World Scientific
www.worldscientific.com

RANDOMIZATIONS OF MODELS
AS METRIC STRUCTURES

ITAÏ BEN YAACOV* and H. JEROME KEISLER[†]

*Université de Lyon, CNRS, Université Claude Bernard Lyon 1,
Institut Camille Jordan, 43 boulevard du 11 novembre 1918,
69622 Villeurbanne Cedex, France

[†]Department of Mathematics, University of Wisconsin,
480 Lincoln Drive, Madison WI 53706, USA
[†]keisler@math.wisc.edu

The notion of a randomization of a first-order structure was introduced by Keisler in the paper Randomizing a Model, Advances in Math. 1999. The idea was to form a new structure whose elements are random elements of the original first-order structure. In this paper we treat randomizations as continuous structures in the sense of Ben Yaacov and Usvyatsov. In this setting, the earlier results show that the randomization of a complete first-order theory is a complete theory in continuous logic that admits elimination of quantifiers and has a natural set of axioms. We show that the randomization operation preserves the properties of being omega-categorical, omega-stable, and stable.

Keywords: Randomization; metric structures; continuous logic; stable theories.

AMS Subject Classification: 03C45, 03C90, 03B50, 03B48

1. Introduction

In this paper, we study randomizations of first-order structures in the setting of continuous model theory. Intuitively, a randomization of a first-order structure $\mathcal{M}$ is a new structure whose elements are random elements of $\mathcal{M}$. In probability theory, one often starts with some structure $\mathcal{M}$ and studies the properties of random elements of $\mathcal{M}$. In many cases, the random elements of $\mathcal{M}$ have properties analogous to those of the original elements of $\mathcal{M}$. With this idea in mind, the paper by Keisler [7] introduced the notion of a randomization of a first-order theory T as a new many-sorted first-order theory. That approach pre-dated the current development of continuous structures in the paper by Ben Yaacov and Usvyatsov [5].

Here, we formally define a randomization of a first-order structure $\mathcal{M}$ as a continuous structure in the sense of [5]. This seems to be a more natural setting for the concept. In this setting, the results of [7] show that if T is the complete theory

of $\mathcal{M}$, the theory T^R of randomizations of $\mathcal{M}$ is a complete theory in continuous logic which admits elimination of quantifiers and has a natural set of axioms.

One would expect that the original first-order theory T and the randomization theory T^R will have similar model-theoretic properties. We show that this is indeed the case for the properties of ω-categoricity, ω-stability, and stability. This provides us with a ready supply of new examples of continuous theories with these properties.

In Sec. 2, we define the randomization theory T^R as a theory in continuous logic, and restate the results we need from [7] in this setting. In Sec. 3, we begin with a proof that a first-order theory T is ω-categorical if and only if T^R is ω-categorical, and then we investigate separable structures.

Section 4 concerns ω-stable theories. In Sec. 4.1, we prove that a complete theory T is ω-stable if and only if T^R is ω-stable. In Sec. 4.2, we extend this result to the case where T has countably many complete extensions.

Section 5 is about stable theories and independence. Section 5.1 contains abstract results on fiber products of measures that will be used later. In Sec. 5.2, we develop some properties of stable formulas in continuous theories. In Sec. 5.3, we prove that a first-order theory T is stable if and only if T^R is stable. We also give a characterization of independent types in T^R.

In [2] another result of this type was proved — that T^R is dependent (does not have the independence property) as a continuous theory if and only if T is dependent as a first-order theory. In this paper we deal only with randomizations of first-order structures, but it should be mentioned that randomizations of continuous structures have recently been developed in [2] and [3].

2. Randomizations

In this section we will restate some notions and results from [7] in the context of continuous structures.

We will assume that the reader is familiar with continuous model theory as it is developed in [5] and [1], including the notions of a structure, pre-structure, signature, theory, and model of a theory. A *pre-model* of a theory is a pre-structure which satisfies each statement in the theory. For both first order and continuous logic, we will abuse notation by treating elements of a structure (or pre-structure) $\mathcal{M}$ as constant symbols outside the signature of $\mathcal{M}$. These constant symbols will be called *parameters* from $\mathcal{M}$.

We will assume throughout this paper that T is a consistent first-order theory with signature L such that each model of T has at least two elements.

The *randomization signature* for L is a two-sorted continuous signature L^R with a sort $\mathbf{K}$ of random elements, and a sort $\mathbf{B}$ of events. L^R has an n-ary function symbol $[\![\varphi(\cdot)]\!]$ of sort $\mathbf{K}^n \to \mathbf{B}$ for each first-order formula φ of L with n free variables, a $[0,1]$-valued unary predicate symbol μ of sort $\mathbf{B}$ for probability, the Boolean operations $\top, \bot, \sqcup, \sqcap, \neg$ of sort $\mathbf{B}$, and distance predicates $d_{\mathbf{K}}$ and $d_{\mathbf{B}}$ for sorts $\mathbf{K}, \mathbf{B}$. All these symbols are 1-Lipschitz with respect to each argument.

We will use $U, V, \ldots$ to denote continuous variables of sort $\mathbf{B}$. We will use $x, y, \ldots$ to denote either first-order variables or continuous variables of sort $\mathbf{K}$, depending on the context. Given a first-order formula φ with n free variables, $\varphi(\overline{x})$ will denote the first-order formula formed by replacing the free variables in φ by first-order variables $\overline{x}$, and $[\![\varphi(\overline{x})]\!]$ will denote the atomic term in L^R formed by filling the argument places of the function symbol $[\![\varphi(\cdot)]\!]$ with continuous variables $\overline{x}$ of sort $\mathbf{K}$.

When working with a structure or pre-structure $(\mathcal{K}, \mathcal{B})$ for L^R, $\mathsf{A}, \mathsf{B}, \mathsf{C}, \ldots$ will denote elements or parameters from $\mathcal{B}$, and $\mathbf{f}, \mathbf{g}, \ldots$ will denote elements or parameters from $\mathcal{K}$. Variables can be replaced by parameters of the same sort. For example, if $\overline{\mathbf{f}}$ is a tuple of elements of $\mathcal{K}$, $[\![\varphi(\overline{\mathbf{f}})]\!]$ will be a constant term of $L^R \cup \overline{\mathbf{f}}$ whose interpretation is an element of $\mathcal{B}$.

We next define the notion of a randomization $(\mathcal{K}, \mathcal{B})$ of $\mathcal{M}$. Informally, the elements $\mathsf{B} \in \mathcal{B}$ are events, that is, measurable subsets of some probability space Ω, and the elements $\mathbf{f} \in \mathcal{K}$ are random elements of $\mathcal{M}$, that is, measurable functions from Ω into $\mathcal{M}$. By "measure" we will always mean "σ-additive measure", unless we explicitly qualify it as "finitely additive measure".

Here is the formal definition. Given a model $\mathcal{M}$ of T, a *randomization of* $\mathcal{M}$ is a pre-structure $(\mathcal{K}, \mathcal{B})$ for L^R equipped with a *finitely additive* measure μ such that:

- $(\mathcal{B}, \mu)$ comes from an atomless finitely additive probability space $(\Omega, \mathcal{B}, \mu)$.
- $\mathcal{K}$ is a set of functions $\mathbf{f} \colon \Omega \to \mathcal{M}$, i.e. $\mathcal{K} \subseteq \mathcal{M}^{\Omega}$.
- For each formula $\psi(\overline{x})$ of L and tuple $\overline{\mathbf{f}}$ in $\mathcal{K}$, we have

$$[\![\psi(\overline{\mathbf{f}})]\!] = \{w \in \Omega : \mathcal{M} \models \psi(\overline{\mathbf{f}}(w))\}.$$

  (It follows that the right side belongs to the set of events $\mathcal{B}$.)
- For each $\mathsf{B} \in \mathcal{B}$ and real $\varepsilon > 0$ there are $\mathbf{f}, \mathbf{g} \in \mathcal{K}$ such that $\mu(\mathsf{B} \triangle [\![\mathbf{f} = \mathbf{g}]\!]) < \varepsilon$, where $\triangle$ is the Boolean symmetric difference operation.
- For each formula $\theta(x, \overline{y})$ of L, real $\varepsilon > 0$, and tuple $\overline{\mathbf{g}}$ in $\mathcal{K}$, there exists $\mathbf{f} \in \mathcal{K}$ such that

$$\mu([\![\theta(\mathbf{f}, \overline{\mathbf{g}})]\!] \triangle [\![(\exists x\, \theta)(\overline{\mathbf{g}})]\!]) < \varepsilon.$$

- On $\mathcal{K}$, the distance predicate $d_{\mathbf{K}}$ defines the pseudo-metric

$$d_{\mathbf{K}}(\mathbf{f}, \mathbf{g}) = \mu[\![\mathbf{f} \neq \mathbf{g}]\!].$$

- On $\mathcal{B}$, the distance predicate $d_{\mathbf{B}}$ defines the pseudo-metric

$$d_{\mathbf{B}}(\mathsf{B}, \mathsf{C}) = \mu(\mathsf{B} \triangle \mathsf{C}).$$

Note that the finitely additive measure μ is determined by the pre-structure $(\mathcal{K}, \mathcal{B})$ via the distance predicates.

A randomization $(\mathcal{K}, \mathcal{B})$ of $\mathcal{M}$ is *full* if in addition

- $\mathcal{B}$ is equal to the set of all events $[\![\psi(\overline{\mathbf{f}})]\!]$ where $\psi(\overline{x})$ is a formula of L and $\overline{\mathbf{f}}$ is a tuple in $\mathcal{K}$.

- $\mathcal{K}$ is *full in* $\mathcal{M}^{\Omega}$, that is, for each formula $\theta(x, \overline{y})$ of L and tuple $\overline{\mathbf{g}}$ in $\mathcal{K}$, there exists $\mathbf{f} \in \mathcal{K}$ such that

$$[\![\theta(\mathbf{f}, \overline{\mathbf{g}})]\!] = [\![(\exists\, x\, \theta)(\overline{\mathbf{g}})]\!].$$

It is shown in [5] that each continuous pre-structure induces a unique continuous structure by identifying elements at distance zero from each other and completing the metrics. It will be useful to consider these two steps separately here. By a *reduced pre-structure* we will mean a pre-structure such that $d_{\mathbf{K}}$ and $d_{\mathbf{B}}$ are metrics. Then every pre-structure $(\mathcal{K}, \mathcal{B})$ induces a unique reduced pre-structure $(\overline{\mathcal{K}}, \overline{\mathcal{B}})$ by identifying elements which are at distance zero from each other. The induced continuous structure is then obtained by completing the metrics, and will be denoted by $(\widehat{\mathcal{K}}, \widehat{\mathcal{B}})$. We say that a pre-structure $(\mathcal{K}, \mathcal{B})$ is *pre-complete* if the reduced pre-structure $(\overline{\mathcal{K}}, \overline{\mathcal{B}})$ is already a continuous structure, that is, $(\widehat{\mathcal{K}}, \widehat{\mathcal{B}}) = (\overline{\mathcal{K}}, \overline{\mathcal{B}})$. We say that $(\mathcal{K}, \mathcal{B})$ is *elementarily pre-embeddable* in $(\mathcal{K}', \mathcal{B}')$ if $(\widehat{\mathcal{K}}, \widehat{\mathcal{B}})$ is elementarily embeddable in $(\widehat{\mathcal{K}'}, \widehat{\mathcal{B}'})$.

Note that for any pre-structure $(\mathcal{K}, \mathcal{B})$, $(\widehat{\mathcal{K}}, \widehat{\mathcal{B}})$ is elementarily equivalent to $(\mathcal{K}, \mathcal{B})$, and $(\widehat{\mathcal{K}}, \widehat{\mathcal{B}})$ is separable and only if $(\mathcal{K}, \mathcal{B})$ is separable.

In [7], a randomization of $\mathcal{M}$ was defined as a three-sorted first-order structure instead of a continuous two-sorted structure, with the value space $[0, 1]$ replaced by any first-order structure $\mathcal{R}$ whose theory is an expansion of the theory of real closed ordered fields which admits quantifier elimination. When $\mathcal{R}$ is the ordered field of reals, such a structure can be interpreted as a randomization in the present sense.

The paper [7] gives axioms for a theory T^R, called the randomization theory of T, in three-sorted first-order logic. We now translate these axioms into a theory in the (two-sorted) continuous logic L^R in the sense of [5], with a connective for each continuous function $[0, 1]^n \mapsto [0, 1]$ which is definable in $\mathcal{R}$. We use φ for arbitrary formulas of L and Φ for arbitrary formulas of the continuous logic L^R. Following [5], we use the notation $\forall\, x(\Phi(x) \leq r)$ for $(\sup_x \Phi(x)) \leq r$, and $\exists\, x(\Phi(x) \leq r)$ for $(\inf_x \Phi(x)) \leq r$. Thus the existential quantifiers are understood in the approximate sense. We also use the notation $\mathsf{U} \doteq \mathsf{V}$ for the statement $d_{\mathbf{B}}(\mathsf{U}, \mathsf{V}) = 0$.

The axioms for the randomization theory T^R of a first-order theory T are as follows:

Validity Axioms:

$$\forall\, \overline{x}([\![\psi(\overline{x})]\!] \doteq \top),$$

where $\forall\, \overline{x}\, \psi(\overline{x})$ is logically valid in first-order logic.

