

Forking and Dividing in Fields with Several Orderings and Valuations

Will Johnson

August 22, 2013

1 Introduction

The theory of fields with n unrelated valuations has a model companion, by the thesis of van den Dries [11]. One can also include orderings and p -valuations. More precisely, suppose that for $1 \leq i \leq n$, the theory T_i is ACVF, RCF, or p CF for some p . Arrange that for $i \neq j$, the languages of T_i and T_j overlap only in the language of rings. Then one forms the theory $\bigcup_{i=1}^n (T_i)_\forall$, whose models are fields¹ K with additional structure making K a model of $(T_i)_\forall$, for each $1 \leq i \leq n$. In van den Dries's notation, $\bigcup_{i=1}^n (T_i)_\forall$ would be denote $((T_1)_\forall, (T_2)_\forall, \dots, (T_n)_\forall)$.

For example, if each T_i is ACVF, then $(T_i)_\forall$ is the theory of valued fields, and

$$((T_1)_\forall, (T_2)_\forall, \dots, (T_n)_\forall)$$

is the theory of fields with n different valuations. If each T_i is RCF, the $(T_i)_\forall$ is the theory of ordered fields, and $((T_1)_\forall, (T_2)_\forall, \dots, (T_n)_\forall)$ is the theory of fields with n orderings. The case of p CF is similar, though for technical reasons one must use the Macintyre language. The T_i can be mixed; for example

$$(ACVF_\forall, RCF_\forall, 3CF_\forall)$$

is the theory of fields with a valuation, an ordering, and a 3-valuation (+ Macintyre predicates). In all these cases, van den Dries proves that a model companion

$$\overline{((T_1)_\forall, (T_2)_\forall, \dots, (T_n)_\forall)}$$

¹Or rather, domains. We're going to sweep this issue under the rug.

exists. In fact, van den Dries’s result is more general than what we have stated, allowing the T_i ’s to be arbitrary theories with quantifier elimination such that the $(T_i)_\forall$ are “t-theories” (Definition III.1.2 in [11]).

We will only consider the case where the T_i are ACVF, RCF, or p CF, however. In these cases, we will prove the following about the model companion $\overline{((T_1)_\forall, \dots, (T_n)_\forall)}$, which we denote T for simplicity:

1. T is NTP_2 , but is never NsOP (obviously), and not NIP if $n > 1$. See Theorems 6.7 and 6.1. If $n = 1$, then T is one of ACVF, RCF, or p CF, which are all known to be NIP .
2. In fact, T is “strong” in the sense of Adler [1], i.e., every type has finite burden. The burden of affine m -space is exactly mn , where n is the number of valuations and orderings. See Theorem 6.7.
3. Forking and dividing agree over sets in the home sort, so every set in the home sort is an “extension base for forking” in the sense of Chernikov and Kaplan [3]. See Theorem 7.5.
4. Forking in the home sort has the following characterization (Theorem 7.10). Suppose $K \models T$, and $A, B, C \subset K$ are subsets of the home sort. For $1 \leq i \leq n$, let K_i be a model of T_i extending the reduct of K to the language of T_i . For example, in the case of n orderings, K_i could be a real closure of K with respect to the i^{th} ordering. Then $A \downarrow_C B$ if and only if $A \downarrow_C B$ holds in K_i for every i . The choice of the K_i does not matter.

In many of the cases, (1) follows from Samaria Montenegro’s recent proof of Conjecture 5.1 in [4], which states that bounded pseudo-real-closed and pseudo- p -closed fields are NTP_2 . This includes the case

$$\overline{(RCF_\forall, RCF_\forall, \dots, RCF_\forall)}$$

of existentially closed fields with n orderings, as well as the case when ever T_i is p CF for some p . I don’t know whether Montenegro has also proven the remaining cases of (1) and (2). As far as I know, (3) and (4) are new. I believe (3) also holds of sets of imaginaries, which would imply that Lascar strong type and compact strong type agree, by [12] Corollary 3.6. But I have not proven this yet. In the case of fields with n orderings, this should follow from elimination of imaginaries, which Montenegro may have proven.

In the case where every T_i is ACVF, the model companion T is the theory of existentially closed fields with n valuations. In this case, the above results can be expressed more cleanly. It turns out (Theorem 4.1) that the model companion is axiomatized by the following axioms: $(K, v_1, v_2, \dots, v_n)$ is a model if and only if K is algebraically closed, each v_i is non-trivial, and v_i and v_j induce different topologies on K , for $i \neq j$. In this case, forking is characterized as follows: $A \downarrow_C B$ holds if and only if it holds in the reduct (K, v_i) , for every i . This holds because each (K, v_i) is already a model of ACVF, and so we can take $K_i = K$ in the statement of (4).

One can also express the axioms of T more concisely if exactly one of the T_i isn't ACVF. If T_1 is RCF and T_2, \dots, T_n are all ACVF, then we are considering the model companion of the theory of ordered fields with $(n-1)$ unrelated valuations. In this case, the model companion is axiomatized by the statement that the field is real closed, the valuations are non-trivial, the valuations induce different topologies from each other, and the valuations induce different topologies from the order topology. Something similar holds with p CF in place of RCF. See Theorem 4.1 for details.

As a concrete example, let K be one of the following fields: $\mathbb{F}_p(t)^{alg}$, \mathbb{Q}^{alg} , $\mathbb{Q}^{alg} \cap \mathbb{R}$, or $\mathbb{Q}^{alg} \cap \mathbb{Q}_p$ for some p . Let R_1, \dots, R_n be valuation rings on K . Then K with the ring structure and with a unary predicate for each R_i is a strong NTP₂ theory, and every set of real elements is an extension base. The same holds for $\bigcap_{i=1}^n R_i$ as a pure ring.

The outline of this paper is as follows. In Section 2, we recall some elementary facts about ACVF, p CF, and RCF which will be needed later. In Section 3, we quickly reprove the main facts needed from Chapters II and III of van den Dries's thesis, arriving at a slightly different way of expressing the axioms of the model companion, and handling the case of positive characteristic, which was not explicitly considered by van den Dries. Section 4 is a digression aimed at proving Theorem 4.1, which drastically simplifies the axioms of the model companions in some cases. I don't know whether Theorem 4.1 is already known, so a proof is included here. In Section 5, we construct some Keisler measures that will be used in the later sections. In Section 6, we determine where the model companion lies in terms of various classification theoretic boundaries, proving that it is NTP₂ and strong, but not NsOP and usually not NIP. In Section 7, we show that forking and dividing agree over sets in the home sort, and we characterize forking in terms of forking in the T_i 's.

2 Various facts about ACVF, p CF, and RCF

Let T be one of ACVF, RCF, or p CF (p -adically closed fields). Work in the usual one-sorted languages with quantifier elimination—for p CF this would be the Macintyre language.

The following fact follows easily from the various quantifier-elimination results:

Fact 2.1. *Let \mathfrak{M} be a model of T , and K be a subfield. Every K -definable set is a positive boolean combination of topologically open sets and affine varieties defined over K . In particular, any K -definable subset of \mathfrak{M}^n has non-empty interior or is contained in a K -definable proper closed subvariety of \mathbb{A}^n .*

Let \mathfrak{M} be a monster model of T .

Definition 2.2. *Let K be a subfield of \mathfrak{M} . Let $D \subset \mathfrak{M}^n$ be a definable set, defined over K . Define the rank $\text{rk}_K D$ to be the supremum of $\text{tr.deg}(\alpha/K)$ as α ranges over D .*

Lemma 2.3. (a) *If $D \subset \mathfrak{M}^n$, then $\text{rk}_K D = n$ if and only if D has non-empty interior.*

(b) *If $D \subset \mathfrak{M}^n$ and $1 \leq k \leq n$, then $\text{rk}_K D \geq k$ if and only if $\text{rk}_K \pi(D) = k$ for one of the (finitely many) coordinate projections $\pi : \mathfrak{M}^n \rightarrow \mathfrak{M}^k$.*

(c) *The rank of D doesn't depend on the choice of K , and rank is definable in families.*

(d) *If $D \subset V$ where V is absolutely irreducible, then $\text{rk} D = \dim V$ if and only if $D(\mathfrak{M})$ is Zariski dense in $V(\mathfrak{M}^{\text{alg}})$.*

Proof. **(a)** If $\text{rk}_K D < n$, then every tuple α from D lives inside an affine K -variety of dimension less than n . By compactness, D is contained in the union of finitely many affine K -varieties of dimension less than n . This union contains the Zariski closure of D , so D is not Zariski dense. This forces D to have no topological interior, because non-empty polydisks in affine space are Zariski dense. Conversely, if D has no interior, then by Fact 2.1, $D \subset V$ for some proper subvariety $V \subsetneq \mathbb{A}^n$ with V defined over K . Then $\text{rk}_K D \leq \dim V < n$.

(b) Clear by properties of rank in pregeometries.

- (c) Combine (a) and (b).
- (d) If $\text{rk } D < \dim V$, then every point in D is contained in an affine K -variety of dimension less than $\dim V$. By compactness, D is contained in the union of finitely many such varieties. This finite union contains the Zariski closure of D , and is strictly smaller than V itself. Conversely, suppose that D is not Zariski dense in V . Let $V' \subset V$ be the Zariski closure of D . As V is absolutely irreducible, $\dim V' < \dim V$. Also, V' is defined over \mathfrak{M} rather than \mathfrak{M}^{alg} , because it is the Zariski closure of a set of \mathfrak{M} -points. Let L be a small subfield of \mathfrak{M} over which V' and D are defined. Then

$$\text{rk}_K D = \text{rk}_L D \leq \text{rk}_L V' \leq \dim V' < \dim V.$$

□

Corollary 2.4. *If $K \leq L$ is an inclusion of small subfields of \mathfrak{M} and α is a finite tuple, we can find $\alpha' \equiv_K \alpha$ with $\text{tr.deg}(\alpha'/L) = \text{tr.deg}(\alpha'/K)$.*

Proof. Let $n = \text{tr.deg}(\alpha/K)$. Let $\Sigma(x)$ be the partial type asserting that $x \equiv_K \alpha$ and that x belongs to no L -variety of dimension less than n . I claim that $\Sigma(x)$ is consistent. Otherwise, there is some formula $\phi(x)$ from $\text{tp}(\alpha/K)$ and some L -varieties V_1, \dots, V_m of dimension less than n , such that $\phi(\mathfrak{M}) \subset \bigcup_{i=1}^m V_i$. But then

$$\text{rk}_K \phi(\mathfrak{M}) = \text{rk}_L \phi(\mathfrak{M}) \leq \max_{1 \leq i \leq m} \dim V_i < n,$$

contradicting the fact that $\alpha \in \phi(\mathfrak{M})$ and $\text{tr.deg}(\alpha/K) \geq n$.

Thus $\Sigma(x)$ is consistent. If α' is a realization, then $\alpha' \equiv_K \alpha$ and

$$\text{tr.deg}(\alpha'/L) \geq n = \text{tr.deg}(\alpha/K) = \text{tr.deg}(\alpha'/K) \geq \text{tr.deg}(\alpha'/L).$$

□

Corollary 2.5. *Let L and L' be two fields satisfying T_\forall , and suppose they share a common subfield K . Then L and L' can be amalgamated over K in a way which makes L and L' be algebraically independent over K .*

Proof. By quantifier elimination, we may as well assume that L and L' and K live inside a monster model $\mathfrak{M} \models T$. By the previous Corollary and compactness, we can extend $\text{tp}(L/K)$ to L' in such a way that any realization is algebraically independent from L' over K . □

Definition 2.6. Let $K \leq L$ be an inclusion of fields. Say that K is nearly closed in L if every $x \in L \cap K^{\text{alg}}$ is in the perfect closure of K .

This is a generalization of K being relatively algebraically closed in L ; in characteristic zero these two concepts are the same. Note that if we embed L into a monster model \mathfrak{M} of ACF, then K is nearly closed in L if and only if $\text{dcl}(K) = \text{acl}(K) \cap \text{dcl}(L)$ if and only if $\text{tp}(L/K)$ is stationary. From this, one gets

Fact 2.7. Let $L \geq K \leq L'$ be (pure) fields. Suppose that K is nearly closed in L or L' . Then there is only one way to amalgamate L and L' over K in such a way that L and L' are algebraically independent over K .

Fact 2.8. If K is nearly closed in L and α is a tuple from L , and V is the variety over K of which α is the generic point, then V is absolutely irreducible.

2.1 Dense formulas

In this section, T continues to be one of ACVF, RCF, or pCF.

Definition 2.9. Let K be a model of T_{\forall} . Let V be an absolutely irreducible affine variety defined over K . Let $\phi(x)$ be a quantifier-free formula with parameters from K , defining a subset of V in any/every model of T extending K . Say that $\phi(x)$ is V -dense if $\text{rk } \phi(\mathfrak{M}) = \dim V$. Here \mathfrak{M} is a monster model of T extending K .

The choice of \mathfrak{M} doesn't matter by quantifier-elimination in T and by Lemma 2.3(c).

Lemma 2.10. Let K be a model of T_{\forall} , L be a model of T extending K , and V be an absolutely irreducible variety defined over K . For a quantifier-free K -formula $\phi(x)$, the following are equivalent:

- (a) $\phi(x)$ is V -dense.
- (b) $\phi(L)$ is Zariski dense in $V(L^{\text{alg}})$.
- (c) We can extend the T_{\forall} -structure on K to the function field $K(V)$ in such a way that the generic point of V in $K(V)$ satisfies $\phi(x)$.

Proof. **(a) \implies (b)** Suppose $\phi(x)$ is V -dense. Let W be the Zariski closure of $\phi(L)$ in $V(L^{alg})$. Then W is defined over L rather than L^{alg} , because W is the Zariski closure of some L -points. Therefore it makes sense to think of W as a definable set. If \mathfrak{M} is a monster model of T extending L , then $\dim V = \text{rk } \phi(\mathfrak{M}) \leq \text{rk } W \leq \dim W \leq \dim V$. Therefore $\dim W = \dim V$. As V is absolutely irreducible, $W = V$.

(b) \implies (a) Let \mathfrak{M} be a monster model of T extending L , and let $n = \dim V$. If $\phi(x)$ is not V -dense, then every element of $\phi(\mathfrak{M})$ has transcendence degree less than n over K . By compactness, $\phi(\mathfrak{M})$ is contained in a finite union of K -definable varieties of dimension less than n . We may assume these varieties are closed subvarieties of V . Of course $\phi(L)$ is also contained in this union, which is clearly a Zariski closed proper subset of V . So $\phi(L)$ is not Zariski dense.

(a) \implies (c) Embed K into a monster model \mathfrak{M} . Let α be a point in $\phi(\mathfrak{M}) \subset V(\mathfrak{M})$ with $\text{tr.deg}(\alpha/K) = \text{rk } \phi(\mathfrak{M}) = \dim V$. Then α is a generic point on V , i.e., $K(\alpha) \cong K(V)$. And α satisfies $\phi(x)$.

(c) \implies (a) Embed $K(V)$ into a monster model \mathfrak{M} . Let α denote the generic point of V , so that $\mathfrak{M} \models \phi(\alpha)$ holds. Clearly $\text{tr.deg}(\alpha/K) = \dim V$. Thus $\text{rk}_K \phi(\mathfrak{M}) \geq \text{tr.deg}(\alpha/K) = \dim V$, implying V -density of $\phi(x)$. □

Lemma 2.11. *Let L be a model of ACVF, and let $V \subset \mathbb{A}^n$ be an irreducible affine variety over L . Suppose $0 \in V$. Let \mathcal{O}_L^n be the closed unit polydisk in \mathbb{A}^n . Then $\mathcal{O}_L^n \cap V$ is Zariski dense in V .*

This Lemma is essentially Lemma 1.1 in [6], but we're going to give a more elementary proof based on the proof of Proposition 4.2.1 in [7].

Proof. Let $L(\alpha)$ be the function field of V , obtained by adding a generic point α of V to the field L . By the implication (c) \implies (b) of Lemma 2.10 applied in the case where $\phi(x)$ is the formula defining $\mathcal{O}_L^n \cap V$, it suffices to extend the valuation on L to $L(\alpha)$ in such a way that every coordinate of α has nonnegative valuation.

Now $L[\alpha]$ is the coordinate ring of V , so the fact that $0 \in V$ implies that there is an L -algebra homomorphism $L[\alpha] \rightarrow L$ sending every coordinate of α to zero. This yields an \mathcal{O}_L -algebra homomorphism $f : \mathcal{O}_L[\alpha] \rightarrow \mathcal{O}_L$

sending every coordinate of α to 0. Let \mathfrak{m} be the maximal ideal of \mathcal{O}_L , and let $\mathfrak{p} = f^{-1}(\mathfrak{m})$. Then \mathfrak{p} is a prime ideal, and $\mathfrak{p} \cap \mathcal{O}_L = \mathfrak{m}$. Also, as f kills the coordinates of α , the coordinates of α live in \mathfrak{p} .

Since $\mathcal{O}_L[\alpha]$ is a domain, there is a valuation v' on $L(\alpha)$, the fraction field of $\mathcal{O}_L[\alpha]$, with the following properties:

- Every element of \mathfrak{p} has positive valuation. In particular, the elements of \mathfrak{m} and the coordinates of α have positive valuation.
- Every element of $\mathcal{O}_L[\alpha] \setminus \mathfrak{p}$ has valuation zero. In particular, the elements of $\mathcal{O}_L^\times = \mathcal{O}_L \setminus \mathfrak{m}$ have valuation zero.

