma(0, \dots, m_\sigma)$, and so the isomorphism f^{-1} shows that \mathcal{A} also satisfies $\mathcal{A} \models \exists y_{n+1} \dots \exists y_{m_\sigma} \gamma_\sigma(\vec{a}, \vec{y})$. Thus \mathcal{A} satisfies $\beta_\alpha(\vec{a})$, as required.

Conversely, suppose $\mathcal{A} \models \beta_\alpha(\vec{a})$. Then for some $\sigma \in H_\alpha$, $\mathcal{A} \models \exists \vec{y} \gamma_\sigma(\vec{a}, \vec{y})$. Fix $b_{n+1}, \dots, b_{m_\sigma}$ such that $\mathcal{A} \models \gamma_\sigma(\vec{a}, \vec{b})$. Define f to be a permutation as above with $f(a_i) = i$ and $f(b_j) = j$. Let \mathcal{A}_f be the corresponding copy of \mathcal{A} , so $\mathcal{A}_f \models \gamma_\sigma(0, \dots, m_\sigma)$ and hence σ is an initial segment of $\Delta(\mathcal{A}_f)$. Because $\sigma \in H_\alpha$, we see that $\Gamma^{\Delta(\mathcal{A}_f)}(\ulcorner \alpha(0, \dots, n) \urcorner) \downarrow = 1$. This means that $\alpha(0, \dots, n)$ must hold in \mathcal{A}_f , and the isomorphism f^{-1} now shows that $\mathcal{A} \models \alpha(a_0, \dots, a_n)$ as well.

All of this was proven under the assumption that $\alpha(\vec{x})$ includes the conditions $x_i \neq x_j$ for $i < j \leq n$. For the general case, one simply expresses α as a finite disjunction of the possibilities, and applies the process above to each one individually, getting a finite disjunction of Σ_1^c formulas. For example, if $n = 2$, then $\alpha(x_0, x_1, x_2)$ is equivalent to the following disjunction.

$$\begin{aligned} & \alpha(x_0, x_0, x_0) \\ \text{or } & (\alpha(x_0, x_1, x_1) \ \& \ x_0 \neq x_1) \\ \text{or } & (\alpha(x_0, x_1, x_0) \ \& \ x_0 \neq x_1) \\ \text{or } & (\alpha(x_0, x_0, x_1) \ \& \ x_0 \neq x_1) \\ \text{or } & (\alpha(x_0, x_1, x_2) \ \& \ x_0 \neq x_1 \ \& \ x_0 \neq x_2 \ \& \ x_1 \neq x_2). \quad \blacksquare \end{aligned}$$

Theorem 2.5 has a natural generalization to theories T that are not computably enumerable. If S is a set that can enumerate T (for example, if S is an axiom set for T), then T is model complete if and only if it is uniformly relatively S -decidable. By this we mean that there is a Turing functional Γ such that, for every model \mathcal{A} of T with domain ω , $\Gamma^{S \oplus \Delta(\mathcal{A})}$ computes $E(\mathcal{A})$. The argument is identical to that in the proof of the theorem, and uses S -effective quantifier elimination down to S -computable infinitary Σ_1 formulas as an intermediate equivalent.

3 Examples

Here we provide some examples of situations where the addition of new symbols to the language can allow a theory to become model complete. Several

examples of model complete theories appear in [2, §3.5]. One of these is the dense linear order with endpoints, in a language with the relation $<$ and two constant symbols to name the endpoints. Without the constants, the theory T of dense linear orders with endpoints would not be model complete: the rational interval $[0, \frac{1}{2}]_{\mathbb{Q}}$ would be a substructure of $[0, 1]_{\mathbb{Q}}$, and both would be models of T , yet the larger structure would satisfy $(\exists x) \frac{1}{2} < x$ and the substructure would not. When T is augmented by sentences in the larger language saying that the constants represent the two endpoints, it becomes model complete.

Similar examples, requiring arbitrarily many constants, are provided by the theory of the usual linear order on $\mathcal{A}_n = \cup_{i=0}^n [2i, 2i+1]_{\mathbb{Q}}$. Here $(2n+2)$ constants are necessary, to name the end points of all the different intervals. The proof is much the same as for a single interval.

A more intriguing example arises when we consider the linear order \mathcal{B} given by the lexicographic order on $\mathbb{Q} \times \{0, 1\}$. This arises often in the literature, either as the *shuffle sum* of countably many copies of the two-element order, or as “ \mathbb{Q} with every point doubled,” or other ways. The theory T of \mathcal{B} is not model complete: the substructure

$$\mathcal{A} = \{(q, k) \in \mathcal{B} : q < 4 \text{ or } q > 5\} \cup \{(4, 0), (5, 1)\}$$

of \mathcal{B} is not an elementary substructure, yet satisfies T , being isomorphic to \mathcal{B} . To make this T model complete, we adjoin a binary relation symbol **Adj** to the language, with axioms stating that **Adj**(x, y) holds just if $x < y$ and $\forall z \neg(x < z < y)$. The theory T^* of the extension \mathcal{B}^* in the new language (with **Adj** ^{\mathcal{B}^*} defined as instructed) includes sentences saying that every element is half of a unique adjacency, and that the pairs satisfying **Adj** are dense with no least or greatest such pair. One then proves that every substructure of \mathcal{B}^* modelling T^* must be elementary, so that T^* is indeed model complete.

