

Existentially Closed Exponential Fields

Levon Haykazyan and Jonathan Kirby

version 0.7, October 28, 2019

Abstract

We characterise the existentially closed models of the theory of exponential fields. They do not form an elementary class, but can be studied using positive logic. We find the amalgamation bases and characterise the types over them. We define a notion of independence and show that independent systems of higher dimension can also be amalgamated. We extend some notions from classification theory to positive logic and position the category of existentially closed exponential fields in the stability hierarchy as NSOP₁ but TP₂.

1 Introduction

An *exponential field* is a field F with a homomorphism E from its additive group $\mathbb{G}_a(F)$ to its multiplicative group $\mathbb{G}_m(F)$. Exponential fields are axiomatised by an inductive theory $T_{\text{E-field}}$.

The classical approach to the model theory of algebraic structures begins by identifying extensions and finding the model companion. Such an approach for exponential algebra was initiated by van den Dries (1984). However there is no model companion which traditionally meant a dead end for model theory. So the model theory of exponentiation developed in other directions, notably Wilkie (1996) proved that the real exponential field is model complete and Zilber (2005) studied exponential fields which are *exponentially closed*, which is similar to existential closedness but only takes account of so-called *strong* extensions.

However, it is possible to study the model theory and stability theory of the existentially closed models of an inductive theory even when there is no model companion. Shelah (1975) called this setting model theory of

Kind 3, where Kind 1 is the usual setting of a complete first-order theory, and Kind 2, also known as *Robinson theories*, is like Kind 3 but where there is amalgamation over every subset of every model. (There is also Kind 4, which is now known as homogeneous abstract elementary classes.) More recently there has been interest in positive model theory, which is only slightly more general than Kind 3.

This is the approach taken in this paper, and it means that we look at embeddings and existentially definable sets instead of elementary embeddings and all first-order definable sets. Following Pillay (2000) we call it the *Category of existentially closed models* rather than Kind 3.

Overview of the paper

In Section 2 we give the necessary background on the category $\mathcal{EC}(T)$ of existentially closed models of an inductive theory T .

In Section 3 we characterise the existentially closed models of $T_{\text{E-field}}$ in Theorem 3.5. They are the exponential fields F where each additively free subvariety V of $\mathbb{G}_a(F)^n \times \mathbb{G}_m(F)^n$ has an exponential point, that is, there is a tuple $\bar{a} \in F$ such that $(\bar{a}, E(\bar{a})) \in V$. Additive freeness of a variety, however, is not a first-order property and we show that the class of existentially closed exponential fields is not axiomatisable by a first-order theory.

The theory $T_{\text{E-field}}$ of exponential fields has neither the joint embedding property nor the amalgamation property. In Section 4 we show that the amalgamation holds if and only if the base is algebraically closed (Theorem 4.3). The absence of the joint embedding property means that the category of exponential fields and embeddings splits into a disjoint union of classes where any two structures in the same class can be jointly embedded into a third one. We show that there are 2^{\aleph_0} classes and characterise the theories of each class (Corollary 4.6).

In Section 5 we define a notion of independence: if F is an algebraically closed exponential field, $A, B, C \subseteq F$, then we say that A is independent of B over C and write $A \downarrow_C B$ if $\langle AC \rangle^{\text{EA}}$ and $\langle BC \rangle^{\text{EA}}$ are field-theoretically independent over $\langle C \rangle^{\text{EA}}$. Here $\langle X \rangle^{\text{EA}}$ is the smallest algebraically closed exponential subfield containing X . We then show that independent systems of higher dimension can be amalgamated (Theorem 5.4). In particular, we have 3-amalgamation for independent systems, which is sometimes known as *the independence theorem*.

In Section 6 we study two tree properties of formulas in $T_{\text{E-field}}$: the tree

property of the second kind, TP_2 , and the 1-strong order property, SOP_1 . Both properties are well-known and extensively studied for complete first-order theories. Our definitions are variants that specialise to the appropriate definitions for model-complete theories. We show that the category $\mathcal{EC}(T_{\text{E-field}})$ has TP_2 (Proposition 6.2) by exhibiting a TP_2 -formula. We then show that the independence notion we defined is well-enough behaved to apply the analogue of a theorem of Chernikov and Ramsey, and deduce that $\mathcal{EC}(T_{\text{E-field}})$ is $NSOP_1$ (Theorem 6.5).

In an appendix, we give the necessary technical details of this generalised stability theory for the category of existentially closed models of an inductive theory. In particular, we adapt the work of Chernikov and Ramsey (2016) on $NSOP_1$ -theories.

Acknowledgement

The first named author would like to thank Rémi Jaoui and Rahim Moosa for many discussions around the topics of this paper.

2 Model-theoretic background

We give some background on the model theory of the category of existentially closed models of a (usually incomplete) inductive first-order theory. More details can be found in Hodges (1993), Pillay (2000) and Ben Yaacov and Poizat (2007). The only novelty in this section is the notion of a JEP-refinement of an inductive theory T (Definition 2.11), which is a useful syntactic counterpart to the choice of a monster model.

Remark. The readers familiar with positive logic will notice that the axioms for exponential fields are not only inductive, but *h-inductive* in the sense of Ben Yaacov and Poizat (2007). So we can treat exponential fields in positive logic, which amounts to working with homomorphisms rather than embeddings. In our case this makes no difference, since homomorphisms of fields are embeddings. This is the case whenever all atomic formulas have positive negations modulo the theory. So the setting here is equivalent to positive model theory, with the extra assumption that all atomic formulas have positive negations modulo the theory. This extra assumption, however, is not essential, and everything in this section applies to a general h-inductive

theory (although the treatment of algebraic and definable closures would require more care).

2.1 Existentially closed models of an inductive theory

Recall that a first-order theory T is *inductive* if the union of a chain of models is a model. Equivalently, T is axiomatised by $\forall\exists$ -sentences. We work in the category $\text{Emb}(T)$ whose objects are models of T and whose arrows are embeddings.

Definition 2.1. A model $M \models T$ is *existentially closed* if for all quantifier-free formulas $\phi(\bar{x}, \bar{y})$ and all \bar{a} in M , if there is an extension $M \subseteq B$ such that $B \models T$ and $B \models \exists \bar{x} \phi(\bar{x}, \bar{a})$ then $M \models \exists \bar{x} \phi(\bar{x}, \bar{a})$.

If T is inductive then for any $A \models T$ it is straightforward to build an extension $A \subseteq M$ by a transfinite induction process such that M is an existentially closed model of T .

Embeddings between structures preserve \exists -formulas. That is if $f : A \rightarrow B$ is an embedding, then for any \exists -formula $\phi(\bar{x})$ and $\bar{a} \in A$ we have

$$A \models \phi(\bar{a}) \implies B \models \phi(f(\bar{a})).$$

If the converse of the above implication holds too, then the embedding is called an *immersion*.

Given $A \models T$ and \bar{a} a tuple from A , the existential type $\text{tp}_{\exists}^A(\bar{a})$ is the set of existential formulas $\phi(\bar{x})$ such that $A \models \phi(\bar{a})$. If $A \subseteq B$ then $\text{tp}_{\exists}^A(\bar{a}) \subseteq \text{tp}_{\exists}^B(\bar{a})$, and if the inclusion is an immersion we have equality.

The following well-known equivalent characterisations of existentially closed models are useful.

Fact 2.2. Let T be an inductive theory and let $M \models T$. Then the following conditions are equivalent.

- (i) M is an existentially closed model of T .
- (ii) For any model $B \models T$ any embedding $f : M \rightarrow B$ is an immersion.
- (iii) For every $\bar{a} \in M$, the \exists -type $\text{tp}_{\exists}^M(\bar{a})$ is maximal, that is, if $B \models T$ and $\bar{b} \in B$ are such that $\text{tp}_{\exists}^M(\bar{a}) \subseteq \text{tp}_{\exists}^B(\bar{b})$ then $\text{tp}_{\exists}^B(\bar{b}) = \text{tp}_{\exists}^M(\bar{a})$.

- (iv) For every $\bar{a} \in M$ and \exists -formula $\phi(\bar{x})$ such that $M \models \neg\phi(\bar{a})$, there is an \exists -formula $\psi(\bar{x})$ such that $M \models \psi(\bar{a})$ and $T \models \neg\exists\bar{x}(\phi(\bar{x}) \wedge \psi(\bar{x}))$.

Proof. The equivalence of (i) and (ii) is (Hodges, 1993, Theorem 8.5.6), and the equivalence of (i) and (iv) is (Hodges, 1985, Corollary 3.2.4). The equivalence of (ii) and (iii) is immediate. \square

We write $\mathcal{EC}(T)$ for the full subcategory of $\text{Emb}(T)$ consisting of the existentially closed models of T and all embeddings between them (which by the above fact are immersions).

The model theory and stability theory of existentially closed models of an inductive T is developed analogously to that of a complete first-order theory T' . Existentially-closed models of T correspond to models of T' ; immersions correspond to elementary embeddings, and models of T correspond to substructures of models of T' .

Note however that the category $\mathcal{EC}(T)$ does not determine the theory T completely. We recall the following well-known equivalence.

Fact 2.3. For two inductive theories T_1 and T_2 the following are equivalent.

- (i) The theories T_1 and T_2 have the same universal consequences.
- (ii) Every model of T_1 embeds in a model of T_2 and vice-versa.
- (iii) The existentially closed models of T_1 and T_2 are the same.

Proof. The equivalence of (i) and (ii) is well known. For the equivalence of (i) and (iii) see (Hodges, 1985, Theorem 3.2.3). \square

Theories satisfying these equivalent conditions are called *companions*. Thus $\mathcal{EC}(T)$ (as a subcategory of all L -structures and embeddings) determines T only modulo companions.