(Indeed, it is a general fact that if S is a domain and \mathfrak{p} is a prime ideal, then there is a valuation on the fraction field of S which assigns a positive valuation to elements of \mathfrak{p} and a vanishing valuation to elements of $S \setminus \mathfrak{p}$. To find such a valuation, take a valuation ring in $\text{Frac}(S)$ dominating the local ring $S_{\mathfrak{p}}$.)

The resulting valuation on $L(\alpha)$ extends the valuation on L , because it assigns positive valuation to elements in \mathfrak{m} , and zero valuation to elements in $\mathcal{O}_L \setminus \mathfrak{m}$. Also, the valuation of any coordinate of α is positive, hence non-negative, so α lives in the closed unit polydisk. \square

Lemma 2.12. *Let V be an absolutely irreducible affine variety over $K \models T_{\forall}$, and let $\phi(x)$ be a quantifier-free K -formula. Let L be a model of T extending K . Suppose $\phi(x)$ defines an open subset of $V(L)$.*

- (a) *If T is ACVF, then $\phi(x)$ is V -dense if and only if $\phi(L)$ is non-empty.*
- (b) *In general, $\phi(x)$ is V -dense if $\phi(L)$ contains a smooth point of V .*

Proof. (a) If $\phi(x)$ is V -dense, then certainly $\phi(L)$ is non-empty. Conversely, suppose $\phi(L)$ is non-empty. Let p be a point in $\phi(L)$ and let U be an open neighborhood of p , with $U \cap V \subset \phi(L)$. There is some L -definable affine transformation f which sends p to the origin and moves U so as to contain the closed unit polydisk. Then $f(U \cap V) = f(U) \cap f(V)$ is Zariski dense in $f(V)$, by Lemma 2.11. So $\phi(L) \supseteq U \cap V$ is Zariski dense in V . Thus $\phi(x)$ is V -dense, by Lemma 2.10.

- (b) If $\phi(x)$ is V -dense, then $\phi(L)$ contains a smooth point of V , because the smooth locus of V is a Zariski dense Zariski open. Conversely, suppose

$\phi(L)$ contains a smooth point p . Note that L is perfect. The tangent space $T_p V$ is L -definable. By Galois descent of vector spaces, there is an L -definable basis of $T_p V$. Therefore, after applying an L -definable change of coordinates, we may assume $T_p V$ is horizontal. By the implicit function theorem, V then looks locally around p like the graph of a function. In particular, the coordinate projection maps a neighborhood of p homeomorphically to an open subset of affine n -space, where $n = \dim V$. By Lemma 2.3, this ensures that any neighborhood of p , such as $\phi(L)$, has rank at least n . So $\phi(x)$ is V -dense. \square

Lemma 2.13. *Let V be an absolutely irreducible affine variety over $K \models T_V$, and let $\phi(x)$ be a quantifier-free K -formula that is V -dense. Then there is a quantifier-free K -formula $\psi(x)$ that is also V -dense, such that in any/every $L \models T$ extending K , $\psi(L)$ is a topologically open subset of $V(L)$, and $\psi(L) \subset \phi(L)$.*

Proof. Choose some monster model $\mathfrak{M} \models T$ extending K and let $\psi(\mathfrak{M})$ pick out the topological interior of $\phi(\mathfrak{M})$ inside $V(\mathfrak{M})$. By quantifier-elimination, we can take $\psi(x)$ to be quantifier-free with parameters from K . It remains to show that $\psi(\mathfrak{M})$ is V -dense. Let $\alpha \in \phi(\mathfrak{M})$ have transcendence degree n over K , where $n = \dim V$. By Fact 2.1, $\phi(\mathfrak{M})$ can be written as a finite union of finite intersections of K -definable opens and varieties. Let X be one of these finite intersections, containing α . So $X = W \cap U$ for some K -variety W and some K -definable open U . As $\alpha \in W$ and α is a generic point on V , we must have $V \subseteq W$. Then

$$\alpha \in V \cap U \subseteq W \cap U \subseteq \phi(\mathfrak{M}).$$

But $V \cap U$ is a relative open in $V(\mathfrak{M})$, so it must be part of $\psi(\mathfrak{M})$. In particular, $\alpha \in \psi(\mathfrak{M})$. As $\text{tr.deg}(\alpha/K) = n$, we conclude that $\psi(x)$ is V -dense. \square

2.2 Forking and Dividing

We continue to work in one of ACVF, RCF, or p CF. Recall that RCF and p CF have definable Skolem functions in the home sort. Thus if S is a subset of the home sort, then $\text{acl}(S) = \text{dcl}(S)$ is a model. In ACVF, $\text{acl}(S)$ is the algebraic closure of S , which is a model unless $\text{acl}(S)$ is trivially valued.

We will always be working in the home sort, rather than working with imaginaries.

Lemma 2.14. *Let S be a set (in the home sort) and let $\phi(x; b)$ be a formula. Then $\phi(x; b)$ forks over S if and only if it divides over S .*

Proof. Indiscernibility over S is the same thing as indiscernibility over $\text{acl}(S)$, so $\phi(x; b)$ divides over S if and only if it divides over $\text{acl}(S)$. Similarly, $\phi(x; b)$ forks over S if and only if it forks over $\text{acl}(S)$. So we may assume $S = \text{acl}(S)$. If T is RCF or p CF, then S is a model, and therefore forking and dividing agree over S by Theorem 1.1 of [3]. If T is ACVF, then forking and dividing agree over all sets, by Corollary 1.3 in [3]. \square

We use \downarrow to denote non-forking or non-dividing, and \downarrow^{ACF} to denote algebraic independence.

Lemma 2.15. *Let \mathfrak{M} be a monster model of T , and let B, C be small subsets of \mathfrak{M} , with B finite. Then we can find a sequence B_0, B_1, B_2, \dots in \mathfrak{M} that is C -indiscernible, such that $B_0 = B$ and $B_i \downarrow_C^{ACF} B_{<i}$ for every i .*

Proof. We may assume that B is ordered as a tuple in such a way that the first k elements of B are a transcendence basis of B over C . Construct a sequence D_0, D_1, \dots of realizations of $\text{tp}(B/C)$ such that $D_i \downarrow_C^{ACF} D_{<i}$ for every i . This is possible by using Corollary 2.4 to extend $\text{tp}(B/C)$ to a type over $CD_{<i}$ having the same transcendence degree over $CD_{<i}$ as over C . Let B_0, B_1, B_2, \dots be a C -indiscernible sequence extracted from D_0, D_1, \dots . Let $\pi(X)$ pick out the first k elements of a tuple X . Then $\pi(D_0) \frown \pi(D_1) \frown \pi(D_2) \frown \dots$ is an algebraically independent sequence of singletons over C . This is part of the EM-type of the D_i over C , so it is also true that $\pi(B_0) \frown \pi(B_1) \frown \pi(B_2) \frown \dots$ is an algebraically independent sequence of singletons over C . Since $D_i \equiv_C B$ for every i , we also have $B_i \equiv_C B$ for every i . Thus $\pi(B_i)$ is a transcendence basis for B_i over C , and we conclude that $B_i \downarrow_C^{ACF} B_{<i}$ for every i . Finally, moving the B_i by an automorphism over C , we may assume that $B_0 = B$. \square

Lemma 2.16. *$A \downarrow_C B$ implies $A \downarrow_C^{ACF} B$.*

Proof. This feels like it should be obvious, but the best proof I could come up with is the following one. By the finite character of forking, we may assume A and B are finite tuples. By Lemma 2.15, produce a sequence

$B_0 = B, B_1, B_2, \dots$ of realizations of $\text{tp}(B/C)$, indiscernible over C , and satisfying $B_i \downarrow_C^{ACF} B_{<i}$ for every i . As $\text{tp}(A/BC)$ doesn't divide over C , we may move the B_i 's by an automorphism over BC , making them be AC -indiscernible. Now the B_i 's are algebraically independent over C , and AC -indiscernible. Since $A \not\downarrow_C^{ACF} B$, we also have $A \not\downarrow_C^{ACF} B_i$ for every i . This contradicts the fact that in ACF, finite tuples have finite preweight. \square

Lastly, we show that dividing is always witnessed by an algebraically independent sequence.

Lemma 2.17. *If a formula $\phi(x; a)$ divides over a set A , then the dividing is witnessed by an A -indiscernible sequence $a = a_0, a_1, a_2, \dots$ such that $a_i \downarrow_A^{ACF} a_{<i}$ for every i .*

Proof. Apply Claim 3.10 of [3] with the abstract independence relation taken to be \downarrow (non-forking). Forking satisfies (1)-(7) of [3] Definition 2.9 by Fact 2.12(5) of [3]. And A is an extension base for forking by Lemma 2.14 above and Theorem 1.1 of [3] (or by Fact 2.14 of [3] in the cases other than pCF). So Claim 3.10 of [3] is applicable. Consequently we get a model M containing A , a global type p extending $\text{tp}(a/M)$, \downarrow -free over A , such that any/every Morley sequence generated by p over M witnesses the dividing of $\phi(x; a)$. Because \downarrow is stronger than Lascar invariance, any such Morley sequence will be M -indiscernible, hence A -indiscernible. Because \downarrow is stronger than algebraic independence (Lemma 2.16), and p is \downarrow -free over A , any Morley sequence a_0, a_1, \dots generated by p will be algebraically independent over A . Specifically, $a_i \models p|_{Ma_{<i}}$, so as p is \downarrow -free over A , $a_i \downarrow_A Ma_{<i}$, and hence $a_i \downarrow_A^{ACF} a_{<i}$. \square

3 The Model Companion

Now we turn our attention to fields with several valuations, several orderings, and several p -valuations. For $1 \leq i \leq n$, let T_i be one of ACVF, RCF, or pCF (in the same languages as in the previous section). Let \mathcal{L}_i denote the language of T_i ; assume that $\mathcal{L}_i \cap \mathcal{L}_j = \mathcal{L}_{rings}$ for $i \neq j$. Let T^0 be $\bigcup_{i=1}^n (T_i)_{\forall}$, the theory that would be denoted $((T_1)_{\forall}, (T_2)_{\forall}, \dots, (T_n)_{\forall})$ in van den Dries's notation. Technically speaking, models of T^0 should be allowed to be domains, rather than fields. However, we will assume that T^0 also includes the field axioms, sweeping domains under the rug.

One essentially knows that T^0 has a model companion T by Chapter III of van den Dries's thesis [11]. We're going to quickly reprove the existence of T in this section, expressing the axioms of the model companion in a more geometric and less syntactic form, and also including the case of positive characteristic explicitly.

3.1 The Axioms

Consider the following axioms that a model K of T^0 could satisfy:

A1: K is existentially closed with respect to finite extensions, i.e., if L/K is a finite algebraic extension and $L \models T^0$, then $L = K$.

A1': For every irreducible polynomial $P(X) \in K[X]$ of degree greater than 1, there is some $1 \leq i \leq n$ such that $P(x) = 0$ has no solution in any/every model of T_i extending $K \upharpoonright \mathcal{L}_i$.

A2(m): Let V be an m -dimensional absolutely irreducible variety over K . For $1 \leq i \leq n$, let $\phi_i(x)$ be a V -dense quantifier-free \mathcal{L}_i -formula with parameters from K . Then $\bigcap_{i=1}^n \phi_i(K) \neq \emptyset$.

A2($\leq m$): A2(m') holds for all $m' \leq m$.

A2: A2(m) holds, for all m

Remark 3.1. For $K \models T^0$, A1 and A1' are equivalent.

Proof. Suppose K satisfies A1, and $P(X) \in K[X]$ is irreducible of degree greater than 1. Suppose that for every $1 \leq i \leq n$, there is a solution α_i of $P(x) = 0$ in a model $M_i \models T_i$ extending $K \upharpoonright \mathcal{L}_i$. Then we can extend the \mathcal{L}_i -structure from K to $K(\alpha) \cong K[X]/P(X)$. Because this holds for every i , we can endow $K[X]/P(X)$ with the structure of a model of T^0 . By A1, $K[X]/P(X)$ must be K , so $P(X)$ has degree 1.

Conversely, suppose K satisfies A1' but not A1. Let L/K be a counterexample to A1, and take some $\alpha \in L \setminus K$. Let $P(X)$ be the irreducible polynomial of α over K . This polynomial must have degree greater than 1. For each i let M_i be a model of T_i extending $L \upharpoonright \mathcal{L}_i$. Then $P(x) = 0$ has a solution in L , hence in M_i , which is a model of T_i extending K . This contradicts A1'. \square

Lemma 3.2. *Let K be a model of T^0 , and $m \geq 1$. The following are equivalent:*

- (a) *For every model L of T^0 extending K , for every tuple α from L with $\text{tr.deg}(\alpha/K) \leq m$, the quantifier-free type $\text{qftp}(\alpha/K)$ is finitely satisfiable in K .*
- (b) *K satisfies A1 and A2($\leq m$).*

Proof. (a) \implies (b) For A1, suppose that L/K is a finite extension, and $L \models T$. If $\alpha \in L$, then α is algebraic over K , so $\text{tr.deg}(\alpha/K) = 0 \leq m$. By (a), the quantifier-free type of α is realized in K . So the irreducible polynomial of α over K has a zero in K , implying $\alpha \in K$. As $\alpha \in L$ was arbitrary, $L = K$.

For A2(m'), let V be an m' -dimensional absolutely irreducible variety over K . For $1 \leq i \leq n$, let $\phi_i(x)$ be a V -dense quantifier-free \mathcal{L}_i -formula with parameters from K . By Lemma 2.10(c), we can extend the \mathcal{L}_i -structure to $K(V)$ in such a way that the generic point satisfies $\phi_i(x)$. Doing this for all i , we make $K(V)$ be a model of T^0 extending K , such that if $\alpha \in K(V)$ denotes the generic point, then $\bigwedge_{i=1}^n \phi_i(\alpha)$ holds. Now $\text{tr.deg}(\alpha/K) = \dim V \leq m$, so by (a), $\text{qftp}(\alpha/K)$ is finitely satisfiable in K . In particular, the formula $\bigwedge_{i=1}^n \phi(x)$ is satisfiable in K , which is the conclusion of A2(m').

(b) \implies (a). Suppose L is a model of T^0 extending K and α is a tuple from L , with $\text{tr.deg}(\alpha/K) \leq m$. By A1, K is relatively algebraically closed in L . Let V be the K -variety of which α is a generic point. Then V is absolutely irreducible, by Fact 2.8. Also, $m' := \dim V = \text{tr.deg}(\alpha/K) \leq m$. Let $\psi(x)$ be a statement in $\text{qftp}(\alpha/K)$. We want to show that ψ is satisfied by an element of K . We may assume that $\psi(x)$ includes the statement that $x \in V$. By Fact 2.1, $\psi(x)$ is a positive boolean combination of statements of the form

- $x \in W$, for some K -definable affine variety W . Since we intersected $\psi(x)$ with V , we may assume $W \subseteq V$.
- $\theta(x)$, where $\theta(x)$ is a quantifier-free \mathcal{L}_i -formula for some i , such that $\theta(L)$ is an open subset of the ambient affine space, for any/every $L \models T_i$ extending $K \upharpoonright \mathcal{L}_i$.

Writing $\psi(x)$ as a disjunction of conjunction of such statements, and replacing $\psi(x)$ by whichever disjunct α satisfies, we may assume that $\psi(x)$ is a

conjunction of such statements. An intersection of K -varieties is a K -variety, and an intersection of open subsets of affine space is an open subset of affine space, so we may assume

$$\psi(x) \equiv "x \in W" \wedge \bigwedge_{i=1}^n \phi_i(x),$$

where W is some K -variety contained in V , and where $\phi_i(x)$ is a quantifier-free \mathcal{L}_i -formula defining an open subset of the ambient affine space, when interpreted in any/every model of T_i extending $K \upharpoonright \mathcal{L}_i$.

Because α satisfies $\psi(x)$, and α is a generic point of V , W must be V . Rewrite ψ as $\bigwedge_{i=1}^n \phi'_i(x)$, where each $\phi'_i(x)$ asserts that $x \in V$ and $\phi_i(x)$ holds. Because K satisfies axiom A2(m'), $\psi(x)$ will be satisfiable in K as long as $\phi'_i(x)$ is V -dense for each i . But note that L provides a way of extending the \mathcal{L}_i -structure from K to $K(\alpha) \cong K(V)$ in such a way that $\phi'_i(\alpha)$ holds, so ϕ'_i is V -dense by Lemma 2.10(c). \square

Theorem 3.3. *The theory T^0 has a model companion T , whose models are exactly the $K \models T^0$ satisfying A1 and A2.*

Proof. It is well known that a model K is existentially closed if and only if for every model L extending K and for every tuple α from L , the quantifier-free type $\text{qftp}(\alpha/K)$ is finitely satisfiable in K . So by Lemma 3.2, a model of T^0 is existentially closed if and only if it satisfies A1 and A2. By basic facts about model companions of $\forall\exists$ -theories, it remains to show that A1 and A2 are first order. For A1, this comes from Remark 3.1, because A1' is first order by quantifier-elimination in the T_i . Axiom A2 is first order by quantifier-elimination in the T_i , by Lemma 2.3(c), and by the fact that absolute irreducibility is definable by a quantifier-free formula in the language of fields (this is well-known and proven in Chapter IV of [11]). \square

Henceforth, we will use T to denote the model companion. Also, we will use T_{\forall} instead of T^0 , sweeping the issue of domains that aren't fields under the rug.