It is natural to ask whether the same could have been accomplished by augmenting the original signature $\{<\}$ by countably many constants, say c_q and d_q for all $q \in \mathbb{Q}$, using c_q and d_q to name the elements $(q, 0)$ and $(q, 1)$ in \mathcal{B} , respectively. If \mathcal{B}^c is the structure \mathcal{B} thus enriched, and T^c is its theory, then \mathcal{B}^c has no substructure except itself. However, consider the structure \mathcal{A}^c of the lexicographic order on $\mathbb{Q} \times \{0, 1\}$ in which $c_q^{\mathcal{A}} = (q, 0)$ for $q \leq 3$, but $c_q^{\mathcal{A}} = (3 + q, 0)$ for $q > 3$ (and each $d_q^{\mathcal{A}}$ is the immediate successor of $c_q^{\mathcal{A}}$, necessarily). This is still a model of T^c , but in the gap between 3 and 6, where no constants were used, the substructure $\mathcal{C}^c \subseteq \mathcal{A}^c$ containing those

elements of \mathcal{A}^c which are $\leq (4, 0)$ and those which are $\geq (5, 1)$ is another model of T^c , but not an elementary substructure of \mathcal{A}^c . So in this situation, constants do not suffice; the new relation symbol **Adj** was necessary to yield model completeness.

4 Relatively Decidable Isomorphism Types

Our ultimate goal is to characterize those theories that are relatively decidable – that is, those theories T such that every model of T with domain ω is relatively decidable. This question remains open, and we discuss it in Section 5. As a step towards that goal, we prove the analogous result here for isomorphism types, characterizing those structures \mathcal{A} such that every copy of \mathcal{A} is relatively decidable. In Corollary 4.3 we will derive a near-effectiveness result regarding those \mathcal{A} for which not every copy is relatively decidable.

Theorem 4.1 *Let \mathcal{A} be a countable structure in a finite relational language, and assume that every structure (on the domain ω) isomorphic to \mathcal{A} is relatively decidable. Then there is a finite tuple $\vec{a} \in A$ such that (\mathcal{A}, \vec{a}) is uniformly relatively decidable.*

The restriction requiring the language to be finite and relational will be removed in Corollary 4.2.

Proof. We assume there is no tuple \vec{a} such that (\mathcal{A}, \vec{a}) is uniformly relatively decidable, and construct $\mathcal{B} \cong \mathcal{A}$ such that $\Phi_e^{\Delta(\mathcal{B})} \neq E(\mathcal{B})$ for every e .

We build \mathcal{B} by constructing a generic permutation g of ω and defining \mathcal{B} with domain ω such that $g : \mathcal{B} \rightarrow \mathcal{A}$ is an isomorphism. We follow the method in Ash, Knight, Manasse and Slaman [1], except we work with finitary formulas rather than in a countable fragment of $\mathcal{L}_{\omega_1, \omega}$. In fact, Corollary 4.3 will show that full-scale forcing is hardly necessary to prove this theorem, but the method is familiar to many readers and will be readily understood.

A condition is a finite partial 1-to-1 function from an initial segment of ω into ω . We define extension by $q \leq p$ if and only if $q \supseteq p$. We expand the language L of \mathcal{A} to L^f by adding a function symbol f to denote the generic. We let $L^f(\mathcal{A})$ denote the further extension when we add constants for the elements of the domain ω of \mathcal{A} . For a sentence $\varphi \in L^f(\mathcal{A})$ and a condition p , we define $p \Vdash \varphi$ as in [1].

Since it is dense to add new elements to the domain and range of a condition, the resulting generic will be a permutation of ω . Given a permutation g , we define the corresponding model \mathcal{B}_g by

$$\mathcal{B}_g \models \varphi(b_1, \dots, b_k) \Leftrightarrow \mathcal{A} \models \varphi(g(b_1), \dots, g(b_k))$$

for all atomic formulas $\varphi(\vec{x})$ and tuples \vec{b} . We use $g(\vec{b})$ to denote the tuple $\langle g(b_1), \dots, g(b_k) \rangle$. Note that by definition, $g : \mathcal{B}_g \rightarrow \mathcal{A}$ is an isomorphism.

In a similar way, each condition p determines a finite part of the atomic diagram of an isomorphic copy of \mathcal{A} . We use $\Delta(p)$ to denote the finite set of atomic and negated atomic sentences already determined by p . For a Turing functional Φ , we write $\Phi^{\Delta(p)}(n) \downarrow$ if the computation with finite oracle $\Delta(p)$ halts and only queries atomic and negated atomic facts determined by p .

For a condition p , we let \vec{d}_p denote the domain of p and \vec{a}_p denote the range p . We say that a structure \mathcal{C} extends p if there is an isomorphism $h : \mathcal{C} \rightarrow \mathcal{A}$ such that $h \upharpoonright \vec{d}_p = p$. Note that h is an isomorphism from (\mathcal{C}, \vec{d}_p) onto (\mathcal{A}, \vec{a}_p) and that if $\Phi^{\Delta(p)}(n) \downarrow$ and \mathcal{C} extends p , then $\Phi^{\Delta(\mathcal{C})}(n) = \Phi^{\Delta(p)}(n)$.

Let p be a condition and let $(\mathcal{E}, \vec{e}) \cong (\mathcal{A}, \vec{a}_p)$. By permuting the domain of \mathcal{E} , there is a structure \mathcal{C} which extends p and satisfies $(\mathcal{C}, \vec{d}_p) \cong (\mathcal{E}, \vec{e})$. That is, we can transform any isomorphic copy of (\mathcal{A}, \vec{a}_p) into a copy of \mathcal{A} which extends p . Moreover, we can find (\mathcal{C}, \vec{d}_p) uniformly in (\mathcal{E}, \vec{e}) and p .