2.2 Amalgamation bases

Definition 2.4. An *amalgamation base* for $\text{Emb}(T)$ is a model $A \models T$ such that given any two models $B_1, B_2 \models T$ and embeddings $f_1 : A \rightarrow B_1$ and $f_2 : A \rightarrow B_2$ there is a model $C \models T$ and embeddings $g_1 : B_1 \rightarrow C$ and $g_2 : B_2 \rightarrow C$ such that $g_1 f_1 = g_2 f_2$.

A is a *disjoint amalgamation base* if, furthermore, we can pick the embeddings g_1 and g_2 in such a way that $g_1(B_1) \cap g_2(B_2) = g_1 f_1(A)$. (Some authors including Hodges call this a *strong amalgamation base*.)

Fact 2.5 ((Hodges, 1993, Corollary 8.6.2)). Every existentially closed model of an inductive theory T is a disjoint amalgamation base for $\text{Emb}(T)$.

However there can be amalgamation bases which are not existentially closed models.

There is also a connection with algebraically closed sets.

Definition 2.6. Let M be an existentially closed model of an inductive theory T and let $A \subseteq M$. Then $\text{acl}_\exists(A)$ is the union of all finite \exists -definable subsets of M using parameters from A . If $A = \text{acl}_\exists(A)$ we say that A is \exists -algebraically closed in M .

It is not immediately obvious that $\text{acl}_\exists(\text{acl}_\exists(A)) = \text{acl}_\exists(A)$ and that $\text{acl}_\exists(A)$ is the same if calculated in an existentially closed extension of M . This follows from the fact that if an existential formula $\phi(x, \bar{a})$ defines a finite set of size n in M , then the formula

$$\exists x_1, \dots, x_{n+1} \left[\bigwedge_{i=1}^{n+1} \phi(x_i, \bar{a}) \wedge \bigwedge_{i \neq j} x_i \neq x_j \right]$$

is false in M . It follows by Fact 2.2(iv) that there must be an existential formula $\chi(\bar{y})$ that implies this and holds of \bar{a} .

However, if A is not an amalgamation base then $\text{acl}_\exists(A)$ does, in general, depend on the choice of existentially closed model in which A embeds.

Fact 2.7 ((Hodges, 1993, Corollary 8.6.8)). Let T be an inductive theory and let A be an amalgamation base for $\text{Emb}(T)$. Then A is a disjoint amalgamation base if and only if A is \exists -algebraically closed (in an existentially closed extension M of A).

Amalgamation bases are also very useful for understanding the existential types which are realised in existentially closed models of T . Indeed, finding the amalgamation bases plays much the same role as proving a quantifier-elimination theorem for a complete first-order theory.

Proposition 2.8. Let M, N be existentially closed models of an inductive theory T , and suppose $\bar{a} \in M$ and $\bar{b} \in N$ are tuples. Suppose also that there are amalgamation bases $A \subseteq M$ and $B \subseteq N$ with $\bar{a} \in A$ and $\bar{b} \in B$, and an isomorphism $\theta : A \rightarrow B$ such that $\theta(\bar{a}) = \bar{b}$. Then $\text{tp}_\exists^M(\bar{a}) = \text{tp}_\exists^N(\bar{b})$.

Proof. We have $\text{tp}_{\exists}^A(\bar{a}) = \text{tp}_{\exists}^B(\bar{b})$. Then by (Hodges, 1993, Theorem 8.6.6), since A and B are amalgamation bases, there is a unique way to extend this type to a maximal \exists -type in T , which by Fact 2.2(iii) must be $\text{tp}_{\exists}^M(\bar{a})$ and $\text{tp}_{\exists}^N(\bar{b})$. \square

So the existential type of a tuple \bar{a} in an existentially closed model is determined by how \bar{a} embeds into an amalgamation base.

2.3 JEP

Definition 2.9. The category $\text{Emb}(T)$ has the *joint embedding property (JEP)* if for any two models of T there is a third in which they can both be embedded. We also say that T has the JEP in that case.

We have the following characterisation of JEP.

Lemma 2.10. The category $\text{Emb}(T)$ has the JEP if and only if for every pair of universal sentences ϕ and ψ , if $T \vdash \phi \vee \psi$ then $T \vdash \phi$ or $T \vdash \psi$.

Proof. An easy argument using the method of diagrams and compactness. See (Hodges, 1985, Exercise 3.2.8). \square

When looking at existentially closed models of T , extending T to an inductive theory with JEP plays the role of choosing a completion of T . However, not any extension will do. For example, if T is the theory of fields, the existentially closed models of T are the algebraically closed fields. The completions are then given by fixing the characteristic. However we could also extend T to the inductive theory T' of orderable fields (fields in which -1 is not a sum of squares). Then the existentially closed models of T' are real-closed fields.

We now make a new definition.

Definition 2.11. An inductive extension T' of an inductive theory T is called a *JEP-refinement* of T if T' has the joint embedding property and every existentially closed model of T' is an existentially closed model of T .

We can use amalgamation bases to find these JEP-refinements.

Lemma 2.12. If A is an amalgamation base for $\text{Emb}(T)$ then $T \cup \text{Th}_{\exists}(A)$ is a JEP-refinement of T .

Proof. It is clear that $T \cup \text{Diag}(A)$ has the JEP, where $\text{Diag}(A)$ is the quantifier-free diagram of A . However $T \cup \text{Diag}(A)$ and $T \cup \text{Th}_{\exists}(A)$ have the same consequences in the language of T . It follows from Lemma 2.10 that $T \cup \text{Th}_{\exists}(A)$ has the JEP too.

Then since $T \cup \text{Th}_{\exists}(A)$ is an extension of T by existential sentences only, it is easy to see that every existentially closed model of $T \cup \text{Th}_{\exists}(A)$ is an existentially closed model of T . \square

Note that each existentially closed model M of T is a model of a unique JEP-refinement of T modulo companions. Indeed $M \models T \cup \text{Th}_{\exists}(M)$ which is a JEP-refinement of T . Conversely if T' is a JEP-refinement of T and $M \models T'$, then for every universal sentence ϕ we have $T' \vdash \phi$ if and only if $M \models \phi$. Indeed if $M \models \phi$, then there is an existential sentence ψ such that $T \vdash \neg\psi \vee \phi$ and $M \models \psi$ (by Fact 2.2). Then $T' \not\vdash \neg\psi$ and therefore, by JEP, $T' \vdash \phi$. This shows that the universal consequences of T' are completely determined by M .

2.4 Monster models

As for complete first-order theories, it is notationally convenient (though not essential) to work inside a *monster model* \mathfrak{M} of an inductive theory T with the JEP, that is, a model of some large cardinality κ which is both κ^+ -universal and strongly κ -homogeneous. By κ^+ -*universal* we mean that every $A \models T$ with $|A| < \kappa^+$ embeds in \mathfrak{M} . By *strongly κ -homogeneous* we mean that if A is an amalgamation base for $\text{Emb}(T)$, and f_1, f_2 are embeddings of A into \mathfrak{M} , then there is an automorphism θ of \mathfrak{M} such that $\theta \circ f_1 = f_2$. Then all models considered are submodels of \mathfrak{M} , and maximal (existential) types are the same as orbits of $\text{Aut}(\mathfrak{M})$. In this setting, monster models are often called universal domains. Pillay (2000) calls them *e-universal domains*. As both universality and homogeneity are important, we prefer the terminology *monster model*.

As for complete first-order theories, the universality and homogeneity properties together are equivalent to a saturation property: any existential type $p(\bar{x})$ using parameters from a set A of cardinality less than κ that has a realisation in an extension of \mathfrak{M} already has a realisation in \mathfrak{M} . It follows that monster models of T are existentially closed.

As usual, the existence of monster models depends on stability-theoretic conditions on T or set-theoretic conditions on κ , for example, that κ is

strongly inaccessible. However since everything could be done in the category $\text{Emb}(T)$ or $\mathcal{EC}(T)$, albeit at some cost in notation, we will not worry about moving outside ZFC like this. The only place we really use monster models in this paper is where we introduce independence relations in the usual way, as certain relations on subsets of the monster model. The alternative of treating them as relations on commuting squares in $\text{Emb}(T)$ seems not to be widely known.

If T does not have the JEP, then by choosing a monster model $\mathfrak{M} \models T$ we are in effect choosing a JEP-refinement of T (modulo companions).

3 Existentially closed exponential fields

In this section we characterise the existentially closed exponential fields. For the basics on exponential rings and fields the reader can consult van den Dries (1984), Macintyre (1996) or Kirby (2013).

Definition 3.1. An *exponential field* (or *E-field* for short) is a field F of characteristic zero, together with a homomorphism E from the additive group $\mathbb{G}_a(F)$ to the multiplicative group $\mathbb{G}_m(F)$.

If the field is algebraically closed we call it an *EA-field*.

The reason for excluding characteristic $p > 0$ is that in that case for any element x we have

$$(E(x) - 1)^p = E(x)^p - 1 = E(px) - 1 = E(0) - 1 = 0,$$

and there are no non-trivial nilpotents, so $E(x) = 1$.

We work in the category of exponential fields and their embeddings. Model theoretically this means that we use the language $L_{\text{E-ring}} = \langle +, -, \cdot, 0, 1, E \rangle$ of E -rings, where E is a unary function symbol, and we look at the class of models of the theory $T_{\text{E-field}}$ axiomatised by

- (i) the axioms of fields of characteristic 0;
- (ii) $\forall x, y [E(x + y) = E(x) \cdot E(y)]$;
- (iii) $E(0) = 1$.