We make several remarks about the axioms:

Remark 3.4. *In the case where T_i is ACVF for $i > 1$, axiom A1 merely says that $K \upharpoonright \mathcal{L}_1$ is a model of T_1 , i.e., is algebraically closed or real closed or p -adically closed.*

Remark 3.5. In Axiom $A2(m)$, it suffices to consider the case of smooth V . If V is not smooth, one can find an open subvariety V' of V which is smooth, and which is isomorphic to an affine variety. (Use the facts that the smooth locus of an irreducible variety is a Zariski dense Zariski open, and that the affine open subsets of a scheme form a basis for its topology.) If $\phi_i(x)$ is V -dense, then $\phi_i(x) \wedge "x \in V'"$ is V' -dense, essentially by Lemma 2.10(b). Then applying the smooth case of $A2(m)$ to V' yields a point in V' satisfying $\bigwedge_{i=1}^n \phi_i(x)$.

Remark 3.6. In Axiom $A2$, it suffices to consider V -dense formulas $\phi_i(x)$ such that $\phi_i(L)$ defines an open subset of $V(L)$ for any/every $L \models T$ extending $K \upharpoonright \mathcal{L}_i$. This follows by Lemma 2.13.

Remark 3.7. We can combine the previous two remarks. Then Lemma 2.12(b), yields the following restatement of $A2(m)$: if V is an absolutely irreducible m -dimensional smooth affine variety defined over K , and if $\phi_i(x)$ is a quantifier-free \mathcal{L}_i -formula over K for each $1 \leq i \leq n$, and if $\phi_i(K_i)$ is a non-empty open subset of $V(K_i)$ for any/every $K_i \models T$ extending $K \upharpoonright \mathcal{L}_i$, then $\bigcap_{i=1}^n \phi_i(K) \neq \emptyset$.

Remark 3.8. If every T_i is ACVF, then $A1$ merely says that K is algebraically closed. Consequently, in Remark 3.7 the K_i can be taken to be K itself. Thus $A2(m)$ ends up being equivalent to the statement that if V is a smooth irreducible m -dimensional affine variety, and $\phi_i(x)$ is a quantifier-free \mathcal{L}_i -formula defining a non-empty open subset of V for $1 \leq i \leq n$, then $\bigcap_{i=1}^n \phi_i(K)$ is non-empty. Even more concisely, this means that for every smooth m -dimensional variety V , the diagonal map $V(K) \rightarrow \prod_{i=1}^n V(K)$ has dense image in the product topology, using the topology from the i^{th} valuation for the i^{th} entry in the product.

In fact, in Section 4, we will see that it suffices to check the case of $V = \mathbb{A}^1$, the affine line(!)

3.2 Quantifier-Elimination up to Algebraic Covers

As in the previous section, T_{\forall} is the theory of fields with $(T_i)_{\forall}$ structure for each $1 \leq i \leq n$, and T is the model companion of T_{\forall} .

Lemma 3.9. Let K be a model of T_{\forall} . Let L and L' be two models of T_{\forall} extending K . Suppose that K is nearly closed in L or L' (Definition 2.6).

Then L and L' can be amalgamated over K , and this can be done in such a way that L and L' are algebraically independent over K .

Proof. For each $1 \leq i \leq n$, we can find some amalgam $M_i \models (T_i)_\forall$ of $L \upharpoonright \mathcal{L}_i$ and $L' \upharpoonright \mathcal{L}_i$ over $K \upharpoonright \mathcal{L}_i$, by Corollary 2.5. The resulting compositums LL' must be isomorphic on the level of fields, by Fact 2.7. Consequently, we can endow the canonical field LL' with a $(T_i)_\forall$ -structure extending those on L and L' , for each i . This gives LL' the structure of a T_\forall -model. And L and L' are algebraically independent inside LL' . \square

Corollary 3.10. *Let K be a model of T_\forall and let L be a model of T extending K . Then K is relatively algebraically closed in L if and only if K satisfies axiom A1. (In particular, this doesn't depend on L .)*

Proof. If K satisfies axiom A1, then obviously K is relatively algebraically closed in L . Conversely, suppose that K is relatively algebraically closed in L but doesn't satisfy A1. Then there is some model L' of T_\forall extending K , with L'/K finite and $L' \neq K$. By Lemma 3.9, we can amalgamate L and L' over K . Embed the resulting compositum LL' in a model M of T . Because T is model-complete, $L \preceq M$. Now choose some $\alpha \in L' \setminus K$. The irreducible polynomial of α over K has a root in M , and hence has a root in L , contradicting the fact that K is relatively algebraically closed in L . \square

Corollary 3.11. *Let K be model of T_\forall , and suppose K satisfies A1. Then the type of K is determined, i.e., if L and L' are two models of T extending K , then K has the same type in L and L' . Equivalently, the diagram of K implies the elementary diagram of K , modulo the axioms of T .*

Proof. By Corollary 3.10, K is relatively algebraically closed in L and L' . So we can amalgamate L and L' over K , by Lemma 3.9. If M is a model of T extending LL' , then by model completeness $L \preceq M \succeq L'$, ensuring that K has the same type in each. \square

Corollary 3.12. *In models of T , field-theoretic algebraic closure agrees with model-theoretic algebraic closure.*

Proof. Let M be a model of T . Let S be a subset of M . Let K be the field-theoretic algebraic closure of S , i.e., the relative algebraic closure of S in M . By Lemma 3.9, we can amalgamate M and a copy M' of M over K in a way that makes M and M' be algebraically independent over K .

Embedding MM' into a model N of T , and using model completeness, we get $M \preceq N \succeq M'$. Now $\text{acl}(S)$ is the same when computed in M , N , or M' . In particular, $\text{acl}(S) \subset M \cap M'$. Since M and M' are algebraically independent over K and K is relatively algebraically closed in each, $M \cap M' = K$. Thus $\text{acl}(S) \subset K$. Obviously $K \subset \text{acl}(S)$. \square

For K a field, let $\text{Abs}(K)$ denote the algebraic closure of the prime field in K .

Corollary 3.13. *Two models $M_1, M_2 \models T$ are elementarily equivalent if and only if $\text{Abs}(M_1)$ and $\text{Abs}(M_2)$ are isomorphic as models of T_\forall .*

Proof. If M_1 and M_2 are elementarily equivalent, we can embed them as elementary substructures into a third model $M_3 \models T$. Then $\text{Abs}(M_1) = \text{Abs}(M_3) = \text{Abs}(M_2)$, so certainly $\text{Abs}(M_1)$ is isomorphic to $\text{Abs}(M_2)$.

Conversely, suppose $\text{Abs}(M_1) \cong \text{Abs}(M_2)$. Then, as $\text{Abs}(M_1)$ is relatively algebraically closed in M_1 and in M_2 , it follows by Corollaries 3.10 and 3.11 that we can amalgamate M_1 and M_2 over $\text{Abs}(M_1)$. Embedding the resulting compositum into a model of T and using model completeness, we get $M_1 \equiv M_2$. \square

Corollary 3.14. *Suppose T_1 isn't ACVF and T_i is ACVF for $i > 1$. Consider the expanded theory where we add in symbols for every zero-definable T_1 -definable function. (This makes sense because if $M \models T$, then $M \upharpoonright \mathcal{L}_1 \models T_1$, by Remark 3.4.) Then T has quantifier-elimination.*

Proof. After adding in these new symbols, a substructure is the same as a subfield K closed under all T_1 -definable functions. As RCF and p CF have definable Skolem functions, this is equivalent to $K \upharpoonright \mathcal{L}_1$ being a model of T_1 , which is equivalent to K satisfying axiom A1, as noted in Remark 3.4. Now apply Corollary 3.11 to get substructure completeness, which is the same thing as quantifier-elimination. \square

I suppose that this also works if more than one of the T_i 's isn't ACVF, though the extra functions would become partial functions.

Without adding in extra symbols, quantifier elimination doesn't work. But we still get quantifier-elimination up to algebraic covers, in a certain sense. The following theorem and its proof won't be used in an essential way in the rest of the paper.

Theorem 3.15. *In T , every formula $\phi(x)$ is equivalent to one of the form*

$$\exists y : P(y, x) = 0 \wedge \psi(y, x), \quad (1)$$

where y is a singleton, $\psi(y, x)$ is quantifier-free, and $P(y, x)$ is a polynomial in y and the coordinates of x , with integer coefficients, monic as a polynomial in y .

Proof. Let $\Sigma(x)$ be the set of all formulas of the form (1). First we observe that $\Sigma(x)$ is closed under disjunction, because

$$(\exists y : P(y, x) = 0 \wedge \psi(y, x)) \vee (\exists y : Q(y, x) = 0 \wedge \psi'(y, x))$$

is equivalent to

$$\exists y : P(y, x)Q(y, x) = 0 \wedge \psi''(y, x),$$

where $\psi''(y, x)$ is the quantifier-free formula

$$(P(y, x) = 0 \wedge \psi(y, x)) \vee (Q(y, x) = 0 \wedge \psi'(y, x)).$$

Now given a formula $\phi(x)$, not quantifier-free, let $\Sigma_0(x)$ be the set of formulas in $\Sigma(x)$ which imply $\phi(x)$, i.e.,

$$\Sigma_0(x) = \{\sigma(x) \in \Sigma(x) : T \vdash \forall x : \sigma(x) \rightarrow \phi(x)\}.$$

Of course $\Sigma_0(x)$ is closed under disjunction. It suffices to show that $\phi(x)$ implies a finite disjunction of formulas in $\Sigma_0(x)$, because then $\phi(x)$ implies and is implied by a formula in $\Sigma_0(x)$.

Suppose for the sake of contradiction that $\phi(x)$ does not imply a finite disjunction of formulas in $\Sigma_0(x)$. Then the partial type

$$\{\phi(x)\} \cup \{\neg\sigma(x) : \sigma(x) \in \Sigma_0(x)\}$$

is consistent with T . Let M be a model of T containing a tuple α realizing this partial type. So $\phi(\alpha)$ holds in M , but not because of any formula of the form (1).

Let R be the ring $\mathbb{Z}[\alpha] \subset M$. Let $K \subset M$ be the smallest perfect field containing R ; note that M itself is perfect so this makes sense. Indeed, if every T_i is ACVF, then M is algebraically closed by Remark 3.4. Otherwise, one of the T_i 's is RCF or p CF, making M be characteristic zero.

Let \overline{K} be the relative algebraic closure of K (or equivalently, α) inside M . By Corollaries 3.10 and 3.11, the diagram of \overline{K} implies the elementary diagram of \overline{K} . In particular, the diagram of \overline{K} implies $\phi(\alpha)$. By compactness, the diagram of L implies $\phi(\alpha)$, for some finite extension L of K . Because K is perfect, $L = K(\beta)$ for some *singleton* β . Multiplying β by an appropriate element from R , we may assume that β is integral over R . Note that L is perfect, because it is an algebraic extension of a perfect field, and in fact L is the smallest perfect field containing α and β .

As the diagram of L implies $\phi(\alpha)$, so does the diagram of $\mathbb{Z}[\alpha, \beta]$, by Lemma 3.16 below. By compactness, there is some true quantifier-free formula $\psi(\beta, \alpha)$ such that

$$T \vdash \forall y, x : \psi(y, x) \rightarrow \phi(x).$$

Let $P(y, x)$ be the polynomial witnessing integrality of β over R . Then clearly

$$T \vdash \forall x : (\exists y : P(y, x) = 0 \wedge \psi(y, x)) \rightarrow \phi(x),$$

so $\exists y : P(y, x) = 0 \wedge \psi(y, x)$ is in $\Sigma_0(x)$, contradicting the fact that it holds of α in M . \square

Lemma 3.16. *Let M be a model of T and R be a subring of M . Let $K \subset M$ be the smallest perfect field containing R . Let α be a tuple from R , and $\phi(x)$ be a formula such that $M \models \phi(\alpha)$. If T and the diagram of K imply $\phi(\alpha)$, then T and the diagram of R imply $\phi(\alpha)$.*

Proof. If not, then there is a model N of T extending R , in which $\phi(\alpha)$ fails to hold. This model N must not satisfy the diagram of K . Now N certainly contains a copy of the pure field K , because the fraction field and perfect closure of a domain are unique. Consequently, there must be at least two ways to extend the T -structure from R to K , one coming from M and one coming from N . But this is absurd, because each valuation/ordering/ p -valuation on R extends uniquely to K , by quantifier elimination in the T_i . \square

3.3 Simplifying the axioms down to curves

Lemma 3.17. *Let K be an \aleph_1 -saturated and \aleph_1 -strongly homogeneous model of T_\forall satisfying axioms A1 and A2(1). Let \mathfrak{M} be a monster model of T extending K . Let S be a countable subset of K and α be a countable tuple from \mathfrak{M} . Then $\text{tp}(\alpha/S)$ is realized in K .*

Proof. Consider the following statements:

- A_k : if α is a finite tuple from \mathfrak{M} , with $tr.deg(\alpha/S) \leq k$, then $qftp(\alpha/S)$ is realized in K .
- B_k : if α is a countable tuple from \mathfrak{M} , with $tr.deg(\alpha/S) \leq k$, then $qftp(\alpha/S)$ is realized in K .
- C_k : if α is a countable tuple from \mathfrak{M} , with $tr.deg(\alpha/S) \leq k$, then $tp(\alpha/S)$ is realized in K .

There are several implications between these statements:

- For each k , A_k implies B_k , by compactness.
- For each k , B_k implies C_k . Indeed, if α is as in C_k , apply B_k to $\alpha' := acl(\alpha S)$ and use Corollary 3.11.
- C_k for all k implies the statement of the Lemma, by compactness.

Finally, observe that C_k and C_j imply C_{k+j} : if α has transcendence degree $k + j$ over S , let β be a subtuple of α with transcendence degree k . Then $tr.deg(\beta/S) \leq k$ and $tr.deg(\alpha/\beta S) \leq j$. By C_k , we can apply an automorphism over S to move β inside K . By C_j applied to $tp(\alpha/\beta S)$, we can then find a further automorphism moving α inside K .

Lemma 3.2 and \aleph_1 -saturation of K imply A_1 . By the above comments, this implies C_1 , which in turn implies C_{1+1}, C_3, C_4, \dots . By compactness, the Lemma is true. \square

Theorem 3.18. *A field $K \models T_\forall$ is existentially closed, i.e., a model of T , if and only if it satisfies A1 and A2(1).*

Proof. If K is existentially closed, then certainly K satisfies A1 and A2(1). Conversely, suppose K satisfies A1 and A2(1). Let K' be an \aleph_1 -saturated \aleph_1 -strongly homogeneous elementary extension of K . As $K \equiv K'$, it suffices to show that $K' \models T$. Let \mathfrak{M} be a monster model of T , extending K' . It suffices to show that $K' \preceq \mathfrak{M}$. It suffices to show that if D is a non-empty K' -definable subset of \mathfrak{M} , then D intersects K' . Let S be a finite subset of K' that D is defined over, and let α be a point in D . By Lemma 3.17, $tp(\alpha/S)$ is realized in K' . Such a realization must live in D . \square

Consequently, in checking the axioms one only needs to consider curves. In fact, one only needs to consider smooth curves, by Remark 3.5.

4 A Special Case

In the case where almost every T_i is ACVF, the axioms can be drastically simplified.

Theorem 4.1. *Suppose T_2, \dots, T_n are all ACVF. A model $K \models T_\forall$ is existentially closed (i.e., a model of T) if and only if the following three conditions hold:*

- $K \upharpoonright \mathcal{L}_1 \models T_1$
- Each valuation v_2, \dots, v_n is non-trivial.
- T_i and T_j do not induce the same topology on K , for $i \neq j$.

For example, if we are considering the theory of ordered valued fields, this says that a model is existentially closed if and only if the field is real closed, the valuation is non-trivial, and the ordering and valuation induce different topologies on K . A field with n valuations is existentially closed if and only if it is algebraically closed and the valuations induce distinct non-discrete topologies on the field. Using this, we can easily see that \mathbb{Q}^{alg} with n distinct valuations is an existentially closed field with n valuations. This surprised me, since I expected the Rumely Local-Global principle to be necessary in the proof.

I have no idea whether Theorem 4.1 is already known. I haven't seen a proof, at least. But I haven't looked very hard either.

In the proof of Theorem 4.1, we will use the following fact:

Fact 4.2. *Let K be a field. Let t_1, \dots, t_n be topologies on K arising from orderings and non-trivial valuations. Suppose that $t_i \neq t_j$ for $i \neq j$. Then the $\{t_i\}$ are independent, i.e., if U_i is a non-empty t_i -open subset of K for each i , then $\bigcap_{i=1}^n U_i$ is non-empty. Equivalently, the diagonal map $K \rightarrow \prod_{i=1}^n K$ has dense image with respect to the product topology, using the topology t_i for the i^{th} term in the product.*

This is Theorem II.1.7 in [11] where it is attributed to A. S. Stone. Apparently a “very short proof” appears in [10].

Also, we will need the following straightforward lemma.

Lemma 4.3. *Let K be a model of T . Let C be an affine smooth curve over K , absolutely irreducible. Let \overline{C} be the canonical smooth projective model (as an*

abstract variety). For each i , let $\phi_i(x)$ be a C -dense quantifier-free \mathcal{L}_i -formula with parameters from K . Then we can find a K -definable rational function $f : \overline{C} \rightarrow \mathbb{P}^1$ which is non-constant, and has the property that the divisor $f^{-1}(0)$ is a sum of distinct points in $\bigcap_{i=1}^n \phi_i(K)$, with no multiplicities. (In particular, the support of the divisor contains no points from $C(K^{alg}) \setminus C(K)$ and no points from $\overline{C} \setminus C$.)