We need to build our generic g so that \mathcal{B}_g satisfies

$$\mathcal{R}_e : \Phi_e^{\Delta(\mathcal{B}_g)} \neq E(\mathcal{B}_g)$$

for each e . We say that a condition p satisfies \mathcal{R}_e if one of the following two conditions holds.

(C1) There is an n such that for every \mathcal{C} extending p , $\Phi_e^{\Delta(\mathcal{C})}(n) \uparrow$.

(C2) There is a formula $\varphi(\vec{x})$ such that either

$$\Phi_e^{\Delta(p)}(\ulcorner \varphi(\vec{d}_p) \urcorner) = 0 \text{ and } p \Vdash \varphi(\vec{d}_p)$$

or

$$\Phi_e^{\Delta(p)}(\ulcorner \varphi(\vec{d}_p) \urcorner) = 1 \text{ and } p \Vdash \neg \varphi(\vec{d}_p)$$

To see why satisfying (C1) or (C2) is sufficient, let g be a generic extending p . If (C1) holds for p , then $\Phi_e^{\mathcal{B}_g}$ is not total and hence satisfies \mathcal{R}_e . If (C2)

holds for p , then $\Phi_e^{\Delta(\mathcal{B}_g)}(\ulcorner \varphi(\vec{d}_p) \urcorner) = 0$ if and only if $\varphi(\vec{d}_p) \in E(\mathcal{B}_g)$, and hence \mathcal{B}_g satisfies \mathcal{R}_e .

To complete our theorem, it suffices to show that the set of conditions satisfying each \mathcal{R}_e is dense. Fixing a condition p , we now show that there exists an extension of p which satisfies \mathcal{R}_e , by considering three different cases and then showing that one of these three cases must hold.

Case 1. There is a structure \mathcal{C} extending p , a formula $\varphi(\vec{y}, \vec{z})$ and a tuple $\vec{c} \in \mathcal{C}$ disjoint from \vec{d}_p such that

$$\Phi_e^{\Delta(\mathcal{C})}(\ulcorner \varphi(\vec{d}_p, \vec{c}) \urcorner) = 0 \text{ and } \mathcal{C} \models \varphi(\vec{d}_p, \vec{c}).$$

In this case, fix an isomorphism $h : \mathcal{C} \rightarrow \mathcal{A}$ such that $h \upharpoonright \vec{d}_p = p$. Let q be a condition extending p such that the elements of \vec{c} are in the domain of q and $q(\vec{c}) = h(\vec{c})$, and furthermore, q agrees with h on a long enough initial segment of ω (i.e. the domain of \mathcal{C}) that every atomic or negated atomic fact queried in the computation from $\Delta(\mathcal{C})$ is determined by q . Thus, we have that $q \leq p$, \mathcal{C} is an extension of q , and $\Phi_e^{\Delta(q)}(\ulcorner \varphi(\vec{d}_p, \vec{c}) \urcorner) = 0$. Since the elements of \vec{d}_p and \vec{c} are contained in \vec{d}_q , we can view $\varphi(\vec{y}, \vec{z})$ as a formula $\varphi(\vec{x})$ such that $\varphi(\vec{d}_p, \vec{c})$ is identical to $\varphi(\vec{d}_q)$. In this notation, we have $\mathcal{C} \models \varphi(\vec{d}_q)$ and $\Phi_e^{\Delta(q)}(\ulcorner \varphi(\vec{d}_q) \urcorner) = 0$.

By the definition of q , the isomorphism $h : \mathcal{C} \rightarrow \mathcal{A}$ satisfies $h \upharpoonright \vec{d}_q = q$. Therefore, since $\mathcal{C} \models \varphi(\vec{d}_q)$, we have $\mathcal{A} \models \varphi(\vec{a}_q)$. Let g be a generic extending q . By the definition of \mathcal{B}_g , g is an isomorphism from \mathcal{B}_g to \mathcal{A} such that $g \upharpoonright \vec{d}_q = q$. Therefore, since $\mathcal{A} \models \varphi(\vec{a}_q)$, we have $\mathcal{B}_g \models \varphi(\vec{d}_q)$.

\mathcal{B}_g is a generic structure, so the fact that $\mathcal{B}_g \models \varphi(\vec{d}_q)$ must be forced. Fix a condition $r \leq q$ such that $r \Vdash \varphi(\vec{d}_q)$. Because $r \leq q$, we maintain $\Phi_e^{\Delta(r)}(\ulcorner \varphi(\vec{d}_q) \urcorner) = 0$, and therefore r is the desired extension of p satisfying \mathcal{R}_e through condition (C2).

Case 2. There is a structure \mathcal{C} extending p , a formula $\varphi(\vec{y}, \vec{z})$ and a tuple $\vec{c} \in \mathcal{C}$ (disjoint from \vec{d}_p) such that

$$\Phi_e^{\Delta(\mathcal{C})}(\ulcorner \varphi(\vec{d}_p, \vec{c}) \urcorner) = 1 \text{ and } \mathcal{C} \models \neg \varphi(\vec{d}_p, \vec{c}).$$

This case proceeds as in the previous case except we work with the formula $\neg \varphi$ in place of φ .

Case 3. There is a condition $q \leq p$ and an n such that $\Phi_e^{\Delta(\mathcal{C})}(n) \uparrow$ for all structures \mathcal{C} extending q . In this case, q is already an extension of p satisfying \mathcal{R}_e through condition (C1).