Atomic $L_{E\text{-ring}}$ -formulas are exponential polynomial equations, so an exponential field F is *existentially closed* if every finite system of exponential polynomial equations and inequations (with coefficients from F) that has a solution in an extension of F , already has a solution in F .

We will repeatedly use the following well known result.

Fact 3.2 (see e.g. Fuchs (1970)). Divisible Abelian groups are injective in the category of Abelian groups. That is, if A, B, Q are Abelian groups, $A \leq B$ and Q is divisible, then any homomorphism $f : A \rightarrow Q$ extends to B .

From this we immediately derive the following result.

Proposition 3.3. Every exponential field extends to an algebraically closed exponential field.

Proof. Let F be an exponential field and let F' be the algebraic closure of F . Then $\mathbb{G}_a(F)$ is a subgroup of $\mathbb{G}_a(F')$ and $\mathbb{G}_m(F)$ is a subgroup of $\mathbb{G}_m(F')$. So we can view E as a homomorphism from $\mathbb{G}_a(F)$ to $\mathbb{G}_m(F')$. But since F' is algebraically closed, its group of units is divisible. Hence, by Fact 3.2, E extends to a homomorphism from $\mathbb{G}_a(F')$ to $\mathbb{G}_m(F')$. \square

It follows that existentially closed exponential fields are algebraically closed fields. The converse of course is not true as being algebraically closed says nothing about the solubility of exponential equations. We will characterise existentially closed exponential rings geometrically, via *Pillay-Pierce*-style axioms.

Definition 3.4. Let F be an exponential field. Let $V \subseteq \mathbb{G}_a(F)^n \times \mathbb{G}_m(F)^n$ be a subvariety which is irreducible over F , and let (\bar{x}, \bar{y}) be a point of V in a field extension of F , which is generic in V over F . We say that V is *additively free* if \bar{x} satisfies no equation of the form $\sum_{i=1}^n m_i x_i = a$ where $a \in F$ and $m_i \in \mathbb{Z}$, not all zero.

Theorem 3.5. Let F be an exponential field. Then F is existentially closed if and only if for every additively free subvariety $V \subseteq \mathbb{G}_a(F)^n \times \mathbb{G}_m(F)^n$ there is a point $\bar{a} \in F$ such that $(\bar{a}, E(\bar{a})) \in V$.

Proof. We first show the left to right implication. Suppose $V \subseteq \mathbb{G}_a(F)^n \times \mathbb{G}_m(F)^n$ is additively free. Let F' be an algebraically closed field extension of F containing a point $(\bar{a}, \bar{b}) \in V$, generic over F , where $\bar{a} = (a_1, \dots, a_n)$ and $\bar{b} = (b_1, \dots, b_n)$. We extend the exponential homomorphism $E : \mathbb{G}_a(F) \rightarrow$

$\mathbb{G}_m(F)$ to $\mathbb{G}_a(F')$ as follows. Define $E'(a_i) = b_i$ for $i = 1, \dots, n$. Since V is additively free and (\bar{a}, \bar{b}) is generic in V over F , the a_i are \mathbb{Q} -linearly independent over F . So this assignment extends uniquely to a homomorphism E' extending E , defined on the subgroup generated by F and \bar{a} . Now since $\mathbb{G}_m(F')$ is divisible we can extend E' to $\mathbb{G}_a(F')$. Thus we have an exponential field extending F where V has an exponential point. By existential closedness of F , there is an exponential point already in F .

Conversely, assume that every additively free subvariety has an exponential point in F . By taking $V \subseteq \mathbb{G}_a \times \mathbb{G}_m$ given by $p(y) = 0$, a single polynomial in the \mathbb{G}_m -coordinate, we see that F is algebraically closed.

To show existential closedness we need to show that every system of exponential equations and inequations that has a solution in an extension of F already has a solution in F . Two easy observations help to simplify this. Firstly, inequations can be reduced to equations as

$$T_{\text{E-field}} \vdash x \neq 0 \leftrightarrow \exists y(xy = 1).$$

Secondly, iterated exponentials can be substituted by additional variables as, for example,

$$T_{\text{E-field}} \vdash f(\bar{x}, E(\bar{x}), E^2(\bar{x})) = 0 \leftrightarrow \exists \bar{y}[f(\bar{x}, \bar{y}, E(\bar{y})) = 0 \wedge \bar{y} = E(\bar{x})].$$

Thus to show that F is existentially closed it is enough to show that every finite system of polynomial equations $P(\bar{x}, \bar{y}) = 0$ that has an exponential solution $P(\bar{a}, E(\bar{a})) = 0$ in an extension $F' \supseteq F$, already has an exponential solution in F . So let P , F' and $\bar{a} = (a_1, \dots, a_n)$ be as above. Let V be the locus of $(\bar{a}, E(\bar{a}))$ over F , that is, the smallest subvariety of $\mathbb{G}_a^n \times \mathbb{G}_m^n$ containing $(\bar{a}, E(\bar{a}))$ which is defined over F . If V is additively free, we will get an exponential point in F by our assumption. However this may not be the case. Without loss of generality assume that a_1, \dots, a_k are \mathbb{Q} -linearly independent over F and a_{k+1}, \dots, a_n are in the \mathbb{Q} -linear span $\langle Fa_1, \dots, a_k \rangle_{\mathbb{Q}}$. Thus there is a n -tuple $\bar{b} \in F$ and an $n \times k$ matrix A of rational numbers such that

$$\begin{bmatrix} a_1 \\ \vdots \\ a_n \end{bmatrix} = A \begin{bmatrix} a_1 \\ \vdots \\ a_k \end{bmatrix} + \bar{b}.$$

Let N be the least common multiple of the denominators of entries of A . Now let V' be the locus of $(\frac{a_1}{N}, \dots, \frac{a_k}{N}, E(\frac{a_1}{N}), \dots, E(\frac{a_k}{N}))$ over F . Then V'

is additively free. By the assumption, there is a point $\bar{c} \in F$ such that $(\bar{c}, E(\bar{c})) \in V'$. Let

$$\bar{d} = A \begin{bmatrix} Nc_1 \\ \vdots \\ Nc_k \end{bmatrix} + \bar{b}.$$

Then $(\bar{d}, E(\bar{d})) \in V(F)$, and, in particular, $P(\bar{d}, E(\bar{d})) = 0$. So F is existentially closed. \square

We can derive a number of consequences from this characterisation.

Corollary 3.6. If F is an existentially closed exponential field, then the homomorphism $E : \mathbb{G}_a(F) \rightarrow \mathbb{G}_m(F)$ is surjective.

Proof. For any $a \in F^\times$ the variety $W(x, y)$ given by $y = ax$ is additively free and therefore has an exponential point. \square

Next we show that \mathbb{Z} is universally definable in any existentially closed exponential field as the multiplicative stabiliser of the kernel.

Corollary 3.7. If F is an existentially closed exponential field, then for every $a \in F$

$$a \in \mathbb{Z} \text{ iff } F \models \forall x [E(x) = 1 \rightarrow E(ax) = 1].$$

Proof. The left to right implication is clear. For the converse assume that $c \in F \setminus \mathbb{Z}$.

If $c \in \mathbb{Q}$, then write $c = \frac{n}{m}$ where $m > 1$ and n, m are coprime. By the previous corollary there is $b \in F$ such that $E(b)$ is a primitive m -th root of unity. Then $E(mb) = 1$ but $E(cmb) = E(nb) = E(b)^n \neq 1$.

If $c \notin \mathbb{Q}$, then pick $d \in F$ that is distinct from 1 and consider the variety $W(x_1, x_2, y_1, y_2)$ defined by the equations

$$cx_1 = x_2, y_1 = 1, y_2 = d.$$

Since $c \notin \mathbb{Q}$, the variety W is additively free and therefore has an exponential point $(a, b, E(a), E(b)) \in W$. But then $E(a) = 1$ and $E(ca) = E(b) = d \neq 1$. \square

Since a countable set cannot be definable in a saturated model of a first-order theory we also have the following consequence.

Corollary 3.8. The class of existentially closed exponential fields is not first-order axiomatisable. Equivalently, the theory $T_{\text{E-field}}$ does not have a model-companion. \square

Remark. Note that the class of existentially closed exponential fields is axiomatisable by an $L_{\omega_1, \omega}$ -sentence.

Several papers of Boris Zilber and of the second author are devoted to the study of *exponentially closed fields*. These are not existentially closed in the sense of this paper, but they are existentially closed within the category of exponential fields and so-called *strong* extensions. To get the JEP we also insist that the exponential fields have the Schanuel property, and usually that they have standard kernel. Then the strong extensions are roughly those which preserve the Schanuel property and the kernel.

The axiom giving the exponential closedness is then quite similar to our axiom giving existential closedness, specifying that certain varieties have exponential points on them. In that case, the varieties are not just additively free but also *multiplicatively free* (to avoid extending the kernel) and *rotund* (to preserve the Schanuel property). These axioms are also $L_{\omega_1, \omega}$ -expressible.

4 Amalgamation bases

In this section we characterise the amalgamation bases for exponential fields. Here and later we will use the independence notion for algebraically closed fields.

Definition 4.1. Let K be an algebraically closed field, and A, B, C subfields of K with $C \subseteq A$ and $C \subseteq B$. We say that A is independent from B over C (with respect to the theory ACF and as subfields of K) and write $A \downarrow_C^{\text{ACF}} B$ if for every finite tuple \bar{a} from A , the transcendence degree satisfies $\text{td}(\bar{a}/B) = \text{td}(\bar{a}/C)$.

The following is a very well-known special case of the uniqueness of non-forking extensions in stable theories.