Proof. Let g be the genus of \overline{C} .

Claim 4.4. *We can find $g + 1$ distinct points p_1, \dots, p_{g+1} in $\bigcap_{i=1}^n \phi_i(K) \subset C(K)$.*

Proof. For each i , choose a model K_i of T_i extending $K \upharpoonright \mathcal{L}_i$. Then $\phi_i(K_i)$ is Zariski dense in $C(K_i^{alg})$. This (easily) implies that $\phi_i(K_i)^{g+1}$ is Zariski dense in $C^{g+1}(K_i^{alg})$. If U denotes the subset of C^{g+1} consisting of (x_1, \dots, x_{g+1}) such that $x_i \neq x_j$ for every i and j , then U is a Zariski dense Zariski open subset of C^{g+1} , because its complement is a closed subvariety of lower dimension. The intersection of a Zariski dense set with a Zariski dense Zariski open is still Zariski dense. So $\phi_i(K_i)^{g+1} \cap U$ is still Zariski dense in C^{g+1} . Let $\psi_i(x_1, \dots, x_{g+1})$ be the following quantifier-free \mathcal{L}_i -formula over K :

$$\bigwedge_{j=1}^{g+1} \phi_i(x_j) \wedge \bigwedge_{j \neq k} x_j \neq x_k.$$

Then $\psi_i(K_i) = \phi_i(K_i)^{g+1} \cap U$ is Zariski dense in $C^{g+1}(K_i^{alg})$, so $\psi_i(-)$ is C^{g+1} -dense. By Axiom A2, it follows that some tuple (p_1, \dots, p_{g+1}) satisfies

$$\bigwedge_{i=1}^n \psi_i(x_1, \dots, x_{g+1}) \equiv \left(\bigwedge_{i=1}^n \bigwedge_{j=1}^{g+1} \phi_i(x_j) \right) \wedge \left(\bigwedge_{j \neq k} x_j \neq x_k \right).$$

Then (p_1, \dots, p_{g+1}) has the desired properties. \square

Now let D be the divisor $\sum_j p_j$ on the curve \overline{C} . By Riemann-Roch, $l(D) \geq \deg D + 1 - g = 2$. The space of global sections of $\mathcal{O}(D)$ is a K -definable vector space of dimension at least 2. Now K is either algebraically closed or has characteristic zero, so K is perfect. Therefore, by Galois descent of vector spaces we know that this vector space has a K -definable basis. Thinking of the sections of $\mathcal{O}(D)$ as functions with poles no worse than D , we can find a non-constant meromorphic function g , with $(g) - D \geq 0$. Then

the divisor of poles of g is a subset of D , so every pole of g has multiplicity 1 and is in $\bigcap_{i=1}^n \psi_i(K)$. Take $f = 1/g$. \square

Proof (of Theorem 4.1). If $K \models T$, then K satisfies Axioms A1 and A2. Axiom A1 implies that K is algebraically closed or real closed or p -closed (Remark 3.4). As K is existentially closed, it is also reasonably clear that all the named valuations must be non-trivial. Consequently $K \upharpoonright \mathcal{L}_1 \models T_1$ and v_2, \dots, v_n are non-trivial. Lastly, suppose T_i and T_j induce the same topology on K for some i . For notational simplicity assume $i = 1$ and $j = 2$. As the topologies are Hausdorff, we can find non-empty U_1 and U_2 with U_1 a T_1 -open, U_2 a T_2 -open, and $U_1 \cap U_2 = \emptyset$. Since the topologies from T_1 and T_2 have a basis of open sets consisting of quantifier-free definable sets, we can shrink U_1 and U_2 a little, and assume U_1 is quantifier-free definable in \mathcal{L}_1 and U_2 is quantifier-free definable in \mathcal{L}_2 . Now U_1 and U_2 are both Zariski dense in the affine line, so the formulas defining U_1 and U_2 are \mathbb{A}^1 -dense. Hence, by Axiom A2, U_1 must intersect U_2 , a contradiction.

The other direction of the theorem is harder. We proceed by induction on n , the number of orderings and valuations. The base case where $n = 1$ is easy/trivial, so suppose $n > 1$. Suppose K satisfies the assumptions of the Theorem. By Fact 4.2, we know that the n different topologies on K^1 are independent. The first bullet point ensures that K satisfies axiom A1. By Theorem 3.18, it suffices to prove axiom A2(1). By Remark 3.7, we merely need to prove the following:

Let C be an absolutely irreducible smooth affine curve defined over K . Let $\phi_1(x)$ be a quantifier-free \mathcal{L}_1 -formula with parameters from K such that $\phi_1(K)$ is a non-empty open subset of C . For $2 \leq i \leq n$, let $\phi_i(x)$ be a quantifier-free \mathcal{L}_i -formula with parameters from K such that $\phi_i(x)$ defines a non-empty open subset of $C(K^{alg})$ with respect to any/every extension of the i th valuation v_i from K to K^{alg} . THEN $\bigcap_{i=1}^n \phi_i(K)$ is non-empty.

Here we are using the facts that $K \upharpoonright \mathcal{L}_1$ is already a model of T_1 , and that for $i > 1$, the field K^{alg} with any extension of v_i will be a model of $T_i = ACVF$.

For $1 < i \leq n$, choose some extension v'_i of the valuation v_i to K^{alg} .

Claim 4.5. K is dense in K^{alg} with respect to the v'_i -adic topology on K^{alg} .

Proof. The claim is trivial if all the T_i are ACVF, in which case $K = K^{alg}$. So we may assume characteristic zero. It suffices to show that K is dense in

every finite Galois extension of K .² Let L/K be a finite Galois extension. We can write L as $K(\zeta)$ for some singleton ζ . Let $P(X) \in K[X]$ be the minimal polynomial of ζ over K . The function $x \mapsto P(x)$ from K to K is finite-to-one, so it has infinite image. As K is a model of ACVF, pCF, or RCF, we see by Fact 2.1 that the image $P(K)$ of this map contains an open subset of K with respect to the T_1 -topology. Because the v_i -adic topology on K is independent from the T_1 -topology on K , we can find elements of $P(K)$ of arbitrarily high v -valuation. By the cofinality of the value groups, for every $\gamma \in v'_i(K^{alg})$, we can find an $x \in K$ with $v_i(P(x)) > \gamma$. Let $\zeta_1, \dots, \zeta_m \in L$ be the conjugates of ζ over K , counted with multiplicities. Then we have just seen that for any $\gamma \in v'_i(K^{alg})$, we can find an $x \in K$ with

$$\gamma < v'_i(P(x)) = \sum_{i=m}^n v'_i(x - \zeta_i).$$

This implies that at least one of the ζ_i 's is in the topological closure of K with respect to v'_i . Consequently, the v'_i -topological closure of K in L must contain $K[\zeta_i]$ for some i . But $K[\zeta_i] = L$, so K is v'_i -dense in L . \square

Now suppose we are given an absolutely irreducible smooth affine curve C defined over K , and we have \mathcal{L}_i -formulas $\phi_i(x)$ with parameters from K , such that $\phi_1(K)$ is a non-empty open subset of $C(K)$, and for $1 < i \leq n$, $\phi_i(K^{alg})$ is a non-empty v'_i -open subset of $C(K^{alg})$. (Here we are interpreting $\phi_i(K^{alg})$ using v'_i .) By the inductive hypothesis, $K \upharpoonright \bigcup_{i < n} \mathcal{L}_i$ is an existentially closed model of $\bigcup_{i < n} (T_i)_\forall$. Applying Lemma 4.3 to it, we can find a K -definable rational function $f : C \rightarrow \mathbb{P}^1$, whose divisor of zeros has no multiplicities and consists entirely of points in $\bigcap_{i < n} \phi_i(K)$ (and no points at infinity and no points in $C(K^{alg}) \setminus C(K)$). Write this divisor as $\sum_{j=1}^m (P_j)$, where the P_j are m distinct points in $\bigcap_{i < n} \phi_i(K)$. Note that m is the degree of f .

Claim 4.6. *There is a T_1 -open neighborhood $U \subset K$ of zero such that for every $y \in U_1$, the divisor $f^{-1}(y)$ consists of j distinct points in $\phi_1(K)$. In particular, it contains no points in $C(K^{alg}) \setminus C(K)$ and no points in $\overline{C} \setminus C$.*

Proof. Because the P_j are distinct, they have multiplicity one, so f does not have a critical point at any of the P_j 's. Consequently, by the implicit

²Note that the value group $v'_i(K)$ is cofinal in $v'_i(K^{alg})$, so e.g. the v_i -adic topology on K is the restriction of the v'_i -adic topology on K^{alg} to K . Various pathologies are thus avoided.

function theorem there is a T_1 -open neighborhood $W_j \subset C(K)$ of P_j such that f induces a T_1 -homeomorphism from W_j to an open neighborhood of 0. By shrinking W_j if necessary, we may assume that $W_j \subset \phi_1(K)$, and that $W_j \cap W_{j'} = \emptyset$ for $j \neq j'$. Now let $U = \bigcap_{j=1}^m f(W_j)$. This is an open neighborhood of 0 in the affine line K^1 . And if $y \in U$, then $f^{-1}(y)$ contains at least one point in each W_j . Since the W_j are distinct, these points are distinct. Since f is a degree- m map, this exhausts the divisor $f^{-1}(y)$. \square

Claim 4.7. *For $1 < i < n$, there is a $\gamma_i \in v_i(K)$ such that if $y \in K^{alg}$ and $v'_i(y) > \gamma_i$, then $f^{-1}(y)$ are all in $\phi_i(K^{alg})$.*

Proof. Use the same argument as Claim 4.6. \square

By Claim 4.5, K is dense in K^{alg} with respect to the v'_n -adic topology. Also, by assumption, $\phi_n(x)$ interpreted in (K^{alg}, v'_n) yields a non-empty v'_n -open subset $W \subset C(K^{alg})$. Since f is finite-to-one, the image $f(W)$ is an infinite subset of $\mathbb{P}^1(K^{alg})$, hence it has non-empty v'_n -interior. Let V be a v'_n -open subset of $\mathbb{P}^1(K^{alg})$ contained in $f(W)$. Now, as K is v'_n -adically dense in K^{alg} , V must intersect K . In particular, $V \cap K$ is a *non-empty* v_n -adic open subset of K . By independence of the topologies, we can find a y in $\mathbb{A}^1(K)$ such that

- y is in U , the T_1 -open neighborhood of 0 from Claim 4.6.
- $v_i(y) > \gamma_i$, for $1 < i < n$, where the γ_i are from Claim 4.7
- y is in $V \cap K$.

Having chosen such a y , we know by Claim 4.6 that $f^{-1}(y)$ consists of j distinct points in $\phi_1(K)$. In particular, each point in $f^{-1}(y)$ is a point of $C(K)$. And by Claim 4.7, each of these points also belongs to $\phi_i(K^{alg})$, hence satisfies $\phi_i(-)$, for $i < n$. Finally, because y is in $V \cap K$, y is in the image of $\phi_n(K^{alg})$ under f . So there is some $x \in \phi_n(K^{alg})$ mapping to y . But we just said that every point in $C(K^{alg})$ mapping to y is already in $C(K)$ and even in $\bigcap_{i < n} \phi_i(K)$. Thus

$$x \in \phi_n(K^{alg}) \cap \bigcap_{i < n} \phi_i(K) = \bigcap_{i=1}^n \phi_i(K).$$

In particular some point in $C(K)$ satisfies $\bigwedge_{i=1}^n \phi_i(x)$, and the theorem is proven. \square

5 Keisler Measures

To establish NTP_2 and analyze forking and dividing in T , we need the following tool.

Theorem 5.1. *Let T be one of the model companions from §3. For each $K \models T_\forall$ that is a perfect field, each formula $\phi(x)$ and each tuple a from K , we can assign a number $P(\phi(a), K) \in [0, 1]$ such that the following conditions hold:*

- *If K is held fixed, the function $P(-, K)$ is a Keisler measure on the space of completions of the quantifier-free type of K . Thus*

$$P(\phi(a), K) + P(\psi(b), K) = P(\phi(a) \wedge \psi(b), K) + P(\phi(a) \vee \psi(b), K)$$

$$P(\neg\phi(a), K) = 1 - P(\phi(a), K)$$

for sentences $\phi(a)$ and $\psi(b)$ over K . And if $\phi(a)$ holds in every model of T extending K , then $P(\phi(a), K) = 1$. For example, if $\phi(x)$ is quantifier-free, then $P(\phi(a), K)$ is 1 or 0 according to whether or not $K \models \phi(a)$. And if K satisfies axiom A1 of §3.1, then $P(\phi(a), K) \in \{0, 1\}$ for every $\phi(a)$, by Corollary 3.11.

- *Isomorphism invariance: if K, L are two perfect fields satisfying T_\forall , and $f : K \rightarrow L$ is an isomorphism of structures, then $P(\phi(a), K) = P(\phi(f(a)), L)$ for every K -sentence $\phi(a)$.*
- *Irrelevance: if $K_0 \subset K$ are perfect fields satisfying T_\forall , and K_0 is relatively algebraically closed in K , and $\phi(a)$ is a formula with parameters from K_0 , then $P(\phi(a), K_0) = P(\phi(a), K)$.*
- *Density: if $K \models T_\forall$ is a perfect field and $\phi(a)$ is a K -formula, and if $M \models \phi(a)$ for at least one $M \models T$ extending K , then $P(\phi(a), K) > 0$. In other words, the associated Keisler measure is spread out throughout the entire Stone space of completions of $\text{qftp}(K)$.*

5.1 The Algebraically Closed Case

We first prove Theorem 5.1 in the case where every T_i is a model of ACVF, i.e., the case of existentially closed fields with n valuations. Define $P(\phi(a), K)$ as follows. Fix some algebraic closure K^{alg} of K . For each $1 \leq i \leq n$, let

v'_i be an extension to K^{alg} of the i^{th} valuation v_i on K . Choose automorphisms $\sigma_1, \dots, \sigma_n \in \text{Gal}(K^{alg}/K)$ randomly with respect to Haar measure on $\text{Gal}(K^{alg}/K)$. Then

$$K_{\sigma_1, \dots, \sigma_n} := (K^{alg}, v'_1 \circ \sigma_1, v'_2 \circ \sigma_2, \dots, v'_n \circ \sigma_n)$$

is a model of T_{\forall} satisfying axiom A1 of §3.1. In particular, whether or not $\phi(a)$ holds in a model of T extending $K_{\sigma_1, \dots, \sigma_n}$ does not depend on the choice of the model, by Corollary 3.11. Define $P(\phi(a), K)$ to be the probability that $\phi(a)$ holds in any/every model of T extending $K_{\sigma_1, \dots, \sigma_n}$. This probability exists, i.e., the relevant event is measurable, because whether or not $\phi(a)$ holds is determined by the behavior of the valuations on some finite Galois extension L/K , by virtue of Theorem 3.15.

Note that the choice of the v'_i does not matter. If v is a valuation on K and w_1 and w_2 are two extensions of v to K^{alg} , then there is a τ in $\text{Gal}(K^{alg}/K)$ such that $w_1 = w_2 \circ \tau$. I believe this is well-known, and at any rate it is an easy consequence of quantifier-elimination in ACVF. From this, it follows that if σ is a randomly chosen element of $\text{Gal}(K^{alg}/K)$, then $w_1 \circ \sigma$ and $w_2 \circ \sigma$ have the same distribution. Consequently the choice of the valuations v'_i does not effect the resulting value of $P(\phi(a), K)$.

So we have a well-defined number $P(\phi(a), K)$, and it is defined canonically. The first two bullet points of Theorem 5.1 are therefore clear. The density part can be seen as follows: suppose $M \models \phi(a)$ for some $M \models T$ extending K . Let K^{alg} be the algebraic closure of K in M . For the v'_i 's, take the restrictions of the valuations on M to K^{alg} . By Theorem 3.15, there is a field $K \leq L \leq K^{alg}$ with L/K a finite Galois extension, such that $\phi(a)$ is implied by T and the diagram of L . Specifically, write $\phi(a)$ as $\exists y : P(y; a) = 0 \wedge \psi(y; a)$, and let L be the splitting field of the polynomial $P(X; a) \in K[X]$. Now with probability $1/[L : K]^n$, every σ_i will restrict to the identity on L . Consequently, $K_{\sigma_1, \dots, \sigma_n}$ will be a model of T_{\forall} extending L , so in any model M of T extending $K_{\sigma_1, \dots, \sigma_n}$, $\phi(a)$ will hold. So $\phi(a)$ holds with probability at least $1/[L : K]^n$, and consequently $P(\phi(a), K) \geq 1/[L : K]^n$.