To finish the proof, we need to show that we must be in one of the three cases above. Assume that none of the three cases apply for some fixed condition p with \mathcal{R}_e . Under this assumption, we now describe a uniform procedure that, given a structure $(\mathcal{E}, \vec{e}) \cong (\mathcal{A}, \vec{a}_p)$, a formula $\varphi(\vec{x}, \vec{e})$ and a tuple \vec{u} from \mathcal{E} , determines whether $\mathcal{E} \models \varphi(\vec{u}, \vec{e})$. (Without loss of generality, \vec{e} and \vec{u} are disjoint.) So, fix (\mathcal{E}, \vec{e}) , $\varphi(\vec{x}, \vec{e})$ and \vec{u} . As noted above, we can use permutations of finite initial segments of ω to transform (\mathcal{E}, \vec{e}) into (\mathcal{C}, \vec{d}_p) such that $(\mathcal{E}, \vec{e}) \cong (\mathcal{C}, \vec{d}_p)$ and \mathcal{C} extends p . Let \vec{c} be the image of \vec{u} under the appropriate permutation. We have reduced our question to the following: describe a procedure that, given a structure \mathcal{C} extending p , a formula $\varphi(\vec{x}, \vec{d}_p)$ and a tuple \vec{c} from \mathcal{C} , determines, uniformly in an oracle for $\Delta(\mathcal{C})$, whether $\mathcal{C} \models \varphi(\vec{c}, \vec{d}_p)$.

For simplicity, we assume that, for each n , the Gödel coding of atomic sentences about elements of the domain ω numbers all sentences about $\{0, \dots, n\}$ before it comes to any sentence about $(n+1)$. We then define the numbers l_n (effectively) so that the restriction $\Delta(\mathcal{C}) \upharpoonright \{0, \dots, n\}$ of the atomic diagram of a structure \mathcal{C} to a finite initial segment $\{0, \dots, n\}$ of its domain is a string in 2^{l_n} .

As described above, we may start with a structure $(\mathcal{C}, \vec{d}_p) \cong (\mathcal{A}, \vec{a}_p)$. Given a formula $\varphi(\vec{x})$ and a tuple \vec{c} from \mathcal{C} , we determine whether $\mathcal{C} \models \varphi(\vec{c})$ by searching for the following, using our $\Delta(\mathcal{C})$ -oracle:

- an $n \geq \max(\vec{c}, \vec{d}_p)$ and a $\sigma \in 2^{l_n}$ such that σ extends the characteristic function of $\Delta(\mathcal{C}) \upharpoonright (\vec{c}, \vec{d}_p)$ and $\Phi_e^\sigma(\ulcorner \varphi(\vec{c}, \vec{d}_p) \urcorner) \downarrow$; and
- a tuple $\vec{b} \in \mathcal{C}^n$ of distinct domain elements of \mathcal{C} which includes all of \vec{c} and all of \vec{d}_p , such that the map $\rho(i) = b_i$ fixes \vec{c} and \vec{d}_p pointwise, and such that applying the map ρ to the elements $\{0, \dots, n\}$ sends σ to $\Delta(\mathcal{C}) \upharpoonright \vec{b}$.

We argue below that such an n , σ and \vec{b} must exist, and that for every such collection of elements, we must have $\Phi_e^\sigma(\ulcorner \varphi(\vec{c}, \vec{d}_p) \urcorner)$ equal to 1 if $\mathcal{C} \models \varphi(\vec{c}, \vec{d}_p)$ and equal to 0 if not. First, we give the intuition for this search.

Each σ that we find represents some atomic diagram on a set of n elements, in the given language. (Of course, it is important for this language

to be finite and relational; otherwise such a σ might need to be infinitely long.) Now Φ_e may converge using oracles σ which have nothing to do with \mathcal{C} or with \mathcal{A} . However, if we also find a tuple \vec{b} in \mathcal{C} and a bijection ρ as described, then we know that the configuration described by σ does actually occur (with the elements \vec{b} , not necessarily with $\{0, \dots, n\}$) in \mathcal{C} . Moreover, it occurs with the particular \vec{c} and \vec{d}_p that matter to us, since ρ fixes these elements. If we extend ρ to a permutation f of ω , then we get a structure $(\mathcal{D}, \vec{c}, \vec{d}_p)$ on the domain ω , defined so that f is an isomorphism from this structure onto $(\mathcal{C}, \vec{c}, \vec{d}_p)$. Thanks to our specifications regarding ρ , the atomic diagram $\Delta(\mathcal{D})$ must restrict to σ . Therefore, $\Phi_e^{\Delta(\mathcal{D})}(\ulcorner \varphi(\vec{c}, \vec{d}_p) \urcorner)$ converges to the value $\Phi_e^\sigma(\ulcorner \varphi(\vec{c}, \vec{d}_p) \urcorner)$ that we found. But by the failure of Cases 1 and 2, this value must be “correct,” in \mathcal{D} , i.e., must tell us accurately whether $\mathcal{D} \models \varphi(\vec{c}, \vec{d}_p)$. Since the isomorphism f from \mathcal{D} onto \mathcal{C} fixes \vec{c} and \vec{d}_p , it also tells us whether $\mathcal{C} \models \varphi(\vec{c}, \vec{d}_p)$, which is what we wanted to know. The last point is that some such n , σ and \vec{b} do exist: if not, then Case 3 would not have failed. Therefore, the search described above must eventually find such an n , σ and \vec{b} , and from them we can decide whether $\mathcal{C} \models \varphi(\vec{c}, \vec{d}_p)$ or not, proving the uniform relative decidability of (\mathcal{A}, \vec{d}_p) , contrary to our original hypothesis.