Fact 4.2. Suppose C is an algebraically closed field and A, B are two field extensions of C . Then there is, up to isomorphism, a unique way to embed A and B into an extension field K such that $A \downarrow_C^{\text{ACF}} B$, and K is generated as a field by (the image of) $A \cup B$.

Theorem 4.3. The amalgamation bases for $\text{Emb}(T_{\text{E-field}})$ are precisely the algebraically closed exponential fields (EA-fields). Furthermore, they are disjoint amalgamation bases.

Proof. Let F be an EA-field, and let $f_1 : F \rightarrow F_1$ and $f_2 : F \rightarrow F_2$ be two E -field extensions of F . Using Fact 4.2 there is an algebraically closed field K and field embeddings $g_1 : F_1 \rightarrow K$ and $g_2 : F_2 \rightarrow K$ such that $g_1 f_1 = g_2 f_2$ and $F_1 \downarrow_F^{\text{ACF}} F_2$. By identifying F, F_1, F_2 with their images in K we can assume that all embeddings f_i, g_i are inclusions.

We would like to extend $E_1 \cup E_2$ to a homomorphism from $\mathbb{G}_a(K)$ to $\mathbb{G}_m(K)$. Note that $E_1 \cup E_2$ extends to a homomorphism from the group $F_1 + F_2$ generated by F_1 and F_2 to $\mathbb{G}_m(K)$. Indeed, it is enough to see that it is well defined. Assume that $a_i, a'_i \in F_i$ and $a_1 + a_2 = a'_1 + a'_2$. But then $a_1 - a'_1 = a'_2 - a_2 \in F_1 \cap F_2$. By ACF-independence, $F_1 \cap F_2 = F$. It follows that $E_1(a_1 - a'_1) = E_2(a'_2 - a_2)$ and therefore

$$\frac{E_1(a_1)}{E_1(a'_1)} = \frac{E_2(a'_2)}{E_2(a_2)}.$$

But this implies that $E_1(a_1)E_2(a_2) = E_1(a'_1)E_2(a'_2)$.

The final step is to extend this homomorphism from $F_1 + F_2$ to $\mathbb{G}_m(K)$ by the divisibility of the latter, using Fact 3.2. So EA-fields are disjoint amalgamation bases.

For the converse, suppose that F is an E-field which is not an EA-field. Let $F_1 = F^{\text{alg}}$. We can extend the exponential map on F to some exponential map E_1 on F_1 , and, in particular, for every a which is algebraic over F we have that $E_1(a)$ is algebraic over F . Now let $a \in F^{\text{alg}} \setminus F$ and let $F_2 = F(a, t)$, with t transcendental over F . Then there is an exponential map E_2 on F_2 extending that on F such that $E_2(a) = t$. Then F_1 and F_2 cannot be amalgamated over F . \square

The fact that algebraically closed exponential fields are disjoint amalgamation bases allows us to characterise acl_{\exists} in existentially closed models.

Definition 4.4. Suppose that F is an EA-field, and $A \subseteq F$. Write $\langle A \rangle_F^{\text{EA}}$ (or just $\langle A \rangle^{\text{EA}}$ if F is clear) for the smallest EA-subfield of F containing A , that is, the intersection of all the EA-subfields of F containing A .

Corollary 4.5. Let F be an existentially closed exponential field and $A \subseteq F$. Then $\text{acl}_{\exists}(A) = \langle A \rangle_F^{\text{EA}}$.

Proof. It is clear that acl_\exists -closed sets are algebraically closed E -fields. The converse follows immediately from Theorem 4.3 and Fact 2.7. \square

This allows us to characterise the JEP-refinements of $T_{\text{E-field}}$.

Corollary 4.6. (i) If F is an EA-field, then $T_{\text{E-field}} \cup \text{Th}_\exists(F)$ is a JEP-refinement of $T_{\text{E-field}}$.

(ii) There are 2^{\aleph_0} many JEP-refinements of $T_{\text{E-field}}$ (modulo companions), corresponding to the minimal EA-fields (those EA-fields F such that $F = \langle 0 \rangle_F^{\text{EA}}$).

Proof. Part (i) follows from Lemma 2.12.

For part (ii), every existentially closed exponential field K is a model of $T_{\text{E-field}} \cup \text{Th}_\exists(F)$ where $F = \langle 0 \rangle_K^{\text{EA}} = \langle 0 \rangle_F^{\text{EA}}$. On the other hand if F_1 and F_2 are non-isomorphic minimal EA-fields, then they cannot be jointly embedded in an E-field. It follows that $T_{\text{E-field}} \cup \text{Th}_\exists(F_1)$ and $T_{\text{E-field}} \cup \text{Th}_\exists(F_2)$ are not companions.

Any minimal EA-field must be countable, so there are at most 2^{\aleph_0} of them. To construct 2^{\aleph_0} non-isomorphic ones, note that for each infinite sequence $(q_n)_{n \in \mathbb{N}, n \geq 2}$ of nonzero rational numbers, there is an E-field where $E(1)$ is transcendental and $E(E(1)^n) = q_n$ for $n \geq 2$. \square

Remark. There is a similarity with the completions of the theory ACFA of existentially closed fields with an automorphism. There the completions are given by specifying the action of the automorphism on the algebraic closure of the empty set. The main difference here is that the underlying field of a minimal EA-field can have any countable transcendence degree, not just 0 as in ACFA.

Finally in this section we can use amalgamation bases to give an easy proof of instability. Of course we prove stronger results later.

Corollary 4.7. Let T be any JEP-refinement of $T_{\text{E-field}}$. Let A be any subset of any model of T . Then there are $2^{\aleph_0 + |A|}$ maximal existential types over A which are realised in $\mathcal{EC}(T)$. In particular, $\mathcal{EC}(T)$ is not stable (in the sense of type-counting).

Proof. Given A and a variable x we can specify that x is transcendental over A and then we can specify the values of $E(x^n)$ in $(A \cup \mathbb{Q}) \setminus \{0\}$ independently for each $n \in \mathbb{N}^+$. Each one of these extends to a distinct maximal existential type. This gives the maximum number of types. \square

5 Independence and higher amalgamation

We now define a notion of independence and proceed to establish higher amalgamation of independent systems.

Definition 5.1. Let F be an EA-field. Let $A, B, C \subseteq F$ be subsets. Recall that $\langle A \rangle^{\text{EA}}$ is the smallest EA-subfield of F containing A , and $\langle AC \rangle^{\text{EA}}$ means $\langle A \cup C \rangle^{\text{EA}}$.

We say that A and B are *independent* over C in F and write $A \perp_C B$ if $\langle AC \rangle^{\text{EA}}$ and $\langle BC \rangle^{\text{EA}}$ are ACF-independent over $\langle C \rangle^{\text{EA}}$.

Remark. This notion of independence is quite weak. For example, it does not look at the behaviour of *logarithms* of elements of A , B , or C , that is, points whose exponentials lie in those subsets. Nor does it look at how A and B might otherwise be related in $\langle ABC \rangle^{\text{EA}}$. However it is appropriate for this setting. Stronger independence notions, suitable for exponentially closed fields, were explored in the thesis of Henderson (2014).

We now introduce the relevant definitions for higher amalgamation. More details can be found for example in Goodrick et al. (2013) or Shelah (1990, Chapter XII).

Let $n \geq 3$ be an integer, which we view as a set, that is, $n = \{0, \dots, n-1\}$. We view $\mathcal{P}(n)$ and $\mathcal{P}^-(n) = \mathcal{P}(n) \setminus n$ as categories where there is a unique morphism from a to b if $a \subseteq b$. We consider functors from these categories to $\text{Emb}(T_{\text{E-field}})$. Such functors are called $\mathcal{P}(n)$ and $\mathcal{P}^-(n)$ -systems respectively. Given a system F and $a \in \mathcal{P}(n)$, we denote by F_a the exponential field associated to a . For $a \subseteq b$ there is an embedding $F_a \rightarrow F_b$. We normally think of this embedding as being an inclusion, that is, that $F_a \subseteq F_b$. We say that F is *independent* if for every $a \subseteq b$ we have

$$F_a \perp_{\bigcup_{c \subseteq a} F_c} \bigcup_{a \not\subseteq d \subseteq b} F_d$$

as subsets of the exponential field F_b . As an example, in an independent $\mathcal{P}(3)$ -system we have, in particular, $F_{\{0,1\}} \perp_{F_{\{0\}}F_{\{1\}}} F_{\{0,2\}}F_{\{1,2\}}$ as subsets of $F_{\{0,1,2\}}$.

If we forget the exponential structure of the given independent system, we end up with an independent system of algebraically closed fields. Such systems – more generally independent systems in stable and simple theories – have been extensively studied. We mention two facts concerning such

systems. Both are special cases for ACF of theorems about independent amalgamations in stable theories.

Fact 5.2 (Folklore, see e.g. de Piro et al. (2006)). Any independent $\mathcal{P}^-(n)$ -system of algebraically closed fields can be completed to an independent $\mathcal{P}(n)$ -system.

Fact 5.3 (Fact XII.2.5 of Shelah (1990)). Let F be an independent $\mathcal{P}(n)$ -system of algebraically closed fields and $t \subseteq n$. For $i = 1, \dots, m$ let $s(i) \in \mathcal{P}(n)$ and let $\bar{a}_i \in F_{s(i)}$. Assume that for some formula ϕ we have $F_n \models \phi(\bar{a}_1, \dots, \bar{a}_m)$. Then there are $\bar{a}'_i \in F_{s(i) \cap t}$ such that $F_n \models \phi(\bar{a}'_1, \dots, \bar{a}'_m)$ and if $s(i) \subseteq t$, then $\bar{a}'_i = \bar{a}_i$.