Boolean Axioms: The usual Boolean algebra axioms in sort $\mathbf{B}$, and the statements

$$\forall\, \overline{x}([\![(\neg\varphi)(\overline{x})]\!] \doteq \neg[\![\varphi(\overline{x})]\!]),$$
$$\forall\, \overline{x}([\![(\varphi \vee \psi)(\overline{x})]\!] \doteq [\![\varphi(\overline{x})]\!] \sqcup [\![\psi(\overline{x})]\!]),$$
$$\forall\, \overline{x}([\![(\varphi \wedge \psi)(\overline{x})]\!] \doteq [\![\varphi(\overline{x})]\!] \sqcap [\![\psi(\overline{x})]\!]).$$

Distance Axioms:

$$\forall x \, \forall y \, d_{\mathbf{K}}(x,y) = 1 - \mu[\![x = y]\!], \quad \forall \mathsf{U} \, \forall \mathsf{V} \, d_{\mathbf{B}}(\mathsf{U},\mathsf{V}) = \mu(\mathsf{U} \triangle \mathsf{V}).$$

Fullness Axioms (*or Maximal Principle*):

$$\forall \overline{y} \, \exists \, x([\![\varphi(x,\overline{y})]\!] \doteq [\![(\exists \, x\varphi)(\overline{y})]\!]).$$

As mentioned previously, the quantifiers should be understood in the approximate sense, so this axiom can also be written in the equivalent form

$$\sup_{\overline{y}} \inf_{x}(d_{\mathbf{B}}([\![\varphi(x,\overline{y})]\!], [\![(\exists \, x\varphi)(\overline{y})]\!])) = 0.$$

Event Axiom:

$$\forall \mathsf{U} \, \exists \, x \, \exists \, y \, (\mathsf{U} \doteq [\![x = y]\!]).$$

Measure Axioms:

$$\mu[\top] = 1 \wedge \mu[\bot] = 0,$$

$$\forall \mathsf{U} \, \forall \mathsf{V} \, (\mu[\mathsf{U}] + \mu[\mathsf{V}] = \mu[\mathsf{U} \sqcup \mathsf{V}] + \mu[\mathsf{U} \sqcap \mathsf{V}]).$$

Atomless Axiom:

$$\forall \mathsf{U} \, \exists \, \mathsf{V} \, (\mu[\mathsf{U} \sqcap \mathsf{V}] = \mu[\mathsf{U}]/2).$$

Transfer Axiom:

$$[\![\varphi]\!] \doteq \top,$$

where $\varphi \in T$.

Theorem 2.1. ([7], Theorem 3.10, restated) *If T is complete, then T^R is complete.*

Proposition 2.2. ([7], Proposition 4.3, restated) *Let T be the complete theory of $\mathcal{M}$. Every randomization $(\mathcal{K}, \mathcal{B})$ of $\mathcal{M}$ is a pre-model of the randomization theory T^R.*

Let $(\mathcal{K}, \mathcal{B})$ and $(\mathcal{K}', \mathcal{B}')$ be pre-models of T^R. We say that $(\mathcal{K}, \mathcal{B})$ *represents* $(\mathcal{K}', \mathcal{B}')$ if their corresponding reduced pre-structures are isomorphic.

Theorem 2.3. (First Representation Theorem) *Let T be the complete theory of $\mathcal{M}$. Every pre-model of T^R is represented by some randomization of $\mathcal{M}$.*

Proof. Let $(\mathcal{K}, \mathcal{B})$ be a pre-model of T^R. It follows from [7], Corollary 6.6 and Theorem 5.7, that the reduced structure $(\overline{\mathcal{K}}, \overline{\mathcal{B}})$ of $(\mathcal{K}, \mathcal{B})$ is isomorphically embeddable in the reduced pre-structure of a randomization of $\mathcal{M}$. Since $(\mathcal{K}, \mathcal{B})$ satisfies the Fullness and Event Axioms, $(\overline{\mathcal{K}}, \overline{\mathcal{B}})$ is isomorphic to the reduced pre-structure of a randomization of $\mathcal{M}$. $\qquad\square$

Definition 2.4. A pre-model $(\mathcal{K}, \mathcal{B})$ of T^R has *perfect witnesses* if the existential quantifiers in the Fullness Axioms and the Event Axiom have witnesses in which the axioms hold exactly rather than merely approximately. That is,

Fullness: For each $\overline{\mathbf{g}}$ in $\mathcal{K}^n$, there exists $\mathbf{f} \in \mathcal{K}$ such that

$$[[\varphi(\mathbf{f}, \overline{\mathbf{g}})]] \doteq [[(\exists\, x\varphi)(\overline{\mathbf{g}})]].$$

Event: For each $\mathsf{B} \in \mathcal{B}$ there exist $\mathbf{f}, \mathbf{g} \in \mathcal{K}$ such that $\mathsf{B} \doteq [[\mathbf{f} = \mathbf{g}]]$.

The first-order models considered in [7] are pre-models with perfect witnesses when viewed as metric structures.

Proposition 2.5. ([7], Proposition 4.3, restated) *Let T be the complete theory of $\mathcal{M}$. Every full randomization $(\mathcal{K}, \mathcal{B})$ of $\mathcal{M}$ has perfect witnesses.*

Theorem 2.6. (Second Representation Theorem) ([7], Theorem 4.5, restated) *Let T be the complete theory of $\mathcal{M}$. Every pre-model of T^R with perfect witnesses is represented by some full randomization of $\mathcal{M}$.*

We now show that every model of T^R has perfect witnesses.

Theorem 2.7. *For any first-order theory T, every pre-complete model of T^R has perfect witnesses. In particular, every model of T^R has perfect witnesses.*

Proof. Let $(\mathcal{K}, \mathcal{B})$ be a pre-complete model of T^R. We first show that $(\mathcal{K}, \mathcal{B})$ has perfect witnesses for the Fullness Axioms. Consider a first-order formula $\varphi(x, \overline{y})$, and let $\overline{\mathbf{g}}$ be a tuple in $\mathcal{K}$. Then for each $n \in \omega$ there exists $\mathbf{f}_n \in \mathcal{K}$ such that

$$d_{\mathbf{B}}([[\varphi(\mathbf{f}_n, \overline{\mathbf{g}})]], [[(\exists\, x\varphi)(\overline{\mathbf{g}})]]) < 2^{-n}.$$

Using the Fullness Axioms, we can get a sequence $\mathbf{h}_n$ in $\mathcal{K}$ such that $\mathbf{h}_1 = \mathbf{f}_1$ and with probability at least $1 - 2^{-2n}$, $\mathbf{h}_{n+1}$ agrees with $\mathbf{h}_n$ when $\varphi(\mathbf{h}_n, \overline{\mathbf{g}})$ holds, and agrees with $\mathbf{f}_{n+1}$ otherwise. Then

$$d_{\mathbf{B}}([[\varphi(\mathbf{h}_n, \overline{\mathbf{g}})]], [[(\exists\, x\varphi)(\overline{\mathbf{g}})]]) < 2^{-n} + 2^{-2n}$$

and

$$d_{\mathbf{K}}(\mathbf{h}_n, \mathbf{h}_{n+1}) < 2^{-n} + 2^{-2n}.$$

Then the sequence $\mathbf{h}_n$ is Cauchy convergent with respect to $d_{\mathbf{K}}$. By pre-completeness, $\mathbf{h}_n$ converges to an element $\mathbf{h} \in \mathcal{K}$ with respect to $d_{\mathbf{K}}$. It follows that

$$d_{\mathbf{B}}([[\varphi(\mathbf{h}, \overline{\mathbf{g}})]], [[(\exists\, x\varphi)(\overline{\mathbf{g}})]]) = 0,$$

as required.

It remains to show that $(\mathcal{K}, \mathcal{B})$ has perfect witnesses for the Event Axiom. Let $\mathsf{B} \in \mathcal{B}$. By the Event Axiom, for each $n \in \omega$ there exist $\mathbf{f}_n, \mathbf{g}_n \in \mathcal{K}$ such that $\mu[\mathsf{B} \triangle [[\mathbf{f}_n = \mathbf{g}_n]]] < 2^{-n}$. Since every model of T has at least two elements, there

must exist $\mathbf{f}', \mathbf{g}' \in \mathcal{K}$ such that $[\![\mathbf{f}' = \mathbf{g}']\!] \doteq \perp$. Taking perfect witnesses for the Fullness Axioms, we can obtain elements $\mathbf{h}_n \in \mathcal{K}$ such that $\mathbf{h}_n$ agrees with $\mathbf{f}'$ on $[\![\mathbf{f}_n = \mathbf{g}_n]\!]$, and $\mathbf{h}_n$ agrees with $\mathbf{g}'$ on $[\![\mathbf{f}_n \neq \mathbf{g}_n]\!]$. Then the sequence $\mathbf{h}_n$ is Cauchy convergent with respect to $d_{\mathbf{K}}$, and hence converges to some $\mathbf{h}$ in $\mathcal{K}$. It follows that $\mathsf{B} \doteq [\![\mathbf{f}' = \mathbf{h}]\!]$. □

Corollary 2.8. *Let T be the complete theory of $\mathcal{M}$. Every pre-complete model of T^R, and hence every model of T^R, is represented by some full randomization of $\mathcal{M}$.*

Theorem 2.9. ([7], Theorems 3.6 and 5.1, restated) *For any first-order theory T, the randomization theory T^R for T admits strong quantifier elimination.*

This means that every formula Φ in the continuous language L^R is T^R-equivalent to a formula with the same free variables and no quantifiers of sort $\mathbf{K}$ or $\mathbf{B}$ (whereas ordinary quantifier elimination for continuous logic means that every formula can be arbitrarily well approximated by quantifier-free formulas). Note that for each first-order formula $\varphi(\overline{x})$ of L, $\mu[\![\varphi(\overline{x})]\!]$ is an atomic formula of L^R which has no quantifiers. The first-order quantifiers within $\varphi(\overline{x})$ do not count as quantifiers in L^R.

By the Event Axiom and Theorem 2.7, in a model of T^R, any element of sort $\mathbf{B}$ is equal to a term $[\![\mathbf{g}_1 = \mathbf{g}_2]\!]$ with parameters $\mathbf{g}_1, \mathbf{g}_2$ of sort $\mathbf{K}$. Therefore, in all discussions of types in the theory T^R, we may confine our attention to types of sort $\mathbf{K}$ over parameters of sort $\mathbf{K}$.

The space of first-order n-types in T will be denoted by $S_n(T)$. If $T = Th(\mathcal{M})$ and A is a set of parameters in $\mathcal{M}$, $S_n(T(A))$ is the space of first-order n-types in T with parameters in A. The space of continuous n-types in T^R with variables of sort $\mathbf{K}$ will be denoted by $S_n(T^R)$. We will use boldface letters $\mathbf{p}, \mathbf{q}, \ldots$ for types of sort $\mathbf{K}$, and boldface $\mathbf{A}$ for a set of parameters of sort $\mathbf{K}$. If $(\mathcal{K}, \mathcal{B})$ is a pre-model of T^R and $\mathbf{A}$ is a set of parameters in $\mathcal{K}$, $S_n(T^R(\mathbf{A}))$ is the space of continuous n-types in T^R with variables of sort $\mathbf{K}$ and parameters from $\mathbf{A}$.

Recall from [5] that for each $\mathbf{p} \in S_n(T^R)$ and formula $\Phi(\overline{x})$ of L^R, we have $(\Phi(\overline{x}))^{\mathbf{p}} \in [0, 1]$.

Let $\mathfrak{R}(S_n(T))$ be the space of regular Borel probability measures on $S_n(T)$. The next corollary follows from quantifier elimination and the axioms of T^R.

Corollary 2.10. *For every $\mathbf{p} \in S_n(T^R)$ there is a unique measure $\nu_{\mathbf{p}} \in \mathfrak{R}(S_n(T))$ such that for each formula $\varphi(\overline{x})$ of L,*

$$\nu_{\mathbf{p}}(\{q : \varphi(\overline{x}) \in q\}) = (\mu[\![\varphi(\overline{x})]\!])^{\mathbf{p}}.$$

Moreover, for each measure $\nu \in \mathfrak{R}(S_n(T))$ there is a unique $\mathbf{p} \in S_n(T^R)$ such that $\nu = \nu_{\mathbf{p}}$.

Similarly for types with infinitely many variables.

We may therefore identify the type space $S_n(T^R)$ with the space $\mathfrak{R}(S_n(T))$.

Remark 2.11. In the special case that $\mathcal{M}$ is the trivial two-element structure with only the identity relation, the above results show that the continuous theory of atomless measure algebras is complete and admits quantifier elimination. (See [7], Sec. 7A.) This fact was given a direct proof in [5].

Remark 2.12. Throughout this paper, we could have worked with a one-sorted randomization theory with only the sort $\mathbf{K}$ instead of a two-sorted theory with sorts $\mathbf{K}$ and $\mathbf{B}$. In this formulation, we would use the results of Sec. 7C of [7], where the sort $\mathbf{B}$ is eliminated. This approach will be taken at the end of Sec. 5.3. The main advantages of the event sort $\mathbf{B}$ are that it allows a nicer set of axioms for T^R, and makes it easier to describe the models of T^R.

3. Separable Structures

In this section we consider small randomizations of $\mathcal{M}$. In order to cover the case that T has finite models, we adopt the following convention:

"*$\mathcal{M}$ is countable*" *means that the cardinality of $\mathcal{M}$ is finite or ω.*

So a complete first-order theory T is ω-categorical if and only if it either has a finite model or has a unique model of cardinality ω up to isomorphism.