Now for the full details. We know $(\mathcal{C}, \vec{d}_p) \cong (\mathcal{A}, \vec{a}_p)$, so fix an isomorphism $f : (\mathcal{C}, \vec{d}_p) \rightarrow (\mathcal{A}, \vec{a}_p)$. Let $q \subseteq f$ be a partial function on an initial segment of ω long enough that $(\vec{c}, \vec{d}_p) \subseteq \text{dom}(q)$. Let $\vec{a} = q(\vec{c})$ in \mathcal{A} and note that q extends p .

Since Case 3 fails, there is some structure \mathcal{D} , isomorphic to \mathcal{A} via an extension of q , such that $\Phi_e^{\Delta(\mathcal{D})}(\ulcorner \varphi(\vec{c}, \vec{d}_p) \urcorner) \downarrow$. Indeed, we may take this \mathcal{D} to be generic, since only a finite initial segment of $\Delta(\mathcal{D})$ is actually used in this computation. By genericity, and since Cases 1 and 2 fail, the output of the computation is the correct answer about whether $\mathcal{D} \models \varphi(\vec{c}, \vec{d}_p)$. Fix a sufficiently large n that, with $\sigma = \Delta(\mathcal{D}) \upharpoonright l_n$, we have $\Phi_e^\sigma(\ulcorner \varphi(\vec{c}, \vec{d}_p) \urcorner) \downarrow$. Now \mathcal{C} and \mathcal{D} are both isomorphic to \mathcal{A} via isomorphisms extending q , so we can take $\rho : \{0, \dots, n\} \rightarrow \mathcal{C}$ to be the restriction of the resulting isomorphism from \mathcal{D} onto \mathcal{C} . The elements \vec{b} in the image of this ρ are the necessary tuple from \mathcal{C} , since ρ , being the restriction of an isomorphism, must send σ to $\Delta(\mathcal{C}) \upharpoonright \vec{b}$. This proves that there do exist an n , σ , and \vec{b} as described in the program for deciding the elementary diagram of \mathcal{C} , so our search there must eventually terminate.

Of course, the n , σ , and \vec{b} found by the program's search are not necessarily the ones determined above using the isomorphism f . We now prove that, for *every* n , σ , and \vec{b} satisfying the conditions in the search, the value of $\Phi_e^\sigma(\ulcorner \varphi(\vec{c}, \vec{d}_p) \urcorner)$ correctly describes whether $\mathcal{C} \models \varphi(\vec{c}, \vec{d}_p)$ or not. The argument was summarized when we gave the intuition for the proof. Define the map $\rho : \{0, \dots, n\} \rightarrow \vec{b}_i$ sending i to the coordinate b_i in \vec{b} . By the conditions in the search, ρ is the identity map on \vec{c} and on \vec{d}_p , and replacing each i by $\rho(i)$ in the formulas named by the Gödel numbers $\leq l_n$ converts σ into $\Delta(\mathcal{C}) \upharpoonright \vec{b}$. Therefore, ρ extends to an isomorphism h from some structure \mathcal{D} onto \mathcal{C} , where σ is an initial segment of $\Delta(\mathcal{D})$, and we may assume this \mathcal{D} to be generic, since only a finite initial segment of $\Delta(\mathcal{D})$ is prescribed. Moreover, $(\mathcal{D}, \vec{c}, \vec{d}_p) \cong (\mathcal{C}, \vec{c}, \vec{d}_p)$ via this h , by our conditions on ρ . It follows that $\Phi_e^{\Delta(\mathcal{D})}(\ulcorner \varphi(\vec{c}, \vec{d}_p) \urcorner) \downarrow$, with at most the initial segment σ of the oracle being used in the computation. If it outputs 0 and $\mathcal{D} \models \varphi(\vec{c}, \vec{d}_p)$, then the generic structure \mathcal{D} , the formula $\varphi(\vec{y}, \vec{z})$, and the tuple \vec{c} would have shown that Case 1 holds, contrary to our assumption. Likewise, if it outputs 1 and $\mathcal{D} \models \neg\varphi(\vec{c}, \vec{d}_p)$, then Case 2 would have held. Therefore, the output of $\Phi_e^\sigma(\ulcorner \varphi(\vec{c}, \vec{d}_p) \urcorner)$ correctly describes whether $\mathcal{D} \models \varphi(\vec{c}, \vec{d}_p)$. But the isomorphism h shows that $\mathcal{D} \models \varphi(\vec{c}, \vec{d}_p)$ if and only if $\mathcal{C} \models \varphi(\vec{c}, \vec{d}_p)$, so the conclusion of our program does correctly decide the truth of $\varphi(\vec{c}, \vec{d}_p)$ in \mathcal{C} , using only the atomic diagram of (\mathcal{C}, \vec{d}_p) as an oracle. Therefore, under the assumption that all three cases fail at p , \mathcal{A} would indeed have a uniformly relatively decidable expansion by constants. \blacksquare

We now remove the restriction to finite relational languages.

Corollary 4.2 *Let \mathcal{A} be a countable structure in a computable language, such that all copies of \mathcal{A} on the domain ω are relatively decidable. Then there is a finite tuple $\vec{a} \in A$ such that (\mathcal{A}, \vec{a}) is uniformly relatively decidable.*

Proof. To apply Theorem 4.1, we need the language to be finite and relational. This is a straightforward application of the well-known theorem of Hirschfeldt, Khossainov, Shore, and Slinko in Appendix A of [7], which takes an arbitrary automorphically nontrivial structure in an arbitrary computable language and produces a directed graph with exactly the same computable-model-theoretic properties. Applying this to the language of \mathcal{A} gives a symmetric irreflexive graph G in the language with equality and a single binary relation symbol, to which we can apply Theorem 4.1.