We can now extend amalgamation to independent $\mathcal{P}^-(n)$ -systems of EA-fields.

Theorem 5.4 (n -amalgamation). Any independent $\mathcal{P}^-(n)$ -system of EA-fields can be completed to an independent $\mathcal{P}(n)$ -system.

Proof. Let F be an independent $\mathcal{P}^-(n)$ -system of EA-fields. For $a \subsetneq n$, let E_a denote the exponentiation on F_a . In this proof it is convenient to use a notation for complements. So we denote $\widehat{i} = n \setminus \{i\}$ and $\widehat{i, j} = n \setminus \{i, j\}$.

By Fact 5.2 there is an algebraically closed field F_n and embeddings (which we think of as inclusions) completing F to an independent $\mathcal{P}(n)$ -system of algebraically closed fields. We need to extend $E_{\widehat{0}} \cup \dots \cup E_{\widehat{n-1}}$ to a homomorphism from $\mathbb{G}_a(F_n)$ to $\mathbb{G}_m(F_n)$. There is a unique way to extend it to the group $F_{\widehat{0}} + \dots + F_{\widehat{n-1}}$ generated by their domains. We show that this is a well-defined homomorphism. Then by Fact 3.2 it can be extended to $\mathbb{G}_a(F_n)$. To show that this map is well-defined reduces to showing that if $a_{\widehat{0}} \in F_{\widehat{0}}, \dots, a_{\widehat{n-1}} \in F_{\widehat{n-1}}$ and $a_{\widehat{0}} + \dots + a_{\widehat{n-1}} = 0$, then it is the case that $E_{\widehat{0}}(a_{\widehat{0}}) \cdots E_{\widehat{n-1}}(a_{\widehat{n-1}}) = 1$.

We prove by induction on k that if $a_{\widehat{0}} \in F_{\widehat{0}}, \dots, a_{\widehat{k}} \in F_{\widehat{k}}$ and $a_{\widehat{0}} + \dots + a_{\widehat{k}} = 0$, then $E_{\widehat{0}}(a_{\widehat{0}}) \cdots E_{\widehat{k}}(a_{\widehat{k}}) = 1$. For $k = 0$ this is clear. So assume it holds for k . Let $a_{\widehat{0}} \in F_{\widehat{0}}, \dots, a_{\widehat{k}} \in F_{\widehat{k}}, a_{\widehat{k+1}} \in F_{\widehat{k+1}}$ be such that

$$a_{\widehat{0}} + \dots + a_{\widehat{k}} + a_{\widehat{k+1}} = 0.$$

By Fact 5.3 there are $b_{\widehat{0, k+1}} \in F_{\widehat{0, k+1}}, \dots, b_{\widehat{k, k+1}} \in F_{\widehat{k, k+1}}$ such that

$$b_{\widehat{0, k+1}} + \dots + b_{\widehat{k, k+1}} + a_{\widehat{k+1}} = 0.$$

But since $b_{\widehat{0,k+1}}, \dots, b_{\widehat{k,k+1}} \in F_{\widehat{k+1}}$ we have

$$\begin{aligned} \prod_{j=0}^k E_{\widehat{j,k+1}}(b_{\widehat{j,k+1}}) \cdot E_{\widehat{k+1}}(a_{\widehat{k+1}}) &= \prod_{j=0}^k E_{\widehat{k+1}}(b_{\widehat{j,k+1}}) \cdot E_{\widehat{k+1}}(a_{\widehat{k+1}}) \\ &= E_{\widehat{k+1}}\left(\sum_{j=0}^k b_{\widehat{j,k+1}} + a_{\widehat{k+1}}\right) \\ &= E_{\widehat{k+1}}(0) = 1. \end{aligned}$$

But also

$$a_{\widehat{0}} - b_{\widehat{0,k+1}} + \dots + a_{\widehat{k}} - b_{\widehat{k,k+1}} = a_{\widehat{k+1}} - a_{\widehat{k+1}} = 0.$$

Therefore by the induction hypothesis we conclude that

$$E_{\widehat{0}}(a_{\widehat{0}}) \cdots E_{\widehat{k}}(a_{\widehat{k}}) = E_{\widehat{0,k+1}}(b_{\widehat{0,k+1}}) \cdots E_{\widehat{k,k+1}}(b_{\widehat{k,k+1}}).$$

Putting these two together we conclude that

$$E_{\widehat{0}}(a_{\widehat{0}}) \cdots E_{\widehat{k}}(a_{\widehat{k}}) E_{\widehat{k+1}}(a_{\widehat{k+1}}) = 1.$$

This completes the induction and taking $k = n - 1$ gives us precisely the statement that we need. \square

Remark. The conclusion of this theorem is sometimes called *n-existence*. Note that the exponential map on F_n is far from uniquely determined, and so the property called *n-uniqueness* does not hold. This is in contrast to the situation for the category of exponentially closed fields (with the Schanuel property, standard kernel, and the countable closure property) as studied in the papers by Zilber (2005) and Bays and Kirby (2018). In that case both *n-existence* and *n-uniqueness* hold for all $n \in \mathbb{N}$, and indeed this is the core of the proof of the existence and uniqueness of models of each uncountable cardinality.

6 Model theoretic tree properties

In this section we introduce and study two properties of formulas: TP_2 and SOP_1 . Both properties have been extensively studied in the literature for complete first-order theories. Our setting of existentially closed models of inductive theories is somewhat more general, however the results generally transfer, with small changes to the definitions and proofs.

6.1 TP₂

Definition 6.1. Let T be an inductive theory with the JEP. An existential formula $\phi(\bar{x}, \bar{y})$ has the *tree property of the second kind* (TP₂ for short) with respect to $\mathcal{EC}(T)$ if there is an amalgamation base $A \models T$, an existential formula $\psi(\bar{y}_1, \bar{y}_2)$ and parameters $(\bar{a}_{i,j})_{i,j < \omega}$ from A such that the following hold:

(i) for all $\sigma \in \omega^\omega$ the set $\{\phi(\bar{x}, \bar{a}_{i,\sigma(i)}) : i < \omega\}$ is consistent, that is, it is realised in some $B \models T$ such that $A \subseteq B$ (or equivalently it is realised in a monster model).

(ii) $\psi(\bar{y}_1, \bar{y}_2)$ implies that $\phi(\bar{x}, \bar{y}_1) \wedge \phi(\bar{x}, \bar{y}_2)$ is inconsistent, that is,

$$T \vdash \neg \exists \bar{x} \bar{y}_1 \bar{y}_2 [\psi(\bar{y}_1, \bar{y}_2) \wedge \phi(\bar{x}, \bar{y}_1) \wedge \phi(\bar{x}, \bar{y}_2)].$$

(iii) for every $i, j, k < \omega$, if $j \neq k$ then $A \models \psi(\bar{a}_{i,j}, \bar{a}_{i,k})$.

In the first-order setting the property TP₂ was introduced by Shelah (1980) and extensively studied by Chernikov and Kaplan (2012), Ben Yaacov and Chernikov (2014) and Chernikov (2014).

Remark. In the setting of a complete first-order theory, we may always take $\psi(\bar{y}_1, \bar{y}_2)$ to be the formula $\neg \exists x [\phi(\bar{x}, \bar{y}_1) \wedge \phi(\bar{x}, \bar{y}_2)]$. So the formula ψ is not mentioned in the definition, which instead simply insists that $\phi(\bar{x}, \bar{a}_{i,j}) \wedge \phi(\bar{x}, \bar{a}_{i,k})$ is inconsistent.

In our setting $\phi(\bar{x}, \bar{a}_{i,j}) \wedge \phi(\bar{x}, \bar{a}_{i,k})$ being inconsistent means that there is an existential formula $\psi(\bar{y}_1, \bar{y}_2)$ satisfied by $\bar{a}_{i,j}, \bar{a}_{i,k}$ that implies that $\phi(\bar{x}, \bar{y}_1) \wedge \phi(\bar{x}, \bar{y}_2)$ is inconsistent. However this formula ψ may be different for different triples (i, j, k) . So the point of asking for ψ explicitly is to have a single one that works for all i, j, k . An alternative approach would be just to stipulate that $\phi(\bar{x}, \bar{a}_{i,j}) \wedge \phi(\bar{x}, \bar{a}_{i,k})$ is inconsistent, but to require the existence of parameters $(\bar{a}_{i,j})_{i,j < \kappa}$ for sufficiently large κ . Then we could use the Erdős-Rado theorem to find a sub-tree of the parameters for which a single formula ψ suffices.

Proposition 6.2. Let T be a JEP-refinement of $T_{\text{E-field}}$. Then the formula $\phi(x, yz) := E(y \cdot x) = z$ has TP₂ with respect to $\mathcal{EC}(T)$.

Proof. Take the formula $\psi(y_1 z_1, y_2 z_2)$ to be $y_1 = y_2 \wedge z_1 \neq z_2$. It is convenient to choose parameters $(\bar{a}_{i,j})_{1 \leq i, j < \omega}$, with indices starting at 1 rather than 0.

Let $F \models T$ be an amalgamation base. So it is algebraically closed, and, in particular, the \mathbb{Q} -linear dimension of F is infinite. Choose $b_i \in F$ for $1 \leq i < \omega$ such that $1, b_1, b_2, \dots$ are \mathbb{Q} -linearly independent. For $1 \leq j < \omega$ choose $c_j \in F$ to be distinct and nonzero. We let $\bar{a}_{i,j} = b_i c_j$. It is clear that if $j \neq k$, then $F \models \psi(b_i c_j, b_i c_k)$. So (iii) holds. Also (ii) holds because the exponential map is a function.