It remains to verify the irrelevance part of Theorem 5.1. Let $K_0 \leq K$ be an inclusion of perfect fields, with K_0 relatively algebraically closed in K . Let $\phi(a)$ be a formula with parameters a from K_0 . As in the previous paragraph, write $\phi(a)$ as $\exists y : P(y; a) \wedge \psi(y; a)$ and let L_0 be the splitting field of $P(y; a)$ over K_0 . At present L_0 is just a pure field. Write $L_0 = K_0(\beta)$ for some singleton $\beta \in L_0$, and let $Q(X)$ be the irreducible polynomial of

β over K_0 . Let $L = L_0K = K(\beta)$; this is a Galois extension of K . There are only finitely many ways of factoring $Q(X)$ in K^{alg} , so in each way of factoring $Q(X)$, the coefficients come from K_0^{alg} . In particular, if $Q(X)$ can be factored over K , the coefficients would belong to $K_0^{alg} \cap K = K_0$. So $Q(X)$ is still irreducible over K . Consequently $[L : K] = \deg Q(X) = [L_0 : K_0]$. Now there is a natural restriction map $\text{Gal}(L/K) \rightarrow \text{Gal}(L_0/K_0)$. It is injective because an element of $\text{Gal}(K(\beta)/K)$ is determined by what it does to β . Since $\text{Gal}(L/K)$ has the same size as $\text{Gal}(L_0/K_0)$, the restriction map must be an isomorphism. Consequently, if τ is chosen from $\text{Gal}(L/K)$ randomly, its restriction to L_0 is a random element of $\text{Gal}(L_0/K_0)$. Consequently, if σ is a random element of $\text{Gal}(K^{alg}/K)$ and σ_0 is a random element of $\text{Gal}(K_0^{alg}/K_0)$, then $\sigma \upharpoonright L_0$ and $\sigma_0 \upharpoonright L_0$ have the same distribution, namely, the uniform distribution on $\text{Gal}(L_0/K_0)$. From this, it follows easily that $P(\phi(a), K) = P(\phi(a), K_0)$.

This completes the proof of Theorem 5.1 when every T_i is ACVF. In the other cases, things are more complicated, though as a consolation all fields are characteristic zero, hence perfect.

A first attempt at defining $P(\phi(a), K)$ is as follows: fix some algebraic closure K^{alg} of K . For each i such that T_i is RCF, let K_i be a real closure of $(K, <_i)$ inside K^{alg} . For each i such that T_i is p CF, let K_i be a p -closure of (K, v_i) inside K^{alg} . For each i such that T_i is ACVF, let K_i be K^{alg} with some valuation extending v_i . In each case, there is a choice, but any two choices are related by an element of $\text{Gal}(K^{alg}/K)$. Now choose $\sigma_1, \dots, \sigma_n \in \text{Gal}(K^{alg}/K)$ randomly. For each i , consider $\sigma_i(K_i)$, which is (usually) a model of T_i extending K . Let K' be the field

$$K' = \bigcap_{i=1}^n \sigma_i(K_i).$$

There is an obvious way to give K' the structure of a T_V -model. If we knew that K' satisfies condition A1 of §3.1 with high probability, we could define $P(\phi(a), K)$ to be the probability that $\phi(a)$ holds in any/every model of T extending K' . Unfortunately, K' usually satisfies condition A1 with probability *zero*. Instead, we'll proceed by repeating the above procedure with K' in place of K , getting a third field K'' . Iterating this, we get an increasing sequence $K \subset K' \subset K'' \subset \dots \subset K^{(n)} \subset \dots$ of T_V -structures on subfields of K^{alg} . The union $K^\infty = \bigcup_{n=1}^\infty K^{(n)}$ does actually turn out to satisfy axiom A1 with probability 1, and we let $P(\phi(a), K)$ be the probability that $\phi(a)$ holds in any/every model of T extending K^∞ .

The rest of this section will make this construction more precise, and verify that it satisfies the requirements of Theorem 5.1.

5.2 The General Case

All fields will be perfect, unless stated otherwise. All models of T_\forall and $(T_i)_\forall$ will be (perfect) fields, unless stated otherwise. Galois extensions needn't be finite Galois extensions.

We start off with some easy but confusing facts that will be needed later.

Lemma 5.2. *Let L/K be a Galois extension of fields, and suppose K has the structure of a $(T_i)_\forall$ model (but L does not). The following are equivalent*

- (a) *For every F , if F is a model of $(T_i)_\forall$ extending K , and F is a subfield of L , then $F = K$.*
- (b) *There is a model $M \models T_i$ extending K , such that $M \cap L = K$.*
- (c) *For every model $M \models T_i$ extending K , $M \cap L = K$.*

Note that it makes sense to talk about whether $M \cap L = K$, because L/K is Galois.

Proof. The equivalence of (b) and (c) follows from quantifier elimination in T_i . Indeed, the statement that $M \cap L = K$ is equivalent to the statement that for each $x \in L \setminus K$, the irreducible polynomial of x over K has no zeros in M . This is a conjunction of first order statements about K , so it holds in one choice of M if and only if it holds in another choice of M .

Suppose (a) holds. Let M be a model of T_i extending K . Taking $F = M \cap L$, (a) implies that $M \cap L = K$. So (a) implies (c).

Conversely, suppose (a) does not hold. Let F witness a contradiction to (a), so $K \subsetneq F \subseteq L$, and F is a model of $(T_i)_\forall$ extending K . Let M be a model of T_i extending F and hence K . Then $M \cap L$ contains F , contradicting (c). \square

Definition 5.3. *Say that K is locally T_i -closed in L if it satisfies the equivalent conditions of the previous lemma.*

Definition 5.4. *Let L/K be a Galois extension of fields, and suppose K has the structure of a $(T_i)_\forall$ -model (but L does not). Let $\mathfrak{C}_i(L/K)$ denote the set of models of $(T_i)_\forall$ which extend K , are subfields of L , and are locally T_i -closed in L .*

The subscript on \mathfrak{C}_i is present so that $\mathfrak{C}_i(L/K)$ will be unambiguous when K is a model of T_\forall , not just a model of $(T_i)_\forall$.

There is a natural action of $\text{Gal}(L/K)$ on $\mathfrak{C}_i(L/K)$.

Lemma 5.5. *Suppose L/K is a Galois extension of fields, and $K \models (T_i)_\forall$.*

- (a) *The action of $\text{Gal}(L/K)$ on $\mathfrak{C}_i(L/K)$ has exactly one orbit.*
- (b) *Suppose K' is a model of $(T_i)_\forall$ extending K , and L' is a field extension of L and K' , with L' Galois over K' . If $F \in \mathfrak{C}_i(L'/K')$, then $F \cap L \in \mathfrak{C}_i(L/L \cap K')$.*

Proof. (a) Note that $\mathfrak{C}_i(L/K)$ is non-empty by a Zorn's lemma argument and condition (a) of Lemma 5.2. Now suppose F and F' are two elements of $\mathfrak{C}_i(L/K)$. By quantifier elimination in T_i , we can amalgamate F and F' over K . Thus, we can find a model $M \models T$ extending F , and an embedding $\iota : F' \rightarrow M$ which is the identity on K . Choosing some way of amalgamating M and L as fields, we get that $\iota(F') \subset L \supset F$, because of L/K being Galois. The compositum $\iota(F')F$ is a subfield of L with a $(T_i)_\forall$ -structure extending that on F and $\iota(F')$, so by local T_i -closedness of $\iota(F')$ and F in L , $\iota(F') = \iota(F')F = F$. It follows that F' and F are isomorphic over K . This isomorphism must extend to an automorphism of L , because L/K is Galois. So some automorphism on L/K maps F' to F (as $(T_i)_\forall$ -structures).

- (b) Let M be a model of T_i extending F . Choose some way of amalgamating M with L' . Then $M \cap L' = F$ by (c) of Lemma 5.2. Therefore, $M \cap L = M \cap L' \cap L = F \cap L$. So by (b) of Lemma 5.2, $F \cap L$ is locally T_i -closed in L . Therefore it is in $\mathfrak{C}_i(L/L \cap K')$. □

Now we turn our attention from T_i to T .

Definition 5.6. *Let $K \models T_\forall$ and let L be a pure field that is a Galois extension of K . Let $\mathfrak{S}(L/K)$ be the set of all $K' \models T_\forall$ extending K , with K' a subfield of L . In other words, an element of $\mathfrak{S}(L/K)$ is a subfield F of L , endowed with a T_\forall -structure, such that $F \supseteq K$ and the structure on F extends the structure on K .*

There is a natural partial order on $\mathfrak{S}(L/K)$ coming from inclusion of substructures. There is also a natural action of $\text{Gal}(L/K)$ on $\mathfrak{S}(L/K)$. One should think of $\mathfrak{S}(L/K)$ as the set of states in a Markov chain, specifically the random process described at the end of the previous section.

Definition 5.7. *Suppose $K \models T_\forall$ and L/K is a Galois extension of K . For $1 \leq i \leq n$, choose some $L_i \in \mathfrak{C}_i(L/K)$. Choose $\sigma_1, \dots, \sigma_n \in \text{Gal}(L/K)$ independently and randomly, using Haar measure on $\text{Gal}(L/K)$. Let F be $\bigcap_{i=1}^n \sigma_i(L_i)$, with the obvious choice of a T_\forall structure. So F is a random variable with values in $\mathfrak{S}(L/K)$. Let $\mu_{L/K}^1$ be the probability distribution on $\mathfrak{S}(L/K)$ obtained in this way. The choice of the L_i 's doesn't matter, by Lemma 5.5(a).*

The superscript 1 is to indicate that this is the first step of the Markov chain.

Lemma 5.8. *Suppose L/K is finite. Then every event (subset of $\mathfrak{S}(L/K)$) which has positive probability with respect to $\mu_{L/K}^1$ has probability at least $1/m^n$, where $m = [L : K]$.*

Proof. The only randomness comes from the σ_i 's. Each element of $\text{Gal}(L/K)$ has an equal probability under Haar measure, and this probability is $1/m$. Since the σ_i 's are chosen independently, each choice of the σ_i 's has probability $1/m^n$ of occurring. \square

Lemma 5.9. *Suppose L/K is finite, and F is a maximal element of $\mathfrak{S}(L/K)$. Then $\mu_{L/K}^1(\{F\}) > 0$.*

Proof. For each i , let M_i be a model of T_i extending $F \upharpoonright \mathcal{L}_i$, and choose some way of amalgamating M_i and L as fields over F . Let $F_i = L \cap M_i$. Of course $F_i \supset F$. By Lemma 5.2(b), $F_i \in \mathfrak{C}_i(L/K)$. Let $F' = \bigcap_{i=1}^n F_i$. Then $F' \in \mathfrak{S}(L/K)$ and F' extends F , so $F = F'$ by maximality of F . Now if we choose $\sigma_1, \dots, \sigma_n \in \text{Gal}(L/K)$ randomly, then $\bigcap_{i=1}^n \sigma_i(F_i)$ is distributed according to $\mu_{L/K}^1$. Since L/K is finite, there is a positive probability that $\sigma_i = 1$ for every i , in which case $\bigcap_{i=1}^n \sigma_i(F_i) = F' = F$. \square

Lemma 5.10. *Let L/K be a Galois extension, and K be a model of T_\forall . Let K' be a model of T_\forall extending K . Let L' be a field extending L and K' , with L' Galois over K' . If $F \in \mathfrak{S}(LK'/K')$ is distributed randomly according to $\mu_{L'/K'}^1$, then $F \cap L$ is distributed randomly according to $\mu_{L/L \cap K'}^1$.*

Proof. For $1 \leq i \leq n$, choose some $F_i \in \mathfrak{C}_i(L'/K')$. By Lemma 5.5(b), $F_i \cap L$ is in $\mathfrak{C}_i(L/L \cap K')$.

Claim 5.11. *If we choose σ from $\text{Gal}(L'/K')$ randomly using Haar measure, then $\sigma \upharpoonright L$ is also randomly distributed in $\text{Gal}(L/L \cap K')$ with respect to Haar measure.*

Proof. It probably suffices to show that if $x \in L$, then the probability that $\sigma \upharpoonright L$ fixes x is the same as the probability that a random element of $\text{Gal}(L/L \cap K')$ fixes x . This is equivalent to showing that x has the same number of conjugates under the action of $\text{Gal}(L/L \cap K')$ as under the action of $\text{Gal}(L'/K')$. If x and y are conjugate over K' , then clearly they are conjugate over $L \cap K'$. Conversely, if they are conjugate over $K' \cap L$ but not over K' , then let $S \subset L$ be the orbit of x under $\text{Gal}(L'/K')$. The code for the finite set S is definable over K' . It is also definable over L , as $S \subset L$. Thus the code for S is in $L \cap K'$. As x and y are conjugate over $L \cap K'$, we conclude that $y \in S \iff x \in S$, a contradiction. \square

From the Claim, we conclude that if the σ_i are distributed randomly from $\text{Gal}(L'/K')$, then $\sigma_i \upharpoonright L$ are distributed randomly from $\text{Gal}(L/L \cap K')$. Taking $F = \bigcap_{i=1}^n \sigma_i(F_i)$, we get F distributed according to $\mu_{L'/K'}^1$. But

$$F \cap L = \bigcap_{i=1}^n (\sigma_i \upharpoonright L)(F_i \cap L)$$

is then distributed according to $\mu_{L/L \cap K'}^1$, because $F_i \cap L \in \mathfrak{C}_i(L/(K' \cap L))$ and $\sigma_i \upharpoonright L$ is distributed according to Haar measure on $\text{Gal}(L/L \cap K')$. \square

Definition 5.12. *Let L/K be a Galois extension, and K be a model of T_V . Define a series of distributions $\{\mu_{L/K}^i\}_{i < \omega}$ on $\mathfrak{S}(L/K)$ as follows:*

- $\mu_{L/K}^0$ assigns probability 1 to $\{K\} \subset \mathfrak{S}(L/K)$.
- $\mu_{L/K}^1$ is as above.
- For $i > 0$, if we choose $F \in \mathfrak{S}(L/K)$ randomly according to $\mu_{L/K}^i$, and then choose $F' \in \mathfrak{S}(L/F) \subset \mathfrak{S}(L/K)$ randomly according to $\mu_{L/F}^1$, then F' is distributed according to $\mu_{L/K}^{i+1}$.

In other words, we are running some kind of Markov chain whose states are the elements of $\mathfrak{S}(L/K)$. The transition probabilities out of the state F are given by $\mu_{L/F}^1$, and $\mu_{L/K}^n$ is the distribution of the Markov chain after n steps. There could be some weird technical problem with Definition 5.12 when L/K is an infinite Galois extension, since the sample space from which the σ_i 's are drawn depends on previous events in the chain. I don't know enough measure theory to verify that this is legitimate, so we'll try to stick to the case where L/K is finite and everything is valid.

Lemma 5.13. *Let L/K be a finite Galois extension, and K be a model of T_\forall . Then $\lim_{i \rightarrow \infty} \mu_{L/K}^i$ exists, and the corresponding distribution on $\mathfrak{S}(L/K)$ is concentrated on the maximal elements of $\mathfrak{S}(L/K)$.*

Proof. The fact that the limit distribution exists is a general fact about Markov chains with finitely many states such that the graph of possible transitions has no cycles other than self-loops.

It remains to check that in the limit, we land in a maximal element of $\mathfrak{S}(L/K)$ with probability one. Let $m = [L : K]$. If $F \in \mathfrak{S}(L/K)$ is not maximal, then the probability of moving from F to some bigger element is positive by Lemma 5.9, and at least $1/m^n$, by Lemma 5.8. The probability of getting stuck at F is therefore bounded above by $\lim_{k \rightarrow \infty} (1 - 1/m^n)^k = 0$. As there are finitely many non-maximal F , we conclude that the probability of getting stuck at any of them is zero. \square

We let $\mu_{L/K}^\infty$ denote the limit distribution on $\mathfrak{S}(L/K)$.

Lemma 5.14. *Let L/K be a finite Galois extension, and K be a model of T_\forall . Then every maximal element of $\mathfrak{S}(L/K)$ has a positive probability with respect to $\mu_{L/K}^\infty$.*

Proof. This follows immediately from Lemma 5.9, and the fact that once the Markov chain reaches a maximal element of $\mathfrak{S}(L/K)$, it must remain there. \square

Lemma 5.15. *Let L/K be a Galois extension, with K a model of T_\forall . Let K' be a model of T_\forall extending K . Let L' be a field extending K' and L , Galois over K' . If F is a random element of $\mathfrak{S}(L'/K')$ distributed according to $\mu_{L'K'/K'}^i$, then $F \cap L$ is distributed according to $\mu_{L/L \cap K'}^i$.*

Proof. We proceed by induction on i . For $i = 0$, F is guaranteed to be K' , and $F \cap L$ is guaranteed to be $K' \cap L$, which agrees with $\mu_{L/L \cap K'}^0$.

For the inductive step, suppose we know the statement of the lemma for μ^i , and prove it for μ^{i+1} . If we let $F \in \mathfrak{S}(L'/K')$ be chosen according to $\mu_{L'/K'}^i$, and we then choose $F' \in \mathfrak{S}(L'/F) \subset \mathfrak{S}(L'/K')$ according to $\mu_{L'/F}^1$, then F' is randomly distributed according to $\mu_{L'/K'}^{i+1}$, by definition of μ^{i+1} . Also, $F' \cap L$ is distributed according to $\mu_{L/L \cap K'}^i$, by the inductive hypothesis. By Lemma 5.10 we know that $F' \cap L$ is distributed according to $\mu_{L/L \cap F}^1$. In particular, the distribution of $F' \cap L$ only depends on $F \cap L$. So if we want to sample $F' \cap L$, we can just choose $F \cap L$ using $\mu_{L/L \cap K'}^i$, and can then choose $F' \cap L$ using $\mu_{L/F \cap L}^1$. This is the recipe for sampling the distribution $\mu_{L/K' \cap L}^{i+1}$. So $F' \cap L$ is indeed distributed according to $\mu_{L/K' \cap L}^{i+1}$. \square

Corollary 5.16. *When L/K and L'/K' are finite Galois extension, Lemma 5.15 holds for $i = \infty$.*

Definition 5.17. *Let $K \models T_{\forall}$ be a perfect field, $\phi(a)$ be a formula in the language of T with parameters a from K . Say that a finite Galois extension L/K determines the truth of $\phi(a)$ if the following holds: whenever M and M' are two models of T extending K , if $M \cap L$ is isomorphic as a model of T_{\forall} to $M' \cap L$, then $[M \models \phi(a)] \iff [M' \models \phi(a)]$. (Note that the isomorphism class of $M \cap L$ doesn't depend on how we choose to form the compositum ML .)*

For every formula $\phi(a)$, there is some finite Galois extension L/K which determines the truth of $\phi(a)$. Namely, use Theorem 3.15 to write $\phi(a)$ in the form $\exists y : P(y; a) = 0 \wedge \psi(y; a)$, and take L to be the splitting field over K of $P(X; a) \in K[X]$.