We first show that the randomization operation preserves ω-categoricity. We will use the following necessary and sufficient condition for a continuous theory to be ω-categorical.

Fact 3.1. ([1], Theorem 13.8) Let U be a complete continuous theory with a countable signature. Then U is ω-categorical if and only if for each $n \geq 1$, every type in $S_n(U)$ is realized in every model of U.

Theorem 3.2. *Suppose T has a countable signature. Then T is ω-categorical if and only if T^R is ω-categorical.*

Proof. Suppose that T is not ω-categorical. Then by the Ryll–Nardzewski Theorem, for some n there is an n-type $q \in S_n(T)$ which is omitted in some countable model $\mathcal{M}$ of T. Then the type $\mathbf{p} \in S_n(T^R)$ such that $\nu_{\mathbf{p}}(\{q\}) = 1$ is omitted in the separable pre-complete model $(\mathcal{M}^{[0,1]}, \mathcal{L})$, of T^R, so T^R is not ω-categorical.

Now suppose T is ω-categorical. Then $T = Th(\mathcal{M})$ for some $\mathcal{M}$. By Ryll–Nardzewski's theorem, for each n the type space $S_n(T)$ is finite. Let $\mathbf{p} \in S_n(T^R)$. By Corollary 2.10, $\nu_{\mathbf{p}} \in \mathfrak{R}(S_n(T))$, and for each formula $\varphi(\overline{x})$ of L,

$$\nu_{\mathbf{p}}(\{q : \varphi(\overline{x}) \in q\}) = (\mu[\![\varphi(\overline{x})]\!])^{\mathbf{P}}.$$

Since $S_n(T)$ is finite, for each $q \in S_n(T)$ there is a formula $\varphi_q(\overline{x})$ of L such that q is the set of T-consequences of φ_q. Let $(\mathcal{K}', \mathcal{B}')$ be a model of T^R. By Corollary 2.8, $(\mathcal{K}', \mathcal{B}')$ is represented by some full pre-complete randomization $(\mathcal{K}, \mathcal{B})$ of $\mathcal{M}$. By Fact 3.1, it suffices to show that $\mathbf{p}$ is realized in $(\mathcal{K}, \mathcal{B})$. By the axioms and pre-completeness, there are events $\mathsf{B}_q, q \in S_n(T)$ in $\mathcal{B}$ such that $\mathsf{B}_q, q \in S_n(T)$ partitions

Ω and $\mu(\mathsf{B}_q) = \nu_{\mathbf{p}}(\{q\})$. Since $(\mathcal{K}, \mathcal{B})$ has perfect witnesses, for each q there is a tuple $\overline{\mathbf{f}}_q$ in $\mathcal{K}$ such that $\mu[\![\varphi_q(\overline{\mathbf{f}}_q)]\!] = 1$. Using the Fullness and Event Axioms, there is a tuple $\overline{\mathbf{f}}$ in $\mathcal{K}$ which agrees with $\overline{\mathbf{f}}_q$ on B_q for each $q \in S_n(T)$. It follows that $\overline{\mathbf{f}}$ realizes $\mathbf{p}$ in $(\mathcal{K}, \mathcal{B})$, as required. $\square$

We next give a method for constructing small models of T^R, and then introduce the notion of a strongly separable model of T^R.

Definition 3.3. Let $(\Omega, \mathcal{B}, \mu)$ be an atomless finitely additive probability space and let $\mathcal{M}$ be a structure for L.

A $\mathcal{B}$-*deterministic element* of $\mathcal{M}$ is a constant function in $\mathcal{M}^\Omega$ (which may be thought of as an element of $\mathcal{M}$).

A $\mathcal{B}$-*simple random element* of $\mathcal{M}$ is a $\mathcal{B}$-measurable function in $\mathcal{M}^\Omega$ with finite range.

A $\mathcal{B}$-*countable random element* of $\mathcal{M}$ is a $\mathcal{B}$-measurable function in $\mathcal{M}^\Omega$ with countable range.

Example 3.4. ([7], Examples 4.6 and 4.11)

(i) The set $\mathcal{K}_S$ of $\mathcal{B}$-simple random elements of $\mathcal{M}$ is full, and $(\mathcal{K}_S, \mathcal{B})$ is a full randomization of $\mathcal{M}$.

(ii) The set $\mathcal{K}_C$ of $\mathcal{B}$-countable random elements of $\mathcal{M}$ is full. If $(\Omega, \mathcal{B}, \mu)$ is a probability space (i.e. is σ-additive), then $(\mathcal{K}_C, \mathcal{B})$ is a full pre-complete randomization of $\mathcal{M}$.

We assume for the rest of this section that the signature L of T is countable

Note that when L is countable, every Borel probability measure on $S_n(T)$ is regular. (This follows from [6], p. 228, and the fact that every open set in $S_n(T)$ is a countable union of compact sets.)

Definition 3.5. Let $([0, 1], \mathcal{L}, \lambda)$ be the natural atomless Borel probability measure on $[0, 1]$. Given a model $\mathcal{M}$ of T, let $(\mathcal{M}^{[0,1]}, \mathcal{L})$ be the pre-structure whose universe of sort $\mathbf{K}$ is the set of $\mathcal{L}$-countable random elements of $\mathcal{M}$. A pre-model $(\mathcal{K}, \mathcal{B})$ of T^R is *strongly separable* if $(\mathcal{K}, \mathcal{B})$ is elementarily pre-embeddable in $(\mathcal{M}^{[0,1]}, \mathcal{L})$ for some countable model $\mathcal{M}$ of T.

Note that $(\mathcal{M}^{[0,1]}, \mathcal{L})$ is only a pre-structure, not a reduced pre-structure, because $\mathcal{L}$ has nonempty null sets.

Corollary 3.6. *If $\mathcal{M}$ is a countable model of T, then $(\mathcal{M}^{[0,1]}, \mathcal{L})$ is a full randomization of $\mathcal{M}$ and is a strongly separable pre-complete model of T^R.*

Proof. By Example 3.4(ii). $\square$

Lemma 3.7. *Consider a first-order L-structure $\mathcal{M}$. Suppose $(\Omega, \mathcal{B}, \mu)$ is a probability space and $\mathcal{A}$ is a subalgebra of $\mathcal{B}$ which is dense with respect to the distance*

predicate $d_{\mathbf{B}}(\mathsf{U}, \mathsf{V}) = \mu(\mathsf{U} \triangle \mathsf{V})$. *Then the set* $\mathcal{K}_S$ *of* $\mathcal{A}$-*simple random elements of* $\mathcal{M}$ *is dense in the set* $\mathcal{K}$ *of* $\mathcal{B}$-*countable random elements of* $\mathcal{M}$ *with respect to the distance predicate* $d_{\mathbf{K}}$. *Therefore* $(\mathcal{K}_S, \mathcal{A})$ *and* $(\mathcal{K}, \mathcal{B})$ *induce the same continuous structure, and* $\widehat{\mathcal{K}}_S = \widehat{\mathcal{K}}$.

Proof. By Example 3.4, $(\mathcal{K}_S, \mathcal{A})$ and $(\mathcal{K}, \mathcal{B})$ are full randomizations of $\mathcal{M}$. Let $\mathbf{f}$ be an element of $\mathcal{K}$ with range $\{a_n : n \in \mathbb{N}\}$. Let $\varepsilon > 0$. Take n such that

$$\mu \left[\!\!\left[\bigvee_{m < n} (\mathbf{f} = a_m) \right]\!\!\right] > 1 - \varepsilon/2.$$

For each $m < n$, let

$$\mathsf{B}_m = \{w : \mathbf{f}(w) = a_m\} \in \mathcal{B}.$$

Then the sets B_m are disjoint and

$$\mu \left[\bigcup_{m < n} \mathsf{B}_m \right] > 1 - \varepsilon/2.$$

For each $m < n$, there is a set $\mathsf{A}_m \in \mathcal{A}$ such that $d_{\mathbf{B}}(\mathsf{A}_m, \mathsf{B}_m) < \varepsilon/(4n^2)$. Let $\mathsf{C}_m = \mathsf{A}_m \backslash \bigcup_{k < m} \mathsf{A}_k$. Then $\mathsf{C}_m \in \mathcal{A}$, the sets C_m are disjoint, and one can check that $d_{\mathbf{B}}(\mathsf{C}_m, \mathsf{B}_m) < \varepsilon/(2n)$. There is an $\mathcal{A}$-simple $\mathbf{g} \in \mathcal{K}_S$ such that $\mathbf{g}(w) = a_m$ whenever $m < n$ and $w \in \mathsf{C}_m$. Then $\mathbf{f}(w) = \mathbf{g}(w)$ whenever $w \in \mathsf{B}_m \cap \mathsf{C}_m$, so

$$\mu[\![\mathbf{f} = \mathbf{g}]\!] > 1 - \varepsilon,$$

and hence $d_{\mathbf{K}}(\mathbf{f}, \mathbf{g}) < \varepsilon$. This shows that $\mathcal{K}_S$ is dense in $\mathcal{K}$.   □

Corollary 3.8. *Every strongly separable pre-model of* T^R *is separable.*

Proof. Let $(\mathcal{K}, \mathcal{B})$ be a strongly separable pre-model of T^R, so that $(\mathcal{K}, \mathcal{B})$ is elementarily pre-embeddable in $(\mathcal{M}^{[0,1]}, \mathcal{L})$ for some countable model $\mathcal{M}$ of T. $\mathcal{L}$ has a countable dense subalgebra $\mathcal{A}$. By Lemma 3.7, the set of $\mathcal{A}$-simple random elements of $\mathcal{M}$ is a countable dense subset of $\mathcal{M}^{[0,1]}$ with respect to $d_{\mathbf{K}}$. Therefore $(\mathcal{M}^{[0,1]}, \mathcal{L})$ is separable, and hence $(\mathcal{K}, \mathcal{B})$ is separable.   □

Lemma 3.9. *Let* T *be complete,* $\mathbf{p} \in S_n(T^R)$, *and* $\nu_{\mathbf{p}}$ *be the corresponding measure on* $S_n(T)$ *defined in Corollary 2.10. Then* $\mathbf{p}$ *is realized in some strongly separable model of* T^R *if and only if there is a countable set* $C \subseteq S_n(T)$ *such that* $\nu_{\mathbf{p}}(C) = 1$.

Proof. Suppose $\mathbf{p}$ is realized in some strongly separable model of T^R. Then $\mathbf{p}$ is realized by an n-tuple $\overline{\mathbf{g}}$ in the pre-complete model $(\mathcal{M}^{[0,1]}, \mathcal{L})$ of T^R for some countable model $\mathcal{M}$ of T. Hence $\nu_{\mathbf{p}}(C) = 1$ where C is the set of types of elements of the range of $\overline{\mathbf{g}}$.

Suppose $\nu_{\mathbf{p}}(C) = 1$ for some countable set $C \subseteq S_n(T)$, and let $\mathcal{M}$ be a countable model of T which realizes each $q \in C$. Then $\mathbf{p}$ is realized in the strongly separable pre-model $(\mathcal{M}^{[0,1]}, \mathcal{L})$ of T^R by any tuple $\overline{\mathbf{g}}$ such that

$$\lambda(\{r : tp^{\mathcal{M}}(\overline{\mathbf{g}}(r)) = q\}) = \nu_{\mathbf{p}}(\{q\})$$

for each $q \in C$. $\square$

Example 3.10. (i) Let T be the first-order theory with countably many independent unary relations $P_0, P_1, \ldots$, and let $\mathbf{p}$ be the type in $S_1(T^R)$ such that the events $[[P_n(x)]], n \in \mathbb{N}$ are independent with respect to $\mathbf{p}$. Then $\mathbf{p}$ is not realized in a strongly separable model of T^R, and hence T^R has separable models which are not strongly separable.

(ii) Let $\mathbb{R}$ be the ordered field of real numbers, as a first order rather than continuous structure. Let $\mathcal{K}$ be the set of all Lebesgue measurable functions from $[0,1]$ into $\mathbb{R}$. By Proposition 4.12 in [7], $(\mathcal{K}, \mathcal{L})$ is a full pre-complete randomization of $\mathbb{R}$. $(\mathcal{K}, \mathcal{L})$ is not separable, and the type of the identity function on $[0,1]$ in $(\mathcal{K}, \mathcal{L})$ is not realized in a strongly separable model.

Lemma 3.11. *Let $\mathcal{M}$ be countable. Then $\mathcal{M}$ is ω-saturated if and only if $(\mathcal{M}^{[0,1]}, \mathcal{L})$ is ω-saturated as a continuous pre-structure.*

Proof. Suppose first that $\mathcal{M}$ is not ω-saturated. Then there is a tuple $\overline{a}$ in $\mathcal{M}$, a countable elementary extension $\mathcal{N}$ of $\mathcal{M}$, and an element $b \in \mathcal{N}$ such that the type of b in $(\mathcal{N}, \overline{a})$ is not realized in $(\mathcal{M}, \overline{a})$. By considering the types of the constant functions from $[0,1]$ to $\overline{a}$ and b, we see that $\mathcal{M}^{[0,1]}$ is not ω-saturated.