To explain more fully: [7] provides an *effective bi-interpretation* between \mathcal{A} and G , in the sense of [9, §5.1]. The construction can also be viewed as a use of *computable functors* between the category of presentations of \mathcal{A} and the category of presentations of G , as defined in [8]; this is shown in [6] to be equivalent to effective bi-interpretation. Using the effective interpretations, we can build from each copy \mathcal{A}' of \mathcal{A} a copy G' of G , with $\Delta(\mathcal{A}') \equiv_T \Delta(G')$, and conversely. It follows that every copy G' of G is relatively decidable: given any sentence $\alpha(\vec{g})$ about elements \vec{g} of G' , we can translate this into a sentence about finitely many elements \vec{a} in the copy \mathcal{A}' of \mathcal{A} built from G' , and then determine the truth of the sentence in \mathcal{A}' by applying the procedure for deciding $E(\mathcal{A}')$ from $\Delta(\mathcal{A}')$ (using $\Delta(G')$ to decide $\Delta(\mathcal{A}')$, since the interpretation is effective). Thus every $G' \cong G$ is relatively decidable, and Theorem 4.1 gives a finite tuple \vec{g} in G such that (G, \vec{g}) is uniformly relatively decidable. This tuple \vec{g} corresponds to a finite tuple \vec{a} from \mathcal{A} (namely, the elements of the tuples in \mathcal{A} that interpret the elements of \vec{g}), and the uniform decision procedure for copies (\mathcal{A}', \vec{a}') of (\mathcal{A}, \vec{a}) begins by building the structure G' corresponding to \mathcal{A}' and identifying the tuple \vec{g}' in G' (using \vec{a}'), which can be done uniformly using the effective interpretation. Then translate the given sentence about (\mathcal{A}', \vec{a}') into an equivalent one about (G', \vec{g}') and apply the uniform decision procedure for copies of (G, \vec{g}) to determine its truth in (G', \vec{g}') .

It remains to consider the automorphically trivial models \mathcal{A} of T . Unsurprisingly, this is simple. By definition, the condition means that there is a finite tuple \vec{a} from \mathcal{A} such that every permutation of ω which fixes \vec{a} pointwise is an automorphism of \mathcal{A} . In this case we take that \vec{a} as our tuple of constants. Now for every $(\mathcal{C}, \vec{c}) \cong (\mathcal{A}, \vec{a})$, we can find an isomorphism from \mathcal{A} onto \mathcal{C} , effectively in $\Delta(\mathcal{C}, \vec{c})$: just map each a_i to c_i , extend this to a permutation of a finite initial segment of ω , and then extend by the identity map to a permutation $f_{\mathcal{C}}$ of all of ω . This $f_{\mathcal{C}}$ is an isomorphism from \mathcal{A} onto \mathcal{C} , computable uniformly in the atomic diagram (in fact, \vec{c} is all that is needed), and so we may use it to decide the truth of formulas $\varphi(\vec{c}, \vec{d})$ in \mathcal{C} by determining the truth of $\varphi(\vec{a}, f(\vec{d}))$ in the relatively decidable structure \mathcal{A} . (The oracle $\Delta(\mathcal{A})$ for running the decision procedure for \mathcal{A} is at hand, since we know both $\Delta(\mathcal{C})$ and the isomorphism f .) ■

The proof in Theorem 4.1 lends itself to an effective construction establishing the following corollary.

Corollary 4.3 *Suppose that \mathcal{A} is a countable structure on the domain ω such that, for every tuple \vec{a} of constants from \mathcal{A} , (\mathcal{A}, \vec{a}) fails to be uniformly relatively decidable. Then there exists a structure $\mathcal{B} \cong \mathcal{A}$ which is low relative to \mathcal{A} and is not relatively decidable: it satisfies $(\Delta(\mathcal{B}))' \leq_T (\Delta(\mathcal{A}))'$ and $E(\mathcal{B}) \not\leq_T \Delta(\mathcal{B})$.*

Proof. This is a finite-injury construction below a $\Delta(\mathcal{A})$ -oracle, using the result (established in the proof of Theorem 4.1) that, for an \mathcal{A} which is not uniformly relatively decidable, the three Cases cannot all fail. Of course, we may assume that \mathcal{A} itself is relatively decidable, since otherwise the corollary is trivial. The goal is to construct a permutation f of ω such that we can pull \mathcal{A} back via f to form the desired structure \mathcal{B} with $f : \mathcal{B} \rightarrow \mathcal{A}$ an isomorphism. At each stage s of the construction, we will compute (using $\Delta(\mathcal{A})$) a finite injective portion p_s of f . These maps p_s will not all be compatible, but with $f = \lim_s p_s$ we will satisfy the theorem. (If the p_s were all compatible, then f would be computable from $\Delta(\mathcal{A})$, and so would $\Delta(\mathcal{B})$. We conjecture that there exist structures \mathcal{A} for which all copies \mathcal{B} with $\Delta(\mathcal{B}) \leq_T \Delta(\mathcal{A})$ are relatively decidable, in which case this would be impossible.)

We mix the requirements \mathcal{R}_e from Theorem 4.1 with standard lowness requirements for f relative to $\Delta(\mathcal{A})$:

$$\mathcal{L}_e : (\exists^\infty s \Phi_{e,s}^{p_s}(e) \downarrow) \implies \Phi_e^f(e) \downarrow,$$

and with surjectivity requirements \mathcal{S}_y , each ensuring that the element y in the domain ω of \mathcal{A} lies in the image of all maps p_s with $s \geq y$. We fix a computable ordering \prec , in order type ω , of the injective elements of $\omega^{<\omega}$.