Lemma 5.18. *Let K be a model of T_{\forall} , M be a model of T extending K , and let L/K be a Galois extension of K . Assume M and L are embedded over K into some bigger field. Then $M \cap L$ is a maximal element of $\mathfrak{S}(L/K)$.*

Proof. Suppose not. Let F be an element of $\mathfrak{S}(L/K)$, strictly bigger than $M \cap L$, and finitely generated over $M \cap L$. Let x be a generator of F over $M \cap L$. If S denotes the set of algebraic conjugates of x over M , then the code for the finite set S is in M , and also in L because $S \subset L$. So the code for S is in $M \cap L$, implying that S is also the set of algebraic conjugates of x over $M \cap L$. Since we're assuming that all fields are perfect, this implies that the

degree of x over M is the same as the degree of x over $M \cap L$. In particular, the irreducible polynomial $P(X)$ of x over $M \cap L$ remains irreducible over M . For $1 \leq i \leq n$, let M_i be a model of T_i extending $M \upharpoonright \mathcal{L}_i$. Let N_i be a model of T_i extending $F \upharpoonright \mathcal{L}_i$. The polynomial $P(X)$ has a zero in F , namely x . Hence it has a zero in $N_i \supseteq F$. As M_i and N_i are two models of T_i extending $M \cap L$ and $P(X)$ is defined over $M \cap L$, it follows from quantifier-elimination in T_i that $P(X)$ also has a zero in M_i .

Now we have a polynomial $P(X)$ of degree > 1 , irreducible over M , such that $P(X)$ has a root in M_i for every i . This contradicts condition A1' of §3.1. \square

Definition 5.19. *Let L/K be a Galois extension, with K a model of T_\forall . Let $\mathfrak{F}(L/K)$ be the set of maximal elements of $\mathfrak{S}(L/K)$.*

By Zorn's lemma, it is clear that every element of $\mathfrak{S}(L/K)$ is bounded above by an element of $\mathfrak{F}(L/K)$, even if L/K is infinite. When L/K is a finite extension, $\mu_{L/K}^\infty$ induces a probability distribution on $\mathfrak{F}(L/K)$. (This probably also works if L/K is infinite, but I'm not sure.)

Remark 5.20. *$\mathfrak{F}(L/K)$ is exactly the set of F of the form $L \cap M$, where M is a model of T extending K . One inclusion is Lemma 5.18. The other inclusion is obvious: if F is a maximal element of $\mathfrak{S}(L/K)$, then letting M be a model of T extending F , and combining M and L into a bigger field in any way we like, we have $F \subseteq M \cap L \in \mathfrak{S}(L/K)$, so maximality of F forces $M \cap L = F$.*

Suppose that L/K determines the truth of $\phi(a)$. Then by Remark 5.20, there must be a uniquely determined map $f_{\phi(a),L}$ from $\mathfrak{F}(L/K)$ to $\{\perp, \top\}$ such that for every $M \models T$ extending K , and every way of forming the compositum ML , the truth of $M \models \phi(a)$ is given by $f_{\phi(a),L}(M \cap L)$.

Another corollary of Remark 5.20 is that if $K \leq L \leq L'$, with L' and L Galois extensions of $K \models T_\forall$, and if $F \in \mathfrak{F}(L'/K)$, then $F = M \cap L'$ for some model M , and hence $F \cap L = M \cap L' \cap L = M \cap L$ is in $\mathfrak{F}(L/K)$.

Finally, we define $P(\phi(a), K)$ to be $\mu_{L/K}^\infty(\{F : f_{\phi(a),L}(F) = \top\})$.

Lemma 5.21. *The choice of L does not matter.*

Proof. If L and L' are two finite Galois extensions of K which determine the truth of $\phi(a)$, then so does their compositum LL' . So we may assume $L \subset L'$. Let $r : \mathfrak{F}(L'/K) \rightarrow \mathfrak{F}(L/K)$ be the restriction map, $F \mapsto F \cap L$.

Claim 5.22. $f_{\phi(a),L'} = f_{\phi(a),L} \circ r$.

Proof. For $F \in \mathfrak{F}(L'/K)$, we will show $f_{\phi(a),L'}(F) = f_{\phi(a),L}(r(F))$. Write F as $M \cap L'$, with M a model of T extending F . Then $f_{\phi(a),L'}(M \cap L') = f_{\phi(a),L'}(F)$ is the truth value of $M \models \phi(a)$. But $M \cap L = F \cap L$, so by definition of $f_{\phi(a),L}$, we also know that $f_{\phi(a),L}(M \cap L) = f_{\phi(a),L}(r(F))$ is the truth value of $M \models \phi(a)$, which is the same thing. So $f_{\phi(a),L}(r(F)) = f_{\phi(a),L'}(F)$. \square

By Corollary 5.16 applied in the $K = K'$ case, if F is a random element of $\mathfrak{F}(L'/K)$ chosen according to $\mu_{L'/K}^\infty$, then $r(F) = F \cap L$ is distributed according to $\mu_{L/K}^\infty$. In particular, the probability of $f_{\phi(a),L'}(F)$ or equivalently of $f_{\phi(a),L}(r(F))$ is the same as the probability of $f_{\phi(a),L}(F')$, with F' chosen directly from $\mu_{L/K}^\infty$. But the former probability is $P(\phi(a), K)$ computed using L' , while the latter is $P(\phi(a), K)$ computed using L . \square

So $P(\phi(a), K)$ is at least a well-defined number. The ‘‘isomorphism invariance’’ part of Theorem 5.1 is clear from the definitions. We need to prove the other conditions of Theorem 5.1.

Lemma 5.23. *For any fixed K , the function $P(-, K)$ is a Keisler measure on the space of completions of the quantifier-free type of K .*

Proof. It suffices to prove the following:

- If $\phi(a)$ and $\psi(b)$ are forced to be equivalent by T and the diagram of K , then $P(\phi(a), K) = P(\psi(b), K)$. This is easy/trivial, because if we choose a finite Galois extension L determining the truth of both $\phi(a)$ and $\psi(b)$, we see that $f_{\phi(a),L} = f_{\psi(b),L}$ by unwinding the definitions.
- $P(\phi(a), K) = 1 - P(\neg\phi(a), K)$, which follows similarly, though it uses the fact that $\mu_{L/K}^\infty$ is concentrated on $\mathfrak{F}(L/K)$.
- If $\phi(a) \wedge \psi(b)$ contradicts $T \cup \text{diag}(K)$, then $P(\phi(a) \vee \psi(b), K) = P(\phi(a), K) + P(\psi(b), K)$. Again, this is not difficult: if L is a field determining the truth of both $\phi(a)$ and $\psi(b)$, then it is clear that

$$f_{\phi(a),L} \wedge f_{\psi(b),L} = f_{\phi(a) \wedge \psi(b),L} = \perp$$

$$f_{\phi(a),L} \vee f_{\psi(b),L} = f_{\phi(a) \vee \psi(b),L}.$$

Consequently, $\{F : f_{\phi(a) \vee \psi(b),L}(F) = \top\}$ is a disjoint union of $\{F : f_{\phi(a),L}(F) = \top\}$ and $\{F : f_{\psi(b),L}(F) = \top\}$, so we reduce to the fact that μ^∞ is a probability distribution.

- $0 \leq P(\phi(a), K) \leq 1$, which is clear from the definition.

Assuming I haven't forgotten any of the axioms of a finitely additive probability distribution, this finishes the lemma. \square

Lemma 5.24. *If $K \subset K'$ are models of T_{\forall} and K is relatively algebraically closed in K' , and $\phi(a)$ is a formula with parameters from K , then $P(\phi(a), K) = P(\phi(a), K')$.*

(This is the “irrelevance” part of Theorem 5.1.)

Proof. Let L be a finite Galois extension of K determining the truth of $\phi(a)$. Let L' be a finite Galois extension of K' determining the truth of $\phi(a)$; we may assume $L' \supseteq L$. (In fact, we can take $L' = LK'$.) Because K is relatively algebraically closed in K' , $L \cap K' = K$. So by Corollary 5.16, if $F \in \mathfrak{F}(L'/K')$ is distributed according to $\mu_{L'/K'}^{\infty}$, then $F \cap L$ is distributed according to $\mu_{L/K}^{\infty}$. Using Lemma 5.14, this implies that $F \cap L \in \mathfrak{F}(L/K)$ for any $F \in \mathfrak{F}(L'/K')$. Let $r : \mathfrak{F}(L'/K') \rightarrow \mathfrak{F}(L/K)$ be the map $F \mapsto F \cap L$. By unwinding the definitions (as in the claim in the proof of Lemma 5.21), one sees that $f_{\phi(a), L'/K'} = f_{\phi(a), L/K} \circ r$. So just like in the proof of Lemma 5.21, we see that for F' chosen randomly from $\mathfrak{F}(L'/K')$ and F chosen randomly from $\mathfrak{F}(L/K)$, the distribution of F and $r(F')$ is the same, and therefore so too is the distribution of

$$f_{\phi(a), L'/K'}(F') = f_{\phi(a), L/K}(r(F'))$$

and

$$f_{\phi(a), L/K}(F).$$

This ensures that $P(\phi(a), K) = P(\phi(a), K')$. \square

Lemma 5.25. *If $K \models T_{\forall}$ and $\phi(a)$ is a K -formula which holds in some model of T extending K , then $P(\phi(a), K) > 0$.*

(This is the “density” part of Theorem 5.1.)

Proof. Let M be the model where $\phi(a)$ holds. Let L be a Galois extension of K determining the truth of $\phi(a)$. Then $L \cap M \in \mathfrak{F}(L/K)$ and $f_{\phi(a), L}(L \cap M) = \top$. By Lemma 5.14, $P(\phi(a), K) > 0$. \square

We have verified each condition of Theorem 5.1, which is now proven.

6 NTP2 and the Independence Property

We show that the model companion T (usually) doesn't have NIP, but always has NTP_2 , the next best possibility. In many cases, perhaps all, these results are already known, though they may not have been written down yet.

6.1 Failure of NIP

If $n = 1$, then $T = T_i$ is one of ACVF, RCF, or $p\text{CF}$, which are all known to be NIP. On the other hand,

Theorem 6.1. *Suppose $n > 1$. Then T has the independence property.*

Other people have already pointed this out, but here is a proof.

Proof. We give a proof which works in characteristic $\neq 2$. It is not hard to modify it to work in characteristic 2.

Claim 6.2. *For each i , we can produce quantifier-free \mathcal{L}_i -formulas $\phi_i(x, y)$ and $\chi_i(y)$ without parameters such that x, y are singletons, and such that if $K_i \models T_i$, then $\chi_i(K_i)$ is a non-empty open set and for every $b \in \chi_i(K_i)$, both square roots of b are present in K_i , and exactly one of them satisfies $\phi_i(x, b)$.*

Proof. If T_i is RCF, let $\chi_i(y)$ say that $y > 0$ and $\phi_i(x, y)$ say that $x > 0$. If T_i is ACVF, let $\chi_i(y)$ say that $v(y - 1/4) > 0$, and $\phi_i(x, y)$ say that $v(x - 1/2) > 0$. Note that if $v(y - 1/4) > 0$ and $x^2 = y$, then $t = x - 1/2$ satisfies

$$t^2 + t + 1/4 - y = (t + 1/2)^2 - y = 0.$$

By Newton polygons, one of the possibilities for t has valuation zero, and the other has valuation $v(y - 1/4) > 0$. If T_i is $p\text{CF}$, the same formulas work as in the case of ACVF. The only thing to check is that if $v(y - 1/4) > 0$ for some $y \in K \models p\text{CF}$, then the two roots of $T^2 + T + (1/4 - y) = 0$ are present in K . If not, then since the two roots have different valuations (in an ambient model of ACVF), there are two different ways to extend the valuation from K to $K[T]/(T^2 + T + (1/4 - y))$, contradicting Henselianity of K . \square

Given the ϕ_i and χ_i from the Claim, let $\chi(y) = \bigwedge_{i=1}^n \chi_i(y)$. Note that $\chi(y)$ defines an infinite subset of any model of T , by condition A2 of §3.1. (Each $\chi_i(-)$ is \mathbb{A}^1 -dense.) If $K \models T$ and $b \in \chi(K)$, then $X^2 - b$ has roots in

K by choice of $\chi_i(-)$ and condition A1' of §3.1. So each element of $\chi(K)$ is a square.

Let $\psi(y)$ assert that $\chi(y)$ holds and there is a square root of y which satisfies exactly one of ϕ_1 and ϕ_2 . Note that if $\chi(b)$ holds, then both square roots of b are present in K , exactly one of them satisfies ϕ_1 , and exactly one of them satisfies ϕ_2 . If we let \oplus denote exclusive-or, then we can write $\psi(y)$ as $\phi_1(\sqrt{y}) \oplus \phi_2(\sqrt{y})$, where the choice of \sqrt{y} doesn't matter.

Let K be a model of T . We will show that $\psi(x+y)$ has the independence property in K . Let a_1, \dots, a_m be any m elements in $\chi(K)$, which as we noted above is an infinite set. We will show that for any subset $S_0 \subset \{1, \dots, m\}$, there is a b in K such that $j \in S_0 \iff K \models \psi(b+a_j)$. It suffices to find such a b in an elementary extension of K , rather than K itself. Let $K' \succeq K$ be an elementary extension containing an element ϵ which is infinitesimal compared to K , with respect to every one of the valuations. That is, for each i such that T_i is valuative, we want $v_i(\epsilon) > v_i(K)$, and for each i such that T_i is RCF, we want $-\alpha <_i \epsilon <_i \alpha$ for every $\alpha >_i 0$ in K . The fact that such an ϵ exists follows by our axiom A2, and isn't hard to show directly.

Note that for $1 \leq j \leq m$, $a_j + \epsilon \in \chi(K')$. (Indeed, for every i , $K' \models \chi_i(a_j + \epsilon)$, because $\chi_i(-)$ defines an open set in a model of T_i , and ϵ is infinitesimal with respect to the prime model of T_i over $K \upharpoonright \mathcal{L}_i$.) Consequently, $\sqrt{a_j + \epsilon} \in K'$ for every $1 \leq j \leq m$. Let L be $K(\sqrt{a_j + \epsilon} : 1 \leq j \leq m) \subset K'$, as a model of T_\forall . Since ϵ is transcendental over K , $\text{Gal}(L/K(\epsilon)) \cong (\mathbb{Z}/2\mathbb{Z})^m$. In particular, for every $S \subset \{1, \dots, m\}$, there is a field automorphism $\sigma_S \in \text{Gal}(L/K(\epsilon))$ which swaps the square roots of $a_j + \epsilon$ if and only if $j \in S$. Let L_S be the T_\forall -model with underlying field L , with the same \mathcal{L}_i -structure as L for $i > 1$, and with the \mathcal{L}_1 -structure obtained by pulling back the \mathcal{L}_1 -structure of L along σ_S . If Δ denotes symmetric difference of sets, then

$$\begin{aligned} \{j : L_S \models \phi_2(\sqrt{a_j + \epsilon})\} &= \{j : L \models \phi_2(\sqrt{a_j + \epsilon})\} \\ \{j : L_S \models \phi_1(\sqrt{a_j + \epsilon})\} &= \{j : L \models \phi_1(\sigma_S(\sqrt{a_j + \epsilon}))\} \\ &= \{j : L \models \phi_1(\sqrt{a_j + \epsilon})\} \Delta S, \end{aligned}$$

where the last equality holds because $L \models \phi_1(\sqrt{a_j + \epsilon}) \iff \neg \phi_1(-\sqrt{a_j + \epsilon})$. Now let K_S be a model of T extending L_S . Since L_S is a model of T_\forall extending

$K, K_S \succeq K$. Also,

$$\begin{aligned} \{j : K_S \models \psi(a_j + \epsilon)\} &= \{j : L_S \models \phi_1(\sqrt{a_j + \epsilon})\} \Delta \{j : L_S \models \phi_2(\sqrt{a_j + \epsilon})\} \\ &= \{j : L \models \phi_1(\sqrt{a_j + \epsilon})\} \Delta \{j : L \models \phi_2(\sqrt{a_j + \epsilon})\} \Delta S \\ &= \{j : K' \models \psi(a_j + \epsilon)\} \Delta S \end{aligned}$$

Therefore, by choosing $S = S_0 \Delta \{j : K' \models \psi(a_j + \epsilon)\}$, we can arrange that

$$\{j : K_S \models \psi(a_j + \epsilon)\} = S_0,$$

i.e., $K_S \models \psi(a_j + \epsilon)$ if and only if $j \in S_0$. Taking b to be $\epsilon \in K_S$, this completes the proof. \square

Because T has the independence property and clearly has the strict order property, the best classification-theoretic property we could hope for T to have is NTP_2 .