Now suppose that $\mathcal{M}$ is ω-saturated. Let T be the complete theory of $\mathcal{M}$ and let $\mathcal{K} = \mathcal{M}^{[0,1]}$. It suffices to show that for each finite tuple $\overline{\mathbf{g}}$ in $\mathcal{K}$, every type $\mathbf{p} \in S_1(T^R(\overline{\mathbf{g}}))$ is realized in $(\mathcal{K}, \mathcal{L})$. Let $\{\overline{b}_n : n < \omega\}$ be the range of $\overline{\mathbf{g}}$ in $\mathcal{M}$ and let

$$\mathsf{B}_n = \{r \in [0,1] : \overline{\mathbf{g}}(r) = \overline{b}_n\} = [[\overline{\mathbf{g}} = \overline{b}_n]].$$

Note that $\mathsf{B}_n, n \in \omega$ is a partition of $[0,1]$ in $\mathcal{L}$, and

$$\sum_n \mu[\mathsf{B}_n] = 1.$$

Since $\mathcal{M}$ is countable and ω-saturated, for each n there are countably many types $q \in S_1(T(\overline{b}_n))$, and each of these types is realized by an element a_q in $\mathcal{M}$. Let $\mathbf{p} \in S_1(T^R(\overline{\mathbf{g}}))$. Let $\mathbf{A}$ be the set of all deterministic elements of $\mathcal{K}$, and extend $\mathbf{p}$ to a type $\mathbf{p}'$ in $S_1(T^R(\mathbf{A} \cup \overline{\mathbf{g}}))$. For each n and first-order formula $\varphi(x, \overline{b}_n)$, $\mu[[[\varphi(x, \overline{b}_n)]] \sqcap \mathsf{B}_n]$ is a continuous formula with parameters in $\mathbf{A} \cup \overline{\mathbf{g}}$. Let

$$\alpha(\varphi, n) = (\mu[[[\varphi(x, \overline{b}_n)]] \sqcap \mathsf{B}_n])^{\mathbf{p}'}$$

and for each $q \in S_1(T(\overline{b}_n))$ let

$$\beta(q, n) = \inf\{\alpha(\varphi, n) : \varphi \in q\}.$$

Since $S_1(T(\overline{b}_n))$ is countable for each n, we have

$$\sum\{\beta(q, n) : q \in S_1(T(\overline{b}_n))\} = \mu[\mathsf{B}_n].$$

There is a function $\mathbf{f} \in \mathcal{K}$ such that for each n and $q \in S_1(T(\overline{b}_n))$,

$$\mu\{r \in \mathsf{B}_n : \mathbf{f}(r) = a_q\} = \beta(q, n).$$

Then for each n and first-order formula $\varphi(x, b_n)$,

$$\mu[[\![\varphi(\mathbf{f}, \overline{\mathbf{g}})]\!] \sqcap \mathsf{B}_n] \geq \alpha(\varphi, n).$$

It follows that $\mathbf{f}$ realizes $\mathbf{p}$ in $(\mathcal{K}, \mathcal{L})$. $\qquad\square$

Theorem 3.12. *Let T be complete. The following are equivalent.*

(i) *T has a countable ω-saturated model.*
(ii) *T^R has a separable ω-saturated model.*
(iii) *Every separable model of T^R is strongly separable.*

Proof. (i) implies (ii) and (iii): Let $\mathcal{M}$ be a countable ω-saturated model of T and let $\mathcal{K} = \mathcal{M}^{[0,1]}$. Then $(\mathcal{K}, \mathcal{L})$ is strongly separable, and hence is separable by Corollary 3.8. By Lemma 3.11, $(\mathcal{K}, \mathcal{L})$ is ω-saturated, so (ii) holds. It follows that every separable model of T^R is elementarily pre-embeddable in the strongly separable pre-model $(\mathcal{K}, \mathcal{L})$, and hence is itself strongly separable. This proves (iii).

(ii) implies (i): Assume that (i) fails. Then for some n, $S_n(T)$ is uncountable. For each $q \in S_n(T)$ there is an n-type $\mathbf{q}'$ in $S_n(T^R)$ such that $\nu_{\mathbf{q}'}(\{q\}) = 1$. Let $(\mathcal{K}, \mathcal{B})$ be an ω-saturated model of T^R. Each n-type $\mathbf{q}'$ is realized by an n-tuple $\overline{\mathbf{f}}_q$ in $(\mathcal{K}, \mathcal{B})$. Suppose $p, q \in S_n(T)$ and $p \neq q$. Then $\nu_{\mathbf{p}'}(\{q\}) = 0$ and $\nu_{\mathbf{q}'}(\{q\}) = 1$. By the First Representation Theorem 2.3, $(\mathcal{K}, \mathcal{B})$ is represented by a randomization of a model of T, so we may assume that $(\mathcal{K}, \mathcal{B})$ is already a randomization of a model of T. It follows that $\overline{\mathbf{f}}_p \neq \overline{\mathbf{f}}_q$ almost everywhere, and $\mu[[\![\overline{\mathbf{f}}_p \neq \overline{\mathbf{f}}_q]\!]] = 1$. Let d_n be the metric on $\mathcal{K}^n$ formed by adding the $d_{\mathbf{K}}$ distances at each coordinate. Then

$$d_n(\overline{\mathbf{f}}_p, \overline{\mathbf{f}}_q) \geq \mu[[\![\overline{\mathbf{f}}_p \neq \overline{\mathbf{f}}_q]\!]] = 1.$$

Since $S_n(T)$ is uncountable, we see that the metric space $(\mathcal{K}^n, d_n)$ is not separable. Therefore $(\mathcal{K}, d_{\mathbf{K}})$ is not separable, so $(\mathcal{K}, \mathcal{B})$ is not separable and (ii) fails.

(iii) implies (i). Again assume (i) fails and $S_n(T)$ is uncountable. Then by enumerating the formulas of L one can construct a measure $\nu \in \mathfrak{R}(S_n(T))$ such that $\nu(\{q\}) = 0$ for each $q \in S_n(T)$. By Corollary 2.10 we may take $\mathbf{p} \in S_n(T^R)$ such that $\nu_{\mathbf{p}} = \nu$. By Lemma 3.9, $\mathbf{p}$ cannot be realized in a strongly separable model, but $\mathbf{p}$ can be realized in a separable model. Therefore (iii) fails. $\qquad\square$

4. ω-Stable Theories

In this section we continue to assume that the signature L of T is countable.

As explained in Sec. 2, when considering types in T^R, we may confine our attention to types of sort $\mathbf{K}$ over parameters of sort $\mathbf{K}$. For each model $(\mathcal{K}, \mathcal{B})$ of T^R and set of parameters $\mathbf{A} \subseteq \mathcal{K}$, the d metric on the type space $S_1(T^R(\mathbf{A}))$ is defined by

$$d(\mathbf{p}, \mathbf{q}) = \inf\{d_{\mathbf{K}}(\mathbf{f}, \mathbf{g}) \colon tp(\mathbf{f}/\mathbf{A}) = \mathbf{p} \text{ and } tp(\mathbf{g}/\mathbf{A}) = \mathbf{q} \text{ in some } (\mathcal{K}', \mathcal{B}') \succ (\mathcal{K}, \mathcal{B})\}.$$

We say that the space $S_1(T^R(\mathbf{A}))$ is *separable* if it is separable with respect to the d metric. Following [1], we say that T^R is *ω-stable* if $S_1(T^R(\mathbf{A}))$ is separable for every model $(\mathcal{K}, \mathcal{B})$ of T^R and countable set $\mathbf{A}$ of parameters in $\mathcal{K}$.

We show that if T has at most countably many complete extensions, then the randomization operation preserves ω-stability. We first take up the case that T is complete.

4.1. *Complete theories*

Theorem 4.1. *A complete theory T is ω-stable if and only if T^R is ω-stable.*

Proof. Assume T^R is ω-stable. Let $\mathcal{M}$ be an arbitrary countable model of T, let $\mathcal{K} = \mathcal{M}^{[0,1]}$, and let $\mathbf{A}$ be the set of deterministic elements of $\mathcal{K}$. Then $\mathbf{A}$ is countable, so $S_1(T^R(\mathbf{A}))$ is separable.

For each $q \in S_1(T(\mathcal{M}))$ let $\mathbf{q}'$ be the type in $S_1(T^R(\mathbf{A}))$ such that for each first-order formula $\varphi(x, \overline{b})$ with parameters in $\mathcal{M}$,

$$(\mu[\![\varphi(x, \overline{b})]\!])^{\mathbf{q}'} = \begin{cases} 1 & \text{if } \varphi(x, \overline{b}) \in q, \\ 0 & \text{otherwise.} \end{cases}$$

Then $q \mapsto \mathbf{q}'$ is a mapping from $S_1(T(\mathcal{M}))$ into $S_1(T^R(\mathbf{A}))$ such that $p \neq q$ implies $d(\mathbf{p}', \mathbf{q}') = 1$. Since $S_1(T^R(\mathbf{A}))$ is separable, it follows that $S_1(T(\mathcal{M}))$ is countable, so T is ω-stable.

Now assume that T is ω-stable. Then T has a countable saturated model $\mathcal{M}_1$, and $\mathcal{M}_1$ has a countable elementary extension $\mathcal{M}_2$ which realizes every type in $S_1(T(\mathcal{M}_1))$. Let $([0,1] \times [0,1], \mathcal{L} \otimes \mathcal{L}, \lambda \otimes \lambda)$ be the natural Borel measure on the unit square. Let $\mathcal{B}_1$ be the subalgebra of $\mathcal{L} \otimes \mathcal{L}$ generated by $\{\mathsf{B} \times [0,1] : \mathsf{B} \in \mathcal{L}\}$, and let $\mathcal{B}_2 = \mathcal{L} \otimes \mathcal{L}$. Let $\mathcal{K}_1$ and $\mathcal{K}_2$ be the sets of $\mathcal{B}_1$-countable random elements of $\mathcal{M}_1$ and $\mathcal{B}_2$-countable random elements of $\mathcal{M}_2$ respectively. Then $(\mathcal{K}_1, \mathcal{B}_1)$ and $(\mathcal{K}_2, \mathcal{B}_2)$ are full randomizations of $\mathcal{M}_1$ and $\mathcal{M}_2$ respectively. By Theorems 2.9 and 2.1, $(\mathcal{K}_1, \mathcal{B}_1)$ is an elementary pre-substructure of $(\mathcal{K}_2, \mathcal{B}_2)$. By Lemma 3.11, $(\mathcal{K}_1, \mathcal{B}_1)$ is ω-saturated. By Lemma 3.7, $\mathcal{K}_2$ has a countable dense subset. It therefore suffices to show that every type in $S_1(T^R(\mathcal{K}_1))$ is realized in $(\mathcal{K}_2, \mathcal{B}_2)$.

Let $\mathbf{p}$ be a type in $S_1(T^R(\mathcal{K}_1))$. Let $\{q_0, q_1, \ldots\}$ be an enumeration of the countable set of types $S_1(T(\mathcal{M}_1))$, and let a_n realize q_n in $\mathcal{M}_2$. It follows from the axioms of T^R that for each n, there is a unique (σ-additive) measure ν_n on $\mathcal{L}$ defined by

$$\nu_n(\mathsf{B}) = \inf\{(\mu[[\![\varphi(x, \overline{\mathbf{g}})]\!]] \sqcap \mathsf{B])^p : \overline{\mathbf{g}} \text{ deterministic}, \varphi(x, \overline{\mathbf{g}}) \in q_n\}.$$

Then ν_n is absolutely continuous with respect to λ, and $\nu(\mathsf{B}) = \sum_n \nu_n(\mathsf{B})$ is a probability measure on $\mathcal{L}$. By the Radon–Nikodym theorem, there is an $\mathcal{L}$-measurable function $f_n : [0, 1] \to [0, 1]$ such that $\nu_n(\mathsf{B}) = \int_{\mathsf{B}} f_n d\lambda$. Define $\mathbf{f} \in \mathcal{K}_2$ as follows. For $(r, s) \in [0, 1] \times [0, 1)$, let $\mathbf{f}(r, s) = a_n$ if and only if

$$\sum_{k<n} f_k(r) \leq s < \sum_{k\leq n} f_k(r).$$

For $r \in [0, 1]$ let $\mathbf{f}(r, 1)$ be some particular element of $\mathcal{M}$, say $\mathbf{f}(r, 1) = a_0$.

Consider a k-tuple $\overline{\mathbf{g}}$ in $\mathcal{K}_1$. Let $\mathcal{M}_1^k = \{\overline{b}_m : m \in \omega\}$. For each $r \in [0, 1]$ we have $\overline{\mathbf{g}}(r) \in \{\overline{b}_m : m < \omega\}$. Let $\mathsf{B}_m = \{r \in [0, 1] : \overline{\mathbf{g}}(r) = \overline{b}_m\}$. Then $\mathsf{B}_m \in \mathcal{L}$. For each m and n we have

$$\nu_n(\mathsf{B}_m) = \int_{\mathsf{B}_m} f_n \, d\lambda = (\lambda \otimes \lambda)\{(r, s) : \mathbf{f}(r, s) = a_n \wedge r \in \mathsf{B}_m\}.$$

By the definition of ν_n, for each first-order formula $\varphi(x, \overline{b}_m) \in q_n$,

$$\nu_n(\mathsf{B}_m) \leq (\mu[[\![\varphi(x, \overline{\mathbf{g}})]\!]] \sqcap \mathsf{B}_m])^p.$$

It follows that $\mathbf{f}$ realizes $\mathbf{p}$ in $(\mathcal{K}_2, \mathcal{B}_2)$. $\square$

Isaac Goldbring has noted that in the above proof,

$$\nu_n(\mathsf{B}) = (\lambda \otimes \lambda)\{(r, s) : r \in \mathsf{B} \wedge \mathbf{f}(r, s) \models q_n\}$$

and hence

$$\nu(\mathsf{B}) = \sum_n \nu_n(\mathsf{B}) = (\lambda \otimes \lambda)(\mathsf{B} \times [0, 1]) = \lambda(\mathsf{B}).$$

Thus ν is the usual measure λ on $\mathcal{L}$ and does not depend on $\mathbf{p}$.