Let $\sigma \in (\omega \setminus \{0\})^{\omega \setminus \{0\}}$. By compactness, it remains to show that for any n the formula $\bigwedge_{i=1}^n E(b_i \cdot x) = c_{\sigma(i)}$ is consistent. But this follows from Theorem 3.5. Indeed, consider the variety $W(x_0, \dots, x_n, y_0, \dots, y_n)$ defined by the equations

$$\begin{aligned} x_1 &= b_1 \cdot x_0, \\ &\vdots \\ x_n &= b_n \cdot x_0, \\ y_1 &= c_{\sigma(1)}, \\ &\vdots \\ y_n &= c_{\sigma(n)}. \end{aligned}$$

We claim that W is additively free. Indeed, assume that for some $m_i \in \mathbb{Z}$ and $d \in F$ we have $\sum_{i=0}^n m_i x_i - d \in I(W)$. Then since $(0, \dots, 0, c_{\sigma(1)}, \dots, c_{\sigma(n)}) \in W$ we conclude that $d = 0$. Also $(1, b_1, \dots, b_n, 1, c_{\sigma(1)}, \dots, c_{\sigma(n)}) \in W$ and hence we conclude that $m_0 + \sum_{i=1}^n m_i b_i = 0$. By the choice of the b_i we must have $m_0 = \dots = m_n = 0$. This shows that W is additively free. Thus it must have an exponential point in some E-field extension of F . The first coordinate of this point realises $\bigwedge_{i=1}^n E(b_i \cdot x) = c_{\sigma(i)}$, as required. \square

As for a complete first-order theory, the property TP_2 for $\text{Emb}(T)$ implies that dividing (appropriately defined) does not have local character. This means that all JEP-refinements of $T_{\text{E-field}}$ are not simple either in the sense of Pillay (2000) or in the weaker sense of Ben-Yaacov (2003). We give the details in the appendix.

6.2 NSOP₁

Definition 6.3. Let T be an inductive theory with the JEP. An existential formula $\phi(\bar{x}, \bar{y})$ has the *1-strong order property* (SOP₁ for short) with respect to $\mathcal{EC}(T)$ if there is an amalgamation base $A \models T$, an existential formula

$\psi(\bar{y}_1, \bar{y}_2)$, and a binary tree of parameters $(\bar{a}_\eta : \eta \in 2^{<\omega})$ from A such that the following hold:

(i) For every branch $\sigma \in 2^\omega$, the set $\{\phi(\bar{x}, \bar{a}_{\sigma|n}) : n < \omega\}$ is consistent, that is, realised in some extension $A \subseteq B$ such that $B \models T$.

(ii) $\psi(\bar{y}_1, \bar{y}_2)$ implies that $\phi(\bar{x}, \bar{y}_1) \wedge \phi(\bar{x}, \bar{y}_2)$ is inconsistent, that is,

$$T \vdash \neg \exists \bar{x} \bar{y}_1 \bar{y}_2 [\psi(\bar{y}_1, \bar{y}_2) \wedge \phi(\bar{x}, \bar{y}_1) \wedge \phi(\bar{x}, \bar{y}_2)].$$

(iii) For every $\eta, \nu \in 2^{<\omega}$, if $\eta \frown 0 \preceq \nu$, then $A \models \psi(\bar{a}_{\eta \frown 1}, \bar{a}_\nu)$.

If no existential formula has SOP_1 , we say that $\mathcal{EC}(T)$ is NSOP_1 .

Here the notation $\eta \frown i$ denotes the sequence η with extra element i at the end and $\eta \preceq \nu$ means that η is an initial segment of ν . As with TP_2 , there is no need for ψ in the full first-order setting where it can always be taken to be $\neg \exists \bar{x} [\phi(\bar{x}, \bar{y}_1) \wedge \phi(\bar{x}, \bar{y}_2)]$.

The property SOP_1 was introduced by Džamonja and Shelah (2004) in the full first-order setting. Its systematic study began in Chernikov and Ramsey (2016) and is presently a very active area. We will use the following version of a theorem proved in Chernikov and Ramsey (2016) for complete first-order theories.

Theorem 6.4. Let T be an inductive theory with the JEP and let \mathfrak{M} be a monster model for T . Assume that there is an $\text{Aut}(\mathfrak{M})$ invariant independence relation \perp on small subsets of \mathfrak{M} which satisfies the following properties for any small existentially closed model M and any small tuples \bar{a}, \bar{b} from \mathfrak{M} :

(i) *Strong finite character:* if $\bar{a} \not\perp_M \bar{b}$, then there is an existential formula $\phi(\bar{x}, \bar{b}, \bar{m}) \in \text{tp}_\exists(\bar{a}/\bar{b}M)$ such that for any \bar{a}' realising ϕ , the relation $\bar{a}' \not\perp_M \bar{b}$ holds;

(ii) *Existence over models:* $\bar{a} \perp_M M$ for any tuple $\bar{a} \in \mathfrak{M}$;

(iii) *Monotonicity:* $\bar{a}\bar{a}' \perp_M \bar{b}\bar{b}'$ implies $\bar{a} \perp_M \bar{b}$;

(iv) *Symmetry:* $\bar{a} \perp_M \bar{b}$ implies $\bar{b} \perp_M \bar{a}$;

(v) *Independent 3-amalgamation:* If $\bar{c}_1 \perp_M \bar{c}_2$, $\bar{b}_1 \perp_M \bar{c}_1$, $\bar{b}_2 \perp_M \bar{c}_2$ and $\bar{b}_1 \equiv_M \bar{b}_2$ then there exists \bar{b} with $\bar{b} \equiv_{\bar{c}_1 M} \bar{b}_1$ and $\bar{b} \equiv_{\bar{c}_2 M} \bar{b}_2$.

Then $\mathcal{EC}(T)$ is NSOP_1 .

Proof. The proof is identical to the proof of (Chernikov and Ramsey, 2016, Proposition 5.3), except the use of Proposition 5.2 in that paper is replaced by Proposition A.7 from the Appendix. \square

Remark. The reader familiar with independence relations for simple theories may observe that the properties given here do not imply either base monotonicity or local character.

Theorem 6.5. The independence \perp for exponential fields defined in Definition 5.1 satisfies the above five conditions, even over any EA-field, not just over existentially closed E-fields. Consequently, if T is any JEP-refinement of $T_{\text{E-field}}$ then $\mathcal{EC}(T)$ is NSOP_1 .

Proof. Let T be a JEP-refinement of $T_{\text{E-field}}$, and let \mathfrak{F} be a monster model for T .

Existence over models, monotonicity, and symmetry follow from the same properties for ACF-independence. Independent 3-amalgamation follows from Theorem 5.4. Although the two statements of independent 3-amalgamation look different, it is well-known and straightforward to prove that they are equivalent.

We prove strong finite character.

Let $F \subseteq \mathfrak{F}$ be a small EA-subfield. Assume that $\bar{a} \not\perp_F \bar{b}$. Then $\langle F\bar{a} \rangle^{\text{EA}}$ and $\langle F\bar{b} \rangle^{\text{EA}}$ are not ACF-independent over F . So there is a finite tuple $\bar{\beta} \in \langle F\bar{b} \rangle^{\text{EA}}$ which is algebraically independent over F , but not algebraically independent over $\langle F\bar{a} \rangle^{\text{EA}}$. Let $q(\bar{y})$ be a polynomial witnessing the algebraic dependence. By dividing through by some coefficient, we may assume that the coefficient of some non-constant term is 1.

Now write $q(\bar{y})$ as $p(\bar{\alpha}, \bar{y})$ where $\bar{\alpha}$ is a finite tuple from $\langle F\bar{a} \rangle^{\text{EA}}$ and $p(\bar{x}, \bar{y}) \in F[\bar{x}, \bar{y}]$.

Let $\psi(\bar{x}, \bar{a}) \in \text{tp}_{\exists}^{\mathfrak{F}}(\bar{\alpha}/F\bar{a})$ be an \exists -formula which exhibits the witnesses which show that $\bar{\alpha}$ is in the EA-closure of $F\bar{a}$. We can assume that for any \bar{a}' and $\bar{\alpha}'$, if $\mathfrak{F} \models \psi(\bar{\alpha}', \bar{a}')$ then $\bar{\alpha}' \in \langle F\bar{a}' \rangle^{\text{EA}}$.

By Corollary 4.5, $\bar{\beta}$ is in the model-theoretic algebraic closure of $F\bar{b}$. Take $\chi(\bar{y}, \bar{b}) \in \text{tp}_{\exists}^{\mathfrak{F}}(\bar{\beta}/F\bar{b})$ to be the \exists -formula defining the smallest finite set containing $\bar{\beta}$ and defined over $F\bar{b}$. Then since $\bar{\beta}$ is (field-theoretically) algebraically independent over F , if $\mathfrak{F} \models \chi(\bar{\beta}', \bar{b})$ then $\bar{\beta}'$ is also algebraically independent over F .

We claim that if \bar{a}' realises the \exists -formula $\phi(\bar{z})$ given by

$$\exists \bar{x}, \bar{y} [\psi(\bar{x}, \bar{z}) \wedge \chi(\bar{y}, \bar{b}) \wedge p(\bar{x}, \bar{y}) = 0],$$

then $\bar{a}' \not\downarrow_F \bar{b}$. Assume that $\phi(\bar{a}')$ holds and let $\bar{\alpha}', \bar{\beta}' \in \mathfrak{F}$ be witnesses for \bar{x} and \bar{y} respectively.