6.2 NTP_2 holds

First we make some elementary remarks about relative algebraic closures.

Lemma 6.3. *Let M be a pure field. Let K be a subfield of M which is nearly closed in M (in the sense of Definition 2.6). Let a and b be two tuples from M such that $a \perp_K^{ACF} b$, i.e., a and b are algebraically independent from each other over K . Then $K(a)$ is nearly closed in $K(a, b)$.*

Proof. Embed M into a monster model $\mathfrak{M} \models ACF$. By the remarks after Definition 2.6, $\text{tp}(a/K)$ and $\text{tp}(b/K)$ are stationary. Since $a \perp_K b$, the type of b over $\text{acl}(K(a))$ is K -definable. Now suppose that some singleton $c \in K(a, b)$ is algebraic over $K(a)$. Write c as $f(a, b)$, for some rational function $f(X, Y) \in K(X, Y)$. Note that $\text{stp}(b/K(a))$ includes the statement $f(a, x) = c$. On the other hand, it does not include $f(a, x) = c'$ for any conjugate $c' \neq c$ of c over $K(a)$. As $\text{stp}(b/K(a))$ is definable over K , ac and ac' cannot have the same type over K . But if c and c' are conjugate over $K(a)$, then ac and ac' have the same type. So c' does not exist, and c has no other conjugates over $K(a)$. Thus $c \in \text{dcl}(K(a))$. So $c^{p^k} \in K(a)$ for some k . As c was an arbitrary element of $K(a, b) \cap K(a)^{\text{alg}}$, we see that $K(a)$ is nearly closed in $K(a, b)$. \square

Lemma 6.4. *Let M be a pure field. Suppose $K_0 \subset K_1 \subset K_2$ are three subfields of M , each nearly closed in M . Let c be tuple from M , possibly infinite. Suppose that $c \downarrow_{K_0}^{ACF} K_2$, i.e., K_2 and c are algebraically independent over K_0 . Then $K_1(c)$ is nearly closed in $K_2(c)$.*

Proof. As in the previous lemma, embed M into a monster model \mathfrak{M} of ACF. Then $c \downarrow_{K_0} K_2$, and by properties of forking, $K_1(c) \downarrow_{K_1} K_2$. By the previous lemma, $K_1(c)$ is nearly closed in $K_2 K_1(c) = K_2(c)$. \square

Now we return to existentially closed fields with valuations and orderings. As always, T is the model companion.

Lemma 6.5. *In a monster model of T , let B be a small set of parameters and a_1, a_2, \dots be a B -indiscernible sequence. Suppose that $B = \text{acl}(B)$ and $a_i = \text{acl}(Ba_i)$ for any/every i . Suppose also that $a_j \downarrow_B^{ACF} a_{<j}$ for every j , i.e., the sequence is algebraically independent over B . Let c be a finite tuple and suppose that a_1, a_2, \dots is quantifier-free indiscernible over cB , i.e., if $i_1 < \dots < i_m$ and $j_1 < \dots < j_m$, then*

$$\text{qftp}(a_{i_1} a_{i_2} \cdots a_{i_m} / cB) = \text{qftp}(a_{j_1} a_{j_2} \cdots a_{j_m} / cB).$$

Let $\phi(x; y)$ be a formula over B such that $\phi(c; a_1)$ holds. Then $\bigwedge_{j=1}^{\infty} \phi(x; a_j)$ is consistent.

Proof. Because a_1, a_2, \dots is B -indiscernible, it suffices to show for each k that $\bigwedge_{j=1}^{\infty} \phi(x; a_j)$ is not k -inconsistent.

First observe that whether or not $c \downarrow_B^{ACF} a_j$ holds depends only on the quantifier-free type of c and a_j over B . In particular, it does not depend on j , by quantifier-free indiscernibility of a_1, a_2, \dots over cB . If $c \not\downarrow_B^{ACF} a_j$ for one j , then this holds for all j . As the a_j are an algebraically independent sequence over B , this contradicts the fact that finite tuples have finite preweight in ACF. So $c \downarrow_B^{ACF} a_j$ for each j . The same argument applied to the sequence $a_1 a_2, a_3 a_4, \dots$ shows that $c \downarrow_B^{ACF} a_1 a_2$. Similarly $c \downarrow_B^{ACF} a_1 a_2 a_3$, and so on, and so $c \downarrow_B^{ACF} a_1 a_2 a_3 \dots$.

Let M be the monster model. Any subset of M closed under $\text{acl}(-)$ is relatively algebraically closed in M , hence nearly closed in M . In particular, if we let $K_0 = B$, $K_1 = B(a_j) = a_j$, and $K_2 = \text{acl}(Ba_1 a_2 \dots)$, then each of K_0, K_1, K_2 is nearly closed in M , and $K_0 \subset K_1 \subset K_2$. By the previous paragraph, $c \downarrow_B^{ACF} K_2$, so by Lemma 6.4, we conclude that $K_1(c)$ is nearly

closed in $K_2(c)$, i.e., $B(a_j, c)$ is nearly closed in $K_2(c)$. Using bars to denote perfect closures, this means that $\overline{B(a_j, c)}$ is relatively algebraically closed in $\overline{K_2(c)}$.

Recall the function $P(-, -)$ from Theorem 5.1. By the “irrelevance” part of that theorem,

$$P(\phi(c; a_j); \overline{B(a_j, c)}) = P(\phi(c; a_j); \overline{K_2(c)}).$$

Now by quantifier-free indiscernibility of a_1, a_2, \dots over cB , we see that $B(a_j, c) \cong B(a_{j'}, c)$ for all j, j' . By the isomorphism-invariance part of Theorem 5.1,

$$P(\phi(c; a_j); \overline{B(a_j, c)}) = P(\phi(c; a_{j'}); \overline{B(a_{j'}, c)})$$

for all j, j' . Consequently, $P(\phi(c; a_j); \overline{K_2(c)})$ does not depend on j .

Now M is a model of T extending $\overline{K_2(c)}$, and in M , $\phi(c; a_1)$ holds. So by the “density” part of Theorem 5.1, $P(\phi(c; a_1); \overline{K_2(c)})$ is some positive number $\epsilon > 0$. Consequently, $P(\phi(c; a_j); \overline{K_2(c)}) = \epsilon > 0$ for every j .

Suppose for the sake of contradiction that $\bigwedge_{j=1}^{\infty} \phi(x; a_j)$ is k -inconsistent for some k . Let N be big enough that $N\epsilon > k$. Let $\psi(x)$ be the statement over K_2 asserting that at least k of $\phi(x; a_1), \dots, \phi(x; a_N)$ hold. By the Keisler measure part of Theorem 5.1, $P(\psi(c); \overline{K_2(c)}) > 0$, and there is a model M' of T extending $\overline{K_2(c)}$ in which ψ holds. In particular, $M' \models \exists x : \psi(x)$. But K_2 is relatively algebraically closed in M , hence satisfies axiom A1 of §3.1 by Corollary 3.10. By Corollary 3.11, the statement $\exists x : \psi(x)$ holds in M if and only if it holds in M' . Consequently, it holds in M , and therefore $\bigwedge_{j=1}^n \phi(x; a_j)$ is not k -inconsistent. \square

Recall from [1] or [2] that the burden of a partial type $p(x)$ is the supremum of κ such that there is an inp-pattern in $p(x)$ of depth κ , that is, an array of formula $\phi_i(x; a_{ij})$ for $i < \kappa$ and $j < \omega$, and some $k_i < \omega$ such that the i th row $\bigwedge_{j < \omega} \phi_i(x; a_{ij})$ is k_i -inconsistent for each i , and such that for any $\eta : \kappa \rightarrow \omega$, the corresponding downwards path $\bigwedge_{i < \kappa} \phi_i(x; a_{i, \eta(i)})$ is consistent with $p(x)$. A theory is NTP_2 if every partial type has burden less than ∞ . A theory is strong if every partial type has burden less than \aleph_0 , roughly. (See [1] for a more precise statement.) At any rate, if every partial type has burden less than \aleph_0 , then the theory is strong. By the submultiplicativity of burden (Theorem 11 in [2]), it suffices to check the burden of the home sort.

Fact 6.6. *If D and E are definable sets, $\text{bdn}(D \times E) \geq \text{bdn}(D) + \text{bdn}(E)$.*

In fact, if $\phi_i(x; a_i)$ is an inp-pattern for D and $\psi_j(y; b_j)$ is an inp-pattern for E , then $\{\phi_i(x; a_i)\} \cup \{\psi_j(y; b_j)\}$ is an inp-pattern for $D \times E$.

In NIP theories, burden is the same thing as dp-rank, which is known to be additive [9]. The theories ACVF, p CF, and RCF are all known to be dp-minimal, i.e., to have dp-rank 1 [5]. One of the descriptions of dp-rank is that a partial type $\Sigma(x)$ over a set C has dp-rank $\geq \kappa$ if and only if there are κ -many mutually indiscernible sequences over C and a realization a of $\Sigma(x)$ such that each sequence is not indiscernible over Ca .

Theorem 6.7. *The model companion T is NTP_2 , and strong. In fact, the burden of affine m -space is exactly mn , where n is the number of valuations and orderings.*

Proof. To show that the burden of \mathbb{A}^m is at least mn , it suffices by Fact 6.6 to show that $\text{bdn}(\mathbb{A}^1) \geq n$. In the case where every T_i is ACVF, one can take $\phi_i(x; y)$ to assert that $v_i(x) = y$, for $1 \leq i \leq n$, and take $a_{i,0}, a_{i,1}, \dots$ to be an increasing sequence in the i th valuation group. Variations on this handle the remaining cases. We leave the details as an exercise to the reader.

For the upper bound, suppose for the sake of contradiction that there is an inp-pattern $\{\phi_i(x; a_{ij})\}_{i < mn+1; 0 \leq j < \omega}$ of depth $mn + 1$, with x a tuple of length m . We may assume that the a_{ij} form a mutually \emptyset -indiscernible array. Extend the array to the left, i.e., let j range over negative numbers. Let B be $\text{acl}(a_{ij} : j < 0)$. From stability theory, one knows that $a_{ij} \downarrow_B^{ACF} a_{i0}a_{i1} \cdots a_{i,j-1}$ for every j . By mutual indiscernibility, each sequence a_{i0}, a_{i1}, \dots is indiscernible over $\{a_{ij} : j < 0\}$, hence over B . In particular, $a_{ij} \equiv_B a_{ij'}$ for $j \neq j'$. For each $i < mn + 1$, let $b_{i,0}$ be an enumeration of $\text{acl}(Ba_{i,0})$. For $j > 0$, choose $b_{i,j}$ such that $a_{i,j}b_{i,j} \equiv_B a_{i,0}b_{i,0}$. Then $b_{i,j}$ is an enumeration of $\text{acl}(Ba_{i,j})$ for every i and every $j \geq 0$. Let $c_{i,j}d_{i,j}$ be a mutually B -indiscernible array extracted from the array $a_{i,j}b_{i,j}$. Then $c_{i,j}d_{i,j} \equiv_B a_{i,0}b_{i,0}$, so $d_{i,j}$ is an enumeration of $\text{acl}(Bc_{i,j})$. Also, because $a_{i,0}, a_{i,1}, \dots$ was already B -indiscernible, we must have

$$c_{i,0}c_{i,1} \cdots \equiv_B a_{i,0}a_{i,1} \cdots$$

for each i . Consequently, $c_{i,j} \downarrow_B^{ACF} c_{i,0} \cdots c_{i,j-1}$. And since $d_{i,j} \subset \text{acl}(Bd_{i,j})$, we also have

$$d_{i,j} \downarrow_B^{ACF} d_{i,0}d_{i,1} \cdots d_{i,j-1},$$

using Corollary 3.12. As $b_{i,0}$ is an enumeration of $\text{acl}(Ba_{i,0})$, the elements of $a_{i,0}$ must actually appear somewhere in $b_{i,0}$. Let π_i be the coordinate projection such that $\pi_i(b_{i,0}) = a_{i,0}$. Hence $c_{i,j} = \pi_i(d_{i,j})$.

Because the a_{ij} formed a mutually \emptyset -indiscernible array, the collective type of all the $c_{i,j}$'s must agree with that of all the $a_{i,j}$'s. Hence $\phi_i(x; c_{i,j})$ is still an inp-pattern of depth $mn + 1$. Let $\psi_i(x; y)$ be $\phi_i(x; \pi_i(y))$. Then $\psi_i(x; d_{i,j})$ is an inp-pattern of depth $mn + 1$. Let c be a realization of $\bigwedge_{i < mn+1} \psi_i(x; d_{i,0})$. Note that c is a tuple of length m .

Let \mathfrak{M} be the ambient monster. For each $1 \leq k \leq n$, let \mathfrak{M}_k be a model of T_k extending $\mathfrak{M} \upharpoonright \mathcal{L}_k$. By quantifier-elimination, the array $\{d_{i,j}\}$ is still mutually B -indiscernible in \mathfrak{M}_i . By additivity of dp-rank and by dp-minimality of the home sort in \mathfrak{M}_k , we know that the dp-rank of $\text{tp}(c/B)$ in \mathfrak{M}_k is at most m . In particular, for each $1 \leq k \leq n$, at most m of the rows in the array $\{d_{i,j}\}$ can fail to be Bc -indiscernible in \mathfrak{M}_k . By the pigeonhole principle, there must be some value of i such that the sequence $d_{i,0}, d_{i,1}, \dots$ is Bc -indiscernible in each of $\mathfrak{M}_1, \mathfrak{M}_2, \dots, \mathfrak{M}_n$. Back in \mathfrak{M} , this means that $d_{i,0}, d_{i,1}, \dots$ is quantifier-free Bc -indiscernible. Since $d_{i,0}, d_{i,1}, \dots$ is also B -indiscernible and B -independent, Lemma 6.5 applies. Consequently, $\bigwedge_{j=0}^{\infty} \psi_i(x; d_{i,j})$ is consistent, because $\psi_i(c; d_{i,0})$ holds. This contradicts the fact that $\{\psi_i(x; d_{i,j})\}$ is an inp-pattern. \square

7 Forking and Dividing

We'll make use of the following general fact:

Fact 7.1. *Let \mathfrak{M} be a monster model of some theory, let $S \subset \mathfrak{M}$ be a small set, and let $\phi(x)$ be a formula with parameters from \mathfrak{M} . Suppose there is a global Keisler measure μ which is Lascar-invariant over S , and suppose $\mu(\phi(x)) > 0$. Then $\phi(x)$ does not fork over S .*

This is probably well-known and probably in [8] somewhere, but here is a proof:

Proof. If $\phi(x)$ forks over S , then $\phi(x)$ implies a disjunction $\psi_1(x) \vee \dots \vee \psi_\ell(x)$ of formulas which divide over S . For some i , we must have $\mu(\psi_i(x)) > 0$. Fix such an i , and let $\epsilon = \mu(\psi_i(x)) > 0$. Write $\psi_i(x)$ as $\chi(x; a)$, where a is the tuple of parameters over which $\psi_i(x)$ is defined. By definition of dividing, there is an S -indiscernible sequence a_1, a_2, \dots with $a_1 = a$, and with $\bigwedge_{j=1}^{\infty} \chi(x; a_j)$ being k -inconsistent for some k . By Lascar-invariance, $\mu(\chi(x; a_j)) = \epsilon$ for every j . Choose N so large that $N\epsilon > k$. Let $\pi(x)$ assert that at least k of the formulas $\chi(x; a_1), \dots, \chi(x; a_N)$ hold. By definition of k -dividing, $\pi(x)$ must

be inconsistent, but by facts from basic probability, $\mu(\pi(x)) > 0$, yielding a contradiction. Specifically, if X_i is the random indicator variable that equals 1 if $\chi(x; a_i)$ holds and 0 otherwise, then the expected value of the random variable X_i is ϵ , and the expected value of the random variable $\sum_{i=1}^N X_i$ is $N\epsilon > k$. So there must be a non-zero probability that $X_i > k$, i.e., that $\pi(x)$ holds. \square

Now we specialize to the theory T under consideration.

Lemma 7.2. *Let \mathfrak{M} be a monster model of T . Let S be a small subset of \mathfrak{M} , and let p be a complete quantifier-free type over \mathfrak{M} which is Lascar-invariant over S . Then there is a Keisler measure μ on $S(\mathfrak{M})$, Lascar-invariant over S , whose support is exactly the set of completions of p .*

This is really just a restatement or special case of Theorem 5.1.

Proof. Let a be a realization of p in some bigger model, and consider the structure $\mathfrak{M}[a]$ generated by \mathfrak{M} and a . The structure of $\mathfrak{M}[a]$ is determined by p . Also, if σ is any Lascar strong automorphism of \mathfrak{M} over S , then $p = \sigma(p)$. This implies that there is a uniquely determined automorphism σ' of $\mathfrak{M}[a]$ extending σ on \mathfrak{M} and fixing a .