Remark 4.2. In [4] it is shown that the theory of atomless measure algebras in continuous logic is ω-stable. In view of Remark 2.11, in the case that $\mathcal{M}$ is the trivial two-element structure, Theorem 4.1 gives another proof of that fact. So Theorem 4.1 can be viewed as a generalization of the result that the theory of atomless measures is ω-stable.

4.2. *Incomplete theories*

Let $S_0(T)$ be the space of complete extensions of T. In this subsection we show that if $S_0(T)$ is countable, then T is ω-stable if and only if T^R is ω-stable.

We first take a brief detour to state a generalization of the Second Representation Theorem 2.6 which shows that every model of T^R can be regarded as a

continuous structure whose elements are random variables taking values in *random models of T*. A *full randomization* $(\mathcal{K}, \mathcal{B})$ of an indexed family $\langle \mathcal{M}(w) : w \in \Omega \rangle$ of models of T is defined in the same way as a full randomization of $\mathcal{M}$ except that $\mathcal{K} \subseteq \Pi_{w \in \Omega} \mathcal{M}(w)$.

Theorem 4.3. ([7], Proposition 5.6 and Theorem 5.7, restated) *A pre-structure $(\mathcal{K}', \mathcal{B}')$ is a pre-model of T^R with perfect witnesses if and only if it can be represented by some full randomization $(\mathcal{K}, \mathcal{B})$ of some indexed family $\langle \mathcal{M}(w) : w \in \Omega \rangle$ of models of T.*

We now introduce countable convex combinations of pre-complete models of T^R. This construction will be used in proving that the randomization operation preserves ω-stability. Let S be a countable set, and let $(S, \mathcal{B}_0, \mu_0)$ be a probability space where $\mathcal{B}_0$ is the power set of S. For each $w \in S$ let $(\mathcal{K}(w), \mathcal{B}(w))$ be a pre-complete model of T^R. We then define $\int_S (\mathcal{K}(w), \mathcal{B}(w)) \, d\mu_0(w)$ to be the pre-structure $(\mathcal{K}, \mathcal{B})$ such that $\mathcal{K} = \Pi_w \mathcal{K}(w)$, $\mathcal{B}$ is the set of events with parameters in $\mathcal{K}$, and for each tuple $\overline{\mathbf{f}}$ in $\mathcal{K}$ and first-order formula $\varphi(\overline{x})$,

$$\mu[\![\varphi(\overline{\mathbf{f}})]\!] = \int_S \mu(w)[\![\varphi(\overline{\mathbf{f}}(w))]\!] \, d\mu_0(w).$$

The proof of the following lemma is routine.

Lemma 4.4. *Let S, μ_0, and $(\mathcal{K}(w), \mathcal{B}(w))$ be as above, and let*

$$(\mathcal{K}, \mathcal{B}) = \int_S (\mathcal{K}(w), \mathcal{B}(w)) \, d\mu_0(w).$$

 (i) *$(\mathcal{K}, \mathcal{B})$ is a pre-complete model of T^R.*
 (ii) *If each $(\mathcal{K}(w), \mathcal{B}(w))$ is separable, then $(\mathcal{K}, \mathcal{B})$ is separable.*
(iii) *If each $(\mathcal{K}(w), \mathcal{B}(w))$ is ω-saturated, then $(\mathcal{K}, \mathcal{B})$ is ω-saturated.*
(iv) *If $S_1(T^R(\mathcal{K}(w)))$ is separable for each $w \in S$, then $S_1(T^R(\mathcal{K}))$ is separable.*

Theorem 4.5. *Suppose that $S_0(T)$ is countable. Then T is ω-stable if and only if T^R is ω-stable.*

Proof. Assume T^R is ω-stable. Let $\mathcal{M}$ be an arbitrary countable model of T, let $\mathcal{K} = \mathcal{M}^{[0,1]}$, and let $\mathbf{A}$ be the set of deterministic elements of $\mathcal{K}$. Then $\mathbf{A}$ is countable, so $S_1(T^R(\mathbf{A}))$ is separable. By the argument used in the proof of Theorem 4.1, it follows that $S_1(T(\mathcal{M}))$ is countable, so T is ω-stable.

Suppose T is ω-stable. Let $S = S_0(T)$. Then each $w \in S$ is a complete ω-stable theory which has a countable ω-saturated model $\mathcal{M}_1(w)$. By Theorem 2.1, each w^R is a complete extension of T^R. By Lemma 3.11 and Corollaries 3.6 and 3.8, w^R has a pre-complete separable ω-saturated model $(\mathcal{K}_1(w), \mathcal{B}_1(w))$. By Theorem 4.1, each w^R is an ω-stable theory. Therefore the type space $S_1(w^R(\mathcal{K}_1(w)))$ is separable.

Let $(\mathcal{K}, \mathcal{B})$ be a separable model of T^R. Let μ_0 be the unique measure in $\mathfrak{R}(S)$ such that for each sentence φ of L and $r \in [0, 1]$, if

$$(\mathcal{K}, \mathcal{B}) \models \mu[\![\varphi]\!] \geq r,$$

then $\mu_0(\{w \in S : w \models \varphi\}) \geq r$. The pre-models $(\mathcal{K}, \mathcal{B})$ and $(\mathcal{K}_1, \mathcal{B}_1)$ assign the same measure to each sentence φ of L. Therefore by quantifier elimination, $(\mathcal{K}_1, \mathcal{B}_1)$ is elementarily equivalent to $(\mathcal{K}, \mathcal{B})$.

By Lemma 4.4, $(\mathcal{K}_1, \mathcal{B}_1) = \int_S (\mathcal{K}_1(w), \mathcal{B}_1(w)) \, d\mu_0(w)$ is an ω-saturated separable pre-complete model of T, and the type space $S_1(T^R(\mathcal{K}_1))$ is separable. Since $(\mathcal{K}, \mathcal{B})$ is separable, $(\mathcal{K}, \mathcal{B})$ is elementarily pre-embeddable in $(\mathcal{K}_1, \mathcal{B}_1)$. Hence the type space $S_1(T^R(\mathcal{K}))$ is separable. This shows that T^R is ω-stable. □

5. Independence and Stability

5.1. *Conditional expectations and fiber products of measures*

We have seen earlier that types in the sense of T^R are measures on the type spaces of T. Here we will give a few results in an abstract measure-theoretic setting that will be useful later on.

Fact 5.1. Let (X, Σ_X) and (Y, Σ_Y) be two measurable spaces, $\pi \colon X \to Y$ a measurable function and let μ a probability measure on X. Let $\hat{\pi}(\mu) = \mu \circ \pi^{-1}$ denote the image measure on Y.

Then for every $f \in L^p(X, \Sigma_X, \mu)$ there is a unique (up to equality a.e.) function $g \in L^p(Y, \Sigma_Y, \hat{\pi}(\mu))$ such that:

$$\int_S g \, d\hat{\pi}(\mu) = \int_{\pi^{-1}(S)} f \, d\mu \quad \forall S \in \Sigma_Y$$

This unique function will be denoted $g = \mathbb{E}^\mu[f|\pi]$. The mapping $\mathbb{E}^\mu[\cdot|\pi]$ has all the usual properties of conditional expectation: it is additive, linear in the sense that $\mathbb{E}^\mu[(h \circ \pi) \cdot f|\pi] = h \cdot \mathbb{E}^\mu[f|\pi]$ for $h \in L^q(Y, \Sigma_Y, \hat{\pi}(\mu))$, satisfies Jensen's inequality and so on.

Proof. Identical to the classical proof of the existence of conditional expectation. Indeed, the classical case is merely the one where the underlying sets X and Y are equal and π is the identity. □

When the mapping $\pi \colon X \to Y$ is clear from the context we may write $\mathbb{E}^\mu[\cdot|Y]$ instead of $\mathbb{E}^\mu[\cdot|\pi]$.

Let us add assumptions. Now X and Y are compact Hausdorff topological spaces, equipped with their respective σ-algebras of Borel sets, and $\pi \colon X \to Y$ is continuous (so in particular Borel measurable). Let $\mathfrak{R}(X)$ denote the set of regular Borel probability measures on X, and similarly $\mathfrak{R}(Y)$ and so on. Since π is continuous the image measure of a regular Borel measure is regular as well, so we have $\hat{\pi} \colon \mathfrak{R}(X) \to \mathfrak{R}(Y)$.

Notation 5.2. For a continuous function $\varphi \colon X \to \mathbb{R}$ and $\mu \in \mathfrak{R}(X)$ let us define

$$\langle \varphi, \mu \rangle = \int \varphi \, d\mu.$$

We equip $\mathfrak{R}(X)$ with the weak-$*$ topology, namely with the minimal topology under which the mapping $\langle \varphi, \cdot \rangle \colon \mu \mapsto \int \varphi \, d\mu$ is continuous for every continuous $\varphi \colon X \to [0,1]$ (equivalently, for every continuous $\varphi \colon X \to \mathbb{R}$). It is a classical fact that this is a compact Hausdorff topology.

Fact 5.3. (Riesz Representation Theorem) The mapping $\mu \mapsto \langle \cdot, \mu \rangle$ defines a bijection between $\mathfrak{R}(X)$ and the positive linear functionals on $C(X, \mathbb{R})$ satisfying $\lambda(1) = 1$.

Proof. Rudin [8, Theorem 2.14]. $\qquad\square$

Equipping the space of positive functionals with the topology of pointwise convergence, this bijection is a homeomorphism, whence the compactness of $\mathfrak{R}(X)$ follows easily.

Let us now consider a "conditional" variant of the functional $\langle \varphi, \cdot \rangle$. For $\nu \in \mathfrak{R}(Y)$ let $\mathfrak{R}_\nu(X)$ denote the fiber above ν. For a continuous $\varphi \colon X \to [0,1]$ and $\mu \in \mathfrak{R}_\nu(X)$ let $\hat{\varphi}_\nu(\mu) = \mathbb{E}^\mu[\varphi | Y]$.

Lemma 5.4. *Let $\varphi \colon X \to [0,1]$ be a continuous function and $\nu \in \mathfrak{R}(Y)$. Then $\hat{\varphi}_\nu \colon \mathfrak{R}_\nu(X) \to L^1(Y, \nu)$ is continuous where $L^1(Y, \nu)$ is equipped with the weak topology. In other words, for every $\psi \in L^\infty(Y, \nu)$ the mapping $\mu \mapsto \int_Y \mathbb{E}^\mu[\varphi | Y] \psi \, d\nu$ is continuous.*

Proof. Let $\psi \in L^\infty(Y, [0,1])$, and let $\varepsilon > 0$ be given. Since ν is regular there is a continuous function $\tilde{\psi} \colon Y \to \mathbb{C}$ which is close in ν to ψ, i.e. such that $\nu\{y \colon |\psi(y) - \tilde{\psi}(y)| > \varepsilon\} < \varepsilon$. Then for every measurable $\chi \colon Y \to [0,1]$:

$$\left| \int_Y \chi \psi \, d\nu - \int_Y \chi \tilde{\psi} \, d\nu \right| < 2\varepsilon.$$

The product $\tilde{\varphi} = \varphi \cdot \tilde{\psi} \circ \pi$ is continuous on X so μ admits a neighborhood

$$U = \{\mu' \in \mathfrak{R}_\nu(X) \colon |\langle \tilde{\varphi}, \mu \rangle - \langle \tilde{\varphi}, \mu' \rangle| < \varepsilon\}.$$

We conclude observing that for every $\mu' \in U$

$$\left| \int_Y \mathbb{E}^\mu[\varphi | Y] \psi \, d\nu - \int_Y \mathbb{E}^{\mu'}[\varphi | Y] \psi \, d\nu \right| < \left| \int_Y \mathbb{E}^\mu[\varphi | Y] \tilde{\psi} \, d\nu - \int_Y \mathbb{E}^{\mu'}[\varphi | Y] \tilde{\psi} \, d\nu \right| + 4\varepsilon$$

$$= |\langle \tilde{\varphi}, \mu \rangle - \langle \tilde{\varphi}, \mu' \rangle| + 4\varepsilon < 5\varepsilon. \qquad\square$$

Sometimes we will find ourselves in a situation where we have some constraints on a measure and we wish to decide whether there is a measure $\mu \in \mathfrak{R}(X)$ satisfying these constraints. This will usually take the following form:

Proposition 5.5. *Let X be a compact Hausdorff space, $A \subseteq C(X, \mathbb{R})$ any subset and $\lambda_0 \colon A \to \mathbb{R}$ any mapping. Then the following are equivalent:*

(1) *There exists a measure $\mu \in \mathfrak{R}(X)$ satisfying $\langle \varphi, \mu \rangle \leq \lambda_0(\varphi)$ for every $\varphi \in A$.*
(2) *Whenever $\{\varphi_i\}_{i < \ell} \subseteq A$ are such that $\sum_{i < \ell} \varphi_i \geq n$ we also have $\sum_{i < \ell} \lambda_0(\varphi_i) \geq n$.*

Proof. One direction is clear so we prove the other. Let S denote the set of all partial functions λ which satisfy the condition in the second item, noticing that it is equivalent to:

$$\sum_{i<\ell} \alpha_i \lambda(\varphi_i) \geq \inf_{x \in X} \sum_{i<\ell} \alpha_i \varphi_i(x) \quad \text{for all } \{(\alpha_i, \varphi_i)\}_{i<\ell} \subseteq \mathbb{R}^+ \times \mathrm{dom}(\lambda).$$

For $\lambda \in S, \psi \in C(X, \mathbb{R})$ and $\beta \in \mathbb{R}$ define $\lambda_{\psi,\beta}$ be the partial functional which coincides with λ except at ψ and satisfying $\lambda_{\psi,\beta}(\psi) = \beta$ (so $\mathrm{dom}(\lambda_{\psi,\beta}) = \mathrm{dom}(\lambda) \cup \{\psi\}$).