Then $\bar{\beta}' \in \langle F\bar{b} \rangle^{\text{EA}}$ and, by the assumption on χ , the tuple $\bar{\beta}'$ is algebraically independent over F . By the assumption on ψ we have $\bar{\alpha}' \in \langle F\bar{a}' \rangle^{\text{EA}}$. So the condition that $p(\bar{\alpha}', \bar{\beta}') = 0$ implies that $\text{td}(\bar{\beta}'/F\alpha') < \text{td}(\bar{\beta}'/F)$. So $\langle F\bar{a}' \rangle^{\text{EA}}$ and $\langle F\bar{b} \rangle^{\text{EA}}$ are not ACF-independent over F , so $\bar{a}' \not\downarrow_F \bar{b}$. So \downarrow has strong finite character, and hence satisfies all the conditions of Theorem 6.4. \square

Appendix A Generalised stability for the category of existentially closed models

In this appendix we give the technicalities of generalised stability theory for the category $\mathcal{EC}(T)$ of an inductive theory T . Everything here applies to positive model theory as well (where atomic formulas need not have negations). These results are well known for complete theories, but need some modifications to work more generally. The main modifications are in the definitions, as we have already done for TP_2 and SOP_1 .

We fix an inductive theory T with the JEP and let \mathfrak{M} denote its monster model. All subsets and tuples are assumed to come from \mathfrak{M} . To simplify the notation we make no distinction between singletons and tuples. For simplicity, the tuples are assumed to be finite, but this is not necessary. The notation $a \equiv_A b$ means that $\text{tp}_{\exists}(a/A) = \text{tp}_{\exists}(b/A)$, or equivalently that there is an automorphism of \mathfrak{M} fixing A pointwise and taking a to b . If I is a linear order, then a sequence $(a_i)_{i \in I}$ is called *indiscernible* over A if for every $i_1 < \dots < i_n$ and $j_1 < \dots < j_n$ we have $a_{i_1} \dots a_{i_n} \equiv_A a_{j_1} \dots a_{j_n}$. The Ramsey method of constructing indiscernibles fails in our setting, but the Erdős-Rado method works to give the following fact:

Fact A.1 (Lemma 3.1 of Pillay (2000)). Fix a set A . If κ is sufficiently large and $(a_i)_{i < \kappa}$ is any sequence, then there is a sequence $(b_i)_{i < \omega}$ indiscernible over A such that for every $n < \omega$ there are $i_1 < \dots < i_n < \kappa$ with

$$b_1 \dots b_n \equiv_A a_{i_1} \dots a_{i_n}.$$

Definition A.2. A partial existential type $\Sigma(x, b)$ *divides* over A if there is a sequence $(b_i)_{i < \omega}$ indiscernible over A in $\text{tp}_{\exists}(b/A)$ such that $\bigcup_{i < \omega} \Sigma(x, b_i)$ is inconsistent.

Note that if $\Sigma(x, b)$ divides over A , then by compactness there is an existential formula $\phi(x, b) \in \Sigma(x, b)$ that divides. Dividing of formulas can be explicitly characterised as follows.

Lemma A.3. A formula $\phi(x, b)$ divides over A if and only if there is $k < \omega$, an existential formula $\psi(y_1, \dots, y_k)$ and a sequence $(b_i)_{i < \omega}$ satisfying the following conditions:

- (i) for each $i < \omega$, we have $b_i \equiv_A b$;
- (ii) $\psi(y_1, \dots, y_k)$ implies that $\phi(x, y_1) \wedge \dots \wedge \phi(x, y_k)$ is inconsistent, i.e.

$$T \vdash \neg \exists x, y_1, \dots, y_k [\psi(y_1, \dots, y_k) \wedge \phi(x, y_1) \wedge \dots \wedge \phi(x, y_k)];$$

- (iii) for each $i_1 < \dots < i_k < \omega$, we have $\mathfrak{M} \models \psi(b_{i_1}, \dots, b_{i_k})$.

Proof. If $(b_i)_{i < \omega}$ is the indiscernible sequence witnessing the dividing, then for some $k < \omega$ and $i_1 < \dots < i_k < \omega$ the formula $\phi(x, b_{i_1}) \wedge \dots \wedge \phi(x, b_{i_k})$ is inconsistent. Therefore there must be an existential formula $\psi(y_1, \dots, y_k)$, satisfied by b_{i_1}, \dots, b_{i_k} , that implies this. By indiscernibility of $(b_i)_{i < \omega}$, the formula ψ is satisfied for every $i_1 < \dots < i_k < \omega$.

Conversely assume that the conditions are satisfied. Then by compactness we can have an arbitrarily long sequence $(b_i)_{i < \kappa}$ satisfying the same property. Hence by Fact A.1 we can extract an indiscernible sequence witnessing the dividing. \square

Definition A.4. Let T be an inductive theory with JEP. Then $\mathcal{EC}(T)$ is called *simple* if dividing has local character. That is, for every A and b there is a subset $A_0 \subseteq A$ of cardinality at most $|T|$ such that $\text{tp}_{\exists}(b/A)$ does not divide over A_0 .

Remark. This definition of simplicity is due to Ben-Yaacov (2003), which develops an independence relation based on non-dividing for simple theories. The approach of Pillay (2000) is slightly different. The notion of dividing there is extended to forking and simplicity is defined as the local character of forking. This condition, however, implies that forking is equivalent to dividing and therefore also implies (but is stronger than) simplicity in the above sense.

Proposition A.5. If $\mathcal{EC}(T)$ is simple then no existential formula has TP_2 .

Proof. Assume that the existential formula $\phi(x, y)$, has TP_2 and let $\psi(y_1, y_2)$ be the formula witnessing that $\phi(x, y_1) \wedge \phi(x, y_2)$ is inconsistent. Fix a regular cardinal λ . We will construct a type over a set of cardinality λ that divides over every subset of smaller cardinality.

By compactness, for arbitrarily large κ we can find parameters $(a_{i,j})_{i,j < \kappa}$ such that

- (i) for every $\sigma \in \kappa^\kappa$ the set $\{\phi(x, a_{i,\sigma(i)}) : i < \kappa\}$ is consistent, and
- (ii) for every $i, j, k < \kappa$, if $j \neq k$, then $\mathfrak{M} \models \psi(a_{i,j}, a_{i,k})$.

Construct a function $\sigma \in \kappa^\lambda$ by induction on i and take $b_i = a_{i,\sigma(i)}$. Assume that $\sigma|_i$ has been constructed. Consider the sequence

$$(\text{tp}_\exists(a_{i,j}/\{b_k : k < i\}))_{j < \kappa}.$$

If κ is large enough, one of these types has to repeat infinitely often. Pick $\sigma(i)$ such that $\text{tp}_\exists(a_{i,\sigma(i)}/\{b_k : k < i\})$ appears infinitely often and set $b_i = a_{i,\sigma(i)}$. Then $\phi(x, b_i)$ divides over $\{b_k : k < i\}$. Now let b realise $\{\phi(x, b_i) : i < \lambda\}$. Then $\text{tp}_\exists(b/\{b_i : i < \lambda\})$ divides over every subset of smaller cardinality. \square

Next we turn our attention to NSOP_1 . We prove Proposition A.7, which is the generalisation of Proposition 5.2 of Chernikov and Ramsey (2016) to our setting, and which is used in Theorem 6.4. The proof is essentially the same. The only differences that are not cosmetic are sidestepping the use of Ramsey's Theorem via Fact A.1, and the use of Skolem functions.

If A is a subset of an existentially closed model M , then by $\text{dcl}_\exists(A)$ we denote the set of all the elements of M which are pointwise existentially definable over A . As with acl_\exists , the operator dcl_\exists is a closure operator. We first show how to add Skolem functions in our setting.

Lemma A.6. Let T be an inductive theory in the language L . Then there is an expansion L' of L and an inductive theory T' extending T such that for every existentially closed model $M' \models T'$ the following hold

- (i) $M'|_L$ is an existentially closed model of T ;
- (ii) if $A \subseteq M'$, then $\text{dcl}_\exists(A)$ is an existentially closed model of T' .

Proof. If we add function symbols witnessing existential quantifiers and the corresponding Skolem axioms in the usual way, then the resulting theory T' will be model complete. So every model of T' will be existentially closed, and the first condition may fail.

We instead add *partial* Skolem functions through relation symbols. More specifically, for each existential formula $\phi(\bar{x}, y)$ and each partition of its free variables into a tuple \bar{x} and a single variable y , we add a new relation symbol $R_\phi(\bar{x}, y)$ together with the following axioms:

- (i) $\forall \bar{x}, y [R_\phi(\bar{x}, y) \rightarrow \phi(\bar{x}, y)];$
- (ii) $\forall \bar{x}, y_1, y_2 [R_\phi(\bar{x}, y_1) \wedge R_\phi(\bar{x}, y_2) \rightarrow y_1 = y_2];$
- (iii) $\forall \bar{x} [\exists y \phi(\bar{x}, y) \leftrightarrow \exists y R_\phi(\bar{x}, y)].$

We iterate this process the usual way and take the limit. Let L' denote this language and T' denote the resulting theory.

In any model M' of T' , the symbol R_ϕ is interpreted by a partial function witnessing $\phi(\bar{x}, y)$. So if N is an extension of $M'|_L$, then this partial function can be extended to witness $\phi(\bar{x}, y)$ in N . Therefore we can interpret the symbols of $L' \setminus L$ in such a way that the resulting structure N' is a model of T' and an extension of M' . (This may not be possible if we add function symbols instead.) Thus, if M' is existentially closed, any existential formula $\psi(\bar{x})$ over M' witnessed by elements of N' is already witnessed by elements of M' . In particular, taking ψ to be an L -formula, we see that $M'|_L$ must be existentially closed too.