Let $\overline{\mathfrak{M}[a]}$ denote the perfect closure of the field of fractions of $\mathfrak{M}[a]$. This is uniquely determined (as a model of T_\forall) by $\mathfrak{M}[a]$, and hence is determined by p . Let μ be the Keisler measure on \mathfrak{M} which assigns to an \mathfrak{M} -formula $\phi(x; b)$ the value

$$P(\phi(a; b); \overline{\mathfrak{M}[a]}),$$

where P is as in Theorem 5.1. By the Keisler measure part of Theorem 5.1, this is a Keisler measure on the space of completions of $\text{qftp}(\mathfrak{M}[a])$. By model completeness, any extension of $\text{qftp}(\mathfrak{M}[a])$ to a complete type must satisfy $\text{tp}(\mathfrak{M})$, so we have a legitimate Keisler measure on the space of extensions of p to complete types over \mathfrak{M} . And if σ is any Lascar strong automorphism over S , then by the ‘‘isomorphism invariance’’ part of Theorem 5.1,

$$P(\phi(a; \sigma(b)); \overline{\mathfrak{M}[a]}) = P(\phi(\sigma'(a); \sigma'(b)); \overline{\mathfrak{M}[a]}) = P(\phi(a, b); \overline{\mathfrak{M}[a]})$$

where σ' is the aforementioned automorphism of $\mathfrak{M}[a]$ extending σ and fixing a . Thus $\mu(\phi(x; b)) = \mu(\phi(x; \sigma(b)))$. We conclude that $\mu(\phi(x; b)) = \mu(\phi(x; b'))$ for any formula $\phi(x; y)$ and any $b, b' \in \mathfrak{M}$ having the same Lascar strong type over S . Finally, if b is a tuple from \mathfrak{M} and $\phi(a; b)$ is a formula which is consistent with p , then $\phi(a; b)$ is also consistent with the diagram of $\mathfrak{M}[a]$, hence has positive probability by the ‘‘density’’ part of Theorem 5.1. \square

Corollary 7.3. *Let \mathfrak{M} be a monster model of T and S be a small subset of \mathfrak{M} . Suppose q is a complete quantifier-free type on \mathfrak{M} which is Lascar invariant over S . Then every complete type on \mathfrak{M} extending q does not fork over S .*

Proof. Let $p(x)$ be a complete type extending $q(x)$. Let $\phi(x)$ be any formula from $p(x)$. Let μ be the Keisler measure from Lemma 7.2. Then μ is Lascar invariant over S , and $\mu(\phi(x))$ is positive because $\phi(x)$ is consistent with $q(x)$. By Fact 7.1, $\phi(x)$ does not fork over S . \square

If M is a model of T and A, B, C are subsets of M , let $A \downarrow_C^{T_i} B$ indicate that $A \downarrow_C B$ holds in any/every model of T_i extending $M \upharpoonright \mathcal{L}_i$.

Lemma 7.4. *Work in a monster model \mathfrak{M} of T . Let a be a finite tuple, and B and C be sets (in the home sort, as always). Suppose $C = \text{acl}(C)$. Suppose $a \downarrow_C^{T_i} B$ holds for every $1 \leq i \leq n$. Then $\text{qftp}(a/BC)$ can be extended to a quantifier-free type $q(x)$ on \mathfrak{M} which is Lascar invariant over C .*

Proof. Let V be the variety over C of which a is a generic point. By Fact 2.8, V is absolutely irreducible.

Let \mathfrak{M}_i be a model of T_i extending $\mathfrak{M} \upharpoonright \mathcal{L}_i$. Within \mathfrak{M}_i , $a \downarrow_C^{T_i} B$. By Adler's characterization of forking in NIP theories (Proposition 2.1 in [8]), there is an \mathcal{L}_i -type $p_i(x)$ on \mathfrak{M}_i which extends the type of a over BC and which is Lascar-invariant over C . The restriction of this \mathcal{L}_i -type to a quantifier-free \mathcal{L}_{rings} -type must say that x lives on V and on no \mathfrak{M}_i -definable proper subvarieties of V . This follows from Lemma 2.16. Let $q_i(x)$ be the set of quantifier-free \mathcal{L}_i -statements in $p_i(x)$ with parameters from \mathfrak{M} . Then $q_i(x)$ is a complete quantifier-free \mathcal{L}_i -type on \mathfrak{M} . Let $q(x)$ be $\bigcup_{i=1}^n q_i(x)$. This is a complete quantifier-free type on \mathfrak{M} ; it is consistent because the $q_i(x)$ all have the same restriction to the language of rings, namely, the generic type of V . Also, $q(x)$ extends $\text{qftp}(a/BC)$, because the \mathcal{L}_i -part of $\text{qftp}(a/BC)$ is present in $p_i(x)$ and $q_i(x)$.

To show Lascar-invariance of $q(x)$ over C , it suffices to show that if I is a C -indiscernible sequence in \mathfrak{M} , a and a' are two elements of I , and $\phi(x; y)$ is a quantifier-free formula, then $\phi(x; a) \in q(x)$ if and only if $\phi(x; a') \in q(x)$. In fact, we only need to consider the case where $\phi(x)$ is a quantifier-free \mathcal{L}_i -formula, for some i . But then

$$\phi(x; a) \in q(x) \iff \phi(x; a) \in p_i(x) \iff \phi(x; a') \in p_i(x) \iff \phi(x; a') \in q(x)$$

where the middle equivalence follows from the fact that $p_i(x)$ is Lascar-invariant, and I is C -indiscernible within \mathfrak{M}_i (by quantifier-elimination in T_i). Thus $q(x)$ is Lascar-invariant over C , as claimed. \square

Theorem 7.5. *Forking and dividing agree over every set (in the home sort).*

Proof. First we show that if a is a finite tuple and B is a set, then $\text{qftp}(a/B)$ does not fork over B . By Lemma 7.4, there is a global quantifier-free type $q(x)$ which is Lascar-invariant over B . By Corollary 7.3, any extension of $q(x)$ to a complete global type does not fork over B . So $\text{qftp}(a/B)$ has a global non-forking extension. Now if a is any small tuple, and B is a set, then $\text{qftp}(a/B)$ does not fork over B , by compactness. Consequently, if a is a small tuple and B is a (small) set, then $\text{qftp}(a'/B)$ does not fork over B , where a' enumerates $\text{acl}(aB)$. By Corollary 3.11, $\text{qftp}(a'/B)$ implies $\text{tp}(a'/B)$, so $\text{tp}(a'/B)$ does not fork over B . By monotonicity, $\text{tp}(a/B)$ does not fork over B . As a and B are arbitrary, every set in the home sort is an extension base for forking in the sense of [3], so by Theorem 1.2 in [3], forking and dividing agree over every set in the home sort. \square

Lemma 7.6. *Let \mathfrak{M} be a monster model of T and $C = \text{acl}(C)$ be a small subset of M . Suppose $p(x)$ is a complete type on C and $q(x)$ is a complete quantifier-free type on M , with $q(x)$ extending the quantifier-free part of $p(x)$. Suppose $q(x)$ is Lascar-invariant over C . Then $q(x) \cup p(x)$ is consistent.*

Proof. Let $\mathfrak{M}[a]$ be the structure obtained by adjoining a realization a of $q(x)$ to \mathfrak{M} . Let W be the variety over \mathfrak{M} of which a is the generic point. By Fact 2.8, W is absolutely irreducible. Moreover, the ACF-theoretic code $[W]$ for W must lie in \mathfrak{M} . By Lascar invariance of $q(x)$, one sees that W is Lascar invariant over C . Consequently, the finite tuple $[W]$ is fixed by every Lascar strong automorphism over C . So $[W] \subset \text{acl}(C) = C$. Consequently, in an ambient model of ACF we have $Cb(\text{stp}(a/\mathfrak{M})) \subset C$, and so $a \downarrow_C^{ACF} \mathfrak{M}$. By Lemma 6.3, $C(a)$ is relatively algebraically closed in $\mathfrak{M}(a)$.

Because the quantifier-free type of a over C is consistent with $p(x)$, there is a model $N \models T$ extending $C[a]$ such that within N , $\text{tp}(a/C) = p(x)$. By Lemma 3.9, we can amalgamate $\mathfrak{M}(a)$ and N over $C(a)$. So there is a model N' of T extending N and $\mathfrak{M}(a)$. In N' , $\text{tp}(a/C) = p(x)$. As $N \preceq N'$, $\text{tp}(a/C) = p(x)$ holds in N' as well. And as $N' \succeq \mathfrak{M}(a)$, $\text{qftp}(a/\mathfrak{M}) = q(x)$. So $q(x) \cup p(x)$ is consistent. \square

Lemma 7.7. *Work in a monster model \mathfrak{M} of T . Let a be a finite tuple, and B and C be sets (in the home sort, as always). Suppose $a \downarrow_C^{T_i} B$ holds for every $1 \leq i \leq n$. Then $a \downarrow_C B$.*

Proof. A type forks/divides over C if and only if it forks/divides over $\text{acl}(C)$, so it suffices to show that $\text{tp}(a/BC)$ does not fork over $\text{acl}(C)$. By monotonicity, it suffices to show that $\text{tp}(a/\text{acl}(BC))$ does not fork over $\text{acl}(C)$. By Claim 3.6 in [3] and Lemma 2.14 above, $a \downarrow_{\text{acl}(C)}^{T_i} \text{acl}(CB)$ for every i . So we may assume that $C = \text{acl}(C) \subset B = \text{acl}(B)$.

Now by Lemma 7.4, there is a global quantifier-free type $q(x)$ extending $\text{qftp}(a/BC) = \text{qftp}(a/B)$, with $q(x)$ Lascar-invariant over C . Clearly $q(x)$ is also Lascar-invariant over B , so by Lemma 7.6, $q(x)$ is consistent with $\text{tp}(a/B)$. Let $p(x)$ be a global complete type extending $q(x) \cup \text{tp}(a/B)$. Then $p(x)$ doesn't fork over C by Corollary 7.3. \square

Let $\text{qftp}^i(a/B)$ denote the quantifier-free \mathcal{L}_i -type of a over B , and let $\text{qftp}^{ACF}(a/B)$ denote the field-theoretic quantifier-free type of a over B .

Lemma 7.8. *Let \mathfrak{M} be a monster model of T , and let $C = \text{acl}(C)$ be a small subset. For each i , let M_i be a model of T_i extending $\mathfrak{M} \upharpoonright \mathcal{L}_i$. For each i , let a_i be a tuple in M_i . Suppose that $\text{qftp}^{ACF}(a_i/C)$ doesn't depend on i . Then we can find a tuple a in \mathfrak{M} such that $\text{qftp}^i(a/C) = \text{qftp}^i(a_i/C)$ for every i .*

Proof. Let $C[a_i]$ denote the subring or subfield of M_i generated by C and a_i . By assumption, $C[a_i]$ is isomorphic to $C[a_{i'}]$ as a ring, for every i and i' . Use these isomorphisms to identify all the $C[a_i]$ with each other, getting a single ring $C[a]$ which is isomorphic to $C[a_i]$ for every i . Use these isomorphisms to move the $(T_i)_\forall$ structure from $C[a_i]$ to $C[a]$. Now $C[a]$ is a model of T_\forall , and $\text{qftp}^i(a/C) = \text{qftp}^i(a_i/C)$, for every i . As $C = \text{acl}(C)$, C is nearly closed in \mathfrak{M} , so by Lemma 3.9, one can embed $C[a]$ and \mathfrak{M} into a bigger model of T . By model completeness and saturation, $\text{tp}(a/C)$ is already realized in \mathfrak{M} . \square

Lemma 7.9. *Let a, B, C be small subsets of a monster model $\mathfrak{M} \models T$. Suppose $a \not\downarrow_C^{T_1} b$. Then $a \not\downarrow_C b$.*

Proof. By Claim 3.6 in [3] applied to both T_1 and T , we may assume $C = \text{acl}(C)$ and $B = \text{acl}(BC)$. By finite character of forking, we may assume a is finite. For every i , let \mathfrak{M}_i be an even more monstrous model of T_i extending

$\mathfrak{M} \upharpoonright \mathcal{L}_i$. As $a \not\downarrow_C^{T_1} B$, within \mathfrak{M}_1 we have $a \not\downarrow_C B$. By Lemma 2.14, some \mathcal{L}_1 -formula $\phi(x; B)$ in $\text{tp}(a/BC)$ divides over C . By quantifier-elimination in T_i , we may assume that $\phi(x; y)$ is a quantifier-free \mathcal{L}_1 -formula. By Lemma 2.17, there is a sequence $B = B_0^1, B_1^1, B_2^1, \dots$ in \mathfrak{M}_1 which is indiscernible over C and algebraically independent over C , and such that $\bigwedge_{j=0}^{\infty} \phi(x; B_j^1)$ is k -inconsistent in \mathfrak{M}_1 , for some k . Thus $\text{qftp}^1(B_j^1/C) = \text{qftp}^1(B/C)$, and in a certain sense

$$\text{qftp}^{ACF}(B_0^1 B_1^1 B_2^1 \dots / C) = \text{qftp}^{ACF}(B/C) \otimes \text{qftp}^{ACF}(B/C) \otimes \dots$$

The right hand side makes sense because C is nearly closed in B (Definition 2.6), so $\text{qftp}^{ACF}(B/C)$ is stationary.

Meanwhile, for $i > 1$, we can apply Lemma 2.15 to \mathfrak{M}_i and $\text{tp}(B/C)$, getting a sequence $B = B_0^i, B_1^i, B_2^i, \dots$ which is indiscernible over C and algebraically independent over C . (Note that Lemma 2.15 is true even without the restriction that B be finite.) So again, we get $\text{qftp}^i(B_j^i/C) = \text{qftp}^i(B/C)$, and

$$\text{qftp}^{ACF}(B_0^i B_1^i B_2^i \dots / C) = \text{qftp}^{ACF}(B/C) \otimes \text{qftp}^{ACF}(B/C) \otimes \dots$$

In particular, $\text{qftp}^{ACF}(B_0^i B_1^i B_2^i \dots / C)$ doesn't depend on i , as i ranges from 1 to n . By Lemma 7.8, we can therefore find a sequence B_0, B_1, \dots in \mathfrak{M} such that

$$\text{qftp}^i(B_0 B_1 \dots / C) = \text{qftp}^i(B_0^i B_1^i B_2^i \dots / C)$$

for every i . In particular, $\text{qftp}^i(B_j/C) = \text{qftp}^i(B_j^i/C) = \text{qftp}^i(B/C)$. Because this holds for all i , $\text{qftp}(B_j/C) = \text{qftp}(B/C)$. Because $B = \text{acl}(B)$, $\text{qftp}(B/C) \vdash \text{tp}(B/C)$ by Corollary 3.11. So $\text{tp}(B_j/C) = \text{tp}(B/C)$ for every j . Also,

$$\text{qftp}^1(B_0 B_1 \dots / C) = \text{qftp}^1(B_0^1 B_1^1 \dots / C)$$

implies that there is an automorphism σ of \mathfrak{M}_1 sending $B_0^1 B_1^1 \dots$ to $B_0 B_1 \dots$. Consequently, $\bigwedge_{j=0}^{\infty} \phi(x; B_j)$ is k -inconsistent in \mathfrak{M}_1 . Clearly it is also k -inconsistent in \mathfrak{M} , because \mathfrak{M} is smaller than \mathfrak{M}_1 . Since B_0, B_1 is a sequence of realizations of $\text{tp}(B/C)$, we conclude that $\phi(x; B)$ divides over C , in \mathfrak{M} . \square

Theorem 7.10. *Let M be a model of T , and let A, B, C be subsets of M (in the home sort). The following are equivalent:*

- $A \downarrow_C B$, i.e., the type of A over BC doesn't fork over C .

- The type of A over BC doesn't divide over C .
- $A \downarrow_C^{T_i} B$ for every $1 \leq i \leq n$.

Proof. The first two bullet points are equivalent by Theorem 7.5. If $A \downarrow_C B$, then by Lemma 7.9 $A \downarrow_C^{T_1} B$. Similarly, $A \downarrow_C^{T_i} B$ for every $1 \leq i \leq n$. Conversely, if $A \downarrow_C^{T_i} B$ for every $1 \leq i \leq n$, then by Lemma 7.7, $a \downarrow_C B$ for every finite subset $a \subset A$. By finite character of forking, $A \downarrow_C B$. \square

References

- [1] Hans Adler. Strong theories, burden, and weight. 2007.
- [2] Artem Chernikov. Theories without the tree property of the second kind. 2012.
- [3] Artem Chernikov and Itay Kaplan. Forking and dividing in ntp2 theories. 2013.
- [4] Artem Chernikov, Itay Kaplan, and Pierre Simon. Groups and fields with ntp2. 2013.
- [5] Alfred Dolich, John Goodrick, and David Lippel. Dp-minimality: Basic facts and examples. 2009.
- [6] Antoine Ducros. Espaces de berkovich, polytopes, squelettes et théorie des modèles. 2012.
- [7] Yuri L. Ershov. *Multi-valued Fields*. 2009.
- [8] Ehud Hrushovski and Anand Pillay. On nip and invariant measures. 2009.
- [9] Itay Kaplan, Alf Onshuus, and Alexander Usvyatsov. Additivity of the dp-rank. 2013.
- [10] Alexander Prestel and Martin Ziegler. Model theoretic methods in the theory of topological fields. *Journal für die reine und angewandte Mathematik*, pages 318 – 341, 1978.

- [11] Laurentius Petrus Dignus van den Dries. *Model Theory of Fields: Decidability and Bounds for Polynomial Ideals*. PhD thesis, 1978.
- [12] Itai Ben Yaacov and Artem Chernikov. An independence theorem for ntp2 theories. 2012.