Naturally p_0 is the empty map. At stage $s + 1$, we have p_s from the preceding stage, and we build p_{s+1} by going through the requirements in order, first \mathcal{L}_0 , then \mathcal{S}_0 , then \mathcal{R}_0 , then \mathcal{L}_1 , and so on up to \mathcal{R}_s . We start with p as the empty string and extend it as we encounter each of these requirements; p_{s+1} will be the ultimate result. For a requirement \mathcal{L}_e , we check whether $\Phi_{e,s}^q(e)$ converges for any of the first (under \prec) s -many injective extensions $q \leq p$ in $\omega^{<\omega}$. If so, then we extend p to the \prec -least such q ; if not, then p stays unchanged. For a requirement \mathcal{S}_y , we leave p unchanged if $y \in \text{range}(p)$, or else extend p to map the next number $|p|$ to y , thus satisfying \mathcal{S}_y .

For a requirement \mathcal{R}_e , we search through the first (under \prec) s extensions $q \leq p$ in $\omega^{<\omega}$, and through all formulas $\alpha(\vec{x})$ and all $\vec{b} \in (\text{dom}(q))^{<\omega}$ with

$\ulcorner \alpha(\vec{b}) \urcorner \leq s$, checking each to see whether $\Phi_e^{\Delta(q^{-1}(\mathcal{A}))}(\ulcorner \alpha(\vec{b}) \urcorner) \downarrow$. (Here $q^{-1}(\mathcal{A})$ is the finite relational structure we get by pulling back the structure \mathcal{A} via the finite map q .) Since we have an oracle for $\Delta(\mathcal{A})$, this is computable. The oracle also allows us to compute $E(\mathcal{A})$, since by assumption \mathcal{A} itself is relatively decidable, and so, for each q where convergence occurs, we may check whether it agrees with the truth in \mathcal{A} of $\alpha(q(\vec{b}))$. If we find any q here for which either

$$\begin{aligned} & [\Phi_e^{\Delta(q^{-1}(\mathcal{A}))}(\ulcorner \alpha(\vec{b}) \urcorner) \downarrow = 0 \text{ and } \mathcal{A} \models \alpha(q(\vec{b}))], \\ \text{or } & [\Phi_e^{\Delta(q^{-1}(\mathcal{A}))}(\ulcorner \alpha(\vec{b}) \urcorner) \downarrow = 1 \text{ and } \mathcal{A} \models \neg \alpha(q(\vec{b}))], \end{aligned}$$

then we set p to equal the q in the least such pair $\langle q, \ulcorner \alpha(\vec{b}) \urcorner \rangle$. If none of the q checked here have either of these properties, then we extend p to the \prec -least element p' such that none of the finitely many q we found with $\Phi_e^{\Delta(q^{-1}(\mathcal{A}))}(\ulcorner \alpha(\vec{b}) \urcorner) \downarrow$ extends p' . This completes the construction.

We claim that, for every requirement \mathcal{P} , the string $p_{s,\mathcal{P}}$ chosen by that requirement at stage s will stabilize on some $p_{\mathcal{P}}$ as $s \rightarrow \infty$. Since $p_{s,\mathcal{P}}$ is an initial segment of p_{s+1} , this will prove that $f = \lim_s p_s$ exists. (The \mathcal{S} -requirements ensure that $|p_s| \geq s$ for all s , so $\text{dom}(f) = \omega$.) By induction we assume stabilization of the string $p_{\mathcal{P}}$ produced by the next-higher-priority requirement. For \mathcal{S} -requirements, the result is clear, and likewise for a requirement \mathcal{L}_e , which will eventually find the \prec -least q extending $p_{\mathcal{P}}$ with $\Phi_e^q(e) \downarrow$ (if any exists), and will choose this q as p_{s+1,\mathcal{L}_e} forever after; or else will choose $p_{\mathcal{P}}$ forever if there is no such q . For a requirement \mathcal{R}_e , if there is any $q \leq p_{\mathcal{P}}$ and any $\alpha(\vec{b})$ that yield a disagreement between $\Phi_e^{\Delta(q^{-1}(\mathcal{A}))}(\ulcorner \alpha(\vec{b}) \urcorner)$ and $E(\mathcal{A})$, then we will eventually find the least such pair and will choose that q as p_{s,\mathcal{R}_e} forever after. If there is no such q and $\alpha(\vec{b})$, then it follows from the proof of Theorem 4.1 that the Case 3 described there must hold: some $n \in \omega$ and $q \leq p$ have $\Phi_e^{\Delta(f^{-1}(\mathcal{A}))}(n) \uparrow$ for every permutation f of ω extending q . In this case, our construction will eventually light on the \prec -least such q and will choose it as p_{s,\mathcal{R}_e} forever after. This completes the induction, allowing us to fix the permutation $f = \lim_s p_s$ and to define \mathcal{B} to be the unique structure with domain ω from which f is an isomorphism onto \mathcal{A} . The \mathcal{S} -requirements ensure that f has image ω , and the \mathcal{L} -requirements show f to be low relative to $\Delta(\mathcal{A})$ in the usual way, so that $(\Delta(\mathcal{B}))' \leq_T (\Delta(\mathcal{A}))'$. Finally, if there were a (least) e with $\Phi_e^{\Delta(\mathcal{B})} = E(\mathcal{B})$, then the corresponding \mathcal{R}_e would never have settled on a limit $p_{\mathcal{R}_e}$: once the next-higher string $p_{\mathcal{P}}$

had stabilized, \mathcal{R}_e would never have found a string giving a disagreement, but also must never have found a $q \leq p_{\mathcal{P}}$ above which, for some $\alpha(\bar{b})$, no convergence ever occurred. (If there were such a string q , \mathcal{R}_e would eventually have chosen it, so that $f \supseteq q$; but then $\Phi_e^{\Delta(\mathcal{B})} = E(\mathcal{B})$ contradicts the claim that no convergence above this q ever occurred.) So indeed our \mathcal{B} fails to be relatively decidable. \blacksquare