Given $\lambda \in S$ and $\psi \in C(X, \mathbb{R})$ one can always find β such that $\lambda_{\psi,\beta} \in S$ as well. For example, $\beta = \sup \psi$ will always do. The least such β is always given by:

$$\tilde{\lambda}(\psi) = \sup \left\{ \inf \left(\psi + \sum_{i<\ell} \alpha_i \varphi_i \right) - \sum_{i<\ell} \alpha_i \lambda(\varphi_i) \colon \{(\alpha_i, \varphi_i)\}_{i<\ell} \subseteq \mathbb{R}^+ \times \mathrm{dom}(\lambda) \right\}.$$

In particular if $\psi \in \mathrm{dom}(\lambda)$ then $\tilde{\lambda}(\psi) \leq \lambda(\psi)$. It is also immediate to check that:

$$\begin{aligned}
\tilde{\lambda}(\alpha) &= \alpha, & \alpha \in \mathbb{R} \\
\tilde{\lambda}(\alpha\psi) &= \alpha\tilde{\lambda}(\psi), & \alpha \in \mathbb{R}^+ \\
\tilde{\lambda}(\psi + \psi') &\geq \tilde{\lambda}(\psi) + \tilde{\lambda}(\psi'), & \\
\tilde{\lambda}(\psi + \psi') &\leq \lambda(\psi) + \lambda(\psi'), & \psi, \psi' \in \mathrm{dom}(\lambda).
\end{aligned}$$

For $\lambda, \lambda' \in S$ say that $\lambda \preceq \lambda'$ if $\mathrm{dom}(\lambda) \subseteq \mathrm{dom}(\lambda')$ and $\lambda(\varphi) \geq \lambda'(\varphi)$ for every $\varphi \in \mathrm{dom}(\lambda)$. By Zorn's Lemma $(S, \preceq)$ admits a maximal element $\lambda \succeq \lambda_0$. By its maximality λ is total and satisfies $\tilde{\lambda}(\varphi) = \lambda(\varphi)$.

It follows that $\lambda(\varphi + \psi) = \lambda(\varphi) + \lambda(\psi)$ and that $\lambda(\alpha\varphi) = \alpha\lambda(\varphi)$, first for $\alpha \geq 0$ and then for every $\alpha \in \mathbb{R}$. Therefore λ is a linear functional, and it is positive as it belongs to S. Since in addition $\lambda(1) = 1$, it is necessarily of the form $\lambda = \langle \cdot, \mu \rangle \colon \varphi \mapsto \langle \varphi, \mu \rangle$ for some measure $\mu \in \mathfrak{R}(X)$. Thus $\langle \cdot, \mu \rangle \succeq \lambda_0$, so μ is as desired. □

One can also prove a conditional variant of Proposition 5.5 but such a result will not be required here. Instead, let us re-phrase Proposition 5.5 a little:

Proposition 5.6. *Let X be a compact Hausdorff space, $A \subseteq C(X, \mathbb{R})$ any subset and $\lambda_0 \colon A \to \mathbb{R}$ any mapping sending $1 \mapsto 1$. Then the following are equivalent:*

(1) *There exists a measure $\mu \in \mathfrak{R}(X)$ satisfying $\langle \mu, \varphi \rangle = \lambda_0(\varphi)$ for every $\varphi \in A$.*
(2) *Whenever $\{(\varphi_i, m_i)\}_{i<\ell} \subseteq A \times Z$ are such that $\sum_{i<\ell} m_i \varphi_i \geq 0$ we also have $\sum_{i<\ell} m_i \lambda_0(\varphi_i) \geq 0$.*

Proof. One direction is clear. For the other observe that if $\varphi, -\varphi \in A$ then necessarily $\lambda_0(-\varphi) + \lambda_0(\varphi) = 0$. We may therefore define λ_1 whose domain is $A \cup -A$ by $\lambda_1(\varphi) = \lambda_0(\varphi), \lambda_1(-\varphi) = -\lambda_0(\varphi)$. Then λ_1 satisfies the conditions of the previous proposition, and the corresponding measure μ is as desired. □

Let us now consider fiber products of measures. Again we consider general measurable spaces (i.e. not necessarily topological) $X = (X, \Sigma_X)$, $Y = (Y, \Sigma_Y)$,

$Z = (Z, \Sigma_Z)$, equipped with measurable mappings $\pi_X \colon X \to Z$, $\pi_Y \colon Y \to Z$. Let $X \times_Z Y$ denote the set theoretic fiber product:

$$X \times_Z Y = \{(x, y) \in X \times Y \colon \pi_X(x) = \pi_Y(y)\}.$$

Let $\Sigma_0 = \{A \times_Z B \colon A \in \Sigma_X, B \in \Sigma_Y\}$ and let $\Sigma_X \otimes_Z \Sigma_Y$ denote the generated σ-algebra on $X \times_Z Y$. The set $X \times_Z Y$ is thereby rendered a measurable space with the canonical mappings to X and Y measurable (it is the measurable space fiber product).

Let μ and ν be probability measures on X and Y, respectively, with the same image measure on Z: $\hat{\pi}_X(\mu) = \hat{\pi}_Y(\nu)$. For $A \in \Sigma_X$ and $B \in \Sigma_Y$ define:

$$(\mu \otimes_Z \nu)_0(A \times_Z B) = \int_Z \mathbb{P}^\mu[A|Z]\mathbb{P}^\nu[B|Z] \, d\hat{\pi}_X(\mu)$$

$$= \int_A \mathbb{P}^\nu[B|Z] \circ \pi_X \, d\mu$$

$$= \int_B \mathbb{P}^\mu[A|Z] \circ \pi_Y \, d\nu.$$

It is not difficult to check that Σ_0 is a semi-ring and that $(\mu \otimes_Z \nu)_0$ defined in this manner is a σ-additive probability measure on Σ_0. By Carathéodory's theorem it extends to a σ-additive probability measure on a σ-algebra containing Σ_0, and in particular to $\Sigma_X \otimes_Z \Sigma_Y$. We will denote this fiber product measure by $\mu \otimes_Z \nu$. Its image measures on X and Y are μ and ν, respectively. If we let $\mathfrak{M}(X)$ denote the collection of probability measures on X and so on we have obtained a mapping:

$$\mathfrak{M}(X) \times_{\mathfrak{M}(Z)} \mathfrak{M}(Y) \to \mathfrak{M}(X \times_Z Y)$$

$$(\mu, \nu) \mapsto \mu \otimes_Z \nu.$$

Let us switch back to the topological setting, where all spaces are compact Hausdorff spaces equipped with the σ-algebras of Borel sets and the mappings are continuous. In particular, we have $\mathfrak{R}(X) \times_{\mathfrak{R}(Z)} \mathfrak{R}(Y) \subseteq \mathfrak{M}(X) \times_{\mathfrak{M}(Z)} \mathfrak{M}(Y)$. (Going through the construction one should be able to verify that if $(\mu, \nu) \in \mathfrak{R}(X) \times_{\mathfrak{R}(Z)} \mathfrak{R}(Y)$ then $\mu \otimes_Z \nu$ is a regular measure on $X \times_Z Y$ as well, but we will not need this observation.) Given a Borel function $\rho \colon X \times_Z Y \to [0, 1]$ we may define:

$$\hat{\rho} \colon \mathfrak{R}(X) \times_{\mathfrak{R}(Z)} \mathfrak{R}(Y) \to [0, 1]$$

$$(\mu, \nu) \mapsto \mathbb{E}^{\mu \otimes_Z \nu}[\rho].$$

Lemma 5.7. *Let X, Y and Z be compact Hausdorff spaces and $X \times_Z Y$ defined as above. Let $\rho \colon X \times_Z Y \to [0, 1]$ be a Borel function.*

Let us fix $\mu \in \mathfrak{R}(X)$, letting $\eta = \hat{\pi}_X(\mu)$, and assume we can find a countable family of Borel subsets $X_i \subseteq X$ such that

(1) $\mu(\bigcup X_i) = 1$.
(2) *For each i there is a continuous function $\rho_i(y) \colon Y \to [0, 1]$ such that $\rho(x, y) = \rho_i(y)$ for every $(x, y) \in X_i \times_Z Y$.*

Then $\hat{\rho}(\mu, \cdot) \colon \mathfrak{R}_\eta(Y) \to [0, 1]$ is continuous.

Proof. We may assume that all the X_i are disjoint. Then:

$$\hat{\rho}(\mu, \nu) = \int \rho \, d(\mu \otimes_Z \nu)$$

$$= \sum_{i \in I} \int_{X_i \times_Z Y} \rho \, d(\mu \otimes_Z \nu)$$

$$= \sum_{i \in I} \int \rho_i \cdot \mathbb{P}^{\mu \otimes_Z \nu}[X_i \times_Z Y | Y] \, d\nu$$

$$= \sum_{i \in I} \int \mathbb{E}^{\nu}[\rho_i | Z] \cdot \mathbb{P}^{\mu \otimes_Z \nu}[X_i \times_Z Y | Z] \, d\eta$$

$$= \sum_{i \in I} \int \mathbb{E}^{\nu}[\rho_i | Z] \cdot \mathbb{P}^{\mu}[X_i | Z] \, d\eta.$$

Fixing $i \in I$, the mapping $\mathfrak{R}_\eta(Y) \to L^1(Z, \Sigma_Z, \eta)$ defined by $\nu \mapsto \mathbb{E}^\nu[\rho_i | Z]$ is continuous in the weak topology by Lemma 5.4. It follows that $\nu \mapsto \int \mathbb{E}^\nu[\rho_i | Z] \cdot \mathbb{P}^\mu[X_i | Z] \, d\eta$ is continuous. Finally, the series above converges absolutely and uniformly so, whence the continuity of $\nu \mapsto \rho(\mu, \nu)$. $\square$

Proposition 5.8. *Let X, Y and Z be compact Hausdorff spaces and $X \times_Z Y$ defined as above. Let $\rho \colon X \times_Z Y \to [0,1]$ be any function. Assume that:*

(1) *The set X is covered by a countable family of Borel sets $X = \bigcup_{i \in \mathbb{N}} X_i$.*
(2) *Each X_i is covered by a (possibly uncountable) family of relatively open subsets $X_i = \bigcup_{j \in J_i} G_{i,j}$.*
(3) *For each pair (i, j) (where $i \in \mathbb{N}$ and $j \in J_i$) there is a continuous function $\rho_{i,j} \colon Y \to [0,1]$ such that $\rho(x, y) = \rho_{i,j}(y)$ for every $(x, y) \in G_{i,j} \times_Z Y$.*

Then ρ is Borel and for every $\eta \in \mathfrak{R}(Z)$ and $\mu \in \mathfrak{R}_\eta(X)$, the mapping $\hat{\rho}(\mu, \cdot) \colon \mathfrak{R}_\eta(Y) \to [0,1]$ is continuous.

Proof. We may assume that the X_i are all disjoint.

Restricted to $G_{i,j} \times_Z Y$, ρ is continuous. Since each $G_{i,j}$ is relatively open in X_i, the set $G_{i,j} \times_Z Y$ is relatively open in $X_i \times_Z Y$. Thus ρ restricted to $X_i \times_Z Y$ is continuous for each i. Moreover, each $X_i \times_Z Y$ is Borel in $X \times_Z Y$. If follows that ρ is Borel.

Now fix $\eta \in \mathfrak{R}(Z)$, $\mu \in \mathfrak{R}_\eta(X)$. Fix again $i \in \mathbb{N}$. Since μ is regular, $\mu(X_i)$ is the supremum of $\mu(K)$ where $K \subseteq X_i$ is compact. Each such K can be covered by finitely many of the $G_{i,j}$. Thus there is a countable family $J_i^0 \subseteq J_i$ such that $\mu(X_i) = \mu(\bigcup_{j \in J_i^0} G_{i,j})$. It follows that

$$\mu \left(\bigcup_{i \in \mathbb{N}, j \in J_i^0} G_{i,j} \right) = 1.$$

Moreover, each $G_{i,j}$ is relatively open in a Borel set and therefore a Borel set itself. The conditions of the lemma are therefore fulfilled.   □

5.2. *Some reminders regarding stable formulas*

Let T denote a first-order theory. Let $\varphi(x, y)$ be a stable formula in T. We may also assume that T eliminates imaginaries.