For the second condition, let $A \subseteq M'$ and $B = \text{dcl}_\exists(A)$. Let $\phi(\bar{b}, y)$ be an existential formula over B and $M' \models \exists y \phi(\bar{b}, y)$. We need to find an element witnessing y in B . Then B will be a substructure of M' and so a model of T'_\forall and by Facts 2.2 and 2.3, also an existentially closed model of T' . Since $B = \text{dcl}_\exists(A)$, there is an existential formula $\psi(\bar{a}, y)$ over A such that $M' \models \forall y (\phi(\bar{b}, y) \leftrightarrow \psi(\bar{a}, y))$. Then there is b' such that $M' \models R_\psi(\bar{a}, b')$. But this b' is \exists -definable over A and so is in B . \square

We write $a \downarrow_C^u b$ to mean that $\text{tp}_\exists(a/Cb)$ is finitely satisfiable in C . Again, we work in a monster model \mathfrak{M} of T .

Proposition A.7. If $\phi(x, y)$ has SOP_1 with the inconsistency witnessed by $\psi(y_1, y_2)$, then there is an existentially closed model M and tuples c_1, c_2, b_1, b_2 such that $c_1 \downarrow_M^u c_2$, $c_1 \downarrow_M^u b_1$, $c_2 \downarrow_M^u b_2$, $b_1 \equiv_M b_2$, and such that

$$\mathfrak{M} \models \phi(b_1, c_1) \wedge \phi(b_2, c_2) \wedge \psi(c_1, c_2).$$

Proof. By Lemma A.6 we may assume that $\text{dcl}_{\exists}(A)$ is an existentially closed model of T for every subset A of \mathfrak{M} .

Let λ be a cardinal. By compactness, for arbitrarily large κ we can find parameters $(a_{\eta})_{\eta \in 2^{<\kappa}}$ such that

- (i) for every $\sigma \in 2^{\kappa}$ the set $\{\phi(x, a_{\sigma|_i}) : i < \kappa\}$ is consistent;
- (ii) for every $\eta, \nu \in 2^{<\kappa}$ such that $\eta \frown 0 \preceq \nu$, $\mathfrak{M} \models \psi(a_{\eta \frown 1}, a_{\nu})$.

Construct a sequence $(\eta_i, \nu_i)_{i < \lambda}$ inductively. Assume $(\eta_j, \nu_j)_{j < i}$ has been constructed. Let $\eta = \bigcup_{j < i} \eta_j$. If κ is large enough, there are $\alpha < \beta < \kappa$ such that

$$a_{\eta \frown 0^{\alpha \frown 1}} \equiv_{\{a_{\eta_j}, a_{\nu_j} : j < i\}} a_{\eta \frown 0^{\beta \frown 1}}.$$

Define $\nu_i = \eta \frown 0^{\alpha \frown 1}$ and $\eta_i = \eta \frown 0^{\beta \frown 1}$.

Since η_i extends η_j for $j < i$, there is b_2 that realises $\{\phi(x, a_{\eta_i}) : i < \lambda\}$. If λ is large enough, by Fact A.1, there is a sequence $(e_i, d_i)_{i < \omega+2}$ indiscernible over b_2 , such that for every $n < \omega$ there are $i_1 < \dots < i_n < \lambda$ with

$$e_1 d_1 \dots e_n d_n \equiv_{b_2} a_{\eta_{i_1}} a_{\nu_{i_1}} \dots a_{\eta_{i_n}} a_{\nu_{i_n}}.$$

Let $M = \text{dcl}_{\exists}(\{e_n d_n : n < \omega\})$, $c_1 = d_{\omega}$ and $c_2 = e_{\omega+1}$. Then $c_1 \downarrow_{\{e_n d_n : n < \omega\}}^u c_2$ and $c_2 \downarrow_{\{e_n d_n : n < \omega\}}^u b_2$ by indiscernibility. It follows that $c_1 \downarrow_M^u c_2$, $c_2 \downarrow_M^u b_2$ and $\mathfrak{M} \models \phi(b_2, c_2)$. Also note that $c_1 c_2 = d_{\omega} e_{\omega+1} \equiv_{a_{\nu_0} a_{\eta_1}}$. But for some α , we have $\nu_0 = 0^{\alpha \frown 1}$ and $0^{\alpha \frown 1} \prec \eta_1$. So $\mathfrak{M} \models \psi(c_1, c_2)$.

It remains to find an appropriate b_1 . We claim that $e_{\omega} \equiv_M d_{\omega}$. Indeed for a fixed $n < \omega$ find $i_0 < \dots < i_n < \lambda$ such that $e_0 d_0 \dots e_n d_n e_{\omega} d_{\omega} \equiv_{b_2} a_{\eta_{i_0}} a_{\nu_{i_0}} \dots a_{\eta_{i_n}} a_{\nu_{i_n}}$. It now follows that

$$\begin{aligned} e_0 d_0 \dots e_{n-1} d_{n-1} e_{\omega} &\equiv a_{\eta_{i_0}} a_{\nu_{i_0}} \dots a_{\eta_{i_{n-1}}} a_{\nu_{i_{n-1}}} a_{\eta_{i_n}} \\ &\equiv a_{\eta_{i_0}} a_{\nu_{i_0}} \dots a_{\eta_{i_{n-1}}} a_{\nu_{i_{n-1}}} a_{\nu_{i_n}} \quad \text{by the choice of } \eta_{i_n} \text{ and } \nu_{i_n} \\ &\equiv e_0 d_0 \dots e_{n-1} d_{n-1} d_{\omega}. \end{aligned}$$

Now let $f \in \text{Aut}(\mathfrak{M}/M)$ such that $f(e_{\omega}) = d_{\omega} = c_1$. Let $b_1 = f(b_2)$. Then $b_1 \equiv_M b_2$ and $\mathfrak{M} \models \phi(b_2, e_{\omega})$, so we also have $\mathfrak{M} \models \phi(b_1, c_1)$. Finally $e_{\omega} \downarrow_{\{e_n d_n : n < \omega\}}^u b_2$ by indiscernibility. Therefore $c_1 \downarrow_{\{e_n d_n : n < \omega\}}^u b_1$ and so $c_1 \downarrow_M^u b_1$. \square

References

- Martin Bays and Jonathan Kirby. Pseudo-exponential maps, variants, and quasiminimality. *Algebra & Number Theory*, 12(3):493–549, 2018.
- Itai Ben Yaacov and Artem Chernikov. An independence theorem for NTP_2 theories. *The Journal of Symbolic Logic*, 79(1):135–153, 2014.
- Itai Ben Yaacov and Bruno Poizat. Fondements de la logique positive. *The Journal of Symbolic Logic*, 72(4):1141–1162, 2007.
- Itay Ben-Yaacov. Simplicity in compact abstract theories. *Journal of Mathematical Logic*, 3(2):169–191, 2003.
- Artem Chernikov. Theories without the tree property of the second kind. *Annals of Pure and Applied Logic*, 165(2):695–723, 2014.
- Artem Chernikov and Itay Kaplan. Forking and dividing in NTP_2 theories. *The Journal of Symbolic Logic*, 77(1):1–20, 2012.
- Artem Chernikov and Nicholas Ramsey. On model-theoretic tree properties. *Journal of Mathematical Logic*, 16(2):1650009, 2016.
- Tristram de Piro, Byunghan Kim, and Jessica Millar. Constructing the hyperdefinable group from the group configuration. *Journal of Mathematical Logic*, 6(2):121–139, 2006.
- Mirna Džamonja and Saharon Shelah. On \triangleleft^* -maximality. *Annals of Pure and Applied Logic*, 125(1-3):119–158, 2004.
- László Fuchs. *Infinite Abelian Groups*. Academic Press Inc., 1970.
- John Goodrick, Byunghan Kim, and Alexei Kolesnikov. Amalgamation functors and boundary properties in simple theories. *Israel Journal of Mathematics*, 193:169–207, 2013.
- Robert Henderson. *Independence in exponential fields*. PhD thesis, University of East Anglia, 2014.
- Wilfrid Hodges. *Building Models By Games*. Cambridge University Press, 1985.

- Wilfrid Hodges. *Model Theory*. Cambridge University Press, 1993.
- Jonathan Kirby. Finitely presented exponential fields. *Algebra and Number Theory*, 7(4):943–980, 2013.
- Angus Macintyre. Exponential algebra. In Aldo Ursini and Paolo Aglianò, editors, *Logic and Algebra*. Marcel Dekker, INC., 1996.
- Anand Pillay. Forking in the category of existentially closed structures. In Angus Macintyre, editor, *Connections between Model Theory and Algebraic and Analytic Geometry*, pages 23–42. Dipartimento di Matematica della Seconda Università di Napoli, 2000.
- Saharon Shelah. The lazy model-theoretician’s guide to stability. *Logique et Analyse (N.S.)*, 18(71-72):241–308, 1975. *Comptes Rendus de la Semaine d’Étude en Théorie des Modèles* (Inst. Math., Univ. Catholique Louvain, Louvain-la-Neuve, 1975).
- Saharon Shelah. Simple unstable theories. *Annals of Mathematical Logic*, 19(3):177–203, 1980.
- Saharon Shelah. *Classification Theory And the Number of Nonisomorphic Models*. North-Holland, revised edition, 1990.
- Lou van den Dries. Exponential rings, exponential polynomials and exponential functions. *Pacific Journal of Mathematics*, 113(1):51 – 66, 1984.
- Alex Wilkie. Model completeness results for expansions of the ordered field of real numbers by restricted pffian functions and the exponential function. *Journal of the American Mathematical Society*, 9(4):1051 – 1094, 1996.
- Boris Zilber. Pseudo-exponentiation on algebraically closed fields of characteristic zero. *Annals of Pure and Applied Logic*, 132(1):67–95, 2005.