5 Model Precompleteness

With the preceding results proven, it remains to ask whether some theorem similar to Theorem 4.1 holds in the situation where every model of a c.e. theory T is relatively decidable. Here we will offer a conjecture about this situation, involving a weakened version of model completeness, along with a reference to the work [5], wherein it has recently been resolved.

Definition 5.1 For a theory T in a language \mathcal{L} , let \mathcal{L}^* be an expansion of \mathcal{L} by finitely many constants. An expansion T^* of T by new formulas in the language \mathcal{L}^* is a *conservative expansion* if, for every \mathcal{L} -formula φ such that $T^* \vdash \varphi$, we already had $T \vdash \varphi$.

For example, the theory T of (ω, S) (in the language with just S and $=$) has an expansion T^* by a constant c_0 and a formula $\forall x c_0 \neq Sx$. Likewise, we have a conservative expansion of the theory of dense linear orders with end points by two constants, along with formulas stating that the two constants are the end points. These both satisfy the following definition.

Definition 5.2 A theory T in a language \mathcal{L} is *model precomplete* if there exist finitely many constants c_1, \dots, c_m not in \mathcal{L} , and a set T^* of formulas in the expanded language \mathcal{L}^* with these constants, such that the theory

$$\text{Cn}(T \cup T^*)$$

is model complete and is a conservative expansion of T . (Model complete theories are model precomplete, by using the empty tuple of constants.)

The theories of the structures \mathcal{A}_n in Section ?? show that it is important to allow an arbitrary finite number of constants. In general, one thinks of the formulas in T^* as definitions of the new constants. One might hope for

T^* to be a finite set (equivalently, a single formula), but here we allow it to be infinite. All examples just mentioned are relatively decidable theories, and it is clear that every model precomplete theory T is relatively decidable, although not necessarily uniformly so.

Conjecture 5.3 *A theory is model precomplete if and only if it is relatively decidable.*

As noted above, the forward direction is immediate. However, after learning of this conjecture, Matthew Harrison-Trainor has now disproven the backwards direction, in a particularly strong way, in [5], by demonstrating that the index set of the relatively decidable c.e. theories is m -complete Π_1^1 . Since model precompleteness is an arithmetic property, this clearly refutes Conjecture 5.3, and shows more broadly that there is no classification of the relatively decidable theories nearly as simple as had been hoped. We salute this result, and its author.

References

- [1] 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.
- [2] C.C. Chang & H.J. Keisler, *Model Theory*, third edition (Amsterdam: Elsevier, 1990).
- [3] S.S. Goncharov, V.S. Harizanov, C. Laskowski, S. Lempp, & C. McCoy; Trivial, strongly minimal theories are model complete after naming constants, *Proceedings of the A.M.S.* **131** (2003) 12, 3901–3912.
- [4] V.S. Harizanov; Pure computable model theory, *Handbook of Recursive Mathematics*, vol. 1 (Amsterdam: Elsevier, 1998), 3–114.
- [5] M. Harrison-Trainor; There is a no simple characterization of the relatively decidable theories, submitted for publication.
- [6] M. Harrison-Trainor, A. Melnikov, R. Miller, & A. Montalbán; Computable functors and effective interpretability, *Journal of Symbolic Logic* **82** (2017) 1, 77–97.

- [7] D.R. Hirschfeldt, B. Khoushainov, R.A. Shore, & A.M. Slinko; Degree spectra and computable dimensions in algebraic structures, *Ann. Pure Appl. Logic* **115** (2002), 71–113.
- [8] R. Miller, B. Poonen, H. Schoutens, & A. Shlapentokh; A computable functor from graphs to fields, *Journal of Symbolic Logic* **83** (2018) 1, 326–348.
- [9] A. Montalbán; Computability theoretic classifications for classes of structures, *Proceedings of ICM 2014* 2 (2014), 79–101.
- [10] R.I. Soare; *Recursively Enumerable Sets and Degrees* (New York: Springer-Verlag, 1987).

DEPARTMENT OF MATHEMATICS
 UNIVERSITY OF SAN FRANCISCO
 2130 FULTON STREET
 SAN FRANCISCO, CA 94117 U.S.A.
E-mail: jcchubb@usfca.edu
Webpage: www.cs.usfca.edu/~jcchubb

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 & COMPUTER SCIENCE
 C.U.N.Y. GRADUATE CENTER
 365 FIFTH AVENUE
 NEW YORK, NEW YORK 10016 U.S.A.
E-mail: Russell.Miller@qc.cuny.edu
Webpage: qcpages.qc.cuny.edu/~rmiller

DEPARTMENT OF MATHEMATICS
 UNIVERSITY OF CONNECTICUT
 341 MANSFIELD ROAD U1009
 STORRS, CT 06269-1009 U.S.A.
E-mail: solomon@math.uconn.edu
Webpage: www.math.uconn.edu/~solomon