Since φ is stable, every φ-type over a model $\hat{p} \in S_\varphi(\mathcal{M})$ is definable by a unique formula $d_{\hat{p}}\varphi(y)$:

$$\hat{p}(x) = \{\varphi(x, b) \leftrightarrow d_{\hat{p}}\varphi(b)\}_{b \in \mathcal{M}}.$$

Let $A \subseteq \mathcal{M}$ denote an algebraically closed set, and we may assume that $\mathcal{M}$ is sufficiently saturated and homogeneous over A. Then for every complete type $p \in S_x(A)$ there exists a unique φ-type $\hat{p} \in S_\varphi(\mathcal{M})$ which is definable over A and compatible with p. The (unique) definition of $\hat{p}$ is also referred to as the φ-definition of p, denoted $d_p\varphi(y)$.

Now let $A \subseteq \mathcal{M}$ be any set, not necessarily algebraically closed, and $p \in S_x(A)$. Let $\hat{P} \subseteq S_\varphi(\mathcal{M})$ denote the set of global φ-types $\hat{p}$ which are definable over $acl(A)$ and compatible with p. Then $\hat{P}$ is non-empty, finite, and $\mathrm{Aut}(\mathcal{M}/A)$ acts transitively on $\hat{P}$. We define a $[0, 1]$-valued continuous predicate $\rho(p, y)$ by:

$$\rho(p, b) = \frac{|\{\hat{p} \in \hat{P} : \hat{p} \vdash \varphi(x, b)\}|}{|\hat{P}|}.$$

We consider $\rho(p, b)$ to be the *probability* that a randomly chosen non-forking φ-extension of p should satisfy $\varphi(x, b)$.

Shelah's notions of local rank $R(\cdot, \varphi)$ and multiplicity $M(\cdot, \varphi)$, both with values in $\mathbb{N} \cup \{\infty\}$, provide a more useful way to calculate $\rho(p, b)$. For our purpose it will be more convenient to define the multiplicity of a partial type $\pi(x)$ at each rank n, denoted $M(\cdot, \varphi, n)$.

(1) If π is consistent, then $R(\pi, \varphi) \geq 0$.
(2) Having defined when $R(\cdot, \varphi) \geq n$, we define $M(\cdot, \varphi, n)$. We say that $M(\cdot, \varphi, n) \geq m$ if there are types $\pi(x) \subseteq \pi_i(x)$ for $i < m$ such that for every $i < j < m$ there is a tuple b_{ij} for which $\pi_i(x) \cup \pi_j(x') \vdash \varphi(x, b_{ij}) \leftrightarrow \neg\varphi(x', b_{ij})$, and in addition $R(\pi_i, \varphi) \geq n$ for all $i < m$.
(3) If $M(\pi, \varphi, n) = \infty$, then $R(\pi, \varphi) \geq n + 1$.

The formula φ is stable if and only if $R(\pi, \varphi)$ is always finite. If $R(\pi, \varphi) = n$, then $M(\pi, \varphi) = M(\pi, \varphi, n)$ is the φ-multiplicity of π. Notice that $n < R(\pi, \varphi) \iff M(\pi, \varphi, n) = \infty$ and $n > R(\pi, \varphi) \iff M(\pi, \varphi, n) = 0$.

If $[\pi] \subseteq S_\varphi(\mathcal{M})$ denoted the (closed) set of global φ-types consistent with π then it is not difficult to check that $R(\pi, \varphi)$ is precisely the Cantor–Bendixson rank of $[\pi]$ in $S_\varphi(\mathcal{M})$. Similarly, $M(\pi, \varphi)$ is the Cantor–Bendixson multiplicity, namely the number of types in $[\pi]$ of rank $R(\pi, \varphi)$. In case $p \in S_x(A)$ is a complete type, then

$\hat{P}$ is precisely the set of φ-types of maximal Cantor–Bendixson rank $n = R(p, \varphi)$ in $[p]$, so $M(p, \varphi) = |\hat{P}|$. The number of $\hat{p} \in \hat{P}$ which contains some instance $\varphi(x, b)$ is then precisely the number of types of rank n in $[p(x) \cup \{\varphi(x, b)\}]$, whereby:

$$\rho(p, b) = \frac{M\big(p \cup \{\varphi(x, b)\}, \varphi, n\big)}{M(p, \varphi, n)}.$$

Finally, let $\xi(x, \bar{a}) \in p$ be such that $M\big(\xi(x, \bar{a}), \varphi, n\big) = M(p, \varphi, n)$, namely, such that $[\xi(x, \bar{a})]$ and $[p]$ contain precisely the same types of rank n. Then we have:

$$\rho(p, b) = \frac{M\big(\xi(x, \bar{a}) \wedge \varphi(x, b), \varphi, n\big)}{M\big(\xi(x, \bar{a}), \varphi, n\big)}.$$

The value $\rho(p, b)$ only depends on $q(y) = tp(b/A)$, so it is legitimate to write $\rho(p, q)$. We may re-write p and q as $p(x, A)$ and $q(y, A)$, where W is a tuple of variables corresponding to A, $p(x, W) \in S_{x,W}(T)$ and $q(y, W) \in S_{y,W}(T)$ are complete types and $p{\restriction}_W = q{\restriction}_W \in S_W(T)$. In other words, $(p, q) \in S_{x,W}(T) \times_{S_W(T)} S_{y,W}(T)$. To avoid this fairly cumbersome notation let us write $\Sigma_{x,y,W}$ for the fiber product $S_{x,W}(T) \times_{S_W(T)} S_{y,W}(T)$.

Conversely, let $(p, q) \in \Sigma_{x,y,W}$ be any pair in the fiber product and let A realize their common restriction to W. Then $\rho\big(p(x, A), q(y, A)\big)$ only depends on (p, q) and not on the choice of A. In other words, $\rho\big(p(x, W), q(y, W)\big)$ makes sense for every $(p, q) \in \Sigma_{x,y,W}$. Henceforth we will therefore consider ρ as a function $\rho \colon \Sigma_{x,y,W} \to [0, 1]$.

The next step is to show that ρ satisfies the assumptions of Proposition 5.8.

We define $\Xi_{n,m}$ as the set of all formulas $\xi(x, \bar{w})$ for which $M\big(\xi(x, \bar{w}), \varphi, n\big) \geq m$ is impossible, namely such that for every choice of parameters $\bar{a}$ in a model of T:

$$M\big(\xi(x, \bar{a}), \varphi, n\big) < m.$$

Clearly, if $\xi \in \Xi_{n,(m+1)}$ then $M\big(\xi(x, \bar{w}), \varphi, n\big) \geq m$ is equivalent to $M\big(\xi(x, \bar{w}), \varphi, n\big) = m$.

Lemma 5.9. *Let $\chi(x, t)$ be any formula. Then the property $M\big(\chi(x, t), \varphi, n\big) \geq m$ is type-definable in t. In other words, there is a partial type $\pi(t)$ such that for every parameter c:*

$$M\big(\chi(x, c), \varphi, n\big) \geq m \Leftrightarrow \pi(c).$$

Proof. Standard.   □

Lemma 5.10. *Assume that $\xi(x, \bar{w}) \in \Xi_{n,(m+1)}$. Then there are formulas $\hat{\xi}_{n,m,\ell}(y, \bar{w})$ for $0 \leq \ell \leq m$ such that:*

(1) *The formulas $\{\hat{\xi}_{n,m,\ell} \colon 0 \leq \ell \leq m\}$ define a partition: one and only one of them holds for every $b, \bar{a}$ in a model of T.*
(2) *Modulo $M\big(\xi(x, \bar{w}), \varphi, n\big) = m$, the formula $\hat{\xi}_{n,m,\ell}(y, \bar{w})$ defines the property $M\big(\xi(x, \bar{w}) \wedge \varphi(x, y), \varphi, n\big) = \ell$.*

Proof. For $0 \leq \ell \leq m$ let $\pi_\ell(\bar{w}, y)$ be the partial type expressing that:

$$M\big(\xi(x, \bar{w}) \wedge \varphi(x, y), \varphi, n\big) \geq \ell \quad \text{and} \quad M\big(\xi(x, \bar{w}) \wedge \neg\varphi(x, y), \varphi, n\big) \geq m - \ell.$$

If $\ell < \ell'$, then $\pi_\ell \wedge \pi_{\ell'}$ imply that $M\big(\xi(x, \bar{w}), \varphi, n\big) \geq \ell' + m - \ell > m$ contradicting the assumption on ξ. Thus $\{\pi_\ell : 0 \leq \ell \leq m\}$ are mutually contradictory. We can therefore find formulas $\hat{\xi}_{n,m,\ell} \in \pi_\ell$ which contradict one another as ℓ varies from 0 to m. We may further replace $\hat{\xi}_{n,m,m}$ with $\neg \bigvee_{\ell < m} \hat{\xi}_{n,m,\ell}$ to achieve a partition. On the other hand, if $M\big(\xi(x, \bar{a}), \varphi, n\big) = m$ and b is arbitrary, then $\pi_\ell(\bar{a}, b)$ holds for precisely one $0 \leq \ell \leq m$, whence the second item. $\qquad\square$

Lemma 5.11. *Assume that $\xi(x, \bar{w}) \in \Xi_{n,(m+1)}$. Then there is a continuous predicate $\hat{\xi}_{n,m}(y, \bar{w})$ taking values in $\{\frac{\ell}{m} : 0 \leq \ell \leq m\}$ such that modulo $M\big(\xi(x, \bar{w}), \varphi, n\big) = m$:*

$$\hat{\xi}_{n,m}(y, \bar{w}) = \frac{M(\xi(x, \bar{w}) \wedge \varphi(x, y), \varphi, n)}{M(\xi(x, \bar{w}), \varphi, n)}.$$

Proof. This is just a re-statement of the previous lemma. $\qquad\square$

We define:

$$P^0_{n,m}(x, W) = \{\neg\xi(x, \bar{w}) : \bar{w} \subseteq W, \xi \in \Xi_{n,m}\}.$$

Lemma 5.12. *Let $p(x, W) = tp(a/A)$ (in a model of T). Then the following are equivalent:*

(1) $M\big(p(x, A), \varphi, n\big) \geq m$.
(2) $p(x, W)$ *contains no member of* $\Xi_{n,m}$.
(3) $p(x, W) \supseteq P^0_{n,m}(x, W)$.

In other words, $P^0_{n,m}(x, W)$ expresses that $M(tp(x/W), \varphi, n) \geq m$.

Proof. The implications (i) $\Rightarrow$ (ii) $\Leftrightarrow$ (iii) are immediate. If $M\big(p(x, A), \varphi, n\big) < m$, then the existence of a certain tree whose leaves satisfy $p(x, A)$ is contradictory. By compactness this is due to some formula $\xi(x, \bar{w}) \in p(x, W)$ in which case $\xi \in \Xi_{n,m}$. $\qquad\square$

For $\xi \in \Xi_{n,(m+1)}$ it is convenient to define:

$$P^\xi_{n,m}(x, W) = P^0_{n,m}(x, W) \cup \{\xi(x, \bar{w})\}.$$

Let also:

$$P_{n,m} = \bigcup_{\xi \in \Xi_{n,(m+1)}} [P^\xi_{n,m}] = [P^0_{n,m}] \backslash [P^0_{n,(m+1)}] \subseteq S_{x,W}(T).$$

This is the collection of types $p(x, W) \in S_{x,W}(T)$ such that $M\big(p(x, W), \varphi, n\big) = m$. While it is in general neither open nor closed, it is locally closed and therefore a Borel set. Thus family $\{P_{n,m} : n < \omega, 1 \leq m < \omega\}$ forms a countable partition of

$S_{x,W}(T)$ into Borel sets. On the other hand, $P_{n,m}$ is a union of relatively clopen subsets of the closed set $[P_{n,m}^0]$. Moreover, let $(p,q) \in \Sigma_{x,y,W}$ and assume that $p(x,W) \models P_{n,m}^\xi(x,W)$, $\xi(x,\bar{w}) \in \Xi_{n,(m+1)}$. Let $\hat{\xi}_{n,m}(y,\bar{w})$ be as in Lemma 5.11. Then $\rho(p,q) = \hat{\xi}_{n,m}(y,\bar{w})^q$.

We have thus checked that $\rho\colon \Sigma_{x,y,W} \to [0,1]$ satisfies the conditions of Proposition 5.8.

5.3. *The main results*

We continue with the same assumptions, namely that T is a classical first-order theory and $\varphi(x,y)$ a stable formula. We shall use boldface characters to denote objects of the randomized realm. Thus a model of T^R may be denoted $\boldsymbol{\mathcal{M}}$, a subset thereof $\mathbf{A}$, a type $\mathbf{p}$, and so on, while objects related to the original theory will be denoted using lightface as usual.

For our purposes here it will be more convenient to consider the theory T^R in a single-sorted language in which we have a predicate symbol $\mathbb{P}[\psi(\bar{x})]$ for every formula $\psi(\bar{x})$ in the original language. Each such predicate is definable in the two-sorted language given earlier by the identity $\mathbb{P}[\psi] = \mu([[\psi]])$. Moreover, T^R admits quantifier elimination in this language as well (for example, since the atomic formulas separate types).

Let us translate the conclusion of Proposition 5.8 applied to ρ to the situation at hand. First of all we know that $\mathfrak{R}\big(S_{x,W}(T)\big) = S_{x,W}(T^R)$ and so on, so:

$$\mathfrak{R}\big(S_{x,W}(T)\big) \times_{\mathfrak{R}\big(S_W(T)\big)} \mathfrak{R}\big(S_{y,W}(T)\big) = S_{x,W}(T^R) \times_{S_W(T^R)} S_{y,W}(T^R).$$

Let $\mathbf{A} \subseteq \boldsymbol{\mathcal{M}} \models T^R$ and let $\mathbf{p}(x,\mathbf{A}) \in S_x(\mathbf{A})$. Let $\mathbf{r}(W) = tp(\mathbf{A})$, so $\mathbf{p}(x,W)$ lies in the fiber above $\mathbf{r}$. Similarly, the fiber of $S_{y,W}(T^R)$ lying above $\mathbf{r}$ can be identified with $S_y(\mathbf{A})$. Then Proposition 5.8 says that the function $\hat{\rho}(\mathbf{p}(x,\mathbf{A}),\cdot)\colon S_y(\mathbf{A}) \to [0,1]$ is continuous. We may therefore identify it with an $\mathbf{A}$-definable predicate:

$$\hat{\rho}\big(\mathbf{p}(x,\mathbf{A}),\mathbf{b}\big) := \hat{\rho}\big(\mathbf{p},tp(\mathbf{b},\mathbf{A})\big).$$

Proposition 5.13. *Assume that* $\mathbf{b} \in \mathbf{A}$. *Then* $\hat{\rho}\big(\mathbf{p}(x,\mathbf{A}),\mathbf{b}\big) = \mathbb{P}[\varphi(x,\mathbf{b})]^{\mathbf{p}(x,\mathbf{A})}$. *In other words,* $\hat{\rho}\big(\mathbf{p}(x,\mathbf{A}),y\big)$ *is the* $\mathbb{P}[\varphi]$*-definition of* $\mathbf{p}(x,\mathbf{A})$.

Proof. Indeed, let $\mathbf{q}(x,W) = tp(\mathbf{b},\mathbf{A})$, and say that $w_0 \in W$ corresponds to $\mathbf{b} \in A$. Then the measure $\mathbf{q}$ is concentrated on those types $q(y,W)$ which satisfy $y = w_0$, and similarly the measure $\mathbf{p} \otimes_W \mathbf{q}$ on $\Sigma_{x,y,W}$ is concentrated on those pairs (p,q) where q is such. For such pairs we have $\rho(p,q) = 1$ if $p \models \varphi(x,w_0)$ and zero otherwise, so

$$\hat{\rho}\big(\mathbf{p}(x,\mathbf{A}),\mathbf{b}\big) = \mathbb{E}^{\mathbf{p} \otimes_W \mathbf{q}}[\rho] = \mathbb{P}[\varphi(x,w_0)]^{\mathbf{p}} = \mathbb{P}[\varphi(x,\mathbf{b})]^{\mathbf{p}(x,\mathbf{A})}. \qquad \square$$

We now have everything we need in order to prove preservation of stability.

Theorem 5.14. *If* $\varphi(x,y)$ *is a stable formula of* T, *then* $\mathbb{P}[\varphi(x,y)]$ *is a stable formula of* T^R. *If* T *is a stable theory then so is* T^R.

Proof. We have shown that if φ is stable, then every $\mathbb{P}[\varphi]$-type is definable, so $\mathbb{P}[\varphi]$ is stable as well. Assume now that T is stable. Then every atomic formula of T^R is stable, so every quantifier-free formula is stable, and by quantifier elimination every formula is stable. Therefore T^R is stable. $\qquad\square$

Given $\mathbf{p}(x, \mathbf{A})$ where $\mathbf{A} \subseteq \mathcal{M}$, define:

$$\hat{\mathbf{p}}(x) = \{\mathbb{P}[\varphi(x, \mathbf{b})] = \hat{\rho}(\mathbf{p}(x, \mathbf{A}), \mathbf{b}) \colon \mathbf{b} \in \mathcal{M}\}.$$

Lemma 5.15. *The set of conditions $\hat{\mathbf{p}}(x)$ is a $\mathbb{P}[\varphi]$-type over $\mathcal{M}$, i.e. $\hat{\mathbf{p}}(x) \in S_{\mathbb{P}[\varphi]}(\mathcal{M})$. Moreover, it is consistent with $\mathbf{p}(x, \mathbf{A})$.*

Proof. It is enough to prove the moreover part. For this matter, it is enough to show that for any finite tuple $\bar{\mathbf{b}} = \mathbf{b}_0, \ldots, \mathbf{b}_{n_1}$ the following is consistent:

$$\mathbf{p}(x, \mathbf{A}) \cup \{\mathbb{P}[\varphi(x, \mathbf{b}_i)] = \hat{\rho}(\mathbf{p}(x, \mathbf{A}), \mathbf{b}_i) \colon i < n_1\}.$$

Let $\mathbf{q}(\bar{y}, W) = tp(\bar{\mathbf{b}}, \mathbf{A})$, $\mathbf{q}_i = tp(\mathbf{b}_i, \mathbf{A})$. Then we need to prove that the following is consistent with T^R:

$$\mathbf{p}(x, W) \cup \mathbf{q}(\bar{y}, W) \cup \{\mathbb{P}[\varphi(x, y_i)] = \hat{\rho}(\mathbf{p}, \mathbf{q}_i) \colon i < n_1\}. \tag{1}$$

Given a sequence of formulas and integer numbers $(\chi_k, m_k)_{k<\ell}$, let $[(\chi_k, m_k)_{k<\ell}]$ denote the formula stating that $\sum_{k<\ell} m_k \mathbb{1}_{\chi_k} \geq 0$ (this is indeed expressible by a first order formula). By Proposition 5.6, in order to show that (1) is consistent with T^R it is enough to check that:

$$T \models [(\varphi(x, y_i), f_i)_{i<n_1}, (\psi_j(x, \bar{w}), g_j)_{j<n_2}, (\chi_k(\bar{y}, \bar{w}), h_k)_{k<n_3}]$$
$$\Downarrow$$
$$\sum_{i<n_1} f_i \hat{\rho}(\mathbf{p}, \mathbf{q}_i) + \sum_{j<n_2} g_j \mathbb{P}[\psi_j(x, \bar{w})]^{\mathbf{P}} + \sum_{k<n_3} h_k \mathbb{P}[\chi_k(\bar{y}, \bar{w})]^{\mathbf{q}} \geq 0.$$

The sum can be rewritten as:

$$\mathbb{E}^{\mathbf{P} \otimes_W \mathbf{q}} \left[\sum_{i<n_1} f_i \rho(p, q{\restriction}_{x, y_i, W}) + \sum_{j<n_2} g_j \mathbb{1}_{\psi_j}(x, \bar{w})^p + \sum_{k<n_3} h_k \mathbb{1}_{\chi_k}(\bar{y}, \bar{w})^q \right].$$

It will therefore be enough to show for every $(p, q) \in \Sigma_{x, \bar{y}, W}$:

$$\sum_{i<n_1} f_i \rho(p, q{\restriction}_{x, y_i, W}) + \sum_{j<n_2} g_j \mathbb{1}_{\psi_j}(x, \bar{w})^p + \sum_{k<n_3} h_k \mathbb{1}_{\chi_k}(\bar{y}, \bar{w})^q \geq 0.$$

We may assume that $q(\bar{y}, W) = tp(\bar{b}, A)$ where $\bar{b}, A \subseteq \mathcal{M} \models T$. Let $\hat{P} \subseteq S_\varphi(\mathcal{M})$ be the set of φ-types compatible with $p(x, A)$ and definable over $acl(A)$. Then the

left-hand side becomes:

$$\sum_{i<n_1} f_i \rho_\varphi(p(x,A), b_i) + \sum_{j<n_2} g_j \mathbb{1}_{\psi_j}(x,\bar w)^p + \sum_{k<n_3} h_k \mathbb{1}_{\chi_k}(\bar y, \bar w)^q$$

$$= \sum_{i<n_1} f_i \frac{\{\hat p \in \hat P \colon \hat p \vdash \varphi(x, b_i)\}}{|\hat P|} + \sum_{j<n_2} g_j \mathbb{1}_{\psi_j}(x,\bar w)^p + \sum_{k<n_3} h_k \mathbb{1}_{\chi_k}(\bar y, \bar w)^q$$

$$= \frac{1}{|\hat P|} \sum_{\hat p \in \hat P} \left(\sum_{i<n_1} f_i \mathbb{1}_\varphi(x, b_i)^{\hat p(x)} + \sum_{j<n_2} g_j \mathbb{1}_{\psi_j}(x,\bar w)^{p(x,W)} \right.$$

$$\left. + \sum_{k<n_3} h_k \mathbb{1}_{\chi_k}(\bar y, \bar w)^{q(\bar y, W)} \right).$$

It will therefore be enough to show that the sum inside the parentheses is positive for every $\hat p \in \hat P$. Since $\hat p$ is compatible with $p(x,A)$, there is a complete type $\tilde r(x, \mathcal{M})$ extending both. Let $r(x, \bar b, A)$ be its restriction to the parameters which interest us, where as usual $r(x, \bar y, W) \in S_{x, \bar y, W}(T)$. Then we have:

$$\sum_{i<n_1} f_i \mathbb{1}_\varphi(x, b_i)^{\hat p(x)} + \sum_{j<n_2} g_j \mathbb{1}_{\psi_j}(x,\bar w)^{p(x,W)} + \sum_{k<n_3} h_k \mathbb{1}_{\chi_k}(\bar y, \bar w)^{q(\bar y, W)}$$

$$= \sum_{i<n_1} f_i \mathbb{1}_\varphi(x, y_i)^r + \sum_{j<n_2} g_j \mathbb{1}_{\psi_j}(x,\bar w)^r + \sum_{k<n_3} h_k \mathbb{1}_{\chi_k}(\bar y, \bar w)^r.$$

Since r is a type of T, this is indeed always positive. $\qquad\square$

Theorem 5.16. *Assume φ is a stable formula of T, and let $\mathbf{A} \subseteq \mathcal{M} \models T^R$. Then every $\mathbb{P}[\varphi]$-type over $\mathbf{A}$ is stationary. If T is stable then every type over $\mathbf{A}$ is stationary.*

Proof. We have shown that if $\mathbf{p}(x) \in S_x(\mathbf{A})$ then there exists a $\mathbb{P}[\varphi]$-type $\hat{\mathbf{p}}(x) \in S_{\mathbb{P}[\varphi]}(\mathcal{M})$ which is definable over $\mathbf{A}$ (rather than merely over $acl(\mathbf{A})$) and compatible with $\mathbf{p}$. $\qquad\square$

Here $\mathbf{A}$ consists of random elements in sorts of T. The theory T^R necessarily introduces new imaginary sorts (even if T admits elimination of imaginaries) and types over elements from these new sorts need not be stationary.

In the course of the proof we have given an explicit characterization of the unique nonforking extension of a type in T^R. Let us restate this characterization in a slightly modified fashion. First of all we observe that the entire development above goes through for a formula of the form $\varphi(x, y, \bar w)$ where $\bar w \subseteq W$ is a fixed tuple of parameter variables (alternatively, we could name the tuple $\bar w$ by new constants). Define $\rho_\varphi \colon \Sigma_{x,y,W} \to [0,1]$ and $\hat\rho_\varphi$ accordingly.

Corollary 5.17. *Let* $\mathbf{c}, \mathbf{b}, \mathbf{A} \subseteq \mathcal{M}$. *Then* $\mathbf{c} \underset{\mathbf{A}}{\downarrow} \mathbf{b}$ *if and only if for every* $\bar{\mathbf{a}} \subseteq \mathbf{A}$ *with corresponding* $\bar{w} \subseteq W$, *and for every formula* $\mathbb{P}[\varphi(x, y, \bar{w})]$:

$$\mathbb{P}[\varphi(\mathbf{c}, \mathbf{b}, \bar{\mathbf{a}})] = \hat{\rho}_\varphi(tp(\mathbf{c}, \mathbf{A}), tp(\mathbf{b}, \mathbf{A})).$$

Acknowledgment

References

1. I. Ben Yaacov, A. Berenstein, C. Ward Henson and A. Usvyatsov, Model theory for metric structures, in *Model Theory with Applications to Algebra and Analysis*, Vol. 2, London Math Society Lecture Note Series, Vol. 350, eds. Z. Chatzidakis, D. Macpherson, A. Pillay and A. Wilkie (Cambridge Univ. Press, 2008), pp. 315–427.
2. I. Ben Yaacov, Continuous and random Vapnik–Chervonenkis classes, *Israel J. Math.*, to appear, arXiv:0802.0068.
3. I. Ben Yaacov, On theories of random variables, arXiv:0901.1584.
4. I. Ben Yaacov, Schrödinger's cat, *Israel J. Math.* **153** (2006) 157–191.
5. I. Ben Yaacov and A. Usvyatsov, Continuous first order logic and local stability, *Trans. Amer. Math. Soc.*, to appear, arXiv:0801.4303.
6. P. R. Halmos, *Measure Theory* (D. Van Nostrand, 1950).
7. H. J. Keisler, Randomizing a model, *Adv. Math.* **143** (1999) 124–158.
8. W. Rudin, *Real and Complex Analysis* (McGraw-Hill, 1966).